

**Complexity Theory, Game  
Theory, and Economics:  
The Barbados Lectures**

**Other titles in Foundations and Trends® in Theoretical Computer Science**

*Coding for Interactive Communication: A Survey*

Ran Gelles

ISBN: 978-1-68083-346-1

*Hashing, Load Balancing and Multiple Choice*

Udi Wieder

ISBN: 978-1-68083-282-2

*Scalable Algorithms for Data and Network Analysis*

Shang-Hua Teng

ISBN: 978-1-68083-130-6

*Communication Complexity (for Algorithm Designers)*

Tim Roughgarden

ISBN: 978-1-68083-114-6

*Quantum Proofs*

Thomas Vidick and John Watrous

ISBN: 978-1-68083-126-9

# Complexity Theory, Game Theory, and Economics: The Barbados Lectures

---

**Tim Roughgarden**  
Columbia University  
New York, NY, USA  
tr@cs.columbia.edu

**now**

the essence of knowledge

Boston — Delft

# Foundations and Trends<sup>®</sup> in Theoretical Computer Science

*Published, sold and distributed by:*

now Publishers Inc.  
PO Box 1024  
Hanover, MA 02339  
United States  
Tel. +1-781-985-4510  
[www.nowpublishers.com](http://www.nowpublishers.com)  
[sales@nowpublishers.com](mailto:sales@nowpublishers.com)

*Outside North America:*

now Publishers Inc.  
PO Box 179  
2600 AD Delft  
The Netherlands  
Tel. +31-6-51115274

The preferred citation for this publication is

T. Roughgarden. *Complexity Theory, Game Theory, and Economics: The Barbados Lectures*. Foundations and Trends<sup>®</sup> in Theoretical Computer Science, vol. 14, no. 3–4, pp. 222–407, 2020.

ISBN: 978-1-68083-655-4

© 2020 T. Roughgarden

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, mechanical, photocopying, recording or otherwise, without prior written permission of the publishers.

Photocopying. In the USA: This journal is registered at the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923. Authorization to photocopy items for internal or personal use, or the internal or personal use of specific clients, is granted by now Publishers Inc for users registered with the Copyright Clearance Center (CCC). The 'services' for users can be found on the internet at: [www.copyright.com](http://www.copyright.com)

For those organizations that have been granted a photocopy license, a separate system of payment has been arranged. Authorization does not extend to other kinds of copying, such as that for general distribution, for advertising or promotional purposes, for creating new collective works, or for resale. In the rest of the world: Permission to photocopy must be obtained from the copyright owner. Please apply to now Publishers Inc., PO Box 1024, Hanover, MA 02339, USA; Tel. +1 781 871 0245; [www.nowpublishers.com](http://www.nowpublishers.com); [sales@nowpublishers.com](mailto:sales@nowpublishers.com)

now Publishers Inc. has an exclusive license to publish this material worldwide. Permission to use this content must be obtained from the copyright license holder. Please apply to now Publishers, PO Box 179, 2600 AD Delft, The Netherlands, [www.nowpublishers.com](http://www.nowpublishers.com); e-mail: [sales@nowpublishers.com](mailto:sales@nowpublishers.com)

**Foundations and Trends<sup>®</sup> in Theoretical  
Computer Science**  
Volume 14, Issue 3–4, 2020  
**Editorial Board**

**Editor-in-Chief**

**Madhu Sudan**  
Harvard University  
United States

**Editors**

Bernard Chazelle  
*Princeton University*

Oded Goldreich  
*Weizmann Institute*

Shafi Goldwasser  
*Massachusetts Institute of Technology and Weizmann Institute*

Sanjeev Khanna  
*University of Pennsylvania*

Jon Kleinberg  
*Cornell University*

László Lovász  
*Eötvös Loránd University*

Christos Papadimitriou  
*University of California, Berkeley*

Peter Shor  
*Massachusetts Institute of Technology*

Eva Tardos  
*Cornell University*

Avi Wigderson  
*AIS, Princeton University*

## Editorial Scope

### Topics

Foundations and Trends® in Theoretical Computer Science publishes survey and tutorial articles in the following topics:

- Algorithmic game theory
- Computational algebra
- Computational aspects of combinatorics and graph theory
- Computational aspects of communication
- Computational biology
- Computational complexity
- Computational geometry
- Computational learning
- Computational Models and Complexity
- Computational Number Theory
- Cryptography and information security
- Data structures
- Database theory
- Design and analysis of algorithms
- Distributed computing
- Information retrieval
- Operations Research
- Parallel algorithms
- Quantum Computation
- Randomness in Computation

### Information for Librarians

Foundations and Trends® in Theoretical Computer Science, 2020, Volume 14, 4 issues. ISSN paper version 1551-305X. ISSN online version 1551-3068. Also available as a combined paper and online subscription.

## Contents

---

<b>Foreword</b>	<b>2</b>
<b>Overview</b>	<b>4</b>
<b>I Solar Lectures</b>	<b>7</b>
<b>1 Introduction, Wish List, and Two-Player Zero-Sum Games</b>	<b>8</b>
1.1 Nash Equilibria in Two-Player Zero-Sum Games . . . . .	8
1.2 Uncoupled Dynamics . . . . .	16
1.3 General Bimatrix Games . . . . .	26
1.4 Approximate Nash Equilibria in Bimatrix Games . . . . .	28
<b>2 Communication Complexity Lower Bound for Computing an Approximate Nash Equilibrium of a Bimatrix Game (Part I)</b>	<b>32</b>
2.1 Preamble . . . . .	32
2.2 Naive Approach: Reduction From DISJOINTNESS . . . . .	34
2.3 Finding Brouwer Fixed Points (The $\epsilon$ -BFP Problem) . . . . .	36
2.4 The End-of-the-Line (EoL) Problem . . . . .	39
2.5 Road Map for the Proof of Theorem 2.1 . . . . .	43
2.6 Step 1: Query Lower Bound for EoL . . . . .	45

2.7	Step 2: Communication Complexity Lower Bound for 2EoL via a Simulation Theorem . . . . .	46
<b>3</b>	<b>Communication Complexity Lower Bound for Computing an Approximate Nash Equilibrium of a Bimatrix Game (Part II)</b>	<b>51</b>
3.1	Step 3: $2\text{EoL} \leq \epsilon\text{-2BFP}$ . . . . .	51
3.2	Step 4: $\epsilon\text{-2BFP} \leq \epsilon\text{-NE}$ . . . . .	59
<b>4</b>	<b>TFNP, PPAD, &amp; All That</b>	<b>67</b>
4.1	Preamble . . . . .	68
4.2	TFNP and Its Subclasses . . . . .	69
4.3	PPAD and Its Complete Problems . . . . .	73
4.4	Are TFNP Problems Hard? . . . . .	78
<b>5</b>	<b>The Computational Complexity of Computing an Approximate Nash Equilibrium</b>	<b>84</b>
5.1	Introduction . . . . .	84
5.2	Proof of Theorem 5.1: An Impressionistic Treatment . . . . .	86
<b>II</b>	<b>Lunar Lectures</b>	<b>97</b>
<b>6</b>	<b>How Computer Science Has Influenced Real-World Auction Design. Case Study: The 2016–2017 FCC Incentive Auction</b>	<b>98</b>
6.1	Preamble . . . . .	98
6.2	Reverse Auction . . . . .	99
6.3	Forward Auction . . . . .	103
<b>7</b>	<b>Communication Barriers to Near-Optimal Equilibria</b>	<b>110</b>
7.1	Welfare Maximization in Combinatorial Auctions . . . . .	110
7.2	Communication Lower Bounds for Approximate Welfare Maximization . . . . .	111
7.3	Lower Bounds on the Price of Anarchy of Simple Auctions	116
7.4	An Open Question . . . . .	124
7.5	Appendix: Proof of Theorem 7.2 . . . . .	124



<b>8</b>	<b>Why Prices Need Algorithms</b>	<b>126</b>
8.1	Markets with Indivisible Items . . . . .	126
8.2	Complexity Separations Imply Non-Existence of Walrasian Equilibria . . . . .	131
8.3	Proof of Theorem 8.5 . . . . .	133
8.4	Beyond Walrasian Equilibria . . . . .	137
<b>9</b>	<b>The Borders of Border's Theorem</b>	<b>140</b>
9.1	Optimal Single-Item Auctions . . . . .	140
9.2	Border's Theorem . . . . .	145
9.3	Beyond Single-Item Auctions: A Complexity-Theoretic Barrier . . . . .	152
9.4	Appendix: A Combinatorial Proof of Border's Theorem . . .	157
<b>10</b>	<b>Tractable Relaxations of Nash Equilibria</b>	<b>159</b>
10.1	Preamble . . . . .	159
10.2	Uncoupled Dynamics Revisited . . . . .	160
10.3	Correlated and Coarse Correlated Equilibria . . . . .	162
10.4	Computing an Exact Correlated or Coarse Correlated Equilibrium . . . . .	165
10.5	The Price of Anarchy of Coarse Correlated Equilibria . . . .	168
	<b>References</b>	<b>172</b>

# Complexity Theory, Game Theory, and Economics: The Barbados Lectures

Tim Roughgarden

*Columbia University, New York, NY, USA; tr@cs.columbia.edu*

---

## ABSTRACT

The goal of this monograph is twofold: (i) to explain how complexity theory has helped illuminate several barriers in economics and game theory, and (ii) to illustrate how game-theoretic questions have led to new and interesting complexity theory, including several very recent breakthroughs.

---

## Foreword

---

This monograph is based on lecture notes from my mini-course “Complexity Theory, Game Theory, and Economics,” taught at the Bellairs Research Institute of McGill University, Holetown, Barbados, February 19–23, 2017, as the 29th McGill Invitational Workshop on Computational Complexity.

The goal of this monograph is twofold:

- (i) to explain how complexity theory has helped illuminate several barriers in economics and game theory; and
- (ii) to illustrate how game-theoretic questions have led to new and interesting complexity theory, including several very recent breakthroughs.

It consists of two five-lecture sequences: the *Solar Lectures*, focusing on the communication and computational complexity of computing equilibria; and the *Lunar Lectures*, focusing on applications of complexity theory in game theory and economics.<sup>1</sup> No background in game theory is assumed.

Thanks are due to many people: Denis Therien and Anil Ada for organizing the workshop and for inviting me to lecture; Omri Weinstein, for giving a guest lecture on simulation theorems in communication

---

<sup>1</sup>Cris Moore: “So when are the *stellar* lectures?”

complexity; Alex Russell, for coordinating the scribe notes; the scribes,<sup>2</sup> for putting together a terrific first draft; and all of the workshop attendees, for making the experience so unforgettable (if intense!). I also thank Yakov Babichenko, Mika Göös, Aviad Rubinstein, Eylon Yogev, and an anonymous reviewer for numerous helpful comments on earlier drafts of this monograph.

The writing of this monograph was supported in part by NSF award CCF-1524062, a Google Faculty Research Award, and a Guggenheim Fellowship. I would be very happy to receive any comments or corrections from readers.

Tim Roughgarden  
Bracciano, Italy  
December 2017  
(Revised December 2019)

---

<sup>2</sup>Anil Ada, Amey Bhangale, Shant Boodaghians, Sumegha Garg, Valentine Kabanets, Antonina Kolokolova, Michal Koucký, Cristopher Moore, Pavel Pudlák, Dana Randall, Jacobo Torán, Salil Vadhan, Joshua R. Wang, and Omri Weinstein.

## **Part II**

# **Lunar Lectures**

# 6

---

## How Computer Science Has Influenced Real-World Auction Design. Case Study: The 2016–2017 FCC Incentive Auction

---

### 6.1 Preamble

Computer science is changing the way auctions are designed and implemented. For over 20 years, the US and other countries have used *spectrum auctions* to sell licenses for wireless spectrum to the highest bidder. What's different this decade, and what necessitated a new auction design, is that in the US the juiciest parts of the spectrum for next-generation wireless applications are already accounted for, owned by over-the-air television broadcasters. This led Congress to authorize the FCC in the fall of 2012 to design a novel auction (the *FCC Incentive Auction*) that would repurpose spectrum—procuring licenses from television broadcasters (a relatively low-value activity) and selling them to parties that would put them to better use (e.g., telecommunication companies who want to roll out the next generation of wireless broadband services). Thus the FCC Incentive Auction is really a *double auction*, comprising two stages: a *reverse auction*, where the government buys back licenses for spectrum from their current owners; and then a *forward auction*, where the government sells the procured licenses to the highest bidder. Computer science techniques played a crucial role in the design of the new reverse auction. The main aspects of the forward auction

have been around a long time; here, theoretical computer science has contributed on the analysis side, and to understanding when and why such forward auctions work well. Sections 6.2 and 6.3 give more details on the reverse and forward parts of the auction, respectively.

The FCC Incentive Auction finished around the end of March 2017, and so the numbers are in. The government spent roughly 10 billion USD in the reverse part of the auction buying back licenses from television broadcasters, and earned roughly 20 billion USD of revenue in the forward auction. Most of the 10 billion USD profit was used to reduce the US debt!<sup>1</sup>

## 6.2 Reverse Auction

### 6.2.1 Descending Clock Auctions

The reverse auction is the part of the FCC Incentive Auction that was totally new, and where computer science techniques played a crucial role in the design. The auction format, proposed by Milgrom and Segal [113], is what's called a *descending clock auction*. By design, the auction is very simple from the perspective of any one participant. The auction is iterative, and operates in rounds. In each round of the auction, each remaining broadcaster is asked a question of the form: "Would you or would you not be willing to sell your license for (say) 1 million dollars?" The broadcaster is allowed to say "no," with the consequence of getting kicked out of the auction forevermore (the station will keep its license and remain on the air, and will receive no compensation from the government). The broadcaster is also allowed to say "yes" and accept the buyout offer. In the latter case, the government will not necessarily buy the license for 1 million dollars—in the next round, the broadcaster might get asked the same question, with a lower buyout price (e.g., 950,000 USD). If a broadcaster is still in the auction when it ends (more on how it ends in a second), then the government does indeed buy their license, at the most recent (and hence lowest) buyout

---

<sup>1</sup>This was the plan all along, which is probably one of the reasons the bill didn't have trouble passing a notoriously partisan Congress. Another reason might be the veto-proof title of the bill: "The Middle Class Tax Relief and Job Creation Act."

offer. Thus all a broadcaster has to do is answer a sequence of “yes/no” questions for some decreasing sequence of buyout offers. The obvious strategy for a broadcaster is to formulate the lowest acceptable offer for their license, and to drop out of the auction once the buyout price drops below this threshold.


The auction begins with very high buyout offers, so that every broadcaster would be ecstatic to sell their license at the initial price. Intuitively, the auction then tries to reduce the buyout prices as much as possible, subject to clearing a target amount of spectrum. Spectrum is divided into channels which are blocks of 6 MHz each. For example, one could target broadcasters assigned to channels 38–51, and insist on clearing 10 out of these 14 channels (60 MHz overall).<sup>2</sup> By “clearing a channel,” we mean clearing it *nationwide*. Of course, in the descending clock auction, bidders will drop out in an uncoordinated way—perhaps the first station to drop out is channel 51 in Arizona, then channel 41 in western Massachusetts, and so on. To clear several channels nationwide without buying out essentially everybody, it was essential for the government to use its power to *reassign* the channels of the stations that remain on the air. Thus while a station that drops out of the auction is guaranteed to retain its license, it is not guaranteed to retain its channel—a station broadcasting on channel 51 before the auction might be forced to broadcast on channel 41 after the auction.

The upshot is that the auction maintains the invariant that the stations that have dropped out of the auction (and hence remain on the air) can be assigned channels so that at most a target number of channels are used (in our example, 4 channels). This is called the *repacking problem*. Naturally, two stations with overlapping broadcasting regions cannot be assigned the same channel (otherwise they would interfere with each other). See Figure 6.1.

---

<sup>2</sup>The FCC Incentive Auction wound up clearing 84 MHz of spectrum (14 channels).





**Figure 6.1:** Different TV stations with overlapping broadcasting areas must be assigned different channels (indicated by shades of gray). Checking whether or not a given subset of stations can be assigned to a given number of channels without interference is an NP-hard problem.

### 6.2.2 Solving the Repacking Problem

Any properly trained computer scientist will recognize the repacking problem as the NP-complete graph coloring problem in disguise.<sup>3</sup> For the proposed auction format to be practically viable, it must quickly solve the repacking problem. Actually, make that thousands of repacking problems every round of the auction!<sup>4</sup>

The responsibility of quickly solving repacking problems fell to a team led by Kevin Leyton-Brown (see [56, 102]). The FCC gave the team a budget of one minute per repacking problem, ideally with most instances solved within one second. The team's approach was to build on

---

<sup>3</sup>The actual repacking problem was more complicated—overlapping stations cannot even be assigned adjacent channels, and there are idiosyncratic constraints at the borders with Canada and Mexico. See Leyton-Brown *et al.* [102] for more details. But the essence of the repacking problem really is graph coloring.

<sup>4</sup>Before the auction makes a lower offer to some remaining broadcaster in the auction, it needs to check that it would be OK for the broadcaster to decline and drop out of the auction. If a station's dropping out would render the repacking problem infeasible, then that station's buyout price remains frozen until the end of the auction.

state-of-the-art solvers for the satisfiability (SAT) problem. As you can imagine, it's straightforward to translate an instance of the repacking problem into a SAT formula (even with the idiosyncratic constraints).<sup>5</sup> Off-the-shelf SAT solvers did pretty well, but still timed out on too many representative instances.<sup>6</sup> Leyton-Brown's team added several new innovations, including taking advantage of problem structure specific to the application and implementing a number of caching techniques (reusing work done solving previous instances to quickly solve closely related new instances). In the end, they were able to solve more than 99% of the relevant repacking problems in under a minute.

Hopefully the high-level point is clear:

without cutting-edge techniques for solving *NP*-complete problems, *the FCC would have had to use a different auction format.*

### 6.2.3 Reverse Greedy Algorithms

One final twist: the novel reverse auction format motivates some basic algorithmic questions (and thus ideas flow from computer science to auction theory and back). We can think of the auction as an algorithm, a heuristic that tries to maximize the value of the stations that remain on the air, subject to clearing the target amount of spectrum. Milgrom and Segal [113] prove that, ranging over all ways of implementing the auction (i.e., of choosing the sequences of descending prices), the corresponding algorithms are exactly the *reverse greedy algorithms*.<sup>7</sup> This result gives

---

<sup>5</sup>A typical representative instance would have thousands of variables and tens of thousands of constraints.

<sup>6</sup>Every time the repacking algorithm fails to find a repacking when one exists, money is left on the table—the auction has to conservatively leave the current station's buyout offer frozen, even though it could have safely lowered it.

<sup>7</sup>For example, Kruskal's algorithm for the minimum spanning tree problem (start with the empty set, go through the edges of the graph from cheapest to most expensive, adding an edge as long as it doesn't create a cycle) is a standard (forward) greedy algorithm. The reverse version is: start with the entire edge set, go through the edges in reverse sorted order, and remove an edge whenever it doesn't disconnect the graph. For the minimum spanning tree problem (and more generally for finding the minimum-weight basis of a matroid), the reverse greedy algorithm is just as optimal as the forward one. In general (and even for e.g. bipartite matching), the reverse version of a good forward greedy algorithm can be bad [50].

the first extrinsic reason to study the power and limitations of reverse greedy algorithms, a research direction explored by Dütting *et al.* [50] and Gkatzelis *et al.* [65].

### 6.3 Forward Auction

Computer science did not have an opportunity to influence the design of the forward auction used in the FCC Incentive Auction, which resembles the formats used over the past 20+ years. Still, the theoretical computer science toolbox turns out to be ideally suited for explaining when and why these auctions work well.<sup>8</sup>

#### 6.3.1 Bad Auction Formats Cost Billions

Spectrum auction design is stressful, because small mistakes can be extremely costly. One cautionary tale is provided by an auction run by the New Zealand government in 1990 (before governments had much experience with auctions). For sale were 10 essentially identical national licenses for television broadcasting. For some reason, lost to the sands of time, the government decided to sell these licenses by running 10 second-price auctions in parallel. A *second-price* or *Vickrey* auction for a single good is a sealed-bid auction that awards the item to the highest bidder and charges her the highest bid by someone else (the second-highest bid overall). When selling a single item, the Vickrey auction is often a good solution. In particular, each bidder has a dominant strategy (always at least as good as all alternatives), which is to bid her true maximum willingness-to-pay.<sup>9,10</sup>

The nice properties of a second-price auction evaporate if many of them are run simultaneously. A bidder can now submit up to one bid in each auction, with each license awarded to the highest bidder (on that license) at a price equal to the second-highest bid (on that license). With multiple simultaneous auctions, it is no longer clear how a bidder

---

<sup>8</sup>Much of the discussion in Sections 6.3.1–6.3.3 is from [136, Lecture 8], which in turn takes inspiration from Milgrom [112].

<sup>9</sup>Intuitively, a second-price auction shades your bid optimally after the fact, so there's no reason to try to game it.

<sup>10</sup>For a more formal treatment of single-item auctions, see Section 9.1.1.

should bid. For example, imagine you want one of the licenses, but only one. How should you bid? One legitimate strategy is to pick one of the licenses—at random, say—and go for it. Another strategy is to bid less aggressively on multiple licenses, hoping that you get one at a bargain price, and that you don't inadvertently win extra licenses that you don't want. The difficulty is trading off the risk of winning too many licenses with the risk of winning too few.

The challenge of bidding intelligently in simultaneous sealed-bid auctions makes the auction format prone to poor outcomes. The revenue in the 1990 New Zealand auction was only \$36 million, a paltry fraction of the projected \$250 million. On one license, the high bid was \$100,000 while the second-highest bid (and selling price) was \$6! On another, the high bid was \$7 million and the second-highest was \$5,000. To add insult to injury, the winning bids were made available to the public, who could then see just how much money was left on the table!

### **6.3.2 Simultaneous Ascending Auctions**

Modern spectrum auctions are based on *simultaneous ascending auctions* (SAAs), following 1993 proposals by McAfee and by Milgrom and Wilson. You've seen—in the movies, at least—the call-and-response format of an ascending single-item auction, where an auctioneer asks for takers at successively higher prices. Such an auction ends when there's only one person left accepting the currently proposed price (who then wins, at this price). Conceptually, SAAs are like a bunch of single-item English auctions being run in parallel in the same room, with one auctioneer per item.

The primary reason that SAAs work better than sequential or sealed-bid auctions is *price discovery*. As a bidder acquires better information about the likely selling prices of licenses, she can implement mid-course corrections—abandoning licenses for which competition is fiercer than anticipated, snapping up unexpected bargains, and rethinking which packages of licenses to assemble. The format typically resolves the miscoordination problems that plague simultaneous sealed-bid auctions.

### 6.3.3 Inefficiency in SAAs

SAAs have two big vulnerabilities. The first problem is *demand reduction*, and this is relevant even when items are substitutes.<sup>11</sup> Demand reduction occurs when a bidder asks for fewer items than she really wants, to lower competition and therefore the prices paid for the items that she gets.

To illustrate, suppose there are two identical items and two bidders. By the *valuation* of a bidder for a given bundle of items, we mean her maximum willingness to pay for that bundle. Suppose the first bidder has valuation 10 for one of the items and valuation 20 for both. The second bidder has valuation 8 for one of the items and does not want both (i.e., her valuation remains 8 for both). The socially optimal outcome is to give both licenses to the first bidder. Now consider how things play out in an SAA. The second bidder would be happy to have either item at any price less than 8. Thus, the second bidder drops out only when the prices of both items exceed 8. If the first bidder stubbornly insists on winning both items, her utility is  $20 - 16 = 4$ . An alternative strategy for the first bidder is to simply concede the second item and never bid on it. The second bidder takes the second item and (because she only wants one license) withdraws interest in the first, leaving it for the first bidder. Both bidders get their item essentially for free, and the utility of the first bidder has jumped to 10.

The second big problem with SAAs is relevant when items can be complements, and is called the *exposure problem*.<sup>12</sup> As an example, consider two bidders and two nonidentical items. The first bidder only wants both items—they are complementary items for the bidder—and

---

<sup>11</sup>Items are substitutes if they provide diminishing returns—having one item only makes others less valuable. For two items  $A$  and  $B$ , for example, the substitutes condition means that a bidder's value for the bundle of  $A$  and  $B$  is at most the sum of her values for  $A$  and  $B$  individually. In a spectrum auction context, two licenses for the same area with equal-sized frequency ranges are usually substitute items.

<sup>12</sup>Items are complements if there are synergies between them, so that possessing one makes others more valuable. With two items  $A$  and  $B$ , this translates to a bidder's valuation for the bundle of  $A$  and  $B$  exceeding the sum of her valuations for  $A$  and  $B$  individually. Complements arise naturally in wireless spectrum auctions, as some bidders want a collection of licenses that are adjacent, either in their geographic areas or in their frequency ranges.

her valuation is 100 for them (and 0 for anything else). The second bidder is willing to pay 75 for either item but only wants one item. The socially optimal outcome is to give both items to the first bidder. But in an SAA, the second bidder will not drop out until the price of both items reaches 75. The first bidder is in a no-win situation: to get both items she would have to pay 150, more than her value. The scenario of winning only one item for a nontrivial price could be even worse. Thus the exposure problem leads to economically inefficient allocations for two reasons. First, an overly aggressive bidder might acquire unwanted items. Second, an overly tentative bidder might fail to acquire items for which she has the highest valuation.

### 6.3.4 When Do SAAs Work Well?

If you ask experts who design or consult for bidders in real-world SAAs, a rough consensus emerges about when they are likely to work well.

**Folklore Belief 1.** *Without strong complements*, SAAs work pretty well. Demand reduction does happen, but it is not a deal-breaker because the loss of efficiency appears to be small.

**Folklore Belief 2.** *With strong complements*, simple auctions like SAAs are not good enough. The exposure problem is a deal-breaker because it can lead to very poor outcomes (in terms of both economic efficiency and revenue).

There are a number of beautiful and useful theoretical results about spectrum auctions in the economics literature, but none map cleanly to these two folklore beliefs. A possible explanation: translating these beliefs into theorems seems to fundamentally involve approximate optimality guarantees, a topic that is largely avoided by economists but right in the wheelhouse of theoretical computer science.

In the standard model of *combinatorial auctions*, there are  $n$  bidders (e.g., telecoms) and  $m$  items (e.g., licenses).<sup>13</sup> Bidder  $i$  has a nonnegative valuation  $v_i(S)$  for each subset  $S$  of items she might receive. Note that, in general, describing a bidder's valuation function requires  $2^m$

---

<sup>13</sup>This model is treated more thoroughly in the next lecture (see Section 7.1).

parameters. Each bidder wants to maximize her utility, which is the value of the items received minus the total price paid for them. From a social perspective, we'd like to award bundles of items  $T_1, \dots, T_n$  to the bidders to maximize the *social welfare*  $\sum_{i=1}^n v_i(T_i)$ .

To make the first folklore belief precise, we need to commit to a definition of “without strong complements” and to a specific auction format. We'll focus on simultaneous first-price auctions (S1As), where each bidder submits a separate bid for each item, for each item the winner is the highest bidder (on that item), and winning bidders pay their bid on each item won.<sup>14</sup> One relatively permissive definition of “complement-free” is to restrict bidders to have *subadditive valuations*. This means what it sounds like: if  $A$  and  $B$  are two bundles of items, then bidder  $i$ 's valuation  $v_i(A \cup B)$  for their union should be at most the sum  $v_i(A) + v_i(B)$  of her valuations for each bundle separately. Observe that subadditivity is violated in the exposure problem example in Section 6.3.3.

We also need to define what we mean by “the outcome of an auction” like S1As. Remember that bidders are strategic, and will bid to maximize their utility (value of items won minus the price paid). Thus we should prove approximation guarantees for the *equilibria* of auctions. Happily, computer scientists have been working hard since 1999 to prove approximation guarantees for game-theoretic equilibria, also known as bounds on *the price of anarchy* [97, 131, 139].<sup>15</sup> In the early days, price-of-anarchy bounds appeared somewhat ad hoc and problem-specific. Fast forwarding to the present, we now have a powerful and user-friendly theory for proving price-of-anarchy bounds, which combine “extension theorems” and “composition theorems” to build up bounds for complex settings (including S1As) from bounds for simple settings.<sup>16</sup>

---

<sup>14</sup>Similar results hold for other auction formats, like simultaneous second-price auctions. Directly analyzing what happens in iterative auctions like SAAs when there are multiple items appears difficult.

<sup>15</sup>See Section 7.3.2 of the next lecture for a formal definition.

<sup>16</sup>We will say more about this theory in Lunar Lecture 10. See also Roughgarden *et al.* [141] for a recent survey.

In particular, Feldman *et al.* [54] proved the following translation of Folklore Belief #1.<sup>17</sup>

**Theorem 6.1** (Feldman *et al.* [54]). *When every bidder has a subadditive valuation, every equilibrium of an S1A has social welfare at least 50% of the maximum possible.*

One version of Theorem 6.1 concerns (mixed) Nash equilibria in the full-information model (in which bidders' valuations are common knowledge), as studied in the Solar Lectures. Even here, the bound in Theorem 6.1 is tight in the worst case [38]. The approximation guarantee in Theorem 6.1 holds more generally for *Bayes-Nash equilibria*, the standard equilibrium notion for games of incomplete information.<sup>18</sup>

Moving on to the second folklore belief, let's now drop the subadditivity restriction. S1As no longer work well.

**Theorem 6.2** (Hassidim *et al.* [78]). *When bidders have arbitrary valuations, an S1A can have a mixed Nash equilibrium with social welfare arbitrarily smaller than the maximum possible.*

Thus for S1As, the perspective of worst-case approximation confirms the dichotomy between the cases of substitutes and complements. But the lower bound in Theorem 6.2 applies only to one specific auction format. Could we do better with a different natural auction format? Folklore Belief #2 asserts the stronger statement that *no* "simple" auction works well with general valuations. This stronger statement can also be translated into a theorem (using nondeterministic communication complexity), and this will be the main subject of Lunar Lecture 7.

---

<sup>17</sup>To better appreciate this result, we note that multi-item auctions like S1As are so strategically complex that they have historically been seen as unanalyzable. For example, we have no idea what their equilibria look like in general. Nevertheless, we can prove good approximation guarantees for them!

<sup>18</sup>In more detail, in this model there is a commonly known prior distribution over bidders' valuations. In a Bayes-Nash equilibrium, every bidder bids to maximize her expected utility given her information at the time: her own valuation, her posterior belief about other bidders' valuations, and the bidding strategies (mapping valuations to bids) used by the other bidders. Theorem 6.1 continues to hold for every Bayes-Nash equilibrium of an S1A, as long as bidders' valuations are independently (and not necessarily identically) distributed.



**Theorem 6.3** [133]. *With general valuations, every simple auction can have an equilibrium with social welfare arbitrarily smaller than the maximum possible.*

The definition of “simple” used in Theorem 6.3 is quite generous: it requires only that the number of strategies available to each player is *sub-doubly-exponential* in the number of items  $m$ . For example, running separate single-item auctions provides each player with only an exponential (in  $m$ ) number of strategies (assuming a bounded number of possible bid values for each item). Thus Theorem 6.3 makes use of the theoretical computer science toolbox to provide solid footing for Folklore Belief #2.

# 7

---

## Communication Barriers to Near-Optimal Equilibria

---

This lecture is about the communication complexity of the welfare-maximization problem in combinatorial auctions and its implications for the price of anarchy of simple auctions. Section 7.1 defines the model, Section 7.2 proves lower bounds for nondeterministic communication protocols, and Section 7.3 gives a black-box translation of these lower bounds to equilibria of simple auctions. In particular, Section 7.3 provides the proof of Theorem 6.3 from last lecture. Section 7.4 concludes with a juicy open problem on the topic.<sup>1</sup>

### 7.1 Welfare Maximization in Combinatorial Auctions

Recall from Section 6.3.4 the basic setup in the study of combinatorial auctions.

1. There are  $k$  players. (In a spectrum auction, these are the telecoms.)
2. There is a set  $M$  of  $m$  items. (In a spectrum auction, these are the licenses.)

---

<sup>1</sup>Much of this lecture is drawn from [137, Lecture 7].

3. Each player  $i$  has a *valuation*  $v_i : 2^M \rightarrow \mathbb{R}_+$ . The number  $v_i(T)$  indicates  $i$ 's value, or willingness to pay, for the items  $T \subseteq M$ . The valuation is the private input of player  $i$ , meaning that  $i$  knows  $v_i$  but none of the other  $v_j$ 's. (I.e., this is a number-in-hand model.) We assume that  $v_i(\emptyset) = 0$  and that the valuations are *monotone*, meaning  $v_i(S) \leq v_i(T)$  whenever  $S \subseteq T$ . (The more items, the better.) To avoid bit complexity issues, we'll also assume that all of the  $v_i(T)$ 's are integers with description length polynomial in  $k$  and  $m$ . We sometimes impose additional restrictions on the valuations to study special cases of the general problem.

Note that we may have more than two players—more than just Alice and Bob. (For example, you might want to think of  $k$  as  $\approx m^{1/3}$ .) Also note that the description length of a player's valuation is exponential in the number of items  $m$ .

In the *welfare-maximization problem*, the goal is to partition the items  $M$  into sets  $T_1, \dots, T_k$  to maximize, at least approximately, the social welfare

$$\sum_{i=1}^k v_i(T_i), \quad (7.1)$$

using communication polynomial in  $k$  and  $m$ . Note this amount of communication is logarithmic in the sizes of the private inputs. Maximizing social welfare (7.1) is the most commonly studied objective in combinatorial auctions, and it is the one we will focus on in this lecture.

## 7.2 Communication Lower Bounds for Approximate Welfare Maximization

This section studies the communication complexity of computing an approximately welfare-maximizing allocation in a combinatorial auction. For reasons that will become clear in Section 7.3, we are particularly interested in the problem's nondeterministic communication complexity.<sup>2</sup>

---

<sup>2</sup>For basic background on nondeterministic multi-party communication protocols, see Kushilevitz and Nisan [98] or Roughgarden [137].

### 7.2.1 Lower Bound for General Valuations

We begin with a result of Nisan [120] showing that, alas, computing even a very weak approximation of the welfare-maximizing allocation requires exponential communication. To make this precise, it is convenient to turn the optimization problem of welfare maximization into a decision problem. In the WELFARE-MAXIMIZATION( $k$ ) problem, the goal is to correctly identify inputs that fall into one of the following two cases:

- (1) Every partition  $(T_1, \dots, T_k)$  of the items has welfare at most 1.
- (0) There exists a partition  $(T_1, \dots, T_k)$  of the items with welfare at least  $k$ .

Arbitrary behavior is permitted on inputs that fail to satisfy either (1) or (0). Clearly, communication lower bounds for WELFARE-MAXIMIZATION( $k$ ) apply to the more general problem of obtaining a better-than- $k$ -approximation of the maximum welfare.<sup>3</sup>

**Theorem 7.1** [120]. *The nondeterministic communication complexity of WELFARE-MAXIMIZATION( $k$ ) is  $\exp\{\Omega(m/k^2)\}$ , where  $k$  is the number of players and  $m$  is the number of items.*

This lower bound is exponential in  $m$ , provided that  $m = \Omega(k^{2+\epsilon})$  for some  $\epsilon > 0$ . Since communication complexity lower bounds apply even to players who cooperate perfectly, this impossibility result holds even when all of the (tricky) incentive issues are ignored.

### 7.2.2 The Multi-Disjointness Problem

The plan for the proof of Theorem 7.1 is to reduce a multi-party version of the DISJOINTNESS problem to the WELFARE-MAXIMIZATION( $k$ ) problem. There is some ambiguity about how to define a version of DISJOINTNESS for three or more players. For example, suppose there are three players, and among the three possible pairings of them, two have disjoint sets while the third have intersecting sets. Should this

---

<sup>3</sup>Achieving a  $k$ -approximation is trivial: every player communicates her value  $v_i(M)$  for the whole set of items, and the entire set of items is awarded to the bidder with the highest value for them.

count as a “yes” or “no” instance? We’ll skirt this issue by worrying only about unambiguous inputs, that are either “totally disjoint” or “totally intersecting.”

Formally, in the MULTI-DISJOINTNESS problem, each of the  $k$  players  $i$  holds an input  $\mathbf{x}_i \in \{0, 1\}^n$ . (Equivalently, a set  $S_i \subseteq \{1, 2, \dots, n\}$ .) The task is to correctly identify inputs that fall into one of the following two cases:

- (1) “Totally disjoint,” with  $S_i \cap S_{i'} = \emptyset$  for every  $i \neq i'$ .
- (0) “Totally intersecting,” with  $\bigcap_{i=1}^k S_i \neq \emptyset$ .

When  $k = 2$ , this is the standard DISJOINTNESS problem. When  $k > 2$ , there are inputs that are neither 1-inputs nor 0-inputs. We let protocols off the hook on such ambiguous inputs—they can answer “1” or “0” with impunity.

The following communication complexity lower bound for MULTI-DISJOINTNESS is credited to Jaikumar Radhakrishnan and Venkatesh Srinivasan in [120]. (The proof is elementary, and for completeness is given in Section 7.5.)

**Theorem 7.2.** *The nondeterministic communication complexity of MULTI-DISJOINTNESS, with  $k$  players with  $n$ -bit inputs, is  $\Omega(n/k)$ .*

This nondeterministic lower bound is for verifying a 1-input. (It is easy to verify a 0-input—the prover just suggests the index of an element  $r$  in  $\bigcap_{i=1}^k S_i$ .)<sup>4</sup>

### 7.2.3 Proof of Theorem 7.1

The proof of Theorem 7.1 relies on Theorem 7.2 and a combinatorial gadget. We construct this gadget using the probabilistic method. Consider  $t$  random partitions  $P^1, \dots, P^t$  of  $M$ , where  $t$  is a parameter to be defined later. By a random partition  $P^j = (P_1^j, \dots, P_k^j)$ , we mean

<sup>4</sup>In proving Theorem 7.1, we’ll be interested in the case where  $k$  is much smaller than  $n$ , such as  $k = \Theta(\log n)$ . Intuition might suggest that the lower bound should be  $\Omega(n)$  rather than  $\Omega(n/k)$ , but this is incorrect—a slightly non-trivial argument shows that Theorem 7.2 is tight for nondeterministic protocols (for all small enough  $k$ , like  $k = O(\sqrt{n})$ ). This factor- $k$  difference won’t matter for our applications, however.

that each of the  $m$  items is assigned to exactly one of the  $k$  players, independently and uniformly at random.

We are interested in the probability that two classes of different partitions intersect: for all  $i \neq i'$  and  $j \neq \ell$ , because the probability that a given item is assigned to  $i$  in  $P^j$  and also to  $i'$  in  $P^\ell$  is  $\frac{1}{k^2}$ , we have

$$\Pr[P_i^j \cap P_{i'}^\ell = \emptyset] = \left(1 - \frac{1}{k^2}\right)^m \leq e^{-m/k^2}.$$

Taking a Union Bound over the  $k$  choices for  $i$  and  $i'$  and the  $t$  choices for  $j$  and  $\ell$ , we have

$$\Pr[\exists i \neq i', j \neq \ell \text{ s.t. } P_i^j \cap P_{i'}^\ell = \emptyset] \leq k^2 t^2 e^{-m/k^2}. \quad (7.2)$$

Call  $P^1, \dots, P^t$  an *intersecting family* if  $P_i^j \cap P_{i'}^\ell \neq \emptyset$  whenever  $i \neq i'$ ,  $j \neq \ell$ . By (7.2), the probability that our random experiment fails to produce an intersecting family is less than 1 provided  $t < \frac{1}{k} e^{m/2k^2}$ . The following lemma is immediate.

**Lemma 7.3.** *For every  $m, k \geq 1$ , there exists an intersecting family of partitions  $P^1, \dots, P^t$  with  $t = \exp\{\Omega(m/k^2)\}$ .*

A simple combination of Theorem 7.2 and Lemma 7.3 now proves Theorem 7.1.

*Proof.* (of Theorem 7.1) The proof is a reduction from MULTI-DISJOINTNESS. Fix  $k$  and  $m$ . (To be interesting,  $m$  should be significantly bigger than  $k^2$ .) Let  $(S_1, \dots, S_k)$  denote an input to MULTI-DISJOINTNESS with  $t$ -bit inputs, where  $t = \exp\{\Omega(m/k^2)\}$  is the same value as in Lemma 7.3. We can assume that the players have coordinated in advance on an intersecting family of  $t$  partitions of a set  $M$  of  $m$  items. Each player  $i$  uses this family and her input  $S_i$  to form the following valuation:

$$v_i(T) = \begin{cases} 1 & \text{if } T \supseteq P_i^j \text{ for some } j \in S_i \\ 0 & \text{otherwise.} \end{cases}$$

That is, player  $i$  is either happy (value 1) or unhappy (value 0), and is happy if and only if she receives all of the items in the corresponding class  $P_i^j$  of some partition  $P^j$  with index  $j$  belonging to its input to MULTI-DISJOINTNESS. The valuations  $v_1, \dots, v_k$  define an input to

WELFARE-MAXIMIZATION( $k$ ). Forming this input requires no communication between the players.

Consider the case where the input to MULTI-DISJOINTNESS is a 1-input, with  $S_i \cap S_{i'} = \emptyset$  for every  $i \neq i'$ . We claim that the induced input to WELFARE-MAXIMIZATION( $k$ ) is a 1-input, with maximum welfare at most 1. To see this, consider a partition  $(T_1, \dots, T_k)$  in which some player  $i$  is happy (with  $v_i(T_i) = 1$ ). For some  $j \in S_i$ , player  $i$  receives all the items in  $P_i^j$ . Since  $j \notin S_{i'}$  for every  $i' \neq i$ , the only way to make a second player  $i'$  happy is to give her all the items in  $P_{i'}^\ell$  in some other partition  $P^\ell$  with  $\ell \in S_{i'}$  (and hence  $\ell \neq j$ ). Since  $P^1, \dots, P^t$  is an intersecting family, this is impossible —  $P_i^j$  and  $P_{i'}^\ell$  overlap for every  $\ell \neq j$ .

When the input to MULTI-DISJOINTNESS is a 0-input, with an element  $r$  in the mutual intersection  $\cap_{i=1}^k S_i$ , we claim that the induced input to WELFARE-MAXIMIZATION( $k$ ) is a 0-input, with maximum welfare at least  $k$ . This is easy to see: for  $i = 1, 2, \dots, k$ , assign the items of  $P_i^r$  to player  $i$ . Since  $r \in S_i$  for every  $i$ , this makes all  $k$  players happy.

This reduction shows that a (deterministic, nondeterministic, or randomized) protocol for WELFARE-MAXIMIZATION( $k$ ) yields one for MULTI-DISJOINTNESS (with  $t$ -bit inputs) with the same communication. We conclude that the nondeterministic communication complexity of WELFARE-MAXIMIZATION( $k$ ) is  $\Omega(t/k) = \exp\{\Omega(m/k^2)\}$ .  $\square$

### 7.2.4 Subadditive Valuations

To an algorithms person, Theorem 7.1 is depressing, as it rules out any non-trivial positive results. A natural idea is to seek positive results by imposing additional structure on players' valuations. Many such restrictions have been studied. We consider here the case of *subadditive* valuations (see also Section 6.3.4 of the preceding lecture), where each  $v_i$  satisfies  $v_i(S \cup T) \leq v_i(S) + v_i(T)$  for every pair  $S, T \subseteq M$ .

Our reduction in Theorem 7.1 easily implies a weaker inapproximability result for welfare maximization with subadditive valuations. Formally, define the WELFARE-MAXIMIZATION(2) problem as that of identifying inputs that fall into one of the following two cases:

- (1) Every partition  $(T_1, \dots, T_k)$  of the items has welfare at most  $k + 1$ .
- (0) There exists a partition  $(T_1, \dots, T_k)$  of the items with welfare at least  $2k$ .

Communication lower bounds for WELFARE-MAXIMIZATION(2) apply also to the more general problem of obtaining a better-than-2-approximation of the maximum social welfare.

**Theorem 7.4** (Dobzinski *et al.* [49]). *The nondeterministic communication complexity of WELFARE-MAXIMIZATION(2) is  $\exp\{\Omega(m/k^2)\}$ , even when all players have subadditive valuations.*

This theorem follows from a modification of the proof of Theorem 7.1. The 0-1 valuations used in that proof are not subadditive, but they can be made subadditive by adding 1 to each bidder's valuation  $v_i(T)$  of each non-empty set  $T$ . The social welfare obtained in inputs corresponding to 1- and 0-inputs of MULTI-DISJOINTNESS become  $k + 1$  and  $2k$ , respectively, and this completes the proof of Theorem 7.4.

There is also a quite non-trivial deterministic and polynomial-communication protocol that guarantees a 2-approximation of the social welfare when bidders have subadditive valuations [52].

### 7.3 Lower Bounds on the Price of Anarchy of Simple Auctions

The lower bounds of the previous section show that every protocol for the welfare-maximization problem that interacts with the players and then explicitly computes an allocation has either a bad approximation ratio or high communication cost. Over the past decade, many researchers have considered shifting the work from the protocol to the players, by analyzing the equilibria of simple auctions. Can such equilibria bypass the communication complexity lower bounds proved in Section 7.2? The answer is not obvious, because equilibria are defined non-constructively, and not through a low-cost communication protocol.

#### 7.3.1 Auctions as Games

What do we mean by a “simple” auction? For example, recall the *simultaneous first-price auctions (S1As)* introduced in Section 6.3.4



of the preceding lecture. Each player  $i$  chooses a strategy  $b_{i1}, \dots, b_{im}$ , with one bid per item.<sup>5</sup> Each item is sold separately in parallel using a “first-price auction”—the item is awarded to the highest bidder on that item, with the selling price equal to that bidder’s bid.<sup>6</sup> The payoff of a player  $i$  in a given outcome (i.e., given a choice of strategy for each player) is then her utility:

$$\underbrace{v_i(T_i)}_{\text{value of items won}} - \underbrace{\sum_{j \in T_i} b_{ij}}_{\text{price paid for them}},$$

where  $T_i$  denotes the items on which  $i$  is the highest bidder (given the bids of the others).

Bidders strategize already in a first-price auction for a single item—a bidder certainly doesn’t want to bid her actual valuation (this would guarantee utility 0), and instead will “shade” her bid down to a lower value. (How much to shade is a tricky question, and depends on what the other bidders are doing.) Thus it makes sense to assess the performance of an auction by its equilibria. As usual, a Nash equilibrium comprises a (randomized) strategy for each player, so that no player can unilaterally increase her expected payoff through a unilateral deviation to some other strategy (given how the other players are bidding).

### 7.3.2 The Price of Anarchy

So how good are the equilibria of various auction games, such as S1As? To answer this question, we use an analog of the approximation ratio, adapted for equilibria. Given a game  $G$  (like an S1A) and a nonnegative maximization objective function  $f$  on the outcomes (like the social welfare), Koutsoupias and Papadimitriou [97] defined the *price of anarchy* (POA) of  $G$  as the ratio between the objective function value of an

<sup>5</sup>To keep the game finite, let’s agree that each bid has to be an integer between 0 and some known upper bound  $B$ .

<sup>6</sup>In the preceding lecture we mentioned the *Vickrey* or *second-price* auction, where the winner does not pay their own bid, but rather the highest bid by someone else (the second-highest overall). We’ll stick with S1As for simplicity, but similar results are known for simultaneous second-price auctions, as well.

optimal solution, and that of the worst equilibrium:

$$\text{PoA}(G) := \frac{f(\text{OPT}(G))}{\min_{\rho \text{ is an equilibrium of } G} f(\rho)},$$

where  $\text{OPT}(G)$  denotes the optimal outcome of  $G$  (with respect to  $f$ ).<sup>7</sup> Thus the price of anarchy of a game quantifies the inefficiency of selfish behavior.<sup>8</sup> The POA of a game and a maximization objective function is always at least 1. We can identify “good performance” of a system with strategic participants as having a POA close to 1.<sup>9</sup>

The POA depends on the choice of equilibrium concept. For example, the POA with respect to approximate Nash equilibria can only be worse (i.e., bigger) than for exact Nash equilibria (since there are only more of the former).

### 7.3.3 The Price of Anarchy of S1As

As we saw in Theorem 6.1 of the preceding lecture, the equilibria of simple auctions like S1As can be surprisingly good.<sup>10</sup> We restate that result here.<sup>11</sup>

**Theorem 7.5** (Feldman *et al.* [54]). *In every S1A with subadditive bidder valuations, the POA is at most 2.*

This result is particularly impressive because achieving an approximation factor of 2 for the welfare-maximization problem with subadditive bidder valuations by any means (other than brute-force search) is not easy (see [52]).

---

<sup>7</sup>If  $\rho$  is a probability distribution over outcomes, as in a mixed Nash equilibrium, then  $f(\rho)$  denotes the expected value of  $f$  w.r.t.  $\rho$ .

<sup>8</sup>Games generally have multiple equilibria. Ideally, we’d like an approximation guarantee that applies to *all* equilibria, so that we don’t need to worry about which one is reached—this is the point of the POA.

<sup>9</sup>One caveat is that it’s not always clear that a system will reach an equilibrium in a reasonable amount of time. A natural way to resolve this issue is to relax the notion of equilibrium enough so that it become relatively easy to reach an equilibrium. See Lunar Lecture 10 for more on this point.

<sup>10</sup>The first result of this type, for simultaneous second-price auctions and bidders with submodular valuations, is due to Christodoulou *et al.* [37].

<sup>11</sup>For a proof, see the original paper [54] or course notes by the author [134, Lecture 17.5].

As mentioned last lecture, a recent result shows that the analysis of [54] is tight.

**Theorem 7.6** (Christodoulou *et al.* [38]). *The worst-case POA of S1As with subadditive bidder valuations is at least 2.*

The proof of Theorem 7.6 is an ingenious explicit construction—the authors exhibit a choice of subadditive bidder valuations and a Nash equilibrium of the corresponding S1A so that the welfare of this equilibrium is only half of the maximum possible. One reason that proving results like Theorem 7.6 is challenging is that it can be difficult to solve for a (bad) equilibrium of a complex game like a S1A.

### 7.3.4 Price-of-Anarchy Lower Bounds from Communication Complexity

Theorem 7.5 motivates an obvious question: can we do better? Theorem 7.6 implies that the analysis in [54] cannot be improved, but can we reduce the POA by considering a different auction? Ideally, the auction would still be “reasonably simple” in some sense. Alternatively, perhaps no “simple” auction could be better than S1As? If this is the case, it’s not clear how to prove it directly—proving lower bounds via explicit constructions auction-by-auction does not seem feasible.

Perhaps it’s a clue that the POA upper bound of 2 for S1As (Theorem 7.5) gets stuck at the same threshold for which there is a lower bound for protocols that use polynomial communication (Theorem 7.4). It’s not clear, however, that a lower bound for low-communication protocols has anything to do with equilibria. Can we extract a low-communication protocol from an equilibrium?

**Theorem 7.7** (Roughgarden [133]). *Fix a class  $\mathcal{V}$  of possible bidder valuations. Suppose that, for some  $\alpha \geq 1$ , there is no nondeterministic protocol with subexponential (in  $m$ ) communication for the 1-inputs of the following promise version of the welfare-maximization problem with bidder valuations in  $\mathcal{V}$ :*

- (1) *Every allocation has welfare at most  $W^*/\alpha$ .*
- (0) *There exists an allocation with welfare at least  $W^*$ .*

Let  $\epsilon$  be bounded below by some inverse polynomial function of  $k$  and  $m$ . Then, for every auction with sub-doubly-exponential (in  $m$ ) strategies per player, the worst-case POA of  $\epsilon$ -approximate Nash equilibria with bidder valuations in  $\mathcal{V}$  is at least  $\alpha$ .

Theorem 7.7 says that lower bounds for nondeterministic protocols carry over to all “sufficiently simple” auctions, where “simplicity” is measured by the number of strategies available to each player. These POA lower bounds follow automatically from communication complexity lower bounds, and do not require any new explicit constructions.

To get a feel for the simplicity constraint, note that S1As with integral bids between 0 and  $B$  have  $(B + 1)^m$  strategies per player—singly exponential in  $m$ . On the other hand, in a “direct-revelation” auction, where each bidder is allowed to submit a bid on each bundle  $S \subseteq M$  of items, each player has a doubly-exponential (in  $m$ ) number of strategies.<sup>12</sup>

The POA lower bound promised by Theorem 7.7 is only for approximate Nash equilibria; since the POA is a worst-case measure and the set of  $\epsilon$ -NE is nondecreasing with  $\epsilon$ , this is weaker than a lower bound for exact Nash equilibria. It is an open question whether or not Theorem 7.7 holds also for the POA of exact Nash equilibria.<sup>13</sup>

Theorem 7.7 has a number of interesting corollaries. First, consider the case where  $\mathcal{V}$  is the set of subadditive valuations. Since S1As have only a singly-exponential (in  $m$ ) number of strategies per player, Theorem 7.7 applies to them. Thus, combining it with Theorem 7.4 recovers the POA lower bound of Theorem 7.6—modulo the exact vs. approximate Nash equilibria issue—and shows the optimality of the upper bound in Theorem 7.5 without an explicit construction. Even more interestingly, this POA lower bound of 2 applies not only to S1As, but more generally to all auctions in which each player has a sub-doubly-exponential number of strategies. Thus, S1As are in fact *optimal* among

---

<sup>12</sup>Equilibria can achieve the optimal welfare in a direct-revelation auction, so some bound on the number of strategies is necessary in Theorem 7.7.

<sup>13</sup>Arguably, Theorem 7.7 is good enough for all practical purposes—a POA upper bound that holds for exact Nash equilibria and does not hold (at least approximately) for approximate Nash equilibria with very small  $\epsilon$  is too brittle to be meaningful.

## 7.3. Lower Bounds on the Price of Anarchy of Simple Auctions 121

the class of all such auctions when bidders have subadditive valuations (w.r.t. the worst-case POA of  $\epsilon$ -approximate Nash equilibria).

We can also take  $\mathcal{V}$  to be the set of all (monotone) valuations, and then combine Theorem 7.7 with Theorem 7.1 to deduce that no “simple” auction gives a non-trivial (i.e., better-than- $k$ ) approximation for general bidder valuations. We conclude that with general valuations, complexity is essential to any auction format that offers good equilibrium guarantees. This completes the proof of Theorem 6.3 from the preceding lecture and formalizes the second folklore belief in Section 6.3.4; we restate that result here.

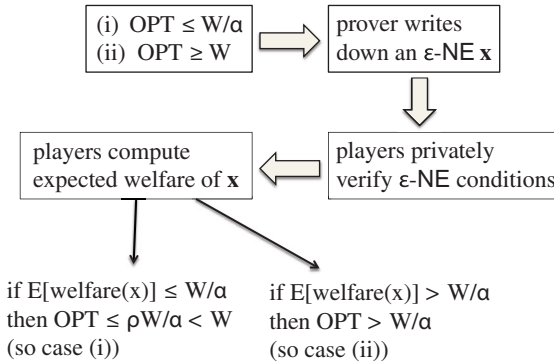
**Theorem 7.8** [133]. *With general valuations, every simple auction can have equilibria with social welfare arbitrarily worse than the maximum possible.*

### 7.3.5 Proof of Theorem 7.7

Presumably, the proof of Theorem 7.7 extracts a low-communication protocol from a good POA bound. The hypothesis of Theorem 7.7 offers the clue that we should be looking to construct a nondeterministic protocol. So what could we use an all-powerful prover for? We’ll see that a good role for the prover is to suggest a Nash equilibrium to the players.

Unfortunately, it can be too expensive for the prover to write down the description of a Nash equilibrium, even in S1As. Recall that a mixed strategy is a distribution over pure strategies, and that each player has an exponential (in  $m$ ) number of pure strategies available in a S1A. Specifying a Nash equilibrium thus requires an exponential number of probabilities. To circumvent this issue, we resort to approximate Nash equilibria, which are guaranteed to exist even if we restrict ourselves to distributions with small descriptions. We proved this for two-player games in Solar Lecture 1 (Theorem 1.15); the same argument works for games with any number of players.

**Lemma 7.9** (Lipton *et al.* [104]). *For every  $\epsilon > 0$  and every game with  $k$  players with strategy sets  $A_1, \dots, A_k$ , there exists an  $\epsilon$ -approximate Nash equilibrium with description length polynomial in  $k$ ,  $\log(\max_{i=1}^k |A_i|)$ , and  $\frac{1}{\epsilon}$ .*



**Figure 7.1:** Proof of Theorem 7.7. How to extract a low-communication nondeterministic protocol from a good price-of-anarchy bound.

In particular, every game with a sub-doubly-exponential number of strategies admits an approximate Nash equilibrium with subexponential description length.

We now proceed to the proof of Theorem 7.7.

*Proof.* (of Theorem 7.7) Fix an auction with at most  $A$  strategies per player, and a value for  $\epsilon = \Omega(1/\text{poly}(k, m))$ . Assume that, no matter what the bidder valuations  $v_1, \dots, v_k \in \mathcal{V}$  are, the POA of  $\epsilon$ -approximate Nash equilibria of the auction is at most  $\rho < \alpha$ . We will show that  $A$  must be doubly-exponential in  $m$ .

Consider the following nondeterministic protocol for verifying a 1-input of the welfare-maximization problem—for convincing the  $k$  players that every allocation has welfare at most  $W^*/\alpha$ . See also Figure 7.1. The prover writes on a publicly visible blackboard an  $\epsilon$ -approximate Nash equilibrium  $(\sigma_1, \dots, \sigma_k)$  of the auction, with description length polynomial in  $k$ ,  $\log A$ , and  $\frac{1}{\epsilon} = O(\text{poly}(k, m))$  as guaranteed by Lemma 7.9. The prover also writes down the expected welfare contribution  $\mathbf{E}[v_i(S)]$  of each bidder  $i$  in this equilibrium.

Given this advice, each player  $i$  verifies that  $\sigma_i$  is indeed an  $\epsilon$ -approximate best response to the other  $\sigma_j$ 's and that her expected welfare is as claimed when all players play the mixed strategies  $\sigma_1, \dots, \sigma_k$ . Crucially, player  $i$  is fully equipped to perform both of these checks

without any communication—she knows her valuation  $v_i$  (and hence her utility in each outcome of the game) and the mixed strategies used by all players, and this is all that is needed to verify her  $\epsilon$ -approximate Nash equilibrium conditions and compute her expected contribution to the social welfare.<sup>14</sup> Player  $i$  accepts if and only if the prover's advice passes these two tests, and if the expected welfare of the equilibrium is at most  $W^*/\alpha$ .

For the protocol correctness, consider first the case of a 1-input, where every allocation has welfare at most  $W^*/\alpha$ . If the prover writes down the description of an arbitrary  $\epsilon$ -approximate Nash equilibrium and the appropriate expected contributions to the social welfare, then all of the players will accept (the expected welfare is obviously at most  $W^*/\alpha$ ). We also need to argue that, for the case of a 0-input—where some allocation has welfare at least  $W^*$ —there is no proof that causes all of the players to accept. We can assume that the prover writes down an  $\epsilon$ -approximate Nash equilibrium and its correct expected welfare  $W$ , as otherwise at least one player will reject. Because the maximum-possible welfare is at least  $W^*$  and (by assumption) the POA of  $\epsilon$ -approximate Nash equilibria is at most  $\rho < \alpha$ , the expected welfare of the given  $\epsilon$ -approximate Nash equilibrium must satisfy  $W \geq W^*/\rho > W^*/\alpha$ . The players will reject such a proof, so we can conclude that the protocol is correct. Our assumption then implies that the protocol has communication cost exponential in  $m$ . Since the cost of the protocol is polynomial in  $k$ ,  $m$ , and  $\log A$ ,  $A$  must be doubly exponential in  $m$ .  $\square$

Conceptually, the proof of Theorem 7.7 argues that, when the POA of  $\epsilon$ -approximate Nash equilibria is small, every  $\epsilon$ -approximate Nash equilibrium provides a privately verifiable proof of a good upper bound on the maximum-possible welfare. When such upper bounds require large communication, the equilibrium description length (and hence the number of available strategies) must be large.

---

<sup>14</sup>These computations may take a super-polynomial amount of time, but they do not contribute to the protocol's cost.

## 7.4 An Open Question

While Theorems 7.4, 7.5, and 7.7 pin down the best-possible POA achievable by simple auctions with subadditive bidder valuations, open questions remain for other valuation classes. For example, a valuation  $v_i$  is *submodular* if it satisfies

$$v_i(T \cup \{j\}) - v_i(T) \leq v_i(S \cup \{j\}) - v_i(S)$$

for every  $S \subseteq T \subset M$  and  $j \notin T$ . This is a “diminishing returns” condition for set functions. Every monotone submodular function is also subadditive, so welfare-maximization with the former valuations is only easier than with the latter.

The worst-case POA of S1As is exactly  $\frac{e}{e-1} \approx 1.58$  when bidders have submodular valuations. The upper bound was proved by Syrgkanis and Tardos [151], the lower bound by Christodoulou *et al.* [38]. It is an open question whether or not there is a simple auction with a smaller worst-case POA. The best lower bound known—for nondeterministic protocols and hence, by Theorem 7.7, for the POA of  $\epsilon$ -approximate Nash equilibria of simple auctions—is  $\frac{2e}{2e-1} \approx 1.23$  [48]. Intriguingly, there is an upper bound (very slightly) better than  $\frac{e}{e-1}$  for polynomial-communication protocols [53]—can this better upper bound also be realized as the POA of a simple auction? What is the best-possible approximation guarantee, either for polynomial-communication protocols or for the POA of simple auctions? Resolving this question would require either a novel auction format (better than S1As), a novel lower bound technique (better than Theorem 7.7), or both.

## 7.5 Appendix: Proof of Theorem 7.2

The proof of Theorem 7.2 proceeds in three easy steps.

**Step 1:** *Every nondeterministic protocol with communication cost  $c$  induces a cover of the 1-inputs of  $M(f)$  by at most  $2^c$  monochromatic boxes.* By “ $M(f)$ ,” we mean the  $k$ -dimensional array in which the  $i$ th dimension is indexed by the possible inputs of player  $i$ , and an array entry contains the value of the function  $f$  on the corresponding joint input.



By a “box,” we mean the  $k$ -dimensional generalization of a rectangle—a subset of inputs that can be written as a product  $A_1 \times A_2 \times \cdots \times A_k$ . By “monochromatic,” we mean a box that does not contain both a 1-input and a 0-input. (Recall that for the MULTI-DISJOINTNESS problem there are also inputs that are neither 1 nor 0—a monochromatic box can contain any number of these.) The proof of this step is the same as the standard one for the two-party case (see e.g. [98]).

**Step 2:** *The number of 1-inputs in  $M(f)$  is  $(k + 1)^n$ .* In a 1-input  $(\mathbf{x}_1, \dots, \mathbf{x}_k)$ , for every coordinate  $\ell$ , at most one of the  $k$  inputs has a 1 in the  $\ell$ th coordinate. This yields  $k + 1$  options for each of the  $n$  coordinates, thereby generating a total of  $(k + 1)^n$  1-inputs.

**Step 3:** *The number of 1-inputs in a monochromatic box is at most  $k^n$ .* Let  $B = A_1 \times A_2 \times \cdots \times A_k$  be a 1-box. The key claim here is: for each coordinate  $\ell = 1, \dots, n$ , there is a player  $i \in \{1, \dots, k\}$  such that, for every input  $\mathbf{x}_i \in A_i$ , the  $\ell$ th coordinate of  $\mathbf{x}_i$  is 0. That is, to each coordinate we can associate an “ineligible player” that, in this box, never has a 1 in that coordinate. This is easily seen by contradiction: otherwise, there exists a coordinate  $\ell$  such that, for every player  $i$ , there is an input  $\mathbf{x}_i \in A_i$  with a 1 in the  $\ell$ th coordinate. As a box,  $B$  contains the input  $(\mathbf{x}_1, \dots, \mathbf{x}_k)$ . But this is a 0-input, contradicting the assumption that  $B$  is a 1-box.

The claim implies the stated upper bound. Every 1-input of  $B$  can be generated by choosing, for each coordinate  $\ell$ , an assignment of at most one “1” in this coordinate to one of the  $k - 1$  eligible players for this coordinate. With only  $k$  choices per coordinate, there are at most  $k^n$  1-inputs in the box  $B$ .

**Conclusion:** Steps 2 and 3 imply that covering the 1s of the  $k$ -dimensional array of the MULTI-DISJOINTNESS function requires at least  $(1 + \frac{1}{k})^n$  1-boxes. By the discussion in Step 1, this implies a lower bound of  $n \log_2(1 + \frac{1}{k}) = \Theta(n/k)$  on the nondeterministic communication complexity of the MULTI-DISJOINTNESS function (and output 1). This concludes the proof of Theorem 7.2.

# 8

---

## Why Prices Need Algorithms

---

You’ve probably heard about “market-clearing prices,” which equate the supply and demand in a market. When are such prices guaranteed to exist? In the classical setting with divisible goods (milk, wheat, etc.), market-clearing prices exist under reasonably weak conditions [6]. But with indivisible goods (houses, spectrum licenses, etc.), such prices may or may not exist. As you can imagine, many papers in the economics and operations research literatures study necessary and sufficient conditions for existence. The punchline of today’s lecture, based on joint work with Inbal Talgam-Cohen [138], is that computational complexity considerations in large part govern whether or not market-clearing prices exist in a market of indivisible goods. This is cool and surprising because the question (existence of equilibria) seems to have nothing to do with computation (cf., the questions studied in the Solar Lectures).

### 8.1 Markets with Indivisible Items

The basic setup is the same as in the preceding lecture, when we were studying price-of-anarchy bounds for simple combinatorial auctions (Section 7.1). To review, there are  $k$  players, a set  $M$  of  $m$  items, and each player  $i$  has a valuation  $v_i : 2^M \rightarrow \mathbb{R}_+$  describing her maximum

willingness to pay for each bundle of items. For simplicity, we also assume that  $v_i(\emptyset) = 0$  and that  $v_i$  is monotone (with  $v_i(S) \leq v_i(T)$  whenever  $S \subseteq T$ ). As in last lecture, we will often vary the class  $\mathcal{V}$  of allowable valuations to make the setting more or less complex.

### 8.1.1 Walrasian Equilibria

Next is the standard definition of “market-clearing prices” in a market with multiple indivisible items.

**Definition 8.1 (Walrasian Equilibrium).** A *Walrasian equilibrium* is an allocation  $S_1, \dots, S_k$  of the items of  $M$  to the players and nonnegative prices  $p_1, p_2, \dots, p_m$  for the items such that:

- (W1) All buyers are as happy as possible with their respective allocations, given the prices: for every  $i = 1, 2, \dots, k$ ,  $S_i \in \operatorname{argmax}_T \{v_i(T) - \sum_{j \in T} p_j\}$ .
- (W2) Feasibility:  $S_i \cap S_j = \emptyset$  for  $i \neq j$ .
- (W3) The market clears: for every  $j \in M$ ,  $j \in S_i$  for some  $i$ .<sup>1</sup>

Note that  $S_i$  might be the empty set, if the prices are high enough for (W1) to hold for player  $i$ . Also, property (W3) is crucial for the definition to be non-trivial (otherwise set  $p_j = +\infty$  for every  $j$ ).

Walrasian equilibria are remarkable: even though each player optimizes independently (modulo tie-breaking) and gets exactly what she wants, somehow the global feasibility constraint is respected.

### 8.1.2 The First Welfare Theorem

Recall from last lecture that the *social welfare* of an allocation  $S_1, \dots, S_k$  is defined as  $\sum_{i=1}^k v_i(S_i)$ . Walrasian equilibria automatically maximize the social welfare, a result known as the “First Welfare Theorem.”

---

<sup>1</sup>The most common definition of a Walrasian equilibrium asserts instead that an item  $j$  is not awarded to any player only if  $p_j = 0$ . With monotone valuations, there is no harm in insisting that every item is allocated.

**Theorem 8.2 (First Welfare Theorem).** *If the prices  $p_1, p_2, \dots, p_m$  and allocation  $S_1, S_2, \dots, S_k$  of items constitute a Walrasian equilibrium, then*

$$(S_1, S_2, \dots, S_k) \in \operatorname{argmax}_{(T_1, T_2, \dots, T_k)} \sum_{i=1}^k v_i(T_i),$$

where  $(T_1, \dots, T_k)$  ranges over all feasible allocations (with  $T_i \cap T_j = \emptyset$  for  $i \neq j$ ).

If one thinks of a Walrasian equilibrium as the natural outcome of a market, then Theorem 8.2 can be interpreted as saying “markets are efficient.”<sup>2</sup> There are many versions of the “First Welfare Theorem,” and all have this flavor.

*Proof.* Let  $(S_1^*, \dots, S_k^*)$  denote a welfare-maximizing feasible allocation. We can apply property (W1) of Walrasian equilibria to obtain

$$v_i(S_i) - \sum_{j \in S_i} p_j \geq v_i(S_i^*) - \sum_{j \in S_i^*} p_j$$

for each player  $i = 1, 2, \dots, k$ . Summing over  $i$ , we have

$$\sum_{i=1}^k v_i(S_i) - \sum_{i=1}^k \left( \sum_{j \in S_i} p_j \right) \geq \sum_{i=1}^k v_i(S_i^*) - \sum_{i=1}^k \left( \sum_{j \in S_i^*} p_j \right). \quad (8.1)$$

Properties (W2) and (W3) imply that the second term on the left-hand side of (8.1) equals the sum  $\sum_{j=1}^m p_j$  of all the item prices. Since  $(S_1^*, \dots, S_k^*)$  is a feasible allocation, each item is awarded at most once and hence the second term on the right-hand side is at most  $\sum_{j=1}^m p_j$ . Adding  $\sum_{j=1}^m p_j$  to both sides gives

$$\sum_{i=1}^k v_i(S_i) \geq \sum_{i=1}^k v_i(S_i^*),$$

which proves that the allocation  $(S_1, \dots, S_k)$  is also welfare-maximizing.  $\square$

---

<sup>2</sup>Needless to say, much blood and ink have been spilled over this interpretation over the past couple of centuries.

### 8.1.3 Existence of Walrasian Equilibria

The First Welfare Theorem says that Walrasian equilibria are great when they exist. But when do they exist?

**Example 8.3.** Suppose  $M$  contains only one item. Consider the allocation that awards the item to the player  $i$  with the highest value for it, and a price that is between player  $i$ 's value and the highest value of some other player (the second-highest overall). This is a Walrasian equilibrium: the price is low enough that bidder  $i$  prefers receiving the item to receiving nothing, and high enough that all the other bidders prefer the opposite. A simple case analysis shows that these are all of the Walrasian equilibria.

**Example 8.4.** Consider a market with two items,  $A$  and  $B$ . Suppose the valuation of the first player is

$$v_1(T) = \begin{cases} 3 & \text{for } T = \{A, B\} \\ 0 & \text{otherwise} \end{cases}$$

and that of the second player is

$$v_2(T) = \begin{cases} 2 & \text{for } T \neq \emptyset \\ 0 & \text{otherwise.} \end{cases}$$

The first bidder is called a “single-minded” or “AND” bidder, and is happy only if she gets both items. The second bidder is called a “unit-demand” or “OR” bidder, and effectively wants only one of the items.<sup>3</sup>

We claim that there is no Walrasian equilibrium in this market. From the First Welfare Theorem, we know that such an equilibrium must allocate the items to maximize the social welfare, which in this case means awarding both items to the first player. For the second player to be happy getting neither item, the price of each item must be at least 2. But then the first player pays 4 and has negative utility, and would prefer to receive nothing.

---

<sup>3</sup>More formally, a *unit-demand* valuation  $v$  is characterized by nonnegative values  $\{\alpha_j\}_{j \in M}$ , with  $v(S) = \max_{j \in S} \alpha_j$  for each  $S \subseteq M$ . Intuitively, a bidder with a unit-demand valuation throws away all her items except her favorite.

These examples suggest a natural question: under what conditions is a Walrasian equilibrium guaranteed to exist? There is a well-known literature on this question in economics (e.g. [74, 92, 111]); here are the highlights.

1. If every player's valuation  $v_i$  satisfies the “gross substitutes (GS)” condition, then a Walrasian equilibrium is guaranteed to exist. We won't need the precise definition of the GS condition in this lecture. GS valuations are closely related to weighted matroid rank functions, and hence are a subclass of the submodular valuations defined at the end of last lecture in Section 7.4.<sup>4</sup> A unit-demand (a.k.a. “OR”) valuation, like that of the second player in Example 8.4, satisfies the GS condition (corresponding to the 1-uniform matroid). It follows that single-minded (a.k.a. “AND”) valuations, like that of the first player in Example 8.4, do not in general satisfy the GS condition (otherwise the market in Example 8.4 would have a Walrasian equilibrium).
2. If  $\mathcal{V}$  is a class of valuations that contains all unit-demand valuations and also some valuation that violates the GS condition, then there is a market with valuations in  $\mathcal{V}$  that does not possess a Walrasian equilibrium.

These results imply that GS valuations are a maximal class of valuations subject to the guaranteed existence of Walrasian equilibria. These results do, however, leave open the possibility of guaranteed existence for classes  $\mathcal{V}$  that contain non-GS valuations but not all unit-demand valuations, and a number of recent papers in economics and operations research have pursued this direction (e.g. [11, 24, 25, 150]). All of the non-existence results in this line of work use explicit constructions, like in Example 8.4.

---

<sup>4</sup>A weighted matroid rank function  $f$  is defined using a matroid  $(E, \mathcal{I})$  and nonnegative weights on the elements  $E$ , with  $f(S)$  defined as the maximum weight of an independent set (i.e., a member of  $\mathcal{I}$ ) that lies entirely in the subset  $S \subseteq E$ .

## 8.2 Complexity Separations Imply Non-Existence of Walrasian Equilibria

### 8.2.1 Statement of Main Result

Next we describe a completely different approach to ruling out the existence of Walrasian equilibria, based on complexity theory rather than explicit constructions. The main result is the following.

**Theorem 8.5** (Roughgarden and Talgam-Cohen [138]). *Let  $\mathcal{V}$  denote a class of valuations. Suppose the welfare-maximization problem for  $\mathcal{V}$  does not reduce to the utility-maximization problem for  $\mathcal{V}$ . Then, there exists a market with all player valuations in  $\mathcal{V}$  that has no Walrasian equilibrium.*

In other words, a necessary condition for the guaranteed existence of Walrasian equilibria is that welfare-maximization is no harder than utility-maximization. This connects a purely economic question (when do equilibria exist?) to a purely algorithmic one.

To fill in some of the details in the statement of Theorem 8.5, by “does not reduce to,” we mean that there is no polynomial-time Turing reduction from the former problem to the latter. By “the welfare-maximization problem for  $\mathcal{V}$ ,” we mean the problem of, given player valuations  $v_1, \dots, v_k \in \mathcal{V}$ , computing an allocation that maximizes the social welfare  $\sum_{i=1}^k v_i(S_i)$ .<sup>5</sup> By “the utility-maximization problem for  $\mathcal{V}$ ,” we mean the problem of, given a valuation  $v \in \mathcal{V}$  and nonnegative prices  $p_1, \dots, p_m$ , computing a utility-maximizing bundle  $S \in \operatorname{argmax}_{T \subseteq M} \{v(T) - \sum_{j \in T} p_j\}$ .

The utility-maximization problem, which involves only one player, can generally only be easier than the multi-player welfare-maximization problem. Thus the two problems either have the same computational complexity, or welfare-maximization is strictly harder. Theorem 8.5 asserts that whenever the second case holds, Walrasian equilibria need not exist.

---

<sup>5</sup>For concreteness, think about the case where every valuation  $v_i$  has a succinct description and can be evaluated in polynomial time. Analogous results hold when an algorithm has only oracle access to the valuations.

### 8.2.2 Examples

Before proving Theorem 8.5, let's see how to apply it. For most natural valuation classes  $\mathcal{V}$ , a properly trained theoretical computer scientist can identify the complexity of the utility- and welfare-maximization problems in a matter of minutes.

**Example 8.6** (AND Valuations). Let  $\mathcal{V}_m$  denote the class of “AND” valuations for markets where  $|M| = m$ . That is, each  $v \in \mathcal{V}_m$  has the following form, for some  $\alpha \geq 0$  and  $T \subseteq M$ :

$$v(S) = \begin{cases} \alpha & \text{if } S \supseteq T \\ 0 & \text{otherwise.} \end{cases}$$

The utility-maximization problem for  $\mathcal{V}_m$  is trivial: for a single player with an AND valuation with parameters  $\alpha$  and  $T$ , the better of  $\emptyset$  or  $T$  is a utility-maximizing bundle. The welfare-maximization problem for  $\mathcal{V}_m$  is essentially set packing and is NP-hard (with  $m \rightarrow \infty$ ).<sup>6</sup> We conclude that the welfare-maximization problem for  $\mathcal{V}$  does not reduce to the utility-maximization problem for  $\mathcal{V}$  (unless  $P = NP$ ). Theorem 8.5 then implies that, assuming  $P \neq NP$ , there are markets with AND valuations that do not have any Walrasian equilibria.<sup>7</sup>

Of course, Example 8.4 already shows, without any complexity assumptions, that markets with AND bidders do not generally have

---

<sup>6</sup>For example, given an instance  $G = (V, E)$  of the INDEPENDENT SET problem, take  $M = E$ , make one player for each vertex  $i \in V$ , and give player  $i$  an AND valuation with parameters  $\alpha = 1$  and  $T$  equal to the edges that are incident to  $i$  in  $G$ .

<sup>7</sup>It probably seems weird to have a conditional result ruling out equilibrium existence. A conditional non-existence result can of course be made unconditional through an explicit example. A proof that the welfare-maximization problem for  $\mathcal{V}$  is NP-hard will generally suggest candidate markets to check for non-existence.

The following analogy may help: consider computationally tractable linear programming relaxations of NP-hard optimization problems. Conditional on  $P \neq NP$ , such relaxations cannot be exact (i.e., have no integrality gap) for all instances. NP-hardness proofs generally suggest instances that can be used to prove directly (and unconditionally) that a particular linear programming relaxation has an integrality gap.



Walrasian equilibria.<sup>8</sup> Our next example addresses a class of valuations for which the status of Walrasian equilibrium existence was not previously known.

**Example 8.7** (Capped Additive Valuations). A *capped additive* valuation  $v$  is parameterized by  $m + 1$  numbers  $c, \alpha_1, \alpha_2, \dots, \alpha_m$  and is defined as

$$v(S) = \min \left\{ c, \sum_{j \in S} \alpha_j \right\}.$$

The  $\alpha_j$ 's indicate each item's value, and  $c$  the "cap" on the maximum value that can be attained. Capped additive valuations were proposed in Lehmann *et al.* [100] as a natural subclass of submodular valuations, and have been studied previously from a welfare-maximization standpoint.

Let  $\mathcal{V}_{m,d}$  denote the class of capped additive valuations in markets with  $|M| = m$  and with  $c$  and  $\alpha_1, \dots, \alpha_m$  restricted to be positive integers between 1 and  $m^d$ . (Think of  $d$  as fixed and  $m \rightarrow \infty$ .) A Knapsack-type dynamic programming algorithm shows that the utility-maximization problem for  $\mathcal{V}_{m,d}$  can be solved in polynomial time (using that  $c$  and the  $\alpha_j$ 's are polynomially bounded). For  $d$  a sufficiently large constant, however, the welfare-maximization problem for  $\mathcal{V}_{m,d}$  is NP-hard (it includes the strongly NP-hard Bin Packing problem). Theorem 8.5 then implies that, assuming  $P \neq NP$ , there are markets with valuations in  $\mathcal{V}_{m,d}$  with no Walrasian equilibrium.

## 8.3 Proof of Theorem 8.5

### 8.3.1 The Plan

Here's the plan for proving Theorem 8.5. Fix a class  $\mathcal{V}$  of valuations, and assume that a Walrasian equilibrium exists in every market with player valuations in  $\mathcal{V}$ . We will show, in two steps, that the welfare-maximization problem for  $\mathcal{V}$  (polynomial-time Turing) reduces to the utility-maximization problem for  $\mathcal{V}$ .

---

<sup>8</sup>Replacing the OR bidder in Example 8.4 with an appropriate pair of AND bidders extends the example to markets with only AND bidders.

**Step 1:** The “fractional” version of the welfare-maximization problem for  $\mathcal{V}$  reduces to the utility-maximization problem for  $\mathcal{V}$ .

**Step 2:** A market admits a Walrasian equilibrium if and only if the fractional welfare-maximization problem has an optimal integral solution. (We’ll only need the “only if” direction.)

Since every market with valuations in  $\mathcal{V}$  admits a Walrasian equilibrium (by assumption), these two steps imply that the integral welfare-maximization problem reduces to utility-maximization.

### 8.3.2 Step 1: Fractional Welfare-Maximization Reduces to Utility-Maximization

This step is folklore, and appears for example in Nisan and Segal [121]. Consider the following linear program (often called the *configuration LP*), with one nonnegative variable  $x_{iS}$  for each player  $i$  and bundle  $S \subseteq 2^M$ :

$$\begin{aligned} \max \quad & \sum_{i=1}^k \sum_{S \subseteq M} v_i(S) x_{iS} \\ \text{s.t.} \quad & \sum_{i=1}^k \sum_{S \subseteq M: j \in S} x_{iS} \leq 1 \quad \text{for } j = 1, 2, \dots, m \\ & \sum_{S \subseteq M} x_{iS} = 1 \quad \text{for } i = 1, 2, \dots, k. \end{aligned}$$

The intended semantics are

$$x_{iS} = \begin{cases} 1 & \text{if } i \text{ gets the bundle } S \\ 0 & \text{otherwise.} \end{cases}$$

The first set of constraints enforces that each item is awarded only once (perhaps fractionally), and the second set enforces that every player receives one bundle (perhaps fractionally). Every feasible allocation induces a 0-1 feasible solution to this linear program according to the intended semantics, and the objective function value of this solution is exactly the social welfare of the allocation.

This linear program has an exponential (in  $m$ ) number of variables. The good news is that it has only a polynomial number of constraints.

This means that the dual linear program will have a polynomial number of variables and an exponential number of constraints, which is right in the wheelhouse of the ellipsoid method.

Precisely, the dual linear program is:

$$\begin{aligned} \min \quad & \sum_{i=1}^k u_i + \sum_{j=1}^m p_j \\ \text{s.t.} \quad & u_i + \sum_{j \in S} p_j \geq v_i(S) \quad \text{for all } i = 1, 2, \dots, k \text{ and } S \subseteq M \\ & p_j \geq 0 \quad \text{for } j = 1, 2, \dots, m, \end{aligned}$$

where  $u_i$  and  $p_j$  correspond to the primal constraints that bidder  $i$  receives one bundle and that item  $j$  is allocated at most once, respectively.

Recall that the ellipsoid method [93] can solve a linear program in time polynomial in the number of variables, as long as there is a polynomial-time *separation oracle* that can verify whether or not a given point is feasible and, if not, produce a violated constraint. For the dual linear program above, this separation oracle boils down to solving the following problem: for each player  $i = 1, 2, \dots, k$ , check that

$$u_i \geq \max_{S \subseteq M} \left[ v_i(S) - \sum_{j \in S} p_j \right].$$

But this reduces immediately to the utility-maximization problem for  $\mathcal{V}$ ! Thus the ellipsoid method can be used to solve the dual linear program to optimality, using a polynomial number of calls to a utility-maximization oracle. The optimal solution to the original fractional welfare-maximization problem can then be efficiently extracted from the optimal dual solution.<sup>9</sup>

---

<sup>9</sup>In more detail, consider the (polynomial number of) dual constraints generated by the ellipsoid method when solving the dual linear program. Form a reduced version of the original primal problem, retaining only the (polynomial number of) variables that correspond to this subset of dual constraints. Solve this polynomial-size reduced version of the primal linear program using your favorite polynomial-time linear programming algorithm.

### 8.3.3 Step 2: Walrasian Equilibria and Exact Linear Programming Relaxations

We now proceed with the second step, which is based on Bikhchandani and Mamer [12] and follows from strong linear programming duality. Recall from linear programming theory (see e.g. [39]) that a pair of primal and dual feasible solutions are both optimal if and only if the “complementary slackness” conditions hold.<sup>10</sup> These conditions assert that every non-zero decision variable in one of the linear programs corresponds to a tight constraint in the other. For our primal-dual pair of linear programs, these conditions are:

- (i)  $x_{iS} > 0$  implies that  $u_i = v_i(S) - \sum_{j \in S} p_j$  (i.e., only utility-maximizing bundles are used);
- (ii)  $p_j > 0$  implies that  $\sum_i \sum_{S: j \in S} x_{iS} = 1$  (i.e., item  $j$  is not fully sold only if it is worthless).

Comparing the definition of Walrasian equilibria (Definition 8.1) with conditions (i) and (ii), we see that a 0-1 primal feasible solution  $\mathbf{x}$  (corresponding to an allocation) and a dual solution  $\mathbf{p}$  (corresponding to item prices) constitute a Walrasian equilibrium if and only if the complementary slackness conditions hold (where  $u_i$  is understood to be set to  $\max_{S \subseteq M} v_i(S) - \sum_{j \in S} p_j$ ). Thus a Walrasian equilibrium exists if and only if there is a feasible 0-1 solution to the primal linear program and a feasible solution to the dual linear problem that satisfy the complementary slackness conditions, which in turn holds if and only if the primal linear program has an optimal 0-1 feasible solution.<sup>11</sup> We conclude that a Walrasian equilibrium exists if and only if the fractional welfare-maximization problem has an optimal integral solution. This completes the proof of Theorem 8.5.

---

<sup>10</sup>If you’ve never seen or have forgotten about complementary slackness, there’s no need to be afraid. To derive them, just write down the usual proof of weak LP duality (which is a chain of inequalities), and back out the conditions under which all the inequalities hold with equality.

<sup>11</sup>This argument re-proves the First Welfare Theorem (Theorem 8.2). It also proves the Second Welfare Theorem, which states that for every welfare-maximizing allocation, there exist prices that render it a Walrasian equilibrium—any optimal solution to the dual linear program furnishes such prices.

## 8.4 Beyond Walrasian Equilibria

For valuation classes  $\mathcal{V}$  that do not always possess Walrasian equilibria, is it possible to define a more general notion of “market-clearing prices” so that existence is guaranteed? For example, what if we use prices that are more complex than item prices? This section shows that complexity considerations provide an explanation of why interesting generalizations of Walrasian equilibria have been so hard to come by.

Consider a class  $\mathcal{V}$  of valuations, and a class  $\mathcal{P}$  of *pricing functions*. A pricing function, just like a valuation, is a function  $p : 2^M \rightarrow \mathbb{R}_+$  from bundles to nonnegative numbers. The item prices  $p_1, \dots, p_m$  used to define Walrasian equilibria correspond to additive pricing functions, with  $p(S) = \sum_{j \in S} p_j$ . The next definition articulates the appropriate generalization of Walrasian equilibria to more general classes of pricing functions.

**Definition 8.8** (Price Equilibrium). A *price equilibrium* (w.r.t. pricing functions  $\mathcal{P}$ ) is an allocation  $S_1, \dots, S_k$  of the items of  $M$  to the players and a pricing function  $p \in \mathcal{P}$  such that:

- (P1) All buyers are as happy as possible with their respective allocations, given the prices: for every  $i = 1, 2, \dots, k$ ,  $S_i \in \operatorname{argmax}_T \{v_i(T) - p(T)\}$ .
- (P2) Feasibility:  $S_i \cap S_j = \emptyset$  for  $i \neq j$ .
- (P3) Revenue maximizing, given the prices:  $(S_1, S_2, \dots, S_k) \in \operatorname{argmax}_{(T_1, T_2, \dots, T_k)} \{\sum_{i=1}^k p(T_i)\}$ .

Condition (P3) is the analog of the market-clearing condition (W3) in Definition 8.1. It is not enough to assert that all items are sold, because with a general pricing function, different ways of selling all of the items can lead to different amounts of revenue. Under conditions (P1)–(P3), the First Welfare Theorem (Theorem 8.2) still holds, with essentially the same proof, and so every price equilibrium maximizes the social welfare.

For which choices of valuations  $\mathcal{V}$  and pricing functions  $\mathcal{P}$  is Definition 8.8 interesting? Ideally, the following properties should hold.

1. Guaranteed existence: for every set  $M$  of items and valuations  $v_1, \dots, v_k \in \mathcal{V}$ , there exists a price equilibrium with respect to  $\mathcal{P}$ .
2. Efficient recognition: there is a polynomial-time algorithm for checking whether or not a given allocation and pricing function constitute a price equilibrium. This boils down to assuming that utility-maximization (with respect to  $\mathcal{V}$  and  $\mathcal{P}$ ) and revenue-maximization (with respect to  $\mathcal{P}$ ) are polynomial-time solvable problems (to check (W1) and (W3), respectively).
3. Markets with valuations in  $\mathcal{V}$  do not always have a Walrasian equilibrium. (Otherwise, why bother generalizing item prices?)

We can now see why there are no known natural choices of  $\mathcal{V}$  and  $\mathcal{P}$  that meet these three requirements. The first two requirements imply that the welfare-maximization problem belongs to  $\text{NP} \cap \text{co-NP}$ . To certify a lower bound of  $W^*$  on the maximum social welfare, one can exhibit an allocation with social welfare at least  $W^*$ . To certify an upper bound of  $W^*$ , one can exhibit a price equilibrium that has welfare at most  $W^*$ —this is well defined by the first condition, efficiently verifiable by the second condition, and correct by the First Welfare Theorem.

Problems in  $(\text{NP} \cap \text{co-NP}) \setminus \text{P}$  appear to be rare, especially in combinatorial optimization. The preceding paragraph gives a heuristic argument that interesting generalizations of Walrasian equilibria are possible only for valuation classes for which welfare-maximization is polynomial-time solvable. For every natural such class known, the linear programming relaxation in Section 8.3 has an optimal integral solution; in this sense, solving the configuration LP appears to be a “universal algorithm” for polynomial-time welfare-maximization. But the third requirement asserts that a Walrasian equilibrium does not always exist in markets with valuations in  $\mathcal{V}$  and so, by the second step of the proof of Theorem 8.5 (in Section 8.3.3), there are markets for which the configuration LP sometimes has only fractional optimal solutions.

The upshot is that interesting generalizations of Walrasian equilibria appear possible only for valuation classes where a non-standard

algorithm is necessary and sufficient to solve the welfare-maximization problem in polynomial time. It is not clear if there are any natural valuation classes for which this algorithmic barrier can be overcome.<sup>12</sup>

---

<sup>12</sup>See [138, Section 5.3.2] for an unnatural such class.

# 9

---

## The Borders of Border's Theorem

---

Border's theorem [16] is a famous result in auction theory about the design space of single-item auctions, and it provides an explicit linear description of the single-item auctions that are “feasible” in a certain sense. Despite the theorem's fame, there have been few generalizations of it. This lecture, based on joint work with Parikshit Gopalan and Noam Nisan [73], uses complexity theory to explain why: if there *were* significant generalizations of Border's theorem, the polynomial hierarchy would collapse!

### 9.1 Optimal Single-Item Auctions

#### 9.1.1 The Basics of Single-Item Auctions

Single-item auctions have made brief appearances in previous lectures; let's now study the classic model, due to Vickrey [154], in earnest. There is a single seller of a single item. There are  $n$  bidders, and each bidder  $i$  has a valuation  $v_i$  for the item (her maximum willingness to pay). Valuations are *private*, meaning that  $v_i$  is known a priori to bidder  $i$  but not to the seller or the other bidders. Each bidder wants to maximize the value obtained from the auction ( $v_i$  if she wins, 0 otherwise) minus



the price she has to pay. In the presence of randomization (either in the input or internal to the auction), we assume that bidders are risk-neutral, meaning they act to maximize their expected utility.

This lecture is our only one on the classical *Bayesian* model of auctions, which can be viewed as a form of average-case analysis. The key assumption is that each valuation  $v_i$  is drawn from a distribution  $F_i$  that is known to the seller and possibly the other bidders. The actual realization  $v_i$  remains unknown to everybody other than bidder  $i$ . For simplicity we'll work with discrete distributions, and let  $V_i$  denote the support of  $F_i$  and  $f_i(v_i)$  the probability that bidder  $i$ 's valuation is  $v_i \in V_i$ . Typical examples include (discretized versions of) the uniform distribution, the lognormal distribution, the exponential distribution, and power-law distributions. We also assume that bidders' valuations are stochastically independent.

When economists speak of an "optimal auction," they usually mean the auction that maximizes the seller's expected revenue with respect to a known prior distribution.<sup>1</sup> Before identifying optimal auctions, we need to formally define the design space. The auction designer needs to decide who wins and how much they pay. Thus the designer must define two (possibly randomized) functions of the bid vector  $\vec{b}$ : an *allocation rule*  $\vec{x}(\vec{b})$  which determines which bidder wins the item, where  $x_i = 1$  and if  $i$  wins and  $x_i = 0$  otherwise, and a *payment rule*  $\vec{p}(\vec{b})$  where  $p_i$  is how much  $i$  pays. We impose the constraint that whenever bidder  $i$  bids  $b_i$ , the expected payment  $\mathbf{E}[p_i(\vec{b})]$  of the bidder is at most  $b_i$  times the probability  $x_i(\vec{b})$  that she wins. (The randomization is over the bids by the other bidders and any randomness internal to the auction.) This participation constraint ensures that a bidder who does not overbid will obtain nonnegative expected utility from the auction. (Without it, an auction could just charge  $+\infty$  to every bidder.) The *revenue* of an auction on the bid vector  $\vec{b}$  is  $\sum_{i=1}^n p_i(\vec{b})$ .

---

<sup>1</sup>One advantage of assuming a distribution over inputs is that there is an unequivocal way to compare the performance of different auctions (by their expected revenues), and hence an unequivocal way to define an optimal auction. One auction generally earns more revenue than another on some inputs and less on others, so in the absence of a prior distribution, it's not clear which one to prefer.

For example, in the *Vickrey* or *second-price auction*, the allocation rule awards the item to the highest bidder, and the payment rule charges the second-highest bid. This auction is (*dominant-strategy*) *truthful*, meaning that for each bidder, truthful bidding (i.e., setting  $b_i = v_i$ ) is a *dominant strategy* that maximizes her utility no matter what the other bidders do. With such a truthful auction, there is no need to assume that the distributions  $F_1, \dots, F_n$  are known to the bidders. The beauty of the Vickrey auction is that it delegates underbidding to the auctioneer, who determines the optimal bid for the winner on their behalf.

A *first-price auction* has the same allocation rule as a second-price auction (give the item to the highest bidder), but the payment rule charges the winner her bid. Bidding truthfully in a first-price auction guarantees zero utility, so strategic bidders will underbid. Because bidders do not have dominant strategies—the optimal amount to underbid depends on the bids of the others—it is non-trivial to reason about the outcome of a first-price auction. The traditional solution is to assume that the distributions  $F_1, \dots, F_n$  are known in advance to the bidders, and to consider Bayes-Nash equilibria. Formally, a *strategy* of a bidder  $i$  in a first-price auction is a predetermined plan for bidding—a function  $b_i(\cdot)$  that maps a valuation  $v_i$  to a bid  $b_i(v_i)$  (or a distribution over bids). The semantics are: “when my valuation is  $v_i$ , I will bid  $b_i(v_i)$ .” We assume that bidders’ strategies are common knowledge, with bidders’ valuations (and hence induced bids) private as usual. A strategy profile  $b_1(\cdot), \dots, b_n(\cdot)$  is a *Bayes-Nash equilibrium* if every bidder always bids optimally given her information—if for every bidder  $i$  and every valuation  $v_i$ , the bid  $b_i(v_i)$  maximizes  $i$ ’s expected utility, where the expectation is with respect to the distribution over the bids of other bidders induced by  $F_1, \dots, F_n$  and their bidding strategies.<sup>2</sup> Note that the set of Bayes-Nash equilibria of an auction generally depends on the prior distributions  $F_1, \dots, F_n$ .

An auction is called *Bayesian incentive compatible (BIC)* if truthful bidding (with  $b_i(v_i) = v_i$  for all  $i$  and  $v_i$ ) is a Bayes-Nash equilibrium.

---

<sup>2</sup>Straightforward exercise: if there are  $n$  bidders with valuations drawn i.i.d. from the uniform distribution on  $[0, 1]$ , then setting  $b_i(v_i) = \frac{n-1}{n} \cdot v_i$  for every  $i$  and  $v_i$  yields a Bayes-Nash equilibrium.

That is, as a bidder, if all other bidders bid truthfully, then you also want to bid truthfully. A second-price auction is BIC, while a first-price auction is not.<sup>3</sup> However, for every choice of  $F_1, \dots, F_n$ , there is a BIC auction that is equivalent to the first-price auction. Specifically: given bids  $a_1, \dots, a_n$ , implement the outcome of the first-price auction with bids  $b_1(a_1), \dots, b_n(a_n)$ , where  $b_1(\cdot), \dots, b_n(\cdot)$  denotes a Bayes-Nash equilibrium of the first-price auction (with prior distributions  $F_1, \dots, F_n$ ). Intuitively, this auction makes the following pact with each bidder: “you promise to tell me your true valuation, and I promise to bid on your behalf as you would in a Bayes-Nash equilibrium.” More generally, this simulation argument shows that for *every* auction  $A$ , distributions  $F_1, \dots, F_n$ , and Bayes-Nash equilibrium of  $A$  (w.r.t.  $F_1, \dots, F_n$ ), there is a BIC auction  $A'$  whose (truthful) outcome (and hence expected revenue) matches that of the chosen Bayes-Nash equilibrium of  $A$ . This result is known as the *Revelation Principle*. This principle implies that, to identify an optimal auction, there is no loss of generality in restricting to BIC auctions.<sup>4</sup>

### 9.1.2 Optimal Auctions

In optimal auction design, the goal is to identify an expected revenue-maximizing auction, as a function of the prior distributions  $F_1, \dots, F_n$ . For example, suppose that  $n = 1$ , and we restrict attention to truthful auctions. The only truthful auctions are take-it-or-leave-it offers (or a randomization over such offers). That is, the selling price must be independent of the bidder’s bid, as any dependence would result in opportunities for the bidder to game the auction. The optimal truthful auction is then the take-it-or-leave-it offer at the price  $r$  that maximizes

$$\underbrace{r}_{\text{revenue of a sale}} \cdot \underbrace{(1 - F(r))}_{\text{probability of a sale}},$$

<sup>3</sup>The second-price auction is in fact *dominant-strategy incentive compatible (DSIC)*—truthful bidding is a dominant strategy for every bidder, not merely a Bayes-Nash equilibrium.

<sup>4</sup>Of course, non-BIC auctions like first-price auctions are still useful in practice. For example, the description of the first-price auction does not depend on bidders’ valuation distributions  $F_1, \dots, F_n$  and can be deployed without knowledge of them. This is not the case for the simulating auction.

where  $F$  denotes the bidder's valuation distribution. Given a distribution  $F$ , it is usually a simple matter to solve for the best  $r$ . An optimal offer price is called a *monopoly price* of the distribution  $F$ . For example, if  $F$  is the uniform distribution on  $[0, 1]$ , then the monopoly price is  $\frac{1}{2}$ .

Myerson [116] gave a complete solution to the optimal single-item auction design problem, in the form of a generic compiler that takes as input prior distributions  $F_1, \dots, F_n$  and outputs a closed-form description of the optimal auction for  $F_1, \dots, F_n$ . The optimal auction is particularly easy to interpret in the symmetric case, in which bidders' valuations are drawn i.i.d. from a common distribution  $F$ . Here, the optimal auction is simply a second-price auction with a reserve price  $r$  equal to the monopoly price of  $F$  (i.e., an eBay auction with a suitably chosen opening bid).<sup>5,6</sup> For example, with any number  $n$  of bidders with valuations drawn i.i.d. from the uniform distribution on  $[0, 1]$ , the optimal single-item auction is a second-price auction with a reserve price of  $\frac{1}{2}$ . This is a pretty amazing confluence of theory and practice—we optimized over the space of all imaginable auctions (which includes some very strange specimens), and discovered that the theoretically optimal auction format is one that is already in widespread use!<sup>7</sup>

Myerson's theory of optimal auctions extends to the asymmetric case where bidders have different distributions (where the optimal auction is no longer so simple), and also well beyond single-item auctions.<sup>8</sup> The books by Hartline [77] and the author [136, Lectures 3 and 5] describe this theory from a computer science perspective.

---

<sup>5</sup>Intuitively, a reserve price of  $r$  acts as an extra bid of  $r$  submitted by the seller. In a second-price auction with a reserve price, the winner is the highest bidder who clears the reserve (if any). The winner (if any) pays either the reserve price or the second-highest bid, whichever is higher.

<sup>6</sup>Technically, this statement holds under a mild "regularity" condition on the distribution  $F$ , which holds for all of the most common parametric distributions.

<sup>7</sup>In particular, there is always an optimal auction in which truthful bidding is a dominant strategy (as opposed to merely being a BIC auction). This is also true in the asymmetric case.

<sup>8</sup>The theory applies more generally to "single-parameter problems." These include problems in which in each outcome a bidder is either a "winner" or a "loser" (with multiple winners allowed), and each bidder  $i$  has a private valuation  $v_i$  for winning (and value 0 for losing).

## 9.2 Border's Theorem

### 9.2.1 Context

Border's theorem identifies a tractable description of *all* BIC single-item auctions, in the form of a polytope in polynomially many variables. (See Section 9.1.1 for the definition of a BIC auction.) This goal is in some sense more ambitious than merely identifying the optimal auction; with this tractable description in hand, one can efficiently compute the optimal auction for any given set  $F_1, \dots, F_n$  of prior distributions.

Economists are interested in Border's theorem because it can be used to extend the reach of Myerson's optimal auction theory (Section 9.1.2) to more general settings, such as the case of risk-averse bidders studied by Maskin and Riley [106]. Matthews [107] conjectured the precise result that was proved by Border [16]. Computer scientists have used Border's theorem for orthogonal extensions to Myerson's theory, like computationally tractable descriptions of the expected-revenue maximizing auction in settings with multiple non-identical items [3, 21]. While there is no hope of deriving a closed-form solution to the optimal auction design problem with risk-averse bidders or with multiple items, Border's theorem at least enables an efficient algorithm for computing a description of an optimal auction (given descriptions of the prior distributions).

### 9.2.2 An Exponential-Size Linear Program

As a lead-in to Border's theorem, we show how to formulate the space of BIC single-item auctions as an (extremely big) linear program. The decision variables of the linear program encode the allocation and payment rules of the auction (assuming truthful bidding, as appropriate for BIC auctions). There is one variable  $x_i(\vec{v}) \in [0, 1]$  that describes the probability (over any randomization in the auction) that bidder  $i$  wins the item when bidders' valuations (and hence bids) are  $\vec{v}$ . Similarly,  $p_i(\vec{v}) \in \mathbb{R}_+$  denotes the expected payment made by bidder  $i$  when bidders' valuations are  $\vec{v}$ .

Before describing the linear program, we need some odd but useful notation (which is standard in game theory and microeconomics).

**Some Notation**

For an  $n$ -vector  $\vec{z}$  and a coordinate  $i \in [n]$ , let  $\vec{z}_{-i}$  denote the  $(n - 1)$ -vector obtained by removing the  $i$ th component from  $\vec{z}$ . We also identify  $(z_i, \vec{z}_{-i})$  with  $\vec{z}$ .

Also, recall that  $V_i$  denotes the possible valuations of bidder  $i$ , and that we assume that this set is finite.

Our linear program will have three sets of constraints. The first set enforces the property that truthful bidding is in fact a Bayes-Nash equilibrium (as required for a BIC auction). For every bidder  $i$ , possible valuation  $v_i \in V_i$  for  $i$ , and possible false bid  $v'_i \in V_i$ ,

$$\begin{aligned}
 & \underbrace{v_i \cdot \mathbf{E}_{\vec{v}_{-i} \sim \vec{F}_{-i}} [x_i(\vec{v})] - \mathbf{E}_{\vec{v}_{-i} \sim \vec{F}_{-i}} [p_i(\vec{v})]}_{\text{expected utility of truthful bid } v_i} \\
 & \geq \underbrace{v_i \cdot \mathbf{E}_{\vec{v}_{-i} \sim \vec{F}_{-i}} [x_i(v'_i, \vec{v}_{-i})] - \mathbf{E}_{\vec{v}_{-i} \sim \vec{F}_{-i}} [p_i(v'_i, \vec{v}_{-i})]}_{\text{expected utility of false bid } v'_i}. \quad (9.1)
 \end{aligned}$$

The expectation is over both the randomness in  $\vec{v}_{-i}$  and internal to the auction. Each of the expectations in (9.1) expands to a sum over all possible  $\vec{v}_{-i} \in \vec{V}_{-i}$ , weighted by the probability  $\prod_{j \neq i} f_j(v_j)$ . Because all of the  $f_j(v_j)$ 's are numbers known in advance, each of these constraints is linear (in the  $x_i(\vec{v})$ 's and  $p_i(\vec{v})$ 's).

The second set of constraints encode the participation constraints from Section 9.1.1, also known as the *interim individually rational (IIR)* constraints. For every bidder  $i$  and possible valuation  $v_i \in V_i$ ,

$$v_i \cdot \mathbf{E}_{\vec{v}_{-i} \sim \vec{F}_{-i}} [x_i(\vec{v})] - \mathbf{E}_{\vec{v}_{-i} \sim \vec{F}_{-i}} [p_i(\vec{v})] \geq 0. \quad (9.2)$$

The final set of constraints assert that, with probability 1, the item is sold to at most one bidder: for every  $\vec{v} \in \vec{V}$ ,

$$\sum_{i=1}^n x_i(\vec{v}) \leq 1. \quad (9.3)$$

By construction, feasible solutions to the linear system (9.1)–(9.3) correspond to the allocation and payment rules of BIC auctions with respect to the distributions  $F_1, \dots, F_n$ . This linear program has an exponential number of variables and constraints, and is not immediately useful.

### 9.2.3 Reducing the Dimension with Interim Allocation Rules

Is it possible to re-express the allocation and payment rules of BIC auctions with a small number of decision variables? Looking at the constraints (9.1) and (9.2), a natural idea is to use only the decision variables  $\{y_i(v_i)\}_{i \in [n], v_i \in V_i}$  and  $\{q_i(v_i)\}_{i \in [n], v_i \in V_i}$ , with the intended semantics that

$$y_i(v_i) = \mathbf{E}_{\vec{v}_{-i}}[x_i(v_i, \vec{v}_{-i})] \quad \text{and} \quad q_i(v_i) = \mathbf{E}_{\vec{v}_{-i}}[p_i(v_i, \vec{v}_{-i})].$$

In other words,  $y_i(v_i)$  is the probability that bidder  $i$  wins when she bids  $v_i$ , and  $q_i(v_i)$  is the expected amount that she pays; these were the only quantities that actually mattered in (9.1) and (9.2). (As usual, the expectation is over both the randomness in  $\vec{v}_{-i}$  and internal to the auction.) In auction theory, the  $y_i(v_i)$ 's are called an *interim allocation rule*, the  $q_i(v_i)$ 's an *interim payment rule*.<sup>9</sup>

There are only  $2 \sum_{i=1}^n |V_i|$  such decision variables, far fewer than the  $2 \prod_{i=1}^n |V_i|$  variables in (9.1)–(9.3). We'll think of the  $|V_i|$ 's (and hence the number of decision variables) as polynomially bounded. For example,  $V_i$  could be the multiples of some small  $\epsilon$  that lie in some bounded range like  $[0, 1]$ .

We can then express the BIC constraints (9.1) in terms of this smaller set of variables by

$$\underbrace{v_i \cdot y_i(v_i) - q_i(v_i)}_{\text{expected utility of truthful bid } v_i} \geq \underbrace{v_i \cdot y_i(v'_i) - q_i(v'_i)}_{\text{expected utility of false bid } v'_i} \quad (9.4)$$

for every bidder  $i$  and  $v_i, v'_i \in V_i$ . Similarly, the IIR constraints (9.2) become

$$v_i \cdot y_i(v_i) - q_i(v_i) \geq 0 \quad (9.5)$$

for every bidder  $i$  and  $v_i \in V_i$ .

Just one problem. What about the feasibility constraints (9.3), which reference the individual  $x_i(\vec{v})$ 's and not merely their expectations?

<sup>9</sup>Auction theory generally thinks about three informational scenarios: *ex ante*, where each bidder knows the prior distributions but not even her own valuation; *interim*, where each bidder knows her own valuation but not those of the others; and *ex post*, where all of the bidders know everybody's valuation. Bidders typically choose their bids at the interim stage.

The next definition articulates what feasibility means for an interim allocation rule.

**Definition 9.1** (Feasible Interim Allocation Rule). An interim allocation rule  $\{y_i(v_i)\}_{i \in [n], v_i \in V_i}$  is *feasible* if there exist nonnegative values for  $\{x_i(\vec{v})\}_{i \in [n], \vec{v} \in \vec{V}}$  such that

$$\sum_{i=1}^n x_i(\vec{v}) \leq 1$$

for every  $\vec{v}$  (i.e., the  $x_i(\vec{v})$ 's constitute a feasible allocation rule), and

$$y_i(v_i) = \underbrace{\sum_{\vec{v}_{-i} \in \vec{V}_{-i}} \left( \prod_{j \neq i} f_j(v_j) \right) \cdot x_i(v_i, \vec{v}_{-i})}_{\mathbf{E}_{\vec{v}_{-i}}[x_i(v_i, \vec{v}_{-i})]}$$

for every  $i \in [n]$  and  $v_i \in V_i$  (i.e., the intended semantics are respected).

In other words, the feasible interim allocation rules are exactly the projections (onto the  $y_i(v_i)$ 's) of the feasible (ex post) allocation rules.

The big question is: how can we translate interim feasibility into our new, more economical vocabulary?<sup>10</sup> As we'll see, Border's theorem [16] provides a crisp and computationally useful solution.

### 9.2.4 Examples

To get a better feel for the issue of checking the feasibility of an interim allocation rule, let's consider a couple of examples. A necessary condition for interim feasibility is that the item is awarded to at most one bidder in expectation (over the randomness in the valuations and internal to the auction):

$$\sum_{i=1}^n \underbrace{\sum_{v_i \in V_i} f_i(v_i) y_i(v_i)}_{\Pr[i \text{ wins}]} \leq 1. \tag{9.6}$$

---

<sup>10</sup>In principle, we know this is possible. The feasible (ex post) allocation rules form a polytope, the projection of a polytope is again a polytope, and every polytope can be described by a finite number of linear inequalities. So the real question is whether or not there's a *computationally useful* description of interim feasibility.



**Table 9.1:** Certifying feasibility of an interim allocation rule is analogous to filling in the table entries while respecting constraints on the sums of certain subsets of entries.

$(\mathbf{v}_1, \mathbf{v}_2)$	$\mathbf{x}_1(\mathbf{v}_1, \mathbf{v}_2)$	$\mathbf{x}_2(\mathbf{v}_1, \mathbf{v}_2)$
(1, 1)		
(1, 2)		
(2, 1)		
(2, 2)		

**Table 9.2:** One solution to Example 9.2.

$(\mathbf{v}_1, \mathbf{v}_2)$	$\mathbf{x}_1(\mathbf{v}_1, \mathbf{v}_2)$	$\mathbf{x}_2(\mathbf{v}_1, \mathbf{v}_2)$
(1, 1)	1	0
(1, 2)	0	1
(2, 1)	3/4	1/4
(2, 2)	1	0

Could this also be a sufficient condition? That is, is every interim allocation rule  $\{y_i(v_i)\}_{i \in [n], v_i \in V_i}$  that satisfies (9.6) induced by a bona fide (ex post) allocation rule?

**Example 9.2.** Suppose there are  $n = 2$  bidders. Assume that  $v_1, v_2$  are independent and each is equally likely to be 1 or 2. Consider the interim allocation rule given by

$$y_1(1) = \frac{1}{2}, \quad y_1(2) = \frac{7}{8}, \quad y_2(1) = \frac{1}{8}, \quad \text{and} \quad y_2(2) = \frac{1}{2}. \quad (9.7)$$

Since  $f_i(v) = \frac{1}{2}$  for all  $i = 1, 2$  and  $v = 1, 2$ , the necessary condition in (9.6) is satisfied. Can you find an (ex post) allocation rule that induces this interim rule? Answering this question is much like solving a Sudoku or KenKen puzzle—the goal is to fill in the table entries in Table 9.1 so that each row sums to at most 1 (for feasibility) and that the constraints (9.7) are satisfied. For example, the average of the top two entries in the first column of Table 9.1 should be  $y_1(1) = \frac{1}{2}$ . In this example, there are a number of such solutions; one is shown in Table 9.2. Thus, the given interim allocation rule is feasible.

**Example 9.3.** Suppose we change the interim allocation rule to

$$y_1(1) = \frac{1}{4}, \quad y_1(2) = \frac{7}{8}, \quad y_2(1) = \frac{1}{8}, \quad \text{and} \quad y_2(2) = \frac{3}{4}.$$

The necessary condition (9.6) remains satisfied. Now, however, the interim rule is not feasible. One way to see this is to note that  $y_1(2) = \frac{7}{8}$  implies that  $x_1(2, 2) \geq \frac{3}{4}$  and hence  $x_2(2, 2) \leq \frac{1}{4}$ . Similarly,  $y_2(2) = \frac{3}{4}$  implies that  $x_2(2, 2) \geq \frac{1}{2}$ , a contradictory constraint.

The first point of Examples 9.2 and 9.3 is that it is not trivial to check whether or not a given interim allocation rule is feasible—the problem corresponds to solving a big linear system of equations and inequalities. The second point is that (9.6) is not a sufficient condition for feasibility. In hindsight, trying to summarize the exponentially many ex post feasibility constraints (9.3) with a single interim constraint (9.6) seems naive. Is there a larger set of linear constraints—possibly an exponential number—that characterizes interim feasibility?

### 9.2.5 Border's Theorem

Border's theorem states that a collection of “obvious” necessary conditions for interim feasibility are also sufficient. To state these conditions, assume for notational convenience that the valuation sets  $V_1, \dots, V_n$  are disjoint.<sup>11</sup> Let  $\{x_i(\vec{v})\}_{i \in [n], \vec{v} \in \vec{V}}$  be a feasible (ex post) allocation rule and  $\{y_i(v_i)\}_{i \in [n], v_i \in V_i}$  the induced (feasible) interim allocation rule. Fix for each bidder  $i$  a set  $S_i \subseteq V_i$  of valuations. Call the valuations  $\cup_{i=1}^n S_i$  the *distinguished* valuations. Consider first the probability, over the random valuation profile  $\vec{v} \sim \vec{F}$  and any coin flips of the ex post allocation rule, that the winner of the auction (if any) has a distinguished valuation. By linearity of expectations, this probability can be expressed in terms of the interim allocation rule:

$$\sum_{i=1}^n \sum_{v_i \in S_i} f_i(v_i) y_i(v_i). \quad (9.8)$$

The expression (9.8) is linear in the  $y_i(v_i)$ 's.

<sup>11</sup>This is without loss of generality, since we can simply “tag” each valuation  $v_i \in V_i$  with the “name”  $i$  (i.e., view each  $v_i \in V_i$  as the set  $\{v_i, i\}$ ).

The second quantity we study is the probability, over  $\vec{v} \sim \vec{F}$ , that there is a bidder with a distinguished valuation. This has nothing to do with the allocation rule, and is a function of the prior distributions only:

$$1 - \prod_{i=1}^n \left( 1 - \sum_{v_i \in S_i} f_i(v_i) \right). \tag{9.9}$$

Because there can only be a winner with a distinguished valuation if there is a bidder with a distinguished valuation, the quantity in (9.8) can only be less than (9.9). Border's theorem asserts that these conditions, ranging over all choices of  $S_1 \subseteq V_1, \dots, S_n \subseteq V_n$ , are also sufficient for the feasibility of an interim allocation rule.

**Theorem 9.4** (Border's theorem [16]). *An interim allocation rule  $\{y_i(v_i)\}_{i \in [n], v_i \in V_i}$  is feasible if and only if for every choice  $S_1 \subseteq V_1, \dots, S_n \subseteq V_n$  of distinguished valuations,*

$$\sum_{i=1}^n \sum_{v_i \in S_i} f_i(v_i) y_i(v_i) \leq 1 - \prod_{i=1}^n \left( 1 - \sum_{v_i \in S_i} f_i(v_i) \right). \tag{9.10}$$

Border's theorem can be derived from the max-flow/min-cut theorem (following [17, 28]); we include the proof in Section 9.4 for completeness.

Border's theorem yields an explicit description as a linear system of the feasible interim allocation rules induced by BIC single-item auctions. To review, this linear system is

$$v_i \cdot y_i(v_i) - q_i(v_i) \geq v_i \cdot y_i(v'_i) - q_i(v'_i) \quad \forall i \text{ and } v_i, v'_i \in V_i \tag{9.11}$$

$$v_i \cdot y_i(v_i) - q_i(v_i) \geq 0 \quad \forall i \text{ and } v_i \in V_i \tag{9.12}$$

$$\sum_{i=1}^n \sum_{v_i \in S_i} f_i(v_i) y_i(v_i) \leq 1 - \prod_{i=1}^n \left( 1 - \sum_{v_i \in S_i} f_i(v_i) \right) \quad \forall S_1 \subseteq V_1, \dots, S_n \subseteq V_n. \tag{9.13}$$

For example, optimizing the objective function

$$\max \sum_{i=1}^n f_i(v_i) \cdot q_i(v_i) \tag{9.14}$$

over the linear system (9.11)–(9.13) computes the expected revenue of an optimal BIC single-item auction for the distributions  $F_1, \dots, F_n$ .

The linear system (9.11)–(9.13) has only a polynomial number of variables (assuming the  $|V_i|$ 's are polynomially bounded), but it does have an exponential number of constraints of the form (9.13). One solution is to use the ellipsoid method, as the linear system does admit a polynomial-time separation oracle [3, 21].<sup>12</sup> Alternatively, Alaei *et al.* [3] provide a polynomial-size extended formulation of the polytope of feasible interim allocation rules (with a polynomial number of additional decision variables and only polynomially many constraints). In any case, we conclude that there is a computationally tractable description of the feasible interim allocation rules of BIC single-item auctions.

### 9.3 Beyond Single-Item Auctions: A Complexity-Theoretic Barrier

Myerson's theory of optimal auctions (Section 9.1.2) extends beyond single-item auctions to all “single-parameter” settings (see footnote 8 for discussion and Section 9.3.1 for two examples). Can Border's theorem be likewise extended? There are analogs of Border's theorem in settings modestly more general than single-item auctions, including  $k$ -unit auctions with unit-demand bidders [3, 21, 28], and approximate versions of Border's theorem exist fairly generally [21, 22]. Can this state-of-the-art be improved upon? We next use complexity theory to develop evidence for a negative answer.

**Theorem 9.5** (Gopalan *et al.* [73]). *(Informal) There is no exact Border's-type theorem for settings significantly more general than the known special cases (unless PH collapses).*

We proceed to defining what we mean by “significantly more general” and a “Border's-type theorem.”

#### 9.3.1 Two Example Settings

The formal version of Theorem 9.5 conditionally rules out “Border's-type theorems” for several specific settings that are representative of what a

---

<sup>12</sup>This is not immediately obvious, as the max-flow/min-cut argument in Section 9.4 involves an exponential-size graph.

9.3. *Beyond Single-Item Auctions: A Complexity-Theoretic Barrier* 153

more general version of Border's theorem might cover. We mention two of these here (more are in [73]).

In a *public project* problem, there is a binary decision to make: whether or not to undertake a costly project (like building a new school). Each bidder  $i$  has a private valuation  $v_i$  for the outcome where the project is built, and valuation 0 for the outcome where it is not. If the project is built, then everyone can use it. In this setting, feasibility means that all bidders receive the same allocation:  $x_1(\vec{v}) = x_2(\vec{v}) = \dots = x_n(\vec{v}) \in [0, 1]$  for every valuation profile  $\vec{v}$ .

In a *matching* problem, there is a set  $M$  of items, and each bidder is only interested in receiving a specific pair  $j, \ell \in M$  of items. (Cf., the AND bidders of the preceding lecture.) For each bidder, the corresponding pair of items is common knowledge, while the bidder's valuation for the pair is private as usual. Feasible outcomes correspond to (distributions over) matchings in the graph with vertices  $M$  and edges given by bidders' desired pairs.

The public project and matching problems are both “single-parameter” problems (i.e., each bidder has only one private parameter). As such, Myerson's optimal auction theory (Section 9.1.2) can be used to characterize the expected revenue-maximizing auction. Do these settings also admit analogs of Border's theorem?

### 9.3.2 Border's-Type Theorems

What do we actually mean by a “Border's-type theorem?” Because we aim to prove impossibility results, we should adopt a definition that is as permissive as possible. Border's theorem (Theorem 9.4) gives a characterization of the feasible interim allocation rules of a single-item auction as the solutions to a finite system of linear inequalities. This by itself is not impressive—the set is a polytope, and as such is guaranteed to have such a characterization. The appeal of Border's theorem is that the characterization uses only the “nice” linear inequalities in (9.10). Our “niceness” requirement is that the characterization use only linear inequalities that can be efficiently recognized and tested. This is a weak necessary condition for such a characterization to be computationally useful.

**Definition 9.6** (Border's-Type Theorem). A *Border's-type theorem* holds for an auction design setting if, for every instance of the setting (specifying the number of bidders and their prior distributions, etc.), there is a system of linear inequalities such that the following properties hold.

1. (Characterization) The feasible solutions of the linear system are precisely the feasible interim allocation rules of the instance.
2. (Efficient recognition) There is a polynomial-time algorithm that can decide whether or not a given linear inequality (described as a list of coefficients) belongs to the linear system.
3. (Efficient testing) The bit complexity of each linear inequality is polynomial in the description of the instance. (The number of inequalities can be exponential.)

For example, consider the original Border's theorem, for single-item auctions (Theorem 9.4). The recognition problem is straightforward: the left-side of (9.10) encodes the  $S_i$ 's, from which the right-hand side can be computed and checked in polynomial time. It is also evident that every inequality in (9.10) has a polynomial-length description.<sup>13</sup>

### 9.3.3 Consequences of a Border's-Type Theorem

The high-level idea behind the proof of Theorem 9.5 is to show that a Border's-type theorem puts a certain computational problem low in the polynomial hierarchy, and then to show that this problem is #P-hard for the public project and matching settings defined in Section 9.3.1.<sup>14</sup> The

---

<sup>13</sup>The characterization in Theorem 9.4 and the extensions in [3, 21, 28] have additional features not required or implied by Definition 9.6, such as a polynomial-time separation oracle (and even a compact extended formulation in the single-item case [3]). The impossibility results in Section 9.3.4 rule out analogs of Border's theorem that merely satisfy Definition 9.6, let alone these stronger properties.

<sup>14</sup>Recall that Toda's theorem [152] implies that a #P-hard problem is contained in the polynomial hierarchy only if PH collapses.

computational problem is: given a description of an instance (including the prior distributions), compute the maximum-possible expected revenue that can be obtained by a feasible and BIC auction.<sup>15</sup>

What use is a Border's-type theorem? For starters, it implies that the problem of testing the feasibility of an interim allocation rule is in  $\text{co-NP}$ . To prove the infeasibility of such a rule, one simply exhibits an inequality of the characterizing linear system that the rule fails to satisfy. Verifying this failure reduces to the recognition and testing problems, which by Definition 9.6 are polynomial-time solvable.

**Proposition 9.7.** *If a Border's-type theorem holds for an auction design setting, then the membership problem for the polytope of feasible interim allocation rules belongs to  $\text{co-NP}$ .*

Combining Proposition 9.7 with the ellipsoid method puts the problem of computing the maximum-possible expected revenue in  $\text{P}^{\text{NP}}$ .

**Theorem 9.8.** *If a Border's-type theorem holds for an auction design setting, then the maximum expected revenue of a feasible BIC auction can be computed in  $\text{P}^{\text{NP}}$ .*

*Proof.* We compute the optimal expected revenue of a BIC auction via linear programming, as follows. The decision variables are the same  $y_i(v_i)$ 's and  $q_i(v_i)$ 's as in (9.11)–(9.13), and we retain the BIC constraints (9.11) and the IIR constraints (9.12). By assumption, we can replace the single-item interim feasibility constraints (9.13) with a linear system that satisfies the properties of Definition 9.6. The maximum expected revenue of a feasible BIC auction can then be computed by optimizing a linear objective function (in the  $q_i(v_i)$ 's, as in (9.14)) subject to these constraints. Using the ellipsoid method [93], this can be accomplished with a polynomial number of invocations of a separation oracle (which either verifies feasibility or exhibits a violated constraint). Proposition 9.7 implies that we can implement this separation oracle

---

<sup>15</sup>Sanity check: this problem turns out to be polynomial-time solvable in the setting of single-item auctions [73].

in co-NP, and thus compute the maximum expected revenue of a BIC auction in  $P^{NP}$ .<sup>16</sup>  $\square$

### 9.3.4 Impossibility Results from Computational Intractability

Theorem 9.8 concerns the problem of computing the maximum expected revenue of a feasible BIC auction, given a description of an instance. It is easy to classify the complexity of this problem in the public project and matching settings introduced in Section 9.3.1 (and several other settings, see [73]).

**Proposition 9.9.** *Computing the maximum expected revenue of a feasible BIC auction of a public project instance is a #P-hard problem.*

Proposition 9.9 is a straightforward reduction from the #P-hard problem of computing the number of feasible solutions to an instance of the KNAPSACK problem.<sup>17</sup>

**Proposition 9.10.** *Computing the maximum expected revenue of a feasible BIC auction of a matching instance is a #P-hard problem.*

Proposition 9.10 is a straightforward reduction from the #P-hard PERMANENT problem.

We reiterate that Myerson's optimal auction theory applies to the public project and matching settings, and in particular gives a polynomial-time algorithm that outputs a description of an optimal auction (for given prior distributions). Moreover, the optimal auction can be implemented as a polynomial-time algorithm. Thus it's not hard to figure out what the optimal auction is, nor to implement it—what's hard is figuring out exactly how much revenue it makes on average!

---

<sup>16</sup>One detail: Proposition 9.7 only promises solutions to the “yes/no” question of feasibility, while a separation oracle needs to produce a violated constraint when given an infeasible point. But under mild conditions (easily satisfied here), an algorithm for the former problem can be used to solve the latter problem as well [144, P.189].

<sup>17</sup>An aside for aficionados of the analysis of Boolean functions: Proposition 9.9 is essentially equivalent to the #P-hardness of checking whether or not given Chow parameters can be realized by some bounded function on the hypercube. See [73] for more details on the surprisingly strong correspondence between Myerson's optimal auction theory (in the context of public projects) and the analysis of Boolean functions.






Figure 9.1: The max-flow/min-cut proof of Border's theorem.

Combining Theorem 9.8 with Propositions 9.9 and 9.10 gives the following corollaries, which indicate that there is no Border's-type theorem significantly more general than the ones already known.

**Corollary 9.11.** *If  $\#P \not\subseteq PH$ , then there is no Border's-type theorem for the setting of public projects.*

**Corollary 9.12.** *If  $\#P \not\subseteq PH$ , then there is no Border's-type theorem for the matching setting.*

### 9.4 Appendix: A Combinatorial Proof of Border's Theorem

*Proof.* (of Theorem 9.4) We have already argued the “only if” direction, and now prove the converse. The proof is by the max-flow/min-cut theorem—given the statement of the theorem and this hint, the proof writes itself.

Suppose the interim allocation rule  $\{y_i(v_i)\}_{i \in [n], v_i \in V_i}$  satisfies (9.10) for every  $S_1 \subseteq V_1, \dots, S_n \subseteq V_n$ . Form a four-layer  $s$ - $t$  directed flow network  $G$  as follows (Figure 9.1(a)). The first layer is the source  $s$ , the last the sink  $t$ . In the second layer  $X$ , vertices correspond to valuation profiles  $\vec{v}$ . We abuse notation and refer to vertices of  $X$  by the corresponding valuation profiles. There is an arc  $(s, \vec{v})$  for every  $\vec{v} \in X$ , with capacity  $\prod_{i=1}^n f_i(v_i)$ . Note that the total capacity of these edges is 1.

In the third layer  $Y$ , vertices correspond to winner-valuation pairs; there is also one additional “no winner” vertex. We use  $(i, v_i)$  to denote the vertex representing the event that bidder  $i$  wins the item and also has valuation  $v_i$ . For each  $i$  and  $v_i \in V_i$ , there is an arc  $((i, v_i), t)$  with capacity  $f_i(v_i)y_i(v_i)$ . There is also an arc from the “no winner” vertex to  $t$ , with capacity  $1 - \sum_{i=1}^n \sum_{v_i \in V_i} f_i(v_i)y_i(v_i)$ .<sup>18</sup>

Finally, each vertex  $\vec{v} \in X$  has  $n + 1$  outgoing arcs, all with infinite capacity, to the vertices  $(1, v_1), (2, v_2), \dots, (n, v_n)$  of  $Y$  and also to the “no winner” vertex.

By construction,  $s$ - $t$  flows of  $G$  with value 1 correspond to ex post allocation rules with induced interim allocation rule  $\{y_i(v_i)\}_{i \in [n], v_i \in V_i}$ , with  $x_i(\vec{v})$  equal to the amount of flow on the arc  $(\vec{v}, (i, v_i))$  times  $(\prod_{i=1}^n f_i(v_i))^{-1}$ .

To show that there exists a flow with value 1, it suffices to show that every  $s$ - $t$  cut has value at least 1 (by the max-flow/min-cut theorem). So fix an  $s$ - $t$  cut. Let this cut include the vertices  $A$  from  $X$  and  $B$  from  $Y$ . Note that all arcs from  $s$  to  $X \setminus A$  and from  $B$  to  $t$  are cut (Figure 9.1(b)). For each bidder  $i$ , define  $S_i \subseteq V_i$  as the possible valuations of  $i$  that are *not* represented among the valuation profiles in  $A$ . Then, for every valuation profile  $\vec{v}$  containing at least one distinguished valuation, the arc  $(s, \vec{v})$  is cut. The total capacity of these arcs is the right-hand side (9.9) of Border's condition.

Next, we can assume that every vertex of the form  $(i, v_i)$  with  $v_i \notin S_i$  is in  $B$ , as otherwise an (infinite-capacity) arc from  $A$  to  $Y \setminus B$  is cut. Similarly, unless  $A = \emptyset$ —in which case the cut has value at least 1 and we're done—we can assume that the “no winner” vertex lies in  $B$ . Thus, the only edges of the form  $((i, v_i), t)$  that are not cut involve a distinguished valuation  $v_i \in S_i$ . It follows that the total capacity of the cut edges incident to  $t$  is at least 1 minus the left-hand side (9.8) of Border's condition. Given our assumption that (9.8) is at most (9.9), this  $s$ - $t$  cut has value at least 1. This completes the proof of Border's theorem.  $\square$

<sup>18</sup>If  $\sum_{i=1}^n \sum_{v_i \in V_i} f_i(v_i)y_i(v_i) > 1$ , then the interim allocation rule is clearly infeasible (recall (9.6)). Alternatively, this would violate Border's condition for the choice  $S_i = V_i$  for all  $i$ .

# 10

---

## Tractable Relaxations of Nash Equilibria

---

### 10.1 Preamble

Much of this monograph is about impossibility results for the efficient computation of exact and approximate Nash equilibria. How should we respond to such rampant computational intractability? What should be the message to economists—should they change the way they do economic analysis in some way?<sup>1</sup>

One approach, familiar from coping with NP-hard problems, is to look for tractable special cases. For example, Solar Lecture 1 proved tractability results for two-player zero-sum games. Some interesting tractable generalizations of zero-sum games have been identified (see [23] for a recent example), and polynomial-time algorithms are also known for some relatively narrow classes of games (see e.g. [90]). Still, for the lion's share of games that we might care about, no polynomial-time algorithms for computing exact or approximate Nash equilibria are known.

---

<sup>1</sup>Recall the discussion in Section 1.1.7 of Solar Lecture 1: a critique of a widely used concept like the Nash equilibrium is not particularly helpful unless accompanied by a proposed alternative.

A different approach, which has been more fruitful, is to continue to work with general games and look for an *equilibrium concept* that is more computationally tractable than exact or approximate Nash equilibria. The equilibrium concepts that we'll consider—the correlated equilibrium and the coarse correlated equilibrium—were originally invented by game theorists, but computational complexity considerations are now shining a much brighter spotlight on them.

Where do these alternative equilibrium concepts come from? They arise quite naturally from the study of uncoupled dynamics, which we last saw in Solar Lecture 1.

## 10.2 Uncoupled Dynamics Revisited

Section 1.2 of Solar Lecture 1 introduced uncoupled dynamics in the context of two-player games. In this lecture we work with the analogous setup for a general number  $k$  of players. We use  $S_i$  to denote the (pure) strategies of player  $i$ ,  $s_i \in S_i$  a specific strategy,  $\sigma_i$  a mixed strategy,  $\vec{s}$  and  $\vec{\sigma}$  for profiles (i.e.,  $k$ -vectors) of pure and mixed strategies, and  $u_i(\vec{s})$  for player  $i$ 's payoff in the outcome  $\vec{s}$ .

### Uncoupled Dynamics ( $k$ -Player Version)

At each time step  $t = 1, 2, 3, \dots$ :

1. Each player  $i = 1, 2, \dots, k$  simultaneously chooses a mixed strategy  $\sigma_i^t$  over  $S_i$  as a function only of her own payoffs and the strategies chosen by players in the first  $t - 1$  time steps.
2. Every player observes all of the strategies  $\vec{\sigma}^t$  chosen at time  $t$ .

“Uncoupled” refers to the fact that each player initially knows only her own payoff function  $u_i(\cdot)$ , while “dynamics” means a process by which players learn how to play in a game.

One of the only positive algorithmic results that we've seen concerned *smooth fictitious play (SFP)*. The  $k$ -player version of SFP is as follows.

### Smooth Fictitious Play ( $k$ -Player Version)

**Given:** parameter family  $\{\eta^t \in [0, \infty) : t = 1, 2, 3, \dots\}$ .

At each time step  $t = 1, 2, 3, \dots$ :

1. Every player  $i$  simultaneously chooses the mixed strategy  $\sigma_i^t$  by playing each strategy  $s_i$  with probability proportional to  $e^{\eta^t \pi_i^t}$ , where  $\pi_i^t$  is the time-averaged expected payoff player  $i$  would have earned by playing  $s_i$  at every previous time step. Equivalently,  $\pi_i^t$  is the expected payoff of strategy  $s_i$  when the other players' strategies  $\vec{s}_{-i}$  are drawn from the joint distribution  $\frac{1}{t-1} \sum_{h=1}^{t-1} \vec{\sigma}_{-i}^h$ .<sup>2</sup>
2. Every player observes all of the strategies  $\vec{\sigma}^t$  chosen at time  $t$ .

A typical choice for the  $\eta_t$ 's is  $\eta_t \approx \sqrt{t}$ .

In Theorem 1.8 in Solar Lecture 1 we proved that, in an  $m \times n$  two-player zero-sum game, after  $O(\log(m+n)/\epsilon^2)$  time steps, the empirical distributions of the two players constitute an  $\epsilon$ -approximate Nash equilibrium.<sup>3</sup> An obvious question is: what is the outcome of a logarithmic number of rounds of smooth fictitious play in a non-zero-sum game? Our communication complexity lower bound in Solar Lectures 2 and 3 implies that it cannot in general be an  $\epsilon$ -approximate Nash equilibrium. Does it have some alternative economic meaning? The answer to this question turns out to be closely related to some classical game-theoretic equilibrium concepts, which we discuss next.

---

<sup>2</sup>Recall from last lecture that for an  $n$ -vector  $\vec{z}$  and a coordinate  $i \in [k]$ ,  $\vec{z}_{-i}$  denotes the  $(k-1)$ -vector obtained by removing the  $i$ th component from  $\vec{z}$ , and we identify  $(z_i, \vec{z}_{-i})$  with  $\vec{z}$ .

<sup>3</sup>Recall the proof idea: smooth fictitious play corresponds to running the vanishing-regret “exponential weights” algorithm (with reward vectors induced by the play of others), and in a two-player zero-sum game, the vanishing-regret guarantee (i.e., with time-averaged payoff at least that of the best fixed action in hindsight, up to  $o(1)$  error) implies the  $\epsilon$ -approximate Nash equilibrium condition.

## 10.3 Correlated and Coarse Correlated Equilibria

### 10.3.1 Correlated Equilibria

The correlated equilibrium is a well-known equilibrium concept defined by Aumann [7]. We define it, then explain the standard semantics, and then offer an example.<sup>4</sup>

**Definition 10.1** (Correlated Equilibrium). A joint distribution  $\rho$  on the set  $S_1 \times \cdots \times S_k$  of outcomes of a game is a *correlated equilibrium* if for every player  $i \in \{1, 2, \dots, k\}$ , strategy  $s_i \in S_i$ , and deviation  $s'_i \in S_i$ ,

$$\mathbf{E}_{\vec{s} \sim \rho}[u_i(\vec{s}) \mid s_i] \geq \mathbf{E}_{\vec{s} \sim \rho}[u_i(s'_i, \vec{s}_{-i}) \mid s_i]. \quad (10.1)$$

Importantly, the distribution  $\rho$  in Definition 10.1 need not be a product distribution; in this sense, the strategies chosen by the players are correlated. The Nash equilibria of a game correspond to the correlated equilibria that are product distributions.

The usual interpretation of a correlated equilibrium involves a trusted third party. The distribution  $\rho$  over outcomes is publicly known. The trusted third party samples an outcome  $\vec{s}$  according to  $\rho$ . For each player  $i = 1, 2, \dots, k$ , the trusted third party privately suggests the strategy  $s_i$  to  $i$ . The player  $i$  can follow the suggestion  $s_i$ , or not. At the time of decision making, a player  $i$  knows the distribution  $\rho$  and one component  $s_i$  of the realization  $\vec{s}$ , and accordingly has a posterior distribution on others' suggested strategies  $\vec{s}_{-i}$ . With these semantics, the correlated equilibrium condition (10.1) requires that every player maximizes her expected payoff by playing the suggested strategy  $s_i$ . The expectation is conditioned on  $i$ 's information— $\rho$  and  $s_i$ —and assumes that other players play their recommended strategies  $\vec{s}_{-i}$ .

Definition 10.1 is a bit of a mouthful. But you are intimately familiar with a good example of a correlated equilibrium that is not a mixed Nash equilibrium—a traffic light! Consider the following two-player game, with each matrix entry listing the payoffs of the row and column players in the corresponding outcome:

---

<sup>4</sup>This section draws from [136, Lecture 13].

	Stop	Go
Stop	0,0	0,1
Go	1,0	-5,-5

This game has two pure Nash equilibria, the outcomes (Stop, Go) and (Go, Stop). Define  $\rho$  by randomizing uniformly between these two Nash equilibria. This is not a product distribution over the game's four outcomes, so it cannot correspond to a Nash equilibrium of the game. It is, however, a correlated equilibrium.<sup>5</sup>

### 10.3.2 Coarse Correlated Equilibria

The outcome of smooth fictitious play in non-zero-sum games relates to a still more permissive equilibrium concept, the *coarse correlated equilibrium*, which was first studied by Moulin and Vial [115].


**Definition 10.2** (Coarse Correlated Equilibrium). A joint distribution  $\rho$  on the set  $S_1 \times \cdots \times S_k$  of outcomes of a game is a *coarse correlated equilibrium* if for every player  $i \in \{1, 2, \dots, k\}$  and every unilateral deviation  $s'_i \in S_i$ ,

$$\mathbf{E}_{\vec{s} \sim \rho}[u_i(\vec{s})] \geq \mathbf{E}_{\vec{s} \sim \rho}[u_i(s'_i, \vec{s}_{-i})]. \quad (10.2)$$

The condition (10.2) is the same as that for the Nash equilibrium (Definition 1.3), except without the restriction that  $\rho$  is a product distribution. In this condition, when a player  $i$  contemplates a deviation  $s'_i$ , she knows only the distribution  $\rho$  and *not* the component  $s_i$  of the realization. That is, a coarse correlated equilibrium only protects against unconditional unilateral deviations, as opposed to the unilateral deviations conditioned on  $s_i$  that are addressed in Definition 10.1. It follows that every correlated equilibrium is also a coarse correlated equilibrium (Figure 10.1).

---

<sup>5</sup>For example, consider the row player. If the trusted third party (i.e., the traffic light) recommends the strategy “Go” (i.e., is green), then the row player knows that the column player was recommended “Stop” (i.e., has a red light). Assuming the column player plays her recommended strategy and stops at the red light, the best strategy for the row player is to follow her recommendation and to go.



**Figure 10.1:** The relationship between Nash equilibria (NE), correlated equilibria (CE), and coarse correlated equilibria (CCE). Enlarging the set of equilibria increases computational tractability but decreases predictive power.

As you would expect,  $\epsilon$ -approximate correlated and coarse correlated equilibria are defined by adding a “ $-\epsilon$ ” to the right-hand sides of (10.1) and (10.2), respectively. We can now answer the question about smooth fictitious play in general games: the time-averaged history of joint play under smooth fictitious play converges to the set of coarse correlated equilibria.

**Proposition 10.3** (SFP Converges to CCE). *For every  $k$ -player game in which every player has at most  $m$  strategies, after  $T = O((\log m)/\epsilon^2)$  time steps of smooth fictitious play, the time-averaged history of play  $\frac{1}{T} \sum_{t=1}^T \bar{\sigma}^t$  is an  $\epsilon$ -approximate coarse correlated equilibrium.*

Proposition 10.3 follows straightforwardly from the definition of  $\epsilon$ -approximate coarse correlated equilibria and the vanishing regret guarantee of smooth fictitious play that we proved in Solar Lecture 1. Precisely, by Corollary 1.11 of that lecture, after  $O((\log m)/\epsilon^2)$  time steps of smooth fictitious play, every player has at most  $\epsilon$  regret (with respect to the best fixed strategy in hindsight, see Definition 1.9 in Solar Lecture 1). This regret guarantee is equivalent to the conclusion of Proposition 10.3 (as you should check).

What about correlated equilibria? While the time-averaged history of play in smooth fictitious play does not in general converge to the set of correlated equilibria, Foster and Vohra [55] and Hart and Mas-Colell [76] show that the time-averaged play of other reasonably simple types



of uncoupled dynamics is guaranteed to be an  $\epsilon$ -correlated equilibrium after a polynomial (rather than logarithmic) number of time steps.

## 10.4 Computing an Exact Correlated or Coarse Correlated Equilibrium

### 10.4.1 Normal-Form Games

Solar Lecture 1 showed that approximate Nash equilibria of two-player zero-sum games can be learned (and hence computed) efficiently (Theorem 1.8). Proposition 10.3 and the extensions in [55, 76] show analogs of this result for approximate correlated and coarse correlated equilibria of general games. Solar Lecture 1 also showed that an exact Nash equilibrium of a two-player zero-sum game can be computed in polynomial time by linear programming (Corollary 1.5). Is the same true for an exact correlated or coarse correlated equilibrium of a general game?

Consider first the case of coarse correlated equilibria, and introduce one decision variable  $x_{\vec{s}}$  per outcome  $\vec{s}$  of the game, representing the probability assigned to  $\vec{s}$  in a joint distribution  $\rho$ . The feasible solutions to the following linear system are then precisely the coarse correlated equilibria of the game:

$$\sum_{\vec{s}} u_i(\vec{s})x_{\vec{s}} \geq \sum_{\vec{s}} u_i(s'_i, \vec{s}_{-i})x_{\vec{s}} \quad \text{for every } i \in [k] \text{ and } s'_i \in S_i \quad (10.3)$$

$$\sum_{\vec{s} \in \vec{S}} x_{\vec{s}} = 1 \quad (10.4)$$

$$x_{\vec{s}} \geq 0 \quad \text{for every } \vec{s} \in \vec{S}. \quad (10.5)$$

Similarly, correlated equilibria are captured by the following linear system:

$$\sum_{\vec{s} : s_i=j} u_i(\vec{s})x_{\vec{s}} \geq \sum_{\vec{s} : s_i=j} u_i(s'_i, \vec{s}_{-i})x_{\vec{s}} \quad \text{for every } i \in [k] \text{ and } j, s'_i \in S_i \quad (10.6)$$

$$\sum_{\vec{s} \in \vec{S}} x_{\vec{s}} = 1 \quad (10.7)$$

$$x_{\vec{s}} \geq 0 \quad \text{for every } \vec{s} \in \vec{S}. \quad (10.8)$$

The following proposition is immediate.

**Proposition 10.4** (Gilboa and Zemel [64]). *An exact correlated or coarse correlated equilibrium of a game can be computed in time polynomial in the number of outcomes of the game.*

More generally, any linear function (such as the sum of players' expected payoffs) can be optimized over the set of correlated or coarse correlated equilibria in time polynomial in the number of outcomes.

For games described in *normal form*, with each player  $i$ 's payoffs  $\{u_i(\vec{s})\}_{\vec{s} \in \vec{S}}$  given explicitly in the input, Proposition 10.4 provides an algorithm with running time polynomial in the input size. However, the number of outcomes of a game scales exponentially with the number  $k$  of players.<sup>6</sup> The computationally interesting multi-player games, and the multi-player games that naturally arise in computer science applications, are those with a *succinct description*. Can we compute an exact correlated or coarse correlated equilibrium in time polynomial in the size of a game's description?

#### 10.4.2 Succinctly Represented Games

For concreteness, let's look at one well-studied example of a class of succinctly represented games: *graphical games* [91, 95]. A graphical game is described by an undirected graph  $G = (V, E)$ , with players corresponding to vertices, and a local payoff matrix for each vertex. The local payoff matrix for vertex  $i$  specifies  $i$ 's payoff for each possible choice of its strategy and the strategies chosen by its neighbors in  $G$ . By definition, the payoff of a player is independent of the strategies chosen by non-neighboring players. When the graph  $G$  has maximum degree  $\Delta$ , the size of the game description is exponential in  $\Delta$  but polynomial in the number  $k$  of players. The most interesting cases are when  $\Delta = O(1)$  or perhaps  $\Delta = O(\log k)$ . In these cases, the number of outcomes (and hence the size of the game's normal-form description) is exponential in the size of the succinct description of the game, and

---

<sup>6</sup>This fact should provide newfound appreciation for the distributed learning algorithms that compute an approximate coarse correlated equilibrium (in Proposition 10.3) and an approximate correlated equilibrium (in [55, 76]), where the total amount of computation is only *polynomial* in  $k$  (and in  $m$  and  $\frac{1}{\epsilon}$ ).

solving the linear system (10.3)–(10.5) or (10.6)–(10.8) does not result in a polynomial-time algorithm.

We next state a result showing that, quite generally, an exact correlated (and hence coarse correlated) equilibrium of a succinctly represented game can be computed in polynomial time. The key assumption is that the following EXPECTED UTILITY problem can be solved in time polynomial in the size of the game’s description.<sup>7</sup>

### The EXPECTED UTILITY Problem

Given a succinct description of a player’s payoff function  $u_i$  and mixed strategies  $\sigma_1, \dots, \sigma_k$  for all of the players, compute the player’s expected utility:

$$\mathbf{E}_{\vec{s} \sim \vec{\sigma}}[u_i(\vec{s})].$$

For most of the succinctly represented multi-player games that come up in computer science applications, the EXPECTED UTILITY problem can be solved in polynomial time. For example, in a graphical game it can be solved by brute force—summing over the entries in player  $i$ ’s local payoff matrix, weighted by the probabilities in the given mixed strategies. This algorithm takes time exponential in  $\Delta$  but polynomial in the size of the game’s succinct representation.

Tractability of solving the EXPECTED UTILITY problem is a sufficient condition for the tractability of computing an exact correlated equilibrium.

**Theorem 10.5** (Papadimitriou and Roughgarden [124] and Jiang and Leyton-Brown [84]). *There is a polynomial-time Turing reduction from the problem of computing a correlated equilibrium of a succinctly described game to the EXPECTED UTILITY problem.*

Theorem 10.5 applies to a long list of succinctly described games that have been studied in the computer science literature, with graphical games serving as one example.<sup>8</sup>

<sup>7</sup>Some kind of assumption is necessary to preclude baking an NP-complete problem into the game’s description.

<sup>8</sup>For the specific case of graphical games, Kakade *et al.* [87] were the first to develop a polynomial-time algorithm for computing an exact correlated equilibrium.

The starting point of the proof of Theorem 10.5 is the exponential-size linear system (10.6)–(10.8). We know that this linear system is feasible (by Nash’s Theorem, since the system includes all Nash equilibria). With exponentially many variables, however, it’s not clear how to efficiently compute a feasible solution. The dual linear system, meanwhile, has a polynomial number of variables (corresponding to the constraints in (10.6)) and an exponential number of inequalities (corresponding to game outcomes). By Farkas’s Lemma—or, equivalently, strong linear programming duality (see e.g. [39])—we know that this dual linear system is infeasible.

The key idea is to run the ellipsoid algorithm [93] on the infeasible dual linear system—called the “ellipsoid against hope” in [124]. A polynomial-time separation oracle must produce, given an alleged solution (which we know is infeasible), a violated inequality. It turns out that this separation oracle reduces to solving a polynomial number of instances of the EXPECTED UTILITY problem (which is polynomial-time solvable by assumption) and computing the stationary distribution of a polynomial number of polynomial-size Markov chains (also polynomial-time solvable, e.g. by linear programming). The ellipsoid against hope terminates after a polynomial number of invocations of its separation oracle, necessarily with a proof that the dual linear system is infeasible. To recover a primal feasible solution (i.e., a correlated equilibrium), one can retain only the primal decision variables corresponding to the (polynomial number of) dual constraints generated by the separation oracle, and solve directly this polynomial-size reduced version of the primal linear system.<sup>9</sup>

## 10.5 The Price of Anarchy of Coarse Correlated Equilibria

### 10.5.1 Balancing Computational Tractability with Predictive Power

We now understand senses in which Nash equilibria are computationally intractable (Solar Lectures 2–5) while correlated equilibria are computationally tractable (Sections 10.3 and 10.4). From an economic

---

<sup>9</sup>As a bonus, this means that the algorithm will output a “sparse” correlated equilibrium, with support size polynomial in the size of the game description.

perspective, these results suggest that it could be prudent to study the correlated equilibria of a game, rather than restricting attention only to its Nash equilibria.<sup>10</sup>

Passing from Nash equilibria to the larger set of correlated equilibria is a two-edged sword. Computational tractability increases, and with it the plausibility that actual game play will conform to the equilibrium notion. But whatever criticisms we had about the Nash equilibrium's predictive power (recall Section 1.1.7 in Solar Lecture 1), they are even more severe for the correlated equilibrium (since there are only more of them). The worry is that games typically have far too many correlated equilibria to say anything interesting about them. Our final order of business is to dispel this worry, at least in the context of price-of-anarchy analyses.

Recall from Lunar Lecture 7 that the *price of anarchy* (POA) is defined as the ratio between the objective function value of an optimal solution, and that of the worst equilibrium:

$$\text{PoA}(G) := \frac{f(\text{OPT}(G))}{\min_{\rho \text{ is an equilibrium of } G} f(\rho)},$$

where  $G$  denotes a game,  $f$  denotes a maximization objective function (with  $f(\rho) = \mathbf{E}_{\vec{s} \sim \rho}[f(\vec{s})]$  when  $\rho$  is a probability distribution), and  $\text{OPT}(G)$  is the optimal outcome of  $G$  with respect to  $f$ . Thus the POA of a game is always at least 1, and the closer to 1, the better.

The POA of a game depends on the choice of equilibrium concept. Because it is defined with respect to the worst equilibrium, the POA only degrades as the set of equilibria grows larger. Thus, the POA with respect to coarse correlated equilibria is only worse (i.e., larger) than that with respect to correlated equilibria, which in turn is only worse than the POA with respect to Nash equilibria (recall Figure 10.1).

The hope is that there's a "sweet spot" equilibrium concept—permissive enough to be computationally tractable, yet stringent enough to allow good worst-case approximation guarantees. Happily, the coarse correlated equilibrium is just such a sweet spot!

<sup>10</sup>This is not a totally unfamiliar idea to economists. According to Solan and Vohra [146], Roger Myerson, winner of the 2007 Nobel Prize in Economics, asserted that "if there is intelligent life on other planets, in a majority of them, they would have discovered correlated equilibrium before Nash equilibrium."

## 10.5.2 Smooth Games and Extension Theorems

After the first ten years of price-of-anarchy analyses (roughly 1999-2008), it was clear to researchers in the area that many such analyses across different application domains share a common architecture (in routing games, facility location games, scheduling games, auctions, etc.). The concept of “proofs of POA bounds that follow the standard template” was made precise in the theory of smooth games [135].<sup>11,12</sup> One can then define the *robust price of anarchy* of a game as the best (i.e., smallest) bound on the game’s POA that can be proved by following the standard template.

The proof template formalized by smooth games superficially appears relevant only for the POA with respect to *pure* Nash equilibria, as the definition involves no randomness (let alone correlation). The good news is that the template’s simplicity makes it relatively easy to use. One would expect the bad news to be that bounds on the POA of more permissive equilibrium concepts require different proof techniques, and that the corresponding POA bounds would be much worse. Happily, this is not the case—every POA bound proved using the canonical template automatically applies not only to the pure Nash equilibria of a game, but more generally to all of the game’s coarse correlated equilibria (and hence all of its correlated and mixed Nash equilibria).<sup>13</sup>

**Theorem 10.6** (Roughgarden [135]). *In every game, the POA with respect to coarse correlated equilibria is bounded above by its robust POA.*

---

<sup>11</sup>The formal definition is a bit technical, and we won’t need it here. Roughly, it requires that the best-response condition is invoked in an equilibrium-independent way and that a certain restricted type of charging argument is used.

<sup>12</sup>There are several important precursors to this theory, including Blum *et al.* [14], Christodoulou and Koutsoupias [36], and Vetta [153]. See [135] for a detailed history.

<sup>13</sup>Smooth games and the “extension theorem” in Theorem 10.6 are the starting point for the modular and user-friendly toolbox for proving POA bounds in complex settings mentioned in Section 6.3.4. Generalizations of this theory to incomplete-information games (like auctions) and to the composition of smooth games (like simultaneous single-item auctions) lead to good POA bounds for simple auctions [151]. (These generalizations also brought together two historically separate subfields of algorithmic game theory, namely algorithmic mechanism design and price-of-anarchy analyses.) See [141] for a user’s guide to this toolbox.

For  $\epsilon$ -approximate coarse correlated equilibria—as guaranteed by a logarithmic number of rounds of smooth fictitious play (Proposition 10.3)—the POA bound in Theorem 10.6 degrades by an additive  $O(\epsilon)$  term.

## References

---

- [1] Aaronson, S., R. Impagliazzo, and D. Moshkovitz (2014). “AM with multiple Merlins”. In: *Proceedings of the 29th IEEE Conference on Computational Complexity (CCC)*. 44–55.
- [2] Adler, I. (2013). “The Equivalence of Linear Programs and Zero-Sum Games”. *International Journal of Game Theory*. 42(1): 165–177.
- [3] Alaei, S., H. Fu, N. Haghpanah, J. D. Hartline, and A. Malekian (2019). “Efficient Computation of Optimal Auctions via Reduced Forms”. *Mathematics of Operations Research*. 44(3): 1058–1086.
- [4] Althöfer, I. (1994). “On Sparse Approximations to Randomized Strategies and Convex Combinations”. *Linear Algebra and Its Applications*. 199(1): 339–355.
- [5] Anshu, A., N. Goud, R. Jain, S. Kundu, and P. Mukhopadhyay (2017). “Lifting Randomized Query Complexity to Randomized Communication Complexity”. Technical Report TR17-054, ECCO.
- [6] Arrow, K. J. and G. Debreu (1954). “Existence of an Equilibrium for a Competitive Economy”. *Econometrica*. 22: 265–290.
- [7] Aumann, R. J. (1974). “Subjectivity and Correlation in Randomized Strategies”. *Journal of Mathematical Economics*. 1(1): 67–96.



- [8] Babichenko, Y. (2016). “Query Complexity of Approximate Nash Equilibria”. *Journal of the ACM*. 63(4): 36.
- [9] Babichenko, Y. and A. Rubinfeld (2017). “Communication complexity of approximate Nash equilibria”. In: *Proceedings of the 49th Annual ACM Symposium on Theory of Computing (STOC)*. 878–889.
- [10] Beame, P., S. Cook, J. Edmonds, R. Impagliazzo, and T. Pitassi (1998). “The Relative Complexity of NP Search Problems”. *Journal of Computer and System Sciences*. 57(1): 3–19.
- [11] Ben-Zwi, O., R. Lavi, and I. Newman (2013). “Ascending Auctions and Walrasian Equilibrium”. Working paper.
- [12] Bikhchandani, S. and J. W. Mamer (1997). “Competitive Equilibrium in an Exchange Economy with Indivisibilities”. *Journal of Economic Theory*. 74: 385–413.
- [13] Bitansky, N., O. Paneth, and A. Rosen (2015). “On the Cryptographic Hardness of Finding a Nash Equilibrium”. In: *Proceedings of the 56th Annual Symposium on Foundations of Computer Science (FOCS)*. 1480–1498.
- [14] Blum, A., M. T. Hajiaghayi, K. Ligett, and A. Roth (2008). “Regret Minimization and the Price of Total Anarchy”. In: *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC)*. 373–382.
- [15] Border, K. C. (1985). *Fixed Point Theorems with Applications to Economics and Game Theory*. Cambridge University Press.
- [16] Border, K. C. (1991). “Implementation of Reduced Form Auctions: A Geometric Approach”. *Econometrica*. 59(4): 1175–1187.
- [17] Border, K. C. (2007). “Reduced form Auctions Revisited”. *Economic Theory*. 31: 167–181.
- [18] Braverman, M., 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 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 970–982.

- [19] Braverman, M., Y. Kun Ko, A. Rubinstein, and O. Weinstein (2017). “ETH Hardness for Densest- $k$ -Subgraph with Perfect Completeness”. In: *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1326–1341.
- [20] Brown, G. W. (1951). “Iterative Solutions of Games by Fictitious Play”. In: *Activity Analysis of Production and Allocation*. Ed. by T. C. Koopmans. *Cowles Commission Monograph No. 13*. Wiley. Chap. XXIV. 374–376.
- [21] Cai, Y., C. Daskalakis, and S. M. Weinberg (2012). “An Algorithmic Characterization of Multi-Dimensional Mechanisms”. In: *Proceedings of the 44th Symposium on Theory of Computing (STOC)*. 459–478.
- [22] Cai, Y., C. Daskalakis, and S. M. Weinberg (2012). “Optimal Multi-Dimensional Mechanism Design: Reducing Revenue to Welfare Maximization”. In: *Proceedings of the 53rd Annual Symposium on Foundations of Computer Science (FOCS)*. 130–139.
- [23] Cai, Y., O. Candogan, C. Daskalakis, and C. H. Papadimitriou (2016). “Zero-sum Polymatrix Games: A Generalization of Min-max”. *Mathematics of Operations Research*. 41(2): 648–655.
- [24] Candogan, O., A. Ozdaglar, and P. Parrilo (2015). “Iterative Auction Design for Tree Valuations”. *Operations Research*. 63(4): 751–771.
- [25] Candogan, O., A. Ozdaglar, and P. Parrilo (2018). “Pricing Equilibria and Graphical Valuations”. *ACM Transactions on Economics and Computation*. 6(1): 2.
- [26] Cesa-Bianchi, N. and G. Lugosi (2006). *Prediction, Learning, and Games*. Cambridge University Press.
- [27] Cesa-Bianchi, N., Y. Mansour, and G. Stolz (2007). “Improved Second-Order Bounds for Prediction with expert advice”. *Machine Learning*. 66(2–3): 321–352.
- [28] Che, Y.-K., J. Kim, and K. Mierendorff (2013). “Generalized Reduced Form Auctions: A Network Flow Approach”. *Econometrica*. 81: 2487–2520.
- [29] Chen, X. and X. Deng (2005). “3-Nash is PPAD-Complete”. Technical Report TR05-134, ECCO.

- [30] Chen, X. and X. Deng (2006). “Settling the Complexity of Two-Player Nash Equilibrium”. In: *Proceedings of the 47th Annual Symposium on Foundations of Computer Science (FOCS)*. 261–270.
- [31] Chen, X. and X. Deng (2009). “On the Complexity of 2D Discrete Fixed Point Problem”. *Theoretical Computer Science*. 410(44): 4448–4456.
- [32] Chen, X., X. Deng, and S.-H. Teng (2006). “Computing Nash Equilibria: Approximation and Smoothed Complexity”. In: *Proceedings of the 47th Annual Symposium on Foundations of Computer Science (FOCS)*. 603–612.
- [33] Chen, X., X. Deng, and S.-H. Teng (2006). “Sparse Games Are Hard”. In: *Proceedings of the Second Annual International Workshop on Internet and Network Economics (WINE)*. 262–273.
- [34] Chen, X., X. Deng, and S.-H. Teng (2009). “Settling the Complexity of Computing Two-Player Nash Equilibria”. *Journal of the ACM*. 56(3): 14. Journal version of [29], [30], [32], and [33].
- [35] Choudhuri, A. R., P. Hubáček, C. Kamath, K. Pietrzak, A. Rosen, and G. N. Rothblum (2019). “Finding a Nash Equilibrium Is No Easier Than Breaking Fiat-Shamir”. In: *Proceedings of the 51st Annual ACM Symposium on Theory of Computing (STOC)*. 1103–1114.
- [36] Christodoulou, G. and E. Koutsoupias (2005). “On the Price of Anarchy and Stability of Correlated Equilibria of Linear Congestion Games”. In: *Proceedings of the 13th Annual European Symposium on Algorithms (ESA)*. 59–70.
- [37] Christodoulou, G., A. Kovács, and M. Schapira (2016). “Bayesian Combinatorial Auctions”. *Journal of the ACM*. 63(2): 11.
- [38] Christodoulou, G., A. Kovács, A. Sgouritsa, and B. Tang (2016). “Tight Bounds for the Price of Anarchy of Simultaneous First Price Auctions”. *ACM Transactions on Economics and Computation*. 4(2): 9.
- [39] Chvátal, V. (1983). *Linear Programming*. Freeman.

- [40] Conitzer, V. 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 (ICML)*.
- [41] Dantzig, G. B. (1951). “A Proof of the Equivalence of the Programming Problem and the Game Problem”. In: *Activity Analysis of Production and Allocation*. Ed. by T. C. Koopmans. *Cowles Commission Monograph No. 13*. Wiley. Chap. XX. 330–335.
- [42] Dantzig, G. B. (1981). “Reminiscences About the Origins of Linear Programming”. Technical Report SOL 81-5, Systems Optimization Laboratory, Department of Operations Research, Stanford University.
- [43] Daskalakis, C. and Q. Pan (2014). “A Counter-example to Karlin’s Strong Conjecture for Fictitious Play”. In: *Proceedings of the 55th Annual Symposium on Foundations of Computer Science (FOCS)*. 11–20.
- [44] Daskalakis, C. and C. H. Papadimitriou (2005). “Three-Player Games are Hard”. Technical Report TR05-139, ECCO.
- [45] Daskalakis, C., P. W. Goldberg, and C. H. Papadimitriou (2006). “The Complexity of Computing a Nash Equilibrium”. In: *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*. 71–78.
- [46] Daskalakis, C., P. W. Goldberg, and C. H. Papadimitriou (2009). “The Complexity of Computing a Nash Equilibrium”. *SIAM Journal on Computing*. 39(1): 195–259. Journal version of [44], [45], and [66].
- [47] Daskalakis, C., P. W. Goldberg, and C. H. Papadimitriou (2009). “The Complexity of Computing a Nash Equilibrium”. *Communications of the ACM*. 52(2): 89–97.
- [48] Dobzinski, S. and J. Vondrák (2013). “Communication Complexity of Combinatorial Auctions with Submodular Valuations”. In: *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1205–1215.
- [49] Dobzinski, S., N. Nisan, and M. Schapira (2010). “Approximation Algorithms for Combinatorial Auctions with Complement-Free Bidders”. *Mathematics of Operations Research*. 35(1): 1–13.

- [50] Dütting, P., V. Gkatzelis, and T. Roughgarden (2017). “The Performance of Deferred-acceptance Auctions”. *Mathematics of Operations Research*. 42(4): 897–914.
- [51] Etessami, K. and M. Yannakakis (2010). “On the Complexity of Nash Equilibria and Other Fixed Points”. *SIAM Journal on Computing*. 39(6): 2531–2597.
- [52] Feige, U. (2009). “On Maximizing Welfare where the Utility Functions are Subadditive”. *SIAM Journal on Computing*. 39(1): 122–142.
- [53] Feige, U. and J. Vondrák (2010). “The Submodular Welfare Problem with Demand Queries”. *Theory of Computing*. 6(1): 247–290.
- [54] Feldman, M., H. Fu, N. Gravin, and B. Lucier (2013). “Simultaneous Auctions Are (Almost) Efficient”. In: *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*. 201–210.
- [55] Foster, D. P. and R. Vohra (1997). “Calibrated Learning and Correlated Equilibrium”. *Games and Economic Behavior*. 21(1–2): 40–55.
- [56] Fréchette, A., N. Newman, and K. Leyton-Brown (2017). “Solving the Station Repacking Problem”. In: *Handbook of Spectrum Auction Design*. Cambridge University Press. Chap. 38. 813–827.
- [57] Freund, Y. and R. E. Schapire (1997). “A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting”. *Journal of Computer and System Sciences*. 55(1): 119–139.
- [58] Freund, Y. and R. E. Schapire (1999). “Adaptive Game Playing Using Multiplicative Weights”. *Games and Economic Behavior*. 29(1–2): 79–103.
- [59] Fudenberg, D. and D. K. Levine (1995). “Consistency and Cautious Fictitious Play”. *Journal of Economic Dynamics and Control*. 19(5): 1065–1089.
- [60] Gale, D., H. W. Kuhn, and A. W. Tucker (1951). “Linear Programming and the Theory of Games”. In: *Activity Analysis of Production and Allocation*. Ed. by T. C. Koopmans. *Cowles Commission Monograph No. 13*. Wiley. Chap. XIX. 317–329.

- [61] Ganor, A., C. S. Karthik, and D. Pálvölgyi (2019). “On Communication Complexity of Fixed Point Computation”. arXiv:1909.10958.
- [62] Garg, S., O. Pandey, and A. Srinivasan (2016). “Revisiting the Cryptographic Hardness of Finding a Nash Equilibrium”. In: *Proceedings of the 36th Annual International Cryptology Conference on Advances in Cryptology (CRYPTO)*. 579–604.
- [63] Geanakoplos, J. (2003). “Nash and Walras Equilibrium via Brouwer”. *Economic Theory*. 21(2/3): 585–603.
- [64] Gilboa, I. and E. Zemel (1989). “Nash and Correlated Equilibria: Some Complexity Considerations”. *Games and Economic Behavior*. 1(1): 80–93.
- [65] Gkatzelis, V., E. Markakis, and T. Roughgarden (2017). “Deferred-Acceptance Auctions for Multiple Levels of Service”. In: *Proceedings of the 18th Annual ACM Conference on Economics and Computation (EC)*. 21–38.
- [66] Goldberg, P. W. and C. H. Papadimitriou (2006). “Reducibility Among Equilibrium Problems”. In: *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*. 61–70.
- [67] Göös, M. (2015). “Lower Bounds for Clique vs. Independent Set”. In: *Proceedings of the 56th Annual Symposium on Foundations of Computer Science (FOCS)*. 1066–1076.
- [68] Göös, M. and T. Pitassi (2018). “Communication Lower Bounds via Critical Block Sensitivity”. *SIAM Journal on Computing*. 47(5): 1778–1806.
- [69] Göös, M. and A. Rubinfeld (2018). “Near-Optimal Communication Lower Bounds for Approximate Nash Equilibria”. In: *Proceedings of the 59th Annual Symposium on Foundations of Computer Science (FOCS)*. 397–403.
- [70] Göös, M., S. Lovett, R. Meka, T. Watson, and D. Zuckerman (2016). “Rectangles are Nonnegative Juntas”. *SIAM Journal on Computing*. 45(5): 1835–1869.
- [71] Göös, M., T. Pitassi, and T. Watson (2017). “Query-to-Communication Lifting for BPP”. In: *Proceedings of the 58th Annual IEEE Symposium on Foundations of Computer Science*. 132–143.

- [72] Göös, M., T. Pitassi, and T. Watson (2018). “Deterministic Communication vs. Partition Number”. *SIAM Journal on Computing*. 47(6): 2435–2450.
- [73] Gopalan, P., N. Nisan, and T. Roughgarden (2018). “Public Projects, Boolean Functions, and the Borders of Border’s Theorem”. *ACM Transactions on Economic and Computation*. 6(3-4): 18.
- [74] Gul, F. and E. Stacchetti (1999). “Walrasian Equilibrium with Gross Substitutes”. *Journal of Economic Theory*. 87: 95–124.
- [75] Hannan, J. (1957). “Approximation to Bayes risk in repeated play”. In: *Contributions to the Theory of Games*. Ed. by M. Dresher, A. W. Tucker, and P. Wolfe. Vol. 3. Princeton University Press. 97–139.
- [76] Hart, S. and A. Mas-Colell (2000). “A Simple Adaptive Procedure Leading to Correlated Equilibrium”. *Econometrica*. 68(5): 1127–1150.
- [77] Hartline, J. D. (2017). “Mechanism Design and Approximation”. Book Draft.
- [78] Hassidim, A., H. Kaplan, Y. Mansour, and N. Nisan (2011). “Non-Price Equilibria in Markets of Discrete Goods”. In: *Proceedings of the 12th Annual ACM Conference on Economics and Computation (EC)*. 295–296.
- [79] Hirsch, M. D., C. H. Papadimitriou, and S. A. Vavasis (1989). “Exponential Lower Bounds for Finding Brouwer Fix Points”. *Journal of Complexity*. 5(4): 379–416.
- [80] Hubáček, P. and E. Yogev (2017). “Hardness of Continuous Local Search: Query Complexity and Cryptographic Lower Bounds”. In: *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1352–1371.
- [81] Hubáček, P., M. Naor, and E. Yogev (2017). “The Journey from NP to TFNP Hardness”. In: *Proceedings of the 8th Conference on Innovations in Theoretical Computer Science (ITCS)*. Article 60.

- [82] Impagliazzo, R. and A. Wigderson (1997). “P = BPP If E Requires Exponential Circuits: Derandomizing the XOR Lemma”. In: *Proceedings of the 29th Annual ACM Symposium on Theory of Computing (STOC)*. 220–229.
- [83] Impagliazzo, R., R. Paturi, and F. Zane (2001). “Which Problems have Strongly Exponential Complexity?” *Journal of Computer and System Sciences*. 63(4): 512–530.
- [84] Jiang, A. X. and K. Leyton-Brown (2015). “Polynomial-Time Computation of Exact Correlated Equilibrium in Compact Games”. *Games and Economic Behavior*. 91: 347–359.
- [85] Johnson, D. S. (2007). “The NP-Completeness Column: Finding Needles in Haystacks”. *ACM Transactions on Algorithms*. 3(2): 24.
- [86] Johnson, D. S., C. H. Papadimitriou, and M. Yannakakis (1988). “How Easy is Local Search?” *Journal of Computer and System Sciences*. 37(1): 79–100.
- [87] Kakade, S., M. Kearns, J. Langford, and L. Ortiz (2003). “Correlated Equilibria in Graphical Games”. In: *Proceedings of the 4th ACM Conference on Electronic Commerce*. 42–47.
- [88] Kalyanasundaram, B. and G. Schnitger (1992). “The Probabilistic Communication Complexity of Set Intersection”. *SIAM Journal on Discrete Mathematics*. 5(4): 545–557.
- [89] Karlin, S. (1959). *Mathematical Methods and Theory in Games, Programming, and Economics*. Addison-Wesley.
- [90] Kearns, M. (2007). “Graphical Games”. In: *Algorithmic Game Theory*. Ed. by N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani. Cambridge University Press. Chap. 7. 159–180.
- [91] Kearns, M., M. L. Littman, and S. Singh (2001). “Graphical Models for Game Theory”. In: *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*. 253–260.
- [92] Kelso, A. S. and V. P. Crawford (1982). “Job Matching, Coalition Formation, and Gross Substitutes”. *Econometrica*. 50(6): 1483–1504.
- [93] Khachiyan, L. G. (1979). “A Polynomial Algorithm in Linear Programming”. *Soviet Mathematics Doklady*. 20(1): 191–194.



- [94] Kjeldsen, T. H. (2001). “John von Neumann’s Conception of the Minimax Theorem: A Journey Through Different Mathematical Contexts”. *Archive for History of Exact Sciences*. 56: 39–68.
- [95] Koller, D. and B. Milch (2003). “Multi-agent Influence Diagrams for Representing and Solving Games”. *Games and Economic Behavior*. 45(1): 181–221.
- [96] Kopparty, S., O. Meir, N. Ron-Zewi, and S. Saraf (2017). “High-Rate Locally Correctable and Locally Testable Codes with Sub-Polynomial Query Complexity”. *Journal of the ACM*. 64(2): 11.
- [97] Koutsoupias, E. and C. H. Papadimitriou (1999). “Worst-Case Equilibria”. In: *Proceedings of the 16th Annual Conference on Theoretical Aspects of Computer Science (STACS)*. 404–413.
- [98] Kushilevitz, E. and N. Nisan (1996). *Communication Complexity*. Cambridge University Press.
- [99] Lautemann, C. (1983). “BPP and the Polynomial Hierarchy”. *Information Processing Letters*. 17(4): 215–217.
- [100] Lehmann, B., D. Lehmann, and N. Nisan (2006). “Combinatorial Auctions with Decreasing Marginal Utilities”. *Games and Economic Behavior*. 55(2): 270–296.
- [101] Lemke, C. E. and J. T. Howson, Jr. (1964). “Equilibrium Points of Bimatrix Games”. *SIAM Journal*. 12(2): 413–423.
- [102] Leyton-Brown, K., P. Milgrom, and I. Segal (2017). “Economics and Computer Science of a Radio Spectrum Reallocation”. *Proceedings of the National Academy of Sciences (PNAS)*. 114(28): 7202–7209.
- [103] Lipton, R. J. and N. E. Young (1994). “Simple Strategies for Large Zero-Sum Games with Applications to Complexity Theory”. In: *Proceedings of the 26th Annual ACM Symposium on Theory of Computing (STOC)*. 734–740.
- [104] Lipton, R. J., E. Markakis, and A. Mehta (2003). “Playing Large Games Using Simple Strategies”. In: *Proceedings of the 4th ACM Conference on Electronic Commerce (EC)*. 36–41.
- [105] Littlestone, N. and M. K. Warmuth (1994). “The Weighted Majority Algorithm”. *Information and Computation*. 108(2): 212–261.

- [106] Maskin, E. and J. Riley (1984). “Optimal Auctions with Risk-Adverse Buyers”. *Econometrica*. 52(6): 1473–1518.
- [107] Matthews, S. A. (1984). “On the Implementability of Reduced Form Auctions”. *Econometrica*. 52(6): 1519–1522.
- [108] McLennan, A. (2018). *Advanced Fixed Point Theory for Economics*. Springer.
- [109] McLennan, A. and R. Tourky (2006). “From Imitation Games to Kakutani”. Unpublished manuscript.
- [110] Megiddo, N. and C. H. Papadimitriou (1991). “On Total Functions, Existence Theorems and Computational Complexity”. *Theoretical Computer Science*. 81(2): 317–324.
- [111] Milgrom, P. (2000). “Putting Auction Theory to Work: The Simultaneous Ascending Auction”. *Journal of Political Economy*. 108(2): 245–272.
- [112] Milgrom, P. (2004). *Putting Auction Theory to Work*. *Churchill Lectures in Economics*. Cambridge University Press.
- [113] Milgrom, P. and I. Segal (2020). “Clock Auctions and Radio Spectrum Reallocation”. *Journal of Political Economy*. To appear.
- [114] Morris, Jr., W. D. (1994). “Lemke Paths on Simple Polytopes”. *Mathematics of Operations Research*. 19(4): 780–789.
- [115] Moulin, H. and J. P. Vial (1978). “Strategically Zero-Sum Games: The Class of Games Whose Completely Mixed Equilibria Cannot Be Improved Upon”. *International Journal of Game Theory*. 7(3–4): 201–221.
- [116] Myerson, R. (1981). “Optimal auction design”. *Mathematics of Operations Research*. 6(1): 58–73.
- [117] Nasar, S. (1998). *A Beautiful Mind: a Biography of John Forbes Nash, Jr., Winner of the Nobel Prize in Economics, 1994*. Simon & Schuster.
- [118] Nash, Jr., J. F. (1950). “Equilibrium Points in  $N$ -Person Games”. *Proceedings of the National Academy of Sciences*. 36(1): 48–49.
- [119] Nash, Jr., J. F. (1951). “Non-Cooperative Games”. *Annals of Mathematics*. 54(2): 286–295.

- [120] Nisan, N. (2002). “The Communication Complexity of Approximate Set Packing and Covering”. In: *Proceedings of the 29th International Colloquium on Automata, Languages and Programming (ICALP)*. 868–875.
- [121] Nisan, N. and I. Segal (2006). “The Communication Requirements of Efficient Allocations and Supporting Prices”. *Journal of Economic Theory*. 129(1): 192–224.
- [122] Papadimitriou, C. H. (1994). “On the Complexity of the Parity Argument and Other Inefficient Proofs of Existence”. *Journal of Computer and System Sciences*. 48(3): 498–532.
- [123] Papadimitriou, C. H. (2007). “The Complexity of Finding Nash Equilibria”. In: *Algorithmic Game Theory*. Ed. by N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani. Cambridge. Chap. 2. 29–51.
- [124] Papadimitriou, C. H. and T. Roughgarden (2008). “Computing Correlated Equilibria in Multi-Player Games”. *Journal of the ACM*. 55(3): 14.
- [125] Pass, R. and M. Venkatasubramanian (2019). “A Round-Collapse Theorem for Computationally-Sound Protocols; or, TFNP is Hard (on Average) in Pessiland”. arXiv:1906.10837.
- [126] Raz, R. and P. McKenzie (1999). “Separation of the Monotone NC Hierarchy”. *Combinatorica*. 19(3): 403–435.
- [127] Raz, R. and A. Wigderson (1994). “Monotone Circuits for Matching Require Linear Depth”. *Journal of the ACM*. 39(3): 736–744.
- [128] Razborov, A. A. (1992). “On the Distributional Complexity of Disjointness”. *Theoretical Computer Science*. 106(2): 385–390.
- [129] Robinson, J. (1951). “An Iterative Method of Solving a Game”. *Annals of Mathematics*: 296–301.
- [130] Rosen, A., G. Segev, and I. Shahaf (2017). “Can PPAD Hardness be Based on Standard Cryptographic Assumptions?” In: *Proceedings of the 15th International Conference on Theory of Cryptography (TCC)*. 173–205.
- [131] Roughgarden, T. (2005). *Selfish Routing and the Price of Anarchy*. MIT Press.

- [132] Roughgarden, T. (2010). “Computing Equilibria: A Computational Complexity Perspective”. *Economic Theory*. 42(1): 193–236.
- [133] Roughgarden, T. (2014). “Barriers to Near-Optimal Equilibria”. In: *Proceedings of the 55th Annual Symposium on Foundations of Computer Science (FOCS)*. 71–80.
- [134] Roughgarden, T. (2014). “CS364B Lecture Notes”. Stanford University.
- [135] Roughgarden, T. (2015). “Intrinsic Robustness of the Price of Anarchy”. *Journal of the ACM*. 62(5): 32.
- [136] Roughgarden, T. (2016). *Twenty Lectures on Algorithmic Game Theory*. Cambridge University Press.
- [137] Roughgarden, T. (2016). “Communication Complexity (for Algorithm Designers)”. *Foundations and Trends in Theoretical Computer Science*. 11(3-4): 217–404.
- [138] Roughgarden, T. and I. Talgam-Cohen (2015). “Why Prices Need Algorithms”. In: *Proceedings of the 16th Annual ACM Conference on Economics and Computation (EC)*. 19–36.
- [139] Roughgarden, T. and É. Tardos (2002). “How Bad Is Selfish Routing?” *Journal of the ACM*. 49(2): 236–259.
- [140] Roughgarden, T. and O. Weinstein (2016). “On the Communication Complexity of Approximate Fixed Points”. In: *Proceedings of the 57th Annual Symposium on Foundations of Computer Science (FOCS)*. 229–238.
- [141] Roughgarden, T., V. Syrgkanis, and É. Tardos (2017). “The Price of Anarchy in Auctions”. *Journal of Artificial Intelligence Research*. 59: 59–101.
- [142] Rubinfeld, A. (2016). “Settling the Complexity of Computing Approximate Two-Player Nash Equilibria”. In: *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science*. 258–265.
- [143] Savani, R. and B. von Stengel (2006). “Hard-to-Solve Bimatrix Games”. *Econometrica*. 74(2): 397–429.
- [144] Schrijver, A. (1986). *Theory of Linear and Integer Programming*. Wiley.

- [145] Shapley, L. S. (1964). “Some Topics in Two-Person Games”. In: *Advances in Game Theory*. Ed. by M. Dresher, L. S. Shapley, and A. W. Tucker. Princeton University Press. 1–28.
- [146] Solan, E. and R. Vohra (2002). “Correlated Equilibrium Payoffs and Public Signalling in Absorbing Games”. *International Journal of Game Theory*. 31(1): 91–121.
- [147] Sperner, E. (1928). “Neuer Beweis für die Invarianz der Dimensionszahl und des Gebietes”. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*. 6(1): 265–272.
- [148] Spielman, D. A. (1997). “The Complexity of Error-Correcting Codes”. In: *Proceedings of the 11th International Symposium on Fundamentals of Computation Theory*. 67–84.
- [149] Spielman, D. A. and S.-H. Teng (2004). “Smoothed Analysis: Why The Simplex Algorithm Usually Takes Polynomial Time”. *Journal of the ACM*. 51(3): 385–463.
- [150] Sun, N. and Z. Yang (2006). “Equilibria and Indivisibilities: Gross Substitutes and Complements”. *Econometrica*. 74(5): 1385–1402.
- [151] Syrgkanis, V. and É. Tardos (2013). “Composable and Efficient Mechanisms”. In: *Proceedings of the 45th ACM Symposium on Theory of Computing (STOC)*. 211–220.
- [152] Toda, S. (1991). “PP is as Hard as the Polynomial-Time Hierarchy”. *SIAM Journal on Computing*. 20(5): 865–877.
- [153] Vetta, A. (2002). “Nash Equilibria in Competitive Societies, with Applications to Facility Location, Traffic Routing and Auctions”. In: *Proceedings of the 43rd Annual Symposium on Foundations of Computer Science (FOCS)*. 416–425.
- [154] Vickrey, W. (1961). “Counterspeculation, Auctions, and Competitive Sealed Tenders”. *Journal of Finance*. 16(1): 8–37.
- [155] Ville, J. (1938). “Sur la Theorie générale des jeux ou intervient l’habileté des joueurs”. Fascicule 2 in Volume 4 of É. Borel, *Traité du Calcul des probabilités et de ses applications*, pages 105–113. Gauthier-Villars.
- [156] von Neumann, J. (1928). “Zur Theorie der Gesellschaftsspiele”. *Mathematische Annalen*. 100: 295–320.
- [157] von Neumann, J. and O. Morgenstern (1944). *Theory of Games and Economic Behavior*. Princeton University Press.

- [158] von Stengel, B. (2007). “Equilibrium Computation for Two-Player Games in Strategic and Extensive Form”. In: *Algorithmic Game Theory*. Ed. by N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani. Cambridge University Press. Chap. 3. 53–78.