

Market Algorithms: Incentives, Learning and Privacy

Jamie Morgenstern

CMU-CS-15-111

April 2015

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Thesis Committee:

Professor Avrim Blum, Chair

Professor Ariel Procaccia

Professor Tuomas Sandholm

Professor Michael Kearns, University of Pennsylvania

*Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy.*

Copyright © 2015 Jamie Morgenstern

This work was supported in part by a Microsoft Women's Scholarship, a Simons Award for Graduate Students in Theoretical Computer Science, and the National Science Foundation under grants CCF-1116892, CCF-1101215, CCF-1331175, and IIS-1065251, as well as an NSF Graduate Research Fellowship. Part of this work was done during an internship at Microsoft Research Redmond, part while visiting the University of Pennsylvania, and part while visiting Stanford University.

The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of any sponsoring organization.

Keywords: Mechanism design, privacy, learning theory, mechanism design without money

To all my friends and family, thank you for your support and encouragement.

Abstract

In this thesis, we study applications of learning theory and differential privacy in the area of mechanism design. Mechanism design aims to optimize over data held by self-interested agents, each of whom will manipulate that data if doing so causes the mechanism to output something more preferred to the agent. Algorithms with learning-theoretic and privacy guarantees are forced to depend upon their inputs in a limited way, suggesting their usefulness in the design of algorithms with limited capacity for manipulation by strategic agents. We explore the particular applications of designing truthful stable matching algorithms, designing simple auctions (using learning theory to choose revenue-optimal auctions, to find equilibrium strategies, and to learn bidder's valuation distributions), and coordinating strategic agents' behavior in a privacy-preserving manner.

Acknowledgments

I stumbled into computer science in college, and for finding my interests there I have many people to thank. Thanks to Robbie Findler for being infinitely patient and encouraging when I knew nothing about programming, to Anne Rogers for inspiring my sharp-witted lecturing style, and to Carly Klivans for first introducing me to the beauty of discrete mathematics. Special thanks to Laci Babai for teaching me that some math takes time, but easy math should be done quickly, as well as my buddies from Discrete Math (Korei, Aaron, we survived and even seem to have made a living out of being precocious!). Extra-special thanks to Dianne Hermann, for letting me think it was a good idea to simultaneously take 199 and 203, allowing me to take the “special section” of 203; to John Boller, Paul Sally, Tam Nguyen and Michael Broshi for putting together an awesome inquiry-based math class that taught me to be a mathematician, and the whole class of students therein (just last week, I used the reverse triangle inequality... some things never change). I would also like to thank Frank Pfenning for taking a chance on me as a summer researcher: I learned a lot about elegance, correctness, and good scientific writing from you, and very much doubt I would have gotten into graduate school without your help.

Avrim has been an awesome advisor, who always has an eye for a beautiful problem (and intuition which leads to beautiful solutions). He has taught me a great deal about theoretical computer science; provided encouragement and advice, and given me a great deal of freedom as a graduate student. He is an exemplary teacher and researcher, and I very much hope I’ve picked up a bit of that from him.

This thesis would not exist if not for the help of my numerous coauthors in graduate school: thank you, Pranjali Awasthi, Avrim Blum, Steven Brams, Simina Brnzei, Ioannis Caragiannis, Nikhil Devanur, Michal Feldman, Deepak Garg, Sampath Kannan, Sean Kennedy, David Kurokawa, John Lai, Omer Lev, Yishay Mansour, Frank Pfenning, Ariel Procaccia, R. Ravi, Ryan Rogers, Aaron Roth, Tim Roughgarden, Ankit Sharma, Or Sheffet, Adam Smith, Vasilis Syrgkanis, Santosh Vempala, Matt Weinberg, Gordon Wilfong, Steven Wu, and Lisa Zhang.

Special thanks to Ariel Procaccia, for having so many great problems to work on, and to Pranjali Awasthi, Or Sheffet, and Ankit Sharma, for being invaluable more-senior students and showing me the ropes. Thanks to Lisa Zhang, Gordon Wilfong, and Sean Kennedy, for a great summer internship at Bell Labs. Thanks to Aaron Roth, for being an excellent collaborator, creative researcher, and mentor; and to Sampath Kannan and Steven Wu (our work from my visit makes up Chapter 5 of this thesis). I particularly appreciate the opportunity to visit Yishay Mansour at Tel Aviv University, Aaron Roth at the University of Pennsylvania, and Tim Roughgarden at Stanford University. Yishay is incredibly good at formulating and solving tough problems; Chapter 4 resulted from Avrim and I visiting him. Chapter 3 came out of an incredibly productive visit with Tim, who has impressive insight and clarity of thought. I also really appreciate the opportunity to intern at Microsoft Research Redmond with Nikhil Devanur: Nikhil has a broad set of tools on which to draw from in formulating and solving problems. The resulting work with Nikhil, Vasilis Syrgkanis and Matt Weinberg formulated Chapter 2. Thank you to Adam Smith, for spending many hours with Avrim, Ankit and myself laboring over the paper that eventually became Chapter 6, and for teaching me so much about privacy and problem-solving. I would also be remiss if I did not thank my committee: thanks to Avrim Blum, Michael Kearns, Ariel Procaccia, and Tuomas Sandholm.

Someone once told me that grad school could be a lonely place: this was never my experience at Carnegie Mellon. Thank you in particular to John Dickerson, Erik Zawadski, Kevin Waugh, Alex Grubb, and Christian Kroer for endless coffees and hallway musings, to Rob Simmons, Chris Martens, David Renshaw, William Lovas, for teaching me about good beer and coffee and teaching and logic.

Personal thanks to Adele Goldberg and Dennis Allison, for being the first computer scientists I ever knew, and encouraging my mathematical interests from before I knew anything about math. Thanks, Mom, for believing I could do anything I put my mind to; to Zayde for always believing there were reasons to break even the best of rules; to Grandma for encouraging me to follow my academic passions, to Aunt Pam, for supporting me no matter what; and to Uncle Scott, for being a fellow academic and pragmatist. Thanks to Jeff and Steph, for being great roommates and friends, especially for feeding me near paper deadlines. Thanks to Tom Murphy VII for always reminding me that tough conversations are never as tough as you expect, and for telling me those IC “fun facts” would be read to the entire department. Thank you to all of my numerous friends, with special thanks to Laura, Michelle, Steph, and Natalie for letting me rant about work during workouts, good and bad. Thanks to Sonia for being a source of reason when I found it difficult to be subjective about my own career (and an excellent friend to boot).

Contents

1	Introduction	1
1.1	Overview of the Thesis	3
2	A Simple Auction with Simple(r) Strategies: the Single-Bid Auction	7
2.1	Introduction	7
2.1.1	Related work	10
2.2	Preliminaries and Notation	12
2.3	Price of Anarchy Upper Bound	15
2.3.1	Almost Tight Lower Bound	19
2.4	Single-Minded Bidders and Restricted Complements	19
2.4.1	Discussion and Future Work	21
2.5	Tighter Upper Bounds for the Single-Bid Auction for Simpler Valuations	22
2.6	Draft Auctions	23
2.6.1	Smoothness of Draft auctions	24
3	A Measure of Simplicity for Auctions: the Fat-Shattering and Pseudo-dimension of Revenue Maximization	29
3.1	Introduction	29
3.2	Related Work	30
3.3	Preliminaries	33
3.3.1	Bayesian Auction Design	33
3.3.2	Sample Complexity, VC Dimension, and the Pseudo-Dimension	34
3.3.3	Applying Pseudo-Dimension to Auction Classes	36
3.4	The sample complexity of some common auctions	36
3.4.1	A Warm-up: the Pseudo-Dimension and Fat-Shattering Dimension of all Incentive-Compatible Mechanisms	36
3.4.2	Pricing, VCG with reserve prices	37
3.4.3	Regular Virtual Valuation Maximizers	40
3.5	Single-Item Auctions	42
3.5.1	t -Level Auctions: The Single-Item Case	42
3.5.2	The Pseudo-Dimension of t -Level Auctions	45
3.5.3	The Representation Error of Single-Item t -Level Auctions	47
3.6	Unbounded MHR Distributions	49
3.7	t -Level Matroid Auctions	51

3.8	Single-parameter t -level auctions	53
3.9	Open Questions	55
4	Learning Valuation Distributions from Partial Observation	57
4.1	Introduction	57
4.1.1	Related Work	58
4.2	Model and Preliminaries	59
4.2.1	A brief primer on the Kaplan-Meier estimator	60
4.2.2	Summary of Main Results	61
4.3	Learning bidders' valuation distributions	61
4.3.1	Subsets	64
4.4	Using the CDF estimates to set optimal reserves	65
4.5	Extensions and Other Models	71
4.5.1	Private value vectors without private randomness	72
4.5.2	Common value vectors with private randomness	73
4.6	Open Questions	74
5	Student-truthful, school-optimal many-to-one stable matching (via differential privacy)	75
5.1	Introduction	75
5.1.1	Our Results and Techniques	77
5.2	Related Work	79
5.2.1	Incentives in Stable Matching	79
5.2.2	Differential Privacy as a Tool for Truthfulness	80
5.3	Preliminaries	81
5.3.1	Many-to-one Matching	81
5.3.2	Differential Privacy Preliminaries	82
5.3.3	Differentially Private Counters	83
5.4	Algorithms Computing Private Matchings are Approximately Truthful	83
5.5	Truthful School-Optimal Mechanism	84
5.6	Conclusions	87
5.7	Privacy Analysis for Counters	88
5.7.1	Composition	88
5.7.2	Details for Counters	89
5.7.3	Counter Privacy under Adaptive Composition	89
5.8	Omitted Proofs	91
5.9	Proofs of Matching Lemmas	92
5.10	A Reference to DA-School	93
5.11	Private Matching Algorithms Must Allow Empty Seats	93
6	Privacy-Preserving Public Information in Coordination Games	95
6.1	Introduction	95
6.1.1	Game Model	96
6.1.2	Information Model	97

6.1.3	Players' Behavior	97
6.1.4	Statement of Main Results	98
6.1.5	Related Work	99
6.2	Privacy-preserving public counters	99
6.3	Resource Sharing	100
6.3.1	Perfect counters and empty counters	101
6.3.2	Privacy-Preserving Counters and Greedy Behavior	102
6.3.3	Privacy-Preserving Counters and Undominated behavior	105
6.4	Resource Sharing with Continuous Resources	106
6.5	Resource Sharing with Future-Dependent Utilities	107
6.6	Private Counters with Smaller Error at Smaller Values	108
6.7	Discussion and Open Problems	111
A	Appendix for Chapter 4	113
A.1	Proofs of inequalities from Chapter 4	113
B	Privacy-Preserving Public Information Appendix	123
B.0.1	Undominated strategic play with Empty Counters: Lower bounds	123
B.0.2	Omitted proofs for Undominated strategies with Privacy-preserving coun- ters	124
B.1	Future-dependent Utilities	125
B.2	Analysis of Private Counters	125
	Bibliography	127

List of Figures

- 3.1 A visualization of the payment rules in Example 3.5.1. Each different line graph corresponds to a different tuple of bids, each corresponding to the examples' different cases. The ticks correspond to thresholds, the nodes to bids. Blue ticks (and nodes) belong to bidder a , green ticks (and nodes) to b , and red ticks (and nodes) to c 44
- 3.2 The successively refined virtual valuation estimates t -level auctions, as t increases. 45

- 5.1 Related works with distributional assumptions on the preferences of one or both sides of the market. Under these assumptions, it is often possible to show that many agents have few (or one) stable partners. If an agent has zero or one stable partner, then she has no incentive to misreport. 80

List of Tables

Chapter 1

Introduction

Classical theoretical computer science has focused on coordinated optimization, where a particular problem has a well-defined input, and a set of optimal solutions for that input, and the goal of the researcher is to design efficient algorithms to construct (near) optimal-solutions for some class of inputs. This setting is well-motivated by settings where a single, central coordinator has full information about the problem parameters, e.g. when an individual wishes to sort his files in order of most to least recently used, or when a single person wishes to compute the most efficient allocation of his jobs to a set of jobs he alone can access. These assumptions fail, however, in a number of important decentralized settings.

The first of these failures arises when the input to our desired computation is held by one or more strategic agents, with preferences over the outcome of this computation. These agents may manipulate the information they give to an algorithm, if they can affect the outcome in a way they prefer. Suppose a system administrator is scheduling jobs for various other parties, some parties may misreport the size or importance of their job so as to expedite the completion of their job over others. In a very different context, bidders in an auction may misreport their values for certain items, in hopes of paying less for those items. The past decade has seen an increased focus on such strategic settings, where agents are optimizing some objective (e.g., their own wait-time) while a central authority is attempting to optimize another objective (e.g., social welfare), where the authority must solicit information from the agents about their objective functions. If algorithm designers wish for an algorithm to perform well in such a setting, it is necessary to design and analyze algorithms in a way which is *incentive-aware*. This can be done either by designing algorithms where no agent can ever manipulate her data to affect the algorithm's outcome in a way which is beneficial to her (designing *truthful* algorithms); or by designing an algorithm and assuming agents will behave strategically in their reporting, then measuring the performance of the algorithm when agents strategically manipulate their data. Designing a truthful mechanism is easy: one which always outputs some fixed outcome which does not depend on the input will always be truthful. On the other hand, it will not be *useful*: such an algorithm has terrible objective value when compared with the optimal solution for the agents' true data. Thus, we will consider **designing strategy-aware algorithms whose performance is measured with respect to the true, unadulterated data**: truthful algorithms will be constrained in their optimization by being truthful; non-truthful mechanism won't have this constraint, but the input to their optimization will be strategically manipulated. In both cases, we will call such

algorithms *mechanisms*.

In other situations, agents might not have intrinsic value for one outcome over another, but might care deeply about the mechanism not revealing too much about their private data. In order to extract the unadulterated data from such agents, the mechanism designer should design algorithms which protect the *privacy* of the individual agents over whose data its computation is run. Anonymizing the agents' data does not, in fact, anonymize the output of the mechanism: simply removing the names from a set of medical records still leaves an abundance of data from which sensitive inferences about individual patients can be made. If the records are purged of names but not mailing addresses of patients, they will still reveal a large amount about household's health records (and thus a large amount about its inhabitants' health). Somewhat more nefarious, even purging all "non-medical" information from the patient records still might leave highly sensitive information about patients behind. For example, suppose there is just one hospital in a region, and the medical records we consider all come from this hospital. Suppose there is only one person who had a broken leg on a particular date at this hospital, and that person also tested positive for HIV. Then, anyone who knows a person whose leg was broken on that date living in this region could discern their friend tested positive for HIV. For this reason, we consider algorithms which satisfy the formal definition of *differential privacy*. Any algorithm A which satisfies this definition has the following remarkable property. Fix any particular individual: any inference one can make about an individual from A 's outcome when that individual provides her data to A one should *similarly likely* to be able to make from A 's outcome when the individual withholds her data. If this is the case, then the individual may as well provide her information to A : any negative (or positive) consequences of the computation with her data will have similar probability of occurring as when the computation is done without her data. Just as in the case of incentive-aware mechanism design, private mechanisms are easy to construct, if the intrinsic quality of the outcome is unimportant. Our goal in these settings will be to designing *privacy-preserving* algorithms which produce output which is *useful* (close in some sense to the non-private computation).

Finally, there are many settings in which the computation's aim is to produce some output which will perform well on *future* data, rather than just on the input to the computation. Over-tailoring the output to the input will usually cause the output to overfit the input and perform poorly on other data, even if that other data is drawn from the same distribution as was the input to the original computation. For example, one might want an algorithm to find a rule for classifying spam email based on previous examples, which would work on future emails and not just the examples given to the algorithm, or one might want to set a price for an ad slot based on previous bidding behavior that gets good revenue in future auctions with different agents participating. We then must assume some relationship between the examples presented to the algorithm (the training data) and the examples on which its performance is important (the test data); for our purposes, we assume the training and test data are drawn from the same distribution. Then, the goal is to construct algorithms whose output performs well on future draws from the distribution, assuming the algorithm was provided enough data and, ideally, to show the necessary amount of input data is small. The area of *machine learning* studies the amount of data necessary to perform such optimizations, and how closely one can optimize with limited amounts of data or computational power.

This goal of constructing algorithms with good generalized performance has philosophical

similarities to the previous two goals: an algorithm which maintains the privacy of agents' data will not overfit the data, in a formal sense, nor will it be able to vary its behavior too much when a single agent manipulates her data. This thesis explores how one can use the tools from machine learning theory and privacy to design new strategy-aware mechanisms, and how the perspective of learning theory can lead to new questions and understanding of standard problems in the area of mechanism design.

1.1 Overview of the Thesis

The first part of this thesis studies connections between machine learning and auction theory. Standard auction theory assumes that agents' preferences are drawn from some distribution, and chooses an auction to optimize (for revenue or some other objective) with respect to those distributions. Prior to the last decade, few had asked how accurately these distributions need to be known to ensure the optimization: more recently, motivated by the Wilson Doctrine [130], a large collection of work has considered detail-free (or detail-limited) auction design. Our work takes a standard machine-learning approach and asks, more precisely, *how much data about bidders' distributions* one needs to learn an accurate estimate of these distributions, or justify such optimizations? The answer to the former question depends upon precisely what form that data takes; the latter question depends upon precisely which set of auctions the optimization is operating over (just as, in standard learning theory, PAC learning's sample complexity depends upon the class from which a function is being chosen).

The sample complexity of learning a nearly-optimal auction from a class of auctions is a formal measure of this class's *simplicity*, a word which previous literature [73] has applied to certain auctions but not others without a mathematical definition. We propose the sample complexity of a class as a formal measure of the simplicity of that class: the smaller a class's sample complexity, the simpler the class of auction. Previous work [55, 107, 126] has also used the word simplicity to refer to a fixed auction (independent of the bidders participating), when the description of that auction is short and easy to parse. Unfortunately, a short description of these auctions has not, in general, implied that the auctions are simple *to play* (e.g., to find equilibria for). We propose a new "simple" auction format, and, using tools from learning theory, show it is formally simpler than those described in the literature. Namely, we show how to compute correlated equilibria of this auction using no-regret learning algorithms.

The second part of this thesis focuses on differential privacy and its applications to mechanism design problems without money. In one case, we study how tools from differential privacy give new results for a long-standing problem in standard mechanism design; in the other, we consider the problem of coordinating behavior of strategic agents where the coordinating entity must preserve the privacy of the choices made by the agents.

In Chapter 2, we study the problem of designing a simple auction for combinatorial bidders, and introduce the single-bid auction. In single-bid auctions, each bidder submits a single real-valued bid for the right to buy items at a price of her bid. We then show single-bid auctions have small price of anarchy: namely, that all correlated equilibria of single-bid auctions achieve welfare which a considerable fraction of the optimal social welfare. Unlike other simple auction formats, such as simultaneous or sequential single-item auctions, bidders can implement no-

regret learning strategies for single-bid auctions in polynomial time. Price of anarchy bounds for correlated equilibria concepts in single-bid auctions therefore have more bite than their counterparts for auctions and equilibria for which learning is not known to be computationally tractable (or worse, known to be computationally *intractable* [31, 45]). This work is joint with Nikhil Devanur, Matthew Weinberg, and Vasilis Syrgkanis, and will appear in EC 2015 [42].

In Chapter 3, we study the problem of finding a formal measure of complexity of a *class* of auctions, when the goal of the mechanism designer is to pick an auction with nearly-optimal revenue from that class. We propose using the *pseudo-dimension* of a class of auctions as a formal measure of its simplicity: smaller pseudo-dimension, in particular, implies a smaller number of samples from each bidder’s distribution necessary to determine which auction from a class is nearly optimal with respect to revenue. We measure the pseudo-dimension of several known classes of auctions (anonymous and nonanonymous pricing, VCG with reserves, Myerson-optimal auctions), and show that our intuition for simplicity often coincides with this formal definition. This approach leads us to the question “Which class of auctions has the ideal tradeoff between expressivity (its ability to nearly optimize revenue) and learnability (its complexity being small enough that polynomially many samples is enough to select a near-optimal auction from the class)?”. We define a new class of auctions, *t*-level auctions, which provides nearly-optimal tradeoff between these two desiderata. This work is joint with Tim Roughgarden and is under submission.

In Chapter 4, we consider a different sort of samples available about bidders’ values: rather than considering that each sample is an *n*-dimensional vector of an independent draw from each bidder’s valuation, we assume a sample is simply the result of a first or second-price auction (e.g., the identity winner of such an auction, and perhaps the price they paid). From this limited information, we ask the information-theoretic question of whether or not one can reconstruct each bidder’s distribution over values to some accuracy, and bound the sample complexity of doing so. This information model is inspired by the idea that the results of previous auctions might be the available source of information about a bidding population. We also show that the sense in which we can approximately learn bidders’ distributions is useful, by showing this approximation is sufficient to set an approximately revenue-optimal reserve for nearly any subset of bidders. This work is joint with Avrim Blum and Yishay Mansour, and is published in AAAI 2015 [21].

In Chapter 5, we present a mechanism for computing asymptotically stable school optimal matchings, while guaranteeing that it is an asymptotic dominant strategy for every student to report their true preferences to the mechanism. Inherent in this work is the ability to coordinate students’ behavior using privacy-preserving information. We design a decentralized version of the celebrated deferred acceptance algorithm which is differentially private in the behavior of the students. This decentralized version’s privacy implies the mechanism is insensitive to the behavior of any one student, leading to our truthfulness guarantee. This is the first setting in which it is known how to achieve nontrivial truthfulness guarantees for students when computing school optimal matchings, assuming *worst-case* preferences (for schools and students) in large markets. This work is joint with Sampath Kannan, Aaron Roth, and Steven Wu, and was published in SODA 2015 [81].

In Chapter 6, we design a mechanism to coordinate player’s choices, for which privacy itself is a constraint. We consider a set of sequential games, where players arrive in some order and

commit to an action based on the choices of previous players. We focus our study on the case where their payoff is determined only by the choices of previous players¹. If agents have perfect information about the decisions of previous players, greedy behavior is a dominant strategy: we show that we can provide privacy-preserving information about previous decisions which is enough to guarantee that greedy behavior still achieves good social welfare. This work is joint with Avrim Blum, Adam Smith, and Ankit Sharma, and was published in ITCS 2015 [23].

¹We do also consider the case where payoffs depend upon the choices of all players, including those in the future, though the results in that setting are weaker.

Chapter 2

A Simple Auction with Simple(r) Strategies: the Single-Bid Auction

2.1 Introduction

A great deal of recent interest in *simple* auction formats [15, 16, 54, 55, 107, 126] has emphasized the importance of the *simplicity* of a combinatorial auction, above and beyond the importance of the auction being truthful. Indeed, much of the work on simple combinatorial auctions considers auction whose bidding languages are not sufficiently expressive to allow for truthful behavior. A natural question to ask regarding these simple auctions is whether some other (approximately) dominant strategy might exist, or, barring such a result, whether or not one could efficiently compute an equilibrium of such mechanisms. In the case of item auctions¹, there is strong negative evidence about either possibility: for simultaneous item auctions, it is hard to even compute a best response in the Bayesian setting [31]². Even in the complete information setting, computing an equilibrium of simultaneous item auctions requires exponential communication for a constant number of subadditive bidders [45]. Random single pricings³, while simple to execute and approximately optimal for both revenue and welfare [13], ignore the possibility of competition between agents. An agent may not, for example, be able to buy an item at *any cost*, in the case of limited supply, if that item is purchased by an agent earlier in the arrival ordering.

Our work explores a new auction format, which we call a single-bid auction, for combinatorial auctions. This format is arguably “simple” in an informal sense, but has an added formal sense in which it is simpler than either sequential or simultaneous single-item auctions: it is possible to compute equilibria of the auction using no-regret learning algorithms. This result alone is not enough to justify the use of single-bid auctions, since many auctions have equilibria which are easy to compute, but allocate items in a way which is terrible for worst-case welfare and revenue (for example, the auction which ignores all bids and allocates the grand bundle of items to bidder 1). We therefore show that single-bid auctions have provably good social welfare for *all*

¹Selling each distinct item separately, either simultaneously or one after the other in some order.

²Cai and Papadimitriou [31] have some discussion of no-regret dynamics in simultaneous second price auctions that indicates that such algorithms may not exist.

³Picking a random single price p , and allowing agents in some order to buy their most-preferred available bundle and pay p for each item.

correlated equilibria: the welfare guarantees combined with the computability of the equilibria provides a strong argument for the practicality of single-bid auctions.

In single-bid auctions, each bidder submits a single real-valued bid for the right to buy items at a fixed price per item of her bid. Contrary to other simple auction formats, the bid space of this auction is small: each bidder submits a single bid rather than m bids. Thus, bidders can implement no-regret learning strategies for single-bid auctions in polynomial time. Price of anarchy bounds for correlated equilibria in single-bid auctions therefore have more bite than their counterparts for auctions for which finding equilibria is not known to be computationally tractable (e.g., for sequential and simultaneous item auctions). We then show that, for any subadditive valuations, the social welfare at equilibrium is an $O(\log m)$ -approximation to the optimal social welfare, where m is the number of items. We also provide tighter approximation results for several subclasses. Our welfare guarantees hold for Nash equilibria and no-regret learning outcomes in both Bayesian and complete information settings via the smooth-mechanism framework. Of independent interest, our techniques show that in a combinatorial auction setting, efficiency guarantees of a mechanism via smoothness for a very restricted class of *cardinality* valuations extend, with a small degradation, to subadditive valuations, the largest complement-free class of valuations.

Our Results. In this chapter, we introduce an extremely simple auction format that we call a *single-bid auction*. In a single-bid auction, bidders submit a single real number as their bid. They are then visited in decreasing order of bids, and each may pay their bid per item for any number of remaining items. Below is a formal description.

1. Initialize, for all $i \in [n]$, $S_i = \emptyset, P_i = 0$. The set of remaining items $I = [m]$.
2. Each bidder $i \in [n]$ submits a sealed bid b_i .
3. Sort bidders in decreasing order according to their bids. Break ties arbitrarily.
4. For $i = 1$ to n :
 5. Let i be the i^{th} highest bidder.
 6. Bidder i chooses any set $X_i \subseteq I$.
 7. Bidder i pays her bid for each item in X_i , i.e., $P_i = b_i |X_i|$.
 8. Update $I = I \setminus X_i$.
9. End For.

Importantly, the strategy space of single-bid auctions is simple enough so that one can efficiently deploy no-regret algorithms. The strategic choices in a single-bid auction consist of making a bid (Step 2) and selecting a set of remaining items to purchase (Step 6). The space of possible bids is just \mathbb{R} , but the space of possible sets to choose is still exponential in m . However, a bidder's dominant strategy in Step 6 is extremely simple: once the auction reaches this phase, bidder i faces an item pricing and may pay b_i for any item in I and has nothing to gain by selecting any set besides $X_i = \operatorname{argmax}_{X \subseteq I} \{v_i(X) - |X|b_i\}$.⁴ In other words, when learning the

⁴For simplicity, we assume throughout the body that bidder i can find X_i in polynomial time. If not, then bidder i has no incentive not to select the best set of items to purchase that she can find computationally efficiently, and our

effectiveness of different strategies, bidder i needs only to learn over different potential bids and not also over potential methods for choosing items to purchase.

The challenge, then, is to show that our auction achieves a good fraction of the optimal social welfare at outcomes of no-regret learning algorithms (or, equivalently, correlated equilibria). Such bounds are called bounds on the *price of anarchy* (PoA), which is the ratio of the optimal social welfare to the welfare at the worst possible equilibrium, for various equilibrium concepts. In a nutshell, we show that for subadditive (a.k.a. complement-free) valuations, the price of anarchy of single-bid auction w.r.t. correlated equilibria is at most $\frac{e}{e-1}H_m$ where m is the number of items and H_m is the m^{th} harmonic number. In comparison, for the same class of valuations, the best deterministic and randomized truthful auctions achieve approximation factors of $O(\sqrt{m})$ and $O(\log m)^5$ respectively, and simultaneous first price auctions have a price of anarchy of 2 w.r.t. Bayes Nash equilibria. On the other hand, *sequential* item auctions could have a price of anarchy of $\Omega(m)$ even for a much simpler class of valuations, namely the union of additive and unit-demand valuations, even for Nash equilibria in the complete information setting. This indicates that *identifying* the right auction is important in order to get a price of anarchy such as $O(\log m)$.

We now present informal statements of the main results of this work.

Theorem 2.1.1 (Informal). *There is a polynomial time no-regret learning algorithm for a bidder participating in a single-bid auction.*

Theorem 2.1.2. *The single-bid auction has a price of anarchy of at most $\frac{e}{e-1}H_m$ w.r.t. coarse correlated equilibria.*

We prove Theorem 2.1.2 by developing a reduction of sorts from proving price of anarchy bounds when bidders have subadditive valuations to proving PoA bounds when bidders have considerably simpler valuations that we call *constraint-homogeneous*. A bidder has constraint-homogeneous valuation if he has an interest set S and the same obtains value v per item in S and 0 per item not in S . This reduction itself may be of independent interest. The proof and formal statement of Theorem 2.1.2 can be found in Section 2.3.

We also provide stronger PoA bounds for restricted classes of valuations, such as unit-demand, concave-symmetric, and k -demand valuations in Section 2.5. We show that, when restricting valuations to have k -restricted complements, for $k \geq 2$, the price of stability (and anarchy) of single-bid auctions is $\Omega(m)$ (a somewhat surprising result, that this is not $O(k)$) in Section 2.4. In Section 2.3.1, we include a lower bound of $\Omega\left(\frac{\log m}{\log \log m}\right)$ on the possible PoA for single-bid auctions when we have additive bidders.

Finally, in Section 2.6 we provide PoA bounds for a sequential format of single-bid auctions which we call **draft auctions**. A draft auction proceeds in rounds: each bidder submits a bid in each round. The highest bidder in each round may pick any of the remaining items, and pays her bid for each item she picks. We show an $O(\log m)$ bound on the price of anarchy with subadditive bidders for draft auctions as well. This offers a significant advantage over *sequential* item auctions, for which the price of anarchy for even simple valuation classes such as the union of additive and unit-demand valuations is $\Omega(m)$ [56].

approximation ratios degrade naturally based on how well the bidders can perform this optimization.

⁵When an upper bound on the largest valuation is known, else the best-known upper bound is $O(\log m \log \log m)$. Assuming $n = \text{poly}(m)$.

2.1.1 Related work

Truthful Auctions. The study of combinatorial auctions has long focused on the design of truthful auctions. Although the VCG mechanism is truthful and gives the socially optimal allocation, it is not computationally efficient. Within the AGT community, this computational barrier has spurred a lively line of research into designing truthful mechanisms that run in polynomial time and approximate the social welfare for various classes of valuations. The state-of-the-art for various instances are: an $O(\log m)$ -approximation for subadditive bidders are subadditive and h is known [13], an $O(\log m \log \log m)$ -approximation when bidders are subadditive [44], an $O(\log m)$ -approximation when bidders are fractionally subadditive [89], and an $e/(e - 1)$ -approximation when bidders have coverage valuations [46]. These mechanisms (and others) are all quite impressive, but have some drawbacks preventing them from being used in practice, such as being non-combinatorial in nature, or having a high probability of completely ignoring many participants.

Comparison to the Random Single Price Mechanism The random single price mechanism of Balcan et al. [13] is truthful and arguably simple: it picks a random real value r from the set $\{\frac{1}{m}, \frac{2}{m}, \dots, h\}$ and asks buyers one at a time to pick their most-preferred bundle B at price r and pay $r|B|$. This truthful mechanism guarantees an $O(\log(mh))$ -approximation to welfare when $h \geq \max_i v_i([m])$ is known and bidders are subadditive (and even a $O(\log m)$ -approximation if $h' = \max_i v_i([m])$ and $h \leq mh'$, since prices need only be as small as h/m^2). While this mechanism is truthful (and therefore requires no learning to find equilibria), it suffers two primary drawbacks that single-bid auctions do not. First, the probability that the random single price mechanism gets welfare of 0 can be as high as $1 - O\left(\frac{1}{\log m}\right)^6$; the single-bid auction has probability $1 - \epsilon$ of returning a correlated equilibrium with an $O(\log m)$ -approximation to welfare (where ϵ is a tunable parameter which affects the running time of the learning procedure). Second, and importantly, the random single price mechanism cannot take advantage of simpler valuation classes to get better approximations to welfare. Even in the case of a single, additive bidder who has value $1/m$ for each item, the random single price mechanism's approximation is $\Theta(\log m)$. In either a single-bidder setting or when bidder's valuations are additive and constant across items, single-bid auctions' price of anarchy is $\Theta(1)$.

Price of Anarchy. More recently, an alternate approach has been to analyze simple auctions that are commonly used in practice, by quantifying the inefficiency of equilibria via the price of anarchy [16, 32, 38, 54, 56, 74, 92, 93, 106, 107, 126]. The dominating theme here has been the emergence of a “smoothness” framework that captures many of the price of anarchy bounds, and allows these bounds to be extended to larger classes of equilibria: Roughgarden [119] to outcomes of learning algorithms and Roughgarden [121] and Syrgkanis [125] to games of incomplete information. Syrgkanis and Tardos [127] give a specialized smoothness framework for auctions with quasi-linear preferences, which we also use. The result most directly comparable

⁶Consider a single buyer and m items; they buyer has value $\frac{2-\epsilon}{m}$ for each item. Then, there is exactly one price at which the buyer will buy any items, namely, at price $1/m$, which is selected with probability at most $\frac{1}{\log m}$.

to ours is that of Feldman et al. [54], which shows that simultaneous item auctions have a constant price of anarchy for subadditive bidders. The ratio is, of course, more desirable than ours. However it is unknown how to compute any of the equilibria for which their PoA guarantees hold (even approximately) in polynomial time. So, without further research, it is unclear whether one should expect bidders in simultaneous item auctions to play an (approximate) equilibrium. In contrast, bidders can reach equilibria of single-bid auctions in polynomial time via distributed no-regret learning, so it is quite reasonable to expect strategic play to approach equilibrium.

Right to Choose Auctions In the economics community the literature on right to choose (RTC) auctions is the closest to our work. Most of this work is empirical, some in the field and others in the lab, and shows that the revenue of RTC auctions is higher than that of other auctions. Among field experiments Ashenfelter and Genesove [5] studied the result of RTC auctions in condominium sales in Miami, which indicated⁷ that the revenue of RTC auctions could be higher than other formats. Alevy et al. [3] studied RTC auctions for water rights sales in Chile and found higher revenue than in the analogous sequential item auction. Laboratory experiments by Eliaz et al. [52], Goeree et al. [66] and Salmon and Iachini [124] all find evidence of higher revenue in RTC auctions under various settings.

Most theoretical work on RTC focuses on very special cases. Harstad [68] finds that revenue equivalence holds between RTC and sequential item auctions, for 2 superadditive bidders. Gale and Hausch [63] has shown that all Bayes-Nash equilibria yield socially optimal allocations for 2 unit-demand bidders. [29] shows that RTC generates more revenue than sequential item auctions, when there are 2 items and many single-minded, risk-averse bidders, each equally likely to prefer either item, whose valuations are drawn i.i.d from a continuous distribution. Yet, it is not clear if RTC auctions always generate a higher revenue than other auctions for a general setting.

Other Work on the Simplicity of Mechanisms Other work has been done on the study of simplified bidding languages (for example, an incomplete list includes Abrams et al. [2], Bichler et al. [17], Blumrosen and Feldman [24], Blumrosen and Nisan [25], Boutilier and Hoos [26], Dütting et al. [48], Ghosh and Sayedi [65], Holzman et al. [75], Kalagnanam and Parkes [80], Milgrom [98], Nisan and Ronen [101], Nisan and Segal [102], Ronen [113]), or more generally, simple mechanisms. Benisch et al. [15] measure the expressiveness of a mechanism in terms of an individual agent's ability to unilaterally distinguish between a collection of outcomes. As one considers the class of all mechanisms of a given expressiveness level, the mechanism with the *best* welfare guarantee in that class will have a (weakly) better welfare guarantee than the *best* welfare guarantee of any mechanism with smaller expressivity.

Equilibrium Computation. We conclude this section by briefly discussing positive and negative results related to equilibrium computation in simple auctions. Lehmann et al. [91] showed how to efficiently compute a pure Nash equilibrium of simultaneous second-price auctions when bidders are submodular. Unfortunately, the equilibrium computed is quite unnatural: it selects a desired winner for each item and asks them to place a large bid on that item, and for all other bidders to bid 0. Even though their construction finds an equilibrium where the large bids are

⁷We find the results inconclusive, due to reasons we cannot go into here.

not “overbids,” it is still clear that this equilibrium is unnatural: it is carefully constructed by a centralized agent with a specific allocation in mind, and it asks bidders to play dominated strategies (why bid 0 if you have any positive value for adding an item?). To our knowledge, there are no other positive results regarding equilibrium computation in simple auctions. On the negative side, recently [31] proved that it is PP-hard⁸ to find an exact Bayes-Nash equilibrium in simultaneous second-price auctions with submodular bidders, and that it is also NP-hard to find an ϵ -Bayes-Nash for some constant ϵ . They further extend their hardness to a notion of ϵ -Bayes-Coarse-Correlated equilibria, and show that this equilibrium is also NP-hard to find. Recently, Dobzinski et al. [45] also show that computing pure Nash equilibria of simultaneous second-price item auctions requires exponential communication. This line of work suggests that simple auctions with strong PoA bounds which do not explicitly consider equilibrium computation may have less bite in a computationally-constrained world. Our work addresses this concern as bidders can run regret-minimization algorithms in polynomial time. A very recent paper of Roughgarden [122] explores formal barriers to obtaining mechanisms with low price of anarchy. He shows, via a reduction to a communication complexity, that no mechanism requiring sub-doubly-exponential communication can have a price of anarchy better than 2 for subadditive bidders (achieved by simultaneous first price auctions [54]). It would be interesting to see if similar techniques can provide formal barriers to designing mechanisms “like” single-bid auctions with constant PoA.

2.2 Preliminaries and Notation

Learnability and correlated equilibria. We begin with a review of standard notions from the online learning literature. Suppose there are N actions and T rounds. An online algorithm A selects an action $a^t \in [N]$ (which is in general randomized and is drawn from a distribution, say x^t) in round t . An adversary selects a reward vector $r^t \in [0, h]^N$, where h or a constant upper-bound on it is assumed to be known; r^t is chosen with the knowledge of x^t but not a^t . A receives reward $r_{a^t}^t$. In the *bandit* setting, this is all A learns, as opposed to the *experts* setting, where A learns the entire reward vector. We now define the *regret* of A .

Definition 2.2.1. We say that algorithm A achieves **regret** $R(T)$ with respect to an action sequence a'_1, \dots, a'_T if, for all reward vectors $r^1, \dots, r^T \in [0, h]^N$,

$$\sum_{t=1}^T \mathbb{E}[r_{a'_t}^t - r_{a^t}^t] \leq R(T).$$

If A achieves regret $R(T)$ with respect to all *fixed* action sequences ($a'_1 = a'_2 = \dots = a'_T$), we say that A achieves **external regret** of $R(T)$. If A achieves regret $R(T)$ with respect to all action sequences $f(a_1), f(a_2), \dots, f(a_T)$ for some $f : [N] \rightarrow [N]$, we say A achieves **swap regret** of $R(T)$.

We say an algorithm is a **no-regret** algorithm if it achieves regret $R(T) = o(T)$. We say that a game is **learnable** if in the setting where the same game is repeated many times, each player has a polynomial time learning algorithm that achieves external/swap regret of $o(T^{1-\delta})$ over the

⁸PP is the class “BPP without the B,” and lies somewhere between the polynomial hierarchy and PSPACE.

set of all his strategies⁹.

The single-bid auction induces a multi-player simultaneous move game among all the bidders, where the strategy of bidder i is his bid b_i . A tuple of bids \mathbf{b} determines the outcome of the auction; player i 's utility is $u_i(\mathbf{b}; v_i) := v_i(S_i(\mathbf{b})) - P_i(\mathbf{b})$ where $S_i(\mathbf{b})$ is the set of items i wins and $P_i(\mathbf{b})$ is her total payment. When v_i is clear from context, we will denote this as $u_i(\mathbf{b})$. Additionally, for any bid tuple \mathbf{b} , we denote with $p_j(\mathbf{b})$ the price that item j was sold at under \mathbf{b} . Finally, let $h = O(\max_i v_i([m]))$ be an upper-bound on the maximum valuation any bidder has for the bundle of all goods, which we assume is known. Players may randomize their strategies, in which case the bids (and everything else that depends on the bids) are random variables.

The standard notion of equilibrium used in such games is that of Nash equilibrium, which says that no player can unilaterally deviate from the equilibrium strategy and gain more utility for himself. We consider the relaxed concept of *correlated equilibrium*: a central mediator suggests a particular strategy to each player, drawn jointly from some distribution. This is a correlated equilibrium if each player, knowing the joint distribution and his suggestion but not the suggestions to others, has no incentive to deviate.

Definition 2.2.2. Correlated equilibrium An α -correlated equilibrium is a joint distribution X over bid vectors \mathbf{b} such that, for each player i , following her suggestion b_i drawn from X is a best-response up to an additive error of α , in expectation over the suggestions \mathbf{b}_{-i} , not known to i and assuming everyone else plays according to their suggestion:

$$\forall i, \forall b'_i, \mathbb{E}_{\mathbf{b} \sim X}[u_i(\mathbf{b}) \mid b_i] \geq \mathbb{E}_{\mathbf{b} \sim X}[u_i(b'_i, \mathbf{b}_{-i}) \mid b_i] - \alpha$$

Note that the deviation b'_i is allowed to depend on the suggestion. In the event that b'_i is independent of b_i for all i , we call X an α -**coarse correlated equilibrium**.

A correlated equilibrium is an equilibrium of the static game in the complete information setting. This means that, even if a player knows the types of all other players, and the joint distribution from which the suggestions are being drawn, he will not deviate from the suggested strategy. The following theorem relates an outcome in the repeated setting when each player employs a no-regret learning algorithm to a correlated equilibrium of the static game.

Theorem 2.2.1 (Foster and Vohra [60], Hart and Mas-Colell [69]). *Suppose that a game is repeated for T rounds and each player employs a no-regret learning algorithm with external regret (resp. swap regret) of at most R . Then the joint distribution over strategy tuples given by the empirical distribution of strategies played by the players in each of the T rounds is an R/T -coarse correlated equilibrium (resp. correlated equilibrium) of the static game.*

Thus, for a learnable game, the empirical distribution over strategies when each player runs a no-regret algorithm converges to a correlated equilibrium.

No-regret learning algorithms. By Theorem 2.2.1, the rate of convergence to a correlated equilibrium will be governed by the regret achieved by each of the players' learning algorithms. In each round, each bidder can compute her payoff from the bid she chooses: it is her utility from that round. On the other hand, a bidder cannot know what items would be available for her if her

⁹We insist that $\delta = \Omega(1)$, so that convergence occurs in polynomially many rounds of running the no-regret algorithms.

bid caused her to be later in the ordering: she cannot compute her payoff for all bids. Therefore, we need algorithms which have low regret in the *bandit* setting, rather than in the experts setting. Previous work has given efficient algorithms with low external and swap regret in the bandit setting.

Theorem 2.2.2 (Auer et al. [7], Blum and Mansour [19]). *There exist efficient algorithms which achieve external regret (resp. swap regret) of at most $\sqrt{hNT \log N}$ (resp. $N\sqrt{hNT \log N}$) in a bandit setting.*

One option for each player is to employ the algorithms as given by the above lemmas over the $O(hm/\epsilon)$ experts in the discretized bid space $(0, \frac{\epsilon}{m}, \frac{2\epsilon}{m}, \dots, \lfloor \frac{hm}{\epsilon} \rfloor \frac{\epsilon}{m}, h)$. We state the convergence rate obtained from such a discretization of bids below.

Corollary 2.2.1. *If each player employs an algorithm with external regret (resp. swap regret) as given by Theorem 2.2.2 on the discretized bid space as mentioned above, then after $O\left(\frac{h^2 m}{\epsilon^3} \log \frac{hm}{\epsilon}\right)$ rounds (resp. $O\left(\frac{h^4 m^3}{\epsilon^5} \log \frac{hm}{\epsilon}\right)$ rounds), the players have reached an ϵ -approximate coarse correlated equilibrium (resp. correlated equilibrium).*

Notice that the above convergence rates are pseudopolynomial (they depend polynomially rather than polylogarithmically on h). Alternatively, we can discretize the bid space as follows: $[0, \frac{h\epsilon}{nm}, \frac{2h\epsilon}{nm}, \dots, \lfloor \frac{nm}{\epsilon} \rfloor \frac{h\epsilon}{nm}, h]$. This reduces the total number of bids in our discretization to $O\left(\frac{nm}{\epsilon}\right)$. Furthermore, each bid $b \in [0, h]$ is within an additive $\frac{h\epsilon}{nm}$ of some bid in the discretized bid-space, therefore this discretization allows us to approach a $\frac{h\epsilon}{n}$ -correlated equilibrium in polynomial time.

Corollary 2.2.2. *If each player employs an algorithm with external regret (resp. swap regret) as given by Theorem 2.2.2 on the discretized bid space as mentioned above, then after $O\left(\frac{n^3 m^2}{h\epsilon^3} \log\left(\frac{nm}{\epsilon}\right)\right)$ rounds (resp. $O\left(\frac{n^5 m^4}{h\epsilon^5} \log\left(\frac{nm}{\epsilon}\right)\right)$ rounds), the players have reached an $\frac{h\epsilon}{n}$ -approximate coarse correlated equilibrium (resp. correlated equilibrium) of the discretized bid space auction.*

The total error of the discretization and approximation to correlated equilibrium is additively $O\left(\frac{h\epsilon}{n}\right)$ per bidder. So, the difference in welfare between this approximate correlated equilibrium and an exact correlated equilibrium is at most $O(h\epsilon)$. Since $h \leq OPT$, this is at most $O(\epsilon OPT)$. Thus, any approximation guarantee we prove for exact correlated equilibria will extend to these learnable approximate equilibria, gaining at most an ϵ factor in the approximation guarantee.

Strategic Play and the Price of Anarchy. Strategic play in many auctions can lead to inefficient allocations of goods; furthermore, it is a priori quite difficult to predict what types of strategic play might arise. In recent years, focus has shifted towards the analysis of simple auctions via the *price of anarchy*: one proves claims of the form “as long as bidders use strategies that form a Nash/correlated/coarse correlated/Bayes-Nash equilibrium, the items are allocated approximately efficiently.” Formally, for a given valuation profile \mathbf{v} , let $SW(OPT(\mathbf{v}))$ be the optimal social welfare, which is the highest social welfare obtainable over all possible allocations of items to bidders. $SW(OPT(\mathbf{v})) := \max \left\{ \sum_{i \in [n]} v_i(S_i) : (S_i)_{i \in [n]} \text{ is a partition of } [m] \right\}$. Let T denote a particular set of equilibria, s an equilibrium in T and $SW(s)$ the social welfare at this

equilibrium. Then the price of anarchy w.r.t equilibria in T is defined as

$$PoA(T) := \max_{s \in T} \frac{SW(\text{OPT}(\mathbf{v}))}{SW(s)}.$$

Smooth Mechanisms. Roughgarden [119] introduced the notion of *smooth games*, which was later extended by Syrgkanis and Tardos [127] to the notion of smooth mechanisms. The smooth mechanism framework provides a method by which to prove Price of Anarchy bounds that hold simultaneously for Nash and correlated equilibria in games of incomplete and complete information. Informally, a mechanism is smooth if, considering any behavior for n bidders, there is *always* a deviation for each bidder, such that each bidder could achieve a reasonable fraction of her value at OPT, while not paying too much. These properties need not be true for *each* bidder's deviation, rather; this need only be true on average over bidders. This definition captures the intuition that, if the social welfare of an equilibrium of a smooth mechanism is low, each bidder has a deviation that could benefit her significantly (thus, such low welfare at equilibrium cannot occur).

Definition 2.2.3 (Syrgkanis and Tardos [127]). A mechanism with allocation rule S and payment rule P is (λ, μ) -smooth for a class of valuations $\mathcal{V} = \times_i \mathcal{V}_i$ if for any valuation profile $\mathbf{v} \in \mathcal{V}$, there exists a mapping $b'_i : [0, h] \rightarrow \Delta([0, h])^{10}$ such that for all $\mathbf{b} \in [0, h]^n$:

$$\sum_i \mathbb{E} [u_i(S(b'_i(b_i), \mathbf{b}_{-i}); v_i)] \geq \lambda SW(\text{OPT}(\mathbf{v})) - \mu \sum_i P_i(\mathbf{b}) \quad (2.1)$$

Theorem 2.2.3 (Syrgkanis and Tardos [127]). *If a mechanism is (λ, μ) -smooth then the price of anarchy w.r.t. mixed Bayes-Nash equilibria of the incomplete information setting and correlated equilibria in the complete information setting is at most $\frac{\max\{1, \mu\}}{\lambda}$. Furthermore, if the mapping b'_i is independent of b_i , then this result holds for coarse correlated equilibria.*

2.3 Price of Anarchy Upper Bound

To prove the upper bound on the price of anarchy of the single-bid auction for subadditive valuations, we will establish that the single-bid auction is a $\left(\frac{e-1}{e \cdot H_m}, 1\right)$ -smooth mechanism, where H_m is the m^{th} harmonic number. Our approach is the following: we first show that the mechanism is $\left(\frac{e-1}{e}, 1\right)$ -smooth for a very restricted class of valuations which we dub constraint-homogeneous valuations (CHV). Each CHV is additive, with value for each individual item either 0 or some value \hat{v} , common for all items. Then we show that smoothness of a mechanism for one class of valuations implies smoothness for a more general class, as long as the latter class can be *approximated* by the former within some factor (we use a non-standard notion of valuation approximation, which we precisely define in Lemma 2.3.2). Moreover, the smoothness property degrades exactly by the factor of approximation. We conclude the proof by showing that subadditive valuations can be *approximated* by CHV within a factor of H_m .

¹⁰A probability distribution over bids.

Definition 2.3.1 (Constraint-Homogeneous Valuation). A valuation on a set of items is *constraint-homogeneous* if it is defined via an interest set S and a per-unit value \hat{v} such that:

$$\forall T \subseteq [m] : v(T) = \hat{v} \cdot |T \cap S| \quad (2.2)$$

Lemma 2.3.1 (Smoothness for Constraint-Homogeneous). *The single-bid auction is a $(1 - \frac{1}{e}, 1)$ -smooth mechanism when players have constraint-homogeneous valuations.*

Proof. Consider a constraint-homogeneous valuation profile $\mathbf{v} = (v_1, \dots, v_n)$ and let S_i^* be the set of items allocated to player i in the welfare maximizing allocation for valuation profile \mathbf{v} . We will show that there exists a randomized deviation B'_i , which does not depend upon the behavior of other agents, such that for any bid profile $\mathbf{b} = (b_1, \dots, b_n)$:

$$\mathbb{E}[u_i(B'_i, \mathbf{b}_{-i})] \geq \left(1 - \frac{1}{e}\right) \hat{v} |S_i^*| - \sum_{j \in S_i^*} p_j(\mathbf{b}). \quad (2.3)$$

where $p_j(\mathbf{b})$ is the price at which item j is sold under bid profile \mathbf{b} , i.e. the bid of the player that acquires it under \mathbf{b} .

Suppose that player i deviates to some deterministic bid $t \in [0, \hat{v}]$. Then for any $j \in S_i^*$, if $t > p_j(\mathbf{b})$, it means that when player i gets to pick his set of items, item j is still available. Thus his utility from such a strategy is lower bounded by:

$$u_i(t, \mathbf{b}_{-i}) \geq \sum_{j \in S_i^*} (\hat{v} - t) \cdot \mathbf{1}\{t > p_j(\mathbf{b})\} \quad (2.4)$$

Thus if B'_i is distributed according to density function $f(t) = \frac{1}{\hat{v}-t}$ and support $[0, (1 - \frac{1}{e})\hat{v}]$ then:

$$\mathbb{E}[u_i(B'_i, \mathbf{b}_{-i})] \geq \sum_{j \in S_i^*} \int_{p_j(\mathbf{b})}^{(1-\frac{1}{e})\hat{v}} (\hat{v} - t) \cdot f(t) \cdot dt = \sum_{j \in S_i^*} \left(\left(1 - \frac{1}{e}\right) \hat{v} - p_j(\mathbf{b}) \right) \quad (2.5)$$

which is exactly the lower bound we wanted to show. Summing the latter lower bound for every player, we get the $(1 - \frac{1}{e}, 1)$ -smoothness property. \blacksquare

We will next show that smoothness for constraint-homogeneous valuations implies smoothness for a much larger class of valuations. We achieve this based on the following re-interpretation of the results in Syrgkanis and Tardos [127]¹¹.

Definition 2.3.2 (Pointwise Valuation Approximation). A valuation class V is pointwise β -approximated by a valuation class V' , if for any valuation $v \in V$, and for any set $S \subseteq [m]$, there exists a valuation $v' \in V'$ such that: $\beta v'(S) \geq v(S)$ and for all $T \subseteq [m]$: $v(T) \geq v'(T)$.

¹¹Hartline [70] gives a special case of this re-interpretation for the mechanism defined by simultaneous single-item auctions, showing how smoothness for additive valuations implies smoothness for unit-demand (and XOS) valuations

Importantly, the valuation v' can depend on S . $\beta v'$ only needs to upper bound v at S , while v' needs to lower bound v everywhere else. This is much weaker than the related notion of approximation by a function class, where for every v we ask for a single v' such that v is sandwiched between $\beta v'$ and v' everywhere. We now show that this notion is useful: if a class can be pointwise approximated by some class which is smooth for a mechanism, then the original class is also smooth for that mechanism.

Lemma 2.3.2 (Extension Lemma). *If a mechanism for a combinatorial auction setting is (λ, μ) -smooth for the class of valuations V' and V is pointwise β -approximated by V' , then it is $(\frac{\lambda}{\beta}, \mu)$ -smooth for the class V .*

Proof. Consider a valuation profile $\mathbf{v} = (v_1, \dots, v_n)$ where each valuation v_i comes from valuation class V . For each player i let S_i^* be her optimal allocation under \mathbf{v} and let $\mathbf{v}^* = (v_1^*, \dots, v_n^*)$ be the valuation profile such that $v_i^* \in V'$ is the valuation that β -approximates v_i for set S_i^* : i.e. $\beta \cdot v_i^*(S_i^*) \geq v_i(S_i^*)$ and for all $T \subseteq [m]$: $v_i(T) \geq v_i^*(T)$. By the first property we get that $\beta \cdot SW(\text{OPT}(\mathbf{v}^*)) \geq SW(\text{OPT}(\mathbf{v}))$. By the second property we get that for all bid profiles \mathbf{b} : $u_i(\mathbf{b}; v_i) \geq u_i(\mathbf{b}; v_i^*)$. Let $b'_i : [0, h] \rightarrow \Delta([0, h])$ be the deviation mapping that is designated by the smoothness property of the mechanism under \mathbf{v}^* . Then for any bid profile \mathbf{b} :

$$\begin{aligned} \sum_i \mathbb{E}[u_i(b'_i(b_i), \mathbf{b}_{-i}; v_i)] &\geq \sum_i \mathbb{E}[u_i(b'_i(b_i), \mathbf{b}_{-i}; v_i^*)] \geq \lambda SW(\text{OPT}(\mathbf{v}^*)) - \mu \sum_i P_i(\mathbf{b}) \\ &\geq \frac{\lambda}{\beta} SW(\text{OPT}(\mathbf{v})) - \mu \sum_i P_i(\mathbf{b}) \end{aligned}$$

which implies the mechanism is smooth for the valuation class V . ■

To conclude our argument that single-bid auctions are smooth for subadditive bidders, we now show that subadditive valuations can be H_m -approximated by constraint-homogeneous valuations.

Lemma 2.3.3 (Constraint-Homogeneous H_m -Approximate Subadditive). *Subadditive valuations can be pointwise H_m -approximated by constraint-homogeneous valuations.*

Proof. Consider a subadditive valuation v , some β , and some set of items $X \subseteq [m]$. Let h_S denote the constraint-homogeneous function $h_S(T) = \frac{v(X)}{|S|\beta} |T \cap S|$. It suffices to find S such that $\beta h_S(X) \geq v(X)$ and also $v(T) \geq h_S(T)$ for all T . We will either find such an S or find an upper bound on β .

Consider $S_1 = X$. Then, $\beta h_{S_1}(X) = \beta h_X(X) = v(X)$, so the first inequality holds. If $v(T) \geq h_{S_1}(T)$ holds for all T , then h_{S_1} pointwise β -approximates v at X . If not, there exists some T_1 such that $v(T_1) < h_{S_1}(T_1)$. Then, since v is monotone, $v(T_1 \cap S_1) \leq v(T_1) < h_{S_1}(T_1) = h_{S_1}(T_1 \cap S_1)$.

Iteratively, consider set $S_i = S_{i-1} \setminus T_{i-1}$. As above, $\beta h_{S_i}(X) = \frac{v(X)}{|S_i|\beta} |X \cap S_i| = v(X)$, so the first condition is satisfied by h_{S_i} for all i . If for some i , $v(T) \geq h_{S_i}(T)$ for all T , then h_{S_i} pointwise β -approximates v at X . If not, then there exists some T_i such that $v(T_i) < h_{S_i}(T_i)$.

After $j \leq m$ iterations, we have either found some h_{S_i} which pointwise β -approximates v at X , or we have constructed a partition T_1, \dots, T_j of X such that for all i

$$v(T_i) < h_{S_i}(T_i) = \frac{v(X)}{\beta|S_i|} |S_i \cap T_i| \leq \frac{v(X)}{\beta|S_i|} |T_i| \quad (2.6)$$

Since v is subadditive: $v(X) \leq \sum_i v(T_i)$. Thus, combining this with Equation 2.6,

$$v(X) < \sum_i \frac{v(X)}{\beta|S_i|} |T_i| = \frac{v(X)}{\beta} \sum_i \frac{|T_i|}{|S_i|}.$$

Thus, $\beta < \sum_i \frac{|T_i|}{|S_i|}$. Now, we simply need to upper-bound $\sum_i \frac{|T_i|}{|S_i|}$ to upper-bound β . Notice that

$$\frac{|T_i|}{|S_i|} = \sum_{t=0}^{|T_i|-1} \frac{1}{|S_i|} \leq \sum_{t=0}^{|T_i|-1} \frac{1}{|S_i| - t}$$

so we have as desired,

$$\beta < \sum_i \frac{|T_i|}{|S_i|} \leq \sum_i \sum_{t=0}^{|T_i|-1} \frac{1}{|S_i| - t} = \sum_{\ell=0}^{m-1} \frac{1}{|X| - \ell} = H_m$$

■

To draw more connections to previous work, when the class V' is the set of general additive valuations, then whether a class V can be pointwise β -approximated by V' is equivalent to asking whether the class V is β -fractionally subadditive, i.e. whether there exist a set of additive valuations indexed by some index set \mathcal{L} such that for any S :

$$\max_{\ell \in \mathcal{L}} v^\ell(S) \leq v(S) \leq \beta \max_{\ell \in \mathcal{L}} v^\ell(S) \quad (2.7)$$

It is known that subadditive valuations are H_m -fractionally subadditive [16, 44], or in other words, that subadditive valuations can be pointwise H_m -approximated by additive valuations. Hence, our Lemma 2.3.3 can be viewed as a strengthening of this result, stating that general additive valuations are not needed and only additive valuations with only one possible non-zero value for each individual item, suffices. This result can be of independent interest in algorithmic and mechanism design questions for sub-additive valuations.

Combining Lemma 2.3.3 with the smoothness of single-bid auctions for constraint-homogeneous valuations (Lemma 2.3.1) and the Extension Lemma (Lemma 2.3.2) we get that the single-bid auction is $\left(\frac{e-1}{e-H_m}, 1\right)$ -smooth for subadditive valuations. Moreover, observing that in all our proofs the smoothness deviation was a fixed strategy and not a mapping, yields our main Theorem 2.1.2.

2.3.1 Almost Tight Lower Bound

This bound on the price of anarchy for single-bid auctions is nearly asymptotically tight, even when restricted to additive bidders.

Theorem 2.3.1. *The price of anarchy of single-bid auctions at pure Nash equilibria is at least $\Omega\left(\frac{\log m}{\log \log m}\right)$, even when all bidders are additive.*

Proof. For the sake of simplicity, assume that ties are broken lexicographically when determining bid order throughout this proof. Consider the following bidders, valuations, and items. Suppose there is a partition of the m items B_0, \dots, B_{k-1} . Let $|B_t| = k^t$. Let bidder 0 have valuation v_0 as follows. For each $j \in B_t$ and each t , $v_0(j) = k^{k-t}$; thus, $v_0(B_t) = k^k$ for each t and $v_0([m]) = k^{k+1}$.

Then, for each $j \in \{0, \dots, k-1\}$, let there be two bidders i_{j1}, i_{j2} with valuations $v_{i_{j1}}(i) = v_{i_{j2}}(i) = \frac{v_0(i)}{k}$ for all $i \in B_j$, and $v_{i_{j1}}(i) = v_{i_{j2}}(i) = 0$ for all $i \notin B_j$. Notice that if each of these “small” bidders bids $\frac{v_0(j)}{k}$, then they are both playing a deterministic best-response, irrespective of bidder 0’s bid.

Given that all of the “small” bidders are bidding $\frac{k^{k-t}}{k}$ for $t \in \{0, \dots, k-1\}$, bidder 0 will bid exactly one of these numbers in equilibrium. Suppose she bids $b_0 = \frac{k^{k-t^*}}{k}$. When she bids b_0 , all items in B_{t^*}, \dots, B_{k-1} will be available for her to purchase. Consider some item $j \in B_t$ for $t > t^* + 1$: it is clear that $v_0(j) = k^{k-t} < k^{k-t^*-1} = \frac{k^{k-t^*}}{k} = b_0$. Thus, bidder 0 will not choose to buy any item in $B_{t^*+1}, \dots, B_{k-1}$. Then, she will buy at most the sets B_{t^*-1}, B_{t^*} , obtaining value $v_0(B_{t^*-1} \cup B_{t^*}) = 2k^k$.

Suppose S_i is the set of items bidder i buys at this equilibrium. We just showed that $v_0(S_0) \leq 2k^k$. For all $i \neq 0$, $v_i(S_i) = \frac{v_0(S_i)}{k}$, so $\sum_i v_i(S_i) \leq 2k^k + (k-2)k^{k-1} \leq 3k^k$, while the optimal social welfare is k^{k+1} . Notice that $m = \sum_{t=0}^k k^t = \Theta(k^{k-1})$. Thus, the price of anarchy is at least $\Omega(k) = \Omega\left(\frac{\log m}{\log \log m}\right)$. ■

2.4 Single-Minded Bidders and Restricted Complements

In this section, we tight upper and lower bounds for the price of anarchy of draft auctions when bidders are single-minded (e.g., each bidder i has a single set S_i for which her value is $v_i \geq 0$, and has value zero for any set S' if $S_i \not\subseteq S'$). We also consider whether similar bounds are possible when bidders have *restricted complementarities* of size at most k .

We begin by presenting the main result, namely, that single-bid and draft auctions have a price of anarchy which scales linearly in the number of items, if bidders’ valuations can have arbitrary complementarities.

Theorem 2.4.1. *The price of stability for draft and single-bid auctions, where bidders have valuations with unrestricted complementarities, is $\Theta(m)$.*

Proof. We begin with proving the lower bound. Consider 2 bidders. One bidder is single-minded, and wants the set of all items $[m]$, for which he has value m . Then, suppose there is a single unit-demand bidder who has value $1 + \epsilon$ for any of the m items. The optimal allocation will give all of $[m]$ to the single-minded bidder, for social welfare of m .

Unfortunately, for the single-minded bidder to take all items at equilibrium, he must pay $1 + \epsilon$ for each of the items. But, since $m(1 + \epsilon) > m$, the single-minded bidder will not be willing to win all items. Thus, the unit-demand bidder gets utility $1 + \epsilon$ and the single-minded bidder gets utility 0.

We now proceed to prove the PoS is at most $O(m)$. Single-minded bidders with desired sets S of size at most k always have a deviation where they can achieve their optimal bundle and pay at most kp_h , where p_h is the highest price paid for any item. Thus, each bidder i deviating to buy S_i by bidding p_h will achieve welfare at least that of OPT and the sum of the prices paid will be at most $m \cdot p_h \leq m \cdot \text{REV}$. Thus, the claim follows. ■

Given the results of Theorem 2.4.1, one might ask whether the price of stability (or anarchy) of single-bid auctions might be bounded by the *degree* of complementarity exhibited by bidders. We now formally define the notion of restricted complementarities which will be used for the remainder of this section.

Definition 2.4.1 (*k*-restricted complementarity valuation). We call a weighted hypergraph \mathcal{G} on m items *k*-restricted if each edge e has endpoints $S_e \subseteq [m]$ and $|S_e| \leq k$, and $w_e \geq 0$. Let the valuation $v_{\mathcal{G}}$ according to \mathcal{G} for a bundle S be defined as the sum of the weights of the edges whose endpoints are all in S , e.g. $v_{\mathcal{G}}(S) = \sum_{e: S_e \subseteq S} w_e$. Consider a valuation function v . We say v has *k*-restricted complementarities if there exists a *k*-restricted hypergraph \mathcal{G} such that $v_{\mathcal{G}} = v$. If $k = 1$, this is equivalent to v having no complementarities, and if $k = m$, any monotone valuation can be encoded this way¹².

A proof identical to the lower bound in Theorem 2.4.1 shows that the price of stability for draft and single-bid auctions with *k*-restricted complement bidders is $\Omega(k)$. One might hope that this bound is tight, and that the price of anarchy is $O(k)$ for draft or one-shot auctions with *k*-restricted complementarities. Unfortunately, at least for single-bid auctions, the answer is much worse: for any $\frac{m}{2} \geq k \geq 2$, the price of stability for single-bid auctions is $\Omega(m)$.

Theorem 2.4.2. *The price of stability for single-bid auctions with at least 2 bidders, one of whom has a $k \leq m/2$ -restricted complementarity valuation, can be as large as $\Omega(m)$.*

Proof. There will be two bidders in our example. Bidder 1, who we will refer to as the “complements bidder”, gets the grand bundle for value of OPT in the optimal allocation. Furthermore, suppose her value comes from $m - k$ equal-weight hyperedges, each of which have k goods associated with them (each has goods $1, \dots, k - 1$ and exactly one of k, \dots, m). Then, her utility for a bundle S which contains $1, \dots, k - 1$ and i other items is

$$v_1(S) = i \cdot \frac{\text{OPT}}{m - k}. \quad (2.8)$$

Then, let there be one other bidder, whose sole interest is good 1, for which she has value $v_2(\{1\}) = \frac{\text{OPT}}{m - k} + \epsilon$. Then, in any equilibrium, either bidder 2 wins item 1, or bidder 1 pays $T = \frac{\text{OPT}}{m - k} + \epsilon$ for item 1; thus she must pay T for each item she buys. If bidder 2 buys item 1, then bidder 2 has zero utility, implying social welfare $\frac{\text{OPT}}{m - k} + \epsilon = O\left(\frac{\text{OPT}}{m}\right)$, or PoA at least $\Omega(m)$.

¹² See Conitzer et al. [40] for a detailed definition.

On the other hand, if buyer 1 chooses to buy item 1, she must be getting value from doing so, so she must be buying at least k items (1 through $k - 1$ and at least 1 other item). But, her value for buying i additional items (on top of the first $k - 1$) is $i \cdot \frac{\text{OPT}}{m-k}$, by Equation 2.8, which is less than $(k + i) \cdot T$, the price she would need to pay to buy those items. Thus, she won't buy item 1 in any equilibrium, implying the price of stability is at least $\Omega(m)$. ■

This proof does not have an immediate extension for proving a stronger than $\Omega(k)$ price of anarchy for draft auctions, since the one bidder with complementarities could, in that case, buy item 1 at the more expensive price, then buy the remaining items at a lower price. If the smaller bidder was unit-demand across the first $k - 1$ bidders, then this would show an $\Omega(k)$ price of stability, but not better. We leave this question as an interesting direction for future work.

2.4.1 Discussion and Future Work

Our work contributes to the recent line of work addressing the design and analysis of simple combinatorial auctions with low price of anarchy. We propose coarse correlated equilibria of single-bid auctions as the first solution concept of a simple auction that both has a low price of anarchy and can be computed in polynomial time. Our work also motivates several directions for future research. First, there are numerous questions related to the analysis of existing simple mechanisms, such as simultaneous item auctions. For example, *is there a poly-time no-regret algorithm for simultaneous item auctions?* It is conceivable that such algorithms exist despite the exponential strategy space and evidence found in Cai and Papadimitriou [31]. If not, *is there a different equilibrium concept for simultaneous item auctions that is well-motivated, poly-time computable, and has low price of anarchy?* Or more generally, *how should one expect bidders to behave in a simultaneous item auction? Can one bound the price of anarchy at this behavior?*

Also motivated is the design of new simple auctions with learnable equilibria and a *constant* price of anarchy. We note that doing so will require very different techniques than the present work in a formal sense. One significant generalization of single-bid auctions is the following: bidders first play an arbitrary game, where each bidder has $\text{poly}(m)$ possible strategies. Then, based on the strategies selected by each bidder, the bidders are visited sequentially (in an order determined by the strategies played) and each offered an item pricing over the remaining items (also determined by the strategies played). Single-bid auctions are an extremely special case of these mechanisms, where the game consists of each bidder simply making a bid, and the item pricing is uniform over the remaining items (and the order of bidders is determined by simply ranking their bids, and the uniform price offered is exactly their bid). Recent work of Braverman et al. [27] shows that no auction of this general format can possibly have a price of anarchy $o(\log m / \log \log m)$ at any equilibrium concept. Therefore, modifications to single-bid auctions such as asking the bidders to report multiple bids, charging non-uniform prices, etc. cannot possibly improve the price of anarchy. We therefore hope that future research will be fruitful in designing novel simple mechanisms achieving low price of anarchy at learnable equilibrium concepts.

2.5 Tighter Upper Bounds for the Single-Bid Auction for Simpler Valuations

In this section we show tighter price of anarchy bounds for two other important classes of valuations: unit-demand and symmetric valuations. A valuation is unit-demand if a player only wants one item and has no value for any extra item. Equivalently if it can be expressed as: $v_i(S) = \max_{j \in S} w_{ij}$ for some $w_{ij} \geq 0$. A valuation is symmetric, if it is a function of the number of items and not of the specific set, i.e. if all items are identical. We will consider the case of concave symmetric valuations, i.e., $v_i(S) = f_i(|S|)$ for some non-decreasing concave function $f_i : \mathbb{N} \rightarrow \mathbb{R}^+$.

We show that both unit-demand valuations and concave symmetric valuations can be pointwise 1-approximated by constraint-homogeneous valuations. As a corollary we get that the price of anarchy of single-bid auctions for this case is at most $\frac{e}{e-1}$.

Theorem 2.5.1. *The class of concave symmetric valuations is pointwise 1-approximated by constraint-homogeneous valuations.*

Proof. Consider a valuation profile v as described in the theorem (i.e. $v(S) = f(|S|)$). Consider a set $S \subseteq [m]$ and let v' be the constraint-homogeneous valuation with interest set S and per-unit valuation $\hat{v}' = \frac{f(|S|)}{|S|}$. By concavity of the function f and since $f(0) = 0$, we know that for any $y > x$, $\frac{f(y)}{y} \leq \frac{f(x)}{x}$. Thus we have that for any $T \subseteq [m]$:

$$v'(T) = \hat{v}' \cdot |T \cap S| = \frac{f(|S|)}{|S|} \cdot |T \cap S| \leq f(|T \cap S|) \leq v(T) \quad (2.9)$$

Additionally, $v'(S) = f(|S|) = v(S)$. ■

Lemma 2.5.1. *The class of unit-demand valuations is 1-approximated by constraint-homogeneous valuations.*

Proof. For each set of items S , let $j(S) = \arg \max_{j \in S} w_{ij}$. Then consider the constraint-homogeneous valuation v' , with $\hat{v}' = w_{ij(S)}$ and interest set $S = \{j\}$. Then: $v'(T) = w_{ij(S)} \cdot \mathbf{1}\{j(S) \in T\} \leq \max_{j \in T} w_{ij}$ and $v'(S) = w_{ij(S)} = v(S)$. ■

Lemma 2.5.2. *The class of k -demand valuations is H_k -approximated by constraint-homogeneous valuations.*

Proof. Consider k -demand valuation v and interest set S . We will construct a constraint-homogeneous v' such that $v'(T) \leq v(T)$ for all T but $v'(S) \geq H_k v(S)$. Let

$$S' = \operatorname{argmax}_{S'' \subseteq S, |S''|=k} v(S'').$$

Then, $v(S') = v(S)$. Now, repeat the proof of Lemma 2.3.3, beginning with set S' instead of X . In the final line of the proof, $|T_i|$ and $|X|$ can be replaced with k , rather than m , implying H_k as an upper bound on β . ■

2.6 Draft Auctions

In this section, we formally define *draft* auctions, a sequential version of single-bid auctions, and prove draft auctions have similar smoothness guarantees as we proved for single-bid auctions. Draft auctions proceed in rounds: each round is a first-price auction in which each bidder submits a bid. The winner in each round chooses some subset of the remaining items, and pays her bid for *each* item. Formally, a draft auction is as follows.

1. Initialize, for all $i \in [n]$, $S_i = \emptyset, P_i = 0$. The set of remaining items $I = [m]$.
2. While $I \neq \emptyset$,
3. Each bidder $i \in [n]$ submits a sealed bid b_i and a set $X_i \subseteq I$.
4. Allocate set X_{i^*} to $i^* = \arg \max_{i \in [n]} \{b_i\}$, i.e., $S_{i^*} = S_{i^*} \cup X_{i^*}$. Break ties arbitrarily.
5. Bidder i^* pays her bid for each item in X_{i^*} , i.e., $P_{i^*} = P_{i^*} + b_{i^*} |X_{i^*}|$.
6. The winner i^* , winning bid b_{i^*} and allocated bundle X_{i^*} is announced.
7. End While.

Suppose each bidder's valuation $v_i \in V_i$ is drawn from a distribution: $v_i \sim D_i$. Bidder i knows v_i but only D_j (rather than v_j) for all $j \neq i$. Then, draft auctions form a sequential game of incomplete information (and, in the case that each D_i is a point mass, a sequential game of complete information). A *strategy* $s_i : \mathcal{V}_i \rightarrow \Delta(B_i)$ of bidder i is a function, from her valuation to a distribution over bid plans $b_i \in B_i$. Each bid plan b_i determines the bid b_{it} that a player makes at some round t and the set X_{it} of items he gets conditional on winning, based on the information h_{it} available to her up to that round. For any given valuation profile \mathbf{v} , a tuple of strategies $\mathbf{s} = \mathbf{s}(\mathbf{v}) = (s_i(v_i))_{i \in [n]}$ determines the outcome of the auction; let $u_i(\mathbf{b}; v_i)$ denote the utility, (or expected utility when b is a distribution over bid plans) obtained by bidder i as a function of the bid plans \mathbf{b} . Recall that for a deterministic profile the utility is $v_i(S_i(\mathbf{b})) - P_i(\mathbf{b})$ where $S_i(\mathbf{b})$ is the set of items i wins and $P_i(\mathbf{b})$ is her total payment. Additionally, for any bid plan \mathbf{b} , we denote with $p_j(\mathbf{b})$ the price that item j was sold at, under bid plan \mathbf{b} . Observe that a bid plan actually also contains information about *what might have happened*, i.e., they specify the result of possible deviations from the actual outcome, which becomes important in the definitions of equilibria. We now define the most basic equilibrium concept, that of a Nash equilibrium.

Definition 2.6.1. A **pure (resp. mixed) Bayes-Nash** equilibrium is a pure (resp. mixed) strategy tuple \mathbf{s} such that no player can unilaterally deviate to obtain a better utility. In other words,

$$\forall i \in [n], \forall v_i \in \mathcal{V}_i, \forall b'_i \in B_i, \mathbb{E}_{\mathbf{v}_{-i}}[u_i(b'_i, \mathbf{s}_{-i}(\mathbf{v}_{-i}); v_i)] \leq \mathbb{E}_{\mathbf{v}_{-i}}[u_i(\mathbf{s}(\mathbf{v}); v_i)],$$

where as is standard, $\mathbf{s}_{-i}(\mathbf{v}_{-i})$ denotes $(s_j(v_j))_{j \in [n], j \neq i}$, the strategy tuple \mathbf{s} restricted to players other than i , and $(b'_i, \mathbf{s}_{-i}(\mathbf{v}_{-i}))$ denotes the tuple where $s_i(v_i)$ is replaced by b'_i in $\mathbf{s}(\mathbf{v})$. Similarly \mathbf{v}_{-i} denotes the tuple of valuations $(v_j)_{j \in [n], j \neq i}$. The expectations are taken over the draw of \mathbf{v}_{-i} .

A Nash equilibrium in sequential games allows for *irrational threats*, where an equilibrium strategy of a bidder could be suboptimal beyond a certain round. A standard refinement of the Nash equilibrium for extensive form games is the *subgame perfect equilibrium*, that allows only for strategies that constitute an equilibrium of any subgame, conditional on any possible history

of play (see [61] for a formal definition and a more comprehensive treatment.) Our results also extend to complete-information correlated equilibria.

Subgame perfect \subseteq Nash \subseteq Correlated Equilibria

The price of anarchy may be defined w.r.t any of these equilibria; larger classes have higher price of anarchy. In the Bayesian setting the price of anarchy is defined as the worst-case ratio of the expectations, over the random values, of the social welfare at the optimum $\mathbb{E}_{\mathbf{v}}[SW(\text{OPT}(\mathbf{v}))]$ and at an equilibrium $\mathbb{E}_{\mathbf{v}}[SW(\mathbf{s}(\mathbf{v}))]$.

2.6.1 Smoothness of Draft auctions

We will show that draft auctions are smooth mechanisms according to the general definition of a smooth mechanism, which has the same implications on the price of anarchy as in Theorem 2.2.3.

Definition 2.6.2 ([127]). A mechanism is (λ, μ) -smooth for a class of valuations $\mathcal{V} = \times_i \mathcal{V}_i$ if for any valuation profile $\mathbf{v} \in \mathcal{V}$, there exists a mapping $b'_i : B_i \rightarrow \Delta(B_i)$ such that for all $\mathbf{b} \in B_1 \times \dots \times B_n$:

$$\sum_i \mathbb{E}[u_i(b'_i(b_i), \mathbf{b}_{-i}; v_i)] \geq \lambda SW(\text{OPT}(\mathbf{v})) - \mu \sum_i P_i(\mathbf{b}) \quad (2.10)$$

There are two main technical hurdles in extending the arguments of smoothness for single-bid auctions to draft auctions. Unlike single-bid auctions, draft auctions proceed in rounds. This means that strategies are functions that map history to bids in each round. Bidders' deviations need to aim for particular items *at their equilibrium prices*. So, a deviating bidder needs to behave as they do in equilibrium (to ensure she faces equilibrium prices) until the right moment, at which point they bid the "right bid", and procure the items they would get in the optimal allocation. The second difficulty is that, unlike in sequential item auctions, a player is not aware, without information about other bidders' strategies, at which step any item is going to be allocated, since this is endogenously chosen by one of his opponents. Thus, deviations of the form: "behave exactly as previously until the optimal item arrives and then deviate to acquire it", will not yield smoothness proofs in the case of draft auctions.¹³ Instead, our deviations for the unit-demand case have a player always attempt to get his optimal item, while it is still available, without changing the observed history when she loses. We show a deviation of the following form does just that: *At each time step, as long as your optimal item is still available, bid the maximum of your equilibrium bid and half your value for your optimal item. If you ever win, buy your optimal item.*

Lemma 2.6.1. *The draft auction for unit-demand bidders is a $(\frac{1}{2}, 2)$ -smooth mechanism.*

Proof of Lemma 2.6.1: Consider a unit-demand valuation profile v (i.e. $v_i(S) = \max_{j \in S} v_{ij}$) and let j_i^* be the item assigned to player i in the optimal matching for valuation profile v . We

¹³Even in the complete information setting, the time at which an item sells is defined by the strategies of other players: using this information to construct a deviation would not fit into the smoothness framework. In the case of mixed strategies, or incomplete information, the time an item sells is a random variable, so such a strategy is not even well-defined.

will show that there exists a deviation mapping $b'_i : B_i \rightarrow B_i$ for each player i , such that for any bid profile b :

$$u_i(b'_i(b_i), \mathbf{b}_{-i}) \geq \frac{1}{2}v_{ij_i^*} - p_{j_i^*}(\mathbf{b}) - P_i(\mathbf{b}). \quad (2.11)$$

Consider the following b'_i : in every auction t , the player bids the maximum of her previous bid b_{it} (conditional on the history) and $\frac{v_{ij_i^*}}{2}$, until j_i^* gets sold. If she ever wins some auction, she picks j_i^* . Suppose that j_i^* was sold at some auction t under strategy profile b . We consider the following two cases separately, which are exhaustive since i drops out after round t at most.

Case 1: i wins an auction $t' \leq t$ in b'_i . If i wins with bid $b_{it'}$ then there must have been her payment under b_i as well, and $P_i(\mathbf{b}) = b_{it'}$. Otherwise it is $b_i^* = \frac{v_{ij_i^*}}{2}$. Therefore her utility is

$$u_i(b'_i, \mathbf{b}_{-i}) \geq v_{ij_i^*} - \max \left\{ \frac{v_{ij_i^*}}{2}, P_i(\mathbf{b}) \right\} \geq v_{ij_i^*} - \frac{v_{ij_i^*}}{2} - P_i(\mathbf{b}) \geq \frac{1}{2}v_{ij_i^*} - p_{j_i^*}(\mathbf{b}) - P_i(\mathbf{b}).$$

Case 2: i does not win any auction in b'_i . In this case, it must be that $p_{j_i^*}(\mathbf{b}) \geq \frac{1}{2}v_{ij_i^*}$ since otherwise i would have won auction t . Her utility in this case utility is zero. Therefore (2.11) holds in this case as well.

Thus we have shown that the deviation b'_i always satisfies (2.11). The smoothness property follows by summing over all players and using the fact that $\sum_i p_{j_i^*}(\mathbf{b}) = \sum_{j \in [m]} p_j(\mathbf{b}) = \sum_i P_i(\mathbf{b})$. \blacksquare

Thus, combining Lemma 2.6.1 and Theorem 2.2.3, we have the following.

Corollary 2.6.2. *The price of anarchy for draft auctions with unit-demand bidders is at most 4.*

We now state that draft auctions are smooth for constraint-homogeneous valuations. This implies that the price of anarchy bounds stated for single-bid auctions hold for draft auctions as well. Just like in the case of single-bid auctions, the need to buy many items adds complexity to the proof of this corollary over the unit-demand setting.

Lemma 2.6.3. *The draft auction is a $(\frac{1}{4}, 2)$ -smooth mechanism when bidders have constraint-homogeneous valuations.*

Before proving Lemma 2.6.3, we state a separate lemma, which shows there always exists a “good” deviation for constraint-homogeneous bidders in draft auctions.

Lemma 2.6.4 (Core Deviation Lemma for Draft Auctions). *Suppose that a player i has a constraint-homogeneous valuation with interest set S and per-unit value \hat{v} . Then in a draft auction there exists a deviation mapping $b'_i : B_i \rightarrow B_i$ such that, for any strategy profile b :*

$$u_i(b'_i(b_i), \mathbf{b}_{-i}; v_i) \geq \frac{1}{2} \frac{\hat{v} \cdot |S|}{2} - \sum_{j \in S} p_j - P_i(\mathbf{b}).$$

We now use Lemma 2.6.4 to prove Lemma 2.6.3.

Proof of Lemma 2.6.3 : Consider a constraint-homogeneous valuation profile v and a bid profile b . Let S_i^* be the units allocated to player i in the optimal allocation for profile v . Also let S_i be the interest set of each player and \hat{v}_i his per-unit value. Consider the alternative valuation profile where each player i has a constraint-homogeneous valuation v'_i with interest set $S'_i = S_i \cap S_i^*$ and per unit value $\hat{v}'_i = \hat{v}_i$.

Observe that for any $T \subseteq [m]$, $v_i(T) \geq v'_i(T)$ and $v_i(S_i^*) = v'_i(S_i^*)$. Thus, for any bid profile b : $u_i(b; v_i) \geq u_i(b; v'_i)$ and $SW(\text{OPT}(v')) \geq SW(\text{OPT}(v))$. Invoking Lemma 2.6.4 on valuations v'_i , we get that there exists a deviation mapping $b'_i : B_i \rightarrow B_i$ for each player i such that for any strategy profile b :

$$\sum_i u_i(b'_i(b_i), \mathbf{b}_{-i}; v_i) \geq \sum_i u_i(b'_i(b_i), \mathbf{b}_{-i}; v'_i) \geq \frac{1}{4} \text{OPT}(v') - 2 \sum_i P_i(\mathbf{b}) \geq \frac{1}{4} \text{OPT}(v) - 2 \sum_i P_i(\mathbf{b}),$$

where we have once again used the fact that $\sum_i p_{j_i^*}(\mathbf{b}) = \sum_{j \in [m]} p_j(\mathbf{b}) = \sum_i P_i(\mathbf{b})$. \blacksquare

Combining Lemma 2.3.3 with Lemma 2.6.3 and Lemma 2.3.2, we get the following efficiency guarantee for draft auctions with subadditive valuations.

Corollary 2.6.5. *The price of anarchy for draft auctions with subadditive bidders is at most $8H_m$.*

The Core Deviation for draft auctions is somewhat more complicated than for single-bid core deviation, because it is multi-stage and needs to mimic a bidder's equilibrium behavior. Just as in the case for single-bid auctions, the key deviation to prove smoothness for draft auctions is to bid the "right price", half of her per-unit value, and then try to acquire the "right number" of those items, which is at least half the number of units in her optimal allocation. However, consider a round where her equilibrium bid is higher than the "right price". If the bidder bids the right price, she may change the history for all the other players and sets the game down an off-equilibrium path. Once a deviation has affected the winning history, the prices in the remaining off-equilibrium subgame are difficult to reason about. Thus, the deviations we consider have a player "mimic" her equilibrium play until she can acquire her optimal number of units at a good price.

To achieve this, the deviation bids the maximum of the original bid and the right price. If the original bid is higher, she follows the original strategy and picks the same set of items¹⁴. If the right price is higher, she then buys sufficient number of items to win the "right number" of units, and drops out of subsequent rounds. The following lemma extends the Core Deviation lemma to the draft auction mechanism.

Definition 2.6.3 (Core Deviation for Draft auctions). The core deviation for draft auctions b'_i for player i with a constraint-homogeneous valuation \widehat{v} with interest set S and per-unit value \widehat{v} is defined as follows.

Let $b_i^* = \frac{\widehat{v}}{2}$. In every auction t , she submits $b'_{it} = \max\{b_i^*, b_{it}\}$. If she wins with bid b_i^* , she buys $s^* - k_{i, < t}$ units of S and drops out. If she wins with a bid of b_{it} , she buys what she did under b_i : k_{it} units together with any other items she was buying under strategy profile b_i at auction t . She continues to bid b'_{it} until she acquires s^* units or the number of units remaining are not sufficient for her to complete s^* units.

The crucial observation is this: as long as the player hasn't already acquired s^* units, she has not affected the game path created by strategy b_i in any way. From the perspective of the other bidders, she behaved exactly as under b_i , by winning at her price under b_i and getting the items she would have got under b_i . If she ever wins at a higher price, she acquires all the units

¹⁴If the deviation were for the bidder to buy all the right number of units when she won because of her equilibrium bid, she might pay too much for them.

needed to reach s^* units in that auction and then drops out. Thus the prices that she faces in all the auctions prior to having won s^* units are the same as the prices under strategy b_i .

The Core Deviation Lemma for draft auctions follows immediately from Lemmas 2.6.6, 2.6.7, and 2.6.8.

Lemma 2.6.6. *If player i wins at least s^* units of S under the Core Deviation for draft auctions b'_i then*

$$u_i(b'_i(b_i), \mathbf{b}_{-i}; v_i) \geq \frac{1}{2}s^*\widehat{v} - P_i(\mathbf{b}).$$

Proof. If player i wins at least s^* units of S under b'_i then the valuation for the items she wins is at least $s^*\widehat{v}$. For the auctions in which she wins with a bid of b_{it} she pays a total amount of at most $P_i(\mathbf{b})$ and for the (at most one) auction she wins with a bid of b_i^* she pays at most $s^*b_i^*$. So her total payment is at most $s^*b_i^* + P_i(\mathbf{b}) = s^*\frac{\widehat{v}}{2} + P_i(\mathbf{b})$. ■

Lemma 2.6.7. *If player i wins fewer than s^* units of S under the Core Deviation for draft auctions b'_i then*

$$u_i(b'_i(b_i), \mathbf{b}_{-i}; v_i) \geq \frac{1}{2}s^*\widehat{v} - \sum_{j \in S} p_j - P_i(\mathbf{b}).$$

Proof. Consider the auction under the original strategy profile \mathbf{b} . Let (by an abuse of notation) $p_1 \leq p_2 \leq \dots \leq p_{|S|}$ be the prices at which the items in S are sold under \mathbf{b} . This is not necessarily the order in which they are sold. We show in Lemma 2.6.8 that, when bidder i wins fewer than s^* units under b'_i , it must be that $p_{s^*} \geq \frac{\widehat{v}}{2}$. Using this we obtain that

$$\sum_{j \in S} p_j \geq \sum_{l=s^*}^{|S|} p_l \geq (|S| - s^* + 1)p_{s^*} \geq s^*p_{s^*} \geq \frac{\widehat{v}}{2}s^*, \quad (2.12)$$

where we also used the simple observation that $s^* \leq \frac{|S|+1}{2}$.

The total payment of player i under b'_i in this case where she wins fewer than $|S|/2$ units of S is at most $P_i(\mathbf{b})$, therefore her utility is (trivially) at least $-P_i(\mathbf{b})$. The lemma now follows from adding the inequalities $u_i(b'_i(b_i), \mathbf{b}_{-i}; v_i) \geq -P_i(\mathbf{b})$ and $0 \geq \frac{\widehat{v}}{2}s^* - \sum_{j \in S} p_j$ (which holds by inequality (2.12)). ■

Lemma 2.6.8. *If player i wins fewer than s^* units of S under the Core Deviation b'_i then the s^* -th lowest price of the units in S under \mathbf{b} , is at least $\widehat{v}/2$.*

Proof. First, observe that if player i was obtaining at least s^* units under \mathbf{b} then she is definitely winning s^* units under b'_i , since she is always bidding at least as high. So, we can assume that under \mathbf{b} player i wins fewer than s^* units.

Recall that $p_1 \leq p_2 \leq \dots \leq p_{|S|}$ are the prices at which the units in S are sold under b . Let P_t be the price of auction t (under b). Let t^* be the first auction that was won at price $P_{t^*} \leq p_{s^*}$ under b but not by bidder i . We know that such an auction must exist; under \mathbf{b} there are s^* units of S that are sold at a price at most p_{s^*} , and since player i wins fewer than s^* of them, some of them are not won by player i .

We now argue that player i is still bidding in auction t^* under b'_i . First of all, she has not won s^* units prior to t^* . The other condition needed for her to be active is that there are at least $s^* - k_{i, < t^*}$ units available for sale in that auction. This follows from the fact that for any auction $t < t^*$ for which $P_t \leq p_{s^*}$, we know that player i was winning under b_i . Thus every unit that was sold prior to t^* at a price of less than or equal to p_{s^*} was sold to player i . There are s^* units sold at a price $\leq p_{s^*}$ and the number of such units sold prior to t^* is at most the number of total units won by bidder i prior to t^* . Thus the number of available units available at t^* is at least: $s^* - k_{i, < t^*}$.

Finally, we argue that $P_{t^*} \geq b_i^*$. Suppose for the sake of contradiction that $P_{t^*} < b_i^*$. Then player i wins auction t^* . Since she was not winning t^* under b_i , it must be that she is winning t^* with a bid of b_i^* . Thus in that auction she will buy every unit needed to reach s^* units. By the analysis in the previous paragraph, we know that there are still enough units available for sale to reach s^* . Thus in this case she will win s^* items, a contradiction with the main assumption of the Lemma. Therefore, $b_i^* \leq P_{t^*}$ and by definition, $P_{t^*} \leq p_{s^*}$ and $b_i^* = \frac{\widehat{v}}{2}$. ■

An easy corollary of the above core deviation lemma is that when all players have constraint-homogeneous valuations, the draft auction is a $(\frac{1}{4}, 2)$ -smooth mechanism, and thus has a price of anarchy of at most 8 for these valuations.

Chapter 3

A Measure of Simplicity for Auctions: the Fat-Shattering and Pseudo-dimension of Revenue Maximization

3.1 Introduction

In this chapter, we explore two natural measures of complexity of different classes of revenue-maximizing auctions for single-parameter settings, borrowed from learning theory, both of which govern the sample complexity of optimizing for revenue over a given class. The results in this section suggest that, in many cases, our intuition for the “simplicity” of a class of auctions directly coincides with a formal notion of simplicity: namely, simpler auctions have smaller pseudo-dimension, or fat-shattering dimension, (and thus sample complexity) than more complex auctions.

This work directly contributes to a recent line of inquiry on the sample complexity of revenue maximization, with many separate investigations of questions of the form

“Suppose bidders are drawn from some distribution \mathcal{D} . How many draws from \mathcal{D} are needed to learn an auction which will approximately maximize our expected revenue on future draws from \mathcal{D} ?”

The answer, of course, depends upon the distribution \mathcal{D} (e.g., whether or not \mathcal{D} is product, whether bidders’ distributions are regular, monotone-hazard rate, continuous, of bounded support), the environment (e.g., whether there is a single item, k -unit, matroid, or downwards closed), the closeness of approximation desired, and, most importantly, upon *what auctions are being decided between*. If the class of auctions is very restricted, few samples are needed to choose the optimal auction from that class, but the optimal auction in that class won’t, in general, have very good revenue. Conversely, very expressive classes of auctions will always contain *some* auction whose revenue guarantee is very good, but knowing *which* auction will have great revenue on future samples will require a large number of samples.

Our work describes how to use the pseudo-dimension and fat-shattering dimension to get distribution-free sample complexity bounds for a number of standard single-dimensional auction classes: pricing and VCG with anonymous and nonanonymous reserves (Section 3.4.2), nonde-

creasing virtual valuation maximizers 3.4.3 in several single-parameter settings (single-item, digital goods, k -unit, and so forth). We then shift our perspective and consider the class of auctions \mathcal{C} as a design parameter that can be selected by the seller. We design a new class of auctions, which we dub t -level auctions, which has nearly ideal tradeoff between ability to maximize revenue (its representation error) and its pseudo-dimension (and thus, its sample complexity and generalization error).

This work's formal definition of simplicity (namely, an auction class's pseudo-dimension) also contributes to recent work on designing simple, or detail-free, auctions for revenue maximization (motivated by the Wilson Doctrine [130], e.g., [35], [73], [36], [41], [123],[131],[10]). While all of these papers' auctions rule out Myerson's auction as a simple auction, none of them gives a measure by which one could directly decide which of the several auctions proposed in this literature was simplest. The work in this chapter offers a formal definition of simplicity for revenue maximization: namely, the pseudo-dimension of the class of auctions being considered. Just as Chapter 2 suggests that the *learnability* of equilibria in non-truthful combinatorial auctions is a reasonable litmus test for their simplicity, this chapter suggests *polynomial pseudo-dimension* of a class of revenue-maximizing auctions as a test of their simplicity.

3.2 Related Work

Traditional auction theory [94, 99, 111, 129] studies the problem of designing a way to sell an item without knowing buyers' willingness to pay. Maximizing the revenue of this selling procedure with zero information about the buyers is impossible, even approximately: if there is a single buyer with some large, unknown value for the item, it is impossible to extract a bounded fraction of that value. The classical study circumvents this issue by modeling buyers' values for the item as *private* draws from *publicly known* distributions; and the goal of the seller is to maximize her *expected* revenue, over the draw of private values from the distributions. The revenue-optimal, incentive-compatible, single item auction in this Bayesian model is known as Myerson's auction [99]. This auction uses detailed information about each bidder's distribution to map values to *virtual* values, and chooses the winner with the highest non-negative virtual value. This approach extends to any single-parameter setting, where each bidder's value can be described by a single number, their value for being included in the winning set (e.g., matroids, k -unit auctions).

In general, the only assumption made about the mapping from values to virtual values is that it is nondecreasing¹, or increasing at a certain rate. With only this assumption, Myerson's auction can be quite complex: winners might have much lower value than losers, the pointwise revenue of a given tuple of valuations may decrease when adding a new bidder, and small changes in the distributions of valuations may lead to large changes in the mappings. This has led the theoretical computer science community to consider a long line of *prior-free* auctions that is, auctions which are constructed with no prior information about the distribution, and work on *simple* auctions for revenue maximization (e.g., Chawla et al. [35], Hartline and Roughgarden [73], Chawla et al. [36], Devanur et al. [41], Roughgarden et al. [123], Yao [131], Babaioff et al. [10]) that relied in a

¹Or, that one has ironed the function so that this is true. See Hartline [72] for more details on ironing.

limited way on the distributions. Our work offers a formal definition of “simplicity” for revenue-maximizing auctions: namely, the pseudo-dimension of the class of auctions. This definition allows the auction designer to trade off between the simplicity, or sample complexity, of the class of auctions, and possible revenue guarantees.

The work in this vein most closely related to ours, Balcan et al. [12] (following Blum et al. [20] and Blum and Hartline [18]) pointed out that statistical learning theory provides a framework for designing near-optimal incentive-compatible auctions. Intuitively, their framework optimizes over a given class of auctions F by first randomly partitioning the bidders into two sets S_1, S_2 , picking the auctions $a_1, a_2 \in F$ which are revenue-optimal for S_1, S_2 , and applying a_1 to S_2 (and a_2 to S_1). This partitioning procedure ensures the incentive-compatibility of the resulting auction. If the number of bidders, or value of the revenue-optimal auction from the class, are large enough, this auction $1 - \epsilon$ -approximates the revenue of the *best* auction $a^* \in \mathcal{C}$, *for these bidders with these values*.

Their work and ours differs in several key ways. First, we assume that buyers are drawn from some fixed, unknown distribution F ; they make no distributional assumption. As a result, we achieve an approximation to the maximum expected revenue achieved by any auction in \mathcal{C} ; they instead approximate the best revenue achievable by any auction in \mathcal{C} *for that instance*. Finally, we give results for arbitrary single-parameter settings, in particular those with small supply, such as single-item and k -unit auctions, for small k . Their work focuses on large or unlimited supply settings. This is perhaps the biggest limitation: their auctions are feasible only in settings where a large number of items may be sold to each of S_1 and S_2 ².

The idea of learning from samples is used in the work of Balcan et al. [12] by the internal randomness of their mechanisms, rather than through an exogenous distribution over inputs (as in this work). Similar to our sample complexity bounds, the sufficient condition on the number of bidders in their work depends polynomially on ϵ^{-1} , the largest possible bid, and a measure of the complexity of the class \mathcal{G} that plays a role analogous to our use of the pseudo-dimension. We consider the fact that our work relies on the pseudo-dimension of auction classes as a feature: the sample complexity bounds follow directly from this measure, once we have done the problem-specific work of bounding the pseudo-dimension and representation error of well-chosen auction classes — rather than having to prove them from scratch.

Other highly relevant work includes that of Elkind [53], who studied learning Myerson-optimal mechanisms for bidders whose valuation distributions have finite support sets of size K : as in the continuous setting, they show the revenue-optimal auction maximizes virtual welfare for single-item auctions. They then construct an efficient learning algorithm for this setting, which roughly reduces to sorting the nK possible bids into an order of allocation precedence which approximately optimizes revenue. While this work shares some of the same goals of this paper: namely, understanding the sample complexity of finding a revenue-optimal auction, our work differs from theirs in several ways. First, our results extend to matroid and even general single-parameter settings. Second, we find a class of auctions with small sample complexity and show it contains a nearly-optimal auction for continuous distributions, even though the exactly-optimal auction may not be contained in that class. Third, our bounds on the sample complexity of learning K -level auctions improves upon theirs (proved from scratch rather than using out-of-

²See Balcan et al. [11] for an analysis of the case with large but limited supply.

the-box learning theory) for bidders with support of size K in the single-item setting by a factor of roughly nK ; moreover, we only need $K^{\frac{1}{\epsilon}} + \log_{1+\epsilon} H$ to approximate the optimal auction, so our work shows one can trade the second factor of K for $\text{poly}(\frac{1}{\epsilon}, H)$. Finally, our multiplicative error guarantees require a careful treatment of near-zero virtual values, which isn't directly implied by their additive approximations.

We next place our work within the context of recent papers studying the sample complexity of learning near-optimal single-parameter auctions. Dhangwatnotai et al. [43] show how, with k bidders drawn from each distribution, using one bidder from each distribution as a reserve achieves a $\frac{1}{4} \frac{k-1}{k}$ -approximation to the optimal auction if the setting is either downwards closed and MHR, or if it is a matroid with regular bidders. Similarly, if $k = \Omega(\text{poly}(\frac{1}{\epsilon}))$, these approximation factors can be improved by using a larger set of bidders to pick a reserve price. Essentially, this work shows that using some number of bidders from each distribution as samples allows one to approximately learn monopoly reserve prices.

Cole and Roughgarden [39] study single-item auctions with n bidders with valuations drawn from non-identical, independent “regular” distributions (see Section 3.3), and prove upper and lower bounds (polynomial in n and ϵ^{-1}) on the sample complexity of learning a $(1-\epsilon)$ -approximate auction. While the formalism in their work is directly inspired by learning theory, no formal connections were offered; in particular, both the upper and lower bounds (which are far from matching) were proved from scratch. Our positive results include single-item auctions as a very special case. Even in this case, for bounded or MHR valuations, our sample complexity upper bounds are much better than those in Cole and Roughgarden [39].

Roughgarden and Schrijvers [120] extend the positive results of Cole and Roughgarden [39] beyond single-item auctions and regular distributions, to position auctions and arbitrary matroid environments and to a wide class of irregular distributions, via a new learning algorithm. While the sample complexity upper bounds proved in the present work certainly cover the settings of Roughgarden and Schrijvers [120] — in addition to arbitrary single-parameter settings and arbitrary bounded irregular distributions — the specialized analysis in the latter offers orthogonal advantages, including computational tractability.

Work done simultaneously to ours studies the problem of optimally setting reserves in a second-price auction [96], bounding the sample complexity of this problem as a function of both the pseudo-dimension and Rademacher complexity. Their paper offers a computationally efficient algorithm for this problem, using a more stylized analysis. However, the work does not suggest how one could use the results therein to design near-optimal auctions, and their work is quite specialized to the particular problem of selecting anonymous reserves.

Huang et al. [78] consider learning the optimal price from samples when there is a single buyer and a single seller; this problem was also studied implicitly in [43]. Our general positive results obviously cover the bounded-valuation and MHR settings in [78], though the specialized analysis in [78] yields better (indeed, almost optimal) sample complexity bounds (as a function of ϵ^{-1} and/or H). Huang et al. [78] also develop a methodology for proving nearly tight lower bounds on sample complexity, a topic barely touched on here.

The main setting studied by Dughmi et al. [47] is that of a single unit-demand player, who will buy at most 1 of m available items. The main result in their work for this multi-parameter setting is negative, in the form of an exponential (in m) lower bound on the sample complexity required to compute a constant-factor approximation of the optimal auction with constant probability.

Their work also considers restricting to a class \mathcal{C} of auctions such that each can be described using c bits (equivalently, $|\mathcal{C}| \leq 2^c$), and note that a $(1 - \epsilon)$ -approximate auction from \mathcal{C} can be learned from a number of samples that is polynomial in H , ϵ^{-1} , and c . Since the pseudo-dimension of a set with finite cardinality $|\mathcal{C}|$ is always at most $\log_2 |\mathcal{C}|$, this positive result can be viewed as a simple special case of the result Theorem 3.3.1 that we employ in our work. They also give sub-exponential sample complexity bounds for some special cases of the problem. Because we consider single-parameter settings, rather than multi-dimensional auctions, the lower bounds in their setting do not apply and we can achieve polynomial sample complexity quite generally.

In Chapter 4, we describe in more detail the problem of learning bidders' valuation distributions from partial observation. That work considers a setting in which samples from each bidder's distribution are not readily available. Instead, the learner has a more limited history, such as the identity of the winner in many previous auctions, from which to learn about bidders. The literature on learning from partial or censored information (e.g., [9, 22, 33, 37, 64, 87]) is morally related to the work in this chapter, though its goal is mostly to reconstruct bidders' distributions from polynomially many *censored* samples, rather than to directly maximize revenue from polynomially many "complete" samples.

3.3 Preliminaries

3.3.1 Bayesian Auction Design

This section reviews useful terminology and notation standard in Bayesian mechanism design and the sample complexity literature. We consider *single-parameter settings* with n bidders, and a collection of feasible winning sets of bidders $\mathcal{X} \subseteq 2^{[n]}$. If, for any $X \in \mathcal{X}$ and $Y \subseteq X$, $Y \in \mathcal{X}$, we say \mathcal{X} is *downward-closed*. If \mathcal{X} is downwards closed and, for two sets $|I_1| < |I_2|$, $I_1, I_2 \in \mathcal{X}$, there is always an augmenting element $i_2 \in I_2 \setminus I_1$ such that $I_1 \cup \{i_2\} \in \mathcal{X}$, \mathcal{X} is called a *matroid*. A simple example of a matroid environment is k -unit auctions with unit-demand bidders (then, \mathcal{X} would contain all subsets of $[n]$ of size at most k). We often illustrate our main ideas using single-item auctions (where \mathcal{X} contains the singletons and empty set).

Each bidder i has some value v_i for the outcome X when $i \in X$ and value 0 if $i \notin X$. A mechanism \mathcal{A} selects a winning set X and a payment p , where p_i is the payment for bidder i . We only consider mechanisms which are ex-post *individually rational (IR)*, so if $i \notin X$, then $p_i = 0$ and if $i \in X$, $p_i \leq v_i$. Bidders' utilities u_i are assumed to be quasi-linear in money: their utility for being in the winning set is $v_i - p_i$, or 0 for $i \notin X$ (since the mechanisms are IR, $p_i > 0$ only if $i \in X$). We will further restrict ourselves by only considering *incentive-compatible (IC)* mechanisms, if $u_i(\mathcal{A}(v_i, b_{-i})) \geq u_i(\mathcal{A}(b_i, b_{-i}))$, for all agents i , all values v_i , and all behavior b_i, b_{-i} . If, for all i and b_{-i} , $\mathcal{A}(b_i, b_{-i}) = (X, p)$ has the property that $i \in X$ if and only if $b_i > p_i$, then \mathcal{A} is clearly IC³. Since we only consider incentive-compatible mechanisms, we assume agents bid truthfully for the remainder of this chapter.

We assume bidders' valuations v are drawn from the continuous joint cumulative distribution F . Except in the extension in Section 3.6, we assume that the support of F is limited to $[1, H]^n$.

³Each agent effectively chooses whether to pay p_i and win, or pay 0 and lose. If her value is above p_i , she's willing to pay p_i , and otherwise she isn't.

It is common to assume that F is product, with $F = F_1 \times F_2 \times \dots \times F_n$ and each $v_i \sim F_i$ drawn independently but not identically; we will mention explicitly where independence is needed for some results. The *virtual value* of bidder i is denoted as $\phi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$. A distribution satisfies the *monotone-hazard rate (MHR)* condition if $f_i(v_i)/(1 - F_i(v_i))$ is nondecreasing; intuitively, if its tails are no heavier than those of an exponential distribution.

The auction which maximizes expected revenue for n Bayesian bidders chooses winners in a way which maximizes the (sum of the) virtual value of the winner(s). This auction is known as Myerson’s auction, which we will refer to as \mathcal{M} . When ϕ_i is not weakly increasing (i.e., is not *regular*), the *ironed* virtual value $\bar{\phi}_i(v_i)$, which is monotone, can be used (for details, see Hartline [72]). For settings where agents are not assumed to be regular, running \mathcal{M} with respect to the ironed virtual valuation functions is the revenue-optimal over all incentive-compatible mechanisms.

VCG refers to the mechanism which allocates to the set X such that $\sum_{i \in X} v_i$ is maximized. Suppose now that we are in a downwards-closed environment. Given n real values q_1, \dots, q_n , VCG with eager reserves refers to the mechanism which picks X such that each bidder $i \in X$ has value $v_i \geq q_i$ which maximizes welfare, e.g. $X = \operatorname{argmax}_{X'} \sum_{i \in X': v_i \geq q_i \forall i \in X'} v_i$; VCG with lazy reserves chooses the set X which maximizes welfare but only allocates to $X' \subseteq X$ where $i \in X'$ only if $v_i \geq q_i$, e.g. $X' = \{i : v_i \geq q_i \wedge i \in X : X = \operatorname{argmax}_{X''} \sum_{i' \in X''} v_{i'}\}$ ⁴. In both cases, only bidders who pass their reserves ever win. VCG with eager reserves maximizes the sum of the values of winning (and reserve-beating) bidders, while VCG with lazy reserves picks a preliminary winning set based on the sum of its (possibly non-reserve-passing) values, then kicks out any bidders who don’t pass their reserves. These two auctions are only equivalent when $q_i = q_j$ for all i, j , when the reserves are *anonymous*, and differ when the reserves are non-anonymous.

Notice that all of these auctions are incentive-compatible, if their payment rule is chosen appropriately: namely, each winner is charged the minimum bid such that they would be included in the winning set.

3.3.2 Sample Complexity, VC Dimension, and the Pseudo-Dimension

For completeness, in this section we review several well-known definitions from learning theory. Suppose there is some domain \mathcal{Q} , and let c be some unknown target function $c : \mathcal{Q} \rightarrow \{0, 1\}$. Assume there is some unknown distribution \mathcal{D} over \mathcal{Q} . We wish to understand how many labeled samples $(x, c(x))$, $x \sim \mathcal{D}$, are necessary and sufficient to be able to output a \hat{c} which agrees with c almost everywhere on \mathcal{D} . The distribution-independent sample complexity of learning c depends fundamentally on the “complexity” of the set of binary functions \mathcal{C} from which we are choosing \hat{c} . We define the relevant complexity measure next.

Let S be a set of m samples from \mathcal{Q} . The set S is said to be *shattered* by \mathcal{C} if, for every subset $T \subseteq S$, there is some $c_T \in \mathcal{C}$ such that $c_T(x) = 1$ if $x \in T$ and $c_T(y) = 0$ if $y \notin T$. That is, ranging over all $c \in \mathcal{C}$ induces all $2^{|S|}$ possible projections onto S . The VC-dimension of \mathcal{C} , denoted $\mathcal{VC}(\mathcal{C})$, is the size of the largest set S which can be shattered by \mathcal{C} .

Let $\operatorname{err}_S(\hat{c}) = (\sum_{x \in S} |c(x) - \hat{c}(x)|)/|S|$ denote the empirical error of \hat{c} on S , and let $\operatorname{err}(\hat{c}) =$

⁴VCG with lazy reserves is only a feasible mechanism in downwards-closed environments.

$\mathbb{E}_{x \sim \mathcal{D}}[|c(x) - \widehat{c}(x)|]$ denote the *true* expected error of \widehat{c} with respect to \mathcal{D} . A key result from learning theory [128] is: for every distribution \mathcal{D} , a sample S of size $\Omega(\epsilon^{-2}(\mathcal{VC}(\mathcal{C}) + \ln \frac{1}{\delta}))$ is sufficient to guarantee that $\text{err}_S(\widehat{c}) \in [\text{err}(\widehat{c}) - \epsilon, \text{err}(\widehat{c}) + \epsilon]$ for every $\widehat{c} \in \mathcal{C}$ with probability $1 - \delta$. In this case, the error on the sample is close to the true error, simultaneously for every hypothesis in \mathcal{C} . In particular, choosing the hypothesis with the minimum sample error minimizes the true error, up to 2ϵ . We say \mathcal{C} is (ϵ, δ) -uniformly learnable with sample complexity m if, given a sample S of size m , with probability $1 - \delta$, for all $c \in \mathcal{C}$, $|\text{err}_S(c) - \text{err}(c)| < \epsilon$: thus, any class \mathcal{C} is (ϵ, δ) -uniformly learnable with $m = \Theta\left(\frac{1}{\epsilon^2} (\mathcal{VC}(\mathcal{C}) + \ln \frac{1}{\delta})\right)$ samples.

Conversely, for every learning algorithm \mathcal{A} that uses fewer than $\frac{\mathcal{VC}(\mathcal{C})}{\epsilon}$ samples, there exists a distribution \mathcal{D}' and a constant q such that, with probability at least q , \mathcal{A} outputs a hypothesis $\widehat{c}' \in \mathcal{C}$ with $\text{err}(\widehat{c}') > \text{err}(\widehat{c}) + \frac{\epsilon}{2}$ for some $\widehat{c} \in \mathcal{C}$. That is, the true error of the output hypothesis is more than $\frac{\epsilon}{2}$ larger the best hypothesis in the class.

To learn real-valued functions, we need a generalization of VC dimension (which concerns binary functions). The *pseudo-dimension* [110] does exactly this. Formally, let $c : \mathcal{Q} \rightarrow [0, H]$ be a real-valued function over \mathcal{Q} , and \mathcal{C} the class we are learning over. Let S be a sample drawn from \mathcal{D} , $|S| = m$, labeled according to c . Both the empirical and true error of a hypothesis \widehat{c} are defined as before, though $|\widehat{c}(x) - c(x)|$ can now take on values in $[0, H]$ rather than in $\{0, 1\}$. Let $(r^1, \dots, r^m) \in [0, H]^m$ be a set of *targets* for S . We say (r^1, \dots, r^m) *witnesses* the shattering of S by \mathcal{C} if, for each $T \subseteq S$, there exists some $c_T \in \mathcal{C}$ such that $f_T(x^i) \geq r^i$ for all $x^i \in T$ and $c_T(x^i) < r^i$ for all $x^i \notin T$. If there exists some \vec{r} witnessing the shattering of S , we say S is *shatterable* by \mathcal{C} . The *pseudo-dimension* of \mathcal{C} , denoted $d_{\mathcal{C}}$, is the size of the largest set S which is shatterable by \mathcal{C} .

There are certain cases in which the pseudo-dimension of a class \mathcal{C} is unbounded, but PAC-style sample complexity bounds are still possible to achieve, through the use of the γ *fat-shattering dimension* [85]. This measure leverages that we need not learn the value of the real-valued functions exactly, but only up to additive error on most of their domain. Again, let S be a set of samples. We say (r^1, \dots, r^m) *witnesses* the γ -shattering of S by \mathcal{C} if, for each $T \subset S$, there exists some $c_T \in \mathcal{C}$ such that $f_T(x^i) \geq r^i + \gamma$ if $x^i \in T$ and $c_T(x^i) \leq r^i - \gamma$ if $x^i \notin T$. If there exists some \vec{r} which witnesses the γ -shattering of S , we say S is γ -shatterable by \mathcal{C} . The γ fat-shattering dimension of \mathcal{C} , denoted $\text{fat}_{\mathcal{C}}(\gamma)$, is the size of the largest set S which is γ -shatterable by \mathcal{C} . Notice that $\text{fat}_{\mathcal{C}}(\gamma)$ is a function which decreases as γ increases, and $\text{fat}_{\mathcal{C}}(0) = d_{\mathcal{C}}$, so the pseudo-dimension is always an upper bound on the fat-shattering dimension.

The sample complexity upper bounds in this chapter are derived from following two theorems, which states that the distribution-independent sample complexity of learning over a class of real-valued functions \mathcal{C} can be bounded in terms of the class's pseudo-dimension or γ fat-shattering dimension. We give slightly tighter bound here than given in Anthony and Bartlett [4], which can be attained by rescaling the functions to $[0, 1]$ and $(\frac{\epsilon}{H}, \delta)$ learning the scaled functions.

Theorem 3.3.1 (Anthony and Bartlett [4]). *Suppose \mathcal{C} is a class of real-valued functions over $[0, H]$ with pseudo-dimension $d_{\mathcal{C}}$. For $\epsilon, \delta \in [0, 1]$, the sample complexity of (ϵ, δ) -uniformly learning f with respect to \mathcal{C} is*

$$m = O\left(\left(\frac{H}{\epsilon}\right)^2 \left(d_{\mathcal{C}} \ln\left(\frac{H}{\epsilon}\right) + \ln\left(\frac{1}{\delta}\right)\right)\right).$$

In the case that the pseudo-dimension is unbounded, or large, one can also use the fat-shattering dimension to achieve similar sample complexity bounds, as in the result stated below.

Theorem 3.3.2 (Anthony and Bartlett [4]). *Suppose \mathcal{C} is a class of real-valued functions over $[0, H]$ with γ fat-shattering dimension $\text{fat}_{\mathcal{C}}(\gamma)$. For $\epsilon, \delta \in [0, 1]$, the sample complexity of (ϵ, δ) -uniformly learning f with respect to \mathcal{C} is*

$$m = O \left(\left(\frac{H}{\epsilon} \right)^2 \left(\text{fat}_{\mathcal{C}} \left(\frac{\epsilon}{H} \right) \ln^2 \left(\frac{H}{\epsilon} \right) + \ln \left(\frac{1}{\delta} \right) \right) \right).$$

As Theorems 3.3.1 and 3.3.2 are true simultaneously for all functions in \mathcal{C} , their guarantees are realized by the learning algorithm that simply outputs the function $c \in \mathcal{C}$ with the smallest empirical error on the sample.

3.3.3 Applying Pseudo-Dimension to Auction Classes

For the remainder of this chapter, we consider classes of truthful auctions \mathcal{C} . When we discuss some auction $\mathcal{A} \in \mathcal{C}$, we treat $\mathcal{A} : [0, H]^n \rightarrow \mathbb{R}$ as the function which maps (truthful) bid tuples to the revenue achieved on them by the auction \mathcal{A} . Then, rather than looking to *minimize* error, we want to *maximize revenue*. In our setting, the guarantee of Theorem 3.3.1 directly implies that, with probability at least $1 - \delta$ (over the m samples), the output of the *empirical revenue maximization* learning algorithm — which returns the auction $\mathcal{A} \in \mathcal{C}$ with the highest average revenue on the samples — chooses an auction with expected revenue (over the true underlying distribution F) is within an additive ϵ of the maximum possible.

3.4 The sample complexity of some common auctions

3.4.1 A Warm-up: the Pseudo-Dimension and Fat-Shattering Dimension of all Incentive-Compatible Mechanisms

In this section, we cut our teeth with the definitions in Section 3.3 using them for a simple task: showing that the class of all auctions is quite complicated. In particular, no finite sample complexity guarantees can be achieved for the class of all auctions (and, if one discretizes the bid-space $[1, H]^n$ into an ϵ -mesh, one will need $\tilde{\Omega} \left(\left(\frac{H}{\epsilon} \right)^n \right)$ samples for good generalization results). Similar results were already known (for example, see [47]), but it is instructive to prove such results using the tools we need for the remainder of this chapter. Nor are these divergent from our intuition: the class of all incentive-compatible auctions is very large. In particular, if we do not restrict the size of their description, they contain any number of hairy rules which can exactly memorize a given set of samples.

We now present the main theorem for this section. We mention that it is not difficult to translate these guarantees to *anonymous* incentive-compatible auctions (whose behavior is blind to the precise identity of any bidder), but present this version for ease of exposition.

Theorem 3.4.1. *The class of all single-item incentive-compatible auctions has infinite pseudo-dimension (and γ fat-shattering dimension). The class of all single-item incentive-compatible*

auctions on the discretized bid-space into an ϵ -mesh has γ fat-shattering dimension $\Omega\left(\left(\frac{H}{\epsilon}\right)^{n-1}\right)$, for $H > 1 + 2\gamma$ and $\gamma < 1$.

Proof. We start by showing that the γ fat-shattering dimension of the non-discretized class is infinite, even with two bidders. We construct a set \mathcal{S} of samples whose size is infinite and can be γ -shattered by the class of all auctions. Let $\mathcal{S} = \{(x, x + 2\gamma) | x \in [1, H - 2\gamma]\}$ be a set of samples, with revenue target $r^x = x + \gamma$ for the sample $(x, x + 2\gamma) \in \mathcal{S}$.

We now show how to γ -shatter this set. Consider some $T \subseteq \mathcal{S}$. Define \mathcal{A}^T as follows. For any $x : (x, x + 2\gamma) \in T$, \mathcal{A}^T will sell to bidder 2 at price $x + 2\gamma$. For any $y : (y, y + 2\gamma) \notin T$ (including those points not in \mathcal{S}), \mathcal{A}^T sells to bidder 2 at price 0. This mechanism is trivially incentive-compatible: it always sells to bidder 2 at a price which is a function of bidder 1's bid alone. Moreover, the revenue target for each sample in T is surpassed by γ , and missed by at least $1 > \gamma$ for each sample not in T . Since this auction is trivially incentive-compatible (bidder 2's price does not depend upon her bid), we have shown it is possible γ -shatter a set of infinite size with the class of all incentive-compatible auctions. Thus, the γ fat-shattering dimension of this class is infinite.

Now, we prove an analogous bound for the discretized bid-space version of this auction with n bidders. Consider the set $\mathcal{S} = \{(\mathbf{v}, H) | \mathbf{v} \in \{k\epsilon | 0 \leq k \leq \lfloor \frac{H}{\epsilon} \rfloor\}^n\}$ with each revenue target being $H - \gamma$ for each sample. An identical argument as from the non-discretized case shows one can shatter \mathcal{S} . In this case, $|\mathcal{S}| = \Omega\left(\left(\frac{H}{\epsilon}\right)^{n-1}\right)$, so the lower bound follows. ■

The lower bound in terms of fat-shattering dimension implies that, without further assumptions on the distribution of bids, no polynomial sample-complexity bounds will be possible for this class. While the particular example is very much not compatible with valuations coming from a product distribution, one can construct a similar example with 2 bidders' values being drawn independently from two continuous, regular distributions: the class of all auctions can still "memorize" the sample to shatter it.

3.4.2 Pricing, VCG with reserve prices

We now show how one can reason about the fat-shattering and pseudo-dimension of several more commonly employed auction classes: those that set a single threshold bid than those that set an individualized threshold bid for each bidder. These classes of auctions are, as one would imagine, strictly simpler than the class of all auctions; the set of auctions which (a) sets a single threshold for all bidders and (b) chooses to sell to some bidder who surpasses this threshold has *constant* pseudo-dimension and thus sample complexity which scales only with $\frac{1}{\epsilon}$: intuitively, there is only one parameter to tune, and a constant number of samples suffices to tune it. The class of auctions which sets a minimum bid for each bidder separately, on the other hand, has pseudo-dimension which grows linearly in n , the number of bidders, where again our intuition correctly places this class as more complicated than anonymous threshold auctions, but substantially simpler than the class of all auctions. We begin our analyses with single-item auctions, before extending the reasoning to k -unit, digital goods, and matroid auctions.

Our analysis holds for any of the following allocation rules according to bids \mathbf{v} , where W is the set of bidders who pass their threshold(s):

1. The winner is selected from W according to some fixed linear ordering over bidders.
2. The winner is selected uniformly at random amongst W .
3. The winner is the highest-valued bidder amongst bidders in W according to \mathbf{v} .
4. The winner is the highest-valued bidder, if she is in W .

Rules 1 and 2 treat the winning threshold(s) as a *price* (or prices, in the case of individualized thresholds). The former has some fixed way to choose who will win, the latter chooses the winner at random. In both cases, the payment is just the threshold (of the winner, in the case of nonanonymous thresholds). Rules 3 and 4 act like VCG with reserves: rule 3 corresponds to *eager* reserves, while rule 4 corresponds to *lazy* reserves. In the single-item setting with a single reserve, these are equivalent: the winner is the highest-valued bidder amongst those in W , if W is nonempty (otherwise, no one wins). When there are multiple reserves, however, these allocation rules are not equivalent: the highest-valued bidder may or may not be in W , even if W is nonempty, because she may not pass *her* reserve, which might be higher than some other (lower-valued) bidder's reserve. That highest-valued bidder is the only bidder that Rule 4 will choose as a winner, while Rule 3 will choose another winner from W , if the highest-valued bidder isn't in W but W is nonempty.

We now present the formal results for single-item auctions, both for anonymous and nonanonymous prices and reserves.

Theorem 3.4.2. *Any class of single-item auctions which sets a single price and uses some fixed allocation rule of type 1, 2, 3, or 4 has pseudo-dimension $\Theta(1)$.*

Proof. Consider a fixed set of samples \mathcal{S} of size m which can be shattered by this class, with witness (r^1, \dots, r^m) . We will prove an upper bound on m . Let \mathcal{A}^x denote the auction whose reserve is x . Any allocation rule of type 1 or 2 has a payment for winning bidder $i \in W$ of the form $p_i(\mathbf{v}) = x$. An allocation rule of type 3 has payment rule for the winner $i \in W$ of the form $p_i(\mathbf{v}) = \max(x, \max_{i' \in W, i' \neq i} \mathbf{v}_{i'})$. In both cases, the payment rule (and thus the revenue) is monotonically nondecreasing in x , so long as $x \leq \max_i \mathbf{v}_i$.

So, consider the class of all auctions, parameterized by x . For a fixed sample $\mathbf{v}^j \in \mathcal{S}$, there are two values $x_1^j \leq x_2^j$ such that $\text{rev}(\mathcal{A}^x, \mathbf{v}^j) < r^j$ for all $x < x_1^j$, $\text{rev}(\mathcal{A}^x, \mathbf{v}^j) \geq r^j$ for all $x \in [x_1^j, x_2^j]$, and $\text{rev}(\mathcal{A}^x, \mathbf{v}^j) = 0$ for $x > x_2^j$. Thus, if we identify the class of all auctions with the real line, a given sample \mathbf{v}^j has some fixed interval $[x_1^j, x_2^j]$ in which it is labeled positive, outside of which it is labeled negative. Varying x can then yield at most $3m$ labelings of all samples, since superimposing all $2m$ of these points breaks the real line into $3m$ contiguous regions (and inside each, the labeling for *all* samples are fixed). Since \mathcal{S} is shatterable, it must be the case that $3m \geq 2^m$, or that $\log(m) \geq m$, implying $m = O(1)$. ■

Oddly enough, the direct analog of the previous proof does not guarantee that Rules 1, 2, or 3 have pseudo-dimension which is bounded like the Rule 4 when shift our attention to nonanonymous reserves. Those rules a more careful analysis, because their winner(s) and revenue depend upon W in a more complicated way than rule 4. Rule 4, for a fixed valuation tuple \mathbf{v} , always has the same winner $i = \text{argmax}_{i'} \mathbf{v}_{i'}$ (if any), who always pays $\max(x_i, \max_{i' \neq i} \mathbf{v}_{i'})$, where x_i is i 's reserve. The other rules don't have this property; the winner for a given set W might be any of the bidders in W . Their pseudo-dimension is also $\tilde{\Theta}(n)$, which is implied by Theorem 3.5.1,

though we do not analyze them directly in this section. We now proceed to prove the result for Rule 4.

Theorem 3.4.3. *The classes of single-item auctions which set an individualized price and uses allocation rule 4, that is, VCG with lazy non-anonymous reserves, has pseudo-dimension $\Theta(n)$ and fat-shattering dimension $\Theta(n)$.*

Proof. We begin by proving the upper bound. Again, consider a sample \mathcal{S} of size m which can be shattered with witnesses (r^1, \dots, r^m) . We will again upper-bound m by counting the number of possible labelings of \mathcal{S} can be achieved by the class of nonanonymous reserve auctions according to allocation rule 4. Let x_i refer to the reserve for bidder i for a fixed auction.

For a fixed sample $\mathbf{v}^j \in \mathcal{S}$, there is some fixed $i^j \in [n]$ such that, for any auction \mathcal{A} in this class, either i^j is the winner or there is no winner (the highest-valued bidder). Furthermore, her payment when she wins is $p_i = \max(x_i, \max_{i' \neq i} \mathbf{v}_{i'})$. Thus, the revenue for any auction on sample \mathbf{v}^j depends only upon the reserve of the highest bidder and the second-highest bid (which is fixed for a fixed sample). So, separate the samples \mathcal{S} into $\mathcal{S}_1, \dots, \mathcal{S}_n$ according to the identity of their highest-valued bidder. The revenue for a set \mathcal{S}_i is now dependent only upon the reserve x_i , and there are at most $|\mathcal{S}_i|$ labelings of the samples \mathcal{S}_i . Since we must be able to shatter each \mathcal{S}_i to shatter all of \mathcal{S} , it must be the case that $|\mathcal{S}_i| \geq 2^{|\mathcal{S}_i|}$, for all i , implying the size of each \mathcal{S}_i is a constant, or that $\sum_i |\mathcal{S}_i| = O(n)$.

We now prove the lower bound by exhibiting a set \mathcal{S} of size n which is shatterable by this class of auctions. For each $j \in [n]$, let $\mathbf{v}^j = \mathbf{e}_j$ be the j -th standard basis vector, and $r^j = 1 - \gamma$. Then, we need to show it is possible to γ -shatter \mathcal{S} . For a given subset $T \subseteq \mathcal{S}$, define \mathcal{A}^T as follows (it will be an auction to positively label T and negatively label all other samples). If $\mathbf{v}^j \in \mathcal{S}$, let $x^j = 1$ and 0 otherwise. Then, the revenue of \mathcal{A}^T on $\mathbf{v}^j \in T$ is 1 and 0 on $\mathbf{v}^{j'} \notin T$, so \mathcal{A}^T γ -separates T from $\mathcal{S} \setminus T$. ■

More General Feasibility Settings for Pricing and Reserves

We now extend these ideas to more general feasibility settings: namely, k -unit auctions, digital goods, and general downwards-closed settings. Our intuition suggests that the problem separates across each bidder in the digital goods case (namely, there is no feasibility constraint), and similarly for k -unit (the feasibility constraint is only easier to satisfy than in the single-item case). However, there are more possible choices for the auction to make when there are multiple items. Namely, in the single item case, there was either one or zero items sold: in the k -unit and digital goods settings, one can up to k (or n) copies, but can also sell fewer copies. This increase in the number of choices is reflected in the fat-shattering dimension of anonymous reserve auctions: in the single-item case, these auctions had constant fat-shattering dimension, while in the k -unit (and digital goods) setting, the fat-shattering dimension grows logarithmically in k (and n , respectively). The fat-shattering dimension of nonanonymous single-item auctions already has linear dependence on n ; a basic covering argument shows that this does result degrades only by a factor of $\frac{H}{\gamma}$ in more general feasibility settings.

Theorem 3.4.4. *The pseudo-dimension of the class of k -unit VCG with lazy anonymous reserves is $\Theta(\log k)$, $\Theta(\log n)$ for digital goods, and $\Theta(\log n)$ for general downwards-closed settings. These lower bounds also hold for γ fat-shattering, so long as $\gamma < \frac{H}{k}$ and $\gamma < \frac{H}{n}$, respectively.*

Proof. The upper bound is analogous to the single-item case. We will argue about the k -unit case, the proof is identical for digital goods and downwards-closed environments, replacing “ k ” by “ n ” or “ $\leq n$ ” everywhere. Consider a set \mathcal{S} with revenue targets (r^1, \dots, r^m) that can be shattered by this class. The only relevant bidders for each sample are those with the k highest bids. Sort the set of $(k+1)m$ relevant bids (the $k+1$ highest bids per sample), along with the mk numbers $\frac{r^j}{t}$ for all $t \in [k]$ and all j . Then, these $(2k+1)m$ numbers divide $[1, H]$ into as many regions. Consider one such region. Any auction with a reserve within this region labels all samples the same way as any auction with a different reserve within the same region. Thus, there are at most $(2k+1)m+1$ distinct labelings of the set of m points by this class of auctions. Since \mathcal{S} is shatterable, it must be the case that $(2k+1)m \geq 2^m$, implying $m = O(\log k)$.

For the lower bound, we construct a set \mathcal{S} and revenue targets (r^1, \dots, r^m) of size $\Omega(\log k)$ that can be γ -shattered by this class. Each sample will be a slightly perturbed version of

$$\left(r + \gamma, \frac{r + \gamma}{2}, \frac{r + \gamma}{3}, \dots, \frac{r + \gamma}{k}, 0, \dots, 0 \right).$$

Since $2^m = O(k)$, it will be possible to encode each subset of samples as one of k bidders. Now, order all subsets of \mathcal{S} as $T_1, T_2, \dots, T_{2^m=k}$. Then, for all samples $s \notin T_1$, subtract some small ϵ from their first coordinate, and similarly for $s \notin T_i$, subtract ϵ from their i th coordinate. Let this set of samples be \mathcal{S} . We now show one can shatter \mathcal{S} . Given some subset T_i , the auction with reserve $x = \frac{r+\gamma}{i}$ will sell to i bidders and therefore surpass the revenue target by γ : on all other samples, it will only sell to $i-1$ bidders and therefore miss the revenue target by $\frac{r+\gamma}{i} \geq \frac{H+\gamma}{k}$ (so long as $Hk \geq \gamma$, this is sufficient to γ -separate T_i from $\mathcal{S} \setminus T_i$). ■

3.4.3 Regular Virtual Valuation Maximizers

One might hope, given that both pricing and welfare maximization subject to reserve prices have $\text{poly}(n)$ pseudo-dimension, that any small amount of structure would be enough to ensure polynomial pseudo-dimension for a class of auctions. In this section, we disabuse ourselves of this hope: even the class which only contains Myerson’s auction for each pair of regular bidders has infinite fat-shattering and pseudo-dimension. This class is much more natural than the class of all auctions: if one only knew that bidders are regular, one natural approach to designing a revenue-optimal auction would be to consider the class of Myerson auctions for any set of regular distributions, and determine which of these had the best revenue on a sample. Unfortunately, since this class has unbounded fat-shattering complexity, one will not be able to guarantee success of this approach with a finite sample⁵.

In this section, we study precisely this class of auctions, the class of Myerson auctions for all regular bidders. Formally, this class contains auctions of the following form. For each bidder i , choose some non-decreasing function $\hat{\phi}_i$. Then, for a tuple \mathbf{v} , the auction allocates to the agent i such that $i = \text{argmax}_{i'} \hat{\phi}_{i'}(\mathbf{v}_{i'})$, if there is some i with $\hat{\phi}_i(v_i) > 0$ (otherwise, choose the empty allocation). We call the set of auctions with this form the class of *regular virtual valuation maximizers*. This name is apt, since there is *some* regular distribution for which each

⁵Running a form of “empirical” Myerson on the set of samples is, however, known to give good revenue guarantees with $\text{poly}(n, \frac{1}{\epsilon})$ samples, for MHR bidders, see Cole and Roughgarden [39] and others

of these $\widehat{\phi}$ is the virtual value function, and the auction corresponding to choosing the correct virtual valuation function for each bidder will be the unique auction which maximizes virtual welfare. This class does not allow for such precise memorization of samples as the class of all auctions. For example, if there are two tuples $\mathbf{v} < \mathbf{v}'$ (where $<$ is coordinate-wise), and the auction sold the item to someone in \mathbf{v} , then it must sell the item to someone in \mathbf{v}' , which was not true for the class of all auctions. That being said, regular virtual valuation maximizers are still quite expressive; they have infinite pseudo-dimension and γ fat-shattering dimension.

Theorem 3.4.5. *The class of single-item regular virtual welfare maximizers has infinite γ fat-shattering and pseudo-dimension.*

Proof. We first prove that there exists a set of samples of infinite size which can be shattered without any margin, even for 2 bidders by constructing such a set. Consider the set $\mathcal{S}_\epsilon = \{(n\epsilon, n\epsilon + \epsilon/2) | n \in \mathbb{N}, n\epsilon < H - \epsilon/2\}$ with witness $n\epsilon + \delta$ for $\delta < \epsilon/2$, for each sample $(n\epsilon, n\epsilon + \epsilon/2) \in \mathcal{S}$. We claim it is possible to shatter \mathcal{S}_ϵ , regardless of ϵ , using virtual welfare maximizers. So, consider some set $T \subseteq \mathcal{S}_\epsilon$. We will show two regular virtual welfare maximizers $\widehat{\phi}_1, \widehat{\phi}_2$ such that the auction \mathcal{A} allocating according to these functions has the property that $\text{rev}(\mathcal{A}, (n\epsilon, n\epsilon + \epsilon/2)) \geq n\epsilon + \delta$ if and only if $(n\epsilon, n\epsilon + \epsilon/2) \in T$.

For each sample $(n\epsilon, n\epsilon + \epsilon/2)$, if $\widehat{\phi}_1(n\epsilon) > \widehat{\phi}_2(n\epsilon + \epsilon/2)$, then the item is sold to agent 1, otherwise, the item is sold to agent 2. Let $\widehat{\phi}_1(n\epsilon) = n\epsilon^{100}$ for each sample. For each n such that $(n\epsilon, n\epsilon + \epsilon/2) \in T$, let $\widehat{\phi}_2$ exceed $n\epsilon^{100}$ for the first time at value $n\epsilon + \delta$. For each n such that $(n\epsilon, n\epsilon + \epsilon/2) \notin T$, let $\widehat{\phi}_2$ exceed $n\epsilon^{100}$ for the first time at value $n\epsilon$. In both cases, agent 2 wins, but her payment in the former case is $n\epsilon + \delta$ and only $n\epsilon$ in the latter case. Thus, this auction hits the revenue targets for all samples in T and misses the revenue targets for all other samples. Finally, both of these partial functions can be completed to nondecreasing functions over $[0, H]$: for $\widehat{\phi}_1$, at each point $n\epsilon$, the valuation is $n\epsilon^{100}$, and one can interpolate in any nondecreasing manner; for $\widehat{\phi}_2$, at either $n\epsilon$ or $n\epsilon + \delta$, the value of $\widehat{\phi}_2$ is set to be $n\epsilon^{100}$, and since $\delta < \epsilon/2$, this does not affect other samples. Then, since $|\mathcal{S}_\epsilon| = \Omega(H/\epsilon)$, there is no finite bound on the number of samples one can shatter.

We now prove the slightly more careful γ fat-shattering bound. Again, consider the case with 2 bidders. Let $\mathcal{S}_\epsilon = \{(n\epsilon, n\epsilon + 2\gamma) | n \in [\frac{H-2\gamma}{\epsilon}]\}$ and consider the witness $r^{n\epsilon} = n\epsilon + \gamma$ for sample $(n\epsilon, n\epsilon + 2\gamma)$. We now proceed to show how to shatter \mathcal{S} with regular virtual welfare maximizing auctions. Let $T \subseteq \mathcal{S}$. We will again define $\widehat{\phi}_1, \widehat{\phi}_2$ such that the revenue from \mathcal{A} on each sample in T will be γ larger than the target, and γ less than the target for each sample not in T . For each n , let $\widehat{\phi}_1(n\epsilon) = n\epsilon$. If $(n\epsilon, n\epsilon + 2\gamma) \in T$, let $\widehat{\phi}_2(n\epsilon + 2\gamma) = n\epsilon$ (and let $n\epsilon + 2\gamma$ be the first point at which $\widehat{\phi}_2$ hits $n\epsilon$). If $(n\epsilon, n\epsilon + 2\gamma) \notin T$, let $\widehat{\phi}_2(n\epsilon + 2\gamma) = n\epsilon - \delta$, for some very small δ . Then, in the former case, the auction sells to bidder 2 and earns revenue $n\epsilon + 2\gamma$; in the latter case, the auction sells to bidder 1 and makes at most $n\epsilon$ revenue. Thus, the revenue target is surpassed by at least γ for all samples in T and missed by at least γ for all other samples. Again, these functions can be completed to nondecreasing functions: $\widehat{\phi}_1$ is obviously nondecreasing already; for $\widehat{\phi}_2$, the constraints are of the form $\widehat{\phi}_2(n\epsilon + 2\gamma) = n\epsilon$ or $\widehat{\phi}_2(n\epsilon + 2\gamma) = n\epsilon - \delta$ and if $\delta < \epsilon$, then $n\epsilon \leq (n+1)\epsilon - \delta$. Thus, these defined values are nondecreasing for adjacent samples. Completing the functions to piecewise constant functions completes the construction. \blacksquare

3.5 Single-Item Auctions

This section focuses on single-item auctions, a simple setting in which many of our key definitions and proof techniques are most easily understood. All of the results of this section will be generalized significantly to matroid environments in Section 3.7 and general single-parameter environments in Section 3.8.

Section 3.5.1 defines the class of t -level single-item auctions, gives an example, and interprets the auctions as approximations to virtual welfare maximizers. Section 3.5.2 proves that the pseudo-dimension of the set of such auctions is $O(nt \log nt)$, which by Theorem 3.3.2 implies good upper bounds on sample complexity provided t is not too large. Section 3.5.3 proves that taking $t = \Omega(\frac{H}{\epsilon})$ yields low representation error.

3.5.1 t -Level Auctions: The Single-Item Case

We now introduce t -level auctions, or \mathcal{C}_t for short. These auctions are a generalization of the idea of running the (welfare-maximizing) VCG mechanism supplemented with non-anonymous reserves; intuitively, one can think of each bidder as having t possible reserves, and the reserve they face depends upon the values of the other bidders. Consider, for each bidder i , t numbers $0 \leq \ell_{i,0} \leq \ell_{i,1} \leq \dots \leq \ell_{i,t-1}$. We refer to these t numbers as *thresholds*. This set of tn numbers defines a particular t -level auction with the following allocation rule. Consider a valuation tuple \mathbf{v} :

1. For each bidder i , let $t_i(v_i)$ denote the index τ of the largest threshold $\ell_{i,\tau}$ that lower bounds v_i (or -1 if $v_i < \ell_{i,0}$). We call $t_i(v_i)$ the *level* of bidder i .
2. Sort the bidders from highest level to lowest level and, within a level, according to a fixed lexicographical order over bidders, or randomly, or from highest valuation to lowest valuation.⁶
3. Award the item to first bidder in this sorted order (unless $t_i = -1$ for every bidder i , in which case there is no sale).

The Payment Rule when Tie-Breaking is by Value

The payment rule is then the unique one that renders truthful bidding a dominant strategy and charges 0 to losing bidders — the winning bidder pays the lowest bid at which she would continue to win. It is important for us to understand this payment rule in detail. Suppose bidder i is the winner. There are three interesting cases. In the first case, i is the only bidder who might be allocated the item (other bidders have level -1), in which case her bid must be at least her lowest threshold. In the second case, there are multiple bidders at her level, so she must bid high enough to be at her level and also to outbid all other bidders at her level. In the final case, she need not compete at her level: she can choose to either pay one level above her competition (in which case the value of her bid doesn't matter), or she can bid at the same level as her highest-level

⁶ Tie-breaking according to valuations is an option, but is only-incentive compatible when bidders' valuations are regular. Moreover, our representation error guarantee only holds for irregular bidders if ties are broken either according to a lexicographical ordering or randomly.

competitors, in which case she needs bid high enough to be at their level and also outbid them. Formally, the payment p of the winner i (if any) is as follows. Let $\bar{\tau}$ denote the highest level τ such that there at least two bidders at or above level τ , and \mathcal{I} be the set of bidders whose level is at least $\bar{\tau}$.

- Monop** If $\bar{\tau} = -1$, then $p_i = \ell_{i,0}$ (she is the only bidder who might win, but needs to be at level 0 to win).
- Mult** If $\bar{\tau} = \bar{t}_i$ then $p_i = \max(\ell_{i,\bar{\tau}}, \max_{j \in \mathcal{I}} v_j)$ (she needs to have the highest bid of those with level \bar{t}_i , and be at level \bar{t}_i).
- Unique** If $\bar{\tau} < \bar{t}_i$, she pays $p_i = \min(\ell_{i,\bar{\tau}+1}, \max(\ell_{i,\bar{\tau}}, \max_{j \in \mathcal{I}} v_j))$ (she either needs to be at level $\bar{\tau} + 1$, in which case she need not compete, or at level $\bar{\tau}$, in which case she needs to compete).

We now give an example of a particular t -level auction, and point out an example of each case of the payment rule.


Example 3.5.1. Consider the following 4-level auction for bidders a, b, c . Let $\ell_{a,\cdot} = [2, 4, 6, 8]$, $\ell_{b,\cdot} = [1.5, 5, 9, 10]$ and $\ell_{c,\cdot} = [1.7, 3.9, 6, 7]$. For example, if bidder a bids less than 2 she is at level -1 , a bid $\in [2, 4)$ puts her at level 0, a bid in $[4, 6)$ at level 1, a bid in $[6, 8)$ at level 2, and a bid above 8 at level 3.

- Monop** If $v_a = 3, v_b < 1.5, v_c < 1.7$, then b, c are at level -1 (to which the item is never allocated). So, a wins and pays 2, the minimum she needs to bid to be at level 0.
- Mult** If $v_a = 9, v_b = 11, v_c < 7$, then a and b are both at level 3, and b has higher valuation, so b will win and pays 10 (the minimum she needs to bid to be at level 3). If v_b, v_c were the same but $v_a = 10.5$, b would need to pay 10.5 to beat v_a .
- Unique** If $v_a = 9, v_b = 8, v_c = 5$, then a is at level 2, b at level 1 and c at level 1, so a will win and pay 6 (her level for a bid of 6 is 2, so she need not bid higher than bidders b, c). If $v_a = 9, v_b = 5, v_c = 4.9$, then bidder a is at level 3 and bidders b, c at level 1: a will win and pay 5 (her level for this bid would be 1 and she needs to beat the bids of b, c). If $v_a = 7, v_b < 5, v_c = 3.9$, then a is at level 2 and bidder c is at level 1, so a wins and pays 4 (she must pay enough to be at the same level as the other bidders).


Figure 3.1 is a visual interpretation of Example 3.5.1.

Remark (Tie-breaking and the payment rule) When multiple agents are at the highest level, the previous description assumes the winner is picked according to value. When ties are broken according to a fixed lexicographical ordering, the payment rule is simpler: the winner pays her threshold for winning (e.g., a simplification of the previous three cases, where no bidder ever has to pay another bidder's value). The corresponding revenue is just the threshold the highest-priority bidder in \mathcal{I} . In the randomized case, the payment rule has whoever wins (according to the random tiebreaking) pay her threshold. The revenue then is the average of the thresholds of the bidders in \mathcal{I} . Either way, a very slightly modified argument proves the pseudo-dimension is the same.


Remark (Connection to virtual valuation functions): The set of t -level auctions are natural interpreted as discrete approximations to virtual welfare maximizers; indeed, our representation error bounds, like Theorem 3.5.2, make precise this intuition.



Monop: only agent a passes her minimum reserve, pays $\ell_{a,0}$



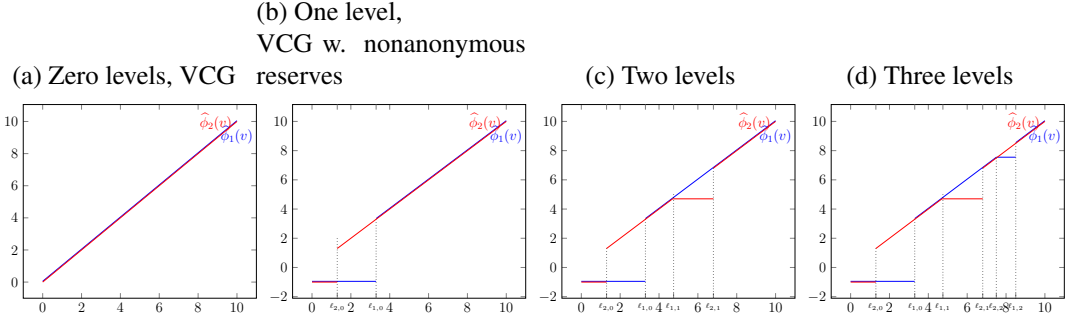
Mult: agents a, b at level 3, so b has to pay $\max(\ell_{b,3}, v_a)$



Unique: a is alone at her level, and will either pay $\ell_{i,\tau+1}$, or need to pay at least $\ell_{i,\tau}$ and compete with other bidders.

Figure 3.1: A visualization of the payment rules in Example 3.5.1. Each different line graph corresponds to a different tuple of bids, each corresponding to the examples' different cases. The ticks correspond to thresholds, the nodes to bids. Blue ticks (and nodes) belong to bidder a , green ticks (and nodes) to b , and red ticks (and nodes) to c .

Figure 3.2: The successively refined virtual valuation estimates t -level auctions, as t increases.



Intuitively, each level corresponds to a constraint of the form “If any bidder has level at least τ , do not sell to any bidder with level less than τ .” We can, roughly, interpret the $\ell_{i,\tau}$ ’s (with fixed τ , ranging over bidders i) as the bidder values that map to some common virtual value. For example, 1-level auctions treat all values below the one level as having negative virtual value, and uses values above that one level as proxies for the true virtual value. 2-level auctions use the next level to refine their estimate of the virtual values when at least one bidder’s virtual value is above the first threshold, and so on. See also Figure 3.2. With this interpretation, it is intuitively clear that as $t \rightarrow \infty$, it is possible to estimate bidders’ virtual valuation functions to arbitrary accuracy (and, thus, to estimate the Myerson-optimal auction to arbitrary accuracy).

3.5.2 The Pseudo-Dimension of t -Level Auctions

This section shows that the pseudo-dimension of the class of t -level single-item auctions with n bidders is $O(nt \log nt)$. Combining this with Theorem 3.3.2 immediately yields sample complexity bounds (parameterized by t) for learning the best such auction from samples.

Theorem 3.5.1. *The pseudo-dimension of the set of single-item t -level auctions with n bidders is $O(nt \log(nt))$.*

Proof. Recall from the definitions (Section 3.3.2) that we need to upper bound the size of every set that is shatterable using t -level auctions. For us, a set is a fixed set of samples $S = (\mathbf{v}^1, \dots, \mathbf{v}^m)$ of size m . Fix also a potential witness $R = (r^1, \dots, r^m)$. Every auction c induces a binary labeling of each of the samples \mathbf{v}^j of S , according to whether c ’s revenue on \mathbf{v}^j is at least r^j or strictly less than r^j . The set S is shattered with witness R if and only if the number of distinct labelings of S , ranging over all t -level auctions, is 2^m (the maximum possible).

We proceed to bounding from above the number of distinct labelings of S induced by t -level auctions (for any potential witness R). We count such labelings in two stages. Note that S involves nm numbers — one valuation v_i^j for each bidder for each sample. A t -level auction involves nt numbers — t thresholds $\ell_{i,\tau}$ for each bidder. Call two t -level auctions with thresholds $\{\ell_{i,\tau}\}$ and $\{\widehat{\ell}_{i,\tau}\}$ *equivalent* if:

1. The relative order of the $\ell_{i,\tau}$ ’s agrees with that of the $\widehat{\ell}_{i,\tau}$, in that both induce the same

permutation of $\{1, 2, \dots, n\} \times \{0, 1, \dots, t-1\}$.

2. Merging the sorted list of the v_i^j 's with the sorted list of the $\ell_{i,\tau}$'s yields the same partition of the v_i^j 's as does merging it with the sorted list of the $\widehat{\ell}_{i,\tau}$'s.

Operationally, the point is that every comparison between two numbers (valuations or thresholds) will be resolved identically by equivalent t -level auctions. Note that this is indeed an equivalence relation.

Using the two defining properties of equivalence, a crude upper bound on the number of equivalence classes of t -level auctions is

$$(nt)! \cdot \binom{nm + nt}{nt} \leq (nm + nt)^{2nt}. \quad (3.1)$$

We now proceed to upper bound the number of distinct binary labelings of S that can be generated by all of the auctions in a single equivalence class C . First, because all comparisons between two numbers (valuations or thresholds) are resolved identically across auctions in C , each bidder i in each sample \mathbf{v}^j of S is assigned a common level (across auctions in C), and in particular the winner (if any) in each sample \mathbf{v}^j is constant across all such auctions. By the same reasoning, the identity of the parameter that gives the winner's payment (either some $\ell_{i,\tau}$ or some v_i^j) is uniquely determined by pairwise comparisons (recall Section 3.5.1) and hence is common across all such auctions. While all of the valuations v_j^i are fixed (independent of the auction), payments of the form $\ell_{i,\tau}$ can vary across auctions in the equivalence class.

For a bidder i and level $\tau \in \{0, 1, 2, \dots, t-1\}$, write $S_{i,\tau} \subseteq S$ for the subset of samples in which bidder i wins and pays the value of the parameter $\ell_{i,\tau}$. The revenue obtained by a t -level auction in the equivalence class C on a sample of $S_{i,\tau}$ is increasing in $\ell_{i,\tau}$ and independent of all other parameters of the auction. Thus, ranging over all t -level auctions of C generates at most $|S_{i,\tau}|$ distinct binary labelings of $S_{i,\tau}$ — the possible subsets of $S_{i,\tau}$ for which an auction meets the corresponding target r^j form a nested collection.

Summarizing, within the equivalence class C of t -level auctions, varying a parameter $\ell_{i,\tau}$ generates at most $|S_{i,\tau}|$ different labelings of the samples $S_{i,\tau}$ and has no effect on the other samples. Since the subsets $\{S_{i,\tau}\}_{i,\tau}$ are disjoint, varying all of the $\ell_{i,\tau}$'s (i.e., ranging over C) generates at most

$$\prod_{i=1}^n \prod_{\tau=0}^{t-1} |S_{i,\tau}| \leq m^{nt} \quad (3.2)$$

distinct labelings of S .

Combining (3.1) and (3.2), the class of all t -level auctions produces at most $(nm + nt)^{3nt}$ distinct labelings of S . Since shattering S requires 2^m distinct labelings, we conclude that

$$2^m \leq (nm + nt)^{3nt}.$$

Solving, we obtain $m = O(nt \log nt)$, as claimed. ■

Remark (Approximate optimality of sample complexity): Given Theorem 3.5.1, one might wonder whether it is possible to learn an approximately optimal auction from a simpler

class (say, with pseudo-dimension $\text{poly}(\log(H, n), \frac{1}{\epsilon})$), allowing for much smaller sample complexity than the results presented in this work. However, lower bounds in Cole and Roughgarden [39] imply that no such class exists: approximate revenue maximization *requires* polynomial dependence on n , the number of bidders, even when bidders' valuations independently drawn from MHR distributions. Thus, any class which is sufficiently expressive to guarantee $1 - \epsilon$ approximately optimal revenue necessarily *must* have fat-shattering complexity which grows polynomially with n .

3.5.3 The Representation Error of Single-Item t -Level Auctions

This section shows there exist t -level auctions whose revenue closely approximates the revenue of the optimal single-item auction. Formally, we show there exists some $t = \text{poly}(\frac{1}{\epsilon}, H)$ such that, for any collection of bidders whose valuations are independent and bounded in $[1, H]$, the class of t -level auctions contains an auction whose expected revenue is at least a $(1 - \epsilon)$ fraction of the optimal auction. The idea is, in effect, to “round” an optimal auction to a t -level auction without losing much revenue.⁷ We accomplish by using levels to approximate each bidder's virtual value: the lowest level equals the bidder's monopoly reserve price, the next $\frac{1}{\epsilon}$ levels are located at the values at which bidder i 's virtual value surpasses multiples of ϵ , and the remaining levels located at those values where bidder i 's virtual value reaches powers of $1 + \epsilon$. The following proof formalizes this idea and quantifies the number of levels needed to execute it.

Theorem 3.5.2. *Provided $t = \Omega(\frac{1}{\epsilon} + \log_{1+\epsilon} H)$, there exists a t -level single-item auction with lexicographical tie-breaking whose expected revenue is at least a $(1 - \epsilon)$ fraction of the optimal expected revenue, if bidders' valuations are product.*

Our proof relies on the following lemma.

Lemma 3.5.1. *Suppose there are n bidders, whose values are bounded in $[0, H]$, and $\mathbb{P}[\max_i v_i > \alpha] \geq \gamma$. Then, there is a t -level auction whose revenue is at least a $1 - \epsilon$ fraction of Myerson's revenue, for $t = O(\frac{1}{\gamma\epsilon} + \log_{1+\epsilon} \frac{H}{\alpha})$.*

Proof. Consider a fixed bidder i . We construct t thresholds for bidder i , and prove that the auction \mathcal{A} that uses these t levels for each bidder closely approximates the revenue of the revenue-optimal auction. Let ϵ' be a parameter to be defined later.

Set $\ell_{i,0} = \phi_i^{-1}(0)$, bidder i 's monopoly reserve.⁸ Then, for $\tau \in [1, \lceil \frac{1}{\gamma\epsilon'} \rceil]$, let $\ell_{i,\tau} = \phi_i^{-1}(\tau \cdot \alpha\gamma\epsilon')$. Finally, for $\tau \in [\lceil \frac{1}{\gamma\epsilon'} \rceil, \lceil \frac{1}{\gamma\epsilon'} \rceil + \lceil \log_{1+\frac{\epsilon}{2}} \frac{H}{\alpha} \rceil]$, let $\ell_{i,\tau} = \phi_i^{-1}(\alpha(1 + \frac{\epsilon}{2})^{\tau - \lceil \frac{1}{\gamma\epsilon'} \rceil})$. Intuitively, these thresholds just bucket bidders by virtual value.

Consider a fixed valuation profile \mathbf{v} . Let i^* denote the winner according to \mathcal{A} , and i' denote the winner according to the optimal auction. Recall from Section 3.3.1 that the optimal auction always awards the item to a bidder with the highest positive ironed virtual valuation (or no one,

⁷This “rounding” is just for the analysis. Recall that our learning algorithm, which does not know the optimal auction, just chooses the auction with the highest average revenue on the samples.

⁸Recall from Section 3.3.1 that ϕ_i denotes the virtual valuation function of bidder i . (From here on, we always mean the ironed version of virtual values.) It is convenient to assume that these functions are strictly increasing (not just nondecreasing); this can be enforced at the cost of losing an arbitrarily small amount of expected revenue.

if no such bidders exist). The definition of the levels immediately imply the following, recalling that ties are broken lexicographically to ensure case 3 holds.⁹

1. \mathcal{A} only allocates to non-negative virtual-valued bidders.
2. If there is no tie (that is, there is a unique bidder at the highest level), \mathcal{A} 's allocation agrees with that of Myerson.
3. When there is a tie at level τ , the virtual value of the winner of \mathcal{A} is close to that of the Myerson's winner:
 - (a) If $\tau \in [0, \lceil \frac{1}{\gamma\epsilon'} \rceil]$ then $\phi_{i'}(\mathbf{v}_{i'}) - \phi_{i^*}(\mathbf{v}_{i^*}) \leq \alpha\gamma\epsilon'$.
 - (b) If $\tau \in [\lceil \frac{1}{\gamma\epsilon'} \rceil, \lceil \frac{1}{\gamma\epsilon'} \rceil + \lceil \log_{1+\frac{\epsilon}{2}} \frac{H}{\alpha} \rceil]$, $\frac{\phi_{i^*}(\mathbf{v}_{i^*})}{\phi_{i'}(\mathbf{v}_{i'})} \geq 1 - \frac{\epsilon}{2}$.

Now, we argue directly about the expected virtual value achieved by \mathcal{A} . Reasoning about and conditioning on case 3a:

$$\begin{aligned}
& \mathbb{E}_{\mathbf{v}}[\phi_{i^*}(\mathbf{v}_{i^*}) | \phi_{i^*}(\mathbf{v}_{i^*}) \in [0, 1]] \mathbb{P}[\phi_{i^*}(\mathbf{v}_{i^*}) \in [0, 1]] \\
&= \mathbb{E}_{\mathbf{v}}[\phi_{i^*}(\mathbf{v}_{i^*}) | \phi_{i^*}(\mathbf{v}_{i^*}) \in [0, 1]] \mathbb{P}[\phi_{i'}(\mathbf{v}_{i'}) \in [0, 1]] \\
&\geq \mathbb{E}_{\mathbf{v}}[\phi_{i'}(\mathbf{v}_{i'}) - \alpha\gamma\epsilon' | \phi_{i'}(\mathbf{v}_{i'}) \in [0, 1]] \mathbb{P}[\phi_{i'}(\mathbf{v}_{i'}) \in [0, 1]] \\
&\geq \mathbb{E}_{\mathbf{v}}[\phi_{i'}(\mathbf{v}_{i'}) | \phi_{i'}(\mathbf{v}_{i'}) \in [0, 1]] \mathbb{P}[\phi_{i'}(\mathbf{v}_{i'}) \in [0, 1]] - \alpha\gamma\epsilon'
\end{aligned} \tag{3.3}$$

where the first inequality comes from the fact that \mathcal{M} allocates to an agent with virtual value in $[0, 1]$ if and only if \mathcal{A} does, the second follows from Fact 3a, and the final from linearity of expectation. Similarly, for case 3b:

$$\begin{aligned}
& \mathbb{E}_{\mathbf{v}}[\phi_{i^*}(\mathbf{v}_{i^*}) | \phi_{i^*}(\mathbf{v}_{i^*}) > 1] \mathbb{P}[\phi_{i^*}(\mathbf{v}_{i^*}) > 1] \\
&= \mathbb{E}_{\mathbf{v}}[\phi_{i^*}(\mathbf{v}_{i^*}) | \phi_{i'}(\mathbf{v}_{i'}) > 1] \mathbb{P}[\phi_{i'}(\mathbf{v}_{i'}) > 1] \\
&\geq \mathbb{E}_{\mathbf{v}}\left[\left(1 - \frac{\epsilon}{2}\right) \phi_{i'}(\mathbf{v}_{i'}) | \phi_{i'}(\mathbf{v}_{i'}) > 1\right] \mathbb{P}[\phi_{i'}(\mathbf{v}_{i'}) > 1] \\
&\geq \left(1 - \frac{\epsilon}{2}\right) \mathbb{E}_{\mathbf{v}}[\phi_{i'}(\mathbf{v}_{i'}) | \phi_{i'}(\mathbf{v}_{i'}) > 1] \mathbb{P}[\phi_{i'}(\mathbf{v}_{i'}) > 1].
\end{aligned} \tag{3.4}$$

Finally, notice that if $\mathbb{P}[\max_i v_i > \alpha] \geq \gamma$, that

$$\mathbb{E}[\text{rev}(\mathcal{M})] \geq \alpha\gamma. \tag{3.5}$$

Then, combining Equations 3.3 and 3.4 along with Fact 1, we have

$$\mathbb{E}_v[\text{rev}(\mathcal{A})] = \mathbb{E}_v[\phi_{i^*}(\mathbf{v}_{i^*})] \geq (1 - \frac{\epsilon}{2}) \mathbb{E}[\text{rev}(\mathcal{M})] - \epsilon' \geq \left(1 - \frac{\epsilon}{2} - \epsilon'\right) \mathbb{E}[\text{rev}(\mathcal{M})]$$

The final inequality follows from Equation 3.5. Setting $\epsilon' = \frac{\epsilon}{2}$ completes the proof. ■

With Lemma 3.5.1 in hand, the proof of Theorem 3.5.2 is immediate.

⁹When bidders are irregular, their virtual valuation functions may be non-monotonic in their values; breaking ties according to value might lead to *lower* virtual value than the lower-bounds the levels are meant to represent, but breaking ties lexicographically implies that the amortized ironed welfare is fixed over ironed intervals. See Hartline [72] for details.

Proof. Lemma 3.5.1 applies, with $\alpha = \gamma = 1$. ■

Corollary 3.5.2. *With probability $1 - \delta$, the empirical revenue maximizer of the class of t -level single-item auctions on a set of samples S of size m has expected revenue which $1 - \epsilon$ -approximates the revenue of Myerson, for $t = O\left(\frac{1}{\epsilon} + \log_{1+\epsilon} H\right)$ and if*

$$m = O\left(\left(\frac{H}{\epsilon}\right)^2 \left(nt \log(nt) \log \frac{H}{\epsilon} + \log \frac{1}{\delta}\right)\right) = \tilde{O}\left(\frac{H^2 n}{\epsilon^3}\right).$$

Remark 3.5.1 (Near-optimality of sample complexity). Can we do better than Theorem 3.5.2? Can we learn an approximately optimal auction from a simpler class — with pseudo-dimension $\text{poly}(\log H, \log n, \frac{1}{\epsilon})$, say — allowing for much smaller sample complexity than achieved here? The answer is negative: lower bounds in Cole and Roughgarden [39] imply that approximate revenue maximization requires sample complexity at least linear in the number of bidder n , even when bidders' valuations independently drawn from MHR distributions. Thus, every class of auctions that is sufficiently expressive to guarantee expected revenue at least $1 - \epsilon$ times optimal must have pseudo-dimension that grows polynomially with n .

3.6 Unbounded MHR Distributions

This section shows how to replace the assumption of bounded valuations by the assumption that each valuation distribution satisfies the monotone hazard rate (MHR) condition, meaning that $\frac{f_i(v_i)}{1-F_i(v_i)}$ is nondecreasing. Our resulting sample complexity bounds depend on the number of bidders n and the error parameter ϵ only. bounded case, following ideas from This extension is based on previous work [30] that effectively reduces the case of MHR valuations to the case of valuations lying in the interval $[\beta\epsilon, 2\beta \log \frac{1}{\epsilon}]$ for a suitable choice of β . Our analysis works with η -truncated t -level auctions, where each t -level auction f is replaced with $f_\eta = \min(f, \eta)$.

Theorem 3.6.1. *Suppose F is a product distribution and each bidder's valuation distribution satisfies the MHR condition. Then, for each $\epsilon > 0$, and each $\hat{\beta} \geq \beta$ such that $\mathbb{P}\left[\max_i v_i \geq \frac{\hat{\beta}}{2}\right] \geq 1 - \frac{1}{\sqrt{\epsilon}} - \epsilon'$, there is a t -level $\left(\hat{\beta} \log \frac{1}{\epsilon'}\right)$ -truncated auction with expected revenue at least $1 - \epsilon$ times that of an optimal auction, where $t = \Theta\left(\frac{1}{\epsilon'} + \log_{1+\epsilon'}\left(\log \frac{1}{\epsilon'}\right)\right)$ and $\epsilon' = O\left(\frac{\epsilon}{\log \frac{1}{\epsilon}}\right)$.*

Before proving Theorem 3.6.1, we quote a key fact about MHR distributions Cai and Daskalakis [30].

Theorem 3.6.2 (Theorem 19 and Lemma 38 of [30]). *Let X_1, \dots, X_n be a collection of independent random variables whose distributions satisfy the MHR condition. Then there exists an anchoring point β such that*

$$\mathbb{P}\left[\max_i X_i \geq \frac{\beta}{2}\right] \geq 1 - \frac{1}{\sqrt{\epsilon}},$$

and for all $\epsilon > 0$,

$$\int_{2\beta \log \frac{1}{\epsilon}}^{\infty} z f_{\max_i X_i}(z) dz \leq 36\beta\epsilon \log \frac{1}{\epsilon}.$$

Now, we proceed to prove Theorem 3.6.1.

Proof of Theorem : thm:levels-rep-mhr Fix ϵ' , to be defined later. Conditioning on all bids being at most $2\widehat{\beta} \log \frac{1}{\epsilon'}$ allows us apply Lemma 3.5.1 as though the valuations are bounded. In particular, since $\mathbb{P}[\max_i v_i \geq \frac{\widehat{\beta}}{2}] \geq 1 - \frac{1}{\sqrt{e}} - \epsilon'$, Lemma 3.5.1 implies for $\alpha = \frac{\widehat{\beta}}{2}$, $\gamma = 1 - \frac{1}{\sqrt{e}} - \epsilon'$ and $H = 2\widehat{\beta} \log \frac{1}{\epsilon'}$, implies the existence of a $t = O\left(\frac{1}{\epsilon'} + \log_{1+\epsilon'}\left(\log \frac{1}{\epsilon'}\right)\right)$ -level truncated¹⁰ auction \mathcal{A} such that:

$$\mathbb{E} \left[\text{rev}(\mathcal{A}) \mid \max_i \mathbf{v}_i \leq 2\widehat{\beta} \log \frac{1}{\epsilon'} \right] \geq (1 - \epsilon') \mathbb{E} \left[\text{rev}(\mathcal{M}) \mid \max_i \mathbf{v}_i \leq 2\widehat{\beta} \log \frac{1}{\epsilon'} \right] \quad (3.6)$$

Thus, we have

$$\begin{aligned} \mathbb{E} [\text{rev}(\mathcal{A})] &\geq \mathbb{E} \left[\text{rev}(\mathcal{A}) \mid \max_i \mathbf{v}_i \leq 2\widehat{\beta} \log \frac{1}{\epsilon'} \right] \mathbb{P} \left[\max_i \mathbf{v}_i \leq 2\widehat{\beta} \log \frac{1}{\epsilon'} \right] \\ &\geq (1 - \epsilon') \mathbb{E} \left[\text{rev}(\mathcal{M}) \mid \max_i \mathbf{v}_i \leq 2\widehat{\beta} \log \frac{1}{\epsilon'} \right] \mathbb{P} \left[\max_i \mathbf{v}_i \leq 2\widehat{\beta} \log \frac{1}{\epsilon'} \right] \\ &\geq (1 - \epsilon') \mathbb{E} [\text{rev}(\mathcal{M})] - 36\widehat{\beta}\epsilon' \log \frac{1}{\epsilon'} \\ &\geq \left(1 - O\left(\epsilon' \log \frac{1}{\epsilon'}\right) \right) \mathbb{E} [\text{rev}(\mathcal{M})] \end{aligned}$$

where the first inequality comes from the fact that \mathcal{A} only sells to agents with non-negative virtual value (so the revenue on a smaller region of bids is only less), the second from Equation 3.6 and probabilities being at most 1, the penultimate from Theorem 3.6.2, and the final from the fact that $\text{rev}(\mathcal{M}) \geq \frac{1 - \frac{1}{\sqrt{e}}}{2} \widehat{\beta}$. Setting $\epsilon = \epsilon' \log \frac{1}{\epsilon'}$ and noticing that this implies $\epsilon' \leq \frac{\epsilon}{\log \frac{1}{\epsilon}}$ yields the desired result. \blacksquare

Thus, we have the following corollary of Theorem 3.3.2 and Theorem 3.6.1.

Corollary 3.6.3. *With probability $1 - \delta$, the empirical revenue maximizer for a sample of size m of the class of t -level η -truncated single-item auctions is a $1 - O(\epsilon)$ -approximation to Myerson for n MHR bidders, for $t = O\left(\frac{1}{\epsilon'} + \log_{1+\epsilon'}\left(\log \frac{1}{\epsilon'}\right)\right)$, $\epsilon' = \frac{\epsilon}{\log \frac{1}{\epsilon}}$ and*

$$m = O\left(\left(\frac{1}{\epsilon'}\right)^2 \left(nt \log(nt) \ln \frac{1}{\epsilon'} + \ln \frac{1}{\delta}\right)\right) = \tilde{O}\left(\frac{n}{\epsilon^3}\right)$$

where η can be learned from the set of m samples.

Proof. We first argue that one can learn some η from the sample. Let $Q(S, g) = \frac{\sum_{\mathbf{v} \in S_{\epsilon'}} \mathbb{I}[\max_i \mathbf{v}_i \geq g]}{|S|}$ and $q(g) = \mathbb{P}[\max_i \mathbf{v}_i \geq g]$ (the empirical and true probability that the maximum bid is at least g , respectively). Given ϵ' , consider a set of samples $S_{\epsilon'}$ of profiles, and compute the largest $\widehat{\beta}$ such that $Q(S_{\epsilon'}, \frac{\widehat{\beta}}{2}) \geq 1 - \frac{1}{\sqrt{e}} - \epsilon'$. Standard VC-bounds imply that, if, for all ρ , $|q(\rho) - Q(S_{\epsilon'}, \rho)| \leq \epsilon'$,

¹⁰Lemma 3.5.1 only implies the existence of such a t -level auction. However, when the bids are all below some η , one can always find an η -truncated auction which is equivalent to each untruncated auction.

with probability $1 - \delta$ if $|S_{\epsilon'}| \geq O\left(\left(\frac{1}{\epsilon'}\right)^2 \ln \frac{1}{\delta}\right)$ ¹¹. In particular, with probability $1 - \delta$, we will have $q(S_{\epsilon'}, \frac{\beta}{2}) \geq 1 - \frac{1}{\sqrt{e}} - \epsilon'$, so it will be the case that $\hat{\beta} \geq \beta$, and also $q(\hat{\beta}) \geq 1 - \frac{1}{\sqrt{e}} - 2\epsilon'$. Then, let $\eta = 2\hat{\beta} \log \frac{1}{\epsilon'}$.

Now, Theorem 3.6.1 implies the existence of a η -truncated t -level auction which $1 - \epsilon'$ -approximates Myerson. The argument is completed using the fact that, if the auctions' values are upper-bounded by η , one can equivalently think of the *values* being upper-bounded by η , so Theorem 3.3.2 implies the sample complexity bound allowing additive error $\epsilon'\eta = \frac{\epsilon'\hat{\beta}}{2} \log \frac{1}{\epsilon'}$. This error is multiplicatively at most $\epsilon = \epsilon' \log \frac{1}{\epsilon'}$, since $\text{rev}(\mathcal{M}) = \Omega\left(\frac{\hat{\beta}}{2}\right)$. ■

3.7 t -Level Matroid Auctions

This section extends the ideas and techniques from Section 3.5.2 to matroid environments. The straightforward generalization of t -level auctions to matroid environments suffices: we order the bidders by level, breaking ties within a level by some fixed linear ordering over agents \succ , and greedily choose winners according to this ordering (subject to feasibility and to bids exceeding the lowest threshold). The next theorem bounds the pseudo-dimension of this more general class of auctions.

Theorem 3.7.1. *The pseudo-dimension of t -level matroid auctions with n bidders is $O(nt \log(nt))$.*

The proof is conceptually similar to that of Theorem 3.5.1, though we require a more general argument. Our proof uses a couple of standard results from learning theory (see e.g. [86] for details). The first, also known as Sauer's Lemma, states that the number of distinct projections of a set S induced by a set system with bounded VC dimension grows only polynomially in $|S|$.

Lemma 3.7.2. *Let \mathcal{C} be a set of functions from \mathcal{Q} to $\{0, 1\}$ with VC dimension d , and $S \subseteq \mathcal{Q}$. Then*

$$|\{S \cap \{x \in \mathcal{Q} : c(x) = 1\} : c \in \mathcal{C}\}| \leq |S|^d.$$

Recall that a *linear separator* in \mathbb{R}^d is defined by coefficients a_1, \dots, a_d , and assigns $x \in \mathbb{R}^d$ the value 1 if $\sum_{i=1}^d a_i x_i \geq 0$ and the value 0 otherwise.

Lemma 3.7.3. *The set of linear separators in \mathbb{R}^d has VC dimension $d + 1$.*

Proof of Theorem 3.7.1: Consider a set of samples S of size m which can be shattered by t -level matroid auctions with revenue targets (r^1, \dots, r^m) . We upper-bound the number of labelings of S possible using t -level auctions, which again yields an upper bound on m .

We partition auctions into equivalence classes, identically to the proof of Theorem 3.5.1. Recall that, across all auctions in an equivalence class, all comparisons between two thresholds or a threshold and a bid are resolved identically. Recall also that the number of equivalence classes is at most $(nm + nt)^{2nt}$. We now upper bound the number of distinct labelings any fixed equivalence class C of auctions can generate.

Consider a class C of equivalent auctions. The allocation and payment rules are more complicated than in the single-item case but still relatively simple. In particular, whether or not a bidder wins depends only on the ordering of the bidders (by level) and the fixed tie-breaking rule

¹¹This can be thought of as a class of binary classifiers with VC-dimension one.

\succ , and thus is a function only of comparisons between bids and thresholds. This implies that, for every sample in S , all auctions in the class C declare the same set of winners. It also implies that the payment of each winning bidder is a fixed threshold $\ell_{i^*,\tau}$, and the identity of this parameter is the same across all auctions in C .

Now, encode each auction $\mathcal{A} \in C$ and sample \mathbf{v}^j as an $nt + 1$ -dimensional vector as follows. Let $x_{i,\tau}^{\mathcal{A}}$ equal the value of $\ell_{i,\tau}$ in the auction \mathcal{A} . Define $x_{nt+1}^{\mathcal{A}} = 1$ for every $\mathcal{A} \in C$. Define $y_{i,\tau}^j = 1$ if bidder i is a winner paying her threshold $\ell_{i,\tau}$ for auctions in C and 0 otherwise. Finally, define $y_{nt+1}^j = -r^j$. The point is that, for every auction \mathcal{A} in the class C and sample \mathbf{v}^j ,

$$x^{\mathcal{A}} \cdot y^j \geq 0$$

if and only if $\text{rev}(\mathcal{A}) \geq r^j$. Thus, the number of distinct labelings of the samples generated by auctions in C is bounded above by the number of distinct sign patterns on m points in \mathbb{R}^{nt+1} generated by all linear separators. (The y^j -vectors are constant across C and can be viewed as m fixed points in \mathbb{R}^{nt+1} ; each auction $\mathcal{A} \in C$ corresponds to the vector $x^{\mathcal{A}}$ of coefficients.) Applying Lemmas 3.7.2 and 3.7.3, t -level matroid auctions can generate at most m^{nt+2} labelings per equivalence class, and hence at most $(nm + nt)^{3nt+2}$ distinct labelings in total. This imposes the restriction

$$2^m \leq (nm + nt)^{3nt+2};$$

solving for m yields the desired bound. ■

We now extend our representation error bound for t -level single-item auctions to t -level matroid auctions.

Theorem 3.7.4. *Consider an arbitrary matroid environment. Suppose F is a production distribution with valuations in $[1, H]$. Provided $t = \Omega\left(\frac{1}{\epsilon} + \log_{1+\epsilon} H\right)$, there exists a t -level matroid auction with expected revenue at least a $1 - \epsilon$ fraction of the optimal expected revenue.*

The key new idea in the proof is to exhibit a bijection between the feasible sets \mathcal{I}^* (our winning set) and \mathcal{I}' (the optimal winning set) such that each bidder from \mathcal{I}^* has a level at least as high as their bijective partner in \mathcal{I}' . To implement this, we use the following property of matroids (e.g. [73, 120]).

Proposition 3.7.5. *Let OPT denote the largest-weight set of a matroid, and let B be any other feasible set such that $|B| = |\text{OPT}|$, and OPT_i, B_i denote the i -th largest element of OPT and B , respectively. Then $w(\text{OPT}_i) \geq w(B_i)$ for all i .*

Proof of Theorem 3.7.4: Define bidders' thresholds exactly as in the proof of Theorem 3.5.2 and let \mathcal{A} denote the corresponding t -level auction. Fix an arbitrary valuation profile \mathbf{v} . Let \mathcal{I}^* denote the set of winning bidders in \mathcal{A} and \mathcal{I}' the set of winning bidders in \mathcal{M} . Recall that the latter is the feasible set that maximizes the sum of virtual valuations. Both sets are maximally independent amongst those bidders with non-negative virtual value (\mathcal{M} , by virtue of being welfare-maximal, and \mathcal{A} , by definition). Then, we claim $|\mathcal{I}^*| = |\mathcal{I}'|$ (if not, by the augmentation property of matroids, the smaller set could be extended to include an element of the larger while maintaining independence, violating their maximality).

Notice that \mathcal{I}^* is lexicographically optimal with respect to the *levels*, rather than the exact weights. Proposition 3.7.5 implies that \mathcal{I}' is also lexicographically optimal with respect to the

levels; thus, the level of the i th largest bidder in \mathcal{I}' has the same level as the i th largest bidder in \mathcal{I}^* . Then, by an accounting argument identical to the one for the single-item case (comparing virtual values for the i th bidder in \mathcal{I}^* to the i th bidder in \mathcal{I}') summing up over all bidders completes the proof. ■

Thus, we have the following corollary about the sample complexity of $1 - \epsilon$ -approximating Myerson in matroid environments with t -level auctions, noting that the maximum revenue is now nH rather than H .

Corollary 3.7.6. *With probability $1 - \delta$, the empirical revenue maximizer for a sample of size m of the class of t -level single-item auctions is a $1 - O(\epsilon)$ -approximation to Myerson for n bidders whose valuations are in $[1, H]$, for $t = O\left(\frac{1}{\epsilon} + \log_{1+\epsilon} H\right)$ and*

$$m = O\left(\left(\frac{Hn}{\epsilon}\right)^2 \left(nt \log(nt) \ln \frac{Hn}{\epsilon} + \ln \frac{1}{\delta}\right)\right) = \tilde{O}\left(\frac{H^2 n^3}{\epsilon^3}\right).$$

3.8 Single-parameter t -level auctions

In this section, we show how to extend the ideas and techniques from Section 3.5.2 to any single-parameter environment which has the empty set as a feasible outcome. With this mild assumption, the results in this section do not require the environment to be a matroid or even downwards-closed. Before we state this result, we need a slight generalization of the t -level auction to this setting. Previously, no t -level auction would allocate to any bidder whose value was below their lowest threshold, and this will not be a possibility in environments which are not downwards-closed. Instead, in this setting, if *any* bidder fails to pass her lowest threshold, we will assume \mathcal{A} will choose the empty outcome. Moreover, a t -level auction will now need more about what the various levels represent: we were implicitly using the k th threshold to correspond to a value where each bidder's virtual value would pass some quantity q_k in other settings.

In this general setting, we will make that connection explicit. There will still be nt numbers which define a particular t -level auction, the t threshold locations per bidder. In addition, we will consider a fixed vector $\Phi \in \mathbb{R}^t$ (*not* parameterizing the auction class) which, for all τ , intuitively assigns an estimate of $\phi_i^{-1}(\ell_{i,\tau})$, which is the same for all bidders i . Formally, Φ_τ will be used to assign a real value to a feasible set $X \in \mathcal{X}$ with a valuation profile \mathbf{v} as follows. Let $e_X = \sum_{i \in X} \Phi_{t_i(v_i)}$, where $t_i(v_i)$ as before is the level agent i 's bid according to \mathbf{v}_i . Then, a particular t -level auction will choose the winning set $X \in \mathcal{X}$ which maximizes e_X (breaking ties in some fixed way which does not depend upon the bids). If, for each bidder i and level τ , the threshold $\ell_{i,\tau}$ is placed exactly at the value at which i 's virtual valuation surpasses Φ_τ , then this auction is approximately optimizing the virtual surplus of the winning set.

Theorem 3.8.1. *The pseudo-dimension of t -level single-parameter auctions with n bidders is $O(nt \log(nt))$.*

These auctions have slightly more complicated payment rules than those for matroids, where each agent i was intuitively competing with (at most) one other bidder for inclusion in the winning set. Now, a bidder i who is in the winning set X will have a payment of the following form. For a fixed assignment of levels to bidders, sort the alternatives according to

their values e_Y for all $Y \in \mathcal{X}$. Let Y be the highest-ranked alternative set which does not contain i . Then, i 's payment will be the threshold corresponding to the minimal τ such that $\sum_{i' \in X, i' \neq i} \Phi_{t_{i'}(v_{i'})} + \Phi_\tau \geq \sum_{i' \in Y} \Phi_{t_{i'}(v_{i'})}$ (namely, the minimal bid which keeps X preferred to Y in terms of the estimated virtual values)¹². While this rule is more complicated, it is *still* the case that, once each bidder is assigned to some level, each of the bidders in the winning set's payment is just one of their thresholds. Thus, the proof of Theorem 3.8.1 is identical to the one of Theorem 3.7.1 without ties.

When considering non-downwards-closed environments, the optimal revenue may be arbitrarily close to 0, making it difficult to argue about multiplicative approximations to Myerson's revenue. Instead, we will give a weaker guarantee, namely, that the empirical revenue maximizer will have expected revenue which is *additively* close to Myerson's revenue guarantee. If one is willing to restrict the environment to be downwards-closed, it is possible to achieve a multiplicative guarantee, since in that case the optimal revenue is at least 1. We now state the Theorem which bounds the representation error of t -level auctions for single-parameter environments.

Theorem 3.8.2. *There is a t -level auction whose expected revenue is within an additive ϵ of Myerson's revenue in any single-parameter setting \mathcal{X} such that $\emptyset \in \mathcal{X}$, for $t = O\left(\frac{Hn^2}{\epsilon}\right)$, for n bidders with valuations bounded in $[0, H]$ when the value distribution is product.*

Proof. Let $t = \frac{Hn^2}{\epsilon} + \frac{Hn}{\epsilon}$. We will begin by defining Φ , the t -dimensional vector corresponding to the estimated virtual values. Let $\Phi_0 = -Hn$ (if any bidder has virtual value $< -Hn$, the virtual value of the set is negative, since virtual values are upper-bounded by values and the value of the remaining set may be at most $H(n-1)$, so in this case one should allocate to \emptyset). Then, let $\Phi_\tau = \Phi_{\tau-1} + \frac{\epsilon}{n}$. Thus, we partition the space of virtual values into additive sections of width $\frac{\epsilon}{n}$.

Then, for each bidder i and τ , let $\ell_{i,\tau} = \phi_i^{-1}(\Phi_\tau)$, the value at which bidder i 's virtual value surpasses Φ_τ . Then, consider a valuation profile \mathbf{v} on which \mathcal{M} and this particular \mathcal{A} disagree on the winning sets $\mathcal{I}^*, \mathcal{I}' \in \mathcal{X}$. Notice that each bidder i 's virtual value is estimated correctly within an additive $\frac{\epsilon}{n}$ by $\Phi_{t_i(v_i)}$ (assuming no bidder has highly negative virtual value, in which case \mathcal{M} and \mathcal{A} both choose outcome \emptyset), and are never overestimated. Thus, it is the case that

$$\sum_{i^* \in \mathcal{I}^*} \phi_{i^*}(\mathbf{v}_{i^*}) \geq \sum_{i^* \in \mathcal{I}^*} \Phi_{t_{i^*}(v_{i^*})} \geq \sum_{i' \in \mathcal{I}'} \Phi_{t_{i'}(v_{i'})} \geq \sum_{i' \in \mathcal{I}'} \phi_{i'}(\mathbf{v}_{i'}) - \epsilon$$

and the claim follows. ■

Thus, we have the following sample complexity result for general single-parameter settings.

Corollary 3.8.1. *With probability $1 - \delta$, the empirical revenue maximizer on m samples S from the class of t -level auctions has true expected revenue within an additive ϵ of Myerson's expected revenue, for the single-parameter environment \mathcal{X} when bidders have valuations in $[0, H]$, for*

¹²Since we assume ties are broken in a way which does not depend on the bids, we can ignore ties in the payment rule, and agents will only ever pay thresholds.

$t = O\left(\frac{Hn^2}{\epsilon}\right)$, and

$$m = O\left(\left(\frac{Hn}{\epsilon}\right)^2 \left(nt \log(nt) \log \frac{Hn}{\epsilon} + \log \frac{1}{\delta}\right)\right) = \tilde{O}\left(\frac{H^3 n^5}{\epsilon^3}\right)$$

3.9 Open Questions

There are several questions we believe are worthy of further study in this domain. First, the results described here are only written in terms of their implications for single-parameter settings (where agents' values can be described by a single real number, and whether or not they are included in the winning set). Our technique of bounding an auction class's sample complexity by analyzing its pseudo-dimension (or fat-shattering dimension) still applies in a setting with multiple parameters. However, the representation error of many classes with small pseudo-dimension may be large in multi-parameter settings (indeed, the work of Dughmi et al. [47] implies that learning an approximately optimal auction for general multi-parameter settings will require exponentially many samples). There may be interesting special cases for which exponentially many samples are not needed. For example, their lower bound considers a buyer with values for different items which are correlated: what if buyers' values for items are independent? What if the market is large, in the sense that no buyer contributes much larger than a $\frac{1}{n}$ fraction of the total expected revenue?

Secondly, we believe understanding the computational complexity of learning $(1-\epsilon)$ -approximately optimal auctions is an interesting one, in general single-parameter settings (for general, not necessarily regular distributions).

Chapter 4

Learning Valuation Distributions from Partial Observation

4.1 Introduction

Auction theory traditionally assumes that bidders' valuation distributions are known to the auctioneer, such as in the celebrated, revenue-optimal Myerson auction [99]. However, this theory does not describe how the auctioneer comes to possess this information. Recently work [39] showed that an approximation based on a finite sample of independent draws from each bidder's distribution is sufficient to produce a near-optimal auction. In this chapter, we study the problem of valuation distribution estimation, which is a trivial task when the observations v consist of independent draws v_i from each bidder i 's distribution. We consider a much more limited observational model: rather than the standard observation v , we consider when observations are simply *outcomes* of an auction on v .

We begin our study of first-price auctions, with bidders' bids drawn independently from their bid distribution. We consider observations which name the winner and her price.¹ We show that, from this information alone, we can reconstruct each bidder's distribution² over bids with polynomially many samples. We then use this tool to develop a method for reconstructing bidders' distributions where each observation is *only* the identity of the bidder. It is clear that this information alone is insufficient to learn anything more than the relative strength of each bidder's distribution. However, if we allow ourselves the ability to *participate* in the auctions, by submitting bids ourselves, we recover our ability to approximate each bidder's distribution over bids using a small number of samples. We consider extensions where different *subsets* of bidders participate in each round, and where bidders' valuations have a common-value component added to their independent private values. We also show that the sense in which we learn bidders' distributions is sufficient to set near-optimal mapping from subsets of bidders to anonymous reserve prices for an auction selling to those bidders. We present this result as an example of the usefulness of these approximate distributions³.

¹Equivalently, the auction could be second-price with the observation being the winner and her bid.

²Approximately, in the range of bids where the bidder can win the auction.

³We in no way mean this would be the most sample-efficient or natural way to construct such a mapping, nor do

We believe these questions are interesting for a number of reasons. First, if a large online advertising firm runs auctions repeatedly, and one of their competitors wishes to understand the bidding behavior of their customers, this competitor would be able to observe the outcome of a large number of these auctions (simply by visiting the webpages which are displaying the winning advertiser's ad). Alternatively, the competitor might compete in the advertising auctions directly to learn more about the advertiser's behavior in the auctions. This competitor might be interested in this information for a number of reasons: she might want to know if the advertisers behave differently in this ad auction than in their ad auction; or she might be doing market research before building her own ad auction. This learning might also be conducted by potential advertisers, who want to learn how to bid in these auctions, or by the company running the auction itself, if the data stored from completed auctions is simply the winner and what was paid, rather than all bids.

4.1.1 Related Work

Problems of reconstructing distributional information from limited or censored observations have been studied in both the medical statistics literature and the manufacturing/operations research literature. In medical statistics, a basic setting where this problem arises is estimating survival rates (the likelihood of death within t years of some medical procedure), when patients are continually dropping out of the study, independently of their time of death. The seminal work in this area is the Kaplan-Meier product-limit estimator [82], analyzed in the limit in the original paper and then for finite sample sizes [59], see also its use for a control problem [64]. In the manufacturing literature, this problem arises when a device, composed of multiple components, breaks down when the first of its components breaks down. From the statistics of when devices break down and which components failed, the goal is to reconstruct the distributions of individual component lifetimes [97, 100]. The methods developed (and assumptions made, and types of results shown) in each literature are different. In our work, we will build on the approach taken by the Kaplan-Meier estimator (described in more detail in Section 4.3), as it is more flexible and better suited to the types of guarantees we wish to achieve, extending it and using it as a subroutine for the kinds of weak observations we work with.

The area of prior-free mechanism design has aimed to understand what mechanisms achieve strong guarantees with limited (or no) information about the priors of bidders, particularly in the area of revenue maximization. There is a large variety truthful mechanisms that guarantee a constant approximation (see, cf, Hartline and Karlin [71]). A different direction is adversarial online setting which minimize the regret with respect to the best single price (see Kleinberg and Leighton [87]), or minimizing the regret for the reserve price of a second price auction [33]. In [33] it was assumed that bidders have an identical bid distribution and the algorithm observes the actual sell price after each auction, and based on this the bidding distribution is approximated.

A recent line of work tries to bridge between the Bayesian setting and the adversarial one, by assuming we observe a limited number of samples. For a regular distribution, as single sample bidders' distributions is sufficient to get a $1/2$ -approximation to the optimal revenue [43], which

we mean that learning the distributions should be thought of as an intermediate step towards a final goal of setting reserves.

follows from an extension of the [28] result that shows the revenue from a second-price auction with $n+1$ (i.i.d) bidders is higher than the revenue from running a revenue-optimal auction with n bidders. Recent work of Cole and Roughgarden [39] analyzes the number of samples necessary to construct a $1 - \epsilon$ -approximately revenue optimal mechanism for asymmetric bidders: they show it is necessary and sufficient to take $\text{poly}(\frac{1}{\epsilon}, n)$ samples from each bidder’s distribution to construct an $1 - \epsilon$ -revenue-optimal auction for bid distributions that are strongly regular. We stress that in this work we make no such assumptions, only that the distributions are continuous.

Chawla et al. [37] design mechanisms which are approximately revenue-optimal and also allow for good *inference*: from a sample of bids made in Bayes-Nash equilibrium, they would like to reconstruct the distribution over values from which bidders are drawn. This learning technique relies heavily on a sample being drawn *unconditionally* from the *symmetric* bid distribution, rather than only seeing the *winner’s identity* from *asymmetric* bid distributions, as we consider in this work.

Most of the focus in the “revenue maximization” literature has a fundamentally different objective than the one in this work. Namely, our primary goal is to reconstruct the bidders’ bid distributions, rather than focusing of the revenue directly. Our work differs from previous work in this space in that it assumes very limited observational information. Rather than assuming all n bids as an observation from a single run of the auction, or even observing only the price, we see only the identity of highest bidder. We do not need to make any regularity assumption on the bid distribution, our methodology handles *any* bounded continuous bid distribution.⁴

4.2 Model and Preliminaries

We assume there are n bidders, and each $i \in [n]$ has some unknown valuation distribution \mathcal{D}_i over the interval $[0, 1]$ ⁵. Each sample $t \in [m]$ refers to a fresh draw $v_i^t \sim \mathcal{D}_i$ for each i . The label of sample t will be denoted $y^t = \text{argmax}_i v_i^t$, the identity of the highest bidder. Our goal is to estimate F_i , the cumulative distribution for \mathcal{D}_i , for each bidder i , up to ϵ additive error for all values in a given range. In Section 4.3.1, we consider a setting in which a subset of bidders S^t participate at time t , and notice that our results extend directly. In Section 4.5, we examine several other extensions and modifications to this basic model.

We consider the problem of finding (sample and computationally) efficient algorithms for constructing an estimate \widehat{F}_i of F_i , the cumulative distribution function, such that for all bidders i and price levels p , $\widehat{F}_i(p) \in \{F_i(p) \pm \epsilon\}$. However, as discussed above, this goal is too ambitious in two ways. First, if the labels contain no information about the value of bids, the best we could hope to learn is the relative probability each person might win, which is insufficient to uniquely identify the CDFs, even without sampling error. We address this issue by allowing, at each time t , our learning algorithm to insert a fake bidder 0 (or reserve) of value $v_0^t = r^t$; the label at time t will be $y^t = \text{argmax}_i v_i^t$ ($y^t = 0$ will refer to a sample where the reserve was not met,

⁴Note that we measure the distance between two distributions using the total variation distance, which is essentially “additive”.

⁵All results can be extended to the case where bidders’ valuations are in $[0, H]$; the sample complexity results will degrade quadratically in H . One of these H s comes from using Chernoff bounds and the other from the number of intervals into which we break up the space of valuations.

or the fake bidder won the auction). The other issue, also described above, is that there will be values below which we simply cannot estimate the F_i s since low-valued bids will not win. In particular, if bids below price p never win, then any two cumulatives F_i, F'_i that agree above p will be statistically indistinguishable. Thus, we will consider a slightly weaker goal. We will guarantee our estimates $\widehat{F}_i(p) \in F_i(p) \pm \epsilon$ for all p where $\mathbb{P}[\text{the winning bid is at most } p] \geq \gamma$. We will let $p = \min_{p'} \{\mathbb{P}[\text{the winning bid is at most } p'] \geq \gamma\}$ be the point down to which we learn each bidder's distribution. Then, our goal is to minimize m , the number of samples necessary to estimate all bidders' distributions in this way, and we hope to have $m \in \text{poly}(n, \frac{1}{\epsilon}, \frac{1}{\gamma})$, with high probability of success over the draw of the sample. One final (and necessary) assumption we will make is that each \mathcal{D}_i has no point masses, and our algorithm will be polynomial in the maximum slope L of the F_i 's.⁶ If $L < \infty$, then in particular, this implies continuity of the PDFs (and thus no point masses), so for the remainder of the chapter we ignore ties.

4.2.1 A brief primer on the Kaplan-Meier estimator

Our work is closely related in spirit to that of the Kaplan-Meier estimator, KM, for survival time; in this section, we describe the techniques used for constructing the KM [82]. This will give some intuition for the estimator we present in Section 4.3. We translate their results into the terminology we use in the auction setting, from the standard terminology used in the survival rate literature. Suppose each sample t is of the following form. Each bidder i draws their bid $b_i^t \sim \mathcal{D}_i$ independently of each other bid. The label $y^t = (\max_i b_i^t, \text{argmax}_i b_i^t)$ consists of the winning bid and the identity of the winner. From this, we would like to reconstruct an estimate \widehat{F}_i of F_i . Given m samples, relabel them so that the winning bids are in increasing order, e.g. $b_{i_1}^1 \leq b_{i_2}^2 \leq b_{i_m}^m$. Here is some intuition behind the KM: $F_i(x) = \mathbb{P}[b_i \leq x] = \mathbb{P}[b_i \leq x | b_i \leq y] \cdot \mathbb{P}[b_i \leq y]$ for $y > x$. Repeatedly applying this, we can see that, for $x < y_1 < y_2 < \dots < y_r$,

$$F_i(x) = \mathbb{P}[b_i \leq x | b_i \leq y_1] \mathbb{P}[b_i \leq y_r] \prod_{t=1}^{r-1} \mathbb{P}[b_i \leq y_t | b_i \leq y_{t+1}] \quad (4.1)$$

Now, we can employ the observation in Equation 4.1, with estimates of such conditional probabilities. Since other players' bids are independent, we can estimate the conditional probabilities as follows:

$$\mathbb{P} \left[b_i \leq b_{i_t}^t | b_i \leq b_{i_{t+1}}^{t+1} \right] \approx \begin{cases} \frac{t-1}{t} & \text{if } i \text{ won sample } t \\ 1 & \text{if } i \text{ lost sample } t \end{cases} \quad (4.2)$$

Thus, combining Equations 4.1 and 4.2, we have the Kaplan-Meier estimator:

$$\text{KM}(x) = \prod_{t: b_j^t \geq x} \left(\frac{t-1}{t} \right)^{\mathbb{I}[i \text{ won sample } t]}$$

⁶We can, alternatively, remove this assumption and change the nature of our CDF's approximation. Rather than saying for each x in a range that $\widehat{F}(x)$ is close to $F(x)$, one can say that there exists some δ such that $\widehat{F}(x)$ is close to $F(x + \delta')$ for some $\delta' \leq \delta$. The sample complexity will then depend upon $\frac{1}{\delta}$ rather than L .

Our estimator uses a similar Bayes-rule product expansion as KM, though it differs in several important ways. First, and most importantly, we do not see the winning bid explicitly; instead, we will just have lower or upper bounds on the highest non-reserve bid (namely, the reserve bid when someone wins or we win, respectively). Secondly, KM generally has no control issue; in our setting, we are *choosing* one of the values which will censor our observation. We need to pick appropriate reserves to get a good estimator (picking reserves that are too high will censor too many observations, only giving us uninformative upper bounds on bids, and reserves that are too low will never win, giving us uninformative lower bounds on bids). Our estimator searches the space $[0, 1]$ for appropriate price points to use as reserves to balance these concerns.

4.2.2 Summary of Main Results

We now summarize the main results of this chapter. We begin by showing how the Kaplan-Meier estimator can be used to reconstruct the CDF of each bidder's bids using polynomially many samples which show the winning bidder, along with the power to *participate* in an auction by bidding some number r . We then mention that this work directly extends to learning bidders' distributions when only certain *subsets* of bidders participate in each round (which could, in particular, model that not all advertising campaigns are interested in all impressions). With these estimators in hand, we then show how to use that estimation to set an approximately revenue-optimal reserve price for any subset of bidders. Finally, we show several extensions of this thinking, to the case where bidders' values are not independent, and where not all subsets of bidders participate in every auction.

4.3 Learning bidders' valuation distributions

In this section, we assume we have the power to insert a reserve price, and observe who won. Using this, we would like to reconstruct the CDFs of each bidder i up to some error, down to some price p_i where i has probability no more than γ of winning at or below p_i , up to additive accuracy ϵ . Our basic plan of attack is as follows. Our algorithm starts by estimating the probability i wins with a bid in some range $[a, a + \delta]$, by setting reserve prices at a and $a + \delta$, and measuring the difference in empirical probability that i wins with the two reserves. It then estimates the probability that no bidder bids above $a + \delta$ (by setting a reserve of $a + \delta$ and observing the empirical probability that no one wins). These together will be enough to estimate the probability that i wins with a bid in that range, conditioned on no one bidding above the range. We then show, for a small enough range, this is a good estimate for the probability i bids in the range, conditioned on no one bidding above the range. Then, we can chain these estimates together to form Kaplan, our estimator.

More specifically, our algorithm begins by partitioning $[0, 1]$ into a collection of intervals. This partition should have the following property. Within each interval $[x, y]$, there should be probability at most β of any person bidding in $[x, y]$, conditioned on no one bidding above y . This won't be possible for the lowest interval, but will be true for the other intervals. Then, the algorithm estimates the probability i will win in $[x, y]$ conditioned on all bidders bidding at most y . Then, our estimate of i 's probability of winning with a bid in $[x, y]$ is a $1 - \beta$ -

multiplicative approximation of i 's probability of bidding in $[x, y]$ (conditioning in both cases on all bidders bidding less than y). Then, the algorithm combines these estimates in a way such that the approximation factors do not blow up to reconstruct the CDF.

Algorithm Kaplan: estimates the CDF of i from samples with reserves

Data: $\epsilon, \gamma, \delta, L$, where L is the Lipschitz constant of the F_i s

Result: \widehat{F}_i , Kaplan estimator

Let $\widehat{F}_i(0) = 0, \widehat{F}_i(1) = 1, k = \frac{2Ln}{\beta\gamma} + 1, \delta' = \frac{\delta}{3k(\log k + 1)}, \beta = \frac{\epsilon\gamma}{32nL}, \alpha = \beta^2/96, \mu = \beta/96,$

$T = \frac{8 \ln 6 / \delta'}{\alpha^2 \gamma^2 (\frac{\mu}{2})^2};$

Let $\ell_1, \dots, \ell_{k'} = \text{Intervals}(\beta, \gamma, T);$

for $t = 2$ **to** $k' - 1$ **do**

 | Let $r_{\ell_\tau, \ell_{\tau+1}} = \text{IWin}(i, \ell_\tau, \ell_{\tau+1}, T);$

for $t = 2$ **to** $k' - 1$ **do**

 | Let $\widehat{F}_i(\ell_\tau) = \prod_{\tau' \geq t+1} (1 - r_{\ell_{\tau'}, \ell_{\tau'+1}});$

Define $\widehat{F}_i(x) = \max_{\ell_\tau \leq x} \widehat{F}_i(\ell_\tau);$

Theorem 4.3.1. *With probability at least $1 - \delta$, Kaplan outputs \widehat{F}_i , an estimate of F_i , with sample complexity*

$$m = O \left(\frac{n^8 L^8 \ln \frac{nL}{\epsilon\gamma} \left(\ln \frac{1}{\delta} + \ln \ln \frac{nL}{\epsilon\gamma} \right)}{\gamma^{10} \epsilon^6} \right)$$

and, for all p where $\mathbb{P}[\exists \text{ j.s.t. } j \text{ wins with a bid } \leq p] \geq \gamma$, if each CDF is L -Lipschitz, the error is at most:

$$F_i(p) - \epsilon \leq \widehat{F}_i(p) \leq F_i(p) + \epsilon.$$

Kaplan calls several other functions, which we will now informally describe, and state several lemmas describing their guarantees (the proofs can be found in Appendix A.1). `Intervals` partitions $[0, 1]$ into small enough intervals such that, conditioned on all bids being in or below that interval, the probability of any bidder bidding within the interval is small (ℓ_2 is close to p_γ , so we need not get a good estimation in $[0, \ell_2]$, and by definition $\ell_1 = 0$). `IWin` estimates the probability i wins in the region $[\ell_\tau, \ell_{\tau+1}]$, conditioned on all bids being at most $\ell_{\tau+1}$.

We now present the lemmas which make up the crux of the proof of Theorem 4.3.1. Lemma 4.3.1 bounds the error of `IWin` when its sample size is T . Lemma 4.3.2 does similarly for `Intervals`. Lemma 4.3.3 states that, if a region $[\ell_\tau, \ell_{\tau+1}]$ is small enough, the probability that i bids in $[\ell_\tau, \ell_{\tau+1}]$ (conditioned on all bids being at most $\ell_{\tau+1}$) is well-approximated by the probability that i wins with a bid in $[\ell_\tau, \ell_{\tau+1}]$ (conditioned on all bids being at most $\ell_{\tau+1}$). In combination, these three imply a guarantee on the sample complexity and accuracy of estimating

$$\mathbb{P}[i \text{ wins in } [\ell_\tau, \ell_{\tau+1}] \mid \max_j b_j \leq \ell_{\tau+1}]$$

which is the key ingredient of the Kaplan estimator.

Algorithm Iwin: Est. $\mathbb{P}[i \text{ wins in } [\ell_\tau, \ell_{\tau+1}] \mid \max_j b_j < \ell_{\tau+1}]$

Data: $i, \ell_\tau, \ell_{\tau+1}, T$

Result: $p_{\ell_\tau, \ell_{\tau+1}}^i$

Let S_{ℓ_τ} be a sample with reserve $\ell_{\tau+1}$ of size T ;

Let $S_{\ell_{\tau+1}}$ be a sample with reserve ℓ_τ of size T ;

Let S_{cond} be a sample with reserve $\ell_{\tau+1}$ of size T ;

Output $p_{\ell_\tau, \ell_{\tau+1}}^i = \frac{\sum_{t \in S_{\ell_\tau}} \mathbb{I}[i \text{ wins on sample } t] - \sum_{t \in S_{\ell_{\tau+1}}} \mathbb{I}[i \text{ wins on sample } t]}{\sum_{t \in S_{\text{cond}}} \mathbb{I}[0 \text{ wins on sample } t]}$;

Lemma 4.3.1. *Suppose, for a fixed interval $[\ell_\tau, \ell_{\tau+1}]$, $\mathbb{P}[i \text{ wins in } [0, \ell_{\tau+1}]] \geq \gamma$. Let $W_{i, \ell_\tau, \ell_{\tau+1}} = \mathbb{P}[i \text{ wins in } [\ell_\tau, \ell_{\tau+1}] \mid \max_j b_j \leq \ell_{\tau+1}]$. Then, with probability at least $1 - 3\delta'$, $\text{Iwin}(i, \ell_\tau, \ell_{\tau+1}, T)$ outputs $p_{\ell_\tau, \ell_{\tau+1}}^i$ such that*

$$(1 - \mu)W_{i, \ell_\tau, \ell_{\tau+1}} - \alpha \leq p_{\ell_\tau, \ell_{\tau+1}}^i \leq (1 + \mu)W_{i, \ell_\tau, \ell_{\tau+1}} + \alpha$$

and uses $3T$ samples, for the values of T, δ' as in Kaplan.

Algorithm Intervals: Partitions bid space to est. f_i

Data: β, γ, T, n, L

Result: $0 = \ell_1 < \dots < \ell_k = 1$

Let $\ell_k = 1, c = k, p_{\ell_c}^i = 1$;

while $p_{\ell_c}^i > \gamma/2$ **do** // Do binary search for the bottom of the next interval

 Let $\hat{\ell}_b = 0$;

while $\text{Inside}(\hat{\ell}_b, \ell_c, T) > \frac{\beta}{48}$ **do** // The interval is too large

 | $\hat{\ell}_b = \frac{\ell_c + \hat{\ell}_b}{2}$;

$\ell_{c-1} = \hat{\ell}_b$;

$c = c - 1$;

 Let S_1 be a sample of size T with reserve ℓ_{c-1} ;

$p_{\ell_c} = \frac{\sum_{t \in S_1} \mathbb{I}[j \geq 1 \text{ wins on sample } t]}{T}$;

Return $0, \ell_c, \dots, \ell_k$;

Lemma 4.3.2. *Let T as in Kaplan. Then, $\text{Intervals}(\beta, \gamma, T, L, n)$ returns $0 = \ell_1 < \dots < \ell_k = 1$ such that*

1. $k \leq \frac{48Ln}{\beta\gamma}$

2. For each $\tau \in [2, k]$, $\mathbb{P}[\max_j b_j \in [\ell_\tau, \ell_{\tau+1}] \mid \max_j b_j \leq \ell_{\tau+1}] \leq \frac{\beta}{16}$

3. $\mathbb{P}[\max_j b_j \in [\ell_1, \ell_2]] \leq \gamma$

with probability at least $1 - 3k \log(k)\delta'$, when bidders' CDFs are L -Lipschitz, using at most $3kT \log k$ samples.

With the guarantee of Lemma 4.3.2, we know that the partition of $[0, 1]$ returned by Intervals is "fine enough". Now, Lemma 4.3.3 shows that, when the partition fine enough, the conditional

probability i wins with a bid in each interval is a good estimate for the conditional probability i bids within that interval.

Lemma 4.3.3. *Suppose that, for bidder i and some $0 \leq \ell_\tau \leq \ell_{\tau+1} \leq 1$,*

$$\mathbb{P}[\max_{j \neq i} b_j \in [\ell_\tau, \ell_{\tau+1}] \mid \max_{j \neq i} b_j < \ell_{\tau+1}] \leq \beta.$$

Then,

$$1 \geq \frac{\mathbb{P}[i \text{ wins in } [\ell_\tau, \ell_{\tau+1}] \mid \max_j b_j < \ell_{\tau+1}]}{\mathbb{P}[i \text{ bids in } [\ell_\tau, \ell_{\tau+1}] \mid \max_j b_j < \ell_{\tau+1}]} \geq 1 - \beta$$

Finally, we observe that F_i can be written as the product of conditional probabilities.

Observation 4.3.1. Consider some set of points $0 < \ell_1 < \dots < \ell_k = 1$. $F_i(\ell_\tau)$ can be rewritten as the following product:

$$F_i(\ell_{\tau-1}) = \prod_{\tau' \geq t} (1 - \mathbb{P}[b_i \in [\ell_{\tau'-1}, \ell_{\tau'}] \mid b_i \leq \ell_{\tau'}])$$

We relegate the formal proof of Theorem 4.3.1 to the full version for reasons of space. We give some intuition for the proof here. With probability $1 - \delta$, `Intervals` returns a good partition, and, for each interval, of which there are at most $k' - 1$, by Lemma 4.3.2, `Iwin` is as accurate as described by Lemma 4.3.1) (which follows from a union bound). Thus, for the remainder of the proof we assume the partition returned by `Intervals` is good and each call to `Iwin` is accurate. Then, by Lemma 4.3.3, the probability that a bidder *wins* with a bid in an interval is a close approximation to the probability she *bids* in that interval (both events are conditioned on all bids being at most the upper bound of the interval). These estimates multiplied together also give good estimates.

4.3.1 Subsets

The argument above extends directly to a more general scenario in which not all bidders necessarily show up each time, and instead there is some distribution over $2^{[k]}$ over which bidders show up each time the auction is run. As mentioned above, this is quite natural in settings where bidders are companies that may or may not need the auctioned resource at any given time, or keyword auctions where there is a distribution over keywords, and companies only participate in the auction of keywords that are relevant to them. To handle this case, we simply apply `Kaplan` to the subset of samples in which bidder i showed up when learning \widehat{F}_i . We then use the fact that, though the distribution over subsets of bidders showing up need not be product, the maximum bid value of the other bidders who show up with bidder i is a random variable that is independent of bidder i 's bid.

Thus, all the above arguments extend directly, with a few very minor modifications. First, the point down to which we can learn each bidder i 's distribution now should depend on the winning bids *when i is participating*. So, define $p_i = \min_p \{\mathbb{P}[\text{the winning bid is at most } p \mid i \text{ shows up}] \geq \gamma\}$: for each bidder i , we will learn their distributions accurately down to p_i . The sample complexity bound of Theorem 4.3.1 is now a sample complexity on observations of bidder i (and so

requires roughly a $1/q$ blowup in total sample complexity to learn the distribution for a bidder that shows up only a q fraction of the time).

4.4 Using the CDF estimates to set optimal reserves

As an example of how one might use these estimated CDFs, we now show how to use the results from the previous section to accomplish a nontrivial task: for each subset $S \subseteq [n]$, we would like to pick a nearly revenue-optimal reserve price r_S for a second-price auction on S , using the `Kaplan` estimators of bidders' CDFs, which we can construct using $\text{poly}(n, \epsilon, \gamma)$ samples.⁷ `Kaplan` ensures we can estimate each bidder's distribution down to the price below which there is small probability γ of that bidder bidding and winning, which we call $p_i(\gamma)$. These estimators can be multiplied together to get accurate estimates of other events, such as all bidders in S bidding at most some amount. With these tools in hand, it is not hard to estimate the reserve from each reserve r : the best reserve according to these estimates will, in turn, be a nearly revenue-optimal reserve. We state the main theorem from this section with its explicit dependence on H , the maximum bid.

Theorem 4.4.1. *Consider the `Kaplan` estimators for $\gamma \leq \frac{\epsilon}{Hn}$ and $\epsilon' = \frac{\epsilon^2 \gamma}{32n^2 H^2}$. From these, for each $S \subseteq [n]$, one can compute \hat{r}_S such that*

$$\mathbb{E}_{v_{S,S}}[\text{Rev}(\text{VCG}_{\hat{r}_S}(S))] \geq \mathbb{E}_{v_S}[\text{Rev}(\text{VCG}_{r_S}(S))] - \epsilon$$

where r_S is the revenue-optimal reserve for S .

We now state a corollary of Theorem 4.3.1 which will allow us to deal in multiplicative approximation rather than additive approximation for the remainder of this section. In particular, it implies that, above the critical points $p_i(\gamma)$, `Kaplan` actually produces CDFs which are *multiplicative approximations*.

Corollary 4.4.1. *For each bidder i , and each $x \geq p_i(\gamma)$, for the `Kaplan-Meier` estimator we have that*

$$\left(1 - \frac{\epsilon'}{\gamma}\right) \leq \frac{\hat{F}_i(y)}{F_i(y)} \leq \left(1 + \frac{\epsilon'}{\gamma}\right).$$

Proof. At each point $y \geq p_i(\gamma)$, agent i has probability at least γ of bidding and winning with a bid at most x . This implies, in particular, that $F_i(y) \geq \gamma$. The claim follows from algebra and Theorem 4.3.1. ■

For the remainder of the argument, let ϵ', r both be fixed and set later.

Lemma 4.4.2. *Fix some $y \geq \max_{i \in S} p_i(\gamma)$ for some subset S . Then, we can construct $\hat{L}(y)$ using `Kaplan` estimators, such that*

$$\left(1 - 2n \frac{\epsilon'}{\gamma}\right) \leq \frac{\mathbb{P}[\text{each bidder in } S \text{ bids at most } y]}{\hat{L}(y)} \leq \left(1 + 2n \frac{\epsilon'}{\gamma}\right)$$

so long as $\frac{\epsilon'}{\gamma} \leq \frac{1}{n^{1+\alpha}}$ for $\alpha > 0, \alpha = \Omega(1)$.

⁷We again stress that this is not the primary goal of constructing these estimators, but merely a showcase of one of the possible uses for them.

Proof. Consider the subset S . Since bidders' bids are independent,

$$\mathbb{P} \left[\max_{i \in S} b_i \leq y \right] = \prod_{i \in S} F_i(y)$$

Now, we construct our estimator, by using the Kaplan estimators to estimate the right-hand side:

$$\widehat{L}(y) = \prod_{i \in S} \widehat{F}_i(y).$$

Since $y \geq \max_i p_i(\gamma)$, by Corollary 4.4.1, we have

$$\left(1 - \frac{\epsilon'}{\gamma}\right)^{|S|} \leq \frac{\widehat{L}(y)}{\mathbb{P}[\max_{i \in S} b_i \leq y]} \leq \left(1 + \frac{\epsilon'}{\gamma}\right)^{|S|} \quad (4.3)$$

Now, we deal with the upper-bound from Equation 4.3 and claim the lower bound's argument is identical. We now express the upper bound as

$$\frac{\widehat{L}(y)}{\mathbb{P}[\max_{i \in S} b_i \leq y]} \leq \left(1 + \frac{\epsilon'}{\gamma}\right)^{|S|} \leq \left(1 + \frac{\epsilon'}{\gamma}\right)^n \leq \left(1 + 2n \frac{\epsilon'}{\gamma}\right)$$

where the final inequality follows from the fact that $\frac{\epsilon'}{\gamma} \leq \frac{1}{n^{1+\alpha}}$ and basic algebra. \blacksquare

Lemma 4.4.3. *For a subset of bidders S , let A be the event where each bidder bids less than y and exactly 1 bidder bids between $y - \rho$ and y . Let $y - \rho \geq \max_{i \in S} p_i(\gamma)$. Then, we can compute \widehat{A} with Kaplan estimators, such that*

$$\left| \mathbb{P}[A] - \widehat{A} \right| \leq 4n^2 \frac{\epsilon'}{\gamma}$$

Proof. Consider the subset S . Since the bids are independent, we know that

$$\begin{aligned} \mathbb{P}[A] &= \mathbb{P} \left[\max_{i \in S} b_i \in [y - \rho, y] \wedge \text{secondmax}_{i \in S} b_i \leq y - \rho \right] \\ &= \sum_{i \in S} \left((F_i(y) - F_i(y - \rho)) \prod_{j \in S \setminus \{i\}} F_j(y - \rho) \right) \end{aligned} \quad (4.4)$$

Then, expressing this term using the Kaplan CDF estimators, let

$$\widehat{A} = \sum_{i \in S} \left(\left(\widehat{F}_i(y) - \widehat{F}_i(y - \rho) \right) \prod_{j \in S \setminus \{i\}} \widehat{F}_j(y - \rho) \right) \quad (4.5)$$

where the final inequality comes from the fact that these probabilities sum to at most 1, since each term describes a disjoint event.

Now, we upper-bound \widehat{A} (the lower bound argument is identical). Combining Equations 4.4 and 4.5, Lemma 4.4.2 and Corollary 4.4.1, we have that

$$\begin{aligned}\widehat{A} &= \sum_{i \in S} \left((F_i(y) - F_i(y - \rho)) \prod_{j \in S \setminus \{i\}} F_j(y - \rho) \right) \\ &\leq \sum_{i \in S} \left(\prod_{j \in S \setminus \{i\}} F_j(y - \rho) \right) \left(\left(1 + \frac{2n\epsilon'}{\gamma}\right) F_i(y) - \left(1 - \frac{2n\epsilon'}{\gamma}\right) F_i(y - \rho) \right) \\ &\leq 4n^2 \frac{\epsilon'}{\gamma} + \sum_{i \in S} \left(F_i(y) \prod_{j \in S \setminus \{i\}} F_j(y - \rho) - F_i(y - \rho) \prod_{j \in S \setminus \{i\}} F_j(y - \rho) \right)\end{aligned}$$

where the final inequality comes from the fact that each $F_j(x) \leq 1$ for each j and x . Thus,

$$|\mathbb{P}[A] - \widehat{A}| \leq 4n^2 \frac{\epsilon'}{\gamma}$$

as desired. ■

Lemma 4.4.4. *Suppose $r \geq p_i(\gamma)$ for all $i \in S$. Then, we can compute $\widehat{P}(r)$ such that*

$$\left| \widehat{P}(r) - \mathbb{P}[\text{the winner pays } r] \right| \leq 4n^2 \frac{\epsilon'}{\gamma}$$

when the reserve is r , using the Kaplan estimator.

Proof. Since the winner pays r only when the highest bid surpasses r and all other bids are below r , the probability can be written as the term from Lemma 4.4.3, setting $y - \rho = r$ and $\rho = 1 - r$. ■

Finally, we state and prove a lemma which allows us to ignore the cases where the winning bid is below the point at which we've learned all bidders' distributions.

Lemma 4.4.5. *Suppose $\gamma \leq \frac{\epsilon}{2n}$. Then,*

$$\mathbb{P}_{S, v_S} \left[\text{winning bid is } \leq \max_{i \in S} p_i \right] \leq \frac{\epsilon}{2H}$$

Proof. We begin by noticing that there are at most n maximizers of $\max_{i \in S} p_i$: each bidder i has some subset $\mathcal{S}_i \subseteq \mathcal{P}(n)$ of subsets for which $i = \operatorname{argmax}_{j \in S} \max_{S \in \mathcal{S}_i} p_j$. Assume without loss of generality that $\mathcal{S}_1, \dots, \mathcal{S}_n$ forms a partition of the power set of $[n]$ (if some $p_i = p_j$, break ties arbitrarily). Notice also that $i \in S$ for all $S \in \mathcal{S}_i$. Then, we can express the term which we wish

to bound by breaking it into these disjoint events:

$$\begin{aligned}
\mathbb{P}_{\mathcal{S}, v_{\mathcal{S}}} \left[\text{winning bid is } \leq \max_{i \in S} p_i \right] &= \sum_{j \in [n]} \mathbb{P}_{\mathcal{S}, v_{\mathcal{S}}} \left[\text{winning bid is } \leq \max_{i \in S} p_i \wedge S \in \mathcal{S}_j \right] \\
&= \sum_{j \in [n]} \mathbb{P}_{\mathcal{S}, v_{\mathcal{S}}} \left[\text{winning bid is } \leq p_j \wedge S \in \mathcal{S}_j \right] \\
&= \sum_{j \in [n]} \mathbb{P}_{\mathcal{S}, v_{\mathcal{S}}} \left[\text{winning bid is } \leq p_j \wedge j \in S \in \mathcal{S}_j \right] \\
&\leq \sum_{j \in [n]} \mathbb{P}_{\mathcal{S}, v_{\mathcal{S}}} \left[\text{winning bid is } \leq p_j \wedge j \in S \right] \\
&\leq \sum_{j \in [n]} \mathbb{P}_{\mathcal{S}, v_{\mathcal{S}}} \left[\text{winning bid is } \leq p_j | j \in S \right] \\
&\leq n\gamma \leq \frac{\epsilon}{2H}
\end{aligned}$$

where the first inequality comes from the fact that the \mathcal{S}_i s form a partition, the second from the definition of \mathcal{S}_i in terms of p_i being maximal, the third from the fact that $i \in S$ for all $S \in \mathcal{S}_i$ by definition, the first inequality from the fact that $\mathbb{P}[A \wedge B] \leq \mathbb{P}[A]$, the second inequality from the fact that $\mathbb{P}[A|B] \leq \mathbb{P}[A \wedge B]$, the penultimate from the definition of p_i and the final from the upper-bound on γ . ■

We now state and prove our final lemma, which shows that we can accurately estimate the revenue of any reserve r on subset S so long as $r \geq \max_{i \in S} p_i$.

Lemma 4.4.6. *Fix some subset S and some $r \geq \max_{i \in S} p_i(\gamma)$. Then, Kaplan estimators with parameters ϵ' and γ are sufficient to compute a \widehat{R}_r such that*

$$|\mathbb{E}[\text{Rev}(\text{VCG}_r(S)) | \text{payment is } > r] \cdot \mathbb{P}[\text{payment is } > r] - \widehat{R}_r| \leq \rho + \frac{H^2}{\rho} \left(8n^2 \frac{\epsilon'}{\gamma} + 4n \frac{\epsilon'}{\gamma} \right)$$

for any fixed $\rho > 0$.

Proof. Let $B(y, y + \rho)$ be the event that the winner of the auction pays between y and $y + \rho$ when the reserve is $r < y$. We will first show we can accurately estimate $\mathbb{P}[B(y, y + \rho)]$ if ρ is small. So, we notice that

$$\begin{aligned}
\mathbb{P}[B(y, y + \rho)] &= \mathbb{P}[\text{two or more bidders bid } \geq y \text{ but at most of them bids } \geq y + \rho] \\
&= \mathbb{P}[\text{two or more bidders bid } \geq y \text{ and none bids } \geq y + \rho] \\
&\quad + \mathbb{P}[\text{two or more bidders bid } \geq y \text{ and exactly one bids } \geq y + \rho]. \quad (4.6)
\end{aligned}$$

Let

$$X = \mathbb{P}[\text{two or more bidders bid } \geq y \text{ and none bids } \geq y + \rho]$$

and

$$Y = \mathbb{P}[\text{two or more bidders bid } \geq y \text{ and exactly one bids } \geq y + \rho].$$

We will manipulate each separately and write them as products and sums of CDFs. First, we can write X as follows:

$$\begin{aligned} X &= \mathbb{P}[\text{all bidders bid at most } y + \rho] - \mathbb{P}[\text{all bidders bid at most } y] \\ &\quad - \mathbb{P}[\text{all but one bidder bid at most } y \text{ and one bids } \in [y, y + \rho]] \\ &= \prod_{i \in S} F_i(y + \rho) - \prod_{i \in S} F_i(y) - \sum_{i \in S} (F_i(y + \rho) - F_i(y)) \prod_{j \in S \setminus \{i\}} F_j(y) \end{aligned}$$

where the first equality follows because all bidders bid at most $y + \rho$ and at least two people bid between y and $y + \rho$ exactly when all bidders bid at most $y + \rho$, and neither exactly 0 nor exactly 1 bidder bids between y and $y + \rho$. Let \hat{X} be X expressed using the Kaplan estimators of the CDFs. By Lemma 4.4.2 and Lemma 4.4.3, we have that

$$|X - \hat{X}| \leq 4n \frac{\epsilon'}{\gamma} (1 + n). \quad (4.7)$$

Now, we argue about Y in the same way:

$$\begin{aligned} Y &= \mathbb{P}[\text{two or more bidders bid } \geq y \text{ and exactly one bids } \geq y + \rho] \\ &= \sum_{i \in S} (1 - F_i(y + \rho)) \left(\prod_{j \in S \setminus \{i\}} F_j(y + \rho) - \prod_{j \in S \setminus \{i\}} F_j(y) \right) \end{aligned}$$

where the second equality holds because exactly one bidder needs to bid between above $y + \rho$, and at least one other bidder needs to bid between y and $y + \rho$. Let \hat{Y} be the representation of Y using Kaplan estimators. Then, using Corollary 4.4.1 Lemma 4.4.2, we have that

$$|Y - \hat{Y}| \leq 4n^2 \frac{\epsilon'}{\gamma} \quad (4.8)$$

Finally, using $\hat{B}(y, y + \rho)$ to represent $\mathbb{P}[B(y, y + \rho)]$ written with the Kaplan estimators from Equation 4.6, combining Equations 4.7, 4.8, it is the case that

$$|\hat{B}(y, y + \rho) - \mathbb{P}[B(y, y + \rho)]| \leq 8n^2 \frac{\epsilon'}{\gamma} + 4n \frac{\epsilon'}{\gamma}. \quad (4.9)$$

We now have the tools to construct our estimator. Let $g_i(t)$ represent the pdf of the distribution over second-highest bids from S evaluated at t . Then,

$$\mathbb{E}[\text{Rev}(\text{VCG}_r(S)) | \text{payment is } > r] \cdot \mathbb{P}[\text{payment is } > r] = \int_r^H t \cdot g_i(t) dt. \quad (4.10)$$

Notice that, since $r + k\rho$ is a lower-bound on the revenue when there are at least two bidders above $r + k\rho$, we can construct an underestimator of the revenue

$$\int_r^H t \cdot g_i(t) dt - \sum_{k=0}^{\frac{H-r}{\rho}} \mathbb{P}[B(r + k\rho, r + (k+1)\rho)] \cdot (r + k\rho) \geq 0 \quad (4.11)$$

but also that this estimator isn't underestimating by more than ρ , since at most one bidder is bidding above $r + (k + 1)\rho$, that

$$\begin{aligned} & \int_r^H t \cdot g_i(t) dt - \sum_{k=0}^{\frac{H-r}{\rho}} \mathbb{P}[B(r + k\rho, r + (k + 1)\rho)] \cdot (r + k\rho) \\ & \leq \rho \sum_{k=0}^{\frac{H-r}{\rho}} \mathbb{P}[B(r + k\rho, r + (k + 1)\rho)] \leq \rho, \end{aligned} \quad (4.12)$$

which follows since the sum in the second-to last term is a sum of probabilities of disjoint events, summing to at most one. Then, let

$$\widehat{R}_r = \sum_{k=0}^{\frac{H-r}{\rho}} \widehat{B}(r + k\rho, r + (k + 1)\rho) \cdot (r + k\rho) \quad (4.13)$$

Thus, we have that

$$\begin{aligned} & |\widehat{R}_r - \mathbb{E}[\text{Rev}(\text{VCG}_r(S)) \mid \text{payment is } > r] \cdot \mathbb{P}[\text{payment is } > r]| \\ & \leq \left| \int_r^H t \cdot g_i(t) dt - \sum_{k=0}^{\frac{H-r}{\rho}} \mathbb{P}[B(r + k\rho, r + (k + 1)\rho)] \cdot (r + k\rho) \right| \\ & \quad + \left| \widehat{R}_r - \sum_{k=0}^{\frac{H-r}{\rho}} \mathbb{P}[B(r + k\rho, r + (k + 1)\rho)] \cdot (r + k\rho) \right| \\ & \leq \rho + \sum_{k=0}^{\frac{H-r}{\rho}} |\widehat{B}(r + k\rho, r + (k + 1)\rho) - \mathbb{P}[B(r + k\rho, r + (k + 1)\rho)]| \cdot (r + k\rho) \\ & \leq \rho + H \sum_{k=0}^{\frac{H-r}{\rho}} |\widehat{B}(r + k\rho, r + (k + 1)\rho) - \mathbb{P}[B(r + k\rho, r + (k + 1)\rho)]| \\ & \leq \rho + H \sum_{k=0}^{\frac{H-r}{\rho}} \left(8n^2 \frac{\epsilon'}{\gamma} + 4n \frac{\epsilon'}{\gamma} \right) \leq \rho + \frac{H^2}{\rho} \left(8n^2 \frac{\epsilon'}{\gamma} + 4n \frac{\epsilon'}{\gamma} \right) \end{aligned}$$

where the first inequality follows from the triangle inequality and Equation 4.10, the second inequality from Equations 4.13 and 4.12, the third follows since $r + k\rho \leq H$, the fourth from Equation 4.9, and the final from the fact that there are at most $\frac{H}{\rho}$ terms in the sum. This completes the proof. \blacksquare

We now have the necessary components to prove the main theorem: since we can accurately estimate the probability each reserve is paid, and the expected revenue when the payment is above the reserve, the theorem follows easily.

Proof of Theorem 4.4.1. We express the revenue of running VCG for the subset S with reserve r as

$$\mathbb{E}[\text{Rev}(\text{VCG}_r(S))] = r \cdot \mathbb{P}[\text{exactly one bidder bids } \geq r] + \int_r^H t \cdot g_i(t) dt. \quad (4.14)$$

Now, define

$$\widehat{\mathbb{E}}[\text{Rev}(r)] = r \cdot \widehat{P}(r) + \widehat{R}_r \quad (4.15)$$

Then, using Lemmas 4.4.4 and 4.4.6, we have that

$$\begin{aligned} \left| \mathbb{E}[\text{Rev}(r)] - \widehat{\mathbb{E}}[\text{Rev}(r)] \right| &\leq r \cdot \frac{4n^2\epsilon'}{\gamma} + \rho + H \left(8n^2 \frac{\epsilon'}{\gamma} + 4n \frac{\epsilon'}{\gamma} \right) \frac{H}{\rho} \\ &\leq H \cdot \frac{4n^2\epsilon'}{\gamma} + \rho + H \left(8n^2 \frac{\epsilon'}{\gamma} + 4n \frac{\epsilon'}{\gamma} \right) \frac{H}{\rho} \\ &\leq \rho + H^2 \frac{16n^2\epsilon'}{\rho\gamma} \end{aligned}$$

Then, setting $\epsilon' = \frac{\epsilon^2\gamma}{64n^2H^2}$, $\rho = \frac{\epsilon}{2}$, this reduces to additive error at most $\frac{\epsilon}{2}$. This is true simultaneously for all $r \geq \max_{i \in S^t} p_i$, so the estimator-optimal reserve will have revenue within $\frac{\epsilon}{2}$ of the optimal reserve above $r \geq \max_{i \in S^t} p_i$. Finally, Lemma 4.4.5 implies that there is at most $\frac{\epsilon}{2H}$ probability of the winning bid being below $\max_{i \in S} p_i$, so at most $\frac{\epsilon}{2}$ revenue is lost by ignoring these events (and using $r \geq \max_{i \in S} p_i$ instead of a smaller reserve). ■

4.5 Extensions and Other Models

So far we have been in the usual model of independent private values. That is, on each run of the auction, bidder i 's value is $v_i \sim \mathcal{D}_i$, drawn independently from the other v_j . We now consider models motivated by settings where we have different items being auctioned on each round, such as different cameras or cars, and these items have observable properties, or features, that affect their value to each bidder.

In the first (easier) model we consider, each bidder i has its own private weight vector $w_i \in R^d$ (which we don't see), and each item is a feature vector $x \in R^d$ (which we do see). The value for bidder i on item x is $w_i \cdot x$, and the winner is the highest bidder $\arg\max_i w_i \cdot x$. There is a distribution \mathcal{P} over items, but no additional private randomness. Our goal, from submitting bids and observing the identity of the winner, is to learn estimates \tilde{w}_i that approximate the true w_i in the sense that for random $x \sim \mathcal{P}$, with probability $\geq 1 - \epsilon$, the \tilde{w}_i correctly predict the winner and how much the winner values the item x up to $\pm\epsilon$.

In the second model we consider, there is a single common vector w , but we reintroduce the distributions \mathcal{D}_i . In particular, the value of bidder i on item x is $w \cdot x + v_i$ where $v_i \sim \mathcal{D}_i$. The " $w \cdot x$ " portion can be viewed as a common value due to the intrinsic worth of the object, and if

$w = \vec{0}$ then this reduces to the setting studied in previous sections. our goal is to learn both the common vector w and each \mathcal{D}_i .

The common generalization of the above two models, with different unknown vectors w_i and unknown distributions \mathcal{D}_i appears to be quite a bit more difficult (in part because the expected value of a draw from \mathcal{D}_i conditioned on bidder i winning depends on the vector x). We leave as an open problem to resolve learnability (positively or negatively) in such a model. We assume that $\|x\|_2 \leq 1$ and $\|w_i\|_2 \leq 1$, and as before, all valuations are in $[0, 1]$.

4.5.1 Private value vectors without private randomness

Here we present an algorithm for the setting where each bidder i has its own private vector $w_i \in R^d$, and its value for an item $x \in R^d$ is $w_i \cdot x$. There is a distribution \mathcal{P} over items, and our goal, from submitting bids and observing the identity of the winner, is to accurately predict the winner and the winning bid. Specifically, we prove the following:

Theorem 4.5.1. *With probability $\geq 1 - \delta$, the algorithm below using sample size*

$$m = O\left(\frac{1}{\epsilon^2} [dn^2 \log(1/\epsilon) + \log(1/\delta)]\right)$$

produces \tilde{w}_i such that on a $1 - \epsilon$ probability mass of $x \sim \mathcal{P}$, $i^ \equiv \operatorname{argmax}_i \tilde{w}_i \cdot x = \operatorname{argmax}_i w_i \cdot x$ (i.e., a correct prediction of the winner), and furthermore*

$$|\tilde{w}_{i^*} \cdot x - w_{i^*} \cdot x| \leq \epsilon.$$

Proof. Our algorithm is simple. We will participate in m auctions using bids chosen uniformly at random from $\{0, \epsilon, 2\epsilon, \dots, 1\}$. We observe the winners, then solve for a consistent set of \tilde{w}_i using linear programming. Specifically, for $t = 1, \dots, m$, if bidder i_t wins item x_t for which we bid b_t , then we have linear inequalities:

$$\begin{aligned} \tilde{w}_{i_t} \cdot x_t &> \tilde{w}_j \cdot x_t \quad (\forall j \neq i_t) \\ \tilde{w}_{i_t} \cdot x_t &> b_t. \end{aligned}$$

Similarly, if we win the item, we have:

$$b_t > \tilde{w}_j \cdot x_t \quad (\forall j).$$

Let \mathcal{P}^* denote the distribution over pairs (x, b) induced by drawing x from \mathcal{P} and b uniformly at random from $\{0, \epsilon, 2\epsilon, \dots, 1\}$ and consider a $(k + 1)$ -valued target function f^* that given a pair (x, b) outputs an integer in $\{0, 1, \dots, n\}$ indicating the winner (with 0 indicating that our bid b wins). By design, the vectors $\tilde{w}_1, \dots, \tilde{w}_n$ solved for above yield the correct answer (the correct highest bidder) on all m pairs (x, b) in our training sample. We argue below that m is sufficiently large so that by a standard sample complexity analysis, with probability at least $1 - \delta$, the true error rate of the vectors \tilde{w}_i under \mathcal{P}^* is at most $\epsilon^2/(1 + \epsilon)$. This in particular implies that for at least a $(1 - \epsilon)$ probability mass of items x under \mathcal{P} , the vectors \tilde{w}_i predict the correct winner for all $\frac{1+\epsilon}{\epsilon}$ bids $b \in \{0, \epsilon, 2\epsilon, \dots, 1\}$ (by Markov's inequality). This implies that for this $(1 - \epsilon)$

probability mass of items x , not only do the \tilde{w}_i correctly predict the winning bidder but they also correctly predict the winning bid value up to $\pm\epsilon$ as desired.

Finally, we argue the bound on m . Any given set of n vectors $\tilde{w}_1, \dots, \tilde{w}_n$ induces a $(n+1)$ -way partition of the $(d+1)$ -dimensional space of pairs (x, b) based on which of $\{0, \dots, n\}$ will be the winner (with 0 indicating that b wins). Each element of the partition is a convex region defined by halfspaces, and in particular there are only $O(n^2)$ hyperplane boundaries, one for each pair of regions. Therefore, the total number of ways of partitioning m data-points is at most $O(m^{(d+1)n^2})$. The result then follows by standard VC upper bounds for desired error rate $\epsilon^2/(1+\epsilon)$. ■

4.5.2 Common value vectors with private randomness

We now consider the case that there is just a single common vector w , but we reintroduce the distributions \mathcal{D}_i . In particular, there is some distribution \mathcal{P} over $x \in R^d$, and the value of bidder i for item x is $w \cdot x + v_i$ where $v_i \sim \mathcal{D}_i$. As before, we assume $\|x\|_2 \leq 1$ and $\|w_i\|_2 \leq 1$, and all valuations are in $[0, 1]$. The goal of the algorithm is to learn both the common vector w and each \mathcal{D}_i . We now show how we can solve this problem by first learning a good approximation \tilde{w} to w which then allows us to reduce to the problem of Section 4.3. In particular, given parameter ϵ' , we learn \tilde{w} such that

$$\Pr_{x \sim \mathcal{P}} (|w \cdot x - \tilde{w} \cdot x| \leq \epsilon') \geq 1 - \epsilon'.$$

Once we learn such a \tilde{w} , we can reduce to the case of Section 4.3 as follows: every time the algorithm of Section 4.3 queries with some reserve bid b , we submit instead the bid $b + \tilde{w} \cdot x$. The outcome of this query now matches the setting of independent private values, but where (due to the slight error in \tilde{w}) after the v_i are each drawn from \mathcal{D}_i , there is some small random fluctuation that is added (and an ϵ' fraction of the time, there is a large fluctuation). But since we can make ϵ' as polynomially small as we want, this becomes a vanishing term in the independent private values analysis. Thus, it suffices to learn a good approximation \tilde{w} to w , which we do as follows.

Theorem 4.5.2. *With probability $\geq 1 - \delta$, the algorithm below using running time and sample size polynomial in $d, n, 1/\epsilon'$, and $\log(1/\delta)$, produces \tilde{w} such that*

$$\Pr_{x \sim \mathcal{P}} [|\tilde{w} \cdot x - w \cdot x| \leq \epsilon'] \geq 1 - \epsilon'.$$

Proof. Let \mathcal{D}_{max} denote the distribution over $\max[v_1, \dots, v_n]$. By performing an additive offset, specifically, by adding a new feature x_0 that is always equal to 1 and setting the corresponding weight w_0 to be the mean value of \mathcal{D}_{max} , we may assume without loss of generality from now on that \mathcal{D}_{max} has mean value 0.⁸

Now, consider the following distribution over labeled examples (x, y) . We draw x at random from \mathcal{P} . To produce the label y , we bid a uniform random value in $[0, 1]$ and set $y = 1$ if we lose and $y = 0$ if we win (we ignore the identity of the winner when we lose). The key point here is that if the highest bidder for some item x bid a value $b \in [0, 1]$, then with probability b we lose and set $y = 1$ and with probability $1 - b$ we win and set $y = 0$. So, $\mathbb{E}[y] = b$. Moreover,

⁸Adding such an x_0 and w_0 has the effect of modifying each v_i to $v_i - E[v_{max}]$. The resulting distributions over $w \cdot x + v_i$ are all the same as before, but now \mathcal{D}_{max} has a zero mean value.

since $b = w \cdot x + v_{max}$, where v_{max} is picked from \mathcal{D}_{max} which has mean value of 0, we have $\mathbb{E}[b|x] = w \cdot x$. So, $\mathbb{E}[y|x] = w \cdot x$.

So, we have examples x with labels in $\{0, 1\}$ such that $\mathbb{E}[y|x] = w \cdot x$. This implies that $w \cdot x$ is the predictor of minimum squared loss over this distribution on labeled examples (in fact, it minimizes mean squared error for every point x). Moreover, any real-valued predictor $h(x) = \tilde{w} \cdot x$ that satisfies the condition that $\mathbb{E}_{(x,y)}[(\tilde{w} \cdot x - y)^2] \leq \mathbb{E}_{(x,y)}[(w \cdot x - y)^2] + \epsilon'^3$ must satisfy the condition:

$$\Pr_{x \sim \mathcal{P}} (|w \cdot x - \tilde{w} \cdot x| \leq \epsilon') \geq 1 - \epsilon'.$$

This is because a predictor that fails this condition incurs an additional squared loss of ϵ'^2 on at least an ϵ' probability mass of the points. Finally, since all losses are bounded (we know all values $w \cdot x$ are bounded since we have assumed all valuations are in $[0, 1]$, so we can restrict to \tilde{w} such that $\tilde{w} \cdot x$ are all bounded), standard confidence bounds imply that minimizing mean squared error over a sufficiently (polynomially) large sample will achieve the desired near-optimal squared loss over the underlying distribution. ■

4.6 Open Questions

There are several interesting lines for future work in this direction. Standard sample complexity results in the auction theory community make the basic assumption that n independent draws (one from each bidder's distribution) make up a sample, and there are several relaxations of this assumption that would be interesting to consider. First, while we considered learning valuation distributions using only the identity of the winning bidder (and the ability to set a reserve), it would be interesting to know whether or not similar results are possible in a second-price auction where the observation is the winner and the price she paid (but it is not possible to set a reserve). It will not be possible to recover anomalously large bids (akin to the impossibility of recovering low bids in general), since there will (almost) never be price-setters whose bid is so high. We also think it would be interesting to understand the connections between approximately learning bidders' CDFS and approximating their virtual valuation function.

Chapter 5

Student-truthful, school-optimal many-to-one stable matching (via differential privacy)

5.1 Introduction

In this chapter, we turn to using differential privacy as a tool for mechanism design. In particular, we compute many-to-one stable matchings in a way which is private, and thus approximately truthful, for one side of the market. It has long been known that a mechanism for selecting an outcome of a game which is differentially private is also approximately truthful [95]¹. In settings where the mechanism aims to optimize for social welfare, and has the ability to assess payments, Huang and Kannan [77] show it is possible to use VCG and the exponential mechanism with appropriate payments to achieve exact truthfulness, differential privacy, and approximate social welfare optimality. The necessary tradeoff between privacy and utility limits the precise usefulness of this approach: the social welfare of the latter mechanism is necessarily worse than the welfare-maximizing allocation by an additive term which is linear n , the number of players. This contrasts with the fact that, in many interesting cases (such as making general statistical queries), the tension between privacy and utility lessens as n grows: queries should release less information about each individual if there are enough other individuals in the population. However, in the standard setting where a mechanism's output is an outcome to the game, that outcome is necessarily either sensitive to each agent's preferences, or unable to closely approximate social welfare.

This dissonance has led to considering private mechanisms whose output are not necessarily outcomes of the game induced by the mechanism, but instead some coordinating information which players can interpret, or prices, some restriction on their allowed actions, or other information from which agents can make their own choices about their own final, individual action. Then, each agent is allowed to choose her behavior with respect to those prices or other information. Her behavior will necessarily be sensitive to her preferences, but *the information provided to all*

¹Indeed, in many interesting settings, a differentially private mechanism can be converted into an exactly truthful mechanism, losing the privacy of the mechanism in the process, see [].

(other) players might not need to depend so heavily on her preferences. This insight led to the investigation of *jointly* differentially private mechanisms [84]: agent i 's behavior or "piece" of the outcome will not be private in i 's preferences, but the $n - 1$ other agent's behavior or "pieces" of the outcome will be. Work on computing equilibria [84, 112] and Walrasian prices [76] in a way which is jointly private has shown that, constrained by this weakened notion of privacy, it is possible to get approximations (in nearness to equilibria and to social welfare) which degrade *sublinearly* in the number of players: thus, as n grows, the guarantees *get better on average for each player*. In most cases, the mechanisms compute some suggested action [84, 112] or prices [76] in a way which is differentially private (in the standard sense), and then allows each agent to choose her behavior. The final outcome is then jointly differentially private (from inspecting $n - 1$ agents' outcomes or behaviors, very little can be learned about the final agents' behavior or preferences).

In this work, we show how to use this perspective to compute many-to-one stable matchings in a way which is (approximately) truthful for one side of the market, using differential privacy. This matching solution concept is used for diverse applications, including matching students to schools [1], and medical residents to hospitals [116, 118]. There are two sides of such a market, which we will without loss of generality refer to as the *students* and the *schools*. The goal is to find a feasible assignment μ of students to schools – each student a can be matched to at most 1 school, but each school u can be potentially matched to up to C_u students, where C_u is the *capacity* of school u . We would like to find a matching that is *stable*. Informally, when each student a has a preference ordering \succ_a over schools, and each school u has a preference ordering \succ_u over students, then an assignment μ forms a stable matching if it is feasible, and there is no student-school pair (a, u) such that they are unmatched ($\mu(a) \neq u$), but such that they would mutually prefer to deviate from the proposed matching μ and match with each other.

The set of stable many-to-one-matchings have a remarkable structural property: there exists a *school optimal* and a *student optimal* stable matching – i.e. a stable matching that all schools simultaneously prefer to all other stable matchings, and a stable matching that all students simultaneously prefer to all other stable matchings. Moreover, these matchings are easy to find, with the school-proposing (respectively, student proposing) version of the Gale-Shapley deferred acceptance algorithm [62]. Unfortunately, the situation is not quite as nice when student incentives are taken into account. Even in the 1-to-1 matching case (i.e. when capacities $C_u = 1$ for all schools), there is no mechanism which makes truthful reporting of one's preferences a dominant strategy for both sides of the market [115]. In the many-to-one matchings case, things are even worse: an algorithm which finds the school optimal stable matching does not incentivize truthful reporting for either the students or the schools [116].

Because of this, a literature has emerged studying the incentive properties of stable matching algorithms under *large market assumptions* (e.g. [79, 88, 90]). In general, this literature has taken the following approach: make restrictive assumptions about the market (e.g. that students preference lists are only of constant length and are drawn uniformly at random), and under those assumptions, prove that an algorithm which computes exactly the school optimal stable matching makes truthful reporting a dominant strategy for a $1 - o(1)$ fraction of student participants (generally even under these assumptions, the schools still have incentive to misreport if the algorithm computes the school optimal stable matching).

In this work, we take a fundamentally different approach. We make absolutely no assump-

tions about student or school preferences, allowing them to be worst-case. We also insist on giving incentive guarantees to *every* student, not just most students. We compute (in a sense to be defined) an *approximately* stable and *approximately* school optimal matching using an algorithm with a particular insensitivity property (differential privacy), and show that truthful reporting is an *approximately* dominant strategy for *every* student in the market. These approximations become perfect as the size of the market grows large. Our notion of a “large market” requires only that the capacity of each school C_u grows with (the square root of) the number of schools, and (logarithmically) with the number of students, and does not require any assumption on how preferences are generated.

5.1.1 Our Results and Techniques

We recall the standard notion of stability in a many-to-one matching market with n students $a_i \in A$ and m schools $u_j \in U$, each with capacity C_j .

Definition 5.1.1. A matching $\mu : A \rightarrow U$ is *feasible* and *stable* if:

1. (Feasibility) For each $u_j \in U$, $|\{i : \mu(a_i) = u_j\}| \leq C_j$
2. (No Blocking Pairs with Filled Seats) For each $a_i \in A$, and each $u_j \in U$ such that $\mu(a_i) \neq u_j$, either $\mu(a_i) \succ_{a_i} u_j$ or for every student $a'_i \in \mu^{-1}(u_j)$, $a'_i \succ_{u_j} a_i$.
3. (No Blocking Pairs with Empty Seats) For every $u_j \in U$ such that $|\mu^{-1}(u_j)| < C_j$, and for every student $a_i \in A$ such that $a_i \succ_{u_j} \emptyset$, $\mu(a_i) \succ_{a_i} u_j$.

Our notion of approximate stability relaxes condition 3. Informally, we still require that there be no blocking pairs among students and *filled* seats, but we allow each school to possibly have a small number of empty seats. We view this as a mild condition, reflecting the reality that schools are not able to perfectly manage yield, and are often willing to accept a small degree of under-enrollment.

Definition 5.1.2 (Approximate Stability). A matching $\mu : A \rightarrow U$ is *feasible* and α -*approximately stable* if it satisfies conditions 1 and 2 (*Feasibility* and *No Blocking Pairs with Filled Seats*) and:

3. (No Blocking pairs with Empty Seats at Under-Enrolled Schools) For every $u_j \in U$ such that $|\mu^{-1}(u_j)| < (1 - \alpha)C_j$, and for every student $a_i \in A$ such that $a_i \succ_{u_j} \emptyset$, $\mu(a_i) \succ_{a_i} u_j$.

We also employ a strong notion of approximate dominant strategy truthfulness, related to first order stochastic dominance – informally, we say that a mechanism is η -approximately dominant strategy truthful if no agent can gain more than η in expectation (measured by *any* cardinal utility function consistent with his ordinal preferences) by misreporting his preferences to the mechanism.

Finally, we, define a notion of school optimality that applies to approximately stable matchings. Informally, we say that an approximately stable matching μ (in the above sense) is *school dominant* if when compared to the school optimal *exactly* stable matching μ' , for every school u_j , every student a_i matched to u_j in μ is strictly preferred by u_j to any student matched to u_j in μ' but not in μ .

We can now give an informal statement of our main result.

Theorem 5.1.1 (Informal). *There is an algorithm for computing feasible and α -approximately stable school dominant matchings that makes truthful reporting an η -approximate dominant*

strategy for every student in the market, under the condition that for every school u , the capacity is sufficiently large, e.g.

$$C_u \geq O\left(\frac{\sqrt{m}}{\eta\alpha} \cdot \text{polylog}(n)\right)$$

Remark 5.1.1. Note that no assumptions are needed about either school or student preferences, which can be arbitrary. The only large market assumption needed is that the capacity C_u of each school is large. If, as the market grows, school capacities grow with the total number of schools at a rate of $\Omega(m^{1/2+\varepsilon})$ for any constant ε , then both η and α can be taken to tend to 0 in the limit.

This result differs from the standard large market results in several ways. First, and perhaps most importantly, the result is worst-case over all possible preferences of both schools and students. Second, the guarantee states that *no* student may substantially gain by misreporting her preferences; previous results [79, 88] show that only a subconstant fraction of students might have (substantial) incentive to deviate. In exchange for these strong guarantees, we relax our notion of stability and school optimality to approximate notions, which can be taken to be exact in the limit as the market grows large (under the condition that school capacities grow at a sufficiently fast rate).

When we do make one of the assumptions on student preferences made in previous work, we get stronger claims than the one above. For example, when the length of the preference lists of students are taken to be bounded, as they are in Immorlica and Mahdian [79] and Kojima and Pathak [88] we can remove our dependence on the number of schools:

Theorem 5.1.2 (Informal). *Under the condition that all students have preference lists over at most k schools (and otherwise prefer to be unmatched), there is an algorithm for computing feasible and α -approximately stable school dominant matchings that makes truthful reporting an η -approximate dominant strategy for every student in the market, under the condition that for every school u , the capacity is sufficiently large:*

$$C_u \geq O\left(\frac{\sqrt{k}}{\eta\alpha} \cdot \text{polylog}(n)\right)$$

Remark 5.1.2. Note that if k is considered to be a constant, then this result requires school capacity to grow only poly-logarithmically with the number of students n .

Our results come from analyzing a differentially private variant of the classic deferred acceptance algorithm. Rather than having schools explicitly *propose* to students, we consider an equivalent variant in which schools u publish a set of “admissions thresholds” which allow any student a who is ranked higher than the current threshold of school u (according to the preferences of u) to enroll. These thresholds naturally induce a matching when each student enrolls at their favorite school, given the thresholds. We first show that if the thresholds are computed under the constraint of differential privacy, then the algorithm is approximately dominant strategy truthful for the students. We then complete the picture by deriving a differentially private algorithm, and showing that with high probability, it produces an approximately stable, school dominant matching.

5.2 Related Work

5.2.1 Incentives in Stable Matching

Stable matching has long been known to be incompatible with truthfulness: no algorithm which produces a stable matching is truthful for both sides of the market [115], though Gale-Shapley is known to be truthful for the side of the market which is proposing in the 1-to-1 setting. Several lines of work have investigated stable matching in *large* markets, where players’ preferences are drawn from some distribution, and considering properties of the market as n , the number of players, grows large. Much of this work reduces to arguing that few players will have more than one stable match; since only those players who have multiple stable matchings ever have incentive to misreport, this directly implies most players will have no incentive to misreport. For a tabular representation of the related work in terms of the expected number of stable matches an individual has under various assumptions, see Table 5.1.

Let \mathcal{D} be a fixed distribution over the set of n women. Consider the following process of generating length- k preference lists over women. Draw some $w_1 \sim \mathcal{D}$, and let w_1 be the first woman in a preference list. Now, let (w_1, \dots, w_{i-1}) be the first $(i-1)$ women, in order, drawn from \mathcal{D} . Draw $w_i \sim \mathcal{D}$ until $w_i \notin \{w_1, \dots, w_{i-1}\}$. We denote such a distribution over preference lists by \mathcal{D}^k . Immorlica and Mahdian [79] prove a generalization of a conjecture of Roth and Peranson [118], showing if the men draw their preference lists according to \mathcal{D}^k , the expected number of women with more than one stable match is $o(n)$ (as n grows, for fixed k). Since it is known that a person has incentive to misreport only if they have more than one stable partner, this implies that only a vanishingly small fraction of the women will have incentive to misreport to any stable matching process. They also show that any stable matching algorithm induces a Nash equilibrium for which a $1 - o(1)$ fraction of players behave truthfully. These results are extended by Kojima and Pathak [88] to the many-to-one matching setting. They show that student-optimal stable matchings, where colleges have arbitrary preferences, and the students have random preference lists of fixed length drawn as above, will have similar a subconstant fraction of schools which have incentive to misreport. Lee [90] considers a slightly different distributional assumption about preferences in the one-to-one setting, where he shows that only a small number of players have large incentive to misreport.

Azevedo and Budish [8] introduce the notion of “strategyproofness in the large”, and show that the Gale-Shapley algorithm satisfies this definition. Roughly, this means that fixing any (constant sized) typespace, and any distribution over that typespace, if player preferences are sampled i.i.d. from the typespace, then for any fixed η , as the number of players n tends to infinity, truthful reporting becomes an η -approximate Bayes Nash equilibrium. These assumptions can be restrictive however – note that this kind of result requires that there are many more players than there are “types” of preferences, which in particular (together with the full support assumption on the type distribution) requires that in the limit, there are infinitely many identical agents of each type. In contrast, our results do not require a condition like this.

To the best of our knowledge, our results are the first to give truthfulness guarantees in settings where both sides of the market have worst-case preference orderings. Unlike some prior work, even without distributional assumptions, we are able to give truthfulness guarantees to *every* student, not only a $1 - o(1)$ fraction of students. Under one of the assumptions used in

Reference	Assumptions	# of possible stable matches
Immorlica and Mahdian [79]	Random, i.i.d. preferences on male side	$1 - o(1)$ -fraction of women have ≤ 1 stable match
Kojima and Pathak [88]	Random, i.i.d. preferences on student side	$1 - o(1)$ -fraction of schools have more than 1 stable set of students
Pittel [109]	Uniform random preferences on both sides	Average rank of partner for side optimized for is $\log(n)$, $\frac{n}{\log(n)}$ for the non-optimized side
[6]	Uniform random preferences on both sides, n men, $n - 1$ women	Average rank for men's match $\frac{n}{3 \log(n)}$, for women's match $3 \log(n)$, in any stable matching

Figure 5.1: Related works with distributional assumptions on the preferences of one or both sides of the market. Under these assumptions, it is often possible to show that many agents have few (or one) stable partners. If an agent has zero or one stable partner, then she has no incentive to misreport.

prior work (namely, that the length of the preference lists is short), our results can be sharpened as well.

5.2.2 Differential Privacy as a Tool for Truthfulness

The study of differentially private algorithms [49] has blossomed in recent years. A comprehensive survey of the work in this area is beyond the scope of this work; here, we mention the work which relates directly to the use of differential privacy in constructing truthful mechanisms.

McSherry and Talwar [95] were the first to identify privacy as a tool for designing approximately truthful mechanisms. Nissim et al. [103] showed how privacy could be used as a tool to design *exactly* truthful mechanisms without needing monetary payments (in certain settings). Huang and Kannan [77] proved that the exponential mechanism, a basic tool in differential privacy introduced in McSherry and Talwar [95] is maximal in distributional range, which implies that there exist payments which make it exactly truthful. Kearns et al. [84] demonstrated a connection between private equilibrium computation and the design of truthful *mediators* (and also showed how to privately compute approximate correlated equilibria in large games). This work was extended by Rogers and Roth [112] who show how to privately compute *Nash* equilibria in large congestion games.

The paper most related to this work is Hsu et al. [76] which shows how to compute approximate Walrasian equilibria privately, when bidders have quasi-linear utility for money and the supply of each good is sufficiently large. In that paper, in the final allocation, every agent is matched to their *approximately* most preferred goods at the final prices. In our setting, there are several significant differences: first, in the Walrasian equilibrium setting, only agents have pref-

erences over goods (i.e. goods have no preferences of their own), but in our setting, both sides of the market have preferences. Second, although there is a conceptual relationship between “threshold scores” in stable matching problems and prices in Walrasian equilibria, the thresholds do not play the role of money in matching problems, and there is no notion of being matched to an “approximately” most preferred school.

5.3 Preliminaries

5.3.1 Many-to-one Matching

A many-to-one stable matching problem consist of m schools $U = \{u_1, \dots, u_m\}$ and n students $A = \{a_1, \dots, a_n\}$. Every student a has a preference ordering \succ_a over all the schools, and each school u has a preference ordering \succ_u over the students. Let \mathcal{P} denote the domain of all preference orderings over schools (so each $\succ_a \in \mathcal{P}$).

It will be useful for us to think of a school u 's ordering over students A as assigning a unique² *score* $\text{score}(u, a)$ to every student, in descending order (for example, these could be student scores on an entrance exam). Every school u has a capacity C_u , the maximum number of students the school can accommodate. A feasible matching μ is a mapping $\mu : A \rightarrow U \cup \emptyset$, which has the property each student a is paired with at most one school $\mu(a)$, and each school u is matched with at most C_u students: $|\mu^{-1}(u)| \leq C_u$. For notational simplicity, we will sometimes simply write $\mu(u)$ to denote the set of students assigned to school u .

A matching is α -approximately *stable* if it satisfies Definition 5.1.2. When computing matchings, it will be helpful for us to think instead about computing *admission thresholds* t_u for each school. A set of admission thresholds $t \in \mathbb{R}_{\geq 0}^m$ induces a matching μ in a natural way: every student $a \in A$ is matched to her most preferred school amongst those whose admissions thresholds are below her score at the school. Formally, for a set of admissions thresholds t , the induced matching μ^t is defined by:

$$\mu^t(a) = \arg \max_{\succ_a} \{u \mid \text{score}(u, a) \geq t_u\}$$

We say that a set of admission thresholds s is feasible and α -approximately stable if its induced matching $\mu^s(a)$ is feasible and α -approximately stable. Note that an α -stable matching is an exactly stable matching in a market in which schools have reduced capacity (where the capacity at each school is reduced by at most a $(1 - \alpha)$ factor).

Remark Definition 5.1.2 also implies that if a school u is under-enrolled by more than αC_u , its admission score $t_u = 0$. This means such a school is very unpopular and could not recruit enough students even without any admission criterion.

We now introduce a notion of approximate school optimality, which our algorithm guarantees.

²It is essentially without loss of generality that students are assigned unique scores. If not, we could break ties by a simple pre-processing step: add noise $\sum_{k=1}^l 2^{-k} b_k$ to each student's score, where each b_k is a random bit; if the scores are integral, the probability of having ties is $1/\text{poly}(n)$ as long as $l \geq O(\log(n))$.

Definition 5.3.1. A matching μ is **school-dominant** if, for each school u , for all $a \in \mu(u) \setminus \mu'(u)$ and all $a' \in \mu'(u) \setminus \mu(u)$, $a \succ_u a'$, where μ' is the school-optimal matching.

In words, a matching μ is school-dominant if for every school u , when comparing the set of students S_1 that u is matched to in μ but not in the school optimal matching μ' , and the set of students S_2 that u is matched to in the school optimal matching μ' , but is not matched to in μ , u strictly prefers every student in S_1 to every student in S_2 . (i.e. compared to the school optimal matching, a school may be matched to *fewer* students, but not to *worse* students.) We note that school-dominance alone is trivial to guarantee: in particular, the empty matching is school dominant. Only together with an upper bound on the number of empty seats allowed per school (for example, as guaranteed by α -approximate stability) is this a meaningful concept.

We want to give mechanisms that make it an *approximately dominant strategy* for students to report truthfully. We have to be careful about what we mean by this, since students a have ordinal preferences \succ_a , rather than cardinal utility functions $v_a : U \rightarrow [0, 1]$. We say that a cardinal utility function v_a is *consistent* with a preference ordering \succ_a if for every $u, u' \in U$, $u \succ_a u'$ if and only if $v_a(u) \geq v_a(u')$. We will say that a mechanism is η -approximately truthful for students if for every student, and every cardinal utility function v_a consistent with truthful \succ_a , truthful reporting is an η -approximate dominant strategy as measured by v_a .

Definition 5.3.2. Consider any randomized mapping $\mathcal{M} : \mathcal{P}^n \rightarrow U^n$. We say that \mathcal{M} is η -approximately dominant strategy truthful (or truthful) if for any vector of student preferences $\succ \in \mathcal{P}^n$, any student a , any utility function $v_a : U \rightarrow [0, 1]$ that is consistent with \succ_a , and any $\succ'_a \neq \succ_a$, we have:

$$\mathbb{E}_{\mu \sim \mathcal{M}(\succ)}[v_a(\mu(a))] \geq \mathbb{E}_{\mu \sim \mathcal{M}(\succ'_a, \succ_{-a})}[v_a(\mu(a))] - \eta.$$

Note that this definition is very strong, since it holds simultaneously for every utility function consistent with student preferences. When $\eta = 0$ it corresponds to first order stochastic dominance.

5.3.2 Differential Privacy Preliminaries

Our tool for obtaining approximate truthfulness is *differential privacy*, which we define in this section. We say that the “private data” of each student a consists of both her preference ordering $\succ_a \in \mathcal{P}$ over the schools and the scores $\text{score}(u, a) \in \mathcal{V}$ assigned by the schools. A private database $D \in (\mathcal{P} \times \mathcal{V})^n$ is a vector of n student profiles, and D and D' are neighboring databases if they differ in no more than one student record. In particular, our matching algorithms take n student profiles as input and produce a set of admission scores as output (i.e., range $\mathcal{R} = \mathcal{V}^m$).

Definition 5.3.3 (Dwork et al. [49]). An (randomized) algorithm $\mathcal{A} : (\mathcal{P} \times \mathcal{V})^n \rightarrow \mathcal{R}$ is (ϵ, δ) -*differentially private* if for every pair of neighboring databases $D, D' \in (\mathcal{P} \times \mathcal{V})^n$ and for every set of subset of outputs $S \subseteq \mathcal{R}$,

$$\Pr[\mathcal{M}(D) \in S] \leq \mathbb{E}(\epsilon) \Pr[\mathcal{M}(D') \in S] + \delta.$$

If $\delta = 0$, we say that \mathcal{M} is ϵ -*differentially private*.

5.3.3 Differentially Private Counters

The central privacy tool in our matching algorithm is the private streaming counter (for a more detailed discussion of differential privacy under continual observation, see Chan et al. [34] and Dwork et al. [50]) proposed by Chan et al. [34] and Dwork et al. [50]. Given a bit stream $\sigma = (\sigma_1, \dots, \sigma_T) \in \{-1, 0, 1\}^T$, a streaming counter $\mathcal{M}(\sigma)$ releases an approximation to $c_\sigma(t) = \sum_{i=1}^t \sigma_i$ at every time step t . Below, we define an accuracy property we will then use to describe the usefulness of these counters.

Definition 5.3.4. A streaming counter \mathcal{M} is (τ, β) -useful if with probability at least $1 - \beta$, for each time $t \in [T]$,

$$|\mathcal{M}(\sigma)(t) - c_\sigma(t)| \leq \tau.$$

For the rest of this chapter, let $\mathbf{Counter}(\varepsilon, T)$ denote the Binary Mechanism of Chan et al. [34], instantiated with parameters ε and T . $\mathbf{Counter}(\varepsilon, T)$ satisfies the following accuracy guarantee (further details may be found in Section 5.7.2). Our mechanism uses m different $\mathbf{Counter}$ s to maintain the counts of temporally enrolled students for all schools. The following theorem allows us to bound the error of each $\mathbf{Counter}$ through the collective sensitivity across all $\mathbf{Counter}$ s.

Theorem 5.3.1. *Suppose we have m bit streams such that the change of an agent's data affects at most k streams, and alters at most c bits in each stream. For any $\beta > 0$, the composition of m distinct $\mathbf{Counter}(\varepsilon/2c\sqrt{2kc\ln(1/\delta)}, T)$ s is (ε, δ) -differentially private, and (α, β) -useful for*

$$\alpha = \frac{16c\sqrt{kc\ln(1/\delta)}}{\varepsilon} \ln\left(\frac{2m}{\beta}\right) \left(\sqrt{\log(T)}\right)^5.$$

5.4 Algorithms Computing Private Matchings are Approximately Truthful

In this section, we prove that that private mechanisms can be used to compute matchings truthfully, motivating our investigation of privacy-preserving stable matching mechanisms. Consider an algorithm M which takes as input student preferences \succ and computes school thresholds s . If M is ε -differentially private, then the algorithm $A(\succ)$ which computes thresholds $s = M(\succ)$ and then outputs the induced matching μ^s is ε -approximately dominant strategy truthful. Note that this guarantee holds independent of stability. For lack of space, we relegate the proof to Appendix 5.8.

Theorem 5.4.1. *Let $M : \mathcal{P}^n \rightarrow \mathbb{R}_{\geq 0}^m$ be any (ε, δ) -differentially private mechanism which takes as input n student profiles and outputs m school thresholds. Let $A : \mathbb{R}_{\geq 0}^m \rightarrow U^n$ be the mechanism which takes as input m school thresholds s , and outputs the corresponding matching $A(s) = \mu^s$. Then the mechanism $A \circ M : \mathcal{P}^n \rightarrow U^n$ is $(\varepsilon + \delta)$ -approximately dominant strategy truthful.*

The intuition behind this theorem is simple: if a mechanism is private, a student's report has almost no effect on the realization of the school thresholds; Given a fixed set of school thresholds, he can do no better than reporting truthfully, which causes him to be matched to his most preferred school that will take him.

Proof of Theorem 5.4.1. Fix any vector of student preferences \succ , any player a , any utility function v_a consistent with \succ_a , and any deviation $\succ'_a \neq \succ_a$. Now consider player a 's utility for truth-telling. For $\varepsilon \leq 1$, we have

$$\begin{aligned}
\mathbb{E}_{\mu \sim A \circ M(\succ)}[v_a(\mu(a))] &= \mathbb{E}_{s \sim M(\succ)}[v_a(\arg \max_{\succ_a} \{u \mid \text{score}(u, a) \geq s_u\})] \\
&= \sum_s \Pr[M(\succ) = s] \cdot v_a(\arg \max_{\succ_a} \{u \mid \text{score}(u, a) \geq s_u\}) \\
&\geq \sum_s \mathbb{E}(-\varepsilon) \Pr[M(\succ'_a, \succ_{-a}) = s] \cdot v_a(\arg \max_{\succ'_a} \{u \mid \text{score}(u, a) \geq s_u\}) - \delta \\
&= \mathbb{E}(-\varepsilon) \mathbb{E}_{s \sim M(\succ'_a, \succ_{-a})}[v_a(\arg \max_{\succ'_a} \{u \mid \text{score}(u, a) \geq s_u\})] - \delta \\
&\geq \mathbb{E}(-\varepsilon) \mathbb{E}_{s \sim M(\succ'_a, \succ_{-a})}[v_a(\arg \max_{\succ'_a} \{u \mid \text{score}(u, a) \geq s_u\})] - \delta \\
&\geq (1 - \varepsilon) \mathbb{E}_{s \sim M(\succ'_a, \succ_{-a})}[v_a(\arg \max_{\succ'_a} \{u \mid \text{score}(u, a) \geq s_u\})] - \delta \\
&\geq \mathbb{E}_{\mu \sim M(\succ'_a, \succ_{-a})}[v_a(\mu(a))] - (\varepsilon + \delta).
\end{aligned}$$

where the first and last equalities follow from the definition of the induced matching μ^s , the first inequality follows from the differential privacy condition, and the second follows from the consistency of v_a with \succ_a . \blacksquare

5.5 Truthful School-Optimal Mechanism

In this section, we present the algorithm which proves our main result Theorem 5.1.1. Algorithm 1 computes an α -approximately stable and school-dominant matching, and enjoys approximate dominant strategy truthfulness for the student side. We assume the reader is familiar with `DA-SCHOOL`, the well-known school-proposing version of the deferred acceptance algorithm [62]. For a brief overview of `DA-SCHOOL` within our context of score thresholds, see Section 5.10. We now state a useful fact about deferred acceptance, whose proof can be found in Appendix 5.9.

Lemma 5.5.1. *Let μ_t be some matching which is an intermediate matching in a run of the school-proposing deferred acceptance algorithm. Then μ_t is school-dominant.*

Our algorithm, `Private-DA-SCHOOL`(ε, δ), is a private version of `DA-SCHOOL`. At each time t , each school will publish a threshold score (initially, for each school, this will be the maximum possible score for that school). Schools will lower their thresholds when they are under capacity; as they do so, some students will tentatively accept admission and some will reject or leave for other schools. Initially, all students will be unmatched. For a given student a , as soon as a school lowers its threshold below the score a has there, a will signal to the mechanism which school is her favorite of those for which her score passes their threshold. Then, as the schools continue to lower their thresholds to fill seats, if a school that a likes better than her current match lowers its threshold below her score, a will inform the mechanism that she wishes to switch to her new favorite.

Each school maintains a private counter of the number of students tentatively matched to the school. We let E be the additive error bound of the counters. The schools will reserve E

number of seats from their initial capacity to avoid being over-enrolled, so the algorithm is run as if the capacity at each school is $C_u - E$. Then each school can be potentially under-enrolled by $2E$ seats, but they would take no more than α fraction of all the seats as long as the capacity $C_u \geq 2E/\alpha$.

Algorithm 1: Private-DA-School(ε, δ)

Input: school capacities $\{C_u\}$, student preferences $\{\succ_a\}$ and scores $\{\text{score}(u, a)\}$, range of scores $[0, J]$
Output: a set of score thresholds $\{t_j\}$
initialize: for each school u_j and each student a_i
 $T = mnJ$, $\varepsilon' = \frac{\varepsilon}{16\sqrt{2m \ln(1/\delta)}}$, $E = \frac{128\sqrt{m \ln(1/\delta)}}{\varepsilon} \ln\left(\frac{2m}{\beta}\right) \left(\sqrt{\log(nT)}\right)^5$,
 $\text{counter}(u_j) = \text{Counter}(\varepsilon', nT)$, $t_j = J$, $\mu(a_i) = \emptyset$, $\hat{C}_{u_j} = C_{u_j} - E$.
while there is some under-enrolled school u_j : $\text{counter}(u_j) < \hat{C}_{u_j}$ and $t_{u_j} > 0$ **do**
 $t_{u_j} = t_{u_j} - 1$
 for all student a_i **do**
 if $\mu(a_i) \neq \text{argmax}_{\succ_{a_i}} \{u_j \mid \text{score}(u_j, a_i) \geq t_{u_j}\}$ **then**
 Send (-1) to $\text{counter}(\mu(a_i))$
 let $\mu(a_i) = \text{argmax}_{\succ_{a_i}} \{u_j \mid \text{score}(u_j, a_i) \geq t_{u_j}\}$
 Send 1 to $\text{counter}(\mu(a_i))$
 Send 0 to all other counters
 else
 Send 0 to all counters
 end if
 end for
end while
return Final threshold scores $\{t_j\}$ and matching μ^t

Now, we state the formal version of Theorem 5.1.1.

Theorem 5.1.1 *Algorithm 1 is (ε, δ) -differentially private, and hence $(\varepsilon + \delta)$ -approximately dominant strategy truthful. With probability at least $1 - \beta$, it outputs an α -approximately stable, school dominant matching, as long as the capacity at each school u satisfies $C_u \geq R = O\left(\frac{\sqrt{m}}{\varepsilon\alpha} \text{polylog}\left(n, m, \frac{1}{\delta}, \frac{1}{\beta}\right)\right)$.*

We prove Theorem 5.1.1 in two parts. Lemma 5.5.2 shows that Private-DA-School(i, s) (ε, δ)-differentially private in the preferences of the students. Lemma 5.5.3 shows that the resulting matching is school-dominant so long as the capacity at each school is large enough. These two together imply Theorem 5.1.1 directly.

Lemma 5.5.2. *Private-DA-School(ε, δ) is (ε, δ) -differentially private.*

We relegate the proof of Lemma 5.5.2 to Appendix 5.8 for lack of space. The intuition behind the proof is as follows. Once a student leaves a school, it will never return to that school; thus, the sensitivity of all m counters to a given student is at most $2m$. The proof formalizes the sense in which Private-DA-School(ε, δ) has sensitivity at most $2m$ to a particular student.

Lemma 5.5.3. *With probability at least $1 - \beta$, $\text{Private-DA-School}(\varepsilon, \delta)$ outputs an α -approximately stable, school-dominant matching, as long as the capacity at each school u satisfies $C_u \geq R = O\left(\frac{\sqrt{m}}{\varepsilon\alpha} \text{polylog}\left(n, m, \frac{1}{\delta}, \frac{1}{\beta}\right)\right)$.*

If the maximum error of each of the collection of counters is bounded by x with probability at least, $1 - \beta$, we need only $C_u \geq O(x)$

Proof of Lemma 5.5.3. We prove that the output thresholds $\{t_{u_j}\}$ induce an α -approximately stable, school-dominant matching μ^t .

We claim that there can be no blocking pairs with filled seats in μ^t . Suppose some student a_i wishes to attend u_j . Then, it is either the case that $\text{score}(u_j, a_i) \geq t_{u_j}$ or $\text{score}(u_j, a_i) < t_{u_j}$. In the first case, a_i cannot block with u_j : she could have gone to u_j and chose a school she preferred to u_j . In the second case, consider some $a_{i'}$ such that $u_j = \mu^t(a_{i'})$; this implies $\text{score}(u_j, a_{i'}) \geq t_{u_j}$. Thus, $\text{score}(u_j, a_{i'}) > \text{score}(u_j, a_i)$, so $a_{i'} \succ_{u_j} a_i$, and a_i doesn't block with $a_{i'}$, so there are no blocking pairs with filled seats.

By Theorem 5.3.1 and union bound, we know that the error of all m counters over all time steps is bounded by E except with probability β , where

$$E = \frac{128\sqrt{m \ln(1/\delta)}}{\varepsilon} \ln\left(\frac{2m}{\beta}\right) \left(\sqrt{\log(nmJ)}\right)^5.$$

So, we condition on the event that all schools' counters are accurate within E throughout the run of $\text{Private-DA-School}(\varepsilon, \delta)$ for the remainder of our argument.

We first claim that no school is over-enrolled in μ^t . Consider the last time u_j lowered its threshold to t_{u_j} . Let n_{u_j} denote the number of students tentatively matched to u_j just prior to the final lowering of t_{u_j} . By definition, u_j only lowers its threshold when $\mathbf{Counter}(u_j) < \widehat{C}_{u_j} = C_{u_j} - E$, so

$$C_{u_j} - E = \widehat{C}_{u_j} > \mathbf{Counter}(u_j) \geq n_{u_j} - E \geq |\mu^t(u_j)| - E - 1$$

where the first equality is by definition, the first inequality comes from the fact that u_j lowered its threshold, the third from the accuracy we've conditioned on from the counters, and the final from the fact that u_j never again lowers its threshold. Thus, $C_{u_j} \geq |\mu^t(u_j)|$, and u_j is not over-enrolled.

Now, we show no school is under-enrolled by more than $2E$, unless $t_{u_j} = 0$. When the algorithm terminates, each school u_j either has a threshold $t_{u_j} = 0$ or

$$|\mu^t(u_j)| + E \geq \mathbf{Counter}(u_j) \geq \widehat{C}_{u_j} = C_{u_j} - E$$

where the first equality comes from the conditional bound on the error of the counters, the second from the fact that the algorithm terminated, and the final one from the definition of \widehat{C}_{u_j} . Thus, $|\mu(u_j)| \geq C_{u_j} - 2E$ whenever $t_{u_j} > 0$, so no school is under-enrolled by more than an α -fraction of its seats so long as $C_{u_j} \geq 2E/\alpha$.

Finally, we show school dominance. We will now show that μ , the matching corresponding to the thresholds output by $\text{Private-DA-School}(\varepsilon, \delta)$, is also achieved by running DA-School on the same instance, and halting early. No school is over-enrolled, by our argument above, at

any point during the run of the algorithm. So, each proposal made by u_j would be a valid proposal to make in DA-School with full capacity. Thus, Private-DA-School(t, e) mimics with each school having made (weakly) fewer proposals than it would have in DA-School. Since each school makes its proposals in the same order (according to \succ_{u_j}), this implies that μ^t is a matching that corresponds to some intermediate point in DA-School using with the same ordering of proposals. Thus, by Lemma 5.5.1, μ is school-dominant (and our argument is entirely parametric in E , so the second part of the claim follows directly). ■

Now, we present the formal version Theorem 5.1.2. Without loss of generality, we can assume algorithm ignores students after they have accepted k or more schools' proposals.

Theorem 5.1.2. *Suppose each student has a preference list of length at most k . Then, School-Proposing($\frac{\varepsilon\sqrt{m}}{2\sqrt{k}}, \delta$) is (ε, δ) -private and thus $\varepsilon + \delta$ -approximately truthful. With probability at least $1 - \beta$, it outputs an α -approximately stable school-dominant matching, as long as the capacity at each school $C_u \geq O\left(\frac{\sqrt{k}}{\alpha\varepsilon} \text{polylog}\left(n, m, \frac{1}{\delta}, \frac{1}{\beta}\right)\right)$.*

We relegate the formal proof of Theorem 5.1.2 to Section 5.8; the intuition is simple enough after the proof of Theorem 5.1.1. When students have much shorter preference lists, this greatly reduces the sensitivity of the counters to a single student's responses (from $\Theta(m)$ to $\Theta(k)$, where k is the maximum length of the preference list). This allows us to more tightly concentrate the error from the counters while maintaining differential privacy.

5.6 Conclusions

In this work, we applied differential privacy as a tool to design a many-to-one stable matching algorithm with strong incentive guarantees for the student side of the market. To the best of our knowledge, our work is the first work to show positive truthfulness results for the non-optimal side of the market, under *worst-case* preferences, for *all* participants on the non-optimal side of the market.

Additionally, although we have not focused on this, our algorithm also provides strong *privacy* guarantees to the students. Each student, upon learning the school thresholds (and hence the school that she herself is matched to) can learn almost nothing about either the preferences or scores of the other students (i.e. almost nothing about the preferences that the other students hold over schools, or the preferences that schools hold over the other students). Here “almost nothing” is the formal guarantee of differential privacy, which in particular implies that for every student a , no matter what her prior belief over the private data of some other student a' is, her posterior belief over a' 's data would be almost the same in the two worlds in which a' participates in the mechanism, and in which she does not. These guarantees might themselves be valuable in settings in which the matching being computed is sensitive – e.g. when computing a matching between patients and drug trials, for example.

5.7 Privacy Analysis for Counters

Theorem 5.7.1 (Chan et al. [34]). *For $\beta > 0$, $\mathbf{Counter}(\varepsilon, T)$ is ε -differentially private with respect to a single bit change in the stream, and (α, β) -useful for*

$$\alpha = \frac{4\sqrt{2}}{\varepsilon} \ln\left(\frac{2}{\beta}\right) \left(\sqrt{\log(T)}\right)^5.$$

Chan et al. [34] show that $\mathbf{Counter}(\varepsilon, T)$ is ε -differentially private with respect to single changes in the input stream, when the stream is generated non-adaptively. For our application, we require privacy to hold for a large number of streams whose joint-sensitivity can nevertheless be bounded, and whose entries can be chosen adaptively. To show that **Counter** is also private in this setting (when ε is set appropriately), we first present a slightly more refined composition theorem.

5.7.1 Composition

An important property of differential privacy is that it degrades gracefully when private mechanisms are composed together, even adaptively. We recall the definition of an adaptive composition experiment due to Dwork et al. [51].

Definition 5.7.1 (Adaptive composition experiment).

- Fix a bit $b \in \{0, 1\}$ and a class of mechanisms \mathcal{M} .
- For $t = 1 \dots T$:
 - The adversary selects two databases $D^{t,0}, D^{t,1}$ and a mechanism $\mathcal{M}_t \in \mathcal{M}$.
 - The adversary receives $y_t = \mathcal{M}_t(D^{t,b})$

The “output” of an adaptive composition experiment is the view of the adversary over the course of the experiment. The experiment is said to be ε -differentially private if

$$\max_{S \subseteq \mathcal{R}} \frac{\Pr[V^0 \in S]}{\Pr[V^1 \in S]} \leq \mathbb{E}(\varepsilon),$$

and (ε, δ) -differentially private if

$$\max_{S \subseteq \mathcal{R}, \Pr[V^0 \in S] \geq \delta} \frac{\Pr[V^0 \in S] - \delta}{\Pr[V^1 \in S]} \leq \mathbb{E}(\varepsilon),$$

where V^0 is the view of the adversary with $b = 0$, V^1 is the view of the adversary with $b = 1$, and \mathcal{R} is the range of outputs.

Any algorithm that can be described as an instance of this adaptive composition experiment (for an appropriately defined adversary) is said to be an instance of the class of mechanisms \mathcal{M} under *adaptive T -fold composition*.

A very useful tool to analyze private algorithms is the following theorem that allows us to analyze the “composition” of private algorithms.

Theorem 5.7.1 (Adaptive Composition Dwork et al. [51]). *Let $\mathcal{A}: \mathcal{U} \rightarrow \mathcal{R}^T$ be a T -fold adaptive composition³ of (ε, δ) -differentially private algorithms. Then \mathcal{A} satisfies $(\varepsilon', T\delta + \delta')$ -differential privacy for*

$$\varepsilon' = \varepsilon \sqrt{2T \ln(1/\delta')} + T\varepsilon(e^\varepsilon - 1).$$

In particular, for any $\varepsilon \leq 1$, if \mathcal{A} is a T -fold adaptive composition of $(\varepsilon/\sqrt{8T \ln(1/\delta)}, 0)$ -differentially private mechanisms, then \mathcal{A} satisfies (ε, δ) -differential privacy.

For a more refined analysis in our setting, we now state a straightforward consequence of a composition theorem of Dwork et al. [51].

Lemma 5.7.2 (Dwork et al. [51]). *Let $\Delta \geq 0$. Under adaptive composition, the class of $\frac{\varepsilon}{\Delta}$ -private mechanisms satisfies ε -differential privacy and the class of $\frac{\varepsilon}{2c\sqrt{2\Delta \ln(1/\delta)}}$ -private mechanisms satisfies (ε, δ) -differential privacy, if the adversary always selects databases satisfying*

$$\text{for all } t \quad |D^{t,0} - D^{t,1}| \leq c, \text{ and also } \sum_{t=1}^T |D^{t,0} - D^{t,1}| \leq \Delta.$$

In other words, the privacy parameter of each mechanism should be calibrated for the total distance between the databases, over the whole composition. This is useful for analyzing the privacy of the counters in our algorithm, which collectively have bounded sensitivity.

5.7.2 Details for Counters

We reproduce Binary mechanism here in order to refer to its internal workings in our privacy proof.

First, it is worth explaining the intuition of the **Counter**. Given a bit stream $\sigma: [T] \rightarrow \{-1, 0, 1\}$, the algorithm releases the counts $\sum_{i=1}^t \sigma(i)$ for each t by maintaining a set of partial sums $\sum[i, j] : \sum_{t=i}^j \sigma(t)$. More precisely, each partial sum has the form $\sum[2^i + 1, 2^i + 2^{i-1}]$, corresponding to powers of 2.

In this way, we can calculate the count $\sum_{i=1}^t \sigma(i)$ by summing at most $\log t$ partial sums: let $i_1 < i_2 \dots < i_m$ be the indices of non-zero bits in the binary representation of t , so that

$$\sum_{i=1}^t \sigma(i) = \sum[1, 2^{i_m}] + \sum[2^{i_m} + 1, 2^{i_m} + 2^{i_m-1}] + \dots + \sum[t - 2^{i_1} + 1, t].$$

Therefore, we can view the algorithm as releasing partial sums of different ranges at each time step t and computing the counts is simply a post-processing of the partial sums. The core algorithm is presented in Algorithm 2.

5.7.3 Counter Privacy under Adaptive Composition

Theorem 5.3.1. *Suppose we have m bit streams such that the change of an agent's data affects at most k streams, and alters at most c bits in each stream. For any $\beta > 0$, the composition of m*

³ See Section 5.7.1 and [51] for further discussion.

Algorithm 2: Counter(ε, T)

Input: A stream $\sigma \in \{-1, 1\}^T$
Output: $B(t)$ as estimate for $\sum_{i=1}^t \sigma(i)$ for each time $t \in [T]$
for all $t \in [T]$ **do**
 Express $t = \sum_{j=0}^{\log t} 2^j \text{Bin}_j(t)$.
 Let $i \leftarrow \min_j \{\text{Bin}_j(t) \neq 0\}$
 $a_i \leftarrow \sum_{j < i} a_j + \sigma(t)$, ($a_i = \sum [t - 2^i + 1, t]$)
 for $0 \leq j \leq i - 1$ **do**
 Let $a_j \leftarrow 0$ and $\hat{a}_j \leftarrow 0$
 end for
 Let $\hat{a}_j = a_j + \text{Lap}(\log(T)/\varepsilon)$
 Let $B(t) = \sum_{i: \text{Bin}_i(t) \neq 0} \hat{a}_i$
end for

distinct Counter $\left(\varepsilon/2c\sqrt{2kc\ln(1/\delta)}, T\right)$ is (ε, δ) -differentially private, and (α, β) -useful for

$$\alpha = \frac{16c\sqrt{kc\ln(1/\delta)}}{\varepsilon} \ln\left(\frac{2m}{\beta}\right) \left(\sqrt{\log(T)}\right)^5.$$

Proof. The composition of m counters is essentially releasing a collection of noisy partial sums adaptively. We need to first frame this setting as an advanced composition experiment defined in Algorithm 5.7.1. First, we treat each segment $\sigma[a, b]$ in a stream as a database. For each such database, we are releasing the sum by adding noise sampled from the Laplace distribution:

$$\text{Lap}\left(\frac{2c\sqrt{2kc\ln(1/\delta)}\log(T)}{\varepsilon}\right),$$

which is $\frac{\varepsilon}{2c\sqrt{2kc\ln(1/\delta)}\log(T)}$ -private mechanism (w.r.t. a single bit change). We know that changing an agent's data changes at most c bits in each stream, and affects at most k streams, and also each bit change can result in $\log(T)$ bits changes across different stream-segment databases. Therefore, we can bound the total distance between all pairs stream-segment databases by

$$\Delta \leq kc\log(T).$$

By Lemma 5.7.2, we know that the composition of all m counters under our condition satisfies (ε, δ) -differential privacy.

Plugging in our choice of ε to the accuracy proof for **Counter** in Chan et al. [34], we obtain our accuracy guarantee by applying union bound. \blacksquare

5.8 Omitted Proofs

Proof of Lemma 5.5.2. `Private-DA-School`(ε, δ) outputs a sequence of sets of thresholds and nothing else. We will construct a mechanism \mathcal{M} , which will output the same sequence of thresholds as `Private-DA-School`(ε, δ), for which it is more obvious to prove (ε, δ)-differential privacy. This will imply (ε, δ)-differential privacy of `Private-DA-School`(ε, δ). Here is the definition of \mathcal{M} :

<p>Algorithm 3: \mathcal{M}</p> <p>Publish threshold $t_{u_j} = J$ for each school u_j; $\varepsilon' = \frac{\varepsilon}{12\sqrt{2m \ln \frac{1}{\delta}}}$; Initialize counter($u_j$) = Counter($\varepsilon', nmJ$); Let $\widehat{C}_{u_j} = C_{u_j} - E$; while <i>there is some under-enrolled school u_j: counter(u_j) $\leq \widehat{C}_{u_j}$ and $t_{u_j} > 0$</i> do Let $t'_{u_j} = t_{u_j} - 1$; Publish thresholds $(t_{u_1}, \dots, t'_{u_j}, \dots, t_{u_m})$; Receive bits $b_{u'_j} \in \{-1, 0, 1\}$ for each u'_j; Send b_{u_j} to counter(u_j);</p>

We define the input bits to the algorithm \mathcal{M} as follows. For a fixed execution of the while loop, we will define the bits $b_{u_{j'}}$ to give to \mathcal{M} . Let u_j be the school which lowered its threshold in this timestep. Let $b_{u_j} = 1$ if and only if, for the unique student a_i such that $\text{score}(u_j, a_i) = t'_{u_j}$, it is true that $u_j = \text{argmax}_{u \in \mathcal{U}} \{u \mid \text{score}(u, a_i) \geq t_{u_j}\}$ (a_i prefers u_j to all other schools for which her score surpasses the threshold). Let $b_{u_{j'}} = -1$ if and only if $b_{u_j} = 1$ and also $u_{j'} = \text{argsecondmax}_{u \in \mathcal{U}} \{u \mid \text{score}(u, a_i) \geq t_{u_j}\}$ (u_j is a_i 's favorite available school and $u_{j'}$ is her second favorite). For all other j'' , let $b_{u_{j''}} = 0$.

Then, there are at most $2m$ nonzero bits sent to \mathcal{M} about a particular student a_i , and at most 2 nonzero bits sent by a particular a_i to any school u_j . These bits are the only interface \mathcal{M} has with private data. Furthermore, \mathcal{M} and `Private-DA-School`(ε, δ) have the same distribution over output data. So it suffices to show that \mathcal{M} is (ε, δ)-differentially private.

Let $f : \{J\}^m \times [n] \rightarrow \{J\}^m$ be the function that, as a function of the previous thresholds and counter values, outputs the new set of thresholds at each time t . Then, the thresholds published by \mathcal{M} are a composition of f , m instantiations of **Counter**(ε', nmJ), and previously computed data. Thus, it suffices to show the composition of the m counters satisfy (ε, δ)-differential privacy. By construction, each school u_j receives at most 2 nonzero bits from a given student, and no student's data creates more than $2m$ nonzero bits in all streams together. By Theorem 5.7.1 and Lemma 5.7.2, the composition of m **Counter**(ε', nmJ) satisfy (ε, δ)-differential privacy when no stream has more than 2 bits affected by a single agent's data and no student has more than $2m$ total nonzero bits in any stream. Thus, \mathcal{M} (and also `Private-DA-School`(ε, δ)) satisfies (ε, δ) differential privacy. ■

Proof of Theorem 5.1.2. We prove (ε, δ)-privacy, which again reduces to proving (ε, δ)-differential

privacy of the set of m counters. By a simple calculation, School-Proposing($\frac{\varepsilon\sqrt{m}}{2\sqrt{k}}, \delta$) uses

$$\varepsilon' = \frac{\varepsilon}{4\sqrt{k \ln \frac{1}{\delta}}}$$

as the privacy parameter for the m counters it uses. Theorem 5.7.1 states that a collection of m Counter(ε', nT) with total sensitivity Δ (and individual sensitivity c) satisfy (ε, δ) -differential privacy so long as

$$\varepsilon' \leq \frac{\varepsilon}{2c\sqrt{\Delta \ln \frac{1}{\delta}}}.$$

Remark 5.1.2 limits the total amount of sensitivity School-Propose will have; a student will be able to affect at most $2k$ bits in the input stream, so $\Delta \leq 2k$, and at most 2 per school, so $c \leq 2$. Thus, it suffices to use privacy parameter

$$\varepsilon' \leq \frac{\varepsilon}{4\sqrt{k \ln \frac{1}{\delta}}},$$

so our algorithm is (ε, δ) -differentially private. Furthermore, Theorem 5.7.1 and a union bound imply the maximum error any one of the counters will have at any time during the execution of the algorithm is

$$E \leq \frac{128\sqrt{k \ln(1/\delta)}}{\varepsilon} \ln\left(\frac{2m}{\beta}\right) \left(\sqrt{\log(nmJ)}\right)^5.$$

with probability $1-\beta$. Thus, by Lemma 5.5.3, we get the desired guarantee for α -approximate stability and school-dominance. ■

5.9 Proofs of Matching Lemmas

We state one more lemma which we will use in the proofs of Lemmas 5.5.1.

Lemma 5.9.1. *Consider a set P_u of proposals made by each school u according to some prefix of school-proposing DA. Let $P \subseteq A \times U$ be the set of proposals made by all schools. Then, the matching μ which results from P is unique (and independent of the order in which proposals are made), assuming students are truthful.*

Proof. Each student ultimately accepts her most preferred proposal among the set of proposals she has received, independent of their ordering. (i.e. admissions thresholds only descend, and she picks her most preferred school amongst those schools with thresholds below her scores). Thus, each school u will be matched to the subset of P_u which finds u to be their favorite offer, independent of the order in which proposals were made. ■

Lemma 5.9.2. *Let μ_t be some matching which is an intermediate matching in a run of the school-proposing deferred acceptance algorithm. Then μ_t is school-dominant.*

Proof. The school-proposing deferred acceptance algorithm is somewhat underspecified. In particular, if multiple schools have space remaining, the *order* in which those schools make proposals isn't predetermined. But, by Lemma 5.9.1 shows that reordering of the same proposals from the schools will arrive at the same matching. Thus, it suffices to show, for a fixed ordering of the entire set of proposals made by DA, that each intermediate matching is school-dominant.

Let t denote the time at which we wish to halt a run of DA. Let $P_{u,t}$ denote the set of proposals which school u has made according some fixed ordering up to time t , and P_u denote set of proposals made by school u according to the entire run of DA. Let μ_t denote the “current” matching according to the first run of DA stopped at time t and μ denote the final outcome of DA.

Consider any school A . Notice that, since $|P_u| \geq |P_{u,t}|$, by the definition of DA,

$$P_{u,t} \subseteq P_u \tag{5.1}$$

since, for a given school, the proposing order is just working down their preference list.

Now consider a particular school u . We must show that for each $a \in \mu_t(u) \setminus \mu(u)$, $a' \in \mu(u) \setminus \mu_t(u)$, $a \succ_u a'$. If a was proposed to by u in P_t and rejects u , then u will be rejected by a when she receives a superset P of proposals. Thus, the only students u has according to μ but not μ_t are students $z \in P_u \setminus P_{u,t}$ (students who are proposed to after time t). But, by the definition of DA, if u proposes to two students a and a' , and proposes to a before a' , $a \succ_u a'$, as desired. ■

5.10 A Reference to DA–School

In this section, we present DA–School, the well-known school-proposing deferred acceptance algorithm [62]. In this setting, schools which are not at capacity *propose* to students one at a time, starting from their favorite students and moving down their preference list. When a student gets a proposal, if she is tentatively matched to some other school, she will reject the offer from whichever school she likes less and accept the offer from the school she likes better. At this point, she is tentatively matched to the school she likes better, and the other school will continue to make proposals to fill the seat offered to her. The version of the algorithm we present here is non-standard – it operates by having each school set an admissions threshold, which it decreases slowly – but is easily seen to be equivalent to the deferred acceptance algorithm. This version of the algorithm will be much more amenable to a private implementation, which we give next. When a school u lowers its threshold t_u below the score of a student a at school u ($\text{score}(u, a)$), we say that school u has *proposed* to student a .

It is well-known that DA–School will output a school-optimal stable matching (in our notation, a 0-approximate school-dominant stable matching) [117], assuming all players are truthful.

5.11 Private Matching Algorithms Must Allow Empty Seats

In this chapter, we gave an algorithm with strong worst-case incentive properties in large markets, without needing to make distributional assumptions about the agents preferences, or requiring

Algorithm 4: DA-SCHOOL, the deferred acceptance algorithm with schools proposing

Input: school capacities $\{C_u\}$, student preferences $\{\succ_a\}$ and scores $\{\text{score}(u, a)\}$, range of scores $[0, J]$

Output: a set of score thresholds $\{t_j\}$

initialize: for each school u_j and each student a_i
 $\text{counter}(u_j) = 0$ $t_j = J$, $\mu(a_i) = \emptyset$

while there is some under-enrolled school u_j : $\text{counter}(u_j) \leq \widehat{C}_{u_j}$ and $t_{u_j} > 0$ **do**
 $t_{u_j} = t_{u_j} - 1$
 for all student a_i **do**
 if $\mu(a_i) \neq \text{argmax}_{\succ_{a_i}} \{u_j \mid \text{score}(u_j, a_i) \geq t_{u_j}\}$ **then**
 $\text{counter}(\mu(a_i)) = \text{counter}(\mu(a_i)) - 1$;
 let $\mu(a_i) = \text{argmax}_{\succ_{a_i}} \{u_j \mid \text{score}(u_j, a_i) \geq t_{u_j}\}$
 $\text{counter}(\mu(a_i)) = \text{counter}(\mu(a_i)) + 1$;
 end if
 end for
end while
return Final threshold scores $\{t_j\}$

any other “large market” condition other than that the capacities of the schools be sufficiently large. However, in exchange, we had to relax our notion of stability to an approximate notion which allows a small number of empty seats per school. We here give an example demonstrating why this relaxation is necessary for any differentially private matching algorithm. An algorithm that must return an -exactly- stable matching must have extremely high sensitivity to the change in preferences of any single agent, if preferences are allowed to be worst case.

Example 5.11.1. Suppose there are n students and 2 schools, H and Y . Suppose, for students $1 \leq a \leq \frac{n}{2}$, $H \succ_a Y$, and for $\frac{n}{2} < a \leq n$, $Y \succ_a H$. Each school has capacity for exactly half of the students: $C_H = C_Y = \frac{n}{2}$. Suppose Y has preference ordering $\succ_Y, s_1 \succ_Y s_2 \succ_Y \dots \succ_Y s_n$; H has preference ordering $s_{\frac{n}{2}+1} \succ_H s_{\frac{n}{2}+2} \succ_H \dots s_n \succ_H s_1 \succ_H \dots \succ_H s_{\frac{n}{2}}$. The school-optimal matching matches students $s_1, \dots, s_{\frac{n}{2}}$ to Y and $s_{\frac{n}{2}+1}, \dots, s_n$ to H . Now consider the market with any single student removed. The school-optimal stable matching changes entirely (i.e. every single student is matched to a different school). For example, if s_1 is removed, Y will admit $s_{\frac{n}{2}+1}$ (who will accept), H will admit s_2 (who will accept), Y will admit $s_{\frac{n}{2}+2}$ and so on. In the end, each student will get her favorite school, and the schools will swap students. The same effect is achieved by having a single student *change* her preferences, by reporting that she prefers to be unmatched than to be matched to her second choice school. This example shows that the exact school-optimal matching is *highly* sensitive to the addition, removal, or alteration of preferences of a single student and hence impossible to achieve under differential privacy. Our algorithms blunt this kind of sensitivity via the use of a small budget of seats that we may leave empty.

Chapter 6

Privacy-Preserving Public Information in Coordination Games

6.1 Introduction

In this chapter, we consider the problem of designing a mechanism to coordinate sequential decisions of strategic agents in a privacy-preserving manner. When agents have very little information about the decisions made by other agents, their behavior may lead very poor outcomes. For example, in examining causes of the recent financial crisis and subsequent recession, the Financial Crisis Inquiry Commission [57, p. 352] concluded that

“The OTC derivatives market’s lack of transparency and of effective price discovery exacerbated the collateral disputes of AIG and Goldman Sachs and similar disputes between other derivatives counterparties.”

Even though regulators have access to detailed confidential information about financial institutions and (indirectly) individuals, current statistics and indices are based only on public data, since disclosures based on confidential information are restricted. However, forecasts based on confidential data can be much more accurate¹, prompting regulators to ask whether aggregate statistics can be economically useful while also providing rigorous privacy guarantees [58].

In this work, we show that such *privacy-preserving public information*, in an interesting class of sequential decision-making games, can achieve (nearly) the best of both worlds. In particular, the goal is to produce information about actions taken by previous agents that can be posted publicly, preserves all agents’ (differential) privacy, and can significantly improve worst-case social-welfare. While our models do not directly speak to the highly complex issues involved in real-world financial decision-making, they do indicate that in settings involving contention for resources and first-mover advantages, privacy-preserving public information can be a significant help in improving social welfare. In the following sections, we describe the game setting and the information model.

¹For example, Oet et al. [105] compared an index based on both public and confidential data with an analogous index based only on publicly available data. The former index would have been a significantly more accurate predictor of financial stress during the recent financial crisis (see Oet et al. [104, Figure 4]). See Flood et al. [58] for further discussion.

6.1.1 Game Model

Consider a setting in which there are m resources and n players. The players arrive online, in an *adversarial* order, one at a time. For ease of exposition, we rename players such that player i is the i th to arrive. Each agent i has a set of available actions A_i and chooses an action $a_i \in A_i$. Each action a_i represents a portfolio over the m resources, e.g. $a_i = (a_{i,1}, \dots, a_{i,m})$, and $a_{i,r}$ represents an amount of investment in resource r of the action a_i . We will assume throughout that $\sum_{r \in [m]} a_{i,r} = 1$ for all i, a_i . For ease of exposition, for the first section of this paper, we will consider the unit-demand case: where each $a_{i,r} \in \{0, 1\}$ (i.e., each agent will choose an action which represents selecting a single resource). We study the continuous version where $a_{i,r}$'s can be fractional, but still sum to 1, in Section 6.4. Furthermore, we do not make the assumption that players have knowledge of their position in the sequence, that is, a player need not know how many players have acted before her.

Each resource r has some non-increasing function $\mathcal{W}_r : \mathbb{Z}^+ \rightarrow \mathbb{R}^+$ indicating the *value*² of this resource to the k th player who chooses it. Therefore, the total utility of player i choosing action a_i is $u_i(a_i, a_{1,\dots,i-1}) = \sum_r a_{i,r} \mathcal{W}_r(x_{i,r})$, where $x_{i,r} = \sum_{j=1}^{i-1} a_{j,r}$ is the usage of r before agent i for each r . So, if agent i chooses an action a_i such that $a_{i,r} = 1$ and k prior agents have chosen r , agent i gets utility $\mathcal{W}_r(k)$.

In this **resource sharing** setting, the utility for a player of choosing a certain resource is a function of the resource and the number of players who have invested in the resource before her (and, importantly, not after her)³.

Illustrative Example For each resource, suppose $\mathcal{W}_r(k) = \mathcal{W}_r(0)/k$, where $\mathcal{W}_r(0)$ is the initial value of resource r . The value of each resource r drops rapidly as a function of the number of players who have chosen it so far. If each player i has *perfect information* about the investment choices made by the players before her, the optimal action for player i is to greedily select the action in A_i of highest utility based on the number of players who have selected each resource so far⁴. As shown in Section 6.3, the resulting social welfare of this behavior is within a factor of 2 of the optimal⁵. In the case where each player has *no information* about other players' behaviors, some particularly disastrous sequences of actions might reasonably occur, leading to very low social welfare. For example, if each player i has access to a common resource r where $\mathcal{W}_r(0) = 1$ and a personal resource r_i where $\mathcal{W}_{r_i}(0) = 1 - \varepsilon$, each might reasonably choose greedily according to $\mathcal{W}(0)$, selecting the resource of highest initial value (in this case, r). This would give social welfare of $\ln(n)$, whereas the optimal assignment would give $n(1 - \varepsilon)$. Without information about the game state, therefore, the players may achieve only a $O\left(\frac{\ln(n)}{n}\right)$ fraction of

²In the event that we evaluate \mathcal{W}_r on some $x \in \mathbb{R}$, if the resource choices are discrete, we will extend the weight functions as $\mathcal{W}_r(x) = \mathcal{W}_r(\max\{\lfloor x \rfloor, 0\})$; if the resource choices are continuous the extension assumed is $\mathcal{W}_r(x) = \mathcal{W}_r(\max\{x, 0\})$.

³In Appendix B.1, we consider a generalization where the utility to a player of investing in a particular resource is a function of the total number of players who have chosen that resource, including those who have invested after her.

⁴And, in general, that will be true for any resource-sharing setting where agents have perfect information about the state of play.

⁵And 4 in the non-unit demand case.

the possible welfare.

6.1.2 Information Model

In resource sharing games, players' decisions about their actions will be best when they know how many players have chosen each resource when they arrive. The mechanisms we consider, therefore, will publicly announce some estimate of these counts. We call this an *announcement mechanism*. We consider the trade-off between the privacy lost by publishing these estimates and the accuracy of the counters in terms of social welfare. We consider three categories of counters for publicly posting the estimate of resource usage: perfect, private and empty counters.

- **Perfect Counters:** At all points, the counters display the exact usage of each resource.
- **Privacy-preserving public counters:** At all points, the counters display an approximate usage of the resources while maintaining privacy for each player. We define the privacy guarantee in Section 6.2.
- **Empty Counters:** At all points, every counter displays the value 0.

6.1.3 Players' Behavior

Each player is a utility-maximizing agent and will choose the resource that, given their beliefs about actions taken by previous players and the publicly displayed counters, gives them maximum value. We analyze the game play under two classes of strategies – greedy and undominated strategies.

1. **Greedy strategy:** Under the greedy strategy, a player has no outside belief about the actions of previous players and chooses the resource that maximizes her utility given the *currently displayed (or announced) values of the counters*. Greedy is a natural choice of strategy to consider since it is the utility-maximizing strategy when the usage counts posted are *perfect*.
2. **Undominated Strategy(UD):** Under undominated strategies, we allow players to have any beliefs about the actions of the previous players that are consistent with the displayed value of the counters⁶, and they are allowed to play *any* undominated strategy a_i according to this belief. A strategy a_i is *undominated* according to a belief, if no other a'_i gets strictly higher utility.⁷

We analyze the social welfare $SW(a) = \sum_i u_i(a)$ generated by an announcement mechanism \mathcal{M} for a set of strategies D and compare it to the optimal social welfare OPT. For a game

⁶As will become clear in Section 6.2, we work with privacy-preserving public counters that display values that can be off from the true usage only in a *bounded* range. Hence with these counters, a player's belief is consistent as long as the belief implies the usage of the resource to be a number that is within the bounded range of the displayed value. Moreover, with empty counters, any belief about the actions of previous players is a consistent belief.

⁷For each counter mechanism we consider, there exists at least one undominated strategy. For example, with perfect counters, the only consistent belief is that the true value is equal to the displayed value and here the greedy strategy is always undominated; moreover, if the counter mechanism has a nonzero probability of outputting the true value, then again the greedy strategy is undominated under the belief that the displayed value is the true value; if the counter mechanism can display values that are arbitrarily off from the true value, then for equal initial values every strategy is undominated.

setting g , constituted of a collection of players $[n]$ and their allowable actions A_i (as defined in Section 6.1.1), $\text{OPT}(g)$ is defined as the optimal social welfare that can be achieved by any allocation of resources to the players, where the space of feasible allocations is determined by the setting g . In the unit-demand setting, $\text{OPT}(g)$ is the maximum weight matching in the bipartite graph $G = (U \cup V, E)$ where U is the set of the n players, V has n vertices for each resource r , one of value $\mathcal{W}_r(k)$ for each $k \in [n]$, and there is an edge between player i and all vertices corresponding to resource r if and only if $r \in A_i$ (Note that the weights are on the vertices in V). The object of our study is $\text{CR}_D(g, \mathcal{M})$, the *worst case competitive ratio* of the optimal social welfare to the welfare achieved under strategy D and counter mechanism \mathcal{M} . As mentioned earlier, D will either be the greedy (GREEDY) or the undominated (UNDOM) strategy, and \mathcal{M} will be either the perfect (\mathcal{M}_{Full}), the privacy-preserving or the empty (\mathcal{M}_\emptyset) counter. When \mathcal{M} uses internal random coins, our results will either be worst-case over all possible throws of the random coins, or will indicate the probability with which the social welfare guarantee holds.

6.1.4 Statement of Main Results

For sequential resource-sharing games, we prove that there exists privacy-preserving counters such that, for all nonincreasing value curves, the greedy strategy following privacy-preserving counters has a competitive ratio *polylogarithmic* in the number of players (Theorem 6.3.6). This should be contrasted with the competitive ratio of 4 achieved by greedy w.r.t. perfect counters (Theorem 6.3.1) and the nearly-linear (in the number of players) competitive ratio of greedy with empty counters (as shown in the illustrative example in Section 6.1.1). Thus, privacy preserving counters perform nearly as well as perfect counters and much better than empty counters.

For the case of undominated strategies, when the marginal values of resources drop slowly, (for example, at a polynomial rate, $\mathcal{W}_r(k) = \mathcal{W}_r(0) / k^p$ for constant $p > 0$), we also give a polylogarithmic bound on the competitive ratio (w.r.t. privacy-preserving counters) (Theorem 6.3.8). With empty counters, the competitive ratio for undominated strategies is unbounded (Theorem 6.3.2) for arbitrary curves and is at least quadratic (in the number of players) if the value curve drops slowly (Theorem 6.3.3). We note here that for many of our positive results for privacy preserving counters state the competitive ratio in terms of parameters of the counter vector α and β (as detailed in Section 6.2) and for a particular implementation of the counter vectors, the values of α and β are mentioned in Section 6.6.

The key privacy tool we use is the differentially private counter under continual observation [50], which we use to publish estimates of the usage of each resource. We improve upon the existing error guarantees of differentially private counters and design a new differentially private counter in Section 6.6. The new counter provides a tighter additive guarantee at the price of introducing a constant multiplicative error.

In Section 6.4, we show these ideas can be extended to continuous actions and in Section 6.5 we show that under certain assumptions, to games where agents' utility depends upon future agents' decisions.

6.1.5 Related Work

A great deal of work has been done at the intersection of mechanism design and privacy; Pai and Roth [108] have an extensive survey. Our work is similar to much of the previous work in that it considers designing mechanisms subject to the constraint of maintaining differential privacy. The focus of our work however is on *how useful information can be provided to players in games of imperfect information* to help achieve a good social objective while respecting the privacy constraint of the players. The work of Kearns et al. [84] is close in spirit to ours. Kearns et al. [84] consider games where players have incomplete information about other players' types and behaviors. They construct a privacy-preserving mechanism which collects information from players, computes an approximate correlated equilibria, and then advises players to play according to this equilibrium. The mechanism is approximately incentive compatible for the players to participate in the mechanism and to follow its suggestions. Several later papers [76, 112] privately compute approximate equilibria in different settings. Our main privacy primitive is the differentially private counters under continual observation [34, 50], also used in much of the related work on private equilibrium computation.

As mentioned in Section 6.1.3, one class of player behavior for which we analyze the games is *greedy*. Our analysis of greedy behavior is in part inspired by the work of Balcan et al. [14], who study best response dynamics with respect to noisy cost functions for potential games. An important distinction between their setting and ours is that the noisy estimates we consider are *estimates of state, not value*, and may for natural value curves be quite far from correct in terms of the *values* of the actions.

6.2 Privacy-preserving public counters

We design announcement mechanisms \mathcal{M}_i which give approximate information about actions made by the previous players to player i . Let Δ_m denote the m -dimensional simplex $\Delta_m = \{a \in [0, 1]^m \mid \|a\|_1 \leq 1\}$. Our collection of mechanisms

$$\mathcal{M}_i : (\Delta_m)^{i-1} \times R \rightarrow [0, n]^m$$

depend upon the actions taken before i (specifically, the usage of each resource by each player), and on internal random coins R . When player i arrives, $m_i(a_1, \dots, a_{i-1}) \sim \mathcal{M}_i(a_1, \dots, a_{i-1})$ is publicly announced. Player i plays according to some strategy $d_i : \Delta_m \rightarrow A_i$, that is $a_i = d_i(m_1, \dots, m_i(a_1, \dots, a_{i-1}))$, a random variable which is a function of this announcement. When it is clear from context, we denote $m_i(a_1, \dots, a_{i-1})$ by m_i . Formally, the counters used in this work satisfy the following notion of privacy.

Definition 6.2.1. An announcement mechanism \mathcal{M} is (ε, δ) -differentially private under adaptive⁸ continual observation in the strategies of players if, for each d , for each player i , each pair of strategies d_i, d'_i , and every $S \subseteq (\Delta_m)^n$:

$$\mathbb{P}[(m_1, \dots, m_n) \in S] \leq e^\varepsilon \mathbb{P}[(m_1, \dots, m_i, m'_{i+1}, \dots, m'_n) \in S] + \delta,$$

⁸Adaptivity is needed in this case because the announcements are arguments to the actions of players: when a particular action changes, this modifies the distribution over the future announcements, which in turn changes the distribution over future selected actions.

where the actions of players up to i are defined as $a_j = d_j(m_1, \dots, m_j)$, $a'_i = d'_i(m_1, \dots, m_i)$, the signals defined as $m_j \sim \mathcal{M}_j(a_1, \dots, a_{j-1})$, $m'_j \sim \mathcal{M}_j(a_1, \dots, a_{i-1}, a'_i, a'_{i+1}, \dots, a'_{j-1})$, and for all $j > i$, the action for j is defined as $a'_j = d_j(m_1, \dots, m_{i-1}, m_i, m'_{i+1}, \dots, m'_j)$.

This definition requires that two worlds which differ in a single player changing her strategy from d_i to d'_i have statistically close joint distributions over all players' announcements (and thus their joint distributions over actions). Note that the distribution of $j > i$'s announcement can change slightly, causing j 's distribution over actions to change slightly, necessitating the cascaded m'_j, a'_j for $j > i$ in our definition. The mechanisms we use maintain approximate use counters for each resource. The values of the counters are *publicly announced* throughout the game play. We now define the notion of accuracy used to describe these counters.

Definition 6.2.2 ((α, β, γ) -accurate counter vector). A set of counters $y_{i,r}$ is defined to be (α, β, γ) -accurate if with probability at least $1 - \gamma$, at all points of time, the displayed value of every counter $y_{i,r}$ lies in the range $[\frac{x_{i,r}}{\alpha} - \beta, \alpha x_{i,r} + \beta]$ where $x_{i,r}$ is the *true* count for resource r prior to the i th update, and is monotonically increasing in the true count.

We refer to a set of $(\alpha, \beta, 0)$ -accurate counters as (α, β) -counters for brevity. It is possible to achieve $\gamma = 0$ (which is necessary for undominated strategies, which assumes the multiplicative and additive bounds on y are worst-case), taking an appropriate loss in the privacy guarantees for the counter (Proposition 6.6.3). Counters satisfying Definitions 6.2.1 and 6.2.2 with $\alpha = 1$ and $\beta = O(\log^2 n)$ ⁹ were given in Chan et al. [34], Dwork et al. [50]; we give a different implementation in Section 6.6 which gives a tighter bound on $\alpha\beta$ by taking α to be a small constant larger than 1. Furthermore, the counters in Section 6.6 are *monotonic* (i.e., the displayed values can only increase as the game proceeds) and we use monotonicity of the counters in some of our results.

In some settings we require counters we a more specific utility guarantee:

Definition 6.2.3 ((α, β, γ) -accurate underestimator). A set of counters $y_{i,r}$ is defined to be (α, β, γ) -accurate underestimator if with probability at least $1 - \gamma$, at all points of time, the displayed value of every counter $y_{i,r}$ lies in the range $[\frac{x_{i,r}}{\alpha} - \beta, x_{i,r}]$ where $x_{i,r}$ is the *true* count for resource r prior to the i th update.

The following observation states that a counter vector can be converted to an undercounter with small loss in accuracy.

Observation 6.2.1. We can convert a (α, β) -counter to an $(\alpha^2, \frac{2\beta}{\alpha})$ -underestimating counter vector.

Proof. We can simply subtract any possible overcounting: $\frac{1}{\alpha}x - \beta \leq y \leq \alpha x + \beta$ implies $y' = \frac{y - \beta}{\alpha} \leq x$ and $\frac{1}{\alpha^2}x - \frac{2\beta}{\alpha} \leq y'$. ■

6.3 Resource Sharing

In this section, we consider resource sharing games – the utility to a player is completely determined by the resource she chooses and the number of players who have chosen that resource before her. This section considers the case where players' actions are discrete: $a_i \in \{0, 1\}^m$

⁹Ignoring the dependence on ϵ, δ .

for all $i, a_i \in A_i$. We defer the analysis of the case where players' actions are continuous to Section 6.4.

6.3.1 Perfect counters and empty counters

Before delving into our main results, we point out that, with perfect counters, greedy is the only undominated strategy, and the competitive ratio of greedy is a constant.

Theorem 6.3.1. *With perfect counters, greedy behavior is dominant-strategy and all other behavior is dominated for any sequential resource-sharing game g ; and $\text{CR}_{\text{GREEDY}}(\mathcal{M}_{\text{Full}}, g) \leq 4$.*

The proof of Theorem 6.3.1 follows from the connection between resource-sharing and online vertex-weighted matching, which we mention below.

Observation 6.3.1. In the setting where $\|a_i\|_1 = 1$ for all $a_i \in A_i$, for all i , full-information, discrete resource-sharing reduces to online, vertex-weighted bipartite matching.

Proof. Construct the following bipartite graph $G = (U, V, E)$ as an instance of online vertex-weighted matching from an instance of the future-independent resource sharing game. For each resource r , create n vertices in V , one with weight $\mathcal{W}_r(t)$ for each $t \in [n]$. As players arrive online, they will correspond to vertices in $u_i \in U$. For each $a_i \in A_i$ corresponding to a set of resources S , u_i is allowed to take any subset of V with a single copy of each $r \in S$. ■

The proof of the social welfare is quite similar to the one-to-one, online vertex-weighted matching proof in Karp et al. [83], with the necessary extension for many-to-one matchings (losing a factor of $1/2$ in extending to the many-to-one setting), as Theorem 6.3.1 does not assume $\|a_i\| = 1$.

Proof of Theorem 6.3.1. Consider any instance of $G = (U, V, E)$, a vertex-weighted bipartite graph. Let μ be the optimal many-to-one matching, which can be applied to nodes in both U and V (where $u \in U$ has potentially multiple neighbors in V). Consider μ' , the greedy many-to-one matching for a particular sequence of arrivals σ .

Consider a particular $u \in U$, and the time it arrives $\sigma(u)$ as μ' progresses. If at least $1/2$ the value of $\mu(u)$ is available at that time, then $w(\mu'(u)) \geq \frac{1}{2}w(\mu(u))$ (since u can be matched to any subset of $\mu(u)$, by the downward closed assumption). If not, then $w(\mu'(\mu(u))) \geq \frac{1}{2}w(\mu(u))$ (at least half the value was taken by others). Thus, we know that, for all u ,

$$w(\mu'(u)) + w(\mu'(\mu(u))) \geq \frac{1}{2}w(\mu(u))$$

summing up over all u , we get

$$\sum_u w(\mu'(u)) + w(\mu'(\mu(u))) = 2w(\mu') \geq \frac{1}{2} \sum_u w(\mu(u)) = \frac{1}{2}w(\mu)$$

Rearranging shows that $w(\mu') \geq \frac{1}{4}w(\mu)$.

Finally, the utility to a player is clearly greatest when they are greedy, so that is a dominant strategy (thus implying any non-greedy strategy is dominated). ■

Recall, from our example in the introduction, that both greedy and undominated strategies can perform poorly with respect to empty counters. We defer the proof of the following results to Appendix B. Recall that \mathcal{M}_\emptyset refers to the empty counter mechanism.

Theorem 6.3.2. *There exist games g whose $\text{CR}_{\text{UNDOM}}(\mathcal{M}_\emptyset, g)$ cannot be bounded by any function of n .*

Theorem 6.3.3. *There exists g such that $\text{CR}_{\text{UNDOM}}(\mathcal{M}_\emptyset, g) \geq \Omega(\frac{n^2}{\log(n)})$, when $\mathcal{W}_r(t) = \frac{\mathcal{W}_r(0)}{t}$.*

6.3.2 Privacy-Preserving Counters and Greedy Behavior

We now present the main theorem of this section: namely, that if a mechanism provides approximate counts of the usage of each resource to each player, and those players choose their action greedily according to that information, this is enough to guarantee social welfare which approximates the optimal social welfare.

Theorem 6.3.4. *With (α, β) -accurate underestimator counter mechanism \mathcal{M} , $\text{CR}_{\text{GREEDY}}(\mathcal{M}, g) = O(\alpha\beta)$ for all resource-sharing games g .*

Before we prove Theorem 6.3.4, we need a way to compare players' utilities with the utility they *think* they get from choosing resources greedily with respect to approximate counters. Let a player's *perceived value* be $\mathcal{W}_r(y_{i,r})$ where r is the resource she chose (the value of a resource if the counter was correct, which may or may not be the *actual* value of the resource). We show for each resource r separately, by arguing that at most $\alpha\beta$ players can see $y_{i,r} = T$ for any T (since the counter for r is an (α, β) undercounter).

Lemma 6.3.5. *Suppose players choose greedily according to a (α, β) -underestimator. Then, the sum of their actual values is at least a $\frac{1}{2\alpha\beta}$ -fraction of the sum of their perceived values.*

Proof. Suppose k players chose a given resource r . For ease of notation, let these be players 1 through k . We wish to bound the ratio

$$\frac{\sum_{i=1}^k \mathcal{W}_r(y_{i,r})}{\sum_{c=1}^k \mathcal{W}_r(c)}.$$

By the definition of $y_{i,r}$ being an (α, β) underestimator, we can bound the desired ratio by rein-

dexing players by the value of the counter $y_{i,r}$ they observed:

$$\begin{aligned}
\frac{\sum_{i=1}^k \mathcal{W}_r(y_{i,r})}{\sum_{c=1}^k \mathcal{W}_r(c)} &\leq \frac{\sum_{x_{i,r}=0}^{k-1} \mathcal{W}_r\left(\frac{x_{i,r}}{\alpha} - \beta\right)}{\sum_{c=1}^k \mathcal{W}_r(c)} \\
&\leq \frac{\sum_{t=0}^{\lceil \frac{k}{\alpha\beta} \rceil} \alpha\beta \mathcal{W}_r((t-1)\alpha\beta)}{\sum_{c=1}^k \mathcal{W}_r(c)} \\
&\leq \frac{\alpha\beta \sum_{t=0}^{\lceil \frac{k}{\alpha\beta} \rceil} \mathcal{W}_r((t-1)\beta)}{\sum_{c=1}^k \mathcal{W}_r(c)} \\
&\leq \frac{\alpha\beta \sum_{t=1}^{\lceil \frac{k}{\alpha\beta} \rceil} \mathcal{W}_r((t-1)\beta)}{\sum_{t=1}^{\lceil \frac{k}{\alpha\beta} \rceil} \mathcal{W}_r((t-1)\beta)} \\
&\leq \alpha\beta
\end{aligned}$$

where the first inequality follows from the value curves being non-increasing and the counters are (α, β) -underestimators, the next term comes from the fact that the value curves are nondecreasing and grouping the values in multiples of $\alpha\beta$, the next from $\alpha > 1$ and nonincreasing value curves, and the penultimate from the fact that the value curves are non-negative. ■

We now have the tools we need to prove Theorem 6.3.4.

Proof of Theorem 6.3.4. The optimal value of the resource-sharing game g , denoted by $\text{OPT}(g)$, is the maximum value matching in the bipartite graph $G = (U \cup V, E)$ where U is the set of the n players and V has n vertices for each resource r , one of value $\mathcal{W}_r(k)$ for each $k \in [n]$. There is an edge between player i and all vertices corresponding to resource r if and only if $r \in A_i$. Note that the values are on the vertices in V .

We now define a complete bipartite graph G' which has the same set of nodes but whose node values differ for some nodes in G . Consider some resource r , and the collection of players who chose r in g . If there were t_k players i who chose resource r when $y_{i,r} = k$, make t_k of the nodes corresponding to r have value $\mathcal{W}_r(k)$. Finally, if there were F_k players who chose resource r , let the remaining $n - F_k$ nodes corresponding to r have value $\mathcal{W}_r(F_k + 1)$. This will be a lower bound on the perceived value $\mathcal{W}_r(y_{i,r})$ since these are underestimators.

We first claim that the perceived utility of players choosing greedily according to the counters is identical to the value of the greedy matching in G' (where nodes arrive in the same order). We prove, in fact, that the corresponding matching will be identical by induction. Since the counters are monotone, earlier copies of a resource appear more valuable. So, when the first player arrives in G' , the most valuable node she has access to is exactly the first node corresponding to the resource she took according to the counters. Now, assume that prior to player i , all players have chosen nodes corresponding to the resource they chose according to the counters. By our induction hypothesis and monotonicity of the counters and value curves, there is a node n_i corresponding to i 's selection r according to counters of value $\mathcal{W}_r(y_{i,r})$, and no higher-valued node corresponding to r . Likewise, for all other resources r' , all nodes corresponding to r' have

value more than $\mathcal{W}_{r'}(y_{i,r'})$. Thus, i will take n_i for value $\mathcal{W}_r(y_{i,r})$. Thus, the value of the greedy matching in G' equals the perceived utility of greedy play according to the counters.

Let $\text{GREEDY}_{\text{COUNTERS}}$ denote the set of actions players make playing greedily with respect to the counters. By Lemma 6.3.5, the social welfare of $\text{GREEDY}_{\text{COUNTERS}}$ is a $\frac{1}{\alpha\beta}$ -fraction of the perceived social welfare. By our previous argument, the perceived social welfare of greedy play according to the counters is the same as the value of the greedy matching in G' . By Theorem 6.3.1, the greedy matching in G' is a 4-approximation to the max-value matching in G' . Finally, since the counters are underestimators, the value of the max-value matching in G' is at least as large as $\text{OPT}(g)$. Thus, that the social welfare of greedy play with respect to counters is a $\frac{1}{4\alpha\beta}$ fraction of the optimal welfare of g . ■

As a corollary of Theorem 6.3.4, the following theorem shows that greedy play with respect to private undercounters has a polylogarithmic competitive ratio.

Theorem 6.3.6. *There exists (ε, δ) -privacy-preserving mechanism \mathcal{M} such that*

$$\text{CR}_{\text{GREEDY}}(\mathcal{M}, g) \leq \min \left(O \left(\frac{\log n \log \frac{nm}{\delta}}{\varepsilon} \right), O \left(\frac{m \log n \log \log \frac{1}{\delta}}{\varepsilon} \right) \right)$$

for all resource-sharing games g .

Proof. In Section 6.6, we prove Corollary 6.6.5 that says that we can achieve an (ε, δ) -differentially private counter vector achieving the better of $(1, O(\frac{(\log n)(\log(nm/\delta))}{\varepsilon}))$ -accuracy and $(\alpha, \tilde{O}_\alpha(\frac{m \log n \log \log(1/\delta)}{\varepsilon}))$ -accuracy for any constant $\alpha > 1$. This along with Theorem 6.3.4 proves the result. ■

Observation 6.3.2 states that players acting greedily according to any estimate that is *deterministically* in between the true count and the values provided by an (α, β) counter vector also achieve similar or better social welfare guarantees. Conversely, Observation 6.3.3 implies that higher accuracy than some (α, β) underestimators alone is not enough to make such a claim due to higher variance.

Observation 6.3.2. Suppose that \mathcal{M} is a (α, β, γ) underestimator, giving estimates $y_{i,r}$. Furthermore, assume each player i is playing greedily with respect to a revised estimate $z_{i,r}$ such that, for each r, i , and value of $z_{i,r}$ is always in the range $[y_{i,r}, x_{i,r}]$. Then, for g , a discrete resource-sharing game, with probability $1 - \gamma$, the ratio of the optimal to the achieved social welfare is $O(\alpha\beta)$.

Proof. Since the $z_{i,r}$'s are deterministically more accurate than the COUNTERS , we have for each i that the utility achieved by greedily choosing according to the estimates $z_{i,r}$ is at least as much as the utility achieved by greedily choosing using $y_{i,r}$. Therefore, summing over all the players, the achieved social welfare is at least as much as it would be if everyone had played greedily according to $y_{i,r}$. Then, applying Theorem 6.3.4, the observation follows. ■

Observation 6.3.3. There exists a resource-sharing game g , such that if the players play greedily according to estimates $z_{i,r}$ that are more accurate than the displayed value only in expectation – specifically for each r, i , and value of $x_{i,r}$, $\mathbb{P}[z_{i,r} < x_{i,r}] \geq 1/2$ and also $\mathbb{E}[|z_{i,r} - x_{i,r}|] = 1$, then the ratio of the optimal to the achieved social welfare can be as bad as $\Omega(\sqrt{n})$.

Proof. Let there be $n + \sqrt{n}$ resources, with resources $r_{1, \dots, \sqrt{n}}^*$ having $\mathcal{W}_{r_f^*}(0) = H$, $\mathcal{W}_{r_f^*}(t) = 0$ for all $t > 0$, and resource r_i such that $\mathcal{W}_{r_i}(t) = H - \varepsilon$ for all t . Player i has access to all resources r_f^* and r_i . Then, $\text{OPT} = H\sqrt{n} + (H - \varepsilon)(n - \sqrt{n}) = Hn - (n - \sqrt{n})\varepsilon$.

Consider the counter vector which is exactly correct with probability $1 - \frac{1}{\sqrt{n}}$ and undercounts by \sqrt{n} with probability $\frac{1}{\sqrt{n}}$ (note that the expected error is just 1 and it undercounts with probability 1). Then, greedy behavior with respect to this counter will (in expectation) have \sqrt{n} players choose r_{*f} for each f , achieving welfare $\sqrt{n}H$. Thus, the competitive ratio is $\Omega(\sqrt{n})$ as $\varepsilon \rightarrow 0$, as desired. ■

6.3.3 Privacy-Preserving Counters and Undominated behavior

We now discuss the performance of *undominated* strategies with respect to private counters. We begin with an illustration of how undominated strategies can perform poorly for arbitrary value curves, as motivation for the restricted class of value curves we consider in Theorem 6.3.8. In the case of greedy players, we were able to avoid the problem of players undervaluing resources rather easily, by forcing the counters to only underestimate $x_{i,r}$. This won't work for undominated strategies: players who know the counts are shaded downward can compensate for that fact in undominated strategies, and only choose lower-valued resources.

Theorem 6.3.7. *For an (ε, δ) -differentially private announcement mechanism \mathcal{M} , there exist games g for which*

$$\text{CR}_{\text{UNDOM}}(g, \mathcal{M}) = \Omega\left(\frac{1}{\delta}\right).$$

Proof. Suppose there are two players 1 and 2, and resources r, r' . Let r have $\mathcal{W}_r(0) = 1$, $\mathcal{W}_r(1) = 0$, and $\mathcal{W}_{r'}(k) = \rho$, for all $k \geq 0$. Furthermore, let player 1 have access only to resource r' but player 2 has access to both r and r' . Player 1 will choose r' . Let player 2's strategy be d_2 , such that if she determines there was nonzero chance that player 1 chose r according to her signal m_2 , she will choose resource r' . This is undominated: if 1 *did* choose r , r' will be more valuable for 2. Thus, if 2 sees any signal that can occur when r is chosen by 1, she will choose r' . The collection of signals 2 can see if 1 chooses r has probability 1 in total. So, because m_2 is (ε, δ) -differentially private in player 1's action, the set of signals *reserved* for the case when 1 chooses r' (that cannot occur when r is chosen by 1) may occur with probability at most δ (they can occur with probability 0 if 1 chose r , implying they can occur with probability at most δ when 1 chooses r'). Thus, with this probability $1 - \delta$, player 2 will choose r' , implying $\mathbb{E}[SW] \leq (1 - \delta)2\rho + \delta(1 + \rho) = \delta + (2 - \delta)\rho$, which for ρ sufficiently small approaches δ , while $1 + \rho$ is the optimal social welfare. ■

Given the above example, we cannot hope to have a theorem as general as Theorem 6.3.4 when analyzing undominated strategies with privacy-preserving counters. Instead, we show that, for a class of well-behaved value curves, we can bound the competitive ratio of undominated strategies. The proof follows a similar outline to that of Theorem 6.3.4, showing that undominated resources r' have utility which is close to the resource of r , the resource chosen by a greedy player.

Theorem 6.3.8. *If each value curve \mathcal{W}_r has the property that*

$$\psi(\alpha, \beta)\mathcal{W}_r(x) \geq \mathcal{W}_r\left(\max\left\{0, \frac{x}{\alpha^2} - \frac{2\beta}{\alpha}\right\}\right)$$

and also

$$\mathcal{W}_r((\alpha^2 x + 2\alpha\beta)) \geq \phi(\alpha, \beta)\mathcal{W}_r(x),$$

then an action profile a of undominated strategies according to (α, β) -underestimator \mathcal{M} has $\text{CR}_{\text{UNDOM}}(g, \mathcal{M}) = O(\psi(\alpha, \beta)\phi(\alpha, \beta))$.

In particular, Theorem 6.3.8 shows that, for games where $\mathcal{W}_r(i) = \frac{\mathcal{W}_r(0)}{g_r(x_{i,r})}$, where g_r is a polynomial, the competitive ratio of undominated strategies degrades gracefully as a function of the maximum degree of those polynomials. Corollary 6.3.9 (whose proof follows from a simple calculation) states this formally.

Corollary 6.3.9. *Suppose for a resource-sharing game g , each resource r has a value curve of the form $\mathcal{W}_r(x) = \frac{\mathcal{W}_r(0)}{g_r(x)}$, where g_r is a monotonically increasing degree- d polynomial and $\mathcal{W}_r(0)$ is some constant. Then,*

$$\text{CR}_{\text{UNDOM}}(g, \mathcal{M}) \leq O(2\alpha^3\beta)^d \leq \min\left(O\left(\left(\frac{\log n \log \frac{nm}{\delta}}{\varepsilon}\right)^d\right), O\left(\left(\frac{m \log n \log \log \frac{1}{\delta}}{\varepsilon}\right)^d\right)\right)$$

with \mathcal{M} providing (α, β) -counters.

6.4 Resource Sharing with Continuous Resources

In this section, we allow agents' choice of resources to be non-discrete. The utility of player i in the continuous model is the following:

$$u_i(a_1, \dots, a_n) = \sum_{r=1}^m \int_{x_{i,r}}^{x_{i,r} + a_{i,r}} v_r(t) dt,$$

where $x_{i,r} = \sum_{i'=1}^{i-1} a_{i',r}$ is the amount already invested in resource r by earlier players.

In this setting, in order to prove a theorem analogous to Theorem 6.3.4 in the discrete setting, we need an analogue to Lemma 6.3.1 that holds in the full-information continuous setting. We no longer have the tight connection between our setting and matching; nonetheless, the fact that the greedy strategy is a 4-approximation to OPT continues to hold.

Lemma 6.4.1. *The greedy strategy for online, continuous, resource-weighted correspondences, where players arrive online and have tuples of allowable volumes of resources, has a competitive ratio of $\frac{1}{4}$.*

Proof Sketch. The proof is identical to the proof of Lemma 6.3.1, with the exception that we no longer want matchings μ, μ' but rather correspondences between continuous regions of the v_r . It is either the case that a player gets utility at least $1/2$ of what she got when everyone was assigned to their optimal action, or $1/2$ of the utility she would have achieved in the optimal setting is captured by some collection of other players. ■

With Lemma 6.4.1, following analysis similar to Theorem 6.3.4, we have the following.

Theorem 6.4.2. *Suppose that \mathcal{M} is an (α, β, γ) -underestimating counter vector. Then, for any continuous, future-independent resource-sharing game g , $\text{CR}_{\text{GREEDY}}(\mathcal{M}, g) = O(\alpha\beta)$.*

6.5 Resource Sharing with Future-Dependent Utilities

The second model of utility we consider is one where the benefit of choosing a resource for a player depends not only the actions of the past players but also on the actions taken by future players. Specifically, all the players who selected a given resource incur the same benefit regardless of the order in which they made the choice. The utility of player i ,

$$u_i(a_1, \dots, a_n) = \sum_{r \in [m]} a_{i,r} \mathcal{W}_r(x_r),$$

where $x_r = \sum_{i'=1}^n a_{i',r}$ is the total utilization of resource r by all players. In this general setting, we investigate the case where value curves do not decrease too quickly¹⁰. Our formal restriction on value curves, which follows, says that the actual welfare from any resource being utilized with x weight is not too much smaller than the integral of \mathcal{W}_r from 0 to x . This will imply that myopic decisions will ultimately have welfare close to the utility an agent believes she will get for her myopic decision, assuming she was the last agent in the system.

Definition 6.5.1 ((w, l) -shallow value curve). A value curve \mathcal{W}_r is (w, l) -shallow if for all $x \leq l$, it is the case that $\mathcal{W}_r(x) \geq \frac{\sum_{t=0}^x \mathcal{W}_r(t)}{wx}$.

For example, the curve $\mathcal{W}_r(x) = 1/x$ is $(\ln n, n)$ -shallow (the total utility for any k bidders choosing the resource is 1, while the perceived total utility would be $\ln k$). Other curves which are polynomial in x are also shallow (bounded by the degree and coefficients of said polynomial). We further restrict our attention to greedy behavior, as the following example shows that undominated behavior can perform quite poorly, even when the value curves decrease slowly.

Example 6.5.1. Suppose there are n players. Consider the case where for every $i \geq 1$, player i is interested in resource r_0 and resource r_i . For every $i \geq 1$, $\mathcal{W}_{r_i}(j) = \frac{(n-i+1)(1-\varepsilon)}{ji}$ (for some small $\varepsilon > 0$). Then, let $\mathcal{W}_{r_0}(j) = \frac{1}{j}$. We claim that there is an undominated strategy game play where every player chooses resource r_0 giving a social welfare of 1, whereas the optimal welfare is achieved by assigning player i resource r_i giving a total welfare of $n(\log(n) - 1)(1 - \varepsilon)$. In particular, the following strategy profile is undominated: for each i , player i believes that every player after her has access only to resource r_i . With this belief, it is easy to see that choosing resource r_0 is an undominated strategy for every player, achieving welfare at most 1.

¹⁰If the curves decrease too quickly, an early player's decision which is independent of later players' behavior cannot always have bounded performance. For example, imagine a setting where player 1 has two possible choices, r, r' , with $\mathcal{W}_r(1) = M$, $\mathcal{W}_r(2) = 0$, and $\mathcal{W}_{r'}(1) = 1$, $\mathcal{W}_{r'}(2) = \delta$. The choice of r will yield 0 welfare when player 2 has only r in her action set (where welfare of $M + 1$ is possible); the choice of r' will yield 0 welfare when player 2 has only r' in her action set (while $M + 1$ welfare is again achievable). Thus, agent 1's choice cannot be independent of agent 2's type or behavior and get any approximation to optimal welfare unless the curves decrease more slowly.

We now state the main theorem of this section, which says that greedy behavior performs well with respect to approximate counters, so long as the resource value curves do not decay too quickly.

Theorem 6.5.1. *Suppose, for a sequential resource-sharing game g , each resource r 's value curve v_r is (w, n) – shallow. Then, in the in the future-dependent setting, $\text{CR}_{\text{GREEDY}}(\mathcal{M}, g) = O(w\alpha\beta)$ for an (α, β) -underestimator \mathcal{M} .*

One might wonder if this polynomial dependence on w is necessary: the following result shows that greedy behavior's performance necessarily decays in this parameter.

Lemma 6.5.2. *Even with perfect counters, there exist sequential resource-sharing games g , where each resource r 's value curve \mathcal{W}_r is (w, n) – shallow, such that in the future-dependent setting, $\text{CR}_{\text{GREEDY}}(\mathcal{M}_{\text{Full}}, g) \geq 2w$.*

Proof. Consider two players and two resources r, r' . Let r have a value curve which is a step function, with $v_r(0) = w$, $v_r(1) = \frac{1}{2}$ and $v_{r'}(0) = w - \varepsilon$. Suppose player one has access to both resources and player two has only resource r as an option. Then, player one will choose r according to greedy, and player two will always select r . The social welfare will be $\text{SW}(\text{GREEDY}) = 1$, whereas OPT is for player 1 to take r' and will have $\text{SW}(\text{OPT}) = 2w - \varepsilon$. As $\varepsilon \rightarrow 0$, this ratio approaches $2w$. ■

Thus, as $w \rightarrow \infty$, the competitive ratio of the greedy strategy is unbounded. Fortunately, the competitive ratio cannot have worse dependence on w , implied by Theorem 6.5.1: we sketch the theorem's proof below.

Proof Sketch of Theorem 6.5.1. According to the greedy strategy, player i chooses the resource in A_i that maximizes $\mathcal{W}_r(x_{i,r} + 1)$, and we say $\mathcal{W}_r(x_{i,r} + 1)$ is her perceived value if r is the resource she chose. Let the sum of the perceived values of all players as the perceived social welfare be denoted by $\text{PSW}(\text{GREEDY})$, we have

$$\begin{aligned} \text{PSW}(\text{GREEDY}) &= \sum_{i \in [n]} \sum_r a_{i,r} \mathcal{W}_r(x_{i,r} + 1) dx = \sum_r \sum_{x=0}^{x_{n,r} + a_{n,r}} \mathcal{W}_r(x) \\ &\leq \sum_r w(x_{n,r} + a_{n,r}) \mathcal{W}_r(x_{n,r} + a_{n,r}) = w \text{SW}(\text{GREEDY}) \end{aligned} \tag{6.1}$$

where the last inequality comes from our assumption about the value curves all being (w, n) – shallow.

The final part of the argument must show that the actual welfare from greedy play with respect to the counters is well-approximated by the perceived welfare with respect to the true counts. Since the counters are accurate within some quantity $\leq n$, this is the case. Following an analysis similar to that of Theorem 6.3.4, we have our result. ■

6.6 Private Counters with Smaller Error at Smaller Values

In this section, we describe a counter for the model of differential privacy under continual observation that has improved guarantees when the value of the counter is small. Recall the basic counter problem: given a stream $\vec{a} = (a_1, a_2, \dots, a_n)$ of numbers $a_i \in [0, 1]$, we wish

to release at every time step t the partial sum $x_t = \sum_{i=1}^t a_i$. We require a generalization, where one maintains a vector of m counters. Each player’s update contribution is now a vector $a_i \in \Delta_m = \{a \in [0, 1]^m \mid \|a\|_1 \leq 1\}$. That is, a player can add non-negative values to all counters, but the total value of her updates is at most 1. The partial sums x_t then lie in $(\mathbb{R}^+)^m$ and have ℓ_1 norm at most t .

Given an algorithm \mathcal{M} , we define the output stream $(y_1, \dots, y_n) = \mathcal{M}(\vec{a})$ where $y_t = \mathcal{M}(t, a_1, \dots, a_{t-1})$ lies in \mathbb{R}^m . We seek counters that are private (Definition 6.2.1) and satisfy a mixed multiplicative and additive accuracy guarantee (Definition 6.2.2). Proofs of all the results in this section can be found in Appendix B.2.

The original works on differentially private counters [34, 50] concentrated on minimizing the additive error of the estimated sums, that is, they sought to minimize $\|x_t - y_t\|_\infty$. Both papers gave a binary tree-based mechanism, which we dub “TreeSum”, with additive error approximately $(\log^2 n)/\varepsilon$. Some of our algorithms use TreeSum, and others use a new mechanism (FTSum, described below) which gets a better additive error guarantee at the price of introducing a small multiplicative error. Formally, they prove:

Lemma 6.6.1. *For every $m \in \mathbb{N}$ and $\gamma \in (0, 1)$: Running m independent copies of TreeSum [34, 50] is $(\varepsilon, 0)$ -differentially private and provides an $(1, C_{tree} \cdot \frac{\log n \log \frac{nm}{\gamma}}{\varepsilon}, \gamma)$ -approximation to partial vector sums, where $C_{tree} > 0$ is an absolute constant.*

Even for $m = 1, \alpha = 1$, this bound is slightly tighter than those in Chan et al. [34] and Dwork et al. [50]; however, it follows directly from the tail bound in Chan et al. [34]. Our new algorithm, FTSum (for Flag/Tree Sum), is described in Algorithm 5. For small m (specifically, when $m = o(\log(n))$), it provides lower additive error at the expense of introducing an arbitrarily small constant multiplicative error.

Lemma 6.6.2. *For every $m \in \mathbb{N}$, $\alpha > 1$ and $\gamma \in (0, 1)$, FTSum (Algorithm 5) is $(\varepsilon, 0)$ -differentially private and $(\alpha, \tilde{O}_\alpha(\frac{m \log \frac{n}{\gamma}}{\varepsilon}), \gamma)$ -approximates partial sums (where $\tilde{O}_\alpha(\cdot)$ hides poly-logarithmic factors in its argument, and treats α as constant).*

FTSum proceeds in two phases. In the first phase, it increments the reported output value only when the underlying counter value has increased significantly. Specifically, the mechanism outputs a public signal, which we will call a “flag”, roughly when the true counter achieves the values $\log n, \alpha \log n, \alpha^2 \log n$ and so on, where α is the desired *multiplicative* approximation. The reported estimate is updated each time a flag is raised (it starts at 0, and then increases to $\log n, \alpha \log n$, etc). The privacy analysis for this phase is based on the “sparse vector” technique of Hardt and Rothblum [67], which shows that the cost to privacy is proportional to the number of times a flag is raised (but not the number of time steps between flags).

When the value of the counter becomes large (about $\frac{\alpha \log^2 n}{(\alpha-1)\varepsilon}$), the algorithm switches to the second phase and simply uses the TreeSum protocol, whose additive error (about $\frac{\log^2 n}{\varepsilon}$) is low enough to provide an α multiplicative guarantee (without need for the extra γ space given by the additive approximation).

If the mechanism were to raise a flag *exactly* when the true counter achieved the values $\log n, \alpha \log n, \alpha^2 \log n$, etc, then the mechanism would provide a $(\alpha, \log n, 0)$ approximation during the first phase, and a $(\alpha, 0, 0)$ approximation thereafter. The rigorous analysis is more complicated, since flags are raised only near those thresholds.

Algorithm 5: FTSum — A Private Counter with Low Multiplicative Error

Input: Stream $\vec{a} = (a_1, \dots, a_n) \in ([0, 1]^m)^n$, parameters $m, n \in \mathbb{N}$, $\alpha > 1$ and $\gamma > 0$

Output: Noisy partial sums $y_1, \dots, y_n \in \mathbb{R}^m$

$k \leftarrow \lceil \log_\alpha \left(\frac{\alpha}{\alpha-1} \cdot C_{tree} \cdot \frac{\log(nm/\gamma)}{\varepsilon} \right) \rceil$;

/ C_{tree} is the constant from Lemma 6.6.1 */*

$\varepsilon' \leftarrow \frac{\varepsilon}{2m(k+1)}$;

for $r = 1$ **to** m **do**

$\text{flag}_r \leftarrow 0$;

$x_{0,r} \leftarrow 0$;

$\tau_r \leftarrow (\log n) + \text{Lap}(2/\varepsilon')$;

for $i = 1$ **to** n **do**

for $r = 1$ **to** m **do**

if $\text{flag}_r \leq k$ **then** (First phase still in progress for counter r)

$x_{i,r} \leftarrow x_{i-1,r} + a_{i,r}$;

$\widetilde{x}_{i,r} \leftarrow x_{i,r} + \text{Lap}(\frac{2}{\varepsilon'})$;

if $\widetilde{x}_{i,r} > \tau_r$ **then** (Raise a new flag for counter r)

$\text{flag}_r \leftarrow \text{flag}_r + 1$;

$\tau_r \leftarrow (\log n) \cdot \alpha^{\text{flag}_r} + \text{Lap}(2/\varepsilon')$;

Release $y_{i,r} = (\log n) \cdot \alpha^{\text{flag}_r - 1}$;

else (Second phase has been reached for counter r)

Release $y_{i,r} = r$ -th counter output from $\text{TreeSum}(\vec{a}, \varepsilon/2)$;

Enforcing Additional Guarantees Finally, we note that it is possible to enforce to additional useful properties of the counter. First, we may insist that the accuracy guarantees be satisfied with probability 1 (that is, set $\gamma = 0$), at the price of increasing the additive term δ in the privacy guarantee:

Proposition 6.6.3. *If \mathcal{M} is (ε, δ) -private and (α, β, γ) -accurate, then one can modify \mathcal{M} to obtain an algorithm \mathcal{M}' with the same efficiency that is $(\varepsilon, \delta + \gamma)$ -private and $(\alpha, \beta, 0)$ -accurate.*

Second, as in [50], we may enforce the requirement that the reported values be monotone, integral values that increase at each time step by at most 1. The idea is to simply report the nearest integral, monotone sequence to the noisy values (starting at 0 and incrementing the reported counter only when the noisy value exceeds the current counter).

Proposition 6.6.4 ([50]). *If \mathcal{M} is (ε, δ) -private and (α, β, γ) -accurate, then one can modify \mathcal{M} to obtain an algorithm \mathcal{M}' which reports monotone, integral values that increase by 0 or 1 at each time step, with the same privacy and accuracy guarantees as \mathcal{M} .*

Corollary 6.6.5. *Algorithm 5 is an (ε, δ) -differentially private vector counter algorithm providing a*

1. $(1, O(\frac{(\log n)(\log(nm/\delta))}{\varepsilon}), 0)$ -approximation (using modified *TreeSum*); or
2. $(\alpha, \tilde{O}_\alpha(\frac{m \log n \log \log(1/\delta)}{\varepsilon}), 0)$ -approximation for any constant $\alpha > 1$ (using *FTSum*).

6.7 Discussion and Open Problems

In this work, we considered how public dissemination of information in sequential games can guarantee a good social welfare while maintaining differential privacy of the players' strategies. We considered two 'extreme' cases – the greedy strategy and the class of all undominated strategies. While analyzing the class of undominated strategies gives guarantees that are robust, in many games that we considered, the competitive ratios were significantly worse than greedy strategies, and in some cases they were unbounded. It is interesting to note that many of the examples in this work that show the poor performance with undominated strategies also hold when the players know their position in the sequence, an assumption we have not made for any of the positive results in this work. It is an interesting direction for future research to consider classes of strategies that more restricted than undominated strategies yet are general enough to be relevant for games where players play with imperfect information.

As mentioned in the introduction, we note here that, while players are making choices subject to approximate information, our results are not a direct extension of the line of thought that approximate information implies approximate optimization. In particular, for greedy strategies, while there may be a bound on the error of the counters, that *does not imply*, for arbitrary value curves, playing greedily according to the counters will be *approximately optimal for each individual*. In particular, consider one resource r with value H for the first 10 investors, and value 0 for the remaining investors, and a second resource r' with value $H/2$ for all investors. With (α, β, γ) , as many as β players might have unbounded ratio between their value for r as r' , but will pick r over r' . The analysis of greedy shows, despite this anomaly, the total social welfare is still well-approximated by this behavior.

All of our results relied on using differentially private counters for disseminating informa-

tion. For the differentially-private counter, a main open question is “what is the optimal trade-off between additive and multiplicative guarantees?”. Furthermore, as part of future research, one can consider other privacy techniques for announcing information that can prove useful in helping players achieve a good social welfare. And more generally, we want to understand what features of games lend themselves to be amenable to public dissemination of information that helps achieve good welfare and simultaneously preserves privacy of the players’ strategies.

Appendix A

Appendix for Chapter 4

A.1 Proofs of inequalities from Chapter 4

<p>Algorithm 6: Estimates $\mathbb{P}[\max_j b_j \geq \ell_\tau \mid \max_j b_j \leq \ell_{\tau+1}]$</p> <p>Data: $\ell_\tau, \ell_{\tau+1}, T$</p> <p>Result: $p_{\ell_\tau, \ell_{\tau+1}}^\epsilon$</p> <p>Let S_1 be a sample of size T with reserve ℓ_τ; Let S_2 be a sample of size T with reserve $\ell_{\tau+1}$;</p> <p>Return $p_{\ell_\tau, \ell_{\tau+1}}^\epsilon = 1 - \frac{\sum_{t \in S_2} \mathbb{I}[0 \text{ wins } t]}{\sum_{t \in S_1} \mathbb{I}[0 \text{ wins } t]}$;</p>
--

Proof of Observation 4.3.1.

$$F_i(\ell_{\tau-1}) = F_i(\ell_\tau)(1 - \mathbb{P}[b_i \geq \ell_{\tau-1} \mid b_i \leq \ell_\tau]) = \prod_{\tau' \geq \tau} (1 - \mathbb{P}[b_i \geq \ell_{\tau'-1} \mid b_i \leq \ell_{\tau'}]) = \prod_{\tau' \geq \tau} (1 - \mathbb{P}[b_i \in [\ell_{\tau'-1}, \ell_{\tau'}] \mid b_i \leq \ell_{\tau'}])$$

■

Lemma A.1.1. *Suppose X is observable and Y is observable, and assume that $\mathbb{P}[Y] \geq \gamma$. Using $2T$ samples, with probability $1 - \delta$, we can estimate $\mathbb{P}[X|Y] = \frac{\mathbb{P}[X \cap Y]}{\mathbb{P}[Y]}$ by \hat{p} such that*

$$\mathbb{P}[X|Y] - \alpha - \mu \leq (1 - \alpha)\mathbb{P}[X|Y] - \beta \leq \hat{p} \leq (1 + \alpha)\mathbb{P}[X|Y] + \beta \leq \mathbb{P}[X|Y] + \alpha + \mu,$$

As a direct corollary, we know that `Inside` is a close approximation to the quantity it estimates.

Corollary A.1.2. *`Inside`($\ell_\tau, \ell_{\tau+1}, T$) outputs an estimator $p_{\ell_\tau, \ell_{\tau+1}}^\epsilon$, such that, for T as in Kaplan,*

$$(1 - \alpha)\mathbb{P}[\max_j b_j \geq \ell_\tau \mid \max_j b_j \leq \ell_{\tau+1}] - \beta \leq p_{\ell_\tau, \ell_{\tau+1}}^\epsilon \leq (1 + \alpha)\mathbb{P}[\max_j b_j \geq \ell_\tau \mid \max_j b_j \leq \ell_{\tau+1}] + \beta$$

and uses $2T$ samples.

Now, we prove Lemma 4.3.3, which is also a corollary of Lemma A.1.1.

Proof of Lemma 4.3.3. Let, for a fixed i , $\ell_\tau, \ell_{\tau+1}$, the event that i bids in $[\ell_\tau, \ell_{\tau+1}]$ be denoted by X , the event that i wins in $[\ell_\tau, \ell_{\tau+1}]$ be denoted by Y , and the event that $\max_j b_j < \ell_{\tau+1}$ be denoted by C .

With this notation, we have an estimate of $\mathbb{P}[Y|C]$ and want an estimate of $\mathbb{P}[X|C]$.

$$\begin{aligned}\mathbb{P}[Y|C] &= \mathbb{P}[X|C] \times \mathbb{P}[Y|C, X] \\ &\geq \mathbb{P}[X|C] \times \mathbb{P}[\text{everyone but } i \text{ bids } < \ell_\tau | C, X] \\ &= \mathbb{P}[X|C] \times \mathbb{P}[\text{everyone but } i \text{ bids } < \ell_\tau | C] \\ &\geq \mathbb{P}[X|C] \times (1 - \beta)\end{aligned}$$

The first equality comes from the fact that $Y \subseteq X$, the next inequality comes from the fact that, conditioned on C and X , everyone but i bids $< \ell_\tau$ is a subset of Y (the times when i will win), the next equality comes from the fact that i 's bid and j 's bid are independent, and the final inequality follows from the assumption $\mathbb{P}[\max_{j \neq i} b_j < \ell_\tau | \max_{j \neq i} b_j < \ell_{\tau+1}] \geq 1 - \beta$. \blacksquare

Fact A.1.3. Suppose $x \geq 0$ and $0 < \eta < \frac{1}{2}$. Then $\frac{x}{1+\eta} \geq (1 - \eta)x$ and $\frac{x}{1-\eta} \leq (1 + 2\eta)x$.

Proof of Fact A.1.3. We prove $\frac{x}{1+\eta} \geq (1 - \eta)x$ first.

$$\frac{x}{1+\eta} = \frac{(1-\eta)x}{1-\eta^2} \geq (1-\eta)x \quad (\text{Since } 1-\eta^2 < 1)$$

Now, we prove $\frac{x}{1-\eta} \leq (1 + 2\eta)x$, for $\eta \leq 1/2$. We have,

$$\frac{x}{1-\eta} = x \sum_{i=0}^{\infty} \eta^i = x \left(1 + \eta \left(\sum_{i=0}^{\infty} \eta^i \right) \right) \leq (1 + 2\eta)x,$$

where the inequality follows from the fact that for $\eta \leq 1/2$ we have $\sum_{i=0}^{\infty} \eta^i = \frac{1}{1-\eta} \leq 2$. \blacksquare

Proof of Theorem 4.3.1. Notice that there are at most k' events each of which happens with probability at most $\delta' = \frac{\delta}{k'}$ (namely, that `Intervals` returns a poor partition, or for each interval, of which there are at most $k' - 1$, by Lemma 4.3.2, that `Iwin` is not accurate as described by Lemma 4.3.1). Thus, by a union bound, none of these events occur with probability $1 - \delta$. Thus, for the remainder of the proof we assume the partition returned by `Intervals` is good and each call to `Iwin` is accurate.

It will suffice to prove, for the lattice points in our discretization, that `Kaplan` provides an ε -approximation to the CDF. This follows because

$$\begin{aligned}F_i(\ell_\tau) - F_i(\ell_{\tau-1}) &= \mathbb{P}[i \text{ bids in } [\ell_{\tau-1}, \ell_\tau]] \\ &= \mathbb{P}[i \text{ bids in } [\ell_{\tau-1}, \ell_\tau] | b_i \leq \ell_\tau] \\ &\leq \mathbb{P}[i \text{ bids in } [\ell_{\tau-1}, \ell_\tau] | \max_j b_j \leq \ell_\tau] \\ &\leq (1 + \beta) \mathbb{P}[i \text{ wins in } [\ell_{\tau-1}, \ell_\tau] | \max_j b_j \leq \ell_\tau] \\ &\leq (1 + \beta)\beta = \beta + \beta^2 \leq \frac{\varepsilon}{2}\end{aligned}$$

where the third and fourth inequality follows from Lemma 4.3.2 and Lemma 4.3.3, and the final one from the fact that $\beta < \frac{\varepsilon}{4}$. Thus, our lattice is fine enough that it suffices to show accuracy of the lattice points. We start by rewriting $F_i(\ell_\tau)$, using Observation 4.3.1:

$$F_i(\ell_\tau) = \prod_{\tau' \geq \tau+1} (1 - \mathbb{P}[b_i \in [\ell_{\tau'-1}, \ell_{\tau'}] \mid b_i \leq \ell_{\tau'}]) \quad (\text{A.1})$$

So, one can compute the probability of bidding at most $\ell_{\tau-1}$ by multiplying together a collection of probabilities of bidding within intervals above ℓ_τ . Let the event $\max_j b_j \leq \ell_{\tau'}$ be denoted $M_{\ell_{\tau'}}$. Now, we can apply Lemma 4.3.2 to imply that, for all τ' ,

$$\mathbb{P}[\max_j b_j \in [\ell_{\tau'-1}, \ell_{\tau'}] \mid M_{\ell_{\tau'}}] \leq \frac{\beta}{16} = \beta'$$

which, by Lemma 4.3.3, implies for all τ' that

$$\begin{aligned} 1 &\geq \frac{\mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] \mid M_{\ell_{\tau'}}]}{\mathbb{P}[i \text{ bids in } [\ell_{\tau'-1}, \ell_{\tau'}] \mid m_{\ell_{\tau'}}]} \\ &= \frac{\mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] \mid M_{\ell_{\tau'}}]}{\mathbb{P}[i \text{ bids in } [\ell_{\tau'-1}, \ell_{\tau'}] \mid i \text{ bids in } [0, \ell_{\tau'}]]} \geq 1 - \beta' \end{aligned} \quad (\text{A.2})$$

where the equality comes from the independence of the bids. Then, combining Equations (A.1) and (A.2), we know

$$\prod_{\tau' \geq \tau+1} (1 - \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] \mid M_{\ell_{\tau'}}]) = F_i(\ell_\tau) \geq \prod_{\tau' \geq \tau+1} (1 - \frac{\mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] \mid M_{\ell_{\tau'}}]}{1 - \beta'})$$

Then, by Fact A.1.3,

$$F_i(\ell_\tau) \in \left[\prod_{\tau' \geq \tau+1} (1 - (1 + 2\beta')\mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] \mid M_{\ell_{\tau'}}]), \prod_{\tau' \geq \tau+1} (1 - \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] \mid M_{\ell_{\tau'}}]) \right]$$

Now, Lemma 4.3.1 states that the result of I_{win} are correct within an additive β and multiplicative α , thus

$$\prod_{\tau' \geq \tau+1} (1 - (1 + \alpha)\mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] \mid M_{\ell_{\tau'}}] - \beta) \leq \widehat{F}_i(\ell_\tau) \leq \prod_{\tau' \geq \tau+1} (1 - (1 - \alpha)\mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] \mid M_{\ell_{\tau'}}])$$

Now, we simply need to look at the potential difference in these terms. We will consider the

lower bound on $F_i(\ell_\tau)$ and upper bound on $\widehat{F}_i(\ell_\tau)$ (the other direction is analogous).

$$\begin{aligned}
& \prod_{\tau' \geq \tau+1} (1 - (1 - \alpha) \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}] + \beta) - \prod_{\tau' \geq t+1} (1 - (1 + 2\beta') \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}]) \\
& \leq \prod_{\tau' \geq \tau+1} (1 - \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}] + \alpha\beta' + \beta) - \prod_{\tau' \geq \tau+1} (1 - \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}] - 2\beta'^2) \\
& \leq \prod_{\tau' \geq \tau+1} (1 - \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}] + \beta'^2) - \prod_{\tau' \geq \tau+1} (1 - \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}] - 2\beta'^2) \\
& \leq \prod_{\tau' \geq \tau+1} (1 - 2\beta'^2)(1 - \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}]) - \prod_{\tau' \geq \tau+1} (1 + 2\beta'^2)(1 - \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}]) \\
& \leq (1 - 2\beta'^2)^k \prod_{\tau' \geq \tau+1} (1 - \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}]) - (1 + 4\beta'^2)^k \prod_{\tau' \geq \tau+1} (1 - \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}]) \\
& \leq (1 - 4k\beta'^2) \prod_{\tau' \geq \tau+1} (1 - \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}]) - (1 + 8k\beta'^2) \prod_{\tau' \geq t+1} (1 - \mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}]) \\
& \leq 12k\beta'^2 \leq 12 \frac{16Ln}{\beta\gamma} \beta'^2 \leq \frac{3Ln\beta}{\gamma} \leq \frac{\varepsilon}{2}
\end{aligned}$$

where the first follows from $\mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | \max_j b_j < \ell_{\tau'}] \leq \beta'$, the second by the definition of $\beta = \frac{\beta'^2}{2}$, $\alpha = \frac{\beta'}{2}$, the third again, by $\mathbb{P}[i \text{ wins in } [\ell_{\tau'-1}, \ell_{\tau'}] | M_{\ell_{\tau'}}] \leq \beta'$, the fourth from $2\beta' < \frac{1}{2}$, the fifth and sixth from basic algebra, the seventh by the bound on $k \leq \frac{16Ln}{\beta\gamma}$, by Lemma 4.3.2, the eighth by $\beta' = \frac{\beta}{16}$, and the ninth by $\beta = \frac{\varepsilon\gamma}{32nL}$.

The sample complexity bound and failure probability follow from Lemmas 4.3.2 and 4.3.1, substituting in for various parameters, since I_{win} is called k times. Thus, in total, there are $\leq 3k \log(k) + 3k$ empirical estimates made, each with probability at most δ' of failure, each with sample size T . \blacksquare

Proof of Lemma 4.3.1. We start by showing that, with no sampling error, the calculation $p_{x,y}^i$ we do is equivalent to $q_{x,y}^i = \mathbb{P}[b_i \in [x, y] \wedge b_i > \max_{j \neq i} b_j | \max_j b_j < y]$. When $x = y$, we will denote this simply as q_x^i (similarly, p_x^i). Similarly, let q_x^0 denote the probability that no one wins when the reserve bidder is set to bid x (and p_x^0 the empirical probability therein).

By definition,

$$\begin{aligned}
q_{\ell_\tau, \ell_{\tau+1}}^i &= \mathbb{P}[b_i \in [\ell_\tau, \ell_{\tau+1}] \wedge b_i > \max_{j \neq i} b_j \mid \max_j b_j < \ell_{\tau+1}] \\
&= \frac{\mathbb{P}[b_i \in [\ell_\tau, \ell_{\tau+1}] \wedge b_i > \max_{j \neq i} b_j \wedge \max_j b_j < \ell_{\tau+1}]}{\mathbb{P}[\max_j b_j < \ell_{\tau+1}]} \\
&= \frac{\mathbb{P}[b_i \in [\ell_\tau, \ell_{\tau+1}] \wedge b_i > \max_{j \neq i} b_j]}{\mathbb{P}[\max_j b_j < \ell_{\tau+1}]} && \text{(i winning in } [\ell_\tau, \ell_{\tau+1}] \text{ implies } \max_j b_j < \ell_{\tau+1}) \\
&= \frac{\mathbb{P}[b_i \geq \ell_\tau \wedge b_i > \max_{j \neq i} b_j] - \mathbb{P}[b_i \geq \ell_{\tau+1} \wedge b_i > \max_{j \neq i} b_j]}{\mathbb{P}[\max_j b_j < \ell_{\tau+1}]} \\
&= \frac{\mathbb{P}[i \text{ wins with reserve } \ell_\tau] - \mathbb{P}[i \text{ wins with reserve } \ell_{\tau+1}]}{\mathbb{P}[\max_j b_j < \ell_{\tau+1}]} && \text{(Assuming no point masses, there are no ties)} \\
&= \frac{q_{\ell_\tau, 1}^i - q_{\ell_{\tau+1}, 1}^i}{q_{\ell_{\tau+1}}^0}
\end{aligned}$$

The final form is identical to the estimated quantity used by I_{win} . It now suffices to now show that each of the three samples give us good estimates of their respective true probabilities. A basic Chernoff bound implies

$$\mathbb{P}[|p_{x,1}^i - q_{x,1}^i| \geq \frac{\beta\gamma(1-\alpha)}{4}] \leq 2e^{-T \frac{1}{8} t_1 \beta^2 \gamma^2 (1-\alpha)^2}.$$

Substituting $T = \frac{8 \ln 6 / \delta'}{\alpha^2 \gamma^2 (\frac{\mu}{2})^2}$, and noting $\alpha < 1 - \alpha$, we have

$$\mathbb{P}[|p_{x,1}^i - q_{x,1}^i| \geq \frac{\beta\gamma(1-\alpha)}{4}] \leq \delta'$$

for each of $x = \ell_\tau, \ell_{\tau+1}$. Similarly,

$$\mathbb{P}[|p_x^0 - q_x^0| > \frac{\alpha\gamma}{2}] \leq 2e^{-\frac{T}{2} \alpha^2 \gamma^2}$$

and substituting for T , we have that $|p_{\ell_{\tau+1}}^0 - q_{\ell_{\tau+1}}^0| \geq \frac{\alpha\gamma}{2}$ with probability at most δ' . Thus, using a union bound, we have that with probability at least $1 - 3\delta'$, for a particular t ,

$$\frac{q_{\ell_\tau, 1}^i - q_{\ell_{\tau+1}, 1}^i - \frac{\beta\gamma(1-\alpha)}{2}}{q_{\ell_{\tau+1}}^0 + \frac{\alpha\gamma}{2}} \leq \frac{p_{\ell_\tau, 1}^i - p_{\ell_{\tau+1}, 1}^i}{p_{\ell_{\tau+1}}^0} \leq \frac{q_{\ell_\tau, 1}^i - q_{\ell_{\tau+1}, 1}^i + \frac{\beta\gamma(1-\alpha)}{2}}{q_{\ell_{\tau+1}}^0 - \frac{\alpha\gamma}{2}} \quad (\text{A.3})$$

Now, it suffices to show that Equation (A.3) implies the relative error stated previously. By assumption, $p_{0, \ell_{\tau+1}}^i > \gamma$. This implies that the probability everyone bids at most $\ell_{\tau+1}$ is at least γ (for a winning bid of $\ell_{\tau+1}$ to win, all bids must be at most $\ell_{\tau+1}$), so

$$q_{\ell_{\tau+1}}^0 \geq \gamma. \quad (\text{A.4})$$

Then,

$$\begin{aligned}
p_{\ell_\tau, \ell_{\tau+1}}^i &= \frac{p_{\ell_\tau, 1}^i - p_{\ell_{\tau+1}, 1}^i}{p_{\ell_{\tau+1}}^0} \\
&\geq \frac{q_{\ell_\tau, 1}^i - q_{\ell_{\tau+1}, 1}^i - \frac{1}{2}\beta\gamma(1-\alpha)}{q_{\ell_{\tau+1}}^0 + \frac{\alpha\gamma}{2}} \\
&\geq \frac{q_{\ell_\tau, 1}^i - q_{\ell_{\tau+1}, 1}^i - \beta\gamma}{q_{\ell_{\tau+1}}^0 + \alpha\gamma} && \text{(Since } \frac{(1-\alpha)}{2} < 1) \\
&\geq \frac{q_{\ell_\tau, 1}^i - q_{\ell_{\tau+1}, 1}^i - \beta\gamma}{q_{\ell_{\tau+1}}^0 + \alpha q_{\ell_{\tau+1}}^0} && \text{(By Eq. (A.4))} \\
&= \frac{q_{\ell_\tau, 1}^i - q_{\ell_{\tau+1}, 1}^i - \beta\gamma}{q_{\ell_{\tau+1}}^0(1+\alpha)} \\
&= \frac{q_{\ell_\tau, \ell_{\tau+1}}^i}{1+\alpha} - \frac{\beta\gamma}{q_{\ell_{\tau+1}}^0(1+\alpha)} \\
&\geq \frac{q_{\ell_\tau, \ell_{\tau+1}}^i}{1+\alpha} - \frac{\beta}{(1+\alpha)} && \text{(By Eq. (A.4))} \\
&\geq \frac{q_{\ell_\tau, \ell_{\tau+1}}^i}{1+\alpha} - \beta \\
&\geq (1-\alpha)q_{\ell_\tau, \ell_{\tau+1}}^i - \beta && \text{(By Fact A.1.3)}
\end{aligned}$$

Now, we prove the upper bound on our estimator.

$$\begin{aligned}
p_{\ell_\tau, \ell_{\tau+1}}^i &= \frac{p_{\ell_\tau, 1}^i - p_{\ell_{\tau+1}, 1}^i}{p_{\ell_{\tau+1}}^0} \\
&\leq \frac{q_{\ell_\tau, 1}^i - q_{\ell_{\tau+1}, 1}^i + \frac{(1-\alpha)\beta\gamma}{2}}{q_{\ell_{\tau+1}}^0 - \frac{\alpha\gamma}{2}} \\
&\leq \frac{q_{\ell_\tau, 1}^i - q_{\ell_{\tau+1}, 1}^i + (1-\alpha)\beta\gamma}{q_{\ell_{\tau+1}}^0 - \frac{\alpha\gamma}{2}} \\
&\leq \frac{q_{\ell_\tau, 1}^i - q_{\ell_{\tau+1}, 1}^i + (1-\alpha)\beta\gamma}{q_{\ell_{\tau+1}}^0 - \frac{\alpha q_{\ell_{\tau+1}, \ell_{\tau+1}}^0}{2}} && \text{(By Eq. (A.4))} \\
&\leq \frac{q_{\ell_\tau, \ell_{\tau+1}}^i}{1 - \frac{\alpha}{2}} + \frac{(1-\alpha)\beta\gamma}{q_{\ell_{\tau+1}}^0 (1 - \frac{\alpha}{2})} \\
&\leq \frac{q_{\ell_\tau, \ell_{\tau+1}}^i}{1 - \frac{\alpha}{2}} + \frac{\beta\gamma}{q_{\ell_{\tau+1}}^0} \\
&\leq \frac{q_{\ell_\tau, \ell_{\tau+1}}^i}{1 - \frac{\alpha}{2}} + \beta && \text{(By Eq. (A.4))} \\
&\leq (1 + 2\frac{\alpha}{2})q_{\ell_\tau, \ell_{\tau+1}}^i + \beta && \text{(By Fact. A.1.3)} \\
&= (1 + \alpha)q_{\ell_\tau, \ell_{\tau+1}}^i + \beta
\end{aligned}$$

Thus, both the upper and lower bounds on the estimator hold with probability $1 - \delta$. \blacksquare

Proof of Lemma 4.3.2. We will show each of the three parts to be true.

1. We start by proving that `Intervals` will output a partition with at most $\frac{24nL}{\beta\gamma}$ intervals.

We claim that each interval is at least $\frac{\beta\gamma}{24nL}$ in length, implying the above bound on the total number of intervals.

Consider some current upper bound for an interval $\ell_{\tau+1}$. If `Intervals` accepts some point ℓ_τ such that $\ell_{\tau+1} - \ell_\tau \geq \frac{\beta\gamma}{24nL}$, then the bound trivially holds.

If this does not hold, `Intervals` tests some point $\widehat{\ell}_\tau$ such that

$$\frac{\beta\gamma}{24nL} \geq \ell_{\tau+1} - \widehat{\ell}_\tau \geq \frac{\beta\gamma}{48nL}$$

since it is doing binary search. We claim `Intervals` will accept $\widehat{\ell}_\tau$; if this is the case, the interval will have length at least $\frac{\beta\gamma}{48nL}$. Notice that

$$\mathbb{P}[\max_j b_j \in [\widehat{\ell}_\tau, \ell_{\tau+1}] \mid \max_j b_j \leq \ell_{\tau+1}] \leq \mathbb{P}[\max_j b_j \in [\ell_{\tau+1} - \frac{\beta\gamma}{24nL}, \ell_{\tau+1}] \mid \max_j b_j \leq \ell_{\tau+1}]$$

so it will suffice to show that `Intervals` would accept the smallest possible value of $\widehat{\ell}_\tau$ (since that region will have the most probability mass). We bound the ratio, for a given $\ell_{\tau+1}$ such that

$$\mathbb{P}[\max_j b_j \in [\ell_{\tau+1} - \frac{\beta\gamma}{24nL}, \ell_{\tau+1}] | \max_j b_j \leq \ell_{\tau+1}] = \frac{\mathbb{P}[\max_j b_j \leq \ell_{\tau+1} - \frac{\beta\gamma}{24nL}]}{\mathbb{P}[\max_j b_j \leq \ell_{\tau+1}]}$$

for some upper point of an interval $\ell_{\tau+1}$ such that $\mathbb{P}[i \text{ wins with a bid } \leq \ell_{\tau+1}] \geq \gamma$. Since F_j is L -Lipschitz for all j ,

$$\mathbb{P}[b_j \leq \ell_{\tau+1}] - \mathbb{P}[b_j \leq \ell_{\tau+1} - \frac{\beta\gamma}{24nL}] \leq L \frac{\beta\gamma}{24Ln} = \frac{\beta\gamma}{24n}.$$

Then, by summing this probability over all n bidders, we have

$$\mathbb{P}[\max_j b_j \leq \ell_{\tau+1}] - \mathbb{P}[\max_j b_j \leq \ell_{\tau+1} - \frac{\beta\gamma}{24nL}] \leq \frac{\beta\gamma}{24}.$$

Rearranging terms, we have

$$\frac{\mathbb{P}[\max_j b_j \leq \ell_{\tau+1} - \frac{\beta\gamma}{24nL}]}{\mathbb{P}[\max_j b_j \leq \ell_{\tau+1}]} \geq 1 - \frac{\beta'\gamma}{\mathbb{P}[\max_j b_j \leq \ell_{\tau+1}]} \geq 1 - \frac{\beta}{24}$$

where the last inequality came from the fact that $\mathbb{P}[i \text{ wins with a bid } \leq \ell_{\tau+1}] \geq \mathbb{P}[\max_j b_j \leq \ell_{\tau+1}] \geq \gamma$. So, `Intervals` will accept $\widehat{\ell}_\tau$ as ℓ_τ , so long as the empirical estimate of `Inside` is correct up to $\beta + \alpha = \frac{\beta}{48}$, which is the case by Corollary A.1.2 with probability $1 - 3\delta'$.

2. We now need to show

$$\mathbb{P}[\max_j b_j \geq \ell_{\tau-1} | \max_j b_j \leq \ell_\tau] \leq \frac{\beta}{16}$$

holds for the lattice points $t > 3$. Since $\mathbb{P}[\max_j b_j \leq \ell_3] \geq \gamma$, by Corollary A.1.2, the accuracy guarantee holds with probability $1 - 3\delta'$ for a fixed t (since $\beta = \frac{\beta^2}{96}$, $\alpha = \frac{\beta}{96}$, and the condition by which $\ell_{\tau-1}$ was accepted was that the empirical estimate of the above quantity was at most $\frac{\beta}{24}$). Thus, with probability $1 - 3k\delta'$, the above holds for all $t > 3$.

3. We begin by showing $\mathbb{P}[\max_j b_j \leq \ell_2] \leq \gamma$ with probability at least $1 - \delta'$. The condition for stopping the search for new interval points is

$$J = \frac{\sum_{t \in S_1} \mathbb{I}[i \text{ wins on sample } t]}{T} \leq \frac{\gamma}{2}$$

where S_1 is a random sample of size T with reserve ℓ_1 . A basic Chernoff bound shows that

$$\mathbb{P}[|J - \mathbb{P}[\max_j b_j \leq \ell_1]| \geq \frac{\gamma}{2}] \leq 2e^{-\frac{T\gamma^2}{2}}$$

which, for $T = \frac{32 \ln \frac{6}{\beta^2 \gamma^2 \delta'}}{\beta^2 \gamma^2 \delta'}$ is at most δ' , so $\mathbb{P}_{S_1}[\mathbb{P}[\max_j b_j \leq \ell_2] \leq \gamma] \geq 1 - \delta'$, as desired.

It remains to sum up the total error probability and sample complexity. The lower bound on the length of each interval also implies a bound on the total number of empirical estimates made to find a fixed ℓ_τ . Formally, the halving algorithm beginning with a search space of size $\ell_{\tau+1} \leq 1$ will halt before the remaining search space has shrunk to $\frac{\beta\gamma}{48Ln}$, which will take at most

$\log \frac{48Ln}{\beta\gamma} = \log(k)$ attempted interval endpoints per accepted interval endpoint. Each of these attempts calls `Inside`, which takes 2 estimates. For each accepted interval, an estimate of the remaining probability mass is done. Thus, in total, there are $2k \log(k) + k$ estimates done by `Intervals`. Each fails with probability at most δ' , so `Intervals` succeeds with probability at least $1 - 3k \log(k)\delta'$ and uses at most $3k \log(k)T$ samples. ■

Appendix B

Privacy-Preserving Public Information Appendix

B.0.1 Undominated strategic play with Empty Counters: Lower bounds

Theorem 6.3.2. *There exist games g whose $\text{CR}_{\text{UNDOM}}(\mathcal{M}_\emptyset, g)$ cannot be bounded by any function of n .*

Proof. Let g be the following game. For each player i , there is a resource r_i such that $v_{r_i}(1) = H$ but $v_{r_i}(> 1) = 0$. Furthermore, let there be some other resource r such that $v_r(1) = 1$. Let A_i contain 2 allowable actions: selecting r_i and selecting r .

OPT in this setting would have each player select r_i , which has $\text{SW}(\text{OPT}) = nH$. On the other hand, we claim it is undominated for each player to select r instead (call this joint action a). If each player were to have a “twin”, then r_i could have already been selected by another player so that i would get more utility from r than r_i . Then, this undominated strategy a has $\text{SW}(a) = n$. Thus, we have a game g for which

$$\text{CR}_{\text{UNDOM}}(g) \geq \frac{nH}{n} = H$$

which, as $H \rightarrow \infty$ is unbounded. ■

The negative result above isn't particularly surprising: if there is some coordination to be done, but there is no coordinator and no information about the target, all is lost. On the other hand, our positive result for undominated strategies (Theorem 6.3.8) in the case of private information relies on a very particular rate of decay of the resources' value. Theorem 6.3.3 show that, even under this stylized assumption where all resources' values shrink slowly, a total lack of information can lead to very poor behaviour in undominated strategies.

Theorem 6.3.3. *There exists g such that $\text{CR}_{\text{UNDOM}}(\mathcal{M}_\emptyset, g) \geq \Omega(\frac{n^2}{\log(n)})$, when $\mathcal{W}_r(t) = \frac{\mathcal{W}_r(0)}{t}$.*

Proof of Theorem 6.3.3. For each player i , let r_i be a resource where $v_{r_i}(1) = n$ (note that this uniquely determines $v_i(c)$ for all c). Let there be another resource r such that $v_r(1) = 1$. Let each A_i contain all resources. Since $\frac{v_{r_i}(1)}{n} = 1$, it is not dominated for player i to select r . Let a denote the joint strategy where each player selects resource r . Thus the social welfare attained

by this strategy profile is $O(\log(n))$, where as the optimal social welfare is n^2 , implying that $\text{CR}_{\text{UNDOM}} \geq \Omega(\frac{n^2}{\log(n)})$. \blacksquare

B.0.2 Omitted proofs for Undominated strategies with Privacy-preserving counters

Proof of Theorem 6.3.8. Consider the optimal allocation and let r_i and z_i denote that the z_i^{th} copy of resource r_i got allocated to player i under the optimal allocation. Now consider any run of the game under undominated strategic play and based on the run, partition all the players into two groups. Group A consists of players i such that $x_{i,r_i} \leq z_i$ (i.e., the copy (or a more valuable copy) of the resource that was allocated to player i was present when the player arrived) and group B consists of all other players.

For the player in group B , the copy of the resource that they received in the optimal allocation was already allocated by the time they arrived in the run of the undominated strategic play. Hence, the total social welfare achieved by the undominated strategic play is at least as much the welfare achieved by group B player under optimal allocation.

Now consider any player i in group A . We show that the resource picked by player i under undominated strategic play gets her a reasonable fraction of the value she would have received under optimal allocation. For any resource r , given the displayed counter value of $y_{i,r}$, by the guarantees of the (α, β) -accuracy guarantee of the counters, we directly argue about the possible range of the consistent beliefs or estimates $\widehat{x}_{i,r}$ by which player i can make her choice.

Specifically, by the bounds on (α, β) -counters, for a given true value x , it must be the case that all announcements $y_{i,r}$ satisfy:

$$\alpha x_{i,r} + \beta \geq y_{i,r} \geq \frac{1}{\alpha} x_{i,r} - \beta$$

Rearranging, we have $y_{i,r} \in [\frac{1}{\alpha} x_{i,r} - \beta, \alpha x_{i,r} + \beta]$. Suppose these bounds are realized; we wish to upper and lower bound $\widehat{x}_{i,r}$ as a function of these announcement values. By the quality of the announcement, we have that $\alpha \widehat{x}_{i,r} + \beta \geq y_{i,r} \geq \frac{1}{\alpha} x_{i,r} - \beta$.

We can similarly upper bound $\widehat{x}_{i,r}$, e.g. $\alpha \widehat{x}_{i,r} + \beta \geq y_{i,r} \geq \frac{1}{\alpha} \widehat{x}_{i,r} - \beta$, which, by the fact that the true count is at least 0, implies $\widehat{x}_{i,r} \in [\max\{0, \frac{x_{i,r}}{\alpha^2} - \frac{2\beta}{\alpha}\}, \alpha^2 x_{i,r} + 2\alpha\beta]$.

Now, suppose player i chose resource r' which was undominated and not r_i which he received in the optimal allocation. Since resource r' is undominated:

$$\mathcal{W}_{r'}(\widehat{x}_{i,r'}) \geq \mathcal{W}_{r_i}(\widehat{x}_{i,r_i}) \tag{B.1}$$

We have

$$\mathcal{W}_{r'}(\widehat{x}_{i,r'}) \leq \mathcal{W}_{r'}\left(\max\left\{0, \frac{x_{i,r'}}{\alpha^2}\right\} - \frac{2\beta}{\alpha}\right) \leq \psi(\alpha, \beta) \mathcal{W}_{r'}(x_{i,r'}) \tag{B.2}$$

where the first inequality came from the lower bound on the counter, and the fact that the valuations are decreasing, and the second from the assumption about \mathcal{W}_r on x and its lower bound. Similarly, we know for each r that

$$\mathcal{W}_{r_i}(\widehat{x}_{i,r_i}) \geq \mathcal{W}_{r_i}(\alpha^2 x_{i,r_i} + 2\alpha\beta) \geq \frac{\mathcal{W}_{r_i}(x_{i,r_i})}{\phi(\alpha, \beta)} \tag{B.3}$$

Combining the three equations above, we have the actual value received by the player i on choosing resource r' , $\mathcal{W}_{r'}(x_{i,r'})$ is at least $\frac{1}{\psi(\alpha,\beta)\phi(\alpha,\beta)}$ fraction of the value $\mathcal{W}_{r_i}(x_{i,r_i})$ that he would receive under the optimal allocation. Therefore, by virtue of partition of the players in groups A and B , we have that social welfare achieved under undominated strategic play is at least $\frac{1}{1+\psi(\alpha,\beta)\phi(\alpha,\beta)}$ fraction of the optimal social welfare. ■

B.1 Future-dependent Utilities

Theorem B.1.1. *There exist sequential resource-sharing games g , where each resource r 's value curve \mathcal{W}_r is (w, n) – shallow, such that in the future-dependent setting, $\text{CR}_{\text{GREEDY}}(\mathcal{M}_{\text{Full}}, g) \geq 2w$.*

Proof. Consider two players and two resources r, r' . Let r have a value curve which is a step function, with $v_r(0) = w$, $v_r(1) = \frac{1}{2}$ and $v_{r'}(0) = w - \varepsilon$. Suppose player one has access to both resources, the other having only resource r as an option. Then, player one will choose r according to greedy, and player two will always select r . The social welfare will be $\text{SW}(\text{GREEDY}) = 1$, whereas OPT is for player 1 to take r' and will have $\text{SW}(\text{OPT}) = 2w - \varepsilon$. As $\varepsilon \rightarrow 0$, this ratio approaches $2w$. ■

B.2 Analysis of Private Counters

Proof of Lemma 6.6.1. We assume the reader is familiar with the TreeSum mechanism. The privacy of this construction follows the same argument as for the original constructions. One can view m independent copies of the TreeSum protocol as a single protocol where the Laplace mechanism is used to release the entire vector of partial sums. Because the ℓ_1 -sensitivity of each partial sum is 1 (since $\|a_t\| \leq 1$), the amount of Laplace noise (per entry) needed to release the m -dimensional vector partial sums case is the same as for a dimensional 1-dimensional counter.

To see why the approximation claims holds, we can apply Lemma 2.8 from [34] (a tail bound for sums of independent Laplace random variables) with $b_1 = \dots = b_{\log n} = \log n / \varepsilon$, error probability $\delta = \gamma / mn$, $\nu = \frac{(\log n) \sqrt{\log(1/\delta)}}{\varepsilon}$ and $\lambda = \frac{(\log n)(\log(1/\delta))}{\varepsilon}$, we get that each individual counter estimate $s_t(j)$ has additive error $O\left(\frac{(\log n)(\log(nm/\gamma))}{\varepsilon}\right)$ with probability at least $1 - \gamma/(mn)$. Thus, all $n \cdot m$ estimates satisfy the bound simultaneously with probability at least $1 - \gamma$. ■

Proof of Lemma 6.6.2. We begin with the proof of privacy. The first phase of the protocol is $\varepsilon/2$ -differentially private because it is an instance of the “sparse vector” technique of Hardt and Rothblum [67] (see also [114, Lecture 20] for a self-contained exposition). The second phase of the protocol is $\varepsilon/2$ -differentially private by the privacy of TreeSum. Since differential privacy composes, the scheme as a whole is ε -differentially private. Note that since we are proving $(\varepsilon, 0)$ -differential privacy, it suffices to consider nonadaptive streams; the adaptive privacy definition then follows [50].

We turn to proving the approximation guarantee. Note that the each of the Laplace noise variables added in phase 1 of the algorithm (to compute $\widetilde{x}_{t,r}$ and τ_j) uses parameter $2/\varepsilon'$. Taking

a union bound over the mn possible times that such noise is added, we see that with probability at least $1 - \gamma/2$, each of these random variables has absolute value at most $O(\frac{\log(mn/\gamma)}{\varepsilon'})$. Since $\frac{2}{\varepsilon'} = O(\frac{mk}{\varepsilon})$ and $k = O(\log \log(\frac{nm}{\gamma}) + \log \frac{1}{\varepsilon})$, we get that each of these noise variables has absolute value $\tilde{O}_\alpha(\frac{m \log(mn/\gamma)}{\varepsilon})$ with probability all but $\gamma/2$. We denote this bound E_1 .

Thus, for each counter, the i -th flag is raised no earlier than when the value of the counter first exceeds $\alpha^i(\log n) - E_1$, and no later than when the counter first exceeds $\alpha^i(\log n) + E_1$. The very first flag might be raised when counter has value 0. In that case, the additive error of the estimate is $\log n$, which is less than E_1 . Hence, the mechanism's estimates during the first phase provide an $(\alpha, E_1, \gamma/2)$ -approximation (as desired).

The flag that causes the algorithm to enter the second phase is supposed to be raised when the counter takes the value $A := \alpha^k(\log n) \geq \frac{\alpha}{\alpha-1} \cdot C_{tree} \cdot \frac{\log(nm/\gamma)}{\varepsilon}$; in fact, the counter could be as small as $A - E_1$. After that point, the additive error is due to the TreeSum protocol and is at most $B := C_{tree} \cdot \log(n) \cdot \log(nm/\gamma)/\varepsilon$ (with probability at least $1 - \gamma/2$) by Lemma 6.6.1. The reported value $y_{i,r}$ thus satisfies

$$y_{i,r} \geq x_{i,r} - B = \frac{1}{\alpha}x_{i,r} + \underbrace{\left(1 - \frac{1}{\alpha}\right)x_{i,r} - B}_{\text{residual error}} .$$

Since $x_{i,r} \geq A - E_1$, the “residual error” in the equation above is at least $(1 - \frac{1}{\alpha})(A - E_1) - B = -(1 - \frac{1}{\alpha})E_1 \geq -E_1$. Thus, the second phase of the algorithm also provides $(\alpha, E_1, \gamma/2)$ -approximation. With probability $1 - \gamma$, both phases jointly provide a (α, E_1, γ) -approximation, as desired. ■

Bibliography

- [1] Atila Abdulkadiroglu, Parag A. Pathak, and Alvin E. Roth. Strategy-proofness versus efficiency in matching with indifferences: Redesigning the nyc high school match. *American Economic Review*, 99(5):1954–78, 2009. doi: 10.1257/aer.99.5.1954. URL <http://www.aeaweb.org/articles.php?doi=10.1257/aer.99.5.1954>.
- [2] Zoe Abrams, Arpita Ghosh, and Erik Vee. Cost of conciseness in sponsored search auctions. In *Internet and Network Economics*, pages 326–334. Springer, 2007.
- [3] Jonathan Alevy, Oscar Cristi, and Oscar Melo. Right-to-choose auctions: A field study of water markets in the limari valley of chile. Framed field experiments, The Field Experiments Website, 2010.
- [4] Martin Anthony and Peter L. Bartlett. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, New York, NY, USA, 1st edition, 1999.
- [5] Orley Ashenfelter and David Genesove. Testing for price anomalies in real-estate auctions. *American Economic Review*, 82(2):501–05, 1992. URL <http://EconPapers.repec.org/RePEc:aea:aecrev:v:82:y:1992:i:2:p:501-05>.
- [6] Itai Ashlagi, Yashodhan Kanoria, and Jacob D Leshno. Unbalanced random matching markets. In *EC*, pages 27–28, 2013.
- [7] Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E. Schapire. The nonstochastic multiarmed bandit problem. *SIAM J. Comput.*, 32(1):48–77, January 2003. ISSN 0097-5397.
- [8] Eduardo Azevedo and Eric Budish. Strategyproofness in the large. Technical report, Working paper, 2012.
- [9] Moshe Babaioff, Liad Blumrosen, Shaddin Dughmi, and Yaron Singer. Posting prices with unknown distributions. In *Innovations in Computer Science (ICS)*. Tsinghua University Press, January 2011.
- [10] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. A simple and approximately optimal mechanism for an additive buyer. *SIGecom Exch.*, 13(2):31–35, January 2015.
- [11] Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. Single price mechanisms for revenue maximization in unlimited supply combinatorial auctions. Technical report, Carnegie Mellon University, 2007.
- [12] Maria-Florina Balcan, Avrim Blum, Jason D Hartline, and Yishay Mansour. Reducing

- mechanism design to algorithm design via machine learning. *Journal of Computer and System Sciences*, 74(8):1245–1270, 2008.
- [13] Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. Item pricing for revenue maximization. In *Proceedings of the 9th ACM conference on Electronic commerce*, pages 50–59. ACM, 2008.
- [14] Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. The price of uncertainty. In *Proceedings of the 10th ACM Conference on Electronic Commerce, EC '09*, pages 285–294, 2009.
- [15] Michael Benisch, Norman Sadeh, and Tuomas Sandholm. A theory of expressiveness in mechanisms. In *Proceedings of the 23rd National Conference on Artificial Intelligence - Volume 1, AAAI'08*, pages 17–23. AAAI Press, 2008.
- [16] Khsipra Bhawalkar and Tim Roughgarden. Welfare guarantees for combinatorial auctions with item bidding. In *SODA*, 2011.
- [17] Martin Bichler, Stefan Schneider, Kemal Guler, and Mehmet Sayal. Compact bidding languages and supplier selection for markets with economies of scale and scope. *European Journal of Operational Research*, 214(1):67–77, 2011.
- [18] Avrim Blum and Jason D. Hartline. Near-optimal online auctions. In *SODA*, pages 1156–1163, Vancouver, BC, Canada., 2005. SIAM.
- [19] Avrim Blum and Yishay Mansour. From external to internal regret. *J. Mach. Learn. Res.*, 8:1307–1324, December 2007. ISSN 1532-4435. URL <http://dl.acm.org/citation.cfm?id=1314498.1314543>.
- [20] Avrim Blum, Vijay Kumar, Atri Rudra, and Felix Wu. Online learning in online auctions. In *Symposium on Discrete Algorithms*, Baltimore, MD, USA., 2003. ACM.
- [21] Avrim Blum, Yishay Mansour, and Jamie Morgenstern. Learning valuation distributions from partial observation. In *AAAI 2015, Papers from the 2005 AAAI Spring Symposium*. AAAI, AAAI Press, 2015.
- [22] Avrim Blum, Yishay Mansour, and Jamie Morgenstern. Learning valuation distributions from partial observation. In *AAAI*, page Forthcoming, Austin, Texas, 2015. AAAI Press.
- [23] Avrim Blum, Jamie Morgenstern, Ankit Sharma, and Adam Smith. Privacy-preserving public information for sequential games. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science, ITCS '15*, pages 173–180, New York, NY, USA, 2015. ACM.
- [24] Liad Blumrosen and Michal Feldman. Implementation with a bounded action space. In *Proceedings of the 7th ACM conference on Electronic commerce*, pages 62–71. ACM, 2006.
- [25] Liad Blumrosen and Noam Nisan. Auctions with severely bounded communication. In *Foundations of Computer Science, 2002. Proceedings. The 43rd Annual IEEE Symposium on*, pages 406–415. IEEE, 2002.
- [26] Craig Boutilier and Holger H Hoos. Bidding languages for combinatorial auctions. In *International Joint Conference on Artificial Intelligence*, volume 17, pages 1211–1217,

2001.

- [27] Mark Braverman, Jieming Mao, and S. Matthew Weinberg. Personal communication, 2015.
- [28] Jeremy Bulow and Paul Klemperer. Auctions vs. negotiations. Technical report, National Bureau of Economic Research, 1994.
- [29] Roberto Burguet. Right to choose in oral auctions. Ufae and iae working papers, Unitat de Fonaments de l'Anàlisi Econòmica (UAB) and Institut d'Anàlisi Econòmica (CSIC), 2007.
- [30] Yang Cai and Constantinos Daskalakis. Extreme-value theorems for optimal multidimensional pricing. In *Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on*, pages 522–531, Palm Springs, CA, USA., Oct 2011. IEEE.
- [31] Yang Cai and Christos H. Papadimitriou. Simultaneous bayesian auctions and computational complexity. In *ACM Conference on Economics and Computation, EC '14, Stanford, CA, USA, June 8-12, 2014*, pages 895–910, 2014.
- [32] Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, and Maria Kyropoulou. On the efficiency of equilibria in generalized second price auctions. In *EC*, 2011.
- [33] Nicol Cesa-Bianchi, Claudio Gentile, and Yishay Mansour. Regret minimization for reserve prices in second-price auctions. In Sanjeev Khanna, editor, *SODA*, pages 1190–1204, New York, NY, USA., 2013. SIAM.
- [34] T-H Hubert Chan, Elaine Shi, and Dawn Song. Private and continual release of statistics. *ACM Transactions on Information and System Security (TISSEC)*, 14(3):26, 2011.
- [35] Shuchi Chawla, Jason D. Hartline, and Robert Kleinberg. Algorithmic pricing via virtual valuations. In *Proceedings of the 8th ACM Conference on Electronic Commerce, EC '07*, pages 243–251, New York, NY, USA, 2007. ACM.
- [36] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the Forty-second ACM Symposium on Theory of Computing, STOC '10*, pages 311–320, New York, NY, USA, 2010. ACM.
- [37] Shuchi Chawla, Jason Hartline, and Denis Nekipelov. Mechanism design for data science. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation, EC '14*, pages 711–712, New York, NY, USA, 2014. ACM.
- [38] George Christodoulou, Annamaria Kovacs, and Michael Schapira. Bayesian Combinatorial Auctions. In *ICALP '08 Proceedings of the 35th international colloquium on Automata, Languages and Programming*, 2008.
- [39] Richard Cole and Tim Roughgarden. The sample complexity of revenue maximization. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing, STOC '14*, pages 243–252, New York, NY, USA, 2014. SIAM.
- [40] Vincent Conitzer, Tuomas Sandholm, and Paolo Santi. Combinatorial auctions with k-wise dependent valuations. In *AAAI*, volume 5, pages 248–254, 2005.

- [41] Nikhil Devanur, Jason Hartline, Anna Karlin, and Thach Nguyen. Prior-independent multi-parameter mechanism design. In *Internet and Network Economics*, pages 122–133. Springer, Singapore, 2011.
- [42] Nikhil Devanur, Jamie Morgenstern, Vasilis Syrgkanis, and S. Matthew Weinberg. Simple auctions with simple strategies. In *Proceedings of the 16th ACM Conference on Economics and Computation*, EC '15, Portland, OR, USA, 2015. ACM.
- [43] Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. In *Proceedings of the 11th ACM Conference on Electronic Commerce*, EC '10, pages 129–138, New York, NY, USA, 2010. ACM.
- [44] Shahar Dobzinski. Two randomized mechanisms for combinatorial auctions. In *Proceedings of the 10th International Workshop on Approximation and the 11th International Workshop on Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 89–103, 2007.
- [45] Shahar Dobzinski, Hu Fu, and Robert Kleinberg. On the complexity of computing an equilibrium in combinatorial auctions. *SODA*, page To appear, 2015.
- [46] Shaddin Dughmi, Tim Roughgarden, and Qiqi Yan. From convex optimization to randomized mechanisms: toward optimal combinatorial auctions. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 149–158. ACM, 2011.
- [47] Shaddin Dughmi, Li Han, and Noam Nisan. Sampling and representation complexity of revenue maximization. In Tie-Yan Liu, Qi Qi, and Yinyu Ye, editors, *Web and Internet Economics*, volume 8877 of *Lecture Notes in Computer Science*, pages 277–291. Springer International Publishing, Beijing, China, 2014.
- [48] Paul Dütting, Felix Fischer, and David C Parkes. Simplicity-expressiveness tradeoffs in mechanism design. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 341–350. ACM, 2011.
- [49] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography*, pages 265–284. Springer, 2006.
- [50] Cynthia Dwork, Moni Naor, Toniann Pitassi, and Guy N Rothblum. Differential privacy under continual observation. In *Proceedings of the 42nd ACM Symposium on Theory of Computing*, pages 715–724. ACM, 2010.
- [51] Cynthia Dwork, Guy N. Rothblum, and Salil P. Vadhan. Boosting and differential privacy. In *FOCS*, pages 51–60, 2010.
- [52] Kfir Eliaz, Theo Offerman, and Andrew Schotter. Creating competition out of thin air: An experimental study of right-to-choose auctions. *Games and Economic Behavior*, 62(2): 383–416, March 2008.
- [53] Edith Elkind. Designing and learning optimal finite support auctions. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 736–745. Society for Industrial and Applied Mathematics, 2007.
- [54] M. Feldman, H. Fu, N. Gravin, and B. Lucier. Simultaneous auctions are (almost) effi-

- cient. In *STOC*, 2013.
- [55] Michal Feldman, Hu Fu, Nick Gravin, and Brendan Lucier. Simultaneous auctions are (almost) efficient. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 201–210. ACM, 2013.
- [56] Michal Feldman, Brendan Lucier, and Vasilis Syrgkanis. Limits of efficiency in sequential auctions. In *Proceedings of the 9th Workshop on Internet and Network Economics*, WINE, 2013.
- [57] Financial Crisis Inquiry Commission. *The Financial Crisis Inquiry Report: Final Report of the National Commission on the Causes of the Financial and Economic Crisis in the United States*. U.S. Government Printing Office, 2011. URL <http://fcic.law.stanford.edu/report>.
- [58] Mark D. Flood, Jonathan Katz, Stephen J. Ong, and Adam Smith. Cryptography and the economics of supervisory information: Balancing transparency and confidentiality. Working Paper #11, Office of Financial Research, US Department of Treasury, August 2013.
- [59] Antonia Foldes and Lidia Rejto. Strong uniform consistency for nonparametric survival curve estimators from randomly censored data. *The Annals of Statistics*, 9(1):122–129, 01 1981. doi: 10.1214/aos/1176345337. URL <http://dx.doi.org/10.1214/aos/1176345337>.
- [60] Dean P Foster and Rakesh V Vohra. Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 21(1):40–55, 1997.
- [61] Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, 1991.
- [62] David Gale and Lloyd S Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, pages 9–15, 1962.
- [63] Ian L. Gale and Donald B. Hausch. Bottom-fishing and declining prices in sequential auctions. *Games and Economic Behavior*, 7(3):318 – 331, 1994.
- [64] Kuzman Ganchev, Yuriy Nevmyvaka, Michael Kearns, and Jennifer Wortman Vaughan. Censored exploration and the dark pool problem. *Communications of the ACM*, 53(5): 99–107, 2010.
- [65] Arpita Ghosh and Amin Sayedi. Expressive auctions for externalities in online advertising. In *Proceedings of the 19th international conference on World wide web*, pages 371–380. ACM, 2010.
- [66] Jacob K. Goeree, Charles R. Plott, and John Wooders. Bidders’ choice auctions: Raising revenues through the right to choose. *Journal of the European Economic Association*, 2 (2-3), 2004. ISSN 1542-4774.
- [67] Moritz Hardt and Guy N. Rothblum. A multiplicative weights mechanism for privacy-preserving data analysis. In *FOCS ’10*, 2010.
- [68] Ronald M. Harstad. Auctioning the right to choose when competition persists. *Decision Analysis*, 7(1):78–85, March 2010. ISSN 1545-8490.

- [69] Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000.
- [70] Jason Hartline. Lecture 16, lecture notes on price of anarchy, northwestern university, 2013.
- [71] Jason Hartline and Anna Karlin. Profit maximization in mechanism design. In *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [72] Jason D Hartline. *Mechanism design and approximation*. Jason Hartline, Chicago, Illinois, 2015.
- [73] Jason D. Hartline and Tim Roughgarden. Simple versus optimal mechanisms. In John Chuang, Lance Fortnow, and Pearl Pu, editors, *ACM Conference on Electronic Commerce*, Stanford, CA, USA., 2009. ACM.
- [74] A. Hassidim, Haim Kaplan, Yishay Mansour, and Noam Nisan. Non-price equilibria in markets of discrete goods. In *EC'11*, 2011.
- [75] Ron Holzman, Noa Kfir-Dahav, Dov Monderer, and Moshe Tennenholtz. Bundling equilibrium in combinatorial auctions. *Games and Economic Behavior*, 47(1):104–123, 2004.
- [76] Justin Hsu, Zhiyi Huang, Aaron Roth, Tim Roughgarden, and Zhiwei Steven Wu. Private matchings and allocations. In *STOC*, pages 21–30. ACM, 2014.
- [77] Zhiyi Huang and Sampath Kannan. The exponential mechanism for social welfare: Private, truthful, and nearly optimal. In *Foundations of Computer Science (FOCS), 2012 IEEE 53rd Annual Symposium on*, pages 140–149. IEEE, 2012.
- [78] Zhiyi Huang, Yishay Mansour, and Tim Roughgarden. Making the most of your samples. *CoRR*, abs/1407.2479:1–3, 2014. URL <http://arxiv.org/abs/1407.2479>. Submitted to EC '15.
- [79] Nicole Immorlica and Mohammad Mahdian. Marriage, honesty, and stability. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 53–62. Society for Industrial and Applied Mathematics, 2005.
- [80] Jayant Kalagnanam and David C Parkes. Auctions, bidding and exchange design. In *Handbook of Quantitative Supply Chain Analysis*, pages 143–212. Springer, 2004.
- [81] Sampath Kannan, Jamie Morgenstern, Aaron Roth, and Zhiwei Steven Wu. *Approximately Stable, School Optimal, and Student-Truthful Many-to-One Matchings (via Differential Privacy)*, chapter 126, pages 1890–1903. ACM-SIAM, 2015.
- [82] Edward L Kaplan and Paul Meier. Nonparametric estimation from incomplete observations. *Journal of the American statistical association*, 53(282):457–481, 1958.
- [83] Richard M. Karp, Umesh V. Vazirani, and Vijay V. Vazirani. An optimal algorithm for on-line bipartite matching. In *STOC '90*, pages 352–358, 1990.
- [84] Michael Kearns, Mallesh Pai, Aaron Roth, and Jonathan Ullman. Mechanism design in large games: Incentives and privacy. In *Proceedings of the 5th conference on Innovations in theoretical computer science*, pages 403–410. ACM, 2014.
- [85] Michael J Kearns and Robert E Schapire. Efficient distribution-free learning of proba-

- bilistic concepts. In *Foundations of Computer Science, 1990. Proceedings., 31st Annual Symposium on*, pages 382–391. IEEE, 1990.
- [86] Michael J. Kearns and Umesh V. Vazirani. *An Introduction to Computational Learning Theory*. MIT Press, Cambridge, MA, 1994.
- [87] Robert D. Kleinberg and Frank Thomson Leighton. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In *FOCS*, pages 594–605, New York, New York, USA., 2003. IEEE Computer Society.
- [88] Fuhito Kojima and Parag A Pathak. Incentives and stability in large two-sided matching markets. *The American Economic Review*, pages 608–627, 2009.
- [89] Piotr Krysta and Berthold Vöcking. Online mechanism design (randomized rounding on the fly). In *Automata, Languages, and Programming*, pages 636–647. Springer, 2012.
- [90] SangMok Lee. Incentive compatibility of large centralized matching markets. *Job market paper*, 2011.
- [91] Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. In *EC*, 2001.
- [92] B. Lucier and A. Borodin. Price of anarchy for greedy auctions. In *SODA*, 2010.
- [93] Brendan Lucier and Renato Paes Leme. GSP auctions with correlated types. In *EC*, 2011.
- [94] Eric Maskin and John Riley. Monopoly with incomplete information. *The RAND Journal of Economics*, 15(2):171–196, 1984.
- [95] Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *Foundations of Computer Science, 2007. FOCS'07. 48th Annual IEEE Symposium on*, pages 94–103. IEEE, 2007.
- [96] Andres Munoz Medina and Mehryar Mohri. Learning theory and algorithms for revenue optimization in second price auctions with reserve. In *Proceedings of The 31st International Conference on Machine Learning*, pages 262–270, 2014.
- [97] Isaac Meilijson. Estimating the lifetime distribution of the parts from the autopsy statistics of the machine. *Journal of Applied Probability*, 18(4):829–838, 1981.
- [98] Paul Milgrom. Simplified mechanisms with an application to sponsored-search auctions. *Games and Economic Behavior*, 70(1):62–70, 2010.
- [99] Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1): 58–73, 1981.
- [100] Arthur Nadas. On estimating the distribution of a random vector when only the smallest coordinate is observable. *Technometrics*, 12(4):923–924, 1970.
- [101] Noam Nisan and Amir Ronen. Computationally feasible vcg mechanisms. *J. Artif. Intell. Res.(JAIR)*, 29:19–47, 2007.
- [102] Noam Nisan and Ilya Segal. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory*, 129(1):192–224, 2006.
- [103] Kobbi Nissim, Claudio Orlandi, and Rann Smorodinsky. Privacy-aware mechanism design. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 774–

789. ACM, 2012.

- [104] Mikhail V. Oet, Timothy Bianco, Dieter Gramlich, and Stephen J. Ong. Safe: An early warning system for systemic banking risk. Working Paper 11-29, Federal Reserve Bank of Cleveland, 2011. URL <http://www.clevelandfed.org/research/workpaper/2011/wp1129.pdf>.
- [105] Mikhail V. Oet, Timothy Bianco, Dieter Gramlich, and Stephen J. Ong. Financial stress index: A lens for supervising the financial system. Working Paper 12-37, Federal Reserve Bank of Cleveland, 2012. URL <http://www.clevelandfed.org/research/workpaper/2012/wp1237.pdf>.
- [106] Renato Paes Leme and Eva Tardos. Pure and bayes-nash price of anarchy for generalized second price auction. In *FOCS*, 2010.
- [107] Renato Paes Leme, Vasilis Syrgkanis, and Éva Tardos. Sequential auctions and externalities. In *SODA*, 2012.
- [108] Mallesh M. Pai and Aaron Roth. Privacy and mechanism design. *CoRR*, abs/1306.2083, 2013.
- [109] Boris Pittel. On likely solutions of a stable marriage problem. *The Annals of Applied Probability*, pages 358–401, 1992.
- [110] David Pollard. *Convergence of stochastic processes*. David Pollard, New Haven, Connecticut, 1984.
- [111] John G Riley and William F Samuelson. Optimal auctions. *The American Economic Review*, 71(3):381–392, 1981.
- [112] Ryan M Rogers and Aaron Roth. Asymptotically truthful equilibrium selection in large congestion games. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 771–782. ACM, 2014.
- [113] Amir Ronen. Mechanism design with incomplete languages. In *Proceedings of the 3rd ACM conference on Electronic Commerce*, pages 105–114. ACM, 2001.
- [114] Aaron Roth. Cis 800/002, fall 2011: The algorithmic foundations of data privacy, 2011.
- [115] Alvin E Roth. The economics of matching: Stability and incentives. *Mathematics of Operations Research*, 7(4):617–628, 1982.
- [116] Alvin E Roth. The evolution of the labor market for medical interns and residents: a case study in game theory. *The Journal of Political Economy*, 92(6):991, 1984.
- [117] Alvin E Roth. The college admissions problem is not equivalent to the marriage problem. *Journal of economic Theory*, 36(2):277–288, 1985.
- [118] Alvin E Roth and Elliott Peranson. The redesign of the matching market for american physicians: Some engineering aspects of economic design. Technical report, National bureau of economic research, 1999.
- [119] T. Roughgarden. Intrinsic robustness of the price of anarchy. In *STOC*, 2009.
- [120] T. Roughgarden and O. Schrijvers. Ironing in the dark. Submitted, 2015.
- [121] Tim Roughgarden. The price of anarchy in games of incomplete information. In *EC*,

2012.

- [122] Tim Roughgarden. Barriers to near-optimal equilibria. In *FOCS*, 2014.
- [123] Tim Roughgarden, Inbal Talgam-Cohen, and Qiqi Yan. Supply-limiting mechanisms. In *Proceedings of the 13th ACM Conference on Electronic Commerce, EC '12*, pages 844–861, New York, NY, USA, 2012. ACM. ISBN 978-1-4503-1415-2.
- [124] Timothy C. Salmon and Michael Iachini. Continuous ascending vs. pooled multiple unit auctions. *Games and Economic Behavior*, 2007.
- [125] Vasilis Syrgkanis. Bayesian games and the smoothness framework. *CoRR*, abs/1203.5155, 2012.
- [126] Vasilis Syrgkanis and Eva Tardos. Bayesian sequential auctions. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 929–944. ACM, 2012.
- [127] Vasilis Syrgkanis and Eva Tardos. Composable and efficient mechanisms. In *STOC*, 2013.
- [128] Vladimir N Vapnik and A Ya Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability & Its Applications*, 16 (2):264–280, 1971.
- [129] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37, 1961.
- [130] Robert Wilson. Game-theoretic analyses of trading processes. In Truman Fassett Bewley, editor, *Advances in Economic Theory*, pages 33–70. Cambridge University Press, 1987.
- [131] Andrew Chi-Chih Yao. An n-to-1 bidder reduction for multi-item auctions and its applications. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 92–109, San Diego, CA, USA., 2015. ACM.