

Hardness of Approximation Between P and NP

Aviad Rubinfeld



Hardness of Approximation Between P and NP

ACM Books

Editor in Chief

M. Tamer Özsu, *University of Waterloo*

ACM Books is a series of high-quality books for the computer science community, published by ACM and many in collaboration with Morgan & Claypool Publishers. ACM Books publications are widely distributed in both print and digital formats through booksellers and to libraries (and library consortia) and individual ACM members via the ACM Digital Library platform.

Hardness of Approximation Between P and NP

Aviad Rubinfeld, *Stanford University*

2019

Conversational UX Design: A Practitioner's Guide to the Natural Conversation Framework

Robert J. Moore, *IBM Research–Almaden*

Raphael Arar, *IBM Research–Almaden*

2019

Heterogeneous Computing: Hardware and Software Perspectives

Mohamed Zahran, *New York University*

2019

Making Databases Work: The Pragmatic Wisdom of Michael Stonebraker

Editor: Michael L. Brodie

2018

The Handbook of Multimodal-Multisensor Interfaces, Volume 2: Signal Processing, Architectures, and Detection of Emotion and Cognition

Editors: Sharon Oviatt, *Monash University*

Björn Schuller, *University of Augsburg and Imperial College London*

Philip R. Cohen, *Monash University*

Daniel Sonntag, *German Research Center for Artificial Intelligence (DFKI)*

Gerasimos Potamianos, *University of Thessaly*

Antonio Krüger, *Saarland University and German Research Center for Artificial Intelligence (DFKI)*

2018

Declarative Logic Programming: Theory, Systems, and Applications

Editors: Michael Kifer, *Stony Brook University*

Yanhong Annie Liu, *Stony Brook University*

2018

The Sparse Fourier Transform: Theory and Practice

Haitham Hassanieh, *University of Illinois at Urbana-Champaign*
2018

The Continuing Arms Race: Code-Reuse Attacks and Defenses

Editors: Per Larsen, *Immunant, Inc.*
Ahmad-Reza Sadeghi, *Technische Universität Darmstadt*
2018

Frontiers of Multimedia Research

Editor: Shih-Fu Chang, *Columbia University*
2018

Shared-Memory Parallelism Can Be Simple, Fast, and Scalable

Julian Shun, *University of California, Berkeley*
2017

Computational Prediction of Protein Complexes from Protein Interaction Networks

Sriganesh Srihari, *The University of Queensland Institute for Molecular Bioscience*
Chern Han Yong, *Duke-National University of Singapore Medical School*
Limsoon Wong, *National University of Singapore*
2017

The Handbook of Multimodal-Multisensor Interfaces, Volume 1: Foundations, User Modeling, and Common Modality Combinations

Editors: Sharon Oviatt, *Incaa Designs*
Björn Schuller, *University of Passau and Imperial College London*
Philip R. Cohen, *Voicebox Technologies*
Daniel Sonntag, *German Research Center for Artificial Intelligence (DFKI)*
Gerasimos Potamianos, *University of Thessaly*
Antonio Krüger, *Saarland University and German Research Center for Artificial Intelligence (DFKI)*
2017

Communities of Computing: Computer Science and Society in the ACM

Thomas J. Misa, Editor, *University of Minnesota*
2017

Text Data Management and Analysis: A Practical Introduction to Information Retrieval and Text Mining

ChengXiang Zhai, *University of Illinois at Urbana-Champaign*
Sean Massung, *University of Illinois at Urbana-Champaign*
2016

An Architecture for Fast and General Data Processing on Large Clusters

Matei Zaharia, *Stanford University*
2016

Reactive Internet Programming: State Chart XML in Action

Franck Barbier, *University of Pau, France*

2016

Verified Functional Programming in Agda

Aaron Stump, *The University of Iowa*

2016

The VR Book: Human-Centered Design for Virtual Reality

Jason Jerald, *NextGen Interactions*

2016

Ada's Legacy: Cultures of Computing from the Victorian to the Digital Age

Robin Hammerman, *Stevens Institute of Technology*

Andrew L. Russell, *Stevens Institute of Technology*

2016

Edmund Berkeley and the Social Responsibility of Computer Professionals

Bernadette Longo, *New Jersey Institute of Technology*

2015

Candidate Multilinear Maps

Sanjam Garg, *University of California, Berkeley*

2015

Smarter Than Their Machines: Oral Histories of Pioneers in Interactive Computing

John Cullinane, *Northeastern University; Mossavar-Rahmani Center for Business and Government, John F. Kennedy School of Government, Harvard University*

2015

A Framework for Scientific Discovery through Video Games

Seth Cooper, *University of Washington*

2014

Trust Extension as a Mechanism for Secure Code Execution on Commodity Computers

Bryan Jeffrey Parno, *Microsoft Research*

2014

Embracing Interference in Wireless Systems

Shyamnath Gollakota, *University of Washington*

2014

Hardness of Approximation Between P and NP

Aviad Rubinfeld

Stanford University

ACM Books #24



Copyright © 2019 by Association for Computing Machinery

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means—electronic, mechanical, photocopy, recording, or any other except for brief quotations in printed reviews—without the prior permission of the publisher.

Designations used by companies to distinguish their products are often claimed as trademarks or registered trademarks. In all instances in which the Association for Computing Machinery is aware of a claim, the product names appear in initial capital or all capital letters. Readers, however, should contact the appropriate companies for more complete information regarding trademarks and registration.

Hardness of Approximation Between P and NP

Aviad Rubinfeld

books.acm.org

<http://books.acm.org>

ISBN: 978-1-94728-723-9 hardcover

ISBN: 978-1-94728-720-8 paperback

ISBN: 978-1-94728-722-2 ePub

ISBN: 978-1-94728-721-5 eBook

Series ISSN: 2374-6769 print 2374-6777 electronic

DOIs:

10.1145/3241304 Book	10.1145/3241304.3241314 Chapter 9
10.1145/3241304.3241305 Preface	10.1145/3241304.3241315 Chapter 10
10.1145/3241304.3241306 Chapter 1	10.1145/3241304.3241316 Chapter 11
10.1145/3241304.3241307 Chapter 2	10.1145/3241304.3241317 Chapter 12
10.1145/3241304.3241308 Chapter 3	10.1145/3241304.3241318 Chapter 13
10.1145/3241304.3241309 Chapter 4	10.1145/3241304.3241319 Chapter 14
10.1145/3241304.3241310 Chapter 5	10.1145/3241304.3241320 Chapter 15
10.1145/3241304.3241311 Chapter 6	10.1145/3241304.3241321 Chapter 16
10.1145/3241304.3241312 Chapter 7	10.1145/3241304.3241322 References/Index/Bios
10.1145/3241304.3241313 Chapter 8	

A publication in the ACM Books series, #24

Editor in Chief: M. Tamer Özsu, *University of Waterloo*

This book was typeset in Arnhem Pro 10/14 and Flama using ZzT_EX.

First Edition

10 9 8 7 6 5 4 3 2 1

Contents

Preface **xiii**

Acknowledgments **xiv**

PART I OVERVIEW 1

Chapter 1 The Frontier of Intractability 3

1.1 PPAD: Finding a Needle You *Know* Is in the Haystack **5**

1.2 Quasi-Polynomial Time and the Birthday Paradox **12**

1.3 Approximate Nash Equilibrium **17**

Chapter 2 Preliminaries 19

Notation **19**

2.1 Nash Equilibrium and Relaxations **19**

2.2 PPAD and END-OF-A-LINE **20**

2.3 Exponential Time Hypotheses **22**

2.4 PCP Theorems **22**

2.5 Learning Theory **24**

2.6 Information Theory **25**

2.7 Useful Lemmata **27**

PART II COMMUNICATION COMPLEXITY 31

Chapter 3 Communication Complexity of Approximate Nash Equilibrium 33

Our Results **34**

3.1 Uncoupled Dynamics **35**

3.2 Techniques **36**

3.3 Additional Related Literature **38**

3.4 Proof Overview **39**

- 3.5 Proofs 44
- 3.6 An Open Problem: Correlated Equilibria in 2-Player Games 60

Chapter 4 Brouwer’s Fixed Point 61

- 4.1 BROUWER with ℓ_∞ 61
- 4.2 Euclidean BROUWER 68

PART III PPAD 77

Chapter 5 PPAD-Hardness of Approximation 79

Chapter 6 The Generalized Circuit Problem 83

Our Results 84

- 6.1 Proof Overview 85
- 6.2 From Brouwer to ϵ -GCIRCUIT 87
- 6.3 GCIRCUIT with Fan-out 2 99

Chapter 7 Many-Player Games 101

Related Works: Tractable Special Cases 101

- 7.1 Graphical, Polymatrix Games 101
- 7.2 Succinct Games 107

Chapter 8 Bayesian Nash Equilibrium 111

Chapter 9 Market Equilibrium 115

- 9.1 Why Are Non-Monotone Markets Hard? 116
- 9.2 High-Level Structure of the Proof 116
- 9.3 Adaptations for Constant Factor Inapproximability 117
- 9.4 Non-Monotone Markets: Proof of Inapproximability 119

Chapter 10 CourseMatch 131

- 10.1 The Course Allocation Problem 131
- 10.2 A-CEEI Is PPAD-Hard 133
- 10.3 A-CEEI \in PPAD 139

PART IV QUASI-POLYNOMIAL TIME 147

Chapter 11 Birthday Repetition 149

- 11.1 Warm-Up: Best ϵ -Nash 150

Chapter 12	Densest k-Subgraph	155
12.1	Construction (and Completeness)	159
12.2	Soundness	160
Chapter 13	Community Detection	177
13.1	Related Works	178
13.2	Overview of Proofs	179
13.3	Hardness of Counting Communities	182
13.4	Hardness of Detecting Communities	183
Chapter 14	VC and Littlestone's Dimensions	187
14.1	Discussion	188
14.2	Techniques	189
14.3	Related Work	190
14.4	Inapproximability of the VC Dimension	191
14.5	Inapproximability of Littlestone's Dimension	200
14.6	Quasi-polynomial Algorithm for Littlestone's Dimension	211
Chapter 15	Signaling	213
15.1	Techniques	214
15.2	Near-Optimal Signaling Is Hard	215
PART V	APPROXIMATE NASH EQUILIBRIUM	221
Chapter 16	2-Player Approximate Nash Equilibrium	223
	Additional Related Work	224
16.1	Technical Overview	225
16.2	END-OF-A-LINE with Local Computation	232
16.3	Holographic Proof	236
16.4	Polymatrix WeakNash	248
16.5	From Polymatrix to Bimatrix	269
	References	273
	Index	291
	Author Biography	303

Preface

This book is based on my Ph.D. thesis (by the same title), which I wrote at the University of California, Berkeley, with the wise guidance of Christos Papadimitriou.

In hindsight, it seems only natural that I ended up working on the complexity of Nash equilibrium, a problem that has been so closely associated with Christos for three decades. But the true story of how I came to work on equilibrium computation is completely different. During my first winter break at Berkeley, Abe Othman contacted me¹ and asked about the complexity of finding an *Approximate Competitive Equilibrium from Equal Incomes* (A-CEEI). At the time, Abe was preparing for the launch of Course Match, a system that implements a heuristic algorithm for A-CEEI to assign students to courses at the Wharton School (University of Pennsylvania). Chapter 10, based on joint work with Abe and Christos, is devoted to the complexity of A-CEEI.

A-CEEI, like the rest of the problems in Part III, belongs to a special class (PPAD) of intractable problems that are *total*: their solution always exists, but it is hard to find it. Given our current understanding of computational complexity, this implies that they are unlikely to be NP-hard. In Part IV we move on to a different category of intractable problems: those that can be solved in *quasi-polynomial* ($n^{\log(n)}$) time; this is typically prohibitively slow to implement even in moderate sized instances, but it is much faster than what we know or conjecture for problems like Satisfiability. This suggests that these problems are also not NP-hard. In both cases, proving intractability for problems that are not NP-hard requires both the right conditional assumptions ($P \neq \text{PPAD}$ and the “Exponential Time Hypothesis”), as well as special techniques.

The most difficult result in this book, for approximate Nash equilibrium (Part V), lies in the intersection of the PPAD problems from Part III and the quasi-polynomial

1. Thanks to Anindya De for making the introduction!

time problems from Part [IV](#). It combines ideas from both parts, together with new and old techniques from across theoretical computer science. It resolves an outstanding open problem that, as a student, I didn't really mean to solve. Well . . . I'm glad I tried!

Acknowledgments

This book is based on my Ph.D. dissertation [[Rubinstein 2017c](#)], completed at University of California, Berkeley, with the wise advice of Christos Papadimitriou. My dissertation, in turn, was based on the following publications:

- “Inapproximability of VC Dimension and Littlestone’s Dimension” [[Manurangsi and Rubinstein 2017](#)].
- “Honest Signaling in Zero-Sum Games Is Hard, and Lying Is Even Harder” [[Rubinstein 2017b](#)].
- “Detecting Communities Is Hard (And Counting Them is Even Harder)” [[Rubinstein 2017a](#)].
- “ETH Hardness for Densest- k -Subgraph with Perfect Completeness” [[Braverman et al. 2017](#)].
- “Communication Complexity of Approximate Nash Equilibria” [[Babichenko and Rubinstein 2017](#)].
- “Settling the Complexity of Computing Approximate Two-Player Nash Equilibria” [[Rubinstein 2016](#)].
- “Inapproximability of Nash Equilibrium” [[Rubinstein 2015](#)].
- “The Complexity of Fairness Through Equilibrium” [[Othman et al. 2016](#)].

I am indebted to Christos and all my co-authors: Yakov Babichenko, Mark Braverman, Young Kun Ko, Pasin Manurangsi, Abraham Othman, and Omri Weinstein.

I am also grateful to Boaz Barak, Shai Ben-David, Jonah Brown-Cohen, Karthik C. S., Yang Cai, Alessandro Chiesa, Paul Christiano, Constantinos Daskalakis, Shaddin Dughmi, Mika Goos, Rishi Gupta, Elad Haramaty, Noam Nisan, Prasad Raghavendra, Tselil Schramm, Madhu Sudan, Luca Trevisan, Michael Viderman, and anonymous reviewers, for commenting on, correcting, and inspiring much of the content of this book.

The writing of this book was supported in part by the MSR Ph.D. Fellowship (and wonderful MSR summer internships).

This has been an amazing journey. I am most grateful for the opportunities I had to share it with my family and friends.

Aviad Rubinstein

Cambridge, MA

March 2019



PART

OVERVIEW

The Frontier of Intractability

The combination of vast amounts of data, unprecedented computing power, and clever algorithms allows today's computer systems to drive autonomous cars, beat the best human players at chess and Go, and live stream videos of cute cats across the world. Yet computers can fail miserably at predicting outcomes of social processes and interactions, elections being only the most salient example. And this is hardly surprising, as there are two potent reasons for such unpredictability: One reason is that human agents are driven by emotions and hormones and billions of neurons interacting with a complex environment, and as a result their behavior is extremely hard to model mathematically. The second reason is perhaps a bit surprising, and certainly closer to the concerns of this book: *Even very simple and idealized models of social interaction, in which agents have extremely clear-cut objectives, can be intractable.*

To have any hope of reasoning about human behavior on national and global scales, we need a mathematically sound theory. Game theory is the mathematical discipline that models and analyzes the interaction between agents (e.g., voters) who have different goals and are affected by each other's actions. The central solution concept in game theory is the *Nash equilibrium*. It has an endless list of applications in economics, politics, biology, etc. (e.g., [[Aumann 1987](#)]).

By Nash's theorem [[Nash 1951](#)], an equilibrium always exists. Furthermore, once at a Nash equilibrium, players have no incentive to deviate. The main missing piece in the puzzle is:

How do players arrive at an equilibrium in the first place?

After many attempts by economists for more than six decades¹ (e.g., [Brown 1951, Foster and Young 2006, Hart and Mas-Colell 2003, Kalai and Lehrer 1993, Lemke and Howson 1964, Robinson 1951, Scarf 1967, Shapley 1964]), we still don't have a satisfying explanation.

About ten years ago, the study of this fundamental question found a surprising answer in computer science: Chen et al. [2009b] and Daskalakis et al. [2009a] proved that finding a Nash equilibrium is computationally intractable.² And if no centralized, specialized algorithm can find an equilibrium, it is even less likely that distributed, selfish agents will naturally converge to one. This casts doubt over the entire solution concept.

For the past decade, the main remaining hope for Nash equilibrium has been *approximation*. The central open question in this field has been:

Is there an efficient algorithm for finding an approximate Nash equilibrium?

We give a strong negative resolution to this question: our main result rules out efficient approximation algorithms for finding Nash equilibria.³

Our theorem is the latest development in a long and intriguing technical story. The first question we computer scientists ask when encountering a new algorithmic challenge is: is it in P, the class of polynomial time (tractable) problems; or is it NP-hard, like Satisfiability and the Traveling Salesperson Problem (where the best known algorithms require exponential time)? Approximate Nash equilibrium falls into neither category; its complexity lies between P and NP-hard—hence the title of our book. Let us introduce two (completely orthogonal) reasons why we do not expect it to be NP-hard.

The first obstacle for NP-hardness is the *totality* of Nash equilibrium. When we say that Satisfiability or the Traveling Salesperson Problems are NP-hard, we formally mean that it is NP-hard to determine *whether* a given formula has a satisfying assignment, or *whether* a given network allows the salesperson to complete her travels within a certain budget. In contrast, by Nash's theorem an equilibrium always exists. Deciding whether an equilibrium exists is trivial, but we still don't know how to find one. This is formally captured by an intermediate complexity class called PPAD.

1. In fact, "more than six decades" is an understatement: Irving Fisher's thesis from 1891 dealt with the closely related question of convergence to market equilibria [Brainard and Scarf 2000].

2. Assuming $P \neq PPAD$; see the discussion in the next section about this assumption.

3. Under a complexity assumption stronger than $P \neq PPAD$ that we call the Exponential Time Hypothesis (ETH) for PPAD; see Section 1.2 for details.

The second obstacle for NP-hardness is an algorithm (the worst kind of obstacle for intractability). The best known algorithms for solving NP-hard or PPAD-hard problems require exponential ($\approx 2^n$) time, and there is a common belief (formulated a few years ago as the “Exponential Time Hypothesis” [Impagliazzo et al. 2001]) that much faster algorithms do not exist. In contrast, an approximate Nash equilibrium can be found in *quasi-polynomial* ($\approx n^{\log n}$) time. Notice that this again places the complexity of approximate Nash equilibrium between the polynomial time solvable problems in P and the exponential time required by NP-hard problems. Therefore, approximate Nash equilibrium is unlikely to be NP-hard—or even PPAD-hard, since we know of no quasi-polynomial algorithm for any other PPAD-hard problem.

As illustrated in the last two paragraphs, approximate Nash equilibrium is very far from our standard notion of intractability, NP-hardness. In some sense, it is one of the computationally easiest problems that we can still prove to be intractable—it lies at the *frontier of our understanding of intractability*. Unsurprisingly, the techniques we had to master to prove the intractability of approximate Nash equilibrium are useful for many other problems. In particular, we also prove hardness of approximation for several other interesting problems that either belong to the class PPAD or have quasi-polynomial time algorithms.

1.1

PPAD: Finding a Needle You *Know* Is in the Haystack

Consider a directed graph G . Each edge contributes to the degree of the two vertices incident on it. Hence the sum of the degrees is twice the number of edges, and in particular it is even. Now, given an odd-degree vertex v , there must exist another odd-degree vertex. But *can you find one?* There is, of course, the trivial algorithm, which simply brute-force enumerates over all the graph vertices.

Now suppose G further has the property⁴ that every vertex has in- and out-degree at most 1. Thus, G is a disjoint union of lines and cycles. If v has an outgoing edge but no incoming edge, it is the beginning of a line. Now the algorithm for finding another odd degree node is a bit more clever—follow the path to the end—but its worst case is still linear in the number of nodes.

In the worst case, both algorithms run in time linear in the number of vertices. When the graph is given explicitly, e.g., as an adjacency matrix, this is quite efficient compared to the size of the input. But what if the graph is given as black-box oracles

4. This guarantee is actually without loss of generality [Papadimitriou 1994], but this is not so important for our purposes.

S and P that return, for each vertex, its successor and predecessor in the graph? The amount of work⁵ required for finding another odd-degree vertex by the exhaustive algorithm, or the path-following algorithm, is still linear in the number of vertices, but now this number is exponential in the natural parameter of the problem, the length of the input to the oracle, which equals the *logarithm* of the number of nodes. In the oracle model, it is not hard to prove that there are no better algorithms.

Let us consider the computational analog of the black-box oracle model: the oracles S and P are implemented as explicit (“white-box”) circuits, which are given as the input to the algorithm. This is the END-OF-A-LINE problem, which is the starting point for most of our reductions. The computational class PPAD (which stands for Polynomial Parity Argument in Directed graphs) is the class of search problems reducible to END-OF-A-LINE.

Definition 1.1 END-OF-A-LINE [Daskalakis et al. 2009a]. Given two circuits S and P , with m input bits and m output bits each, such that $P(0^m) = 0^m \neq S(0^m)$, find an input $x \in \{0, 1\}^m$ such that $P(S(x)) \neq x$ or $S(P(x)) \neq x \neq 0^m$.

How hard is END-OF-A-LINE? We believe that it is likely to be very hard, almost as hard as the black-box variant. This belief is backed by relativized separations [Morioka 2003] and cryptographic hardness [Bitansky et al. 2015, Garg et al. 2016, Hubáček and Yogev 2017], as well as zero algorithmic progress, even on special cases, since it was first introduced in Papadimitriou [1994]. More importantly, we are embarrassingly bad at gaining insight to a circuit’s functionality by looking at its structure (with some notable exceptions in cryptography [Barak 2004])— hence it seems reasonable to conjecture that the white-box variant is easier than the black-box problem.

Even more embarrassing is our inability to prove computational intractability. Essentially all “proofs” of computational intractability (including the ones in this book) are *conditional*; i.e., we assume that some problem is hard (e.g., END-OF-A-LINE or SATISFIABILITY), and reduce it to a new problem we want to “prove” is hard. The intractability of the new problem only holds conditioned on the intractability of the original hard problem. Without first resolving whether $P = NP$, we cannot prove that END-OF-A-LINE is indeed hard. Furthermore, even merely proving that END-OF-A-LINE is NP-hard would already imply unlikely consequences like $NP = \text{coNP}$ [Megiddo and Papadimitriou 1991]. The ultimate reason we believe that END-OF-A-LINE is intractable is that we have no idea how to prove it.

5. Here *work* is measured by the number of oracle calls rather than running time; indeed this model will be the starting point of our reductions for query and communication complexity.

Once we believe that END-OF-A-LINE is indeed intractable, any problem that is PPAD-complete, i.e., “at least as hard as END-OF-A-LINE,” is also intractable. The celebrated works of [Chen et al. \[2009b\]](#) and [Daskalakis et al. \[2009a\]](#) prove that finding an exact Nash equilibrium is PPAD-complete. In Section 1.3 we describe some variants of approximate Nash equilibrium that are also PPAD-complete.

We conclude this section with a few problems (other than variants of Nash equilibrium) that we show are also PPAD-hard:

- Finding an approximate fixed point of an implicit function; Brouwer’s fixed point theorem, which guarantees the existence of a fixed point, lays the mathematical foundation for the rest of our PPAD-complete problems.
- Finding an approximate market equilibrium, which is the conceptual foundation of neoclassical economics.
- Finding an Approximate Competitive Equilibrium from Equal Incomes (ACEEI), which is an algorithmic problem of practical interest due to its use for allocating classes to students (CourseMatch).

1.1.1 Brouwer’s Fixed Point

Brouwer’s fixed point theorem, together with its generalizations (in particular, Kakutani’s fixed point theorem), is the basis for many of the equilibrium concepts in game theory and economics. It states that any continuous function f from a compact convex set (in particular, the n -dimensional hypercube) to itself has a fixed point, i.e., a point x^* such that $f(x^*) = x^*$. Just like in the case of a Nash equilibrium and the odd-degree vertex, the existence does not come with an algorithm for finding the fixed point. Brouwer eventually became so unsatisfied with the non-constructive nature of his proof that he founded intuitionism, a formalism of mathematics mandating constructive proofs [[Iemhoff 2016](#)].

Before we discuss the computational tractability of finding a fixed point, there are two subtle but important technicalities that we have to clear. The first is that finding an exact fixed point may be “obviously” intractable when all the fixed points are irrational. Fortunately there is a natural notion⁶ of approximate fixed point: find x such that $f(x) \approx x$. The second issue is that we want a finite description of the input to our algorithm. Naively, one can define the value of the function on a grid, but then the set is no longer convex and a fixed point might not exist at all (see, e.g., [Roughgarden and Weinstein \[2016\]](#)). It turns out that a better way to define

6. There is also a stricter notion that requires x to be close to an x^* for which $f(x^*) = x^*$ exactly. See, e.g., [Etessami and Yannakakis \[2010\]](#).

the input is as an arithmetic circuit: this gives a local and easy-to-verify guarantee that the function is indeed continuous.

Admittedly, there is something dissatisfying about proving computational hardness of “circuit problems”: we are going to prove that this problem is PPAD-hard, i.e., no easier than END-OF-A-LINE— but the reason we believe in the first place that END-OF-A-LINE is intractable is that we don’t know how to gain insight to the functionality of the circuit from staring at its structure. Nevertheless, we begin with the complexity of Brouwer’s fixed point because: this is a fundamental question; it is the starting point of many of our reductions; and, as we discuss later, even the black-box hardness in this case is highly non-trivial.

We know of two algorithms for computing an approximate Brouwer fixed point: there is Scarf’s algorithm [Scarf 1967], but its running time may be exponential in the description of the function; the same holds for brute-force search.

On the complexity side, Daskalakis et al. [2009a] proved that even in three dimensions, finding an exponentially close approximation (x such that $\|f(x) - x\|_\infty \leq 2^{-n}$, where n is the size of the input) is PPAD-complete; Chen and Deng [2009] proved the same problem continues to be hard even with only two dimensions. Notice that with a constant number of dimensions, any weaker approximation desideratum becomes trivial for brute-force search. Chen et al. [2009b] showed that in n dimensions, polynomial approximations are also PPAD-complete. We will later prove that even constant approximations are PPAD-complete, i.e., there exists some absolute constant $\varepsilon > 0$ such that it’s hard to find an x for which $\|f(x) - x\|_\infty \leq \varepsilon$; this is already a significant improvement over the existing state of the art.

We can prove an even stronger form of inapproximability for finding a Brouwer fixed point: so far, we characterized the approximate fixed point with ℓ_∞ -norm, which means that $f(x)$ and x must be close in every coordinate. Our main result in this regard (Theorem 4.2) is that even if we only require that $f(x)$ and x are close in *most* of the coordinates, finding such x and $f(x)$ remains PPAD-complete.

1.1.2 Market Equilibrium

Supply and demand is central to our modern understanding of economics: when demand exceeds supply, raising the price makes it less attractive to consumers and more attractive to producers, thereby reducing demand and increasing supply. Vice versa if the supply exceeds the demand. This simple idea marks the birth of neoclassical economics [Jevons 1866, Menger 1871, Walras 1874], and continues to inspire us today when we reason about free market economies. Since the consumption and production of one good depends on other goods, we are interested in a *market equilibrium*: a vector of prices where the supply and demand of every good matches.

The supply and demand argument has many weak spots, one of which is Giffen goods [Marshall 1895]: Consider a poor 19th-century English family consuming 10,000 calories a day from meat and bread. They prefer to eat meat, but the bread is cheaper—so they spend their daily income on 6 loaves of bread and 4 pounds of meat (each provides 1,000 calories). What happens if the price of bread increases? They have less free money to spend on the preferred luxury good (meat), which *increases* their demand for the Giffen good (bread). The existence of Giffen goods in practice has been contentiously debated for over a century [Edgeworth 1909, Stigler 1947], with evidence ranging from econometric analysis of historical market data to experiments in lab rats [Battalio et al. 1991].

Despite counterintuitive issues like Giffen goods, Arrow and Debreu proved that, under quite general conditions, a market equilibrium always exists [Debreu and Arrow 1954]. In particular, in the Arrow-Debreu model, agents sell the goods they own and use the revenue to buy other goods they want; this is in contrast to Fisher markets, where agents with a monetary budget purchase from a centralized seller. Market equilibria in both models exist, but can we find them? As in the case of Nash equilibrium, this question is of particular importance because if a centralized, omniscient algorithm cannot compute an equilibrium, it is hard to expect a distributed market with selfish agents to converge to one. In the words of Kamal Jain [Jain 2004, Nisan 2009b]:

If your laptop cannot find it, neither can the market.

In the same paper, Jain also gave an algorithm for computing Arrow-Debreu’s equilibrium when consumers’ utilities are linear. Since then, there has been a long sequence of algorithms for computing or approximating market equilibria (e.g., [Birnbbaum et al. 2011, Codenotti et al. 2005, Cole and Fleischer 2008, Devanur et al. 2008, Garg et al. 2015, Garg and Kapoor 2006, Jain 2004, Jain and Vazirani, 2010]), for certain models and utility functions, as well as intractability results in other settings [Chen et al. 2009a, Deng and Du 2008, Garg et al. 2017, Vazirani and Yannakakis 2011]. In particular, Chen et al. [2013] consider a setting of markets that exhibit *non-monotonicity*: when the price of one product increases, its demand increases. Giffen goods, described above, are one example of non-monotone markets; Chen et al. [2013] construct examples in Arrow-Debreu markets when the increased price of a product increases the revenue of the seller, who now has more available money and demands more of the same good. Chen et al. show that for essentially any class of utility function that exhibits some non-monotonicity, computing the Arrow-Debreu market equilibrium is PPAD-hard.

We extend the PPAD-hardness proof of Chen et al. [2013] and show that even approximate market equilibrium is PPAD-hard (Theorem 9.1). We note that although

our inapproximability factor is stronger than that showed by Chen et al., the results are incomparable as ours only holds for the stronger notion of “tight” approximate equilibrium, by which we mean the more standard definition that bounds the two-sided error of the market equilibrium. Chen et al., in contrast, prove that even if we allow arbitrary excess supply, finding a $(1/n)$ -approximate equilibrium is PPAD-hard. Furthermore, for the interesting case of constant elasticity of substitution (CES) utilities with parameter $\rho < 0$, they show that there exist markets where every $(1/2)$ -tight equilibrium requires prices that are doubly exponentially large (and thus require an exponential-size representation). Indeed, for a general non-monotone family of utility functions, the problem of computing a (tight or not) approximate equilibrium may not belong to PPAD. Nevertheless, the important family of additively separable, concave piecewise-linear utilities is known to satisfy the non-monotone condition [Chen et al. 2013], and yet the computation of (exact) market equilibrium is in PPAD [Vazirani and Yannakakis 2011]. Therefore, we obtain as a corollary that computing an ε -tight approximate equilibrium for Arrow-Debreu markets with additively separable, concave piecewise-linear utilities is PPAD-complete.

1.1.3 A-CEEI (CourseMatch)

University courses have limited capacity, and some are more popular than others. This creates an interesting allocation problem. Imagine that each student has ordered all the possible schedules—bundles of courses—from most desirable to least desirable, and the capacities of the classes are known. What is the best way to allocate seats in courses to students? There are several desiderata for a course allocation mechanism:

Fairness. In what sense is the mechanism “fair”?

Efficiency. Are seats in courses allocated to the students who want them the most?

Feasibility. Are any courses oversubscribed?

Truthfulness. Are students motivated to honestly report their preferences to the mechanism?

Computational efficiency. Can the allocation be computed from the data in polynomial time?

Competitive Equilibrium from Equal Incomes (CEEI) [Foley 1967, Thomson and Varian 1985, Varian 1974] is a venerable mechanism with many attractive prop-

erties: In CEEI all agents are allocated the same amount of “funny money”; next they declare their preferences, and then a price equilibrium is found that clears the market. The market clearing guarantees Pareto efficiency and feasibility. The mechanism has a strong, albeit technical, *ex post* fairness guarantee that emerges from the notion that agents who miss out on a valuable, competitive item will have extra funny money to spend on other items at equilibrium. Truthfulness is problematic—as usual with market mechanisms—but potential incentives for any individual agent to deviate are mitigated by the large number of agents. However, CEEI only works when the resources to be allocated are divisible and the utilities are relatively benign. This restriction has both benefits and drawbacks. It ensures computational feasibility, because CEEI can be computed in polynomial time with a linear or convex program, depending on the utilities involved [Devanur et al. 2008, Ghodsi et al. 2011, Varian 1974]; on the other hand, it is easy to construct examples in which a CEEI does not exist when preferences are complex or the resources being allocated are not divisible. Indeed, both issues arise in practice in a variety of allocation problems, including shifts to workers, landing slots to airplanes, and the setting that we focus on here, courses to students [Budish 2011, Varian 1974].

It was shown in Budish [2011] that an approximation to a CEEI solution, called A-CEEI, exists even when the resources are indivisible and agent preferences are arbitrarily complex, as required by the course allocation problems one sees in practice. The approximate solution guaranteed to exist is approximately fair (in that the students are given almost the same budget), and approximately Pareto efficient and feasible (in that all courses are filled close to capacity, with the possible exception of courses with more capacity than popularity). This result seems to be wonderful news for the course allocation problem. However, there is a catch: Budish’s proof is non-constructive, as it relies on Kakutani’s fixed point theorem.

A heuristic search algorithm for solving A-CEEI was introduced in Othman et al. [2010]. The algorithm resembles a traditional tâtonnement process, in which the prices of courses that are oversubscribed are increased and the prices of courses that are undersubscribed are decreased. A modified version of this algorithm that guarantees courses are not oversubscribed is currently used by the Wharton School (University of Pennsylvania) to assign their MBA students to courses [Budish et al. 2014]. While it has been documented that the heuristic algorithm often produces much tighter approximations than the theoretical bound, on some instances it fails to find even the guaranteed approximation [Budish 2011 (Section 9)].

Thus A-CEEI is a problem where practical interest motivates theoretical inquiry. We have a theorem that guarantees the existence of an approximate equilibrium—the issue is finding it. Can the heuristic algorithms currently used to assign

Wharton MBAs to their courses be replaced by a fast and rigorous algorithm for finding an approximate CEEI? Theorem 10.3 answers this question in the negative, showing that computing an A-CEEI is PPAD-complete.

1.2 Quasi-Polynomial Time and the Birthday Paradox

The following bilinear optimization meta-problem captures a wide range of applications, from areas like statistics Sparse Principle Component Analysis (Sparse PCA), graph theory (Clique), and game theory (Nash equilibrium):

$$\max_{(x,y) \in \mathcal{X}} x^\top Ay. \quad (1.1)$$

For all the applications we consider, once we fix some y^* , finding the best feasible x that maximizes $x^\top Ay^*$ is a tractable problem. (Similarly, if we were given a good x^* , finding a matching y is easy.) But optimizing x and y simultaneously is NP-hard. What about approximations?

Caratheodory's theorem states that a point v in the convex hull of n points in \mathbb{R}^d can be written as a convex combination of $d + 1$ points. In general, $d + 1$ points are necessary, but this number can be reduced drastically if we are willing to settle for approximation. In particular, Barman⁷ [2015] proves an approximate Caratheodory's theorem that requires only $r = O(p/\varepsilon^2)$ points to express a point \hat{v} such that $\|\hat{v} - v\|_p < \varepsilon$, assuming the n points belong to a unit ℓ_p -ball. In particular, \hat{v} can be written as an average over a multi-set of r out of the n points.

Viewing the columns of A as n vectors in \mathbb{R}^n , Barman observes that (after proper normalization) the point $v = Ay^*$ is in their convex hull. If we only want to approximately solve the bilinear optimization (1.1), we drastically reduce the dimension of the problem by enumerating over all choices of \hat{v} , and for each \hat{v} solving the optimization problem $A\hat{v}$. It turns out that in order to guarantee a good additive approximation of (1.1), we need to set $p \approx \log n$. The dominant term in the running time of this meta-algorithm is the enumeration over all choices of \hat{v} (all multi-sets of the columns of A), which takes approximately $n^r = n^{O(\frac{\log n}{\varepsilon^2})}$, i.e., quasi-polynomial time.

The above meta-algorithm provides evidence that the approximate variant of all those problems is much easier than solving them exactly: in particular we believe that NP-hard (respectively, PPAD-hard) problems like 3-SAT (resp. END-OF-A-LINE)

7. Both the approximate Caratheodory's theorem and the resulting algorithm have been described by many researchers (e.g., [Arora et al. 2012, Lipton et al. 2003]); however, the presentation in Barman [2015] is our favorite.

require approximately 2^n time. This belief is formulated by the *Exponential Time Hypothesis*, or ETH [Impagliazzo et al. 2001] (resp. ETH for PPAD [Babichenko et al. 2016]). The reasons we believe in ETH are similar to the ones outlined in the previous section for our belief that END-OF-A-LINE is intractable: on one hand, we have a combination of unconditional lower bounds in analogous models and little or no progress on algorithms; on the other hand, we are “forced” to believe it because we have no idea how to prove it (in particular, ETH implies the much weaker $P \neq NP$).

However, quasi-polynomial time is still both unrealistic to implement in practice, and does not meet our gold standard of polynomial time in theory. The main question we address in this section is:

Can the quasi-polynomial time be improved to polynomial time?

For Sparse PCA this is indeed possible [Alon et al. 2013, Asteris et al. 2015, Chan et al. 2016]. But for several other problems we can prove that, assuming ETH, quasi-polynomial time is actually necessary:

- Finding a k -subgraph that is almost a clique; this is one of the most fundamental approximation problems in theoretical computer science.
- Finding and counting stable communities in a friendship graph; this problem received a lot of attention in recent years with the emergence of online social networks.
- Finding an approximately optimal signaling scheme; this resolves a recent open question by Dughmi.
- Computing the VC and Littlestone’s Dimensions of a binary concept class, two of the most fundamental quantities in learning theory.

A common approach to all our proofs is the *birthday repetition* framework due to Aaronson et al. [2014]: construct a reduction from 3-SAT to any of the above problems, with reduction size $N \approx 2^{\sqrt{n}}$. Then, assuming ETH, one needs approximately $N^{\log N} \approx 2^n$ time to solve the problem on the larger instance. A key step in the reduction is to consider subsets of \sqrt{n} variables; then by the *birthday paradox* any two subsets are likely to share at least one variable (hence the name “birthday repetition”).

1.2.1 Densest k -Subgraph

k -CLIQUE is one of the most fundamental problems in computer science: given a graph, decide whether it has a fully connected induced subgraph on k vertices. Since

it was proven NP-complete by Karp [1972], extensive research has investigated the complexity of its relaxations.

We consider two natural relaxations of k -CLIQUE that have received significant attention from both algorithmic and complexity communities: The first one is to relax the “ k ” requirement, i.e., looking for a smaller subgraph: Given an n -vertex graph G containing a clique of size k , find a clique of size at least δk for some parameter $0 < \delta < 1$.

The second natural relaxation is to relax the “Clique” requirement, replacing it with the more modest goal of finding a subgraph that is almost a clique: Given an n -vertex graph G containing a clique of size k , find an induced subgraph of G of size k with (edge) density at least $1 - \varepsilon$, for some parameter $0 < \varepsilon < 1$.

The first relaxation has been a motivating example throughout a long line of research that laid the foundations for NP-hardness of approximation [Arora and Safra 1998, Arora et al. 1998, Feige et al. 1996, Håstad 1999, Khot 2001, Zuckerman 2007]. In particular, we now know that it is NP-hard to distinguish between a graph that has a clique of size k , and a graph whose largest induced clique is of size at most $k' = \delta k$, where $\delta = 1/n^{1-\varepsilon}$ [Zuckerman 2007]. Until our work, the computational complexity of the second relaxation remained largely open. There are a couple of (very different) quasi-polynomial algorithms that guarantee finding a $(1 - \varepsilon)$ -dense k -subgraph in every graph containing a k -clique: the meta-algorithm by Barman, which we outlined above, and an older algorithm due to Feige and Seltser [1997], but nothing non-trivial was known about hardness of approximation.

In Chapter 12 we prove that, assuming ETH, even if one makes *both relaxations* the problem remains intractable. In particular, even if the graph contains a clique of size k , it takes quasi-polynomial time to find a $(1 - \varepsilon)$ -dense δk -subgraph, for constant $\varepsilon > 0$ and $\delta = o(1)$.

1.2.2 Community Detection

Identifying communities is a central graph-theoretic problem with important applications to sociology and marketing (when applied to social networks), biology and bioinformatics (when applied to protein interaction networks), and more (see, e.g., Fortunato’s classic survey [Fortunato 2010]). Defining what exactly is a *community* remains an interesting problem on its own (see Arora et al. [2012] and Borgs et al. [2016] for excellent treatment from a theoretical perspective). Ultimately, there is no single “right” definition, and the precise meaning of community should be different for social networks and protein interaction networks.

In Chapter 13 we focus on the algorithmic questions arising from one of the simplest and most canonical definitions, which has been considered by several

theoretical computer scientists [Arora et al. 2012, Balcan et al. 2013, Braverman et al. 2017, Mishra et al. 2008].

Definition 1.2 (α, β) -community. Given an undirected graph $G = (V, E)$, an (α, β) -community is a subset $S \subseteq V$ that satisfies:

Strong ties inside the community. For every $v \in S$, $|\{v\} \times S \cap E| \geq \alpha \cdot |S|$; and

Weak ties to nodes outside the community. For every $u \notin S$, $|\{u\} \times S \cap E| \leq \beta \cdot |S|$.

This problem has been considered by several researchers before: Mishra et al. [2008] gave a polynomial-time algorithm for finding (α, β) -communities that contain a vertex with very few neighbors outside the community. Balcan et al. [2013] give a polynomial-time algorithm for enumerating (α, β) -communities in the special case where the degree of every node is $\Omega(n)$. Arora et al. [2012] consider several semi-random models where the edges inside the community are generated at random, according to the expected degree model. For the general case, the latter paper by Arora et al. gave a simple quasi-polynomial ($n^{O(\log n)}$) time for detecting (α, β) -communities whenever $\alpha - \beta$ is at least some positive constant. (Their algorithm is essentially identical to the meta-algorithm for bilinear optimization that we outlined above.)

We show that, for every constants $\alpha > \beta \in (0, 1]$, community detection requires quasi-polynomial time (assuming ETH). For example, when $\alpha = 1$ and $\beta = 0.01$, this means that we can hide a clique C , such that every single vertex not in C is connected to at most 1% of C . Our main result is actually a much stronger inapproximability: even in the presence of a $(1, o(1))$ -community, finding any $(\beta + o(1), \beta)$ -community is hard.

Unlike all quasi-polynomial approximation schemes mentioned above, Arora et al.'s algorithm has the unique property that it can also *exactly count* all the (α, β) -communities. Our second result is that counting even the number of $(1, o(1))$ -communities requires quasi-polynomial time. A nice feature of this result is that we can base it on the much weaker #ETH assumption, which asserts that counting the satisfying assignment for a 3-SAT instance requires time $2^{\Omega(n)}$. (Note, for example, that #ETH is likely to be true even if $P = NP$.)

1.2.3 VC and Littlestone's Dimensions

A common and essential assumption in learning theory is that the concepts we want to learn come from a nice, simple concept class, or (in the agnostic case) that they can at least be approximated by a concept from a simple class. When the

concept class is sufficiently simple, there is hope for good (i.e., sample-efficient and low-error) learning algorithms.

There are many different ways to measure the *simplicity* of a concept class. The most influential measure of simplicity is the VC Dimension [Vapnik and Chervonenkis 1971], which captures learning in the Probably Almost Correct (PAC) model. We also consider Littlestone’s Dimension [Littlestone 1987], which corresponds to minimizing mistakes in online learning (see Section 2.5 for definitions). When either dimension is small, there are algorithms that exploit the simplicity of the class, to obtain good learning guarantees.

In Chapter 14 we consider the algorithmic challenge of computing either dimension. In particular, we study the most optimistic setting, where the entire universe and concept class are given as explicit input (a binary matrix whose (x, c) -th entry is 1 iff element x belongs to concept c). In this model, both dimensions can be computed in quasi-polynomial time. Interestingly, the algorithm does not go through the bilinear optimization problem; instead, it exploits the fact that for concept class \mathcal{C} , both dimensions are bounded by $\log |\mathcal{C}|$. Two decades ago, it was shown that quasi-polynomial time is indeed necessary for both dimensions [Frances and Litman 1998, Papadimitriou and Yannakakis 1996]. The computational intractability of computing the (VC, Littlestone’s) dimension of a concept class suggests that even in cases where a simple structure exists, it may be inaccessible to computationally bounded algorithms.

Theorems 14.1 and 14.2 extend the results of Frances and Litman [1998] and Papadimitriou and Yannakakis [1986] to show that the VC and Littlestone’s Dimensions cannot even be *approximately* computed in polynomial time.

1.2.4 Signaling

Many classical questions in economics involve extracting information from strategic agents. Lately, there has been growing interest within algorithmic game theory in *signaling*: the study of how to *reveal information* to strategic agents (see, e.g., [Cheng et al. 2015, Dughmi 2014, Dughmi et al. 2013, Emek et al. 2014, Miltersen and Sheffet 2012] and references therein). Signaling has been studied in many interesting economic and game theoretic settings. Among them, ZERO-SUM SIGNALING proposed by Dughmi [2014] stands out as a canonical problem that cleanly captures the computational nature of signaling. In particular, focusing on zero-sum games clears away issues of equilibrium selection and computational tractability of finding an equilibrium.

Definition 1.3 ZERO-SUM SIGNALING [Dughmi 2014]. Alice and Bob play a Bayesian zero-sum game where the payoff matrix M is drawn from a publicly known prior. The signaler

Sam privately observes the state of nature (i.e., the payoff matrix), and then publicly broadcasts a signal $\varphi(M)$ to both Alice and Bob. Alice and Bob Bayesian-update their priors according to $\varphi(M)$'s and play the Nash equilibrium of the expected game; but they receive payoffs according to the true M . Sam's goal is to design an efficient signaling scheme φ (a function from payoff matrices to strings) that maximizes Alice's expected payoff.

Dughmi's main result proves that assuming the hardness of the PLANTED CLIQUE problem, there is no additive Fully Polynomial Time Approximation Scheme (FPTAS) for ZERO-SUM SIGNALING. The main open question left by Dughmi [2014] is whether an additive PTAS exists. Here we answer this question in the negative: we prove that assuming the ETH [Impagliazzo et al. 2001], obtaining an additive- ϵ -approximation (for some constant $\epsilon > 0$) requires quasi-polynomial time ($n^{\tilde{\Omega}(\log n)}$). This result is tight thanks to a recent quasi-polynomial ($n^{\frac{\log n}{\text{poly}(\epsilon)}}$) time algorithm by Cheng et al. [2015]. Another important advantage of our result is that it replaces the hardness of PLANTED CLIQUE with a more believable worst-case hardness assumption (see, e.g., the discussion in Braverman et al. [2015]).

1.3 Approximate Nash Equilibrium

The main result in this book rules out the PTAS (polynomial time approximation schemes) for two-player Nash equilibrium. Consider a game between two players, each choosing between a large number (n) of actions. The inputs to the algorithm are two $n \times n$ matrices with entries in $[0, 1]$. The goal is to find, for every constant ϵ , an ϵ -approximate Nash equilibrium; i.e., a mixed strategy for each player, such that either player can gain at most ϵ by deviating to a different strategy.

This has been the central open question in equilibrium computation for the past decade. There were good reasons to be hopeful: there was a quasi-polynomial time [Lipton et al. 2003], a series of improved approximation ratios [Bosse et al. 2010, Daskalakis et al. 2007, 2009b, Kontogiannis et al. 2009, Tsaknakis and Spirakis 2008], and several approximation schemes for special cases [Alon et al. 2013, Barman 2015, Daskalakis and Papadimitriou 2009, Kannan and Theobald 2007]. Our main result settles this question in the negative:

Theorem 1.1 Main theorem. There exists a constant $\epsilon > 0$ such that, assuming ETH for PPAD, finding an ϵ -approximate Nash equilibrium in a two-player $n \times n$ game requires time $T(n) = n^{\log^{1-o(1)} n}$.

We supplement Theorem 1.1 with a series of other hardness of approximation results for Nash equilibrium in related settings, further establishing the point that even approximate Nash equilibria are intractable. First, we consider different

sources of complexity. Our main result shows intractability of approximate Nash equilibria when the complexity of the game arises from each player choosing among many actions. In Theorems 5.1 and 5.2, we prove that in games where each player only has two actions, the complexity can arise from a large number of players; finding an approximate Nash equilibrium in n -player, binary action games is PPAD-hard (settling another open question from Daskalakis's thesis [Daskalakis 2008]). Alternatively, even if there are only two players and the number of actions is constant, a different source of complexity can be the players' uncertainty; finding an approximate Bayesian Nash equilibrium in such incomplete information games is also PPAD-hard (Corollary 8.1).

We also prove intractability in different models: query complexity, communication complexity, and uncoupled dynamics (settling a long list of open questions from Babichenko [2012, 2016], Chen et al. [2017], Fearnley et al. [2013], Hart and Mansour [2010], Hart and Nisan [2013], Nisan [2009a], Roughgarden and Weinstein [2016]). The main advantage of these results is that they are *unconditional*, i.e., they do not rely on complexity assumptions such as ETH for PPAD, or $P \neq NP$. In particular, in the setting where each player knows her own utility function, even computationally unbounded players have to communicate almost all their private information in order to reach even an approximate equilibrium.

Preliminaries

Notation

We use $\mathbf{0}_n$ (respectively $\mathbf{1}_n$) to denote the length- n vectors whose value is 0 (1) in every coordinate. For vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, we let

$$\|\mathbf{x} - \mathbf{y}\|_2 \triangleq \sqrt{\mathbb{E}_{i \in [n]} (x_i - y_i)^2}$$

denote the *normalized 2-norm*. Unless specified otherwise, when we say that \mathbf{x} and \mathbf{y} are Δ -close (or Δ -far), we mean Δ -close in normalized 2-norm. Similarly, for a binary string $\pi \in \{0, 1\}^n$, we denote

$$|\pi| \triangleq \mathbb{E}_{i \in [n]} [\pi_i].$$

For a graph $G = (V, E)$ and $S \subseteq V$, we use $\text{den}(S) \in [0, 1]$ to denote the density of subgraph S ,

$$\text{den}(S) := \frac{|(S \times S) \cap E|}{|S \times S|}.$$

2.1 Nash Equilibrium and Relaxations

A mixed strategy of player i is a distribution x_i over i 's set of actions, A_i . We say that a vector of mixed strategies $\mathbf{x} \in \times_j \Delta A_j$ is a *Nash equilibrium* if every strategy a_i in the support of every x_i is a best response to the actions of the mixed strategies of the rest of the players, x_{-i} . Formally, for every $a_i \in \text{supp}(x_i)$,

$$\mathbb{E}_{a_{-i} \sim x_{-i}} [u_i(a_i, a_{-i})] = \max_{a' \in A_i} \mathbb{E}_{a_{-i} \sim x_{-i}} [u_i(a', a_{-i})].$$

Equivalently, \mathbf{x} is a Nash equilibrium if each mixed strategy x_i is a best response to x_{-i} :

$$\mathbb{E}_{\mathbf{a} \sim \mathbf{x}} [u_i(\mathbf{a})] = \max_{x'_i \in \Delta A_i} \mathbb{E}_{\mathbf{a} \sim (x'_i; x_{-i})} [u_i(\mathbf{a})].$$

Each of those equivalent definitions can be generalized to include approximation in a different way.

Definition 2.1 ϵ -approximate Nash Equilibrium. We say that \mathbf{x} is an ϵ -Approximate Nash Equilibrium (ϵ -ANE) if each x_i is an ϵ -best response to x_{-i} :

$$\mathbb{E}_{\mathbf{a} \sim \mathbf{x}}[u_i(\mathbf{a})] \geq \max_{x'_i \in \Delta A_i} \mathbb{E}_{\mathbf{a} \sim (x'_i; x_{-i})}[u_i(\mathbf{a})] - \epsilon.$$

On the other hand, we generalize the first definition of Nash equilibrium in the following stricter definition:

Definition 2.2 ϵ -well-supported Nash Equilibrium. \mathbf{x} is an ϵ -Well-Supported Nash Equilibrium (ϵ -WSNE) if every a_i in the support of x_i is an ϵ -best response to x_{-i} : for every $a_i \in \text{supp}(x_i)$,

$$\mathbb{E}_{a_{-i} \sim x_{-i}}[u_i(a_i, a_{-i})] \geq \max_{a' \in A_i} \mathbb{E}_{a_{-i} \sim x_{-i}}[u_i(a', a_{-i})] - \epsilon.$$

WeakNash

We can further relax the (already more lenient) notion of ϵ -ANE by requiring that the ϵ -best response condition only hold for most of the players (rather than all of them).

Definition 2.3 (ϵ, δ) -WeakNash [Babichenko et al. 2016]. We say that \mathbf{x} is an (ϵ, δ) -WeakNash if for a $(1 - \delta)$ -fraction of i 's, x_i is an ϵ -best mixed response to x_{-i} :

$$\Pr_i \left[\mathbb{E}_{\mathbf{a} \sim \mathbf{x}}[u_i(\mathbf{a})] \geq \max_{x'_i \in \Delta A_i} \mathbb{E}_{\mathbf{a} \sim (x'_i; x_{-i})}[u_i(\mathbf{a})] - \epsilon \right] \geq 1 - \delta.$$

Definition 2.4 (ϵ, δ) -well-supported WeakNash. \mathbf{x} is an (ϵ, δ) -well-supported WeakNash if for a $(1 - \delta)$ -fraction of i 's, every a_i in the support of x_i is an ϵ -best response to x_{-i} :

$$\Pr_i \left[\forall a_i \in \text{supp}(x_i) \mathbb{E}_{a_{-i} \sim x_{-i}}[u_i(a_i, a_{-i})] \geq \max_{a' \in A_i} \mathbb{E}_{a_{-i} \sim x_{-i}}[u_i(a', a_{-i})] - \epsilon \right] \geq 1 - \delta.$$

2.2 PPAD and END-OF-A-LINE

The END-OF-A-LINE of a problem considers an implicitly represented, exponential-size, directed graph whose vertices have in- and out-degree at most 1 (this is without loss of generality). The special vertex $\mathbf{0}_n$ has in-degree 0, and the goal is to find another odd-degree vertex. The graph is a union of lines and cycles, so in particular the line starting at $\mathbf{0}_n$ ends with another odd-degree vertex.

The graph is implicitly defined with functions S, P that, for each vertex, give its Successor (outgoing neighbor) and Predecessor (incoming neighbor). In the computational variant of END-OF-A-LINE, S, P are given as explicit circuits, whereas in the query complexity variant they are given as oracles. There is also a communication complexity variant, whose definition is more involved and is deferred to Section 3.4.

In the formal definition of the computational variant, we also have to consider the case that S, P are inconsistent, i.e., for some $u \neq v$ we have $S(u) = v$, but $P(v) \neq u$; we also allow the algorithm to succeed by finding such an inconsistency. (In the oracle model we can explicitly require that there are no inconsistencies.)

Definition 2.5 END-OF-A-LINE. The input to the problem is functions $S, P : \{0, 1\}^n \rightarrow \{0, 1\}^n$, such that $S(\mathbf{0}_n) = \mathbf{0}_n \neq P(\mathbf{0}_n)$. The output is another odd-degree vertex $\mathbf{0}_n \neq v \in \{0, 1\}^n$ such that $P(S(v)) \neq S(P(v))$.

The computational complexity class PPAD is defined as the class of all total search problems reducible to END-OF-A-LINE.

MEMBERSHIP END-OF-A-LINE

The following variant of END-OF-A-LINE is equivalent and more convenient for some of our reductions. In particular, the problem is restricted to a subset of the vertices. The restricted vertex-set is defined implicitly via a membership function $T : \{0, 1\}^n \rightarrow \{0, 1\}$. Abusing notation, let T also denote the restricted set of vertices whose T -value is 1. We think of S and P as only applying to vertices in T , and connecting them to other vertices in T . Formally, we also allow the algorithm to return any violations.

Definition 2.6 MEMBERSHIP END-OF-A-LINE. The input to the problem is functions $S, P : \{0, 1\}^n \rightarrow \{0, 1\}^n$ and $T : \{0, 1\}^n \rightarrow \{0, 1\}$, such that $S(\mathbf{0}_n) = \mathbf{0}_n \neq P(\mathbf{0}_n)$ and $T(\mathbf{0}_n), T(S(\mathbf{0}_n)) = 1$. The output is another odd-degree vertex $\mathbf{0}_n \neq v \in \{0, 1\}^n$ such that $T(v) = 1$ and v satisfies at least one of the following:

End-of-a-line. $P(S(v)) \neq S(P(v))$; or

Boundary condition. $T(S(v)) = 0$ or $T(P(v)) = 0$.

Lemma 2.1 END-OF-A-LINE is equivalent to MEMBERSHIP END-OF-A-LINE.

Proof Given an instance of END-OF-A-LINE, we can construct an equivalent instance of MEMBERSHIP END-OF-A-LINE by setting $T \equiv 1$. In the other direction, we can add self-loops to every vertex v such that $T(v) = 0$ (i.e., $P(v) = S(v) = v$); this guarantees that v is never a solution to the new END-OF-A-LINE instance. ■

We will be interested in restricted (but equally hard) variants of MEMBERSHIP END-OF-A-LINE. For example, in Section 16.2 we define LOCAL END-OF-A-LINE where, among other restrictions, T, S, P are AC^0 circuits. In particular, in Chapter 3, we will consider a variant where each vertex is a priori restricted to have at most two potential incoming/outgoing neighbors, and the functions S, P merely specify which neighbor is chosen. We then abuse notation and let S, P output just a single bit.

2.3 Exponential Time Hypotheses

Our quasi-polynomial hardness results are conditional on the following hypotheses. We begin with the “plain” ETH:

Hypothesis 2.1 Exponential time hypothesis (ETH) [Impagliazzo et al. 2001]. 3SAT takes time $2^{\Omega(n)}$.

Since a Nash equilibrium always exists, we are unlikely to have a reduction (even of subexponential size) from 3SAT to Nash equilibrium. Instead, we need to assume the following analogue of ETH for the total class PPAD:

Hypothesis 2.2 ETH for PPAD [Babichenko et al. 2016]. Solving END-OF-A-LINE requires time $2^{\tilde{\Omega}(n)}$.¹

In Section 13.3 we will prove a quasi-polynomial lower bound on the running time for counting the number of communities in a social network. This result is also conditional, but requires the following much weaker #ETH assumption:

Hypothesis 2.3 #ETH [Dell et al. 2014]. Given a 3SAT formula, counting the number of satisfying assignments takes time $2^{\Omega(n)}$.

2.4 PCP Theorems

2.4.1 2CSP and the PCP Theorem

In the 2CSP problem (Constraint Satisfaction Problem), we are given a graph $G = (V, E)$ on $|V| = n$ vertices, where each of the edges $(u, v) \in E$ is associated with some constraint function $\psi_{u,v} : \Sigma \times \Sigma \rightarrow \{0, 1\}$ that specifies a set of legal “colorings” of u and v , from some finite alphabet Σ (2 in the term “2CSP” stands for the “arity” of

1. As usual, n is the size of the description of the instance, i.e., the size of the circuits S and P .

each constraint, which always involves two variables). Let us denote by ψ the entire **2CSP** instance, and define by $\text{OPT}(\psi)$ the optimum (maximum) fraction of satisfied constraints in the associated graph G , over all possible assignments (colorings) of V .

The starting point of some of our reductions is the following version of the Probabilistically Checkable Proof (PCP) theorem, which asserts that it is NP-hard to distinguish a **2CSP** instance whose value is 1, and one whose value is $1 - \eta$, where η is some small constant:

Theorem 2.1 PCP Theorem [Dinur 2007]. Given a **3SAT** instance φ of size n , there is a polynomial time reduction that produces a **2CSP** instance ψ , with size $|\psi| = n \cdot \text{polylog } n$ variables and constraints, and constant alphabet size such that

- (Completeness) If $\text{OPT}(\varphi) = 1$ then $\text{OPT}(\psi) = 1$.
- (Soundness) If $\text{OPT}(\varphi) < 1$ then $\text{OPT}(\psi) < 1 - \eta$, for some constant $\eta = \Omega(1)$.
- (Graph) The constraint graph is d -regular, for some constant d , and bipartite.

See, e.g., the full version of Braverman et al. [2017] or Aaronson et al. [2014] for derivations of this formulation of the PCP theorem.

Notice that since the size of the reduction is near linear, ETH implies that solving the above problem requires near exponential time.

Corollary 2.1 Let ψ be as in Theorem 2.1. Then assuming ETH, distinguishing between $\text{OPT}(\psi) = 1$ and $\text{OPT}(\psi) < 1 - \eta$ requires time $2^{\tilde{\Omega}(|\psi|)}$.

Label Cover

Definition 2.7 LABEL COVER. LABEL COVER is a maximization problem, and a special case of **2CSP**. The input is a bipartite graph $G = (A, B, E)$, alphabets Σ_A, Σ_B , and a projection $\pi_e : \Sigma_A \rightarrow \Sigma_B$ for every $e \in E$.

The output is a labeling $\varphi_A : A \rightarrow \Sigma_A, \varphi_B : B \rightarrow \Sigma_B$. Given a labeling, we say that a constraint (or edge) $(a, b) \in E$ is *satisfied* if $\pi_{(a,b)}(\varphi_A(a)) = \varphi_B(b)$. The *value of a labeling* is the fraction of $e \in E$ that are satisfied by the labeling. The value of the instance is the maximum fraction of constraints satisfied by any assignment.

We often encounter an assignment that only labels a subset of $A \cup B$ but leaves the rest unlabeled. We refer to such assignment as a *partial assignment* to an instance; more specifically, for any $V \subseteq A \cup B$, a V -partial assignment (or partial assignment on V) is a function $\phi : V \rightarrow \Sigma$. For notational convenience, we sometimes write Σ^V to denote the set of all functions from V to Σ .

Theorem 2.2 Moshkovitz-Raz PCP [Moshkovitz and Raz 2010 (Theorem 11)]. For every n and every $\epsilon > 0$ (in particular, ϵ may be a function of n), solving 3SAT on inputs of size n can be reduced to distinguishing between the case that a (d_A, d_B) -bi-regular instance of LABEL COVER, with parameters $|A| + |B| = n^{1+o(1)} \cdot \text{poly}(1/\epsilon)$, $|\Sigma_A| = 2^{\text{poly}(1/\epsilon)}$, and $d_A, d_B, |\Sigma_B| = \text{poly}(1/\epsilon)$, is completely satisfiable, versus the case that it has value at most ϵ .

Counting the number of satisfying assignments is even harder. The following hardness is well-known, and we sketch its proof only for completeness:

Fact 2.1 There is a linear-time reduction from #3SAT to counting the number of satisfying assignments of a LABEL COVER instance.

Proof Construct a vertex in A for each variable and a vertex in B for each clause. Set $\Sigma_A \triangleq \{0, 1\}$ and let $\Sigma_B \triangleq \{0, 1\}^3 \setminus (000)$ (i.e., Σ_B is the set of satisfying assignments for a 3SAT clause, after applying negations). Now if variable x appears in clause C , add a constraint that the assignments to x and C are consistent (taking into account the sign of x in C). Notice that: (i) any assignment to A corresponds to a unique assignment to the 3SAT formula; and (ii) if the 3SAT formula is satisfied, this assignment uniquely defines a satisfying assignment to B . Therefore there is a one-to-one correspondence between satisfying assignments to the 3SAT formula and to the instance of LABEL COVER. ■

2.5 Learning Theory

For a universe (or ground set) \mathcal{U} , a concept C is simply a subset of \mathcal{U} and a concept class \mathcal{C} is a collection of concepts. For convenience, we sometimes relax the definition and allow the concepts to not be subsets of \mathcal{U} ; all definitions here extend naturally to this case.

The VC and Littlestone's Dimensions can be defined as follows.

Definition 2.8 VC Dimension [Vapnik and Chervonenkis 1971]. A subset $S \subseteq \mathcal{U}$ is said to be *shattered* by a concept class \mathcal{C} if, for every $T \subseteq S$, there exists a concept $C \in \mathcal{C}$ such that $T = S \cap C$.

The VC Dimension $\text{VC-dim}(\mathcal{C}, \mathcal{U})$ of a concept class \mathcal{C} with respect to the universe \mathcal{U} is the largest d such that there exists a subset $S \subseteq \mathcal{U}$ of size d that is shattered by \mathcal{C} .

Definition 2.9 Mistake tree and Littlestone's Dimension [Littlestone 1987]. A depth- d instance-labeled tree of \mathcal{U} is a full binary tree of depth d such that every internal node of the tree is assigned an element of \mathcal{U} . For convenience, we will identify each node in the tree canonically by a binary string s of length at most d .

A depth- d mistake tree (aka shattered tree [Ben-David et al. 2009]) for a universe \mathcal{U} and a concept class \mathcal{C} is a depth- d instance-labeled tree of \mathcal{U} such that, if we let $v_s \in \mathcal{U}$ denote the element assigned to the vertex s for every $s \in \{0, 1\}^{$d$$, then, for every leaf $\ell \in \{0, 1\}^d$, there exists a concept $C \in \mathcal{C}$ that agrees with the path from root to it, i.e., that, for every $i < d$, $v_{\ell_{\leq i}} \in C$ iff $\ell_{i+1} = 1$ where $\ell_{\leq i}$ denotes the prefix of ℓ of length i .

The Littlestone's Dimension $L\text{-dim}(\mathcal{C}, \mathcal{U})$ of a concept class \mathcal{C} with respect to the universe \mathcal{U} is defined as the maximum d such that there exists a depth- d mistake tree for \mathcal{U}, \mathcal{C} .

An equivalent formulation of Littlestone's Dimension is through mistakes made in online learning, as stated below. This interpretation will be useful in the proof of Theorem 14.4.

Definition 2.10 Mistake bound. An online algorithm \mathcal{A} is an algorithm that, at time step i , is given an element $x_i \in \mathcal{U}$ and the algorithm outputs a prediction $p_i \in \{0, 1\}$ whether x is in the class. After the prediction, the algorithm is told the correct answer $h_i \in \{0, 1\}$. For a sequence $(x_1, h_1), \dots, (x_n, h_n)$, a *prediction mistake* of \mathcal{A} is defined as the number of incorrect predictions, i.e., $\sum_{i \in [n]} \mathbf{1}[p_i \neq h_i]$. The *mistake bound* of \mathcal{A} for a concept class \mathcal{C} is defined as the maximum prediction mistake of \mathcal{A} over all the sequences $(x_1, h_1), \dots, (x_n, h_n)$, which corresponds to a concept $C \in \mathcal{C}$ (i.e., $h_i = \mathbf{1}[x_i \in C]$ for all $i \in [n]$).

Theorem 2.3 Littlestone [1987]. For any universe \mathcal{U} and any concept class \mathcal{C} , $L\text{-dim}(\mathcal{C}, \mathcal{U})$ is equal to the minimum mistake bound of \mathcal{C}, \mathcal{U} over all online algorithms.

The following facts are well-known and follow easily from the above definitions.

Fact 2.2 For any universe \mathcal{U} and concept class \mathcal{C} , we have

$$\text{VC-dim}(\mathcal{C}, \mathcal{U}) \leq L\text{-dim}(\mathcal{C}, \mathcal{U}) \leq \log |\mathcal{C}|.$$

Fact 2.3 For any two universes $\mathcal{U}_1, \mathcal{U}_2$ and any concept class \mathcal{C} ,

$$L\text{-dim}(\mathcal{C}, \mathcal{U}_1 \cup \mathcal{U}_2) \leq L\text{-dim}(\mathcal{C}, \mathcal{U}_1) + L\text{-dim}(\mathcal{C}, \mathcal{U}_2).$$

2.6 Information Theory

In this section, we introduce information-theoretic quantities used mostly in Chapter 12. For a more thorough introduction, the reader should refer to Cover and Thomas [2012]. Unless stated otherwise, all log's in this book paper are base-2.

Definition 2.11 Let μ be a probability distribution on sample space Ω . The *Shannon entropy* (or just *entropy*) of μ , denoted by $H(\mu)$, is defined as $H(\mu) := \sum_{\omega \in \Omega} \mu(\omega) \log \frac{1}{\mu(\omega)}$.

Definition 2.12 Binary entropy function. For $p \in [0, 1]$, the binary entropy function is defined as follows (with a slight abuse of notation): $H(p) := -p \log p - (1-p) \log(1-p)$.

Fact 2.4 Concavity of binary entropy. Let μ be a distribution on $[0, 1]$, and let $p \sim \mu$. Then $H(\mathbb{E}_\mu[p]) \geq \mathbb{E}_\mu[H(p)]$.

For a random variable A , we shall write $H(A)$ to denote the entropy of the induced distribution on the support of A . We use the same abuse of notation for other information-theoretic quantities appearing later in this section.

Definition 2.13 The *conditional entropy* of a random variable A conditioned on B is defined as

$$H(A|B) = \mathbb{E}_b(H(A|B=b)).$$

Fact 2.5 Chain rule.

$$H(AB) = H(A) + H(B|A).$$

Fact 2.6 Conditioning decreases entropy. $H(A|B) \geq H(A|BC)$.

Another measure we will use (briefly) in our proof is that of *mutual information*, which informally captures the correlation between two random variables.

Definition 2.14 Conditional mutual information. The *mutual information* between two random variables A and B , denoted by $I(A; B)$, is defined as

$$I(A; B) := H(A) - H(A|B) = H(B) - H(B|A).$$

The *conditional mutual information* between A and B given C , denoted by $I(A; B|C)$, is defined as

$$I(A; B|C) := H(A|C) - H(A|BC) = H(B|C) - H(B|AC).$$

The following is a well-known fact on mutual information.

Fact 2.7 Data processing inequality. Suppose we have the following Markov Chain:

$$X \rightarrow Y \rightarrow Z$$

where $X \perp Z | Y$. Then $I(X; Y) \geq I(X; Z)$ or, equivalently, $H(X|Y) \leq H(X|Z)$.

Mutual information is related to the Kullback Leiber Divergence similarity measure (also known as relative entropy).

Definition 2.15 Kullback-Leiber Divergence. Given two probability distributions μ_1 and μ_2 on the same sample space Ω such that $(\forall \omega \in \Omega)(\mu_2(\omega) = 0 \Rightarrow \mu_1(\omega) = 0)$, the *Kullback-Leibler Divergence* between them is defined as

$$D_{\text{KL}}(\mu_1 \parallel \mu_2) = \sum_{\omega \in \Omega} \mu_1(\omega) \log \frac{\mu_1(\omega)}{\mu_2(\omega)}.$$

The connection between the mutual information and the Kullback-Leibler divergence is provided by the following fact.

Fact 2.8 For random variables A , B , and C we have

$$I(A; B|C) = \mathbb{E}_{b,c} [D_{\text{KL}}(A_{bc} \parallel A_c)].$$

2.7 Useful Lemmata

2.7.1 Concentration

Lemma 2.2 Chernoff bound. Let X_1, \dots, X_n be independent and identically distributed (i.i.d.) random variables taking value from $\{0, 1\}$ and let p be the probability that $X_i = 1$; then, for any $\delta > 0$, we have

$$\Pr \left[\sum_{i=1}^n X_i \geq (1 + \delta)np \right] \leq \begin{cases} 2^{-\delta^2 np/3} & \text{if } \delta < 1, \\ 2^{-\delta np/3} & \text{otherwise.} \end{cases}$$

2.7.2 Pseudorandomness

Theorem 2.4 k -wise independence Chernoff bound [Schmidt et al. 1995 (Theorem 5.I)]. Let $x_1 \dots x_n \in [0, 1]$ be k -wise independent random variables, and let $\mu \triangleq \mathbb{E}[\sum_{i=1}^n x_i]$ and $\delta \leq 1$. Then

$$\Pr \left[\left| \sum_{i=1}^n x_i - \mu \right| > \delta \mu \right] \leq e^{-\Omega(\min\{k, \delta^2 \mu\})}.$$

2.7.3 λ -Biased Sets

Definition 2.16 λ -biased sets. Let \mathcal{G} be a finite field, and $t > 0$ an integer. A multiset $S \subseteq \mathcal{G}^t$ is λ -biased if for every non-trivial character χ of \mathcal{G}^t ,

$$\left| \mathbb{E}_{\mathbf{y} \sim S} [\chi(\mathbf{y})] \right| \leq \lambda.$$

Lemma 2.3 [Azar et al. 1998 (Theorem 3.2)]. A randomly chosen multi-set $S \subseteq \mathcal{G}^t$ of cardinality $\Theta(t \log |\mathcal{G}|/\lambda^2)$ is λ -biased with high probability.

For many applications, an explicit construction is necessary. In our case, however, we can enumerate over all sets S of sufficient cardinality in quasi-polynomial time.² The following Sampling Lemma due to Ben-Sasson et al. [2003] allows us to estimate the average of any function over \mathcal{G}^t using only one line and $(1 + o(1)) \log_2 |\mathcal{G}^t|$ randomness:

Lemma 2.4 Sampling Lemma [Ben-Sasson et al. 2003 (Lemma 4.3)]. Let $B : \mathcal{G}^t \rightarrow [0, 1]$. Then, for any $\epsilon > 0$,

$$\Pr_{\substack{\mathbf{x} \in \mathcal{G}^t \\ \mathbf{y} \in S}} \left[\left| \mathbb{E}_{\beta \in \mathcal{G}} [B(\mathbf{x} + \beta \mathbf{y})] - \mathbb{E}_{\mathbf{z} \in \mathcal{G}^t} [B(\mathbf{z})] \right| > \epsilon \right] \leq \left(\frac{1}{|\mathcal{G}|} + \lambda \right) \frac{\mathbb{E}_{\mathbf{z} \in \mathcal{G}^t} [B(\mathbf{z})]}{\epsilon^2}.$$

2.7.4 Partitions

Given a 2CSP formula, we provide a few techniques to deterministically partition n variables to approximately \sqrt{n} subsets of approximately \sqrt{n} variables each, so that the number of constraints between every pair of partitions is approximately as expected.

Greedy Partition

Lemma 2.5 Let $G = (V, E)$ be a d -regular graph and $n \triangleq |V|$. We can partition V into n/k disjoint subsets $\{S_1, \dots, S_{n/k}\}$ of size at most $2k$ such that:

$$\forall i, j \left| (S_i \times S_j) \cap E \right| \leq 8d^2k^2/n. \quad (2.1)$$

Proof We assign vertices to subsets iteratively, and show by induction that we can always maintain (2.1) and the bound on the subset size. Since the average set size is less than k , we have by Markov's inequality that at each step less than half of the subsets are full. The next vertex we want to assign, v , has neighbors in at most d subsets. By our induction hypothesis, each S_i is of size at most $2k$, so on average over $j \in [n/k]$, it has less than Adk^2/n neighbors in each S_j . Applying Markov's inequality again, S_i has at least $8d^2k^2/n$ neighbors in less than a $(1/2d)$ -fraction of subsets S_j . In total, we ruled out less than half of the subsets for being full, and less than half of the subsets for having too many neighbors with subsets that contain neighbors

2. Note that we need an ϵ -biased set for a large field $\mathcal{G} = \mathbb{F}_{2^\ell}$. Such constructions are not as common in the literature, which mostly focuses on the field \mathbb{F}_2 . To the best of our knowledge, existing explicit constructions for larger fields require much larger cardinality. Nevertheless, for our modest pseudo-randomness desiderata, we could actually use the explicit construction from Alon et al. [1992]. For ease of presentation, we prefer to brute-force derandomize the construction from Azar et al. [1998].

of v . Therefore there always exists some subset S_i to which we can add v while maintaining the induction hypothesis. ■

Derandomized Partition

We use Chernoff bound with $\Theta(\log n)$ -wise independent variables to deterministically partition variables into subsets of cardinality $\approx \sqrt{n}$. Our (somewhat naive) deterministic algorithm for finding a good partition takes quasi-polynomial time ($n^{O(\log n)}$), which is negligible with respect to the sub-exponential size ($N = 2^{\tilde{O}(\sqrt{n})}$) of our reduction.³

Lemma 2.6 Let $G = (A, B, E)$ be a bipartite (d_A, d_B) -bi-regular graph, and let $n_A \triangleq |A|, n_B \triangleq |B|$; set also $n \triangleq n_B + n_A$ and $\rho \triangleq \sqrt{n} \log n$. Let $T_1, \dots, T_{n_B/\rho}$ be an arbitrary partition of B into disjoint subsets of size ρ . There is a quasi-polynomial deterministic algorithm (alternatively, linear-time randomized algorithm) that finds a partition of A into $S_1, \dots, S_{n_A/\rho}$, such that:

$$\forall i \quad \left| |S_i| - \rho \right| < \rho/2, \quad (2.2)$$

and

$$\forall i, j \quad \left| |(S_i \times T_j) \cap E| - \frac{d_A \rho^2}{n_B} \right| < \frac{d_A \rho^2}{2n_B}. \quad (2.3)$$

Proof Suppose that we place each $a \in A$ into a uniformly random S_i . By Chernoff bound and union bound, (2.2) and (2.3) hold with high probability. Now, by Chernoff bound for k -wise independent variables (Theorem 2.4), it suffices to partition A using a $\Theta(\log n)$ -wise independent distribution. Such distribution can be generated with a sample space of $n^{O(\log n)}$ (e.g., [Alon et al. 1986]). Therefore, we can enumerate over all possibilities in quasi-polynomial time. By the probabilistic argument, we will find at least one partition that satisfies (2.2) and (2.3). ■

2.7.5 How to Catch a Far-from-Uniform Distribution

The following lemma due to Daskalakis and Papadimitriou [2009] implies that

Lemma 2.7 Lemma 3 in the full version of Daskalakis and Papadimitriou [2009]. Let $\{a_i\}_{i=1}^n$ be real numbers satisfying the following properties for some $\theta > 0$: (1) $a_1 \geq a_2 \geq \dots \geq a_n$; (2) $\sum a_i = 0$; (3) $\sum_{i=1}^{n/2} a_i \leq \theta$. Then $\sum_{i=1}^n |a_i| \leq 4\theta$.

3. Do not confuse this with the quasi-polynomial lower bound ($N^{\tilde{O}(\log N)}$) we obtain for the running time of the community detection problem.

2.7.6 Simulation Theorem

Let $D : \{0, 1\}^N \rightarrow \{0, 1\}$ be a decision problem. We consider the following query complexity model (called also *decision tree complexity*). Each query is an index $k \in [N]$ and the answer is the k -th bit of the input. The randomized query complexity of D is denoted by $\text{BPP}_\delta^{\text{dt}}(D)$, where δ is the allowed probability of error.

We also consider the following communication complexity model. Here, for every $k \in [N]$, Alice holds a vector $\alpha_k \in \{0, 1\}^M$ and Bob holds an index $\beta_k \in [M]$, for some $M = \text{poly}(N)$. Their goal is to compute D for the input $(\alpha_1(\beta_1), \dots, \alpha_N(\beta_N))$. The standard bounded error two-party probabilistic communication complexity of the simulated problem D is denoted by $\text{BPP}_\delta^{\text{cc}}(\text{SIM}-D)$.

To “lift” from query complexity hardness to communication complexity, we use the following recent *simulation theorem* for BPP, due to [Göös et al. \[2017\]](#).


Theorem 2.5 BPP simulation theorem, [[Göös et al. 2017](#)]. There exists $M = \text{poly}(N)$ such that for any constants $0 < \delta < 1/2$,

$$\text{BPP}_\delta^{\text{cc}}(\text{SIM}-D) = \Omega \left(\text{BPP}_\delta^{\text{dt}}(D)(\log N) \right).$$



PART

**COMMUNICATION
COMPLEXITY**



Communication Complexity of Approximate Nash Equilibrium

The main motivation for studying the complexity of approximate Nash equilibrium is the insight about the relevance of Nash equilibrium as a predictive solution concept: if specialized algorithms cannot compute an (approximate) equilibrium, it is unreasonable to expect selfish agents to “naturally” converge to one. (See also discussions in the introduction, as well as [Daskalakis et al. \[2009a\]](#), [Hart and Mansour \[2010\]](#), [Nisan \[2009b\]](#).) Although extremely useful and the main focus of this book, lower bounds on computational complexity suffer from an obvious caveat: we actually don’t know how to truly *prove* any computational lower bounds: All our computational lower bounds inherently rely on complexity assumptions (such as $\text{NP} \neq \text{P}$ or $\text{PPAD} \neq \text{P}$); even though these assumptions are widely accepted by computer scientists, they make these theorems less accessible to game theorists and economists. For example, it is not clear how they relate to the uncoupled dynamics model studied in game theory [[Babichenko 2012](#), [Hart and Mas-Colell 2003, 2006](#)].

In this part of the book, we prove *unconditional* lower bounds on the communication complexity of approximate Nash equilibrium. In the *communication complexity* model, each player knows her own utility function, and we restrict the amount of information exchanged between players in order to agree on an approximate Nash equilibrium. The players in this model are unreasonably powerful beings with unlimited computational power. In this sense, obtaining lower bounds even in this setting is more convincing than our computational complexity results.

Furthermore, our lower bounds on communication complexity translate immediately to the *uncoupled dynamics* model mentioned above (see also Subsection 3.1). The tradeoff is that the computational complexity lower bounds we can prove are stronger. Take two-player games with N actions, for example. The main result in this book is that no polynomial time algorithm can find an approximate equilibrium. In the communication complexity model, per contra, it is trivial for both players to send their entire $N \times N$ utility matrix; hence the most we can hope for is a polynomial lower bound.

Indeed, our communication complexity results do not directly fit into the “between P and NP” theme of this book. However, we chose to include them because they provide further evidence that real players may not converge to a Nash equilibrium. More importantly, en route to obtaining our communication complexity lower bounds, we develop a construction of a hard-to-find Brouwer fixed point. This construction will be useful in Chapters 6 and 16.

Our Results

We study both two-player games with a large number (N) of actions, and two-action games with a large number (n) of players. The trivial solution of communicating every player’s entire utility function in normal form requires $O(N^2)$ and $O(n2^n)$ communication, respectively.¹ For constant approximation, no non-trivial lower bounds were previously known for the general model, and even for the restricted case of randomized query complexity (see Subsection 3.3) both settings were stated as open questions in Fearnley et al. [2013] and Babichenko [2016], Chen et al. [2017], Hart and Nisan [2013], respectively. For n -player, Hart and Mansour [2010] gave an $\exp(n)$ lower bound on the communication complexity of exact Nash equilibrium in n -player games as also $\exp(n)$,² and even for an approximate parameter of $1/\text{poly}(n)$ the problem was open [Nisan 2009a].

For two-player games, we prove a polynomial lower bound on the communication complexity:

Theorem 3.1 There exists a constant $\epsilon > 0$ such that the randomized communication complexity (BPP^{cc}) of ϵ -Nash equilibrium in two-player $N \times N$ games is at least N^ϵ .

1. Since we care about ϵ -approximate equilibrium for constant ϵ , it suffices to represent the utility with constant precision.

2. The unique Nash equilibrium in the game considered by Hart and Mansour [2010] requires probabilities that are *doubly exponentially* small. Hence their lower bound is exponential in the number of players, but only polynomial in the size of the description of the equilibrium; see Nisan [2009a].

For n -player games, we consider a two-party communication problem where the set of players $[n]$ is divided into two disjoint subsets $[n] = n^A \cup n^B$. Alice holds the utilities of the players in n^A , and Bob holds the utilities of the players in n^B . In particular, this communication problem is easier than the n -parties communication problem where each player holds his own utility function. Furthermore, our negative result holds for the notion of *weak* approximate Nash equilibrium [Babichenko et al. 2016], which in particular implies the same negative result for the standard notion of approximate Nash equilibrium (see also Definition 2.3).

Theorem 3.2 There exists a constant $\epsilon > 0$ such that the randomized communication complexity (BPP^{cc}) of (ϵ, ϵ) -weak approximate Nash equilibrium in n -player binary-action games is at least $2^{\epsilon n}$.

3.1 Uncoupled Dynamics

An underlying assumption of the Nash equilibrium solution is that players predict correctly the (mixed) action of their opponents (or alternatively predict correctly their expected payoff at each action). One justification for this problematic assumption, which appears in the seminal work of John Nash [Nash 1951], is that *in some scenarios* players may *learn* the behavior of their opponents in cases where the game is played repeatedly. This idea led to an extensive study of learning dynamics and their convergence to Nash equilibrium (see, e.g., [Young 2004, Hart and Mas-Colell 2013, Kalai and Lehrer 1993]). One natural, and general, class of adaptive dynamics is that of *uncoupled dynamics* [Hart and Mas-Colell 2003, 2006], where it is assumed that players do not know the utilities of their opponents (but observe their past behavior). The question on the *existence* of uncoupled dynamics that lead to Nash equilibrium is quite well understood [Babichenko 2012, Foster and Young 2006, Germano and Lugosi 2007, Hart and Mas-Colell 2006]. Several uncoupled dynamics that converge to approximate Nash equilibrium (or pure Nash equilibrium [Young 2009]) are known. All these dynamics are based on an exhaustive search principle, where at the moment a player realizes she is acting sub-optimally she updates her action to a random one (rather than to an optimal one or a better one). One criticism of these dynamics is that convergence to equilibrium may take an unreasonably long time in large games where the exhaustive search is done over a large space. This led to the study of the *rate of convergence* of uncoupled dynamics. As pointed out by Conitzer and Sandholm [2004], for every solution concept (in particular equilibria solutions), the (randomized) communication complexity of a solution is identical (up to a logarithmic factor) to the rate of convergence

by any (randomized) uncoupled dynamics to the solution. This observation initiated the communication complexity study in games. As was mentioned above, the communication complexity, and thus also the rate of convergence of uncoupled dynamics, was known only for exact or pure Nash equilibrium. The question on the rate of convergence of uncoupled dynamics to approximate Nash equilibrium was an open question. Given the fact that all known positive results introduce dynamics that converge to *approximate* Nash equilibrium, this question is central. Our results for communication complexity resolve this open question, yielding the following negative results for uncoupled dynamics:

Corollary 3.1 *Uncoupled dynamics.* There exists a constant $\epsilon > 0$ such that any uncoupled dynamics require:

2-player. At least $\text{poly}(N)$ rounds to converge to an ϵ -Nash equilibrium in two-player $N \times N$ games.

n -player. At least $2^{\Omega(n)}$ rounds to converge to an ϵ -Nash equilibrium (or even (ϵ, ϵ) -weak approximate Nash equilibrium) in n -player binary-action games.

3.2 Techniques

Proving communication complexity lower bounds for Nash equilibrium is notoriously challenging for two reasons. The first reason, as is common in hardness of Nash equilibrium in other models, is *totality*: there always exists at least one (exact) equilibrium, and the proof of existence induces a non-trivial (yet inefficient) algorithm for finding it. In order to construct hard instances we must carefully hide the equilibrium (we can't just remove it), and make sure that the above algorithm is indeed inefficient for our instances.

Another reason for the communication complexity of approximate equilibrium being an open question for a long time is the fact that there exist efficient *non-deterministic* communication protocols ($\text{polylog}(N)$ for two-player, $\text{poly}(n)$ for n -player): verification of equilibrium (exact or approximate) requires only constant communication, and small-representation approximate equilibria always exist (e.g., by [Lipton et al. \[2003\]](#)). Therefore, the communication complexity lower bounds for approximate equilibria, as we prove in this chapter, show an exponential gap between the non-deterministic and randomized (or even deterministic) communication complexity of a total search problem. We are certainly not the

first to show such separations (see, e.g., [Karchmer et al. 1995, Raz and McKenzie 1999, Raz and Wigderson 1990]).³ But such separations are still not very common in communication complexity, and for a good reason: for decision problems, they are impossible! The deterministic communication complexity is upper-bounded by the product of the non-deterministic and co-non-deterministic communication complexities [Aho et al. 1983].

The first ingredient in our proof is a construction of a special continuous function $f : [0, 1]^n \rightarrow [0, 1]^n$ whose approximate fixed points are hard to find. The construction is inspired by that of Hirsch et al. [1989], and the main new ingredient is the use of error-correcting codes to replace the ℓ_∞ inapproximability with ℓ_2 inapproximability. The construction appears in Chapter 4. The second ingredient in our proof is the *simulation theorems* for randomized communication complexity due to Anshu et al. [2017] and Göös et al. [2017].

The main steps in our proofs are as follows. First, we prove a randomized *query complexity* hardness result for the problem of finding the end of a line in a particular constant-degree graph. Then we use a *simulation theorem* of Anshu et al. [2017] and Göös et al. [2017] to “lift” this query complexity hardness result to randomized communication complexity hardness. We use the construction in Chapter 4 to embed this line as a continuous Lipschitz function $f : [0, 1]^n \rightarrow [0, 1]^n$. Finally, we build on ideas from Babichenko [2016], McLennan and Tourky [2005], and Shmaya [2012] to construct a two-player (respectively n -player) “imitation game” that simulates the short communication protocol for the computation of f , as well as a fixed point verification. In particular, every (approximate) Nash equilibrium of the game corresponds to an approximate fixed point of f , which in turn corresponds to an end of a line.

Since in this chapter we are proving unconditional intractability results, we have the privilege of reasoning about an explicit distribution of hard instances. In particular, it suffices to begin with the END-OF-THE-LINE special case of the END-OF-A-LINE problem, where the graph consists of just one line—and we want to find *the* end of that line. This hardly changes the proof, but it makes the notation a little simpler. For example, it suffices to prove that a *decision* problem (find the most significant bit of the end of the line) is hard. Furthermore, our hardness now holds for the interesting case where the game has a unique Nash equilibrium.

3. It is interesting to remark that our result is arguably the first example of a *natural* problem that exhibits such a gap: To the best of our knowledge, approximate Nash equilibrium is the first problem that is not defined in order to exhibit a gap, but rather happens to have one.

3.3 Additional Related Literature

Pure Nash Equilibrium. The communication complexity of *pure* Nash equilibrium has been studied before: in two-player $N \times N$ games it is $\text{poly}(N)$ [Conitzer and Sandholm 2004], and in n -player games it is $\exp(n)$ [Hart and Mansour 2010].

Approximation Protocols. For two-player $N \times N$ games and $\epsilon \approx 0.382$, Czumaj et al. [2015] show that $\text{polylog}(N)$ communication is sufficient for computing an ϵ -approximate Nash equilibrium (improving over a protocol for $\epsilon \approx 0.438$ due to Goldberg and Pastink [2014]).

Query Complexity. There are several interesting results on the more restricted model of query complexity of approximate Nash equilibria, where the algorithm is assumed to have black-box access to the normal form representation of the utility function. Note that our communication complexity lower bounds immediately extend to this model as well.

Hart and Nisan [2013] prove that any deterministic algorithm needs to query at least an exponential number of queries to compute any ϵ -well-supported Nash equilibrium—and even any ϵ -correlated equilibrium. Babichenko [2016] showed that any randomized algorithm requires an exponential number of queries to find an ϵ -well-supported Nash equilibrium. Chen et al. [2017] extended Babichenko’s result to an almost-exponential ($2^{\Omega(n/\log n)}$) lower bound on the query complexity of ϵ -approximate Nash equilibrium. Note that our lower bound here is not only bigger (saving the $\log n$ factor), but also holds for the more general notion of weak approximate Nash equilibrium, and in the more general model of communication complexity.

Goldberg and Roth [2016] give a polynomial upper bound on the query complexity of ϵ -WSNE for any family of games that have *any* concise representation. This result is to be contrasted with (a) Babichenko’s query complexity lower bound, which uses a larger family of games, and (b) our lower bounds on the *computational complexity* of succinctly represented games (Theorem 5.1).

A much older yet very interesting and closely related result is that of Hirsch et al. [1989]. They show that any deterministic algorithm for computing a Brouwer fixed point in the oracle model must make an exponential -in the dimension n and the approximation ϵ - number of queries for values of the function. Our construction here builds upon and improves over Hirsch et al. [1989] by working with the ℓ_2 -norm instead of the ℓ_∞ -norm.

Correlated Equilibrium. For the related notion of correlated equilibrium, in n -player games with a constant number of actions, it is known that even exact

correlated equilibrium can be computed using only $\text{poly}(n)$ -communication (see [Hart and Mansour 2010, Jiang and Leyton-Brown 2015, Papadimitriou and Roughgarden 2008]). Interestingly, for exact correlated equilibria, there is an exponential gap between the above communication protocol and the *query complexity* lower bound of Babichenko and Barman [2015] and Hart and Nisan [2013]. Goldberg and Roth [2016] characterize the query complexity of approximate *coarse* correlated equilibrium in games with many players. Further discussion on correlated equilibria appears in Section 3.6.

Communication Complexity of Finding Fixed Points. For the related problem of finding a fixed point, Roughgarden and Weinstein [2016] study the communication complexity of an approximate fixed point of the decomposition. Namely, Alice holds a Lipschitz function $f : A \rightarrow B$ and Bob holds a Lipschitz function $g : B \rightarrow A$, where A and B are compact convex sets, and their goal is to compute a fixed point of the decomposition $g \circ f$. Roughgarden and Weinstein prove that the following version of this problem is communicationally hard: find an approximate fixed point of $g \circ f$ on a grid of A , when it is promised that such an approximate fixed point *on the grid* exists (the problem is not total).

Complexity of Equilibrium and Price of Anarchy. As discussed earlier, the main motivation for studying the (communication) complexity of Nash equilibrium is understanding its relevance as a predictive solution concept. This is a good place to mention a recent work of Roughgarden [2014], which highlights another important motivation for studying the complexity of Nash equilibrium: understanding the *quality* of equilibria. The *Price of Anarchy (PoA)* of a game is the ratio between the *social welfare* (sum of players' utilities) in an optimum strategy profile, and the social welfare in the worst Nash equilibrium of that game. Roughgarden [2014] provides the following widely applicable recipe for lower bounds on PoA: if a Nash equilibrium can be found efficiently (in particular, via the non-deterministic protocol due to Lipton et al. [2003]), but approximating the optimal social welfare requires a higher communication complexity (even for non-deterministic protocols, e.g., by reduction from set disjointness), then clearly not all Nash equilibria yield high social welfare.

3.4 Proof Overview

The formal proofs appear in Section 3.5. Below we present the main ideas of the proof. As mentioned in the Introduction, the proof consists of four main steps. Below we present the ideas of each step.

Query Complexity of End-of-the-Line

Our proof starts with the following *query complexity* hardness result (Lemma 3.2): There exists a constant degree graph $G = (V, E)$ with $2^{\Theta(n)}$ vertices, such that finding the end of a line in G requires $2^{\Omega(n)}$ queries. In fact, we prove the hardness result for directed graph G where each vertex has outgoing and incoming degree 2. Therefore, the successor and predecessor of each vertex are binary variables. In particular, for each $v \in V$, the information about its role in the line can be represented using only three bits, which we denote $I(v) \triangleq (T(v), P(v), S(v)) \in \{0, 1\}^3$:

- (a) Whether the line goes through v , which is denoted by $T(v)$,
- (b) Who is the successor of v (if v is on the line), which is denoted by $S(v)$,
- (c) Who is the predecessor of v (if v is on the line), which is denoted by $P(v)$.

Lemma 3.1 QUERY END-OF-A-LINE; informal. Finding an end of a line with high probability requires $2^{\Omega(n)}$ queries to I .

From Query Complexity to Communication Complexity

We use a recent *simulation theorem* to “lift” our randomized query complexity lower bound to a randomized communication complexity bound.

The simulated communicationally hard problem has the following form. For each $v \in V$, Alice holds a triplet of vectors $\alpha_{T,v}, \alpha_{S,v}, \alpha_{P,v} \in \{0, 1\}^M$ where $M = 2^{O(n)}$, and Bob holds a reasonably small input which is just a triplet of indexes $\beta_{T,v}, \beta_{S,v}, \beta_{P,v} \in [M]$. $T(v)$ is given by the decomposition $T(v) = \alpha_{T,v}(\beta_{T,v})$ (similarly for the successor and predecessor). The simulation theorem of Göös et al. [2017] and Anshu et al. [2017] now implies:

Corollary 3.2 CC(SIMULATION END-OF-A-LINE); informal. Finding an end of a line requires $2^{\Omega(n)}$ bits of communication.

Embedding as a Continuous Function

Our next step is to reduce the problem of finding an end of a line to that of finding a Brouwer fixed point. Here, we use the construction from Chapter 4.

We embed the vertices of the discrete graph G in the continuous space $[-1, 2]^{\Theta(n)}$. Specifically, we embed each vertex v of G into a point \mathbf{x}_v in $[-1, 2]^{\Theta(n)}$ and each edge (v, w) in G into a (continuous) path in $[-1, 2]^{\Theta(n)}$ that connects the corresponding points \mathbf{x}_v and \mathbf{x}_w . In particular, we construct a continuous Lipschitz function $f : [-1, 2]^{\Theta(n)} \rightarrow [-1, 2]^{\Theta(n)}$ such that:

1. The computation of f can be done using *local information about I* . Namely, for points that are close to \mathbf{x}_v it is sufficient to know $I(v)$ to compute f . For points that are close to a path that corresponds to the edge (v, w) but far from the points $\mathbf{x}_v, \mathbf{x}_w$, it is sufficient to know whether (v, w) is an edge in the line (in particular, knowing either $I(u)$ or $I(v)$ suffices). For points that are far from all paths (v, w) , f does not depend on I at all (and thus can be computed without any communication).
2. Any (normalized) $\|\cdot\|_2$ -approximate fixed point of f can be mapped (efficiently) back to an end of some line in I .

Property 1 induces the following efficient communication protocol for computing $f(\mathbf{x})$: Bob finds v such that \mathbf{x} is close to \mathbf{x}_v , and sends $\beta_{T,v}, \beta_{S,v}, \beta_{T,v}$; Alice replies with $I(v) = (\alpha_{T,v}(\beta_{T,v}), \alpha_{T,v}(\beta_{T,v}), \alpha_{T,v}(\beta_{T,v}))$, and they each use $I(v)$ to locally compute $f(\mathbf{x})$. (Similarly, if \mathbf{x} is close to the path corresponding to edge (v, w) , they use a similar protocol to compute $I(v)$ and $I(w)$.)

By Property 2, we have:

Corollary 3.3 CC(BROUWER); informal. Finding a (normalized) $\|\cdot\|_2$ -approximate fixed point of f requires $2^{\Omega(n)}$ bits of communication.

Two-Player Game

Naively thinking, we would like to design a game where Alice chooses a point $\mathbf{x} \in [-1, 2]^{\Theta(n)}$ (on the ε -grid) and Bob chooses a point $\mathbf{y} \in [-1, 2]^{\Theta(n)}$ (on the ε -grid). Alice's utility will be given by $-\|\mathbf{x} - \mathbf{y}\|_2^2$, and Bob's utility will be given by $-\|\mathbf{y} - f(\mathbf{x})\|_2^2$. Then, by applying similar arguments to those in Babichenko [2016], McLennan and Tourky [2005], Rubinstein [2016], and Shmaya [2012], we can deduce that every approximate Nash equilibrium corresponds to an approximate fixed point, and thus also to an end of a line.

However, the above idea is obviously incorrect because Bob's utility depends on f , whereas in the communication problem his utility should depend on the β s only. Our key idea is to use the fact that f can be computed locally to design a somewhat similar game where similar phenomena to those in the "naive" game will occur in approximate equilibria.

Bob doesn't know f , but to compute $f(\mathbf{x})$ he should only know the local information about the vertex (or vertices) that corresponds to \mathbf{x} . We incentivize Alice

4. Note that here it is crucial that we use the normalized $\|\cdot\|_2$ to obtain payoffs bounded in $[-9, 0]$; using the non-normalized $\|\cdot\|_2$ we get payoffs in $[-\sqrt{n}, 0]$.

and Bob to combine their private information about the corresponding vertex (or vertices) by the following utilities structure.

- Alice's first component of utility is given by $-\|\mathbf{x} - \mathbf{y}\|_2^2$. As in the "naive" game, in any approximate Nash equilibrium Alice will play points in the ϵ -cube of the ϵ -grid that contains $\mathbb{E}[\mathbf{y}]$ with probability close to one.
- Bob tries to guess the vertex v (or the vertices v, w) that corresponds to the point \mathbf{x} . Since \mathbf{x} (almost always) belongs to the same ϵ -cube, in any (approximate) Nash equilibrium, his guess should be correct (with high probability). In addition, Bob should announce the β indexes β_T, β_S , and β_P of v (of v and w). Again, we incentivize him to do so by defining that he should "guess" also these β indexes, and in an (approximate) equilibrium his guess should be correct with high probability (w.h.p.).
- We want Alice to announce $I(v)$ (similarly for w in the case of two vertices). Thus, we ask her to guess the decomposition $\alpha_{v,B}(\beta^B)$ where v^B and β^B are the announced v and β by Bob. In (approximate) equilibrium, since Bob announces the correct v and β (w.h.p.), this incentivizes her to announce the correct $I(v)$ (w.h.p.).
- Now Bob uses the local information of $I(v)$ (and $I(w)$) to compute f . Namely, his last utility component is defined by $-\|\mathbf{y} - f_{I^A(v), I^A(w)}(\mathbf{x})\|_2^2$ where $f_{I^A(v), I^A(w)}$ is Bob's "estimation" of f under the relevant local information announced by Alice. In (approximate) equilibrium Alice announces the correct local information (w.h.p.); thus Bob computes f correctly (w.h.p.).

Summarizing, the (approximate) equilibrium analysis of the presented game is similar to the analysis of the naive game, because in (approximate) equilibrium f is computed correctly (w.h.p.). But unlike the naive game, here Alice's utility depends only on the α s and Bob's utility only on the β s.

***n*-Player Game: ϵ -WSNE**

The n -player game reduction is based on the same ideas as the two-player reduction. For clarity, we present first the idea of a reduction that proves the following weaker result:

There exists a constant $\epsilon > 0$ such that the communication complexity of ϵ -well-supported approximate Nash equilibrium in n -player games with a constant number of actions for each player is at least 2^{cn} for some constant c .

After that, we explain how we can strengthen this result in two aspects: first to improve the constant-number-of-action to binary-action, second to improve the ϵ -well-supported Nash equilibrium to (ϵ, ϵ) -weak approximate equilibrium.

The idea is to replace a single player—Alice—who chooses \mathbf{x} in the ϵ -grid of $[-1, 2]^{\Theta(n)}$ by a population of $\Theta(n)$ players $\{p_{x_i}\}_{i \in \Theta(n)}$; each player p_{x_i} in the population is responsible for the i th coordinate of \mathbf{x} . The payoff of player p_{x_i} is given by $-|x_i - y_i|^2$. This incentivizes player p_{x_i} to play a single action, or two adjacent actions, in the ϵ -grid of the segment $[-1, 2]$ (in every WSNE). By looking at the action profile of all p_{x_i} players, we get the same phenomenon as in the two-player case: every point x in the support of Alice's players belongs to the same ϵ -cube of the ϵ -grid.

Now, we replace the guess of $v \in \{0, 1\}^{\Theta(n)}$, which is done by Bob, by population of size $\Theta(n)$ where again each player is responsible to a single coordinate of v . Again, in a WSNE all players will guess correctly.

Similarly for the guess of β : we think of $\beta \in [M]^3$ as an element of $\{0, 1\}^{3 \log M}$ and we construct a population of $3 \log M$ players; each controls a single bit.

Similarly for Alice's guesses of $I^A(v)$ and $I^A(w)$: we construct 6 players, and each chooses a bit.

Finally, we again replace the choice of $\mathbf{y} \in [-1, 2]^{\Theta(n)}$ by a population of $\Theta(n)$ players p_{y_i} . Each is responsible to a single coordinate. The payoff of player p_{y_i} is given by $-|y_i - [f_{I^A(v), I^A(w)}(\mathbf{x})]_i|^2$. The analysis of this game is very similar to the two-player game analysis.

n -Player Game: (ϵ, ϵ) -Weak ANE and Binary Actions

In the above reduction, the \mathbf{x} -type (and \mathbf{y} -type) players have $3/\epsilon$ actions each. To construct a binary action game, we use the technique of Babichenko [2016]. We replace each such player by a population of $3/\epsilon$ players, each located at a point in the ϵ -grid of the segment $[-1, 2]$. A player that is located at $j \in [-1, 2]$ (on the ϵ -grid) has to choose between the two points j or $j + \epsilon$. In a WSNE, all players located to the left of y_i will choose $j + \epsilon$, and all players located to the right of y_i will choose j .

More tricky is to generalize this reduction to weak approximate equilibria. Recall that in weak approximate equilibria, a constant fraction of players may play an arbitrary suboptimal action. Here we take into account both;

1. Players that are not ϵ -best replying, and
2. Players that are ϵ -best replying, but put small positive weight on the inferior action (among the two) and the realization of their mixed action turns out to be the inferior action.

In order to be immune from this small constant fraction of irrational players, we use error-correcting codes.⁵ Let $E_\beta: \{0, 1\}^{3 \log M} \rightarrow \{0, 1\}^{\Theta(3 \log M)}$ be a good binary error-correcting code. Instead of having a population of size $3 \log M$ that tries to guess β , we replace it by a population of size $\Theta(3 \log M)$ where each player tries to guess his bit in the *encoding* of β . Now even if a small constant fraction of players will act irrationally, the *decoding* of the action profile of the β -type players will turn out to be β . We use the same idea for all types of populations (x-type, y-type, v -type, and I -type). This idea completes the reduction for weak approximate equilibria.

3.5 Proofs

In Subsection 3.5.1 we prove a randomized query lower bound for the end-of-the-line problem. In Subsection 3.5.2 we show how the lower bounds of Subsection 3.5.1 can be “lifted” to get a hard problem in the randomized communication complexity models. In Subsections 3.5.3, 3.5.4, and 3.5.5 we reduce the communicationally hard end-of-any-line problem to the approximate Nash equilibrium problem.

3.5.1 A Randomized Query Complexity Lower Bound

Let G be a *directed* graph with the vertices $V = \{0, 1\}^n \times \{0, 1\}^n \times [n + 1]$. Each vertex (v_1, v_2, k) , where $v_1, v_2 \in \{0, 1\}^n$ and $k \in [n]$, has two outgoing edges to the vertices $(v_1, v_2^{k+1}(0), k + 1)$ and $(v_1, v_2^{k+1}(1), k + 1)$, where $v^j(0) = (v_1, \dots, v_{j-1}, 0, v_{j+1}, \dots, v_n)$. We call $(v_1, v_2^{k+1}(0), k + 1)$ the *0-successor* of v , and $(v_1, v_2^{k+1}(1), k + 1)$ the *1-successor* of v . Each vertex $v = (v_1, v_2, n + 1)$ has a single outgoing edge to the vertex $(v_2, v_1, 0)$. Note that the incoming degree of each vertex $v = (v_1, v_2, k) \in V$ is at most two. For $k = 1$ there is a single incoming edge from $(v_2, v_1, n + 1)$. For $k > 1$ there are two incoming edges from $(v_1, v_2^k(0), k - 1)$ and from $(v_1, v_2^k(1), k - 1)$. We call $(v_1, v_2^k(0), k - 1)$ the *0-predecessor* of v , and $(v_1, v_2^k(1), k - 1)$ the *1-predecessor* of v .

We define the QUERY END-OF-THE-LINE to be the problem of finding the end of a line in G that starts at the point $\mathbf{0}_{2n+1}$. More formally, we represent a line in G by a triple $I(v) \triangleq (T(v), S(v), P(v))$ where $T(v) \in \{0, 1\}$ indicates whether the line goes through v , $S(v) \in \{0, 1\}$ indicates who is the successor of v , and $P(v) \in \{0, 1\}$ indicates who is the predecessor of v (here we use the fact that each vertex has outgoing and incoming degree of at most two). Throughout the chapter, we slightly abuse notation and use $S(v)/P(v)$ to refer both to the bits and to the corresponding

5. In fact, we use error-correcting codes even earlier, in Rubinfeld’s [2016] modified construction of a hard Brouwer function.

vertices (i.e., the $S(v)/P(v)$ -successor/predecessor of v). The end of the line is the vertex v^* such that $T(v^*) = 1$ but $T(S(v^*)) = 0$.

Definition 3.1 QUERY END-OF-THE-LINE.

Input: A line $I = (T, S, P)$ over the graph G that starts at the point $\mathbf{0}_{2n+1}$.

Output: The first bit ($[v^*]_1$) of the end-of-the-line vertex.

Queries: Each query is a vertex $v \in V$. The answer is the triplet of bits $I(v) = (T(v), S(v), P(v)) \in \{0, 1\}^3$.

Lemma 3.2 Randomized query complexity. For every constant $\delta < \frac{1}{2}$, $\text{BPP}_\delta^{\text{dt}}(\text{QUERY END-OF-THE-LINE}) = \Omega(2^n)$.

Proof By Yao's Minmax Theorem it is sufficient to introduce a distribution over paths such that every deterministic query algorithm requires $\Omega(2^n)$ queries to determine the first bit of the end-of-line vertex with probability of at least $1 - \delta$. We choose a permutation π over $\{0, 1\}^n \setminus \{0_n\}$ uniformly at random, and set $\pi(0) \triangleq \mathbf{0}_n$. π induces a line of length $\Theta(2^n \cdot n)$ over G , starting at $\mathbf{0}_{2n+1}$, ending at $(\pi(2^n - 1), \pi(2^n - 1), 0)$, and where two consecutive vertices $v = \pi(i)$ and $w = \pi(i + 1)$ are mapped to the following line of $n + 1$ edges:

$$\begin{aligned} (v, v, 0) &\rightarrow \dots \rightarrow (v, (w_{[1,k]}, v_{[k+1,n]}), k) \rightarrow \\ &\rightarrow \dots \rightarrow (v, w, n) \rightarrow (w, w, 0). \end{aligned}$$

Here $(w_{[1,k]}, v_{[k+1,n]})$ denotes the vector with the first k coordinates as in w and the last $n - k$ coordinates as in v .

The information of a single query of QUERY END-OF-THE-LINE (for the above class of lines) can be extracted from $\pi(i - 1)$, $\pi(i)$, and $\pi(i + 1)$. Therefore QUERY END-OF-THE-LINE is at least as hard as the problem of finding the first bit of the last element in a random permutation, where each query returns the previous, the current, and the next vertices. Conditioning on the answers to k queries $\pi(q_1 - 1)$, $\pi(q_1)$, $\pi(q_1 + 1)$, \dots , $\pi(q_k - 1)$, $\pi(q_k)$, $\pi(q_k + 1)$, the last element of the permutation is still uniformly random across all vertices that are not $\pi(q_1)$, \dots , $\pi(q_k)$, $\pi(q_1 - 1)$, \dots , $\pi(q_k - 1)$, $\pi(q_1 + 1)$, \dots , $\pi(q_k + 1)$. This proves that the latter problem requires $\Omega(2^n)$ queries. ■

3.5.2 Communicationally Hard, Discrete End-of-Any-Line Problem

In order to use a simulation theorem (Theorem 2.5) for randomized communication complexity, we define the following simulation variant of QUERY END-OF-THE-LINE:

Definition 3.2 SIMULATION END-OF-THE-LINE. We let $N = 2^n \cdot 2^n \cdot (n + 1) \cdot 3$.

Input: For each $v \in \{0, 1\}^n \times \{0, 1\}^n \times [n + 1]$, Alice receives three vectors $\alpha_v^T, \alpha_v^S, \alpha_v^P \in \{0, 1\}^M$, and Bob receives three indices $\beta_v^T, \beta_v^S, \beta_v^P \in [M]$.

We define

$$T(v) = \alpha_v^T(\beta_v^T), \quad S(v) = \alpha_v^S(\beta_v^S), \quad \text{and} \quad P(v) = \alpha_v^P(\beta_v^P). \quad (3.1)$$

We simulate the problem QUERY END-OF-THE-LINE; therefore we restrict attention to inputs such that (T, S, P) that are defined in (3.1) meet all the requirements of QUERY END-OF-THE-LINE.

Output: The first bit ($[v^*]_1$) of a non-trivial end or start of a line $(v^*, v^*, 0) \neq \mathbf{0}_{2n+1}$.

Applying the randomized simulation theorem (Theorem 2.5) to the query complexity lower bound (Lemma 3.2) gives a lower bound on the randomized communication complexity of a discrete end-of-line problem SIMULATIONEND-OF-THE-LINE.

Corollary 3.4 $\text{BPP}_{0.3}^{\text{cc}}(\text{SIMULATIONEND-OF-THE-LINE}) = \Omega(2^n)$.

3.5.3 Embedding a Line as a Local Lipschitz Function

We embed I as a Euclidean-norm hard continuous function à la Section 4.2. Below, we recall some of the properties of the construction that will be useful for our reduction.

It will be more convenient to think of G as a graph over $\{0, 1\}^{2n+\log(n+1)}$.

Let $m = \Theta(2n + \log(n + 1)) = \Theta(n)$ and let $E: \{0, 1\}^{2n+\log(n+1)} \rightarrow \{0, 1\}^m$ denote the encoding function of a good binary error-correcting code. We embed the discrete graph G into the continuous cube $[-1, 2]^{4m}$.

The vertex v is embedded to the point $(E(v), E(v), \mathbf{0}_m, \mathbf{0}_m) \in [-1, 2]^{4m}$, which is called *the embedded vertex*.

For two vertices $v, w \in V$ such that (v, w) is an edge in the graph G , we define five vertices:

$$\begin{aligned} \mathbf{x}^1(v, w) &\triangleq (E(v), E(v), \mathbf{0}_m, \mathbf{0}_m) \\ \mathbf{x}^2(v, w) &\triangleq (E(v), E(v), \mathbf{1}_m, \mathbf{0}_m) \\ \mathbf{x}^3(v, w) &\triangleq (E(v), E(w), \mathbf{1}_m, \mathbf{0}_m) \\ \mathbf{x}^4(v, w) &\triangleq (E(v), E(w), \mathbf{0}_m, \mathbf{0}_m) \\ \mathbf{x}^5(v, w) &\triangleq (E(w), E(w), \mathbf{0}_m, \mathbf{0}_m). \end{aligned}$$

Note that $\mathbf{x}^1(v, w)$ is the embedded vertex v , $\mathbf{x}^5(v, w)$ is the embedded vertex w .

The line that connects the points $\mathbf{x}^i(v, w)$ and $\mathbf{x}^{i+1}(v, w)$ is called a *Brouwer line segment*. The union of these four Brouwer line segments is called the *embedded edge* (v, w) . It is not hard to check that non-consecutive Brouwer line segments are $\Omega(1)$ -far one from the other, and in particular it implies that non-consecutive embedded edges are sufficiently far one from the other.

The following proposition shows that the END-OF-THE-LINE problem can be reduced to the problem of finding an approximate fixed point of a continuous Lipschitz function, when the function is “local” in the following sense: every edge in G is embedded as a path in the continuous hypercube (as described above). For points close to the embedding of an edge, f depends only on the “local behavior” of the lines I at the endpoints of this edge; for all other points, f is independent of the lines I .

Proposition 3.1 Theorem 4.2 and Fact 4.2. There exist constants $\delta, h > 0$ such that given a line $I = (T, S, P)$ over G there exists a function $f = f(I) : [-1, 2]^{4m} \rightarrow [-1, 2]^{4m}$ with the following properties:

1. $\|f(\mathbf{x}) - \mathbf{x}\|_2 > \delta$ for every \mathbf{x} that is not h -close to the embedded edge of the end of the line (i.e., the embedding of the edge $(P(v^*), v^*)$).
2. f is $O(1)$ -Lipschitz in $\|\cdot\|_2$ norm.
3. f is *local* in the sense that it can be defined as an interpolation between a few (in fact, 64) functions, $\{f_{I_1, I_2} : [-1, 2]^{4m} \rightarrow [-1, 2]^{4m}\}_{I_1, I_2 \in \{0, 1\}^3}$, that do not depend on the lines I and such that:
 - (a) If the first m -tuple of coordinates of \mathbf{x} is $6h$ -close to the encoded vertex $E(v)$, but the second m -tuple of coordinates of \mathbf{x} is $6h$ -far from any encoded vertex $E(w)$, then $f_{I(v), I_2}(\mathbf{x}) = f(\mathbf{x})$ for every $I_2 \in \{0, 1\}^3$.
 - (b) If the second m -tuple of coordinates of \mathbf{x} is $6h$ -close to the encoded vertex $E(w)$, but the first m -tuple of coordinates of \mathbf{x} is $6h$ -far from any encoded vertex $E(v)$, then $f_{I_1, I(w)}(\mathbf{x}) = f(\mathbf{x})$ for every $I_1 \in \{0, 1\}^3$.
 - (c) If the first m -tuple of coordinates of \mathbf{x} is $6h$ -close to the encoded vertex $E(v)$, and the second m -tuple of coordinates of \mathbf{x} is $6h$ -close to the encoded vertex $E(w)$, then $f_{I(v), I(w)}(\mathbf{x}) = f(\mathbf{x})$.
 - (d) If none of the above conditions are satisfied, then $f_{I_1, I_2}(\mathbf{x}) = f(\mathbf{x})$ for every $I_1, I_2 \in \{0, 1\}^3$.

3.5.4 Two-Player Game

Theorem 3.3 Theorem 3.1, restated. There exists a constant $\epsilon > 0$ such that the communication complexity of ϵ -Nash equilibrium in two-player $N \times N$ games is at least N^ϵ .

We construct a two-player game between Alice and Bob of size $N_A \times N_B$ for

$$N_A \triangleq (3/\epsilon)^{4m} \cdot 2^3 = 2^{\Theta(n)}$$

$$N_B \triangleq (3/\epsilon)^{4m} \cdot (2^{2n+\log(n+1)})^2 \cdot M^3 = 2^{\Theta(n)},$$

such that Alice's utility depends on $\{\alpha_v^T, \alpha_v^S, \alpha_v^P\}_v$ only, Bob's utility depends on $\{\beta_v^T, \beta_v^S, \beta_v^P\}_v$ only, and all ϵ^4 -approximate Nash equilibria of the game correspond to a δ -fixed point of f from Proposition 3.1. By property 1 in Proposition 3.1, any fixed point of f corresponds to a non-trivial end or start of a line in I .

3.5.4.1 The Game

In this subsection we construct our reduction from SIMULATION END-OF-THE-LINE to the problem of finding an ϵ -WSNE.

Strategies. Recall that δ is the desired approximation parameter for a Brouwer fixed point in the construction of Proposition 3.1. We let ϵ be a sufficiently small constant; in particular, $\epsilon = O(\delta)$ (this will be important later for Inequality (3.10)).

Each of Alice's actions corresponds to an ordered tuple $(\mathbf{x}, I_v^A, I_w^A)$, where:

- $\mathbf{x} \in [-1, 2]^{4m}$, where the interval $[-1, 2]$ is discretized into $\{-1, -1 + \epsilon, \dots, 2 - \epsilon, 2\}$;
- $I_v^A \triangleq (t_v^A, s_v^A, p_v^A) \in \{0, 1\}^3$ and $I_w^A \triangleq (t_w^A, s_w^A, p_w^A) \in \{0, 1\}^3$.

Each of Bob's actions corresponds to an ordered tuple $(\mathbf{y}, v^B, w^B, \beta_v^B, \beta_w^B)$, where:

- $\mathbf{y} \in [-1, 2]^{4m}$, where the interval $[-1, 2]$ is discretized into $\{-1, -1 + \epsilon, \dots, 2 - \epsilon, 2\}$;
- $v^B, w^B \in \{0, 1\}^{2n+\log(n+1)}$ are vertices in the graph G ;
- $\beta_v^B = (\beta_v^{B,T}, \beta_v^{B,S}, \beta_v^{B,P}) \in [M]^3$ and $\beta_w^B = (\beta_w^{B,T}, \beta_w^{B,S}, \beta_w^{B,P}) \in [M]^3$ are triples of indexes.

Utilities. Alice's and Bob's utilities decompose as

$$U^A \triangleq U_{\text{IMITATION}}^A + U_{\text{GUESSV}}^A + U_{\text{GUESSW}}^A,$$

$$U^B \triangleq U_{\text{BROUWER}}^B + U_{\text{GUESSV}}^B + U_{\text{GUESSW}}^B.$$

The first component of Alice's utility depends only on the first components of her and Bob's strategies; it is given by:

$$U_{\text{IMITATION}}^A(\mathbf{x}; \mathbf{y}) \triangleq -\|\mathbf{x} - \mathbf{y}\|_2^2.$$

Given the first component $\mathbf{x} \in [-1, 2]^{4m}$ of Alice's strategy, we define two decoding functions $D_v, D_w : [-1, 2]^{4m} \rightarrow \{0, 1\}^n$ as follows. Let $R_v(\mathbf{x}) \in \{0, 1\}^m$ be the rounding of the first m -tuple of coordinates of \mathbf{x} to $\{0, 1\}^m$; let $D_v(\mathbf{x}) = E^{-1}(R_v(\mathbf{x})) \in \{0, 1\}^{2n+\log(n+1)}$, where E^{-1} denotes the decoding of the error-correcting code from Subsection 3.5.3. We define $D_w(\mathbf{x}) \in \{0, 1\}^{2n+\log(n+1)}$ analogously with respect to the second m -tuple of coordinates of \mathbf{x} . The second component of Bob's utility is now given by $U_{\text{GUESSV}}^B = 1$ iff he guesses correctly the vertex $D_v(\mathbf{x})$, and the corresponding β operation on this vertex. Namely, $U_{\text{GUESSV}}^B(v^B, \beta_v^B; \mathbf{x}) = 1$ if $v^B = D_v(\mathbf{x})$ and $\beta_v^B = (\beta_{D_v(\mathbf{x})}^T, \beta_{D_v(\mathbf{x})}^S, \beta_{D_v(\mathbf{x})}^P)$, and $U_{\text{GUESSV}}^B(v^B, \beta_v^B; \mathbf{x}) = 0$ otherwise. We similarly define Bob's third component U_{GUESSW}^B with respect to $D_w(\mathbf{x})$.

Note that Bob knows the indexes $\beta_v^T, \beta_v^S, \beta_v^P$ (for every v); thus to achieve $U_{\text{GUESS}}^B = 1$ Bob needs to guess correctly only the vertices $D_v(\mathbf{x}), D_w(\mathbf{x})$ and announce the corresponding triplet of β indexes.

Going back to Alice, the second component of her utility is given by $U_{\text{GUESSV}}^A = 1$ iff she guesses correctly the triplet $I(v^B) = (T(v^B), S(v^B), P(v^B))$ when the calculation of T, S, P is done by the decomposition of $\alpha(\beta^B)$. Namely, $U_{\text{GUESSV}}^A(I_v^A; v^B, \beta^B) = 1$ if $I_v^A = (\alpha_{v^B}^T(\beta_{v^B}^B, T), \alpha_{v^B}^S(\beta_{v^B}^B, S), \alpha_{v^B}^P(\beta_{v^B}^B, P))$, and $U_{\text{GUESSV}}^A(I_v^A; v^B, \beta^B) = 0$ otherwise. We similarly define Alice's third component U_{GUESSW}^B .

Finally, the first component of Bob's utility is given by:

$$U_{\text{BROUWER}}^B(\mathbf{y}; \mathbf{x}, e^A) \triangleq -\|f_{I_v^A, I_w^A}(\mathbf{x}) - \mathbf{y}\|_2^2,$$

where the function f_{I_1, I_2} is as defined in Proposition 3.1.

3.5.4.2 Analysis of Game

In this subsection, we prove the reduction from SIMULATION END-OF-THE-LINE to finding an ϵ^4 -ANE. The proof proceeds via a sequence of lemmas that establish the structure of any ϵ^4 -ANE.

Lemma 3.3 In every ϵ^4 -ANE $(\mathcal{A}; \mathcal{B})$, it holds that $\|\mathbf{x} - \mathbb{E}_{\mathbf{y} \sim \mathcal{B}}[\mathbf{y}]\|_2^2 = O(\epsilon^2)$ with probability of at least $1 - \epsilon^2$ (where the probability is taken over \mathcal{A}).

Proof We denote $\mathbb{E}_i(\mathcal{B}) = \mathbb{E}_{\mathbf{y} \sim \mathcal{B}}[y_i]$, $\mathbb{E}(\mathcal{B}) = (\mathbb{E}_1(\mathcal{B}), \dots, \mathbb{E}_n(\mathcal{B}))$ as the vector of expectations, and $\text{Var}(\mathcal{B}) = (\text{Var}_{\mathbf{y} \sim \mathcal{B}}[y_1], \dots, \text{Var}_{\mathbf{y} \sim \mathcal{B}}[y_n])$ as the vector of variances. For every \mathbf{x} we can rewrite

$$\begin{aligned}
U_{\text{IMITATION}}^A(\mathbf{x}, \mathcal{B}) &= -\mathbb{E}_{\mathbf{y} \sim \mathcal{B}} \|\mathbf{x} - \mathbf{y}\|_2^2 \\
&= -\frac{1}{4m} \sum_{i \in [4m]} \mathbb{E}_{\mathbf{y} \sim \mathcal{B}} [(x_i - y_i)^2] \\
&= -\frac{1}{4m} \sum_{i \in [4m]} [(x_i - \mathbb{E}_i(\mathcal{B}))^2 + \text{Var}_{\mathbf{y} \sim \mathcal{B}}[y_i]] \\
&= -\|\mathbf{x} - \mathbb{E}(\mathcal{B})\|_2^2 - \|\text{Var}(\mathcal{B})\|_2^2.
\end{aligned} \tag{3.2}$$

Since the variance of the y_i 's, as well as that of U_{GUESSV}^A and U_{GUESSW}^A , does not depend on \mathbf{x} , Alice's best response to \mathcal{B} is

$$\mathbf{x}^* = \left([\mathbb{E}_1(\mathcal{B})]_\epsilon, \dots, [\mathbb{E}_n(\mathcal{B})]_\epsilon \right)$$

where $[\cdot]_\epsilon$ denotes the rounding to the closest ϵ integer multiplication. \mathbf{x}^* yields a payoff of at least

$$U_{\text{IMITATION}}^A(\mathbf{x}^*, \mathcal{B}) \geq -\frac{\epsilon^2}{4} - \|\text{Var}(\mathcal{B})\|_2^2.$$

Note that in every ϵ^4 -ANE Alice assigns a probability of at most $1 - \epsilon^2$ to actions that are ϵ^2 -far from optimal. By Equation (3.2) this implies that the probability that Alice choosing a vector \mathbf{x} that satisfies $\|\mathbf{x} - \mathbb{E}(\mathcal{B})\|_2^2 \geq \epsilon^2 + \frac{\epsilon^2}{4}$ is at most ϵ^2 . ■

Lemma 3.4 In every ϵ^4 -ANE $(\mathcal{A}; \mathcal{B})$, if the first m -tuple of coordinates of $\mathbb{E}_{\mathbf{y} \sim \mathcal{B}}[\mathbf{y}]$ is $6h$ -close to the binary encoding $E(v)$ of a vertex v , then

$$v^B = v \quad \beta_v^B = \left(\beta_v^T, \beta_v^S, \beta_v^P \right) \tag{3.3}$$

with probability of at least $1 - O(\epsilon^4)$ (where the probability is taken over \mathcal{B}).

Proof By Lemma 3.3 and the triangle inequality, with probability of at least $1 - \epsilon^2$, the first m -tuple of \mathbf{x} is $O(h)$ -close to $E(v)$. Rounding to $R_v(\mathbf{x}) \in \{0, 1\}^m$ can at most double the distance to $E(v)$ in each coordinate. Therefore, the Hamming distance of $R_v(\mathbf{x})$ and $E(v)$ is $O(h)$. Hence $R_v(\mathbf{x})$ is correctly decoded as $D_v(\mathbf{x}) = v$, with probability of at least $1 - \epsilon^2$.

Since v^B, β_v^B do not affect $U_{\text{BROUWER}}^B + U_{\text{GUESSW}}^B$, Bob's utility from guessing $v^B = v$, and $\beta_v^B = (\beta_v^T, \beta_v^S, \beta_v^P)$ is at least $1 - \epsilon^2$, whereas his utility from guessing any

other guess is at most ϵ^2 . Therefore, Bob assigns probability at least $1 - \epsilon^4/(1 - 2\epsilon^2)$ to actions that satisfy (3.3). ■

A similar lemma holds for the second m -tuple of \mathbf{x} and the vertex w :

Lemma 3.5 In every ϵ^4 -ANE $(\mathcal{A}; \mathcal{B})$, if the second m -tuple of coordinates of $\mathbb{E}_{\mathbf{y} \sim \mathcal{B}}[\mathbf{y}]$ is $6h$ -close to the binary encoding $E(w)$ of a vertex w , then

$$w^B = w \quad \beta_w^B = (\beta_w^T, \beta_w^S, \beta_w^P)$$

with probability of at least $1 - O(\epsilon^4)$ (where the probability is taken over \mathcal{B}).

Since Alice receives the correct v^B and β^B , we also have:

Lemma 3.6 In every ϵ^4 -ANE $(\mathcal{A}; \mathcal{B})$, if the first m -tuple of coordinates of $\mathbb{E}_{\mathbf{y} \sim \mathcal{B}}[\mathbf{y}]$ is $6h$ -close to the binary encoding $E(v)$ of a vertex v , then

$$I_v^A = (\alpha_v^T(\beta_v^T), \alpha_v^S(\beta_v^S), \alpha_v^P(\beta_v^P))$$

with probability $1 - O(\epsilon^4)$ (where the probability is taken over \mathcal{A} and \mathcal{B}).

Proof Follows immediately from Lemma 3.4 and the fact that I_v^A does not affect $U_{\text{IMITATION}}^A + U_{\text{GUESSW}}^A$. ■

A similar lemma holds for the second m -tuple of \mathbf{x} and the vertex w :

Lemma 3.7 In every ϵ^4 -ANE $(\mathcal{A}; \mathcal{B})$, if the second m -tuple of coordinates of $\mathbb{E}_{\mathbf{y} \sim \mathcal{B}}[\mathbf{y}]$ is $6h$ -close to the binary encoding $E(w)$ of a vertex w , then

$$I_w^A = (\alpha_w^T(\beta_w^T), \alpha_w^S(\beta_w^S), \alpha_w^P(\beta_w^P))$$

with probability $1 - O(\epsilon^4)$ (where the probability is taken over \mathcal{A} and \mathcal{B}).

Lemma 3.8 In every ϵ^4 -ANE $(\mathcal{A}; \mathcal{B})$, $f_{I_v^A, I_w^A}(\mathbf{x}) = f(\mathbf{x})$ with probability $1 - O(\epsilon^2)$.

Proof Follows immediately from Lemmas 3.6 and 3.7 and the “locality” condition in Proposition 3.1. ■

The following corollary completes the analysis of the 2-player game.

Corollary 3.5 In every ϵ^4 -ANE $(\mathcal{A}; \mathcal{B})$, $\|\mathbb{E}_{\mathbf{x}' \sim \mathcal{A}}[\mathbf{x}'] - f(\mathbb{E}_{\mathbf{x}' \sim \mathcal{A}}[\mathbf{x}'])\|_2 < \delta$.

Proof We recall that in Lemma 3.3 we have proved that

$$\|\mathbf{x} - \mathbb{E}_{\mathbf{y} \sim \mathcal{B}}[\mathbf{y}]\|_2^2 = O(\epsilon^2) \tag{3.4}$$

with probability $1 - O(\epsilon^2)$. This also implies that \mathbf{x} is, with high probability, close to its expectation:

$$\begin{aligned} \left\| \mathbf{x} - \mathbb{E}_{\mathbf{x}' \sim \mathcal{A}}[\mathbf{x}'] \right\|_2^2 &\leq \left(\left\| \mathbf{x} - \mathbb{E}_{\mathbf{y} \sim \mathcal{B}}[\mathbf{y}] \right\|_2 + \left\| \mathbb{E}_{\mathbf{x}' \sim \mathcal{A}}[\mathbf{x}'] - \mathbb{E}_{\mathbf{y} \sim \mathcal{B}}[\mathbf{y}] \right\|_2 \right)^2 \\ &\leq 2 \left\| \mathbf{x} - \mathbb{E}_{\mathbf{y} \sim \mathcal{B}}[\mathbf{y}] \right\|_2^2 + 2 \left\| \mathbb{E}_{\mathbf{x}' \sim \mathcal{A}}[\mathbf{x}'] - \mathbb{E}_{\mathbf{y} \sim \mathcal{B}}[\mathbf{y}] \right\|_2^2 \\ &\leq 2 \left\| \mathbf{x} - \mathbb{E}_{\mathbf{y} \sim \mathcal{B}}[\mathbf{y}] \right\|_2^2 + 2 \mathbb{E}_{\mathbf{x}' \sim \mathcal{A}} \left[\left\| \mathbf{x}' - \mathbb{E}_{\mathbf{y} \sim \mathcal{B}}[\mathbf{y}] \right\|_2^2 \right] \\ &= O(\epsilon^2), \end{aligned} \tag{3.5}$$

with probability $1 - O(\epsilon^2)$. Where the first inequality follows from the triangle inequality, the second follows from the Arithmetic-Mean Geometric-Mean inequality (AM-GM inequality), the third follows from convexity of $\|\cdot\|_2^2$, and the last follows from Lemma 3.3.

Using that f is $O(1)$ -Lipschitz together with Equation (3.5), we get that

$$\left\| f(\mathbf{x}) - f(\mathbb{E}_{\mathbf{x}' \sim \mathcal{A}}[\mathbf{x}']) \right\|_2^2 = O(\epsilon^2) \tag{3.6}$$

with probability $1 - O(\epsilon^2)$.

By Lemma 3.8 we know that $f_{I_v^A, I_w^A}(\mathbf{x}) = f(\mathbf{x})$ with probability $1 - O(\epsilon^2)$, which implies that

$$\left\| \mathbb{E}_{\mathbf{x}' \sim \mathcal{A}}[f_{I_v^A, I_w^A}(\mathbf{x}')] - \mathbb{E}_{\mathbf{x}' \sim \mathcal{A}}[f(\mathbf{x}')] \right\|_2^2 = O(\epsilon^2). \tag{3.7}$$

Using similar arguments to those of Lemma 3.3 we can show that

$$\left\| \mathbf{y} - \mathbb{E}_{\mathbf{x}' \sim \mathcal{A}}[f_{I_v^A, I_w^A}(\mathbf{x}')] \right\|_2^2 = O(\epsilon^2) \tag{3.8}$$

with probability $1 - O(\epsilon^2)$. As in the derivation of Equation (3.5), this implies:

$$\left\| \mathbf{y} - \mathbb{E}_{\mathbf{y}' \sim \mathcal{B}}[\mathbf{y}'] \right\|_2^2 = O(\epsilon^2) \tag{3.9}$$

with probability $1 - O(\epsilon^2)$.

With probability $1 - O(\epsilon^2)$, Inequalities (3.5), (3.4), (3.9), (3.8), (3.7), (3.6) hold simultaneously. In such a case, by the triangle inequality and by applying the inequalities in the exact order above, we have

$$\left\| \mathbb{E}_{\mathbf{x}' \sim \mathcal{A}}[\mathbf{x}'] - f(\mathbb{E}_{\mathbf{x}' \sim \mathcal{A}}[\mathbf{x}']) \right\|_2^2 = O(\epsilon^2) < \delta^2. \tag{3.10}$$

■

Proof Proof of Theorem 3.1. Any communication protocol that solves the ϵ^4 -Nash equilibrium problem in games of size $N \times N$ for $N = 2^{\Theta(n)}$ induces a communication protocol for the problem SIMULATION END-OF-THE-LINE: Alice constructs her utility in the above presented game using her private information of the α s, Bob constructs his utility using the β s. They implement the communication protocol to find an ϵ^4 -Nash equilibrium, and then both of them know $\mathbb{E}_{\mathbf{x} \sim \mathcal{A}}[\mathbf{x}]$, which is a δ -approximate fixed point of f (by Corollary 3.5). Using D_v , they decode the vertex v^* and they know the first coordinate of v^* .

Using Corollary 3.4, we deduce that the communication complexity of ϵ^4 -Nash equilibrium in games of size $2^{\Theta(n)} \times 2^{\Theta(n)}$ is at least $2^{\Omega(n)}$. ■

3.5.5 n -player Game

Theorem 3.4 Theorem 3.2, restated. There exists a constant $\epsilon > 0$ such that the communication complexity of (ϵ, ϵ) -weak approximate Nash equilibrium in n -player binary-action games is at least $2^{\epsilon n}$.

The proof follows similar lines to those in the proof of Theorem 3.1. Rather than two players whose actions correspond to $\Theta(n)$ -long vectors, we have a player for each bit of (an encoding of) those vectors. We construct a game with $8m'$ -players for $m' = \Theta(n)$ such that Alice holds the utility function of (the first) $3m'$ players, Bob holds the utilities of (the last) $5m'$ players, Alice's players' utilities depend only on the α s, Bob's utilities depend only on the β s, and every $(\epsilon^5/82, \epsilon^5/82)$ -weak approximate Nash equilibrium corresponds to a δ -fixed point of the function f from Proposition 3.1.

Players and Actions. In subsection 3.5.4 we used error-correcting code to encode vertices that are deduced from \mathbf{x} and \mathbf{y} . Here, since we consider *weak* approximate equilibria, we should add additional encodings for $I_v^A, I_v^B, v^B, w^B, \beta_v^B$ and β_w^B . Since we want to use the same number of players for each of the above objects, it will be convenient to encode them in the same space $\{0, 1\}^{m'}$. We let the following be encoding functions of binary error-correcting codes with constant (relative) distance:

- $E_I : \{0, 1\}^3 \rightarrow \{0, 1\}^{m'}$.
- $E_u : \{0, 1\}^{2n+\log(n+1)} \rightarrow \{0, 1\}^{m'}$.
- $E_\beta : \{0, 1\}^{3 \log M} \rightarrow \{0, 1\}^{m'}$ (note that $3 \log M = \Theta(n)$).

Let E and m denote encoding function and block length of the error-correcting code from Subsection 3.5.3, i.e.:

- $E : \{0, 1\}^{2n+\log(n+1)} \rightarrow \{0, 1\}^m$.

For vectors $\mathbf{x}, \mathbf{y} \in [-1, 2]^{4m}$, we use $(\frac{3}{\epsilon} - 1)$ bits to encode each continuous coordinate (up to precision ϵ) in unary encoding. We choose m' such that $m' = 4(\frac{3}{\epsilon} - 1)m$, so the encoding of each of \mathbf{x}, \mathbf{y} also takes m' bits. (For E_β , we must also have $m' > 3 \log M$.) Here and henceforth, ϵ is a sufficiently small constant, satisfying $\epsilon = \Theta(\delta)$.

Instead of having a single player, Alice, with actions $(\mathbf{x}, I_v^A, I_w^A) \in \{-1, -1 + \epsilon, \dots, 2 - \epsilon, 2\}^m \times \{0, 1\}^3 \times \{0, 1\}^3$, we replace her by $3m'$ players with binary actions. We have three types of Alice players:

- **x-type players.** Player \mathbf{x}_j^i chooses one of the actions $a_j^i \in \{j, j + \epsilon\}$ for every $i \in [4m]$ and $j \in \{-1, -1 + \epsilon, \dots, 2 - 2\epsilon, 2 - \epsilon\}$. Note that the total number of x-type players is $4m(\frac{3}{\epsilon} - 1) = m'$.
- **I^v -type players.** Player I_i^v chooses a bit 0 or 1 for every $i \in [m']$.
- **I^w -type players.** Similarly player I_i^w chooses a bit 0 or 1 for every $i \in [m']$.

In the communication problem, we assume that Alice knows the utilities of all the above players.

Instead of having a single player, Bob, with actions $(\mathbf{y}, v^B, w^B, \beta_v^B, \beta_w^B) \in \{-1, -1 + \epsilon, \dots, 2 - \epsilon, 2\}^m \times \{0, 1\}^{2n+\log(n+1)} \times \{0, 1\}^{2n+\log(n+1)} \times [M]^3 \times [M]^3$, we replace him by $5m'$ players with binary actions. We have five types of players:

- **y-type players.** Player \mathbf{y}_j^i chooses one of the actions $b_j^i \in \{j, j + \epsilon\}$ for every $i \in [4m]$ and $j \in \{-1, -1 + \epsilon, \dots, 2 - 2\epsilon, 2 - \epsilon\}$.
- **v -type players.** Player v_i chooses a bit 0 or 1 for every $i \in [m']$.
- **w -type players.** Similarly, player w_i chooses a bit 0 or 1 for every $i \in [m']$.
- **β^v -type players.** Player β_i^v chooses a bit 0 or 1 for every $i \in [m']$.
- **β^w -type players.** Similarly player β_i^w chooses a bit 0 or 1 for every $i \in [m']$.

In the communication problem, we assume that Bob knows the utilities of all the above players.

Utilities. Before getting to the description of the utilities we define the notions of *realized number* and *realized point* by a set of players. For every $i \in [m]$, for simplicity of notation we add a dummy player \mathbf{x}_2^i who has a single action $a_2^i = 2$. Given an action profile $a^i = (a_{-1}^i, a_{-1+\epsilon}^i, \dots, a_2^i)$ of the players $\{\mathbf{x}_j^i\}_j$, the *realized number*

$r(a^i) \in [-1, 2]$ is defined to be the minimal j such that $a_j^i = j$. Note that $r(a_i)$ is well defined because the last player x_2^i plays 2. Given an action profile $a = (a_j^i)_{i,j}$ of all x -type players, we denote by $r(a) = (r(a_i))_i \in [-1, 2]^m$ the realized point. We similarly define the realized point of y -type players.

The utilities are defined similarly to the two-player case with the following differences:

1. x -type/ y -type players' utilities are defined with respect to the *realized* points of the opponents. In addition, the player that is responsible to the i -th coordinate of the point pays the distance from the i -th coordinate of the opponent's point (or the i -th coordinate of the f operation of the opponent's point).
2. For all other types, the i -th player chooses the value of the i -th bit in the (alleged) codeword in $\{0, 1\}^{m'}$.

Formally the payoffs are defined as follows:

- For x -type players, $U^{x_j^i}(a_j^i; b^i) \triangleq -|a_j^i - r(b^i)|^2$, where we recall that player x_j^i is allowed to choose only $a_{i,j} = j$ or $a_{i,j} = j + \epsilon$, and b^i is the profile of action played by players $\{y_j^i\}_j$.
- For a v -type player v_i , we define $U^{v_i}(v_i; a) = 1$ iff he announces the bit $[E_u(D_v(r(a)))_i]$ (where the decoding function D_v is as defined in Subsection 3.5.4.1). Otherwise, $U^{v_i}(v_i; a) = 0$. Namely, the i -th player tries to guess the i -th coordinate of the encoded vector $E_u(v) \in \{0, 1\}^{m'}$, where v is computed using the decoding operation D_v on the realized point $r(a) \in [-1, 2]^{4m}$. We similarly define the utility of a w -type player.
- For a β^v -type player β_i^v , we define $U^{\beta_i^v}(\beta_i^v; a) = 1$ iff he announces the bit $[E_\beta(\beta_{D_v(r(a))}^S)_i]$. Namely, the i -th player tries to guess the i -th coordinate of the encoded vector $E_\beta(\beta_v^S)$, where v , as in the previous bullet, is computed using decoding. We similarly define the utilities of β^w -type players.
- For a I^v -type player I_i^v , we define $U^{I_i^v}(I_i^v, \beta^v) = 1$ iff she announces the bit $[E_u(\alpha_v^T([\bar{\beta}]_1), \alpha_v^S([\bar{\beta}]_2), \alpha_v^P([\bar{\beta}]_3))]_i$ where \bar{v} is the decoded vertex announced by v -type players and $\bar{\beta}$ is the decoded vector of indexes announced by β^v -type players. We similarly define the utilities of I^w -type players.
- For y -type players, $U_{y_j^i} = -|b_j^i - f_{\bar{I}^v, \bar{I}^w}(r(a))|^2$, where \bar{I}^v and \bar{I}^w are the decoding of the vertices announced by I^v -type and I^w -type players. We recall that the function f_{I^v, I^w} is defined in Proposition 3.1.

3.5.5.1 Analysis of Game

We analyze $(\bar{\epsilon}, \bar{\epsilon})$ -weak approximate equilibria for $\bar{\epsilon} = \epsilon^5/82$. The analysis of the game follows the same sequence of lemmas as the analysis in the two-player case (Subsection 3.5.4.2). The analogue of Lemma 3.3 is the following.

Lemma 3.9 In every $(\bar{\epsilon}, \bar{\epsilon})$ -weak approximate equilibrium $(\mathcal{A}, \mathcal{B})$, the realized point by the x-type players $r(a)$ satisfies

$$\left\| r(a) - \mathbb{E}_{b \sim \mathcal{B}}[r(b)] \right\|_2^2 \leq \epsilon^2 \quad (3.11)$$

with high probability⁶ (the probability is over the mixed strategy of the x-type players).

Proof We say that player \mathbf{x}_j^i 's action j is *wrong* if $\mathbb{E}_{b^i \sim \mathcal{B}}[r(b^i)] \geq j + \epsilon$; similarly, we say that action $j + \epsilon$ is *wrong* if $\mathbb{E}_{b^i \sim \mathcal{B}}[r(b^i)] \leq j$. Note that if for some coordinate i , no player \mathbf{x}_j^i plays a wrong action, then the realized number $r_i(a_i)$ is ϵ -close to $\mathbb{E}_{b^i \sim \mathcal{B}}[r(b^i)]$. We show that indeed in an $(\bar{\epsilon}, \bar{\epsilon})$ -weak approximate equilibrium we will have many such coordinates i .

Recall that player \mathbf{x}_j^i 's utility when she plays j is given by

$$\begin{aligned} u(j) &\triangleq \mathbb{E}_{b^i \sim \mathcal{B}}[U^{\mathbf{x}_j^i}(j; b^i)] \\ &= \mathbb{E}_{b^i \sim \mathcal{B}}[-|j - r(b^i)|^2] \\ &= -\left|j - \mathbb{E}_{b^i \sim \mathcal{B}}[r(b^i)]\right|^2 - \text{Var}_{b^i \sim \mathcal{B}}[r(b^i)]. \end{aligned}$$

Similarly, when she plays $j + \epsilon$ her utility is given by

$$\begin{aligned} u(j + \epsilon) &\triangleq \mathbb{E}_{b^i \sim \mathcal{B}}[U^{\mathbf{x}_j^i}(j + \epsilon; b^i)] \\ &= \mathbb{E}_{b^i \sim \mathcal{B}}[-|j + \epsilon - r(b^i)|^2] \\ &= -\left|j + \epsilon - \mathbb{E}_{b^i \sim \mathcal{B}}[r(b^i)]\right|^2 - \text{Var}_{b^i \sim \mathcal{B}}[r(b^i)]. \end{aligned}$$

6. Here and throughout this section, we use “with high probability” to mean with probability approaching 1 as n grows (in fact, with an exponential dependence); in particular, the probability is approaching 1 faster than any polynomial in ϵ .

When j is wrong (i.e., $\mathbb{E}_{b^i \sim \mathcal{B}}[r(b^i)] \geq j + \epsilon$), the difference in the utilities $u(j + \epsilon) - u(j)$ is given by

$$\begin{aligned} u(j + \epsilon) - u(j) &= -\left(\mathbb{E}_{b^i \sim \mathcal{B}}[r(b^i)] - j - \epsilon\right)^2 + \left(\mathbb{E}_{b^i \sim \mathcal{B}}[r(b^i)] - j\right)^2 \\ &= \left(2\mathbb{E}_{b^i \sim \mathcal{B}}[r(b^i)] - 2j - \epsilon\right) \epsilon \geq \epsilon^2. \end{aligned}$$

When $j + \epsilon$ is wrong ($\mathbb{E}_{b^i \sim \mathcal{B}}[r(b^i)] \leq j$), the difference in the utilities $u(j) - u(j + \epsilon)$ is given by

$$\begin{aligned} u(j) - u(j + \epsilon) &= -\left(j - \mathbb{E}_{b^i \sim \mathcal{B}}[r(b^i)]\right)^2 + \left(j + \epsilon - \mathbb{E}_{b^i \sim \mathcal{B}}[r(b^i)]\right)^2 \\ &= \left(2j - 2\mathbb{E}_{b^i \sim \mathcal{B}}[r(b^i)] + \epsilon\right) \epsilon \geq \epsilon^2. \end{aligned}$$

Therefore, player x_j^i can always increase her payoff by at least ϵ^2 by deviating from a wrong action. Note that if player x_j^i is $\bar{\epsilon}$ -best replying, she assigns a probability of at most $\bar{\epsilon}/\epsilon^2$ to a wrong action. In addition, the fraction of x -type players that are not $\bar{\epsilon}$ -best replying is at most $8\bar{\epsilon}$ (because we have 8 types of players of equal cardinality). Therefore, in the expected fraction of x -type players, playing a wrong action is at most $8\bar{\epsilon} + 2\bar{\epsilon}/\epsilon^2 < 2.5\bar{\epsilon}/\epsilon^2$. Therefore, with high probability over x -type players' mixed strategies, at most a $3\bar{\epsilon}/\epsilon^2$ -fraction play a wrong action (e.g., by Chernoff bound). Therefore the fraction of coordinates $i \in [4m]$ where at least one player x_j^i plays a wrong action is at most $9\bar{\epsilon}/\epsilon^3$ (because we have $3/\epsilon$ players in each coordinate). So in a $(1 - 9\bar{\epsilon}/\epsilon^3)$ fraction of coordinates, we have $|r_i(a_i) - \mathbb{E}_{b^i \sim \mathcal{B}}[r(b^i)]| \leq \epsilon$, which implies

$$\begin{aligned} \left\|r(a) - \mathbb{E}_{b \sim \mathcal{B}}[r(b)]\right\|_2^2 &= \frac{1}{4m} \sum_i \left|r(a_i) - \mathbb{E}_{b \sim \mathcal{B}}[r(b^i)]\right|^2 \\ &\leq \left(1 - \frac{9\bar{\epsilon}}{\epsilon^3}\right) \epsilon^2 + \frac{9\bar{\epsilon}}{\epsilon^3} 3^2 < \frac{82\bar{\epsilon}}{\epsilon^3} = \epsilon^2. \quad \blacksquare \end{aligned}$$

The analogue of Lemma 3.4 is the following.

Lemma 3.10 In every $(\bar{\epsilon}, \bar{\epsilon})$ -weak approximate equilibrium $(\mathcal{A}, \mathcal{B})$, if the first m -tuple of coordinates of $\mathbb{E}_{b \sim \mathcal{B}}[r(b)]$ is $6h$ -close to the binary encoding $E(v)$ of a vertex v , then

1. The decoding of the action profile of the v -type players is v with probability $1 - o(\epsilon)$.
2. The decoding of the action profile of the β^v -type players is $(\beta_v^T, \beta_v^S, \beta_v^P)$ with probability $1 - o(\epsilon)$.

Proof Whenever (3.11) holds, $D_v(r(a)) = v$. In particular, for each $i \in [m']$, $[E_u(D_v(r(a)))]_i = [E_u(v)]_i$ with high probability. Therefore, by playing the action $[E_u(v)]_i$ player v_i has expected utility of $1 - o(1)$, whereas by playing the action $1 - [E_u(v)]_i$ his expected utility is $o(1)$.

Every player that is $\bar{\epsilon}$ -best replying assigns probability of at least $1 - O(\bar{\epsilon})$ to the correct bit. In addition, we have at most an $8\bar{\epsilon}$ fraction of v -type players who are not $\bar{\epsilon}$ -best replying (because we have 8 types of players of equal cardinality). Therefore the expected fraction of v -type players who play the wrong bit is $O(\bar{\epsilon})$. By Chernoff bound, it also holds that with high probability at most an $O(\bar{\epsilon})$ -fraction of v -type players play the wrong bit. Whenever this is the case, v is indeed decoded correctly.

We similarly prove the second claim in the lemma for β^v -type players. \blacksquare

In a similar way we can show that analogues of Lemmas 3.5, 3.6, 3.7, and 3.8 hold for the n -player game. In particular,

Lemma 3.11 In every $(\bar{\epsilon}, \bar{\epsilon})$ -weak approximate equilibrium $(\mathcal{A}, \mathcal{B})$, $f_{I_v^A, I_w^A}(\mathbf{x}) = f(\mathbf{x})$ with high probability.

Now we get to the analogue of the last Corollary 3.5.

Corollary 3.6 In every $(\bar{\epsilon}, \bar{\epsilon})$ -weak approximate equilibrium $(\mathcal{A}, \mathcal{B})$, the expectation of the realized point $\mathbb{E}_{a \sim \mathcal{A}}[r(a)]$ is a δ -approximate equilibrium of f ; i.e.,

$$\left\| \mathbb{E}_{a \sim \mathcal{A}}[r(a)] - f(\mathbb{E}_{a \sim \mathcal{A}}[r(a)]) \right\|_2 \leq \delta.$$

Proof The proof is similar to the proof of Corollary 3.5. We recall that in Lemma 3.9 we have proved that

$$\left\| r(a) - \mathbb{E}_{b \sim \mathcal{B}}[r(b)] \right\|_2^2 \leq \epsilon^2 \tag{3.12}$$

with high probability. This, in particular, implies that $r(a)$ is, with high probability, close to its expectation:

$$\left\| r(a) - \mathbb{E}_{a' \sim \mathcal{A}}[r(a')] \right\|_2^2 \tag{3.13}$$

$$\begin{aligned} &\leq 2 \left\| r(a) - \mathbb{E}_{b \sim \mathcal{B}}[r(b)] \right\|_2^2 + 2 \left\| \mathbb{E}_{a' \sim \mathcal{A}}[r(a')] - \mathbb{E}_{b \sim \mathcal{B}}[r(b)] \right\|_2^2 \\ &\leq 2 \left\| r(a) - \mathbb{E}_{b \sim \mathcal{B}}[r(b)] \right\|_2^2 + 2 \mathbb{E}_{a' \sim \mathcal{A}} \left[\left\| r(a') - \mathbb{E}_{b \sim \mathcal{B}}[r(b)] \right\|_2^2 \right] \\ &= O(\epsilon^2), \end{aligned} \tag{3.14}$$

with high probability. Where the first inequality follows from the triangle inequality, the second follows from convexity, and the last is Lemma 3.9.

Using the $O(1)$ -Lipschitzness of f we deduce that

$$\left\| f(r(a)) - f(\mathbb{E}_{a' \sim \mathcal{A}}[r(a')]) \right\|_2^2 = O(\epsilon^2) \quad (3.15)$$

with high probability.

Using similar arguments to those of Lemma 3.9, we can show that

$$\left\| r(b) - \mathbb{E}_{a' \sim \mathcal{A}}[f_{\overline{I^v}, \overline{I^w}}(r(a'))] \right\|_2^2 = O(\epsilon^2) \quad (3.16)$$

with high probability, where we recall that $\overline{I^v}, \overline{I^w}$ denote the decoded line information of the action profile played by the I^v and I^w -type players. By an analogous argument to (3.13),

$$\left\| r(b) - \mathbb{E}_{b' \sim \mathcal{B}}[r(b')] \right\|_2^2 = O(\epsilon^2) \quad (3.17)$$

with high probability.

By Lemma 3.11,

$$\left\| \mathbb{E}_{a' \sim \mathcal{A}}[f_{\overline{I^v}, \overline{I^w}}(r(a'))] - \mathbb{E}_{a' \sim \mathcal{A}}[f(r(a'))] \right\|_2^2 = O(\epsilon^2). \quad (3.18)$$

By Equations (3.13), (3.12), (3.17), (3.16), (3.18), (3.15) (applied exactly in this order) and the triangle inequality, we get

$$\left\| \mathbb{E}_{a' \sim \mathcal{A}}[r(a')] - f(\mathbb{E}_{a' \sim \mathcal{A}}[r(a')]) \right\|_2^2 = O(\epsilon^2) < \delta^2. \quad (3.19)$$

■

Proof Proof of Theorem 3.2. Any communication protocol that solves the $(\epsilon^5/82, \epsilon^5/82)$ -weak approximate Nash equilibrium problem in $\Theta(n)$ -player games with binary actions induces a communication protocol for the problem SIMULATION END-OF-THE-LINE: Alice constructs the utilities of her players using her private information of the α s, Bob constructs his utility using the β s. They implement the communication protocol to find an $(\epsilon^5/82, \epsilon^5/82)$ -weak approximate Nash equilibrium, and then both of them know $\mathbb{E}_{a' \sim \mathcal{A}}[r(a)]$, which is a δ -approximate fixed point of f (by Corollary 3.6). Finally, they round and decode the approximate fixed point to recover the end of the line.

Using Corollary 3.4, we deduce that the communication complexity of the $(\epsilon^5/82, \epsilon^5/82)$ -weak approximate Nash equilibrium problem in $\Theta(n)$ -player games with binary actions is at least $2^{\Omega(n)}$. ■

3.6 An Open Problem: Correlated Equilibria in 2-Player Games

As mentioned in Subsection 3.3, it is known that for n -player, $O(1)$ -action games, even *exact* correlated equilibrium can be found with $\text{poly}(n)$ deterministic communication complexity (see [Hart and Mansour 2010, Jiang and Leyton-Brown 2015, Papadimitriou and Roughgarden 2008]).

In two-player $N \times N$ games, for approximate correlated equilibrium with constant value of approximation, to the best of our knowledge, no non-trivial results are known (neither positive nor negative). Does a $\text{polylog}(N)$ communication protocol for approximate correlated equilibrium exist? Is there a $\text{poly}(N)$ communication lower bound? For small values of approximation, recently Ganor and Karthik [2017] have shown that $1/N$ -correlated equilibrium requires $\text{poly}(N)$ communication.

Brouwer's Fixed Point

Brouwer's fixed point theorem guarantees that any continuous function $f : [0, 1]^n \rightarrow [0, 1]^n$ has a fixed point, i.e., a point $\mathbf{x} \in [0, 1]^n$ such that $f(\mathbf{x}) = \mathbf{x}$. In this section we construct continuous functions $f : [0, 1]^n \rightarrow [0, 1]^n$ where even finding an approximate fixed point (i.e., \mathbf{x} such that $\mathbf{x} \approx f(\mathbf{x})$) is hard.

In particular, we reduce the END-OF-A-LINE problem to the problem of finding an approximate fixed point. Our reductions require only local, computationally efficient, black-box access to the END-OF-A-LINE instance. Thus they immediately apply both to the computational setting (defined in Section 2.2) and the query-complexity model (which is the focus of Chapter 3).

We present two reductions. We begin with an easier construction, where we think of the error $\mathbf{x} - f(\mathbf{x})$ in terms of ℓ_∞ -norm; it is a good starting point for introducing the main ideas, and the special structure of the construction is also useful for the results in Chapter 6. We then add error-correcting codes to obtain a stronger construction for ℓ_2 -norm.

4.1 BROUWER with ℓ_∞

Below we state and prove the hardness of finding an approximate fixed point in ℓ_∞ -norm. For the results in Chapter 6 we require a stronger characterization of the construction, which we omit from the theorem statement for simplicity (see Fact 4.1 for details).

Theorem 4.1 ℓ_∞ BROUWER. There exists a constant $\epsilon > 0$ such that the following holds. Given (local, black-box access to) an instance of END-OF-A-LINE, we can construct a computationally efficient, continuous function $f : [0, 1]^{O(n)} \rightarrow [0, 1]^{O(n)}$, such that:

- f is $O(1)$ -Lipschitz in ℓ_∞ -norm; and
- given \mathbf{x} for which $\|f(\mathbf{x}) - \mathbf{x}\|_\infty < \epsilon$, we can efficiently reconstruct a solution to the END-OF-A-LINE instance.

Proof In the first step (Subsection 4.1.2), we embed the END-OF-A-LINE problem (over $\{0, 1\}^n$) as a collection H of vertex-disjoint paths over the $(2n + 1)$ -dimensional hypercube graph. Given H , our second step (Subsection 4.1.3) is to construct a continuous mapping $f: [0, 1]^{2n+2} \rightarrow [0, 1]^{2n+2}$ whose fixed points correspond to ends of paths in H . This step generalizes a construction of Hirsch et al. [1989] for embedding a single path. ■

4.1.1 Preliminaries: ℓ_∞ -Norm Geometry

Throughout this section, we work with the ℓ_∞ -norm. This has some implications that may contradict our geometric intuition. For example: in an ℓ_∞ -norm world, a circle is a square.

ℓ_∞ -norm interpolation. Given coordinates $x, y \geq 0$, we define the ℓ_∞ -norm angle¹ that point (x, y) forms with the x -axis (in the xy -plane) as

$$\theta_\infty(x, y) = \frac{y}{x + y}.$$

The ℓ_∞ -norm angle is useful for interpolation. Given the values of $f: [0, 1]^n \rightarrow [0, 1]^n$ on two neighboring facets of the hypercube, we can extend f to all points of the hypercube by *angular interpolation*: interpolate according to the ℓ_∞ -norm angle $\theta_\infty(x_i, x_j)$ where x_i and x_j are the respective distances from the two facets. When f is defined on two opposite facets, we can simply use *Cartesian interpolation*, which again means to interpolate according to the distance from each facet.

ℓ_∞ -norm local polar coordinates. Given a point $\mathbf{z} \in \mathbb{R}^n$, we define a new *local ℓ_∞ -norm polar coordinate* system around \mathbf{z} . Every $\mathbf{x} \in \mathbb{R}^n$ is transformed into $\langle r, \mathbf{p} \rangle_{\mathbf{z}} \in \mathbb{R} \times \mathbb{R}^n$ where $r = \|\mathbf{x} - \mathbf{z}\|$ is the ℓ_∞ -norm radius, and $\mathbf{p} = (\mathbf{x} - \mathbf{z})/r$ is the ℓ_∞ -norm unit vector that points from \mathbf{z} in the direction of \mathbf{x} .

4.1.2 Embedding the END-OF-A-LINE Graph as Paths in $\{0, 1\}^{2n+1}$

Our first step in the reduction is to embed an END-OF-A-LINE graph $G_{S,P}$ as vertex-disjoint paths on the $(2n + 1)$ -dimensional hypercube graph. We construct a collection H of vertex-disjoint paths and cycles over the $(2n + 1)$ -dimensional hypercube graph, such that there is a 1-to-1 correspondence between starting and end points of paths in H and starting and end points of lines in $G_{S,P}$.

In order to construct our embedding, we divide the $2n + 1$ coordinates as follows: the first n coordinates store the current vertex \mathbf{u} , the next n coordinates for

1. Our ℓ_∞ -norm angle was called *unit* in Hirsch et al. [1989].

the next vertex in the line, \mathbf{v} , and finally, the last coordinate b stores a compute-next vs copy bit. When $b = 0$, the path proceeds to update $\mathbf{v} \leftarrow S(\mathbf{u})$, bit-by-bit. When this update is complete, the value of b is changed to 1. Whenever $b = 1$, the path proceeds by copying $\mathbf{u} \leftarrow \mathbf{v}$ bit-by-bit, and then changes that value of b again. Finally, when $\mathbf{u} = \mathbf{v} = S(\mathbf{u})$ and $b = 0$, the path reaches an end point. For example, the edge $\mathbf{u} \rightarrow \mathbf{v}$ maps into the path:

$$(\mathbf{u}, \mathbf{u}, 0) \rightarrow \dots \rightarrow (\mathbf{u}, \mathbf{v}, 0) \rightarrow (\mathbf{u}, \mathbf{v}, 1) \rightarrow \dots \rightarrow (\mathbf{v}, \mathbf{v}, 1) \rightarrow (\mathbf{v}, \mathbf{v}, 0).$$

Notice that the paths in H do not intersect. Furthermore, given a vector in $\mathbf{p} \in \{0, 1\}^{2n+1}$, we can efficiently output whether \mathbf{p} belongs to a path in H , and if so, which are the previous and consecutive vectors in the path. Finding a starting or end point of any path in H (other than $\mathbf{0}_{2n+1}$) is therefore equivalent to finding an odd-degree vertex (other than $\mathbf{0}_n$) in $G_{S,P}$.

4.1.3 Continuous Mapping on $[0, 1]^{2n+2}$

In order to construct a hard instance of Brouwer function, we use techniques introduced by Hirsch et al. [1989]. The continuous Brouwer function is denoted by $f: [0, 1]^{2n+2} \rightarrow [0, 1]^{2n+2}$, while the associated displacement function is denoted by $g(\mathbf{x}) \triangleq f(\mathbf{x}) - \mathbf{x}$. The following lemma (proven below) completes the proof of Theorem 4.1.

Lemma 4.1 The displacement g satisfies:

1. g is $O(1)$ -Lipschitz (thus, f is also $O(1)$ -Lipschitz).
2. $\|g(\mathbf{x})\|_\infty = \Omega(1)$ for every \mathbf{x} that does not correspond to a starting or end point in H .
3. The value of g at each point \mathbf{x} can be computed (efficiently) with local, black-box access to S, P .

4.1.3.1 Overview of the Construction

The domain of f is the $2n + 2$ -dimensional (solid) hypercube. The hypercube is divided into subcubes, of side length h (we fix $h = 1/4$). We define f separately on each subcube such that it agrees on the intersections (no interpolation is needed in this sense).

The last $((2n + 2)$ -th) dimension is special; we use “up” (resp. “down”) to refer to the positive (negative) $(2n + 2)$ -th direction. All the action takes place in the second-from-bottom $(2n + 1)$ -dimensional layer of subcubes; this layer is called the *slice*. Within the slice, we also ignore the subcubes that are near the boundary of the

hypercube (those compose the *frame*); we are left with the subcubes in the center of the slice, which we call the *picture*. We identify between the vertices of the $(2n + 1)$ -dimensional hypercube graph (over which H was defined) and the 2^{2n+1} subcubes of the picture.

The subset of subcubes into which we embed each path or cycle from H is called a *tube*. The *home subcube*, the subcube that corresponds to the $\mathbf{0}_{2n+1}$ -vertex, is special: all the flow from all subcubes that do not belong to any tube leads to this subcube.

Below we define the displacement in the following regions, and argue that it satisfies the desiderata of Lemma 4.1:

- inside the picture, but not in any tube;
- inside a tube; and
- outside the picture.

4.1.3.2 Default Displacement

Most of the slice has the same *default displacement*: directly upward, i.e., $g(\mathbf{x}) = \delta \xi_{2n+2}$, where ξ_{2n+2} is the $(2n + 2)$ -unit vector, and $\delta > 0$ is a small constant. Formally,

Fact 4.1 $g(\mathbf{x}) = \delta \xi_{2n+2}$, for every \mathbf{x} such that at least one of the following holds:

1. \mathbf{x} lies on a corner, i.e., the intersection of two or more facets of a subcube;
2. \mathbf{x} lies on an outer facet of a tube subcube, i.e., a facet other than the two facets that continue the path; or
3. \mathbf{x} lies in a subcube that does not belong to any tube.

Intuitively, Property 2 implies that *all subcubes—whether they belong to the tube or not—look the same from the outside* (except for the two facets that continue the path). In particular, the displacement on both sides of each facet is the same; so if the displacement is $O(1)$ -Lipschitz on each subcube, it is also $O(1)$ -Lipschitz on the entire hypercube.

Property 1, stating that *all corners look the same*, is key to the sampling gadgets in Chapter 6, because it liberates us from having to disambiguate the position of a point near the corners (that is, deciding exactly to which subcube it belongs).

4.1.3.3 Displacement at a Tube

The mapping is defined so that in the center of the tube, the flow goes along the direction of the path; slightly outside the center, the flow points toward the center of the tube; further away from the center, the flow goes *against* the direction of the

path; at the outer boundary of the tube, as we previously described, the flow goes upward.

We first define g on facets. Let $\langle r, \mathbf{p} \rangle_{\mathbf{z}}$ be a point on the facet centered at \mathbf{z} , and suppose that the tube enters the subcube through \mathbf{z} , advancing in the positive i -th coordinate. We define

$$g(\langle r, \mathbf{p} \rangle_{\mathbf{z}}) \triangleq \begin{cases} \delta \xi_i & r = 0 \\ -\delta \mathbf{p} & r = h/8 \\ -\delta \xi_i & r = h/4 \\ \delta \xi_{2n+2} & r = h/2. \end{cases} \quad (4.1)$$

(Recall that h is the subcube side length, and δ is some small constant.) Notice that at each r , the displacement g is $O(1)$ -Lipschitz and has magnitude $\|g(\mathbf{x})\|_\infty = \Omega(1)$ (thus satisfying the first two desiderata of Lemma 4.1).

For $r \in (0, h/8)$, interpolate between $\delta \xi_i$ and $-\delta \mathbf{p}$ (Hirsch et al. [1989] call this *radial interpolation*), and similarly for $r \in (h/8, h/4)$ and $r \in (h/4, h/2)$. See also the illustration in Figure 4.1. It is easy to see that the $O(1)$ -Lipschitz property is preserved. Notice also that ξ_i is orthogonal to \mathbf{p} and ξ_{2n+2} ; this guarantees that the interpolation does not lead to cancellation, i.e., we still have $\|g(\mathbf{x})\|_\infty = \Omega(1)$.

In the last couple of paragraphs we defined g on two facets for each subcube that belongs to the tubes; for all other points in the tubes we interpolate (angular interpolation) between those two facets: Consider a point \mathbf{x} in the tube, and assume (w.l.o.g.) that $x_i, x_j > 1/2$, and suppose that the value of $f(\cdot)$ on the $y_i = 1/2$ and $y_j = 1/2$ facets of the subcube containing \mathbf{x} is determined by (4.1). Let

$$\mathbf{x}^i = \left(\mathbf{x}_{-i,j}, \frac{1}{2}, \max\{x_i, x_j\} \right)$$

$$\mathbf{x}^j = \left(\mathbf{x}_{-i,j}, \max\{x_i, x_j\}, \frac{1}{2} \right)$$

denote the corresponding “ ℓ_∞ -norm projections” to the respective $y_i = 1/2$ and $y_j = 1/2$ facets. We set

$$g(\mathbf{x}) = \theta_\infty\left(x_i - \frac{1}{2}, x_j - \frac{1}{2}\right) g(\mathbf{x}^i) + \left(1 - \theta_\infty\left(x_i - \frac{1}{2}, x_j - \frac{1}{2}\right)\right) g(\mathbf{x}^j).$$

Notice that \mathbf{x}^i and \mathbf{x}^j are at the same distance from the respective facet centers, i.e., they correspond to the same r . For each case of (4.1), the (i, j) -components of the displacements at \mathbf{x}^i and \mathbf{x}^j are orthogonal, and for the rest of the components they are aligned. Therefore, when we interpolate between $g(\mathbf{x}^i)$ and $g(\mathbf{x}^j)$ there is again no cancellation, i.e., $\|g(\mathbf{x})\|_\infty = \Omega(\|g(\mathbf{x}^i)\|_\infty) = \Omega(1)$. Finally, recall that the displacement on each facet is $O(1)$ -Lipschitz, and the displacements agree




Figure 4.1 An illustration of the displacement on a facet between two subcubes in a tube; the direction of the path is into the paper. In the center (red), the displacement points into the paper; slightly further outside (blue), the displacement points toward the center; further outside (yellow), the displacement points out of the paper; finally, in the outer layer (green), the displacement points in the special $2n + 2$ dimension.

on the intersection of the facets. Therefore the interpolated displacement is $O(1)$ -Lipschitz over the entire subcube by a triangle-inequality argument.

The home subcube is defined using (4.1) as if the tube enters from above, i.e., coming down the $(2n + 2)$ -dimension, and exits through another facet (in one of the first $(2n + 1)$ dimensions) in the direction of the path (here again we have $\|g(\mathbf{x})\|_\infty = \Omega(1)$). For all other starting and end points, we define $g(\mathbf{x}) = \delta \mathbf{x}_{2n+2}$ on the facet opposite the one that continues the tube, and interpolate between the opposite facets using Cartesian interpolation. Notice that this gives a fixed point when the interpolation cancels the default displacement at the opposite facet, with the displacement $-\delta \xi_{2n+2}$ at the point on the tube facet that is at distance $h/8$ above the path.

4.1.3.4 Outside the Picture

For all points in the frame and below the slice, the displacement points directly upward, i.e., $g(\mathbf{x}) = \delta \xi_{2n+2}$. Moving above the slice, let $\mathbf{z}[\text{top}]$ be the point on the top

facet of the hypercube that is directly above the center of the home subcube. For all points $\langle r, \mathbf{p} \rangle_{\mathbf{z}[\text{top}]}$ on the top facet of the hypercube, define the displacement as follows:

$$g\left(\langle r, \mathbf{p} \rangle_{\mathbf{z}[\text{top}]}\right) = \begin{cases} -\delta \xi_{2n+2} & r = 0 \\ -\delta \mathbf{p} & r \geq h/8 \end{cases}$$

and interpolate for $r \in (0, h/8)$. Notice that this displacement is $O(1)$ -Lipschitz and has $\Omega(1)$ magnitude for each r , and this is preserved after interpolation.

Notice that the definition of g on the slice from the previous subsection implies that all the points in the top facet of the slice, except for the top of the home subcube, point directly upward. Let $\mathbf{z}[\text{home}]$ denote the center of the top facet of the home subcube. We therefore have that for any $\langle r, \mathbf{p} \rangle_{\mathbf{z}[\text{home}]}$ in the top facet of the slice,

$$g\left(\langle r, \mathbf{p} \rangle_{\mathbf{z}[\text{home}]}\right) = \begin{cases} -\delta \xi_{2n+2} & r = 0 \\ -\delta \mathbf{p} & r = h/8 \\ \delta \xi_{2n+2} & r \geq h/4 \end{cases}$$

where we again interpolate radially for r in $(0, h/8)$ and $(h/8, h/4)$.

Finally, to complete the construction above the slice, simply interpolate (using Cartesian interpolation) between the top facets of the slice and the hypercube. See also the illustration in Figure 4.2.




Figure 4.2 An illustration of the displacement outside the picture.

4.2 Euclidean BROUWER

Theorem 4.2 **Euclidean BROUWER.** There exist constants $\delta, h > 0$ such that the following holds. Given (local, black-box access to) an instance $I = (T, S, P)$ of MEMBERSHIP END-OF-A-LINE, we can construct a computationally efficient, continuous function $f : [0, 1]^{O(n)} \rightarrow [0, 1]^{O(n)}$, such that

- f is $O(1)$ -Lipschitz in ℓ_2 -norm; and
- given \mathbf{x} for which $\|f(\mathbf{x}) - \mathbf{x}\|_2 < \epsilon$, we can efficiently reconstruct a solution to the MEMBERSHIP END-OF-A-LINE instance.

4.2.1 Discrete Embedding of a Graph in the Euclidean Space

Each vertex v is embedded to the point $(E(v), E(v), \mathbf{0}_m, \mathbf{0}_m) \in [-1, 2]^{4m}$, which is called the *embedded vertex*.

For every edge (v, w) we define five vertices:

$$\begin{aligned} \mathbf{x}^1(v, w) &\triangleq (E(v), E(v), \mathbf{0}_m, \mathbf{0}_m) \\ \mathbf{x}^2(v, w) &\triangleq (E(v), E(v), \mathbf{1}_m, \mathbf{0}_m) \\ \mathbf{x}^3(v, w) &\triangleq (E(v), E(w), \mathbf{1}_m, \mathbf{0}_m) \\ \mathbf{x}^4(v, w) &\triangleq (E(v), E(w), \mathbf{0}_m, \mathbf{0}_m) \\ \mathbf{x}^5(v, w) &\triangleq (E(w), E(w), \mathbf{0}_m, \mathbf{0}_m). \end{aligned}$$

The vertices $\mathbf{x}^i(v, w)$ are called *Brouwer vertices*. Note that $\mathbf{x}^1(v, w)$ is the embedded vertex v , and $\mathbf{x}^5(v, w)$ is the embedded vertex w . The line that connects the points $\mathbf{x}^i(v, w)$ and $\mathbf{x}^{i+1}(v, w)$ is called a *Brouwer line segment*. The union of these four Brouwer line segments is called the *embedded edge* (v, w) .

4.2.2 The Function f

We set h to be a sufficiently small constant such that the \sqrt{h} neighborhood of any two Brouwer vertices will not intersect and such that the $3h$ neighborhood of any two Brouwer line segments will not intersect—unless they share the same common Brouwer vertex. We take δ to be a constant arbitrarily smaller than h ($\delta = h^3$ suffices). We define a displacement function $g : [-1, 2]^{4m} \rightarrow [-\delta, \delta]^{4m}$ and $f(\mathbf{x}) \triangleq \mathbf{x} + g(\mathbf{x})$. In order to satisfy properties (1)–(3) of Proposition 3.1, we should define g such that:

1. $\|g(\mathbf{x})\|_2 = \Omega(\delta)$ for every \mathbf{x} that is not $2\sqrt{h}$ -close to the Brouwer line segments of any non-trivial end or starting of a line.
2. g is $O(1)$ -Lipschitz.

3. g is defined “locally,” which will allow us to generate the class of functions $\{f_{I_1, I_2}\}$.

We think of the $4m$ coordinates as partitioned into four parts. The first m -tuple of coordinates represent the current vertex in the line; the second m -tuple represent the next vertex in the line. We think of the third m -tuple as all being equal to a single bit that switches between computing the next vertex, and copying from the second to first m -tuple. Finally, the last m coordinates represent a special default direction in which the displacement points when far from all Brouwer line segments (similarly to the single special coordinate in Hirsch et al. [1989]).

We consider a path starting at $(\mathbf{0}_{3m}, 2 \cdot \mathbf{1}_m)$, i.e., the concatenation of 0 on the first $3m$ coordinates, and 2 on the last m coordinates. The path first goes to $(\mathbf{0}_{4m})$ (in a straight line), and thereafter the last m coordinates remain constantly 0 (note that every Brouwer vertex has $\mathbf{0}_m$ in its last m -tuple). The first $3m$ coordinates follow the line according to the embedding in Subsection 3.5.3. This path corresponds to the line starting at $\mathbf{0}_{2n+1}$; for any additional line starting at vertex u , we have another path starting at $(E(u), E(u), \mathbf{0}_{2m})$.

We say that a point \mathbf{x} is in the *picture* if $\frac{1}{m} \sum_{i=3m+1}^{4m} x_i < 1/2$. We construct g separately inside and outside the picture (and make sure that the construction agrees on the hyperplane $\frac{1}{m} \sum_{i=3m+1}^{4m} x_i = 1/2$).

Truncation. In order for $g(\cdot)$ to be a displacement function, we must ensure that it never sends any points outside the hypercube, i.e., $\forall \mathbf{x} \in [-1, 2]^{4m}$, we require that also $\mathbf{x} + g(\mathbf{x}) \in [-1, 2]^{4m}$. Below, it is convenient to first define an *untruncated* displacement function $\hat{g} : [-1, 2]^{4m} \rightarrow [-\delta, \delta]^{4m}$ that is not restricted by the above condition. We then truncate each coordinate to fit in $[-1, 2]$: $[g(\mathbf{x})]_i = \max\{-1, \min\{2, x_i + [\hat{g}(\mathbf{x})]_i\}\} - x_i$. It is clear that if $\hat{g}(\cdot)$ is $(M - 1)$ -Lipschitz, then $g(\cdot)$ is M -Lipschitz. It is, however, important to make sure that the magnitude of the displacement is not compromised. Typically, some of the coordinates may need to be truncated, but we design the displacement so that most coordinates, say 99%, are not truncated. If $\hat{g}(\mathbf{x})$ has a non-negligible component in at least 5% of the coordinates, then in total $g(\mathbf{x})$ maintains a non-negligible magnitude.

4.2.2.1 Inside the Picture

The line $\mathbf{0} = v_0, v_1, \dots, v^*$ is embedded to a path in $[-1, 2]^{4m}$ that goes in straight lines through the following sequence of Brouwer vertices:

$$\begin{aligned} (\mathbf{0}_{3m}, 2 \cdot \mathbf{1}_m), (\mathbf{0}_{4m}) &= \mathbf{x}^1(v_0, v_1), \mathbf{x}^2(v_0, v_1), \dots, \mathbf{x}^5(v_0, v_1) \\ &= \mathbf{x}^1(v_1, v_2), \mathbf{x}^2(v_1, v_2), \dots, \mathbf{x}^5(P(v^*), v^*). \end{aligned}$$

Similarly, if I contains another line u, \dots, w , it is embedded as a path through:

$$(E(u), E(u), \mathbf{0}_{2m}) = \mathbf{x}^1(u, S(u)), \dots, \mathbf{x}^5(P(w), w) = (E(w), E(w), \mathbf{0}_{2m}),$$

and analogously for cycles.

Now we *cut the corners* of this path as follows: For two consecutive Brouwer vertices \mathbf{s}, \mathbf{y} in the embedded path we let $\mathbf{z}_{(\mathbf{s} \rightarrow \mathbf{y})}^1$ be the point in the Brouwer line segment $[\mathbf{s}, \mathbf{y}]$ that is exactly \sqrt{h} -far from \mathbf{s} . Similarly, $\mathbf{z}_{(\mathbf{s} \rightarrow \mathbf{y})}^2$ is the point in $[\mathbf{s}, \mathbf{y}]$ that is exactly \sqrt{h} -far from \mathbf{y} ; for three consecutive Brouwer vertices $\mathbf{s} \rightarrow \mathbf{y} \rightarrow \mathbf{t}$, the path after “cutting the corners” goes in straight lines through

$$\dots, \mathbf{z}_{(\mathbf{s} \rightarrow \mathbf{y})}^1, \mathbf{z}_{(\mathbf{s} \rightarrow \mathbf{y})}^2, \mathbf{z}_{(\mathbf{y} \rightarrow \mathbf{t})}^1, \mathbf{z}_{(\mathbf{y} \rightarrow \mathbf{t})}^2, \dots$$

instead of going through $\mathbf{s} \rightarrow \mathbf{y} \rightarrow \mathbf{t}$.

First, for all points inside the picture that are $3h$ -far from the embedded path after cutting the corners, we use the same *default displacement*, which points in the positive special direction: $\hat{g}(\mathbf{x}) = (\mathbf{0}_{3m}, \delta \cdot \mathbf{1}_m)$. Because \mathbf{x} is inside the picture, the truncated displacement $g(\mathbf{x})$ is close to $\hat{g}(\mathbf{x})$, and therefore satisfies $\|g(\mathbf{x})\|_2 = \Omega(\delta)$.

Now we define the displacement $3h$ -close to the embedded path in two regions:

1. for points that are $3h$ -close to a segment of the form $[\mathbf{z}_{(\mathbf{s} \rightarrow \mathbf{y})}^1, \mathbf{z}_{(\mathbf{s} \rightarrow \mathbf{y})}^2]$ but (approximately²) \sqrt{h} -far from both Brouwer vertices \mathbf{s}, \mathbf{y} ;
2. for the remaining points, those that are $3h$ -close to a segment of the form $[\mathbf{z}_{(\mathbf{s} \rightarrow \mathbf{y})}^2, \mathbf{z}_{(\mathbf{y} \rightarrow \mathbf{t})}^1]$ and (approximately²) \sqrt{h} -close to the Brouwer vertex \mathbf{y} .

We make sure that the definitions agree on the interface between the two regions, as well as on the interface with the points that receive the default displacement.

4.2.2.2 Close to the Path but Far from a Brouwer Vertex

On the Brouwer line segment, the displacement points in the direction of the path; at distance h from the Brouwer line segment, the displacement points in toward the Brouwer line segment; at distance $2h$ from the Brouwer line segment, the displacement points against the direction of the path; at distance $3h$, the displacement points in the default direction.

Formally, let $\sigma_{(\mathbf{s} \rightarrow \mathbf{t})}(\mathbf{x})$ denote the magnitude of the component of $\mathbf{x} - \mathbf{s}$ in the direction of line $(\mathbf{s} \rightarrow \mathbf{t})$,

2. It will be more convenient to set the threshold of points \mathbf{x} that are “far”/“close” from/to a Brouwer vertex using the expression $\sigma_{(\mathbf{s} \rightarrow \mathbf{y})}(\mathbf{x})$ that is defined below. $\sigma_{(\mathbf{s} \rightarrow \mathbf{y})}(\mathbf{x})$ is closely related to the distance of \mathbf{x} from the points \mathbf{s}, \mathbf{y} but is not precisely the distance.

$$\sigma_{(\mathbf{s} \rightarrow \mathbf{t})}(\mathbf{x}) \triangleq \frac{(\mathbf{t} - \mathbf{s})}{\|\mathbf{s} - \mathbf{t}\|_2} \cdot (\mathbf{x} - \mathbf{s}),$$

where \cdot denotes the (in-expectation) dot product. Let $\mathbf{z} = \mathbf{z}(\mathbf{x})$ be the point nearest to \mathbf{x} on the Brouwer line segment; notice that \mathbf{z} satisfies

$$\mathbf{z} = \sigma_{(\mathbf{s} \rightarrow \mathbf{t})}(\mathbf{x})\mathbf{t} + (1 - \sigma_{(\mathbf{s} \rightarrow \mathbf{t})}(\mathbf{x}))\mathbf{s}.$$

For points near the Brouwer line segment ($\|\mathbf{x} - \mathbf{z}\|_2 \leq 3h$), but far from its end points ($\sigma_{(\mathbf{s} \rightarrow \mathbf{t})}(\mathbf{x}) \in [\sqrt{h}, 1 - \sqrt{h}]$), we define the displacement:

$$\hat{g}(\mathbf{x}) \triangleq \begin{cases} \delta \frac{(\mathbf{t} - \mathbf{s})}{\|\mathbf{t} - \mathbf{s}\|_2} & \|\mathbf{x} - \mathbf{z}\|_2 = 0 \\ \delta \frac{(\mathbf{z} - \mathbf{x})}{h} & \|\mathbf{x} - \mathbf{z}\|_2 = h \\ \delta \frac{(\mathbf{s} - \mathbf{t})}{\|\mathbf{t} - \mathbf{s}\|_2} & \|\mathbf{x} - \mathbf{z}\|_2 = 2h \\ \delta(\mathbf{0}_{3m}, \mathbf{1}_m) & \|\mathbf{x} - \mathbf{z}\|_2 = 3h. \end{cases} \quad (4.2)$$

At intermediate distances from the Brouwer line segment, we interpolate: at distance $\|\mathbf{x} - \mathbf{z}\|_2 = \frac{1}{3}h$, for example, we have $\hat{g}(\mathbf{x}) = \frac{2}{3}\delta \frac{(\mathbf{t} - \mathbf{s})}{\|\mathbf{t} - \mathbf{s}\|_2} + \frac{1}{3}\delta \frac{(\mathbf{z} - \mathbf{x})}{h}$. Notice that $(\mathbf{t} - \mathbf{s})$ is orthogonal to both $(\mathbf{z} - \mathbf{x})$ and $(\mathbf{0}_{3m}, \mathbf{1}_m)$, so the interpolation does not lead to cancellation. Also, every point \mathbf{z} on the Brouwer line segment is $\Omega(1)$ -far in every coordinate from $\{-1, 2\}$, so the truncated displacement $g(\mathbf{x})$ still satisfies $\|g(\mathbf{x})\|_2 = \Omega(\delta)$. For each case in (4.2), $\hat{g}(\cdot)$ is either constant, or (in the case of $\|\mathbf{x} - \mathbf{z}\|_2 = h$) $O(\delta/h)$ -Lipschitz ($\frac{(\mathbf{z} - \mathbf{x})}{h}$ is $O(1/h)$ -Lipschitz because two ‘‘antipodal’’ points at distance $2h$ have opposite directions, both pointing parallel to the Brouwer line segment); by choice of $\delta \ll h$, it follows that $\hat{g}(\cdot)$ is in particular $O(1)$ -Lipschitz. Furthermore, notice that $\|\mathbf{x} - \mathbf{z}\|_2$ is 1-Lipschitz, so after interpolating for intermediate distances, $\hat{g}(\cdot)$ continues to be $O(1)$ -Lipschitz. Notice also that at distance $3h$ the displacement defined in (4.2) agrees with the displacements for points far from every Brouwer line segment, so Lipschitz continuity is preserved.

4.2.2.3 Close to the Path and a Brouwer Vertex

Let L_y be the line that connects the points $\mathbf{z}_{(\mathbf{s} \rightarrow \mathbf{y})}$ and $\mathbf{z}_{(\mathbf{y} \rightarrow \mathbf{t})}$. Given \mathbf{x} , we let \mathbf{z} be the closest point to \mathbf{x} on L_y .

Our goal is to interpolate between the line displacement for $(\mathbf{s} \rightarrow \mathbf{y})$ (which is defined up to $\sigma_{(\mathbf{s} \rightarrow \mathbf{y})}(\mathbf{x}) = 1 - \sqrt{h}$), and the line displacement for $(\mathbf{y} \rightarrow \mathbf{t})$ (which begins at $\sigma_{(\mathbf{y} \rightarrow \mathbf{t})}(\mathbf{x}) = \sqrt{h}$). Let $\Delta_{(\mathbf{s} \rightarrow \mathbf{y})}(\mathbf{x}) \triangleq \sigma_{(\mathbf{s} \rightarrow \mathbf{y})}(\mathbf{x}) - (1 - \sqrt{h})$, and $\Delta_{(\mathbf{y} \rightarrow \mathbf{t})}(\mathbf{x}) \triangleq \sqrt{h} - \sigma_{(\mathbf{y} \rightarrow \mathbf{t})}(\mathbf{x})$. We set our interpolation parameter $\tau = \tau(\mathbf{x}) \triangleq \frac{\Delta_{(\mathbf{y} \rightarrow \mathbf{t})}(\mathbf{x})}{\Delta_{(\mathbf{y} \rightarrow \mathbf{t})}(\mathbf{x}) + \Delta_{(\mathbf{s} \rightarrow \mathbf{y})}(\mathbf{x})}$, and set

$$\mathbf{z} \triangleq \tau \mathbf{z}_{(\mathbf{s} \rightarrow \mathbf{y})} + (1 - \tau) \mathbf{z}_{(\mathbf{y} \rightarrow \mathbf{t})}. \quad (4.3)$$

For points \mathbf{x} near \mathbf{y} such that $\Delta_{(\mathbf{s} \rightarrow \mathbf{y})}(\mathbf{x}), \Delta_{(\mathbf{y} \rightarrow \mathbf{t})}(\mathbf{x}) \geq 0$, we can now define the displacement analogously to (4.2):

$$\hat{g}(\mathbf{x}) \triangleq \begin{cases} \delta \cdot \left[\tau \frac{(\mathbf{y}-\mathbf{s})}{\|\mathbf{y}-\mathbf{s}\|_2} + (1-\tau) \frac{(\mathbf{t}-\mathbf{y})}{\|\mathbf{t}-\mathbf{y}\|_2} \right] & \|\mathbf{x}-\mathbf{z}\|_2 = 0 \\ \delta \frac{(\mathbf{z}-\mathbf{x})}{h} & \|\mathbf{x}-\mathbf{z}\|_2 = h \\ \delta \cdot \left[\tau \frac{(\mathbf{s}-\mathbf{y})}{\|\mathbf{y}-\mathbf{s}\|_2} + (1-\tau) \frac{(\mathbf{y}-\mathbf{t})}{\|\mathbf{t}-\mathbf{y}\|_2} \right] & \|\mathbf{x}-\mathbf{z}\|_2 = 2h \\ \delta(\mathbf{0}_{3m}, \mathbf{1}_m) & \|\mathbf{x}-\mathbf{z}\|_2 \geq 3h \end{cases}. \quad (4.4)$$

At intermediate distances, interpolate according to $\|\mathbf{x}-\mathbf{z}\|_2$ (see Figure 4.3).

Notice that for each fixed choice of $\tau \in [0, 1]$ (and \mathbf{z}), \hat{g} is $O(\delta/h) = O(1)$ -Lipschitz. Furthermore, $\Delta_{(\mathbf{s} \rightarrow \mathbf{y})}$ and $\Delta_{(\mathbf{y} \rightarrow \mathbf{t})}$ are 1-Lipschitz in \mathbf{x} . For any $\mathbf{z} \in L_{\mathbf{y}}$, $\Delta_{(\mathbf{y} \rightarrow \mathbf{t})}(\mathbf{z}) + \Delta_{(\mathbf{s} \rightarrow \mathbf{y})}(\mathbf{z}) = \sqrt{h}$. For general \mathbf{x} , we have

$$\Delta_{(\mathbf{y} \rightarrow \mathbf{t})}(\mathbf{x}) + \Delta_{(\mathbf{s} \rightarrow \mathbf{y})}(\mathbf{x}) \geq \Delta_{(\mathbf{y} \rightarrow \mathbf{t})}(\mathbf{z}) + \Delta_{(\mathbf{s} \rightarrow \mathbf{y})}(\mathbf{z}) - 2\|\mathbf{x}-\mathbf{z}\|_2 = \sqrt{h} - 2\|\mathbf{x}-\mathbf{z}\|_2; \quad (4.5)$$

so τ is $O(1/\sqrt{h})$ -Lipschitz whenever $\|\mathbf{x}-\mathbf{z}\|_2 < 3h$, and otherwise has no effect on $\hat{g}(\mathbf{x})$. We conclude that \hat{g} is $O(1)$ -Lipschitz when interpolating across different values of τ . At the interface with (4.2) τ is 1 (0 near $\mathbf{z}_{(\mathbf{y} \rightarrow \mathbf{t})}$), so (4.2) and (4.4) are equal. Therefore \hat{g} is $O(1)$ -Lipschitz on all of $[-1, 2]^{4m}$.

To lower bound the magnitude of the displacement, we argue that $(\mathbf{z}-\mathbf{x})$ is orthogonal to $\left[\tau \frac{(\mathbf{y}-\mathbf{s})}{\|\mathbf{y}-\mathbf{s}\|_2} + (1-\tau) \frac{(\mathbf{t}-\mathbf{y})}{\|\mathbf{t}-\mathbf{y}\|_2} \right]$. First, observe that we can restrict our attention to the component of $(\mathbf{z}-\mathbf{x})$ that belongs to the plane defined by $\mathbf{s}, \mathbf{y}, \mathbf{t}$ (in which \mathbf{z} also lies). Let $P_{\mathbf{s}, \mathbf{y}, \mathbf{t}}(\mathbf{x})$ denote the projection of \mathbf{x} to this plane. We can write points in this plane in terms of their $\Delta(\cdot) \triangleq (\Delta_{(\mathbf{s} \rightarrow \mathbf{y})}(\cdot), \Delta_{(\mathbf{y} \rightarrow \mathbf{t})}(\cdot))$ values. (Recall that $(\mathbf{s} \rightarrow \mathbf{y})$ and $(\mathbf{y} \rightarrow \mathbf{t})$ are orthogonal.)

First, observe that $\Delta(\mathbf{z}_{(\mathbf{s} \rightarrow \mathbf{y})}) = (0, \sqrt{h})$, $\Delta(\mathbf{z}_{(\mathbf{y} \rightarrow \mathbf{t})}) = (\sqrt{h}, 0)$ and $\Delta(\mathbf{y}) = (\sqrt{h}, \sqrt{h})$. Notice also that

$$\left[\tau \frac{(\mathbf{y}-\mathbf{s})}{\|\mathbf{y}-\mathbf{s}\|_2} + (1-\tau) \frac{(\mathbf{t}-\mathbf{y})}{\|\mathbf{t}-\mathbf{y}\|_2} \right] = \left[\tau \frac{(\mathbf{y}-\mathbf{z}_{(\mathbf{s} \rightarrow \mathbf{y})})}{\sqrt{h}} + (1-\tau) \frac{(\mathbf{z}_{(\mathbf{y} \rightarrow \mathbf{t})}-\mathbf{y})}{\sqrt{h}} \right].$$

Putting those together, we have that

$$\begin{aligned} & \Delta \left(\left[\tau \frac{\mathbf{y}}{\|\mathbf{y}-\mathbf{s}\|_2} + (1-\tau) \frac{\mathbf{t}}{\|\mathbf{t}-\mathbf{y}\|_2} \right] \right) \\ & - \Delta \left(\left[\tau \frac{\mathbf{s}}{\|\mathbf{y}-\mathbf{s}\|_2} + (1-\tau) \frac{\mathbf{y}}{\|\mathbf{t}-\mathbf{y}\|_2} \right] \right) = (\tau, 1-\tau). \end{aligned} \quad (4.6)$$




Figure 4.3 The figure (not drawn to scale) shows some of the important points near a Brouwer vertex y : There is an incoming Brouwer line segment from s through $z_{(s \rightarrow y)}$, and an outgoing Brouwer line segment to t through $z_{(y \rightarrow t)}$. For each point x between the dashed lines, we assign a point z on the line L_y as in (4.3), and define the displacement according to (4.4). Outside the dashed lines (including at y itself), we use the default displacement $\delta(\mathbf{0}_{3m}, \mathbf{1}_m)$.

For z , we have

$$\Delta(\mathbf{z}) = \tau \Delta(\mathbf{z}_{(s \rightarrow y)}) + (1 - \tau) \Delta(\mathbf{z}_{(y \rightarrow t)}) = \sqrt{h}(1 - \tau, \tau).$$

Finally, for $P_{s,y,t}(\mathbf{x})$, we can write

$$\begin{aligned} \Delta(P_{s,y,t}(\mathbf{x})) &= (\Delta_{(y \rightarrow t)}(\mathbf{x}), \Delta_{(s \rightarrow y)}(\mathbf{x})) \\ &= \frac{1}{\Delta_{(y \rightarrow t)}(\mathbf{x}) + \Delta_{(s \rightarrow y)}(\mathbf{x})} (1 - \tau, \tau). \end{aligned}$$

Therefore $\Delta(\mathbf{z}) - \Delta(P_{s,y,t}(\mathbf{x}))$ is orthogonal to (4.6).

4.2.2.4 Close to an End-of-Any-Line

Close to the non-trivial end or start of any line, we don't have to be as careful with defining the displacement: any Lipschitz extension of the displacement we defined everywhere else would do, since here we are allowed (in fact, expect) to have fixed points.

For concreteness, let $(\mathbf{s} \rightarrow \mathbf{t})$ be the last Brouwer line segment in a path. In (4.2), we defined the displacement for points \mathbf{x} such that $\sigma_{(\mathbf{s} \rightarrow \mathbf{t})}(\mathbf{x}) \leq 1 - \sqrt{h}$. For points such that $\sigma_{(\mathbf{s} \rightarrow \mathbf{t})}(\mathbf{x}) = 1$ (i.e., at the hyperplane through \mathbf{t} and perpendicular to $(\mathbf{s} \rightarrow \mathbf{t})$), we simply set the default displacement $\hat{g}(\mathbf{x}) \triangleq \delta(\mathbf{0}_{3m}, \mathbf{1}_m)$. For intermediate values of $\sigma_{(\mathbf{s} \rightarrow \mathbf{t})}(\mathbf{x}) \in [1 - \sqrt{h}, 1]$, we simply interpolate according to $\sigma_{(\mathbf{s} \rightarrow \mathbf{t})}(\mathbf{x})$. Notice that this induces a fixed point for some intermediate point, since for \mathbf{x} directly “above” the Brouwer line segment, $\delta \frac{\mathbf{z} - \mathbf{x}}{h}$ perfectly cancels $\delta(\mathbf{0}_{3m}, \mathbf{1}_m)$. Define the displacement analogously at the (non-trivial) start of a path.

4.2.2.5 Outside the Picture

The displacement outside the picture is constructed by interpolating the displacement at $\frac{1}{m} \sum_{i=3m+1}^{4m} x_i = 1/2$, and the displacement at points in the “top” of the hypercube, where $x_i = 2$ for every i in the last m coordinates. The former displacement, where $\mathbb{E}_{i \in \{3m+1, \dots, 4m\}} x_i = 1/2$, is defined to match the displacement inside the picture. Namely, it is the default displacement everywhere except near the first Brouwer line segment that goes “down” from $\mathbf{s} = (\mathbf{0}_{3m}, 2 \cdot \mathbf{1}_m)$ to $\mathbf{t} = (\mathbf{0}_{4m})$. Near this line, it is defined according to (4.2). (Notice that $\|\mathbf{t} - \mathbf{s}\|_2 = 1$.)

Formally, let $\mathbf{z}_{1/2} = (\mathbf{0}_{3m}, \frac{1}{2} \cdot \mathbf{1}_m)$; for \mathbf{x} on the boundary of the picture, we have:

$$\hat{g}(\mathbf{x}) \triangleq \begin{cases} \delta(\mathbf{0}_{3m}, -\mathbf{1}_m) & \|\mathbf{x} - \mathbf{z}_{1/2}\|_2 = 0 \\ \delta \frac{(\mathbf{z}_{1/2} - \mathbf{x})}{h} & \|\mathbf{x} - \mathbf{z}_{1/2}\|_2 = h \\ \delta(\mathbf{0}_{3m}, \mathbf{1}_m) & \|\mathbf{x} - \mathbf{z}_{1/2}\|_2 \geq 2h. \end{cases} \quad (4.7)$$

For points \mathbf{x} such that $\sum_{i=3m+1}^{4m} x_i$ is very close to 2, the displacement $\delta(\mathbf{0}_{3m}, \mathbf{1}_m)$ is not helpful because it points outside the hypercube, i.e., it would get completely erased by the truncation. Instead, we define the displacement as follows:

$$\hat{g}(\mathbf{x}) \triangleq \begin{cases} \delta(\mathbf{0}_{3m}, -\mathbf{1}_m) & \|\mathbf{x} - \mathbf{z}_2\|_2 = 0 \\ \delta \frac{(\mathbf{z}_1 - \mathbf{x})}{h} & \|\mathbf{x} - \mathbf{z}_2\|_2 \geq h, \end{cases} \quad (4.8)$$

where $\mathbf{z}_2 = (\mathbf{0}_{3m}, 2 \cdot \mathbf{1}_m)$. When $\theta \triangleq \sum_{i=3m+1}^{4m} x_i \in (1/2, 2)$, we interpolate between (4.7) and (4.8) according to $\frac{\theta - 1/2}{3/2}$.

4.2.3 Locally Computing the Brouwer Function

The function f defined above is *local* in two different ways: First, in order to compute $f(\mathbf{x})$ at a point \mathbf{x} that is close to the embedding of one or a few vertices in the MEMBERSHIP END-OF-A-LINE instance I , we only need to understand I at those vertices. The second type of locality observes that in order to compute just a single coordinate $f_i(\mathbf{x})$, we only need partial information about \mathbf{x} .

Locality in the MEMBERSHIP END-OF-A-LINE Instance

Fact 4.2 The function f is *local* in the sense that there exists a class of functions $\{f_{I_1, I_2} : [-1, 2]^{4m} \rightarrow [-1, 2]^{4m}\}$ that do not depend on I , and f can be defined as an interpolation between these functions such that:

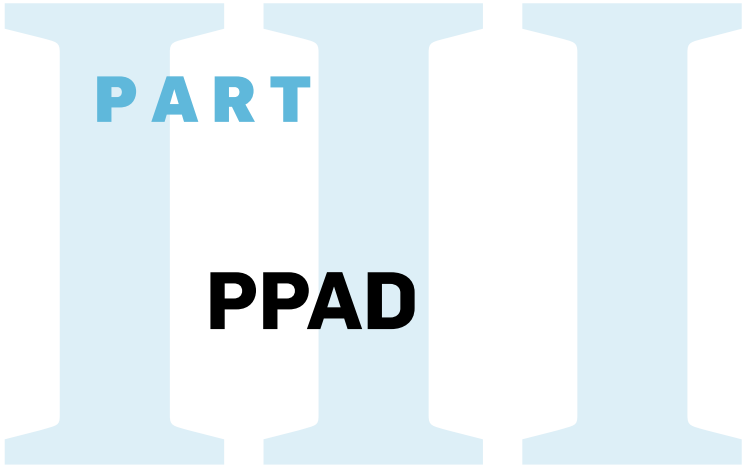
1. If the first m -tuple of coordinates of \mathbf{x} is $12\sqrt{h}$ -close to the encoded vertex $E(v)$, but the second m -tuple of coordinates of \mathbf{x} is $12\sqrt{h}$ -far from any encoded vertex $E(w)$, then $f_{I(v), I_2}(\mathbf{x}) = f(\mathbf{x})$ for every I_2 .
2. If the second m -tuple of coordinates of \mathbf{x} is $12\sqrt{h}$ -close to the encoded vertex $E(w)$, but the first m -tuple of coordinates of \mathbf{x} is $12\sqrt{h}$ -far from any encoded vertex $E(v)$, then $f_{I_1, I(w)}(\mathbf{x}) = f(\mathbf{x})$ for every I_1 .
3. If the first m -tuple of coordinates of \mathbf{x} is $12\sqrt{h}$ -close to the encoded vertex $E(v)$, and the second m -tuple of coordinates of \mathbf{x} is $12\sqrt{h}$ -close to the encoded vertex $E(w)$, then $f_{I(v), I(w)}(\mathbf{x}) = f(\mathbf{x})$.
4. If none of the above conditions are satisfied, then $f_{I_1, I_2}(\mathbf{x}) = f(\mathbf{x})$ for every I_1, I_2 .

Locality in Single Coordinates

In order to compute all of $f(\mathbf{x})$ exactly, we essentially need to know \mathbf{x} in every coordinate. However, in order to compute $f_i(\mathbf{x})$ (the i -th coordinate of $f(\mathbf{x})$), it suffices to know x_i and that \mathbf{x} is one of the following:

1. inside the picture, but far from every Brouwer line segment
2. close to some point \mathbf{z} on Brouwer line segment $(\mathbf{s} \rightarrow \mathbf{t})$ (but far from \mathbf{s} and \mathbf{t})
 - we also need to know $s_i, t_i, z_i, \|\mathbf{x} - \mathbf{z}\|_2$, and $\|\mathbf{t} - \mathbf{s}\|_2$
3. close to some point \mathbf{z} on line $L_{\mathbf{y}}$ for Brouwer vertex \mathbf{y} on the intersection of Brouwer lines $(\mathbf{s} \rightarrow \mathbf{y})$ and $(\mathbf{y} \rightarrow \mathbf{t})$
 - we also need to know $s_i, y_i, t_i, z_i, \|\mathbf{x} - \mathbf{z}\|_2, \|\mathbf{y} - \mathbf{s}\|_2, \|\mathbf{t} - \mathbf{y}\|_2$, and α
4. outside the picture
 - we also need to know $\mathbb{E}_{i \in \{3m+1, \dots, 4m\}} x_i$ and $\|\mathbf{x} - \mathbf{z}\|_2$, where \mathbf{z} is the $(\mathbb{E}_{i \in \{3m+1, \dots, 4m\}} x_i)$ -weighted average of $\mathbf{z}_{1/2}$ and \mathbf{z}_2

By Lipschitz continuity, if we only want to compute $f_i(\mathbf{x})$ to within $\pm O(\epsilon)$, it suffices to know all the quantities above to within $\pm \epsilon$. Furthermore, at distance $\pm \epsilon$ near interfaces between the different cases (inside/outside the picture, close to 0/1/2 lines), we can use the wrong displacement, and still be within $\pm O(\epsilon)$ of $f_i(\mathbf{x})$.

A large, light blue, stylized number '111' is centered on the page. The number is composed of three vertical bars with horizontal caps at the top and bottom, giving it a blocky, architectural appearance. The word 'PART' is written in a smaller, blue, sans-serif font across the top of the first bar. The word 'PPAD' is written in a larger, bold, black, sans-serif font across the middle of the second bar.

PART

PPAD

PPAD-Hardness of Approximation

In this part we introduce and prove our PPAD-hardness of approximation results, in particular for finding Nash equilibria. As we discussed in the introduction, for a constant number of players, we are unlikely to prove PPAD-hardness since a quasi-polynomial time approximation algorithm exists [Lipton et al. 2003]. Our main focus in this part is games with a large number of players. For such games, there is a question of representation: the normal form representation is exponential in the number of the players. Instead, we consider three natural and well-studied classes of many-player games that have succinct representations:

Definition 5.1

Polymatrix games. In a polymatrix game [Yanovskaya 1968], each pair of players simultaneously plays a separate two-player game. Every player has to play the same strategy in every two-player subgame, and her utility is the sum of her subgame utilities. The game is given in the form of the payoff matrix for each two-player game.

Graphical games. In a graphical game [Kearns 2007], the utility of each player depends only on the action chosen by a few other players. This game now naturally induces a directed graph: we say that $(i, j) \in E$ if the utility of player j depends on the strategy chosen by player i . When the maximal incoming degree is bounded, the game has a representation polynomial in the number of players and strategies.

Succinct games. Most generally, in a succinct game [Schoenebeck and Vadhan 2012]¹ the normal form representation is succinctly described by some small circuit.

1. Related notions for two-player games have also been considered by Feigenbaum et al. [1995] and Fortnow et al. [2008].

Any of the above restrictions suffices to guarantee that the game has a succinct representation. Our main result in this part is that even for games that satisfy all restrictions simultaneously, finding an ϵ -approximate Nash equilibrium is PPAD-complete.

Theorem 5.1 There exists a constant $\epsilon > 0$, such that given a degree 3, bipartite, polymatrix game where each player has two actions, finding an ϵ -approximate Nash equilibrium is PPAD-complete.

The notion of ϵ -approximate Nash equilibrium used in Theorem 5.1 requires that *every* player plays ϵ -optimally. The most interesting question left open in this part is whether the equilibrium computation problem remains PPAD-hard even if we only require that *most* of the players play ϵ -optimally. This is the “PCP Conjecture for PPAD”:

Conjecture 5.1 PCP for PPAD. There exist constants $\epsilon, \delta > 0$ such that given a degree 3, bipartite, polymatrix game where each player has two actions, finding an (ϵ, δ) -WeakNash is PPAD-complete.

While proving (or disproving) the “PCP Conjecture for PPAD” remains open (see additional discussion in Subsection 16.1.1), we do prove in this section that for the more general class of succinct games, finding an (ϵ, δ) -WeakNash is indeed PPAD-hard.

Theorem 5.2 There exist constants $\epsilon, \delta > 0$ such that finding an (ϵ, δ) -WeakNash is PPAD-hard for succinct multiplayer games where each player has a constant number of actions.

Besides Theorem 5.2, all our results in this part rely on the hardness of approximation for the generalized circuit problem. Generalized circuits are similar to standard algebraic circuits, the main difference being that generalized circuits contain cycles, which allow them to verify fixed points of continuous functions. A generalized circuit induces a constraint satisfaction problem, ϵ -GCIRCUIT [Chen et al. 2009b]: find an assignment for the values on the lines of the circuit that simultaneously ϵ -approximately satisfies all the constraints imposed by the gates (see Chapter 6 for a formal definition). ϵ -GCIRCUIT was implicitly proven PPAD-complete for exponentially small ϵ by Daskalakis et al. [2009a], and explicitly for polynomially small ϵ by Chen et al. [2009b]. Here we prove that it continues to be PPAD-complete for some constant ϵ .

Theorem 5.3 Generalized circuit. There exists a constant $\epsilon > 0$ such that ϵ -GCIRCUIT with fan-out 2 is PPAD-complete.

This result, in turn, builds on the hardness of finding an ℓ_∞ -approximate fixed point, which we proved in Theorem 4.1.

We note that except for the hardness of the course allocation problem (Chapter 10), all the problems we consider in this part were previously known to be PPAD-hard for polynomial approximation factors.

The Generalized Circuit Problem

Generalized circuits are similar to the standard algebraic circuits, the main difference being that generalized circuits contain cycles, which allow them to verify fixed points of continuous functions. We restrict the class of generalized circuits to include only a particular list of gates described below. Formally,

Definition 6.1 Generalized circuits, [Chen et al. 2009b]. A *generalized circuit* \mathcal{S} is a pair (V, \mathcal{T}) , where V is a set of nodes and \mathcal{T} is a collection of gates. Every gate $T \in \mathcal{T}$ is a 5-tuple $T = G(\zeta \mid v_1, v_2 \mid v)$, in which $G \in \{G_\zeta, G_{\times\zeta}, G_=, G_+, G_-, G_<, G_\vee, G_\wedge, G_-\}$ is the type of the gate; $\zeta \in \mathbb{R} \cup \{nil\}$ is a real parameter; $v_1, v_2 \in V \cup \{nil\}$ are the first and second input nodes of the gate; and $v \in V$ is the output node.

The collection \mathcal{T} of gates must satisfy the following important property: For every two gates $T = G(\zeta \mid v_1, v_2 \mid v)$ and $T' = G'(\zeta' \mid v'_1, v'_2 \mid v')$ in \mathcal{T} , $v \neq v'$.

Alternatively, we can think of each gate as a constraint on the values on the incoming and outgoing wires. We are interested in the following constraint satisfaction problem: given a generalized circuit, find an assignment to all the wires that simultaneously satisfies all the gates. When every gate computes a continuous function of the incoming wires (with inputs and output in $[0, 1]$), a solution must exist by Brouwer's fixed point theorem.

In particular, we are interested in the approximate version of this CSP, where we must approximately satisfy every constraint.

Definition 6.2 Given a generalized circuit $\mathcal{S} = (V, \mathcal{T})$, we say that an assignment $\mathbf{x}: V \rightarrow [0, 1]$ *ϵ -approximately satisfies* \mathcal{S} if for each of the gates shown in Table 6.1, \mathbf{x} satisfies the corresponding constraints.

Given a generalized circuit $\mathcal{S} = (V, \mathcal{T})$, ϵ -GCIRCUIT is the problem of finding an assignment that ϵ -approximately satisfies it.

Table 6.1 Gates for the ϵ -GCIRCUIT Problem

Gate	Constraint
$G_{\zeta}(\alpha \mid a)$	$\mathbf{x}[a] = \alpha \pm \epsilon$
$G_{\times\zeta}(\alpha \mid a \mid b)$	$\mathbf{x}[b] = \alpha \cdot \mathbf{x}[a] \pm \epsilon$
$G_{=}(\mid a \mid b)$	$\mathbf{x}[b] = \mathbf{x}[a] \pm \epsilon$
$G_{+}(\mid a, b \mid c)$	$\mathbf{x}[c] = \min(\mathbf{x}[a] + \mathbf{x}[b], 1) \pm \epsilon$
$G_{-}(\mid a, b \mid c)$	$\mathbf{x}[c] = \max(\mathbf{x}[a] - \mathbf{x}[b], 0) \pm \epsilon$
$G_{<}(\mid a, b \mid c)$	$\mathbf{x}[c] = \begin{cases} 1 \pm \epsilon & \mathbf{x}[a] < \mathbf{x}[b] - \epsilon \\ 0 \pm \epsilon & \mathbf{x}[a] > \mathbf{x}[b] + \epsilon \end{cases}$
$G_{\vee}(\mid a, b \mid c)$	$\mathbf{x}[c] = \begin{cases} 1 \pm \epsilon & \mathbf{x}[a] = 1 \pm \epsilon \text{ or } \mathbf{x}[b] = 1 \pm \epsilon \\ 0 \pm \epsilon & \mathbf{x}[a] = 0 \pm \epsilon \text{ and } \mathbf{x}[b] = 0 \pm \epsilon \end{cases}$
$G_{\wedge}(\mid a, b \mid c)$	$\mathbf{x}[c] = \begin{cases} 1 \pm \epsilon & \mathbf{x}[a] = 1 \pm \epsilon \text{ and } \mathbf{x}[b] = 1 \pm \epsilon \\ 0 \pm \epsilon & \mathbf{x}[a] = 0 \pm \epsilon \text{ and or } \mathbf{x}[b] = 0 \pm \epsilon \end{cases}$
$G_{-}(\mid a \mid b)$	$\mathbf{x}[b] = \begin{cases} 1 \pm \epsilon & \mathbf{x}[a] = 0 \pm \epsilon \\ 0 \pm \epsilon & \mathbf{x}[a] = 1 \pm \epsilon \end{cases}$

Note: Where G_{ζ} and $G_{\times\zeta}$ also take a parameter $\alpha \in [0, 1]$.

Brittle Comparators. Intuitively, in order for (approximate) solutions to the circuit problem to correspond to (approximate) equilibria, all our gates should implement continuous (Lipschitz) functions. The gate $G_{<}(\mid a, b \mid c)$, for example, approximates that the function

$$c(a, b) = \begin{cases} 1 & a < b \\ 0 & a \geq b, \end{cases}$$

which is not continuous. To overcome this problem, [Daskalakis et al. \[2009a\]](#) defined the *brittle comparator*: when a is (ϵ -) larger than b , it outputs 0; when b is (ϵ -) larger than a , it outputs 1. However, when a and b are (ϵ -approximately) equal, its behavior is undefined.

Brittleness introduces difficulties in the transition from continuous to discrete solutions. This challenge is overcome by an averaging gadget, which is described in detail in Section 6.2.

Our Results

ϵ -GCIRCUIT was implicitly proven PPAD-complete for exponentially small ϵ by [Daskalakis et al. \[2009a\]](#), and explicitly for polynomially small ϵ by [Chen et al. \[2009b\]](#). Here we prove that it continues to be PPAD-complete for some constant ϵ .

Theorem 6.1 Generalized circuit; Theorem 5.3 restated. There exists a constant $\epsilon > 0$ such that ϵ -GCIRCUIT with fan-out 2 is PPAD-complete.

6.1 Proof Overview

The key idea that enables us to improve over previous hardness of approximation for ϵ -GCIRCUIT (and Nash equilibrium) [Chen et al. 2009b, Daskalakis et al. 2009a] is our particular choice of hard fixed point instance in Section 4.1. The first advantage is that our instance is simply harder to approximate: finding an ϵ -approximate fixed point (i.e., \mathbf{x} such that $\|f(\mathbf{x}) - \mathbf{x}\|_\infty \leq \epsilon$) is PPAD-hard for $\epsilon = \Omega(1)$ (as opposed to $\epsilon = 1/\exp(n)$ for Daskalakis et al. [2009a] and $\epsilon = 1/\text{poly}(n)$ in Chen et al. [2009b]).

A Simpler Averaging Gadget. Our construction of hard instances of Brouwer functions, as do the ones from previous works, partitions the (continuous) hypercube into subcubes, and defines the function separately on each subcube. When we construct a circuit that approximately simulates such a Brouwer function, we have a problem near the facets of the subcubes: using *approximate gates* and *brittle comparators* (see Table 6.1), one cannot determine to which subcube the input belongs. This is the most challenging part of our reduction, as was also the case in Chen et al. [2009b] and Daskalakis et al. [2009a].

Originally, Daskalakis et al. [2009a] tackled this obstacle by approximating $f(\mathbf{x})$ as the average over a ball around \mathbf{x} . The key observation is that even if \mathbf{x} is close to a facet between subcubes, most of the points in its neighborhoods will be sufficiently far. Yet if f is Lipschitz they are mapped approximately to the same point as \mathbf{x} . This works fine in $O(1)$ dimensions, but then the inapproximability parameter is inherently exponentially small (in constant dimensions, it is easy to construct a $1/\text{poly}(n)$ -net over the unit hypercube). For $\text{poly}(n)$ dimensions, the (discretization of the) ball around \mathbf{x} contains exponentially many points.

Chen et al. [2009b] overcome this problem using *equiangle sampling*: consider many translations of the input vector by adding small multiples of the all-ones vector; compute the displacement for each translation, and average. Since each translation may be close to a facet in a different dimension, Chen et al. consider a polynomial number of translations. Thus, all translations must be polynomially close to each other—otherwise they will be too far to approximate the true input.

We avoid this problem by observing another nice property of Hirsch et al.’s [1989] construction: when the input vector lies near two or more facets, the displacement is (approximately) the same, regardless of the subcube. Once we rule out such points, it suffices to sample only a constant number of points (as at most one of them may be too close to a facet). See also the illustration in Figure 6.1.

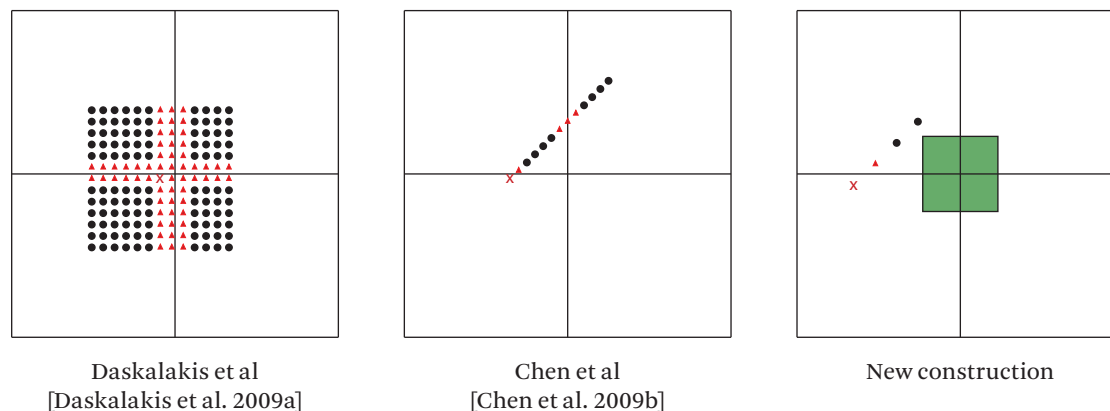


Figure 6.1 A comparison of the averaging gadgets of [Daskalakis et al. \[2009a\]](#) and [Chen et al. \[2009b\]](#), and ours. The red x is the point whose displacement we would like to estimate using imprecise gates and brittle comparators. Points that are too close to a facet between subcubes are denoted by red triangles, while points that are sufficiently far are denoted by black circles. Finally, our construction has a “safe” zone (shaded green) around the corner where we don’t need to parse the subcube; thus we only need to avoid one facet.

Completing the Proof. Given a point $\mathbf{x}' \approx \mathbf{x}$ that is safely in the interior of one subcube, we can parse the corresponding binary vector, use logical operator gates to simulate the END-OF-A-LINE circuit, and then approximately compute $f(\mathbf{x}')$. This is tedious, but mostly straightforward.

One particular challenge that nevertheless arises is preventing the error from accumulating when concatenating approximate gates. Of course this is more difficult in our setting where each gate may err by a constant $\epsilon > 0$. Fortunately, the definition of ϵ -GCIRCUIT provides logical operator gates that round the output to $\{0, 1\}$ before introducing new error. As long as the inputs are unambiguous bits, approximate logical operator gates can be concatenated without accumulating errors.

In order to carry out the reduction to Nash equilibrium (Section 7.1), we must first ensure that every gate in our generalized circuit has a constant fan-out (Section 6.3). We can replace each logical operator gate with a binary tree of fan-out 2, alternating negation gates (that do not accumulate error). Given an arithmetic gate with large fan-out, we convert its output to unary representation¹ using a constant

1. Unary representation of numbers with constant precision is prevalent throughout our implementation of the generalized circuit. We prefer unary representation over binary, because in the former at most one bit can be ambiguous due to the use of brittle comparators.

number of (fan-out 2) gates. Then we copy the unary representation using a binary tree of negation gates. Finally, we convert each copy back to a real number using a constant number of gates.

6.2 From Brouwer to ϵ -GCIRCUIT

In this section we prove a slightly easier version of Theorem 5.3, for a generalized circuit with unbounded fan-out. We reduce to constant fan-out in the next section.

Proposition 6.1 There exists a constant $\epsilon > 0$ such that ϵ -GCIRCUIT is PPAD-complete.

Proof We continue to denote the hard Brouwer function by $f: [0, 1]^{2n+2} \rightarrow [0, 1]^{2n+2}$, and its associated displacement by $g(\mathbf{y}) = f(\mathbf{y}) - \mathbf{y}$. We design a generalized circuit \mathcal{S} that computes f and verifies that the output is equal to the input. We show that every ϵ -approximate solution to \mathcal{S} corresponds to an $O(\epsilon^{1/4})$ -approximate fixed point of f .

Recall that the construction from Section 4.1 divides the hypercube into equal-sized subcubes (of length $1/4$). Furthermore, all the paths in H are embedded in the 2^{2n+1} subcubes that belong to the picture. For ease of exposition, we present a construction that only works for points in the picture, i.e., $\mathbf{y} \in [1/4, 3/4]^{2n+1} \times [1/4, 1/2]$. It is straightforward to use the same ideas to extend the circuit to deal with all $\mathbf{y} \in [0, 1]^{2n+2}$.

The most challenging part of the construction is the extraction of the information about the local subcube: Is it part of a tube? If so, which are the entrance and exit facets? This is done by extracting the binary representation of the current subcube, and feeding it to the (Boolean) circuit that computes H (recall that H is our collection of paths and cycles from Subsection 4.1.2). Notice that whenever we have valid logic bits, i.e., $\mathbf{x}[b] < \epsilon$ or $\mathbf{x}[b] > 1 - \epsilon$, we can perform logic operations on them without increasing the error.

Once we know the behavior of the path on the current subcube, we simply have to locally implement the mapping from the previous section, for which we have a closed form description, using the available gates in the definition of generalized circuits. Since this definition does not include multiplication and division, we implement multiplication and division in Algorithms 6.2 and 6.3 in Subsection 6.2.1.

Our construction has four parts: (1) equiangle sampling segment, (2) computing the displacement, (3) summing the displacement vectors, and (4) closing the loop. The first part contains a new idea: using only a constant size sample. The second part is a more technical but straightforward description of the implementation

Algorithm 6.1 IF-ELSE($| a, b, c | d$)

1. $G_{-}(| a | \bar{a})$ $\triangleright \bar{a}$ is the negation of a
 2. $G_{-}(| b, \bar{a} | b')$ $\triangleright b'$ is (approximately) equal to b iff $a = 1$
 3. $G_{-}(| \bar{a} | \bar{\bar{a}})$ $\triangleright \bar{\bar{a}}$ is the rounding of a to $\{0, 1\}$
 4. $G_{-}(| c, \bar{\bar{a}} | c')$ $\triangleright c'$ is (approximately) equal to c iff $a = 0$
 5. $G_{-}(| b', c' | d)$
-

of the closed-form mapping by approximate gates. The third and fourth parts are essentially identical to [Chen et al. \[2009b\]](#). ■

6.2.1 Subroutines

In this subsection we show how to implement a few useful subroutines using the gates in the definition of ϵ -GCIRCUIT.

6.2.1.1 If-Else

We begin by describing how to implement a simple if-else. Similar ideas can be used to implement more involved cases such as (4.1).

Claim 6.1 In any ϵ -approximate solution to IF-ELSE($| a, b, c | d$),

$$\mathbf{x}[d] = \begin{cases} \mathbf{x}[c] \pm O(\epsilon) & \text{if } \mathbf{x}[a] < \sqrt{\epsilon} \\ \mathbf{x}[b] \pm O(\epsilon) & \text{if } \mathbf{x}[a] > 1 - \sqrt{\epsilon}. \end{cases}$$

Proof By definition of G_{-} , we have that

$$\mathbf{x}[\bar{a}] = \begin{cases} 1 \pm \epsilon & \text{if } \mathbf{x}[a] < \sqrt{\epsilon} \\ 0 \pm \epsilon & \text{if } \mathbf{x}[a] > 1 - \sqrt{\epsilon}. \end{cases}$$

Therefore by definition of G_{-} ,

$$\mathbf{x}[b'] = \begin{cases} 0 \pm O(\epsilon) & \text{if } \mathbf{x}[a] < \sqrt{\epsilon} \\ \mathbf{x}[b] \pm O(\epsilon) & \text{if } \mathbf{x}[a] > 1 - \sqrt{\epsilon}. \end{cases}$$

Similarly,

$$\mathbf{x}[c'] = \begin{cases} \mathbf{x}[c] \pm O(\epsilon) & \text{if } \mathbf{x}[a] < \sqrt{\epsilon} \\ 0 \pm O(\epsilon) & \text{if } \mathbf{x}[a] > 1 - \sqrt{\epsilon}. \end{cases}$$

Finally, the claim follows by definition of G_{+} . ■

Algorithm 6.2 MULTIPLY($| a, b | c$)

1. $G_\zeta(0 | h_0)$
 2. **for each** $k \in [1/\sqrt{\epsilon}]$:
 - (a) $G_\zeta(k\sqrt{\epsilon} | \zeta_k), G_{<}(| a, \zeta_k | \bar{a}_k)$
 - ▷ # The vector (\bar{a}_k) is the unary representation of a :

$$\sum_{k:\mathbf{x}[\bar{a}_k] < \epsilon} \sqrt{\epsilon} = \max_{k:\mathbf{x}[\bar{a}_k] < \epsilon} k\sqrt{\epsilon} = \mathbf{x}[a] \pm O(\sqrt{\epsilon}),$$
 - (b) $G_{\times\zeta}(\sqrt{\epsilon} | b | d_k)$,
 - ▷ # The vector (d_k) is simply, equal to $b \cdot \sqrt{\epsilon}$ everywhere:

$$\sum_{k:\mathbf{x}[\bar{a}_k] < \epsilon} \mathbf{x}[d_k] = \left(\sum_{k:\mathbf{x}[\bar{a}_k] < 1-\epsilon} \sqrt{\epsilon} \right) \cdot \mathbf{x}[b] \pm O(\sqrt{\epsilon}) = \mathbf{x}[a] \cdot \mathbf{x}[b] \pm O(\sqrt{\epsilon})$$
 - (c) $G_{-}(| d_k, \bar{a}_k | e_k)$
 - ▷ # The vector (e_k) is $b \cdot \sqrt{\epsilon}$ only when $(\bar{a}_k) < \epsilon$:

$$\sum_{k:\mathbf{x}[\bar{a}_k] < \epsilon} \mathbf{x}[e_k] = \mathbf{x}[a] \cdot \mathbf{x}[b] \pm O(\sqrt{\epsilon})$$
 - (d) $G_{+}(| h_{k-1}, e_k | h_k)$
 - ▷ Finally, we sum the e_k 's to get $a \cdot b$:

$$\mathbf{x}\left[h_{1/\sqrt{\epsilon}}\right] = \mathbf{x}[a] \cdot \mathbf{x}[b] \pm O(\sqrt{\epsilon})$$
 3. $G_{=}(| h_{1/\sqrt{\epsilon}} | c)$
-

6.2.1.2 Multiply

Claim 6.2 In any ϵ -approximate solution to MULTIPLY($| a, b | c$),

$$\mathbf{x}[c] = \mathbf{x}[a] \cdot \mathbf{x}[b] \pm O(\sqrt{\epsilon}).$$

Proof For any k , the first gate implies that

$$\mathbf{x}[\zeta_k] = k\sqrt{\epsilon} \pm \epsilon.$$

The second gate thus gives

$$\mathbf{x}[\bar{a}_k] = \begin{cases} 0 \pm \epsilon & \text{if } \mathbf{x}[a] > k\sqrt{\epsilon} + O(\epsilon) \\ 1 \pm \epsilon & \text{if } \mathbf{x}[a] < k\sqrt{\epsilon} - O(\epsilon). \end{cases} \quad (6.1)$$

Notice that the above equation is ambiguous for at most one value of k . In particular,

$$\sum_k \left(1 - \mathbf{x}[\bar{a}_k]\right) \sqrt{\epsilon} = \mathbf{x}[a] \pm O(\sqrt{\epsilon}). \quad (6.2)$$

We also have

$$\mathbf{x}[d_k] = \mathbf{x}[b] \cdot \sqrt{\epsilon} \pm \epsilon.$$

The subtraction gate zeros $\mathbf{x}[d_k]$ for all k such that $\mathbf{x}[a] < k\sqrt{\epsilon} - O(\epsilon)$, and has negligible effect for k such that $\mathbf{x}[a] > k\sqrt{\epsilon} + O(\epsilon)$:

$$\mathbf{x}[e_k] = \begin{cases} \mathbf{x}[b] \cdot \sqrt{\epsilon} \pm 2\epsilon & \text{if } \mathbf{x}[a] > k\sqrt{\epsilon} + O(\epsilon) \\ 0 \pm 2\epsilon & \text{if } \mathbf{x}[a] < k\sqrt{\epsilon} - O(\epsilon). \end{cases}$$

The sum of the $\mathbf{x}[e_k]$'s satisfies:

$$\sum_k \mathbf{x}[e_k] = \mathbf{x}[a] \cdot \mathbf{x}[b] \pm O(\sqrt{\epsilon}),$$

where we have an error of $\pm O(\sqrt{\epsilon})$ arising from aggregating $\pm 2\epsilon$ for $1/\sqrt{\epsilon}$ distinct k 's, and another $\pm O(\sqrt{\epsilon})$ from (6.2).

By induction, each h_k is approximately equal to the sum of the first $\mathbf{x}[e_j]$'s:

$$\mathbf{x}[h_k] = \sum_{j=1}^k \mathbf{x}[e_j] \pm k\epsilon.$$

In particular, we have

$$\begin{aligned} \mathbf{x}[h_{1/\sqrt{\epsilon}}] &= \sum_k \mathbf{x}[e_k] \pm \sqrt{\epsilon} \\ &= \mathbf{x}[a] \cdot \mathbf{x}[b] \pm O(\sqrt{\epsilon}). \quad \blacksquare \end{aligned}$$

6.2.1.3 Divide

Claim 6.3 In any ϵ -approximate solution to $\text{DIVIDE}(|a, b|c)$,

$$\mathbf{x}[c] \cdot \mathbf{x}[b] = \mathbf{x}[a] \pm O(\sqrt{\epsilon}).$$

Notice that for Algorithm 6.3, in any ϵ -approximate solution, $\mathbf{x}[c] = \mathbf{x}[a]/\mathbf{x}[b] \pm O(\sqrt{\epsilon})/\mathbf{x}[b]$; when $\mathbf{x}[b]$ and ϵ are bounded away from 0, this is only a constant factor increase in the error.

Proof For each k , we have

$$\mathbf{x}[b_k] = k\sqrt{\epsilon} \cdot \mathbf{x}[b] \pm \epsilon.$$

Thus also

$$\mathbf{x}[d_k] = \begin{cases} 1 \pm \epsilon & \text{if } \mathbf{x}[a] > k\sqrt{\epsilon} \cdot \mathbf{x}[b] + O(\epsilon) \\ 0 \pm \epsilon & \text{if } \mathbf{x}[a] < k\sqrt{\epsilon} \cdot \mathbf{x}[b] - O(\epsilon). \end{cases}$$

Notice that $\mathbf{x}[d_k]$ is ambiguous for at most k . Furthermore, aggregating the $\pm\epsilon$ error over $1/\sqrt{\epsilon}$ distinct k 's, we have:

$$\sum \mathbf{x}[d_k] \cdot \sqrt{\epsilon} \mathbf{x}[b] = \mathbf{x}[a] \pm O(\sqrt{\epsilon}).$$

Algorithm 6.3 **DIVIDE**($| a, b | c$)

1. $G_{\zeta}(0 | | h_0)$
 2. **for each** $k \in [1/\sqrt{\epsilon}]$:
 - (a) $G_{\times\zeta}(k\sqrt{\epsilon} | b | b_k), G_{<}(| b_k, a | d_k)$
 - ▷ The vector (d_k) is the unary representation of a/b :

$$\left(\sum_{k:\mathbf{x}[d_k] > \epsilon} \sqrt{\epsilon} \right) \cdot \mathbf{x}[b] = \left(\max_{k:\mathbf{x}[d_k] > \epsilon} k\sqrt{\epsilon} \right) \cdot \mathbf{x}[b] = \mathbf{x}[a] \pm O(\sqrt{\epsilon})$$
 - (b) $G_{\times\zeta}(\sqrt{\epsilon} | d_k | e_k)$
 - ▷ The vector (e_k) is a $(\sqrt{\epsilon})$ -scaled version of (d_k) :

$$\left(\sum \mathbf{x}[e_k] \right) \cdot \mathbf{x}[b] = \mathbf{x}[a] \pm O(\sqrt{\epsilon})$$
 - (c) $G_{+}(| h_{k-1}, e_k | h_k)$
 - ▷ Finally, we sum the e_k 's:

$$\mathbf{x}[h_{1/\sqrt{\epsilon}}] \cdot \mathbf{x}[b] = \mathbf{x}[a] \pm O(\sqrt{\epsilon})$$
 3. $G_{=}(| h_{1/\sqrt{\epsilon}} | c)$
-

$\mathbf{x}[e_k]$'s are a step closer to what we need:

$$\mathbf{x}[e_k] = \mathbf{x}[d_k]\sqrt{\epsilon} \pm \epsilon,$$

and therefore also

$$\sum \mathbf{x}[e_k] \cdot \mathbf{x}[b] = \mathbf{x}[a] \pm O(\sqrt{\epsilon}). \quad (6.3)$$

Finally, by induction

$$\mathbf{x}[h_{1/\sqrt{\epsilon}}] = \sum \mathbf{x}[e_k] \pm \sqrt{\epsilon},$$

and the claim follows by plugging into (6.3). ■

6.2.1.4 Max

Claim 6.4 In any ϵ -approximate solution to $\text{MAX}(| a_1, \dots, a_n | b)$,

$$\mathbf{x}[b] = \max \mathbf{x}[a_i] \pm O(\sqrt{\epsilon}).$$

Proof Similarly to (6.1), we have that for each i, k

$$\mathbf{x}[c_{k,i}] = \begin{cases} 0 \pm \epsilon & \text{if } \mathbf{x}[a_i] < k\sqrt{\epsilon} - O(\epsilon) \\ 1 \pm \epsilon & \text{if } \mathbf{x}[a_i] > k\sqrt{\epsilon} + O(\epsilon). \end{cases}$$

Algorithm 6.4 $\text{MAX}(|a_1, \dots, a_n| b)$

1. $G_\zeta(0 \parallel h_0)$
 2. **for each** $k \in [1/\sqrt{\epsilon}]$:
 - (a) $G_\zeta(k\sqrt{\epsilon} \parallel \zeta_k)$
 - (b) $G_\zeta(0 \parallel d_{k,0})$
 - (c) **for each** $i \in [n]$:
 - (i) $G_{<}(|\zeta_k, a_i| c_{k,i})$
 - ▷ The vector $(c_{k,i})_k$ is the unary representation of a_i :

$$\forall i \left(\max_{k: x[c_{k,i}] > \epsilon} k\sqrt{\epsilon} \right) = \mathbf{x}[a_i] \pm O(\sqrt{\epsilon})$$
 - (ii) $G_{\vee}(|d_{k,i-1}, c_{k,i}| d_{k,i})$
 - ▷ The vector $(d_{k,n})$ is the unary representation of $\max a_i$:

$$\left(\max_{k: x[d_{k,n}] > \epsilon} k\sqrt{\epsilon} \right) = \max \mathbf{x}[a_i] \pm O(\sqrt{\epsilon})$$
 - (d) $G_{\times\zeta}(\sqrt{\epsilon} \parallel d_{k,n} \parallel e_k)$
 - ▷ The vector (e_k) is a $(\sqrt{\epsilon})$ -scaled version of (d_k) :

$$\left(\sum \mathbf{x}[e_k] \right) = \max \mathbf{x}[a_i] \pm O(\sqrt{\epsilon})$$
 - (e) $G_{+}(|h_{k-1}, e_k| h_k)$
 - ▷ Finally, we sum the e_k 's:

$$\mathbf{x} \left[h_{1/\sqrt{\epsilon}} \right] = \max \mathbf{x}[a_i] \pm O(\sqrt{\epsilon})$$
3. $G_{=}(|h_{1/\sqrt{\epsilon}}| b)$
-

For each k , taking OR of all the $c_{k,i}$'s gives (approximately) 1 iff any of the $\mathbf{x}[a_i]$'s is sufficiently large; in particular if the maximum is:

$$\mathbf{x}[d_{k,n}] = \begin{cases} 0 \pm \epsilon & \text{if } \max_i \mathbf{x}[a_i] < k\sqrt{\epsilon} - O(\epsilon) \\ 1 \pm \epsilon & \text{if } \max_i \mathbf{x}[a_i] > k\sqrt{\epsilon} + O(\epsilon). \end{cases}$$

Therefore also (similarly to (6.2)):

$$\sum_k \mathbf{x}[d_{k,n}] \sqrt{\epsilon} = \max_i \mathbf{x}[a_i] \pm O(\sqrt{\epsilon}).$$

The $\mathbf{x}[e_k]$'s take care of scaling by $\sqrt{\epsilon}$:

$$\sum_k \mathbf{x}[e_k] = \max_i \mathbf{x}[a_i] \pm O(\sqrt{\epsilon}).$$

Finally, by induction,

$$\mathbf{x}[h_{1/\sqrt{\epsilon}}] = \max \mathbf{x}[a_i] \pm O(\sqrt{\epsilon}). \quad \blacksquare$$

Algorithm 6.5 INTERPOLATE($a, w_a, b, w_b \mid c$)

1. $G_{\times\zeta}(1/2 \mid w_a \mid w_{a/2})$ and $G_{\times\zeta}(1/2 \mid w_b \mid w_{b/2})$
 \triangleright We divide by 2 before adding in order to stay in $[0, 1]$
 2. $G_+(\mid w_{a/2}, w_{b/2} \mid w_{a/2+b/2})$
 \triangleright Add the weights
 3. DIVIDE($\mid w_{a/2}, w_{a/2+b/2} \mid \bar{w}_a$) and DIVIDE($\mid w_{b/2}, w_{a/2+b/2} \mid \bar{w}_b$)
 $\triangleright \bar{w}_a$ and \bar{w}_b are the normalized weights:

$$\mathbf{x}[\bar{w}_a] \cdot (\mathbf{x}[w_a] + \mathbf{x}[w_b]) = \mathbf{x}[w_a] \pm O(\sqrt{\epsilon})$$
 4. MULTIPLY($\mid \bar{w}_a, a \mid c_a$) and MULTIPLY($\mid \bar{w}_b, b \mid c_b$)
 $\triangleright c_a$ and c_b are the a and b components, respectively, of c :

$$\mathbf{x}[c_a] = \mathbf{x}[w_a] \cdot \mathbf{x}[a] / (\mathbf{x}[w_a] + \mathbf{x}[w_b]) \pm O(\sqrt{\epsilon}) / (\mathbf{x}[w_a] + \mathbf{x}[w_b])$$
 5. $G_+(\mid c_a, c_b \mid c)$
 \triangleright Finally, c is the interpolation of a and b :

$$\mathbf{x}[c] = (\mathbf{x}[w_a] \cdot \mathbf{x}[a] + \mathbf{x}[w_b] \cdot \mathbf{x}[b]) / (\mathbf{x}[w_a] + \mathbf{x}[w_b]) \pm O(\sqrt{\epsilon}) / (\mathbf{x}[w_a] + \mathbf{x}[w_b])$$
-

6.2.1.5 Interpolate

Claim 6.5 In any ϵ -approximate solution to INTERPOLATE($a, w_a, b, w_b \mid c$),

$$\mathbf{x}[c](\mathbf{x}[w_a] + \mathbf{x}[w_b]) = (\mathbf{x}[w_a] \cdot \mathbf{x}[a] + \mathbf{x}[w_b] \cdot \mathbf{x}[b]) \pm O(\sqrt{\epsilon}).$$

Proof By Claim 6.3, we have

$$\begin{aligned} \mathbf{x}[\bar{w}_a] \cdot (\mathbf{x}[w_a] + \mathbf{x}[w_b]) &= \mathbf{x}[w_a] \pm O(\sqrt{\epsilon}) \\ \mathbf{x}[\bar{w}_b] \cdot (\mathbf{x}[w_a] + \mathbf{x}[w_b]) &= \mathbf{x}[w_b] \pm O(\sqrt{\epsilon}). \end{aligned}$$

Therefore, by Claim 6.2,

$$\begin{aligned} \mathbf{x}[c_a] \cdot (\mathbf{x}[w_a] + \mathbf{x}[w_b]) &= \mathbf{x}[w_a] \cdot \mathbf{x}[c] \pm O(\sqrt{\epsilon}) \\ \mathbf{x}[c_b] \cdot (\mathbf{x}[w_a] + \mathbf{x}[w_b]) &= \mathbf{x}[w_b] \cdot \mathbf{x}[c] \pm O(\sqrt{\epsilon}). \end{aligned}$$

The claim follows by definition of G_+ . ■

6.2.2 Equiangle Sampling Segment

The first information we require in order to compute the Hirsch et al. mapping $f(\mathbf{y})$ is about the subcube to which \mathbf{y} belongs: Is it part of the tube? If so, which are the entrance and exit facets? In order to answer those questions, we extract the binary representation of the cube. Recall that our circuit uses brittle comparators; thus

when \mathbf{y} is close to a facet between subcubes, the behavior of the brittle comparators may be unpredictable. We start with the easy case, where \mathbf{y} is actually far from every facet:

Definition 6.3 We say that \mathbf{y} is an *interior point* if for every i , $|y_i - 1/2| > \epsilon$; otherwise, we say that \mathbf{y} is a *boundary point*.

A very nice property of the Hirsch et al. construction is that whenever \mathbf{y} is at the intersection of two or more facets, the displacement is the same: $g(\mathbf{y}) = \delta \xi_{2n+2}$. Thus, by the Lipschitz property of g , whenever \mathbf{y} is close to the intersection of two or more facets, the displacement is approximately $\delta \xi_{2n+2}$. For such \mathbf{y} 's, we don't care to which subcube they belong.

Definition 6.4 We say that \mathbf{y} is a *corner point* if there exist distinct $i, j \in [2n + 2]$ such that $|y_i - 1/2| < \epsilon^{1/4}$ and $|y_j - 1/2| < \epsilon^{1/4}$.

(Notice that \mathbf{y} may be an interior point and a corner point at the same time.)

We still have a hard time handling \mathbf{y} 's that are neither an interior point nor a corner point. To mitigate the effect of such \mathbf{y} 's we use an equiangle averaging scheme. Namely, we consider the set:

$$E^\epsilon(\mathbf{y}) = \{\mathbf{y}^l = \mathbf{y} + (6l \cdot \epsilon)\mathbf{1} : 0 \leq l < 1/\sqrt{\epsilon}\}$$

where $\mathbf{1}$ denotes the all-ones vector. Notice that since g is λ -Lipschitz for constant λ , $g(\mathbf{y}^l)$ will be approximately the same for all $\mathbf{y}^l \in E^\epsilon(\mathbf{y})$.

Fact 6.1 If any $\mathbf{y}^l \in E^\epsilon(\mathbf{y})$ is not a corner point, then at most one $\mathbf{y}^l \in E^\epsilon(\mathbf{y})$ is a boundary point.

Proof For each dimension, at most one element in $E^\epsilon(\mathbf{y})$ can be ϵ -close to the $(1/2)$ facet. Thus if two elements in $E^\epsilon(\mathbf{y})$ are boundary points, it must be because of distinct dimensions—and therefore every \mathbf{y}^l is a corner point. ■

Given input \mathbf{y} , we compute the displacement $g(\cdot)$ separately and in parallel for each $\mathbf{y}^l \in E^\epsilon$, and average at the end. Since at most one \mathbf{y}^l is a boundary point, this will incur an error of at most $\sqrt{\epsilon}$.

In the generalized circuit we can construct E^ϵ using $(1/\sqrt{\epsilon})$ auxiliary nodes and G_ζ and G_+ gates:

$$\mathbf{x} \left[y_i^l \right] = \min \{ \mathbf{x} [y_i^0] + (6l \cdot \epsilon), 1 \} \pm 2\epsilon.$$

6.2.3 Computing the Displacement

For each $\mathbf{y}^l \in E^\epsilon$, we construct a disjoint circuit that approximates the displacement $g(\mathbf{y}^l)$. In the description of the circuit below we omit the index l .

Lemma 6.1 The circuit below $O(\sqrt{\epsilon})$ -approximately simulates the computation of the Hirsch et al. displacement:

1. Whenever $(\mathbf{x}[y_i])_{i \in [2n+2]}$ is an interior point,

$$\mathbf{x}[g_i^+] - \mathbf{x}[g_i^-] = g_i(\mathbf{x}[\mathbf{y}]) \pm O(\epsilon^{1/4}).$$

2. Furthermore, whenever $(\mathbf{x}[y_i])_{i \in [2n+2]}$ is a corner point,

$$\mathbf{x}[g_{2n+2}^+] - \mathbf{x}[g_{2n+2}^-] = \delta \pm O(\sqrt{\epsilon}),$$

and $\forall i < 2n + 2$:

$$\mathbf{x}[g_i^+] - \mathbf{x}[g_i^-] = 0 \pm O(\sqrt{\epsilon}).$$

Proof We construct the circuit in five stages: (1) given \mathbf{y} , we extract \mathbf{b} , that is the binary representation of the corresponding subcube in $\{0, 1\}^{2n+2}$; (2) we then compute whether \mathbf{b} belongs to a path in H , and if so which are the previous and next vertices; (3) we compute the centers of the coordinate systems corresponding to the entrance and exit facets, and label them \mathbf{z}^{in} and \mathbf{z}^{out} ; (4) we project \mathbf{y} to each facet, and transform this projection to the local polar coordinate systems, $(r^{in}, \mathbf{p}^{in})$; and (5) finally, we use all the information above to compute the displacement $\mathbf{g} = g(\mathbf{y})$.

The correctness of Lemma 6.1 follows from Claims 6.6–6.12.

Extract $\mathbf{b} \in \{0, 1\}^{2n+2}$

Our first step is to extract the binary vector \mathbf{b} , which represents the subcube to which \mathbf{y} belongs. In other words, we want b_i to be the indicator of $y_i < 1/2$. We do that by adding the following gadgets: $G_{\zeta}(1/2 \parallel c_{1/2})$ and, for each i , $G_{<}(y_i, c_{1/2} \mid b_i)$. Observe that now

$$\mathbf{x}[b_i] = \begin{cases} 0 \pm \epsilon & \mathbf{x}[y_i] < \mathbf{x}[c_{1/2}] - \epsilon \\ 1 \pm \epsilon & \mathbf{x}[y_i] > \mathbf{x}[c_{1/2}] + \epsilon. \end{cases}$$

Claim 6.6 If $\mathbf{x}[\mathbf{y}]$ is an interior point, $\mathbf{x}[\mathbf{b}]$ is the correct representation (up to ϵ error) of the corresponding bits in $\{0, 1\}^{2n+2}$.

Neighbors in H

Given $\mathbf{x}[\mathbf{b}]$, we can construct, using G_{\wedge} 's and G_{-} 's and a polynomial number of unused nodes, the circuits S^H and P^H that give the next and previous vertex visited

by our collection of paths, H . The output of each circuit is represented by $2n + 2$ unused nodes $\{P_i^H(\mathbf{b})\}$ and $\{S_i^H(\mathbf{b})\}$.

Recall that H is defined in $\{0, 1\}^{2n+1}$, so the last input bit is simply ignored (inside the picture it is always 0); the last output bit is used as follows. Our convention is that starting points and end points correspond to $P^H(\mathbf{b}) = \mathbf{b}$ and $S^H(\mathbf{b}) = \mathbf{b}$, respectively, and likewise for points that do not belong to any path. An exception to this is the $\mathbf{0}$ starting point, which will correspond to $P^H(\mathbf{0}) = (0^{2n+1}; 1)$: This is in accordance with the Hirsch et al. construction, where the home subcube is constructed as if it continues a path from the subcube above it.

Claim 6.7 If $\mathbf{x}[\mathbf{b}]$ is an ϵ -approximate binary vector, i.e., $\mathbf{x}[\mathbf{b}] \in ([0, \epsilon] \cup [1 - \epsilon, 1])^{2n+2}$, then $\mathbf{x}[P^H(\mathbf{b})]$ and $\mathbf{x}[S^H(\mathbf{b})]$ correctly represent (up to ϵ error) the previous vertex and next vertex in H .

Entrance and Exit Facets

Let $b_i^{+in} = b_i \wedge \neg P_i^H(\mathbf{b})$, i.e., b_i^{+in} is 1 if the path enters the current subcube via the positive i -th direction; define b_i^{-in} analogously. Let b_i^{in} denote the OR of b_i^{+in} and b_i^{-in} .

The center of the entrance facet is constructed via G_ζ , $G_{\times\zeta}$, G_+ , and G_- according to the formula:

$$z_i^{in} = \begin{cases} 1/2 - h/2 & b_i^{in} = 0 \text{ AND } b_i = 0 \\ 1/2 + h/2 & b_i^{in} = 0 \text{ AND } b_i = 1 \\ 1/2 & b_i^{in} = 1 \end{cases}$$

Construct \mathbf{z}^{out} analogously.

Notice that if we know on which coordinate the path enters, in $\{0, 1\}^{2n+2}$ it has only one possible direction; in the Hirsch et al. hypercube this corresponds to always entering from the center (i.e., from the $y_i = 1/2$ facet). Also, if \mathbf{b} corresponds to a non-trivial starting point, $\mathbf{b}^{in} = \mathbf{0}$ and \mathbf{z}^{in} is simply the center of the subcube (and similarly for \mathbf{b}^{out} , \mathbf{z}^{out} when \mathbf{b} is an end point).

Claim 6.8 If $\mathbf{x}[\mathbf{b}]$, $\mathbf{x}[P^H(\mathbf{b})]$, and $\mathbf{x}[S^H(\mathbf{b})]$ are ϵ -approximate binary vectors, then $\mathbf{x}[\mathbf{z}^{in}]$ and $\mathbf{x}[\mathbf{z}^{out}]$ are $O(\epsilon)$ -approximations to the centers of the entrance facet and exit facets, respectively.

Max-Norm Polar Coordinates

We are now ready to compute the local max-norm polar coordinates of the projections of \mathbf{y} on the entrance and exit facets.

The max-norm radius is given by

$$r^{in} = \max_{i:b_i^{in}=0} |z_i^{in} - y_i|.$$

Finding the maximum of a (length $2n + 1$) vector requires some care when on each gate we can incur a constant error; the details are described in Algorithm 6.4.

The direction (max-norm) unit-vector, \mathbf{p} , is given by

$$p_i^{in} = (z_i^{in} - y_i) / r^{in}.$$

Division is computed using DIVIDE, introducing an error of $O(\sqrt{\epsilon}/r^{in})$; this approximation suffices because for $r^{in} < h/8$, we multiply p_i^{in} by r^{in} when we interpolate. Also, we will use two nodes for each p_i^{in} to represent the positive and negative values. We do the same for $(r^{out}, \mathbf{p}^{out})$.

Claim 6.9 If $\mathbf{x}[\mathbf{y}]$ is an interior point, then $\mathbf{x}[r^{in}]$ and $\mathbf{x}[r^{out}]$ are $O(\sqrt{\epsilon})$ -approximations to the distances of $\mathbf{x}[\mathbf{y}]$ from $\mathbf{x}[\mathbf{z}^{in}]$ and $\mathbf{x}[\mathbf{z}^{out}]$, respectively. Furthermore, $\mathbf{x}[\mathbf{p}^{in}]$ and $\mathbf{x}[\mathbf{p}^{out}]$ are $O(\sqrt{\epsilon}/\mathbf{x}[r^{in}])$ and $O(\sqrt{\epsilon}/\mathbf{x}[r^{out}])$ to the unit-length vectors that point from $\mathbf{x}[\mathbf{y}]$ in the directions of $\mathbf{x}[\mathbf{z}^{in}]$ and $\mathbf{x}[\mathbf{z}^{out}]$, respectively.

The Final Displacement

Given \mathbf{p}^{in} and \mathbf{b}^{in} , we can compute \mathbf{g}^{in} for the special values of r^{in} . Recall that

$$\mathbf{g}^{in}(r^{in}, \mathbf{p}^{in}) = \begin{cases} \delta(\mathbf{b}^{+in} - \mathbf{b}^{-in}) & r^{in} = 0 \\ -\delta\mathbf{p}^{in} & r^{in} = h/8 \\ \delta(\mathbf{b}^{-in} - \mathbf{b}^{+in}) & r^{in} = h/4 \\ \delta\xi_{2n+2} & r^{in} = h/2 \end{cases}$$

We use Algorithm 6.5 to interpolate for intermediate values of r^{in} . We also need to interpolate between \mathbf{g}^{in} and \mathbf{g}^{out} . The ratio at which we interpolate is exactly the ratio between the distance of \mathbf{y} from the entrance and exit facets. We label the positive and negative components of this last interpolation $\mathbf{g}^{(interior)+}$ and $\mathbf{g}^{(interior)-}$, respectively.

When \mathbf{y} is close to both facets, the interpolation may be inaccurate; however, in this case it is a corner point. (We remark that this seems to be the only part of the proof that requires us to set the threshold for a corner point and the final error at $\Theta(\epsilon^{1/4})$ rather than $\Theta(\sqrt{\epsilon})$; this issue may be avoidable by a more careful construction of Algorithms 6.3 and 6.5.)

Claim 6.10 If $\mathbf{x}[\mathbf{y}]$ is an interior point of an intermediate subcube in the tube, and it is not a corner point, then $(\mathbf{x}[\mathbf{g}^{(interior)+}] - \mathbf{x}[\mathbf{g}^{(interior)-}])$ is an $O(\epsilon^{1/4})$ -approximation of the Hirsch et al. displacement $g(\mathbf{x}[\mathbf{y}])$.

Corner Points. We must ensure that if \mathbf{y} is a corner point, we set $\mathbf{g}^{(corner)+} = \delta\xi_{2n+2}$ and $\mathbf{g}^{(corner)-} = \mathbf{0}$: check over all pairs of coordinates whether $|\mathbf{x}[y_i] - 1/2| < 2\epsilon^{1/4}$ and $|\mathbf{x}[y_j] - 1/2| < 2\epsilon^{1/4}$. Let z be the variable representing the OR of those $\binom{2n+2}{2}$ indicators. Interpolate (e.g., using Algorithm 6.5) between the $(\mathbf{g}^{(interior)+}, \mathbf{g}^{(interior)-})$ we computed earlier and $\delta\xi_{2n+2}$ with weights z and $\neg z$. Label the result of the interpolation $(\mathbf{g}^{(tube)+}, \mathbf{g}^{(tube)-})$.

We remark that whenever z is ambiguous, i.e., the second smallest $|\mathbf{x}[y_i] - 1/2|$ is very close to $2\epsilon^{1/4}$, then we cannot predict the value of $\mathbf{x}[z]$; it can take any value in $[0, 1]$. Nevertheless, in this case $\mathbf{x}[\mathbf{y}]$ is not a corner point, thus for most $\mathbf{y}^l \in E^\epsilon$, $\mathbf{x}[\mathbf{y}^l]$ is an interior point. This means that by Claim 6.10, we would compute the (approximately) correct interior displacement $(\mathbf{x}[\mathbf{g}^{(interior)+}] - \mathbf{x}[\mathbf{g}^{(interior)-}])$. Since $\mathbf{x}[\mathbf{y}]$ is close to a corner point, this value is very close to $\delta\xi_{2n+2} = (\mathbf{x}[\mathbf{g}^{(corner)+}] - \mathbf{x}[\mathbf{g}^{(corner)-}])$. Therefore, although we don't know the value of $\mathbf{x}[z]$, we use it to interpolate between two (approximately) equal vectors—so the result is guaranteed to be (approximately) correct regardless of the value of $\mathbf{x}[z]$.

Claim 6.11 If $\mathbf{x}[\mathbf{y}]$ is a corner point, then $(\mathbf{x}[\mathbf{g}^{(tube)+}] - \mathbf{x}[\mathbf{g}^{(tube)-}])$ is an $O(\sqrt{\epsilon})$ -approximation of $\delta\xi_{2n+2}$, and thus also an $O(\epsilon^{1/4})$ -approximation of the Hirsch et al. displacement $g(\mathbf{x}[\mathbf{y}])$. Furthermore, Claim 6.10 continues to hold for $(\mathbf{g}^{(tube)+}, \mathbf{g}^{(tube)-})$.

Start/End Subcubes and Subcubes Outside the Tube. For start/end subcubes (except the home subcube), we use a slightly different (Cartesian) interpolation that yields a fixed point in the center of the subcube, and a displacement of $\delta\xi_{2n+2}$ on all facets but the exit/entrance facet, respectively. For subcubes in the picture but outside the tube, we again set $\mathbf{g} = \delta\xi_{2n+2}$.

Notice that we can infer the type of subcube from the following two-bits vector:

$$T = (\vee_i b_i^{in}, \vee_i b_i^{out}).$$

If $T = (0, 0)$, the subcube is outside the tube; when $T = (0, 1)$, we are at a start subcube, while $T = (1, 0)$ corresponds to an end subcube; $T = (1, 1)$ is an intermediate subcube in the tube. Finally, interpolate between the displacement for each type of subcube using T and $\neg T$; label the result of the interpolation $(\mathbf{g}^+, \mathbf{g}^-)$.

Claim 6.12 If $\mathbf{x}[\mathbf{y}]$ is either an interior point or a corner point of any subcube in the slice, then $(\mathbf{x}[g^+] - \mathbf{x}[g^-])$ is an $O(\epsilon^{1/4})$ -approximation of the Hirsch et al. displacement $g(\mathbf{x}[\mathbf{y}])$. ■

6.2.4 Summing the Displacement Vectors

We are now ready to average over the displacement vectors we computed for each \mathbf{y}^l . Using $G_{\times\zeta}$ and G_+ we have that

$$\mathbf{x}[g_i^+] = \sum_{l=1}^{1/\sqrt{\epsilon}} (\sqrt{\epsilon} \mathbf{x}[g_i^{l+}]) \pm O(\sqrt{\epsilon}) \quad \text{and} \quad \mathbf{x}[g_i^-] = \sum_{l=1}^{1/\sqrt{\epsilon}} (\sqrt{\epsilon} \mathbf{x}[g_i^{l-}]) \pm O(\sqrt{\epsilon}).$$

Lemma 6.2 For every input $\mathbf{x}[\mathbf{y}]$ and every $i \in [2n + 2]$,

$$\mathbf{x}[g_i^+] - \mathbf{x}[g_i^-] = g_i(\mathbf{x}[\mathbf{y}]) \pm O(\epsilon^{1/4}).$$

Proof By Fact 6.1, every $\mathbf{y}^l \in E^\epsilon$, except at most one, is either an interior point or a corner point. By Lemma 6.1, for all those \mathbf{y}^l , $(\mathbf{x}[g_i^{l+}] - \mathbf{x}[g_i^{l-}])$ is at most $O(\epsilon^{1/4})$ -far from the right displacement. The single point, which is neither an interior point nor a corner point, increases the error by at most $O(\sqrt{\epsilon})$, as does the summation above. Because g is λ -Lipschitz for constant λ , the error added due to the distance between the \mathbf{y}^l 's is again at most $O(\sqrt{\epsilon})$. ■

6.2.5 Closing the Loop

Finally, for each $i \in [2n + 2]$: $G_+(|y_i^1, g_i^+ | y_i^1)$, $G_+(|y_i^1, g_i^- | y_i^1)$ and $G_=(|y_i^1 | y_i^1)$.

6.3 GCIRCUIT with Fan-out 2

In the previous section, we proved that ϵ -GCIRCUIT is PPAD-complete for some constant $\epsilon > 0$. Each generalized circuit gate has fan-in at most 2, which would eventually correspond to a bound on the *incoming* degree of each player in the graphical game. In order to bound the total degree (as well as for the application to A-CEEI), we need to also bound fan-out of each gate.

Proof Proof of Theorem 5.3. We present a black-box reduction that converts a general ϵ' -GCIRCUIT instance to an instance of ϵ -GCIRCUIT with fan-out 2, for $\epsilon' = \Theta(\sqrt{\epsilon})$. Daskalakis et al. [2009a] bound the fan-out of the generalized circuit by introducing a binary tree of $G_=($ gates. Unfortunately, this blows up the error: each $G_=($ gate introduces an additive error of ϵ , so the increase is proportional to the depth of the tree, i.e., $\Theta(\epsilon \cdot \log n)$. While this was acceptable for Daskalakis et al. [2009a] who used an exponentially small ϵ , it clearly fails in our setting.

Algorithm 6.6 REAL2UNARY($| a | b_1, \dots, b_{4/\epsilon'}$)

1. $G_{=}(| a | c_0)$
 2. **foreach** $k \in [4/\epsilon']$:
 - (a) $G_{=}(| c_{k-1} | c_k)$
 - ▷ The c_k 's are simply copies of a , ensuring that each gate has fan-out at most 2:

$$\forall k \mathbf{x}[c_k] = \mathbf{x}[a] \pm (k + 1)\epsilon$$
 - (b) $G_{\zeta}(k\epsilon'/4 || \zeta_k)$
 - (c) $G_{<}(| \zeta_k, c_k | b_k)$
 - ▷ The vector (b_k) is the unary representation of a :

$$\forall i \left(\max_{k: \mathbf{x}[b_k] > \epsilon} k\epsilon'/4 \right) = \mathbf{x}[a] \pm \epsilon'/2$$
-

Algorithm 6.7 UNARY2REAL($| b_1, \dots, b_{4/\epsilon'} | a'$)

1. $G_{\zeta}(0 || d_0)$
 2. **foreach** $k \in [4/\epsilon']$:
 - (a) $G_{\times\zeta}(\epsilon'/4 | b_k | c_k)$
 - ▷ The sum of the c_k 's is approximately a :

$$\sum \mathbf{x}[c_k] = \mathbf{x}[a] \pm (\epsilon'/2 + \epsilon \cdot 4/\epsilon')$$
 - (b) $G_{+}(| d_{k-1}, c_k | d_k)$
 3. $G_{=}(| d_k | a')$
 - ▷ a' approximately recovers the original a :

$$\mathbf{x}[a'] = \mathbf{x}[a] \pm \epsilon'$$
-

We overcome this obstacle by resorting to logical gates ($G_{<}$, G_{\vee} , G_{\wedge} , and in particular G_{-}). Recall that the logical gates are at most ϵ' -far from the correct Boolean value. Therefore, concatenating multiple logical gates does not amplify the error. In particular, if any logical gate has a large fan-out, we can distribute its output using a binary tree of G_{-} gates (we use trees of even depth). When the gate is arithmetic (G_{ζ} , $G_{\times\zeta}$, $G_{=}$, G_{+} , or G_{-}), we convert its output to unary representation over $\Theta(\epsilon')$ logical gates (Algorithm 6.6). Then, we copy the unary representation using trees of G_{-} gates. Finally, we use $G_{\times\zeta}$ and G_{+} gates to convert each copy of the unary representation back to a real value in $[0, 1]$ (Algorithm 6.7).

It is interesting to note that for constant ϵ and ϵ' , the unary representation has constant size, so the number of new gates is proportional to the original fan-out (i.e., the number of leaves of the binary tree that copies the unary representation). In particular, this reduction increases the size of the circuit by a factor of $\Theta(1/\epsilon')$. ■

Many-Player Games

In this chapter we prove our hardness results for games with a large number of players. We begin in Section 7.1 with our result for the more restricted class of graphical and polymatrix games, and then in Section 7.2 prove a stronger hardness of approximation for the more general class of succinct games. The two results follow relatively easily from Theorems 5.3 and 4.1, respectively.

Related Works: Tractable Special Cases

For games with many players and a constant number of strategies, PTAS were given for the special cases of anonymous games by [Daskalakis and Papadimitriou \[2015\]](#) and polymatrix games on a tree by [Barman et al. \[2015\]](#). Finally, let us return to the more general class of succinct n -player games, and mention an approximation algorithm due to [Goldberg and Roth \[2016\]](#); their algorithm runs in exponential time, but uses only a polynomial number of oracle queries.

7.1 Graphical, Polymatrix Games

Theorem 7.1 Theorem 5.1 restated. There exists a constant $\epsilon > 0$, such that given a degree 3, bipartite, polymatrix game where each player has two actions, finding an ϵ -approximate Nash equilibrium is PPAD-complete.

The proof proceeds in two steps. First, we reduce to the problem of finding an ϵ -well-supported Nash equilibrium in a polymatrix (degree 3, bipartite) graphical game. This reduction is implicit in [Daskalakis et al. \[2009a\]](#).¹ In the second step we use the fact that our graphical games have a constant degree to reduce to the more lenient notion of ϵ -approximate Nash equilibrium.

1. In [Daskalakis et al. \[2009a\]](#), polymatrix games are called *games with additive utility functions*.

7.1.1 Hardness of Well-Supported Nash Equilibrium

Proposition 7.1 Essentially Daskalakis et al. [2009a]. 2ϵ -GCIRCUIT with fan-out 2 is polynomial-time reducible to finding an ϵ -well-supported Nash equilibrium in a graphical (bipartite, degree 3), polymatrix binary action game.

The reduction closely follows the lines of Daskalakis et al. [2009a]. We replace each gate in the generalized circuit with a *subgame* over a constant number of players. Each subgame has one or more designated *input players*, a single *output player*, and potentially some *auxiliary players*. The utility of the input players does not depend on the subgame at all. The guarantee is that in any ϵ -WSNE of the subgame, the mixed strategy profile $O(\epsilon)$ -approximately implements the gate, that is, the probability that the output player assigns to strategy 1 is (approximately) equal to the gate function applied to the corresponding probabilities for the input players.

Finally, we concatenate all the subgames together, i.e., we identify between the output player of one gadget and an input player of (at most two) other gadgets. Any ϵ -WSNE in the resulting large game $O(\epsilon)$ -approximately implements all the gates in the generalized circuit—so we can extract valid assignments to all the generalized circuit lines from the probability that each output player assigns to strategy 1.

In the gadget constructions below, every input and output player has degree 1 (i.e., it only interacts with one other player, in the sense of Definition 5.1). Each player participates in at most one gadget as an output player, and at most two gadgets as an input player (since the circuit has maximum fan-out 2). Therefore, the total degree of each player is ≤ 3 .

Below we construct the gadgets for each of the gates. By the above argument, Lemmata 7.1–7.4 together imply Proposition 7.1.

Lemma 7.1 $G_{\times\zeta}$, G_ζ , G_+ , $G_=-$ gadgets. Let $\alpha, \beta > 0$. Let v_{IN1} , v_{IN2} , v_{OUT} , w be players in a graphical game, and suppose that the payoffs of v_{OUT} and w are as follows.

Payoff for v_{OUT}:		w plays 0	w plays 1
v_2 plays 0		0	1
v_{OUT} plays 1		1	0

Payoffs for w:		v_{IN1} plays 0	v_{IN1} plays 1
game with v_{IN1} :	w plays 0	0	α
	w plays 1	0	0

		v_{IN2} plays 0	v_{IN2} plays 1
game with v_{IN2} :	w plays 0	0	β
	w plays 1	0	0

		v_{OUT} plays 0	v_{OUT} plays 1
game with v_{OUT} :	w plays 0	0	0
	w plays 1	0	1

Then, in every ϵ -WSNE,

$$\mathbf{p}[v_{\text{OUT}}] = \min(\alpha \mathbf{p}[v_{\text{IN1}}] + \beta \mathbf{p}[v_{\text{IN2}}], 1) \pm \epsilon, \quad (7.1)$$

where $\mathbf{p}[u]$ denotes the probability that u assigns to strategy 1.

Notice that each of $G_{\times\zeta}$, G_ζ , G_+ , G_- can be implemented using (7.1): for $G_{\times\zeta}$ set $\alpha = \zeta$, $\beta = 0$; for G_ζ set the same α , β , and modify w 's game with v_{IN1} as if the latter always plays strategy 1; etc.

Proof Proof of Lemma 7.1. Assume by contradiction that (7.1) is violated:

- If $\mathbf{p}[v_{\text{OUT}}] > \alpha \mathbf{p}[v_{\text{IN1}}] + \beta \mathbf{p}[v_{\text{IN2}}] + \epsilon$, then in every ϵ -WSNE $\mathbf{p}[w] = 1$. But then v_{OUT} is incentivized to play strategy 0, which contradicts $\mathbf{p}[v_{\text{OUT}}] > \epsilon$.
- Similarly, if $\mathbf{p}[v_{\text{OUT}}] < \alpha \mathbf{p}[v_{\text{IN1}}] + \beta \mathbf{p}[v_{\text{IN2}}] - \epsilon$, then in every ϵ -WSNE $\mathbf{p}[w] = 0$. Therefore, $\mathbf{p}[v_{\text{OUT}}] \geq 1 - \epsilon$. ■

Lemma 7.2 G_+ , G_- gadgets. Let v_{IN1} , v_{IN2} , v_{OUT} , w be players in a graphical game, and suppose that the payoffs of v_{OUT} and w are as follows.

Payoff for v_{OUT}:		w plays 0	w plays 1
v_2 plays 0		0	1
v_{OUT} plays 0		1	0

Payoffs for w:		v_{IN1} plays 0	v_{IN1} plays 1
game with v_{IN1} :	w plays 0	0	1
	w plays 1	0	0

		v_{IN2} plays 0	v_{IN2} plays 1
game with v_{IN2} :	w plays 0	1	0
	w plays 1	0	0
		v_{OUT} plays 0	v_{OUT} plays 1

game with v_{OUT} :	w plays 0	0	0
	w plays 1	0	1

Then, in every ϵ -WSNE,

$$\mathbf{p}[v_{\text{OUT}}] = \max(\mathbf{p}[v_{\text{IN1}}] - \mathbf{p}[v_{\text{IN2}}], 0) \pm \epsilon, \tag{7.2}$$

where $\mathbf{p}[u]$ denotes the probability that u assigns to strategy 1.

Notice that each of $G_{>}, G_{<}$ can be implemented using (7.2).

Proof Proof of Lemma 7.2. Assume by contradiction that (7.2) is violated:

- If $\mathbf{p}[v_{\text{OUT}}] > \mathbf{p}[v_{\text{IN1}}] - \mathbf{p}[v_{\text{IN2}}] + \epsilon$, then in every ϵ -WSNE $\mathbf{p}[w] = 1$. But then v_{OUT} is incentivized to play strategy 0, which contradicts $\mathbf{p}[v_{\text{OUT}}] > \epsilon$.
- Similarly, if $\mathbf{p}[v_{\text{OUT}}] < \mathbf{p}[v_{\text{IN1}}] - \mathbf{p}[v_{\text{IN2}}] - \epsilon$, then in every ϵ -WSNE $\mathbf{p}[w] = 0$. Therefore, $\mathbf{p}[v_{\text{OUT}}] \geq 1 - \epsilon$. ■

Lemma 7.3 $G_{<}$ gadget. Let $v_{\text{IN1}}, v_{\text{IN2}}, v_{\text{OUT}}, w$ be players in a graphical game, and suppose that the payoffs of v_{OUT} and w are as follows.

Payoff for v_{OUT}:		w plays 0	w plays 1
	v_{OUT} plays 0	0	1
	v_{OUT} plays 1	1	0

Payoffs for w:		v_{IN1} plays 0	v_{IN1} plays 1
game with v_{IN1} :	w plays 1	0	0
	w plays 0	1	1

		v_{IN2} plays 0	v_{IN2} plays 1
game with v_{IN2} :	w plays 0	1	0
	w plays 1	0	0

Then, in every ϵ -WSNE,

$$\mathbf{p}[v_{\text{OUT}}] = \begin{cases} \geq 1 - \epsilon & \mathbf{p}[v_{\text{IN1}}] < \mathbf{p}[v_{\text{IN2}}] - \epsilon \\ \leq \epsilon & \mathbf{p}[v_{\text{IN1}}] > \mathbf{p}[v_{\text{IN2}}] + \epsilon, \end{cases} \tag{7.3}$$

where $\mathbf{p}[u]$ denotes the probability that u assigns to strategy 1.

Proof The payoff for player w for action 0 is $1 + \mathbf{p}[v_{\text{IN}1}] - \mathbf{p}[v_{\text{IN}2}]$, whereas for action 1 it is $1 - \mathbf{p}[v_{\text{IN}1}] + \mathbf{p}[v_{\text{IN}2}]$. Therefore, in every ϵ -WSNE,

$$\mathbf{p}[w] = \begin{cases} \leq \epsilon & \mathbf{p}[v_{\text{IN}1}] < \mathbf{p}[v_{\text{IN}2}] - \epsilon \\ \geq 1 - \epsilon & \mathbf{p}[v_{\text{IN}1}] > \mathbf{p}[v_{\text{IN}2}] + \epsilon. \end{cases}$$

(7.3) follows immediately because player v_{OUT} wants to play the opposite of w . ■

Lemma 7.4 G_{\vee}, G_{\wedge} gadgets. Let $v_{\text{IN}1}, v_{\text{IN}2}, v_{\text{OUT}}, w$ be players in a graphical game, and suppose that the payoffs of v_{OUT} and w are as follows.

Payoff for v_{OUT}:		w plays 0	w plays 1
v_{OUT} plays 0		0	1
v_{OUT} plays 1		1	0

Payoffs for w:		$v_{\text{IN}0}$ plays 1	$v_{\text{IN}1}$ plays 1
game with $v_{\text{IN}1}$:	w plays 0	1	0
	w plays 1/2	0	1

		$v_{\text{IN}2}$ plays 0	$v_{\text{IN}2}$ plays 1
game with $v_{\text{IN}2}$:	w plays 0	1	0
	w plays 1/2	0	1

Then, in every ϵ -WSNE,

$$\mathbf{p}[v_{\text{OUT}}] = \begin{cases} \geq 1 - \epsilon & \mathbf{p}[v_{\text{IN}1}] + \mathbf{p}[v_{\text{IN}2}] < 2/3 - \epsilon \\ \leq \epsilon & \mathbf{p}[v_{\text{IN}1}] + \mathbf{p}[v_{\text{IN}2}] > 2/3 + \epsilon, \end{cases} \quad (7.4)$$

where $\mathbf{p}[u]$ denotes the probability that u assigns to strategy 1.

Notice that (7.4) implements a G_{\vee} (OR) gadget when the input values are approximately Boolean. A G_{\wedge} (AND) gadget can be implemented analogously by interchanging the 1 and 1/2 values in w 's payoff matrices.

Proof Proof of Lemma 7.4. The payoff for player w for action 0 is $\mathbf{p}[v_{\text{IN}1}] + \mathbf{p}[v_{\text{IN}2}]$, whereas for action 1 it is $(1 - \mathbf{p}[v_{\text{IN}1}])/2 + (1 - \mathbf{p}[v_{\text{IN}2}])/2 = 1 - (\mathbf{p}[v_{\text{IN}1}] + \mathbf{p}[v_{\text{IN}2}])/2$. In particular, the latter payoff from action 1 is larger whenever $\mathbf{p}[v_{\text{IN}1}] + \mathbf{p}[v_{\text{IN}2}] < 2/3$.

Therefore, in every ϵ -WSNE

$$\mathbf{p}[w] = \begin{cases} \leq \epsilon & \mathbf{p}[v_{\text{IN}1}] + \mathbf{p}[v_{\text{IN}2}] < 2/3 - \epsilon \\ \geq 1 - \epsilon & \mathbf{p}[v_{\text{IN}1}] + \mathbf{p}[v_{\text{IN}2}] > 2/3 + \epsilon. \end{cases}$$

(7.4) follows immediately because player v_{OUT} wants to play the opposite of w . ■

7.1.2 From $\sqrt{\epsilon}$ -WSNE to $\Theta(\epsilon)$ -ANE

The reduction above shows hardness for the slightly stronger notion (therefore weaker hardness) of ϵ -WSNE. [Daskalakis et al. \[2009a\]](#) show a reduction from $\sqrt{\epsilon}$ -WSNE to $\Theta(\epsilon)$ -ANE for games with a constant number of players. It is easy to see that the same reduction holds for graphical games with constant degree. We sketch the proof below.

Lemma 7.5 Essentially [Daskalakis et al. \[2009a\]](#). Given an ϵ -ANE of a graphical game with payoffs in $[0, 1]$ and incoming degree d_{in} , we can construct in polynomial time a $\sqrt{\epsilon} \cdot (\sqrt{\epsilon} + 1 + 4d_{in})$ -WSNE.

Proof Let V be the set of players, where each $v \in V$ has utility U^v and action set A^v . Let $\mathbf{x} = (x_a^v) \in \Delta(\times_v A^v)$ be an ϵ -ANE. Let $U_a^v(\mathbf{x}^{-v})$ denote the expected payoff for playing a when all other players play according to \mathbf{x} . Note that $U_a^v(\mathbf{x}^{-v}) = U_a^v(\mathbf{x}^{\mathcal{N}_{in}(v)})$ depends only on the distributions of the players in the incoming neighborhood of v , which we denote $\mathcal{N}_{in}(v)$. Finally, let $U_{\max}^v(\mathbf{x}^{-v}) = \max_{a \in A^v} U_a^v(\mathbf{x}^{-v})$.

Let $k = k(\epsilon) > 0$ be some large number to be specified later. We construct our new approximate equilibrium by taking, for each player, only the strategies that are within ϵk of the optimum:

$$\hat{x}_a^v = \begin{cases} \frac{x_a^v}{1-z^v} & U_a^v(\mathbf{x}^{-v}) \geq U_{\max}^v(\mathbf{x}^{-v}) - \epsilon k, \\ 0 & \text{otherwise,} \end{cases}$$

where z^v is the total probability that player v assigns to strategies that are more than ϵk away from the optimum.

The above division is well defined because for $k > 1$, z^v is bounded away from 1. Moreover, the following claim from [Daskalakis et al. \[2009a\]](#) formalizes the intuition that when k is sufficiently large, the total weight on actions removed is small, so $\hat{\mathbf{x}}^v$ is close to \mathbf{x}^v :

Claim 7.1 Claim 6 of [Daskalakis et al. \[2009a\]](#).

$$\forall v \in V \sum_{a \in A^v} |\hat{x}_a^v - x_a^v| \leq \frac{2}{k-1}$$

Now, the total change to the expected payoff to player v for each action a is bounded by the total change in mixed strategies of its *incoming neighbors*:

$$\begin{aligned} |U_a^v(\mathbf{x}^{-v}) - U_a^v(\hat{\mathbf{x}}^{-v})| &= |U_a^v(\mathbf{x}^{\mathcal{N}_{in}(v)}) - U_a^v(\hat{\mathbf{x}}^{\mathcal{N}_{in}(v)})| \\ &\leq \sum_{\mathbf{a} \in A^{\mathcal{N}(v)}} |\hat{x}_{\mathbf{a}}^{\mathcal{N}_{in}(v)} - x_{\mathbf{a}}^{\mathcal{N}_{in}(v)}| \leq \sum_{w \in \mathcal{N}(v)} \sum_{a \in A^w} |\hat{x}_a^w - x_a^w| \leq \frac{2d_{in}}{k-1}. \end{aligned}$$

It follows that \hat{x}_a^v is a $(k\epsilon + \frac{4d_{in}}{k-1})$ -WSNE:

$$U_a^v(\hat{\mathbf{x}}^{-v}) \geq U_a^v(\mathbf{x}^{-v}) - \frac{2d_{in}}{k-1} \geq U_{\max}^v(\mathbf{x}^{-v}) - \epsilon k - \frac{2d_{in}}{k-1} \geq U_{\max}^v(\hat{\mathbf{x}}^{-v}) - \epsilon k - \frac{4d_{in}}{k-1}.$$

Finally, take $k = 1 + 1/\sqrt{\epsilon}$ to get that

$$k\epsilon + \frac{4d_{in}}{k-1} \leq \sqrt{\epsilon} \cdot (\sqrt{\epsilon} + 1 + 4d_{in}). \quad \blacksquare$$

7.2 Succinct Games

Theorem 7.2 Theorem 5.2 restated. There exist constants $\epsilon, \delta > 0$, such that finding an (ϵ, δ) -WeakNash is PPAD-hard for succinct multiplayer games where each player has a constant number of actions.

Proof Let f be the hard function guaranteed by Theorem 4.1. We construct a game with two groups of n players each. The action set of each player corresponds to $\{0, 1/k, 2/k, \dots, 1\}$ for a sufficiently large constant $k > 0$. We denote the choice of strategies for the first group $\mathbf{a} \triangleq (a_1, \dots, a_n)$, and $\mathbf{b} \triangleq (b_1, \dots, b_n)$ for the second group.

Each player (A, i) in the first group attempts to imitate the behavior of the corresponding player in the second group. Her utility is given by

$$u_i(a_i, b_i) \triangleq -|a_i - b_i|^2.$$

The second group players attempt to imitate the value of f , when applied to the vector of actions taken by all the players in the first group. The utility of the j -th player, (B, j) , is

$$v_j(b_j, \mathbf{a}) \triangleq -|f_j(\mathbf{a}) - b_j|^2,$$

where f_j denotes the j -th output of f .

Observe that the expected utility for (A, i) is given by:

$$\mathbb{E}[u_i(a_i, b_i)] = -|a_i - \mathbb{E}(b_i)|^2 - \text{Var}(b_i).$$

For any value of $\mathbb{E}(b_i)$, player (A, i) has an action $\alpha_i \in [\mathbb{E}(b_i) - 1/2k, \mathbb{E}(b_i) + 1/2k]$. Her utility when playing α_i is lower bounded by:

$$\mathbb{E}[u_i(\alpha_i, b_i)] \geq -\frac{1}{4k^2} - \text{Var}(b_i).$$

On the other hand, for any $\widehat{a}_i \notin \{\alpha_i, \alpha_i + 1/k\}$,

$$\mathbb{E}[u_i(\widehat{a}_i, b_i)] \leq -\frac{1}{k^2} - \text{Var}(b_i).$$

Therefore in every $(\delta, \delta/k^2)$ -WeakNash, it holds for all but a 2δ -fraction of i 's that player (A, i) assigns probability at least $(1 - O(\delta))$ to strategies $\{\alpha_i, \alpha_i + 1/k\}$. In particular, with high probability over the randomness of the players, it holds that all but an $O(\delta)$ -fraction of the players play one of those strategies. Therefore with high probability

$$\|\mathbf{a} - \mathbb{E}(\mathbf{b})\|_2 = O(\sqrt{\delta} + 1/k). \quad (7.5)$$

Similarly, player (B, j) (the j -th player in the second group) has an action $\beta_j \in [\mathbb{E}(f_j(\mathbf{a})) - 1/2k, \mathbb{E}(f_j(\mathbf{a})) + 1/2k]$. Therefore in every $(\delta, \delta/k^2)$ -WeakNash, it holds for all but a 2δ -fraction of j 's that player (B, j) assigns probability at least $(1 - O(\delta))$ to strategies $\{\beta_j, \beta_j + 1/k\}$. In particular, with high probability over the randomness of the players, it holds that all but an $O(\delta)$ -fraction of the players play one of those strategies. Therefore with high probability

$$\|\mathbf{b} - \mathbb{E}(f(\mathbf{a}))\|_2 = O(\sqrt{\delta} + 1/k).$$

Now, by the Lipschitz condition on f , whenever (7.5) holds,

$$\|f(\mathbf{a}) - f(\mathbb{E}(\mathbf{b}))\|_2 = O(\sqrt{\delta} + 1/k).$$

Therefore $\mathbb{E}(\mathbf{a})$ and $\mathbb{E}(\mathbf{b})$ are both solutions to the $\Theta(\sqrt{\delta} + 1/k)$ -approximate Euclidean BROUWER instance. ■

7.2.1 Binary Actions

Corollary 7.1 There exist constants $\epsilon, \delta > 0$, such that finding an (ϵ, δ) -WeakNash is PPAD-hard for succinct multiplayer games where each player has two actions.

Proof We replace each player (A, i) (respectively, (B, j)) in the proof of Theorem 5.2 with $k + 1$ players, denoted: $(A, i, 0), (A, i, 1/k), \dots, (A, i, 1)$. Each new player has two actions: $\{+, -\}$. Given a $(k + 1)$ -tuple of actions for players $\{(A, i, \cdot)\}$, we define the *realized value* $r_i \in [0, 1]$ to be

$$r_i \triangleq \max\{x \text{ s.t. } (A, i, x) \text{ plays action } +\}.$$

Let the realized value q_j for players $(B, j, 0), \dots, (B, j, 1)$ be defined analogously.

We let player (A, i, x) 's utility be:

$$u_{i,x}(+, \mathbf{b}) \triangleq - \left| \left(x + \frac{1}{k} \right) - q_i \right|^2$$

$$u_{i,x}(-, \mathbf{b}) \triangleq - \left| \left(x - \frac{1}{k} \right) - q_i \right|^2.$$

Similarly, for player (B, j, y) , we have

$$u_{j,y}(\pm, \mathbf{a}) \triangleq - \left| \left(y \pm \frac{1}{k} \right) - f_i(\mathbf{r}) \right|^2,$$

where $\mathbf{r} \triangleq (r_1, \dots, r_n)$ is the vector of realized values on the first group.

For any (A, i, x) , we have

$$\mathbb{E}[u_{i,x}(\pm, \mathbf{b})] = - \left| \left(x \pm \frac{1}{k} \right) - \mathbb{E}[q_i] \right|^2 - \text{Var}[q_i].$$

Subtracting her two possible payoffs, we have

$$\begin{aligned} \mathbb{E}[u_{i,x}(+, \mathbf{b}) - u_{i,x}(-, \mathbf{b})] &= \left| \left(x - \mathbb{E}[q_i] \right) - \frac{1}{k} \right|^2 - \left| \left(x - \mathbb{E}[q_i] \right) + \frac{1}{k} \right|^2 \\ &= -\frac{4}{k} \left(x - \mathbb{E}[q_i] \right). \end{aligned}$$

In particular, if $\mathbb{E}[q_i] > x + k\sqrt{\epsilon}$, player (A, i, x) must assign probability at most $\sqrt{\epsilon}$ to action $\{-\}$ in order to play ϵ -optimally. For any i such that all players (A, i, \cdot) use ϵ -optimal mixed strategies, we have that

$$|\mathbb{E}[q_i] - \mathbb{E}[r_i]| = O(k\sqrt{\epsilon} + 1/k). \quad (7.6)$$

In any (ϵ, δ) -WeakNash, for all but a $2\delta k$ -fraction of i 's it holds that all players (A, i, \cdot) use ϵ -optimal mixed strategies; thus (7.6) holds for all but a $2\delta k$ -fraction of i 's. Therefore,

$$\|\mathbb{E}[\mathbf{q}] - \mathbb{E}[\mathbf{r}]\|_2 = O(\sqrt{\delta k} + k\sqrt{\epsilon} + 1/k).$$

By the Lipschitz condition on f , the latter also implies

$$\|f(\mathbb{E}[\mathbf{q}]) - f(\mathbb{E}[\mathbf{r}])\|_2 = O(\sqrt{\delta k} + k\sqrt{\epsilon} + 1/k).$$

Similarly, for every j such that all players (B, j, \cdot) use ϵ -optimal mixed strategies, we have that

$$|\mathbb{E}[q_j] - f_j(\mathbb{E}[\mathbf{r}])| = O(k\sqrt{\epsilon} + 1/k).$$

Thus in any (ϵ, δ) -WeakNash,

$$\|\mathbb{E}[\mathbf{q}] - f(\mathbb{E}[\mathbf{r}])\|_2 = O(\sqrt{\delta k} + k\sqrt{\epsilon} + 1/k).$$

Therefore $\mathbb{E}(\mathbf{q})$ and $\mathbb{E}(\mathbf{r})$ are both solutions to the $\Theta(\sqrt{\delta} + 1/k)$ -approximate Euclidean BROUWER instance. ■

Bayesian Nash Equilibrium

In this short chapter we supplement our hardness of approximation of Nash equilibrium in many-player games and many-actions games, by showing a different reason of complexity: players' uncertainty. Specifically, we prove that finding an approximate Bayesian Nash equilibrium in two-player games with incomplete information is PPAD-complete, even when the players have only a constant number of actions.

In a game with incomplete information, each player i has a type t_i known only to her, and the players' types $\mathbf{t} = (t_1, t_2)$ are drawn from a joint distribution that is known to everyone. The payoff for player i is a function $u_i(\mathbf{a}, t_i)$ of her own type and all the players' actions.

Definition 8.1 (e.g., [Singh et al. \[2004\]](#)). In a *Bayesian Nash equilibrium*, for every player i and every type t_i , the mixed strategy $x_i(t_i)$ must be a best response in expectation over the other players' types and actions:

$$\mathbb{E}_{\mathbf{t}|t_i} \left[\mathbb{E}_{\mathbf{a} \sim (x_i(t_i), x_{-i}(t_{-i}))} [u_i(\mathbf{a}; t_i)] \right] \geq \max_{x'_i(t_i) \in \Delta A_i} \mathbb{E}_{\mathbf{t}|t_i} \left[\mathbb{E}_{\mathbf{a} \sim (x'_i(t_i), x_{-i}(t_{-i}))} [u_i(\mathbf{a}; t_i)] \right].$$

Similarly, in an ϵ -*approximate Bayesian Nash equilibrium*, for every player i and every type t_i , the mixed strategy $x_i(t_i)$ must be an ϵ -best response in expectation over the other players' types and actions:

$$\mathbb{E}_{\mathbf{t}|t_i} \left[\mathbb{E}_{\mathbf{a} \sim (x_i(t_i), x_{-i}(t_{-i}))} [u_i(\mathbf{a}; t_i)] \right] \geq \max_{x'_i(t_i) \in \Delta A_i} \mathbb{E}_{\mathbf{t}|t_i} \left[\mathbb{E}_{\mathbf{a} \sim (x'_i(t_i), x_{-i}(t_{-i}))} [u_i(\mathbf{a}; t_i)] \right] - \epsilon.$$

Before we prove our main corollary for incomplete information games, it is helpful to prove the following slightly weaker statement, for two players with many strategies.

Lemma 8.1 There exists a constant $\epsilon > 0$, such that given a two-player game with incomplete information where each player has n actions, finding an ϵ -approximate Bayesian Nash equilibrium is PPAD-complete.

Proof We reduce from a bipartite polymatrix game, and let the type set for each of the two players in the incomplete information game correspond to one side of the bipartite graph. The utility of player i on edge (i, j) of the polymatrix game depends on her identity (i) , as well as the identity (j) of the vertex on the other side of that edge. We use the types of the incomplete information game to encode i . We encode j using the strategies of the second player in the incomplete information game.

In more detail, consider a bipartite polymatrix game for which it is PPAD-hard to compute a 4ϵ -approximate Nash equilibrium. Use an affine transformation to change all the payoffs from $[0, 1]$ to $[1/2, 1]$. It is PPAD-hard to find a 2ϵ -approximate Nash equilibrium in the transformed game.

We now construct the two-player incomplete information game: As we hinted before, we let the type set of each player correspond to the vertices on one side of the bipartite graphical game. Player i has $|T_i|$ types and $2|T_i|$ strategies, where each strategy corresponds to a pair of a vertex and a strategy for that vertex. If a player plays a strategy whose vertex does not match her type, her payoff is 0. Therefore in every ϵ -approximate Bayesian Nash equilibrium, every player, on every type, plays the two strategies that correspond to her type with probability at least $1 - 2\epsilon$.

Let the joint distribution over types be as follows: pick a random edge in the bipartite graph, and assign the types corresponding to its vertices. Whenever both players play strategies that match their respective types, their payoffs are the payoffs in the (transformed) bimatrix game on that edge. In every ϵ -approximate Bayesian Nash equilibrium, every player, on every type, plays a mixed strategy which is the ϵ -best response. Since the other player plays strategies that correspond to the correct vertex with probability at least $1 - 2\epsilon$, the same mixture must be a 2ϵ -best response for the vertex player in the bipartite game. ■

In order to prove the main corollary, we need to reduce the number of actions in the above construction. Observe that we don't need each player to choose an action that uniquely identifies her type. Rather, it suffices to specify which neighbor of the other player's vertex is chosen. This can be done concisely via a coloring of the vertices such that every pair of vertices of distance exactly two have different colors; i.e., a coloring of the square of the polymatrix game's graph. The squared graph has degree $3 \cdot (3 - 1) = 6$, and therefore we can efficiently find a 7-coloring. It

suffices for each player to specify one of 7 colors, together with one of 2 strategies for the vertex with that color. Therefore we can encode this choice using only 14 strategies.

Corollary 8.1 There exists a constant $\epsilon > 0$, such that given a two-player game with incomplete information where each player has 14 actions, finding an ϵ -approximate Bayesian Nash equilibrium is PPAD-complete.

Market Equilibrium

In this chapter we formally introduce our result for non-monotone markets, discuss it, and compare it to the original result of [Chen et al. \[2013\]](#). A sketch of the proof appears in the next section.

Intuitively, a market is monotone if increasing the price of some good, while fixing the rest of the prices, never increases the excess demand for that good. Formally, we have the following definition by Chen et al.:

Definition 9.1 [Chen et al. \[2013\]](#). Let M be a market with $k \geq 2$ goods. We say that M is *non-monotone* at price vector $\boldsymbol{\pi}$ if there exist $c > 0$, and a good g_1 such that:

- the excess demand $Z_1(y_1, \dots, y_k)$ is a continuous function (rather than correspondence) over $\mathbf{y} \in B(\boldsymbol{\pi}, c)$;
- $Z_1(\boldsymbol{\pi}) > 0$;
- the partial derivative of $\partial Z_1 / \partial y_1$ exists and is continuous over $B(\boldsymbol{\pi}, c)$; and
- $\partial Z_1 / \partial y_1(\boldsymbol{\pi}) > 0$.

We say that a market M is non-monotone if there exists such a rational price vector $\boldsymbol{\pi} \geq \mathbf{0}$, and $Z_1(\boldsymbol{\pi})$ is *moderately computable*; i.e., for any $\gamma > 0$, $Z_1(\boldsymbol{\pi})$ can be approximated to within γ in time polynomial in $1/\gamma$.

In general, we want to talk about non-monotone families of utility functions, i.e., ones that support non-monotone markets. Formally,

Definition 9.2 [Chen et al. \[2013\]](#). We say that a family \mathcal{U} of utility functions is *non-monotone* if:

- \mathcal{U} is countable;
- if $u: [0, \infty)^k \rightarrow \mathbb{R}$ is in \mathcal{U} , then so is $u'(x_1, \dots, x_m) = a \cdot u(x_{l_1}/b_1, \dots, x_{l_k}/b_k)$ for any indices $l_1, \dots, l_k \in [m]$ and positive (rational) a, b_1, \dots, b_k ;
- $u(\mathbf{x}) = g^*(x_i)$ is in \mathcal{U} for some strictly increasing $g: [0, \infty) \rightarrow \mathbb{R}$; and
- there exists a non-monotone market $M_{\mathcal{U}}$ with utilities from \mathcal{U} .

We need to include one more definition: that of ϵ -tight market equilibrium.

Definition 9.3 A price vector π is an ϵ -tight approximate market equilibrium of M if there exists a $\mathbf{z} \in Z(\pi)$ (the excess demand at π) such that for every good j , $|z_j| \leq \epsilon W_j$, where W_j is the sum of the endowments of good j .

Our main result for non-monotone market equilibria is now formally defined:

Theorem 9.1 Non-monotone market. Let \mathcal{U} be any non-monotone family of utility functions. There exists a constant $\epsilon_{\mathcal{U}} > 0$ such that given a market M where the utility of each trader is either linear or taken from \mathcal{U} , finding an $\epsilon_{\mathcal{U}}$ -tight approximate market equilibrium is PPAD-hard.

9.1 Why Are Non-Monotone Markets Hard?

Before delving into the details of the construction, we attempt to reach some intuition: why should we expect equilibrium computation to be hard in non-monotone markets? Probably the most intuitive algorithm for finding market equilibrium is via tâtonnement: raise the prices of over-demanded goods, and decrease the prices of under-demanded goods. For many markets, the tâtonnement process is also computationally efficient [Codenotti et al. 2005]. One obvious problem is that when the market is non-monotone, the tâtonnement step actually takes us further away from equilibrium. However, the non-monotonicity is only local: if we set the (relative) price of the non-monotone good high enough, even the most enthusiastic traders can only afford a small amount.

The “real” reason that tâtonnement fails to converge efficiently for non-monotone markets is a little more subtle. What happens when the demand for the non-monotone good g increases by a factor of $(1 + \delta)$ for some small δ ? The tâtonnement increases the price of g , which further increases the demand. Eventually, the price is high enough, and the demand is reduced; but due to the non-monotonicity we may have to increase the price by a larger factor, i.e., $(1 + \delta')$ for $\delta' > \delta$. Now, another trader with a positive endowment of g has increased her spending budget by $(1 + \delta')$, further increasing the demand for yet another good (by a larger factor). Thus a small change in the demand for one good may cause a much larger change in the demand for another good. Exploiting this “butterfly effect” lies at the heart of Chen et al.’s construction.

9.2 High-Level Structure of the Proof

Our reduction from polymatrix games to non-monotone markets closely follows the footsteps of Chen et al. [2013]. To gain some intuition, consider two goods g_{2i-1} and g_{2i} for each player i , corresponding to her two available strategies (soon

each of those goods will become a subset of goods). Let $\pi(g_{2i-1})$ and $\pi(g_{2i})$ denote their corresponding prices; those prices correspond to the probabilities that player i assigns to her respective strategies. For every $i, j \in [n]$, we add a trader who is interested in selling g_{2i-1} and buying g_{2j-1} (and similarly for $(2i, 2j-1)$, $(2i-1, 2j)$, and $(2i, 2j)$). This trader has an endowment of g_{2i-1} that is proportional to $P_{2i-1, 2j-1}$, the utility of player j in the bimatrix game with player i , when they both play the first strategy. Qualitatively, if the price $\pi(g_{2i-1})$ is high (player i assigns a high probability to her first strategy), and $P_{2i-1, 2j-1}$ is high (player j receives a high utility from playing her first strategy to i 's first strategy), then the demand for good g_{2j-1} is high—implying a high price in every approximate market equilibrium (i.e., player j indeed assigns a high probability to her first strategy).

In order to turn this qualitative intuition into a proof, we use a more complex construction. The main difficulty comes from the need to amplify the effect of a higher income for one trader on the incomes of other traders. To this end we consider, for each $i \in [n]$, two sequences of goods: $g_{2i-1} = g_{2i-1,0}, g_{2i-1,1}, \dots, g_{2i-1,4t} = h_{2i-1}$ and $g_{2i} = g_{2i,0}, g_{2i,1}, \dots, g_{2i,4t} = h_{2i}$. The trader mentioned in the previous paragraph actually owns $P_{2i-1, 2j-1}$ units of good h_{2i-1} ; she is still interested in good g_{2j-1} . Now we construct (in Lemma 9.7) a chain of gadgets that use copies of the non-monotone markets in \mathcal{U} to amplify the small gap guaranteed between $\pi(g_{2j-1})$ and $\pi(g_{2j})$ to a larger gap between $\pi(h_{2j-1})$ and $\pi(h_{2j})$.

Additionally, we want to bound the range that these prices may take. In Lemma 9.4 we use a price-regulating gadget [Chen et al. 2009a, Vazirani and Yannakakis 2011] to control the relative prices of $\pi(g_{2i-1, j})$ and $\pi(g_{2i, j})$. In Lemma 9.6 we show that the sums $\pi_{i, j} = \pi(g_{2i-1, j}) + \pi(g_{2i, j})$ are approximately equal. Finally, in Subsection 9.4.5 we combine these lemmata to formalize a quantitative version of the qualitative intuition described above.

9.3 Adaptations for Constant Factor Inapproximability

As mentioned in the introduction, Theorem 9.1 has a weakness in comparison to the results of Chen et al. [2013]: it only applies to *tight* approximate market equilibrium.

Maintaining the constant hardness of approximation through most of Chen, Paparas, and Yannakakis's proof is rather straightforward, but there are a few hurdles along the way. To understand the first obstacle, we must consider a subtle issue of normalization. Chen et al. normalize the bimatrix game between every pair of players to have an average value of $1/2$. While this does not change the absolute utility gained from any deviation, the relative utility from deviation is now divided

by a factor of $\Theta(n)$. In contrast, in Theorem 5.1 we prove hardness for a constant ϵ when normalizing with respect to a constant degree (3), i.e., each player participates in only a constant number of bimatrix games. We overcome this difficulty by using a different normalization: only edges (i.e., bimatrix games) that belong to the game graph will have an average utility of $1/2$, while the utility on other edges remains 0. Since we proved hardness for a degree 3 graphical game, the normalization only costs us a constant factor.

More serious complications arise when trying to prove Chen, Paparas, and Yannakakis's Lemma 3¹ for a constant ϵ . This lemma says that certain prices (in fact, these are sums of prices), denoted $\pi_{i,0}$ for $i \in [n]$, are approximately equal. A key step in the proof of Chen et al. [2013] is to show, roughly, that in every $\epsilon(n)$ -approximate market equilibrium,

$$\pi_{i,0} \geq \frac{1}{n} \sum_{j \in [n]} \pi_{j,0} - O(\epsilon(n)).$$

When $\epsilon(n)$ is polynomially small, this immediately implies that $\min_{i \in [n]} \pi_{i,0}$ is within $O(\epsilon(n))$ of the average, and therefore it must also be that $\max_{i \in [n]} \pi_{i,0}$ is within $O(n \cdot \epsilon(n))$ of the average. When taking a larger $\epsilon(n)$, this reasoning breaks. The first modification we make to overcome this obstacle is to require $\epsilon(n)$ -tight approximate market equilibrium. This gives a two-sided bound:

$$\left| \pi_{i,0} - \frac{1}{n} \sum_{j \in [n]} \pi_{j,0} \right| = O(\epsilon(n)). \quad (9.1)$$

A second issue that arises in the same inequality is that with our new normalization, which depends on the game graph G , we can only prove that $\pi_{i,0}$ approximates the values of its neighbors, denoted $\mathcal{N}_G(i)$. In other words, (9.1) becomes

$$\left| \pi_{i,0} - \frac{1}{|\mathcal{N}_G(i)|} \sum_{j \in \mathcal{N}_G(i)} \pi_{j,0} \right| = O(\epsilon(n)). \quad (9.2)$$

In order to relate $\pi_{i,0}$ to other $\pi_{j,0}$'s (i.e., the corresponding values of the neighboring vertices), we consider T consecutive applications of (9.2): $\pi_{i,0}$ is $O(T \cdot \epsilon)$ -close to the expectation over $\pi_{j,0}$ where j is taken from the distribution of a T -steps random walk on G starting from i . For example, if G is a constant degree expander, the random walk converges in $O(\log n)$ steps, yielding a $(1/\log n)$ -inapproximability result.

1. Compare with our Lemma 9.6. The reader may also want to refer to Lemma 6 in the full version of Chen et al. [2013].

Achieving Constant Hardness

Finally, in order to achieve hardness for a constant ϵ , we want a graph with constant mixing time—and this clearly cannot be done with a constant degree.² Instead, in Section 9.4.2 we construct a normalized game whose graph has a constant mixing time, each vertex has degree $O(\sqrt{n})$, and yet approximating Nash equilibrium is hard for a constant ϵ . In short, we take n copies of the original n -player game (our new game has n^2 players). For any pair of players that play a (non-trivial) bimatrix game in the original game, we have a copy of the same bimatrix game between all $\binom{n}{2}$ pairs of their respective copies. We also add a trivial bimatrix game between every pair of players that belong to the same copy of the original game. In Section 9.4.2 we argue that these newly added trivial edges are only a constant fraction of all edges in the new game graph, yet this graph has a constant mixing time.

9.4 Non-Monotone Markets: Proof of Inapproximability

In this section we prove our main inapproximability result for non-monotone markets (Theorem 9.1).

9.4.1 Normalized Polymatrix Games

We identify n -player, 2-strategy polymatrix graphical games with $2n \times 2n$ matrices by letting the (i, j) -th 2×2 block correspond to the payoff matrix of player j in the bimatrix game with player i .

Given a game \mathcal{G} , let \mathbf{P}' be the $2n \times 2n$ induced payoff matrix. We normalize \mathbf{P}' as follows:

$$P_{2i,2j-1} = \begin{cases} 1/(2\Delta) + (P'_{2i,2j-1} - P'_{2i,2j})/(2\Delta) & (i, j) \in E, \\ 0 & \text{otherwise,} \end{cases} \quad (9.3)$$

where E is the edge set for the graphical game³ and Δ is the maximum degree. We define $P_{2i,2j}$, $P_{2i-1,2j-1}$, $P_{2i-1,2j}$ analogously. Notice that \mathbf{P} and \mathbf{P}' have the same ϵ -WSNE, up to the normalization by the degree Δ . In particular, finding an (ϵ/Δ) -WSNE in \mathbf{P} continues to be PPAD-complete.

Observe that in this formulation, finding an ϵ -WSNE is equivalent to finding a vector $\mathbf{x} \in [0, 1]^{2n}$ s.t. $x_{2i-1} + x_{2i} = 1$ and

$$\mathbf{x}^\top \cdot \mathbf{P}_{2i-1} > \mathbf{x}^\top \cdot \mathbf{P}_{2i} + \epsilon \implies x_{2i} = 0,$$

$$\mathbf{x}^\top \cdot \mathbf{P}_{2i-1} < \mathbf{x}^\top \cdot \mathbf{P}_{2i} - \epsilon \implies x_{2i-1} = 0.$$

2. In fact, it seems that a graph where the random walks starting from any pair of neighbors converge in constant time would suffice. We do not know whether such graphs can be constructed with constant degree.

3. Notice that this definition allows self-loops in the game graph.

9.4.2 Games on Graphs with a Constant Mixing Time

Given the correspondence defined above between n -player games and $2n \times 2n$ matrices, we see that the structure of the game graph plays a non-trivial role in the construction. In particular, adding trivial edges between vertices, i.e., adding zero-utility bimatrix games between players, has no effect on the utility of the players, but changes the corresponding normalized matrix. For reasons that will become clear much later in the proof, we would like our game graph to have a constant mixing time.

Indeed, a trivial candidate with very fast mixing is the complete graph. However, such a blowup in the degree would dilute our inapproximability factor in the *normalized game*. Instead, we consider n copies (v^1, \dots, v^n) of each player $v \in V$ in the original game. If players u and v play a bimatrix game $\mathcal{G}_{u,v}$ in the original game \mathcal{G} , then for every $i, j \in [n]$, we construct the same bimatrix game $\mathcal{G}_{u,v}$ between u^i and v^j . Our game graph now consists of n^2 vertices, each with degree⁴ $3n$. Finally, within each copy V^i , we add trivial edges between all the vertices not otherwise connected (including self-loops). Normalize this game using (9.3). We use $\bar{\mathcal{G}}$ to denote the new game and \bar{G} for the new game graph; we henceforth let $\Delta = 4n - 3$ denote the degree. In the next two lemmata we show that this game satisfies the two properties we need: finding an ϵ -WSNE of $\bar{\mathcal{G}}$ is PPAD-complete, and the mixing time of \bar{G} is constant.

Lemma 9.1 Given an ϵ -WSNE in $\bar{\mathcal{G}}$, we can (efficiently) construct a $(4\epsilon/3)$ -WSNE for \mathcal{G} .

Proof For each player v , we take the average of the mixed strategies of v^1, \dots, v^n . The utility of v is the same as the average of utilities of v^1, \dots, v^n , and if v has a $(4\epsilon/3)$ -improving deviation, then at least one of the copies v^i has an ϵ -improving deviation. (The $(4/3)$ factor comes from the change in the degree.) ■

Lemma 9.2 Let $\pi_{v^i, T}$ be the distribution of a random walk on \bar{G} after T steps, starting from v^i , and let π^* be the uniform distribution on the vertices of \bar{G} . Then

$$\|\pi_{v^i, T} - \pi^*\|_1 \leq \left(\frac{1}{4}\right)^{T/2} + \left(\frac{3}{4}\right)^{T/2} = 2^{-O(T)}.$$

Proof At each step of the random walk, there is a constant probability (greater than $3/4$) of walking on a non-trivial edge, which takes us to another (uniformly random) copy of the original game; thereafter the copy of the game remains uniformly random. Similarly, at each step there is a constant probability (greater than $1/4$) of moving

4. Theorem 5.1 promises a graphical game of degree at most 3. It is not hard to extend to a 3-regular graph game with only a constant loss in the approximation factor.

to a vertex within the same copy (again, uniformly random). Thus conditioned on having walked on a non-trivial edge, and then on an edge within the same copy, the distribution is uniform. Since all vertices have the same degree, this is also the stationary distribution, and we never leave it. ■

For simplicity, in the following we redefine n to be the size of \overline{G} (and hence $\Delta \approx 4\sqrt{n}$).

9.4.3 Construction

Let N be a sufficiently large constant, and let $t = \log N$. Note that N depends on the parameters of the non-monotone market in \mathcal{U} , but not on the size n of our construction. We use the notation $O_N(\cdot)$ to denote the asymptotic behavior when N goes to infinity (but arbitrarily slower than n). We prove that it is PPAD-hard to find an η -tight approximate market clearing equilibrium for $\eta = N^{-8}\epsilon$, where ϵ is the inapproximability factor from Lemma 9.1.

For each vertex $i \in [n]$ we construct a series of $4t + 1$ gadgets $\mathcal{R}_{i,j}$, for $j \in [0 : 4t]$. Each gadget is composed of the following:

Main goods. $g_{2i-1,j}$ and $g_{2i,j}$ are the main goods in the reduction. They are used to encode the weights assigned to strategies x_{2i-1} and x_{2i} , respectively.

Non-monotone gadget. For each $j \in [4t]$, we include additional goods $s_{i,j,3}, \dots, s_{i,j,k}$ and a non-monotone gadget

$$\mathbf{NM}(\mu, \gamma, g_{2i-1,j}, g_{2i,j}, s_{i,j,3}, \dots, s_{i,j,k}).$$

This means that we scale the non-monotone market guaranteed to exist in \mathcal{U} according to parameters γ and μ such that when all the prices are approximately the same, the excess demand of $g_{2i-1,j}$ increases linearly with its price. Formally, we have the following lemma by Chen et al.

Lemma 9.3 Lemma 3.1 of [Chen et al. \[2013\]](#). There exist two (not necessarily rational) positive constants c and d with the following property. Given any $\gamma > 0$, one can build a market M_γ with utilities from \mathcal{U} and goods $g_{2i-1,j}, g_{2i,j}, s_{i,j,3}, \dots, s_{i,j,k}$ in polynomial time in $1/\gamma$ such that the following is true.

Let $f_{\gamma,\mu}(x)$ denote the excess demand function of $g_{2i-1,j}$ when the price of $g_{2i-1,j}$ is $1+x$, and the prices of all other $k-1$ goods are $1-x$. Then $f_{\gamma,\mu}$ is well defined over $[-c, c]$ with $|f_{\gamma,\mu}(0)| \leq \mu\gamma$ and its derivative $f'_{\gamma,\mu}(0) = d > 0$. For any $x \in [-c, c]$, $f_{\gamma,\mu}(x)$ also satisfies

$$|f_{\gamma,\mu}(x) - f_{\gamma,\mu}(0) - \mu dx| \leq |\mu x/D| \quad \text{where } D = \max\{20, 20/d\}.$$

Finally, we would like to set $\mu = \Delta/d$; in particular, this would imply that $f_{\gamma, \mu}(x) \approx \Delta x$. However, as mentioned above, d may be irrational. Instead, let d^* be a positive rational constant that satisfies

$$1 - 1/D \leq d^* \cdot d \leq 1.$$

We set the parameters $\mu = d^* \Delta$ and $\gamma = 1/N^6$.

Price-regulating gadget. For $j \in [4t]$, we include a price-regulating gadget

$$\mathbf{PR}(\tau, \alpha_j, g_{2i-1,j}, g_{2i,j}, s_{i,j,3}, \dots, s_{i,j,k}),$$

whereas for $j = 0$, we don't have goods $s_{i,0,3}, \dots, s_{i,0,k}$, and simply include the gadget

$$\mathbf{PR}(\tau, \alpha_0, g_{2i-1,0}, g_{2i,0}).$$

The parameters are set to $\tau = N \Delta$ and $\alpha_i = 2^i/N^5$. Notice that $\alpha_0 = N^{-5}$ and $\alpha_{4t} = N^{-1} = \beta$.

This gadget ensures that in any approximate equilibrium, the price ratio $\pi(g_{2i-1,j})/\pi(g_{2i,j})$ is always in the range $\left[\frac{1-\alpha_j}{1+\alpha_j}, \frac{1+\alpha_j}{1-\alpha_j}\right]$. Furthermore, within each gadget $\mathcal{R}_{i,j}$, the prices of all the goods besides $g_{2i-1,j}$ are exactly equal:

$$\pi(g_{2i,j}) = \pi(s_{i,j,3}) = \dots = \pi(s_{i,j,k}).$$

More specifically, we have two traders T_1 and T_2 as follows. T_1 has $(k-1)\tau$ endowment of $g_{2i-1,j}$; T_2 has τ endowment of each of the other goods in the gadget ($g_{2i,j}$ and $s_{i,j,3}, \dots, s_{i,j,k}$). Their utilities are defined as

$$u_1 = (1 + \alpha_j)x(g_{2i-1,j}) + (1 - \alpha_j) \left(x(g_{2i,j}) + \sum_{l=3}^k x(s_{i,j,l}) \right)$$

$$u_2 = (1 - \alpha_j)x(g_{2i-1,j}) + (1 + \alpha_j) \left(x(g_{2i,j}) + \sum_{l=3}^k x(s_{i,j,l}) \right).$$

In particular, T_1 and T_2 do not trade whenever

$$\pi(g_{2i-1,j})/\pi(g_{2i,j}) \in \left(\frac{1-\alpha_j}{1+\alpha_j}, \frac{1+\alpha_j}{1-\alpha_j} \right).$$

Auxiliary goods. For $j = 0$, we also include an auxiliary good AUX_i . Its eventual purpose is to disentangle the price of $g_{2i-1,0}$ and $g_{2i,0}$ from the utility that the actions of player i cause to other players.

Single-Minded Traders Graph

We connect the groups of goods $(\mathcal{R}_{i,j}$'s) using the following single-minded traders. We use $(w, g_1 : g_2)$ to denote a trader with endowment w of good g_1 who only wants good g_2 . Similarly, we use $(w, g_1, g_2 : g_3)$ to denote a trader who has an endowment w of each of g_1 and g_2 , and only wants g_3 .

1. For each $i \in [n]$ and $j \in [0 : 4t - 1]$, we add two traders from $\mathcal{R}_{i,j}$ to $\mathcal{R}_{i,j+1}$: $(\Delta, g_{2i-1,j} : g_{2i-1,j+1})$ and $(\Delta, g_{2i,j} : g_{2i,j+1})$. These traders help propagate price discrepancies from $g_{i,0}$ to $g_{i,4t}$.
2. Recall that we use g_i as short for $g_{i,0}$ and h_i for $g_{i,4t}$. For each pair $(i, j) \in E$ we add the following four traders: $(\Delta P_{2i-1,2j-1}, h_{2i-1} : g_{2j-1})$, $(\Delta P_{2i,2j-1}, h_{2i} : g_{2j-1})$, $(\Delta P_{2i-1,2j}, h_{2i-1} : g_{2j})$, $(\Delta P_{2i,2j}, h_{2i} : g_{2j})$. Since \mathbf{P} is normalized, we have

$$\Delta P_{2i-1,2j-1} + \Delta P_{2i-1,2j} = \Delta P_{2i,2j-1} + \Delta P_{2i,2j} = 1.$$

These traders will enforce the approximate Nash equilibrium.

3. Connect the auxiliary goods: We let

$$r_{2j-1} = 2\Delta - \Delta \sum_{i \in \mathcal{N}(j)} (P_{2i-1,2j-1} + P_{2i,2j-1}) > 0$$

$$r_{2j} = 2\Delta - \Delta \sum_{i \in \mathcal{N}(j)} (P_{2i-1,2j} + P_{2i,2j}) > 0;$$

note that $r_{2j-1} + r_{2j} = 2\Delta$.

We add the following traders: $((1 - \beta)r_{2j-1}, \text{AUX}_j : g_{2j-1})$, $((1 - \beta)r_{2j}, \text{AUX}_j : g_{2j})$, and $((1 - \beta)\Delta, g_{2j-1}, g_{2j} : \text{AUX}_j)$.

Notice that the economy graph is strongly connected (because G is strongly connected); therefore an equilibrium always exists [Maxfield 1997]. The supplies and demands for each good are summarized in Table 9.1.

9.4.4 Structure of a Market Equilibrium

We now prove some properties that every η -tight approximate equilibrium π must satisfy. Recall that $\eta = N^{-8}\epsilon$, where ϵ is the inapproximability factor for the polymatrix game.

Table 9.1 Goods and traders

Good		
[Total Supply]	Supplied By	Demanded By
g_{2i-1}, g_{2i}	PR trader;	PR traders;
$[N\Delta(1 + o_N(1))]$	$(\Delta, g_{2i-1} : g_{2i-1,1}),$ $(\Delta, g_{2i} : g_{2i,1});$ $((1 - \beta)\Delta, g_{2i-1}, g_{2i} : AUX_i)$	$(\Delta P_{2j-1,2i-1}, h_{2j-1} : g_{2i-1}),$ \vdots $(\Delta P_{2j,2i}, h_{2j} : g_{2i});$ $((1 - \beta)r_{2i-1}, AUX_i : g_{2i-1}),$ $((1 - \beta)r_{2i}, AUX_i : g_{2i})$
h_{2i-1}	PR trader;	PR traders;
$[(k - 1)N\Delta(1 + o_N(1))]$	NM traders; $(\Delta P_{2i-1,2j-1}, h_{2i-1} : g_{2j-1}),$ $(\Delta P_{2i-1,2j}, h_{2i-1} : g_{2j})$	NM traders; $(\Delta, g_{2i-1,4i-1} : h_{2i-1})$
h_{2i}	PR trader;	PR traders;
$[N\Delta(1 + o_N(1))]$	NM traders; $(\Delta P_{2i,2j-1}, h_{2i} : g_{2j-1}),$ $(\Delta P_{2i,2j}, h_{2i} : g_{2j})$	NM traders; $(\Delta, g_{2i,4i-1} : h_{2i})$
$g_{2i-1,j}$	PR trader;	PR traders;
$[(k - 1)N\Delta(1 + o_N(1))]$	NM traders; $(\Delta, g_{2i-1,j} : g_{2i-1,j+1})$	NM traders; $(\Delta, g_{2i-1,j-1} : g_{2i-1,j})$
$g_{2i,j}$	PR trader;	PR traders;
$[N\Delta(1 + o_N(1))]$	NM traders; $(\Delta, g_{2i,j} : g_{2i,j+1})$	NM traders; $(\Delta, g_{2i,j-1} : g_{2i,j})$
$s_{i,j,l}$	PR trader;	PR traders;
$[N\Delta(1 + o_N(1))]$	NM traders;	NM traders;
AUX_i	$((1 - \beta)r_{2i-1}, AUX_i : g_{2i-1}),$ $((1 - \beta)r_{2i}, AUX_i : g_{2i})$	$((1 - \beta)\Delta, g_{2i-1}, g_{2i} : AUX_i)$
$[N\Delta(1 + o_N(1))]$		

We begin with the application of the price regulating markets:

Lemma 9.4 For every $i \in [n]$ and $j \in [0 : 4t]$,

$$\frac{1 - \alpha_j}{1 + \alpha_j} \leq \frac{\pi(g_{2i-1}, j)}{\pi(g_{2i}, j)} \leq \frac{1 + \alpha_j}{1 - \alpha_j}$$

and

$$\pi(g_{2i}, j) = \pi(s_{i,j,3}) = \dots = \pi(s_{i,j,k}).$$

Proof Follows from the construction of the price-regulating markets. For more details see the proof⁵ of Lemma 6 in the full version of [Chen et al. \[2013\]](#), or previous works that use similar gadgets [[Chen et al. 2009a](#), [Vazirani and Yannakakis 2011](#)]. ■

We henceforth use $\pi_{i,j}$ to denote the sum of the (i, j) -th main goods: $\pi_{i,j} = \pi(g_{2i-1}, j) + \pi(g_{2i}, j)$.

Lemma 9.5

$$(1 - O_N(\eta))\pi_{i,0}/2 \leq \pi(\text{AUX}_i) \leq (1 + O_N(\eta))\pi_{i,0}/2.$$

Proof The total supply of AUX_i is $2(1 - \beta)\Delta$, yet the demand from the single-minded trader $((1 - \beta)\Delta, g_{2i-1}, g_{2i} : \text{AUX}_i)$ is $(1 - \beta)\Delta \frac{\pi_{i,0}}{\pi(\text{AUX}_i)}$. (For the upper bound we use the fact that π is a *tight* approximate market equilibrium.) ■

We are now ready to prove that the cost of every pair of main goods is approximately the same. Let $\delta = N^2\eta$.

Lemma 9.6 Let $\pi_{\max} = \max_{i,j} \pi_{i,j}$ and $\pi_{\min} = \min_{i,j} \pi_{i,j}$; then

$$\pi_{\max}/\pi_{\min} \leq 1 + O_N(\delta).$$

Proof The proof of this lemma is the main obstacle that requires the tightness assumption of the market equilibrium, as well as our bound on the mixing time from Lemma 9.2.

Recall that by Lemma 9.4, the prices of all the goods in each gadget $\mathcal{R}_{i,j}$ are approximately equal. Thus, using our bound on the clearing error, we have that for each $(i, j) \in [n] \times [0 : 4t]$,

$$\begin{aligned} & \left| \text{total spent on } \mathcal{R}_{i,j} - \text{total worth of } \mathcal{R}_{i,j} \right| \\ & \leq O_N(\eta \cdot kN\Delta)\pi_{i,j} = O_N(\eta \cdot N\Delta)\pi_{i,j}. \end{aligned} \tag{9.4}$$

5. In their statement, Chen et al. require an ϵ -additively approximate equilibrium, and for a much smaller ϵ . However, their proof continues to hold with our parameters.

By Walras's Law, the traders within each $\mathcal{R}_{i,j}$ (i.e., the price-regulating and non-monotone gadgets) contribute the same to both quantities in (9.4). Similarly, by Lemma 9.5, the auxiliary traders contribute $O_N(\eta\Delta\pi_{i,0})$ -approximately the same (for $j = 0$). Therefore the money spent on $\mathcal{R}_{i,j}$ by the single-minded traders is approximately the same as the total worth of endowments in $\mathcal{R}_{i,j}$ of single-minded traders:

Comment 9.1

$$\left| \sum_{u \in \mathcal{N}^-(v)} w_{u,v} \pi_u - \sum_{u \in \mathcal{N}^+(v)} w_{v,u} \pi_v \right| \leq O((\eta'/\log n) \cdot k\pi_v)$$

- For each $(i, j) \in [n] \times [4t]$, the restriction of (9.4) to the single-minded traders gives

$$|\Delta\pi_{i,j-1} - \Delta\pi_{i,j}| = O_N(\eta \cdot N\Delta)\pi_{i,j}. \quad (9.5)$$

- Similarly, for each group $\mathcal{R}_{i,0}$, we have

$$\left| \sum_{l \in \mathcal{N}(i)} \pi_{l,4t} - \Delta\pi_{i,0} \right| = O_N(\eta \cdot N\Delta)\pi_{i,0}. \quad (9.6)$$

Applying (9.5) inductively, we have that for any $i \in [n]$ and for any $j, l \in [0: 4t]$,

$$|\pi_{i,j} - \pi_{i,l}| = O_N(\eta \cdot Nt)\pi_{i,l}.$$

Combining with (9.6) we have

$$\left| \pi_{i,0} - \frac{1}{\Delta} \sum_{l \in \mathcal{N}_{GG}(i)} \pi_{l,0} \right| = O(\eta \cdot Nt)\pi_{i,0}.$$

Thus for each i , $\pi_{i,0}$ is $O(\eta \cdot Nt)$ -approximately equal to the average of its neighbors in \bar{G} . Repeating this argument T times, we have that $\pi_{i,0}$ is $O_N(T\eta \cdot Nt)$ -approximately equal to the expectation over a T -step random walk in \bar{G} starting from i . By Lemma 9.2, after $T = O(\log \delta)$ steps the random walk δ -approximately converges to the uniform distribution, and we have

$$\begin{aligned} \left| \pi_{i,0} - \frac{1}{n} \sum_{l \in [n]} \pi_{l,0} \right| &= O_N(T\eta \cdot Nt) \frac{1}{n} \sum_{l \in [n]} \pi_{l,0} + \delta \max_{l \in [n]} \pi_{l,0} \\ &= O_N(\delta) \max_{l \in [n]} \pi_{l,0}. \quad \blacksquare \end{aligned}$$

Finally, we have the following lemma, which describes the action of the non-monotone gadgets.

Lemma 9.7 Lemma 6 of [Chen et al. \[2013\]](#).

$$\frac{1 + \alpha_{j-1}}{\pi(g_{2i-1, j-1})} = \frac{1 - \alpha_{j-1}}{\pi(g_{2i, j-1})} \implies \frac{1 + \alpha_j}{\pi(g_{2i-1, j})} = \frac{1 - \alpha_j}{\pi(g_{2i, j})} \quad \text{and}$$

$$\frac{1 - \alpha_{j-1}}{\pi(g_{2i-1, j-1})} = \frac{1 + \alpha_{j-1}}{\pi(g_{2i, j-1})} \implies \frac{1 - \alpha_j}{\pi(g_{2i-1, j})} = \frac{1 + \alpha_j}{\pi(g_{2i, j})}$$

Proof The demand for $g_{2i-1, j}$ and $g_{2i, j}$ comes from three sources: the single-minded traders, $(\Delta, g_{2i-1, j-1} : g_{2i-1, j})$ and $(\Delta, g_{2i, j-1} : g_{2i, j})$; the non-monotone gadget; and the price regulating gadget. Assume without loss of generality that the first premise holds, i.e., $\frac{1 + \alpha_{j-1}}{\pi(g_{2i-1, j-1})} = \frac{1 - \alpha_{j-1}}{\pi(g_{2i, j-1})}$. When the prices of $g_{2i-1, j}$ and $g_{2i, j}$ are equal, the demand from $(\Delta, g_{2i-1, j-1} : g_{2i-1, j})$ is larger since she has more income from $g_{2i-1, j-1}$. In order to account for this difference, $\pi(g_{2i-1, j})$ must be higher— but then the demand from the traders in the non-monotone market increases. Thus we have to further increase $\pi(g_{2i-1, j})$, until we reach the threshold of the price regulating traders: $(1 + \alpha_j)/(1 - \alpha_j)$.

Formally, normalize π such that $\pi_{i, j} = \pi(g_{2i-1, j}) + \pi(g_{2i, j}) = 2$. Thus by Lemma 9.6, $\pi_{i, j-1}$ is also $O_N(\delta)$ -close to 2. Let $f(x)$ denote the excess demand from the traders in the non-monotone gadget when $\pi(g_{2i-1, j}) = 1 + x$ and $\pi(g_{2i-1, j}) = \pi(s_{i, j, 3}) = \dots = \pi(s_{i, j, k}) = 1 - x$ (recall from Lemma 9.4 that the latter prices are always equal to each other). By Lemma 9.3, we have that $|f(0)| \leq \mu\gamma$, and for all $x \in [-c, c]$:

$$|f(x) - f(0) - \mu dx| \leq |\mu x/D|.$$

Now, let $\pi(g_{2i-1, j-1}) = 1 + y$; notice that by Lemma 9.6, $y = \alpha_{j-1} \pm O_N(\delta)$. Let $h(x, y)$ denote the excess demand from all traders besides the two that belong to the price regulating gadget. Then,

$$h(x, y) = f(x) + \frac{\Delta(1 + y)}{1 + x} - \Delta = f(x) - \frac{\Delta x}{1 + x} + \frac{\Delta y}{1 + x}.$$

For small x , we show that $f(x) \approx \Delta x/(1 + x)$. More precisely,

$$\begin{aligned} |f(x) - \Delta x/(1 + x)| &\leq |f(x) - \mu dx| + |\mu dx - \Delta x| + \Delta|x - x/(1 + x)| \\ &\leq |f(0)| + 2|\mu x/D| + 2\Delta x^2 \\ &\leq \Delta \cdot (d^* \cdot \gamma + 2x/20 + 2x^2) \\ &\leq \frac{\Delta y}{3}, \end{aligned}$$

where the first inequality follows from the triangle inequality; the second follows by application of Lemma 9.3 for the first difference and the definitions of μ and d^* for the second; the third inequality applies Lemma 9.3 again; and finally the last inequality holds because for sufficiently large constant N , the parameters γ and x are sufficiently small.

Therefore, the excess demand must be balanced by the demand from the price regulating traders, implying that indeed $\frac{1+\alpha_j}{\pi(g_{2i-1,j})} = \frac{1-\alpha_j}{\pi(g_{2i,j})}$. ■

9.4.5 From Market Equilibrium to Nash Equilibrium

To complete the proof of Theorem 9.1, we must construct an ϵ -WSNE from any η -tight approximate market equilibrium.

Proof For each $i \in [n]$, let $\theta_i = (\pi(h_{2i-1}) + \pi(h_{2i}))/2$. We define

$$x_{2i-1} = \frac{\pi(h_{2i-1})/\theta_i - (1-\beta)}{2\beta} \quad \text{and} \quad x_{2i} = \frac{\pi(h_{2i})/\theta_i - (1-\beta)}{2\beta}. \quad (9.7)$$

Observe that $x_{2i-1} + x_{2i} = 1$.

Suppose that

$$\mathbf{x}^\top \cdot \mathbf{P}_1 \geq \mathbf{x}^\top \cdot \mathbf{P}_2 + \epsilon.$$

We show that this forces $x_1 = 1$ and $x_2 = 0$; by the discussion in Section 9.4.1 this implies that \mathbf{x} is indeed an ϵ -WSNE.

The following traders spend money on g_1 :

1. $(\Delta P_{i,1}, h_i; g_1)$ traders: for each $i \in \mathcal{N}_{\overline{G}}(1)$, there is one. The total money these traders spend on g_1 is

$$\sum \Delta P_{i,1} \cdot \pi(h_1) = \sum \Delta P_{i,1} (1 - \beta + 2\beta \cdot x_i) \theta_{\lceil i/2 \rceil}$$

2. $(\Delta P_{i,2}, h_i; g_2)$ traders: for each $i \in \mathcal{N}_{\overline{G}}(2)$, there is a one. The total money these traders spend on g_2 is

$$\sum \Delta P_{i,2} \cdot \pi(h_2) = \sum \Delta P_{i,2} (1 - \beta + 2\beta \cdot x_i) \theta_{\lceil i/2 \rceil}$$

3. $((1-\beta)r_1, \text{AUX}_1; g_1)$ and $((1-\beta)r_2, \text{AUX}_1; g_2)$ traders

Let M_1 be the total amount that these traders spend on g_1 . Then

$$M_1 = \sum_{i \in \mathcal{N}_{\overline{G}}(1)} \Delta P_{i,1} (1 - \beta + 2\beta \cdot x_i) \theta_{\lceil i/2 \rceil} + (1-\beta)r_1 \pi(\text{AUX}_1).$$

Normalizing the prices such that $\frac{1}{n} \sum_{i=1}^n \theta_i = 1/\Delta$, this means that

$$M_1 \geq 2(1 - \beta) + 2\beta \mathbf{x}^\top \cdot \mathbf{P}_1 - O_N(\delta).$$

Similarly,

$$M_2 \leq 2(1 - \beta) + 2\beta \mathbf{x}^\top \cdot \mathbf{P}_2 + O_N(\delta).$$

Therefore,

$$M_1 \geq M_2 + 2\beta\epsilon - O_N(\delta) = M_2 + \Theta_N(\beta\epsilon),$$

so the difference between the demands for g_1 and g_2 from these traders is

$$\frac{M_1}{\pi(g_1)} - \frac{M_2}{\pi(g_2)} \geq \frac{M_2 + \Theta_N(\beta\epsilon)}{\pi(g_1)} - \frac{M_2(1 + \alpha_0)}{\pi(g_1)(1 - \alpha_0)} = \Theta_N(\beta\epsilon).$$

Thus the price-regulating traders T_1 and T_2 must have different demands for g_1 and g_2 —but this can only happen when

$$\frac{1 + \alpha_0}{\pi(g_1)} = \frac{1 - \alpha_0}{\pi(g_2)}.$$

Therefore, by consecutive applications of Lemma 9.7,

$$\frac{1 + \beta}{\pi(h_1)} = \frac{1 - \beta}{\pi(h_2)}.$$

Finally, by (9.7) this implies that $x_1 = 1$ and $x_2 = 0$. ■

10 CourseMatch

In this chapter we characterize the computational complexity of computing an approximate competitive equilibrium from equal incomes (A-CEEI). CEEI is a ubiquitous mechanism for fair allocation of divisible goods. In settings with indivisible goods, a CEEI often does not exist, but a theorem of [Budish \[2011\]](#) guarantees an approximate CEEI (see definition below) with favorable economic properties. This mechanism is used in practice to allocate seats in university courses to students [[Budish et al. 2014](#)]. Budish’s existence proof does not come with an efficient algorithm for finding an A-CEEI. In applications a heuristic search algorithm without provable convergence guarantees is used. Understanding the computational complexity of A-CEEI is thus an important problem with both theoretical and practical implications.

Theorem 10.1 Theorem 10.3, informal statement. The problem of finding an A-CEEI as guaranteed by [Budish \[2011\]](#) is PPAD-complete.

10.1 The Course Allocation Problem

Even though the A-CEEI and the existence theorem in [Budish \[2011\]](#) are applicable to a broad range of allocation problems, we shall describe our results in the language of the course allocation problem.

We are given a set of M courses with integer capacities (the supply) $(q_j)_{j=1}^M$, and a set of N students, where each student i has a set $\Psi_i \subseteq 2^M$ of permissible course bundles, with each bundle containing at most $k \leq M$ courses. The set Ψ_i encodes both scheduling constraints (e.g., courses that meet at the same time) and any constraints specific to student i (e.g., prerequisites).

Each student i has a strict ordering over her permissible schedules, denoted by \succsim_i . We allow arbitrarily complex preferences—in particular, students may regard courses as substitutes or complements. More formally:

Definition 10.1 Course allocation problem. The input to a course allocation problem consists of:

- for each student i a set of course bundles $(\Psi_i)_{i=1}^N$;

- the students' reported preferences, $(\preceq_i)_{i=1}^N$; and
- the course capacities, $(q_j)_{j=1}^M$.

The output to a course allocation problem consists of:

- prices for each course $(p_j^*)_{j=1}^M$,
- allocations for each student $(x_i^*)_{i=1}^N$, and
- budgets for each student $(b_i^*)_{i=1}^N$.

How is an allocation evaluated? The *clearing error* of a solution to the allocation problem is the \mathcal{L}_2 norm of the length- M vector of seats oversubscribed in any course, or undersubscribed seats in courses with positive price.

Definition 10.2 The *clearing error* α of an allocation is

$$\alpha \equiv \sqrt{\sum_j z_j^2},$$

where z_j is given by

$$z_j = \begin{cases} \sum_i x_{ij}^* - q_j & \text{if } p_j^* > 0; \\ \max\left[\left(\sum_i x_{ij}^* - q_j\right), 0\right] & \text{if } p_j^* = 0. \end{cases}$$

We can now define the notion of *approximate CEEI*. The quality of approximation is characterized by two parameters: α , the clearing error (how far is our solution from a true competitive equilibrium?), and β , the bound on the difference in budgets (how far from equal are the budgets?). Informally, α can be thought of as the approximation loss on efficiency, and β can be thought of as the approximation loss on fairness.

Definition 10.3 An allocation is an (α, β) -*CEEI* if:

1. Each student is allocated their most preferred affordable bundle. Formally,

$$\forall i : x_i^* = \arg \max_{\preceq_i} \left[x_i \in \Psi_i : \sum_j x_{ij} p_j^* \leq b_i^* \right].$$

2. Total clearing error is at most α .
3. Every budget $b_i^* \in [1, 1 + \beta]$.

In [Budish \[2011\]](#) it is proved that an (α, β) -approximate CEEI always exists, for some quite favorable (and as we shall see, essentially optimal) values of α and β :

Theorem 10.2 Budish [2011]. For any input preferences, there exists an (α, β) -CEEI with $\alpha = \sqrt{kM/2}$ and any $\beta > 0$.

Recall that k is the maximum bundle size.

The bound of $\alpha = \sqrt{kM/2}$ means that, for a large number of students and course capacities, the market-clearing error converges to zero quite fast as a fraction of the endowment. It is also shown in Budish [2011] that the mechanism that allocates courses according to such an A-CEEI satisfies attractive criteria of approximate fairness, approximate truthfulness, and approximate Pareto efficiency. The reader may consult Budish [2011] for the precise definitions of the economic properties of the A-CEEI mechanism.

10.1.1 Our Results

We are now ready to formally state our result for the complexity of A-CEEI:

Theorem 10.3 Computing a $(\sqrt{\frac{kM}{2}}, \beta)$ -CEEI is PPAD-complete, for some small constant $\beta > 0$.

In Section 10.2, we prove that computing a $(\sqrt{\frac{kM}{2}}, \beta)$ -CEEI is PPAD-hard; this is accomplished by a reduction from ϵ -GCIRCUIT. Then, in Section 10.3, we prove that the same problem also belongs to the class PPAD; this proof mostly follows along the lines of Budish’s existence proof [Budish 2011], but certain probabilistic arguments must be constructively derandomized.

10.2 A-CEEI Is PPAD-Hard

Informally, in this section we provide a construction demonstrating that it is possible to define a set of courses, students, and preferences such that the price of the courses in an A-CEEI simulates the various “basic circuit functions” (e.g., an OR gate) that, when combined and wired together, are the necessary building blocks sufficient to emulate any continuous function. Therefore, any algorithm capable of solving A-CEEI in polynomial time would also suffice to solve ϵ -GCIRCUIT, and hence any problem in PPAD, in polynomial time as well.

Overview of the Reduction

We shall reduce ϵ -GCIRCUIT with fan-out 2 to the problem of finding an (α, β) -CEEI, with approximation parameters $\alpha = \Theta(N/M)$ and $\epsilon = \epsilon(\beta)$. (Note that, by increasing N , we can make α arbitrarily large as a function of M ; in particular, $\alpha > \sqrt{kM/2}$.)

We will construct gadgets (that is, small sets of courses, students, capacities and preferences) for the various types of gates in the generalized circuit problem. Each

gadget that we construct has one or more dedicated “input course,” a single “output course,” and possibly some “interior courses.” An output course of one gadget can (and will) be an input to another. The construction will guarantee that in any A-CEEI the price of the output course will be approximately equal to the gate applied to the prices of the input courses.

10.2.1 The NOT Gate (G_-) Gadget

To illustrate what needs to be done, we proceed to construct a gadget for the G_- gate; in particular, this implements a logical NOT.

Lemma 10.1 NOT gate (G_-) gadget.

Let $n_x > 4\alpha$ and suppose that the economy contains the following courses:

- c_x (the “input course”);
- c_{1-x} with capacity $q_{1-x} = n_x/2$ (the “output course”);

and the following set of students:

- n_x students interested only in the schedule $\{c_x, c_{1-x}\}$;

and suppose further that at most $n_{1-x} = n_x/4$ other students are interested in course c_{1-x} .

Then in any (α, β) -CEEI,

$$p_{1-x}^* \in [1 - p_x^*, 1 - p_x^* + \beta].$$

Proof Observe that:

- If $p_{1-x}^* > 1 - p_x^* + \beta$, then none of the n_x students will be able to afford the bundle $\{c_x, c_{1-x}\}$, and therefore there will be at most $n_{1-x} = n_x/4$ students enrolled in the c_{1-x} —much less than the capacity $n_x/2$. Therefore $z_{1-x} \geq n_x/4$.
- On the other hand, if $p_{1-x}^* < 1 - p_x^*$, then all n_x students can afford the bundle $\{c_x, c_{1-x}\}$ —therefore the course will be overbooked by $n_x/2$; thus, $z_{1-x} \geq n_x/2$.

Therefore if $p_{1-x}^* \notin [1 - p_x^*, 1 - p_x^* + \beta]$, then $\|z\|_2 \geq n_x/4 > \alpha$ —a contradiction to (α, β) -CEEI. ■

Similarly, in Section 10.2.3, we construct gadgets that simulate all the gates of the generalized circuit problem.

10.2.2 Course-Size Amplification

In the next subsection, we will construct gadgets that compute all the gates necessary for the circuit in the reduction from ϵ -GCIRCUIT. What happens when we try to concatenate them to form a circuit? Recall the penultimate sentence in the statement of Lemma 10.1: It says that the output course's price continues to behave like the output of the simulated gate, as long as there are not too many additional students that try to take the output course. (If there are more students, they may raise the price of the course beyond what we expect.) In particular, *the number of additional students that may want the output course is smaller than the number of students that want the input course.*

If we concatenated the gadgets without change, we would need to have larger course sizes as we increase the depth of the simulated circuit. This increase in course size is exponential in the depth of the circuit. Things get even worse—since we reduce from *generalized* circuits, our gates form cycles. If the course size must increase at every gate, it would have to be infinite!

To overcome this problem, we construct a special $G_{=}$ gadget that (approximately) preserves the price from the input course, but is robust to twice as many additional students:

Lemma 10.2 Course-size amplification gadget. Let $n_x \geq 100\alpha$ and suppose that the economy contains the following courses:

- c_x (the “input course”)
- for $i = 1, \dots, 10$, c_i with capacities $q_i = 0.5 \cdot n_x$ (“interior courses”);
- $c_{x'}$ with capacity $q_{x'}$, s.t. $q_x \leq q_{x'} \leq 4n_x$ (“output course”);

and the following sets of students:

- n_x students interested in schedules $(\{c_x, c_i\})_{i=1}^{10}$ (in this order);
- $n_i = 0.49 \cdot n_x$ students ($\forall i$) interested in schedules $(\{c_{x'}, c_i\}, \{c_i\}, \{c_{i+1}\}, \dots, \{c_{10}\})$ (in this order);

and suppose further that at most $n_{x'} = 2n_x$ other students are interested in course $c_{x'}$.

Then in any (α, β) -CEEI,

$$p_{x'}^* \in [p_x^* - \beta, p_x^* + \beta].$$

In particular, notice that the price of $c_{x'}$ is guaranteed to approximate the price of c_x , even in the presence of additional $n_{x'} = 2n_x$ students—twice as many students as we added to c_x .

Proof We start by proving that all the c_i 's simulate NOT gadgets simultaneously, i.e., for every i and every (α, β) -CEEI, $p_i^* \in [1 - p_x^*, 1 - p_x^* + \beta]$.

- If $p_i^* > 1 - p_x^* + \beta$, assume w.l.o.g. that it is the first such i , i.e., $p_j^* \leq 1 - p_x^* + \beta < p_i^*$ for every $j < i$.

None of the n_x students can afford buying both c_x and c_i . Furthermore, for every $j < i$, none of the n_j students will prefer c_i over c_j . Therefore at most n_i students will take this course: $z_i^* \geq 0.01n_x$.

- If, on the other hand, $p_i^* < 1 - p_x^*$, then all n_x students will buy course c_i or some previous course c_j (for $j \leq i$); additionally, for every $j \leq i$, each of the n_j corresponding students will buy some course c_k for $j \leq k \leq i$. Therefore the total overbooking of courses $1, \dots, i$ will be at least $\sum_{j \leq i} z_j^* \geq n_x \cdot (1 - 0.01i)$ —a contradiction to (α, β) -CEEI.

Now that we have established that $p_i^* \in [1 - p_x^*, 1 - p_x^* + \beta]$, we shall prove the main claim, i.e., that $p_{x'}^* \in [p_x^* - \beta, p_x^* + \beta]$.

- If $p_{x'}^* > p_x^* + \beta$, then none of the n_i students, for any n_i , can afford buying both $c_{x'}$ and c_i . Therefore, even in the presence of additional $n_{x'} = 2n_x$ students who want to take $c_{x'}$, the course will be undersubscribed by $z_{x'}^* \geq q_{x'} - n_{x'} = 2n_x$.
- If $p_{x'}^* < p_x^* - \beta$, then all n_i students, for each i , can afford to buy their top schedule—both $\{c_i, c_{x'}\}$. Therefore $c_{x'}$ will be oversubscribed by at least $z_{x'}^* \geq 0.9 \cdot n_x$ —a contradiction to (α, β) -CEEI. ■

Finally, given an instance of ϵ -GCIRCUIT with fan-out 2, we can use the gadgets we constructed in Lemmata 10.1–10.2 to construct an instance of (α, β) -CEEI that simulates the generalized circuit. We concatenate gadgets by identifying the output course of one gadget with the input course of the next two gadgets. In particular, after each gate gadget, we insert a series of course-size amplifying gadgets. Each amplifying gadget doubles the number of additional students that the gadget can tolerate, so a constant number of amplifying gadgets suffice; thus the blowup in error is also constant. As for the size of the reduction, each gadget introduces a constant number of new courses, and $\Theta(\alpha)$ new students; thus $M = \Theta(|V|)$ and $N = \Theta(\alpha \cdot |V|)$, where $|V|$ is the number of gates in the generalized circuit.

10.2.3 Additional Gate Gadgets

In this section we construct the rest of the gate gadgets, completing the proof of PPAD-hardness.

In the lemma below we construct gadgets for a slightly modified set of gates. In particular, instead of implementing gates G_ζ and $G_{\times\zeta}$ from Definition 6.2, we only consider the special cases corresponding to $\zeta = \frac{1}{2}$ (denoted $G_{\frac{1}{2}}$ and $G_{/2}$, respectively).

Lemma 10.3 Let $n_x \geq 2^8 \cdot \alpha$ and suppose that the economy has courses c_x and c_y . Then for any of the functions f listed below, we can add a course c_z , and at most n_x students interested in each of c_x and c_y , such that in any (α, β) -CEEI $p_z^* \in [f(p_x^*, p_y^*) - 2\beta, f(p_x^*, p_y^*) + 2\beta]$.

1. HALF: $f_{G_{/2}}(x) = x/2$
2. VALUE: $f_{G_{\frac{1}{2}}} \equiv \frac{1}{2}$
3. SUM: $f_{G_+}(x, y) = \min(x + y, 1)$
4. DIFF: $f_{G_-}(x, y) = \max(x - y, 0)$
5. LESS: $f_{G_<}(x, y) = \begin{cases} 1 & x > y + \beta \\ 0 & y > x + \beta \end{cases}$
6. AND: $f_{G_\wedge}(x, y) = \begin{cases} 1 & (x > \frac{1}{2} + \beta) \wedge (y > \frac{1}{2} + \beta) \\ 0 & (x < \frac{1}{2} - \beta) \vee (y < \frac{1}{2} - \beta) \end{cases}$
7. OR: $f_{G_\vee}(x, y) = \begin{cases} 1 & (x > \frac{1}{2} + \beta) \vee (y > \frac{1}{2} + \beta) \\ 0 & (x < \frac{1}{2} - \beta) \wedge (y < \frac{1}{2} - \beta) \end{cases}$

In particular, $p_z^* \in [f(p_x^*, p_y^*) - 2\beta, f(p_x^*, p_y^*) + 2\beta]$ in every (α, β) -CEEI even if up to $n_z \leq n_x/2^8$ additional students (beyond the ones specified in the proofs below) are interested in course c_z .

Proof HALF $G_{/2}$

Let c_z have capacity $q_z = n_x/8$, let $n_z = q_z/2$, and consider three auxiliary courses c_1, c_2 , and $c_{\bar{x}}$ of capacities $q_1 = q_2 = q_z$ and $q_{\bar{x}} = n_x/2$. Using Lemma 10.1, add n_x students that will guarantee $p_{\bar{x}} \in [1 - p_x^*, 1 - p_x^* + \beta]$. Additionally, consider $n_{\bar{x}} = n_x/4$ students with preference list $(\{c_z, c_1, c_{\bar{x}}\}, \{c_z, c_2, c_{\bar{x}}\}, \{c_1, c_2, c_{\bar{x}}\})$ (in this order), then:

- If the total price $p_i^* + p_j^*$ of any pair $i, j \in \{1, 2, z\}$ is less than $p_x^* - \beta$, then all $n_{\bar{x}}$ students will be able to afford some subset in their preference list, leaving a total overbooking of at least $z_z^* + z_1^* + z_2^* \geq 2n_{\bar{x}} - 3q_z = n_x/8$, which violates the (α, β) -CEEI conditions.
- If the total price of any of the pairs above (wlog, $p_1^* + p_2^*$) is greater than $p_x^* + \beta$, then none of the $n_{\bar{x}}$ students will be able to afford the subset $\{c_1, c_2, c_{\bar{x}}\}$.

Therefore the number of students taking c_z will be at least the sum of students taking c_1 or c_2 . Therefore, even after taking into account n_z additional students, we have that $z_z^* + z_1^* + z_2^* \geq q_z - n_z = n_x/16$.

VALUE $G_{\frac{1}{2}}$

Similarly to the HALF gadget, consider two auxiliary courses c_1 and c_2 , and let n_x students have preferences $(\{c_z, c_1\}, \{c_z, c_2\}, \{c_1, c_2\})$. Then, following the argument for the HALF gadget, it is easy to see that $p_z^* \in [\frac{1}{2}, \frac{1}{2} + \beta]$ in any (α, β) -CEEI, with $n_z = n_x/8$.

DIFF G_-

Let $c_{\bar{x}}$ be a course with price $p_{\bar{x}}^* \in [1 - p_x^*, 1 - p_x^* + \beta]$, $q_{\bar{x}} = n_x/2$, and consider $n_{\bar{x}} = n_x/4$ students willing to take $\{c_{\bar{x}}, c_y, c_z\}$. Then it is easy to see that

$$\begin{aligned} p_z^* &\in [1 - p_{\bar{x}}^* - p_y^*, 1 - p_{\bar{x}}^* - p_y^* + \beta] \\ &\subseteq [p_x^* - p_y^* - \beta, p_x^* - p_y^* + \beta] \end{aligned}$$

with $n_z = n_x/16$.

SUM G_+

Concatenating NOT and DIFF gadgets, we have:

$$\begin{aligned} p_{\bar{x}}^* &\in [1 - p_x^*, 1 - p_x^* + \beta] \\ p_{\bar{z}}^* &\in [p_{\bar{x}}^* - p_y^* - \beta, p_{\bar{x}}^* - p_y^* + \beta] \\ p_z^* &\in [1 - (p_{\bar{x}}^* - p_y^* + \beta), 1 - (p_{\bar{x}}^* - p_y^* - \beta) + \beta] \\ &\subseteq [p_x^* + p_y^* - 2\beta, p_x^* + p_y^* + 2\beta] \end{aligned}$$

for $n_z = n_x/2^8$.

LESS $G_<$

Let $c_{\bar{x}}$ be a course with price $p_{\bar{x}}^* \in [1 - p_x^*, 1 - p_x^* + \beta]$, $q_{\bar{x}} = n_x/2$; let $q_z = n_x/8$ and $n_z = n_x/16$. Consider $n_x/4$ students wishing to take $(\{c_{\bar{x}}, c_y\}, \{c_z\})$, in this order:

- If $p_y^* > p_x^* + \beta$, then $p_{\bar{x}}^* + p_y^* > 1 + \beta$, and therefore none of the $n_x/4$ students will be able to afford the first pair; they will all try to sign up for c_z , which will be overbooked unless $p_z^* > 1$.
- If $p_x^* > p_y^* + \beta$, then all $n_x/4$ students will sign up for the first pair, forcing $p_z^* = 0$ in any (α, β) -CEEI.

AND G_{\wedge}

Let $c_{\frac{1}{2}}$ be a course with price $p_{\frac{1}{2}}^* \in [\frac{1}{2}, \frac{1}{2} + \beta]$ and $n_{\frac{1}{2}} = n_x/8$, as guaranteed by gadget VALUE; let $q_z = n_x/32$ and $n_z = n_x/64$. Consider $n_x/16$ students wishing to take $(\{c_x, c_{\frac{1}{2}}\}, \{c_y, c_{\frac{1}{2}}\}, \{c_z\})$, in this order.

- If $(p_x^* > \frac{1}{2} + \beta) \wedge (p_y^* > \frac{1}{2} + \beta)$, then the $n_x/16$ students can afford neither pair. They will all try to sign up for c_z , forcing $p_z^* > 1$, in any (α, β) -CEEI.
- If $(x < \frac{1}{2} - \beta) \vee (y < \frac{1}{2} - \beta)$, then the $n_x/16$ students can afford at least one of the pairs and will register for those courses. Thus $p_z^* = 0$.

OR G_{\vee}

Similar to the AND gadget; students will want $(\{c_x, c_y, c_{\frac{1}{2}}\}, \{c_z\})$, in this order. ■

10.3 A-CEEI ∈ PPAD

In this section we establish that computing a $(\frac{\sqrt{\sigma M}}{2}, \beta)$ -CEEI is in PPAD, for $\sigma = \min\{2k, M\}$. We follow the steps of the existence proof in Budish [2011], and show that each one can be carried out either in polynomial time, or through a fixed point. One difficulty is that certain steps of Budish's proof are randomized and must be constructively derandomized in polynomial time.

Remark 10.1 We assume that the student preferences (\succsim_i) are given in the form of an ordered list of all the bundles in Ψ_i (i.e., all the bundles that student i prefers over the empty bundle). In particular, we assume that the total number of permissible bundles is polynomial.

Remark 10.2 In fact, we prove that the following, slightly more general problem, is in PPAD: Given any $\beta, \epsilon > 0$ and initial approximate-budgets vector $\mathbf{b} \in [1, 1 + \beta]^N$, find a $(\frac{\sqrt{\sigma M}}{2}, \beta)$ -CEEI with budgets \mathbf{b}^* such that $|b_i - b_i^*| < \epsilon$ for every i .

Our proof will follow the steps of the existence proof by Budish [2011]. We will use the power of PPAD to solve the Kakutani problem, and derandomize the other nonconstructive ingredients.

10.3.1 Preliminaries

Our algorithm receives as input an economy $((q_j)_{j=1}^M, (\Psi_i)_{i=1}^N, (\succsim_i)_{i=1}^N)$, parameters $\beta, \epsilon > 0$, and an initial approximate-budgets vector $\mathbf{b} \in [1, 1 + \beta]^N$. We denote $\bar{\beta} = \min\{\beta, \epsilon\}/2$.

We will consider M -dimensional price vectors in $\mathcal{P} = [0, 1 + \beta + \epsilon]^M$. In order to define a price adjustment function, we consider an enlargement $\tilde{\mathcal{P}} = [-1, 2 +$

$\beta + \epsilon]^M$, as well as a truncation function $t : \tilde{\mathcal{P}} \rightarrow \mathcal{P}$ (whose j -th coordinate is given by $t_j(\tilde{p}_j) = \min\{\max\{\tilde{p}_j, 0\}, 1 + \beta + \epsilon\}$).

For each student i , we denote her demand at prices $\tilde{\mathbf{p}}$ with budget b_i by

$$d_i(\tilde{\mathbf{p}}, b_i) = \max_{(\succeq_i)} \{x' \in \Psi_i : \tilde{\mathbf{p}} \cdot x' \leq b_i\}.$$

Given the total demand of all the students, we can define the excess demand to be:

$$\mathbf{z}(\tilde{\mathbf{p}}, \mathbf{b}) = \sum_{i=1}^N d_i(\tilde{\mathbf{p}}, b_i) - \mathbf{q}.$$

A key ingredient to the analysis is the budget-constraint hyperplanes. These are the hyperplanes in price space along which a student can exactly afford a specific bundle. For each student i and bundle x , the corresponding *budget-constraint hyperplane* is defined as $H(i, x) = \{\tilde{\mathbf{p}} \in \mathcal{P} : \tilde{\mathbf{p}} \cdot x = b_i\}$.

10.3.2 Deterministically Finding a “General Position” Perturbation (Step 1)

It is convenient to assume that the budget-constraint hyperplanes are in “general position,” i.e., there is no point $\tilde{\mathbf{p}} \in \mathcal{P}$ at which any subset of linearly dependent budget-constraint hyperplanes intersect (in particular, no more than M hyperplanes intersect at any point). In the existence proof, this is achieved by assigning a small random reverse tax $\tau_{i,x} \in (-\epsilon, \epsilon)$ for each student i and bundle x ; i 's modified cost for bundle x at prices $\tilde{\mathbf{p}}$ becomes $\tilde{\mathbf{p}} \cdot x - \tau_{i,x}$. Given taxes $\tau = (\tau_{i,x})_{i \in \mathcal{S}, x \in \Psi_i}$, we redefine $d_i(\tilde{\mathbf{p}}, b_i, \tau)$, $\mathbf{z}(\tilde{\mathbf{p}}, \mathbf{b}, \tau)$, and $H(i, x, \tau_{i,x})$ analogously.

In this section, we show how to deterministically choose these taxes.

Lemma 10.4 There exists a polynomial-time algorithm that finds a vector of taxes $\tau = (\tau_{i,x})_{i \in \mathcal{S}, x \in \Psi_i}$ such that the following desiderata are satisfied:

1. $-\epsilon < \tau_{i,x} < \epsilon$ (taxes are small),
2. $\tau_{i,x} > \tau_{i,x'}$ if $x \succ_i x'$ (taxes prefer more-preferred bundles),
3. $1 \leq \min_{i,x} \{b_i + \tau_{i,x}\} \leq \max_{i,x} \{b_i + \tau_{i,x}\} \leq 1 + \beta$ (inequality bound is preserved),
4. $b_i + \tau_{i,x} \neq b_{i'} + \tau_{i',x'}$ for $(i, x) \neq (i', x')$ (no two perturbed prices are equal),
5. there is no price $\tilde{\mathbf{p}} \in \mathcal{P}$ at which any subset of linearly dependent budget-constraint hyperplanes intersect.¹

1. The original existence proof of Budish [2011] requires only that no more than M hyperplanes intersect at any point; this causes problems in the conditional expectation argument [Budish 2011 (Step 5)].

Proof Assume w.l.o.g. that \mathbf{b} is rounded to the nearest integer multiple of $\bar{\beta}M^{-M}$: otherwise we can include this rounding in the taxes.

We proceed by induction on the pairs (i, x) of students and bundles: at each step let $\tau_{i,x}$ be much smaller in absolute value than all the taxes introduced so far. (For each i , we consider the (i, x) 's either in the order \succ_i or in the reverse order, maintaining Desideratum 2 depending on the sign of $\tau_{i,x}$.)

More precisely, if (i, x) is the v^{th} pair to be considered, then we set

$$\tau_{i,x} \in \pm \bar{\beta}M^{-2vM},$$

where the sign is chosen such that Desideratum 3 in the statement of the lemma is preserved.

Now, assume by contradiction that there exists a k -tuple $H(i_1, x_1, \tau_{i_1, x_1}), \dots, H(i_k, x_k, \tau_{i_k, x_k})$ of hyperplanes that intersect at price vector $\bar{\mathbf{p}}$, and such that the x_i 's are linearly dependent. (Note that the latter holds, in particular, for every $(M+1)$ -tuple.)

Assume further, wlog, that this is the first such k -tuple, with respect to the order of the induction. In particular, this means that $\{x_1, \dots, x_{k-1}\}$ are linearly independent. Now consider the system

$$(x_1^T \quad \dots \quad x_{k-1}^T)(\alpha) = (x_k).$$

Notice that it has rank $k-1$. We can now take $k-1$ linearly independent rows j_1, \dots, j_{k-1} such that the following system has the same unique solution α :

$$\begin{pmatrix} x_{1,j_1} & \dots & x_{k-1,j_1} \\ \vdots & \ddots & \\ x_{1,j_{k-1}} & & x_{k-1,j_{k-1}} \end{pmatrix} (\alpha) = \begin{pmatrix} x_{k,j_1} \\ \vdots \\ x_{k,j_{k-1}} \end{pmatrix}.$$

Denote

$$X = \begin{pmatrix} x_{1,j_1} & \dots & x_{k-1,j_1} \\ \vdots & \ddots & \\ x_{1,j_{k-1}} & & x_{k-1,j_{k-1}} \end{pmatrix}.$$

Since X is a square matrix of full rank, it is invertible, so we have that

$$\alpha = X^{-1} \begin{pmatrix} x_{k,j_1} \\ \vdots \\ x_{k,j_{k-1}} \end{pmatrix}.$$

Now, recall that

$$X^{-1} = \frac{1}{\det X} \begin{pmatrix} X_{1,1} & \cdots & X_{k-1,1} \\ \vdots & \ddots & \\ X_{1,k-1} & & X_{k-1,k-1} \end{pmatrix},$$

where $X_{i,j}$ is the (i, j) -cofactor of X . Finally, since X is a Boolean matrix, its determinant and all of its cofactors are integers of magnitude less than $(k-1)^{k-1} \leq M^M$. The entries of α are therefore rational fractions with numerators and denominators of magnitude less than M^M .

Now, by our assumption by contradiction, k hyperplanes intersect at $\tilde{\mathbf{p}}$:

$$\begin{pmatrix} x_1 \\ \vdots \\ x_k \end{pmatrix} \cdot \tilde{\mathbf{p}} = \begin{pmatrix} b_{i_1} + \tau_{i_1, x_1} \\ \vdots \\ b_{i_k} + \tau_{i_k, x_k} \end{pmatrix}.$$

Therefore,

$$b_{i_k} + \tau_{i_k, x_k} = x_k \cdot \tilde{\mathbf{p}} = \sum_{l=1}^{k-1} \alpha_l (x_l \cdot \tilde{\mathbf{p}}) = \sum_{l=1}^{k-1} \alpha_l (b_{i_l} + \tau_{i_l, x_l}). \quad (10.1)$$

However, if (i_k, x_k) is the v th pair added by the induction, then the following is an integer:

$$\sum_{l=1}^{k-1} (\det(X) \cdot \alpha_l) \cdot \frac{M^{2(v-1)M}}{\beta} (b_{i_l} + \tau_{i_l, x_l}).$$

By our assumption that all the budgets are rounded, $\frac{M^M}{\beta} \cdot b_{i_k}$ is also an integer. Yet $|\det(X) \cdot \frac{M^{2(v-1)M}}{\beta} \cdot \tau_{i_k, x_k}| \leq |\frac{M^{(2v-1)M}}{\beta} \cdot \tau_{i_k, x_k}| \leq M^{-M}$ is not an integer. This yields a contradiction to Equation (10.1). ■

10.3.3 Finding a Fixed Point (Steps 2–4)

This subsection describes the price adjustment correspondence of [Budish \[2011\]](#), and is brought here mostly for completeness.

We first define the price adjustment function:

$$f(\tilde{\mathbf{p}}) = t(\tilde{\mathbf{p}}) + \frac{1}{2N} \mathbf{z}(t(\tilde{\mathbf{p}}); \mathbf{b}, \tau).$$

Observe that if $\tilde{\mathbf{p}}^*$ is a fixed point $\tilde{\mathbf{p}}^* = f(\tilde{\mathbf{p}}^*)$ of f , then its truncation $t(\tilde{\mathbf{p}}^*) = \mathbf{p}^*$ defines an exact competitive equilibrium.² Yet we know that the economy may not have an exact equilibrium—and indeed f is discontinuous at the budget constraint hyperplanes, and so it is not guaranteed to have a fixed point.

Instead, we define an upper hemicontinuous, set-valued “convexification” of f :

$$F(\mathbf{p}) = \text{co} \{ \mathbf{y} : \exists \text{ a sequence } \mathbf{p}^w \rightarrow \mathbf{p}, \mathbf{p} \neq \mathbf{p}^w \in \mathcal{P} \text{ such that } f(\mathbf{p}^w) \rightarrow \mathbf{y} \}.$$

The correspondence F is upper hemicontinuous, non-empty, and convex; therefore, by Kakutani’s fixed point theorem it has a fixed point (i.e., a price vector that satisfies $\tilde{\mathbf{p}}^* \in F(\tilde{\mathbf{p}}^*)$).

By Papadimitriou [1994], finding a Kakutani fixed point of F is in PPAD.

Working with Finite Precision

To be rigorous, we need to complete a few subtle numerical details about finding a fixed point of F . We round all price vectors to the nearest integer multiple of $\delta := (\bar{\beta} M^{\frac{1}{2} - 2(v_{\max} + 1)M})$ (this precision suffices to implement the algorithm in Lemma 10.4).

At any point on the δ -grid, the price of any bundle is an integer multiple of δ , so any budget-constraint hyperplane that does not contain \mathbf{p} must be at (L_1) distance at least δ . In particular, this means that every $\delta/2$ -approximate fixed point of F is also an exact fixed point. Finally, we can use the PPAD algorithm of Papadimitriou [1994] to find a $\delta/2$ -approximate fixed point.

There is also an issue of computing the correspondence F . From the proof of Papadimitriou [1994] it follows that it suffices to compute just a single point in $F(\mathbf{p})$ for every \mathbf{p} . This is important because the number of points in $F(\mathbf{p})$ on the δ -grid may be exponential. As we mentioned earlier, every budget-constraint hyperplane that does not contain $t(\mathbf{p})$ must be at least δ -far. Therefore, we can take any point \mathbf{p}' whose truncation $t(\mathbf{p}')$ is at distance $\delta/2$ from $t(\mathbf{p})$, and does not lie on any hyperplanes. (\mathbf{p}' will not be on the δ -grid.) Because no budget-constraint hyperplanes lie between $t(\mathbf{p}')$ and $t(\mathbf{p})$, it follows that $t(\mathbf{p}) + \frac{1}{2N} \mathbf{z}(t(\mathbf{p}'); \mathbf{b}, \tau) \in F(\mathbf{p})$.

10.3.4 From a Fixed Point to an Approximate CEEI (Steps 5–9)

Lemma 10.5 Given a fixed point \mathbf{p}^* of F , we can find in polynomial time a vector of prices $\mathbf{p}^{\phi'}$ such that $\|\mathbf{z}(\mathbf{p}^{\phi'}, \mathbf{b}, \tau)\|_2 \leq \frac{\sqrt{\sigma M}}{2}$.

2. See Budish [2011, Appendix A, Step 2] for more details.

Proof We use the method of conditional expectation to derandomize Step 8 of [Budish \[2011\]](#).

Recall from the previous subsection that there exists a neighborhood around \mathbf{p}^* that does not intersect any budget-constraint hyperplanes (beyond those that contain \mathbf{p}^*). Let $1, \dots, L'$ be the indices of students whose budget-constraint hyperplanes intersect at \mathbf{p}^* . For student $i \in [L']$, let w_i be the number of corresponding hyperplanes $H(i, x_i^1, \tau_{i,x_i^1}), \dots, H(i, x_i^{w_i}, \tau_{i,x_i^{w_i}})$ intersecting at \mathbf{p}^* , and assume wlog that the superindices of $x_i^1, \dots, x_i^{w_i}$ are ordered according to \succsim_i .

Let d_i^0 be agent i 's demand when prices are slightly perturbed from \mathbf{p}^* such that all x_i^j 's are affordable. Such a perturbation exists and is easily computable because the hyperplanes are linearly independent.³ Similarly, let d_i^1 denote agent i 's demand when $x_i^2, \dots, x_i^{w_i}$ are affordable, but x_i^1 is not, and so on. Finally, let $z_{S \setminus [L']}(\mathbf{p}^*, \mathbf{b}, \tau) = d_{S \setminus [L']}(\mathbf{p}^*, \mathbf{b}, \tau) - \mathbf{q}$ be the market clearing error when considering the rest of the students. (The demand of $S \setminus [L']$ is constant in the small neighborhood \mathbf{p}^* that does not intersect any additional hyperplanes.)

By Lemma 3 of [Budish \[2011\]](#), there exist distributions a_i^f over d_i^f :

$$a_i^f \in [0, 1] \forall i \in [L'], \forall f \in \{0\} \cup [w_i]$$

$$\sum_{f=0}^{w_i} a_i^f = 1 \forall i \in [L'],$$

such that the clearing error of the *expected demand* is 0:

$$z_{S \setminus [L']}(\mathbf{p}^*, \mathbf{b}, \tau) + \sum_{i=1}^{L'} \sum_{f=0}^{w_i} a_i^f d_i^f = 0.$$

We first find such a_i^f in polynomial time using linear programming.

The existence proof then considers, for each i , a random vector $\Theta_i = (\Theta_i^1, \dots, \Theta_i^{w_i})$: the vectors are independent and in any realization θ_i satisfy $\sum_{f=0}^{w_i} \theta_i^f = 1$, while the variables each have support $\text{supp}(\Theta_i^f) = \{0, 1\}$, and expectation $\mathbb{E}[\Theta_i^f] = a_i^f$.

3. This perturbation eventually guarantees the existence of $\mathbf{p}^{\phi'}$. As we mentioned in Footnote 1, [Budish](#) does not require that the hyperplanes are linearly independent, so $\mathbf{p}^{\phi'}$ may not exist. However, it seems that $\mathbf{p}^{\phi'}$ is not actually crucial to the overall existence proof. In particular, as [Budish](#) points out, even if it exists it may be infeasible (i.e., require negative prices), so the final solution uses \mathbf{p}^* instead.

By Lemma 4 of Budish [2011], the *expected clearing error* is bounded by:

$$\mathbb{E}_{\Theta_1 \dots \Theta_{L'}} \left\| \sum_{i=1}^{L'} \sum_{f=0}^{w_i} (a_i^f - \theta_i^f) d_i^f \right\|_2^2 = \sum_{i=1}^{L'} \mathbb{E}_{\Theta_i} \left\| \sum_{f=0}^{w_i} (a_i^f - \theta_i^f) d_i^f \right\|_2^2 \leq \frac{\sigma M}{4}.$$

We now proceed by induction on the students. For each i , if the conditional expectation on $(\hat{\theta}_j)_{j < i}$ satisfies

$$\mathbb{E}_{\Theta_1 \dots \Theta_{L'}} \left[\left\| \sum_{i=1}^{L'} \sum_{f=0}^{w_i} (a_i^f - \theta_i^f) d_i^f \right\|_2^2 \mid \hat{\theta}_1, \dots, \hat{\theta}_{i-1} \right] \leq \frac{\sigma M}{4},$$

then at least one $\hat{\theta}_i$ must also satisfy the above bound. We can find such $\hat{\theta}_i$ in polynomial time by computing the conditional expectation for every feasible $\hat{\theta}_i$:

$$\begin{aligned} \mathbb{E}_{\Theta_{i+1} \dots \Theta_{L'}} \left[\left\| \sum_{j=1}^{L'} \sum_{f=0}^{w_j} (a_j^f - \theta_j^f) d_j^f \right\|_2^2 \mid \hat{\theta}_1, \dots, \hat{\theta}_i \right] &= \sum_{j=1}^i \left\| \sum_{f=0}^{w_j} (a_j^f - \hat{\theta}_j^f) d_j^f \right\|_2^2 \\ &\quad + \sum_{j=i+1}^{L'} \mathbb{E}_{\Theta_j} \left\| \sum_{f=0}^{w_j} (a_j^f - \theta_j^f) d_j^f \right\|_2^2 \\ &\quad + \sum_{j \neq h \leq i} \sum_{f=0}^{w_j} \sum_{g=0}^{w_h} (a_j^f - \hat{\theta}_j^f) (a_h^g - \hat{\theta}_h^g). \end{aligned}$$

Finally, the choice of $(\hat{\theta}_i)_{i=1}^{L'}$ induces the promised price vector $\mathbf{p}^{\phi'}$. ■

Completing the Proof of A-CEEI ∈ PPAD

The chosen $(\hat{\theta}_i)_{i=1}^{L'}$ define an allocation \mathbf{x}^* with bounded clearing error. We now follow step 9 of Budish [2011] in order to define budgets \mathbf{b}^* such that \mathbf{x}^* is the preferred consumption by all the students at price \mathbf{p}^* .

We define, for every i , $b_i^* = b_i + \tau_{i, x_i^*}$. For $i > L'$ we have $x_i^* = d_i(\mathbf{p}^*, b_i, \tau_i)$. By Desideratum 2 of Lemma 10.4, every bundle that student i prefers over x_i^* had a greater tax and was still unaffordable at \mathbf{p}^* ; it now costs more than $b_i + \tau_{i, x_i^*}$.

For $i \leq L'$, notice that every bundle x_i^\perp that i prefers over x_i^* and was exactly affordable at \mathbf{p}^* with taxes τ and budget \mathbf{b} , x_i^\perp must cost strictly more than i 's new budget b_i^* . Therefore, $(\mathbf{x}^*, \mathbf{b}^*, \mathbf{p}^*)$ is a $(\frac{\sqrt{\sigma M}}{2}, \beta)$ -CEEI. This completes the proof that finding a $(\frac{\sqrt{\sigma M}}{2}, \beta)$ -CEEI is in PPAD, and thus, together with section 10.2, also the proof of Theorem 10.3. ■



PART

**QUASI-POLYNOMIAL
TIME**

11 Birthday Repetition

This part of the book deals with several fundamental problems that admit quasi-polynomial ($n^{\log n}$) time algorithms. What can we learn from such an algorithm? On one hand, assuming the Exponential Time Hypothesis (ETH, see Hypothesis 2.1), it means that they are not NP-hard. On the other hand, it does not meet our gold standard of efficiency, polynomial time (P). Furthermore, in practical applications the logarithmic factor in the exponent is still prohibitive.

The approach we take in this book is inspired by the *birthday repetition* meta-reduction due to Aaronson et al. [2014]. The birthday repetition is best explained as a game¹ between two provers (Alice and Bob) and a verifier. The omniscient but untrusted provers want to convince the verifier that a certain 3-SAT formula is (approximately) satisfiable. The provers agree in advance on an assignment, and are then placed in separate rooms. The verifier asks Alice for the assignments of 3 variables on one randomly chosen clause, and asks Bob for the assignment of one of the variables in the same clause. It is not hard to see that if at most 90% of the clauses are satisfiable, the provers have probability at most 90% to succeed in sending assignments that both satisfy Alice's clause and agree on Bob's variable. Otherwise, the verifier can detect that Alice and Bob are lying and the 3-SAT formula is not satisfiable. But this requires coordinating the challenge-messages sent to Alice and Bob.

In the birthday-repetition version of the same two-prover one-verifier game, the verifier sends Alice a random selection of \sqrt{n} clauses, and sends Bob an independently random subset of \sqrt{n} variables. By the birthday paradox, we expect that one of Bob's variables appears in one of Alice's clauses. Thus their probability of tricking the verifier into believing that a far-from-satisfiable formula is satisfiable is not much higher than in the original (no-repetition) two-prover one-verifier game. The main advantage in this approach is that the challenges sent to Alice and Bob are independent; this is sometimes called a *free game*. This is particularly helpful when

1. The sense in which we use the word “game” here is unrelated to the game theoretic “game.”

reducing to bilinear optimization problems such as Densest k -Subgraph or Nash equilibrium (recall (1.1) in the introduction).

The other interesting feature of the birthday repetition game is that it blows up the number of possible questions to Alice and Bob, and (more importantly) the length of their answers. In particular, given a list of \sqrt{n} variables, Bob has $N \triangleq 2^{\sqrt{n}}$ distinct choices of truth assignments to those variables. Therefore, if Alice and Bob can devise an (approximately) optimal strategy much faster than $N^{\log N} = 2^n$, they would violate ETH. Hence with this approach we can obtain quasi-polynomial lower bounds on the running time that almost exactly match the running time of the best known algorithms.

There is one more property that is common to all our quasi-polynomial lower bounds: while they are all based on the birthday repetition approach, each problem introduces new obstacles and requires new ideas.

11.1 Warm-Up: Best ϵ -Nash

In order to introduce the “birthday repetition” framework, we begin with a particularly simple application. As we have already discussed, proving hardness of total problems is notoriously difficult both conceptually and technically. Often, however, there are related decision problems that admit much easier proofs of intractability. For example, [Gilboa and Zemel \[1989\]](#) proved that deciding whether a Nash equilibrium with certain welfare guarantees exists is NP-hard.

By looking at the analogous question for approximate Nash equilibrium, we circumvent the difficulty of totality. But the quasi-polynomial algorithm of [Lipton et al. \[2003\]](#) can approximately² solve this question as well. Hence, assuming ETH, this problem is still not NP-hard.

[Braverman et al. \[2015\]](#) showed a nearly matching quasi-polynomial lower bound on the running time for approximating the best ϵ -Nash equilibrium. The theorem below obtains slightly better parameters. More importantly, our proof is significantly simpler.³

2. This is a bi-criteria approximation: in quasi-polynomial time the algorithm can distinguish between a game that has a high-payoff exact (or $\epsilon/2$ -approximate) Nash equilibrium, and a game where every ϵ -approximate Nash equilibrium has low payoff. Notice, however, that this does not imply that the algorithm can approximate (to within an additive ϵ) the value of the best ϵ -approximate Nash equilibrium.

3. Briefly, in our proof we focus on a single partition of the variables, where it is guaranteed that every pair of subsets induces approximately the right number of constraints. In contrast,

Theorem 11.1 There exists a constant $\epsilon > 0$ such that, given a two-player, N -strategy game with utilities in $[0, 1]$ distinguishing between the following requires $N^{\tilde{\Omega}(\log N)}$ time (assuming the Exponential Time Hypothesis):

Completeness. The game has an exact Nash equilibrium with expected payoff > 0.99 for both players.

Soundness. In every ϵ -approximate Nash equilibrium, both players have expected payoff < 0.01 .

The rest of this section is devoted to the proof of Theorem 11.1. We first prove a small ($O(\epsilon)$) additive gap between completeness and soundness. This is the main step that exhibits the birthday repetition. We then achieve a large gap by a simple amplification step, where we introduce a new action for each player, and a spurious low-payoff pure equilibrium: in the completeness case, the players don't want to deviate from the high-payoff mixed equilibrium; but for the soundness, they are lured by unilateral deviations into the low-payoff equilibrium.

11.1.1 Constructing the First Gap

Lemma 11.1 There exist absolute constants $\epsilon, \delta > 0$ such that the following holds. Given a two-player, N -strategy game with utilities in $[0, 1]$ distinguishing between the following requires $N^{\tilde{\Omega}(\log N)}$ time (assuming the Exponential Time Hypothesis):

Completeness. The game has an exact Nash equilibrium with expected payoff c for both players.

Soundness. In every ϵ -Nash equilibrium, at least one player has expected payoff $s \leq c - \delta$.

Proof We reduce from a bipartite, constant alphabet 2CSP. By the PCP Corollary 2.1, distinguishing between a completely satisfiable instance and one where only $(1 - \eta)$ -fraction (for some constant $\eta > 0$) of the clauses are satisfiable requires time $2^{\tilde{\Omega}(n)}$, assuming ETH. We construct a game where, at every approximate Nash equilibrium, Alice's and Bob's respective mixed strategies encode an assignment (or distribution over assignments) to the 2CSP.⁴ Alice's and Bob's expected payoff will depend on the number of satisfied clauses.

Braverman et al. [2015] enumerate over all the subsets of size $\approx \sqrt{n}$. Then it is only true that *most* pairs of subsets induce the right number of constraints, which somewhat complicates the proof.

4. Note that so far *every* approximate equilibrium encodes *some* assignment (or distribution over assignments), even in the soundness case.

Construction

At a high level, our construction is composed of three subgames. The *main subgame* will test the 2CSP assignment; in this subgame we perform the “birthday repetition.” The remaining two *auxiliary subgames* enforce the structure of (approximate) equilibria; in particular, they force the players to assign (approximately) uniform probabilities in the main subgame to (almost) all 2CSP variables; this step uses a construction due to Althofer [1994].

The players’ final payoff is a weighted average of the payoffs in the three subgames. Since we want to make sure that the structural constraints on the equilibria are enforced, we place only ϵ weight on the main subgame, and weight $\frac{1-\epsilon}{2}$ on each of the auxiliary subgames.

Main Subgame. We partition Alice’s and Bob’s respective sets of variables into \sqrt{n} subsets $\{S_i\}$ and $\{T_j\}$, respectively. By Lemma 2.5, we can guarantee that each subset has at most $2\sqrt{n}$ variables, and that the number of constraints between any pair S_i, T_j is at most 8 times the expectation.

In the main subgame, each player chooses an index $i, j \in [\sqrt{n}]$ and a partial assignment for the corresponding respective subsets $\sigma \in \Sigma_A^{S_i}, \tau \in \Sigma_B^{T_j}$. If the partial assignments jointly satisfy all the induced constraints, both players receive payoff 1; otherwise they receive payoff 0.

Auxiliary Subgames. In addition to the choices of i, σ in the main subgame, Alice chooses a subset $C_A \subset [\sqrt{n}]$ of cardinality $|C_A| = \sqrt{n}/2$. In the first auxiliary subgame, Alice wants to use the set C_A to “catch” Bob’s index j : Alice has payoff 1 if $j \in C_A$, and 0 otherwise. Bob, on the other hand, tries to “escape”: his payoff is 0 if $j \in C_A$, and 1 otherwise.

The second auxiliary subgame is defined analogously (Bob chooses a set C_B of cardinality $|C_B| = \sqrt{n}/2$ and tries to catch Alice’s index i).

Analysis

Game Size. Each player has an action set of size $N \leq \sqrt{n} \cdot \Sigma^{2\sqrt{n}} \cdot \binom{\sqrt{n}}{\sqrt{n}/2} = 2^{O(\sqrt{n})}$. Therefore, assuming ETH, finding an (approximately) satisfying assignment to the LABEL COVER instance requires time $2^{\Omega(n)} = N^{\tilde{\Omega}(\log N)}$.

Completeness. If the 2CSP has a satisfying assignment, the players can play the following high-welfare equilibrium: in the main subgame, Alice and Bob choose i, j uniformly at random, and σ, τ as the respective restrictions of the satisfying

assignments. In the auxiliary subgames, Alice and Bob choose C_A, C_B uniformly at random.

Because i, j are chosen uniformly at random, they have no incentives to deviate from their choice of C_A, C_B ; similarly, since C_A, C_B are chosen uniformly at random (and the payoff in the main subgame is always 1), neither player has an incentive to deviate from the choice of i, j . Therefore this is indeed a Nash equilibrium.

Finally, the expected payoff for each player in each auxiliary subgame is $1/2$ (they have probability $1/2$ of winning), and 1 in the main subgame; in total it is $c \triangleq 1/2 + \epsilon/2$.

Soundness

First, we claim that in every ϵ -approximate Nash equilibrium (ϵ -ANE), the players' respective distributions over i, j are $O(\epsilon)$ -close to uniform. Assume by contradiction that Alice's distribution over i is 36ϵ -far from uniform. By Lemma 2.7, Bob can pick a set $C_B \in \binom{[n]}{\lfloor n/2 \rfloor}$ such that $i \in C_B$ with probability $1/2 + 9\epsilon$. Therefore Bob can guarantee a payoff of at least $1/2 + 9\epsilon$ from the second auxiliary subgame (without affecting his payoffs from the other subgames). Therefore at any ϵ -ANE, he guarantees himself more than $1/2 + 6\epsilon$ (otherwise, having a 3ϵ -improving deviation on the auxiliary subgame implies a $> \epsilon$ -deviation in total). Therefore Alice's payoff on this subgame is less than $1/2 - 6\epsilon$, whereas she can guarantee herself a payoff of $1/2$ by picking i uniformly at random. Thus she can improve her total payoff by more than 2ϵ , which is a violating deviation even after we subtract ϵ for potential loss of payoff due to changing her strategy in the main subgame.

Any mixed strategy profile induces a distribution over assignments to the original 2CSP instance: for each set S_i , we take the distribution over $\Sigma_A^{S_i}$ induced by Alice's mixed strategy restricted to actions where she picks index i (if she never picks i , choose an arbitrary assignment for S_i). Alice's expected payoff in the main subgame is equal to the probability that an assignment drawn from the induced distribution satisfies the constraints on a random pair (S_i, T_j) , where i and j are drawn according to Alice and Bob's respective mixed strategies. Since those strategies are $O(\epsilon)$ -close to uniform, Alice's expected payoff in the main subgame is within $O(\epsilon)$ of the probability that the constraints between a uniformly random (S_i, T_j) are satisfied. If the value of the 2CSP is $1 - \eta$, then the probability that any consistent assignment (in particular, one drawn from the induced distribution) satisfies a uniformly random (S_i, T_j) is at most $1 - \eta/8$; this is true because, by our construction of the partition, the η -fraction of unsatisfied constraints cannot concentrate on less than $\eta/8$ -fraction of pairs. Therefore Alice's (respectively Bob's) expected

payoff in the main subgame is at most $1 - \eta/8 + O(\epsilon)$. Since the auxiliary subgames are 1-sum games, at least one of Alice or Bob must have total expected payoff at most $1/2 + \epsilon/2(1 - \eta/8 + O(\epsilon)) < c - \eta\epsilon/16 + O(\epsilon^2)$. Choosing ϵ sufficiently small compared to η (yet still constant) completes the proof. ■

11.1.2 Gap Amplification

Proof Proof of Theorem 11.1. Consider the game from Lemma 11.1. First, we scale all the payoffs by a 10^{-4} -factor and shift them so that they are between $1 - 10^{-4}$ and 1. Now, by Lemma 11.1, it is hard to distinguish between the game having an exact Nash equilibrium with expected payoff c' and every ϵ' -Nash equilibrium having at least one player with expected payoff $< c' - \delta'$. We choose ϵ sufficiently small; in particular, $\epsilon < \epsilon'/4$, $\delta'/4 < 10^{-4}$.

We add one new action to each player. When both players play old actions, their payoffs are as in the game from Lemma 11.1 (with modified payoffs). When Alice plays the new action and Bob plays an old action, Alice's payoff is $c' - \delta'/2$, and Bob's payoff is 0. Similarly, when Alice plays an old action and Bob plays his new action, Alice's payoff is 0 and Bob's payoff is $c' - \delta'/2$. Finally, when they both play their new actions, their payoff is 10^{-3} .

Notice that playing the new action is an (exact) pure Nash equilibrium. Notice further that when the 2CSP is satisfied, the original equilibrium from the completeness of Lemma 11.1 is still a Nash equilibrium since the average payoff to both players from that equilibrium is higher than what they would get by deviating to the new strategy.

To complete the proof, we need to argue for soundness. In particular, we claim that if the value of the 2CSP is only $1 - \eta$, then every ϵ -approximate Nash equilibrium is 10^{-3} -close (in total variation distance) to the new strategy pure equilibrium. Assume by contradiction that this is not the case, i.e., there is an ϵ -approximate equilibrium where (wlog) Alice assigns probability more than 10^{-3} to old strategies. First, notice that Bob must assign at least $1/2$ probability to old strategies: whenever Bob plays the new strategy, Alice has a large incentive to deviate to her new strategy. Therefore, by a symmetric argument, Alice must also assign at least $1/2$ probability to old actions.

Finally, the mixed strategy profile restricted to the old strategies must be a 4ϵ -approximate Nash equilibrium of the scaled old game. By Lemma 11.1, the expected payoff to one of the players is at most $c' - \delta'$. But then that player has at least $\delta'/4 > \epsilon$ incentive to deviate to her new strategy. ■

Densest k -Subgraph

k -CLIQUE is one of the most fundamental problems in computer science: given a graph, decide whether it has a fully connected induced subgraph on k vertices. Since it was proven NP-complete by Karp [1972], extensive research has investigated the complexity of relaxed versions of this problem.

This work focuses on two natural relaxations of k -CLIQUE that have received significant attention from both algorithmic and complexity communities: The first one is to relax “ k ,” i.e., to look for a smaller subgraph:

Problem 12.1 Approximate max clique, informal. Given an n -vertex graph G , decide whether G contains a clique of size k , or all induced cliques of G are of size at most δk for some $1 > \delta(n) > 0$.

The second natural relaxation is to relax the “Clique” requirement, replacing it with the more modest goal of finding a subgraph that is almost a clique:

Problem 12.2 Densest k -subgraph with perfect completeness, informal. Given an n -vertex graph G containing a clique of size k , find an induced subgraph of G of size k with (edge) density at least $(1 - \varepsilon)$, for some $1 > \varepsilon > 0$. (More modestly, given an n -vertex graph G , decide whether G contains a clique of size k , or all induced k -subgraphs of G have density at most $(1 - \varepsilon)$.)

Today, after a long line of research [Arora and Safra 1998, Arora et al. 1998, Feige et al. 1996, Håstad 1999, Khot 2001, Zuckerman 2007], we have a solid understanding of the inapproximability of Problem 12.1. In particular, we know that it is NP-hard to distinguish between a graph that has a clique of size k , and a graph whose largest induced clique is of size at most $k' = \delta k$ for $\delta = 1/n^{1-\varepsilon}$ [Zuckerman 2007]. The computational complexity of the second relaxation (Problem 12.2) remains largely open. There are a couple of (very different) quasi-polynomial algorithms that guarantee finding a $(1 - \varepsilon)$ -dense k subgraph in every graph containing a k -clique [Barman 2015, Feige and Seltser 1997], suggesting that this problem is not NP-hard.

Yet we know neither polynomial-time algorithms, nor general impossibility results for this problem.

In this work we provide strong evidence that the aforementioned quasi-polynomial time algorithms for Problem 12.2 [Barman 2015, Feige and Seltser 1997] are essentially tight, assuming the (deterministic) *Exponential Time Hypothesis* (ETH) (Hypothesis 2.1). In fact, we show that under ETH, both parameters of the above relaxations are simultaneously hard to approximate:

Theorem 12.1 Main result. There exists a universal constant $\varepsilon > 0$ such that, assuming the (deterministic) Exponential Time Hypothesis, distinguishing between the following requires time $n^{\tilde{\Omega}(\log n)}$, where n is the number of vertices of G .

Completeness. G has an induced k -clique.

Soundness. Every induced subgraph of G size $k' = k \cdot 2^{-\Omega\left(\frac{\log n}{\log \log n}\right)}$ has density at most $1 - \varepsilon$.

Our result has implications for two major open problems whose computational complexity remained elusive for more than two decades: the (general) DENSEST k -SUBGRAPH problem, and the PLANTED CLIQUE problem.

Densest k -Subgraph

The DENSEST k -SUBGRAPH problem, $\text{DkS}(\eta, \varepsilon)$, is the same as (the decision version of) Problem 12.2, except that in the “completeness” case, G has a k -subgraph with density η , and in the “soundness” case, every k -subgraph is of density at most ε , where $\eta \gg \varepsilon$. Since Problem 12.2 is a special case of this problem, our main theorem can also be viewed as a new inapproximability result for $\text{DkS}(1, 1 - \varepsilon)$. We remark that the aforementioned quasi-polynomial algorithms for the “perfect completeness” regime completely break in the sparse regime, and indeed it is believed that $\text{DkS}(n^{-\alpha}, n^{-\beta})$ (for $k = n^\varepsilon$) in fact requires much more than quasi-polynomial time [Bhaskara et al. 2012]. The best to-date approximation algorithm for DENSEST k -SUBGRAPH, due to Bhaskara et al., is guaranteed to find a k -subgraph whose density is within an $\sim n^{1/4}$ -multiplicative factor of the densest subgraph of size k [Bhaskara et al. 2010], and thus $\text{DkS}(\eta, \varepsilon)$ can be solved efficiently whenever $\eta \gg n^{1/4} \cdot \varepsilon$ (this improved upon a previous $n^{1/3-\delta}$ -approximation of Feige et al. [2001]).

Several inapproximability results for DENSEST k -SUBGRAPH were known against specific classes of algorithms [Bhaskara et al. 2012] or under incomparable assumptions of Unique Games with expansion [Raghavendra and Steurer 2010] and hardness of random k -CNF [Alon et al. 2011, Feige 2002]. The most closely related

result is by Khot [2006], who shows that the DENSEST k -SUBGRAPH problem has no PTAS unless SAT can be solved in time 2^{n^ϵ} , as opposed to $2^{n^{1/2+\epsilon}}$ in this book. While Khot’s work uses a slightly weaker assumption, an important advantage of our work is simplicity: our construction is very simple, especially in contrast to Khot’s reduction.

We stress that the results of Khot [2006], as well as other aforementioned works, focus on the sub-constant density regime, i.e., they show hardness for distinguishing between a graph where every k -subgraph is sparse, and one where every k -subgraph is even sparser. In contrast, our result has perfect completeness and provides the first *additive* inapproximability for DENSEST k -SUBGRAPH—the best one can hope for as per the upper bound of Barman [2015].

Planted Clique

The PLANTED CLIQUE problem is a special case of our problem, where the inputs come from a specific distribution ($G(n, p)$ versus $G(n, p) +$ “a planted clique of size k ,” where p is some constant).¹ The *Planted Clique Conjecture* [Alon et al. 1998, 2007, Dekel et al. 2010, Feige and Krauthgamer 2000, Jerrum 1992, Kucera 1995] asserts that distinguishing between the aforementioned cases for $p = 1/2$, $k = o(\sqrt{n})$ cannot be done in polynomial time, and has served as the underlying hardness assumption in a variety of recent applications including machine-learning and cryptography (e.g., Alon et al. 2007, Berthet and Rigollet 2013) that inherently use the average-case nature of the problem, as well as in reductions to worst-case problems (e.g., Alon et al. 2011, Badanidiyuru et al. 2016, Balcan et al. 2013, Chen et al. 2015, Hazan and Krauthgamer 2011, Koiran and Zouzias 2011).

The main drawback of average-case hardness assumptions is that many average-case instances (even those of worst-case-hard problems) are in fact tractable. In recent years, the centrality of the planted clique conjecture inspired several works that obtain lower bounds in restricted models of computation [Barak et al. 2016, Deshpande and Montanari 2015, Feldman et al. 2013, Hopkins et al. 2016, Meka et al. 2015]. Nevertheless, a general lower bound for the average-case planted clique problem appears out of reach for existing techniques. Therefore, an important potential application of our result is replacing average-case assumptions such as the planted-clique conjecture, in applications that do not inherently rely on the

1. PLANTED CLIQUE typically refers to $p = 1/2$, while our hardness result is analogous to $p = 1 - \delta$, for a small constant $\delta > 0$. Nevertheless, in almost all applications of PLANTED CLIQUE, hardness for any constant p suffices.

distributional nature of the inputs (e.g., when the ultimate goal is to prove a worst-case hardness result). In such applications, there is a good chance that planted clique hardness assumptions can be replaced with a more “conventional” hardness assumption, such as the ETH, even when the problem has a quasi-polynomial algorithm. Recently, such a replacement of the planted-clique conjecture with ETH was obtained for the problem of finding an approximate Nash equilibrium with approximately optimal social welfare [Braverman et al. 2015].

We also remark that, while showing hardness for PLANTED CLIQUE from worst-case assumptions seems beyond the reach of current techniques, our result can also be seen as circumstantial evidence that this problem may indeed be hard. In particular, any polynomial time algorithm (if it exists) would have to inherently use the (rich and well-understood) structure of $G(n, p)$.

Follow-up Work by Manurangsi

The soundness in our result was greatly improved by Manurangsi [2017], who showed that even $\text{DkS}(1, n^{-1/\text{poly} \log \log n})$ may be intractable; i.e., in the NO case the maximal density is almost inverse-polynomial.

Techniques

Our simple construction is inspired by the “birthday repetition” technique: given a 2CSP (e.g., 3COL), we have a vertex for each $\tilde{\Omega}(\sqrt{n})$ -tuple of variables and assignments (respectively, 3COL vertices and colorings). We connect two vertices by an edge whenever their assignments are consistent and satisfy all 2CSP constraints induced on these tuples. In the completeness case, a clique consists of choosing all the vertices that correspond to a fixed satisfying assignment. In the soundness case (where the value of the 2CSP is low), the “birthday paradox” guarantees that most pairs of vertices (i.e., two $\tilde{\Omega}(\sqrt{n})$ -tuples of variables) will have a significant intersection (nonempty CSP constraints), thus resulting in lower densities whenever the 2CSP does not have a satisfying assignment. In the language of two-prover games, the intuition here is that the verifier has a “constant chance at catching the players in a lie if they are trying to cheat” in the game while not satisfying the CSP.

While our construction is simple, analyzing it is intricate. The main challenge is to rule out a “cheating” dense subgraph that consists of different assignments to the same variables (inconsistent colorings of the same vertices in 3COL). Intuitively, this is similar in spirit to *proving a parallel repetition theorem where the provers can answer some questions multiple times, and completely ignore other questions*. Continuing with the parallel repetition metaphor, notice that the challenge is doubled: in addition to a cheating prover correlating her answers (the standard

obstacle to parallel repetition), each prover can now also correlate which questions she chooses to answer. Our argument follows by showing that a sufficiently large subgraph must accumulate many non-edges (violations of either 2CSP or consistency constraints). To this end we introduce an information theoretic argument that carefully counts the entropy of choosing a random vertex in the dense subgraph.

We note that our entropy-based argument is completely different from all other known applications of “birthday repetition” to other problems. The main reason is that enforcing consistency is much more difficult in the case of DENSEST k -SUBGRAPH than in other applications because the problem formulation is so simple. In fact, even the follow-up work on the same problem by Manurangsi [2017] used a completely different (and quite elegant) argument that is based on counting small bi-cliques in any given subgraph.

12.1 Construction (and Completeness)

12.1.1 Construction

Let ψ be the 2CSP instance produced by the reduction in Theorem 2.2, i.e., a constraint graph over n variables with alphabet A of constant size. We construct the following graph $G_\psi = (V, E)$:

- Let $\rho := \sqrt{n} \log \log n$ and $k := \binom{n}{\rho}$.
- Vertices of G_ψ correspond to all possible assignments (colorings) to all ρ -tuples of variables in ψ , i.e., $V = [n]^\rho \times A^\rho$. Each vertex is of the form $v = (y_{x_1}, y_{x_2}, \dots, y_{x_\rho})$ where $\{x_1, \dots, x_\rho\}$ are the chosen variables of v , and y_{x_i} is the corresponding assignment to variable x_i .
- If $v \in V$ violates any 2CSP constraints, i.e., if there is a constraint on (x_i, x_j) in ψ that is not satisfied by (y_{x_i}, y_{x_j}) , then v is an isolated vertex in G_ψ .
- Let $u = (y_{x_1}, y_{x_2}, \dots, y_{x_\rho})$ and $v = (y'_{x'_1}, y'_{x'_2}, \dots, y'_{x'_\rho})$. $(u, v) \in E$ iff:
 - (u, v) does not violate any consistency constraints: for every shared variable x_i , the corresponding assignments agree, $y_{x_i} = y'_{x_i}$; and
 - (u, v) also does not violate any 2CSP constraints: for every 2CSP constraint on (x_i, x'_j) (if it exists), the assignment $(y_{x_i}, y'_{x'_j})$ satisfies the constraint.

Notice that the size of our reduction (number of vertices of G_ψ) is $N = \binom{n}{\rho} \cdot |A|^\rho = 2^{\tilde{O}(\sqrt{n})}$.

Completeness. If $\text{OPT}(\psi) = 1$, then G_ψ has a k -clique: Fix a satisfying assignment for ψ , and let S be the set of all vertices that are consistent with this assignment. Notice that $|S| = \binom{n}{\rho} = k$. Furthermore, its vertices do not violate any consistency constraints (since they agree with a single assignment) or **2CSP** constraints (since we started from a satisfying assignment).

12.2 Soundness

Suppose that $\text{OPT}(\psi) < 1 - \eta$, and let $\epsilon_0 > 0$ be some constant to be determined later. We shall show that for any subset S of size $k' \geq k \cdot |V|^{-\epsilon_0 / \log \log |V|}$, $\text{den}(S) < 1 - \delta$, where δ is some constant depending on η . The remainder of this section is devoted to proving the following theorem:

Theorem 12.2 If $\text{OPT}(\psi) < 1 - \eta$, then $\forall S \subset V$ of size $k' \geq k \cdot |V|^{-\epsilon_0 / \log \log |V|}$, $\text{den}(S) < 1 - \delta$ for some constant δ .

12.2.1 Setting up the Entropy Argument

Fix some subset S of size k' , and let $v \in_R S$ be a uniformly chosen vertex in S (recall that v is a vector of ρ coordinates, corresponding to labels for a subset of ρ chosen variables). For $i \in [n]$, let X_i denote the indicator variable associated with v such that $X_i = 1$ if the i -th variable appears in v and 0 otherwise. We let Y_i represent the coloring assignment (label) for the i -th variable whenever $X_i = 1$. Throughout the proof, let

$$W_{i-1} = X_{<i}, Y_{<i}$$

denote the i -th prefix corresponding to v . We can write:

$$\begin{aligned} H(Y_i | W_{i-1}, X_i) &= \Pr[X_i = 0] \cdot H(Y_i | W_{i-1}, X_i = 0) \\ &\quad + \Pr[X_i = 1] \cdot H(Y_i | W_{i-1}, X_i = 1) \\ &= \Pr[X_i = 1] \cdot H(Y_i | W_{i-1}, X_i = 1) \end{aligned}$$

since $H(Y_i | W_{i-1}, X_i = 0) = 0$. Notice that since (XY) and v determine each other, and v was uniform on a set of size $|S| = k'$, we have

Observation 12.1 $H(XY) = \log k'$.

Thus, in total, the choice of challenge and the choice of assignments should contribute $\log k'$ to the entropy of v . If much of the entropy comes from the assignment

distribution (conditioned on the fixed challenge variables), we will show that S must have many consistency violations, implying that S is sparse. If, on the other hand, almost all the entropy comes from the challenge distribution, we will show that this implies many CSP constraint violations (implied by the soundness assumption). From now on, we denote

$$\alpha_i := H(X_i | X_{<i}, Y_{<i}) \quad \text{and} \quad \beta_i := H(Y_i | X_{\leq i}, Y_{<i}).$$

When conditioning on the i -th prefix, we shall write $\alpha_i(w_{i-1}) := H(X_i | X_{<i}, Y_{<i} = w_{i-1})$, and similarly for $\beta_i(\cdot)$. Also for brevity, we denote

$$q_i := \Pr[X_i = 1] \quad \text{and} \quad q_i(w_{i-1}) := \Pr[X_i = 1 | w_{i-1}].$$

Prefix Graphs

The consistency constraints induce, for each i , a graph over the prefixes: the vertices are the prefixes, and two prefixes are connected by an edge if their labels are consistent. (We can ignore the **2CSP** constraints for now—the prefix graph will be used only in the analysis of the consistency constraints.) Formally,

Definition 12.1 Prefix graph. For $i \in [n + 1]$ let the i -th prefix graph, $G_i = (V_i, E_i)$, be defined over the prefixes of length $i - 1$ as follows. We say that w_{i-1} is a neighbor of σ_{i-1} if they do not violate any consistency constraints. Namely, for all $j < i$, if $X_j = 1$ for both w_{i-1} and σ_{i-1} , then w_i and σ_i assign the same label Y_j .

In particular, we will heavily use the following notation: let $\mathcal{N}(w_{i-1})$ be the *prefix neighborhood* of w_{i-1} ; i.e., it is the set of all prefixes (of length $i - 1$) that are consistent with w_{i-1} . For technical issues of normalization, we let $w_{i-1} \in \mathcal{N}(w_{i-1})$, i.e., all the prefixes have self-loops.

Notice that G_{n+1} is defined over the vertices of S (the original subgraph). The set of edges on S is contained in the set of edges of G_{n+1} , since in the latter we only remove pairs that violated consistency constraints (recall that we ignore the **2CSP** constraints).

Unless stated otherwise, we always think of prefixes as weighted by their probabilities. Naturally, we also define the weighted degree and weighted edge density of the prefix graph.

Definition 12.2 Prefix degree and density. The *prefix degree* of w_{i-1} is given by:

$$\deg(w_{i-1}) = \sum_{\sigma_{i-1} \in \mathcal{N}(w_{i-1})} \Pr[\sigma_{i-1}].$$

Similarly, we define the *prefix density* of G_i as:

$$\text{den}(G_i) = \sum_{w_{i-1}} \sum_{\sigma_{i-1} \in \mathcal{N}(w_{i-1})} \Pr[w_{i-1}] \cdot \Pr[\sigma_{i-1}].$$

When it is clear from the context, we henceforth drop the *prefix* qualification, and simply refer to the *neighborhood* or *degree*, etc., of w_{i-1} .

Notice that in G_{n+1} , the probabilities are uniformly distributed. In particular, $\text{den}(G_{n+1}) \geq \text{den}(S)$, since, as we mentioned earlier, the set of edges in S is contained in that of G_{n+1} . Finally, observe also that because we accumulate violations, the density of the prefix graphs is monotonically non-increasing with i .

Observation 12.2

$$\text{den}(G_1) \geq \dots \geq \text{den}(G_{n+1}) \geq \text{den}(S).$$

Useful Approximations

We use the following bounds on α_i and β_i many times throughout the proof:

Fact 12.1

$$\alpha_i = \mathbb{E} [H(q_i(w_{i-1}))] \leq H(\mathbb{E}[q_i(w_{i-1})]) = H(q_i).$$

Fact 12.2

$$\beta_i = \mathbb{E} [\beta_i(w_{i-1})] \leq \mathbb{E} [q_i(w_{i-1}) \cdot \log |A|] = q_i \log |A|.$$

Proof The bound on α_i follows from concavity of entropy (Fact 2.4). For the second bound, observe that β_i is maximized by spreading q_i mass uniformly over alphabet A . ■

We also recall some elementary approximations to logarithms and entropies that will be useful in the analysis. The proofs are deferred to the appendix.

Fact 12.3 For $k = \binom{n}{\rho}$, then,

$$\log k = nH\left(\frac{\rho}{n}\right) \pm O(\log n) = \left(\frac{1}{2} - o(1)\right)\rho \log n.$$

More useful to us will be the following bounds on $\log k'$:

Fact 12.4 Let $\epsilon_1 \geq 5\epsilon_0$, and k, k', V, n, ρ as specified in the construction. Then,

$$\log k' \geq \max \left\{ \log k, nH\left(\frac{\rho}{n}\right) \right\} - \underbrace{\epsilon_1 \log k / \log n}_{\approx \frac{\epsilon_1}{2} \cdot \rho}.$$

This means that most indices i should contribute roughly $H\left(\frac{\rho}{n}\right)$ entropy to the choice of v .

We will also need the following bound, which relates the entropies of a very biased coin and a slightly less biased one:

Fact 12.5 Let $1/n \ll |\nu| \ll 1$. Then

$$\begin{aligned} H\left(\frac{1+\nu}{n}\right) &= H\left(\frac{1}{n}\right) - \frac{\nu}{n} \log \frac{1}{n} - (\log e) \frac{\nu^2}{2n} \\ &\quad + O(n^{-2}) + O\left(\frac{\nu^3}{n}\right). \end{aligned}$$

A Useful Lemma: Bias Implies Less Entropy

In Fact 12.1 we saw that, always, $\alpha_i \leq H(q_i)$. Equality happens only if the q_i -mass is evenly distributed across all prefixes. We argue that if q_i is far from evenly distributed, then the inequality is also far from tight. In particular:

Claim 12.1 Let $B \subset V_i$ be a subset of prefixes such that for some $0 < a < b < 1$,

1. $\sum_{w_{i-1} \in B} \Pr[w_{i-1}] \leq b$; but also
2. $\sum_{w_{i-1} \in B} \Pr[w_{i-1}] q_i(w_{i-1}) > a$.

Then $\alpha_i \leq H(q_i) - q_i \mathcal{D}_{\text{KL}}(a \| b)$.

Proof Abusing notation, let $B(\cdot)$ be the indicator variable for $W_{i-1} \in B$. By the data processing inequality (Fact 2.7),

$$\begin{aligned} \alpha_i &= H(X_i | W_{i-1}) \\ &\leq H(X_i | B(W_{i-1})) \\ &= H(X_i) - I(X_i; B(W_{i-1})). \end{aligned} \tag{12.1}$$

Since we can write mutual information as expected Kullbeck-Leibler Divergence (Fact 2.8), and Kullbeck-Leibler Divergence is non-negative, we get

$$\begin{aligned} I(X_i; B(W_{i-1})) &= \mathbb{E}_{x_i} [\mathcal{D}_{\text{KL}}(B(W_{i-1}) | x_i \| B(W_{i-1}))] \\ &\geq q_i \mathcal{D}_{\text{KL}}(\Pr[B(W_{i-1}) = 1 | x_i = 1] \| B(W_{i-1}) = 1) \\ &\geq q_i \mathcal{D}_{\text{KL}}(a \| b), \end{aligned}$$

where the second inequality follows from the premise assumptions that $\Pr[B(W_{i-1})] \leq b$ and $\Pr[B(W_{i-1}) = 0 | x_i = 1] \geq a$.

Plugging into (12.1), we have:

$$\alpha_i \leq H(q_i) - q_i D_{\text{KL}}(a \| b). \quad \blacksquare \quad (12.2)$$

12.2.2 Consistency Violations

In this section, we show that if the total entropy contribution of the assignments ($\sum_i \beta_i$) is large, there are many consistency violations between vertices, which lead to constant density loss. First, we show that if $\sum_i \beta_i > 5\varepsilon_1 \log k / \log n$, then at least a constant fraction of such entropy is concentrated on *good variables* that contribute to both “types” of entropy.

Definition 12.3 Good variables. We say that an index i is *good* if

- $\alpha_i \geq H(q_i) - 2q_i \log |A|$;
- $\beta_i \geq \frac{1}{2}\varepsilon_1 q_i$,

where ε_1 is a constant to be determined later in the proof.

Claim 12.2 For any constant ε_1 , if $\sum_i \beta_i > 5\varepsilon_1 \log k / \log n$,

$$\sum_{\text{good } i\text{'s}} q_i^2 \geq \left(\frac{1}{5}\varepsilon_1 \rho\right)^2 / (n \log^2 |A|) = \Omega(\rho^2/n).$$

Proof We want to show that many of the indices i have both a large α_i and a large β_i simultaneously. Let $\iota \subseteq [n]$ denote the set of i such that $\alpha_i + \beta_i < H(q_i) - q_i \log |A|$. We can write

$$\sum_{i \in [n]} (\alpha_i + \beta_i) = \sum_{i \in \iota} (\alpha_i + \beta_i) + \sum_{i \notin \iota} (\alpha_i + \beta_i).$$

Using Facts 12.1 and 12.2, we have

$$\sum_{i \in [n]} (\alpha_i + \beta_i) \leq \sum_{i \in \iota} (H(q_i) - \beta_i) + \sum_{i \notin \iota} (H(q_i) + \beta_i). \quad (12.3)$$

Because the subgraph is of size k' , from the expansion of $\log k'$ (Fact 12.4),

$$\begin{aligned} \sum_{i \in [n]} (\alpha_i + \beta_i) &\geq nH\left(\frac{\rho}{n}\right) - \varepsilon_1 \log k / \log n \\ &\geq \sum H(q_i) - \varepsilon_1 \log k / \log n, \end{aligned}$$

where the second inequality follows from the concavity of entropy. Plugging into (12.3), we have

$$\begin{aligned} \sum_{i \notin t} \beta_i &\geq \sum_{i \in t} \beta_i - \epsilon_1 \log k / \log n \\ &= \left(\sum_i \beta_i - \sum_{i \notin t} \beta_i \right) - \epsilon_1 \log k / \log n. \end{aligned}$$

Rearranging, we get

$$\sum_{i \notin t} \beta_i \geq \frac{1}{2} \sum_i \beta_i - \epsilon_1 \log k / \log n. \quad (12.4)$$

For all the i 's in the LHS summation, $\alpha_i \geq H(q_i) - 2q_i \log |A|$ by Fact 12.2. From now on, we will consider only i 's that satisfy this condition. Now, using the premise on $\sum_i \beta_i$ and (12.4), we have:

$$\begin{aligned} \sum_{i: \alpha_i \geq H(q_i) - 2q_i \log |A|} \beta_i &\geq (5/2 - 1)\epsilon_1 \log k / \log n \\ &\geq 0.7\epsilon_1\rho, \end{aligned}$$

where the second inequality follows from our approximation for $\log k$ (Fact 12.3).

We want to further restrict our attention to i 's for which β_i is at least $\frac{1}{2}\epsilon_1 q_i$ (aka good i 's). Note that the above inequality can be decomposed to

$$\sum_{\text{good } i\text{'s}} \beta_i + \sum_{\substack{i: \alpha_i \geq H(q_i) - 2q_i \log |A| \\ \beta_i < \frac{1}{2}\epsilon_1 q_i}} \beta_i \geq 0.7\epsilon_1\rho.$$

Now via a simple sum bound,

$$\sum_{\substack{i: \alpha_i \geq H(q_i) - 2q_i \log |A| \\ \beta_i < \frac{1}{2}\epsilon_1 q_i}} \beta_i \leq \frac{1}{2}\epsilon_1 \sum_i q_i = \frac{1}{2}\epsilon_1\rho.$$

Rearranging, we get

$$\sum_{\text{good } i\text{'s}} \beta_i \geq \frac{1}{5}\epsilon_1\rho.$$

By Cauchy-Schwartz we have:

$$\sum_{\text{good } i\text{'s}} \beta_i^2 \geq \left(\frac{1}{5}\epsilon_1\rho \right)^2 / n.$$

Finally, since $\beta_i \leq q_i \log |A|$,

$$\sum_{\text{good } i\text{'s}} q_i^2 \geq \left(\frac{1}{5}\varepsilon_1\rho\right)^2 / (n \log^2 |A|). \quad \blacksquare$$

In the same spirit, we now define a notion of a “good” prefix. Intuitively, conditioning on a good prefix leaves a significant amount of entropy on the i -th index. We also require that a good prefix has a high prefix degree; that is, it has many neighbors it could potentially lose when revealing the i -th label.

Definition 12.4 Good prefixes. We say w_{i-1} is a *good* prefix if:

- i is good;
- $\sum_{\sigma_{i-1} \in \mathcal{N}(w_{i-1})} q_i(\sigma_{i-1}) \Pr[\sigma_{i-1}] \geq (1 - \varepsilon_2)q_i$;
- $\beta_i(w_{i-1}) \geq \varepsilon_3 q_i(w_{i-1})$,

for $\varepsilon_3 = (\varepsilon_4 + \kappa) \log |A|$, with ε_4 an arbitrarily small constant that denotes the fraction of assignments that disagree with the majority of the assignments, $\kappa = \Theta(1/\log |A|)$, and ε_2 a constant that satisfies $\delta = \left(\frac{\varepsilon_2}{|A|^{2/\varepsilon_2}}\right)^4$, with $\text{den}(S) = 1 - \delta$.

In the following claim, we show that these prefixes contribute some constant fraction of entropy, assuming that our subset is dense.

Claim 12.3 If $\text{den}(S) > 1 - \delta$, where $\delta = \left(\frac{\varepsilon_2}{|A|^{2/\varepsilon_2}}\right)^4$ and $\varepsilon_1 \geq 4\varepsilon_2 \log |A| + 8\varepsilon_3$, then for every good index i , it holds that

$$\sum_{\text{good } w_{i-1}\text{'s}} \Pr[w_{i-1}] \beta_i(w_{i-1}) \geq \beta_i/4.$$

Proof We begin by proving that most prefixes satisfy the degree condition of Definition 12.4. Let w_{i-1} be *popular* if i is a good variable and its degree in the prefix graph G_i is at least $\text{deg}(w_{i-1}) := \sum_{\sigma_{i-1} \in \mathcal{N}(w_{i-1})} \Pr[\sigma_{i-1}] \geq 1 - \sqrt{\delta}$. Recall that $\text{den}(G_i) \geq \text{den}(S) \geq (1 - \delta)$ (by Observation 12.2). Thus by Markov inequality, at most $\sqrt{\delta}$ -fraction of the prefixes are unpopular.

We now argue that:

$$\sum_{\text{unpopular } w_{i-1}\text{'s}} \Pr[w_{i-1}] q_i(w_{i-1}) \leq \varepsilon_2 q_i. \quad (12.5)$$

Otherwise, by Claim 12.1, $\alpha_i \geq H(q_i) - q_i D_{\text{KL}}(\varepsilon_2 \| \sqrt{\delta})$. On the other hand, recall that since i is good, $\alpha_i \geq H(q_i) - 2q_i \log |A|$. Recall also that $\delta = \left(\frac{\varepsilon_2}{|A|^{2/\varepsilon_2}}\right)^4$, and therefore $D_{\text{KL}}(\varepsilon_2 \| \sqrt{\delta}) \geq 2 \log |A|$. Thus, we get a contradiction.

Inequality (12.5) implies that even if the assignment is uniform over the alphabet, the contribution to $\sum \beta_i$ from unpopular prefixes is small:

$$\begin{aligned} \sum_{\text{unp.}} \Pr[w_{i-1}] \beta_i(w_{i-1}) &\leq \sum_{\text{unp.}} \Pr[w_{i-1}] q_i(w_{i-1}) \log |A| \\ &\leq \varepsilon_2 q_i \log |A| \\ &\leq \frac{1}{4} \varepsilon_1 q_i \leq \frac{1}{2} \beta_i, \end{aligned}$$

where the first inequality follows from Fact 12.2, the second from (12.5), the third from our setting of $\varepsilon_1 \geq 4\varepsilon_2 \log |A|$, and the fourth from $\beta_i \geq \frac{1}{2}\varepsilon_1 q_i$ since i is good. Therefore,

$$\begin{aligned} \sum_{\text{pop.}} \Pr[w_{i-1}] \beta_i(w_{i-1}) &= \beta_i - \sum_{\text{unp.}} \Pr[w_{i-1}] \beta_i(w_{i-1}) \\ &\geq \beta_i/2. \end{aligned}$$

Using a similar argument, we show that for any popular w_{i-1} , most of the q_i mass is concentrated on its neighbors. Consider any popular w_{i-1} , and let $\mathcal{N}^C(w_{i-1})$ denote the complement of $\mathcal{N}(w_{i-1})$. Then we can rewrite α_i as:

$$\begin{aligned} \alpha_i &= \sum_{\sigma_{i-1} \in \mathcal{N}(w_{i-1})} \Pr[\sigma_{i-1}] \alpha_i(\sigma_{i-1}) \\ &\quad + \sum_{\sigma_{i-1} \in \mathcal{N}^C(w_{i-1})} \Pr[\sigma_{i-1}] \alpha_i(\sigma_{i-1}). \end{aligned}$$

Notice that since w_{i-1} is popular, $\mathcal{N}^C(w_{i-1})$ has measure at most $\sqrt{\delta}$. Thus, if an ε_2 -fraction of the q_i mass is concentrated on $\mathcal{N}^C(w_{i-1})$, Claim 12.1 implies:

$$\alpha_i \leq H(q_i) - q_i \mathbf{D}_{\text{KL}}(\varepsilon_2 \| \sqrt{\delta}),$$

which (as in (12.5)) would yield a contradiction to i being a good variable. Therefore every popular prefix also satisfies the q_i -weighted condition on the degree:

$$\sum_{\sigma_{i-1} \in \mathcal{N}(w_{i-1})} \Pr[\sigma_{i-1}] q_i(\sigma_{i-1}) \geq (1 - \varepsilon_2) q_i. \quad (12.6)$$

Recall that a prefix w_{i-1} is good if it also satisfies $\beta_i(w_{i-1}) \geq \varepsilon_3 \cdot q_i(w_{i-1})$. Fortunately, prefixes that violate this condition (i.e., those with small $\beta_i(w_{i-1})$) cannot account for much of the weight on β_i :

$$\sum_{\beta_i(w_{i-1}) < \varepsilon_3 q_i(w_{i-1})} \Pr[w_{i-1}] \beta_i(w_{i-1}) \leq \varepsilon_3 q_i.$$

Since i is good and $\varepsilon_1 \geq 8\varepsilon_3$, this implies:

$$\sum_{\text{good } w_{i-1}'\text{'s}} \Pr[w_{i-1}] \beta_i(w_{i-1}) \geq \beta_i/2 - \varepsilon_3 q_i \geq \beta_i/4$$

since

$$\varepsilon_3 q_i \leq \frac{1}{8} \varepsilon_1 q_i \leq \frac{1}{4} \beta_i$$

where the last inequality follows from i being good. ■

Corollary 12.1 For every good index i ,

$$\sum_{\text{good } w_{i-1}'\text{'s}} \Pr[w_{i-1}] q_i(w_{i-1}) \geq \frac{\varepsilon_1}{8 \log |A|} q_i.$$

Proof

$$\begin{aligned} \sum_{\text{good}} \Pr[w_{i-1}] q_i(w_{i-1}) &\geq \sum_{\text{good}} \Pr[w_{i-1}] \beta_i / \log |A| \\ &\geq \beta_i / (4 \log |A|) \\ &\geq \frac{\varepsilon_1}{8 \log |A|} q_i, \end{aligned}$$

where the first inequality follows by Fact 12.2, the second by Claim 12.3, and the last by definition of good i 's. ■

With Claim 12.2 and Corollary 12.1, we are ready to prove the main lemma of this section:

Lemma 12.1 Labeling entropy bound. If $\sum_i H(Y_i | X_{\leq i}, Y_{< i}) > \frac{5\varepsilon_1 \log k}{\log n}$, then $\text{den}(S) < 1 - \delta$.

Proof Assume for a contradiction that $\text{den}(S) \geq 1 - \delta$. For prefix w_{i-1} , let $\mathcal{D}_{w_{i-1}}$ denote the induced distribution on labels to the i -th variable, conditioned on w_{i-1} and $x_i = 1$. (If $q_i(w_{i-1}) = 0$, take an arbitrary distribution.) After revealing each variable i , the loss in prefix density is given by the probability of “fresh violations”: the sum over

all prefix edges (w_{i-1}, σ_{i-1}) of the probability that they assign different labels to the i -th variable:

$$\begin{aligned} \text{den}(G_i) - \text{den}(G_{i+1}) &= \sum_{w_{i-1}} \sum_{\sigma_{i-1} \in \mathcal{N}(w_{i-1})} \dots & (12.7) \\ &\left(\Pr[w_{i-1}] \Pr[\sigma_{i-1}] q_i(w_{i-1}) q_i(\sigma_{i-1}) \right) \Pr_{\substack{Y_i \sim \mathcal{D}_{w_{i-1}} \\ Y'_i \sim \mathcal{D}_{\sigma_{i-1}}}} [Y_i \neq Y'_i]. \end{aligned}$$

We now lower-bound $\Pr_{\mathcal{D}_{w_{i-1}} \times \mathcal{D}_{\sigma_{i-1}}} [Y_i \neq Y'_i]$ for good w_{i-1} (notice that we assume nothing about σ_{i-1}). A simple calculation shows that for $\kappa < 1/2$, if

$$\begin{aligned} \beta_i(w_{i-1}) &\geq \\ &\left(\kappa \log |A| - \kappa \log \kappa - (1 - \kappa) \log(1 - \kappa) \right) q_i(w_{i-1}), \end{aligned}$$

then the probability mass (under $\mathcal{D}(w_{i-1})$) on the most common label is at most $1 - \kappa$. Observe that this probability is an upper bound on $\Pr_{\mathcal{D}_{w_{i-1}} \times \mathcal{D}_{\sigma_{i-1}}} [Y_i = Y'_i]$. For good w_{i-1} , we indeed have

$$\beta_i(w_{i-1}) \geq \varepsilon_3 q_i(w_{i-1}) \geq \left(\varepsilon_4 \log |A| - \varepsilon_4 \log \varepsilon_4 - (1 - \varepsilon_4) \log(1 - \varepsilon_4) \right) q_i(w_{i-1}),$$

where the second inequality follows from the choice of ε_4 . Therefore

$$\Pr_{\mathcal{D}_{w_{i-1}} \times \mathcal{D}_{\sigma_{i-1}}} [Y_i \neq Y'_i] \geq \varepsilon_4.$$

We now have, for every good index i ,

$$\begin{aligned} \text{den}(G_i) - \text{den}(G_{i+1}) &\geq \sum_{\text{good } w_{i-1}'\text{s}} \sum_{\sigma_{i-1} \in \mathcal{N}(w_{i-1})} \varepsilon_4 \\ &\cdot \left(\Pr[w_{i-1}] \Pr[\sigma_{i-1}] q_i(w_{i-1}) q_i(\sigma_{i-1}) \right) \\ &\geq \varepsilon_4 q_i(1 - \varepsilon_2) \sum_{\text{good } w_{i-1}'\text{s}} \Pr[w_{i-1}] q_i(w_{i-1}) \\ &\geq \frac{\varepsilon_1 \varepsilon_4}{10 \log |A|} q_i^2, \end{aligned}$$

where the first inequality follows by Eq. (12.7); the second by definition of good prefix; and the last by Corollary 12.1 and $\varepsilon_2 < 0.2$.

Finally, summing over all good i 's, we get a negative density for S , which is, of course, a contradiction. By Observation 12.2 we have:

$$\begin{aligned}
1 - \text{den}(S) &\geq \text{den}(G_1) - \text{den}(G_{n+1}) \\
&= \sum_i \text{den}(G_i) - \text{den}(G_{i+1}) \\
&\geq \sum_{\text{good } i\text{'s}} \text{den}(G_i) - \text{den}(G_{i+1}) \\
&\geq \sum_{\text{good } i\text{'s}} \left(\frac{\varepsilon_1 \varepsilon_4}{10 \log |A|} \right) q_i^2 \\
&\geq \left(\frac{\varepsilon_1^3 \varepsilon_4}{250 \log^3 |A|} \right) \rho^2 / n = \Omega(\rho^2 / n),
\end{aligned}$$

where the last inequality follows by Claim 12.2. ■

12.2.3 2CSP Violation

Intuitively, if $\sum_i H(X_i | X_{<i}, Y_{<i})$ is large, then the subgraph approximately corresponds to assignments to all subsets in $\binom{[n]}{\rho}$. More specifically, in this section we show that most of the constraints appear approximately as frequently as we expect. Since in any assignment a constant fraction of them must be violated, this implies (eventually) that a constant fraction of the edges have a violated constraint.

First, we show that most of the variables appear approximately as frequently as we expect by showing that most of them are “typical.”

Definition 12.5 Typical variables. Prefix w_{i-1} is *typical* if

$$(1 - \varepsilon_5) \cdot \rho / n < \Pr[X_i = 1 | w_{i-1}] < (1 + \varepsilon_5) \cdot \rho / n,$$

where ε_5 is some constant such that $\left(\frac{\log e}{8}\right) \varepsilon_5^4 > 14\varepsilon_1$.

Similarly, we say that variable x_i is *typical* if

$$\sum_{\text{typical } w_{i-1}\text{'s}} \Pr[w_{i-1}] \geq 1 - \varepsilon_5.$$

Claim 12.4 If $\sum_i H(X_i | X_{<i}, Y_{<i}) \geq \left(1 - \frac{6\varepsilon_1}{\log n}\right) \log k = \log k - \Theta(\rho)$, then all but at most $\varepsilon_5 n$ variables are typical.

Proof Assume by contradiction that there are $\varepsilon_5 n$ atypical variables. That is, $\varepsilon_5 n / 2$ variables x_i appear with probability at least $(1 + \varepsilon_5) \cdot \rho / n$ (or at most $(1 - \varepsilon_5) \cdot \rho / n$) for an $(\varepsilon_5 / 2)$ -fraction of the prefixes w_{i-1} . Now, subject only to this constraint and

maintaining the correct expected number of variables in each vertex, the entropy is maximized by spreading the $(\varepsilon_5^3/4)$ -loss in frequency evenly across all other prefixes and variables. That is, on the atypical prefixes, labels are assigned with probability $(1 + \varepsilon_5)\rho/n$, and with probability $\left(1 - \frac{\varepsilon_5^3/4}{1 - \varepsilon_5^2/4}\right)\rho/n$ on the rest. Thus,

$$\begin{aligned} \sum_i H(X_i|X_{<i}, Y_{<i}) &< \frac{\varepsilon_5^2}{4}n \cdot H((1 + \varepsilon_5)\rho/n) \\ &\quad + \left(1 - \frac{\varepsilon_5^2}{4}\right)nH\left(\left(1 - \frac{\varepsilon_5^3/4}{1 - \varepsilon_5^2/4}\right)\rho/n\right). \end{aligned}$$

Recall from Fact 12.5 the expansion of the entropy function:

$$H\left(\frac{1+v}{n}\right) = H\left(\frac{1}{n}\right) - \frac{v}{n} \log \frac{1}{n} - \left(\frac{\log e}{2}\right) \frac{v^2}{n} + O(n^{-2}) + O\left(\frac{v^3}{n}\right).$$

Therefore,

$$\begin{aligned} \sum_i H(X_i|X_{<i}, Y_{<i}) &< \frac{\varepsilon_5^2}{4}n \left[H\left(\frac{\rho}{n}\right) - \varepsilon_5 \frac{\rho}{n} \log \frac{\rho}{n} - \left(\frac{\log e}{2}\right) \frac{\rho}{n} \cdot \varepsilon_5^2 + O\left(\left(\frac{\rho}{n}\right)^2\right) + O\left(\frac{\rho}{n} \varepsilon_5^3\right) \right] \\ &\quad + \left(1 - \frac{\varepsilon_5^2}{4}\right)n \left[H\left(\frac{\rho}{n}\right) + \left(\frac{\varepsilon_5^3/4}{1 - \varepsilon_5^2/4}\right) \frac{\rho}{n} \log \frac{\rho}{n} + O\left(\left(\frac{\rho}{n}\right)^2\right) + O\left(\frac{\rho}{n} \varepsilon_5^6\right) \right] \\ &= n \left[H\left(\frac{\rho}{n}\right) - \left(\frac{\log e}{8}\right) \frac{\rho}{n} \cdot \varepsilon_5^4 + O\left(\left(\frac{\rho}{n}\right)^2\right) + O\left(\frac{\rho}{n} \varepsilon_5^5\right) \right]. \end{aligned}$$

Recall that $-2 \log \frac{\rho}{n} < \log n$. Thus for $\left(\frac{\log e}{8}\right) \varepsilon_5^4 > 14\varepsilon_1$, we have that

$$\begin{aligned} \left(\frac{\log e}{8}\right) \frac{\rho}{n} \cdot \varepsilon_5^4 - O\left(\left(\frac{\rho}{n}\right)^2\right) - O\left(\frac{\rho}{n} \varepsilon_5^5\right) &> \frac{\rho}{n} \cdot 12\varepsilon_1 > -\frac{\rho}{n} \log \frac{\rho}{n} \cdot 24\varepsilon_1 / \log n \\ &> (12\varepsilon_1 / \log n) H\left(\frac{\rho}{n}\right) \end{aligned}$$

and therefore,

$$\sum_i H(X_i|X_{<i}, Y_{<i}) < (1 - 12\varepsilon_1 / \log n)nH\left(\frac{\rho}{n}\right) < (1 - 6\varepsilon_1 / \log n) \log k,$$

where the second inequality follows from Fact 12.3. Thus we have reached a contradiction. Notice that the $\left(\frac{\log e}{8}\right) \frac{\rho}{n} \cdot \varepsilon_5^4$ term of missing entropy is symmetric (but not

the negligible higher order terms); i.e., the same derivation can be used to show a contradiction when many variables appear with probability less than $(1 - \varepsilon_5)\rho/n$. ■

Definition 12.6 Let $\mathcal{I}(u, v)$ be defined as the number of (i, j) pairs such that

- In the original **2CSP** instance ψ , there exists an edge (constraint) between typical variables x_i and x_j .
- $X_i = 1$ for u and $X_j = 1$ for v .
- u_{i-1} and v_{j-1} are typical prefixes, where u_{i-1} denotes the prefix represented by u for $X_{<i}, Y_{<i}$, similarly for v_{j-1} .

Intuitively, $\mathcal{I}(u, v)$ is the number of “tests” of **2CSP**-constraints between vertices u, v , when restricting to typical prefixes and variables. We now use the properties of typical prefixes and constraints to show that $\mathcal{I}(u, v)$ behaves “nicely.”

Claim 12.5 $\mathbb{E}_{u,v}[\mathcal{I}(u, v)] \geq (1 - \varepsilon_7)\rho^2/n$ and $\mathbb{E}_{u,v}[\mathcal{I}^2(u, v)] \leq (1 + 2\varepsilon_7)d^4(\mathbb{E}_{u,v}[\mathcal{I}(u, v)])^2$, where ε_7 is some constant $\varepsilon_7 \geq 6\varepsilon_5 + \Theta(\varepsilon_5^2)$.

Proof For any $i, j \in [n]$, we say that $i \in \mathcal{N}^{2CSP}(j)$ if there is a constraint on (x_i, x_j) . For the proof of this claim, we also abuse notation and denote $i \in v$ when i is typical, v_{i-1} is a typical prefix, and $X_i = 1$ for v . We also say that $i \in \mathcal{N}(u)$ if i is a typical variable, $i \in \mathcal{N}^{2CSP}(j)$, and $j \in u$ (for some $j \in [n]$). (Do not confuse this notation with prefix neighborhood in the prefix graph.) We can now lower bound the expectation of $\mathcal{I}(u, v)$ as:

$$\mathbb{E}_{u,v}[\mathcal{I}(u, v)] \geq \mathbb{E}_u \left[\sum_{i \in \mathcal{N}(u)} \Pr_v[i \in v] \right].$$

Notice that this bound may not be tight since any $i \in v$ can potentially have d neighbors in u . Thus our upper bound is:

$$\mathbb{E}_{u,v}[\mathcal{I}(u, v)] \leq d \cdot \mathbb{E}_u \left[\sum_{i \in \mathcal{N}(u)} \Pr_v[i \in v] \right].$$

By definition of typical variables, for each typical $i, i \in v$ with probability at least $(1 - \varepsilon_5)^2\rho/n$; thus,

$$\mathbb{E}_{u,v}[\mathcal{I}(u, v)] \geq \mathbb{E}_u \left[\sum_{i \in \mathcal{N}(u)} (1 - \varepsilon_5)^2\rho/n \right] = (1 - \varepsilon_5)^2\rho/n \cdot \mathbb{E}_u[|\mathcal{N}(u)|]. \quad (12.8)$$

All but $\varepsilon_5 n$ variables are typical, so all but $2\varepsilon_5 n$ variables are typical and have at least one typical neighbor. We restrict our attention to the set of such variables and fix

one typical neighbor for each; this neighbor appears in u with probability at least $(1 - \varepsilon_5)^2 \rho/n$. Therefore,

$$\mathbb{E}_u[|\mathcal{N}(u)|] \geq (1 - 2\varepsilon_5)n \cdot ((1 - \varepsilon_5)^2 \rho/n) \geq (1 - 4\varepsilon_5)\rho. \quad (12.9)$$

Combining (12.8) and (12.9), we get the desired bound:

$$\mathbb{E}_{u,v}[\mathcal{I}(u, v)] \geq ((1 - \varepsilon_5)^2 \rho/n)(1 - 4\varepsilon_5)\rho \geq (1 - \varepsilon_7)\rho^2/n. \quad (12.10)$$

Similarly, for the variance we have

$$\begin{aligned} \mathbb{E}_{u,v}[\mathcal{I}^2(u, v)] &\leq d^2 \cdot \mathbb{E}_{u,v} \left(\sum_{i \in v \cap \mathcal{N}(\text{left}(u))} 1 \right)^2 \\ &= d^2 \cdot \mathbb{E}_{u,v} \left[\sum_{i \neq j \in v \cap \mathcal{N}(u)} 1 + \sum_{i \in v \cap \mathcal{N}(u)} 1 \right] \\ &\leq d^2 \cdot \mathbb{E}_u \left[2 \sum_{i < j \in \mathcal{N}(u)} \Pr_v[i \in v] \Pr_v[j \in v \mid i \in v] \right] + d^2 \cdot \mathbb{E}_{u,v}[\mathcal{I}(u, v)]. \end{aligned}$$

Since for every prefix, each variable receives a typical assignment with probability at most $(1 + \varepsilon_5) \cdot \rho/n$, we have that

$$\begin{aligned} \mathbb{E}_{u,v}[\mathcal{I}^2(u, v)] &\leq 2d^2 \cdot \mathbb{E}_u \left[\sum_{i < j \in \mathcal{N}(u)} ((1 + \varepsilon_5) \cdot \rho/n)^2 \right] + d^2 \cdot \mathbb{E}_{u,v}[\mathcal{I}(u, v)] \\ &\leq ((1 + \varepsilon_5) \cdot \rho/n)^2 \cdot 2d^2 \cdot \mathbb{E}_u \binom{|\mathcal{N}(u)|}{2} + d^2 \cdot \mathbb{E}_{u,v}[\mathcal{I}(u, v)]. \quad (12.11) \end{aligned}$$

We would like to bound $\mathbb{E}_u \binom{|\mathcal{N}(u)|}{2}$.

$$\begin{aligned} \mathbb{E}_u \binom{|\mathcal{N}(u)|}{2} &= \sum_{i < j} \sum_{k \in \mathcal{N}^{2CSP}(i)} \sum_{l \in \mathcal{N}^{2CSP}(j)} \Pr_u[k \in u] \Pr_u[l \in u \mid k \in u] \\ &= \sum_{i < j} \sum_{\substack{k \in \mathcal{N}^{2CSP}(i) \\ l \in \mathcal{N}^{2CSP}(j) \\ \text{and } k < l}} \Pr_u[k \in u] \Pr_u[l \in u \mid k \in u] \quad (12.12) \end{aligned}$$

$$+ \sum_{i < j} \sum_{\substack{k \in \mathcal{N}^{2CSP}(i) \\ l \in \mathcal{N}^{2CSP}(j) \\ \text{and } k > l}} \Pr_u[l \in u] \Pr_u[k \in u \mid l \in u] \quad (12.13)$$

$$+ \sum_{i < j} \sum_{k \in \mathcal{N}^{2CSP}(i) \cap \mathcal{N}^{2CSP}(j)} \Pr_u[k \in u]. \quad (12.14)$$

For the first two summands, we can use the condition on the prefixes to conclude that

$$(12.12) + (12.13) \leq \binom{n}{2} d^2 ((1 + \varepsilon_5) \cdot \rho/n)^2,$$

whereas to bound the third summand, we first change the order of summation:

$$(12.14) = \sum_k \Pr_u[k \in u] \cdot \left| \left\{ (i, j) : i \neq j \text{ and } k \in \mathcal{N}^{2CSP}(i) \cap \mathcal{N}^{2CSP}(j) \right\} \right| \\ \leq ((1 + \varepsilon_5) \cdot \rho) \binom{d}{2} = O(\rho).$$

Summing the last two inequalities, we have

$$2 \cdot \mathbb{E}_u \binom{|\mathcal{N}(u)|}{2} \leq d^2 ((1 + \varepsilon_5) \cdot \rho)^2 + O(\rho) \leq (1 + \varepsilon_5)^3 d^2 \rho^2.$$

Plugging back into (12.11):

$$\mathbb{E}_{u,v}[\mathcal{I}^2(u, v)] \leq (1 + \varepsilon_5)^5 d^4 \rho^4 / n^2 + d^2 \cdot \mathbb{E}_{u,v}[\mathcal{I}(u, v)].$$

Using (12.10) and the fact that $\rho = \sqrt{n} \log \log n \gg \sqrt{n}$, this gives

$$\mathbb{E}_{u,v}[\mathcal{I}^2(u, v)] \leq \frac{d^4 (1 + \varepsilon_5)^5}{1 - \varepsilon_7} (\mathbb{E}_{u,v}[\mathcal{I}(u, v)])^2 + d^2 \cdot \mathbb{E}_{u,v}[\mathcal{I}(u, v)] \\ \leq (1 + 2\varepsilon_7) d^4 (\mathbb{E}_{u,v}[\mathcal{I}(u, v)])^2. \quad \blacksquare$$

It will also be convenient to count the number of tests between a pair of variables.

Definition 12.7 For any pair of typical $(i, j) \in \psi$, let $\mathcal{I}^\top(i, j)$ be defined as the number of $(u, v) \in (S \times S)$ pairs such that

- $X_i = 1$ for u and $X_j = 1$ for v .
- u_{i-1} and v_{j-1} are typical prefixes, where u_{i-1} denotes the prefix represented by u for $X_{<i}, Y_{<i}$; similarly for v_{j-1} .

We now have two ways to count the total number of tests between typical prefixes to typical variables:

Observation 12.3 $\sum_{(u,v) \in (S \times S)} \mathcal{I}(u, v) = \sum_{(i,j) \in \psi} \mathcal{I}^\top(i, j).$

Furthermore, since i and j are typical, the number of tests between also behaves “nicely”:

Observation 12.4 For every typical $(i, j) \in \psi$, we have $\mathcal{I}^\top(i, j) \in |S|^2 \rho^2 / n^2 [(1 - \varepsilon_5)^4, (1 + \varepsilon_5)^2].$

Proof

$$\begin{aligned} \mathcal{I}^\top(i, j) &= \sum_{\text{typical } u_{i-1}\text{'s}} |S| \cdot \Pr[u_{i-1}] \Pr[X_i = 1 \mid u_{i-1}] \sum_{\text{typical } v_{j-1}\text{'s}} |S| \cdot \Pr[v_{j-1}] \Pr[X_j = 1 \mid v_{j-1}] \\ &\in |S|^2 \rho^2 / n^2 [(1 - \varepsilon_5)^4, (1 + \varepsilon_5)^2]. \quad \blacksquare \end{aligned}$$

Armed with these Claims 12.4 and 12.5 and Observations 12.3 and 12.4, we are now ready to prove the main lemma of this section. Recall that the soundness of the 2CSP we started with is $1 - \eta$ for a small constant η .

Lemma 12.2 If $\sum_i H(X_i \mid X_{<i}, Y_{<i}) \geq \left(1 - \frac{6\varepsilon_1}{\log n}\right) \log k$, then $\delta(S) < 1 - \delta$, where $\delta < \frac{\varepsilon_6^2}{d^4(1+2\varepsilon_7)}$ and $\varepsilon_6 = (\eta/2 - \varepsilon_5)(1/|A|^2) \frac{(1-\varepsilon_5)^4}{(1+\varepsilon_5)^2}$.

Proof Let the *mode assignment* be the assignment $\mathcal{A}: [n] \rightarrow A$, which assigns to each variable x_i its most common typical assignment (i.e., assignment after a typical prefix), breaking ties arbitrarily. In particular, at least $1/|A|$ of the typical assignments for x_i are equal to $\mathcal{A}(i)$. Of course, this assignment cannot satisfy more than a $(1 - \eta)$ -fraction of the constraints in the original 2CSP; after removing the $\varepsilon_5 n$ atypical variables, $(\eta/2 - \varepsilon_5)dn$ constraints out of the $dn/2$ constraints must still be unsatisfied.

Recall that the number of tests for each constraint over typical variables, $\mathcal{I}^\top(i, j)$, is approximately the same for every pair of (i, j) —up to a $\frac{(1-\varepsilon_5)^4}{(1+\varepsilon_5)^2}$ -multiplicative factor (Observation 12.4). Therefore, the total fraction of tests over unsatisfied constraints, out of all tests, is approximately proportional to the fraction of unsatisfied constraints:

$$\begin{aligned} \sum_{\substack{\text{typical,} \\ \text{unsatisfied } (i, j)\text{'s}}} \mathcal{I}^\top(i, j) &\geq \frac{(1 - \varepsilon_5)^4}{(1 + \varepsilon_5)^2} \cdot \frac{|\{\text{typical, unsatisfied } (i, j)\text{'s}\}|}{|\{\text{typical } (i, j) \in \psi\}|} \cdot \sum_{(i, j) \in \psi} \mathcal{I}^\top(i, j) \\ &\geq \frac{(1 - \varepsilon_5)^4}{(1 + \varepsilon_5)^2} \cdot \frac{(\eta/2 - \varepsilon_5)dn}{dn/2} \cdot \sum_{(i, j) \in \psi} \mathcal{I}^\top(i, j) \\ &= \frac{(1 - \varepsilon_5)^4}{(1 + \varepsilon_5)^2} \cdot (\eta - 2\varepsilon_5) \cdot \sum_{(u, v) \in (S \times S)} \mathcal{I}(u, v). \quad (\text{Observation 12.3}) \end{aligned}$$

For each such pair (i, j) , on at least a $1/|A|^2$ -fraction of the tests both variables receive the mode assignment, so the constraint is violated.² Thus the total number of violations is at least $\varepsilon_6 \sum_{(u,v) \in (S \times S)} \mathcal{I}(u, v)$ (where $\varepsilon_6 = (\eta/2 - \varepsilon_5)(1/|A|^2) \frac{(1-\varepsilon_5)^4}{(1+\varepsilon_5)^2}$).

Finally, we show that so many violations cannot concentrate on less than a δ -fraction of the pairs $u, v \in S$; otherwise:

$$\begin{aligned} \sum_{(u,v) \in (S \times S) \setminus E} \mathcal{I}^2(u, v) &\geq \frac{1}{\delta |S|^2} \left(\sum_{(u,v) \in (S \times S) \setminus E} \mathcal{I}(u, v) \right)^2 && \text{(Cauchy-Schwartz)} \\ &\geq \frac{1}{\delta |S|^2} \left(\varepsilon_6 \sum_{(u,v) \in (S \times S)} \mathcal{I}(u, v) \right)^2 \\ &= \frac{|S|^2 \varepsilon_6^2}{\delta} (\mathbb{E}_{u,v}[\mathcal{I}(u, v)])^2; \end{aligned}$$

yet by Claim 12.5,

$$\sum_{(u,v) \in (S \times S) \setminus E} \mathcal{I}^2(u, v) \leq \sum_{(u,v) \in S \times S} \mathcal{I}^2(u, v) \leq (1 + 2\varepsilon_7) d^4 |S|^2 (\mathbb{E}_{u,v}[\mathcal{I}(u, v)])^2.$$

Thus we have a contradiction since $d^4(1 + 2\varepsilon_7) < \varepsilon_6^2/\delta$ by our setting of δ . Therefore we have **2CSP**-violations in more than a δ -fraction of the pairs $u, v \in S$. ■

With Lemma 12.1 and Lemma 12.2, we can now complete the proof of Theorem 12.2.

Proof of Theorem 12.2 Recall that $\sum_i \alpha_i + \beta_i = \log k' \geq \left(1 - \frac{\varepsilon_1}{\log n}\right) \log k$ by Fact 12.4. If $\sum_i \beta_i > \left(\frac{5\varepsilon_1}{\log n}\right) \log k$, then by Lemma 12.1, $\delta(S) < 1 - \delta$. Otherwise, if $\sum_i \alpha_i > \left(1 - \frac{6\varepsilon_1}{\log n}\right) \log k$, by Lemma 12.2, $\delta(S) < 1 - \delta$. ■

2. We remark that a more careful analysis of the expected number of violations would allow one to save an $|A|^2$ -factor in the value of ε_6 . Since it does not qualitatively affect the result, we opt for the simpler analysis.

Community Detection

Identifying communities is a central graph-theoretic problem with important applications to sociology and marketing (when applied to social networks), biology and bioinformatics (when applied to protein interaction networks), and more (see, e.g., Fortunato's classic survey [Fortunato 2010]). Defining what exactly is a *community* remains an interesting problem on its own (see Arora et al. [2012] and Borgs et al. [2016] for excellent treatment from a theoretical perspective). Ultimately, there is no single “right” definition, and the precise meaning of community should be different for social networks and protein interaction networks.

Here we focus on the algorithmic questions arising from one of the simplest and most canonical definitions, which has been considered by several theoretical computer scientists [Arora et al. 2012, Balcan et al. 2013, Mishra et al. 2008] (see Section 13.1 for further discussion):

Definition 13.1 ((α, β) -community). Given an undirected graph $G = (V, E)$, an (α, β) -community is a subset $S \subseteq V$ that satisfies:

Strong ties inside the community. For every $v \in S$, $|\{v\} \times S \cap E| \geq \alpha \cdot |S|$; and

Weak ties to nodes outside the community. For every $u \notin S$, $|\{u\} \times S \cap E| \leq \beta \cdot |S|$.

Arora et al. [2012, Theorem 3.1] gave a simple quasi-polynomial ($n^{O(\log n)}$) time for detecting (α, β) -communities whenever $\alpha - \beta$ is at least some positive constant. Similar to the meta-algorithm we described in the introduction, Arora et al.'s algorithm enumerates over $O(\log n)$ -tuples of vertices. For each tuple, consider the set of vertices that are neighbors of an $(\alpha + \beta)/2$ -fraction of the tuple; test whether this candidate set is indeed a community.

Here we show that, for every constant $\alpha > \beta \in (0, 1]$, community detection requires quasi-polynomial time (assuming ETH). For example, when $\alpha = 1$ and $\beta = 0.01$, this means that we can hide a clique C , such that every single vertex not in C is

connected to at most 1% of C . Our main result is actually a much stronger inapproximability: even in the presence of a $(1, o(1))$ -community, finding any $(\beta + o(1), \beta)$ -community is hard.

Theorem 13.1 For every n there exists an $\epsilon = \epsilon(n) = o(1)$ such that, assuming ETH, distinguishing between the following requires time $n^{\tilde{\Omega}(\log n)}$:

Completeness. G contains a $(1, \epsilon)$ -community; and

Soundness. G does not contain a $(\beta + \epsilon, \beta)$ -community for any $\beta \in [0, 1]$.

Unlike many related quasi-polynomial approximation schemes mentioned above, Arora et al.’s algorithm has the unique property that it can also *exactly count* all the (α, β) -communities. Our second result is that counting even the number of $(1, o(1))$ -communities requires quasi-polynomial time. A nice feature of this result is that we can base it on the much weaker #ETH assumption, which asserts that counting the satisfying assignment for a 3SAT instance requires time $2^{\Omega(n)}$. (Note, for example, that #ETH is likely to be true even if $P = NP$.)

Theorem 13.2 For every n there exists an $\epsilon = \epsilon(n) = o(1)$ such that, assuming #ETH, counting $(1, \epsilon)$ -communities requires time $n^{\log^{1-o(1)} n}$.

13.1 Related Works

The most closely related work is a reduction by Balcan et al. [2013, Theorem 5.3] from PLANTED CLIQUE to finding $(1, 1 - \gamma)$ -communities, for some small (unspecified) constant $\gamma > 0$. Note that our inapproximability in Theorem 13.1 is much stronger in all parameters; furthermore, although formally incomparable, our ETH assumption is preferable over the average-case hardness assumption of PLANTED CLIQUE.

Algorithms for Special Cases

Mishra et al. [2008] gave a polynomial-time algorithm for finding (α, β) -communities that contain a vertex with very few neighbors outside the community. Balcan et al. [2013] give a polynomial-time algorithm for enumerating (α, β) -communities in the special case where the degree of every node is $\Omega(n)$.

Arora et al. [2012] consider several semi-random models where the edges inside the community are generated at random, according to the expected degree model. (In fact, their quasi-polynomial time algorithm is also stated in this setting, but only their “Gap Assumption,” which is equivalent to $\alpha - \beta = \Omega(1)$, is used in the analysis.)

Stochastic Block Model

Variants of the community detection problem on graphs generated by different stochastic models are extremely popular (see, e.g., [Bandeira et al. 2016, Banks et al. 2016, Chen et al. 2016, Florescu and Perkins 2016, Hajek et al. 2016, Makarychev et al. 2016, Moitra et al. 2016, Mossel and Xu 2016, Tremblay et al. 2016] for papers in conference proceedings from June 2016). Perhaps the most influential is the *Stochastic Block Model* [Holland et al. 1983]: The graph is partitioned into two disjoint communities; the edges within each community are present with probability α , independently, whereas edges between communities are present with probability β . Hence this model can also be seen as a special case of the (α, β) -Community Detection problem.

Stochastic models are extremely helpful in physics, for example, because atoms' interactions obey simple mathematical formulas with high precision. Unfortunately, for applications such as social networks, existing models do not describe human behavior with atomic precision, hence casting a shadow over the applicability of algorithms that work on ideal stochastic models. Recent works [Makarychev et al. 2016, Moitra et al. 2016] attempt to bridge the gap from ideal model to practice by showing that certain SDP-based algorithms continue to work in a particular semi-random model where a restricted adversary is allowed to modify the random input graph. These success stories beg the question of how strong one can make the adversary. This chapter illuminates some of the computational barriers.

Alternative Approaches to Modeling Communities

As we mentioned above, there are many different definitions of “communities” in networks. For in-depth discussion of different definitions, see Arora et al. [2012] or Borgs et al. [2016]. As pointed out by the latter, for some definitions even verifying that a candidate subset is a community is intractable.

There is also an important literature on axiomatic approaches to the related problem of clustering (e.g., [Ben-David and Ackerman 2008, Kleinberg 2002, van Laarhoven and Marchiori 2014]); note that while clustering typically aims to partition a set of nodes, our main focus is on detecting just a single community; in particular, different communities may intersect.

13.2 Overview of Proofs

A good starting point for the technical discussion is the reduction from 3SAT to the related problem of DENSEST- k -SUBGRAPH we describe in Chapter 12. Let us recall

the main two ingredients in that reduction: “birthday repetition” [Aaronson et al. 2014] and the “FGLSS graph” [Feige et al. 1996].

“Birthday repetition”. Starting with an instance of LABEL COVER (see definition in Section 2.4), the reduction considers a mega-variable for every ρ -tuple of variables, for $\rho \approx \sqrt{n}$. By the birthday paradox, almost every pair of ρ -tuples of variables intersect, inducing a consistency constraint on the two mega-assignments. Similarly, we expect to see some LABEL COVER edges in the union of the two ρ -tuples, inducing an additional LABEL COVER constraint between the two mega assignments. Notice that we have $\binom{n}{\rho} \approx 2^{\sqrt{n}}$ mega variables, and the alphabet size is also approximately $N = 2^{\sqrt{n}}$. Therefore, assuming ETH, finding an approximately satisfying assignment for the mega-variables requires time $2^{\Omega(n)} \approx N^{\log N}$.

FGLSS. Similarly to the classic reduction by Feige et al. [1996] for the CLIQUE problem in Chapter 12, we construct a vertex for each mega assignment to each mega variable, and draw an edge between two vertices if the induced assignments do not violate any consistency or LABEL COVER constraints. Notice that if the LABEL COVER instance has a satisfying assignment, then the graph contains a clique of size $\binom{n}{\rho}$ where each mega variable receives the mega assignment induced by the globally satisfying assignment. On the other hand, any subgraph that corresponds to a consistent assignment that violates many constraints must be missing most of its edges.

Unfortunately, this simple reduction is still quite far from working for the COMMUNITY DETECTION problem. Below we describe some of the obstacles and outline how we overcome them.

Completeness

Surprisingly, the main problem with using the same reduction for COMMUNITY DETECTION is the completeness: even if the LABEL COVER instance has a satisfying assignment, the resulting graph has no (α, β) -communities, for any constants $\alpha > \beta$! Observe, in particular, that the clique that corresponds to the satisfying assignment does not satisfy the weak ties condition. For any vertex v in that clique, consider any vertex v' that corresponds to changing the assignment to just one variable x_i in v 's assignment. If v agrees with the assignments of all other vertices in the clique, v' agrees with almost all of them—except for the negligible fraction that cover x_i or its neighbors in the LABEL COVER graph.

To overcome this problem of vertices that are “just outside the community,” we use error-correcting codes. Namely, we encode each assignment as a low-degree bivariate polynomial over finite field \mathcal{G} of size $|\mathcal{G}| \approx \sqrt{n}$. Now vertices correspond to low-degree assignments to rows/columns of the polynomial. This guarantees that the assignments induced by every two vertices are far. If v agrees with all other vertices in the community, then almost all of those vertices disagree with v' .

Soundness

The main challenge for soundness is ruling out communities that do not correspond to a single, globally consistent assignment to the LABEL COVER instance. The key idea is to introduce auxiliary vertices that punish such communities by violating the weak ties desideratum.

Let us begin with the reduction to the counting variant (Theorem 13.2), which is easier, mostly because we are not concerned with approximation (i.e., we only have to show that subsets that are exactly $(1, \epsilon)$ -communities correspond to satisfying assignments). Here we further simplify matters by sketching a construction with weighted edges. The full reduction (Section 13.3) uses unweighted edges and is only slightly more involved. Consider, for every $g \in \mathcal{G}$, an auxiliary vertex that is ϵ -connected to all proper vertices that do not correspond to assignments to the g -th row/column. Now if a $(1, \epsilon)$ -community C does not contain a vertex with assignment to the g -th row/column, the auxiliary vertex must simultaneously: (i) belong to C so as not to violate the weak ties desideratum; yet (ii) it cannot belong to C because all its edges have weight ϵ (this would violate the strong ties desideratum). Therefore every $(1, \epsilon)$ -community assigns values to every row/column in \mathcal{G}^2 .

The reduction we described above suffices to show that (assuming ETH) deciding whether the graph contains a $(1, \epsilon)$ -community also requires quasi-polynomial time. To get the stronger statement of Theorem 13.1 we must rule out even $(\beta, \beta + \epsilon)$ -communities in case the LABEL COVER instance is far from satisfiable. In particular, we need to show that subsets that do not correspond to unique, consistent assignments are never $(\beta, \beta + \epsilon)$ -communities. Instead of a single column/row, we let each proper vertex correspond to a subset of $t \approx \log n$ columns/rows. Instead of a single $g \in \mathcal{G}$, each auxiliary vertex corresponds to subset $H \subset \mathcal{G}$ of size $|H| = |\mathcal{G}|/2$. We draw an edge between an auxiliary vertex and a proper vertex if the indices of all t columns/rows are contained in H ; if they are picked randomly this only happens with polynomially small probability. If, however, a β -fraction of the community is restricted to a small subset $R \subset \mathcal{G}$, then there are auxiliary vertices for $H \supseteq R$ that connect to all those nodes and violate the weak ties desideratum. Roughly, we show that at least a $(1 - \beta)$ -fraction of the vertices have assignments that are “well

spread" over \mathcal{G}^2 , and among those assignments there are many violations of the LABEL COVER constraints.

13.3 Hardness of Counting Communities

Theorem 13.3 There exists an $\epsilon(n) = o(1)$ such that, assuming #ETH, counting $(1, \epsilon)$ -communities requires time $n^{\log^{1-o(1)} n}$.

Construction

Begin with an instance (A, B, E, π) of LABEL COVER of size $n = n_A + n_B$ where $n_A \triangleq |A|$ and $n_B \triangleq |B|$. Let \mathcal{G} be a finite field of size \sqrt{n}/ϵ^3 , and let $\mathcal{F} \subset \mathcal{G}$ be an arbitrary subset of size $|\mathcal{F}| = \sqrt{n}$. We identify between $A \cup B$ and points in \mathcal{F}^2 ; we also identify between a subset of \mathcal{G} and $\Sigma_A \cup \Sigma_B$. Thus there is a one-to-one correspondence between a subset of assignments to $P_{\mathcal{F}}: \mathcal{F}^2 \rightarrow \mathcal{G}$ and assignments to the LABEL COVER instance. We can extend any such $P_{\mathcal{F}}$ to an individual-degree- $(|\mathcal{F}| - 1)$ polynomial $P: \mathcal{G}^2 \rightarrow \mathcal{G}$. In the other direction, we think of each low-individual-degree polynomial $P: \mathcal{G}^2 \rightarrow \mathcal{G}$ as a (possibly invalid) assignment to the LABEL COVER instance.

For every $g \in \mathcal{G}$, and degree- $(|\mathcal{F}| - 1)$ polynomials $p_1, p_2: \mathcal{G} \rightarrow \mathcal{G}$ such that $p_1(g) = p_2(g)$, we construct $1/\epsilon$ vertices $\{v_{g,p_1,p_2,i}\}_{i=1}^{1/\epsilon} \subset V$ in the communities graph. Each vertex naturally induces an assignment (p_1, p_2) on $(\mathcal{G} \times \{g\}) \cup (\{g\} \times \mathcal{G})$. We draw an edge between two vertices in V if they agree on the intersection of their lines, and if their induced assignments satisfy all the LABEL COVER constraints.

For every $g \in \mathcal{G}$ and $i \in [1/\epsilon]$, we also add two identical auxiliary vertices $u_{g,i}$, which are connected to every $v_{g',p_1,p_2,i}$ for $g' \neq g$ (but not to each other).

Completeness

For each assignment to the LABEL COVER instance, we construct a $(1, \epsilon)$ -community by taking the induced assignment $P_{\mathcal{F}}: \mathcal{F}^2 \rightarrow \mathcal{G}$ and extending it to an individual-degree- $(|\mathcal{F}| - 1)$ polynomial $P: \mathcal{G}^2 \rightarrow \mathcal{G}$. Let C be all the vertices $v_{g,p_1,p_2,i}$ such that p_1, p_2 are the restrictions of P to $(\mathcal{G} \times \{g\}), (\{g\} \times \mathcal{G})$. This correspondence is one-to-one and we need to show that the resulting C is actually a $(1, \epsilon)$ -community.

Because all the vertices correspond to a consistent satisfying assignment, C is a clique. Let $v_{g,q_1,q_2,i} \notin C$; $w \log q_1$ disagrees with the restriction of P to $(\mathcal{G} \times \{g\})$. Since both q_1 and the restriction of P are degree- $(|\mathcal{F}| - 1)$ polynomials, they must disagree on all but at most $(|\mathcal{F}| - 1)$ elements of \mathcal{G} . For all other $h \in \mathcal{G}$, the vertex $v_{g,q_1,q_2,i}$ does not share edges with any $v_{h,p_1,p_2,j} \in C$. Therefore, $v_{g,q_1,q_2,i}$ has edges to less than an $(|\mathcal{F}|/|\mathcal{G}|)$ -fraction of vertices in C . Finally, every auxiliary vertex $u_{g,i}$

has edges to a $\frac{|\mathcal{G}|-1}{|\mathcal{G}|} \cdot \epsilon < \epsilon$ -fraction of the vertices in ϵ . Therefore, C is a $(1, \epsilon)$ -community.

13.3.1 Soundness

Structure of $(1, \epsilon)$ -Communities

Claim 13.1 Every $(1, \epsilon)$ -community C contains exactly $1/\epsilon$ vertices $\{v_{g,p_1,p_2,i}\}_{i=1}^{1/\epsilon}$ for each g .

Proof First, observe that C cannot contain any auxiliary vertices: if C contains one copy of $u_{g,i}$, it must also contain the other; but they don't have an edge between them, so they cannot both belong to a $(1, \epsilon)$ -community.

Now, assume by contradiction that for some $g \in \mathcal{G}$, C does not contain any vertices with assignments for $(\mathcal{G} \times \{g\}) \cup (\{g\} \times \mathcal{G})$. Then every vertex in C is connected to (both copies of) $u_{g,i}$, for some $i \in [1/\epsilon]$. Therefore there is at least one $i \in [1/\epsilon]$ such that $u_{g,i}$ is connected to an ϵ -fraction of the vertices in C . But this is a contradiction since $u_{g,i} \notin C$.

If we ignore the auxiliary vertices (which, as we argued, C does not contain), the different vertices $v_{g,p_1,p_2,i}$ that correspond to the same assignment to the same lines (i.e., if we only change i) are indistinguishable. Therefore if C contains one of them, it must contain all of them (hence, at least $1/\epsilon$ vertices for each g).

Finally, since C is a clique, it cannot contain vertices that disagree on any assignments. (In particular, it cannot contain more than $1/\epsilon$ vertices for each g .) ■

Completing the Proof

Proof Proof of soundness. By Claim 13.1, every $(1, \epsilon)$ -community C contains exactly $1/\epsilon$ vertices $\{v_{g,p_1,p_2,i}\}_{i=1}^{1/\epsilon}$ for each g . Furthermore, since C is a clique, all the induced assignments agree on all the intersections. So every $(1, \epsilon)$ -community corresponds to a unique consistent assignment to the LABEL COVER instance. Finally, appealing again to the fact that C is a clique, this assignment must also satisfy all the LABEL COVER constraints. ■

13.4 Hardness of Detecting Communities

Theorem 13.4 There exists an $\epsilon(n) = o(1)$ such that, assuming ETH, distinguishing between the following requires time $n^{\Omega(\log n)}$:

Completeness. G contains a $(1, \epsilon)$ -community; and

Soundness. G does not contain a $(\beta + \epsilon, \beta)$ -community for any $\beta \in [0, 1]$.

The rest of this section is devoted to the proof of Theorem 13.4. Our starting point is the LABEL COVER of Moshkovitz-Raz (Theorem 2.2). We compose the birthday repetition technique of Aaronson et al. [2014] with a bi-variate low-degree encoding. We then encode this as a graph à la FGLSS [Feige et al. 1996]. We add auxiliary vertices to ensure that any $(\beta + \epsilon, \beta)$ -community corresponds, approximately, to a uniform distribution over the variables.

Construction

Begin with a (d_A, d_B) -bi-regular instance (A, B, E, π) of LABEL COVER of size $n = n_A + n_B$ where $n_A \triangleq |A|$ and $n_B \triangleq |B|$. Let $\rho \triangleq \sqrt{n} \log n$; let \mathcal{G} be a finite field of size $\rho/\epsilon^3 = \tilde{O}(\rho)$, and let $\mathcal{F} \subset \mathcal{G}$ be an arbitrary subset of size $|\mathcal{F}| = 2\rho$. Let $\mathcal{F}_A, \mathcal{F}_B \subset \mathcal{F}$ be disjoint subsets of size $n_A/\rho, n_B/\rho$, respectively. By Lemma 2.6, we can partition A and B into subsets $X_1, \dots, X_{|\mathcal{F}_A|}$ and $Y_1, \dots, Y_{|\mathcal{F}_B|}$ of size at most $|\mathcal{F}|$ such that between every two subsets there are approximately $\frac{d_A \rho^2}{n_B} = \frac{d_B \rho^2}{n_A}$ constraints. For $i \in \mathcal{F}_A$, we think of the points $\{i\} \times \mathcal{F} \subset \mathcal{G}^2$ as representing assignments to variables in X_i ; for $j \in \mathcal{F}_B$, we think of $\mathcal{F} \times \{j\} \subset \mathcal{G}^2$ as representing assignments to variables in Y_j . Notice that each point in \mathcal{F}^2 may represent an assignment to both a vertex from A and a vertex from B , to one of them, or to neither. In particular, any assignment $P: \mathcal{G}^2 \rightarrow \mathcal{G}$ induces an assignment for the LABEL COVER instance; note that since $|\mathcal{G}| > |\Sigma_A| |\Sigma_B|$, one value $P(f_1, f_2) \in \mathcal{G}$ suffices to describe assignments to both $a \in A$ and $b \in B$.

Let $t \triangleq \log n \cdot \left(\frac{|\mathcal{G}|}{|\mathcal{F}_A|} + \frac{|\mathcal{G}|}{|\mathcal{F}_B|} \right) = \text{polylog}(n)$. We say that a subset $S \in \binom{\mathcal{G}}{t}$ is *balanced* if: $|S \cap \mathcal{F}_A| = \frac{|\mathcal{F}_A|}{|\mathcal{G}|} \cdot t$ and $|S \cap \mathcal{F}_B| = \frac{|\mathcal{F}_B|}{|\mathcal{G}|} \cdot t$. For every balanced subset S , consider $2t$ polynomials $q_\ell: \mathcal{G} \rightarrow \mathcal{G}$ of degree at most $|\mathcal{F}| - 1$, representing an assignment¹ $Q: (S \times \mathcal{G}) \cup (\mathcal{G} \times S) \rightarrow \mathcal{G}$. For balanced S and $2t$ -tuple of polynomials (q_ℓ) , we construct a corresponding vertex $v_{S, (q_\ell)}$ in the communities graph. Let V denote the set of vertices defined so far. For $g \in \mathcal{G}$ we abuse notation and say that $g \in v_{S, (q_\ell)}$ if $g \in S$. We construct an edge in the communities graph between two vertices in V if their assignments agree on the variables in their intersection, and their induced assignments to $A \cup B$ satisfy all the LABEL COVER constraints.

Additionally, for every $H \subset \mathcal{G}$ of size $|H| = |\mathcal{G}|/2$, define $|V|^2$ identical auxiliary vertices u_H in the communities graph. We draw an edge between auxiliary vertex u_H and vertex $v_{S, (q_\ell)}$ if $S \subset H$. Similarly, for every $H_A \subset \mathcal{F}_A$ of size $|H_A| = |\mathcal{F}_A|/2$, we define $|V|^2$ identical auxiliary vertices u_{H_A} with edges to every vertex $v_{S, (q_\ell)}$ such

1. We will only consider polynomials that correspond to a consistent assignment Q ; i.e., for each point in $S \times S$ we expect the two corresponding polynomials to agree with each other.

that $(S \cap \mathcal{F}_A) \subset H_A$. For $H_B \subset \mathcal{F}_B$ of size $|H_B| = |\mathcal{F}_B|/2$, we draw edges between u_{H_B} and $v_{S, (q_\ell)}$ such that $(S \cap \mathcal{F}_B) \subset H_B$.

Completeness

Suppose that the LABEL COVER instance has a satisfying assignment. Let $\mathcal{Z} \subseteq \mathcal{G}^2$ denote the subset of points that correspond to at least one variable in A or B . Let $P_{\mathcal{Z}} : \mathcal{Z} \rightarrow \mathcal{G}$ be the induced function on \mathcal{Z} that corresponds to the satisfying assignment, and let $P : \mathcal{G}^2 \rightarrow \mathcal{G}$ be the extension of $P_{\mathcal{Z}}$ by setting $P(f_1, f_2) = 0$ for $(f_1, f_2) \in \mathcal{F}^2 \setminus \mathcal{Z}$ (this choice is arbitrary), and then extending to an $(|\mathcal{F}| - 1)$ -individual-degree polynomial over all of \mathcal{G}^2 .

Let C be the set of vertices that correspond to restrictions of P to balanced sets, i.e.,

$$C = \{v_{S, (P|_S)} : S \text{ is balanced}\},$$

where $P|_S$ denotes the restriction of P to $(S \times \mathcal{G}) \cup (\mathcal{G} \times S)$. Since all those vertices correspond to a consistent satisfying assignment, C is a clique.

For any vertex $v_{S, (q_\ell)} \notin C$, at least one of the polynomials, q_{ℓ^*} , disagrees with the restriction of P to the corresponding line. Since both q_{ℓ^*} and the restriction of P to that line are degree- $(|\mathcal{F}| - 1)$ polynomials, they must disagree on at least $\left(1 - \frac{|\mathcal{F}|}{|\mathcal{G}|}\right)$ -fraction of the coordinates. The probability that a random balanced set S' is contained in the $O(\epsilon^3)$ -fraction of coordinates where they do agree is smaller than ϵ (and in fact polynomially small in n). Therefore $v_{S, (q_\ell)}$ has inconsistency violations with all but (less than) an ϵ -fraction of the vertices in C .

For any auxiliary vertex u_{H_A} , the probability that a random vertex $v_{S, (P|_S)} \in C$ is connected to u_{H_A} is $2^{-|S \cap \mathcal{F}_A|} < 1/n$, and similarly for u_{H_B} and u_H . Therefore, every auxiliary vertex is connected to less than a $(1/n)$ -fraction of the vertices in C .

13.4.1 Soundness

Lemma 13.1 If the LABEL COVER instance has value at most ϵ^3 , then there are no $(\beta + \epsilon, \beta)$ -communities.

Auxiliary vertices.

Claim 13.2 Every $(\beta + \epsilon, \beta)$ -community does not contain any auxiliary vertices.

Proof There are $|V|^2$ identical copies of each auxiliary vertex. Since they are identical, any community must either contain all of them, or none of them. If the community contains all $|V|^2$ copies, then it has a vast majority of auxiliary vertices, so none of them can have edges to an ϵ -fraction of the community. ■

List decoding.

Claim 13.3 The vertices in any $(\beta + \epsilon, \beta)$ -community C induce at most $4/\epsilon$ different assignments for each variable.

Proof Suppose by contradiction that this is not the case. Then, wlog, there is a line $\{g_1\} \times \mathcal{G}$ that receives at least $2/\epsilon$ different assignments from vertices in C . Every two assignments agree on at most $|\mathcal{F}|$ points (g_1, g') on the line, so in total there are at most $2|\mathcal{F}|/\epsilon^2$ points where at least two assignments agree. Let $R \subseteq \mathcal{G}$ denote the set of g' such that no two assignments agree on (g_1, g') ; we have that $|R| \geq |\mathcal{G}| - 2|\mathcal{F}|/\epsilon^2 \geq |\mathcal{G}|/2$. Therefore, by the weak ties property, for at most a β -fraction of the vertices $v_{S, (q_\ell)} \in C$, $S \cap R = \emptyset$.

Consider the remaining $(1 - \beta)$ -fraction of vertices in C . Suppose that v assigns a value to some (g_1, g') for $g' \in R$: this value can only agree with one of the $2/\epsilon$ different assignments to (g_1, g') . Therefore, in expectation, each of the $2/\epsilon$ vertices that assign different values for (g_1, g') is connected to at most a $(\beta + \epsilon/2)$ -fraction of the vertices in C . This is a contradiction to C being a $(\beta + \epsilon, \beta)$ -community. ■

Completing the proof.

Proof Proof of Lemma 13.1. Suppose that at most an ϵ^3 -fraction of the LABEL COVER constraints can be satisfied by any single assignment, and assume by contradiction that C is a $(\beta + \epsilon, \beta)$ -community. By Claim 13.3, C induces at most $4/\epsilon$ assignments on each variable, so at most an $O(\epsilon)$ -fraction of the constraints are satisfied by any pair of assignments.

By Markov's inequality, for at least half of the subsets $X_i \subset A$, only an $O(\epsilon)$ -fraction of the constraints that depend on X_i are satisfied. By Claim 13.2 at least $(1 - \beta)$ -fraction of the vertices in C assign values to at least one such X_i . Consider any such vertex $v_{S, (q_\ell)}$ where $S \ni i$. By construction of the partitions (Lemma 2.6), each X_i shares approximately the same number of constraints with each Y_j . Therefore, for all but an $O(\epsilon)$ -fraction of Y_j 's, X_i and Y_j observe a violation—for all the assignments given by vertices in C to the variables in Y_j . In other words, $v_{S, (q_\ell)}$ cannot have edges to any vertex $v_{T, (r_\ell)}$ such that $T \ni j$, for a $(1 - O(\epsilon))$ -fraction of $j \in [n_B/k_B]$. Finally, applying Claim 13.2 again, at most a β fraction of vertices in C do not contain any of those j 's. This is a contradiction to $v_{S, (q_\ell)}$ having edges to a $(\beta + \epsilon)$ -fraction of the vertices in C . ■

VC and Littlestone's Dimensions

A common and essential assumption in learning theory is that the concepts we want to learn come from a nice, simple concept class, or (in the agnostic case) they can at least be approximated by a concept from a simple class. When the concept class is sufficiently simple, there is hope for good (i.e., sample-efficient and low-error) learning algorithms.

There are many different ways to measure the *simplicity* of a concept class. The most influential measure of simplicity is the VC Dimension, which captures learning in the PAC model. We also consider Littlestone's Dimension [Littlestone 1987], which corresponds to minimizing mistakes in online learning (see Section 2.5 for definitions). When either dimension is small, there are algorithms that exploit the simplicity of the class, to obtain good learning guarantees.

Two decades ago, it was shown that (under appropriate computational complexity assumptions) neither dimension can be computed in polynomial time [Frances and Litman 1998, Papadimitriou and Yannakakis 1996]; and these impossibility results hold even in the most optimistic setting where the entire universe and concept class are given as explicit input (a binary matrix whose (x, c) -th entry is 1 iff element x belongs to concept c). The computational intractability of computing the (VC, Littlestone's) dimension of a concept class suggests that even in cases where a simple structure exists, it may be inaccessible to computationally bounded algorithms (see discussion below).

In this work we extend the results of Frances and Litman [1998] and Papadimitriou and Yannakakis [1996] to show that the VC and Littlestone's Dimensions cannot even be *approximately* computed in polynomial time. We don't quite prove that those problems are NP-hard: both dimensions can be computed (exactly) in quasi-polynomial ($n^{O(\log n)}$) time, hence it is very unlikely that either

problem is NP-hard. Nevertheless, assuming the randomized Exponential Time Hypothesis (ETH),¹ we prove essentially tight quasi-polynomial lower bounds on the running time that hold even against approximation algorithms.

Theorem 14.1 Hardness of approximating VC Dimension. Assuming randomized ETH, approximating the VC Dimension to within a $(1/2 + o(1))$ -factor requires $n^{\log^{1-o(1)} n}$ time.

Theorem 14.2 Hardness of approximating Littlestone's Dimension. There exists an absolute constant $\varepsilon > 0$ such that, assuming randomized ETH, approximating Littlestone's Dimension to within a $(1 - \varepsilon)$ -factor requires $n^{\log^{1-o(1)} n}$ time.

14.1 Discussion

As we mentioned before, the computational intractability of computing the (VC, Littlestone's) dimension of a concept class suggests that even in cases where a simple structure exists, it may be inaccessible to computationally bounded algorithms. We note, however, that it is not at all clear that any particular algorithmic applications are immediately intractable as a consequence of our results.

Consider, for example, the adversarial online learning zero-sum game corresponding to Littlestone's Dimension: At each iteration, Nature presents the learner with an element from the universe; the learner attempts to classify the element, and loses a point for every wrong classification; at the end of the iteration, the correct (binary) classification is revealed. Littlestone's Dimension is equal to the worst case loss of the learner before learning the exact concept. (See Section 2.5 for a more detailed definition.)

What can we learn from the fact that Littlestone's Dimension is hard to compute? The first observation is that there is no efficient learner that can commit to a concrete mistake bound. But this does not rule out a computationally efficient learner that plays optimal strategy and makes at most as many mistakes as the unbounded learner. We can, however, conclude that Nature's task is computationally intractable! Otherwise, we could efficiently construct an entire worst-case mistake tree (for a concept class \mathcal{C} , any mistake tree has at most $|\mathcal{C}|$ leaves, requiring $|\mathcal{C}| - 1$ oracle calls to Nature).

On a philosophical level, we think it is interesting to understand the implications of an intractable, adversarial Nature. Perhaps this is another evidence that the mistake bound model is too pessimistic?

1. The randomized ETH (rETH) postulates that there is no $2^{o(n)}$ -time Monte Carlo algorithm that solves 3SAT correctly with probability at least $2/3$ (i.e., $3\text{SAT} \notin \text{BPTIME}(2^{o(n)})$).

Also, the only algorithm we know for computing the optimal learner’s decision requires computing Littlestone’s Dimension. We think that it is an interesting open question whether an approximately optimal computationally efficient learner exist.

Finally, let us note that in the other direction, computing Littlestone’s Dimension exactly implies an exactly optimal learner. However, since the learner has to compute Littlestone’s Dimension many times, we have no evidence that an approximation algorithm for Littlestone’s Dimension would imply any guarantee for the learner.

Finally, we remark that for either problem (VC or Littlestone’s Dimension), we are not aware of any non-trivial approximation algorithms.

14.2 Techniques

Our reductions in this chapter are also inspired by “birthday repetition,” but surprising challenges arise.

VC Dimension. The first challenge we have to overcome in order to adapt this framework to hardness of approximation of the VC Dimension is that the number of concepts involved in shattering a subset S is $2^{|S|}$. Therefore any inapproximability factor we prove on the size of the shattered set of elements “goes in the exponent” of the size of the shattering set of concepts. Even a small constant factor gap in the VC Dimension requires proving a polynomial factor gap in the number of shattering concepts (obtaining polynomial gaps via “birthday repetition” for simpler problems is an interesting open problem [Manurangsi 2017, Manurangsi and Raghavendra 2016]). Fortunately, having a large number of concepts is also an advantage: we use each concept to test a different set of 3-COLOR constraints chosen independently at random; if the original instance is far from satisfied, the probability of passing all $2^{\Theta(|S|)}$ tests should now be doubly exponentially small $(2^{-2^{\Theta(|S|)}})!$ More concretely, we think of half of the elements in the shattered set as encoding an assignment, and the other half as encoding which tests to run on the assignments.

Littlestone’s Dimension. Our starting point is the reduction for the VC Dimension outlined in the previous paragraph. Without going into the exact definition of Littlestone’s Dimension, recall that it corresponds to an online learning model. If the test-selection elements arrive before the assignment-encoding elements, the adversary can adaptively tailor his assignment to pass the specific test selected in the previous steps. To overcome this obstacle, we introduce a special gadget that

forces the assignment-encoding elements to arrive first; this makes the reduction to Littlestone's Dimension somewhat more involved. Note that there is a reduction by [Frances and Litman \[1998\]](#) from the VC Dimension to Littlestone's Dimension. Unfortunately, their reduction is not (approximately) gap-preserving, so we cannot use it directly to obtain [Theorem 14.2](#) from [Theorem 14.1](#).

14.3 Related Work

The study of the computational complexity of the VC Dimension was initiated by [Linial et al. \[1991\]](#), who observed that it can be computed in quasi-polynomial time. [Papadimitriou and Yannakakis \[1996\]](#) proved that it is complete for the class LOGNP, which they define in the same paper. [Frances and Litman \[1998\]](#) reduced the problem of computing the VC Dimension to that of computing Littlestone's Dimension, hence the latter is also LOGNP-hard. (It follows as a corollary of our [Theorem 14.1](#) that, assuming ETH, solving any LOGNP-hard problem requires quasi-polynomial time.)

Both problems were also studied in an implicit model, where the concept class is given in the form of a Boolean circuit that takes as input an element x and a concept c and returns 1 iff $x \in c$. Observe that in this model even computing whether either dimension is 0 or not is already NP-hard. Schaefer proved that the VC Dimension is Σ_3^P -complete [[Schaefer 1999](#)], while Littlestone's Dimension is PSPACE-complete [[Schaefer 2000](#)]. [Mossel and Umans \[2002\]](#) proved that the VC Dimension is Σ_3^P -hard to approximate to within a factor of almost 2; can be approximated to within a factor slightly better than 2 in AM; and is AM-hard to approximate to within $n^{1-\varepsilon}$.

Another line of related work in the implicit model proves computational intractability of PAC learning (which corresponds to the VC Dimension). Such intractability has been proved either from cryptographic assumptions (e.g., [[Feldman et al. 2006](#), [Kalai et al. 2008](#), [Kearns and Valiant 1994](#), [Kharitonov 1993, 1995](#), [Klivans 2016](#), [Klivans and Sherstov 2009](#)]) or from average case assumptions (e.g., [[Daniely 2016](#), [Daniely and Shalev-Shwartz 2016](#)]). [Blum \[1994\]](#) showed a “computational” separation between PAC learning and online mistake bound (which correspond to the VC Dimension and Littlestone's Dimension, respectively): if one-way function exists, then there is a concept class that can be learned by a computationally bounded learner in the PAC model, but not in the mistake-bound model.

Recently, [Bazgan et al. \[2016\]](#) introduced a generalization of the VC Dimension, which they call Partial VC Dimension, and proved that it is NP-hard to approximate (even when given an explicit description of the universe and concept class).

It is interesting to note that in contrast to other problems discussed in this part of the book, VC's and Littlestone's Dimensions do not naturally fit into the bilinear optimization framework, and they can be computed *exactly* in quasi-polynomial time with completely different algorithms: For the VC Dimension, it suffices to simply enumerate over all $\log |\mathcal{C}|$ -tuples of elements (where \mathcal{C} denotes the concept class and $\log |\mathcal{C}|$ is the trivial upper bound on the VC dimension) [Linial et al. 1991]. Littlestone's Dimension can be computed in quasi-polynomial time via a recursive “divide and conquer” algorithm (see our Theorem 14.5).

14.4 Inapproximability of the VC Dimension

In this section, we present our reduction from LABEL COVER to the VC Dimension, stated more formally below. We note that this reduction, together with Moshkovitz-Raz PCP (Theorem 2.2), with parameter $\delta = 1/\log n$ gives a reduction from 3SAT of size n to VC Dimension of size $2^{n^{1/2+o(1)}}$ with gap $1/2 + o(1)$, which immediately implies Theorem 14.1.

Theorem 14.3 For every $\delta > 0$, there exists a randomized reduction from a bi-regular Label Cover instance $\mathcal{L} = (A, B, E, \Sigma, \{\pi_e\}_{e \in E})$ to a ground set \mathcal{U} and a concept class \mathcal{C} such that, if $n \triangleq |A| + |B|$ and $r \triangleq \sqrt{n}/\log n$, then the following conditions hold for every sufficiently large n .

- (Size) The reduction runs in time $|\Sigma|^{O(|E| \text{poly}(1/\delta)/r)}$ and $|\mathcal{C}|, |\mathcal{U}| \leq |\Sigma|^{O(|E| \text{poly}(1/\delta)/r)}$.
- (Completeness) If \mathcal{L} is satisfiable, then $\text{VC-dim}(\mathcal{C}, \mathcal{U}) \geq 2r$.
- (Soundness) If $\text{val}(\mathcal{L}) \leq \delta^2/100$, then $\text{VC-dim}(\mathcal{C}, \mathcal{U}) \leq (1 + \delta)r$ with high probability.

14.4.1 A Candidate Reduction (and Why It Fails)

To best understand the intuition behind our reduction, we first describe a simpler candidate reduction and explain why it fails, which will lead us to the eventual construction. In this candidate reduction, we start by evoking Lemma 2.6 to partition the vertices $A \cup B$ of the Label Cover instance $\mathcal{L} = (A, B, E, \Sigma, \{\pi_e\}_{e \in E})$ into $U_1 \triangleq (S_1, T_1), \dots, U_r$ where $r = \sqrt{n}/\log n$. We then create the universe \mathcal{U} and the concept class \mathcal{C} as follows:

- We make each element in \mathcal{U} correspond to a partial assignment to U_i for some $i \in [r]$, i.e., we let $\mathcal{U} = \{x_{i, \sigma_i} \mid i \in [r], \sigma_i \in \Sigma^{U_i}\}$. In the completeness case, we expect to shatter the set of size r that corresponds to a satisfying

assignment $\sigma^* \in \Sigma^{A \cup B}$ of the Label Cover instance \mathcal{L} , i.e., $\{x_{i, \sigma^*|_{U_i}} \mid i \in [r]\}$. As for the soundness, our hope is that, if a large set $S \subseteq \mathcal{U}$ gets shattered, then we will be able to decode an assignment for \mathcal{L} that satisfies many constraints, which contradicts our assumption that $\text{val}(\mathcal{L})$ is small. Note that the number of elements used in this candidate reduction is at most $r \cdot |\Sigma|^{O(|E| \text{poly}(1/\delta)r)} = 2^{\tilde{O}(\sqrt{n})}$ as desired.

- As stated above, the intended solution for the completeness case is $\{x_{i, \sigma^*|_{U_i}} \mid i \in [r]\}$, meaning that we must have at least one concept corresponding to each subset $I \subseteq [r]$. We will try to make our concepts “test” the assignment; for each $I \subseteq [r]$, we will choose a set $T_I \subseteq A \cup B$ of $\tilde{O}(\sqrt{n})$ vertices and “test” all the constraints within T_I . Before we specify how T_I is picked, let us elaborate what “test” means: for each T_I -partial assignment ϕ_I that does not violate any constraints within T_I , we create a concept C_{I, ϕ_I} . This concept contains x_{i, σ_i} if and only if $i \in I$ and σ_i agrees with ϕ_I (i.e., $\phi_I|_{T_I \cap U_i} = \sigma_i|_{T_I \cap U_i}$). Recall that, if a set $S \subseteq \mathcal{U}$ is shattered, then each $T \subseteq S$ is an intersection between S and C_{I, ϕ_I} for some I, ϕ_I . We hope that the I 's are different for different T so that many different tests have been performed on S .
- Finally, let us specify how we pick T_I . Assume without loss of generality that r is even. We randomly pick a perfect matching over $[r]$, i.e., we pick a random permutation $\pi_I : [r] \rightarrow [r]$ and let $(\pi_I(1), \pi_I(2)), \dots, (\pi_I(r-1), \pi_I(r))$ be the chosen matching. We pick T_I such that all the constraints in the matchings, i.e., constraints between $U_{\pi_I(2i-1)}$ and $U_{\pi_I(2i)}$ for every $i \in [r/2]$, are included. More specifically, for every $i \in [r]$, we include each vertex $v \in U_{\pi_I(2i-1)}$ if at least one of its neighbors lie in $U_{\pi_I(2i)}$ and we include each vertex $u \in U_{\pi_I(2i)}$ if at least one of its neighbors lie in $U_{\pi_I(2i-1)}$. (By Lemma 2.6, for every pair in the matching the size of the intersection is at most $\frac{2|E|}{r^2}$, so each concept contains assignments to at most $\frac{2|E|}{r}$ variables; so the total size of the concept class is at most $2^r \cdot |\Sigma|^{\frac{2|E|}{r}}$.)

Even though the above reduction has the desired size and completeness, it unfortunately fails in the soundness. Let us now sketch a counterexample. For simplicity, let us assume that each vertex in $T_{[r]}$ has a unique neighbor in $T_{[r]}$. Note that, since $T_{[r]}$ has quite small size (only $\tilde{O}(\sqrt{n})$), almost all the vertices in $T_{[r]}$ satisfy this property w.h.p., but assuming that all of them satisfy this property makes our life easier.

Pick an assignment $\tilde{\sigma} \in \Sigma^V$ such that none of the constraints in $T_{[r]}$ is violated. From our unique neighbor assumption, there is always such an assignment. Now, we claim that the set $S_{\tilde{\sigma}} \triangleq \{x_{i, \tilde{\sigma}|_{U_i}} \mid i \in [r]\}$ gets shattered. This is because, for every

subset $I \subseteq [r]$, we can pick another assignment σ' such that σ' does not violate any constraint in $T_{[r]}$ and $\sigma'|_{U_i} = \bar{\sigma}|_{U_i}$ if and only if $i \in I$. This implies that $\{x_i, \bar{\sigma}_i \mid i \in I\} = S \cap C_{[r], \sigma'}$ as desired. Note here that such σ' exists because, for every $i \notin I$, if there is a constraint from a vertex $a \in U_i$ to another vertex $b \in T_{[r]}$, then we can change the assignment to a in such a way that the constraint is not violated; by doing this for every $i \notin I$, we have created the desired σ' . As a result, $\text{VC-dim}(\mathcal{C}, \mathcal{U})$ can still be as large as r even when the value of \mathcal{L} is small.

14.4.2 The Final Reduction

In this subsection, we will describe the actual reduction. To do so, let us first take a closer look at the issue with the above candidate reduction. In the candidate reduction, we can view each $I \subseteq [r]$ as being a seed used to pick a matching. Our hope was that many seeds participate in shattering some set S , and that this means that S corresponds to an assignment of high value. However, the counterexample showed that in fact only one seed ($I = [r]$) is enough to shatter a set. To circumvent this issue, we will not use the subset I as our seed anymore. Instead, we create r new elements y_1, \dots, y_r , which we will call *test selection elements*, to act as seeds; namely, each subset $H \subseteq \mathcal{Y}$ will now be a seed. The benefit of this is that, if $S \subseteq \mathcal{Y}$ is shattered and contains test selection elements y_{i_1}, \dots, y_{i_t} , then at least 2^t seeds must participate in the shattering of S . This is because, for each $H \subseteq \mathcal{Y}$, the intersection of S with any concept corresponding to H , when restricted to \mathcal{Y} , is always $H \cap \{y_{i_1}, \dots, y_{i_t}\}$. Hence, each subset of $\{y_{i_1}, \dots, y_{i_t}\}$ must come from a different seed.

The only other change from the candidate reduction is that each H will test multiple matchings rather than one matching. This is due to a technical reason: we need the number of matchings, ℓ , to be large in order to get the approximation ratio down to $1/2 + o(1)$; in our proof, if $\ell = 1$, then we can only achieve a factor of $1 - \varepsilon$ to some $\varepsilon > 0$. The full details of the reduction are shown in Algorithm 14.1.

Before we proceed to the proof, let us define some additional notation that will be used throughout.

- Every assignment element of the form x_{i, σ_i} is called an *i -assignment element*; we denote the set of all i -assignment elements by \mathcal{X}_i , i.e., $\mathcal{X}_i = \{x_{i, \sigma_i} \mid \sigma_i \in \Sigma^{U_i}\}$. Let \mathcal{X} denote all the assignment elements, i.e., $\mathcal{X} = \bigcup_i \mathcal{X}_i$.
- For every $S \subseteq \mathcal{U}$, let $I(S)$ denote the set of all $i \in [r]$ such that S contains an i -assignment element, i.e., $I(S) = \{i \in [r] \mid S \cap \mathcal{X}_i \neq \emptyset\}$.
- We call a set $S \subseteq \mathcal{X}$ *non-repetitive* if, for each $i \in [r]$, S contains at most one i -assignment element, i.e., $|S \cap \mathcal{X}_i| \leq 1$. Each non-repetitive set S canonically

Algorithm 14.1 Reduction from Label Cover to VC Dimension

Input: A bi-regular Label Cover instance $\mathcal{L} = (A, B, E, \Sigma, \{\pi_e\}_{e \in E})$ and a parameter $\delta > 0$.

Output: A ground set \mathcal{U} and a concept class \mathcal{C} .

The procedure to generate $(\mathcal{U}, \mathcal{C})$ works as follows:

1. Let r be $\sqrt{n}/\log n$ where $n = |A| + |B|$. Use Lemma 2.6 to partition $A \cup B$ into r blocks U_1, \dots, U_r .
2. For convenience, we assume that r is even. Moreover, for $i \neq j \in [r]$, let $\mathcal{N}_i(j) \subseteq U_i$ denote the set of all vertices in U_i with at least one neighbor in U_j (w.r.t. the graph (A, B, E)). We also extend this notation naturally to a set of j 's; for $J \subseteq [r]$, $\mathcal{N}_i(J)$ denotes $\bigcup_{j \in J} \mathcal{N}_i(j)$.
3. The universe \mathcal{U} consists of two types of elements, as described below.
 - (a) *Assignment elements*: for every $i \in [r]$ and every partial assignment $\sigma_i \in \Sigma^{U_i}$, there is an assignment element x_{i, σ_i} corresponding to it. Let \mathcal{X} denote all the assignment elements, i.e., $\mathcal{X} = \{x_{i, \sigma_i} \mid i \in [r], \sigma_i \in \Sigma^{U_i}\}$.
 - (b) *Test selection elements*: there are r test selection elements, which we will call y_1, \dots, y_r . Let \mathcal{Y} denote the set of all test selection elements.
4. The concepts in \mathcal{C} are defined by the following procedure.
 - (a) Let $\ell \triangleq 80/\delta^3$ be the number of matchings to be tested.
 - (b) For each $H \subseteq \mathcal{Y}$, we randomly select ℓ permutations $\pi_H^{(1)}, \dots, \pi_H^{(\ell)} : [r] \rightarrow [r]$; this gives us ℓ matchings (i.e., the t -th matching is $(\pi_H^{(t)}(1), \pi_H^{(t)}(2)), \dots, (\pi_H^{(t)}(r-1), \pi_H^{(t)}(r))$). For brevity, let us denote the set of (up to ℓ) elements that i is matched with in the matchings by $M_H(i)$. Let $T_H = \bigcup_i \mathcal{N}_i(M_H(i))$.
 - (c) For every $I \subseteq [r]$, $H \subseteq \mathcal{Y}$ and for every partial assignment $\sigma_H \in \Sigma^{T_H}$ that does not violate any constraints, we create a concept C_{I, H, σ_H} such that each $x_{i, \sigma_i} \in \mathcal{X}$ is included in C_{I, H, σ_H} if and only if $i \in I$ and σ_i is consistent with σ_H , i.e., $\sigma_i|_{\mathcal{N}_i(M_H(i))} = \sigma_H|_{\mathcal{N}_i(M_H(i))}$, whereas $y_i \in \mathcal{Y}$ is included in C_{I, H, σ_H} if and only if $y_i \in H$.

induces a partial assignment $\phi(S) : \bigcup_{i \in I(S)} U_i \rightarrow \Sigma$. This is the unique partial assignment that satisfies $\phi(S)|_{U_i} = \sigma_i$ for every $x_{i, \sigma_i} \in S$.

- Even though we define each concept as C_{I, H, σ_H} where σ_H is a partial assignment to a subset $T_H \subset A \cup B$, it will be more convenient to view each concept as $C_{I, H, \sigma}$ where $\sigma \in \Sigma^V$ is the assignment to the entire Label Cover instance. This is just a notational change: the actual definition of the concept does not depend on the assignment outside T_H .

- For each $I \subseteq [r]$, let U_I denote $\bigcup_{i \in I} U_i$. For each $\sigma_I \in \Sigma^{U_I}$, we say that (I, σ_I) passes $H \subseteq \mathcal{Y}$ if σ_I does not violate any constraint within T_H . Denote the collection of H 's that (I, σ_I) passes by $\mathcal{H}(I, \sigma_I)$.
- Finally, for any non-repetitive set $S \subseteq \mathcal{X}$ and any $H \subseteq \mathcal{Y}$, we say that S passes H if $(I(S), \phi(S))$ passes H . We write $\mathcal{H}(S)$ as a shorthand for $\mathcal{H}(I(S), \phi(S))$.

The output size of the reduction and the completeness follow almost immediately from definition.

Output Size of the Reduction. Clearly, the size of \mathcal{U} is $\sum_{i \in [r]} |\Sigma|^{U_i} \leq r \cdot |\Sigma|^{n/r} \leq |\Sigma|^{O(|E| \text{poly}(1/\delta)/r)}$. As for $|\mathcal{C}|$, note first that the number of choices for I and H are both 2^r . For fixed I and H , Lemma 2.6 implies that, for each matching $\pi_H^{(t)}$, the number of vertices from each U_i with at least one constraint to the matched partition in $\pi_H^{(t)}$ is at most $O(|E|/r^2)$. Since there are ℓ matchings, the number of vertices in $T_H = \mathcal{N}_1(M_H(1)) \cup \dots \cup \mathcal{N}_r(M_H(r))$ is at most $O(|E|\ell/r)$. Hence, the number of choices for the partial assignment σ_H is at most $|\Sigma|^{O(|E| \text{poly}(1/\delta)/r)}$. In total, we can conclude that \mathcal{C} contains at most $|\Sigma|^{O(|E| \text{poly}(1/\delta)/r)}$ concepts.

Completeness. If \mathcal{L} has a satisfying assignment $\sigma^* \in \Sigma^V$, then the set $S_{\sigma^*} = \{x_i, \sigma_i^* \mid i \in [r]\} \cup \mathcal{Y}$ is shattered because, for any $S \subseteq S_{\sigma^*}$, we have $S = S_{\sigma^*} \cap C_{I(S), S \cap \mathcal{Y}, \sigma^*}$. Hence, $\text{VC-dim}(\mathcal{C}, \mathcal{U}) \geq 2r$.

The rest of this section is devoted to the soundness analysis.

14.4.3 Soundness

In this subsection, we will prove the following lemma, which, combined with the completeness and output size arguments above, implies Theorem 14.3. For brevity, we will assume throughout this subsection that r is sufficiently large, and leave it out of the lemmas' statements.

Lemma 14.1 Let $(\mathcal{C}, \mathcal{U})$ be the output from the reduction in Algorithm 14.1 on input \mathcal{L} . If $\text{val}(\mathcal{L}) \leq \delta^2/100$, then $\text{VC-dim}(\mathcal{C}, \mathcal{U}) \leq (1 + \delta)r$ with high probability.

At a high level, the proof of Lemma 14.1 has two steps:

1. Given a shattered set $S \subseteq \mathcal{U}$, we extract a maximal non-repetitive set $S^{\text{no-rep}} \subseteq S$ such that $S^{\text{no-rep}}$ passes many ($\geq 2^{|S| - |S^{\text{no-rep}}|}$) H 's. If $|S^{\text{no-rep}}|$ is small, the trivial upper bound of 2^r on the number of different H 's implies that $|S|$ is also small. As a result, we are left to deal with the case that $|S^{\text{no-rep}}|$ is large.
2. When $|S^{\text{no-rep}}|$ is large, $S^{\text{no-rep}}$ induces a partial assignment on a large fraction of vertices of \mathcal{L} . Since we assume that $\text{val}(\mathcal{L})$ is small, this partial assignment must violate many constraints. We will use this fact to argue that, with high

probability, $S^{\text{no-rep}}$ only passes very few H 's, which implies that $|S|$ must be small.

The two parts of the proof are presented in Sections 14.4.3.1 and 14.4.3.2 respectively. We then combine them in Section 14.4.3.3 to prove Lemma 14.1.

14.4.3.1 Part I: Finding a Non-Repetitive Set That Passes Many Tests

The goal of this subsection is to prove the following lemma, which allows us, given a shattered set $S \subseteq \mathcal{U}$, to find a non-repetitive set $S^{\text{no-rep}}$ that passes many H 's.

Lemma 14.2 For any shattered $S \subseteq \mathcal{U}$, there is a non-repetitive set $S^{\text{no-rep}}$ of size $|I(S)|$ such that $|\mathcal{H}(S^{\text{no-rep}})| \geq 2^{|S| - |I(S)|}$.

We will start by proving the following lemma, which will be a basis for the proof of Lemma 14.2.

Lemma 14.3 Let $C, C' \in \mathcal{C}$ correspond to the same H (i.e., $C = C_{I, H, \sigma}$ and $C' = C_{I', H, \sigma'}$ for some $H \subseteq \mathcal{Y}$, $I, I' \subseteq [r]$, $\sigma, \sigma' \in \Sigma^V$).

For any subset $S \subseteq \mathcal{U}$ and any maximal non-repetitive subset $S^{\text{no-rep}} \subseteq S$ such that $I(S^{\text{no-rep}}) = I(S)$, if $S^{\text{no-rep}} \subseteq C$ and $S^{\text{no-rep}} \subseteq C'$, then $S \cap C = S \cap C'$.

The most intuitive interpretation of this lemma is as follows. Recall that if S is shattered, then, for each $\tilde{S} \subseteq S$, there must be a concept $C_{I_{\tilde{S}}, H_{\tilde{S}}, \sigma_{\tilde{S}}}$ such that $\tilde{S} = S \cap C_{I_{\tilde{S}}, H_{\tilde{S}}, \sigma_{\tilde{S}}}$. The above lemma implies that, for each $\tilde{S} \supseteq S^{\text{no-rep}}$, $H_{\tilde{S}}$ must be different. This means that at least $2^{|S| - |S^{\text{no-rep}}|}$ different H 's must be involved in shattering S . Indeed, this will be the argument we use when we prove Lemma 14.2.

Proof Proof of Lemma 14.3. Let $S, S^{\text{no-rep}}$ be as in the lemma statement. Suppose for the sake of contradiction that there exists $H \subseteq \mathcal{Y}$, $I, I' \subseteq [r]$, $\sigma, \sigma' \in \Sigma^V$ such that $S^{\text{no-rep}} \subseteq C_{I, H, \sigma}$, $S^{\text{no-rep}} \subseteq C_{I', H, \sigma'}$ and $S \cap C_{I, H, \sigma} \neq S \cap C_{I', H, \sigma'}$.

First, note that $S \cap C_{I, H, \sigma} \cap \mathcal{Y} = S \cap H \cap \mathcal{Y} = S \cap C_{I', H, \sigma'} \cap \mathcal{Y}$. Since $S \cap C_{I, H, \sigma} \neq S \cap C_{I', H, \sigma'}$, we must have $S \cap C_{I, H, \sigma} \cap \mathcal{X} \neq S \cap C_{I', H, \sigma'} \cap \mathcal{X}$. Assume w.l.o.g. that there exists $x_{i, \sigma_i} \in (S \cap C_{I, H, \sigma}) \setminus (S \cap C_{I', H, \sigma'})$.

Note that $i \in I(S) = I(S^{\text{no-rep}})$ (where the equality follows by maximality of $S^{\text{no-rep}}$). Thus there exists $\sigma'_i \in \Sigma^{U_i}$ such that $x_{i, \sigma'_i} \in S^{\text{no-rep}} \subseteq C_{I, H, \sigma} \cap C_{I', H, \sigma'}$. Since x_{i, σ'_i} is in both $C_{I, H, \sigma}$ and $C_{I', H, \sigma'}$, we have $i \in I \cap I'$ and

$$\sigma|_{\mathcal{N}_i(M_H(i))} = \sigma'_i|_{\mathcal{N}_i(M_H(i))} = \sigma'|_{\mathcal{N}_i(M_H(i))}. \quad (14.1)$$

However, since $x_{i,\sigma_i} \in (S \cap C_{I,H,\sigma}) \setminus (S \cap C_{I',H,\sigma'})$, we have $x_{i,\sigma_i} \in C_{I,H,\sigma} \setminus C_{I',H,\sigma'}$. This implies that

$$\sigma|_{\mathcal{N}_i(M_H(i))} = \sigma_i|_{\mathcal{N}_i(M_H(i))} \neq \sigma'|_{\mathcal{N}_i(M_H(i))},$$

which contradicts (14.1). ■

In addition to the above lemma, we will also need the following observation, which states that, if a non-repetitive $S^{\text{no-rep}}$ is contained in a concept C_{I,H,σ_H} , then $S^{\text{no-rep}}$ must pass H . This observation follows definitions.

Observation 14.1 If a non-repetitive set $S^{\text{no-rep}}$ is a subset of some concept C_{I,H,σ_H} , then $H \in \mathcal{H}(S^{\text{no-rep}})$.

With Lemma 14.3 and Observation 14.1 ready, it is now easy to prove Lemma 14.2.

Proof Proof of Lemma 14.2. Pick $S^{\text{no-rep}}$ to be any maximal non-repetitive subset of S such that $I(S^{\text{no-rep}}) = I(S)$. Clearly, $|S^{\text{no-rep}}| = |I(S)|$. To see that $|\mathcal{H}(S^{\text{no-rep}})| \geq 2^{|S|-|I(S)|}$, consider any \tilde{S} such that $S^{\text{no-rep}} \subseteq \tilde{S} \subseteq S$. Since S is shattered, there exists $I_{\tilde{S}}, H_{\tilde{S}}, \sigma_{\tilde{S}}$ such that $S \cap C_{I_{\tilde{S}}, H_{\tilde{S}}, \sigma_{\tilde{S}}} = \tilde{S}$. Since $\tilde{S} \supseteq S^{\text{no-rep}}$, Observation 14.1 implies that $H_{\tilde{S}} \in \mathcal{H}(S^{\text{no-rep}})$. Moreover, from Lemma 14.3, $H_{\tilde{S}}$ is distinct for every \tilde{S} . As a result, $|\mathcal{H}(S^{\text{no-rep}})| \geq 2^{|S|-|I(S)|}$ as desired. ■

14.4.3.2 Part II: No Large Non-Repetitive Set Passes Many Tests

The goal of this subsection is to show that, if $\text{val}(\mathcal{L})$ is small, then w.h.p. (over the randomness in the construction) every large non-repetitive set passes only few H 's. This is formalized as Lemma 14.4 below.

Lemma 14.4 If $\text{val}(\mathcal{L}) \leq \delta^2/100$, then, with high probability, for every non-repetitive set $S^{\text{no-rep}}$ of size at least δr , $|\mathcal{H}(S^{\text{no-rep}})| \leq 100n \log |\Sigma|$.

Note that the mapping $S^{\text{no-rep}} \mapsto (I(S^{\text{no-rep}}), \phi(S^{\text{no-rep}}))$ is a bijection from the collection of all non-repetitive sets to $\{(I, \sigma_I) \mid I \subseteq [r], \sigma_I \in \Sigma^{U_I}\}$. Hence, the above lemma is equivalent to the following.

Lemma 14.5 If $\text{val}(\mathcal{L}) \leq \delta^2/100$, then, with high probability, for every $I \subseteq [r]$ of size at least δr and every $\sigma_I \in \Sigma^{U_I}$, $|\mathcal{H}(I, \sigma_I)| \leq 100n \log |\Sigma|$.

Here we use the language in Lemma 14.5 instead of Lemma 14.4 as it will be easier for us to reuse this lemma later. To prove the lemma, we first need to bound

the probability that the assignment σ_I does not violate any constraint induced by a random matching. More precisely, we will prove the following lemma.

Lemma 14.6 For any $I \subseteq [r]$ of size at least δr and any $\sigma_I \in \Sigma^{U_I}$, if $\pi : [r] \rightarrow [r]$ is a random permutation of $[r]$, then the probability that σ_I does not violate any constraint in $\bigcup_{i \in [r]} \mathcal{N}_i(M(i))$ is at most $(1 - 0.1\delta^2)^{\delta r/8}$ where $M(i)$ denotes the index that i is matched with in the matching $(\pi(1), \pi(2)), \dots, (\pi(r-1), \pi(r))$.

Proof Let p be any positive odd integer such that $p \leq \delta r/2$ and let $i_1, \dots, i_{p-1} \in [r]$ be any $p-1$ distinct elements of $[r]$. We will first show that conditioned on $\pi(1) = i_1, \dots, \pi(p-1) = i_{p-1}$, the probability that σ_I violates a constraint induced by $\pi(p), \pi(p+1)$ (i.e., in $\mathcal{N}_{\pi(p)}(\pi(p+1)) \cup \mathcal{N}_{\pi(p+1)}(\pi(p))$) is at least $0.1\delta^2$.

To see that this is true, let $I_{\geq p} = I \setminus \{i_1, \dots, i_{p-1}\}$. Since $|I| \geq \delta r$, we have $|I_{\geq p}| = |I| - p + 1 \geq \delta r/2 + 1$. Consider the partial assignment $\sigma_{\geq p} = \sigma_I|_{U_{I_{\geq p}}}$. Since $\text{val}(\mathcal{L}) \leq 0.01\delta^2$, $\sigma_{\geq p}$ can satisfy at most $0.01\delta^2|E|$ constraints. From Lemma 2.6, we have, for every $i \neq j \in I_{\geq p}$, the number of constraints between U_i and U_j is at least $|E|/r^2$. Hence, there are at most $0.01\delta^2 r^2$ pairs of $i < j \in I_{\geq p}$ such that $\sigma_{\geq p}$ does not violate any constraint between U_i and U_j . In other words, there are at least $\binom{|I_{\geq p}|}{2} - 0.01\delta^2 r^2 \geq 0.1\delta^2 r^2$ pairs $i < j \in I_{\geq p}$ such that $\sigma_{\geq p}$ violates some constraints between U_i and U_j . Now, if $\pi(p) = i$ and $\pi(p+1) = j$ for some such pair i, j , then $\phi(S^{\text{no-rep}})$ violates a constraint induced by $\pi(p), \pi(p+1)$. Thus, we have

$$\begin{aligned} \Pr\left[\sigma_I \text{ doesn't violate a constraint induced by } \pi(p), \pi(p+1) \mid \bigwedge_{t=1}^{p-1} \pi(t) = i_t\right] \\ \leq 1 - 0.1\delta^2. \end{aligned} \quad (14.2)$$

Let E_p denote the event that σ_I does not violate any constraints induced by $\pi(p)$ and $\pi(p+1)$. We can now bound the desired probability as follows.

$$\begin{aligned} \Pr\left[\sigma_I \text{ doesn't violate any constraint in } \bigcup_{i \in [r]} \mathcal{N}_i(M(i))\right] &\leq \Pr\left[\bigwedge_{\text{odd } p \in [\delta r/2+1]} E_p\right] \\ &= \prod_{\text{odd } p \in [\delta r/2+1]} \Pr\left[E_p \mid \bigwedge_{\text{odd } t \in [p-1]} E_t\right] \\ &\stackrel{\text{(from (14.2))}}{\leq} \prod_{\text{odd } p \in [\delta r/2+1]} (1 - 0.1\delta^2) \\ &\leq (1 - 0.1\delta^2)^{\delta r/4-1}, \end{aligned}$$

which is at most $(1 - 0.1\delta^2)^{\delta r/8}$ for sufficiently large r (i.e., $r \geq 8/\delta$). ■

We can now prove our main lemma.

Proof Proof of Lemma 14.5. For a fixed $I \subseteq [r]$ of size at least δr and a fixed $\sigma_I \in \Sigma^{U_I}$, Lemma 14.6 tells us that the probability that σ_I does not violate any constraint induced by a single matching is at most $(1 - 0.1\delta^2)^{\delta r/8}$. Since for each $H \subseteq \mathcal{Y}$ the construction picks ℓ matchings at random, the probability that (I, σ_I) passes each H is at most $(1 - 0.1\delta^2)^{\delta \ell r/8}$. Recall that we pick $\ell = 80/\delta^3$; this gives the following upper bound on the probability:

$$\begin{aligned} \Pr[(I, \sigma_I) \text{ passes } H] &\leq (1 - 0.1\delta^2)^{\delta \ell r/8} = (1 - 0.1\delta^2)^{10r/\delta^2} \\ &\leq \left(\frac{1}{1 + 0.1\delta^2} \right)^{10r/\delta^2} \leq 2^{-r} \end{aligned} \quad (14.3)$$

where the last inequality comes from Bernoulli's inequality.

Inequality (14.3) implies that the expected number of H 's that (I, σ_I) passes is less than 1. Since the matchings M_H are independent for all H 's, we can apply Chernoff bound, which implies that

$$\Pr\left[|\mathcal{H}(I, \sigma_I)| \geq 100n \log |\Sigma|\right] \leq 2^{-10n \log |\Sigma|} = |\Sigma|^{-10n}.$$

Finally, note that there are at most $2^r |\Sigma|^n$ different (I, σ_I) 's. By union bound, we have

$$\begin{aligned} \Pr\left[\exists I \subseteq [r], \sigma_I \in \Sigma^{U_I} \text{ s.t. } |I| \geq \delta r \text{ AND } |\mathcal{H}(I, \sigma_I)| \geq 100n \log |\Sigma|\right] &\leq (2^r |\Sigma|^n) (|\Sigma|^{-10n}) \\ &\leq |\Sigma|^{-8n}, \end{aligned}$$

which concludes the proof. \blacksquare

14.4.3.3 Putting Things Together

Proof Proof of Lemma 14.1. From Lemma 14.4, every non-repetitive set $S^{\text{no-rep}}$ of size at least δr , $|\mathcal{H}(S^{\text{no-rep}})| \leq 100n \log |\Sigma|$. Conditioned on this event happening, we will show that $\text{VC-dim}(\mathcal{U}, \mathcal{C}) \leq (1 + \delta)r$.

Consider any shattered set $S \subseteq \mathcal{U}$. Lemma 14.2 implies that there is a non-repetitive set $S^{\text{no-rep}}$ of size $|I(S)|$ such that $|\mathcal{H}(S^{\text{no-rep}})| \geq 2^{|S| - |I(S)|}$. Let us consider two cases:

1. $|I(S)| \leq \delta r$. Since $\mathcal{H}(S^{\text{no-rep}}) \subseteq \mathcal{P}(\mathcal{Y})$, we have $|S| - |I(S)| \leq |\mathcal{Y}| = r$. This implies that $|S| \leq (1 + \delta)r$.
2. $|I(S)| > \delta r$. From our assumption, $|\mathcal{H}(S^{\text{no-rep}})| \leq 100n \log |\Sigma|$. Thus, $|S| \leq |I(S)| + \log(100n \log |\Sigma|) \leq r + o(r)$; when r is sufficiently large, the latter expression is at most $(1 + \delta)r$.

Hence, for sufficiently large r , $\text{VC-dim}(\mathcal{U}, \mathcal{C}) \leq (1 + \delta)r$ with high probability. \blacksquare

14.5 Inapproximability of Littlestone's Dimension

We next proceed to Littlestone's Dimension. The main theorem of this section is stated below. Again, note that this theorem, together with Theorem 2.2, implies Theorem 14.2.

Theorem 14.4 There exists $\varepsilon > 0$ such that there is a randomized reduction from a bi-regular Label Cover instance $\mathcal{L} = (A, B, E, \Sigma, \{\pi_e\}_{e \in E})$ to a ground set \mathcal{U} and a concept class \mathcal{C} such that, if $n \triangleq |A| + |B|$, $r \triangleq \sqrt{n}/\log n$ and $k \triangleq 10^{10}|E| \log |\Sigma|/r^2$, then the following conditions hold for every sufficiently large n .

- (Size) The reduction runs in time $2^{rk} \cdot |\Sigma|^{O(|E|/r)}$ and $|\mathcal{C}|, |\mathcal{U}| \leq 2^{rk} \cdot |\Sigma|^{O(|E|/r)}$.
- (Completeness) If \mathcal{L} is satisfiable, then $\text{L-dim}(\mathcal{C}, \mathcal{U}) \geq 2rk$.
- (Soundness) If $\text{val}(\mathcal{L}) \leq 0.001$, then $\text{L-dim}(\mathcal{C}, \mathcal{U}) \leq (2 - \varepsilon)rk$ with high probability.

14.5.1 Why the VC Dimension Reduction Fails for Littlestone's Dimension

It is tempting to think that, since our reduction from the previous section works for the VC Dimension, it may also work for Littlestone's Dimension. In fact, thanks to Fact 2.2, completeness for that reduction even translates for free to Littlestone's Dimension. Alas, the soundness property does not hold. To see this, let us build a depth- $2r$ mistake tree for \mathcal{C}, \mathcal{U} , even when $\text{val}(\mathcal{L})$ is small, as follows.

- We assign the test-selection elements to the first r levels of the tree, one element per level. More specifically, for each $s \in \{0, 1\}^{<r}$, we assign $y_{|s|+1}$ to s .
- For every string $s \in \{0, 1\}^r$, the previous step of the construction gives us a subset of \mathcal{Y} corresponding to the path from root to s ; this subset is simply $H_s = \{y_i \in \mathcal{Y} \mid s_i = 1\}$. Let T_{H_s} denote the set of vertices tested by this seed H_s . Let $\phi_s \in \Sigma^V$ denote an assignment that satisfies all the constraints in T_{H_s} . Note that, since T_{H_s} is of small size (only $\tilde{O}(\sqrt{n})$), even if $\text{val}(\mathcal{L})$ is small, ϕ_s is still likely to exist (and we can decide whether it exists or not in time $2^{\tilde{O}(\sqrt{n})}$).

We then construct the subtree rooted at s that corresponds to ϕ_s by assigning each level of the subtree $x_{i, \phi_s|_{U_i}}$. Specifically, for each $t \in \{0, 1\}^{\geq r}$, we assign $x_{|t|-r+1, \phi_{t_{\leq r}}|_{U_{|t|-r+1}}}$ to node t of the tree.

It is not hard to see that the constructed tree is indeed a valid mistake tree. This is because the path from root to each leaf $l \in \{0, 1\}^{2r}$ agrees with $C_{I(l), H_{l_{\leq r}}, \phi_{l_{\leq r}}}$ (where $I(l) = \{i \in [r] \mid l_i = 1\}$).

14.5.2 The Final Reduction

The above counterexample demonstrates the main difference between the two dimensions: order does not matter in VC Dimension, but it does in Littlestone's Dimension. By moving the test-selection elements up the tree, the tests are chosen before the assignments, which allows an adversary to "cheat" by picking different assignments for different tests. We would like to prevent this, i.e., we would like to make sure that, in the mistake tree, the upper levels of the tree are occupied with the assignment elements whereas the lower levels are assigned test-selection elements. As in the VC Dimension argument, our hope here is that, given such a tree, we should be able to decode an assignment that passes many different tests. Indeed, we will tailor our construction to achieve such property.

Recall that, if we use the same reduction as VC Dimension, then, in the completeness case, we can construct a mistake tree in which the first r layers consist solely of assignment elements and the rest of the layers consist of only test-selection elements. Observe that there is no need for different nodes on the r -th layer to have subtrees composed of the same set of elements; the tree would still be valid if we made each test-selection element only work with a specific $s \in \{0, 1\}^r$ and created concepts accordingly. In other words, we can modify our construction so that our test-selection elements are $\mathcal{Y} = \{y_{I,i} \mid I \subseteq [r], i \in [r]\}$ and the concept class is $\{C_{I,H,\sigma_H} \mid I \subseteq [r], H \subseteq \mathcal{Y}, \sigma_H \in \Sigma^{T_H}\}$ where the condition that an assignment element lies in C_{I,H,σ_H} is the same as in the VC Dimension reduction, whereas for $y_{I',i}$ to be in C_{I,H,σ_H} , we require not only that $i \in H$ but also that $I = I'$. Intuitively, this should help us, since each $y_{I,i}$ is now only in a small fraction ($\leq 2^{-r}$) of concepts; hence, one would hope that any subtree rooted at any $y_{I,i}$ cannot be too deep, which would indeed imply that the test-selection elements cannot appear in the first few layers of the tree.

Alas, for this modified reduction, it is not true that a subtree rooted at any $y_{I,i}$ has small depth; specifically, we can bound the depth of a subtree $y_{I,i}$ by the log of the number of concepts containing $y_{I,i}$ plus one (for the first layer). Now, note that $y_{I,i} \in C_{I',H,\sigma_H}$ means that $I' = I$ and $i \in H$, but there still can be as many as $2^{r-1} \cdot |\Sigma|^{T_H} = |\Sigma|^{O(|E|/r)}$ such concepts. This gives an upper bound of $r + O(|E| \log |\Sigma|/r)$ on the depth of the subtree rooted at $y_{I,i}$. However, $|E| \log |\Sigma|/r = \Theta(\sqrt{n} \log n) = \omega(r)$; this bound is meaningless here since, even in the completeness case, the depth of the mistake tree is only $2r$.

Fortunately, this bound is not useless after all: if we can keep this bound but make the intended tree depth much larger than $|E| \log |\Sigma|/r$, then the bound will indeed imply that no $y_{I,i}$ -rooted tree is deep. To this end, our reduction will have one more parameter $k = \Theta(|E| \log |\Sigma|/r)$ where $\Theta(\cdot)$ hides a large constant and

the intended tree will have depth $2rk$ in the completeness case; the top half of the tree (first rk layers) will again consist of assignment elements and the rest of the tree comprises the test-selection elements. The rough idea is to make k “copies” of each element: the assignment elements will now be $\{x_{i,\sigma_i,j} \mid i \in [r], \sigma_i \in \Sigma^{U_i}, j \in [k]\}$ and the test-selection elements will be $\{y_{I,i,j} \mid I \subseteq [r] \times [k], j \in [k]\}$. The concept class can then be defined as $\{C_{I,H,\sigma_H} \mid I \subseteq [r] \times [k], H \subseteq [r] \times [k], \sigma_H \in \Sigma^{T_H}\}$ naturally, i.e., H is used as the seed to pick the test set T_H , $y_{I',i,j} \in C_{I,H,\sigma_H}$ iff $I' = I$ and $(i,j) \in H$, whereas $x_{i,\sigma_i,j} \in C_{I,H,\sigma_H}$ iff $(i,j) \in I$ and $\sigma_i|_{(I,\sigma_I)} = \sigma_H|_{(I,\sigma_I)}$. For this concept class, we can again bound the depth of a $y_{I,i}$ -rooted tree to be $rk + O(|E| \log |\Sigma|/r)$; this time, however, rk is much larger than $|E| \log |\Sigma|/r$, so this bound is no more than, say, $1.001rk$. This is indeed the desired bound, since this means that, for any depth- $1.999rk$ mistake tree, the first $0.998rk$ layers must consist solely of assignment elements.

Unfortunately, the introduction of copies in turn introduces another technical challenge: it is not true any more that a partial assignment to a large set only passes a few tests w.h.p. (i.e., an analogue of Lemma 14.5 does not hold). By Inequality (14.3), each H is passed with probability at most 2^{-r} , but now we want to take a union bound, and there are $2^{rk} \gg 2^r$ different H 's. To circumvent this, we will define a map $\tau : \mathcal{P}([r] \times [k]) \rightarrow \mathcal{P}([r])$ and use $\tau(H)$ to select the test instead of H itself. The map τ we use in the construction is the *threshold projection* where i is included in H if and only if, for at least half of $j \in [k]$, H contains (i,j) . To motivate our choice of τ , recall that our overall proof approach is to first find a node that corresponds to an assignment to a large subset of the Label Cover instance, then argue that it can pass only a few tests, which we hope would imply that the subtree rooted there cannot be too deep. For this implication to be true, we need the following to also hold: for any small subset $\mathcal{H} \subseteq \mathcal{P}([r])$ of $\tau(H)$'s, we have that $\text{L-dim}(\tau^{-1}(\mathcal{H}), [r] \times [k])$ is small. This property indeed holds for our choice of τ (see Lemma 14.13).

With all the moving parts explained, we state the full reduction formally in Algorithm 14.2.

Similar to our VC Dimension proof, we will use the following notation:

- For every $i \in [r]$, let $\mathcal{X}_i \triangleq \{x_{i,\sigma_i,j} \mid \sigma_i \in \Sigma^{U_i}, j \in [k]\}$; we refer to these elements as the i -assignment elements. Moreover, for every $(i,j) \in [r] \times [k]$, let $\mathcal{X}_{i,j} \triangleq \{x_{i,\sigma_i,j} \mid \sigma_i \in \Sigma^{U_i}\}$; we refer to these elements as the (i,j) -assignment elements.
- For every $S \subseteq \mathcal{U}$, let $I(S) = \{i \in [r] \mid S \cap \mathcal{X}_i \neq \emptyset\}$ and $IJ(S) = \{(i,j) \in [r] \times [k] \mid S \cap \mathcal{X}_{i,j} \neq \emptyset\}$.
- A set $S \subseteq \mathcal{X}$ is *non-repetitive* if $|S \cap \mathcal{X}_{i,j}| \leq 1$ for all $(i,j) \in [r] \times [k]$.

Algorithm 14.2 Reduction from Label Cover to Littlestone's Dimension

Input: A bi-regular Label Cover instance $\mathcal{L} = (A, B, E, \Sigma, \{\pi_e\}_{e \in E})$.

Output: A ground set \mathcal{U} and a concept class \mathcal{C} .

The procedure to generate $(\mathcal{U}, \mathcal{C})$ works as follows:

1. Let $r, U_1, \dots, U_r, \mathcal{N}$ be defined in the same manner as in Reduction 14.1 and let $k \triangleq 10^{10}|E| \log |\Sigma|/r^2$.
2. The universe \mathcal{U} consists of two types of elements, as described below.
 - (a) *Assignment elements*: For every $i \in [r]$, every partial assignment $\sigma_i \in \Sigma^{U_i}$, and every $j \in [k]$, there is an assignment element $x_{i, \sigma_i, j}$ corresponding to it. Let \mathcal{X} denote all the assignment elements, i.e., $\mathcal{X} = \{x_{i, \sigma_i, j} \mid i \in [r], \sigma_i \in \Sigma^{U_i}, j \in [k]\}$.
 - (b) *Test-selection elements*: There are $rk(2^{r^k})$ test-selection elements, which we will call $y_{I, i, j}$ for every $i \in [r]$, $j \in [k]$, $I \subseteq [r] \times [k]$. Let \mathcal{Y} denote the set of all test-selection elements. Let \mathcal{Y}_i denote $\{y_{I, i, j} \mid I \subseteq [r] \times [k], j \in [k]\}$. We call the elements of \mathcal{Y}_i *i-test-selection elements*.
3. The concepts in \mathcal{C} are defined by the following procedure.
 - (a) Let $\ell \triangleq 1000$ be the number of matchings to be tested.
 - (b) For each $\tilde{H} \subseteq [r]$, we randomly select ℓ permutations $\pi_{\tilde{H}}^{(1)}, \dots, \pi_{\tilde{H}}^{(\ell)} : [r] \rightarrow [r]$; this gives us ℓ matchings (i.e., the t -th matching is $(\pi_{\tilde{H}}^{(t)}(1), \pi_{\tilde{H}}^{(t)}(2)), \dots, (\pi_{\tilde{H}}^{(t)}(r-1), \pi_{\tilde{H}}^{(t)}(r))$). Denote the set of elements that i is matched with in the matchings by $M_{\tilde{H}}(i)$. Let $T_{\tilde{H}} = \bigcup_i \mathcal{N}_i(M_{\tilde{H}}(i))$.
 - (c) Let $\tau : \mathcal{P}([r] \times [k]) \rightarrow \mathcal{P}([r])$ denote the *threshold projection* operation where each $i \in [r]$ is included in $\tau(H)$ if and only if H contains at least half of the i -test-selection elements, i.e., $\tau(H) = \{i \in [r] \mid |H \cap \mathcal{Y}_i| \geq k/2\}$.
 - (d) For every $I \subseteq [r] \times [k]$, $H \subseteq [r] \times [k]$ and for every partial assignment $\sigma_{\tau(H)} \in \Sigma^{T_{\tau(H)}}$ that does not violate any constraints, we create a concept $C_{I, H, \sigma_{\tau(H)}}$ such that each $x_{i, \sigma_i, j} \in \mathcal{X}$ is included in $C_{I, H, \sigma_{\tau(H)}}$ if and only if $(i, j) \in I$ and σ_i is consistent with $\sigma_{\tau(H)}$, i.e., $\sigma_i|_{\mathcal{N}_i(M_{\tau(H)}(i))} = \sigma_{\tau(H)}|_{\mathcal{N}_i(M_{\tau(H)}(i))}$, whereas each $y_{I', i, j} \in \mathcal{Y}$ is included in $C_{I, H, \sigma_{\tau(H)}}$ if and only if $(i, j) \in H$ and $I' = I$.

- We say that S passes \tilde{H} if the following two conditions hold:
 - For every $i \in [r]$ such that $S \cap \mathcal{X}_i \neq \emptyset$, all i -assignment elements of S are consistent on $T_{\tilde{H}}|_{U_i}$, i.e., for every $(i, \sigma_i, j), (i, \sigma'_i, j') \in S$, we have $\sigma_i|_{U_i} = \sigma'_i|_{U_i}$.

- The canonically induced assignment on $T_{\tilde{H}}$ does not violate any constraint (note that the previous condition implies that such assignment is unique).

We use $\mathcal{H}(S)$ to denote the collection of all seeds $\tilde{H} \subseteq [r]$ that S passes.

We also use the following notation for mistake trees:

- For any subset $S \subseteq \mathcal{U}$ and any function $\rho : S \rightarrow \{0, 1\}$, let $\mathcal{C}[\rho] \triangleq \{C \in \mathcal{C} \mid \forall a \in S, a \in C \Leftrightarrow \rho(a) = 1\}$ be the collection of all concepts that agree with ρ on S . We sometimes abuse the notation and write $\mathcal{C}[S]$ to denote the collection of all the concepts that contain S , i.e., $\mathcal{C}[S] = \{C \in \mathcal{C} \mid S \subseteq C\}$.
- For any binary string s , let $\text{pre}(s) \triangleq \{\emptyset, s_{\leq 1}, \dots, s_{\leq |s|-1}\}$ denote the set of all proper prefixes of s .
- For any depth- d mistake tree \mathcal{T} , let $v_{\mathcal{T},s}$ denote the element assigned to the node $s \in \{0, 1\}^{\leq d}$, and let $P_{\mathcal{T},s} \triangleq \{v_{\mathcal{T},s'} \mid s' \in \text{pre}(s)\}$ denote the set of all elements appearing from the path from root to s (excluding s itself). Moreover, let $\rho_{\mathcal{T},s} : P_{\mathcal{T},s} \rightarrow \{0, 1\}$ be the function corresponding to the path from root to s , i.e., $\rho_{\mathcal{T},s}(s') = s_{|s'|+1}$ for every $s' \in \text{pre}(s)$.

Output Size of the Reduction The output size of the reduction follows immediately from a similar argument as in the VC Dimension reduction. The only difference here is that there are 2^{rk} choices for I and H , instead of 2^r choices as in the previous construction.

Completeness. If \mathcal{L} has a satisfying assignment $\sigma^* \in \Sigma^V$, we can construct a depth- rk mistake tree \mathcal{T} as follows. For $i \in [r]$, $j \in [k]$, we assign $x_{i,\sigma_i^*,j}$ to every node in the $((i-1)k + j)$ -th layer of \mathcal{T} . Note that we have so far assigned every node in the first rk layers. For the rest of the vertices, if s lies in layer $rk + (i-1)k + j$, then we assign $y_{I(\rho_{\mathcal{T},s}^{-1}(1)),i,j}$ to it. It is clear that, for a leaf $s \in \{0, 1\}^{rk}$, the concept $C_{I(\rho_{\mathcal{T},s}^{-1}(1)),H_{\mathcal{T},s},\sigma^*}$ agrees with the path from root to s where $H_{\mathcal{T},s}$ is defined as $\{(i, j) \mid y_{I(\rho_{\mathcal{T},s}^{-1}(1)),i,j} \in \rho_{\mathcal{T},s}^{-1}(1)\}$. Hence, $\text{L-dim}(\mathcal{C}, \mathcal{U}) \geq 2rk$.

14.5.3 Soundness

Next, we will prove the soundness of our reduction, stated more precisely below. Note that this lemma, together with completeness and output size properties we argue above, implies Theorem 14.4 with $\varepsilon = 0.001$.

Lemma 14.7 Let $(\mathcal{C}, \mathcal{U})$ be the output from the reduction in Algorithm 14.2 on input \mathcal{L} . If $\text{val}(\mathcal{L}) \leq 0.001$, then $\text{L-dim}(\mathcal{C}, \mathcal{U}) \leq 1.999rk$ with high probability.

Roughly speaking, the overall strategy of our proof of Lemma 14.7 is as follows:

1. First, we will argue that any subtree rooted at any test-selection element must be shallow (of depth $\leq 1.001rk$). This means that, if we have a depth- $1.999rk$ mistake tree, then the first $0.998rk$ levels must be assigned solely assignment elements.
2. We then argue that, in this $0.998rk$ -level mistake tree of assignment elements, we can always extract a leaf s such that the path from root to s indicates inclusion of a large non-repetitive set. In other words, the path to s can be decoded into a (partial) assignment for the Label Cover instance \mathcal{L} .
3. Let the leaf from the previous step be s and the non-repetitive set be $S^{\text{no-rep}}$. Our goal now is to show that the subtree rooted at s must have small depth. We start working toward this by showing that, with high probability, there are few tests that agree with $S^{\text{no-rep}}$. This is analogous to Part II of the VC Dimension proof.
4. With the previous steps in mind, we only need to argue that, when $|\mathcal{H}(S^{\text{no-rep}})|$ is small, the Littlestone's Dimension of all the concepts that contain $S^{\text{no-rep}}$ (i.e., $\text{L-dim}(\mathcal{C}[S^{\text{no-rep}}], \mathcal{U})$) is small. Thanks to Fact 2.3, it is enough for us to bound $\text{L-dim}(\mathcal{C}[S^{\text{no-rep}}], \mathcal{X})$ and $\text{L-dim}(\mathcal{C}[S^{\text{no-rep}}], \mathcal{Y})$ separately. For the former, our technique from the second step also gives us the desired bound; for the latter, we prove that $\text{L-dim}(\mathcal{C}[S^{\text{no-rep}}], \mathcal{Y})$ is small by designing an algorithm that provides correct predictions on a constant fraction of the elements in \mathcal{Y} .

Let us now proceed to the details of the proofs.

14.5.3.1 Part I: Subtree of a Test-Selection Assignment Is Shallow

Lemma 14.8 For any $y_{I,i,j} \in \mathcal{Y}$, $\text{L-dim}(\mathcal{C}[\{y_{I,i,j}\}], \mathcal{U}) \leq rk + (4|E|\ell/r) \log |\Sigma| \leq 1.001rk$.

Note that the above lemma implies that, in any mistake tree, the depth of the subtree rooted at any vertex s assigned to some $y_{I,i,j} \in \mathcal{Y}$ is at most $1 + 1.001rk$. This is because every concept that agrees with the path from the root to s must be in $\mathcal{C}[\{y_{I,i,j}\}]$, which has depth at most $1.001rk$.

Proof Proof of Lemma 14.8. Consider any $C_{I',H,\sigma_{\tau(H)}} \in \mathcal{C}[\{y_{I,i,j}\}], \mathcal{U}$. Since $y_{I,i,j} \in C_{I',H,\sigma_{\tau(H)}}$, we have $I = I'$. Moreover, from Lemma 2.6, we know that $|\mathcal{N}_i(M_{\tau(H)}(i))| \leq 4|E|\ell/r^2$, which implies that $|T_{\tau(H)}| \leq 4|E|\ell/r$. This means that there are only at most $|\Sigma|^{4|E|\ell/r}$ choices of $\sigma_{\tau(H)}$. Combined with the fact that there

are only 2^{rk} choices of H , we have $|\mathcal{C}[\{y_{I,i,j}\}]| \leq 2^{rk} \cdot |\Sigma|^{4|E|\ell/r}$. Fact 2.2 then implies the lemma. \blacksquare

14.5.3.2 Part II: Every Deep Mistake Tree Contains a Non-Repetitive Set

The goal of this part of the proof is to show that, for mistake tree \mathcal{X}, \mathcal{C} of depth slightly less than rk , there exists a leaf s such that the corresponding path from root to s indicates an inclusion of a large non-repetitive set; in our notation, this means that we would like to identify a leaf s such that $IJ(\rho_{\mathcal{T},s}^{-1}(1))$ is large. Since we will also need a similar bound later in the proof, we will prove the following lemma, which is a generalization of the stated goal that works even for the concept class $\mathcal{C}[S^{\text{no-rep}}]$ for any non-repetitive $S^{\text{no-rep}}$. To get back the desired bound, we can simply set $S^{\text{no-rep}} = \emptyset$.

Lemma 14.9 For any non-repetitive set $S^{\text{no-rep}}$ and any depth- d mistake tree \mathcal{T} of $\mathcal{X}, \mathcal{C}[S^{\text{no-rep}}]$, there exists a leaf $s \in \{0, 1\}^d$ such that $|IJ(\rho_{\mathcal{T},s}^{-1}(1)) \setminus IJ(S^{\text{no-rep}})| \geq d - r$.

The proof of this lemma is a double counting argument where we count a specific class of leaves in two ways, which ultimately leads to the above bound. The leaves that we focus on are the leaves $s \in \{0, 1\}^d$ such that, for every (i, j) such that an (i, j) -assignment element appears in the path from root to s but not in $S^{\text{no-rep}}$, the first appearance of (i, j) -assignment element in the path is included. In other words, for every $(i, j) \in IJ(P_{\mathcal{T},s}) \setminus IJ(S^{\text{no-rep}})$, if we define $u_{i,j} \triangleq \inf_{s' \in \text{pre}(s), v_{\mathcal{T},s'} \in \mathcal{X}_{i,j}} |s'|$, then $s_{u_{i,j}}$ must be equal to 1. We call these leaves the *good* leaves. Denote the set of good leaves of \mathcal{T} by $\mathcal{G}_{\mathcal{T}, S^{\text{no-rep}}}$.

Our first way of counting is the following lemma. Informally, it asserts that different leaves agree with different sets $\tilde{H} \subseteq [r]$. This can be thought of as an analogue of Lemma 14.3 in our proof for the VC Dimension. Note that this lemma immediately gives an upper bound of 2^r on $|\mathcal{G}_{\mathcal{T}, S^{\text{no-rep}}}|$.

Lemma 14.10 For any depth- d mistake tree \mathcal{T} of $\mathcal{X}, \mathcal{C}[S^{\text{no-rep}}]$ and any different $s_1, s_2 \in \mathcal{G}_{\mathcal{T}, S^{\text{no-rep}}}$, if C_{I_1, H_1, σ_1} agrees with s_1 and C_{I_2, H_2, σ_2} agrees with s_2 for some $I_1, I_2, H_1, H_2, \sigma_1, \sigma_2$, then $\tau(H_1) \neq \tau(H_2)$.

Proof Suppose for the sake of contradiction that there exist $s_1 \neq s_2 \in \mathcal{G}_{\mathcal{T}, S^{\text{no-rep}}}, H_1, H_2, I_1, I_2, \sigma_1, \sigma_2$ such that C_{I_1, H_1, σ_1} and C_{I_2, H_2, σ_2} agree with s_1 and s_2 respectively, and $\tau(H_1) = \tau(H_2)$. Let s be the common ancestor of s_1, s_2 , i.e., s is the longest string in $\text{pre}(s_1) \cap \text{pre}(s_2)$. Assume w.l.o.g. that $(s_1)_{|s|+1} = 0$ and $(s_2)_{|s|+1} = 1$. Consider the node $v_{\mathcal{T},s}$ in tree \mathcal{T} where the paths to s_1, s_2 split; suppose that this is $x_{i, \sigma_i, j}$. Therefore $x_{i, \sigma_i, j} \in C_{I_2, H_2, \sigma_2} \setminus C_{I_1, H_1, \sigma_1}$.

We now argue that there is some $x_{i,\sigma'_i,j}$ (with the same i, j but a different assignment σ'_i) that is in both concepts, i.e., $x_{i,\sigma'_i,j} \in C_{I_2,H_2,\sigma_2} \cap C_{I_1,H_1,\sigma_1}$. We do this by considering two cases:

- If $(i, j) \in IJ(S^{\text{no-rep}})$, then there is $x_{i,\sigma'_i,j} \in S^{\text{no-rep}} \subseteq C_{I_1,H_1,\sigma_1}, C_{I_2,H_2,\sigma_2}$ for some $\sigma'_i \in \Sigma^{U_i}$.
- Suppose that $(i, j) \notin IJ(S^{\text{no-rep}})$. Since s_1 is a good leaf, there is some $t \in \text{pre}(s)$ such that $v_{\mathcal{T},t} = x_{i,\sigma'_i,j}$ for some $\sigma'_i \in \Sigma^{U_i}$ and t is included by the path (i.e., $s_{|t|+1} = 1$). This also implies that $x_{i,\sigma'_i,j}$ is in both C_{I_1,H_1,σ_1} and C_{I_2,H_2,σ_2} .

Now, since both $x_{i,\sigma_i,j}$ and $x_{i,\sigma'_i,j}$ are in the concept C_{I_2,H_2,σ_2} , we have $(i, j) \in I_2$ and

$$\sigma_i|_{\mathcal{N}_i(M_{\tau(H_1)})} = \sigma_2|_{\mathcal{N}_i(M_{\tau(H_1)})} = \sigma'_i|_{\mathcal{N}_i(M_{\tau(H_1)})}. \quad (14.4)$$

On the other hand, since C_{I_1,H_1,σ_1} contains $x_{i,\sigma'_i,j}$ but not $x_{i,\sigma_i,j}$, we have $(i, j) \in I_1$ and

$$\sigma_i|_{\mathcal{N}_i(M_{\tau(H_2)})} \neq \sigma_1|_{\mathcal{N}_i(M_{\tau(H_2)})} = \sigma'_i|_{\mathcal{N}_i(M_{\tau(H_2)})}, \quad (14.5)$$

which contradicts (14.4) since $\tau(H_1) = \tau(H_2)$. ■

Next, we will present another counting argument that gives a lower bound on the number of good leaves, which, together with Lemma 14.10, yields the desired bound.

Proof Proof of Lemma 14.9. For any depth- d mistake tree \mathcal{T} of $\mathcal{C}[S^{\text{no-rep}}]$, \mathcal{X} , let us consider the following procedure, which recursively assigns a weight λ_s to each node s in the tree. At the end of the procedure, all the weight will be propagated from the root to good leaves.

1. For every non-root node $s \in \{0, 1\}^{\geq 1}$, set $\lambda_s \leftarrow 0$. For root $s = \emptyset$, let $\lambda_\emptyset \leftarrow 2^d$.
2. While there is an internal node $s \in \{0, 1\}^{< d}$ such that $\lambda_s > 0$, do the following:
 1. Suppose that $v_s = x_{i,\sigma_i,j}$ for some $i \in [r]$, $\sigma_i \in \Sigma^{U_i}$ and $j \in [k]$.
 2. If so far no (i, j) -element has appeared in the path or in $S^{\text{no-rep}}$, i.e., $(i, j) \notin IJ(P_{\mathcal{T},s}) \cup IJ(S^{\text{no-rep}})$, then $\lambda_{s1} \leftarrow \lambda_s$. Otherwise, set $\lambda_{s0} = \lambda_{s1} = \lambda_s/2$.
 3. Set $\lambda_s \leftarrow 0$.

The following observations are immediate from the construction:

- The total of λ 's over all the tree, $\sum_{s \in \{0,1\}^{\leq d}} \lambda_d$, always remains 2^d .

- At the end of the procedure, for every $s \in \{0, 1\}^{\leq d}$, $\lambda_s \neq 0$ if and only if $s \in \mathcal{G}_{\mathcal{T}, S^{\text{no-rep}}}$.
- If $s \in \mathcal{G}_{\mathcal{T}, S^{\text{no-rep}}}$, then $\lambda_s = 2^{|IJ(\rho_{\mathcal{T},s}^{-1}(1)) \setminus IJ(S^{\text{no-rep}})|}$ at the end of the execution.

Note that the last observation comes from the fact that λ always get divided in half when moving down one level of the tree unless we encounter an (i, j) -assignment element for some i, j that has never appeared in the path or in $S^{\text{no-rep}}$ before. For any good leaf s , the set of such (i, j) is exactly the set $IJ(\rho_{\mathcal{T},s}^{-1}(1)) \setminus IJ(S^{\text{no-rep}})$.

As a result, we have $2^d = \sum_{s \in \mathcal{G}_{\mathcal{T}}} 2^{|IJ(\rho_{\mathcal{T},s}^{-1}(1)) \setminus IJ(S^{\text{no-rep}})|}$. Since Lemma 14.10 implies that $|\mathcal{G}_{\mathcal{T},s}| \leq 2^r$, we can conclude that there exists $s \in \mathcal{G}_{\mathcal{T}, S^{\text{no-rep}}}$ such that $|IJ(\rho_{\mathcal{T},s}^{-1}(1)) \setminus IJ(S^{\text{no-rep}})| \geq d - r$ as desired. ■

14.5.3.3 Part III: No Large Non-Repetitive Set Passes Many Tests

The main lemma of this subsection is the following, which is analogous to Lemma 14.4.

Lemma 14.11 If $\text{val}(\mathcal{L}) \leq 0.001$, then, with high probability, for every non-repetitive set $S^{\text{no-rep}}$ of size at least $0.99rk$, $|\mathcal{H}(S^{\text{no-rep}})| \leq 100n \log |\Sigma|$.

Proof For every $I \subseteq [r]$, let $U_I \triangleq \bigcup_{i \in I} U_i$. For every $\sigma_I \in \Sigma^{U_I}$ and every $\tilde{H} \subseteq \mathcal{Y}$, we say that (I, σ_I) passes \tilde{H} if σ_I does not violate any constraint in $T_{\tilde{H}}$. Note that this definition and the way the test is generated in the reduction is the same as that of the VC Dimension reduction. Hence, if we can apply Lemma 14.5 with $\delta = 0.99$, this implies the following: with high probability, for every $I \subseteq [r]$ of size at least $0.99r$ and every $\sigma_I \in \Sigma^{U_I}$, $|\mathcal{H}(I, \sigma_I)| \leq 100n \log |\Sigma|$. Conditioned on this event happening, we will show that, for every non-repetitive set $S^{\text{no-rep}}$ of size at least $0.99rk$, $|\mathcal{H}(S^{\text{no-rep}})| \leq 100n \log |\Sigma|$.

Consider any non-repetitive set $S^{\text{no-rep}}$ of size $0.99rk$. Let $\sigma_{I(S^{\text{no-rep}})}$ be an assignment on $U_{I(S^{\text{no-rep}})}$ such that, for each $i \in I(S^{\text{no-rep}})$, we pick one $x_{i, \sigma_i, j} \in S^{\text{no-rep}}$ (if there is more than one such x , pick one arbitrarily) and let $\sigma_{I(S^{\text{no-rep}})}|_{U_i} = \sigma_i$. It is obvious that $\mathcal{H}(S^{\text{no-rep}}) \subseteq \mathcal{H}(I(S^{\text{no-rep}}), \sigma_{I(S^{\text{no-rep}})})$. Since $S^{\text{no-rep}}$ is non-repetitive and of size at least $0.99rk$, we have $|I(S^{\text{no-rep}})| \geq 0.99r$, which means that $|\mathcal{H}(I(S^{\text{no-rep}}), \sigma_{I(S^{\text{no-rep}})})| \leq 100n \log |\Sigma|$ as desired. ■

14.5.3.4 Part IV: A Subtree Containing $S^{\text{no-rep}}$ Must Be Shallow

In this part, we will show that, if we restrict ourselves to only concepts that contain some non-repetitive set $S^{\text{no-rep}}$ that passes few tests, then the Littlestone's Dimension of this restricted concept class is small. Therefore when we build a tree for the

whole concept class \mathcal{C} , if a path from root to some node indicates an inclusion of a non-repetitive set that passes few tests, then the subtree rooted at this node must be shallow.

Lemma 14.12 For every non-repetitive set $S^{\text{no-rep}}$,

$$\text{L-dim}(\mathcal{C}[S^{\text{no-rep}}], \mathcal{U}) \leq 1.75rk - |S^{\text{no-rep}}| + r + 1000k\sqrt{r} \log(|\mathcal{H}(S^{\text{no-rep}})| + 1).$$

We prove the above lemma by bounding $\text{L-dim}(\mathcal{C}[S^{\text{no-rep}}], \mathcal{X})$ and $\text{L-dim}(\mathcal{C}[S^{\text{no-rep}}], \mathcal{Y})$ separately, and combining them via Fact 2.3. First, we can bound $\text{L-dim}(\mathcal{C}[S^{\text{no-rep}}], \mathcal{X})$ easily by applying Lemma 14.9 coupled with the fact that $|IJ(S^{\text{no-rep}})| = |S^{\text{no-rep}}|$ for every non-repetitive $S^{\text{no-rep}}$. This gives the following corollary.

Corollary 14.1 For every non-repetitive set $S^{\text{no-rep}}$,

$$\text{L-dim}(\mathcal{C}[S^{\text{no-rep}}], \mathcal{X}) \leq rk - |S^{\text{no-rep}}| + r.$$

We will next prove the following bound on $\text{L-dim}(\mathcal{C}[S^{\text{no-rep}}], \mathcal{Y})$. Note that Corollary 14.1, Lemma 14.13, and Fact 2.3 immediately imply Lemma 14.12.

Lemma 14.13 For every non-repetitive set $S^{\text{no-rep}}$,

$$\text{L-dim}(\mathcal{C}[S^{\text{no-rep}}], \mathcal{Y}) \leq 0.75rk + 500k\sqrt{r} \log(|\mathcal{H}(S^{\text{no-rep}})| + 1).$$

The overall outline of the proof of Lemma 14.13 is that we will design a prediction algorithm whose mistake bound is at most $0.75rk + 1000k\sqrt{r} \log |\mathcal{H}(S^{\text{no-rep}})|$. Once we design this algorithm, Theorem 2.3 immediately implies Lemma 14.13. To define our algorithm, we will need the following lemma, which is a general statement that says that, for a small collection of H 's, there is some $\tilde{H}^* \subset [r]$ that agrees with almost half of every H in the collection.

Lemma 14.14 Let $\mathcal{H} \subseteq \mathcal{P}([r])$ be any collection of subsets of $[r]$; there exists $\tilde{H}^* \subseteq [r]$ such that, for every $\tilde{H} \in \mathcal{H}$, $|\tilde{H}^* \Delta \tilde{H}| \leq 0.5r + 1000\sqrt{r} \log(|\mathcal{H}| + 1)$, where Δ denotes the symmetric difference between two sets.

Proof We use a simple probabilistic method to prove this lemma. Let \tilde{H}^r be a random subset of $[r]$ (i.e., each $i \in [r]$ is included independently with probability 0.5). We will show that, with non-zero probability, $|\tilde{H}^r \Delta \tilde{H}| \leq 0.5r + 1000\sqrt{r} \log(|\mathcal{H}| + 1)$ for all $\tilde{H} \in \mathcal{H}$, which immediately implies that a desired \tilde{H}^* exists.

Fix $\tilde{H} \in \mathcal{H}$. Observe that $|\tilde{H}^r \Delta \tilde{H}|$ can be written as $\sum_{i \in [r]} \mathbf{1}[i \in (\tilde{H}^r \Delta \tilde{H})]$. For each i , $\mathbf{1}[i \in (\tilde{H}^r \Delta \tilde{H})]$ is a 0, 1 random variable with mean 0.5 independent of other $i' \in [r]$. Applying Chernoff bound here yields

$$\Pr\left[|\tilde{H}^r \Delta \tilde{H}| > 0.5r + 1000\sqrt{r} \log(|c\mathcal{H}|+1)\right] \leq 2^{-\log^2(|\mathcal{H}|+1)} \leq \frac{1}{|\mathcal{H}|+1}.$$

Hence, by union bound, we have

$$\Pr\left[\exists \tilde{H} \in \mathcal{H}, |\tilde{H}^r \Delta \tilde{H}| > 0.5r + 1000\sqrt{r} \log(|c\mathcal{H}|+1)\right] \leq \frac{|\mathcal{H}|}{|\mathcal{H}|+1} < 1.$$

In other words, $|\tilde{H}^r \Delta \tilde{H}| \leq 0.5r + 1000\sqrt{r} \log(|\mathcal{H}|+1)$ for all $\tilde{H} \in \mathcal{H}$ with non-zero probability as desired. \blacksquare

We also need the following observation, which is an analogue of Observation 14.1 in the VC Dimension proof; it follows immediately from definition of $\mathcal{H}(S)$.

Observation 14.2 If a non-repetitive set $S^{\text{no-rep}}$ is a subset of some concept $C_{I,H,\sigma_{\tau(H)}}$, then $\tau(H) \in \mathcal{H}(S^{\text{no-rep}})$.

With Lemma 14.14 and Observation 14.2 in place, we are now ready to prove Lemma 14.13.

Proof Proof of Lemma 14.13. Let $\tilde{H}^* \subseteq [r]$ be the set guaranteed by applying Lemma 14.14 with $\mathcal{H} = \mathcal{H}(S^{\text{no-rep}})$. Let $H^* \triangleq \tilde{H}^* \times [k]$.

Our prediction algorithm will be very simple: it always predicts according to H^* ; i.e., on an input² $y \in \mathcal{Y}$, it outputs $\mathbf{1}[y \in H^*]$. Consider any sequence $(y_1, h_1), \dots, (y_w, h_w)$ that agrees with a concept $C_{I,H,\sigma_{\tau(H)}} \in \mathcal{C}[S^{\text{no-rep}}]$. Observe that the number of incorrect predictions of our algorithm is at most $|H^* \Delta H|$.

Since $C_{I,H,\sigma_{\tau(H)}} \in \mathcal{C}[S^{\text{no-rep}}]$, Observation 14.2 implies that $\tau(H) \in \mathcal{H}(S^{\text{no-rep}})$. This means that $|\tau(H) \Delta \tilde{H}^*| \leq 0.5r + 1000\sqrt{r} \log(|\mathcal{H}|+1)$. Now, let us consider each $i \in [r] \setminus (\tau(H) \Delta \tilde{H}^*)$. Suppose that $i \in \tau(H) \cap \tilde{H}^*$. Since $i \in \tau(H)$, at least $k/2$ elements of \mathcal{Y}_i are in H and, since $i \in \tilde{H}^*$, we have $\mathcal{Y}_i \subseteq H^*$. This implies that $|(H^* \Delta H) \cap \mathcal{Y}_i| \leq k/2$. A similar bound can also be derived when $i \notin \tau(H) \cap \tilde{H}^*$. As a result, we have

2. We assume w.l.o.g. that input elements are distinct; if an element appears multiple times, we know the correct answer from its first appearance and can always correctly predict it afterward.

$$\begin{aligned}
|H^* \Delta H| &= \sum_{i \in [r]} |(H^* \Delta H) \cap Y_i| \\
&= \sum_{i \in \tau(H) \Delta \tilde{H}^*} |(H^* \Delta H) \cap Y_i| + \sum_{i \in [r] \setminus (\tau(H) \Delta \tilde{H}^*)} |(H^* \Delta H) \cap Y_i| \\
&\leq (|\tau(H) \Delta \tilde{H}^*|)(k) + (r - |\tau(H) \Delta \tilde{H}^*|)(k/2) \\
&\leq 0.75rk + 500k\sqrt{r} \log(|\mathcal{H}| + 1),
\end{aligned}$$

concluding our proof of Lemma 14.13. ■

14.5.3.5 Putting Things Together

Proof Proof of Lemma 14.7. Suppose for the sake of contradiction that $\text{val}(\mathcal{L}) \leq 0.001$ but $\text{L-dim}(\mathcal{C}, \mathcal{U}) \geq 1.999rk$. Consider any depth- $1.999rk$ mistake tree \mathcal{T} of \mathcal{C}, \mathcal{U} . From Lemma 14.8, we know that no test-selection element is assigned to any node in the first $1.999rk - 1.001rk - 1 \geq 0.997rk$ levels. In other words, the tree induced by the first $0.997rk$ levels is simply a mistake tree of \mathcal{C}, \mathcal{X} . By Lemma 14.9 with $S^{\text{no-rep}} = \emptyset$, there exists $s \in \{0, 1\}^{0.997rk}$ such that $|IJ(\rho_{\mathcal{T},s}^{-1}(1))| \geq 0.997rk - r \geq 0.996rk$.

Since $|IJ(\rho_{\mathcal{T},s}^{-1}(1))| \geq 0.996rk$, there exists a non-repetitive set $S^{\text{no-rep}} \subseteq \rho_{\mathcal{T},s}^{-1}(1)$ of size $0.996rk$. Consider the subtree rooted at s . This is a mistake tree of $\mathcal{C}[\rho_{\mathcal{T},s}], \mathcal{U}$ of depth $1.002rk$. Since $S^{\text{no-rep}} \subseteq \rho_{\mathcal{T},s}^{-1}(1)$, we have $\mathcal{C}[\rho_{\mathcal{T},s}] \subseteq \mathcal{C}[S^{\text{no-rep}}]$. However, this implies

$$\begin{aligned}
1.002rk &\leq \text{L-dim}(\mathcal{C}[\rho_{\mathcal{T},s}], \mathcal{U}) \\
&\leq \text{L-dim}(\mathcal{C}[S^{\text{no-rep}}], \mathcal{U}) \\
&\text{(from Lemma 14.12)} \leq 1.75rk - 0.996rk + r + 100k\sqrt{r} \log(|\mathcal{H}(S^{\text{no-rep}})| + 1) \\
&\text{(from Lemma 14.11)} \leq 0.754rk + r + 100k\sqrt{r} \log(100n \log |\Sigma| + 1) \\
&= 0.754rk + o(rk),
\end{aligned}$$

which is a contradiction when r is sufficiently large. ■

14.6 Quasi-polynomial Algorithm for Littlestone's Dimension

Theorem 14.5 Quasi-polynomial time algorithm for Littlestone's dimension. Littlestone's Dimension can be computed (exactly) in time

$$O(|\mathcal{C}| \cdot (2|\mathcal{U}|)^{\text{L-dim}(\mathcal{C}, \mathcal{U})}) \leq |\mathcal{C}| \cdot |\mathcal{U}|^{O(\log |\mathcal{C}|)}.$$

Proof We prove by induction on $L\text{-dim}(\mathcal{C}, \mathcal{U})$. Assume by induction that the bound on the running time holds for all concept classes of smaller dimension. Enumerate over all possible roots $x \in \mathcal{U}$ for the mistake tree. For each root x , partition \mathcal{C} into $\mathcal{C}_{x=1}$, concepts that include x , and $\mathcal{C}_{x=0}$ (this takes time $O(|\mathcal{C}|)$). Now any mistake tree rooted at x has depth exactly $1 + \min\{L\text{-dim}(\mathcal{C}_{x=0}, \mathcal{U}), L\text{-dim}(\mathcal{C}_{x=1}, \mathcal{U})\}$. We attempt to compute both $L\text{-dim}(\mathcal{C}_{x=0}, \mathcal{U})$, $L\text{-dim}(\mathcal{C}_{x=1}, \mathcal{U})$ until one of the computations terminates. Since one of $\mathcal{C}_{x=0}, \mathcal{C}_{x=1}$ has a lower Littlestone's Dimension, it follows by our induction hypothesis that we can compute it in time $O(|\mathcal{C}| \cdot (2|\mathcal{U}|)^{L\text{-dim}(\mathcal{C}, \mathcal{U})-1})$. Multiplying by 2 because we're computing for two classes and by $|\mathcal{U}|$ because we're performing this computation for every root x , the total upper bound on the running time follows. ■

15 Signaling

Many classical questions in economics involve extracting information from strategic agents. Lately, there has been growing interest within algorithmic game theory in *signaling*: the study of how to *reveal information* to strategic agents (see, e.g., [Cheng et al. 2015, Dughmi 2014, Dughmi et al. 2013, Emek et al. 2014, Miltersen and Sheffet 2012], and references therein). Signaling has been studied in many interesting economic and game theoretic settings. Among them, ZERO-SUM SIGNALING proposed by Dughmi [2014] stands out as a canonical problem that cleanly captures the computational nature of signaling. In particular, focusing on zero-sum games clears away issues of equilibrium selection and computational tractability of finding an equilibrium.

Definition 15.1 ZERO-SUM SIGNALING [Dughmi 2014]. Alice and Bob play a Bayesian zero-sum game where the payoff matrix M is drawn from a publicly known prior. The signaler Sam privately observes the state of nature (i.e., the payoff matrix), and then publicly broadcasts a signal $\varphi(M)$ to both Alice and Bob. Alice and Bob Bayesian-update their priors according to $\varphi(M)$'s and play the Nash equilibrium of the expected game; but they receive payoffs according to the true M . Sam's goal is to design an efficient signaling scheme φ (a function from payoff matrices to strings) that maximizes Alice's expected payoff.

Dughmi's main result proves that assuming the hardness of the PLANTED CLIQUE problem, there is no additive FPTAS for ZERO-SUM SIGNALING. The main open question left by Dughmi [2014] is whether there exists an additive PTAS. Here we answer this question in the negative: we prove that assuming the Exponential Time Hypothesis (ETH), obtaining an additive- ϵ -approximation (for some constant $\epsilon > 0$) requires quasi-polynomial time ($n^{\tilde{\Omega}(\lg n)}$). This result is tight thanks to a recent quasi-polynomial ($n^{\frac{\lg n}{\text{poly}(\epsilon)}}$) time algorithm by Cheng et al. [2015]. Another important advantage of our result is that it replaces the hardness of PLANTED CLIQUE with a more believable worst-case hardness assumption (see, e.g., the discussion in Braverman et al. [2015]).

Theorem 15.1 There exists a constant $\epsilon > 0$, such that assuming ETH, approximating ZERO-SUM SIGNALING with payoffs in $[-1, 1]$ to within an additive ϵ requires time $n^{\tilde{\Omega}(\lg n)}$.

NP-Hard Variants

This chapter is based on Rubinstein [2017b], where we also proved two NP-hardness of approximation results: first, we show that obtaining a multiplicative approximation is NP-hard. Then we introduce a new variant of the signaling problem where the signaler may lie, i.e., he commits to one signaling scheme, but uses a different scheme. In the lying variant, even an additive $1 - 2^{-\Omega(n)}$ -approximation is NP-hard.

Concurrent Work of Bhaskar, et al.

In independent concurrent work by Bhaskar et al. [2016], quasi-polynomial time hardness for additive approximation of ZERO-SUM SIGNALING was obtained assuming the hardness of the PLANTED CLIQUE problem (among other interesting results¹ about network routing games and security games). Although we are not aware of a formal reduction, hardness of PLANTED CLIQUE is a qualitatively stronger assumption than ETH in the sense that it requires average case instances to be hard. Hence, in this respect, our result is stronger.

15.1 Techniques

Our starting point for this reduction is “birthday repetition.” We reduce from a 2-ary constraint satisfaction problem (2-CSP) over n variables to a distribution over N zero-sum $N \times N$ games, with $N = 2^{\Theta(\sqrt{n})}$. Alice and Bob’s strategies correspond to assignments to tuples of \sqrt{n} variables. By the birthday paradox, the two \sqrt{n} -tuples chosen by Alice and Bob share a constraint with constant probability. If a constant fraction of the constraints are unsatisfiable, Alice’s payoff will suffer with constant probability. Assuming ETH, approximating the value of the CSP requires time $2^{\tilde{\Omega}(n)} = N^{\tilde{\Omega}(\lg N)}$.

The Challenge. The main difficulty is that once the signal is public, the zero-sum game is tractable. Thus we would like to force the signaling scheme to output a satisfying assignment. Furthermore, if the scheme would output partial assignments

1. For zero-sum games, Bhaskar et al. [2016] also rule out an additive FPTAS assuming $P \neq NP$. This result follows immediately from the NP-hardness for multiplicative approximation in Rubinstein [2017b].

on different states of nature (aka different zero-sum games in the support), it is not clear how to check consistency between different signals. Thus we would like each signal to contain an entire satisfying assignment. The optimal scheme may be very complicated and even require randomization, yet by an application of the Caratheodory Theorem the number of signals is, wlog, bounded by the number of states of nature [Dughmi 2014]. If the state of nature can be described using only $\lg N = \tilde{\Theta}(\sqrt{n})$ bits,² how can we force the scheme to output an entire assignment?

To overcome this obstacle, we let the state of nature contain a partial assignment to a random \sqrt{n} -tuple of variables. We then check the consistency of Alice's assignment with nature's assignment, Bob's assignment with nature's assignment, and Alice and Bob's assignments with each other; let $\tau^{A,Z}$, $\tau^{B,Z}$, $\tau^{A,B}$ denote the outcomes of those consistency checks, respectively. Alice's payoff is given by:

$$U = \delta\tau^{A,Z} - \delta^2\tau^{B,Z} + \delta^3\tau^{A,B}$$

for some small constant $\delta \in (0, 1)$. Now, both Alice and Bob want to maximize their chances of being consistent with nature's partial assignment, and the signaling scheme gains by maximizing $\tau^{A,B}$.

Of course, if nature outputs a random assignment, we have no reason to expect that it can be completed to a full satisfying assignment. Instead, the state of nature consists of N assignments, and the signaling scheme helps Alice and Bob play with the assignment that can be completed.

15.2 Near-Optimal Signaling Is Hard

Theorem 15.2 There exists a constant $\epsilon > 0$ such that, assuming ETH, approximating ZERO-SUM SIGNALING with payoffs in $[-1, 1]$ to within an additive ϵ requires time $n^{\tilde{\Omega}(\lg n)}$.

Construction Overview

Our reduction begins with a 2CSP ψ over n variables from alphabet Σ . We partition the variables into n/k disjoint subsets $\{S_1, \dots, S_{n/k}\}$, each of size at most $2k$ for $k = \sqrt{n}$ such that every two subsets share at most a constant number of constraints.

Nature chooses a random subset S_i from the partition, a random assignment $\vec{u} \in \Sigma^{2k}$ to the variables in S_i , and an auxiliary vector $\hat{b} \in \{0, 1\}^{\Sigma \times [2k]}$. As mentioned earlier in this chapter, \vec{u} may not correspond to any satisfying assignment. Alice and

2. In other words, N , the final size of the reduction, is an upper bound on the number of states of nature.

Bob participate in one of $|\Sigma|^{2k}$ subgames; for each $\vec{v} \in \Sigma^{2k}$, there is a corresponding subgame where all the assignments are XOR-ed with \vec{v} . The goal of the auxiliary vector \hat{b} is to force Alice and Bob to participate in the right subgame, i.e., the one where the XOR of \vec{v} and \vec{u} can be completed to a full satisfying assignment. In particular, the optimum signaling scheme reveals partial information about \hat{b} in a way that guides Alice and Bob to participate in the right subgame. The scheme also outputs the full satisfying assignment, but reveals no information about the subset S_i chosen by nature.

Each player has $(|\Sigma|^{2k} \times 2) \times (n/k \times \binom{n/k}{n/2k}) \times |\Sigma|^{2k} = 2^{\Theta(\sqrt{n})}$ strategies. The first $|\Sigma|^{2k}$ strategies correspond to a Σ -ary vector \vec{v} that the scheme will choose after observing the random input. The signaling scheme forces both players to play (w.h.p.) the strategy corresponding to \vec{v} by controlling the information that corresponds to the next 2 strategies. Namely, for each $\vec{v}' \in \Sigma^{2k}$, there is a random bit $b(\vec{v}')$ such that each player receives a payoff of 1 if they play $(\vec{v}', b(\vec{v}'))$ and 0 for $(\vec{v}', 1 - b(\vec{v}'))$. The b 's are part of the state of nature, and the optimal signaling scheme will reveal only the bit corresponding to the special \vec{v} . Since there are $|\Sigma|^{2k}$ bits, nature cannot choose them independently, as that would require $2^{|\Sigma|^{2k}}$ states of nature. Instead we construct a pairwise independent distribution.

The next n/k strategies correspond to the choice of a subset S_i from the specified partition of variables. The $\binom{n/k}{n/2k}$ strategies that follow correspond to a gadget due to Althofer [1994] whereby each player forces the other player to randomize (approximately) uniformly over the choice of subset.

The last $|\Sigma|^{2k}$ strategies correspond to an assignment to S_i . The assignment to each S_i is XOR-ed entry-wise with \vec{v} . Then, the players are paid according to checks of consistency between their assignments, and a random assignment to a random S_i picked by nature. (The scheme chooses \vec{v} so that nature's random assignment is part of a globally satisfying assignment.) Each player wants to pick an assignment that passes the consistency check with nature's assignment. Alice also receives a small bonus if her assignment agrees with Bob's; thus her payoff is maximized when there exists a globally satisfying assignment.

See formal construction below, as well as summary in Table 15.1.

Formal Construction

Let ψ be a 2CSP- d over n variables from alphabet Σ , as guaranteed by Theorem 2.1. In particular, ETH implies that distinguishing between a completely satisfiable instance and $(1 - \eta)$ -satisfiable requires time $2^{\tilde{\Omega}(n)}$. By Lemma 2.5, we can (deterministically and efficiently) partition the variables into n/k subsets $\{S_1, \dots, S_{n/k}\}$ of size at most $2k = 2\sqrt{n}$, such that every two subsets share at most $8d^2k^2/n = O(1)$ constraints.

Table 15.1 Variables in proof of Theorem 15.2

Variable	Role in Reduction	Chosen by . . .
Nature		
$i \in [n/k]$	Specifies subset S_i	Uniformly at random
$\vec{u} \in \Sigma^{2k}$	Specifies (shifted) assignment for S_i	Uniformly at random
$\hat{b} \in \{0, 1\}^{\Sigma \times [2k]}$	Force Alice and Bob into right subgame	Uniformly at random
Alice^a (Expected Behavior)		
\vec{v}^A, c^A	Force Alice into right subgame	$c^A = b(\vec{v}^A)$
$j^A \in [n/k]$	Specifies subset S_{j^A}	Uniformly at random
$T^A \subset [n/k]$	Force Bob to pick j^B at random	Random subset of size $n/2k$
$\vec{w}^A \in \Sigma^{2k}$	Assignment for S_{j^A}	$\vec{\alpha}$ restricted to S_{j^A}
Signaler (Completeness)		
$\vec{\alpha} \in \Sigma^n$	Assignment to entire 2CSP	Satisfying assignment
$\vec{v} \in \Sigma^{2k}$	Correct shift to Nature's assignment	$\vec{v} \oplus_{\Sigma} \vec{u}$ is $\vec{\alpha}$ restricted to S_i
$b(\vec{v}) \in \{0, 1\}$	Force Alice and Bob into \vec{v} -subgame	Extracted from \hat{b}

a. Bob's variables are defined analogously to Alice's.

States of Nature. Nature chooses a state $(\hat{b}, i, \vec{u}) \in \{0, 1\}^{\Sigma \times [2k]} \times [n/k] \times \Sigma^{2k}$ uniformly at random. For each \vec{v} , $b(\vec{v})$ is the XOR of bits from \hat{b} that correspond to entries of \vec{v} :

$$\forall \vec{v} \in \Sigma^{2k} \quad b(\vec{v}) \triangleq \left(\bigoplus_{(\sigma, \ell): [\vec{v}]_{\ell} = \sigma} [\hat{b}]_{(\sigma, \ell)} \right).$$

Notice that the $b(\vec{v})$'s are pairwise independent and each marginal distribution is uniform over $\{0, 1\}$.

Strategies. Alice and Bob each choose a strategy $(\vec{v}, c, j, T, \vec{w}) \in \Sigma^{2k} \times \{0, 1\} \times [n/k] \times \binom{[n/k]}{[n/2k]} \times \Sigma^{2k}$. We use \vec{v}^A, c^A , etc. to denote the strategy Alice plays, and similarly \vec{v}^B, c^B , etc. for Bob. For $\sigma, \sigma' \in \Sigma$, we denote $\sigma \oplus_{\Sigma} \sigma' \triangleq \sigma + \sigma' \pmod{|\Sigma|}$, and for vectors $\vec{v}, \vec{v}' \in \Sigma^{2k}$, we let $\vec{v} \oplus_{\Sigma} \vec{v}' \in \Sigma^{2k}$ denote the entry-wise \oplus_{Σ} .

Payoffs. Consider state of nature (\hat{b}, i, \vec{u}) and players' strategies $(\vec{v}^A, c^A, j^A, T^A, \vec{w}^A)$ and $(\vec{v}^B, c^B, j^B, T^B, \vec{w}^B)$.

When $\vec{v}^A = \vec{v}^B = \vec{v}$, we set $\tau^{A,Z} = 1$ if assignments \vec{w}^A and $(\vec{v} \oplus_{\Sigma} \vec{u})$ to subsets S_{j^A} and S_i , respectively, satisfy all the constraints in ψ that are determined by $(S_i \cup S_{j^A})$, and $\tau^{A,Z} = 0$ otherwise. Similarly, $\tau^{B,Z} = 1$ iff \vec{w}^B and $(\vec{v} \oplus_{\Sigma} \vec{u})$ satisfy

the corresponding constraints in ψ ; and $\tau^{A,B}$ checks \vec{w}^A and \vec{w}^B . When $\vec{v}^A \neq \vec{v}^B$, we set $\tau^{A,Z} = \tau^{B,Z} = \tau^{A,B} = 0$.

We decompose Alice's payoff as:

$$U^A \triangleq U_b^A + U_{\text{Althofer}}^A + U_\psi^A,$$

where

$$U_b^A \triangleq \mathbf{1}\{c^A = b(\vec{v}^A)\} - \mathbf{1}\{c^B = b(\vec{v}^B)\},$$

$$U_{\text{Althofer}}^A \triangleq \mathbf{1}\{j^B \in T^A\} - \mathbf{1}\{j^A \in T^B\},$$

and

$$U_\psi^A \triangleq \delta\tau^{A,Z} - \delta^2\tau^{B,Z} + \delta^3\tau^{A,B}, \quad (15.1)$$

for a sufficiently small constant $0 < \delta \ll \sqrt{\eta}$.

Completeness

Lemma 15.1 If ψ is satisfiable, there exists a signaling scheme and a mixed strategy for Alice that guarantees expected payoff $\delta - \delta^2 + \delta^3$.

Proof Fix a satisfying assignment $\vec{\alpha} \in \Sigma^n$. Given state of nature (\hat{b}, i, \vec{u}) , let \vec{v} be such that $(\vec{v} \oplus_\Sigma \vec{u}) = [\vec{\alpha}]_{\mathcal{S}_i}$. The scheme outputs the signal $(\vec{v}, b(\vec{v}), \vec{\alpha})$. Alice's mixed strategy sets $(\vec{v}^A, c^A) = (\vec{v}, b(\vec{v}))$, picks j^A and T^A uniformly at random, and sets $\vec{w}^A = [\vec{\alpha}]_{\mathcal{S}_{j^A}}$.

Because Bob has no information about $b(\vec{v}')$ for any $\vec{v}' \neq \vec{v}$, he has probability $1/2$ of losing whenever he picks $\vec{v}^B \neq \vec{v}$, i.e., $\mathbb{E}[U_b^A] \geq \frac{1}{2} \Pr[\vec{v}^B \neq \vec{v}]$. Furthermore, because Alice chooses T^A and j^A uniformly, $\mathbb{E}[U_{\text{Althofer}}^A] = 0$.

Since $\vec{\alpha}$ completely satisfies ψ , we have that $\tau^{A,Z} = 1$ as long as $\vec{v}^B = \vec{v}$ (regardless of the rest of Bob's strategy). Bob's goal is thus to maximize $\mathbb{E}[\delta^2\tau^{B,Z} - \delta^3\tau^{A,B}]$. Notice that \vec{w}^A and $(\vec{v} \oplus_\Sigma \vec{u})$ are two satisfying partial assignments to uniformly random subsets from the partition. In particular, they are both drawn from the same distribution, so we have that for any mixed strategy that Bob plays, $\mathbb{E}[\tau^{B,Z}] = \mathbb{E}[\tau^{A,B}]$. Therefore Alice's payoff is at least

$$(\delta - \delta^2 + \delta^3) \Pr[\vec{v}^B = \vec{v}] + \frac{1}{2} \Pr[\vec{v}^B \neq \vec{v}] \geq \delta - \delta^2 + \delta^3. \quad \blacksquare$$

Soundness

Lemma 15.2 If at most a $(1 - \eta)$ -fraction of the constraints are satisfiable, Alice's maxmin payoff is at most $\delta - \delta^2 + (1 - \Omega_\eta(1))\delta^3$, for any signaling scheme.

Proof Fix any mixed strategy by Alice; we show that Bob can guarantee a payoff of at least $-(\delta - \delta^2 + (1 - \Omega_\eta(1))\delta^3)$. On any signal, Bob chooses (\vec{v}^B, c^B) from the same distribution that Alice uses for (\vec{v}^A, c^A) . He chooses j^B uniformly, and picks T^B so as to minimize $\mathbb{E}[U_{\text{Althofer}}^A]$. Finally, for each j^B , he draws \vec{w}^B from the same marginal distribution that Alice uses for \vec{w}^A conditioning on $j^A = j^B$ (and uniformly at random if Alice never plays $j^A = j^B$). By symmetry, $\mathbb{E}[U_b^A] = 0$ and $\mathbb{E}[U_{\text{Althofer}}^A] \leq 0$.

In this paragraph, we use Althofer's gadget to argue that, wlog, Alice's distribution over the choice of j^A is approximately uniform. In Althofer's gadget, Alice can guarantee an (optimal) expected payoff of 0 by randomizing uniformly over her choice of j^A and T^A . By Lemma 2.7, if Alice's marginal distribution over the choice of j^A is $8\delta^2$ -far from uniform (in total variation distance), then Bob can guess that j^A is in some subset $T^B \in \binom{[n/k]}{n/2k}$ with advantage (over guessing at random) of at least $2\delta^2$. Therefore $\mathbb{E}[U_{\text{Althofer}}^A] \leq -2\delta^2$; but this would imply $\mathbb{E}[U^A] \leq -2\delta^2 + \mathbb{E}[U_\psi^A] \leq \delta - 2\delta^2 + \delta^3$. So henceforth we assume wlog that Alice's marginal distribution over the choice of j^A is $O(\delta^2)$ -close to uniform (in total variation distance).

Since Alice's marginal distribution over j^A is $O(\delta^2)$ -close to uniform, we have that Bob's distribution over (j^B, \vec{w}^B) is $O(\delta^2)$ -close to Alice's distribution over (j^A, \vec{w}^A) . Therefore $\mathbb{E}[\tau^{B,Z}] \geq \mathbb{E}[\tau^{A,Z}] - O(\delta^2)$, and so we also get:

$$\mathbb{E}[U^A] \leq \mathbb{E}[U_\psi^A] \leq \delta - \delta^2 + \delta^3 \mathbb{E}[\tau^{A,B}] + O(\delta^4). \quad (15.2)$$

Bounding $\mathbb{E}[\tau^{A,B}]$. To complete the proof, it remains to show an upper bound on $\mathbb{E}[\tau^{A,B}]$. In particular, notice that it suffices to bound the probability that Alice's and Bob's induced assignments agree. Intuitively, if they gave assignments to uniformly random (and independent) subsets of variables, the probability that their assignments agree cannot be much higher than the value of the 2CSP; below we formalize this intuition.

By the premise, any assignment to all variables violates at least an η -fraction of the constraints. In particular, this is true in expectation for assignments drawn according to Alice's and Bob's mixed strategy. This is a bit subtle: in general, it is possible that Alice's assignment alone doesn't satisfy many constraints and neither does Bob's, but when we check constraints between Alice's and Bob's assignments everything is satisfied (for example, think of the 3-Coloring CSP, where Alice colors all her vertices blue, and Bob colors all his vertices red). Fortunately, this subtlety is irrelevant for our construction, since we explicitly defined Bob's mixed strategy so that conditioned on each set S_j of variables, Alice and Bob have the same distribution over assignments.

The expected number of violations between pairs directly depends on the value of the 2CSP. To bound the *probability* of observing at least one violation, recall that every pair of subsets shares at most a constant number of constraints, so this probability is within a constant factor of the expected number of violations. In particular, an $\Omega(\eta)$ -fraction of the pairs of assignments chosen by Alice and Bob violate ψ .

Finally, Alice doesn't choose j^A uniformly at random; but her distribution is $O(\delta^2)$ -close to uniform. Therefore, we have $\mathbb{E}[\tau^{A,B}] \leq 1 - \Omega(\eta) + O(\delta^2)$. Plugging into (15.2) completes the proof. ■



PART

**APPROXIMATE
NASH EQUILIBRIUM**

2-Player Approximate Nash Equilibrium

For the past decade, the central open problem in equilibrium computation has been whether two-player Nash equilibrium admits a PTAS. We had good reasons to be hopeful: there was a series of improved approximation ratios [Bosse et al. 2010; Daskalakis et al. 2007, 2009b; Kontogiannis et al. 2009; Tsaknakis and Spirakis 2008] and several approximation schemes for special cases [Alon et al. 2013, Barman 2015, Daskalakis and Papadimitriou 2009, Kannan and Theobald 2007]. Yet most interesting are two inefficient algorithms for two-player Nash equilibrium:

- the classic Lemke-Howson algorithm [Lemke and Howson 1964] finds an exact Nash equilibrium in exponential time; and
- a simple algorithm by Lipton et al. [2003] finds an ϵ -approximate Nash equilibrium in time $n^{O(\log n)}$.

Although the Lemke-Howson algorithm takes exponential time, it has a special structure that places the problem inside the complexity class PPAD [Papadimitriou 1994]. Proving hardness for problems in PPAD is notoriously challenging because they are *total*, i.e., they always have a solution, so the standard techniques from NP-hardness do not apply. By now, though, we know that exponential and polynomial approximations for two-player Nash equilibrium are PPAD-complete [Chen et al. 2009b, Daskalakis et al. 2009a]. However, ϵ -approximation for two-player Nash equilibrium is unlikely to have the same fate: otherwise, the quasi-polynomial algorithm of Lipton et al. [2003] would refute the Exponential Time Hypothesis for PPAD. Thus the strongest hardness result we can hope to prove (given our current understanding of complexity¹) is a quasi-polynomial hardness that sits inside PPAD:

1. Given this current understanding, refuting ETH for PPAD seems unlikely: there are matching black-box lower bounds [Beame et al. 1998, Hirsch et al. 1989]. Recall that the NP-analogue ETH

Theorem 16.1 Theorem 1.1 restated. There exists a constant $\epsilon > 0$ such that, assuming ETH for PPAD, finding an ϵ -approximate Nash equilibrium in a two-player $n \times n$ game requires time $T(n) = n^{\log^{1-o(1)} n}$.

Additional Related Work

Previous attempts to show lower bounds for approximate Nash in two-player games have mostly focused on limited models of computation [Daskalakis and Papadimitriou 2009] and lower bounding the support required for obtaining approximate equilibria [Althofer 1994; Anbalagan et al. 2013, 2015; Feder et al. 2007] (in contrast, Lipton et al.'s algorithm runs in quasi-polynomial time because there exist approximate equilibria with support size at most $O(\frac{\log n}{\epsilon^2})$).

Best Nash Equilibrium. Hazan and Krauthgamer [2011] showed that finding an ϵ -approximate Nash equilibrium with ϵ -optimal welfare is as hard as the PLANTED CLIQUE problem; Austrin et al. [2013] later showed that the optimal-welfare constraint can be replaced by other decision problems. Braverman et al. [2015] recently showed that the hardness PLANTED CLIQUE can be replaced by the Exponential Time Hypothesis, the NP-analog of the ETH for PPAD we use here. (See also Theorem 11.1.)

Multiplicative Hardness of Approximation. Daskalakis [2013] and our recent work [Rubinfeld 2015] show that finding an ϵ -relative well-supported Nash equilibrium in two-player games is PPAD-hard. The case of ϵ -relative approximate Nash equilibrium is still open: our main theorem implies that it requires at least quasi-polynomial time, but it is not known whether it is PPAD-hard, or even if it requires a large support (see also discussion in Babichenko et al. [2016]).

Approximation Algorithms. The state of the art for games with arbitrary payoffs is ≈ 0.339 for two-player games due to Tsaknakis and Spirakis [2008] and $0.5 + \epsilon$ for polymatrix games due to Deligkas et al. [2017]. For two-player games, PTAS have been given for the special cases of constant rank games by Kannan and Theobald [2007], small-probability games by Daskalakis and Papadimitriou [2009], positive semi-definite games by Alon et al. [2013], and sparse games by Barman [2015].

[Impagliazzo et al. 2001] is widely used (e.g., [Aaronson et al. 2014, Braverman et al. 2015, Cygan et al. 2016, Karpinski and Schudy 2010, Lokshtanov et al. 2011] as well as the previous part of this book), often in stronger variants such as SETH [Calabro et al. 2009, Impagliazzo and Paturi 2001] and NSETH [Carmosino et al. 2016].

16.1 Technical Overview

Given an END-OF-A-LINE instance of size n , we construct a two-player $N \times N$ game for $N = 2^{n^{1/2+o(1)}}$ whose approximate equilibria correspond to solutions to the END-OF-A-LINE instance. Thus, assuming the “ETH for PPAD,” finding an approximate equilibrium requires time $2^n = N^{\log^{1-o(1)} N}$.

The main steps of the final construction are: (i) reducing END-OF-A-LINE to a new discrete problem that we call LOCAL END-OF-A-LINE; (ii) reducing LOCAL END-OF-A-LINE to a problem of finding an approximate Brouwer fixed point (this step uses our reduction from Section 4.2); (iii) reducing from Brouwer fixed point to finding an approximate Nash equilibrium in a multiplayer game over $n^{1/2+o(1)}$ players with $2^{n^{1/2+o(1)}}$ actions each; and (iv) reducing to the two-player game.

The main novelty in the reduction is the use of techniques such as error-correcting codes and probabilistically checkable proofs (PCPs) inside PPAD. In particular, the way we use PCPs in our proof is very unusual.

Constructing the First Gap: Showing Hardness of Euclidean BROUWER

The first step in all known PPAD-hardness results for (approximate) Nash equilibrium is reducing END-OF-A-LINE to the problem of finding an (approximate) Brouwer fixed point of a continuous, Lipschitz function $f: [0, 1]^n \rightarrow [0, 1]^n$. Even Theorem 4.1, which shows hardness of approximation in ℓ_∞ -norm, does not suffice for our purposes. In particular, it only implies that it is hard to find an \mathbf{x} such that $f(\mathbf{x})$ is approximately equal to \mathbf{x} on every coordinate. The first step in our proof is to strengthen this result to obtain hardness of approximation with respect to 2-norm (Theorem 4.2). Now, even finding an \mathbf{x} such that $f(\mathbf{x})$ is approximately equal to \mathbf{x} on most of the coordinates is already PPAD-hard.

Theorem 4.1 was obtained by adapting a construction due to Hirsch et al. [1989] via the use of *error-correcting code*. The first obstacle to using PCP-like techniques for problems in PPAD is totality: problems in PPAD always have a solution. For NP-hard problems, the PCP verifier expects the proof to be encoded in some error-correcting code. If the proof is far from any codeword, the verifier detects that (with high probability), and immediately rejects. For problems in PPAD (more generally, in TFNP),² this is always tricky because it is not clear what it means “to reject.” Hirsch et al.’s original construction has the following useful property: for the vast majority of \mathbf{x} ’s (in particular, all \mathbf{x} ’s far from the embedding of the paths), the displacement $f(\mathbf{x}) - \mathbf{x}$ is the same default displacement. Thus, when an \mathbf{x} is too

2. TFNP, a superclass of PPAD, is the class of total search problems whose solutions can be verified in polynomial time.

far from any codeword to faithfully decode it, we can simply apply the default displacement.

The Main Challenge: Locality

Our ultimate goal is to construct a two-player game that simulates the Brouwer function from Theorem 4.2. This is done via an *imitation gadget*: Alice’s mixed strategy induces a point $\mathbf{x}^{(A)} \in [0, 1]^n$; Bob’s strategy induces $\mathbf{x}^{(B)} \in [0, 1]^n$; Alice wants to minimize $\|\mathbf{x}^{(A)} - \mathbf{x}^{(B)}\|_2$, whereas Bob wants to minimize $\|f(\mathbf{x}^{(A)}) - \mathbf{x}^{(B)}\|_2$. Alice and Bob are both satisfied at a fixed point, where $\mathbf{x}^{(A)} = \mathbf{x}^{(B)} = f(\mathbf{x}^{(A)})$. (Recall also our construction of communicationally hard games for Theorem 3.1.)

The main obstacle is that we want to incentivize Bob to minimize $\|f(\mathbf{x}^{(A)}) - \mathbf{x}^{(B)}\|_2$ via *local constraints* (payoffs—each depends on one pure strategy), while $f(\mathbf{x}^{(A)})$ has a *global dependency* on Alice’s entire mixed strategy.

Our goal is thus to construct a hard Brouwer function that can be *locally computed*. How local does the computation need to be? In a game of size $2^{\sqrt{n}} \times 2^{\sqrt{n}}$, each strategy can faithfully store information about \sqrt{n} bits. Specifically, our construction will be $n^{1/2+o(1)}$ -local.

We haven’t yet defined exactly what it means for our construction to be “ $n^{1/2+o(1)}$ -local”; the exact formulation is quite cumbersome as the query access needs to be partly adaptive, robust to noise, etc. Eventually (Section 16.4), we formalize the “locality” of our Brouwer function via a statement about multiplayer games. On a high level, however, our goal is to show that for any $j \in \{1, \dots, n\}$, the j -th output $f_j(\mathbf{x})$ can be approximately computed, with high probability, by accessing \mathbf{x} at only $n^{1/2+o(1)}$ coordinates.

This is a good place to note that achieving any sense of “local computation” in our setting is surprising, even if we consider just the error-correcting encoding for our Brouwer function: in order to maintain constant relative distance, an average bit of the output must depend on a constant fraction of the input bits!

LOCAL END-OF-A-LINE

In order to introduce locality, we go back to the END-OF-A-LINE problem. “Wishful thinking”: imagine that we could replace the arbitrary predecessor and successor circuits in END-OF-A-LINE with NC^0 (constant depth and constant fan-in) circuits $S^{\text{LOCAL}}, P^{\text{LOCAL}} : \{0, 1\}^n \rightarrow \{0, 1\}^n$, so that each output bit only depends on a constant number of input bits. Imagine further that we had the guarantee that for each input, the outputs of $S^{\text{LOCAL}}, P^{\text{LOCAL}}$ differ from the input on just a constant number of bits. Additionally, it would be really nice if we had a succinct pointer that immediately

told us which bits are about to be replaced. (We later call this succinct pointer the *counter*, because it also cycles through its possible values in a fixed order.)

Suppose all our wishes came true, and furthermore the hard Brouwer function from Theorem 4.2 used a *linear* error-correcting code. Then, we could use the encoding of the counter, henceforth $C(u)$, to read only the bits that are about to be replaced, and the inputs that determine the new values of those bits. Thus, using only local access to a tiny fraction of the bits ($|C(u)| + O(1)$), we can construct a difference vector $u - S^{\text{LOCAL}}(u)$ (which is 0 almost everywhere). As we discussed above, the encodings $E(u)$, $E(S^{\text{LOCAL}}(u))$ must differ on a constant fraction of the bits—but because the code is linear, we can also locally construct the difference vector $E(u) - E(S^{\text{LOCAL}}(u)) = E(u - (S^{\text{LOCAL}}(u)))$. Given $E(u) - E(S^{\text{LOCAL}}(u))$, we can locally compute any bit of $E(S^{\text{LOCAL}}(u))$ by accessing only the corresponding bit of $E(u)$.

Back to reality: unfortunately we do not know of a reduction to such a restricted variant of END-OF-A-LINE. Surprisingly, we can almost do that. The problem LOCAL END-OF-A-LINE (formally defined in Section 16.2) satisfies all the guarantees defined above, is linear-time reducible from END-OF-A-LINE, but has one caveat: it is only defined on a strict subset V^{LOCAL} of the discrete hypercube ($V^{\text{LOCAL}} \subsetneq \{0, 1\}^n$). Verifying that a vertex belongs to V is quite easy: it can be done in AC^0 . Let us take a brief break to acknowledge this new insight about the canonical problem of PPAD:

Theorem 16.2 The predecessor and successor circuits of END-OF-A-LINE are, wlog, AC^0 circuits.

The class AC^0 is quite restricted, but the outputs of its circuits are not local functions of the inputs. Now, we want to represent u in a way that will make it possible to locally determine whether $u \in V^{\text{LOCAL}}$ or not. To this end we augment the linear error-correcting encoding $E(u)$ with a *probabilistically checkable proof* (PCP) $\pi(u)$ of the statement ($u \in V^{\text{LOCAL}}$).

Our Holographic Proof System

Some authors distinguish between PCPs and holographic proofs:³ a PCP verifier has unrestricted access to the instance, and queries the proof locally; whereas the holographic proof verifier has restricted, local access to both the proof and (an

3. In a nutshell, PCPs or holographic proofs are proofs that can be verified “locally” (with high probability) by reading only a small (random) portion of the proof; see, e.g., Arora and Barak [2009, Chapter 18] for many more details.

error-correcting encoding of) the instance. In this sense, what we actually want is a holographic proof.

We construct a holographic proof system with some very unusual properties. We are able to achieve these properties thanks to our modest locality desideratum: $n^{1/2+o(1)}$, as opposed to the typical $\text{polylog}(n)$ or $O(1)$. We highlight here a few of these properties; see Section 16.3 for details.

Local proof-construction. The most surprising property of our holographic proof system is that the proof $\pi(u)$ can be constructed from local access to the encoding $E(u)$. In particular, note that we can locally compute $E(S^{\text{LOCAL}}(u))$ because $E(\cdot)$ is linear—but $\pi(\cdot)$ is not. Once we obtain $E(S^{\text{LOCAL}}(u))$, we can use local proof construction to compute $\pi(S^{\text{LOCAL}}(u))$ locally.

Very low random-bit complexity. Our verifier is only allowed to use $(1/2 + o(1)) \log_2 n$ random bits; this is much lower even than the $\log_2 n$ bits necessary to choose one entry at random. In related works, similar random-bit complexity was achieved by bundling the entries together via “birthday repetition.” To some extent, something similar happens here, but our locality is already $n^{1/2+o(1)}$ so no bundling (or repetition) is necessary. To achieve nearly optimal random-bit complexity, we use λ -biased sets over large finite fields together with the Sampling Lemma of Ben-Sasson et al. [2003].

Tolerant verifier. Typically, a verifier must reject (with high probability) whenever the input is far from valid, but it is allowed to reject even if the input is off by only one bit. Our verifier, however, is *required* to accept (with high probability) inputs that are close to valid proofs. (This is related to the notion of “tolerant testing,” which was defined in Parnas et al. [2006] and discussed in Guruswami and Rudra [2005] for locally testable codes.)

Local decoding. We make explicit use of the property that our holographic proof system is also a locally decodable code. While the relations between PCPs and locally testable codes have been heavily explored (see, e.g., Goldreich’s survey [Goldreich 2010]), the connection to locally decodable codes is not as immediate. Nevertheless, related ideas of Locally Decode/Reject Codes [Moshkovitz and Raz 2010] and decodable PCP [Dinur and Harsha 2013] have been used before in order to facilitate composition of tests (our holographic proof system, in contrast, is essentially composition-free). Fortunately, as noted by Dinur and Harsha [2013], many constructions of PCPs are already implicitly locally decodable.

Robust everything. Ben-Sasson et al. [2006] introduce a notion of *robust soundness*, where on an invalid proof, the string read by the verifier must be far from any acceptable string. (Originally the requirement is far in expectation, but we want far with high probability.) Another way of looking at the same requirement is that even if a malicious prover *adaptively* changes a small fraction of the bits queried by the verifier, the test is still sound. In this sense, we require that all our guarantees, not just soundness, continue to hold (with high probability) even if a malicious entity adaptively changes a small fraction of the bits queried by the verifier.

How is local proof-construction possible? At a high level, our holographic proof system expects an encoding of u as a low-degree t -variate polynomial, and a few more low-degree t -variate polynomials that encode the proof of $u \in V^{\text{LOCAL}}$. (This is essentially the standard “arithmetization,” dating back at least to Babai and Fortnow [1991] and Shamir [1992], although our construction is most directly inspired by Polishchuk and Spielman [1994] and Spielman [1995].) In our actual proof, t is a small super-constant, e.g., $t \triangleq \sqrt{\log n}$; but for our exposition here, let us consider $t = 2$, i.e., we have bivariate polynomials.

The most interesting part of the proof verification is testing that a certain low-degree polynomial $\Psi: \mathcal{G}^2 \rightarrow \mathcal{G}$, for some finite field \mathcal{G} of size $|\mathcal{G}| = \Theta(n^{1/2+o(1)})$, is identically zero over all of \mathcal{F}^2 , for some subset $\mathcal{F} \subsetneq \mathcal{G}$ of cardinality $|\mathcal{F}| = |\mathcal{G}|/\text{polylog}(n)$. This can be done by expecting the prover to provide the following low-degree polynomials:

$$\begin{aligned}\Psi'(x, y) &\triangleq \sum_{f_i \in \mathcal{F}} \Psi(x, f_i) y^i \\ \Psi''(x, y) &\triangleq \sum_{f_j \in \mathcal{F}} \Psi'(f_j, y) x^j.\end{aligned}$$

Then, $\Psi''(x, y) = \sum_{f_i, f_j \in \mathcal{F}} \Psi(f_j, f_i) x^j y^i$ is the zero polynomial if and only if Ψ is indeed identically zero over all of \mathcal{F}^2 . $\Psi(x, y)$ can be computed by accessing $E(u)$ on just a constant number of entries. Thus, computing $\Psi'(x, y)$ requires $\Psi(x, f_i)$ for all $f_i \in \mathcal{F}$, so a total of $\Theta(n^{1/2+o(1)})$ queries to $E(u)$. However, computing even one entry of $\Psi''(\cdot)$ requires $\Omega(n)$ queries to $E(u)$. The crucial observation is that we don’t actually need the prover to provide Ψ'' . Instead, it suffices that the prover provide Ψ' , and the verifier checks that $\sum_{f_j \in \mathcal{F}} \Psi'(f_j, y) x^j = 0$ for sufficiently many (x, y) .

Putting It All Together Via Polymatrix Games

The above arguments suffice to construct a hard Brouwer function (in the sense of Theorem 4.2) that can be computed “ $n^{1/2+o(1)}$ -locally.” We formalize this statement in terms of approximate Nash equilibria in a polymatrix game.

Definition 16.1 Polymatrix games. In a polymatrix game, each pair of players simultaneously plays a separate two-player subgame. Every player has to play the same strategy in every two-player subgame, and her utility is the sum of her subgame utilities. The game is given in the form of the payoff matrix for each two-player subgame.

We construct a bipartite polymatrix game between $n^{1/2+o(1)}$ players with $2^{n^{1/2+o(1)}}$ actions each. By “bipartite,” we mean that each player on Alice’s side only interacts with players on Bob’s side and vice versa. The important term here is “polymatrix”: it means that when we compute the payoffs in each subgame, they can only depend on the $n^{1/2+o(1)}$ coordinates described by the two players’ strategies. It is in this sense that we guarantee “local computation.”

The mixed strategy profile \mathcal{A} of all the players on Alice’s side of the bipartite game induces a vector $\mathbf{x}^{(\mathcal{A})} \in [0, 1]^m$, for some $m = n^{1+o(1)}$. The mixed strategy profile \mathcal{B} of all the players on Bob’s side induces a vector $\mathbf{x}^{(\mathcal{B})} \in [0, 1]^m$. Our main technical result is:

Proposition 16.1 Informal. If all but an ϵ -fraction of the players play ϵ -optimally, then $\|\mathbf{x}^{(\mathcal{A})} - \mathbf{x}^{(\mathcal{B})}\|_2^2 = O(\epsilon)$ and $\|f(\mathbf{x}^{(\mathcal{A})}) - \mathbf{x}^{(\mathcal{B})}\|_2^2 = O(\epsilon)$.

Each player on Alice’s side corresponds to one of the PCP verifier’s random string. Her strategy corresponds to an assignment to the bits queried by the verifier given this random string. On Bob’s side, we consider a partition of $\{1, \dots, m\}$ into $n^{1/2+o(1)}$ tuples of $n^{1/2+o(1)}$ indices each. Each player on Bob’s side assigns values to one such tuple.

On each two-player subgame, the player on Alice’s side is incentivized to imitate the assignment of the player on Bob’s side on the few coordinates where they intersect. The player on Bob’s side uses Alice’s strategy to locally compute $f_j(\mathbf{x}^{(\mathcal{A})})$ on a few j ’s in his $(n^{1/2+o(1)})$ -tuple of coordinates. This computation may be inaccurate, but we can guarantee that for most coordinates it is approximately correct most of the time.

From Polymatrix to Bimatrix

The final reduction from the polymatrix game to the two-player game follows more or less from known techniques for hardness of Nash equilibria [Althofer 1994, Babichenko et al. 2016, Daskalakis et al. 2009a]. We let each of Alice and Bob

control one side of the bipartite polymatrix game. In particular, each strategy in the two-player game corresponds to picking a player of the polymatrix game, and a strategy for that player. We add a gadget due to Althofer [1994] to guarantee that Alice and Bob mix approximately uniformly across all their players. See Section 16.5 for details.

16.1.1 The PCP Conjecture for PPAD

Recall our conjecture from Chapter 5:

Conjecture 16.1 PCP for PPAD; [Babichenko et al. 2016]. There exist constants $\epsilon, \delta > 0$ such that finding an (ϵ, δ) -WeakNash in a bipartite, degree-three polymatrix game with two actions per player is PPAD-complete.

The main original motivation was an approach to prove our main theorem given this conjecture. As pointed out by Babichenko et al. [2016], it turns out that resolving this conjecture would also have interesting consequences for relative approximations of two-player Nash equilibrium, as well as applications to inapproximability of market equilibrium.

More importantly, this question is interesting in its own right: how far can we extend the ideas from the PCP Theorem (for NP) to the world of PPAD? The PCP $[r(n), q(n)]$ characterization [Arora and Safra 1998] is mainly concerned with two parameters: $r(n)$, the number of random bits, and $q(n)$, the number of bits read from the proof. A major tool in all proofs of the PCP Theorem is *verifier composition*: in the work of Polishchuk and Spielman [1994, Spielman 1995], for example, it is first shown that $\text{NP} \subseteq \text{PCP}[O(\log n), n^{1/2+o(1)}]$, and then via composition it is eventually obtained that $\text{NP} = \text{PCP}[O(\log n), O(1)]$. In some *informal sense*, one may think of our main technical result as something analogous⁴ to $\text{PPAD} \subseteq \text{PCP}[(1/2 + o(1)) \log_2 n, n^{1/2+o(1)}]$. Furthermore, our techniques in Section 16.3 build on many existing ideas from the PCP literature [Babai et al. 1991, Ben-Sasson et al. 2003, Ben-Sasson et al. 2006, Polishchuk and Spielman 1994, Spielman 1995] that have been used to show similar statements for NP. It is thus natural to ask: Is there a sense in which our “verifier” can be composed? Can such composition eventually resolve the PCP Conjecture for PPAD?

More generally, some of the tools we use here, even as simple as error-correcting codes, have been the basic building blocks in hardness of approximation for

4. We stress that our analogy is very loose. For example, we are not aware of any formal extension of PCP to function problems, and it is well known that $\text{NP} \subseteq \text{PCP}[(1/2 + o(1)) \log_2 n, n^{1/2+o(1)}]$.

decades, yet to the best of our knowledge have not been used before for any problem in PPAD. We hope to see other applications of similar ideas in this regime.

16.1.2 Organization

In Section 16.2 we introduce the restricted variant LOCAL END-OF-A-LINE and prove that it is also PPAD-complete. In Section 16.3 we construct our holographic proof system. In Section 16.4 we bring together ideas from Sections 4.2, 16.2, and 16.3 to prove our hardness for polymatrix games of subexponential size. Finally, in Section 16.5 we reduce from polymatrix to two-player games.

Small Constants

Our proof uses several arbitrary small constants that satisfy:

$$0 < \epsilon_{\text{PRECISION}} \ll \epsilon_{\text{NASH}} \ll \delta \ll h \ll \epsilon_{\text{COMPLETE}} \ll \epsilon_{\text{SOUND}} \ll \epsilon_{\text{DECODE}} \ll 1.$$

By \ll we mean that we pick them such that $\epsilon_{\text{PRECISION}}$ is arbitrarily smaller than any polynomial in ϵ_{NASH} , and ϵ_{NASH} is arbitrarily smaller than any polynomial in δ , etc. Although their significance will be fully understood later in the chapter, we briefly mention that $\epsilon_{\text{PRECISION}}$ is the precision with which the players can specify real values; ϵ_{NASH} is the approximation factor of Nash equilibrium in Proposition 16.2; δ and h are parameters of the Brouwer function construction; and finally, $\epsilon_{\text{COMPLETE}}$, ϵ_{SOUND} , ϵ_{DECODE} are parameters of our holographic proof system.

16.2 END-OF-A-LINE with Local Computation

In this section we introduce a local variant of END-OF-A-LINE with very simple successor and predecessor circuits.

Definition 16.2 LOCAL END-OF-A-LINE. The problem LOCAL END-OF-A-LINE is similar to END-OF-A-LINE, but the graph is defined on a subset $V^{\text{LOCAL}} \subseteq \{0, 1\}^n$. The input consists of a membership circuit $M_{V^{\text{LOCAL}}}: \{0, 1\}^n \rightarrow \{0, 1\}$ that determines whether a string $u \in \{0, 1\}^n$ corresponds to a vertex of the graph (i.e., $M_{V^{\text{LOCAL}}}(u) = \begin{cases} 1 & u \in V^{\text{LOCAL}} \\ 0 & u \notin V^{\text{LOCAL}} \end{cases}$), a special vertex $u_0 \in V^{\text{LOCAL}}$, and successor and predecessor circuits $S^{\text{LOCAL}}, P^{\text{LOCAL}}: \{0, 1\}^n \rightarrow \{0, 1\}^n$ with the promise that every output bit depends only on a constant number of input bits (i.e., S^{LOCAL} and P^{LOCAL} are in NC^0). Furthermore, we require that the outputs of S^{LOCAL} and P^{LOCAL} are identical to the respective inputs, except at a constant number of coordinates. Similarly to END-OF-A-LINE, we are guaranteed that $P^{\text{LOCAL}}(u_0) = u_0 \neq S^{\text{LOCAL}}(u_0)$.

The goal is to find an input $u \in V^{\text{LOCAL}}$ that satisfies any of the following:

- End-of-a-line: $P^{\text{LOCAL}}(S^{\text{LOCAL}}(u)) \neq u$ or $S^{\text{LOCAL}}(P^{\text{LOCAL}}(u)) \neq u \neq u_0$; or
- Boundary conditions: $S^{\text{LOCAL}}(u) \notin V^{\text{LOCAL}}$ or $P^{\text{LOCAL}}(u) \notin V^{\text{LOCAL}}$.

Notice that each vertex $u \in V^{\text{LOCAL}}$ only differs from $S^{\text{LOCAL}}(u)$ and $P^{\text{LOCAL}}(u)$ at a constant number of bits, and the value of each of those only depends on a constant number of bits of u . We will refer to all those bits as the *critical bits* of u , and denote their cardinality q_{CRITICAL} .

We now reduce END-OF-A-LINE to LOCAL END-OF-A-LINE, proving that the latter is PPAD-complete. In the following sections we use this reduction “white-box,” and revisit some of its specific properties.

Theorem 16.3 There is a linear-time reduction from END-OF-A-LINE to LOCAL END-OF-A-LINE.

Notice that there is a trivial reduction in the other direction: add self-loops to all strings not in V^{LOCAL} .

Proof Proof of Theorem 16.3. Given circuits $S, P: \{0, 1\}^n \rightarrow \{0, 1\}^n$, we construct new circuits $S^{\text{LOCAL}}, P^{\text{LOCAL}}: \{0, 1\}^m \rightarrow \{0, 1\}^m$ that satisfy Definition 16.2. We assume wlog that the circuits have fan-out 2 (otherwise, replace gates with larger fan-outs with a binary tree of equality gates; this only blows up the size of the circuit by a constant factor). We set $m = 4(|S| + |P|)$, where by $|S|$ and $|P|$ we mean the number of lines (i.e., number of inputs + number of logic gates) in each circuit, respectively. We think of the m bits as representing two copies of each circuit (S_1, P_1 and S_2, P_2), with each line represented by two bits, representing three possible states: value 0, value 1, and inactive.

The path begins with circuit S_1 containing the computation from $\mathbf{0}_n$ to $S(\mathbf{0}_n)$, and circuit P_1 containing the reverse computation back to $\mathbf{0}_n$; the lines of circuits S_2, P_2 are inactive. This is the special vertex, u_0 .

Over the next n steps, the output bits of S_1 are copied (one-by-one) to the input bits of S_2 (and the corresponding lines are activated). Over the next $|S| - n$ steps, the values on the lines of S_2 are updated—one-by-one, starting from the inputs and propagating to the outputs in a fixed order—until all the output bits are activated and set to $S(S(\mathbf{0}_n))$. Then, over the next $|P|$ steps, the input of P_2 is set to $S(S(\mathbf{0}_n))$, the values on the gates are updated (propagating in a fixed order from output to input), and finally the output of P_2 is set to $S(\mathbf{0}_n)$.

Notice that so far, if we want to trace back a step (a.k.a. implement P^{LOCAL}), we simply need to deactivate the last activated line. We now start erasing values when going forward, but we must do so carefully so that we can reconstruct them when

going backward. We proceed by deactivating the lines in S_1 : line by line, in reverse order from outputs to inputs. Indeed, in order to go back we can reconstruct the output of any gate from its inputs. If we want to reconstruct an input to S_1 , this is no problem as it is saved as the output of P_1 . Once we finish deactivating S_1 , we deactivate P_1 , also in reverse order. Notice again that the input to P_1 is saved as the input to S_2 and output of P_2 . Now, we copy the values from S_2 and P_2 to S_1 and P_1 in the same order that we used to deactivate the lines, starting from the outputs of S_1 , to inputs, to outputs of P_1 , to inputs. Then we deactivate the lines in S_2 and P_2 ; again, we use the same order in which they were activated, starting from inputs of S_2 , to outputs, to inputs of P_2 , to outputs. So now we have S_1 as a witness to the computation of $S(S(\mathbf{0}_n))$ from $S(\mathbf{0}_n)$ and P_1 goes in the reverse order, while all the lines of S_2 and P_2 are inactive.

We repeat the process from the last two paragraphs to find $S(S(S(\mathbf{0}_n)))$, etc.

We let V^{LOCAL} be the set of strings that correspond to a legal partial computation as described above. $M_{V^{\text{LOCAL}}}$ verifies that a string corresponds to one of the following scenarios:

- S_1 holding the computation from some $x \in \{0, 1\}^n$ to $S(x)$, P_1 holding the computation from $S(x)$ back to x , and S_2 and P_2 holding part of the computation from $S(x)$ to $S(S(x))$ and back to $S(x)$;
- S_2 and P_2 holding the computation from some x to $S(x)$ and back to x , and P_1 and S_1 in the process of erasing the computation from x to $P(x)$ and back to x ;
- S_2 and P_2 holding the computation from some x to $S(x)$ and back to x , and S_1 and P_1 in the process of copying it; or
- S_1 and P_1 holding the computation from some x to $S(x)$ and back to x , and S_2 and P_2 in the process of erasing the same computation.

Notice that the joint active/inactive state of all the lines always belongs to one of $O(m)$ different states (that can be described by $\log_2 m + O(1)$ bits). Verifying that a string corresponds to one of the above scenarios is equivalent to checking that the joint active/inactive state is valid, and that the values on all the lines satisfy all of the following *local conditions*:

- The i -th input of S_1 is either inactive, equal to the i -th output of P_1 , or, if the latter is inactive (during the copy phase), equal to the i -th input of S_2 .

- The i -th input of P_1 is either inactive, equal to the i -th output of S_1 , or, if the latter is inactive (during the erase phase), equal to the i -th input of S_2 .
- The output of every gate in S_1 , P_1 is either inactive, equal to the gate applied to its inputs, or, if either of the inputs is inactive (during the copy phase), equal to the respective line in S_2 , P_2 .
- The i -th input of S_2 is either inactive, equal to the i -th output of P_2 , or, if the latter is inactive (during the compute phase), equal to the i -th input of P_2 .
- The i -th input of P_2 is either inactive, equal to the i -th output of S_2 , or, if the latter is inactive (during the erase phase), equal to the i -th input of P_1 .
- The output of every gate in S_2 , P_2 is either inactive, equal to the gate applied to its inputs, or, if either of the inputs is inactive (during the erase phase), equal to the respective line in S_1 , P_1 .

Notice that each line participates in at most three local constraints. Note also that each of those conditions only depends on a constant number of bits so $M_{V^{\text{LOCAL}}}$, which simply needs to compute their AND, is in AC^0 (this proves Theorem 16.2).

In every step a line is either activated or deactivated, so there are no fixed points; but fixed points of the original circuit lead to violations of the form $P^{\text{LOCAL}}(S^{\text{LOCAL}}(u)) \neq u$ and $S^{\text{LOCAL}}(P^{\text{LOCAL}}(u)) \neq u \neq u_0$. In particular, every solution to the new LOCAL END-OF-A-LINE instance corresponds to a valid solution to the original END-OF-A-LINE instance.

Also, notice that in every step we change at most two bits (and at least one), and deciding the next value of each bit can be done by looking only at a constant number of bits (conditioned on being a legal vertex in V^{LOCAL}): Which is the next line to be activated/deactivated? If activated, what are the inputs to the corresponding gate? ■

16.2.1 The LOCAL END-OF-A-LINE Counter

Consider the instance produced in the reduction above. Notice that joint active/inactive state of all the lines rotates among $O(m)$ possible vectors. That is, there is a $(\log m + O(1))$ -bit vector, henceforth called the *counter*, that describes precisely which lines should be active. Furthermore, given the counter of some $u \in V^{\text{LOCAL}}$, it is easy to compute the counter of $S^{\text{LOCAL}}(u)$, i.e., we always know which line should be activated/deactivated next—regardless of the values on the lines. Similarly, for all $u \neq u_0 \in V^{\text{LOCAL}}$, it is also easy to compute the counter of

$P^{\text{LOCAL}}(u)$. An additional useful property is that, given the counter of u , we also know the coordinates of the critical bits of u .

16.3 Holographic Proof

In this section we construct our holographic proof system for statements of the form $u \in V^{\text{LOCAL}}$. Note that we use the term “verifier” loosely, as our verifier’s responsibilities extend far beyond checking the validity of a proof—it is also expected to perform local decoding, “local proof-construction,” etc.

Proposition 16.2 For sufficiently small constants $\epsilon_{\text{DECODE}} \gg \epsilon_{\text{SOUND}} \gg \epsilon_{\text{COMPLETE}} > 0$, the following holds. Fix any instance $(S^{\text{LOCAL}}, P^{\text{LOCAL}}, V^{\text{LOCAL}})$ of LOCAL END-OF-A-LINE, generated via the reduction in Section 16.2 from an instance (S, P) END-OF-A-LINE. Given a vertex $u \in V^{\text{LOCAL}}$, there is a polynomial time deterministic algorithm that takes as input (S, P) and u , and outputs a *holographic proof* $\Pi(u) = (E(u), C(u), \pi(u))$ where:

- $E(u)$ is an *encoding* of u with a *linear* error-correcting code of constant relative distance.
- $C(u)$ is an encoding of the *counter* of u (with a good error-correcting code).
- $\pi(u)$ is a *proof* that $u \in V^{\text{LOCAL}}$.
- Let $n \triangleq |u|$; then the total length of $\Pi(u)$ is $n^{1+o(1)}$.
- Let $t \triangleq \sqrt{\log n}$ (in particular, $t = \omega(1)$, but also $t = o(\frac{\log n}{\log \log n})$). Let \mathcal{G} be a finite field of size $O(n^{1/t+o(1)})$. $E(u)$ and $\pi(u)$ can be written as functions over domain \mathcal{G}^t . In particular, we can talk about accessing $E(u)$, $\pi(u)$, or $\Pi(u)$ at a point $\mathbf{g} \in \mathcal{G}^t$. ($C(u)$ is much shorter and the verifier reads it entirely.)
- There is a quasi-polynomial⁵ time probabilistic verifier such that:
- **Local Access.** The verifier’s access to $\Pi(u)$ is restricted as follows:

Querying Subspaces. The verifier reads $n^{o(1)}$ axis-parallel $(t/2)$ -dimensional subspaces of $\Pi(u)$. Each subspace is defined by a restriction of a $(t/2)$ -tuple of coordinates. The tuple of coordinates that is restricted is always one of a constant number of possibilities (either $\{2, \dots, t/2 + 1\}$, or the union of any two of: $\{1, \dots, t/4\}$, $\{t/4 +$

5. The verifier of our holographic proof actually runs in polynomial time, except for the derandomization of the construction of λ -biased sets. As we discuss in Section 2.7.3, this could also be obtained (with slightly worse parameters) in deterministic polynomial time, but quasi-polynomial time suffices for the purpose of our main theorem.

$1, \dots, t/2\}, \{t/2 + 1, \dots, 3t/4\}$, and $\{3t/4 + 1, \dots, t\}$). We denote the union of all the subspaces read by the verifier G_{ALL} . It will be useful to decompose $G_{\text{ALL}} = G_{\text{PCP}} \cup G_{\text{LTC}} \cup G_{\text{CRITICAL}}$; another subset of interest is $G_{\text{SAMPLE}} \subset G_{\text{PCP}}$.

Partially Adaptive. $G_{\text{SAMPLE}}, G_{\text{PCP}}, G_{\text{LTC}}$ are chosen non-adaptively, i.e., they depend only on the verifier's internal randomness. G_{CRITICAL} is a union of affine translations of a subspace that is also chosen non-adaptively, but the affine translations depend on $C(u)$.

Randomness. The verifier uses only $(1/2 + o(1)) \log_2 n$ bits of randomness to decide which subspaces to query.

Good Sample. For any function $B: \mathcal{G}^t \rightarrow [0, 1]$, and any $G \in \{G_{\text{SAMPLE}}, G_{\text{PCP}}, G_{\text{LTC}}, G_{\text{CRITICAL}}\}$, with probability $1 - o(1)$ over the verifier's randomness,

$$|\mathbb{E}_{\mathbf{g} \in G}[B(\mathbf{g})] - \mathbb{E}_{\mathbf{h} \in \mathcal{G}^t}[B(\mathbf{h})]| = o(1).$$

- **Soundness.** If the verifier is given a string Π' that is ϵ_{SOUND} -far from $\Pi(u)$ for every $u \in V^{\text{LOCAL}}$, the verifier rejects with probability $1 - o(1)$.

Robust Soundness. Furthermore, with probability $1 - o(1)$, the bits read by the verifier are $\epsilon_{\text{SOUND}}^2$ -far from any string that would make the verifier accept. In other words, the verifier continues to reject with probability $1 - o(1)$ even in the presence of an adaptive adversary, who observes the queries made by the verifier, and then corrupts an $\epsilon_{\text{SOUND}}^2/2$ -fraction of each of $G_{\text{PCP}}, G_{\text{LTC}}$.

- **Completeness.** If the verifier is given access to a string $\widehat{\Pi}$ that is $\epsilon_{\text{COMPLETE}}$ -close to a valid $\Pi(u)$ for some $u \in V^{\text{LOCAL}}$, then the verifier accepts with probability $1 - o(1)$.

Robust Completeness. Furthermore, the verifier continues to accept with probability at least $1 - o(1)$ even in the presence of an adaptive adversary, who observes the queries made by the verifier, and then corrupts a $\sqrt{\epsilon_{\text{COMPLETE}}}$ -fraction of $G_{\text{PCP}}, G_{\text{LTC}}$.

- **Decoding.** If the verifier is given access to a string $\widehat{\Pi}$ that is ϵ_{DECODE} -close to a valid $\Pi(u)$ for some $u \in V^{\text{LOCAL}}$;

Error-Correction on a Sample. With probability at least $1 - o(1)$, the verifier correctly decodes the entries of $E(u)$ and $\pi(u)$ on G_{SAMPLE} . In particular, with probability $1 - o(1)$, the verifier can estimate the distance between $\widehat{\Pi}$ and $\Pi(u)$ to within an additive error of $\pm o(1)$.

Error-Correction on the Critical Bits. The verifier can *adaptively* decode and correct any q_{CRITICAL} -tuple of symbols from the proof with success probability at least $1 - o(1)$. Let $G_{\text{CRITICAL}} \subset \mathcal{G}^t$ denote the union of all $O(|\mathcal{G}|^2)$ axis-parallel subspaces queried in this process.

Proof-Construction. Given access as above to the counter $C(u)$ and the linear code $E(u)$ (but not to the proof $\pi(u)$), the verifier can output (with probability at least $1 - o(1)$) the correct values of $\pi(u)$ on G_{SAMPLE} .

Robust Decoding. Both types of error-correction and the proof-construction continue to hold even in the presence of an adaptive adversary, who observes the queries made by the verifier, and then corrupts a $\sqrt{\epsilon_{\text{DECODE}}}$ -fraction of each of $G_{\text{SAMPLE}}, G_{\text{PCP}}, G_{\text{LTC}}, G_{\text{CRITICAL}}$.

The rest of this section is devoted to the proof of Proposition 16.2.

16.3.1 From LOCAL END-OF-A-LINE to Coloring Graphs

Let $t \triangleq \sqrt{\log n}$. Set $\ell \triangleq \frac{\log_2 n}{t-1} = \Theta(\sqrt{\log n})$, and $\ell' \triangleq \ell + c \log \log n$ for some sufficiently large constant c . Let $p(z)$ be an \mathbb{F}_2 -irreducible polynomial of degree ℓ' , and let $\mathcal{G} \triangleq \mathbb{F}_2[z]/\langle p(z) \rangle$ denote the field of size $2^{\ell'}$. We will focus our attention on a subset $\mathcal{F} \subset \mathcal{G}$ of polynomials (modulo $p(z)$) of degree ℓ ; notice that $|\mathcal{G}|/|\mathcal{F}| = \text{poly } \log n$.

Let α be a generator of \mathcal{G} , and let $\mathcal{E} \triangleq \{1, \alpha, \dots, \alpha^{5t\ell}\}$. We define the *Extended Shuffle- \mathcal{F} -Exchange graph* on $\mathcal{F}^{t-1} \times \mathcal{E}$ to be the directed graph such that each vertex $(x_1, \dots, x_{t-1}, \alpha^i) \in \mathcal{F}^{t-1} \times (\mathcal{E} \setminus \{\alpha^{5t\ell}\})$ has a *static edge* to vertex $(x_1, \dots, x_{t-1}, \alpha^{i+1})$, a *shuffle edge* to vertex $(x_2, \dots, x_{t-1}, x_1, \alpha^{i+1})$, and ℓ *exchange edges* to all vertices of the form $(x_1 + [z^j], \dots, x_{t-1}, \alpha^{i+1})$, for every $j \in \{0, \dots, \ell - 1\}$. Here $[z^j]$ is the element of \mathcal{G} representing the equivalence class $\{z^j + q(z)p(z) : q(z) \in \mathbb{F}_2[z]\}$. Notice that $x \in \mathcal{F}$ iff it can be written as $x = \sum_{j \in S} [z^j]$ for some subset $S \subseteq \{0, \dots, \ell - 1\}$. In particular, for any $(x_1, \dots, x_{t-1}), (y_1, \dots, y_{t-1}) \in \mathcal{F}^{t-1}$, the Extended Shuffle- \mathcal{F} -Exchange graph contains a path of length $(t-1)(\ell+1)$ from $(x_1, \dots, x_{t-1}, 1)$ to $(y_1, \dots, y_{t-1}, \alpha^{(t-1)(\ell+1)})$.

For each $\alpha^i \in \mathcal{E}$, we call the set of vertices $\{(x_1, \dots, x_{t-1}, \alpha^i)\}$ a *layer*; in particular, $\{(x_1, \dots, x_{t-1}, 1)\}$ form the *first layer*, and $\{(x_1, \dots, x_{t-1}, \alpha^{5t\ell})\}$ form the *last layer*.

We encode each vertex $u \in V^{\text{LOCAL}}$ as a partial coloring of the Extended Shuffle- \mathcal{F} -Exchange graph, representing the assignments to the lines in the four circuits S_1, P_1, S_2, P_2 . Recall that $|\mathcal{F}| = 2^\ell = n^{1/(t-1)}$. We associate each line of S_1, P_1, S_2, P_2 (including input, gates, and output lines) with a vector $\mathbf{x} \in \mathcal{F}^{t-1}$ arbitrarily. For the first- and last-layer vertices $(\mathbf{x}, 1)$ and $(\mathbf{x}, \alpha^{5t\ell})$, we assign colors that represent the assignments (from $\{0, 1, \perp\}$, where \perp represents an inactive line) to the corresponding

line. Recall from the “local conditions” in Section 16.2 that every line only participates in three constraints.

We would like the inputs to the line in $(x_1, \dots, x_{t-1}, \alpha^{5t\ell})$ to be propagated by the Extended Shuffle- \mathcal{F} -Exchange graph from the first layer. One important feature of the Extended Shuffle- \mathcal{F} -Exchange graph is that we can do exactly that: By using standard packet-routing techniques, we can route the assignments from the first layer to the last layer so that every vertex on the way needs to remember at most three assignments [Leighton 1992 (Theorem 3.16)]. Fix one such routing. Our coloring now assigns the same color from $\{0, 1, \perp\} \times \{\perp\}^2$ to each of the corresponding vertices on the first and last layers, and propagates this assignment through the middle layers according to the routing scheme; each vertex has a color in $\{0, 1, \perp\}^3$.

It is important to note that the middle layers do not perform any computation—they only copy symbols according to a fixed routing. In particular, if we change one entry in the representation of LOCAL END-OF-A-LINE vertex u , we can locally compute the difference of the old and new assignments for every middle layer node, even if we don’t know the old assignment itself (which may correspond to up to three values routed through that node).

Notice that in order to verify that $u \in V^{\text{LOCAL}}$, we need to check that the active/inactive pattern matches the counter, and the “local conditions.” Given a coloring of the Extended Shuffle- \mathcal{F} -Exchange graph, checking the local conditions reduces to checking that the color for each first-layer vertex is equal to the color on the respective vertex on the last layer, and that the color of every vertex not in the first layer is computed correctly from its incoming neighbors.

16.3.2 Arithmetization

We now want to represent the verification of $u \in V^{\text{LOCAL}}$ as constraints on polynomials over \mathcal{G} (do not confuse with polynomials over \mathbb{F}_2 !).

We represent each vertex $u \in V^{\text{LOCAL}}$ as two polynomials: $T_u: \mathcal{F}^{t-1} \times \mathcal{E} \rightarrow \mathcal{C}$ and $T_{\text{COUNTER}}: \mathcal{F}^{t-1} \rightarrow \mathcal{C}$, where $\mathcal{C} \subset \mathcal{F}$ is a subset of \mathcal{F} of constant size. For each $(x_1, \dots, x_{t-1}, \alpha^i) \in \mathcal{F}^{t-1} \times \mathcal{E}$, we set $T_u(x_1, \dots, x_{t-1}, \alpha^i)$ to be equal to the color assigned to the respective vertex. $T_{\text{COUNTER}}(x_1, \dots, x_{t-1})$ is assigned 1 if the corresponding line should be active, and 0 otherwise. Additionally, we construct a polynomial $T_{S,P}: \mathcal{F}^{t-1} \times \mathcal{E} \rightarrow \mathcal{C}'$, for $\mathcal{C} \subset \mathcal{C}' \subset \mathcal{F}$ of size $|\mathcal{C}'| = \Theta(\log^{3/2} n)$; $T_{S,P}(\cdot)$ represents the instance of END-OF-A-LINE: for each $(x_1, \dots, x_{t-1}, \alpha^{5t\ell}) \in \mathcal{F}^{t-1} \times \{\alpha^{5t\ell}\}$, it specifies the gate functionality, whereas for $(x_1, \dots, x_{t-1}, \alpha^i)$ ($i < 5t\ell$) it specifies the routing through this vertex. Note that each vertex takes its inputs from at most 3 out of $\Theta(\sqrt{\log n})$ incoming edges.

Notice that $T_u, T_{\text{COUNTER}}, T_{S,P}$ are all polynomials of degree at most $|\mathcal{F}|, |\mathcal{E}|$ in their respective variables. The verifier will expect low-degree extensions of $T_u, T_{S,P}$ on all of \mathcal{G}^t , and analogously for T_{COUNTER} over \mathcal{G}^{t-1} .

In the following, keep in mind that T_{COUNTER} is constructed implicitly (once we decode all $\log_2 n + O(1)$ bits of the counter), and $T_{S,P}$ is part of the input of the problem, so we only need to locally decode and verify the validity of T_u .

First, we construct a polynomial

$$\Psi(\mathbf{x}) \triangleq \psi(T_u(\mathbf{x})) \triangleq \prod_{c \in \mathcal{C}} (T_u(\mathbf{x}) - c), \quad (16.1)$$

such that $\psi(T_u(\mathbf{x})) = 0$ iff $T_u(\mathbf{x}) \in \mathcal{C}$.

Similarly, we would like to construct a polynomial that verifies that the colors are propagated correctly through the middle layers; for the last layer, it should also verify that the gate functionality is implemented correctly. First, we describe the edges of the Extended Shuffle- \mathcal{F} -Exchange graph with the following affine transformations:

$$\begin{aligned} \rho_{\text{STATIC}}(x_1, \dots, x_{t-1}, \alpha^{i+1}) &\triangleq (x_1, \dots, x_{t-1}, \alpha^i) \\ \rho_{\text{SHUFFLE}}(x_1, \dots, x_{t-1}, \alpha^{i+1}) &\triangleq (x_{t-1}, x_1, \dots, x_{t-2}, \alpha^i) \\ \rho_0(x_1, \dots, x_{t-1}, \alpha^{i+1}) &\triangleq (x_1 + [z^0], \dots, x_{t-1}, \alpha^i) \\ &\vdots \\ \rho_{\ell-1}(x_1, \dots, x_{t-1}, \alpha^{i+1}) &\triangleq (x_1 + [z^{\ell-1}], \dots, x_{t-1}, \alpha^i) \end{aligned}$$

(So each vertex $\mathbf{x} \in \mathcal{F}^{t-1} \times (\mathcal{E} \setminus \{1\})$ has incoming edges from $\rho_{\text{STATIC}}(\mathbf{x}), \rho_{\text{SHUFFLE}}(\mathbf{x}), \rho_0(\mathbf{x}), \dots, \rho_{\ell-1}(\mathbf{x})$.)

Now we can define a polynomial

$$\begin{aligned} \Phi(\mathbf{x}) \triangleq \phi \left(T_u(\mathbf{x}), T_u(\rho_{\text{STATIC}}(\mathbf{x})), T_u(\rho_{\text{SHUFFLE}}(\mathbf{x})), T_u(\rho_0(\mathbf{x})), \right. \\ \left. \dots, T_u(\rho_{\ell-1}(\mathbf{x})), T_{(S,P)}(\mathbf{x}) \right), \end{aligned} \quad (16.2)$$

where for $\mathbf{x} \in \mathcal{F}^{t-1} \times (\mathcal{E} \setminus \{1, \alpha^{5t\ell}\})$, $\phi(\dots) = 0$ iff the colors are propagated correctly to the corresponding vertex, and for $\mathbf{x} \in \mathcal{F}^t \times \{\alpha^{5t\ell}\}$, $\phi(\dots)$ further checks that the gate functionality is implemented correctly.

Finally, we add a third polynomial over \mathcal{F}^{t-1} ,

$$\Xi(\mathbf{x}) \triangleq \xi \left(T_u(x_1, \dots, x_{t-1}, 1), T_u(x_1, \dots, x_{t-1}, \alpha^{5t\ell}), T_{\text{COUNTER}}(x_1, \dots, x_{t-1}) \right), \quad (16.3)$$

which is zero iff $T_u(x_1, \dots, x_{t-1}, 1) = T_u(x_1, \dots, x_{t-1}, \alpha^{5t\ell})$ and the corresponding line is active or inactive as dictated by $T_{\text{COUNTER}}(x_1, \dots, x_{t-1})$.

We now have that $u \in v^{\text{LOCAL}}$ is equivalent to AND of the following conditions:

- values-in-domain: $\Psi(\mathbf{x}) = 0$ for all $\mathbf{x} \in \mathcal{F}^{t-1} \times \mathcal{E}$;
- value-propagation and computation: $\Phi(\mathbf{x}) = 0$ for all $\mathbf{x} \in \mathcal{F}^{t-1} \times (\mathcal{E} \setminus \{1\})$; and
- active-vs-inactive: $\Xi(\mathbf{x}) = 0$ for all $\mathbf{x} \in \mathcal{F}^{t-1}$.

Notice also that ψ, ξ are constant total degree polynomials, whereas ϕ has total degree polylog n ; thus Ψ, Φ, Ξ are of degrees $|\mathcal{F}| \cdot \text{polylog } n, |\mathcal{E}| \cdot \text{polylog } n \leq |\mathcal{G}|/2$ in their respective variables.

16.3.3 PCP

Our verifier should test that Ψ, Φ, Ξ are zero everywhere on the respective domains $(\mathcal{F}^{t-1} \times \mathcal{E}, \mathcal{F}^{t-1} \times (\mathcal{E} \setminus \{1\}), \mathcal{F}^{t-1})$. We describe the test for Ψ ; the corresponding tests for Φ and Ξ follow analogously.

In order to test that $\Psi(\mathbf{x})$ is zero over all of $\mathcal{F}^{t-1} \times \mathcal{E}$, consider the following polynomials:

$$\Psi'(x_1, \dots, x_{t-1}, \alpha^j) \tag{16.4}$$

$$\triangleq \sum_{\substack{f_{i_t/2+1}, \dots, f_{i_{t-1}} \in \mathcal{F} \\ j' \in \{1, \dots, 5t\ell\}}} (\Psi(x_1, \dots, x_{t/2}, f_{i_t/2+1}, \dots, f_{i_{t-1}}, \alpha^{j'})) \prod_{k=t/2+1}^{t-1} x_k^{i_k-1} \cdot (\alpha^j)^{j'}$$

$$\Psi''(x_1, \dots, x_{t-1}, \alpha^j) \tag{16.5}$$

$$\triangleq \sum_{f_1, \dots, f_{i_t/2} \in \mathcal{F}} \Psi'(f_1, \dots, f_{i_t/2}, x_{t/2+1}, \dots, x_{t-1}, \alpha^j) \prod_{k=1}^{t/2} x_k^{i_k-1}.$$

In particular, observe that

$$\Psi''(x_1, \dots, x_{t-1}, \alpha^j) = \sum_{\substack{f_{i_1}, \dots, f_{i_{t-1}} \in \mathcal{F} \\ j' \in \{1, \dots, 5t\ell\}}} \Psi(f_{i_1}, \dots, f_{i_{t-1}}, \alpha^{j'}) \cdot \prod_{k=1}^{t-1} x_k^{i_k-1} \cdot (\alpha^j)^{j'}.$$

The main point of this construction is that Ψ'' is the zero polynomial if and only if $\Psi(\mathbf{x}) = 0$ for every $\mathbf{x} \in \mathcal{F}^{t-1} \times \mathcal{E}$. Now, by definition, Ψ'' has degrees at most $|\mathcal{F}|, |\mathcal{E}|$ in each variable, respectively, so it suffices to test that it is indeed zero at one random point from \mathcal{G}^t .

The verifier should also ensure that Ψ' and Ψ'' are constructed correctly. This is done by picking a uniformly random $\mathbf{x} \in \mathcal{G}^t$ and verifying that (16.4) and (16.5) hold.

If Ψ is constructed correctly, then Ψ' also has low-degree in each variable (compared to $|\mathcal{G}|$), so it suffices to test on a single $\mathbf{x} \in \mathcal{G}^t$. Similarly, the verifier should also check that Ψ is constructed correctly as in (16.1); thanks to the low-degree guarantee, a single $\mathbf{x} \in \mathcal{G}^t$ suffices here as well.

Notice that the verifier does not need to receive Ψ'' explicitly. (This is crucial later for the proof construction!) Testing that Ψ'' is constructed correctly at $(x_1, \dots, x_{t-1}, \alpha^j)$ and that it is zero at the same point is equivalent to simply testing that

$$\sum_{f_{i_1}, \dots, f_{i_{t/2}} \in \mathcal{F}} \Psi'(f_{i_1}, \dots, f_{i_{t/2}}, x_{t/2+1}, \dots, x_{t-1}, \alpha^j) \prod_{k=1}^{t/2} x_k^{i_k-1} = 0. \quad (16.6)$$

We also need to test that all polynomials are indeed low-degree. We do this later, in Section 16.3.7.

Define and test Φ', Ξ' analogously to Ψ' . We conclude that if $T_u, \Psi, \Phi, \Xi, \Psi', \Phi', \Xi'$ are all indeed of low degree as promised and the verifier accepts with probability greater than ϵ , then T_u, T_{COUNTER} represent a real vertex $u \in V^{\text{LOCAL}}$. Notice that all of our tests choose $\mathbf{x} \in \mathcal{G}^{t/2}$ with uniform marginal probabilities.

For simplicity of notation, we henceforth treat $\Xi(\cdot)$ as a t -variate polynomial (with the t -th variable being a dummy variable).

16.3.4 The Final Encoding

Our proof $\Pi(u)$ consists of the following strings.

- $C(u)$ consists of a good encoding of the “counter” of u . (Notice that we require a succinct encoding that can be fully decoded, and then the entire T_{COUNTER} can be computed locally.)
- $E(u)$ consists of the polynomial $T_u: \mathcal{G}^t \rightarrow \mathcal{G}$.
- $\pi(u)$ consists of the polynomials $\Psi, \Psi', \Phi, \Phi', \Xi, \Xi': \mathcal{G}^t \rightarrow \mathcal{G}$.

Although they have different lengths (in particular, $C(u)$ is much shorter), we think of them as having equal “weights,” i.e., we allow only a constant fraction of errors in each. Also, we think of symbols in larger alphabets as encoded with a good binary error-correcting code.

16.3.5 Local Decoding and Error Correction

When a typical PCP verifier is given a proof that is wrong even at a single bit, it is already allowed to reject the proof. For our purpose, we want a more lenient verifier that, when given a proof that is wrong at a small fraction of the bits, corrects

those bits and accepts. Furthermore, we want to ensure that our verifier only uses very little randomness: only $(1/2 + o(1)) \log_2 n$ random bits. A second goal of this subsection is to establish the desideratum that the verifier can adaptively decode an arbitrary q_{CRITICAL} -tuple of entries from the proof.

We describe how to locally decode and correct Ψ ; $T_u, \Psi', \Phi, \Xi, \Phi', \Xi'$ follow with minor modifications, whereas the counter has a succinct representation and can be fully decoded. Specifically, the verifier needs to decode Ψ in order to test (16.1), (16.4). For (16.1) the verifier only needs to decode Ψ at one uniformly random point, so this is the easy case. For (16.4), the verifier wants to decode an entire axis-parallel $(t/2)$ -dimensional affine subspace, $Q_{(x_1, \dots, x_{t/2})} \triangleq \{(x_1, \dots, x_{t/2})\} \times \mathcal{G}^{t/2}$, for a uniformly random choice of $(x_1, \dots, x_{t/2}) \in \mathcal{G}^{t/2}$.

Let $S \subset \mathcal{G}^{t/2}$ be a $(1/(t^2 \log |\mathcal{G}|))$ -biased set of cardinality $|S| = \text{poly}(t, \log |\mathcal{G}|)$ (as guaranteed by Section 2.7.3). We pick a uniformly random $\mathbf{x} = (x_1, \dots, x_{t/2}) \in \mathcal{G}^{t/2}$, and a uniformly random $\mathbf{y} = (y_1, \dots, y_{t/2}) \in S$. We then consider the $(t/2 + 1)$ -dimensional affine subspace $R_{\mathbf{x}, \mathbf{y}} \triangleq \bigcup_{\beta \in \mathcal{G}} Q_{(\mathbf{x} + \beta \mathbf{y})}$. If $\widehat{\Psi}$ is $O(\epsilon_{\text{DECODE}})$ -close to Ψ , we have by Lemma 2.4 that with probability $1 - o(1)$, the restriction of $\widehat{\Psi}$ to $R_{\mathbf{x}, \mathbf{y}}$, denoted $\widehat{\Psi}|_{R_{\mathbf{x}, \mathbf{y}}}$, is also $O(\epsilon_{\text{DECODE}})$ -close to $\Psi|_{R_{\mathbf{x}, \mathbf{y}}}$. Whenever this is the case, the verifier correctly decodes Ψ on $R_{\mathbf{x}, \mathbf{y}}$.

Notice that so far, our verifier queries $|\mathcal{G}|^{t/2+1} = n^{1/2+o(1)}$ symbols, and uses $\log_2 |\mathcal{G}^{t/2}| + \log_2 |S| = (1/2 + o(1)) \log_2 n$ random bits.

The local decoding for $T_u, \Psi', \Phi, \Xi, \Phi', \Xi'$ follows with minor modifications that we now describe. Notice that we can reuse the same random bits from the decoding of Ψ to decode each of the other polynomials. For (16.5) the verifier wants to decode Ψ' (equivalently for Φ', Ξ') on a subspace that fixes the second half of the coordinates, $Q'_{(x_{t/2+1}, \dots, x_t)} \triangleq \mathcal{G}^{t/2} \times \{(x_{t/2+1}, \dots, x_t)\}$.

For (16.1) and (16.3), the verifier also needs to decode one or two entries of T_u , but this is no harder than decoding an entire plane. For (16.2), there is some subtlety: $\Phi(\mathbf{x})$ is a function of the values of T_u on a super-constant number $(\ell + 3 = \Theta(\sqrt{\log n}))$ of points $\mathbf{x}, \rho_{\text{STATIC}}(\mathbf{x}), \rho_{\text{SHUFFLE}}(\mathbf{x}), \rho_0(\mathbf{x}), \dots, \rho_{\ell-1}(\mathbf{x})$. Observe that $\mathbf{x}, \rho_{\text{STATIC}}(\mathbf{x}), \rho_0(\mathbf{x}), \dots, \rho_{\ell-1}(\mathbf{x})$ all belong to the same $(t/2)$ -dimensional subspace (only the first and last coordinates change); thus we only need to apply subspace-decoding twice.

Notice that our decoding mechanism also suffices to locally correct errors (with probability at least $(1 - o(1))$, and assuming that we start with a function that is close to low-degree). Furthermore, by Lemma 2.4, the distance $|\widehat{\Psi}|_{R_{\mathbf{x}, \mathbf{y}}} - \Psi|_{R_{\mathbf{x}, \mathbf{y}}}|$ is within $\pm o(1)$ of the global distance $|\widehat{\Psi} - \Psi|$. Handling the rest of the polynomials as described in the previous paragraphs, we can thus estimate $|\widehat{\Pi}|_{E, \pi} - \Pi(u)|_{E, \pi}|$ to within $\pm o(1)$; for the bits that correspond to the counter, we can compute the distance exactly.

Decoding a Particular Subspace

Above, we described how to locally decode a random subspace. Suppose instead that we want to locally decode $\Psi|_{Q(w_1, \dots, w_{t/2})}$ for some particular $\mathbf{w} \in \mathcal{G}^{t/2}$ (specifically, when $Q(w_1, \dots, w_{t/2})$ contains information about one of the critical bits). Since we did not waste any random bits on picking \mathbf{w} , we can afford to pick \mathbf{y} uniformly from all of $\mathcal{G}^{t/2}$. On the other hand, since \mathbf{w} is arbitrary, there may still be an $\Theta(\epsilon_{\text{DECODE}})$ chance that $|\widehat{\Psi}|_{R_{\mathbf{w}, \mathbf{y}}} - \Psi|_{R_{\mathbf{w}, \mathbf{y}}}|$ is too large ($> 1/2$) to guarantee correct decoding. Instead, we pick a third vector $\mathbf{z} \in S$, and query $\widehat{\Psi}$ on the $(t/2 + 2)$ -dimensional affine subspace $P_{\mathbf{w}, \mathbf{y}, \mathbf{z}} \triangleq \bigcup_{\beta, \gamma \in \mathcal{G}} Q(\mathbf{w} + \beta(\mathbf{y} + \gamma\mathbf{z}))$; notice that $P_{\mathbf{w}, \mathbf{y}, \mathbf{z}} \triangleq \bigcup_{\gamma \in \mathcal{G}} R_{\mathbf{w}, \mathbf{y} + \gamma\mathbf{z}}$. In expectation, for each $(\mathbf{y} + \gamma\mathbf{z})$, the restriction $\widehat{\Psi}|_{R_{\mathbf{w}, \mathbf{y} + \gamma\mathbf{z}}}$ is at distance at most $O(\epsilon_{\text{DECODE}})$ from $\Psi|_{R_{\mathbf{w}, \mathbf{y} + \gamma\mathbf{z}}}$. Therefore, by Lemma 2.4, with probability $1 - o(1)$, we have that $|\widehat{\Psi}|_{P_{\mathbf{w}, \mathbf{y}, \mathbf{z}}} - \Psi|_{P_{\mathbf{w}, \mathbf{y}, \mathbf{z}}}| = O(\epsilon_{\text{DECODE}})$.

16.3.6 Local Proof-Construction

We now want to establish the following unusual property of our holographic proof scheme: we can, with high probability, locally construct any subspace $Q(x_1, \dots, x_{t/2})$ of the proof $\pi(u)$ by reading only a small part of the encoding $E(u)$ (in particular, we need to decode a constant number of $(t/2 + 1)$ -dimensional subspaces from $E(u)$). Here, it is important that $Q(x_1, \dots, x_{t/2})$ is defined via a restriction of the *first* $t/2$ coordinates, but we assume access to decoded subspaces of $E(u)$ with restriction of any $(t/2)$ -tuple of coordinates.

Recall that in order to compute $\Psi(\mathbf{x})$ or $\Xi(\mathbf{x})$, it suffices to know the values of T_u at a constant number of locations ((16.1) and (16.3)), so here this property trivially holds. For $\Phi(\mathbf{x})$, we need to know T_u at $\Theta(\sqrt{\log n})$ locations (16.2); however, as we argue in Section 16.3.5, they belong to two axis-parallel subspaces, so it again suffices to decode T_u on a constant number of subspaces.

In order to compute $\Psi'(\mathbf{x})$, $\Phi'(\mathbf{x})$, $\Xi'(\mathbf{x})$, we need to correctly decode the values of Ψ , Φ , Ξ on (almost) an entire axis-parallel subspace $Q(x_1, \dots, x_{t/2})$. The values of Ψ , Φ , Ξ are not given as part of $E(u)$, but (as we argued in the previous paragraph) we can reconstruct them from T_u . Specifically, in order to know the values of Ψ or Ξ on $Q(x_1, \dots, x_{t/2})$, we simply need to decode T_u on $Q(x_1, \dots, x_{t/2})$. In order to construct Φ on an entire subspace, observe that for every $(x_1, \dots, x_{t/2}, f_{i_{t/2+1}}, \dots, f_{i_{t-1}}, \alpha^j) \in Q(x_1, \dots, x_{t/2})$, all the vectors $\rho_{\text{STATIC}}(x_1, \dots, x_{t/2}, f_{i_{t/2+1}}, \dots, f_{i_{t-1}}, \alpha^j)$, $\rho_0(x_1, \dots, x_{t/2}, f_{i_{t/2+1}}, \dots, f_{i_{t-1}}, \alpha^j)$, \dots , $\rho_{\ell-1}(x_1, \dots, x_{t/2}, f_{i_{t/2+1}}, \dots, f_{i_{t-1}}, \alpha^j)$ belong to the same subspace $Q(x_1, \dots, x_{t/2})$. Furthermore, all the vectors $\rho_{\text{SHUFFLE}}(x_1, \dots, x_{t/2}, f_{i_{t/2+1}}, \dots, f_{i_{t-1}}, \alpha^j) = (f_{i_{t-1}}, x_1, \dots, x_{t/2}, f_{i_{t/2+1}}, \dots, f_{i_{t-2}}, \alpha^j)$ belong to the subspace $Q_{(x_1, \dots, x_{t/2})}^{\text{SHUFFLE}} \triangleq \mathcal{G} \times (x_1, \dots, x_{t/2}) \times \mathcal{G}^{t/2-1}$. Therefore the value of $\Phi'(\mathbf{x})$ can also be computed from the values of T_u on a constant number of subspaces.

16.3.7 Local Robust Testing

So far, we assumed that we are given functions that are low-degree (Section 16.3.3), or approximately low-degree (Section 16.3.5). However, we must verify that this is indeed the case.

We will use results for locally testing tensor codes [Viderman 2015]. In order to describe them, let us briefly introduce some minimal background and notation. Given linear codes C_1, C_2 with respective generator matrices M_1, M_2 , their *tensor product*, denoted $C_1 \otimes C_2$, is the linear code generated by the tensor product of the matrices $M_1 \otimes M_2$. We say that $C_1 \otimes C_2$ is a *tensor code*. In general, we can talk about the tensor product of k codes, $C_1 \otimes \dots \otimes C_k$, which is defined recursively. We say that $C_1 \otimes \dots \otimes C_k$ is the k -th power of C (denoted C^k), if $C_i = C$ for all $i \in [k]$. In our case, for example, the code \mathcal{C} of all $(|\mathcal{G}|/2)$ -individual-degree polynomials $T: \mathcal{G}^t \rightarrow \mathcal{G}$ is the t -th power of the code of all $(|\mathcal{G}|/2)$ -degree polynomials $t: \mathcal{G} \rightarrow \mathcal{G}$.

Alternatively, we can think of any product code as a function from some product domain to some range. The (axis-parallel) *hyperplane tester* picks a uniformly random index $i \in [k]$, and a uniformly random value v_i from the i -th domain. It reads the assignment of the function for all vectors that have v_i in their i -th coordinate, and tests whether they form (or are close to) a codeword from $C_1 \otimes \dots \otimes C_{i-1} \otimes C_{i+1} \otimes \dots \otimes C_k$. Going back to our setting, this could mean selecting a random $i \in \{1, \dots, t\}$ and a random value $v_i \in \mathcal{G}$, reading $T(\mathbf{x})$ for all \mathbf{x} such that $x_i = v_i$, and verifying that the restricted $(t-1)$ -variate function is close to a low-degree polynomial.

To optimize query complexity, it will be more convenient for us to think of the code of $(t/4)$ -variate $(|\mathcal{G}|/2)$ -individual-degree polynomials $T_{t/4}: \mathcal{G}^{t/4} \rightarrow \mathcal{G}$, henceforth denoted $\mathcal{C}_{1/4}$. The code \mathcal{C} can now be written $\mathcal{C} = (\mathcal{C}_{1/4})^4$. The hyperplane tester now chooses a random $i \in \{1, 2, 3, 4\}$ and a random value $v_i \in \mathcal{G}^{t/4}$, and reads $T(\mathbf{x})$ for all \mathbf{x} whose restriction to the i -th $(t/4)$ -tuple of coordinates is equal to v_i .

We can now apply the following theorem due to Viderman [2015]:

Theorem 16.4 [Viderman 2015 (Theorem 4.4)]. Let $\mathcal{C} = C_1 \otimes \dots \otimes C_k$ be the tensor product of $k > 2$ linear codes of relative distances $\delta_1, \dots, \delta_k$. Let M be a string that is ϵ -far from the tensor code \mathcal{C} . Then the restriction of M to a random hyperplane drawn from the above distribution is, in expectation, at least $\left(\epsilon \cdot \frac{\prod_{j=1}^k \delta_j}{2k^2}\right)$ -far from the restricted code $C_1 \otimes \dots \otimes C_{i-1} \otimes C_{i+1} \otimes \dots \otimes C_k$.

Let $S_{1/4} \subset \mathcal{G}^{t/4}$ be a $(1/(t^2 \log |\mathcal{G}|))$ -biased set of cardinality $|S_{1/4}| = \text{poly}(t, \log |\mathcal{G}|)$ (for example, take the restriction of each $\mathbf{y} \in S$ to its first $t/4$ entries). Sample $\mathbf{x} \in \mathcal{G}^{t/4}$ and $\mathbf{y} \in S_{1/4}$ uniformly at random. Run the hyperplane tester for $(\mathcal{C}_{1/4})^4$ on each alleged polynomial $4|\mathcal{G}|$ times, taking all the possibilities for $i \in \{1, 2, 3, 4\}$ and

$v_i = (\mathbf{x} + \beta\mathbf{y})$ for all $\beta \in \mathcal{G}$. By Theorem 16.4, if the polynomial is ϵ_{SOUND} -far from $(\mathcal{C}_{1/4})^4$, then the expected distance on each of the $4|\mathcal{G}|$ tests is $(\epsilon_{\text{SOUND}} \cdot \frac{\prod_{j=1}^k \delta_j}{2k^2}) = \epsilon_{\text{SOUND}}/512$ (since $k = 4$ and $\delta_j \geq 1/2$). By Lemma 2.4, the average over all the tests is, with probability $1 - o(1)$, within $\pm o(1)$ of this expectation.

So far the tester requires $(1/4 + o(1)) \log_2 n$ random bits to sample a hyperplane, but has prohibitively high query complexity: $n^{3/4+o(1)}$. In order to drive down the query complexity, we can recurse on the test using fresh $(1/4 + o(1)) \log_2 n$ random bits: we re-partition $(\mathcal{C}_{1/4})^3$ as a tensor product of three linear codes, and reapply Videman's theorem. Our query complexity is now $n^{1/2+o(1)}$, the total random-bits complexity is $(1/2 + o(1)) \log_2 n$, and the robustness guarantee is $(\epsilon_{\text{SOUND}} \cdot \frac{2^{-4}}{2 \cdot 4^2} \cdot \frac{2^{-3}}{2 \cdot 3^2}) - o(1) > \epsilon_{\text{SOUND}}/10^5$.

Finally, our tester accepts iff the average distance across all $(t/2)$ -dimensional subspaces from the nearest low-degree polynomial is less than $2\sqrt{\epsilon_{\text{COMPLETE}}} < \epsilon_{\text{SOUND}}/(2 \cdot 10^5)$. Thus our tester accepts with high probability whenever $\widehat{\Psi}$ is $\epsilon_{\text{COMPLETE}}$ -close to the true Ψ , and rejects with high probability whenever $\widehat{\Psi}$ is ϵ_{SOUND} -far from any low-degree polynomial. (Testing $\widehat{T}_u, \widehat{\Phi}, \widehat{\Xi}, \widehat{\Psi}', \widehat{\Phi}', \widehat{\Xi}'$ follows analogously.)

16.3.8 Summary

The verifier, on input $\widehat{\Pi} = (\widehat{E}, \widehat{C}, \widehat{\pi})$, does the following.

Randomness

In a preprocessing step, the verifier fixes in advance a $(1/(t^2 \log |\mathcal{G}|))$ -biased set S of cardinality $|S| = \text{poly}(t, \log |\mathcal{G}|)$. Let $S_{1/4}$ denote the restriction of S to the first $t/4$ coordinates.

The tester then:

- draws \mathbf{x} uniformly at random from $\mathcal{G}^{t/2}$, and \mathbf{y} uniformly at random from S ; and
- reusing the same randomness, draws $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{G}^{t/4}$ and $\mathbf{y}_1, \mathbf{y}_2 \in S_{1/4}$.

Queries

The verifier reads \widehat{T}_u on all vectors $\mathbf{z} \in \mathcal{G}^t$ such that one of $\{(z_1, \dots, z_{t/2}), (z_2, \dots, z_{t/2+1})\}$ is of the form $(\mathbf{x} + \beta\mathbf{y})$ for some $\beta \in \mathcal{G}$. It also reads $\widehat{\Psi}, \widehat{\Phi}, \widehat{\Xi}, \widehat{\Psi}', \widehat{\Phi}', \widehat{\Xi}'$ on all $\mathbf{z} \in \mathcal{G}^t$ such that $(z_1, \dots, z_{t/2}) = (\mathbf{x} + \beta\mathbf{y})$ for some $\beta \in \mathcal{G}$ (denote those vectors $R_{\mathbf{x},\mathbf{y}}$). Similarly, it also reads Ψ', Φ', Ξ' on all $\mathbf{z} \in \mathcal{G}^t$ such that $(z_{t/2+1}, \dots, z_t) = (\mathbf{x} + \beta\mathbf{y})$ (denote those vectors $R'_{\mathbf{x},\mathbf{y}}$). The union of all those subspaces (i.e., $R_{\mathbf{x},\mathbf{y}} \cup R'_{\mathbf{x},\mathbf{y}}$) forms G_{PCP} , while $G_{\text{SAMPLE}} \triangleq R_{\mathbf{x},\mathbf{y}}$.

For each choice of $i_1 \neq i_2 \in \{1, 2, 3, 4\}$ and each $\beta_1, \beta_2 \in \mathcal{G}$, the verifier reads each polynomial on all vectors $\mathbf{z} \in \mathcal{G}^t$ such that the i_1 -th $(t/4)$ -tuple of \mathbf{z} is equal to $\mathbf{x}_1 + \beta_1 \mathbf{y}_1$, and the i_2 -th $(t/4)$ -tuple of \mathbf{z} is equal to $\mathbf{x}_2 + \beta_2 \mathbf{y}_2$. The union of all those subspaces forms G_{LTC} .

Finally, the verifier also reads the entire encoding $C(u)$ of the counter of u . Given $C(u)$, the verifier also picks q_{CRITICAL} vectors $\mathbf{w}^1, \dots, \mathbf{w}^{q_{\text{CRITICAL}}} \in \mathcal{G}^{t/2}$ that correspond to each of the critical bits, and reads T_u on all vectors $\mathbf{z} \in \mathcal{G}^t$ such that $(z_1, \dots, z_{t/2}) = (\mathbf{w}^i + \beta(\mathbf{x} + \beta \mathbf{y}))$ for some $\beta, \gamma \in \mathcal{G}$ and $i \in \{1, \dots, q_{\text{CRITICAL}}\}$. Those are G_{CRITICAL} .

Test

The verifier uses the hyperplane tester from Section 16.3.7 to verify that all the polynomials are indeed close to low-degree polynomials. (If not, it suffices to return an error message and abort.)

Correct & Decode

For each subspace queried, the verifier finds the unique low-degree polynomials $T_u, \Psi, \Phi, \Xi, \Psi', \Phi', \Xi'$ that are close to the values read.

Verify

The verifier checks that the corrected polynomials $T_u, \Psi, \Phi, \Xi, \Psi', \Phi', \Xi'$ satisfy (16.1), (16.2), (16.3), and (16.4) on all of $R_{\mathbf{x}, \mathbf{y}}$. It also checks that (16.6) (and its analogues for Φ', Ξ') are satisfied on all of $R'_{\mathbf{x}, \mathbf{y}}$.

Estimate Distance

The verifier also computes the Hamming distance between the values of $\widehat{T}_u, \widehat{\Psi}, \widehat{\Phi}, \widehat{\Xi}, \widehat{\Psi}', \widehat{\Phi}', \widehat{\Xi}'$ and $T_u, \Psi, \Phi, \Xi, \Psi', \Phi', \Xi'$ on $R_{\mathbf{x}, \mathbf{y}}$.

Extrapolate

Given $C(u)$, the verifier can locally reconstruct $C(S^{\text{LOCAL}}(u))$. Additionally, after the verifier has decoded the values of T_u at the critical bits, it also knows the current and new values of all the entries that change between u and $S^{\text{LOCAL}}(u)$. Since $E(\cdot)$ is a linear encoding, the verifier can also completely construct the difference vector $E(S^{\text{LOCAL}}(u)) - E(u)$. In other words, given the value of $T_u(\mathbf{x})$ for any \mathbf{x} , the verifier can locally compute $T_{S^{\text{LOCAL}}(u)}(\mathbf{x})$.

The verifier can also use the ideas from Section 16.3.6 to construct the values of all the polynomials $(\Psi, \Phi, \Xi, \Psi', \Phi', \Xi')$ for $S^{\text{LOCAL}}(u)$, on all of $R_{\mathbf{x}, \mathbf{y}}$.

For $P^{\text{LOCAL}}(u)$, the verifier must first check that u is not the special vertex u_0 . This is easy to do since T_u and the respective T_{u_0} are both low-degree polynomials and must differ (if $u \neq u_0$) on a constant fraction of their entries. If indeed $u \neq u_0$,

recovering information about $P^{\text{LOCAL}}(u)$ is completely analogous to the procedure for $S^{\text{LOCAL}}(u)$. Otherwise, we simply have $P^{\text{LOCAL}}(u_0) = u_0$.

16.4 Polymatrix WeakNash

In this section we prove our main technical result:

Theorem 16.5 There exists a constant $\epsilon_{\text{NASH}} > 0$, such that there is a $2^{n^{1/2+o(1)}}$ -time reduction from END-OF-A-LINE to finding an $(\epsilon_{\text{NASH}}, \epsilon_{\text{NASH}})$ -WELL-SUPPORTED-WEAKNASH of a complete bipartite polymatrix game between $n^{1/2+o(1)}$ players with $2^{n^{1/2+o(1)}}$ actions each; the payoffs in each bimatrix subgame are in $[0, 1/n_B]$, $[0, 1/n_A]$, where n_A, n_B denote the number of players on each side of the bipartite graph.

We construct a subexponential polymatrix game that implements a variant of the hard Brouwer function constructed in Section 4.2. The new Brouwer function is very similar, but we make the following modifications:

- We use the vertices of a LOCAL END-OF-A-LINE instance instead of END-OF-A-LINE.
- Instead of encoding each LOCAL END-OF-A-LINE vertex u with an arbitrary error-correcting code, the Brouwer vertices correspond to the holographic proofs $\Pi(u)$ from Section 16.3.
- For convenience of notation, instead of maintaining m coordinates (which we expect to be identical anyway) for the auxiliary compute-vs-copy bit and the special direction, we keep just one coordinate for each, but give each a constant relative weight.
- In particular, the “first” Brouwer line segment goes from $\mathbf{z}_2 \triangleq (\Pi(u_0), \Pi(u_0), 0, 2)$ to $(\Pi(u_0), \Pi(u_0), 0, 0)$.

The rest of the proof proceeds as follows: In Section 16.4.1, we define the strategy space for the players and relate it to the holographic proof system from Section 16.3. In Section 16.4.2, we introduce some important notation used throughout this section. In the next two subsections (16.4.3 and 16.4.4), we define the players’ payoffs, and implement an imitation gadget, assuming that the Brouwer function f can be locally computed. In Section 16.4.5, we translate the PCP guarantees from Proposition 16.2 to guarantees about approximate equilibria of our polymatrix game. This completes the setup of the argument. The rest of this section (Sections 16.4.6 to 16.4.9) shows that we can indeed locally compute f in the different regions where it is defined (outside the picture, close to a Brouwer line segment, etc.).

16.4.1 Players and Strategies

Let \mathcal{R}_{PCP} denote the set of random strings used by the PCP verifier in Proposition 16.2. For each $r \in \mathcal{R}_{\text{PCP}}$ we construct a player (A, r) on one side of the bipartite game. On the opposite side, we construct a player (B, \mathbf{q}) for every $\mathbf{q} \in \mathcal{G}^{t/2}$. We refer to the respective sides of the bipartite game as Alice’s players and Bob’s players. (In Section 16.5 we will construct a bimatrix game between Alice and Bob where each player “controls” the vertices on her or his side of the polymatrix game.)

Each player has $(3/\epsilon_{\text{PRECISION}} + 1)^{n^{1/2+o(1)+2}}$ actions. We think of each action as an ordered tuple of:

- two vectors in $[-1, 2]^{n^{1/2+o(1)}}$, where the interval $[-1, 2]$ is discretized into $\{-1, -1 + \epsilon_{\text{PRECISION}}, \dots, 2 - \epsilon_{\text{PRECISION}}, 2\}$; and
- two more variables in $[-1, 2]$.

Recall that each Brouwer vertex corresponds to a pair (u, v) of either identical or consecutive vertices from the LOCAL END-OF-A-LINE instance. For player (A, r) , the two vectors allegedly correspond to bits read from two holographic proofs à la Section 16.3, $(\Pi(u), \Pi(v))$ by the PCP verifier with random string r . For player (B, \mathbf{q}) the vectors represent (a binary encoding of) the counters $C(u), C(v)$, and the entries of $E(u), \pi(u), E(v), \pi(v)$ on a $(t/2)$ -dimensional subspace $G^{(B, \mathbf{q})} \subset \mathcal{G}^t$ to be specified below.

The additional two variables represent the compute-vs-copy bit and the special direction.

In addition to the assignments described above, each player on Alice’s side has an additional choice among $(3/\epsilon_{\text{PRECISION}} + 1)^{n^{1/2+o(1)}}$ actions. These additional actions allegedly correspond to $2q_{\text{CRITICAL}}$ (affine, $(t/2 + 2)$ -dimensional) subspaces from which we can decode the q_{CRITICAL} critical bits of each of the vertices. (We know which entries are represented by reading the “counter” part of the proof.)

Bob’s Subspace $G^{(B, \mathbf{q})}$

Let $G_{\text{ALL}}^{(A, r)} \subset \mathcal{G}^t$ denote the set of entries of $E(\cdot), \pi(\cdot)$ read by the PCP verifier with random string r . Recall that the PCP verifier picks a set of $(t/2)$ -dimensional axis-parallel affine subspaces; each subspace is an affine shift of a linear subspace, spanned by one of a constant number of $(t/2)$ -tuples of standard basis vectors. We pick $G^{(B, \mathbf{q})}$ so that it is spanned by vectors that are linearly independent of all those tuples. In particular, $G^{(B, \mathbf{q})}$ intersects each of (A, r) ’s subspaces at exactly one point.

For any $(t/2)$ -tuple of standard basis vectors of \mathcal{G}^t , a uniformly random $(t/2)$ -tuple of new vectors in \mathcal{G}^t completes it to a basis with high probability. To see this, consider adding the new vectors one by one; at each step, there is at least a $(1 - 1/|\mathcal{G}|)$ -probability that it is linearly independent of the previously chosen vectors; the claim follows by union bound and $t = o(|\mathcal{G}|)$. Therefore, a random $(t/2)$ -tuple of vectors will also, with high probability, be simultaneously linearly independent of each of the relevant tuples of standard basis vectors. Fix any such $(t/2)$ -tuple (there are also efficient ways to deterministically find such a tuple, but in quasi-polynomial time we can simply brute-force enumerate through all tuples).

Finally, we let $G^{(B, \mathbf{0}_{t/2})}$ be the linear subspace spanned by those vectors; for general $\mathbf{q} \in \mathcal{G}^{t/2}$, we let $G^{(B, \mathbf{q})}$ be the same subspace shifted by $(\mathbf{q}, \mathbf{0}_{t/2})$.

16.4.2 Notation

We introduce some necessary additional notation for formally handling vectors in this section. The notation for (A, r) 's assignment to the critical bits is analogous in spirit but more subtle to formalize; we come back to it in Section 16.4.3.

Definition 16.3 We say that a player is *happy* if every strategy on her support is ϵ_{NASH} -optimal.

Rounding

For $x \in \mathbb{R}$, let $v(x) \triangleq \begin{cases} 1 & x > 1/2 \\ 0 & \text{otherwise} \end{cases}$ denote the rounding of x to $\{0, 1\}$; for $\mathbf{x} \in \mathbb{R}^n$, we compute $v(\mathbf{x})$ coordinate-wise.

Coordinates

Let $M \triangleq \{1, \dots, 2m + 2\}$ denote the set of all coordinates. We partition M into: $M_1 \triangleq \{1, \dots, m\}$ (the first proof), $M_2 \triangleq \{m + 1, \dots, 2m\}$ (the second proof), $M_3 \triangleq \{2m + 1\}$ (the compute-vs-copy bit), and $M_4 \triangleq \{2m + 2\}$ (the special direction). We will also use the abbreviations $M_{1,2} \triangleq M_1 \cup M_2$ and $M_{3,4} \triangleq M_3 \cup M_4$.

Fix some player (A, r) on Alice's side; we further detail (A, r) 's strategy space as follows. Let $G_{\text{SAMPLE}}^{(A, r)}, G_{\text{PCP}}^{(A, r)}, G_{\text{LTC}}^{(A, r)} \subset \mathcal{G}^t$ denote the sets of entries of $E(\cdot), \pi(\cdot)$ read by the PCP verifier with random string r from the alleged proofs. Let $I_{\text{SAMPLE}}^{(A, r)}, I_{\text{PCP}}^{(A, r)}, I_{\text{LTC}}^{(A, r)} \subset M_{1,2}$ denote the indices of the corresponding bits, respectively. For player (B, \mathbf{q}) on Bob's side, let $I^{(B, \mathbf{q})} \subset M_{1,2}$ denote the indices of the bits corresponding to $G^{(B, \mathbf{q})}$. We also let $K \subset M_{1,2}$ (K for Kounter) denote the set of indices that correspond to the encodings of counters in both proofs. (Notice that K is the same for every player on both sides.)

Finally, we will be particularly interested in the following sets:

$$J^{(A,r)} \triangleq K \cup I_{\text{SAMPLE}}^{(A,r)} \cup M_{3,4}$$

$$J^{(B,\mathbf{q})} \triangleq K \cup I^{(B,\mathbf{q})} \cup M_{3,4}.$$

Partial Vectors

Player (A, r) 's strategy \mathbf{a} naturally corresponds to a *partial vector* $\mathbf{x}_{\text{ALL}}^{(\mathbf{a})} \in ([-1, 2] \cup \{\perp\})^M$. For $i \in K \cup I_{\text{PCP}}^{(A,r)} \cup I_{\text{LTC}}^{(A,r)} \cup M_{3,4}$, we set the i -th coordinate of $\mathbf{x}_{\text{ALL}}^{(\mathbf{a})}$ to the value that \mathbf{a} assigns to the corresponding coordinate; for all other i 's, we let $\mathbf{x}_{\text{ALL}}^{(\mathbf{a})} \upharpoonright_i \triangleq \perp$.

In particular, we are interested in the partial vector $\mathbf{x}^{(\mathbf{a})} \in ([-1, 2] \cup \{\perp\})^M$ that describes (A, r) 's partial assignment to $J^{(A,r)}$. For $i \in J^{(A,r)}$, we set the i -th coordinate of $\mathbf{x}^{(\mathbf{a})}$ to the value that \mathbf{a} assigns to the corresponding coordinate; for all other i 's, we let $\mathbf{x}^{(\mathbf{a})} \upharpoonright_i \triangleq \perp$.

For action \mathbf{b} taken by player (B, \mathbf{q}) , we can simply define one partial vector $\mathbf{x}^{(\mathbf{b})} \in ([-1, 2] \cup \{\perp\})^M$, where $\mathbf{x}^{(\mathbf{b})} \upharpoonright_i \triangleq \perp$ for $i \notin J^{(B,\mathbf{q})}$.

Partial Norms

All norms in this section are 2-norm: we henceforth drop the 2 subscript from $\|\cdot\|_2$. Instead, we will use the subscript to denote the subset of coordinates over which we want to compute the norm. For example,

$$\|\mathbf{x}\|_K^2 \triangleq \mathbb{E}_{i \in K} [x_i^2].$$

In fact, we are often interested in expectation with respect to a non-uniform distribution; for example:

$$\|\mathbf{x}\|_{J^{(A,r)}}^2 \triangleq \frac{1}{4} \mathbb{E}_{i \in I_{\text{SAMPLE}}^{(A,r)}} [x_i^2] + \frac{1}{4} \mathbb{E}_{i \in K} [x_i^2] + \frac{1}{2} \mathbb{E}_{i \in M_{3,4}} [x_i^2]$$

$$\|\mathbf{x}\|_{J^{(B,\mathbf{q})}}^2 \triangleq \frac{1}{4} \mathbb{E}_{i \in I^{(B,\mathbf{q})}} [x_i^2] + \frac{1}{4} \mathbb{E}_{i \in K} [x_i^2] + \frac{1}{2} \mathbb{E}_{i \in M_{3,4}} [x_i^2].$$

The distribution is set such that each part (e.g., K , $I_{\text{PCP}}^{(A,r)}$, $I^{(B,\mathbf{q})}$, or M_1, \dots, M_4) receives an equal weight (up to constant factors). The distribution will henceforth be implicit, and only the subset of coordinates will be explicit.

This notation also allows us to talk about distances between partial vectors, e.g.,

$$\|\mathbf{x}^{(\mathbf{a})} - \mathbf{x}^{(\mathbf{b})}\|_{J^{(A,r)} \cap J^{(B,\mathbf{q})}}^2 \triangleq \frac{1}{4} \|\mathbf{x}^{(\mathbf{a})} - \mathbf{x}^{(\mathbf{b})}\|_K^2$$

$$+ \frac{1}{2} \|\mathbf{x}^{(\mathbf{a})} - \mathbf{x}^{(\mathbf{b})}\|_{M_{3,4}}^2 + \frac{1}{4} \|\mathbf{x}^{(\mathbf{a})} - \mathbf{x}^{(\mathbf{b})}\|_{I_{\text{SAMPLE}}^{(A,r)} \cap I^{(B,\mathbf{q})}}^2.$$

Aggregate Vectors

Let \mathcal{A} denote the joint (product) distribution over all Alice's players' mixed strategies, and define \mathcal{B} analogously for Bob's players. We define $\mathbf{x}^{(A)}, \mathbf{x}^{(B)} \in [-1, 2]^M$ to be the weighted entry-wise average of $\mathbf{x}^{(a)}, \mathbf{x}^{(b)}$, respectively, where for each coordinate we only take the expectation over vectors for which it is assigned a real (non- \perp) value.

16.4.3 Alice's Imitation Gadget

Our first goal is to incentivize player (A, r) to imitate $\mathbf{x}^{(B)}$ on her assigned coordinates. We break up her utility on each bimatrix subgame as follows:⁶

$$U^{(A,r)} \triangleq U_{\text{SAMPLE}}^{(A,r)} + U_{\text{PCP}}^{(A,r)} + U_{\text{LTC}}^{(A,r)} + U_{\text{COUNTER}}^{(A,r)} + U_{M_{3,4}}^{(A,r)} + U_{\text{CRITICAL}}^{(A,r)}.$$

For each player (B, \mathbf{q}) on Bob's side, we simply define the first five parts to incentivize copying (B, \mathbf{q}) 's strategy on each shared coordinate. Let \mathbf{a}, \mathbf{b} denote (A, r) and (B, \mathbf{q}) respective strategies; then

$$U_{\text{SAMPLE}}^{(A,r)} \triangleq - \left\| \mathbf{x}_{\text{ALL}}^{(a)} - \mathbf{x}^{(b)} \right\|_{I_{\text{SAMPLE}}^{(A,r)} \cap I^{(B,\mathbf{q})}}^2$$

$$U_{\text{PCP}}^{(A,r)} \triangleq - \left\| \mathbf{x}_{\text{ALL}}^{(a)} - \mathbf{x}^{(b)} \right\|_{I_{\text{PCP}}^{(A,r)} \cap I^{(B,\mathbf{q})}}^2$$

$$U_{\text{LTC}}^{(A,r)} \triangleq - \left\| \mathbf{x}_{\text{ALL}}^{(a)} - \mathbf{x}^{(b)} \right\|_{I_{\text{LTC}}^{(A,r)} \cap I^{(B,\mathbf{q})}}^2$$

$$U_{\text{COUNTER}}^{(A,r)} \triangleq - \left\| \mathbf{x}_{\text{ALL}}^{(a)} - \mathbf{x}^{(b)} \right\|_K^2$$

$$U_{M_{3,4}}^{(A,r)} \triangleq - \left\| \mathbf{x}_{\text{ALL}}^{(a)} - \mathbf{x}^{(b)} \right\|_{M_{3,4}}^2.$$

Notice that

$$\begin{aligned} & \frac{1}{8} U_{\text{PCP}}^{(A,r)} + \frac{1}{8} U_{\text{LTC}}^{(A,r)} + \frac{1}{4} U_{\text{COUNTER}}^{(A,r)} + \frac{1}{2} U_{M_{3,4}}^{(A,r)} = \\ & - \left\| \mathbf{x}_{\text{ALL}}^{(a)} - \mathbf{x}^{(b)} \right\|_{(K \cup I_{\text{PCP}}^{(A,r)} \cup I_{\text{LTC}}^{(A,r)} \cup M_{3,4}) \cap I^{(B,\mathbf{q})}}^2, \end{aligned}$$

6. For convenience of notation, we define bimatrix subgames with payoffs in $[-54/n_B, 0]$, $[-36/n_A, 0]$. Payoffs $[0, 1/n_B]$, $[0, 1/n_A]$ can be obtained by scaling and shifting.

whereas

$$\frac{1}{4}U_{\text{SAMPLE}}^{(A,r)} + \frac{1}{4}U_{\text{COUNTER}}^{(A,r)} + \frac{1}{2}U_{M_{3,4}}^{(A,r)} = -\|\mathbf{x}^{(\mathbf{a})} - \mathbf{x}^{(\mathbf{b})}\|_{J^{(A,r)} \cap J^{(B,\mathbf{q})}}^2.$$

As we mentioned earlier, we also want (A, r) to encode the critical bits for each LOCAL END-OF-A-LINE vertex. The indices of this encoding depend on both:

- the indices of the critical bits—which depend on the “counter” of the two LOCAL END-OF-A-LINE vertices, allegedly encoded in both \mathbf{a} and \mathbf{b} ; and
- the directions in $\mathcal{G}^{t/2}$ that the PCP verifier uses to locally decode those bits—which depend on the random string r .

We let $\mathcal{G}_{\text{CRITICAL}}^{(A,r)}(\mathbf{a}) \subset \mathcal{G}^t$ denote the set of additional entries that encode the critical bits, as induced by (A, r) 's strategy $\mathbf{x}^{(\mathbf{a})} \upharpoonright_K$ and random string r ; let $I_{\text{CRITICAL}}^{(A,r)}(\mathbf{a}) \subset M_{1,2}$ denote the corresponding bits' indices. In order to define (A, r) 's utilities, it is actually more convenient to use the critical bits determined by (B, \mathbf{q}) 's \mathbf{b} , henceforth denoted $I_{\text{CRITICAL}}^{(A,r)}(\mathbf{b})$. (The superscript remains (A, r) because we still decode those bits according to random string r .) This added level of indirection prevents (A, r) from manipulating her encoding of the counter to increase her utility from the critical bits. In other words, \mathbf{a} includes an assignment to q_{CRITICAL} affine $(t/2 + 2)$ -dimensional subspaces of \mathcal{G}^t . On each subgame, we subjectively interpret \mathbf{a} 's assignment as a partial vector $\mathbf{x}_{\text{CRITICAL}}^{(\mathbf{a})}$ with values in $[-1, 2]$ on $I_{\text{CRITICAL}}^{(A,r)}(\mathbf{b})$. In particular, we define

$$U_{\text{CRITICAL}}^{(A,r)} \triangleq -\|\mathbf{x}_{\text{CRITICAL}}^{(\mathbf{a})} - \mathbf{x}^{(\mathbf{b})}\|_{I_{\text{CRITICAL}}^{(A,r)}(\mathbf{b}) \cap I^{(B,\mathbf{q})}}^2.$$

For some choices of \mathbf{b} , the induced assignment to the counter may be ambiguous or just far from every codeword. However, in that case we are probably far from every Brouwer line segment, so we don't care what (A, r) assigns to the critical bits.

Claim 16.1 For every happy player (A, r) on Alice's side and any action \mathbf{a} in her support, for each $I \in \{I_{\text{SAMPLE}}^{(A,r)}, I_{\text{PCP}}^{(A,r)}, I_{\text{LTC}}^{(A,r)}\}$

$$\|\mathbf{x}_{\text{ALL}}^{(\mathbf{a})} - \mathbf{x}^{(B)}\|_I^2 = O(\epsilon_{\text{NASH}}).$$

Proof Recall that (A, r) 's total utility from her assignment to $I_{\text{LTC}}^{(A,r)}$ is given by

$$\begin{aligned}
 U_{\text{LTC}}^{(A,r)} &= -\mathbb{E}_{\mathbf{q} \in \mathcal{G}^{t/2}} \mathbb{E}_{\mathbf{b} \sim \mathcal{B}[(B, \mathbf{q})]} \left[\left\| \mathbf{x}_{\text{ALL}}^{(\mathbf{a})} - \mathbf{x}^{(\mathbf{b})} \right\|_{I_{\text{LTC}}^{(A,r)} \cap I^{(B, \mathbf{q})}}^2 \right] \\
 &= -\mathbb{E}_{i \in I_{\text{LTC}}^{(A,r)}} \left[\left(\mathbf{x}_{\text{ALL}}^{(\mathbf{a})} \mid_i - \mathbb{E}_{\mathbf{b} \sim \mathcal{B}[(B, \mathbf{q}(i))]} \left[\mathbf{x}^{(\mathbf{b})} \mid_i \right] \right)^2 \right] \\
 &\quad - \mathbb{E}_{i \in I_{\text{LTC}}^{(A,r)}} \left[\text{Var}_{\mathbf{b} \sim \mathcal{B}[(B, \mathbf{q}(i))]} \left[\mathbf{x}^{(\mathbf{b})} \mid_i \right] \right], \tag{16.7}
 \end{aligned}$$

where $(B, \mathbf{q}(i))$ is the player on Bob's side for which $i \in I^{(B, \mathbf{q}(i))}$ (i.e., $\mathbf{q}(i)$ denotes the value $\mathbf{q} \in \mathcal{G}^{t/2}$ such that bit i is part of the representation of some $\mathbf{g} \in G^{(B, \mathbf{q})}$).

Notice that the second term of (16.7) does not depend on (A, r) 's strategy, so she simply wants to maximize the first term. For every i , she can approximate $\mathbb{E}_{\mathbf{b} \sim \mathcal{B}[(B, \mathbf{q}(i))]} \left[\mathbf{x}^{(\mathbf{b})} \mid_i \right]$ to within $\pm \epsilon_{\text{PRECISION}}$. Therefore, if \mathbf{a} is an ϵ_{NASH} -optimal strategy, we have

$$U_{\text{PCP}}^{(A,r)} \geq -\mathbb{E}_{i \in I_{\text{PCP}}^{(A,r)}} \left[\text{Var}_{\mathbf{b} \sim \mathcal{B}[(B, \mathbf{q}(i))]} \left[\mathbf{x}^{(\mathbf{b})} \mid_i \right] \right] - \underbrace{\epsilon_{\text{NASH}} - \epsilon_{\text{PRECISION}}^2}_{=O(\epsilon_{\text{NASH}})}.$$

An analogous argument works for $I_{\text{SAMPLE}}^{(A,r)}, I_{\text{PCP}}^{(A,r)}$. ■

The following claims follow along the same lines:

Claim 16.2 For every happy player (A, r) on Alice's side and any action \mathbf{a} in her support,

$$\left\| \mathbf{x}^{(\mathbf{a})} - \mathbf{x}^{(\mathcal{B})} \right\|_K^2 = O(\epsilon_{\text{NASH}}).$$

Claim 16.3 For every happy player (A, r) on Alice's side and any action \mathbf{a} in her support,

$$\left\| \mathbf{x}^{(\mathbf{a})} - \mathbf{x}^{(\mathcal{B})} \right\|_{J^{(A,r)}}^2 = O(\epsilon_{\text{NASH}}).$$

Corollary 16.1 Let $(\mathcal{A}, \mathcal{B})$ be an $(\epsilon_{\text{NASH}}, \epsilon_{\text{NASH}})$ -Well-Supported-WeakNash; then

$$\left\| \mathbf{x}^{(\mathcal{A})} - \mathbf{x}^{(\mathcal{B})} \right\|_M^2 = O(\epsilon_{\text{NASH}}).$$

16.4.4 Bob's Imitation Gadget

We would also like to have the players on Bob's side minimize $\|f(\mathbf{x}^{(\mathcal{A})}) - \mathbf{x}^{(\mathcal{B})}\|_{J^{(B, \mathbf{q})}}^2$. In order to implement this in a polymatrix game, we want to locally compute $f(\cdot)$ given access only to the partial information provided by player (A, r) 's strategy \mathbf{a} . In

Section 4.2.3 we argued that we can locally compute $f_i(\mathbf{x})$ using only approximate, partial information about \mathbf{x} . The main goal of this section is to argue that for most $i \in J^{(A,r)}$, we can compute this partial information from \mathbf{a} , approximately and with high probability over r . Before we do that, however, let us assume that we have some estimate $f^{(\mathbf{a})} \upharpoonright_i$ and derive the utility of player (B, \mathbf{q}) .

For any action \mathbf{a} that (A, r) plays, we define a partial target vector $f^{(\mathbf{a})} \in [-1, 2] \cup \{\perp\}^M$ that we would like (B, \mathbf{q}) to imitate. In particular, (B, \mathbf{q}) is incentivized to play a strategy that is close to $f^{(\mathbf{a})}$ on the coordinates where both $f^{(\mathbf{a})}$ and $\mathbf{x}^{(\mathbf{b})}$ are defined. We define

$$U^{(B, \mathbf{q})} \triangleq - \left\| f^{(\mathbf{a})} - \mathbf{x}^{(\mathbf{b})} \right\|_{J^{(A,r)} \cap J^{(B, \mathbf{q})}}^2.$$

Similarly to $\mathbf{x}^{(A)}$ and $\mathbf{x}^{(B)}$, let $f^{(A)} \in [-1, 2]^M$ denote the weighted entry-wise average of $f^{(\mathbf{a})}$, where for each i we take an average over $f^{(\mathbf{a})} \upharpoonright_i$ for all \mathbf{a} in the support of (A, r) , for r such that $i \in J^{(A,r)}$.

Claim 16.4 For every happy player (B, \mathbf{q}) on Bob's side and any action \mathbf{b} in his support,

$$\left\| \mathbf{x}^{(\mathbf{b})} - f^{(A)} \right\|_{J^{(B, \mathbf{q})}}^2 = O(\epsilon_{\text{NASH}}).$$

Proof Analogous to Claim 16.1. ■

Corollary 16.2 Let $(\mathcal{A}, \mathcal{B})$ be an $(\epsilon_{\text{NASH}}, \epsilon_{\text{NASH}})$ -Well-Supported-WeakNash; then

$$\left\| \mathbf{x}^{(B)} - f^{(A)} \right\|_M^2 = O(\epsilon_{\text{NASH}}).$$

For the rest of this section, our goal is to prove that $f^{(A)}$ is also close to $f(\mathbf{x}^{(B)})$ —this would prove that $\mathbf{x}^{(B)}$ is an approximate fixed point of f .

16.4.5 Reading the Holographic Proofs

In this subsection we translate the desiderata from Proposition 16.2 to guarantees about approximate equilibria in our game. Our first task in approximating $f(\mathbf{x}^{(B)})$ is deciding whether $\mathbf{x}^{(B)}$ is close to a Brouwer line segment, close to a Brouwer vertex, or far from both. Close to any Brouwer vertex, both restrictions $\mathbf{x}^{(B)} \upharpoonright_{M_1}$ and $\mathbf{x}^{(B)} \upharpoonright_{M_2}$ are close to some $\Pi(u), \Pi(v)$, for some $u, v \in V^{\text{LOCAL}}$; close to a Brouwer line, at least one of $\mathbf{x}^{(B)} \upharpoonright_{M_1}$ and $\mathbf{x}^{(B)} \upharpoonright_{M_2}$ is close to some $\Pi(u)$.

The first step of this first task is to decide whether $\mathbf{x}^{(B)} \upharpoonright_{M_1}$ is close to some $\Pi(u)$. Formally, we have a *game verifier* that encourages rounding assignments in

$[-1, 2]^M$ to $\{0, 1\}^M$ and making sure they satisfy the PCP verifier. Given that player (A, r) chooses strategy \mathbf{a} , we say that the game verifier *accepts* $\mathbf{a} \upharpoonright_{M_1}$ if:

- $\mathbf{x}^{(\mathbf{a})} \upharpoonright_{M_1}$ is close to binary, i.e.,

$$\left\| \mathbf{x}^{(\mathbf{a})} - v(\mathbf{x}^{(\mathbf{a})} \upharpoonright_{M_1}) \right\|_{J(A,r)}^2 < \sqrt{\epsilon_{\text{COMPLETE}}};$$

- and the PCP verifier accepts the rounded bits it reads from the first proof, $v(\mathbf{x}_{\text{ALL}}^{(\mathbf{a})} \upharpoonright_{M_1})$.

If either of these does not hold, we say that the game verifier *rejects* $\mathbf{a} \upharpoonright_{M_1}$.

The main idea in all the lemmata below is that by the Good Sample desideratum in Proposition 16.2, $G_{\text{SAMPLE}}^{(A,r)}$, $G_{\text{PCP}}^{(A,r)}$, $G_{\text{LTC}}^{(A,r)}$, and $G_{\text{CRITICAL}}^{(A,r)}(\mathbf{a})$ are all representing samples of \mathcal{G}^t . Therefore, if $\mathbf{x}^{(\mathcal{B})} \upharpoonright_{M_1}$ is close to $\{0, 1\}^{M_1}$, then we expect that $\mathbf{x}^{(\mathcal{B})} \upharpoonright_{M_1 \cap I_{\text{PCP}}^{(A,r)}}$ is also close to $\{0, 1\}^{M_1 \cap I_{\text{PCP}}^{(A,r)}}$. By Claim 16.1, we can then also expect that $\mathbf{x}_{\text{ALL}}^{(\mathbf{a})} \upharpoonright_{M_1 \cap I_{\text{PCP}}^{(A,r)}}$ is close to the *same* vector in $\{0, 1\}^{M_1 \cap I_{\text{PCP}}^{(A,r)}}$. Furthermore, if $\mathbf{x}_{\text{ALL}}^{(\mathbf{a})} \upharpoonright_{M_1 \cap I_{\text{PCP}}^{(A,r)}}$ is ϵ -close to some binary vector $v(\mathbf{x}^{(\mathcal{B})} \upharpoonright_{M_1 \cap I_{\text{PCP}}^{(A,r)}})$, then its rounding, $v(\mathbf{x}_{\text{ALL}}^{(\mathbf{a})} \upharpoonright_{M_1 \cap I_{\text{PCP}}^{(A,r)}})$, is also $O(\epsilon)$ -close to $v(\mathbf{x}^{(\mathcal{B})} \upharpoonright_{M_1 \cap I_{\text{PCP}}^{(A,r)}})$: On each coordinate i , if the two binary vectors disagree (if $|v(\mathbf{x}^{(\mathcal{B})} \upharpoonright_i) - v(\mathbf{x}_{\text{ALL}}^{(\mathbf{a})} \upharpoonright_i)| = 1$), then we must have $|v(\mathbf{x}^{(\mathcal{B})} \upharpoonright_i) - \mathbf{x}^{(\mathcal{B})} \upharpoonright_i| + |\mathbf{x}^{(\mathcal{B})} \upharpoonright_i - \mathbf{x}_{\text{ALL}}^{(\mathbf{a})} \upharpoonright_i| \geq 1/2$. Finally, if $v(\mathbf{x}_{\text{ALL}}^{(\mathbf{a})} \upharpoonright_{M_1 \cap I_{\text{PCP}}^{(A,r)}})$ is close to $v(\mathbf{x}^{(\mathcal{B})} \upharpoonright_{M_1 \cap I_{\text{PCP}}^{(A,r)}})$ for most \mathbf{a} , then we can argue that reading $v(\mathbf{x}_{\text{ALL}}^{(\mathbf{a})} \upharpoonright_{M_1 \cap I_{\text{PCP}}^{(A,r)}})$ is almost like sampling $v(\mathbf{x}^{(\mathcal{B})} \upharpoonright_{M_1})$, and then letting an adaptive adversary corrupt a small fraction of the entries. Therefore, we can use the guarantees of the PCP verifier with robust soundness/decoding/completeness.

16.4.5.1 Soundness

Lemma 16.1 Let $(\mathcal{A}, \mathcal{B})$ be an $(\epsilon_{\text{NASH}}, \epsilon_{\text{NASH}})$ -Well-Supported-WeakNash. If $\|\mathbf{x}^{(\mathcal{B})} - \Pi(u)\|_{M_1}^2 > 4\epsilon_{\text{SOUND}}$ for every $u \in V^{\text{LOCAL}}$, then for a $(1 - O(\epsilon_{\text{NASH}}))$ -fraction of r 's, the game verifier rejects $\mathbf{a} \upharpoonright_{M_1}$ for every strategy \mathbf{a} in the support of (A, r) .

Proof Suppose that $\|\mathbf{x}^{(\mathcal{B})} - \Pi(u)\|_{M_1}^2 > 4\epsilon_{\text{SOUND}}$. Then by triangle inequality we have that at least one of the following holds:

$$\left\| \mathbf{x}^{(\mathcal{B})} - v(\mathbf{x}^{(\mathcal{B})}) \right\|_{M_1}^2 > \epsilon_{\text{SOUND}}^3, \quad (16.8)$$

$$\left\| v(\mathbf{x}^{(\mathcal{B})}) - \Pi(u) \right\|_{M_1}^2 > \epsilon_{\text{SOUND}}, \quad (16.9)$$

Far from $\{0, 1\}^{M_1}$

Assume first that (16.8) holds. We can further break (16.8) into its K -component and (E, π) -component:

$$\left\| \mathbf{x}^{(B)} - v(\mathbf{x}^{(B)}) \right\|_{M_1}^2 = \frac{1}{2} \left\| \mathbf{x}^{(B)} - v(\mathbf{x}^{(B)}) \right\|_{M_1 \cap K}^2 + \frac{1}{2} \left\| \mathbf{x}^{(B)} - v(\mathbf{x}^{(B)}) \right\|_{M_1 \setminus K}^2. \quad (16.10)$$

When restricted to a random $r \in \mathcal{R}_{\text{PCP}}$, we can learn the first term exactly, and estimate the second term via the proxy:

$$\left\| \mathbf{x}^{(B)} - v(\mathbf{x}^{(B)}) \right\|_{M_1 \cap I_{\text{SAMPLE}}^{(A,r)}}^2.$$

By the Good Sample guarantee from Proposition 16.2, the $(M_1 \cap I_{\text{SAMPLE}}^{(A,r)})$ -restricted distance between $\mathbf{x}^{(B)}$ and $v(\mathbf{x}^{(B)})$ concentrates (to within $\pm o(1)$, with high probability) around its expectation (the last term of (16.10)). Therefore, with probability $(1 - o(1))$ over the choice of $r \in \mathcal{R}_{\text{PCP}}$,

$$\left\| \mathbf{x}^{(B)} - v(\mathbf{x}^{(B)}) \right\|_{J^{(A,r)}}^2 \geq \frac{1}{4} \left\| \mathbf{x}^{(B)} - v(\mathbf{x}^{(B)}) \right\|_{\underbrace{M_1 \cap J^{(A,r)}}_{=M_1 \cap (I_{\text{SAMPLE}}^{(A,r)} \cup K)}}^2 = \Omega(\epsilon_{\text{SOUND}}^3).$$

For any such r , if player (A, r) is also happy, we have by Claim 16.3 that for every strategy \mathbf{a} in her support,

$$\left\| \mathbf{x}^{(\mathbf{a})} - v(\mathbf{x}^{(\mathbf{a})}) \right\|_{J^{(A,r)}}^2 = \Omega(\epsilon_{\text{SOUND}}^3) \gg \sqrt{\epsilon_{\text{COMPLETE}}}. \quad (16.11)$$

Close to $\{0, 1\}^{M_1}$, but Far from Every $\Pi(u)$

Alternatively, assume that $\mathbf{x}^{(B)}|_{M_1}$ is close to $\{0, 1\}^{M_1}$, but (16.9) holds: we have a binary vector $v(\mathbf{x}^{(B)})|_{M_1}$ that is ϵ_{SOUND} -far from a valid proof.

We assume wlog that (16.8) is false, namely

$$\left\| \mathbf{x}^{(B)} - v(\mathbf{x}^{(B)}) \right\|_{M_1}^2 \leq \epsilon_{\text{SOUND}}^3.$$

By the Good Sample guarantee in Proposition (16.2), we have that for $(1 - o(1))$ -fraction of r 's, also

$$\left\| \mathbf{x}^{(B)} - v(\mathbf{x}^{(B)}) \right\|_{M_1 \cap I_{\text{PCP}}^{(A,r)}}^2 = O(\epsilon_{\text{SOUND}}^3). \quad (16.12)$$

By Claim 16.1, for every happy (A, r) and every strategy \mathbf{a} in her support,

$$\left\| \mathbf{x}_{\text{ALL}}^{(\mathbf{a})} - \mathbf{x}^{(B)} \right\|_{M_1 \cap I_{\text{PCP}}^{(A,r)}}^2 = O(\epsilon_{\text{SOUND}}^3).$$

Therefore, combining with (16.12), we have:

$$\left\| v(\mathbf{x}_{\text{ALL}}^{(\mathbf{a})}) - v(\mathbf{x}^{(\mathcal{B})}) \right\|_{M_1 \cap I_{\text{PCP}}^{(A,r)}}^2 = O(\epsilon_{\text{SOUND}}^3),$$

and analogous arguments hold for $M_1 \cap K$ and $M_1 \cap I_{\text{LTC}}^{(A,r)}$. Therefore, by the Robust Soundness guarantee in Proposition 16.2, the game verifier rejects every $\mathbf{a} \upharpoonright_{M_1}$ for all but an $o(1)$ -fraction of happy (A, r) 's. ■

16.4.5.2 Completeness

Lemma 16.2 Let $(\mathcal{A}, \mathcal{B})$ be an $(\epsilon_{\text{NASH}}, \epsilon_{\text{NASH}})$ -Well-Supported-WeakNash. If $\|\mathbf{x}^{(\mathcal{B})} - \Pi(u)\|_{M_1}^2 \leq \frac{1}{4}\epsilon_{\text{COMPLETE}}$ for some $u \in V^{\text{LOCAL}}$, then for a $(1 - O(\epsilon_{\text{NASH}}))$ -fraction of r 's, and every strategy \mathbf{a} in the support of (A, r) , the game verifier accepts $\mathbf{a} \upharpoonright_{M_1}$.

Proof First, notice that the premise implies that

$$\left\| v(\mathbf{x}^{(\mathcal{B})}) - \Pi(u) \right\|_{M_1}^2 \leq \epsilon_{\text{COMPLETE}} \quad (16.13)$$

(since for any $i \in M_1$, $\|v(\mathbf{x}^{(\mathcal{B})}) - \Pi(u)\|_i^2 \leq 4\|\mathbf{x}^{(\mathcal{B})} - \Pi(u)\|_i^2$).

Furthermore, we also have that

$$\left\| \mathbf{x}^{(\mathcal{B})} - v(\mathbf{x}^{(\mathcal{B})}) \right\|_{M_1}^2 \leq \left\| \mathbf{x}^{(\mathcal{B})} - \Pi(u) \right\|_{M_1}^2 \leq \frac{1}{4}\epsilon_{\text{COMPLETE}}.$$

Therefore, by the Good Sample guarantee from Proposition 16.2, we also have that for a $(1 - o(1))$ -fraction of r 's,

$$\left\| \mathbf{x}^{(\mathcal{B})} - v(\mathbf{x}^{(\mathcal{B})}) \right\|_{M_1 \cap I_{\text{PCP}}^{(A,r)}}^2 = O(\epsilon_{\text{COMPLETE}}).$$

Combining with Claim 16.1, we get that whenever (A, r) is also happy,

$$\left\| \mathbf{x}_{\text{ALL}}^{(\mathbf{a})} - v(\mathbf{x}^{(\mathcal{B})}) \right\|_{M_1 \cap I_{\text{PCP}}^{(A,r)}}^2 = O(\epsilon_{\text{COMPLETE}}).$$

Thus, as in (16.13), we have

$$\left\| v(\mathbf{x}_{\text{ALL}}^{(\mathbf{a})}) - v(\mathbf{x}^{(\mathcal{B})}) \right\|_{M_1 \cap I_{\text{PCP}}^{(A,r)}}^2 = O(\epsilon_{\text{COMPLETE}}).$$

Analogous arguments hold for $M_1 \cap K$ and $M_1 \cap I_{\text{LTC}}^{(A,r)}$. Therefore, for a $(1 - O(\epsilon_{\text{NASH}}))$ -fraction of (A, r) 's, sampling $v(\mathbf{x}_{\text{ALL}}^{(\mathbf{a})} \upharpoonright_{M_1})$ satisfies the distance criteria for the Completeness and Robust Completeness in Proposition 16.2. ■

16.4.5.3 Decoding

Lemma 16.3 Let $(\mathcal{A}, \mathcal{B})$ be an $(\epsilon_{\text{NASH}}, \epsilon_{\text{NASH}})$ -Well-Supported-WeakNash. If $\|\mathbf{x}^{(\mathcal{B})} - \Pi(u)\|_{M_1}^2 \leq \frac{1}{4}\epsilon_{\text{DECODE}}$ for some $u \in V^{\text{LOCAL}}$, then for a $(1 - O(\epsilon_{\text{NASH}}))$ -fraction of r 's, and every strategy \mathbf{a} in the support of (A, r) :

Error Correction. The game verifier can compute the restriction of $\Pi(u)$ to $M_1 \cap J^{(A,r)}$ correctly.

Critical Bits. The induced assignment on the critical bits is approximately correct:

$$\left\| \mathbf{x}_{\text{CRITICAL}}^{(\mathbf{a})} - \Pi(u) \right\|_{M_1 \cap I_{\text{CRITICAL}}^{(A,r)}(\mathbf{a})}^2 = O(\epsilon_{\text{DECODE}}).$$

Extrapolation. The game verifier can also compute the restrictions of $\Pi(S^{\text{LOCAL}}(u))$ and $\Pi(P^{\text{LOCAL}}(u))$ to $M_1 \cap J^{(A,r)}$ correctly.

Proof By analogous argument to Lemma 16.2 (replace $\epsilon_{\text{COMPLETE}}$ with ϵ_{DECODE}), we have that, for a $(1 - O(\epsilon_{\text{NASH}}))$ -fraction of (A, r) 's, sampling $\nu(\mathbf{x}_{\text{ALL}}^{(\mathbf{a})} |_{M_1})$ satisfies the distance criteria for the Decoding and Robust Decoding in Proposition 16.2. Thus the decoding of $\Pi(u)$ on $M_1 \cap J^{(A,r)}$ follows from the Error-Correction on a Sample desideratum of Proposition 16.2.

Critical Bits

By the premise, we have that

$$\left\| \mathbf{x}^{(\mathcal{B})} - C(u) \right\|_{M_1 \cap K}^2 = O(\epsilon_{\text{DECODE}}). \quad (16.14)$$

Thus, by Corollary 16.2, also

$$\left\| f^{(A)} - C(u) \right\|_{M_1 \cap K}^2 = O(\epsilon_{\text{DECODE}}).$$

By Claims 16.2 and 16.4, for every ϵ_{NASH} -optimal \mathbf{a} and \mathbf{b} we also have that

$$\left\| \mathbf{x}^{(\mathbf{a})} - C(u) \right\|_{M_1 \cap K}^2 = O(\epsilon_{\text{DECODE}}). \quad (16.15)$$

$$\left\| \mathbf{x}^{(\mathbf{b})} - C(u) \right\|_{M_1 \cap K}^2 = O(\epsilon_{\text{DECODE}}). \quad (16.16)$$

Therefore, we have that for every ϵ_{NASH} -optimal \mathbf{a} and \mathbf{b} , the counter's encoding $C(u)$ is decoded correctly; and in particular,

$$M_1 \cap I_{\text{CRITICAL}}^{(A,r)}(\mathbf{a}) = M_1 \cap I_{\text{CRITICAL}}^{(A,r)}(\mathbf{b}). \quad (16.17)$$

Let $U_{\text{CRITICAL}}^{(A,r)}|_{M_1}$ denote the portion of $U_{\text{CRITICAL}}^{(A,r)}$ that is derived from assignments to the critical bits in the first proof. For every happy player (A, r) , and every strategy \mathbf{a} in her support, we have

$$\begin{aligned} U_{\text{CRITICAL}}^{(A,r)}|_{M_1} &= -\mathbb{E}_{\mathbf{q} \in \mathcal{G}^{l/2}} \mathbb{E}_{\mathbf{b} \sim \mathcal{B}(B, \mathbf{q})} \left[\left\| \mathbf{x}_{\text{CRITICAL}}^{(\mathbf{a})} - \mathbf{x}^{(\mathbf{b})} \right\|_{M_1 \cap I_{\text{CRITICAL}}^{(A,r)}(\mathbf{b}) \cap I(B, \mathbf{q})}^2 \right] \\ &= -\mathbb{E}_{\mathbf{q} \in \mathcal{G}^{l/2}} \mathbb{E}_{\mathbf{b} \sim \mathcal{B}(B, \mathbf{q})} \left[\left\| \mathbf{x}_{\text{CRITICAL}}^{(\mathbf{a})} - \mathbf{x}^{(\mathbf{b})} \right\|_{M_1 \cap I_{\text{CRITICAL}}^{(A,r)}(\mathbf{a}) \cap I(B, \mathbf{q})}^2 \right] \pm O(\epsilon_{\text{NASH}}) \\ &= -\left\| \mathbf{x}_{\text{CRITICAL}}^{(\mathbf{a})} - \mathbf{x}^{(B)} \right\|_{M_1 \cap I_{\text{CRITICAL}}^{(A,r)}(\mathbf{a})}^2 \\ &\quad - \mathbb{E}_{i \in M_1 \cap I_{\text{CRITICAL}}^{(A,r)}(\mathbf{a})} \left[\text{Var}_{\mathbf{b} \sim \mathcal{B}(B, \mathbf{q}(i))} [\mathbf{x}^{(\mathbf{b})} | i] \right] \pm O(\epsilon_{\text{NASH}}), \end{aligned}$$

where the second equality follows from (16.17). Therefore, for every ϵ_{NASH} -optimal strategy \mathbf{a} ,

$$\left\| \mathbf{x}_{\text{CRITICAL}}^{(\mathbf{a})} - \mathbf{x}^{(B)} \right\|_{M_1 \cap I_{\text{CRITICAL}}^{(A,r)}(\mathbf{a})}^2 = O(\epsilon_{\text{PRECISION}}^2 + \epsilon_{\text{NASH}}) = O(\epsilon_{\text{NASH}}). \quad (16.18)$$

By the Good Sample guarantee from Proposition 16.2, we have that for a $(1 - o(1))$ -fraction of r 's,

$$\left\| \mathbf{x}^{(B)} - \Pi(u) \right\|_{M_1 \cap I_{\text{CRITICAL}}^{(A,r)}(\mathbf{a})}^2 = O(\epsilon_{\text{DECODE}}). \quad (16.19)$$

Combining with (16.18), we have that for a $(1 - O(\epsilon_{\text{NASH}}))$ -fraction of r 's,

$$\left\| \mathbf{x}_{\text{CRITICAL}}^{(\mathbf{a})} - \Pi(u) \right\|_{M_1 \cap I_{\text{CRITICAL}}^{(A,r)}(\mathbf{a})}^2 = O(\epsilon_{\text{DECODE}}).$$

Extrapolation

Notice that (16.19) in particular implies

$$\left\| \mathbf{x}^{(B)} - v(\mathbf{x}^{(B)}) \right\|_{M_1 \cap I_{\text{CRITICAL}}^{(A,r)}(\mathbf{a})}^2 = O(\epsilon_{\text{DECODE}}).$$

Combining with (16.18), we have that

$$\left\| \mathbf{x}_{\text{CRITICAL}}^{(\mathbf{a})} - v(\mathbf{x}^{(B)}) \right\|_{M_1 \cap I_{\text{CRITICAL}}^{(A,r)}(\mathbf{a})}^2 = O(\epsilon_{\text{DECODE}}).$$

Therefore, for a $(1 - O(\epsilon_{\text{NASH}}))$ -fraction of (A, r) 's, sampling $v(\mathbf{x}_{\text{CRITICAL}}^{(\mathbf{a})}|_{M_1})$ satisfies the distance criteria for the Decoding and Robust Decoding in Proposition 16.2. Therefore, by the Error-Correction on the Critical Bits desideratum, for a $(1 -$

$O(\epsilon_{\text{NASH}})$ -fraction of (A, r) 's, the PCP verifier correctly decodes the assignment to all the critical bits.

From the assignment to the critical bits, the PCP verifier can reconstruct the complete difference vector $E(u) - E(S^{\text{LOCAL}}(u))$. So now, accessing $E(S^{\text{LOCAL}}(u))$ is equivalent to accessing $E(u)$ (with the same error guarantees). Using the local proof construction guarantee from Proposition 16.2, the PCP verifier can correctly compute $\pi(S^{\text{LOCAL}}(u))$ on $I_{\text{SAMPLE}}^{(A,r)}$. Finally, recall that once the PCP verifier has successfully decoded $C(u)$, it can also locally compute $C(S^{\text{LOCAL}}(u))$. ■

16.4.6 Default Displacement

The simplest case is when the restrictions of $\mathbf{x}^{(B)}$ to both proofs are far from any $\Pi(u), \Pi(v)$. In this case $\mathbf{x}^{(B)}$ is far from all Brouwer line segments, and $f^{(A)}$ needs to apply the default displacement.

Lemma 16.4 If $(\mathcal{A}, \mathcal{B})$ is an $(\epsilon_{\text{NASH}}, \epsilon_{\text{NASH}})$ -Well-Supported-WeakNash, and both $\|\mathbf{x}^{(B)} - \Pi(u)\|_{M_1}^2 > 4\epsilon_{\text{SOUND}}$ and $\|\mathbf{x}^{(B)} - \Pi(v)\|_{M_2}^2 > 4\epsilon_{\text{SOUND}}$ for every $u, v \in V^{\text{LOCAL}}$, then $\|f^{(A)} - f(\mathbf{x}^{(B)})\|_M^2 = O(\epsilon_{\text{NASH}})$.

Proof By Lemma 16.1, for a $(1 - O(\epsilon_{\text{NASH}}))$ -fraction of r 's and every strategy \mathbf{a} in their support, the game verifier rejects both proofs. In this case, for all those \mathbf{a} 's, $f^{(a)}$ implements the default displacement. If \mathbf{a} is also an ϵ_{NASH} -optimal strategy, then $\|\mathbf{x}^{(a)} - \mathbf{x}^{(B)}\|_{J(A,r)}^2 = O(\epsilon_{\text{NASH}})$, in which case also

$$\|f^{(a)} - f(\mathbf{x}^{(B)})\|_{J(A,r)}^2 = O(\epsilon_{\text{NASH}}).$$

Since the rest of Alice's players can have at most an $O(\epsilon_{\text{NASH}})$ effect, we have that

$$\|f^{(A)} - f(\mathbf{x}^{(B)})\|_M^2 = O(\epsilon_{\text{NASH}}). \quad \blacksquare$$

The following cases are handled with minor modifications:

- $\mathbf{x}^{(B)}$ is outside the picture and far from the first line $(\mathbf{z}_2 \rightarrow (\Pi(u_0), \Pi(u_0), 0, 0))$.
- $\mathbf{x}^{(B)}$ is inside the picture, both $\mathbf{x}^{(B)}|_{M_1}$ and $\mathbf{x}^{(B)}|_{M_2}$ are close to valid proofs, but $\mathbf{x}^{(B)}|_{M_{3,4}}$ doesn't match any Brouwer line segment or vertex.
- $\mathbf{x}^{(B)}$ is inside the picture, both $\mathbf{x}^{(B)}|_{M_1}$ and $\mathbf{x}^{(B)}|_{M_2}$ are close to valid proofs $\Pi(u)$ and $\Pi(v)$, but u and v are not consecutive vertices in the LOCAL END-OF-A-LINE graph.

16.4.7 Close to a Line (1)

Now suppose that $\mathbf{x}^{(B)}$ is close to some Brouwer line segment ($\mathbf{s} \rightarrow \mathbf{t}$) from $\mathbf{s} = (\Pi(u), \Pi(u), 0, 0)$ to $\mathbf{t} = (\Pi(u), \Pi(v), 0, 0)$, for $u, v \in V^{\text{LOCAL}}$ such that v is the LOCAL END-OF-A-LINE-successor of u . (The case of a line from $(\Pi(u), \Pi(v), 1, 0)$ to $(\Pi(v), \Pi(v), 0, 1)$ follows with minor modifications.) Now, the line ($\mathbf{s} \rightarrow \mathbf{t}$) consists of points of the form $\beta\mathbf{s} + (1 - \beta)\mathbf{t}$ (for $\beta \in [0, 1]$). From player (A, r) 's assignment to M_1 , we can locally decode and verify $\Pi(u)$ on $I_{\text{PCP}}^{(A,r)} \cup I_{\text{LTC}}^{(A,r)} \cup K$. Furthermore, we can reconstruct both $\Pi(u)$ and $\Pi(v)$ on $I_{\text{SAMPLE}}^{(A,r)} \cup K$. Let $\mathbf{s}^{(a)}, \mathbf{t}^{(a)}$ denote the locally reconstructed restrictions of \mathbf{s}, \mathbf{t} , respectively, to $J^{(A,r)}$.

If all the tests passed, then we want to locally apply the displacement close to a line, as defined in (4.2). The assignments of $\mathbf{x}^{(a)}, \mathbf{s}^{(a)}, \mathbf{t}^{(a)}$ also induce a partial vector $\mathbf{z}^{(a)}$, which is the point closest to $\mathbf{x}^{(a)}$ on the line segment ($\mathbf{s}^{(a)} \rightarrow \mathbf{t}^{(a)}$). We can also use $\|\mathbf{x}^{(a)} - \mathbf{z}^{(a)}\|_{J^{(A,r)}}$, $\|\mathbf{t}^{(a)} - \mathbf{s}^{(a)}\|_{J^{(A,r)}}$ as estimates of $\|\mathbf{x} - \mathbf{z}\|_M$, $\|\mathbf{t} - \mathbf{s}\|_M$.

Lemma 16.5 Let (A, \mathcal{B}) be an $(\epsilon_{\text{NASH}}, \epsilon_{\text{NASH}})$ -Well-Supported-WeakNash. Suppose that $\mathbf{x}^{(B)}$ is somewhat close to some Brouwer line segment,

$$\epsilon_{\text{COMPLETE}}^2 < \min_{\beta} \left\| \mathbf{x}^{(B)} - (\Pi(u), \beta\Pi(u) + (1 - \beta)\Pi(v), 0, 0) \right\|_M^2 < \epsilon_{\text{DECODE}}, \quad (16.20)$$

but far from its end points,

$$\begin{aligned} \left\| \mathbf{x}^{(B)} - (\Pi(u), \Pi(u), 0, 0) \right\|_M &> 2\sqrt{h} \\ \left\| \mathbf{x}^{(B)} - (\Pi(u), \Pi(v), 0, 0) \right\|_M &> 2\sqrt{h}, \end{aligned} \quad (16.21)$$

where $u, v \in V^{\text{LOCAL}}$ satisfy $v = S^{\text{LOCAL}}(u)$ and $u = P^{\text{LOCAL}}(v)$.

Then $\|f^{(A)} - f(\mathbf{x}^{(B)})\|_M^2 = O(\epsilon_{\text{NASH}})$.

Proof By Lemma 16.3, for a $(1 - O(\epsilon_{\text{NASH}}))$ -fraction of r 's, and every strategy \mathbf{a} in the support of (A, r) , the game verifier can compute $\Pi(u) \upharpoonright_{I_{\text{SAMPLE}}^{(A,r)} \cup K}$ and $\Pi(v) \upharpoonright_{I_{\text{SAMPLE}}^{(A,r)} \cup K}$ correctly; let $\widehat{\Pi}(u)$ and $\widehat{\Pi}(v)$ denote the respective results of those computations.

Consider

$$\beta_{(\mathbf{s} \rightarrow \mathbf{t})}^{(a)} \triangleq \arg \min_{\beta} \left\| \mathbf{x}^{(a)} - \left(\beta\widehat{\Pi}(u) + (1 - \beta)\widehat{\Pi}(v) \right) \right\|_{M_2 \cap (I_{\text{SAMPLE}}^{(A,r)} \cup K)}^2 \quad (16.22)$$

and

$$\mathbf{z}^{(\beta_{(\mathbf{s} \rightarrow \mathbf{t})}^{(a)})} \triangleq \left(\Pi(u), \beta_{(\mathbf{s} \rightarrow \mathbf{t})}^{(a)} \Pi(u) + (1 - \beta_{(\mathbf{s} \rightarrow \mathbf{t})}^{(a)}) \Pi(v), 0, 0 \right) \in [-1, 2]^M.$$

By (16.20), we have

$$\left\| \mathbf{x}^{(\mathcal{B})} - \mathbf{z}^{\left(\beta_{(\mathbf{s} \rightarrow \mathbf{t})}^{(\mathbf{a})}\right)} \right\|_M^2 > \epsilon_{\text{COMPLETE}}^2. \quad (16.23)$$

By the Good Sample guarantee from Proposition 16.2, whenever $\Pi(u) \upharpoonright_{I_{\text{SAMPLE}}^{(A,r)} \cup K}$ and $\Pi(v) \upharpoonright_{I_{\text{SAMPLE}}^{(A,r)} \cup K}$ are indeed decoded correctly,

$$\left\| \mathbf{x}^{(\mathbf{a})} - \mathbf{z}^{(\mathbf{a})} \right\|_{J^{(A,r)}}^2 > \epsilon_{\text{COMPLETE}}^2 - O(\epsilon_{\text{NASH}}) \gg (3h)^2.$$

Therefore, for all those \mathbf{a} 's, we have that $f^{(\mathbf{a})}$ correctly applies the default displacement. \blacksquare

Lemma 16.6 Let $(\mathcal{A}, \mathcal{B})$ be an $(\epsilon_{\text{NASH}}, \epsilon_{\text{NASH}})$ -Well-Supported-WeakNash. Suppose that $\mathbf{x}^{(\mathcal{B})}$ is close to some Brouwer line segment,

$$\min_{\beta} \left\| \mathbf{x}^{(\mathcal{B})} - (\Pi(u), \beta\Pi(u) + (1-\beta)\Pi(v), 0, 0) \right\|_M^2 \leq \epsilon_{\text{COMPLETE}}^2, \quad (16.24)$$

but far from its end points,

$$\begin{aligned} \left\| \mathbf{x}^{(\mathcal{B})} - (\Pi(u), \Pi(u), 0, 0) \right\|_M &> 2\sqrt{h} \\ \left\| \mathbf{x}^{(\mathcal{B})} - (\Pi(u), \Pi(v), 0, 0) \right\|_M &> 2\sqrt{h}, \end{aligned} \quad (16.25)$$

where $u, v \in V^{\text{LOCAL}}$ satisfy $v = S^{\text{LOCAL}}(u)$ and $u = P^{\text{LOCAL}}(v)$.

Then $\|f^{(\mathcal{A})} - f(\mathbf{x}^{(\mathcal{B})})\|_M^2 = O(\epsilon_{\text{NASH}})$.

Proof We show that $f^{(\mathbf{a})}$ correctly implements the displacement close to a line for $\mathbf{s} = (\Pi(u), \Pi(u), 0, 0)$ and $\mathbf{t} = (\Pi(u), \Pi(v), 0, 0)$.

By Lemma 16.2, for a $(1 - O(\epsilon_{\text{NASH}}))$ -fraction of r 's, and every strategy \mathbf{a} in the support of (A, r) , the game verifier accepts $\mathbf{a} \upharpoonright_{\{1, \dots, m\}}$. Furthermore, by Lemma 16.3, it can compute $\Pi(u) \upharpoonright_{I_{\text{SAMPLE}}^{(A,r)} \cup K}$ and $\Pi(v) \upharpoonright_{I_{\text{SAMPLE}}^{(A,r)} \cup K}$ correctly; we again denote the partial proofs locally computed by the game verifier $\widehat{\Pi}(u)$ and $\widehat{\Pi}(v)$. Denote the induced partial vectors $\mathbf{s}^{(\mathbf{a})}$ and $\mathbf{t}^{(\mathbf{a})}$. Recall that for every ϵ -optimal \mathbf{a} , $\|\mathbf{x}^{(\mathbf{a})} - \mathbf{x}^{(\mathcal{B})}\|_{J^{(A,r)}}^2 = O(\epsilon_{\text{NASH}})$, and also, $\|\mathbf{s}^{(\mathbf{a})} - \mathbf{t}^{(\mathbf{a})}\|_{J^{(A,r)}} - \|\mathbf{s} - \mathbf{t}\|_M|^2 = o(1)$.

Let

$$\beta_{(\mathbf{s} \rightarrow \mathbf{t})}^{(\mathbf{a})} \triangleq \frac{(\mathbf{t}^{(\mathbf{a})} - \mathbf{s}^{(\mathbf{a})}) \cdot (\mathbf{x}^{(\mathbf{a})} - \mathbf{s}^{(\mathbf{a})})}{\|\mathbf{s}^{(\mathbf{a})} - \mathbf{t}^{(\mathbf{a})}\|_{J^{(A,r)}}^2},$$

where the (\cdot) denotes a $J^{(A,r)}$ -restricted dot-product.

For every ϵ -optimal \mathbf{a} , $\|\mathbf{x}^{(\mathbf{a})} - \mathbf{x}^{(\mathbf{B})}\|_{J^{(A,r)}}^2 = O(\epsilon_{\text{NASH}})$; for most of them $\Pi(u)|_{J^{(A,r)} \cup K_{\text{SAMPLE}}}$ and $\Pi(v)|_{J^{(A,r)} \cup K_{\text{SAMPLE}}}$ are also computed correctly, so $\mathbf{s}^{(\mathbf{a})}, \mathbf{t}^{(\mathbf{a})}$ are the correct restrictions of \mathbf{s}, \mathbf{t} to $J^{(A,r)}$. Then, by the Good Sample guarantee from Proposition 16.2, we have that with probability $(1 - o(1))$, the $J^{(A,r)}$ -restricted dot-product is a good approximation. Namely,

$$\left| (\mathbf{t}^{(\mathbf{a})} - \mathbf{s}^{(\mathbf{a})}) \cdot (\mathbf{x}^{(\mathbf{B})}|_{J^{(A,r)}} - \mathbf{s}^{(\mathbf{a})}) - (\mathbf{t} - \mathbf{s}) \cdot (\mathbf{x}^{(\mathbf{B})} - \mathbf{s}) \right| = o(1)$$

and

$$\left| (\mathbf{t}^{(\mathbf{a})} - \mathbf{s}^{(\mathbf{a})}) \cdot (\mathbf{x}^{(\mathbf{a})} - \mathbf{x}^{(\mathbf{B})}|_{J^{(A,r)}}) \right| = O\left(\|\mathbf{x}^{(\mathbf{a})} - \mathbf{x}^{(\mathbf{B})}\|_{J^{(A,r)}}\right) = O(\sqrt{\epsilon_{\text{NASH}}}),$$

and thus also

$$\left| (\mathbf{t}^{(\mathbf{a})} - \mathbf{s}^{(\mathbf{a})}) \cdot (\mathbf{x}^{(\mathbf{a})} - \mathbf{s}^{(\mathbf{a})}) - (\mathbf{t} - \mathbf{s}) \cdot (\mathbf{x}^{(\mathbf{B})} - \mathbf{s}) \right| = O(\sqrt{\epsilon_{\text{NASH}}}).$$

Because $\Pi(\cdot)$ is a constant relative distance error-correcting code, $\|\mathbf{s} - \mathbf{t}\|_M = \Theta(1)$. Therefore,

$$\left| \beta_{(\mathbf{s} \rightarrow \mathbf{t})}^{(\mathbf{a})} - \beta_{(\mathbf{s} \rightarrow \mathbf{t})}(\mathbf{x}^{(\mathbf{B})}) \right|^2 = O(\epsilon_{\text{NASH}}).$$

In particular, let

$$\begin{aligned} \mathbf{z}^{(\mathbf{a})} &\triangleq \beta_{(\mathbf{s} \rightarrow \mathbf{t})}^{(\mathbf{a})} \mathbf{s}^{(\mathbf{a})} + (1 - \beta_{(\mathbf{s} \rightarrow \mathbf{t})}^{(\mathbf{a})}) \mathbf{t}^{(\mathbf{a})} \\ \mathbf{z} &\triangleq \beta_{(\mathbf{s} \rightarrow \mathbf{t})}(\mathbf{x}^{(\mathbf{B})}) \mathbf{s} + (1 - \beta_{(\mathbf{s} \rightarrow \mathbf{t})}(\mathbf{x}^{(\mathbf{B})})) \mathbf{t}; \end{aligned}$$

then for a $(1 - O(\epsilon_{\text{NASH}}))$ -fraction of r 's, we have

$$\left\| \mathbf{z}^{(\mathbf{a})} - \mathbf{z} \right\|_{J^{(A,r)}}^2 = O(\epsilon_{\text{NASH}}).$$

By the triangle inequality, also:

$$\left| \left\| \mathbf{x}^{(\mathbf{a})} - \mathbf{z}^{(\mathbf{a})} \right\|_{J^{(A,r)}} - \left\| \mathbf{x}^{(\mathbf{B})} - \mathbf{z} \right\|_{J^{(A,r)}} \right|^2 = O(\epsilon_{\text{NASH}}),$$

and thus by the Good Sample guarantee, also:

$$\left| \left\| \mathbf{x}^{(\mathbf{a})} - \mathbf{z}^{(\mathbf{a})} \right\|_{J^{(A,r)}} - \left\| \mathbf{x}^{(\mathbf{B})} - \mathbf{z} \right\|_M \right|^2 = O(\epsilon_{\text{NASH}}).$$

Finally, whenever \mathbf{a} is ϵ_{NASH} -optimal and satisfies all the above, we have by Lipschitz continuity that:

$$\left\| f(\mathbf{a}) - f(\mathbf{x}^{(\mathbf{B})}) \right\|_{J^{(A,r)}}^2 = O(\epsilon_{\text{NASH}}). \quad \blacksquare$$

16.4.8 Close to a Line (2)

There are a few different scenarios where we expect the game verifier to accept both proofs:

- $\mathbf{x}^{(B)}$ may still be far from every Brouwer line segment; e.g., because the corresponding vertices are not neighbors in the LOCAL END-OF-A-LINE graph, or the values on $M_{3,4}$ don't match.
- $\mathbf{x}^{(B)}$ may be close to the “first” Brouwer line segment, where only the special direction (M_4) changes.
- $\mathbf{x}^{(B)}$ may be close to a Brouwer line segment where only the compute-vs-copy bit (M_3) changes.
- $\mathbf{x}^{(B)}$ may be close to a Brouwer vertex.

In the first case, simply apply the default displacement. In this subsection we briefly describe the two cases corresponding to $\mathbf{x}^{(B)}$ close to a single Brouwer line segment. Finally, the case where $\mathbf{x}^{(B)}$ is close to a Brouwer vertex is deferred to Section 16.4.9.

Close to the First Line

Near the “first” Brouwer line segment, locally computing the displacement is relatively simple. First observe that this case is easy to recognize by local access: by the Good Sample guarantee from Proposition 16.2, we can estimate the Hamming distance of the first $2m + 1$ coordinates to $(\Pi(u_0), \Pi(u_0), 0)$. Furthermore, we know $\mathbf{s} = \mathbf{z}_2$ and $\mathbf{t} = (\Pi(u_0), \Pi(u_0), 0, 0)$ exactly, and therefore also $\|\mathbf{t} - \mathbf{s}\|_M = 1$. Finally, let $\mathbf{z}^{(a)}$ be equal to $(\Pi(u_0), \Pi(u_0), 0)$ on its first $2m + 1$ coordinates, and $\mathbf{x}^{(a)}|_{M_4}$ on the last one.

Close to a Line that Updates the Auxiliary Compute-Vs-Copy Bit

Suppose that $\mathbf{x}^{(B)}$ is close to a Brouwer line segment from $\mathbf{s} = (\Pi(u), \Pi(v), 0, 0)$ to $\mathbf{t} = (\Pi(u), \Pi(v), 1, 0)$, for $u, v \in V^{\text{LOCAL}}$ such that v is the LOCAL END-OF-A-LINE-successor of u . (The case of a Brouwer line segment from $(\Pi(v), \Pi(v), 1, 0)$ to $(\Pi(v), \Pi(v), 0, 0)$ follows with minor modifications.)

By Lemma 16.3, we can locally decode both $\Pi(u)$ and $\Pi(v)$ (and verify that they are valid proofs of consecutive vertices in the LOCAL END-OF-A-LINE graph). We can therefore construct partial vectors $\mathbf{s}^{(a)}$, $\mathbf{t}^{(a)}$ that are equal, for a $(1 - O(\epsilon_{\text{NASH}}))$ -fraction of (A, r) 's, to the restrictions of the true \mathbf{s} , \mathbf{t} to $J^{(A,r)}$. Finally, let

$$\mathbf{z}^{(a)} \triangleq (\mathbf{x}^{(a)}|_{M_3})\mathbf{t}^{(a)} + (1 - \mathbf{x}^{(a)}|_{M_3})\mathbf{s}^{(a)}.$$

16.4.9 Close to a Vertex

In this subsection we consider the case where $\mathbf{x}^{(\mathcal{B})}$ is close to a Brouwer vertex representing two different LOCAL END-OF-A-LINE vertices. In particular, we assume that it is close to a Brouwer vertex of the form $\mathbf{y} = (\Pi(u), \Pi(v), 1, 0)$. The case of $(\Pi(u), \Pi(v), 0, 0)$ follows with minor modifications; we will return to the cases of $(\Pi(u), \Pi(u), 0, 0)$ and $(\Pi(v), \Pi(v), 1, 0)$ in a couple of paragraphs.

After the game verifier accepts both proofs and that v is the successor of u in the LOCAL END-OF-A-LINE graph, we are assured that $\mathbf{x}^{(\mathcal{B})}$ is indeed close to some Brouwer vertex \mathbf{y} . Furthermore, we know that there is an incoming Brouwer line segment from $\mathbf{s} = (\Pi(u), \Pi(v), 1, 0)$ and an outgoing Brouwer line segment to $\mathbf{t} = (\Pi(v), \Pi(v), 1, 0)$. We can locally compute $\mathbf{s}, \mathbf{y}, \mathbf{t}$ on $J^{(A,r)}$ with high probability; denote the resulting partial vectors $\mathbf{s}^{(a)}, \mathbf{y}^{(a)}, \mathbf{t}^{(a)}$.

Alternatively, consider the case where $\mathbf{y} = (\Pi(v), \Pi(v), 1, 0)$ (similarly for $(\Pi(u), \Pi(u), 0, 0)$), with an incoming Brouwer line segment from $\mathbf{s} = (\Pi(P^{\text{LOCAL}}(v)), \Pi(v), 1, 0)$ and outgoing Brouwer line segment to $\mathbf{t} = (\Pi(v), \Pi(v), 0, 0)$. By the Error-Correction guarantee in Lemma 16.3, we can (with high probability over r) locally decode partial vectors $\mathbf{y}^{(a)}$ and $\mathbf{t}^{(a)}$ on $J^{(A,r)}$; furthermore, by the Extrapolation guarantee in the same lemma, we can locally compute $\mathbf{s}^{(a)}$.

It is left to show that we can also locally compute the more involved construction of displacement next to a Brouwer vertex.

We know $\|\mathbf{y} - \mathbf{t}\|_M = 1/2$ exactly, and we can estimate $\|\mathbf{s} - \mathbf{y}\|_M$ from $\|\mathbf{s}^{(a)} - \mathbf{y}^{(a)}\|_{J^{(A,r)}}$. Now, we can locally compute $\mathbf{z}_{(\mathbf{s} \rightarrow \mathbf{y})}$ and $\mathbf{z}_{(\mathbf{y} \rightarrow \mathbf{t})}$ on the coordinates $J^{(A,r)}$ for which we know the value of $\mathbf{s}, \mathbf{y}, \mathbf{t}$; denote those partial vectors $\mathbf{z}_{(\mathbf{s} \rightarrow \mathbf{y})}^{(a)}$ and $\mathbf{z}_{(\mathbf{y} \rightarrow \mathbf{t})}^{(a)}$, respectively.

Recall that

$$\Delta_{(\mathbf{s} \rightarrow \mathbf{y})}(\mathbf{x}) = \frac{(\mathbf{y} - \mathbf{s})}{\|\mathbf{s} - \mathbf{y}\|_M} \cdot (\mathbf{x} - \mathbf{s}) - (1 - \sqrt{h})$$

$$\Delta_{(\mathbf{y} \rightarrow \mathbf{t})}(\mathbf{x}) = \sqrt{h} - \frac{(\mathbf{t} - \mathbf{y})}{\|\mathbf{y} - \mathbf{t}\|_M} \cdot (\mathbf{x} - \mathbf{y}).$$

We can use $\mathbf{x}^{(a)}, \mathbf{s}^{(a)}, \mathbf{y}^{(a)}, \mathbf{t}^{(a)}$ to locally compute estimates $\Delta_{(\mathbf{s} \rightarrow \mathbf{y})}^{(a)}(\mathbf{x})$ and $\Delta_{(\mathbf{y} \rightarrow \mathbf{t})}^{(a)}(\mathbf{x})$, and therefore also $\alpha^{(a)} \approx \alpha(\mathbf{x}^{(\mathcal{B})})$ and $\mathbf{z}^{(a)} \triangleq \alpha^{(a)} \mathbf{z}_{(\mathbf{s} \rightarrow \mathbf{y})}^{(a)} + (1 - \alpha^{(a)}) \mathbf{z}_{(\mathbf{y} \rightarrow \mathbf{t})}^{(a)}$.

Lemma 16.7 Let $(\mathcal{A}, \mathcal{B})$ be an $(\epsilon_{\text{NASH}}, \epsilon_{\text{NASH}})$ -Well-Supported-WeakNash, and let $\|\mathbf{x}^{(\mathcal{B})} - (\Pi(u), \Pi(v), 1, 0)\|_M \leq 2\sqrt{h}$ for $u, v \in V^{\text{LOCAL}}$ such that $v = S^{\text{LOCAL}}(u)$ and $u = P^{\text{LOCAL}}(v)$. Then $\|f^{(a)} - f(\mathbf{x}^{(\mathcal{B})})\|_M^2 = O(\epsilon_{\text{NASH}}/h)$.

Proof By Lemmata 16.2 and 16.3, we have that for a $(1 - O(\epsilon_{\text{NASH}}))$ -fraction of (A, r) 's and every strategy \mathbf{a} in their supports, the game verifier recognizes that $\mathbf{x}^{(\mathbf{a})}$ is close to some Brouwer vertex $\mathbf{y} \triangleq (\Pi(u), \Pi(v), 1, 0)$. In particular, by the Error-Correction desideratum in Lemma 16.3, the game verifier can construct $\mathbf{y}^{(\mathbf{a})}, \mathbf{s}^{(\mathbf{a})}, \mathbf{t}^{(\mathbf{a})} \in \{0, 1, \perp\}^M$ that, for a $(1 - O(\epsilon_{\text{NASH}}))$ -fraction of (A, r) 's and every strategy \mathbf{a} in their supports, are equal to $\mathbf{y}, \mathbf{s}, \mathbf{t}$ on $J^{(A, r)}$.

We also know $\|\mathbf{y} - \mathbf{t}\|_M = 1/2$ exactly, and by the Good Sample guarantee from Proposition 16.2,

$$\left| \|\mathbf{y}^{(\mathbf{a})} - \mathbf{t}^{(\mathbf{a})}\|_{J^{(A, r)}} - \|\mathbf{y} - \mathbf{t}\|_M \right|^2 = o(1).$$

Furthermore, as we argue in the proof of Lemma 16.6,

$$\begin{aligned} \left| \beta_{(\mathbf{s} \rightarrow \mathbf{y})}^{(\mathbf{a})} - \beta_{(\mathbf{s} \rightarrow \mathbf{y})}(\mathbf{x}^{(\mathcal{B})}) \right|^2 &= O(\epsilon_{\text{NASH}}) \\ \left| \beta_{(\mathbf{y} \rightarrow \mathbf{t})}^{(\mathbf{a})} - \beta_{(\mathbf{y} \rightarrow \mathbf{t})}(\mathbf{x}^{(\mathcal{B})}) \right|^2 &= O(\epsilon_{\text{NASH}}). \end{aligned} \quad (16.26)$$

Let $\Delta_{(\mathbf{s} \rightarrow \mathbf{y})}^{(\mathbf{a})} \triangleq \beta_{(\mathbf{s} \rightarrow \mathbf{y})}^{(\mathbf{a})} - (1 - \sqrt{h})$, and $\Delta_{(\mathbf{y} \rightarrow \mathbf{t})}^{(\mathbf{a})} \triangleq \sqrt{h} - \beta_{(\mathbf{y} \rightarrow \mathbf{t})}^{(\mathbf{a})}$. If either quantity is negative, continue with the displacement close to a line as in Section 16.4.7. Recall that f is $O(1)$ -Lipschitz on $[-1, 2]^M$, and in particular near the interface between the line and vertex displacements (the hyperplanes defined by $\beta_{(\mathbf{s} \rightarrow \mathbf{y})}(\mathbf{x}) = 0$ and $\beta_{(\mathbf{y} \rightarrow \mathbf{t})}(\mathbf{x}) = 0$). Therefore, whenever all the parameters are computed approximately correctly, the displacement is also approximately correct—even if near the interface we use the vertex displacement instead of the line displacement or vice versa. We henceforth focus on $\Delta_{(\mathbf{s} \rightarrow \mathbf{y})}^{(\mathbf{a})}, \Delta_{(\mathbf{y} \rightarrow \mathbf{t})}^{(\mathbf{a})} \geq 0$.

Define

$$\alpha^{(\mathbf{a})} \triangleq \frac{\Delta_{(\mathbf{y} \rightarrow \mathbf{t})}^{(\mathbf{a})}}{\Delta_{(\mathbf{y} \rightarrow \mathbf{t})}^{(\mathbf{a})} + \Delta_{(\mathbf{s} \rightarrow \mathbf{y})}^{(\mathbf{a})}},$$

and finally also

$$\begin{aligned} \mathbf{z}^{(\mathcal{B})} &\triangleq \alpha(\mathbf{x}^{(\mathcal{B})})\mathbf{z}_{(\mathbf{s} \rightarrow \mathbf{y})} + (1 - \alpha(\mathbf{x}^{(\mathcal{B})}))\mathbf{z}_{(\mathbf{y} \rightarrow \mathbf{t})} \\ \mathbf{z}^{(\mathbf{a})} &\triangleq \alpha^{(\mathbf{a})}\mathbf{z}_{(\mathbf{s} \rightarrow \mathbf{y})}^{(\mathbf{a})} + (1 - \alpha^{(\mathbf{a})})\mathbf{z}_{(\mathbf{y} \rightarrow \mathbf{t})}^{(\mathbf{a})}. \end{aligned} \quad (16.27)$$

We now consider two different cases depending on the value of $\|\mathbf{x}^{(\mathcal{B})} - \mathbf{z}^{(\mathcal{B})}\|_M$:

Case $\|\mathbf{x}^{(\mathcal{B})} - \mathbf{z}^{(\mathcal{B})}\|_M > 10h$:

The challenge is that when $\|\mathbf{x}^{(\mathcal{B})} - \mathbf{z}^{(\mathcal{B})}\|_M$ is huge, $\Delta_{(\mathbf{y} \rightarrow \mathbf{t})} + \Delta_{(\mathbf{s} \rightarrow \mathbf{y})}$ may be very small, which could lead to $\alpha^{(\mathbf{a})}$ being far from the true $\alpha(\mathbf{x}^{(\mathcal{B})})$. Fortunately, in this

case the true $\hat{g}(\mathbf{x}^{(B)})$ is just the default displacement, $\delta(\mathbf{0}_{2m+1}, \mathbf{1})$ —so we only have to argue that $f^{(\mathbf{a})}$ also applies the default displacement (for most players (A, r)).

Observe that if $\alpha(\mathbf{x}^{(B)}) = 1/2$, the point on L_y that is closest to $\mathbf{x}^{(B)}$ is indeed $\mathbf{z}^{(B)}$. In general, this is not true, but $\|\mathbf{x}^{(B)} - \mathbf{z}^{(B)}\|_M$ is at most $\sqrt{2}$ -times larger than the distance from $\mathbf{x}^{(B)}$ to L_y . In particular, for any $\mathbf{z}^{(\alpha^{(\mathbf{a})})} \triangleq \alpha^{(\mathbf{a})}\mathbf{z}_{(s \rightarrow y)} + (1 - \alpha^{(\mathbf{a})})\mathbf{z}_{(y \rightarrow t)}$, we have that $\|\mathbf{x}^{(B)} - \mathbf{z}^{(\alpha^{(\mathbf{a})})}\|_M > 7h$.

Therefore by the Good Sample guarantee from Proposition 16.2 and Claim 16.3, for a $(1 - O(\epsilon_{\text{NASH}}))$ -fraction of (A, r) 's and every strategy \mathbf{a} in their supports, we have

$$\|\mathbf{x}^{(\mathbf{a})} - \mathbf{z}^{(\mathbf{a})}\|_{J(A,r)} > 7h - O(\epsilon_{\text{NASH}}).$$

For all those \mathbf{a} 's, $f^{(\mathbf{a})}$ correctly implements the default displacement.

Case $\|\mathbf{x}^{(B)} - \mathbf{z}^{(B)}\|_M \leq 10h$:

From (16.26),

$$\begin{aligned} \left| \Delta_{(s \rightarrow y)}^{(\mathbf{a})} - \Delta_{(s \rightarrow y)}(\mathbf{x}^{(B)}) \right|^2 &= O(\epsilon_{\text{NASH}}) \\ \left| \Delta_{(y \rightarrow t)}^{(\mathbf{a})} - \Delta_{(y \rightarrow t)}(\mathbf{x}^{(B)}) \right|^2 &= O(\epsilon_{\text{NASH}}). \end{aligned}$$

Plugging into (4.5), we have that

$$\Delta_{(s \rightarrow y)}^{(\mathbf{a})} + \Delta_{(y \rightarrow t)}^{(\mathbf{a})} \geq \sqrt{h} - O(h),$$

and therefore also

$$\left| \alpha^{(\mathbf{a})} - \alpha(\mathbf{x}^{(B)}) \right|^2 = O(\epsilon_{\text{NASH}}/h).$$

Similarly, also⁷

$$\left\| \mathbf{z}^{(\mathbf{a})} - \mathbf{z}^{(B)} \right\|_{J(A,r)}^2 = O(\epsilon_{\text{NASH}}/h).$$

Finally, whenever \mathbf{a} is ϵ_{NASH} -optimal and satisfies all the above, we have

$$\left\| f^{(\mathbf{a})} - f(\mathbf{x}^{(B)}) \right\|_{J(A,r)}^2 = O(\epsilon_{\text{NASH}}/h). \quad \blacksquare$$

7. In fact, we actually have $\|\mathbf{z}^{(\mathbf{a})} - \mathbf{z}^{(B)}\|_{J(A,r)}^2 = O(\epsilon_{\text{NASH}})$, because α interpolates between $\mathbf{z}_{(s \rightarrow y)}$ and $\mathbf{z}_{(y \rightarrow t)}$, which are already at distance $O(\sqrt{h})$ from each other.

16.5 From Polymatrix to Bimatrix

In this section we complete the proof of our main result, Theorem 1.1, by reducing from multiplayer polymatrix games (à la Theorem 16.5) to two-player games.

16.5.1 From $(\sqrt{\epsilon} + \epsilon, \delta)$ -Well-Supported-WeakNash to (ϵ, δ) -WeakNash

Lemma 16.8 Consider a complete bipartite polymatrix game such that every two adjacent vertices play a bimatrix subgame with payoffs in $[0, 1/n_B]$, $[0, 1/n_A]$, where n_A, n_B are respectively the numbers of players on the two sides of the bipartite game. Given an (ϵ, δ) -WeakNash, we can construct in polynomial time a $(\sqrt{\epsilon} \cdot (\sqrt{\epsilon} + 5), \delta)$ -Well-Supported-WeakNash.

Proof Let $(V_A; V_B)$ be the sets of players, where each $v \in V_A \cup V_B$ has utility U^v and action set \mathcal{S}^v . Let $\mathbf{x} = (x_s^v) \in \Delta(\times_{v \in V_A \cup V_B} \mathcal{S}^v)$ be an (ϵ, δ) -WeakNash. Finally, let $U_{\max}^v(\mathbf{x}^{-v}) = \max_{s \in \mathcal{S}^v} U_s^v(\mathbf{x}^{-v})$. Since this is a polymatrix game, we can write

$$\forall v \in V_A \quad U_s^v(\mathbf{x}) = \sum_{u \in V_B} U_s^{v,u}(\mathbf{x}^u),$$

$$\forall v \in V_B \quad U_s^v(\mathbf{x}) = \sum_{u \in V_A} U_s^{v,u}(\mathbf{x}^u).$$

Let $k = k(\epsilon) > 0$ be some large number to be specified later. We construct our new approximate equilibrium as follows. For each of the $(1 - \delta)$ players who play ϵ -optimally in \mathbf{x} , we take only the strategies that are within ϵk of the optimum:

$$\hat{x}_s^v = \begin{cases} \frac{x_s^v}{1-z^v} & \text{if } U_s^v(\mathbf{x}^{-v}) \geq U_{\max}^v(\mathbf{x}^{-v}) - \epsilon k \\ 0 & \text{otherwise} \end{cases}$$

where z^v is the total probability that player v assigns to strategies that are more than ϵk away from the optimum.

The above division is well defined because for $k > 1$ and v who plays ϵ -optimally, z^v is bounded away from 1. Moreover, the following claim from [Daskalakis et al. \[2009a\]](#) formalizes the intuition that when k is sufficiently large, the total weight on actions removed is small, so $\hat{\mathbf{x}}^v$ is close to \mathbf{x}^v :

Claim 16.5 ([Daskalakis et al. 2009a](#), Claim 6)

$$\forall v \in V_A \cup V_B \quad \sum_{s \in \mathcal{S}^v} \left| \hat{x}_s^v - x_s^v \right| \leq \frac{2}{k-1}.$$

Now, the total change to the expected payoff of player v for each action s is bounded by the total change in mixed strategies of its neighbors. In the following

let $v \in V_A$; the analogous argument for the utility of players in V_B follows with minor modifications.

$$\begin{aligned} \left| U_s^v(\mathbf{x}^{-v}) - U_s^v(\hat{\mathbf{x}}^{-v}) \right| &\leq \sum_{u \in V_B} \left| U_s^{v,u}(\mathbf{x}^u) - U_s^{v,u}(\hat{\mathbf{x}}^u) \right| \\ &\leq \frac{1}{n_B} \cdot \sum_{u \in V_B} \sum_{s \in S^u} \left| \hat{x}_s^u - x_s^u \right| \leq \frac{2}{k-1}. \end{aligned}$$

It follows that $\hat{\mathbf{x}}$ is a $(k\epsilon + \frac{2}{k-1}, \delta)$ -Well-Supported-WeakNash:

$$U_s^v(\hat{\mathbf{x}}^{-v}) \geq U_s^v(\mathbf{x}^{-v}) - \frac{2}{k-1} \geq U_{\max}^v(\mathbf{x}^{-v}) - \epsilon k - \frac{2}{k-1} \geq U_{\max}^v(\hat{\mathbf{x}}^{-v}) - \epsilon k - \frac{4}{k-1}.$$

Finally, take $k = 1 + 1/\sqrt{\epsilon}$ to get that

$$k\epsilon + \frac{4}{k-1} \leq \sqrt{\epsilon} \cdot (\sqrt{\epsilon} + 5). \quad \blacksquare$$

16.5.2 From (ϵ, δ) -WeakNash to $\Theta(\epsilon, \delta)$ -ANE in a Bimatrix Game

Lemma 16.9 Suppose we are given a complete bipartite polymatrix game between $n^{1/2+o(1)}$ vertices with $2^{n^{1/2+o(1)}}$ actions each; the payoffs in each bimatrix game are in $[0, 1/n_B], [0, 1/n_A]$, where n_A, n_B denote the number of players on each side of the bipartite graph. Then we can construct a bimatrix game of size $2^{n^{1/2+o(1)}}$ such that every ϵ -ANE of the new bimatrix game corresponds to a $(\delta, \epsilon_{\text{POLYMATRIX}})$ -WeakNash of the polymatrix game, for sufficiently small $\epsilon = \Theta(\delta^2 \cdot \epsilon_{\text{POLYMATRIX}}^2)$.

Proof The two players play three games simultaneously: the main game, which is the heart of the reduction; and two games based on a construction due to [Althofer \[1994\]](#), which impose structural properties of any approximate Nash equilibrium.

Main game. We let each of the two players “control” the vertices on one side of the bipartite graphical game.

Alice’s actions correspond to a choice of vertex on her side of the polymatrix game, and an action for that vertex. Similarly, Bob’s actions correspond to a choice of vertex on his side, and a choice of action for that vertex. The utility for each player is the utility to the corresponding vertex from the induced subgame, scaled by a factor of $\lambda \cdot n_B$ or $\lambda \cdot n_A$ for Alice or Bob, respectively. Here $\lambda \gg \epsilon$ is a small constant to be defined later.

Althofer Games. In addition to the main game, we use a construction due to [Althofer \[1994\]](#) to introduce two auxiliary games that force players to spread their mixed strategies approximately evenly across all vertices. Althofer’s original game

is a one-sum game with payoffs in $\{0, 1\}$ and size $k \times \binom{k}{k/2}$. In each column, exactly half of the entries are 1's, and the rest are 0's. For example, for $k = 4$, the payoff matrices are given by:

$$R = 1 - C = \begin{pmatrix} 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 & 1 & 0 \end{pmatrix}.$$

The value of this game is $1/2$.

For our purposes, we consider two instantiations of Althofer's gadget: a "primal" Althofer game of size $n_A \times \binom{n_A}{n_A/2}$, and a "dual" Althofer game of size $\binom{n_B}{n_B/2} \times n_B$. In any (approximate) Nash, both Alice and Bob must mix (approximately) evenly among (almost) all of their actions.

The Final Construction

Finally, we compose all three games together by identifying Alice's choice of vertex in the main game with her choice of row in the primal Althofer game, and Bob's choice of vertex in the main game with his choice of column in the dual Althofer game.

Analysis. Alice's and Bob's respective mixed strategies induce a mixed strategy profile $(\mathcal{A}, \mathcal{B})$ on their vertices: Alice's vertex (A, i) 's action is drawn from Alice's action conditioned on picking (A, i) , and analogously for Bob's (B, j) . For any vertex that is never picked by its player, fix an arbitrary strategy. Our goal is to show that $(\mathcal{A}, \mathcal{B})$ is a $(\delta, \epsilon_{\text{POLYMATRIX}})$ -WeakNash of the polymatrix game.

Let $U^{(A,i)}(\mathcal{A}[(A, i)], \mathcal{B})$ denote (A, i) 's expected utility when all players draw their strategies according to $(\mathcal{A}, \mathcal{B})$, and analogously for $U^{(B,j)}(\mathcal{B}[(B, j)], \mathcal{A})$.

By Lemma 2.7, in every λ -ANE (and in particular in any ϵ -ANE), Alice's and Bob's respective marginal distributions over their vertices are $O(\lambda)$ -close to uniform. Therefore, Alice's expected utilities from the main game satisfy

$$\left| U_{\text{MAIN}}^A - \lambda \mathbb{E}_{i \in [n_A]} U^{(A,i)}(\mathcal{A}[(A, i)], \mathcal{B}) \right| = O(\lambda^2).$$

Suppose by contradiction that Alice and Bob are in an ϵ -ANE in the bimatrix game, but a δ -fraction of Alice's vertices have an $\epsilon_{\text{POLYMATRIX}}$ -deviating strategy in the polymatrix game. Let $(\widehat{\mathcal{A}}, \mathcal{B})$ denote the deviating mixed strategy profile (notice that the deviation of a player on Alice's side does not affect the utilities of other players on the same side). $(\widehat{\mathcal{A}}, \mathcal{B})$ induces, without changing the marginal distributions over

vertices, a profile of mixed strategies for Alice and Bob; Alice's new expected utility from the main game, $\widehat{U}_{\text{MAIN}}^A$, also satisfies

$$\left| \widehat{U}_{\text{MAIN}}^A - \lambda \mathbb{E}_{i \in [n_A]} U^{(A,i)}(\widehat{\mathcal{A}}[(A,i)], \mathcal{B}) \right| = O(\lambda^2).$$

Therefore,

$$\begin{aligned} \widehat{U}_{\text{MAIN}}^A &\geq \lambda \mathbb{E}_{i \in [n_A]} U^{(A,i)}(\widehat{\mathcal{A}}[(A,i)], \widehat{\mathcal{B}}) - O(\lambda^2) \\ &\geq \lambda \mathbb{E}_{i \in [n_A]} U^{(A,i)}(\mathcal{A}[(A,i)], \mathcal{B}) + \lambda \delta \cdot \epsilon_{\text{POLYMATRIX}} - O(\lambda^2) \\ &\geq U_{\text{MAIN}}^A + \lambda \delta \cdot \epsilon_{\text{POLYMATRIX}} - O(\lambda^2). \end{aligned}$$

Set $\lambda \delta \cdot \epsilon_{\text{POLYMATRIX}} - O(\lambda^2) > \epsilon$. Since Alice and Bob have not changed their marginal vertex utilities, this is an ϵ -deviating strategy. ■

References

- S. Aaronson, R. Impagliazzo, and D. Moshkovitz. 2014. AM with multiple merlins. In *Proceedings of the IEEE Twenty-Ninth Conference on Computational Complexity*, June 2014, Vancouver, British Columbia, pp. 44–55. DOI: [10.1109/CCC.2014.13](https://doi.org/10.1109/CCC.2014.13). [13](#), [23](#), [149](#), [180](#), [184](#), [224](#)
- A. V. Aho, J. D. Ullman, and M. Yannakakis. 1983. On notions of information transfer in VLSI circuits. In *Proceedings of the Fifteenth Annual ACM Symposium on Theory of Computing*, April 1983, Boston, Massachusetts, pp. 133–139. DOI: [10.1145/800061.808742](https://doi.org/10.1145/800061.808742). [37](#)
- N. Alon, L. Babai, and A. Itai. 1986. A fast and simple randomized parallel algorithm for the maximal independent set problem. *Journal of Algorithms*, 7(4): 567–583. DOI: [10.1016/0196-6774\(86\)90019-2](https://doi.org/10.1016/0196-6774(86)90019-2). [29](#)
- N. Alon, O. Goldreich, J. Håstad, and R. Peralta. 1992. Simple construction of almost k -wise independent random variables. *Random Structures and Algorithms*, 3(3): 289–304. DOI: [10.1002/rsa.3240030308](https://doi.org/10.1002/rsa.3240030308). [28](#)
- N. Alon, M. Krivelevich, and B. Sudakov. 1998. Finding a large hidden clique in a random graph. In *Proceedings of the Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, January 1998, San Francisco, California, pp. 594–598. DOI: [10.1002/\(SICI\)1098-2418\(199810/12\)13:3/4<457::AID-RSA14>3.0.CO;2-W](https://doi.org/10.1002/(SICI)1098-2418(199810/12)13:3/4<457::AID-RSA14>3.0.CO;2-W). [157](#)
- N. Alon, A. Andoni, T. Kaufman, K. Matulef, R. Rubinfeld, and N. Xie. 2007. Testing k -wise and almost k -wise independence. In *Proceedings of the 39th Annual ACM Symposium on the Theory of Computing*, June 2007, San Diego, California, pp. 496–505. DOI: [10.1145/1250790.1250863](https://doi.org/10.1145/1250790.1250863). [157](#)
- N. Alon, S. Arora, R. Manokaran, D. Moshkovitz, and O. Weinstein. 2011. Inapproximability of densest κ -subgraph from average case hardness. Unpublished manuscript. <http://www.cs.tau.ac.il/~nogaa/PDFS/dks8.pdf>. [156](#), [157](#)
- N. Alon, T. Lee, A. Shraibman, and S. Vempala. 2013. The approximate rank of a matrix and its algorithmic applications: Approximate rank. In *Proceedings of the Symposium on Theory of Computing Conference*, June 2013, Palo Alto, California, pp. 675–684. DOI: [10.1145/2488608.2488694](https://doi.org/10.1145/2488608.2488694). [13](#), [17](#), [223](#), [224](#)

- I. Althofer. 1994. On sparse approximations to randomized strategies and convex combinations. *Linear Algebra and Its Applications*, 199: 339–355. DOI: [10.1016/0024-3795\(94\)90357-3](https://doi.org/10.1016/0024-3795(94)90357-3). 152, 216, 224, 230, 231, 270
- Y. Anbalagan, S. Norin, R. Savani, and A. Vetta. 2013. Polylogarithmic supports are required for approximate well-supported Nash equilibria below $2/3$. In *Proceedings of the Ninth International Conference on Web and Internet Economics*, December 2013, Cambridge, Massachusetts, pp. 15–23. DOI: [10.1007/978-3-642-45046-4_2](https://doi.org/10.1007/978-3-642-45046-4_2). 224
- Y. Anbalagan, H. Huang, S. Lovett, S. Norin, A. Vetta, and H. Wu. 2015. Large supports are required for well-supported Nash equilibria. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, August 2015, Princeton, New Jersey, pp. 78–84. DOI: [10.4230/LIPIcs.APPROX-RANDOM.2015.78](https://doi.org/10.4230/LIPIcs.APPROX-RANDOM.2015.78). 224
- A. Anshu, N. B. Goud, R. Jain, S. Kundu, and P. Mukhopadhyay. 2017. Lifting randomized query complexity to randomized communication complexity. arXiv preprint arXiv:1703.07521. 37, 40
- S. Arora and B. Barak. 2009. *Computational Complexity—A Modern Approach*. Cambridge University Press. 227
- S. Arora and S. Safra. 1998. Probabilistic checking of proofs: A new characterization of NP. *Journal of the ACM*, 45(1): 70–122. DOI: [10.1145/273865.273901](https://doi.org/10.1145/273865.273901). 14, 155, 231
- S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. 1998. Proof verification and the hardness of approximation problems. *Journal of the ACM*, 45(3): 501–555. DOI: [10.1145/278298.278306](https://doi.org/10.1145/278298.278306). 14, 155
- S. Arora, R. Ge, S. Sachdeva, and G. Schoenebeck. 2012. Finding overlapping communities in social networks: toward a rigorous approach. In *Proceedings of the ACM Conference on Electronic Commerce*, June 2012, Valencia, Spain, pp. 37–54. DOI: [10.1145/222901.2229020](https://doi.org/10.1145/222901.2229020). 12, 14, 15, 177, 178, 179
- M. Asteris, D. S. Papailiopoulos, A. Kyrillidis, and A. G. Dimakis. 2015. Sparse PCA via bipartite matchings. In *Proceedings of the Twenty-Eighth International Conference on Neural Information Processing Systems*, December 2015, Montreal, Quebec, pp. 766–774. <http://papers.nips.cc/paper/5901-sparse-pca-via-bipartite-matchings>. 13
- R. J. Aumann. 1987. Game theory. In J. Eatwell, M. Milgate, and P. Newman, eds., *The New Palgrave: A Dictionary of Economics*. Palgrave Macmillan. 3
- P. Austrin, M. Braverman, and E. Chlamtac. 2013. Inapproximability of NP-complete variants of Nash equilibrium. *Theory of Computing*, 9: 117–142. DOI: [10.4086/toc.2013.v009a003](https://doi.org/10.4086/toc.2013.v009a003). 224
- Y. Azar, R. Motwani, and J. Naor. 1998. Approximating probability distributions using small sample spaces. *Combinatorica*, 18(2): 151–171. DOI: [10.1007/PL00009813](https://doi.org/10.1007/PL00009813). 27, 28
- L. Babai and L. Fortnow. 1991. Arithmetization: A new method in structural complexity theory. *Computational Complexity*, 1: 41–66. DOI: [10.1007/BF01200057](https://doi.org/10.1007/BF01200057). 229
- L. Babai, L. Fortnow, L. A. Levin, and M. Szegedy. 1991. Checking computations in polylogarithmic time. In *Proceedings of the Twenty-Third Annual ACM Symposium*

- on Theory of Computing*, May 1991, New Orleans, Louisiana, pp. 21–31. DOI: [10.1145/103418.103428](https://doi.org/10.1145/103418.103428). 231
- Y. Babichenko. 2012. Completely uncoupled dynamics and Nash equilibria. *Games and Economic Behavior*, 76(1): 1–14. DOI: [10.1016/j.geb.2012.06.004](https://doi.org/10.1016/j.geb.2012.06.004). 18, 33, 35
- Y. Babichenko. 2016. Query complexity of approximate Nash equilibria. *Journal of the ACM*, 63(4): 36:1–36:24. DOI: [10.1145/2908734](https://doi.org/10.1145/2908734). 18, 34, 37, 38, 41, 43
- Y. Babichenko and S. Barman. 2015. Query complexity of correlated equilibrium. *ACM Transactions on Economics and Computation*, 3(4): 22. DOI: [10.1145/2785668](https://doi.org/10.1145/2785668). 39
- Y. Babichenko and A. Rubinstein. 2017. Communication complexity of approximate Nash equilibria. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, June 2017, Montreal, Quebec, pp. 878–889. DOI: [10.1145/3055399.3055407](https://doi.org/10.1145/3055399.3055407). xiv
- Y. Babichenko, C. H. Papadimitriou, and A. Rubinstein. 2016. Can almost everybody be almost happy? In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, January 2016, Cambridge, Massachusetts, pp. 1–9. DOI: [10.1145/2840728.2840731](https://doi.org/10.1145/2840728.2840731). 13, 20, 22, 35, 224, 230, 231
- A. Badanidiyuru, C. H. Papadimitriou, A. Rubinstein, L. Seeman, and Y. Singer. 2016. Locally adaptive optimization: Adaptive seeding for monotone submodular functions. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, January 2016, Arlington, VA, pp. 414–429. DOI: [10.1137/1.9781611974331.ch31](https://doi.org/10.1137/1.9781611974331.ch31). 157
- M. Balcan, C. Borgs, M. Braverman, J. T. Chayes, and S. Teng. 2013. Finding endogenously formed communities. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, January 2013, New Orleans, Louisiana, pp. 767–783. DOI: [10.1137/1.9781611973105.55](https://doi.org/10.1137/1.9781611973105.55). 15, 157, 177, 178
- A. S. Bandeira, N. Boumal, and V. Voroninski. 2016. On the low-rank approach for semidefinite programs arising in synchronization and community detection. In *Proceedings of the Twenty-Ninth Conference on Learning Theory*, June 2016, New York, pp. 361–382. <http://jmlr.org/proceedings/papers/v49/bandeira16.html>. 179
- J. Banks, C. Moore, J. Neeman, and P. Netrapalli. 2016. Information-theoretic thresholds for community detection in sparse networks. In *Proceedings of the Twenty-Ninth Conference on Learning Theory*, June 2016, New York, pp. 383–416. <http://jmlr.org/proceedings/papers/v49/banks16.html>. 179
- B. Barak. 2004. *Non-Black-Box Techniques in Cryptography*. PhD thesis, Weizmann Institute of Science. DOI: [10.1007/11753728_1](https://doi.org/10.1007/11753728_1). 6
- B. Barak, S. B. Hopkins, J. Kelner, P. K. Kothari, A. Moitra, and A. Potechin. 2016. A nearly tight sum-of-squares lower bound for the planted clique problem. In *Proceedings of the IEEE Fifty-Seventh Annual Symposium on Foundations of Computer Science*, October 2016, New Brunswick, New Jersey, pp. 428–437. DOI: [10.1109/FOCS.2016.53](https://doi.org/10.1109/FOCS.2016.53). 157
- S. Barman. 2015. Approximating Nash equilibria and dense bipartite subgraphs via an approximate version of Caratheodory’s theorem. In *Proceedings of the Forty-Seventh*

- Annual ACM Symposium on Theory of Computing*, June 2015, Portland, Oregon, pp. 361–369. DOI: [10.1145/2746539.2746566](https://doi.org/10.1145/2746539.2746566). 12, 17, 155, 156, 157, 223, 224
- S. Barman, K. Ligett, and G. Piliouras. 2015. Approximating Nash equilibria in tree polymatrix games. In *Algorithmic Game Theory—Eighth International Symposium*, September 2015, Saarbrücken, Germany, pp. 285–296. DOI: [10.1007/978-3-662-48433-3_22](https://doi.org/10.1007/978-3-662-48433-3_22). 101
- R. C. Battalio, J. H. Kagel, and C. A. Kogut. 1991. Experimental confirmation of the existence of a Giffen good. *The American Economic Review*, 81(4): 961–970. <https://www.jstor.org/stable/2006656>. 9
- C. Bazgan, F. Foucaud, and F. Sikora. 2016. On the approximability of partial VC dimension. In *Proceedings of the Tenth International Conference on Combinatorial Optimization and Applications*, December 2016, Hong Kong, pp. 92–106. DOI: [10.1007/978-3-319-48749-6_7](https://doi.org/10.1007/978-3-319-48749-6_7). 190
- P. Beame, S. A. Cook, J. Edmonds, R. Impagliazzo, and T. Pitassi. 1998. The relative complexity of NP search problems. *Journal of Computer System Sciences*, 57(1): 3–19. DOI: [10.1006/jcss.1998.1575](https://doi.org/10.1006/jcss.1998.1575). 223
- S. Ben-David and M. Ackerman. 2008. Measures of clustering quality: A working set of axioms for clustering. In *Proceedings of the Twenty-Second Annual Conference on Neural Information Processing Systems*, December 2018, Vancouver, British Columbia, pp. 121–128. <http://papers.nips.cc/paper/3491-measures-of-clustering-quality-a-working-set-of-axioms-for-clustering>. 179
- S. Ben-David, D. Pál, and S. Shalev-Shwartz. 2009. Agnostic online learning. In *Proceedings of the Twenty-Second Conference on Learning Theory*, June 2009, Montreal, Quebec. <http://www.cs.mcgill.ca/~colt2009/papers/032.pdf#page=1>. 25
- E. Ben-Sasson, M. Sudan, S. P. Vadhan, and A. Wigderson. 2003. Randomness-efficient low degree tests and short PCPs via epsilon-biased sets. In *Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing*, June 2003, San Diego, California, pp. 612–621. DOI: [10.1145/780542.780631](https://doi.org/10.1145/780542.780631). 28, 228, 231
- E. Ben-Sasson, O. Goldreich, P. Harsha, M. Sudan, and S. P. Vadhan. 2006. Robust PCPs of proximity, shorter PCPs, and applications to coding. *SIAM Journal on Computing*, 36(4): 889–974. DOI: [10.1137/S0097539705446810](https://doi.org/10.1137/S0097539705446810). 229, 231
- Q. Berthet and P. Rigollet. 2013. Complexity theoretic lower bounds for sparse principal component detection. In *Proceedings of the Twenty-Sixth Annual Conference on Learning Theory*, June 2013, Princeton University, New Jersey, pp. 1046–1066. <https://arxiv.org/abs/1304.0828>. 157
- U. Bhaskar, Y. Cheng, Y. K. Ko, and C. Swamy. 2016. Hardness results for signaling in Bayesian zero-sum and network routing games. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, July 2016, Maastricht, The Netherlands, pp. 479–496. DOI: [10.1145/2940716.2940753](https://doi.org/10.1145/2940716.2940753). 214
- A. Bhaskara, M. Charikar, E. Chlamtac, U. Feige, and A. Vijayaraghavan. 2010. Detecting high log-densities: An $O(n^{1/4})$ approximation for densest k -subgraph. In *Proceedings*

- of the Forty-Second ACM Symposium on Theory of Computing*, June 2010, Cambridge, Massachusetts, pp. 201–210. DOI: [10.1145/1806689.1806718](https://doi.org/10.1145/1806689.1806718). 156
- A. Bhaskara, M. Charikar, A. Vijayaraghavan, V. Guruswami, and Y. Zhou. 2012. Polynomial integrality gaps for strong SDP relaxations of densest k -subgraph. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*, January 2012, Kyoto, Japan, pp. 388–405. <http://dl.acm.org/citation.cfm?id=2095116.2095150>. 156
- B. E. Birnbaum, N. R. Devanur, and L. Xiao. 2011. Distributed algorithms via gradient descent for fisher markets. In *Proceedings of the Twelfth ACM Conference on Electronic Commerce (EC-2011)*, June 2011, San Jose, California, pp. 127–136. DOI: [10.1145/1993574.1993594](https://doi.org/10.1145/1993574.1993594). 9
- N. Bitansky, O. Paneth, and A. Rosen. 2015. On the cryptographic hardness of finding a nash equilibrium. In *Proceedings of the IEEE Fifty-Sixth Annual Symposium on Foundations of Computer Science*, October 2015, Berkeley, California, pp. 1480–1498. DOI: [10.1109/FOCS.2015.94](https://doi.org/10.1109/FOCS.2015.94). 6
- A. Blum. 1994. Separating distribution-free and mistake-bound learning models over the Boolean domain. *SIAM Journal on Computing*, 23(5): 990–1000. DOI: [10.1137/S009753979223455X](https://doi.org/10.1137/S009753979223455X). 190
- C. Borgs, J. T. Chayes, A. Marple, and S. Teng. 2016. An axiomatic approach to community detection. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, January 2016, Cambridge, Massachusetts, pp. 135–146. DOI: [10.1145/2840728.2840748](https://doi.org/10.1145/2840728.2840748). 14, 177, 179
- H. Bosse, J. Byrka, and E. Markakis. 2010. New algorithms for approximate Nash equilibria in bimatrix games. *Theoretical Computer Science*, 411(1): 164–173. DOI: [10.1016/j.tcs.2009.09.023](https://doi.org/10.1016/j.tcs.2009.09.023). 17, 223
- W. C. Brainard and H. Scarf. 2000. How to compute equilibrium prices in 1891. Cowles Foundation Discussion Papers 1272, Cowles Foundation for Research in Economics, Yale University. <http://cowles.yale.edu/sites/default/files/files/pub/d12/d1272.pdf>. 4
- M. Braverman, Y. Kun-Ko, and O. Weinstein. 2015. Approximating the best Nash equilibrium in $n^{O(\log n)}$ -time breaks the exponential time hypothesis. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, January 2015, San Diego, California, pp. 970–982. DOI: [10.1137/1.9781611973730.66](https://doi.org/10.1137/1.9781611973730.66). 17, 150, 151, 158, 213, 224
- M. Braverman, Y. Kun-Ko, A. Rubinfeld, and O. Weinstein. 2017. ETH hardness for densest- k -subgraph with perfect completeness. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, January 2017, Barcelona, Spain, pp. 1326–1341. DOI: [10.1137/1.9781611974782.86](https://doi.org/10.1137/1.9781611974782.86). xiv, 15, 23
- G. W. Brown. 1951. Iterative solutions of games by fictitious play. In *Activity Analysis of Production and Allocation*. T. C. Koopmans (ed.) New York: Wiley, pp. 374–376. 4
- E. Budish. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6): 1061–1103. DOI: [10.1086/664613](https://doi.org/10.1086/664613). 11, 131, 132, 133, 139, 140, 142, 143, 144, 145

- E. Budish, G. P. Cachon, J. Kessler, and A. Othman. 2014. Course Match: A large-scale implementation of approximate competitive equilibrium from equal incomes for combinatorial allocation. Working paper. DOI: [10.1287/opre.2016.1544](https://doi.org/10.1287/opre.2016.1544). 11, 131
- C. Calabro, R. Impagliazzo, and R. Paturi. 2009. The complexity of satisfiability of small depth circuits. In *Parameterized and Exact Computation, 4th International Workshop*, September 2009, Copenhagen, Denmark, pp. 75–85. DOI: [10.1007/978-3-642-11269-0_6](https://doi.org/10.1007/978-3-642-11269-0_6). 224
- M. L. Carmosino, J. Gao, R. Impagliazzo, I. Mihajlin, R. Paturi, and S. Schneider. 2016. Nondeterministic extensions of the strong exponential time hypothesis and consequences for non-reducibility. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, January 2016, Cambridge, Massachusetts, pp. 261–270. DOI: [10.1145/2840728.2840746](https://doi.org/10.1145/2840728.2840746). 224
- S. O. Chan, D. Papailiopoulos, and A. Rubinfeld. 2016. On the approximability of sparse PCA. In *Proceedings of the Twenty-Ninth Conference on Learning Theory*, June 2016, New York, pp. 623–646. <http://jmlr.org/proceedings/papers/v49/chan16.html>. 13
- W. Chen, F. Li, T. Lin, and A. Rubinfeld. 2015. Combining traditional marketing and viral marketing with amphibious influence maximization. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, June 2015, Portland, Oregon, pp. 779–796. DOI: [10.1145/2764468.2764480](https://doi.org/10.1145/2764468.2764480). 157
- X. Chen and X. Deng. 2009. On the complexity of 2D discrete fixed point problem. *Theoretical Computer Science*, 410(44): 4448–4456. DOI: [10.1016/j.tcs.2009.07.052](https://doi.org/10.1016/j.tcs.2009.07.052). 8
- X. Chen, D. Dai, Y. Du, and S.-H. Teng. 2009a. Settling the complexity of Arrow-Debreu equilibria in markets with additively separable utilities. In *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science*, October 2009, Atlanta, Georgia, pp. 273–282. DOI: [10.1109/FOCS.2009.29](https://doi.org/10.1109/FOCS.2009.29). 9, 117, 125
- X. Chen, X. Deng, and S.-H. Teng. 2009b. Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM*, 56(3). DOI: [10.1145/1516512.1516516](https://doi.org/10.1145/1516512.1516516). 4, 7, 8, 80, 83, 84, 85, 86, 88, 223
- X. Chen, D. Paparas, and M. Yannakakis. 2013. The complexity of non-monotone markets. In *Symposium on Theory of Computing Conference*, June 2013, Palo Alto, California, pp. 181–190. DOI: [10.1145/3064810](https://doi.org/10.1145/3064810). 9, 10, 115, 116, 117, 118, 121, 125, 127
- X. Chen, Y. Cheng, and B. Tang. 2017. Well-supported versus approximate Nash equilibria: Query complexity of large games. <https://dblp.uni-trier.de/rec/bibtex/conf/innovations/ChenCT17> 18, 34, 38
- Y. Chen, G. M. Kamath, C. Suh, and D. Tse. 2016. Community recovery in graphs with locality. In *Proceedings of the Thirty-Third International Conference on Machine Learning*, June 2016, New York City, pp. 689–698. <http://jmlr.org/proceedings/papers/v48/chena16.html>. 179
- Y. Cheng, H. Y. Cheung, S. Dughmi, E. Emamjomeh-Zadeh, L. Han, and S. Teng. 2015. Mixture selection, mechanism design, and signaling. In *Proceedings of the IEEE Fifty-*

- Sixth Annual Symposium on Foundations of Computer Science*, October 2015, Berkeley, California, pp. 1426–1445. DOI: [10.1109/FOCS.2015.91](https://doi.org/10.1109/FOCS.2015.91). 16, 17, 213
- B. Codenotti, B. McCune, and K. Varadarajan. 2005. Market equilibrium via the excess demand function. In *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing*, May 2005, New York, pp. 74–83. DOI: [10.1145/1060590.1060601](https://doi.org/10.1145/1060590.1060601). 9, 116
- R. Cole and L. Fleischer. 2008. Fast-converging tâtonnement algorithms for one-time and ongoing market problems. In *Proceedings of the Fortieth Annual ACM Symposium on the Theory of Computing*, May 2008, Victoria, British Columbia, pp. 315–324. DOI: [10.1145/1374376.1374422](https://doi.org/10.1145/1374376.1374422). 9
- V. Conitzer and T. Sandholm. 2004. Communication complexity as a lower bound for learning in games. In *Proceedings of the Twenty-First International Conference on Machine Learning*, July 2004, Banff, Alberta, p. 24. DOI: [10.1145/1015330.1015351](https://doi.org/10.1145/1015330.1015351). 35, 38
- T. M. Cover and J. A. Thomas. 2012. *Elements of Information Theory*. John Wiley & Sons. 25
- M. Cygan, M. Pilipczuk, and M. Pilipczuk. 2016. Known algorithms for edge clique cover are probably optimal. *SIAM Journal on Computing*, 45(1): 67–83. DOI: [10.1137/130947076.224](https://doi.org/10.1137/130947076.224)
- A. Czumaj, A. Deligkas, M. Fasoulakis, J. Fearnley, M. Jurdzinski, and R. Savani. 2015. Distributed methods for computing approximate equilibria. *CoRR*, abs/1512.03315. <https://dblp.uni-trier.de/rec/bibtex/journals/algoritmica/CzumajDFEJS19>. DOI: [10.1007/978-3-662-54110-4_2](https://doi.org/10.1007/978-3-662-54110-4_2). 38
- A. Daniely. 2016. Complexity theoretic limitations on learning halfspaces. In *Proceedings of the Forty-Eighth Annual ACM SIGACT Symposium on Theory of Computing*, June 2016, Cambridge, Massachusetts, pp. 105–117. DOI: [10.1145/2897518.2897520](https://doi.org/10.1145/2897518.2897520). 190
- A. Daniely and S. Shalev-Shwartz. 2016. Complexity theoretic limitations on learning DNF's. In *Proceedings of the Twenty-Ninth Conference on Learning Theory*, June 2016, New York, pp. 815–830. <http://jmlr.org/proceedings/papers/v49/daniely16.html>. 190
- C. Daskalakis. 2008. *The Complexity of Nash Equilibria*. PhD thesis, University of California, Berkeley. <http://www.eecs.berkeley.edu/Pubs/TechRpts/2008/EECS-2008-107.pdf> 18
- C. Daskalakis. 2013. On the complexity of approximating a Nash equilibrium. *ACM Transactions on Algorithms*, 9(3): 23. DOI: [10.1145/2483699.2483703](https://doi.org/10.1145/2483699.2483703). 224
- C. Daskalakis and C. H. Papadimitriou. 2009. On oblivious PTAS's for Nash equilibrium. In *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*, May 31–June 2, 2009, Bethesda, Maryland, pp. 75–84. Full version available at <http://arxiv.org/abs/1102.2280>. DOI: [10.1145/1536414.1536427](https://doi.org/10.1145/1536414.1536427). 17, 29, 223, 224
- C. Daskalakis and C. H. Papadimitriou. 2015. Approximate Nash equilibria in anonymous games. *J. Economic Theory*, 156: 207–245. DOI: [10.1016/j.jet.2014.02.002](https://doi.org/10.1016/j.jet.2014.02.002). 101
- C. Daskalakis, A. Mehta, and C. H. Papadimitriou. 2007. Progress in approximate nash equilibria. In *Proceedings of the Eighth ACM Conference on Electronic Commerce*, June 2007, San Diego, California, pp. 355–358. DOI: [10.1145/1250910.1250962](https://doi.org/10.1145/1250910.1250962). 17, 223

- C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. 2009a. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1): 195–259. DOI: [10.1137/070699652](https://doi.org/10.1137/070699652). 4, 6, 7, 8, 33, 80, 84, 85, 86, 99, 101, 102, 106, 223, 230, 269
- C. Daskalakis, A. Mehta, and C. H. Papadimitriou. 2009b. A note on approximate Nash equilibria. *Theoretical Computer Science*, 410(17): 1581–1588. DOI: [10.1016/j.tcs.2008.12.031](https://doi.org/10.1016/j.tcs.2008.12.031). 17, 223
- G. Debreu and K. J. Arrow. 1954. Existence of an equilibrium for a competitive economy. *Econometrica*, 22(3): 265–90. DOI: [10.2307/1907353](https://doi.org/10.2307/1907353). 9
- Y. Dekel, O. Gurel-Gurevich, and Y. Peres. 2010. Finding hidden cliques in linear time with high probability. *CoRR*, abs/1010.2997. <https://dblp.org/rec/bibtex/journals/cpc/DekelGP14>. DOI: [10.1017/S096354831300045X](https://doi.org/10.1017/S096354831300045X). 157
- A. Deligkas, J. Fearnley, R. Savani, and P. G. Spirakis. 2017. Computing approximate Nash equilibria in polymatrix games. *Algorithmica*, 77(2): 487–514. DOI: [10.1007/s00453-015-0078-7](https://doi.org/10.1007/s00453-015-0078-7). 224
- H. Dell, T. Husfeldt, D. Marx, N. Taslaman, and M. Wahlen. 2014. Exponential time complexity of the permanent and the Tutte polynomial. *ACM Transactions on Algorithms*, 10(4): 21:1–21:32. DOI: [10.1145/2635812](https://doi.org/10.1145/2635812). 22
- X. Deng and Y. Du. 2008. The computation of approximate competitive equilibrium is PPA-hard. *Information Processing Letters*, 108(6): 369–373. DOI: [10.1016/j.ipl.2008.07.011](https://doi.org/10.1016/j.ipl.2008.07.011). 9
- Y. Deshpande and A. Montanari. 2015. Improved sum-of-squares lower bounds for hidden clique and hidden submatrix problems. *CoRR*, abs/1502.06590. <https://dblp.org/rec/bibtex/conf/colt/DeshpandeM15>. 157
- N. R. Devanur, C. H. Papadimitriou, A. Saberi, and V. V. Vazirani. Nov. 2008. Market equilibrium via a primal–dual algorithm for a convex program. *Journal of the ACM*, 55(5): 22:1–22:18. DOI: [10.1145/1411509.1411512](https://doi.org/10.1145/1411509.1411512). 9, 11
- I. Dinur. 2007. The PCP theorem by gap amplification. *Journal of the ACM*, 54(3): 12. DOI: [10.1145/1236457.1236459](https://doi.org/10.1145/1236457.1236459). 23
- I. Dinur and P. Harsha. 2013. Composition of low-error 2-query PCPs using decodable PCPs. *SIAM Journal on Computing*, 42(6): 2452–2486. DOI: [10.1137/100788161](https://doi.org/10.1137/100788161). 228
- S. Dughmi. 2014. On the hardness of signaling. In *Proceedings of the Fifty-Fifth IEEE Annual Symposium on Foundations of Computer Science*, October 2014, Philadelphia, Pennsylvania, pp. 354–363. DOI: [10.1109/FOCS.2014.45](https://doi.org/10.1109/FOCS.2014.45). 16, 17, 213, 215
- S. Dughmi, N. Immorlica, and A. Roth. 2013. Constrained signaling for welfare and revenue maximization. *SIGecom Exchanges*, 12(1): 53–56. DOI: [10.1145/2509013.2509022](https://doi.org/10.1145/2509013.2509022). 16, 213
- F. Y. Edgeworth. 1909. Review of free trade in being. *Economic Journal*, 19: 104–105. 9
- Y. Emek, M. Feldman, I. Gamzu, R. P. Leme, and M. Tennenholtz. 2014. Signaling schemes for revenue maximization. *ACM Transactions on Economics and Computing*, 2(2): 5. DOI: [10.1145/2594564](https://doi.org/10.1145/2594564). 16, 213

- K. Etessami and M. Yannakakis. 2010. On the complexity of Nash equilibria and other fixed points. *SIAM Journal on Computing*, 39(6): 2531–2597. DOI: [10.1137/080720826](https://doi.org/10.1137/080720826). 7
- J. Fearnley, M. Gairing, P. W. Goldberg, and R. Savani. 2013. Learning equilibria of games via payoff queries. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, June 2013, Philadelphia, Pennsylvania, pp. 397–414. DOI: [10.1145/2492002.2482558](https://doi.org/10.1145/2492002.2482558). 18, 34
- T. Feder, H. Nazerzadeh, and A. Saberi. 2007. Approximating Nash equilibria using small-support strategies. In *Proceedings of the Eighth ACM Conference on Electronic Commerce*, June 2007, San Diego, California, pp. 352–354. DOI: [10.1145/1250910.1250961](https://doi.org/10.1145/1250910.1250961). 224
- U. Feige. 2002. Relations between average case complexity and approximation complexity. In *Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing*, Montreal Quebec, May 2002, New York, pp. 534–543. DOI: [10.1145/509907.509985](https://doi.org/10.1145/509907.509985). 156
- U. Feige and R. Krauthgamer. 2000. Finding and certifying a large hidden clique in a semirandom graph. *Random Structures and Algorithms*, 16(2): 195–208. DOI: [10.1002/\(SICI\)1098-2418\(200003\)16:2<195::AID-RSA5>3.0.CO;2-A](https://doi.org/10.1002/(SICI)1098-2418(200003)16:2<195::AID-RSA5>3.0.CO;2-A). 157
- U. Feige and M. Seltser. 1997. *On the densest k -subgraph problem*. Citeseer. DOI: [10.1.1.37.9962](https://doi.org/10.1.1.37.9962). 14, 155, 156
- U. Feige, S. Goldwasser, L. Lovász, S. Safra, and M. Szegedy. 1996. Interactive proofs and the hardness of approximating cliques. *Journal of the ACM (JACM)*, 43(2): 268–292. DOI: [10.1145/226643.226652](https://doi.org/10.1145/226643.226652). 14, 155, 180, 184
- U. Feige, G. Kortsarz, and D. Peleg. 2001. The dense k -subgraph problem. *Algorithmica*, 29(3): 410–421. DOI: [10.1007/s004530010050](https://doi.org/10.1007/s004530010050). 156
- J. Feigenbaum, D. Koller, and P. W. Shor. 1995. A game-theoretic classification of interactive complexity classes. In *Proceedings of the Structure in Complexity Theory Conference*, June 1995, Minneapolis, Minnesota, pp. 227–237. DOI: [10.1109/SCT.1995.514861](https://doi.org/10.1109/SCT.1995.514861). 79
- V. Feldman, P. Gopalan, S. Khot, and A. K. Ponnuswami. 2006. New results for learning noisy parities and halfspaces. In *Proceedings of the Forty-Seventh Annual IEEE Symposium on Foundations of Computer Science*, October 2006, Berkeley, California, pp. 563–574. DOI: [10.1109/FOCS.2006.51](https://doi.org/10.1109/FOCS.2006.51). 190
- V. Feldman, E. Grigorescu, L. Reyzin, S. Vempala, and Y. Xiao. 2013. Statistical algorithms and a lower bound for detecting planted cliques. In *Symposium on Theory of Computing Conference*, June 2013, Palo Alto, California, pp. 655–664. DOI: [10.1145/2488608.2488692](https://doi.org/10.1145/2488608.2488692). 157
- L. Florescu and W. Perkins. 2016. Spectral thresholds in the bipartite stochastic block model. In *Proceedings of the 29th Conference on Learning Theory*, June 2013, New York, pp. 943–959. <http://jmlr.org/proceedings/papers/v49/florescu16.html>. 179
- D. Foley. 1967. Resource allocation and the public sector. *Yale Economic Essays*, 7(1): 45–98. 10

- L. Fortnow, R. Impagliazzo, V. Kabanets, and C. Umans. 2008. On the complexity of succinct zero-sum games. *Computational Complexity*, 17(3): 353–376. DOI: [10.1109/CCC.2005.18](https://doi.org/10.1109/CCC.2005.18). 79
- S. Fortunato. 2010. Community detection in graphs. *Physics Reports*, 486: 75–174. DOI: [10.1016/j.physrep.2009.11.002](https://doi.org/10.1016/j.physrep.2009.11.002). 14, 177
- D. P. Foster and H. P. Young. Sept. 2006. Regret testing: learning to play Nash equilibrium without knowing you have an opponent. *Theoretical Economics*, 1(3): 341–367. <http://econtheory.org/ojs/index.php/te/article/view/20060341>. 4, 35
- M. Frances and A. Litman. 1998. Optimal mistake bound learning is hard. *Information and Computation*, 144(1): 66–82. DOI: [10.1006/inco.1998.2709](https://doi.org/10.1006/inco.1998.2709). 16, 187, 190
- A. Ganor and Karthik. 2017. Communication complexity of correlated equilibrium in two-player games. arXiv preprint arXiv:1704.01104. 60
- J. Garg, R. Mehta, M. A. Sohoni, and V. V. Vazirani. 2015. A complementary pivot algorithm for market equilibrium under separable, piecewise-linear concave utilities. *SIAM Journal on Computing*, 44(6): 1820–1847. DOI: [10.1137/140971002](https://doi.org/10.1137/140971002). 9
- J. Garg, R. Mehta, V. V. Vazirani, and S. Yazdanbod. 2017. Settling the complexity of Leontief and PLC exchange markets under exact and approximate equilibria. In *Proceedings of the Forty-Ninth Annual ACM SIGACT Symposium on Theory of Computing*, June 2017, Montreal, Quebec, pp. 890–901. DOI: [10.1145/3055399.3055474](https://doi.org/10.1145/3055399.3055474). 9
- R. Garg and S. Kapoor. 2006. Auction algorithms for market equilibrium. *Math. Oper. Res.*, 31(4): 714–729. DOI: [10.1287/moor.1060.0216](https://doi.org/10.1287/moor.1060.0216). 9
- S. Garg, O. Pandey, and A. Srinivasan. 2016. Revisiting the cryptographic hardness of finding a nash equilibrium. In *Advances in Cryptology: Proceedings of the Thirty-Sixth Annual International Cryptology Conference*, August 2016, Santa Barbara, California, Part II, pp. 579–604. DOI: [10.1007/978-3-662-53008-5_20](https://doi.org/10.1007/978-3-662-53008-5_20). 6
- F. Germano and G. Lugosi. 2007. Global Nash convergence of Foster and Young’s regret testing. *Games and Economic Behavior*, 60(1): 135–154. DOI: [10.1016/j.geb.2006.06.001](https://doi.org/10.1016/j.geb.2006.06.001). 35
- A. Ghodsi, M. Zaharia, B. Hindman, A. Konwinski, S. Shenker, and I. Stoica. 2011. Dominant resource fairness: Fair allocation of multiple resource types. In *Proceedings of the Eighth USENIX Conference on Networked Systems Design and Implementation*, March 30–April 1, 2011, Boston, Massachusetts, pp. 323–336. <http://dl.acm.org/citation.cfm?id=1972457.1972490>. 11
- I. Gilboa and E. Zemel. 1989. Nash and correlated equilibria: Some complexity considerations. <https://www.sciencedirect.com/science/article/pii/0899825689900067> 150
- P. W. Goldberg and A. Pastink. 2014. On the communication complexity of approximate Nash equilibria. *Games and Economic Behavior*, 85: 19–31. DOI: [10.1007/978-3-642-33996-7_17](https://doi.org/10.1007/978-3-642-33996-7_17). 38
- P. W. Goldberg and A. Roth. 2016. Bounds for the query complexity of approximate equilibria. *ACM Transactions on Economics and Computing*, 4(4): 24:1–24:25. DOI: [10.1145/2956582](https://doi.org/10.1145/2956582). 38, 39, 101

- O. Goldreich, ed. 2010. *Property Testing—Current Research and Surveys* [outgrowth of a workshop at the Institute for Computer Science (ITCS) at Tsinghua University, January 2010], volume 6390 of *Lecture Notes in Computer Science*. Springer. DOI: [10.1007/978-3-642-16367-8](https://doi.org/10.1007/978-3-642-16367-8). 228
- M. Göös, T. Pitassi, and T. Watson. 2017. Query-to-communication lifting for BPP. DOI: [10.1007/s00037-018-0175-5](https://doi.org/10.1007/s00037-018-0175-5). 30, 37, 40
- V. Guruswami and A. Rudra. 2005. Tolerant locally testable codes. In *Approximation, Randomization and Combinatorial Optimization, Algorithms and Techniques, Proceedings of the Eighth International Workshop on Approximation Algorithms for Combinatorial Optimization Problems, and Ninth International Workshop on Randomization and Computation*, August 2005, Berkeley, California, pp. 306–317. DOI: [10.1007/11538462_26](https://doi.org/10.1007/11538462_26). 228
- B. E. Hajek, Y. Wu, and J. Xu. 2016. Semidefinite programs for exact recovery of a hidden community. In *Proceedings of the Twenty-Ninth Conference on Learning Theory*, June 2016, New York, pp. 1051–1095. <http://jmlr.org/proceedings/papers/v49/hajek16.html>. 179
- S. Hart and Y. Mansour. 2010. How long to equilibrium? The communication complexity of uncoupled equilibrium procedures. *Games and Economic Behavior*, 69(1): 107–126. DOI: [10.1016/j.geb.2007.12.002](https://doi.org/10.1016/j.geb.2007.12.002). 18, 33, 34, 38, 39, 60
- S. Hart and A. Mas-Colell. 2003. Uncoupled dynamics do not lead to nash equilibrium. *American Economic Review*, 93(5): 1830–1836. DOI: [10.1257/1000282803322655581](https://doi.org/10.1257/1000282803322655581). 4, 33, 35
- S. Hart and A. Mas-Colell. 2006. Stochastic uncoupled dynamics and Nash equilibrium. *Games and Economic Behavior*, 57(2): 286–303. DOI: [10.1016/j.geb.2005.09.007](https://doi.org/10.1016/j.geb.2005.09.007). 33, 35
- S. Hart and A. Mas-Colell. 2013. *Simple adaptive strategies: From regret-matching to uncoupled dynamics*, volume 4. World Scientific. 35
- S. Hart and N. Nisan. 2013. The query complexity of correlated equilibria. *CoRR*, abs/1305.4874. <https://dblp.uni-trier.de/rec/bibtex/journals/geb/HartN18>. DOI: [10.1016/j.geb.2016.11.003](https://doi.org/10.1016/j.geb.2016.11.003). 18, 34, 38, 39
- J. Håstad. 1999. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Mathematica*, 182(1): 105–142. DOI: [10.1007/BF02392825](https://doi.org/10.1007/BF02392825). 14, 155
- E. Hazan and R. Krauthgamer. 2011. How hard is it to approximate the best Nash equilibrium? *SIAM Journal on Computing*, 40(1): 79–91. DOI: [10.1137/090766991](https://doi.org/10.1137/090766991). 157, 224
- M. D. Hirsch, C. H. Papadimitriou, and S. A. Vavasis. 1989. Exponential lower bounds for finding Brouwer fix points. *J. Complexity*, 5(4): 379–416. DOI: [10.1109/SFCS.1987.24](https://doi.org/10.1109/SFCS.1987.24). 37, 38, 62, 63, 65, 69, 85, 223, 225
- P. W. Holland, K. B. Laskey, and S. Leinhardt. 1983. Stochastic blockmodels: First steps. *Social Networks*, 5: 109–137. DOI: [10.1016/0378-8733\(83\)90021-7](https://doi.org/10.1016/0378-8733(83)90021-7). 179

- S. B. Hopkins, P. Kothari, A. H. Potechin, P. Raghavendra, and T. Schramm. 2016. On the integrality gap of degree-4 sum of squares for planted clique. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, January 2016, Arlington, Virginia, pp. 1079–1095. DOI: [10.1137/1.9781611974331.ch76](https://doi.org/10.1137/1.9781611974331.ch76). 157
- P. Hubáček and E. Yogev. 2017. Hardness of continuous local search: Query complexity and cryptographic lower bounds. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, January 2017, Barcelona, Spain, pp. 1352–1371. DOI: [10.1137/1.9781611974782.88](https://doi.org/10.1137/1.9781611974782.88). 6
- R. Iemhoff. 2016. Intuitionism in the philosophy of mathematics. In E. N. Zalta, ed., *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2016. 7
- R. Impagliazzo and R. Paturi. 2001. On the complexity of k -SAT. *Journal of Computer System Sciences*, 62(2): 367–375. DOI: [10.1006/jcss.2000.1727](https://doi.org/10.1006/jcss.2000.1727). 224
- R. Impagliazzo, R. Paturi, and F. Zane. 2001. Which problems have strongly exponential complexity? *Journal of Computer System Sciences*, 63(4): 512–530. DOI: [10.1006/jcss.2001.1774](https://doi.org/10.1006/jcss.2001.1774). 5, 13, 17, 22, 224
- K. Jain. 2004. A polynomial time algorithm for computing the Arrow-Debreu market equilibrium for linear utilities. In *Proceedings of the Forty-Fifth Symposium on Foundations of Computer Science*, October 2004, Rome, Italy, pp. 286–294. DOI: [10.1109/FOCS.2004.6](https://doi.org/10.1109/FOCS.2004.6). 9
- K. Jain and V. V. Vazirani. 2010. Eisenberg-Gale markets: Algorithms and game-theoretic properties. *Games and Economic Behavior*, 70(1): 84–106. DOI: [10.1016/j.geb.2008.11.011](https://doi.org/10.1016/j.geb.2008.11.011). 9
- M. Jerrum. 1992. Large cliques elude the metropolis process. *Random Struct. Algorithms*, 3(4): 347–360. DOI: [10.1002/rsa.3240030402](https://doi.org/10.1002/rsa.3240030402). 157
- W. S. Jevons. 1866. Brief account of a general mathematical theory of political economy. *History of Economic Thought Articles*, 29: 282–287. <http://EconPapers.repec.org/RePEc:hay:hetart:jevons1866>. 8
- A. X. Jiang and K. Leyton-Brown. 2015. Polynomial-time computation of exact correlated equilibrium in compact games. *Games and Economic Behavior*, 91: 347–359. DOI: [10.1016/j.geb.2013.02.002](https://doi.org/10.1016/j.geb.2013.02.002). 39, 60
- A. T. Kalai, A. R. Klivans, Y. Mansour, and R. A. Servedio. 2008. Agnostically learning halfspaces. *SIAM Journal on Computing*, 37(6): 1777–1805. DOI: [10.1137/060649057](https://doi.org/10.1137/060649057). 190
- E. Kalai and E. Lehrer. 1993. Rational learning leads to Nash equilibrium. *Econometrica*, 61(5): 1019–1045. DOI: [10.2307/2951492](https://doi.org/10.2307/2951492). 4, 35
- R. Kannan and T. Theobald. 2007. Games of fixed rank: A hierarchy of bimatrix games. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, January 2007, New Orleans, Louisiana, pp. 1124–1132. <http://dl.acm.org/citation.cfm?id=1283383.1283504>. DOI: [10.1145/1283383.1283504](https://doi.org/10.1145/1283383.1283504). 17, 223, 224

- M. Karchmer, R. Raz, and A. Wigderson. 1995. Super-logarithmic depth lower bounds via the direct sum in communication complexity. *Computational Complexity*, 5(3-4): 191–204. DOI: [10.1109/SCT.1991.160273](https://doi.org/10.1109/SCT.1991.160273). 37
- R. M. Karp. 1972. Reducibility among combinatorial problems. In *Proceedings of a Symposium on the Complexity of Computer Computations*, March 1972, Yorktown Heights, New York, pp. 85–103. <http://www.cs.berkeley.edu/~luca/cs172/karp.pdf>. DOI: [10.1007/978-3-540-68279-0_8](https://doi.org/10.1007/978-3-540-68279-0_8). 14, 155
- M. Karpinski and W. Schudy. 2010. Faster algorithms for feedback arc set tournament, kemeny rank aggregation and betweenness tournament. In *Algorithms and Computation: Proceedings of the Twenty-First International Symposium*, December 2010, Jeju Island, Korea, *Part I*, pp. 3–14. DOI: [10.1007/978-3-642-17517-6_3](https://doi.org/10.1007/978-3-642-17517-6_3). 224
- M. Kearns. 2007. *Graphical games*, chapter 7, pp. 159–180. Cambridge University Press. 79
- M. J. Kearns and L. G. Valiant. 1994. Cryptographic limitations on learning Boolean formulae and finite automata. *Journal of the ACM*, 41(1): 67–95. DOI: [10.1145/174644.174647](https://doi.org/10.1145/174644.174647). 190
- M. Kharitonov. 1993. Cryptographic hardness of distribution-specific learning. In *Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing*, May 1993, San Diego, California, pp. 372–381. DOI: [10.1145/167088.167197](https://doi.org/10.1145/167088.167197). 190
- M. Kharitonov. 1995. Cryptographic lower bounds for learnability of Boolean functions on the uniform distribution. *Journal of Computer System Sciences*, 50(3): 600–610. DOI: [10.1006/jcss.1995.1046](https://doi.org/10.1006/jcss.1995.1046). 190
- S. Khot. 2001. Improved inapproximability results for MaxClique, chromatic number and approximate graph coloring. In *Proceedings of the Forty-Second Annual Symposium on Foundations of Computer Science*, October 2001, Las Vegas, Nevada, pp. 600–609. DOI: [10.1109/SFCS.2001.959936](https://doi.org/10.1109/SFCS.2001.959936). 14, 155
- S. Khot. 2006. Ruling out PTAs for graph min-bisection, dense k -subgraph, and bipartite clique. *SIAM Journal on Computing*, 36(4): 1025–1071. DOI: [10.1109/FOCS.2004.59](https://doi.org/10.1109/FOCS.2004.59). 157
- J. M. Kleinberg. 2002. An impossibility theorem for clustering. In *Proceedings of the Advances in Neural Information Processing Systems 15*, December 2002, Vancouver, British Columbia, pp. 446–453. <http://papers.nips.cc/paper/2340-an-impossibility-theorem-for-clustering>. 179
- A. R. Klivans. 2016. Cryptographic hardness of learning. In *Encyclopedia of Algorithms*, pp. 475–477. Springer. DOI: [10.1007/978-1-4939-2864-4_96](https://doi.org/10.1007/978-1-4939-2864-4_96). 190
- A. R. Klivans and A. A. Sherstov. 2009. Cryptographic hardness for learning intersections of halfspaces. *Journal of Computer System Sciences*, 75(1): 2–12. DOI: [10.1016/j.jcss.2008.07.008](https://doi.org/10.1016/j.jcss.2008.07.008). 190
- P. Koiran and A. Zouzias. 2011. On the certification of the restricted isometry property. *CoRR*, abs/1103.4984. <http://arxiv.org/abs/1103.4984>. 157

- S. C. Kontogiannis, P. N. Panagopoulou, and P. G. Spirakis. 2009. Polynomial algorithms for approximating Nash equilibria of bimatrix games. *Theoretical Computer Science*, 410(17): 1599–1606. DOI: [10.1016/j.tcs.2008.12.033](https://doi.org/10.1016/j.tcs.2008.12.033). 17, 223
- L. Kucera. 1995. Expected complexity of graph partitioning problems. *Discrete Applied Mathematics*, 57(2–3): 193–212. DOI: [10.1016/0166-218X\(94\)00103-K](https://doi.org/10.1016/0166-218X(94)00103-K). 157
- F. T. Leighton. 1992. *Introduction to Parallel Algorithms and Architectures: Array, Trees, Hypercubes*. Morgan Kaufmann Publishers Inc. 239
- C. E. Lemke and J. T. Howson. 1964. Equilibrium points of bimatrix games. *Journal of the Society for Industrial and Applied Mathematics*, 12(2): 413–423. DOI: [10.1137/0112033](https://doi.org/10.1137/0112033). 4, 223
- N. Linial, Y. Mansour, and R. L. Rivest. 1991. Results on learnability and the Vapnik-Chervonenkis dimension. *Information and Computation*, 90(1): 33–49. DOI: [10.1016/0890-5401\(91\)90058-A](https://doi.org/10.1016/0890-5401(91)90058-A). 190, 191
- R. J. Lipton, E. Markakis, and A. Mehta. 2003. Playing large games using simple strategies. In *Proceedings of the Fourth ACM Conference on Electronic Commerce*, June 2003, San Diego, California, pp. 36–41. DOI: [10.1145/779928.779933](https://doi.org/10.1145/779928.779933). 12, 17, 36, 39, 79, 150, 223, 224
- N. Littlestone. 1987. Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm. *Machine Learning*, 2(4): 285–318. DOI: [10.1007/BF00116827](https://doi.org/10.1007/BF00116827). 16, 24, 25, 187
- D. Lokshantov, D. Marx, and S. Saurabh. 2011. Lower bounds based on the exponential time hypothesis. *Bulletin of the EATCS*, 105: 41–72. <http://albcm.lsi.upc.edu/ojs/index.php/beatcs/article/view/96>. 224
- K. Makarychev, Y. Makarychev, and A. Vijayaraghavan. 2016. Learning communities in the presence of errors. In *Proceedings of the Twenty-Ninth Conference on Learning Theory*, June 2016, New York, pp. 1258–1291. <http://jmlr.org/proceedings/papers/v49/makarychev16.html>. 179
- P. Manurangsi. 2017. Almost-polynomial ratio ETH-hardness of approximating densest k -subgraph. In *Proceedings of the Forty-Ninth Annual ACM Symposium on Theory of Computing*. DOI: [10.1145/3055399.3055412](https://doi.org/10.1145/3055399.3055412). 158, 159, 189
- P. Manurangsi and P. Raghavendra. 2016. A birthday repetition theorem and complexity of approximating dense CSPs. *CoRR*, abs/1607.02986. <https://dblp.org/rec/bibtex/conf/icalp/ManurangsiR17>. DOI: [10.4230/LIPIcs.ICALP.2017.78](https://doi.org/10.4230/LIPIcs.ICALP.2017.78). 189
- P. Manurangsi and A. Rubinfeld. 2017. Inapproximability of VC dimension and Littlestone’s dimension. In *Proceedings of the Thirtieth Conference on Learning Theory*, July 2017, Amsterdam, The Netherlands, pp. 1432–1460. <http://proceedings.mlr.press/v65/manurangsi17a.html>. xiv
- A. Marshall. 1895. *Principles of Economics*. Number III. Macmillan. 9
- R. R. Maxfield. 1997. General equilibrium and the theory of directed graphs. *Journal of Mathematical Economics*, 27(1): 23–51. DOI: [10.1016/0304-4068\(95\)00763-6](https://doi.org/10.1016/0304-4068(95)00763-6). 123

- A. McLennan and R. Tourky. 2005. From imitation games to Kakutani. Manuscript, available at <http://www.econ.umn.edu/mclennan/Papers/papers.html>. 37, 41
- N. Megiddo and C. H. Papadimitriou. 1991. On total functions, existence theorems and computational complexity. *Theoretical Computer Science*, 81(2): 317–324. DOI: [10.1016/0304-3975\(91\)90200-L](https://doi.org/10.1016/0304-3975(91)90200-L). 6
- R. Meka, A. Potechin, and A. Wigderson. 2015. Sum-of-squares lower bounds for planted clique. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing*, June 2015, Portland, Oregon, pp. 87–96. DOI: [10.1145/2746539.2746600](https://doi.org/10.1145/2746539.2746600). 157
- C. Menger. 1871. *Principles of Economics*. Ludwig von Mises Institute. 8
- P. B. Miltersen and O. Sheffet. 2012. Send mixed signals: earn more, work less. In *ACM Conference on Electronic Commerce*, June 2012, Valencia, Spain, pp. 234–247. DOI: [10.1145/2229012.2229033](https://doi.org/10.1145/2229012.2229033). 16, 213
- N. Mishra, R. Schreiber, I. Stanton, and R. E. Tarjan. 2008. Finding strongly knit clusters in social networks. Volume 5, pp. 155–174. <https://dblp.uni-trier.de/rec/bibtex/journals/im/MishraSST08> DOI: [10.1080/15427951.2008.10129299](https://doi.org/10.1080/15427951.2008.10129299). 15, 177, 178
- A. Moitra, W. Perry, and A. S. Wein. 2016. How robust are reconstruction thresholds for community detection? In *Proceedings of the Forty-Eighth Annual ACM SIGACT Symposium on Theory of Computing*, June 2016, Cambridge, Massachusetts, pp. 828–841. DOI: [10.1145/2897518.2897573](https://doi.org/10.1145/2897518.2897573). 179
- T. Morioka. 2003. The relative complexity of local search heuristics and the iteration principle. *Electronic Colloquium on Computational Complexity (ECCC)*, (051). <http://eccc.hpi-web.de/eccc-reports/2003/TR03-051/index.html>. 6
- D. Moshkovitz and R. Raz. 2010. Two-query PCP with subconstant error. *Journal of the ACM*, 57(5): 29:1–29:29. DOI: [10.1145/1754399.1754402](https://doi.org/10.1145/1754399.1754402). 24, 228
- E. Mossel and C. Umans. 2002. On the complexity of approximating the VC dimension. *Journal of Computer System Sciences*, 65(4): 660–671. DOI: [10.1016/S0022-0000\(02\)00022-3](https://doi.org/10.1016/S0022-0000(02)00022-3). 190
- E. Mossel and J. Xu. 2016. Density evolution in the degree-correlated stochastic block model. In *Proceedings of the Twenty-Ninth Conference on Learning Theory*, June 2016, New York, pp. 1319–1356. <http://jmlr.org/proceedings/papers/v49/mossel16.html>. <https://dblp.uni-trier.de/rec/bibtex/conf/colt/MosselX16>. 179
- J. Nash. 1951. Non-cooperative games. *The Annals of Mathematics*, 54: 286–295. DOI: [10.2307/1969529](https://doi.org/10.2307/1969529). 3, 35
- N. Nisan. 2009a. Communication complexity of mixed-Nash equilibria. <https://agtb.wordpress.com/2009/09/28/communication-complexity-of-mixed-nash-equilibria/>. 18, 34
- N. Nisan. 2009b. Economists and complexity. <https://agtb.wordpress.com/2009/11/27/economists-and-complexity/>. 9, 33
- A. Othman, T. Sandholm, and E. Budish. 2010. Finding approximate competitive equilibria: efficient and fair course allocation. In *Ninth International Conference on Autonomous*

- Agents and Multiagent Systems*, May 2010, Toronto, Canada, *Volume 1–3*, pp. 873–880. DOI: [10.1145/1838206.1838323](https://doi.org/10.1145/1838206.1838323). 11
- A. Othman, C. H. Papadimitriou, and A. Rubinstein. 2016. The complexity of fairness through equilibrium. *ACM Transactions on Economics and Computing*, 4(4): 20. DOI: [10.1145/2956583](https://doi.org/10.1145/2956583). xiv
- C. H. Papadimitriou. 1994. On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer System Sciences*, 48(3): 498–532. DOI: [10.1016/80022-0000\(05\)80063-7](https://doi.org/10.1016/80022-0000(05)80063-7). 5, 6, 143, 223
- C. H. Papadimitriou and T. Roughgarden. 2008. Computing correlated equilibria in multi-player games. *Journal of the ACM*, 55(3), Article 14. DOI: [10.1145/1379759.1379762](https://doi.org/10.1145/1379759.1379762). 39, 60
- C. H. Papadimitriou and M. Yannakakis. 1986. A note on succinct representations of graphs. *Information and Control*, 71(3): 181–185. DOI: [10.1016/S0019-9958\(86\)80009-2](https://doi.org/10.1016/S0019-9958(86)80009-2). 16
- C. H. Papadimitriou and M. Yannakakis. 1996. On limited nondeterminism and the complexity of the V-C dimension. *Journal of Computer System Sciences*, 53(2): 161–170. DOI: [10.1006/jcss.1996.0058](https://doi.org/10.1006/jcss.1996.0058). 16, 187, 190
- M. Parnas, D. Ron, and R. Rubinfeld. 2006. Tolerant property testing and distance approximation. *Journal of Computer System Sciences*, 72(6): 1012–1042. DOI: [10.1016/j.jcss.2006.03.002](https://doi.org/10.1016/j.jcss.2006.03.002). 228
- A. Polishchuk and D. A. Spielman. 1994. Nearly linear size holographic proofs. In *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing*, May 1994, Montreal, Quebec, pp. 194–203. DOI: [10.1145/195058.195132](https://doi.org/10.1145/195058.195132). 229, 231
- P. Raghavendra and D. Steurer. 2010. Graph expansion and the unique games conjecture. In *Proceedings of the Forty-Second ACM Symposium on Theory of Computing*, June 2010, Cambridge, Massachusetts, pp. 755–764. DOI: [10.1145/1806689.1806792](https://doi.org/10.1145/1806689.1806792). 156
- R. Raz and P. McKenzie. 1999. Separation of the monotone NC hierarchy. *Combinatorica*, 19(3): 403–435. DOI: [10.1007/s004930050062](https://doi.org/10.1007/s004930050062). 37
- R. Raz and A. Wigderson. 1990. Monotone circuits for matching require linear depth. In *Proceedings of the Twenty-Second Annual ACM Symposium on Theory of Computing*, May 1990, Baltimore, Maryland, pp. 287–292. DOI: [10.1145/100216.100253](https://doi.org/10.1145/100216.100253). 37
- J. Robinson. 1951. An iterative method of solving a game. *Annals of Mathematics*, 54: 296–301. DOI: [10.2307/1969530](https://doi.org/10.2307/1969530). 4
- T. Roughgarden. 2014. Barriers to near-optimal equilibria. In *Fifty-Fifth IEEE Annual Symposium on Foundations of Computer Science*, October 2014, Philadelphia, Pennsylvania, pp. 71–80. DOI: [10.1109/FOCS.2014.16](https://doi.org/10.1109/FOCS.2014.16). 39
- T. Roughgarden and O. Weinstein. 2016. On the communication complexity of approximate fixed points. In *Electronic Colloquium on Computational Complexity (ECCC)*, volume 23, p. 55. DOI: [10.1109/FOCS.2016.32](https://doi.org/10.1109/FOCS.2016.32). 7, 18, 39
- A. Rubinstein. 2015. Inapproximability of Nash equilibrium. In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*, June 2015, Portland, Oregon, pp. 409–418. DOI: [10.1145/2746539.2746578](https://doi.org/10.1145/2746539.2746578). xiv, 224

- A. Rubinstein. 2016. Settling the complexity of computing approximate two-player Nash equilibria. In *Proceedings of the IEEE Fifty-Seventh Annual Symposium on Foundations of Computer Science*, October 2016, New Brunswick, New Jersey, pp. 258–265. DOI: [10.1109/FOCS.2016.35](https://doi.org/10.1109/FOCS.2016.35). xiv, 41, 44
- A. Rubinstein. 2017a. Detecting communities is hard and counting them is even harder. <https://dblp.uni-trier.de/rec/bibtex/conf/innovations/Rubinstein17>. DOI: [10.4230/LIPIcs.ITCS.2017.42](https://doi.org/10.4230/LIPIcs.ITCS.2017.42). xiv
- A. Rubinstein. 2017b. Honest signaling in zero-sum games is hard, and lying is even harder. In *Forty-Fourth International Colloquium on Automata, Languages, and Programming*, July 2017, Warsaw, Poland, pp. 77:1–77:13. DOI: [10.4230/LIPIcs.ICALP.2017.77](https://doi.org/10.4230/LIPIcs.ICALP.2017.77). xiv, 214
- A. Rubinstein. 2017c. *Hardness of Approximation Between P and NP*. PhD thesis, University of California at Berkeley, USA. <https://www2.eecs.berkeley.edu/Pubs/TechRpts/2017/EECS-2017-146.pdf>. xiv
- H. Scarf. 1967. The approximation of fixed points of a continuous mapping. *SIAM Journal of Applied Mathematics*, 15: 1328–1343. DOI: [10.1137/01151116](https://doi.org/10.1137/01151116). 4, 8
- M. Schaefer. 1999. Deciding the Vapnik-Červonenkis dimension is Σ_3^P -complete. *Journal of Computer System Sciences*, 58(1): 177–182. DOI: [10.1006/jcss.1998.1602](https://doi.org/10.1006/jcss.1998.1602). 190
- M. Schaefer. 2000. Deciding the k -dimension is PSPACE-complete. In *Proceedings of the Fifteenth Annual IEEE Conference on Computational Complexity*, July 2000, Florence, Italy, pp. 198–203. DOI: [10.1109/CCC.2000.856750](https://doi.org/10.1109/CCC.2000.856750). 190
- J. P. Schmidt, A. Siegel, and A. Srinivasan. 1995. Chernoff-Hoeffding bounds for applications with limited independence. *SIAM J. Discrete Math.*, 8(2): 223–250. DOI: [10.1137/S089548019223872X](https://doi.org/10.1137/S089548019223872X). 27
- G. Schoenebeck and S. P. Vadhan. 2012. The computational complexity of Nash equilibria in concisely represented games. *ACM Transactions on Computation Theory*, 4(2): 4. DOI: [10.1145/2189778.2189779](https://doi.org/10.1145/2189778.2189779). 79
- A. Shamir. 1992. IP = PSPACE. *Journal of the ACM*, 39(4): 869–877. DOI: [10.1145/146585.146609](https://doi.org/10.1145/146585.146609). 229
- L. S. Shapley. 1964. Some topics in two-person games. *Annals of Mathematical Studies*, 52: 1–28. <https://apps.dtic.mil/dtic/tr/fulltext/u2/407345.pdf>. 4
- E. Shmaya. 2012. Brouwer implies Nash implies Brouwer. <http://theoryclass.wordpress.com/2012/01/05/brouwer-implies-nash-implies-brouwer/>. 37, 41
- S. P. Singh, V. Soni, and M. P. Wellman. 2004. Computing approximate Bayes-Nash equilibria in tree-games of incomplete information. In *Proceedings of the Fifth ACM Conference on Electronic Commerce*, May 2004, New York, pp. 81–90. DOI: [10.1145/988772.988785](https://doi.org/10.1145/988772.988785). 111
- D. A. Spielman. 1995. *Computationally Efficient Error-Correcting Codes and Holographic Proofs*. PhD thesis, Massachusetts Institute of Technology. 229, 231
- G. J. Stigler. 1947. Notes on the history of the Giffen paradox. *Journal of Political Economy*, 55(2): 152–156. <https://www.jstor.org/stable/1825304>. 9

- W. Thomson and H. Varian. 1985. Theories of justice based on symmetry. *Social Goals and Social Organizations: Essays in Memory of Elisha Pazner*. Cambridge University Press. 10
- N. Tremblay, G. Puy, R. Gribonval, and P. Vandergheynst. 2016. Compressive spectral clustering. In *Proceedings of the Thirty-Third International Conference on Machine Learning*, June 2016, New York City, pp. 1002–1011. <http://jmlr.org/proceedings/papers/v48/tremblay16.html>. 179
- H. Tsaknakis and P. G. Spirakis. 2008. An optimization approach for approximate Nash equilibria. *Internet Mathematics*, 5(4): 365–382. DOI: [10.1080/15427951.2008.10129172](https://doi.org/10.1080/15427951.2008.10129172). 17, 223, 224
- T. van Laarhoven and E. Marchiori. 2014. Axioms for graph clustering quality functions. *Journal of Machine Learning Research*, 15(1): 193–215. <http://dl.acm.org/citation.cfm?id=2627441>. 179
- V. N. Vapnik and A. Y. Chervonenkis. 1971. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability & Its Applications*, 16(2): 264–280. DOI: [10.1137/1116025](https://doi.org/10.1137/1116025). 16, 24
- H. Varian. 1974. Equity, envy, and efficiency. *Journal of Economic Theory*, 9(1): 63–91. DOI: [10.1016/0022-0531\(74\)90075-1](https://doi.org/10.1016/0022-0531(74)90075-1). 10, 11
- V. V. Vazirani and M. Yannakakis. 2011. Market equilibrium under separable, piecewise-linear, concave utilities. *Journal of the ACM*, 58(3): 10. DOI: [10.1145/1970392.1970394](https://doi.org/10.1145/1970392.1970394). 9, 10, 117, 125
- M. Viderman. 2015. A combination of testability and decodability by tensor products. *Random Struct. Algorithms*, 46(3): 572–598. DOI: [10.1002/rsa.20498](https://doi.org/10.1002/rsa.20498). 245
- L. Walras. 1874. *Éléments d'économie politique pure, ou, Théorie de la richesse sociale*. L. Corbaz Lausanne. 8
- E. B. Yanovskaya. 1968. Equilibrium points in polymatrix games. *Litovskii Matematicheskii Sbornik*, 8: 381–384. DOI: [10.1155/2014/937070](https://doi.org/10.1155/2014/937070). 79
- H. P. Young. 2004. *Strategic learning and its limits*. Oxford University Press. 35
- H. P. Young. 2009. Learning by trial and error. *Games and economic behavior*, 65(2): 626–643. DOI: [10.1016/j.geb.2008.02.011](https://doi.org/10.1016/j.geb.2008.02.011). 35
- D. Zuckerman. 2007. Linear degree extractors and the inapproximability of Max Clique and chromatic number. *Theory of Computing*, 3(1): 103–128. DOI: [10.4086/toc.2007.v003a006](https://doi.org/10.4086/toc.2007.v003a006). 14, 155

Index

- 2-player games, correlated equilibria in, 60
- 2ϵ -GCIRCUIT. *See* Generalized circuit
- 2CSP problem
 - birthday repetition, 151–154
 - densest k -subgraph, 158–160
 - partitions, 28
 - PCP theorems, 22–23
 - signaling, 215–217, 219–220
 - violations, 170–176
- 3-COLOR constraints in VC Dimension, 189
- 3SAT
 - community detection, 178, 180
 - exponential time hypothesis, 22–24
 - VC Dimension, 191
- (α, β) -communities, 15, 177, 183
- ℓ_∞ Brouwer, 61–62
- ℓ_∞ -norm geometry, 62
- ϵ^4 -ANE, 49–51
- ϵ -approximate Bayesian Nash equilibrium, 111
- ϵ -approximate Nash equilibrium
 - birthday repetition, 153–154
 - defined, 20
- ϵ GCIRCUIT. *See* Generalized circuit
- ϵ -Nash in birthday repetition, 150–151
- $(\bar{\epsilon}, \bar{\epsilon})$ -weak approximate equilibrium (A, B), 57–58
- (ϵ, ϵ) -Weak ANE, 43–44
- (ϵ, δ) -WeakNash, 20
- (ϵ, δ) -WeakNash to $\Theta(\epsilon \cdot \delta)$ -ANE in bimatrix games, 270–271
- ϵ -well-supported Nash Equilibrium
 - defined, 20
 - communication complexity of approximate Nash equilibrium, 42–44
 - many-player games, 102–104
 - non-monotone markets, 120
- (ϵ, δ) -well-supported WeakNash, 20
- $((\sqrt{\epsilon} + \epsilon, \delta))$ -well-supported WeakNash to (ϵ, δ) -WeakNash, 269–270
- $\sqrt{\epsilon}$ -WSNE to $\Theta(\epsilon)$ -ANE, 106–107
- λ -biased sets, 27
- A-CEEI. *See* Approximate CEEI
- Adaptations for constant factor inapproximability in market equilibrium, 117–119
- Aggregate vectors in 2-player approximate Nash equilibrium, 252
- Althofer gadget, 219, 231, 270–271
- AM-GM inequality (Arithmetic-Mean Geometric-Mean inequality), 52
- Angular interpolation in ℓ_∞ -norm geometry, 62
- Approximate CEEI (A-CEEI), 132–133
 - A-CEEI \in PPAD
 - course allocation problem, 139–145
 - general position perturbation, 140–142
 - preliminaries, 139–140
 - A-CEEI problem
 - PPAD completeness, 10–12, 131
 - PPAD-hardness, 133–139

- Approximate gates in generalized circuit problem, 85
- Approximate Nash equilibrium, 4
 - 2-player approximate. *See* 2-player approximate Nash equilibrium
 - communication complexity, 33–44
 - intractability, 17–18
- Approximation algorithms, 224
- Approximation protocols, 38
- Approximations in densest k-subgraph, 162–163
- Arithmetic-Mean Geometric-Mean inequality (AM-GM inequality), 52
- Arithmetization in 2-player approximate Nash equilibrium, 239–241
- Arrow-Debreu model, 9
- Assignment elements
 - Littlestone’s Dimension, 201–203, 205–206
 - VC Dimension, 193–195
- Auxiliary goods in non-monotone markets, 122
- Auxiliary players in many-player games, 102
- Auxiliary subgames in birthday repetition, 152
- Auxiliary vertices in hardness of counting communities, 185–186
- Averaging gadgets in generalized circuit problem, 85–86

- Balanced subsets in hardness of detecting communities, 184
- Bayesian Nash equilibrium, 111–113
- Bimatrix games in 2-player approximate Nash equilibrium, 230–231
- Binary actions
 - communication complexity of approximate Nash equilibrium, 43–44
 - succinct games, 108–110
- Binary encoding in approximate Nash equilibrium, 50–51, 57
- Binary entropy function, 26
- Bipartite polymatrix game, 112

- Birthday paradox, 13
- Birthday repetition, 149
 - ϵ -Nash, 150–151
 - first gap, 151–154
 - framework, 13
 - gap amplification, 154
- Black-box oracle model, 6
- Boundary conditions
 - 2-player approximate Nash equilibrium, 233
 - END-OF-A-LINE, 21–22
- Boundary points in generalized circuit problem, 94
- Bounding in near-optimal signaling, 219–220
- BPP simulation theorem, 30
- $\text{BPP}_{0.3}^{\text{cc}}$ (SIMULATION-OF-THE-LINE) corollary, 46
- Brittle comparators, 84–85
- Brouwer line segments
 - embedding lines, 47
 - Euclidean Brouwer, 68
- Brouwer vertices, 68
- Brouwer with ℓ_∞ , 61–62
- Brouwer’s fixed point theorem, 61
 - Brouwer with ℓ_∞ , 61–62
 - continuous mapping, 63–67
 - Euclidean Brouwer, 68–75
 - PPAD, 7–8
- Budget-constraint hyperplane in course allocation problem, 140

- Candidate reductions in VC Dimension, 191–193
- Caratheodory Theorem
 - description, 12
 - signaling, 215
- Cartesian interpolation in ℓ_∞ norm geometry, 62
- Cauchy-Schwartz for densest k-subgraph, 165
- CC(BROUWER) informal corollary, 41
- CC(SIMULATION END-OF-A-LINE) informal corollary, 40

- CEEI (Competitive Equilibrium from Equal Incomes)
 - approximate, 132–133
 - course match, 10–12
- Chain rule, 26
- Circuit problems
 - generalized. *See* Generalized circuit problem
 - hardness, 8
- Clearing errors in course allocation problem, 132, 144–145
- Coloring graphs in 2-player approximate Nash equilibrium, 238–239
- Communication complexity of approximate Nash equilibrium, 33–34
 - proof overview, 39–44
 - related literature, 38–39
 - results, 34–35
 - techniques, 36–37
 - uncoupled dynamics, 35–36
- Communicationally hard Discrete End-of-Any-Line problem, 45
- Community detection, 177–178
 - alternative approaches, 179
 - hardness of counting communities, 182–183
 - hardness of detecting communities, 183–186
 - proofs, 179–182
 - quasi-polynomial time, 14–15
 - related literature, 178–179
- Competitive Equilibrium from Equal Incomes (CEEI)
 - approximate, 132–133
 - course match, 10–12
- Completeness
 - 2-player approximate Nash equilibrium, 237, 258
 - birthday repetition, 150–153
 - community detection, 178, 180–181
 - densest k -subgraph, 156–157, 160
 - hardness of counting communities, 182–183
 - hardness of detecting communities, 183, 185
 - Littlestone’s Dimension, 200–202, 204
 - near-optimal signaling, 218
 - PCP Theorem, 23
 - VC Dimension, 191–192, 195–196
- Concavity of binary entropy, 26
- Concentration lemma, 27
- Conditional entropy, 26
- Conditional mutual information, 26
- Conditional proofs of hardness, 6
- Consistency violations in densest k -subgraph, 164–170
- Constant factor inapproximability in market equilibrium, 117–119
- Constant hardness in market equilibrium, 119
- Constraint Satisfaction Problem, 22–23
- Continuous function, embedding as, 40–41
- Continuous mapping in Brouwer’s fixed point theorem, 63–67
- Coordinates in 2-player approximate Nash equilibrium, 250–251
- Corner points in generalized circuit problem, 94, 97–99
- Correlated equilibria
 - 2-player games, 60
 - communication complexity of approximate Nash equilibrium, 38–39
- Counters in 2-player approximate Nash equilibrium, 227, 236
- Counting communities, hardness of, 182–183
- Course allocation problem, 131–133
 - A-CEEI, 10–12, 133–139
 - A-CEEI \in PPAD, 139–145
 - course-size amplification, 135–136
 - finite precision, 143
 - fixed points, 142–145
 - gate gadgets, 136–139
- Course-size amplification in course allocation problem, 135–136
- Critical bits in 2-player approximate Nash equilibrium, 233, 238, 259–261

- Data processing inequality, 26
- Decision problem in communication
 - complexity of approximate Nash equilibrium, 37
- Decision tree complexity, 30
- Decoding
 - 2-player approximate Nash equilibrium, 228, 237–238, 242–244, 259–261
 - action profiles, 44
 - community detection, 186
- Default displacement
 - 2-player approximate Nash equilibrium, 261
 - continuous mapping in Brouwer’s fixed point theorem, 64
 - Euclidean Brouwer, 70, 73–74
- Densest k -subgraph, 155–157
- 2CSP violations, 170–176
- consistency violations, 164–170
- construction, 159–160
- entropy argument, 160–164
- PLANTED CLIQUE problem, 157–158
- quasi-polynomial time, 13–14
- soundness, 160–176
- techniques, 158–159
- Depth- d mistake trees
 - description, 25
 - Littlestone’s Dimension, 204, 206–208
- Derandomized partitions, 29
- Detecting communities, hardness of, 183–186
- DIFF gadget, 138
- Discrete embedding of graphs in Euclidean space, 68
- Discrete End-of-Any-Line problem, 45
- Displacement
 - 2-player approximate Nash equilibrium, 261
 - continuous mapping in Brouwer’s fixed point theorem, 64–67
 - Euclidean Brouwer, 70, 73–74
 - generalized circuit problem, 95–99
- Distance in 2-player approximate Nash equilibrium, 247
- Divide subroutine, 90–91
- Doubly exponentially small probabilities, 34n
- Embedded edges
 - Brouwer line segments, 47
 - Euclidean Brouwer, 68
- Embedded vertices
 - Euclidean Brouwer, 68
 - as local Lipschitz function, 46
- Embedding END-OF-A-LINE graph as paths in $\{0,1\}^{2n+1}$, 62
- Embedding lines
 - as continuous function, 40–41
 - as local Lipschitz function, 46–47
- END-OF-A-LINE problem
 - 2-player approximate Nash equilibrium, 225, 232–236
 - Brouwer’s fixed point theorem, 62
 - embedding lines as local Lipschitz function, 47
 - hardness, 6–7
 - PPAD, 20–22
 - query complexity, 40
- End subcubes in generalized circuit problem, 98
- Entrance facets in generalized circuit problem, 96
- Entropy
 - defined, 26
 - densest k -subgraph, 160–164, 168, 171
- Equiangle sampling in generalized circuit problem, 85, 87, 93–94
- Error correction in 2-player approximate Nash equilibrium, 242–243, 259–260
 - code, 225–227, 236
 - on critical bits, 238
 - on samples, 237
- ETH (Exponential Time Hypothesis), 5
 - densest k -subgraph, 156
 - quasi-polynomial time, 13, 22
 - signaling, 213–216
- #ETH hypotheses, 22
- Euclidean Brouwer, 68
 - 2-player approximate Nash equilibrium, 225–226

- discrete embedding of graphs in
 - Euclidean space, 68
- function f , 68–74
- locally computing, 74–75
- Exchange edges in 2-player approximate Nash equilibrium, 238
- Exit facets in generalized circuit problem, 96
- Expected clearing error in course allocation problem, 145
- Exponential Time Hypothesis (ETH), 5
 - densest k -subgraph, 156
 - quasi-polynomial time, 13, 22
 - signaling, 213–216
- Extended Shuffle-F-Exchange graphs, 238–239
- Extrapolation in 2-player approximate Nash equilibrium, 247–248, 260–261
- Fan-out 2, generalized circuit problem with, 99–100
- Far-from-uniform distributions, 29
- FGLSS in community detection, 180
- Finding fixed points in communication complexity, 39
- Finite precision in course allocation problem, 143
- First gap in birthday repetition, 151–154
- Fixed points
 - Brouwer's. *See* Brouwer's fixed point theorem
 - communication complexity, 39
 - course allocation problem, 142–145
- FPTAS (Fully Polynomial Time Approximation Scheme)
 - signaling, 213
 - ZERO-SUM SIGNALING problem, 17
- Frames in displacement function, 64
- Free games, 149
- Fully Polynomial Time Approximation Scheme (FPTAS)
 - signaling, 213
 - ZERO-SUM SIGNALING problem, 17
- Function f in Euclidean Brouwer, 68–74
- Gadgets
 - 2-player approximate Nash equilibrium, 226, 252–255, 271
 - course allocation problem, 134–139
 - generalized circuit problem, 84–86, 95
 - many-player games, 102–105
 - market equilibrium, 117
 - near-optimal signaling, 219
 - non-monotone markets, 121–122, 125–127
- Game size in birthday repetition, 152
- Game theory in Nash equilibrium, 3
- Games on graphs with constant mixing time in non-monotone markets, 120–121
- Gaps
 - 2-player approximate Nash equilibrium, 225–226
 - birthday repetition, 154
- Gates
 - course allocation problem, 136–139
 - generalized circuit problem, 84
- Generalized circuit problem, 83–84
 - 2ϵ -GCIRCUIT, 102
 - ϵ GCIRCUIT, 86–87
 - in course allocation problem, 135–136
 - brittle comparators, 84
 - course allocation problem, 135
 - with fan-out 2, 99–100
 - proof overview, 85–87
 - results, 84–85
 - subroutines, 88–99
- Giffen goods, 9
- Global dependency in 2-player approximate Nash equilibrium, 226
- Good prefixes in densest k -subgraph, 165–166
- Good Sample guarantee, 257, 260, 263–265
- Good samples in 2-player approximate Nash equilibrium, 237
- Good variables in densest k -subgraph, 164
- Goods in non-monotone markets, 124
- Graphical games
 - defined, 79
 - polymatrix, 101–107

- Graphs in 2-player approximate Nash equilibrium, 238–239
- Greedy partitions, 28–29
- HALF gadget in course allocation problem, 137–138
- Hamming distance in 2-player approximate Nash equilibrium, 247
- Hardness of approximation
 - 2-player approximate Nash equilibrium, 224
 - Littlestone’s Dimension, 188
 - VC Dimension, 188
- Hardness of counting communities in community detection, 182–183
- Hardness of detecting communities in community detection, 183–186
- Hardness of well-supported Nash equilibrium in many-player games, 102–105
- Holographic proofs, 227–229, 236–248, 255–261
- Home subcubes in displacement function, 64
- Hyperplane testers, 245
- i*-assignment elements
 - Littlestone’s Dimension, 203
 - VC Dimension, 193
- Imitation gadgets in 2-player approximate Nash equilibrium, 252–255
- Imitation game in communication
 - complexity of approximate nash equilibrium, 37
- Inapproximability
 - community detection, 178
 - densest *k*-subgraph, 156–157
 - generalized circuit problem, 85
 - Littlestone’s Dimension, 200–211
 - non-monotone markets, 117–129
 - VC Dimension, 191–199
- Incoming neighbors in many-player games, 106–107
- Information theory, 25–26
- Input players in many-player games, 102
- Interior points in generalized circuit problem, 94
- Interpolate subroutine, 93
- Interpolation
 - ℓ_∞ -norm geometry, 62
 - continuous mapping in Brouwer’s fixed point theorem, 64–67
- Intractability
 - approximate Nash equilibrium, 17–18
 - Brouwer’s fixed point theorem, 7–8
 - market equilibrium, 8–10
 - overview, 3–5
 - PPAD, 5–12
 - quasi-polynomial time and birthday paradox, 12–17
- Iw-type players, 54
- k*-CLIQUE problem, 13–14
- Kakutani’s fixed point theorem, 11
- Kullback Leiber Divergence
 - defined, 27
 - densest *k*-subgraph, 163
- LABEL COVER problem
 - birthday repetition, 152
 - community detection, 180–181
 - hardness of counting communities, 182
 - hardness of detecting communities, 184–186
 - Littlestone’s Dimension, 200, 202–203
 - PCP theorems, 23–24
 - VC Dimension, 191–192, 194–195
- Labeling entropy bound lemma, 168–169
- Layers in 2-player approximate Nash equilibrium, 238
- Learning theory, 24–25
- Lemke-Howson algorithm, 223
- Linear error-correcting code in 2-player approximate Nash equilibrium, 227, 236
- Lipschitz function
 - 2-player approximate Nash equilibrium, 225, 264
 - embedding lines as, 46–47

- List decoding in hardness of counting communities, 186
- Littlestone's Dimension
 - defined, 24–25
 - discussion, 188–189
 - inapproximability, 200–211
 - quasi-polynomial algorithm, 211–212
 - quasi-polynomial time, 16
 - related literature, 190–191
 - soundness, 204–211
 - techniques, 189–190
- Local ℓ_∞ -norm polar coordinate system, 62
- Local access in 2-player approximate Nash equilibrium, 236–237
- Local conditions in 2-player approximate Nash equilibrium, 234
- Local decoding in 2-player approximate Nash equilibrium, 228, 242–243
- LOCAL END-OF-A-LINE problem, 266
 - 2-player approximate Nash equilibrium, 225–227, 232–236, 238–239, 248, 253, 262
 - defined, 22
- Local Lipschitz function, embedding lines as, 46–47
- Local proof-construction in 2-player approximate Nash equilibrium, 228, 244
- Local robust testing in 2-player approximate Nash equilibrium, 245–246
- Locality
 - 2-player approximate Nash equilibrium, 225–226
 - Euclidean Brouwer, 75
- LOGNP class in VC Dimension, 190
- Lower bounds in communication complexity of approximate Nash equilibrium, 33–34
- Main game in 2-player approximate Nash equilibrium, 270
- Main goods in non-monotone markets, 121
- Main subgame in birthday repetition, 152
- Many-player games, 101
 - $\sqrt{\epsilon}$ -WSNE to $\Theta(\epsilon)$ -ANE, 106–107
 - graphical, polymatrix games, 101–107
 - hardness of well-supported Nash equilibrium, 102–105
 - succinct games, 107–110
- Market equilibrium, 115–116
 - adaptations for constant factor inapproximability, 117–119
 - high-level structure of proof, 116–117
 - non-monotone markets, 116, 119–129
 - PPAD, 8–10
- Markov's inequality for hardness of counting communities, 186
- Max-norm polar coordinates in generalized circuit problem, 96–97
- Max subroutine, 91–92
- MEMBERSHIP END-OF-A-LINE problem
 - defined, 21
 - Euclidean Brouwer, 68, 75
 - variants, 22
- Minmax Theorem, 45
- Mistake bound, defined, 25
- Mistake trees
 - defined, 24
 - Littlestone's Dimension, 204, 206–208
- Mode assignments in densest k-subgraph, 175
- Moderately computable price vectors, 115
- Monotone markets, 115
- Moshkovitz-Raz PCP theorem
 - defined, 24
 - VC Dimension, 191
- Multiplicative hardness of approximation in 2-player approximate Nash equilibrium, 224
- Multiply subroutine, 89–90
- Mutual information, conditional, 26
- n -player game in communication complexity of approximate Nash equilibrium, 42–44, 53–59
- Nash equilibrium, 3–4
 - 2-player approximate. *See* 2-player approximate Nash equilibrium
 - Bayesian, 111–113
 - intractability, 17–18

- Nash equilibrium (*continued*)
 - relaxations, 19–20
- Near-optimal signaling
 - completeness, 218
 - construction, 215–218
 - soundness, 218–220
- Neighborhoods in densest k -subgraph, 162
- Neighbors 95–96, 106–107
- Non-deterministic communication
 - protocols, 36
- Non-monotone markets, 9–10
 - construction, 121–123
 - defined, 115
 - games on graphs with constant mixing time, 120–121
 - hardness, 116
 - from market equilibrium to Nash equilibrium, 128–129
 - normalized polymatrix games, 119
 - structure of market equilibrium, 123–128
- Non-repetitive sets
 - Littlestone’s Dimension, 206–210
 - VC Dimension, 193–194, 196–199
- Normalized 2-norm, 19
- Normalized games in non-monotone markets
 - inapproximability, 120
 - polymatrix, 119
- NOT gate (G_-) gadget in course allocation problem, 134
- Notation
 - 2-player approximate Nash equilibrium, 250–252
 - overview, 19
 - PCP Theorem, 23
- NP-hardness
 - Nash equilibrium, 4–5
 - signaling, 214
 - VC Dimension and Littlestone’s Dimension, 188
- $O(1)$ -action games, correlated equilibria in, 60
- $O(1)$ -Lipschitzness of f , 59
- Output players in many-player games, 102
- PAC learning in VC Dimension, 190
- Partial assignments in LABEL COVER problem, 23
- Partial norms in 2-player approximate Nash equilibrium, 251
- Partial VC Dimension, 190
- Partial vectors in 2-player approximate Nash equilibrium, 251
- Partially adaptive samples in 2-player approximate Nash equilibrium, 237
- Partitions, 28
 - derandomized, 29
 - greedy, 28–29
- Payoffs in signaling, 217–218
- PCP. *See* Probabilistically checkable proof (PCP)
- Pictures in displacement function, 64
- Planted Clique Conjecture, 157
- PLANTED CLIQUE problem
 - 2-player approximate Nash equilibrium, 224
 - community detection, 178
 - densest k -subgraph, 157–158
 - quasi-polynomial time, 17
 - signaling, 213–214
- Players in 2-player approximate Nash equilibrium, 249–250
- PoA (Price of Anarchy) game, 39
- Polymatrix games
 - 2-player approximate Nash equilibrium, 230–231
 - defined, 79
- Polymatrix to bimatrix, 269–272
- Polymatrix WeakNash, 248–269
- Polynomial lower bounds in communication complexity of approximate Nash equilibrium, 34
- Polynomial time approximation schemes (PTAS)
 - 2-player approximate Nash equilibrium, 224
 - Nash equilibrium, 17–18

- signaling, 213
- Polynomial upper bound in communication
 - complexity of approximate Nash equilibrium, 38
- PPAD (Polynomial Parity Argument in Directed graphs)
 - A-CEEI, 10–12
 - Brouwer’s fixed point theorem, 7–8
 - END-OF-A-LINE, 20–22
 - exponential time hypothesis, 22
 - market equilibrium, 8–10
 - overview, 5–7
- PPAD-completeness, 223–224
- PPAD-hardness
 - 2-player approximate Nash equilibrium, 225–226
 - A-CEEI, 133–139
 - of approximation, 79–81
- Predecessors of vertices in communication
 - complexity of approximate Nash equilibrium, 44
- Prediction mistakes, 25
- Prefixes in densest k-subgraphs
 - consistency violations, 165–167
 - degree, 161
 - density, 162
 - graphs, 161–162
- Price of Anarchy (PoA) game, 39
- Price-regulating gadgets, 122
- Probabilistically checkable proof (PCP)
 - 2-player approximate Nash equilibrium, 241–242
 - 2CSP, 22–23
 - holographic proofs, 227–229
 - LABEL COVER problem, 23–24
 - PPAD conjecture, 231–232
- Pseudorandomness theorem, 27
- PSPACE-completeness, 190
- PTAS (polynomial time approximation schemes)
 - 2-player approximate Nash equilibrium, 224
 - Nash equilibrium, 17–18
 - signaling, 213
- Pure Nash equilibrium, 38
- Quality of equilibria in Nash equilibrium, 39
- Quasi-polynomial time
 - community detection, 14–15
 - densest k-subgraph, 13–14
 - Littlestone’s Dimension, 211–212
 - Nash equilibrium, 5
 - overview, 12–13
 - signaling, 16–17
 - VC and Littlestone’s dimensions, 15–16
- Queries
 - 2-player approximate Nash equilibrium, 246–247
 - approximate Nash equilibrium, 37
 - END-OF-THE-LINE, 40
 - literature, 38–39
 - subspaces, 236
- Query complexity, 38
- QUERY END-OF-A-LINE lemma, 40
- QUERY END-OF-THE-LINE problem, 44–46
- Radial interpolation in continuous mapping
 - in Brouwer’s fixed point theorem, 64
- Random-bit complexity in 2-player approximate Nash equilibrium, 228
- Randomized query complexity lemma, 45
- Randomized query complexity lower bound, 44–45
- Randomness in 2-player approximate Nash equilibrium, 237, 246
- Rate of convergence in uncoupled dynamics, 35–36
- Real2Unary algorithm, 100
- Realized numbers, 54–55
- Realized points, 54–55
- Realized value in succinct games, 108–109
- Reject codes in 2-player approximate Nash equilibrium, 228
- Relaxations
 - k-CLIQUE problem, 14

- Relaxations (*continued*)
 - Nash equilibrium, 19–20
- Robust Soundness guarantee, 258
- Robustness in 2-player approximate Nash equilibrium
 - completeness, 237
 - decoding, 238
 - soundness, 229, 237
 - testing, 245–246
- Rounding in 2-player approximate Nash equilibrium, 250
- Sampling Lemma, 28
- Satisfied constraints in LABEL COVER problem, 23
- Scarf's algorithm, 8
- Shannon entropy, 26
- Shattered subsets, 24
- Shuffle edges in 2-player approximate Nash equilibrium, 238
- Signaling
 - bounding, 219–220
 - completeness, 218
 - near-optimal, 215–220
 - quasi-polynomial time, 16–17
 - techniques, 214–215
 - ZERO-SUM SIGNALING problem, 213–214
- Simplicity, measuring, 187
- SIMULATION END-OF-THE-LINE
 - communication complexity of approximate Nash equilibrium, 53, 59
 - defined, 45–46
 - reduction, 48–50
- Simulation theorems
 - communication complexity of approximate Nash equilibrium, 37, 40
 - overview, 30
- Single-minded traders graphs, 123–124
- Slice layers in continuous mapping in Brouwer's fixed point theorem, 63–64
- Social welfare in Nash equilibrium, 39
- Soundness
 - 2-player approximate Nash equilibrium, 237, 256–258
 - birthday repetition, 151, 153–154
 - community detection, 178, 181–182
 - densest k-subgraph, 156, 160–176
 - entropy argument, 160–164
 - hardness of counting communities, 183, 185–186
 - hardness of detecting communities, 183
 - Littlestone's Dimension, 200, 204–211
 - near-optimal signaling, 218–220
 - PCP Theorem, 23
 - robust, 229
 - VC Dimension, 191, 195–199
- Start subcubes in generalized circuit problem, 98
- States of nature, signaling, 217
- Static edges in 2-player approximate Nash equilibrium, 238
- Stochastic block model, 179
- Strategies
 - 2-player approximate Nash equilibrium, 249–250
 - communication complexity of approximate Nash equilibrium, 48
- Strong ties inside communities, 15, 177
- Subcubes
 - displacement function, 64
 - generalized circuit problem, 98
- Subgames
 - birthday repetition, 152
 - many-player games, 102
- Subspaces in 2-player approximate Nash equilibrium, 244, 249–250
- Subtrees in Littlestone's Dimension, 208–209
- Successors of vertices in communication complexity of approximate Nash equilibrium, 44
- Succinct games
 - defined, 79
 - many-player games, 107–110

- Summing displacement vectors in generalized circuit problem, 99
- Tâtonnement process, 116
- Tensor product in 2-player approximate Nash equilibrium, 245
- Test-selection elements
 - Littlestone's Dimension, 200, 203, 205–206
 - VC Dimension, 193–194
- Testing 2-player approximate Nash equilibrium, 245–247
- TFNP in 2-player approximate Nash equilibrium, 225
- Threshold projections in Littlestone's Dimension, 202–203
- Tolerant verifier in 2-player approximate Nash equilibrium, 228
- Totality
 - communication complexity of approximate Nash equilibrium, 36
 - Nash equilibrium, 4
- Tractable special cases, 101
- Traders in non-monotone markets, 124
- Truncation in Euclidean Brouwer, 69
- Tubes in displacement function, 64–66
- Two-player game in communication complexity of approximate Nash equilibrium, 41–42, 47–48
- Type players, 54
- Typical variables in densest k -subgraph, 170–172
- Unary2Real algorithm, 100
- Unconditional lower bounds in communication complexity of approximate Nash equilibrium, 33–34
- Uncoupled dynamics in communication complexity of approximate Nash equilibrium, 34–36
- University course matches, 10–12
- Untruncated displacement function in Euclidean Brouwer, 69
- Utilities in communication complexity of approximate Nash equilibrium, 48–49, 54–55
- v -type players, 54–55
- VC Dimension, 187–188
 - defined, 24
 - discussion, 188–189
 - inapproximability, 191–199
 - quasi-polynomial time, 16
 - reductions, 191–195
 - related literature, 190–191
 - soundness, 195–199
 - techniques, 189
- Verifiers
 - 2-player approximate Nash equilibrium, 247
 - PCP Theorem, 231
- Vertices in Euclidean Brouwer, 68
- w -type players, 54
- Walras's Law, 126
- Weak approximate Nash equilibria, 53
- Weak ties to nodes outside communities, 15, 177
- WeakNash, 20
- WELL-SUPPORTED-WEAKNASH, 248
- x -type players, 54–55
- y -type players, 54–55
- Yao's Minmax Theorem, 45
- zero-sum game in Littlestone's Dimension, 188
- ZERO-SUM SIGNALING problem
 - quasi-polynomial time, 16–17
 - signaling, 213–214

Author Biography

Aviad Rubinstein



Aviad Rubinstein is an Assistant Professor of Computer Science at Stanford University. Before coming to Stanford, he received his Ph.D. from the University of California, Berkeley, and spent one year as a Rabin Postdoctoral Fellow at Harvard University.