NORTHWESTERN UNIVERSITY

Strategic Computation via Non-Revelation Mechanism Design

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Electrical Engineering and Computer Science

By

Darrell Hoy

EVANSTON, ILLINOIS

August 2015

ABSTRACT

Strategic Computation via Non-Revelation Mechanism Design

Darrell Hoy

Systems with strategic agents provide a challenge for the techniques of Computer Science: the behavior and objectives of agents are effectively already incorporated into the system, and so must be accommodated in analysis or design. This thesis studies how to incorporate the incentives and objectives of strategic agents into the analysis of non-revelation auctions. A central theme is understanding when simple auctions, which require little knowledge of the details of the agents, perform nearly as well as optimal auctions which require much more knowledge of the agents and the setting.

The center of our analysis is the simplest of auctions: the first-price auction, where the highest bidder wins and pays her bid. We will show that the first-price auction behaves well in many settings: for risk-neutral and risk-averse bidders, when the designer wants to maximize revenue or welfare, and in asymmetric as well as symmetric settings. We will also show that the first-price auction is an archetypal auction for analysis: we can reduce the analysis of other auctions, theoretical or empirical, to the analysis of the first-price auction.

Acknowledgments

I want to first thank my advisor Jason Hartline, for his advising, patience, unending commitment to his students, and for setting a great example of how to pursue research topics and agendas and ideas.

Without my collaborators, these works would be shells of themselves if they even existed. I owe a lot to each of them: Umang Bhaskar, Chien-Chung Huang, Lisa Fleischer, Jason Hartline, Evdokia Nikolova, Hu Fu, Kamal Jain, Chris Wilkens, Sam Taggart, Vasilis Syrgkanis, Denis Nekipelov, Brendan Lucier and Nicole Immorlica. I want to particularly thank Kamal Jain, Brendan Lucier and Nicole Immorlica for hosting me for internships, and Lisa Fleischer, who introduced me to Algorithmic Game Theory while an undergraduate at Dartmouth.

I have benefited immensely from my communities, both while at Northwestern and while visiting at Harvard. I want to particularly thank Greg Stoddard, Manolis Pountourakis, Nima Haghpanah and Sam Taggart for their camaraderie throughout both Northwestern in Evanston and Northwestern in Boston. I want to thank Alan and Laura Clark, who graciously welcomed me to Chicago, and left just as I did.

I want to thank everyone in the Harvard community for being so welcoming to the Northwestern crowd, particularly David Parkes, Yiling Chen, Drew Fudenberg, Bo Wagonner, CJ Ho and Rafael Frongillo.

I want to thank Michelle Zagardo, who has joined me for almost all of this grand adventure, making it so much more enjoyable, fulfilling and attainable.

Finally, I owe so many thanks to my parents for their support and encouragement and everything else throughout and before these five years.

Dedication

To the memory of my grandparents, Herb and Tommye Sauer, who through their own lifelong commitment to teaching and learning instilled academic and supportive values in me and in so many others they touched inside and outside their family.

Table of Contents

ABSTRACT 3

Acknowledgments 4

Dedication 6

List of Figures 10

Chapter 1. Introduction 12

1.1. Truthfulness in Auctions 15

1.2. Risk Aversion 17

1.3. Our Contributions 18

1.4. Organization 21

Chapter 2. Model & Background 22

2.1. Mechanism Design & Auction Theory 23

2.2. The Price of Anarchy & Robust Analysis 36

2.3. Risk-Aversion 38

Chapter 3. Price of Anarchy Bounds for Auction Welfare & Revenue 42

3.1. Introduction 43

3.2. Preliminaries 49

3.3. Single-Item First Price Auction 52

3.4. Revenue Approximation 58
3.5. Beyond Single-Item 65
3.6. Beyond Pay-Your-Bid Auctions 76
3.7. Conclusion 84

Chapter 4. Empirical Price of Anarchy Bounds 86
4.1. Introduction 87
4.2. Preliminaries 92
4.3. Price of Anarchy from Data 96
4.4. Data Analysis 107

Chapter 5. Risk-Averse Bidders 111
5.1. Introduction 112
5.2. Preliminaries 120
5.3. The Optimal Auctions 124
5.4. An Upper Bound on Two-Priced Expected Payment 127
5.5. Approximation Mechanisms and a Payment Identity 136
5.6. First Price Auctions with CARA Agents 142
5.7. Conclusions 148

Chapter 6. Simplifying Strategic Behavior: The Utility-Target Auction 151
6.1. Introduction 152
6.2. Definitions and Preliminaries 159
6.3. Quasi-Truthful Bidding 162
6.4. Static Equilibrium Analysis 164
6.5. Utility-Target Auctions for Sponsored Search 169

6.6. Dynamic Analysis 173

6.7. Conclusion and Open Questions 184

References 187

Appendix A. Revenue Covering 194

A.1. Single-item First-Price Auction Lowerbounds 194

A.2. Framework Proofs 195

A.3. Revenue Extension Proofs 196

A.4. Revenue Covering Proofs 199

A.5. All-Pay Auctions 203

Appendix B. Price of Anarchy from Data 207

Appendix C. The Utility Target Auction 211

C.1. Selected Proofs 211

Appendix D. Risk Averse bidders 232

D.1. Proofs from Section 5.3 232

D.2. Proofs from Section 5.4 234

D.3. Proofs from Section 5.5 237

List of Figures

2.1 The revenue curve plots the revenue from setting a price to sell to bidders in the top q fraction of bidders. 30

2.2 A comparison between the revenue curves and revenue from all agents and only agents with positive virtual values. 33

2.3 A visual depiction of the proof of the Bulow-Klemperer Theorem 2.9 for $n = 2, 3, 4$. 34

2.4 A concave utility function, corresponding to a risk-averse bidder 39

3.1 A depiction of the optimization problem a bidder faces in the first-price auction. 53

3.2 The indifference curve and expected threshold for a bidder in the first-price auction. 54

3.3 Value covering and virtual value covering in the first-price auction. 55

3.4 The effect of a reserve price on the first-price expected thresholds and value covering. 61

3.5 When bids are only allowed in multiples of δ, the expected threshold expands, but only by at most an additive amount of δ. 75

4.1 The price-per-click allocation rule in advertising auctions. 98

4.2 A comparison of full-information and Bayesian price-per-click allocation rules for GSP. 99

4.3 A comparison of expected thresholds in full-information and Bayesian settings for GSP. 100

4.4 The effect of the revenue covering parameter on price of anarchy bounds. 102

4.5 Allocation and threshold curves for a subset of six advertisers for a specific keyword. 108

4.6 Average bids, quality factors and revenue contribution for six advertisers on a specific keyword. 108

5.1 The capacitated utility function, $U_C(w) = \min(w, C)$. 114

5.2 Lower bounds on the capacitated allocation rule in two-priced auctions. 126

5.3 An example where the two-price capacitated auction has revenue 1.55 times the optimal risk-neutral revenue. 131

5.4 Depictions and bounds on the payments in two-priced auctions. 133

5.5 Breakdown of the expected payment upper bound in a two-priced auction. 134

A.1 The threshold in the GFP auction. 200

CHAPTER 1

Introduction

With the growth of the Internet, more and more systems are being developed that involve many agents beyond the control of the specific user or designer of the system. The performance of Yelp's rating system, Wikipedia's moderation system or Google's advertising auction depend crucially on how users behave within the systems.

Yet systems involving strategic agents provide a challenge for the techniques of Computer Science: the behavior and objectives of agents are a key component of the system and so must be accommodated in analysis or design. Incorporating strategic behavior of users is more challenging than incorporating properties of a circuit or other component we control, because we rarely have as precise an understanding of the incentives of an agent as we do of the properties and behavior of a circuit.

A mechanism is a strategic system wherein the real inputs to the algorithm are in the minds of the strategic agents. Traditional revelation mechanism design focuses on incentivizing the agents to reveal this information, so that the designer can easily solve her problem with the actual inputs.

However, revelation mechanisms are rarely found in practice, and so this thesis focuses on understanding non-revelation mechanisms for resource allocation problems (*auctions*): where the agents are incentivized to reveal *something* about their information, not all of their information. The private information in this setting is the value the agents place on the resource the designer controls. Auctions are used for the allocation of many resources:

advertising space, flowers, licenses to drill for oil or use wireless spectrum to name a few. They also offer a simple set of incentives from which to explore the nature of strategic computing: the bidders in the auction want the resource; the designer or auctioneer wants money (revenue maximizer), or wants to make sure the best person gets the resource (welfare maximizer).

Much of our focus will center around the simplest of non-revelation auctions: the first-price auction. In the (sealed-bid) first-price auction, agents submit their bids; the highest bid wins and pays her bid.

This thesis focuses primarily on auctions in the Bayesian setting, where the private information that agents have is modeled as being drawn from distributions, and we expect agents to best respond to the expected actions from other bidders. In (Bayes-Nash) equilibrium, each agent best responds to the expected actions from other agents, taking expectation over the private information of the other agents.

Optimal Auctions. Consider a designer trying to maximize revenue. If the designer (rather unrealistically) knows exactly the bidder who values the resource the most and how much they value it at, then the designer can simply charge that amount.

That process however is very unrealistic: at best, the designer might know the distributions of values the agents have for the resource, or something about the distributions of values. If all of the agents have independent and identically distributed private valuations for the resource, Myerson [1981] showed that the auctioneer need only know one number about the setting: the optimal reserve price. The optimal auction comes from running a first-price auction with that reserve: solicit bids, and sell to the bidder who offers the highest bid that is at least as high as the optimal reserve.

Simple vs. Optimal. Even that single number is not essential: Bulow and Klemperer [1996] showed that adding one more bidder to the auction is better than setting the optimal reserve price. If there are at least two bidders, this implies that the revenue from the first-price auction is at least half of the revenue of the first-price auction with the optimal reserve price.

When the setting is even a little more realistic, however, both the optimal auction and the behavior of bidders in the first-price auction become much more complicated: equilibria of the first price auction become very hard to characterize[1], and the optimal auction needs to know the exact distributions of values from the agents rather than only the optimal reserve price.

A core theme of this thesis is to push our understanding of the gap between simple and optimal auctions into more settings: in the asymmetric settings mentioned above (Chapter 3), and in settings with risk-averse bidders (Chapter 5). This is particularly important because it is rare that a designer knows the exact setting she is designing for, the exact distributions of values from the agents. Designing the perfect mechanism for the setting that a designer thinks is the right setting may be useless if the setting is actually a little bit different, and the mechanism's performance is not robust. Wilson [1987] advocated for this direction, in what is now known as the "Wilson Doctrine": that as the exact details of the decision making behavior from the agents are unknowable, we must focus on understanding how robust mechanisms and auctions are to changes in the details of the setting, not mechanisms that inherently depend and rely on the details.

[1] Vickrey [1961] posed the question of solving for equilibrium with two agents with private values drawn uniformly from distinct intervals; the problem was not solved until a half century later by Kaplan and Zamir [2012] with a very thorough case-based analysis.

1.1. Truthfulness in Auctions

The traditional approach in mechanism design is to incentivize the strategic agents to report their true inputs. The Vickrey-Clarke-Groves Mechanism, developed for single-item auctions by Vickrey [1961], and then later generalized [Clarke, 1971; Groves, 1973], has a truthful equilibrium that maximizes welfare in any quasi-linear setting.

Truthful auctions are very prevalent in theoretical work, because the assumption of truthful reporting sidesteps a lot of the issues in design with strategic agents. If the agents are truthful, the designer knows exactly the inputs and exactly the problem that needs to be solved, and so can simply solve it with traditional algorithms.

In practice, truthful auctions are rarely found: they suffer when agents do not know their exact values for an outcome, they are often sensitive to collusion [Ausubel and Milgrom, 2006], and as we will later see, simple versions of truthful auctions often suffer when bidders are risk-averse. They can also suffer if the agents are wary of the auctioneer learning their true information, for instance to take advantage of their situation the next time the auction is run.

But among all of these negatives, there is a strong justification for the prevalence of truthful auctions in theoretical analysis: the *revelation principle*, which tells us that any non-truthful equilibrium of an auction can be mapped to a truthful equilibrium of a different auction, that has the same outcomes and revenue and welfare.

Lemma 1.1 (The Revelation Principle [Myerson, 1981])**.** *For any non-truthful Bayes-Nash equilibrium of an auction $\mathcal{A}$, there is an auction $\mathcal{A}'$ with a truthful equilibrium that implements the same outcomes and payments.*

The proof is simple: ask agents for their private information; then act as they would have acted in the original auction with their private information.

As a result, when looking for the optimal auction for a given setting, we can restrict our analysis to only the truthful auctions. However, the revelation mapping says nothing about the simplicity of the revelation construction relative to the original non-revelation mechanism. This construction relies fundamentally on the designer's ability to simulate exactly how agents will make decisions. While this is a part of an agents private information, it is often not private information that can be easily reported: "how do you make decisions?", "what is your risk attitude?" are both questions that may be much harder to answer for some agents rather than "how much would you like to pay for this?".

Aside from the practical implementation details, the revelation mapping says nothing about the robustness of the truthful equilibrium. It says nothing about the case that agents only approximately best respond, or treat uncertainty differently than other agents.

This thesis focuses primarily on developing a robust, theoretical understanding of non-revelation auctions. We do this both because they are more prevalent in the real world, but also because non-truthful auctions are a closer model for the more general systems with strategic agents. Even when real world auctioneers run "truthful" auctions, they can become non-truthful if the setting changes even a little bit: for instance, if bidders are bidding in many auctions over the course of a day or a week. For example, Facebook runs the "truthful" Vickrey-Clarke-Groves auction [Varian and Harris, 2014]. However, if agents have a budget constraint over the course of the day or a week, they may prefer to bid non-truthfully in the auction and so the auction should be treated as non-truthful.

1.2. Risk Aversion

Traditionally, auction theory assumes that agents are indifferent to uncertainty, that they seek to maximize their expected wealth. This assumption leads to very tractable and elegant results: like Myerson's theorem of revenue equivalence between the first- and second-price auction:

Theorem 1.2 (Symmetric Revenue Equivalence [Myerson, 1981])**.** *For symmetric, risk-neutral, single-item environments, the revenues of Bayes-Nash equilibria of the first- and second-price auctions are identical.*

When bidders are risk-averse however, revenue equivalence breaks down. A risk averse bidder will bid higher in the first-price auction than a risk-neutral bidder will, to guarantee a larger probability of winning (at the cost of paying more when she wins), leading to revenue *dominance* rather than equivalence.

If we want to understand how to do computation subject to the laws of rational agents, then we must build theories and build an understanding that is robust to these behaviors and details of decision making. A theme explored in Chapter 5 is trying to understand what features of our understanding of risk-neutral strategic design carry over to the risk-averse setting.

The form of our results are generally comparisons between a simple auction that does not depend on the details of the risk-attitudes of the bidders, with the optimal auction that does depend on the details of the risk-attitudes of the bidders. This form of results aligns well with the Wilson doctrine.

1.3. Our Contributions

Theoretical Welfare and Revenue Analyses

In Chapter 3, we develop tools for welfare and revenue analysis of auctions in asymmetric settings. Much of our contributions are based on a very simple analysis of a very simple auction: the first-price auction, in single-item environments. We begin a simple proof of the following theorem in Bayesian settings:

Theorem 1.3 (informal, [Syrgkanis, 2012])**.** *For risk-neutral, asymmetric single-item environments, the first price auction has good welfare.*

The objective of welfare is a special one, because it is one in which the objectives of each bidder are totally aligned with the objective of the designer.

For revenue, the incentives are much less aligned between bidders and the designer. To adapt our approach to revenue, we use Myerson's [1981] analysis of revenue to effectively measure the alignment of incentives. When agents have positively aligned incentives, we can reduce the analysis to welfare analysis. We then show that setting the right reserve prices or ensuring agents face sufficient competition suffices to mitigate the impact of agents with negatively aligned incentives.

Theorem 1.4 (informal)**.** *For risk-neutral, asymmetric, regular, single-item environments, the first price auction has good revenue with either a) the right reserve prices set, or b) with sufficient competition.*

Notably, this result is the first revenue approximation result of a non-truthful auction *without* needing to solve for equilibrium.

The first-price auction works well for revenue and welfare because of a nice property of the auction: that if it is challenging for a bidder to win in the auction, it is for a good reason: other bidders must already be bidding (and hence paying) a lot.

Property 1.1 (informal)**.** An auction is *revenue-covered* if, when it is challenging for bidders to win, the auction has good revenue.

Theorem 1.5 (informal)**.** *For any auction that (approximately) satisfies revenue covering, the welfare of the auction is (approximately) good, and the revenue is (approximately) good, if the auction mitigates the impact of certain agents.*

With this theorem, the only work needed to analyze the welfare of an auction becomes analyzing the relationship between the challenge of winning in the auction and the revenue of the auction. We do this analysis for a number of auctions, including pay-your-bid multi-item auctions; all-pay single- and multi-item auctions; and the Generalized First Price position auction.

For revenue, Theorem 1.5 gives a bound for the positively aligned agents. The caveat is then that the auction must find a way to eliminate the impact of negatively aligned agents. Traditionally this is done with reserve prices, but sufficient competition from similar agents is also enough to give an approximation bound.

Price of Anarchy from Data

In Chapter 4, we build on the theoretical foundation in Chapter 3 with a key observation:

Observation 1.1 (informal)**.** We can empirically estimate the relationship between how hard it is to win and the expected revenue (revenue covering).

That is, if we are observing the data from an auction, we can measure an empirical relationship between the challenge of winning and revenue, and then use the entire theoretical framework of Chapter 3 to prove empirical price of anarchy bounds on the welfare of the auction.

Effectively, the data that we need to measure the relationship is the revenue, and the data needed to understand each bidder's optimization problem. If there is enough data available that a bidder can bid intelligently in the auction, then the designer can analyze the empirical price of anarchy.

This can be done for any auction, particularly for auctions in which we cannot theoretically bound the expected revenue against the challenge of winning. We analyze data from the Generalized Second Price auction that is run on Microsoft's BingAds platform, and find a range of values, indicating that for some keywords the auction is performing very well, for others there is a possibility for a large loss of welfare.

Risk-Averse Bidders

In Chapter 5, we consider the impact of risk-averse bidders on the performance of simple auctions. We show that the first-price auction does well when bidders risk-preferences are one of two types: "capacitated", or have constant absolute risk aversion (CARA).

For capacitated bidders, we show that the revenue of the first price auction with at least two bidders is a constant approximation to the revenue of the optimal auction.

For CARA bidders, we show the first-price auction is approximately optimal relative to the first-price auction with the optimal reserve price. Namely, we show that the Bulow-Klemperer approximation result for the first-price auction extends to the case of CARA bidders, relative to the first-price auction with the optimal reserve. Our result makes use of a generalization of Myerson's virtual-value based characterization of revenue to the CARA setting.

The Utility Target Auction

In Chapter 6, we flip the nature of strategic behavior, and consider modifications in the rules of an auction so as to support simpler strategic behavior. Notably, we take a revelation principle like approach, and develop a quasi-truthful auction that has many of the features of a non-truthful auction.

In position auctions, the most-natural extension of the first-price auction can lack equilibria in full-information settings. One approach to ensure equilibria is to solicit bids from each agent for each possible position. However, as the number of possible outcomes rises, the strategic complexity for agents becomes very high. We instead have agents report their valuation function — which in the case of position auctions, can be very concise, like "value-per-click" — and a single strategic bid, their "utility-target". The utility-target auction has good equilibrium performance, and in full-information settings, we show it has quasi-truthful equilibria.

1.4. Organization

In Chapter 2, I introduce the fundamentals of the settings that we will be operating in. I recommend reading Chapter 2 before other chapters, and Chapter 3 before Chapter 4; all other chapters are fairly self-contained.

CHAPTER 2

Model & Background

In this section, I give an introduction to single-parameter Bayesian mechanism design, with a focus on understanding the behavior of simple auctions relative to optimal auctions, and understanding the role of risk attitudes.

This model is the primary model used throughout this thesis. Chapters 3, 4, 5 all use the single-parameter Bayesian Model, while Chapter 6 only differs in focusing on full-information (non-Bayesian) settings.

The sections regarding risk-aversion are important for Chapter 5, and can be skipped for readers interested only in the other chapters.

For a more thorough introduction to Bayesian mechanism design, the reader is encouraged to see Hartline [2014].

2.1. Mechanism Design & Auction Theory

A *mechanism* is an algorithm where the input comes from each of n strategic agents, who have preferences over the outcome. In an *auction*, the outcome consists of a feasible allocation $\mathbf{x} = (x_1, \ldots, x_n)$ of a resource to the agents, and payments $\mathbf{p} = (p_1, \ldots, p_n)$ made by the agents to the designer. In such settings, we will treat the auction as consisting of an allocation rule $\tilde{\mathbf{x}} : \mathrm{A}^n \to \mathbb{R}^n$ and payment rule $\tilde{\mathbf{p}} : \mathrm{A}^n \to \mathbb{R}^n$ as mapping actions to feasible allocations and payments respectively[1].

Agent Preferences. Each agent i chooses an action $a_i \in \mathrm{A}_i$, based on her own preferences over the outcomes. We assume agents have *single-dimensional* preferences over the resource, with a private value $v_i \in V_i$ for being allocated the resource. As we focus on Bayesian environments, we model each agent's value as being drawn independently from some known distribution, $v_i \sim F_i$. This is known as the *Independent Private Values* (IPV) setting. The symmetry of a setting refers generally to whether or not the value distributions are identical across agents.

We call $v_i x_i - p_i$ the wealth created for agent i and we assume that agents are expected-utility maximizers over wealths. An agent's utility for allocation x_i and payment p_i is $U_i(v_i x_i - p_i)$, and the agent's utility when the actions $\mathbf{a}$ are played is

$$\tilde{u}_i(\mathbf{a}) = U_i(v_i \tilde{x}_i(\mathbf{a}) - \tilde{p}_i(\mathbf{a})). \tag{2.1}$$

When there is uncertainty over the actions that other agents will play, then the agent will seek to maximize her expected utility,

[1] I will use boldface ($\mathbf{a}$) to denote a vector, or a function that produces a vector. Subscripts will refer to the elements in the vector: a_i is the ith element of $\mathbf{a}$, and $\mathbf{a}_{-i}$ is the vector consisting of all but the ith element of $\mathbf{a}$.

$$\tilde{u}_i(a_i) = \mathbf{E}_{\mathbf{a}_{-i}}\left[U_i(v_i\tilde{x}_i(\mathbf{a}) - \tilde{p}_i(\mathbf{a}))\right]. \tag{2.2}$$

If $U_i(w) = w$, then we say the agent is *risk-neutral.* We primarily focus on the risk-neutral case, but see Section 2.3 and Chapter 5 for discussions of the case when agents' utilities are non-linear.

Feasibility Environments. The feasibility environment describes the set of allowed allocations by the auction. We will primarily operate in the single-item feasibility environment, where one resource can be allocated to at most one agent, but many of the results will be generalized to the following environments:

- Multi-item - assignments of m identical items to at most m of the n agents.
- Position Auctions - assignment of m ordered positions to at most m of the agents. Used primarily in Internet advertising auctions, where the better the position, the higher the probability is that a searcher will click on the advertisement.
- Combinatorial - assignments of m non-identical objects to at most m of the n agents, with a set system denoting feasible allocations.
- Matroid - A combinatorial feasibility environment, where the set system of feasible allocations forms a matroid. This includes the multi-item environment, as well as other systems, like spanning trees in a graph where the base resource is the edges in the graph.

We represent the set of feasible outcomes as X.

2.1.1. Design Objectives

The designer of the mechanism has some objective over the outcomes and the agents. The most common objectives are *revenue*, the sum of payments to the designer, and *welfare*, the utility of the agents plus the payments to the designer.

2.1.2. Timeline of an Auction

The timeline of an auction is as follows:

(1) Agents realize private values, $\mathbf{v} \sim \mathbf{F}$.

(2) Agents report actions $\mathbf{a} = \mathbf{s}(\mathbf{v})$ in accordance with strategy profile $\mathbf{s}$.

(3) The auction chooses allocations $\tilde{\mathbf{x}}(\mathbf{a})$, and charges payments $\tilde{\mathbf{p}}(\mathbf{a})$ to the bidders.

There are three important stages in the auction: before private values are realized (the *ex-ante* stage), after private values have been realized, before actions have been reported (the *interim* stage), and after the payments and allocations have been announced (the *ex-post* stage).

Individual Rationality. Depending on the circumstance of the auction, agents may or may not be able to walk away. For instance, if a mechanism instructs a bidder to pay $1 million for an ad impression, the agent will likely refuse, and leave the auction.

We call this property *individual rationality* [IR]. There are three variants of individual rationality, corresponding to the expected utilities in the three stages of the mechanism (ex-ante, interim, ex-post).

2.1.3. Strategies & Equilibrium

A *strategy profile* $\mathbf{s} : V^n \to \mathrm{A}^n$ maps values of agents to actions.

If agents are acting according to their own incentives and values, what should we expect them to do? Nash [1951] introduced what we now call the Nash Equilibrium: a steady state in which all agents are best-responding to the actions of other bidders. We focus primarily on Bayes-Nash equilibrium, where the best responding happens over the uncertainty created by other bidders random draws of valuations.

Definition 2.1. A strategy profile $\mathbf{s}$ is a Bayes-Nash Equilibrium (BNE) if for each agent i and every realized value $v_i \sim F_i$ and every alternate action a_i',

$$\tilde{u}_i(s_i(v_i)) \geq \tilde{u}_i(a_i'). \tag{2.3}$$

2.1.4. Types of Auctions

The majority of this thesis is focused on one of the simplest of auctions: the *first-price* auction (FPA). In single item settings, the first-price auction solicits bids; the highest bidder wins and pays her bid. In more complicated feasibility environments, the first-price auction generalizes to the *pay-your-bid* mechanisms, in which agents submit bids, the auction chooses a feasible allocation, and charges each agent who is allocated the resource her bid. Generally the allocation is picked so as to either maximize or approximately maximize the sum of the bids.

In the *second-price* auction (SPA), the auction solicits bids; the highest bidder wins and pays the *second-highest* bid.

In the *all-pay* auction (AP), the auction solicits bids; the highest bidder wins and everyone pays her bid.

Auctions can also have *reserve* prices, either individual by bidder or common across all bidders. With a reserve price, the designer only considers bids above the reserve price.

2.1.5. Bayesian Auction Theory

Given a strategy profile $\mathbf{s}$, we often consider the expected allocation and payment an agent faces from choosing an action $a_i \in \mathrm{A}_i$, with expectation taken with respect to other agents' values and actions induced by $\mathbf{s}$. We treat $\mathbf{s}$ implicitly and write $\tilde{x}_i(a_i) = \mathbf{E}_{\mathbf{v}_{-i}}[\tilde{x}_i(a_i, \mathbf{s}_{-i}(\mathbf{v}_{-i}))]$, with $\tilde{p}_i(a_i)$ and $\tilde{u}_i(a_i)$ defined analogously.

Given $\mathbf{s}$ implicitly, we also consider values as inducing payments and allocations. We write $\mathbf{x}(\mathbf{v}) = \tilde{\mathbf{x}}(\mathbf{s}(\mathbf{v}))$ and $\mathbf{p}(\mathbf{v}) = \tilde{\mathbf{p}}(\mathbf{s}(\mathbf{v}))$, respectively. Furthermore, we can denote agent i's interim allocation probability and payment by $x_i(v_i) = \tilde{x}_i(s_i(v_i))$ and $p_i(v_i) = \tilde{p}_i(s_i(v_i))$. We define $\mathbf{u}(\mathbf{v})$ and $u_i(v_i)$ analogously. In general, we use a tilde to denote outcomes induced by actions, and omit the tilde when indicating outcomes induced by values. We refer to $\tilde{\mathbf{x}}$ as the *bid allocation rule*, to distinguish it from $\mathbf{x}$, the *allocation rule*. We adopt a similar convention with other notation.

Myerson [1981] showed that the value-based allocation rule of an auction is all that is needed to characterize the expected payments in an auction (up to a constant factor):

Lemma 2.1 (Myerson [1981])**.** *For any single-parameter, risk-neutral auction, the expected payment for agent i in BNE is:*

$$p_i(v_i) = v_i x_i(v_i) - \int_0^{v_i} x_i(z)\,\mathrm{d}z + C. \tag{2.4}$$

Because the payment is decided entirely by the allocation rule, two auctions with the same allocation rule must have the same revenue: up to the constant factor C in Equation (2.4)[2]. This property is known as revenue equivalence.

[2]We generally assume that this constant is 0: if C is positive, then this will violate interim IR, as an agent with no value for the item will be forced to pay; if C is negative, then the auctioneer is giving away free money to all the agents.

Corollary 2.2 (Revenue Equivalence)**.** *For risk-neutral, single-item environments, the revenues of auctions with the same allocation rules are identical.*

In symmetric settings, the highest bidder will win in Bayes-Nash Equilibrium of either the first- or second-price auctions [Vickrey, 1961]. Thus the allocation rules and hence the revenues will be the same:

Corollary 2.3 (Symmetric Revenue Equivalence)**.** *For symmetric, risk-neutral, single-item environments, the revenues of the first- and second-price auctions are equivalent.*

Note the structure of the payment rule: each bidders payment is influenced by the allocation rule of all of the agents with smaller values. Myerson [1981] gives an amortized analysis of this impact of a bidder on higher valued agents.

Lemma 2.4. *The revenue in any BNE* $\mathbf{s}$ *of auction* $\mathcal{A}$ *satisfies*

$$\textsc{Rev}(\mathcal{A}(\mathbf{s})) = \mathbf{E}_{\mathbf{v}\sim\mathbf{F}}\left[\sum_i \phi_i(v_i)x_i(v_i)\right], \tag{2.5}$$

where $\phi_i(v_i)$ *is the virtual value of the agent, defined as*

$$\phi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}. \tag{2.6}$$

The virtual value quantity plays the same role for revenue as the value does for the objective of welfare: the revenue is measured as the expected virtual value of the agents served. Note that the virtual value will always be smaller than the value. It can also be either positive or negative, indicating that for some bidders, the auction would prefer to completely exclude them from allocation.

We will break down the revenue of an auction into the revenue gained by serving agents with positive virtual values, and the revenue lost due to serving agents with negative virtual values.

Definition 2.2. For any BNE $\mathbf{s}$ of auction $\mathcal{A}$, let $\textsc{Rev}^+(\mathcal{A})$ and $\textsc{Rev}^-(\mathcal{A})$ be the virtual surplus from agents with positive and negative virtual values, respectively. Thus,

$$\textsc{Rev}^+(\mathcal{A}) = \sum_i \mathbf{E}_{v_i}\left[\max(0, \phi_i(v_i))x_i(v_i)\right],$$

$$\textsc{Rev}^-(\mathcal{A}) = -\sum_i \mathbf{E}_{v_i}\left[\min(0, \phi_i(v_i))x_i(v_i)\right],$$

and $\textsc{Rev}(\mathcal{A}) = \textsc{Rev}^+(\mathcal{A}) - \textsc{Rev}^-(\mathcal{A})$.

Quantiles and Revenue Curves. We will often refer to bidders of a given type by their quantile within their distribution, $1 - F_i(v)$. Let the inverse be $v(q) = F_i^{-1}(1 - q)$. We will lightly overload the value-based bid and allocation rules, and let $x_i(q) = x_i(v_i(q))$ and $p_i(q) = p_i(v_i(q))$.

Consider posting a price of $v(q)$ to a single bidder with value distribution F_i. The revenue is $qv(q)$, as a bidder has value at least $v(q)$ with probability q. We refer to the plot of this as the revenue curve of the distribution: see Figure 2.1 for an illustration of a revenue curve.

Lemma 2.5. *The revenue in any BNE* $\mathbf{s}$ *of auction* $\mathcal{A}$ *satisfies*

$$\textsc{Rev}(\mathcal{A}(\mathbf{s})) = \sum_i \int_0^1 -x_i'(q)R_i(q)\,\mathrm{d}q, \tag{2.7}$$

where the revenue curve $R_i(q) = qv_i(q)$ *is the revenue from posting a price* $v_i(q) = F_i^{-1}(1-q)$ *to a single agent with value drawn from* F_i*.*

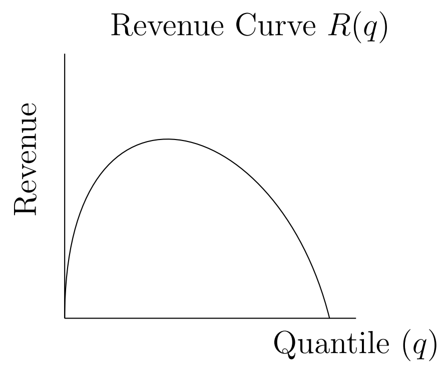

Figure 2.1. The revenue curve plots the revenue from setting a price to sell to bidders in the top q fraction of bidders.

Note that because a larger quantile means a smaller value, the allocation is decreasing in the quantile, hence $x_i'(q) < 0$.

The virtual value quantity $\phi(v)$ is the derivative of the allocation rule:

$$\begin{aligned} R'(q) &= v_i(q) - (-v_i'(q)q) \\ &= v_i - \frac{1 - F_i(v_i)}{f_i(v_i)} \\ &= \phi_i(v_i). \end{aligned}$$

Note again that as quantile is decreasing in value, $v_i'(q) < 0$.

Definition 2.3. A distribution F is *regular* if its revenue curve is concave, equivalently if the virtual value $\phi(v_i) = v_i - \frac{1-F(v_i)}{f(v_i)}$ is monotone non-decreasing in v.

Since the virtual value does not depend on any property of the allocation rule, the optimization problem of choosing the best allocation rule for revenue becomes easy: simply allocate to the agent with the highest non-negative virtual value.

In symmetric, regular settings, the agent with the highest virtual value is also the agent with the highest value, so the optimal auction has a very simple form: run a second-price auction with reserve set at $\phi^{-1}(0)$. Setting that reserve eliminates any agents with negative virtual value from being allocated, so $\textsc{Rev}^{-} = 0$. A first price auction will induce the same allocation rule, so by revenue equivalence (Lemma 2.2), a first-price auction with the same reserve will also be optimal.

Theorem 2.6 ([Myerson, 1981])**.** *In a single-item auction, with regular, risk-neutral bidders, the optimal auction allocates to the bidder with the highest non-negative virtual value.*

Corollary 2.7. *In a symmetric single-item auction with regular, risk-neutral bidders, the optimal auction is implemented by either a first or second price auction with reserves* $\mathbf{r}$ *set to satisfy* $\phi_i^{-1}(r_i) = 0$ *for each bidder.*

In asymmetric settings, the optimal auction for revenue does not have quite as simple a form. The underlying allocation rule is still the same: allocate to the agent with the highest non-negative virtual value, but this allocation rule will no longer be implemented by a first or second price auction.

Lemma 2.8. *For asymmetric, regular environments, the optimal auction is to run the Vickrey-Clarke-Groves mechanism on virtual values of agents in place of the values of the agents.*

Note that this mechanism is now much more complicated than in symmetric settings: the designer must know exactly the distribution of values of each agent so as to calculate the virtual values. In the symmetric setting, the designer needed only to know or learn one number, the optimal reserve price to set.

2.1.6. Simple vs. Optimal

In practice most auctions are simple: they do not allocate by virtual values, and oftentimes they lack even the reserve price that is a fundamental piece of Myerson's optimal auctions. Bulow and Klemperer [1996] offered some theoretical justification for this in a very elegant result: in symmetric risk-neutral settings, it is better to add one more agent to the auction than implement the optimal reserve price. Thus, instead of expending effort to learn the optimal reserve, the auctioneer can simply expend that effort trying to attract more bidders.

We will make strong use of a slightly differently phrased version of the result:

Theorem 2.9 (Bulow and Klemperer [1996]). *For regular, symmetric, risk-neutral, single-item environments, the revenue in the first- or second-price auctions with n agents satisfies*

$$\textsc{Rev}(\text{FPA}) = \textsc{Rev}(\text{SPA}) \geq \frac{n-1}{n}\textsc{Rev}(\textsc{Opt}). \tag{2.8}$$

Later, we will generalize Theorem 2.9 to the first-price auction in asymmetric settings (Theorem 3.12), and with risk-averse bidders (Theorem 5.19).

I provide a proof to illustrate the underlying concepts for the simple symmetric and risk-neutral case: both extensions will rely on tweaking the intuition developed through these proofs.

Our proof will build on the split of revenue into the revenue due to agents with positive and negative virtual values. Let $R^+(q) = \int_0^q \max(0, \phi_i(z))\,\mathrm{d}z$ be the revenue curve from positive virtual-valued agents (analogous to $\textsc{Rev}^+$ from Definition 2.2). For regular

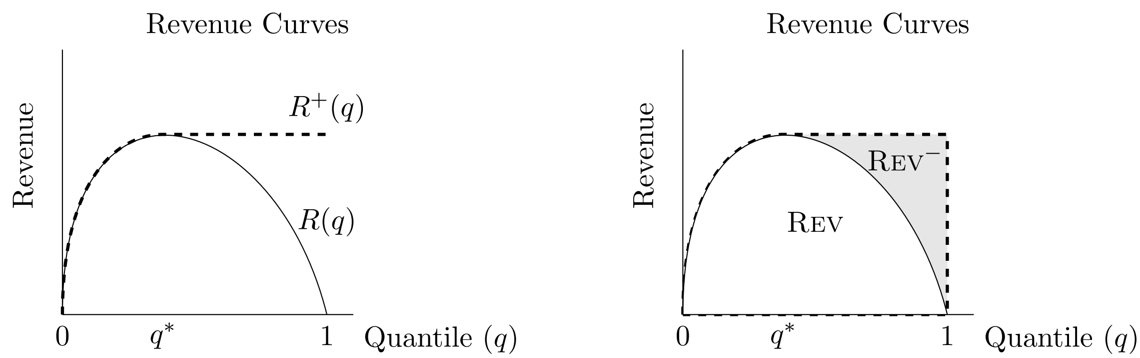

(a) A comparison of the revenue curve from positive virtual valued agents $R^+(q)$ versus the normal revenue curve $R(q)$.

(b) A comparision of REV and REV$^-$. By the concavity of R, REV$^- \le$ REV.

Figure 2.2. A comparison between the revenue curves and revenue from all agents and only agents with positive virtual values.

environments, the monotonicity of the revenue curve then gives

$$R^+(q) = \max_{q' \le q} R(q') \tag{2.9}$$

See Figure 2.2a for a comparison of the normal revenue curve and the revenue curve from positive virtual valued agents.

We first give the proof for $n = 2$ to build intuition, and then provide the general proof. [3]

PROOF FOR $n = 2$. For the two agent setting, the allocation rule by quantile is $x(q) = (1-q)$. The expected revenue is then the area underneath the revenue curves: REV(FPA) $= 2\int_0^1 -x'(q)R(q)\,\mathrm{d}q = 2\int_0^1 R(q)\,\mathrm{d}q$, and REV$^+$(FPA) $= 2\int_0^1 R^+(q)\,\mathrm{d}q$. See Figure 2.2a for a comparison of the two revenue curves. The revenue lost due to negative virtual valued agents is then the area between the curves, REV$^-$(FPA) $= 2\int_0^1 (R^+(q) - R(q))\,\mathrm{d}q$.

[3]The geometric proof for two bidders is essentially the same proof as the two-bidder case by Dhangwatnotai et al. [2010]; the generalization to the revenue curve versus losing probability plot is new.

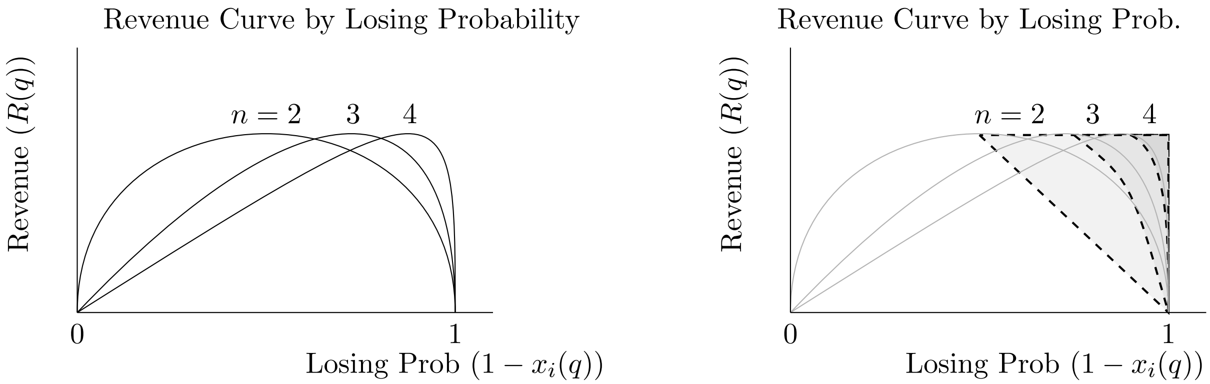

(a) A plot of the revenue curve $R(q)$ against $1 - x_i(q)$, for n = 2,3, and 4.

(b) The concavity of $R(q)$ in q results in an upper bound on $\textsc{Rev}^-$ which shrinks as n grows.

Figure 2.3. A visual depiction of the proof of the Bulow-Klemperer Theorem 2.9 for $n = 2, 3, 4$.

By the concavity of $R(q)$ in q, $\textsc{Rev}^-(\text{FPA}) \leq \textsc{Rev}(\text{FPA})$, so

$$\textsc{Rev}(\text{FPA}) \leq \frac{1}{2}\textsc{Rev}^+(\text{FPA}) = \frac{1}{2}\textsc{Rev}^+(\textsc{Opt}).$$

See Figure 2.2b for an illustration. □

The geometric interpretation for two agents relied on the simplicity of the quantile-allocation rule: the allocation rule corresponded to a uniform distribution over the quantiles. For n agents, the allocation rule corresponds to the maximum of $n - 1$ uniformly drawn quantiles, so $x_i(q) = (1 - q)^{n-1}$ and $x_i'(q) = -(n-1)(1-q)^{n-2}$. Instead of the uniform distribution, the distribution is weighted towards smaller quantiles, which serves to further shrink $\textsc{Rev}^-(\text{FPA})$.

We can generalize the geometric intuition for two agents to n agents by plotting the probability of losing for an agent $(1 - x(q))$ against their revenue curve $(R(q))$. See Figure 2.3a.

Proof for $n \geq 2$. For the n agents case, we have

$$\textsc{Rev}(\text{FPA}) = n \int_0^1 -x_i'(q) R(q)\, \mathrm{d}q \tag{2.10}$$

$$= n \int_0^1 \left(\frac{\partial}{\partial q} (1 - x_i(q)) \right) R(q)\, \mathrm{d}q \tag{2.11}$$

This integral is exactly the area underneath the parametric curve $(1 - x_i(q), R(q))$ from $q = 0$ to $q = 1$, pictured in Figure 2.3a for $n = 1$, 2, and 3.

Let q^* be the optimal quantile. Using $x_i'(q) = -(n-1)(1-q)^{n-2}$, we have

$$\textsc{Rev}(\textsc{Opt}) = n \int_0^1 -x_i'(q) R^+(q)\, \mathrm{d}q \tag{2.12}$$

$$= n \int_0^{q^*} (n-1)(1-q)^{n-2} R(q)\, \mathrm{d}q + (1-q^*)^{n-1} R(q^*) \tag{2.13}$$

By concavity of $R(q)$ in q, for any quantile between q^* and 1,

$$R(q) \geq \frac{1-q}{1-q^*} R(q^*). \tag{2.14}$$

Plotting the right-hand side of Equation (2.14) against the losing probability results in a lower bound that gains more curvature as the number of agents grows - see Figure 2.3b.

For the first price revenue we then have

$$\begin{aligned}
\textsc{Rev}(\text{FPA}) &= \int_0^{q^*} (n-1)(1-q)^{n-2} R(q)\, \mathrm{d}q + \int_{q^*}^1 (n-1)(1-q)^{n-2} R(q)\, \mathrm{d}q \\
&\geq \int_0^{q^*} (n-1)(1-q)^{n-2} R(q)\, \mathrm{d}q + \int_{q^*}^1 (n-1)(1-q)^{n-2} \frac{1-q}{1-q^*} R(q^*)\, \mathrm{d}q \\
&= \int_0^{q^*} (n-1)(1-q)^{n-2} R(q)\, \mathrm{d}q + \frac{n-1}{n} R(q^*)(1-q)^{n-1} \\
&\geq \frac{n-1}{n} \textsc{Rev}(\textsc{Opt}).
\end{aligned}$$

□

The optimal auction for risk-neutral agents has two primary features: (1) agents are picked in order of virtual values, and (2) no agents with negative virtual values are allocated. In symmetric, regular environments, the first feature is easily attainable because a larger value leads to a larger bid and a larger virtual value. So picking agents by bid or by value is equivalent to picking them by virtual value. In the symmetric setting, Theorem 2.9 is fundamentally about the value of the reserve, about bounding the impact of agents with negative virtual values.

In the asymmetric setting, the first or second price auctions differ from the optimal auction for both positive and negative virtual values. Hartline and Roughgarden [2009] generalized Theorem 2.9 for asymmetric settings in which each agent has a duplicate bidder with value drawn from the same distribution. They also isolated the effect of ranking by value in place of virtual value, and showed a constant approximation for $\text{VCG}_{\mathbf{r}}$ with the right reserve prices set.

Theorem 2.10 (Hartline and Roughgarden [2009]). *For regular, asymmetric, matroid, risk-neutral environments, the revenue in the truthful equilibrium of* VCG *with duplicate bidders is at least half the revenue of the optimal auction without duplicates.*

In Chapter 3, we will provide analogous generalizations of the Bulow-Klemperer result for the first-price and all-pay auctions in asymmetric settings.

2.2. The Price of Anarchy & Robust Analysis

Traditional equilibrium analysis comes from first characterizing equilibrium, then proving properties of that equilibrium. Indeed, this is how the results so far have been shown:

Theorems 2.6, 2.9 and 2.10 all stem from the assumption that either agents bid truthfully or play symmetric strategies.

Koutsoupias and Papadimitriou [1999] introduced a style of analysis for understanding the effect of selfish behavior on the efficiency of outcomes, the *price of anarchy.*

Definition 2.4. The Price of Anarchy of a game $\mathcal{G}$, for (maximization) objective $\mathcal{O}$ is

$$\textsc{PoA}(\mathcal{G}) = \max_{\text{settings,Eq.}} \frac{\mathcal{O}(\textsc{Opt})}{\mathcal{O}(\mathcal{G})},$$

Smooth Games & Mechanisms. Roughgarden [2009] and Syrgkanis and Tardos [2013] provide a robust approach to proving price of anarchy bounds for the objective of welfare in games like auctions: that agents have a potential strategy not based on the actions of the other agents which guarantees them a good fraction of utility *if* there is a good strategy for them. This property is called *smoothness.* A core principle of smoothness is that the precise manner or assumptions used in proving the smoothness property dictate how broadly the result extends: whether it holds for instance when a number of auctions are run simultaneously.

In this thesis, we will make extensive use of price of anarchy based approaches to understanding equilibria. In Chapter 3, we will use smoothness type techniques specifically applied to the Bayesian setting to derive simpler proofs of welfare bounds that have been proven via smoothness, and to build the framework for our revenue approximation results. In Chapter 4, we will present a refinement of the price of anarchy for empirical settings, using the approach of worst-case analysis for robust inference of auction performance.

2.3. Risk-Aversion

We generally assume in Auction Theory that agents are risk-neutral, that they try only to maximize their expected payoffs. This enables simple and elegant analysis, as many problems within mechanism design become linear and tractable. For example, in a symmetric, single-item environment, Myerson's payment identity (Lemma 2.1) and revenue equivalence tell us that the first and second price auction will have the same revenue — as a designer, we can then choose whichever auction we prefer.

But when agents are risk averse, this foundation of auction theory no longer applies: the first- and second-price auction are *not* revenue equivalent, and the first-price auction usually dominates the revenue of the second price auction.

In reality people care about risk and uncertainty. Most people will prefer a certain \$1 million rather than a 50% shot at \$3 million, even though the certain outcome is worse in expectation. If trying to drive to the airport on time, a traveler will try and maximize the probability she arrives on time, not just minimize expected travel time. If a buyer for a company has approval to pay up to \$1000 for a good, then she will be much more willing to pay \$1000 all the time rather than \$1500 half the time, free the other half of the time.

In the single-item auction example, if bidders are risk-averse, then the first-price auction will likely get much more revenue than the second-price auction as bidders are willing to trade off a higher payment for a higher probability of winning. This behavior of revenue dominance happens in many settings: later, we show it for two types of risk-averse bidders.

There are many ways of modeling the way that agents react to uncertainty. The most standard approach, introduced by Von Neumann and Morgenstern [1944], is to assume that agents have a utility function $U : \mathbf{O} \to \mathbb{R}$ which maps outcomes in some space $\mathbf{O}$ to a

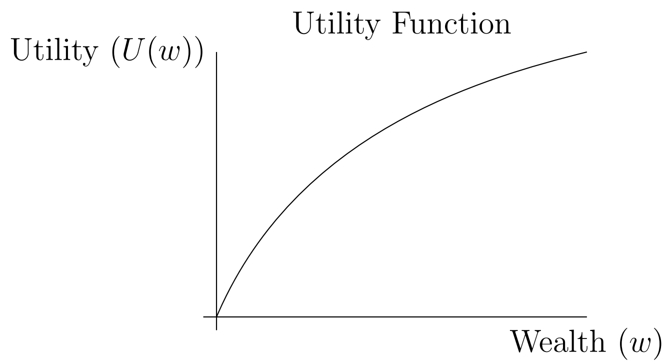

Figure 2.4. A concave utility function, corresponding to a risk-averse bidder

cardinal value, the "utility" for that outcome, and that the agents act so as to maximize their expected utility.

We will assume agents preferences are over the resulting wealth of outcomes: the value created by the allocation, less any payment. So the problem of such an agent is the following:

$$\max_{a \in \mathrm{A}} \mathbf{E}\left[U(v(o(a)) - \tilde{p}_i(a))\right] \tag{2.15}$$

If the utility function of a bidder is concave, the agent is risk-averse: the utility of the average of two outcomes will be higher than the average utility of the outcomes. Likewise, risk-seeking behavior corresponds to a convex utility function.

The more concave the utility function, the more risk averse the bidders are. The Arrow-Pratt measure of absolute risk-aversion [Arrow, 1971; Pratt, 1964]:

$$A(w) = -\frac{U''(w)}{U'(w)} \tag{2.16}$$

measures this at the wealth level w. When $A(w)$ is constant, then the agent exhibits *constant absolute risk aversion* (CARA), which means that an agents preferences are not affected by a shift in wealth: if the agent prefers a certain \$4 to a 50% \$10, 50% \$0 , then the bidder will also prefer a certain \$104 to 50% \$110, 50% \$100.

We focus on the case that agents are expected utility maximizers, for tractability and as a foundational point for analysis. Indeed, when agents are not expected utility maximizers, many things get even more complicated: for instance, Nash's foundational theorem that a mixed nash equilibrium always exists no longer holds [Fiat and Papadimitriou, 2010].

2.3.1. Risk Averse Auction Theory

Auction analysis and design become significantly more complicated when bidders are risk-averse. In first-price auctions, bidders with concave utilities over wealth bid higher than they would have if risk-neutral, leading to increased revenue. For the second-price auction however, risk-aversion has no effect - leading to revenue dominance of the first-price auction over the second-price, rather than the revenue equivalence of symmetric risk-neutral settings [Riley and Samuelson, 1981; Holt, 1980; Maskin and Riley, 1984; Matthews, 1983].

For CARA bidders, revenue dominance comes at no cost, as the bidders are indifferent between first- and second-price auctions, even though they pay more in expectation in the first-price auction. Thus the increased revenue to the auctioneer is in effect free money — revenue comes not at the expense of bidders' utilities, but from giving them a payment scheme that they are much happier with.

The optimal auction for revenue is also much more complicated, and is found usually via optimal control [Matthews, 1983; Maskin and Riley, 1984]. The resulting auctions are crucially dependent on the exact details of the utility function and distributions of the agents, violating the Wilson Doctrine.

In Chapter 5, we focus on developing simple-vs-optimal results for risk-averse bidders which are particularly important given the challenge risk-aversion offers to the analyst and designer.

CHAPTER 3

Price of Anarchy Bounds for Auction Welfare & Revenue

In this chapter, we develop tools for welfare and revenue analysis in asymmetric settings. Much of our contributions are based on a very simple analysis of a very simple auction: the single-item, first-price auction. We then show that analysis of much more complicated mechanisms can be reduced to the analysis of the first-price auction: and later, in Chapter 4, we show that we can reduce the econometric analysis of other auctions to the theoretical analysis of the simple first-price auction.

3.1. Introduction

The first step of a classical microeconomic analysis is to solve for equilibrium. Consequently, such analysis is restricted to settings for which equilibrium is analytically tractable; these settings are often disappointingly idealistic. Methods from the price of anarchy provide an alternative approach. Instead of solving for equilibrium, properties of equilibrium can be quantified from consequences of best response. These methods have been primarily employed for analyzing social welfare. While welfare is a fundamental economic objective, there are many other properties of economic systems that are important to understand. This chapter gives methods for analyzing the price of anarchy for revenue.

Equilibrium requires that each agent's strategy be a best response to the strategies of others. A typical price-of-anarchy analysis obtains a bound on the social welfare (the sum of the revenue and all agent utilities) from a lower bound on an agent's utility implied by best response. Notice that the agents themselves are each directly attempting to optimize a term in the objective. This property makes social welfare special among objectives. Can simple best-response arguments be used to quantify and compare other objectives? This chapter considers the objective of revenue, i.e., the sum of the agent payments. Notice that each agent's payment appears negatively in her utility and, therefore, she prefers smaller payments; collectively the agents prefer smaller revenue.

The agenda of this chapter parallels a recent trend in mechanism design. Mechanism design looks at identifying a mechanism with optimal performance in equilibrium. Optimal mechanisms tend to be complicated and impractical; consequently, a recent branch of mechanism design has looked at quantifying the loss between simple mechanisms and optimal mechanisms. These simple (designed) mechanisms have carefully constructed equilibria (typically, the truthtelling equilibrium). The restriction to truthtelling equilibrium, though convenient in theory, is problematic in practice [Ausubel and Milgrom, 2006]. In particular, this truthtelling equilibrium is specific to an ideal agent model and tends to be especially non-robust to out-of-model phenomena. The price of anarchy literature instead considers the analysis of the performance of simple mechanisms absent a carefully constructed equilibrium.

As an example, consider the single-item first-price auction, in which agents place sealed bids, the auctioneer selects the highest bidder to win, and the winner pays her bid. The fundamental tradeoff faced by the agents in selecting a bidding strategy is that higher bids correspond to higher chance of winning (which is good) but higher payments (which is bad). This first-price auction is the most fundamental auction in practice and it is the role of auction theory to understand its performance. When the agents' values for the item are drawn independently and identically, first-price equilibria are well-behaved: the symmetry of the setting enables the easy solving for equilibrium [Krishna, 2009], the equilibrium is unique [Chawla and Hartline, 2013; Lebrun, 2006; Maskin and Riley, 2003], and the highest valued agent always wins (i.e., the social welfare is maximized). When the agents' values are non-identically distributed, analytically solving for equilibrium is notoriously difficult. For example, Vickrey [1961] posed the question of solving for equilibrium with two agents with values drawn uniformly from distinct intervals; this problem was finally resolved half a century later by Kaplan and Zamir [2012].

Price-of-anarchy style analysis allows us to make general statements about behavior in equilibrium without needing an analytical characterization of equilibrium. For example, a recent analysis of Syrgkanis and Tardos [2013] shows that the first-price auction's social welfare in equilibrium is at least an $e/(e-1) \approx 1.58$ approximation to the optimal social welfare, and moreover, this bound continues to hold if multiple items are sold simultaneously by independent first-price auctions. Importantly, this price-of-anarchy analysis sidesteps the intractability of solving for equilibrium and instead derives its bounds from simple best-response arguments.

3.1.1. Methods

Our analysis begins with the single-item, first-price auction. We break the analysis of welfare into two parts: first, relating an agents contribution to equilibrium welfare to her expected threshold prices for allocation, next relating the expected threshold prices for allocation to the revenue of the auction for any actions, not only in BNE.

First-Price Revenue. We then use the characterization of revenue in Bayes-Nash equilibrium of Myerson [1981] to re-purpose our welfare analysis for revenue. The same covering condition that holds for bidders' values also holds for their (positive) virtual values: if a bidder has a positive virtual value, her contribution to virtual welfare in equilibrium and expected prices for additional allocation combine to cover an $(e-1)/e$ fraction of her contribution to virtual welfare in the optimal mechanism. If the revenue of the mechanism covers the the prices agents see for additional allocation, then the virtual welfare from positive virtual valued agents approximates the revenue of the revenue-optimal mechanism.

If the impact of negative virtual-valued agents is small enough, equilibrium revenue will then approximate the optimal revenue. One such approach is to set monopoly reserve prices; another approach is to ensure agents face competition from agents of their same type.

General Auction Reduction. We extend our analysis to general auctions in two steps. First, we translate the payments in any auction into *equivalent bids*: the first-price bids or payments if the payment rule of the mechanism used first-price semantics. This allows us to reduce the optimization problem a bidder faces into the same problem a bidder in the first price auction faces. From this viewpoint we show that in a Bayes-Nash equilibrium of *any auction* an agent's contribution to welfare and her expected prices for additional allocation combine to approximate an $(e-1)/e$ fraction of her contribution to welfare in the optimal mechanism. Intuitively, either an agent's utility and hence contribution to welfare is high, or the price she has to pay for additional allocation is high relative to her value.

This leaves only the correspondence between the revenue of the general mechanism and the prices agents see for allocation.

Revenue Covering. If for any auction, the prices agents see for allocation can be bounded relative to the expected revenue, then we immediately get welfare and positive virtual surplus bounds, proportional to the bound of prices for allocation to revenue.

If the prices agents see for additional allocation correspond directly or approximately to the revenue of the mechanism, then combining across all agents implies that the welfare and revenue in Bayes-Nash equilibrium combine to approximate the welfare of the optimal mechanism. With reserve prices, considering only the agents with values above their reserve gives an approximation result to the optimal auction subject to the same reserves.

3.1.2. Results

For single-item and single-dimensional matroid auctions (where the feasibility constraint is given by a matroid set system), we give welfare and revenue price of anarchy results with both first-price and all-pay payment semantics. The first-price variants of these auctions (a) solicit bids, (b) choose an outcome to optimize the sum of the reported bids of served agents, and (c) charge the agents that are served their bids. These first-price results are compatible with reserve prices. The all-pay variants of these auctions (a) solicit bids, (b) choose an outcome to optimize the sum of the reported bids of served agents, and (c) charge all agents their bids.

Welfare. In first-price auctions, we give a simple proof that the price of anarchy for welfare is at most $e/(e-1)$. This result also extend to the generalized first-price position auction. For all-pay auctions in the above environments, we show the price of anarchy for welfare is at most $2e/(e-1)$. These proofs are new, though the results have appeared before [Syrgkanis and Tardos, 2013].

Revenue. For first-price auctions with monopoly reserves in regular, single-parameter environments, we show that the price of anarchy for revenue is at most $2e/(e-1)$. The same bound holds in the generalized first-price position auction with monopoly reserves. If instead of reserves each bidder must compete with at least one duplicate bidder, the price of anarchy for revenue in first-price auctions is at most $3e/(e-1)$; in all-pay auctions, at most $4e/(e-1)$.

Simultaneous Composition. We also show via an extension theorem that the above bounds hold when auctions are run simultaneously if agents are *unit-demand* and *single-valued* across the outcomes of the auctions.

Beyond the single-item first-price auction, we reduce the problem of analyzing the price of anarchy for welfare and positive virtual surplus to analyzing one property of the rules of the mechanism: *revenue covering*, corresponding to the relationship between the price of allocation and the expected revenue of the auction.

If a mechanism is approximately revenue covered, then we immediately gain price of anarchy results for both the welfare and the positive virtual surplus of the mechanism that differ from the first-price results only in the revenue covering approximation factor.

3.1.3. Related Work

Understanding welfare in games without solving for equilibrium is a central theme in the smooth games framework of Roughgarden [2009] and the smooth mechanisms extension of Syrgkanis and Tardos [2013]. A core principle of smoothness is that the precise manner or assumptions used in proving the smoothness property dictate how broadly the result extends. One view of our work is that we limit our proofs in just the right way that allows for extensions to revenue approximations.

Our framework refines the smoothness framework for Bayesian games in three notable ways. First, we decompose smoothness into two components, separating the specifics of a mechanism from the actions of a best-responding agent in any auction equilibrium. Second, because we focus on the optimization problem that individual bidders are facing, we can attain results that only hold for certain bidders — for instance, bidders with values above their reserve prices. Third, we only consider the Bayesian setting, which allows us to use the BNE characterization of Myerson [1981] to approximate revenue and relate other auctions to the first-price auction via equivalent bids.

A number of papers have looked at revenue guarantees for the welfare-optimal Vickrey-Clarke-Groves (VCG) mechanism in asymmetric settings. Hartline and Roughgarden [2009] show that VCG with monopoly reserves or duplicate bidders achieves revenue that is a constant approximation to the revenue optimal auction. Dhangwatnotai et al. [2010] show that the single-sample mechanism, essentially VCG using a single value from the distribution as a reserve, achieves approximately optimal revenue in broader settings. Roughgarden et al. [2012] showed that in broader environments, including matching settings, limiting the supply of items in relation to the number of bidders gives a constant approximation to the optimal auction.

In the economics literature, Kirkegaard [2009] shows that understanding the ratios of expected payoffs in equilibria of asymmetric auctions can be easier than understanding equilibrium and lead to insights about equilibria. Kirkegaard [2012] shows that some properties of distributions can be used to compare revenue of the first price auction to revenue of the second price auction. Lebrun [2006] and Maskin and Riley [2003] establish equilibrium uniqueness in the asymmetric setting with some assumptions on the distributions of agents.

3.2. Preliminaries

Bayesian Mechanisms. This chapter considers mechanisms for n single-dimensional agents with linear utility. Each agent has a private value for service, v_i, drawn independently from a distribution F_i over valuation space V_i. We write $\mathbf{F} = \prod_i F_i$ and $\mathbf{V} = \prod_i V_i$ to denote the joint value distribution and space of valuation profiles, respectively. A *mechanism* consists of a bid allocation rule $\tilde{\mathbf{x}}$ and a payment rule $\tilde{\mathbf{p}}$, mapping actions of agents to allocations and payments respectively. Each agent i draws their private value v_i from F_i and selects an action according to some strategy $s_i : V_i \to \mathrm{A}_i$, where A_i is the set of possible actions for i. We write $\mathbf{s} = (s_1, \ldots, s_n)$ to denote the vector of agents' strategies. Given the actions $\mathbf{a} = (a_1, \ldots, a_n)$ selected by each agent, the mechanism computes $\tilde{\mathbf{x}}(\mathbf{a})$ and $\tilde{\mathbf{p}}(\mathbf{a})$. Each agent's utility is $\tilde{u}_i(\mathbf{a}) = v_i\tilde{x}_i(\mathbf{a}) - \tilde{p}_i(\mathbf{a})$.

Mechanisms typically operate with constraints on permissible allocations. A *feasibility environment* specifies the set of feasible allocation vectors. Mechanisms for a feasibility environment must choose only allocations from the feasible set. The simplest example is a single-item auction, in which at most one person at a time can be served. This chapter assumes feasibility environments are *downward-closed*: if $(x_1, \ldots, x_k, \ldots, x_n)$ is feasible, so is $(x_1, \ldots, 0, \ldots, x_n)$. We will often consider the special case of *matroid* environments, in which the set of feasible allocations correspond to the independent sets of a matroid set system.

Given a strategy profile $\mathbf{s}$, we often consider the expected allocation and payment an agent faces from choosing some action $a_i \in \mathrm{A}_i$, with expectation taken with respect to other agents' values and actions induced by $\mathbf{s}$. We treat $\mathbf{s}$ as implicit and write $\tilde{x}_i(a_i) = \mathbf{E}_{\mathbf{v}_{-i}}[\tilde{x}_i(a_i, \mathbf{s}_{-i}(\mathbf{v}_{-i}))]$, with $\tilde{p}_i(a_i)$ and $\tilde{u}_i(a_i)$ defined analogously. Given $\mathbf{s}$ implicitly, we also consider values as inducing payments and allocations. We write $\mathbf{x}(\mathbf{v}) = \tilde{\mathbf{x}}(\mathbf{s}(\mathbf{v}))$ and $\mathbf{p}(\mathbf{v}) = \tilde{\mathbf{p}}(\mathbf{s}(\mathbf{v}))$, respectively. Furthermore, we can denote agent i's interim allocation probability and payment by $x_i(v_i) = \tilde{x}_i(s_i(v_i))$ and $p_i(v_i) = \tilde{p}_i(s_i(v_i))$. We define $u(\mathbf{v})$ and $u_i(v_i)$ similarly. In general, we use a tilde to denote outcomes induced by actions, and omit the tilde when indicating outcomes induced by values. We refer to $\tilde{\mathbf{x}}$ as the *bid allocation rule*, to distinguish it from $\mathbf{x}$, the *allocation rule*. We adopt a similar convention with other notation.

Bayes-Nash Equilibrium. A strategy profile $\mathbf{s}$ is in *Bayes-Nash equilibrium* (BNE) if for all agents i, $s_i(v_i)$ maximizes i's interim utility, taken in expectation with respect to other agents' value distributions $\mathbf{F}_{-i}$ and their actions induced by $\mathbf{s}$. That is, for all i, v_i, and alternative actions a': $\mathbf{E}_{\mathbf{v}_{-i}}[\tilde{u}_i(\mathbf{s}(\mathbf{v}))] \geq \mathbf{E}_{\mathbf{v}_{-i}}[\tilde{u}_i(a', \mathbf{s}_{-i}(\mathbf{v}_{-i}))]$.

Myerson [1981] characterizes the interim allocation and payment rules that arise in BNE. These results are summarized in the following theorem.

Theorem 3.1 (Myerson, 1981)**.** *For any mechanism and value distribution* $\mathbf{F}$ *in BNE,*

(1) *(monotonicity) The interim allocation rule $x_i(v_i)$ for each agent is monotone non-decreasing in v_i.*

(2) *(payment identity) The interim payment rule satisfies $p_i(v_i) = v_i x_i(v_i) - \int_0^{v_i} x_i(z)dz$.*

(3) *(revenue equivalence) Mechanisms and equilibria which result in the same interim allocation rule* $\mathbf{x}$ *must therefore have the same interim payments as well.*

Mechanism Design Objectives. We consider the problem of maximizing two primary objectives in expectation at BNE: utilitarian welfare and revenue. The revenue of a mechanism $\mathcal{M}$ is the total payment of all agents. Mechanism $\mathcal{M}$'s expected revenue for $\mathbf{v} \sim \mathbf{F}$ in a given Bayes-Nash equilibrium $\mathbf{s}$ is denoted $\textsc{Rev}(\mathcal{M}) = \mathbf{E}_v[\sum_i p_i(\mathbf{v})]$. The welfare of a mechanism $\mathcal{M}$ is the total utility of all participants including the auctioneer; denoted $\textsc{Welfare}(\mathcal{M}) = \textsc{Rev}(\mathcal{M}) + \mathbf{E}_v[\sum_i u_i(\mathbf{v})] = \mathbf{E}_v[v_i x_i(\mathbf{v})]$. We will also refer to welfare as "surplus."

Our welfare benchmark is the mechanism that always serves the highest valued feasible agents. That is, we seek to approximate $\textsc{Welfare}(\textsc{Opt}) = \mathbf{E}_{\mathbf{v}}[\max_{\mathbf{x}^*} \sum_i v_i x_i^*]$. This can be implemented via the Vickrey-Clarke-Groves (VCG) mechanism. We measure a mechanism $\mathcal{M}$'s welfare performance by the *Bayesian price of anarchy for welfare*, given by

$$\max_{\mathbf{F}, \mathbf{s} \in \mathrm{BNE}(\mathcal{M}, \mathbf{F})} \frac{\textsc{Welfare}(\textsc{Opt})}{\textsc{Welfare}(\mathcal{M})},$$

where $\mathrm{BNE}(\mathcal{M}, \mathbf{F})$ is the set of BNE for $\mathcal{M}$ under value distribution $\mathbf{F}$.

For revenue, we will make extensive use of the characterization of revenue in Myerson [1981] that follows from Theorem 3.1:

Lemma 3.2. *In BNE, the ex ante expected payment of an agent is*

$$\mathbf{E}_{v_i}\left[p_i(v_i)\right] = \mathbf{E}_{v_i}\left[\phi_i(v_i) x_i(v_i)\right],$$

where $\phi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$ *is the* virtual value *for value* v_i*. It follows that* $\textsc{Rev}(\mathcal{M}) = E_{\mathbf{v}}[\sum_i p_i(\mathbf{v})] = \mathbf{E}_{\mathbf{v}}[\sum_i \phi_i(v_i)x_i(\mathbf{v})]$.

Using this result, Myerson [1981] derives the revenue-optimal mechanism for any value distribution $\mathbf{F}$. This mechanism is parameterized by the value distribution $\mathbf{F}$, and the optimality is in expectation over $\mathbf{v} \sim \mathbf{F}$. We specifically consider distributions with no point masses where $\phi_i(v_i)$ is monotone in v_i for each i. Such distributions are said to be *regular*. If each agent has a regular distribution, then the revenue-optimal mechanism selects the allocation which maximizes $\sum_i \phi_i(v_i)x_i(\mathbf{v})$. We will seek to characterize the *Bayesian price of anarchy for revenue*,

$$\max_{\mathbf{F}\in\mathcal{R},\mathbf{s}\in\text{BNE}(\mathcal{M},\mathbf{F})} \frac{\textsc{Rev}(\textsc{Opt}_{\mathbf{F}})}{\textsc{Rev}(\mathcal{M})},$$

where $\mathcal{R}$ is the set of regular distributions and $\textsc{Opt}_{\mathbf{F}}$ is the Bayesian revenue-optimal mechanism for value distribution $\mathbf{F}$.

3.3. Single-Item First Price Auction

We begin by analyzing the welfare of the first-price auction, showing that it always approximates the welfare of the welfare-optimal mechanism. Subsequent results will use this proof as a template.

Theorem 3.3. *The welfare in any BNE of the first price auction is at least an* $\frac{e}{e-1}$*-approximation to the welfare of the welfare-optimal mechanism.*

Our proof proceeds in two steps. First, we analyze the interim optimization problem faced by every bidder. We quantify an intuitively obvious tradeoff: either that bidder can obtain high expected utility, or the threshold bid below which they go unallocated tends

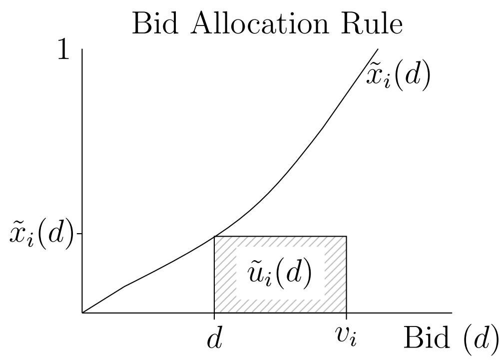

Figure 3.1. For any bid d, the area of a rectangle between $(d, \tilde{x}_i(d))$ and $(v_i, 0)$ on the bid allocation rule is the expected utility $\tilde{u}_i(d)$. The BNE bid b_i is chosen to maximize this area.

to be high. Second, we note that these threshold bids are a lower bound on revenue. This implies a tradeoff between revenue (seller welfare) and utilities (buyer welfare):

$$\textsc{Util}(\text{FPA}) + \textsc{Rev}(\text{FPA}) \geq \tfrac{e-1}{e}\textsc{Welfare}(\textsc{Opt}). \tag{3.1}$$

This in turn will imply the theorem.

What to bid? We now develop ideas needed to make this analysis formal. Consider the optimization problem faced by a bidder i with value v_i in the first price auction. Bidder i's expected utility for a possible bid d is $\tilde{u}_i(d) = (v_i - d)\tilde{x}_i(d)$, where $\tilde{x}_i(d)$ is the interim bid allocation rule faced by the bidder. Let b_i be her best response bid given her value v_i. That is, b_i maximizes $\tilde{u}_i(d)$. If we plot the bid allocation rule $\tilde{x}_i(d)$ for any alternate bid d, then $\tilde{u}_i(b_i)$ is precisely the area of the rectangle in the lower right; see Figure 3.1.

When other bidders have realized values and submitted bids, bidder i wins only if her bid exceeds both her reserve and the bids of other agents. Consequently the price a bidder must pay to win is $\tau_i(\mathbf{v}_{-i}) = \max_{j \neq i} b_j(v_j)$; we will formally refer to it as her *threshold bid*.

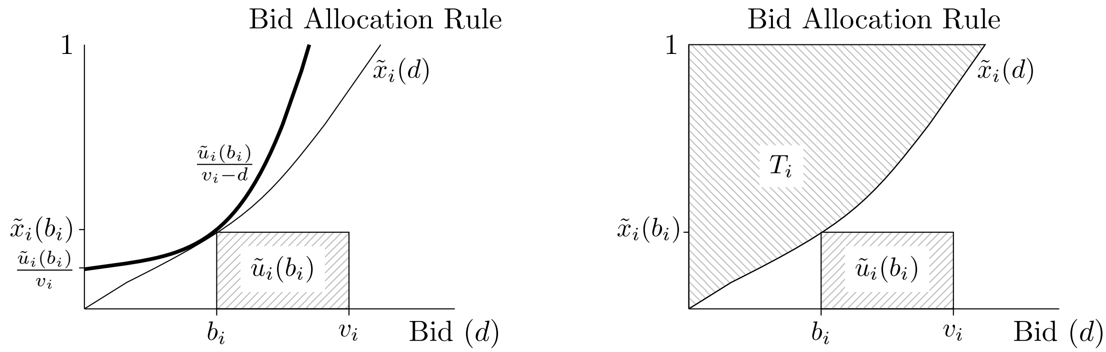

(a) As b_i is a best-response to the actions of other agents, the indifference curve $\tilde{u}_i(b_i)/(v_i - d)$ upper bounds $\tilde{x}_i(d)$.

(b) The expected threshold T_i is the area above the allocation rule.

Figure 3.2. The indifference curve and expected threshold for a bidder in the first-price auction.

As we are in the Bayesian setting, a bidder is not reacting to this threshold, but is acting in expectation over the types and actions of her competitors. Consequently, each allocation probability has a threshold price a bidder must pay to secure that chance of being allocated. It is therefore convenient to refer to thresholds in these terms. Let $\tau_i(x)$ refer to the smallest bid that achieves allocation of at least x, hence $\tau_i(x) = \min\{\ b \mid \tilde{x}_i(b) \geq x\}$. Note that $\tau_i(x)$ is the price an agent faces to attain allocation x. Also, note that $\tau_i(x)$ is effectively the inverse of the cumulative distribution function of the highest bids from all other agents.

The expected threshold is $T_i = \int_0^1 \tau_i(z)\ dz$. This is illustrated in Figure 3.2b. The expected threshold is the quantity we will be using to relate welfare contributions in equilibrium and the optimal allocation.

Relating Contributions to First-Price and Optimal Welfare: We will now approximate each bidder's contribution to the optimal welfare individually, using the bidder's utility in the first-price auction and a fraction of the revenue in the first-price auction.

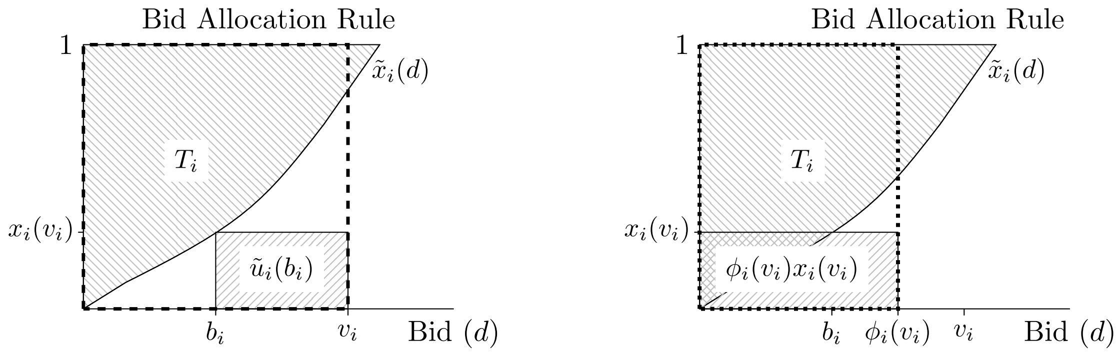

(a) Lemma 3.4 shows the shaded areas cover a $(e-1)/e$ fraction of the dashed box, bidder i's value v_i and hence maximum contribution to the optimal welfare.

(b) Lemma 3.6 shows the shaded areas cover an $(e-1)/e$ fraction of i's virtual value $\phi_i(v_i)$.

Figure 3.3. Value covering and virtual value covering in the first-price auction.

In these terms, the steps to prove Theorem 3.3 are:

(1) *Value Covering*: Each bidder's utility in the FPA and expected threshold together approximate her value. (Lemma 3.4)

(2) *Revenue Covering*: The revenue of the FPA approximates the expected thresholds of all agents. (Lemma 3.5)

The final approximation result follows by summing the value covering inequality across agents, taking expectation over values, and combining with revenue covering.

Lemma 3.4 (Value Covering)**.** *In any BNE of* FPA*, for any bidder i with value v_i,*

$$u_i(v_i) + T_i \geq \tfrac{e-1}{e} v_i. \tag{3.2}$$

Proof. We will prove value covering in two steps: first, by developing a lower bound $\underline{T}$ on the expected threshold T_i; second, by optimizing to get the worst such bound. The proof can also be done using a modification of the first-price bid deviation approach of Syrgkanis and Tardos [2013].

Lowerbounding T. In best responding, bidder i chooses an action which maximizes her utility. If b_i is a best response bid, then for any alternate bid d, $\tilde{u}_i(b_i) \geq (v_i - d)\tilde{x}_i(d)$, hence $\tilde{x}_i(d) \leq \frac{\tilde{u}_i(b_i)}{v_i - d}$. The upperbound $\frac{\tilde{u}_i(b_i)}{v_i - d}$ is an indifference curve for bidder i; it is the alternate bid allocation rule that would lead to her being indifferent over all reasonable bids (see Figure 3.2a). The area above the indifference curve gives a lower bound on the expected threshold; call the lower bound $\underline{T}_i = \int_0^1 \max(0, v_i - u_i(v_i)/z)\ dz$.

Worst-case $\underline{T}_i$. Evaluating the integral for $\underline{T}_i$ gives $\underline{T}_i = v_i - u_i(v_i)(1 - \ln \frac{u_i(v_i)}{v_i})$, hence $u_i + \underline{T}_i = v_i + u_i \ln \frac{u_i(v_i)}{v_i}$. Holding v_i fixed and minimizing with respect to $u_i(v_i)$ yields a minimum at $u_i(v_i) = v_i/e$, hence $u_i(v_i) + \underline{T}_i \geq \frac{e-1}{e} v_i$ and as desired

$$u_i(v_i) + T_i \geq \tfrac{e-1}{e} v_i. \tag{3.3}$$

Note that this analysis depended only on the fact that bidder i was best responding to a bid distribution - this will allow us to generalize the lemma later. □

We now show that in the first price auction the expected revenue is greater than the expected threshold, which we can then combine with value covering to give a welfare approximation result. While value covering depended only on equilibrium bidding behavior, revenue covering will only depend on the form of the first price auction, and will thus hold for arbitrary (not necessarily BNE) bidding strategies.

Lemma 3.5 (Revenue Covering)**.** *For any agent i,* $\textsc{Rev}(\text{FPA}) \geq T_i$*.*

Proof. The revenue of a first price auction is the expected highest bid, and T_i is the expected highest bid from all agents except i. The result follows. $\square$

We now combine value and revenue covering to attain an approximation to the optimal welfare.

Proof of Theorem 3.3. First, for each agent i with value v_i, combining value and revenue covering gives

$$u_i(v_i) + \textsc{Rev}(FPA) \geq \frac{e-1}{e} v_i.$$

Let $\mathbf{x}^*(\mathbf{v})$ be the optimal allocation. For any agent, $x_i^*(v_i) \leq 1$, hence

$$u_i(v_i) + x_i^*(v_i)\textsc{Rev}(FPA) \geq x_i^*(v_i)u_i(v_i) + x_i^*(v_i)\textsc{Rev}(FPA) \geq x_i^*(v_i)\frac{e-1}{e} v_i.$$

Summing over all agents and taking expectation over values gives

$$\textsc{Util}(\text{FPA}) + \textsc{Rev}(\text{FPA}) \geq \frac{e-1}{e}\textsc{Welfare}(\textsc{Opt}).$$

As $\textsc{Welfare}(\text{FPA}) = \textsc{Util}(\text{FPA}) + \textsc{Rev}(\text{FPA})$, $\textsc{Welfare}(\text{FPA})$ is then an $e/(e-1)$ approximation to $\textsc{Opt}$. $\square$

3.3.1. Welfare Lower Bounds

The approximation results we have given in this section for the single-item first-price auction are not known to be tight. For welfare with no reserves, the price of anarchy can be as bad as 1.15; we give such an example in Appendix A.1. Note the large gap between this lower bound and the upper bound of $\frac{e}{e-1} \approx 1.58$ from Theorem 3.3 and Syrgkanis and Tardos [2013].

Beyond a single auction, Christodoulou et al. [2013] have shown that the $\frac{e}{e-1}$ bound is tight for the simultaneous composition of item auctions when bidders have submodular valuations.

3.4. Revenue Approximation

Our welfare result hinged on the complementary relationship between the utility of a bidder and the bids of other bidders in the mechanism. Using this relationship to directly bound revenue is not as straightforward. The results of Myerson [1981], however, provide another method of accounting for each bidder's impact on revenue, their *virtual value.* Using virtual values in place of utilities allows us to adapt our method for proving welfare results to the objective of revenue.

3.4.1. Revenue

The welfare of a mechanism can be expressed as the expected total value of agents who are served. Myerson [1981] demonstrated a similar characterization of revenue in terms of the expected total virtual value, reducing the problem of revenue maximiation to welfare maximization. We adopt a similar approach, using virtual values to reuse our tools from welfare analysis.

We will begin by showing the analogue of value covering, *virtual* value covering, in which each bidder's positive contribution to equilibrium virtual welfare and expected threshold bid combine to approximate her contribution to the optimal revenue, which by Myerson [1981] is the optimal virtual welfare.

Lemma 3.6 (Virtual Value Covering)**.** *In any BNE of* FPA*, for any bidder i with value v_i such that $\phi_i(v_i) \geq 0$,*

$$\phi_i(v_i)x_i(v_i) + T_i \geq \tfrac{e-1}{e}\phi_i(v_i). \tag{3.4}$$

Proof. First, note that the virtual value is defined as $\phi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$. The second term is always nonnegative, so $\phi_i(v_i) \leq v_i$. Moreover, note that $v_i x_i(v_i) \geq u_i(v_i)$, so

$$\tfrac{1}{\phi_i(v_i)}\left(\phi_i(v_i)x_i(v_i) + T_i\right) \geq \tfrac{1}{v_i}\left(v_i x_i(v_i) + T_i\right) \geq \tfrac{1}{v_i}\left(u_i(v_i) + T_i\right) \geq \tfrac{e-1}{e}.$$

See Figure 3.3b for an illustration. Intuitively, value covering captured the idea that the expected threshold made up the difference between an agent's utility and their value. The difference between virtual surplus and virtual value is proportionally smaller, so the expected threshold can cover that gap as well. □

This approximation holds for agents with positive virtual values. To mitigate the impact of agents with negative virtual values, one approach is to implement reserve prices.

3.4.2. Reserve Prices

We consider first price auctions with a reserve set to exclude exactly the agents with negative virtual values. Under the assumption that bidders' value distributions are regular, it suffices to set monopoly reserves at $r_i = \phi_i^{-1}(0)$.

However, introducing reserves to a first price auction affects the relationship between threshold bids and revenue, and therefore revenue covering does not hold. Introducing reserves to an auction inflates the threshold bids an agent sees - when all other bidders are bidding low, the reserve binds as a minimum bid the agent must place. Unlike other agents' bids, this reserve does not directly correspond to revenue - in fact, holding a bid distribution fixed and introducing reserves will actually lower the revenue of a first-price auction.

This inflation only occurs when all other bidders bid below the reserve. When this is not the case, an agent's threshold bid still corresponds to revenue. In other words, a weaker version of revenue covering still holds for the auction with reserves.

To capture this notion, we will define exactly the portion of the expected threshold still corresponding to bids. As before, $\tau_i(x)$ refers to the smallest bid that achieves allocation of at least x. Now, however, note that below $\tilde{x}_i(r_i)$ no longer corresponds to the inverse cumulative distribution function of the highest bids from all other agents. Above this point, however, the correspondence to bids remains the same as without reserves. For any bid y, we therefore define the *expected threshold above* y to be $T_i^y = \int_{\tilde{x}_i(y)}^1 \tau_i(z)\, dz$. With $y = r_i$, the expected threshold above r_i is precisely the portion of the expected threshold generated by bids. See Figure 3.4a for an illustration.

We now prove a weaker notion of value covering using $T_i^{r_i}$ and the payment of an agent to make up for the loss of revenue corresponding to threshold bids below the reserve. See Figure 3.4b for an illustration.

Lemma 3.7 (Value Covering with Reserves)**.** *In any BNE of* $\text{FPA}_{\mathbf{r}}$*, for any bidder i with value $v_i \geq r_i$,*

$$v_i x_i(v_i) + T_i^{r_i} \geq \tfrac{e-1}{e} v_i. \tag{3.5}$$

Proof. In BNE, if your value is above the reserve, you will bid at least the reserve. Then $T_i^{r_i} + p_i(v_i) = T_i^{r_i} + b_i(v_i)x_i(v_i) \geq T_i^{r_i} + r_i\tilde{x}_i(r_i) = T_i$. The proof mimics that of Lemma 3.4, with $T_i^{r_i} + p_i(v_i)$ in place of T_i. $\square$

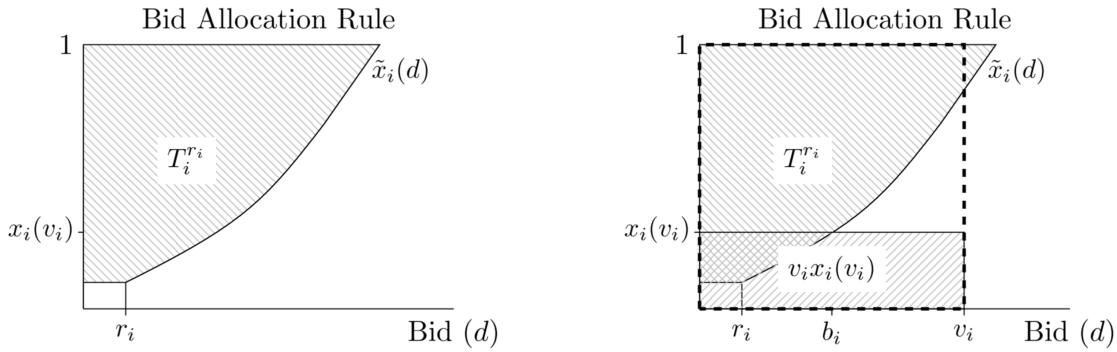

(a) In a first-price auction with reserve r_i, the threshold above r_i, $T_i^{r_i}$, only includes the thresholds when greater than r_i, which corresponds to the case that the threshold is from a bid from another agent rather than the reserve.

(b) Lemma 3.7 shows that $v_i x_i(v_i)$ and $T_i^{r_i}$ cover an $(e-1)/e$ fraction of i's value v_i, because $v_i x_i(v_i) - u_i(v_i) = b_i x_i(v_i)$ covers $T_i - T_i^{r_i}$

Figure 3.4. The effect of a reserve price on the first-price expected thresholds and value covering.

As previously mentioned, the thresholds for bidder i above r_i correspond to bids from other agents. It follows that this portion of i's expected threshold corresponds to revenue. We can formalize this with the following lemma:

Lemma 3.8 (Revenue Covering with Reserves)**.** *For any agent i,* $\textsc{Rev}(\text{FPA}_{\mathbf{r}}) \geq T_i^{r_i}$.

Just as in the no-reserves case, value covering and revenue covering with reserves combine to give a welfare approximation result relative to the welfare-optimal auction with those same reserves as our benchmark. The weaker version of value covering requires the use of an extra copy of the auction revenue, adding a factor of two to the bound.

Theorem 3.9. *The welfare in any BNE of the first price auction with reserves* $\mathbf{r}$ *is at least an* $\frac{2e}{e-1}$*-approximation to the welfare of the welfare-optimal mechanism with those same reserves,* $\textsc{Opt}_{\mathbf{r}}$.

Proof. Take any bidder i with $v_i \geq r_i$. Combining value covering and revenue covering with reserves gives

$$v_i x_i(v_i) + \textsc{Rev}(\text{FPA}_{\mathbf{r}}) \geq \frac{e-1}{e} v_i. \tag{3.6}$$

Let $\mathbf{x}^*(\mathbf{v})$ be the optimal allocation. As before, $x_i^*(v_i) \leq 1$, hence

$$v_i x_i(v_i) + x_i^*(v_i)\textsc{Rev}(\text{FPA}_{\mathbf{r}}) \geq x_i^*(v_i) v_i x_i(v_i) + x_i^*(v_i)\textsc{Rev}(\text{FPA}_{\mathbf{r}}) \geq x_i^*(v_i)\frac{e-1}{e} v_i. \tag{3.7}$$

For any agents with $v_i < r_i$, (3.7) holds as well, as $x_i^*(v_i) = 0$ for such agents. We can therefore sum (3.7) over all agents and take expectation over values to get $\textsc{Util}(\text{FPA}_{\mathbf{r}}) + \textsc{Rev}(\text{FPA}_{\mathbf{r}}) \geq \frac{e-1}{e}\textsc{Welfare}(\textsc{Opt}_{\mathbf{r}})$. As $\textsc{Welfare}(\text{FPA}_{\mathbf{r}}) = \textsc{Util}(\text{FPA}_{\mathbf{r}}) + \textsc{Rev}(\text{FPA}_{\mathbf{r}})$, $\textsc{Welfare}(\text{FPA}_{\mathbf{r}})$ is then an $2e/(e-1)$ approximation to $\textsc{Opt}_{\mathbf{r}}$. □

To derive a revenue result, we can make the necessary modification to virtual value covering to get:

Lemma 3.10 (Virtual Value Covering with Reserves)**.** *In any BNE of* $\text{FPA}_{\mathbf{r}}$*, for any bidder i with value $v_i \geq r_i$ such that $\phi_i(v_i) \geq 0$,*

$$\phi_i(v_i) x_i(v_i) + T_i^{r_i} \geq \tfrac{e-1}{e}\phi_i(v_i). \tag{3.8}$$

Proof. The lemma follows from value covering with reserves exactly as Lemma 3.6 follows from Lemma 3.4. □

With regular value distributions, adding the monopoly reserves $r_i^* = \phi_i^{-1}(0)$ to the auction excludes exactly the agents with negative virtual values. It follows that for such reserves, (3.8) holds whenever $v_i \geq r_i^*$.

In regular environments the optimal mechanism for revenue allocates the item to the agent with the highest positive virtual value. Approximating the optimal revenue is therefore equivalent to approximating this agent's expected virtual value. By adapting revenue covering and virtual value covering to first price auction with reserves, though, we are able to treat this quantity much as we treated welfare, yielding the following:

Theorem 3.11. *The revenue in any BNE of the first price auction with monopoly reserves (*$\mathrm{FPA}_{\mathbf{r}^*}$*) and agents with regularly distributed values is at least a* $2e/(e-1)$ *approximation to the revenue of the optimal auction.*

Proof. The proof mirrors that of Theorem 3.9. First, take any bidder i with $v_i \geq r_i^*$. Combining Lemmas 3.8 and 3.10 yields:

$$\phi_i(v_i)x_i(v_i) + \textsc{Rev}(\mathrm{FPA}_{\mathbf{r}}) \geq \frac{e-1}{e}\phi_i(v_i). \tag{3.9}$$

Let $\mathbf{x}^*(\mathbf{v})$ be the optimal allocation. As before, $x_i^*(v_i) \leq 1$, hence

$$\phi_i(v_i)x_i(v_i) + x_i^*(v_i)\textsc{Rev}(FPA) \geq x_i^*(v_i)\frac{e-1}{e}\phi_i(v_i). \tag{3.10}$$

For any agents with $v_i < r_i$, (3.10) holds as well, as $x_i^*(v_i) = 0$ for such agents. We can therefore sum (3.10) over all agents and take expectation over values to get $\textsc{Rev}(\mathrm{FPA}_{\mathbf{r}^*}) + \textsc{Rev}(\mathrm{FPA}_{\mathbf{r}^*}) \geq \frac{e-1}{e}\textsc{Rev}(\textsc{Opt})$, yielding the desired result. $\square$

3.4.3. Duplicate bidders

Another approach to mitigating the impact of negative virtual-valued agents is to ensure each agent faces adequate competition. Bulow and Klemperer [1996] show that this intuition guarantees approximately optimal revenue in regular, symmetric, single-item settings.

We show the same intuition holds for first-price auctions in asymmetric settings: if each bidder must compete with at least $k-1$ other bidders with values drawn from her same distribution, revenue is approximately optimal compared to the revenue-optimal mechanism (including the duplicate bidders). We say such a setting satisfies *k-duplicates.*

Theorem 3.12. *The revenue in any BNE of the first price auction with k-duplicates (FPA_k) and agents with regularly distribution values is at least a $\frac{k}{k-1}\frac{2e}{e-1}$ approximation to the revenue of the optimal auction.*

The proof is included in Appendix A.3.

3.4.4. Revenue Lower Bounds

For revenue, the approximation ratio can be at least as bad as 2, using the same lower bound Hartline and Roughgarden [2009] show for VCG with monopoly reserves. The example has two bidders, one with deterministic value 1, the other with value drawn according to the equal revenue distribution with support over $[1, H]$ for some large H. With a light perturbation of the distribution the monopoly price for the second bidder is 1. Assuming ties go to bidder 2, an equilibrium exists where both agents bid 1, giving revenue of 1. The optimal auction however can set a reserve of H for the second bidder and sell to the first bidder at price 1 if the reserve is not met, achieving a revenue of 2 as H grows.

3.5. Beyond Single-Item

Two main ideas drove the single-item welfare (Theorem 3.3) and revenue (Theorem 3.11) results. The first idea, value covering (resp. virtual value covering), captured the tradeoff between an agent's threshold bid and their utility (resp. virtual surplus). This idea depends only on a bidder's interim optimization problem, which is the same in every pay-your-bid auction. We can extend the single-item proof to get:

Lemma 3.13 (Pay-Your-Bid Value Covering)**.** *In any BNE of a pay-your-bid auction, for any bidder i with value v_i,*

$$u_i(v_i) + T_i \geq \tfrac{e-1}{e} v_i. \tag{3.11}$$

The second idea, revenue covering, captured the correspondence between threshold bids and mechanism revenue. With a general feasibility constraint, revenue covering in the single-item sense may not hold. However, a weaker, parameterized notion might, and when it does we will be able to derive similar, parameterized results.

Definition 3.1 (μ-Revenue Covering)**.** A mechanism $\mathcal{M}$ is *μ-revenue covered* if for any (implicit) distribution of actions and feasible allocation $\mathbf{x}'$,

$$\mu \textsc{Rev}(\mathcal{M}) \geq \sum\nolimits_i T_i x_i'.$$

There are two key differences between Definition 3.1 and its single-item counterpart. First, the additional parameter μ allows the relationship between revenue and thresholds to be weaker. Second, we require a sum on the right-hand side because multiple agents might be feasibly allocated.

As in the single-item setting, revenue covering and value covering imply approximation results. We can combine revenue and value covering using the logic that drove the proof of Theorem 3.3. Summing inequality (3.11) over all bidders, applying revenue covering, and taking expectations with respect to $\mathbf{v}$ yields the following:

Theorem 3.14. *The welfare of any μ-revenue covered pay-your-bid mechanism is a $\mu\frac{e}{e-1}$-approximation to the welfare of any other mechanism.*

Note that a tighter analysis of value-covering making use of the parameter μ can give a bound of $\frac{\mu}{1-e^{-\mu}}$ (see Syrgkanis and Tardos [2013]). We do not include the analysis because it is not extend to results for revenue.

The virtual value covering results for single-item auctions with and without reserves hold also without modification in general pay-your-bid auctions. They require only value covering, pay-your-bid semantics, and Myerson's virtual value characterization, all of which are agnostic to feasibility constraints. For example:

Lemma 3.15. *In any BNE of any pay-your-bid auction with reserves $\mathbf{r}$, for any bidder i with value $v_i \geq r_i$ and $\phi_i(v_i) \geq 0$,*

$$\phi_i(v_i)x_i(v_i) + T_i^{r_i} \geq \tfrac{e-1}{e}\phi_i(v_i).$$

Because of Lemma 3.15, any μ-revenue covered mechanism achieves approximately optimal revenue. As with Theorem 3.14, we can sum, apply revenue covering, and take expectations to yield:

Theorem 3.16. *The revenue of any μ-revenue covered pay-your-bid mechanism with regular bidders and monopoly reserves is a $\mu\frac{e}{e-1}$-approximation to the welfare of any other mechanism.*

We now show how to derive revenue covering results for environments beyond single-item auctions. We will see in Section 3.5.1 that the pay-your-bid auction with a matroid feasibility constraint is 1-revenue covered like its single-item counterpart, but beyond matroid feasibility constraints may not be revenue covered. In combinatorial settings, using a greedy algorithm instead of the optimal algorithm gives a *better* revenue covering result (Section 3.5.2). In Section 3.5.3 we discuss the revenue covering of the generalized-first-price position auction, and in Section 3.5.4 we discuss pay-your-bid auctions with a discretized bid space. In all cases, welfare and revenue results follow as corollaries.

3.5.1. First Price Matroid Auctions

In the single-item setting, both the revenue and the losers' critical bids were the bid of the unique winner. In matroid settings, multiple winners complicates this relationship. Using standard matroid properties, we show how to untangle the relationship between critical bids and winners' payments. We get as a corollary approximation guarantees identical to those achieved in the single-item case. In particular, we derive:

Lemma 3.17. *The first-price auction is 1-revenue covered in any matroid feasibility environment.*

As corollaries, we have:

Theorem 3.18. *The welfare of the first price matroid auction is at least an $e/(e-1)$-approximation to that of any other mechanism.*

Theorem 3.19. *For the first price matroid auction with monopoly reserves and regular bidders, the revenue of any BNE is at least a $2e/(e-1)$-approximation to that of any other mechanism.*

We first develop the matroid-specific tools we need to relate bids (and therefore revenue) to the thresholds agents see. To do so, we use a property related to a result from Talwar [2003] for VCG in matroid environments.

Lemma 3.20. *For any strategy profile* $\mathbf{s}$*, value profile* $\mathbf{v}$ *and feasible allocation* $\mathbf{x}'$*,*

$$\sum_i s_i(v_i)x_i(\mathbf{v}) \geq \sum_i \tau_i(\mathbf{v}_{-i})x'_i. \tag{3.12}$$

The proof is based on the following matroid property:

Lemma 3.21 (Replacement Property)**.** *Let S_1 and S_2 be independent sets of size k in a matroid $\mathcal{M}$. Then there is a bijective function $f : S_2 \setminus S_1 \to S_1 \setminus S_2$ such that, for every $i \in S_2 \setminus S_1$, the set $(S_1 \setminus \{f(i)\}) \cup \{i\}$ is independent in $\mathcal{M}$.*

Proof of Lemma 3.20. Because subsets of feasible allocations are feasible, threshold bids are nonnegative, so we only need consider allocations $\mathbf{x}'$ which are bases. Let S and S' be sets served by $\mathbf{x}$ and $\mathbf{x}'$, respectively. Since bids are nonnegative, it follows that S and S' are the same size. By Lemma 3.21, there exists a bijection f from $S' \setminus S$ to $S \setminus S'$ with the replacement property in the lemma. For each $i \in S' \setminus S$, $s_{f(i)}(v_{f(i)}) \geq \tau_i(\mathbf{v}_{-i})$, as if i bids above $s_{f(i)}$, then $(S \setminus \{f(i)\}) \cup \{i\}$ would be optimal and therefore i would be allocated in

BNE. For each $i \in S' \cap S$, i was served in $x(v)$, it must be that $s_i(v_i) \geq \tau_i(\mathbf{v}_{-i})$. The result follows by summing over i. □

Note that this proof extends to any auction in a matroid environment where agents submit bids, and the mechanism selects a basis maximizing the sum of selected bids. We will make use of this lemma once more when discussing all-pay auctions.

Having established a relationship between threshold bids and bids made by the basis selected by the mechanism, proving revenue covering is simply a matter of taking expectations.

Lemma 3.22. *The first-price matroid auction is 1-revenue covered.*

Proof. Consider some alternate allocation $\mathbf{x}'$ and action profile $\mathbf{a}$. By the mechanism's payment scheme and Lemma 3.20,

$$\text{Rev}(\mathcal{M}) = \mathbf{E}_v\left[\sum\nolimits_i s_i(v_i)x_i(v)\right] \geq \mathbf{E}_v\left[\sum\nolimits_i \tau_i(\mathbf{v}_{-i})x'_i\right] = \sum\nolimits_i \mathbf{E}_v\left[\tau_i(\mathbf{v}_{-i})\right]x'_i.$$

The result follows from the observation that $\mathbf{E}_v[\tau_i^b(\mathbf{v}_{-i})] = T_i$. □

Having proven revenue covering, value covering, and virtual value covering with reserves, Theorems 3.18 and 3.19 follow from the analysis in the previous section.

3.5.2. Combinatorial Pay-Your-Bid Auctions

In this section we consider the *single-minded combinatorial auction* setting. We analyze two different pay-your-bid mechanisms: one which allocates to the bid-maximizing set of bidders, and one which allocates greedily. The former mechanism can fall short of optimal by a factor of m, the number of items. It is known from Lucier and Borodin [2010], though, that the former mechanism has a price of anarchy of $\sqrt{m}\frac{e}{e-1}$. We frame these two results in terms of revenue covering - the former mechanism is not μ-revenue covered for any $\mu < m$, whereas the greedy mechanism is $\sqrt{m}$-revenue covered. (As a corollary, the revenue of the greedy mechanism with monopoly reserves is within a factor of $(\sqrt{m}+1)\frac{e}{e-1}$ of optimal, when bidders have regular distributions.) This comparative analysis suggests the power of revenue covering as a design objective.

Single-Minded Combinatorial Auctions. In a single-minded combinatorial auction feasibility environment, there are m indivisible items. Each bidder i wishes to acquire a set of items $\mathcal{S}_i$ - she receives value v_i for receiving any superset of $\mathcal{S}_i$, and value 0 otherwise. A $0-1$ allocation vector $\mathbf{x}$ is feasible if and only if it is possible to simultaneously allocate $\mathcal{S}_i$ to each i with $x_i = 1$.

Highest-Bids-Win Mechanism. The *highest-bids-win mechanism* allocates to a feasible set of agents which maximizes the sum of bids of winners. This is the optimal allocation in the absence of incentives. With incentives, the following example shows that this allocation scheme performs poorly under pay-your-bid semantics:

Consider a setting with m items and $m+2$ bidders. The first m bidders each want a single item - bidder j wants item j, each with a value for allocation of 1, deterministically. The final two bidders, meanwhile, each want the grand bundle of all m items, with values of $1+\epsilon$, again deterministically. With appropriate tiebreaking, it is a BNE for each of the first m agents to bid 0, while the final two bidders bid $1+\epsilon$. The optimal social welfare and revenue are both attained by selling to the first m bidders, for welfare and revenue of m, yielding a price of anarchy of m for both welfare and revenue.

This equilibrium also shows that the highest-bids-win mechanism isn't μ-revenue covered for any $\mu < m$. First, note that the total revenue of the mechanism is $1+\epsilon$. Next, consider the feasible allocation $\mathbf{x}'$ which allocates the first m bidders. For these bidders, they must bid at least $1+\epsilon$ to get allocated, at which point they get allocated with probability 1. It follows that for such bidders, $T_i = 1+\epsilon$, and hence $\sum_i T_i x'_i = (1+\epsilon)m$, m times the mechanism's revenue.

Greedy Allocation Mechanism. The *greedy allocation mechanism* sorts bidders in non-increasing order of $s_i/\sqrt{|\mathcal{S}_i|}$, where s_i is the bid of agent i. In this order, it then adds bidders to the winning set whenever doing so would maintain feasibility. The greedy allocation rule is suboptimal in the absence of incentives. With pay-your-bid semantics, however, it is $\sqrt{m}$-revenue covered. In other words, the greedy rule possesses better incentive properties than highest-bids-win. The welfare result of Lucier and Borodin [2010] and a revenue result (with monopoly reserves and regular bidders) follow as a corollary.

The proof proceeds in three steps. First, we note that, absent any inefficiency caused from incentives, the greedy allocation rule only loses a factor of $\sqrt{m}$ compared to the highest-bids-win rule. Formally:

Lemma 3.23. *Let* $\mathbf{x}(\mathbf{s})$ *be the allocation selected by the greedy allocation rule, and* $\mathbf{x}'$ *be any other feasible allocation. Then*

$$\sum\nolimits_i s_i \mathbf{x}(\mathbf{s}) \geq \sqrt{m} \sum\nolimits_i s_i \mathbf{x}'.$$

Refer to Lehmann et al. [2002] for a proof.

Next, we note that the greedy allocation rule lacks a pathology that ruined the welfare under highest-bids-win in our example. In the example, the high bids of the $(m+1)$st and $(m+2)$nd bidders discouraged participation from the other bidders - individually, a bidder would have to bid $1+\epsilon$ to win. As a group, though, the losing bidders could have all bid slightly more and won. The difficulty of allocation was not accurately reflected by the we threshold bids. The greedy allocation rule lacks this problem. Formally:

Lemma 3.24. *Let* $\mathbf{s}$ *be a profile of bids, let* $\tau_i(\mathbf{s}_{-i})$ *be the critical price of bidder* i *under the greedy allocation rule, and let* $\mathbf{x}(\mathbf{s})$ *be the corresponding allocation function. Further let* $\mathbf{s}'$ *be any profile of bids such that if* $x_i(\mathbf{s}) = 1$*, then* $s'_i = s_i$*, and if* $x_i(\mathbf{s}) = 0$*, then* $s'_i \leq \tau_i(\mathbf{s}_{-i})$*. Then* $\mathbf{x}(\mathbf{s}) = \mathbf{x}(\mathbf{s}')$*.*

Proof. Imagine changing $\mathbf{s}$ to $\mathbf{s}'$ by increasing one loser's bid at a time. Each time we increase a bid, say, of bidder i, two things remain true: (1) i still loses - as long as $s'_i \leq \tau_i(\mathbf{s}_{-i})$, i is after some other winning agent in the ordering, who wins some item in $\mathcal{S}_i$. (2) the threshold of every other losing agent i' remains unchanged - i' still wins if and only if $s_{i'}$ high enough to precede in the ordering all winners who get items in their set, which is not affected by i's bid as long as i still loses. □

Finally, we note that these two properties are enough to guarantee $\sqrt{m}$-revenue covering.

Lemma 3.25. *The pay-your-bid greedy combinatorial auction is* $\sqrt{m}$*-revenue covered.*

Proof. We argue that for any strategy profile $\mathbf{s}$, value profile $\mathbf{v}$, and alternate allocation $\mathbf{x}'$,

$$\sum_i s_i(v_i) x_i(\mathbf{v}) \geq \tfrac{1}{\sqrt{m}} \sum_i \tau_i(\mathbf{v}_{-i}) x'_i.$$

Taking the expectation of both sides yields the desired inequality.

Let $\mathbf{s}'$ be a vector of bids where losers under $\mathbf{s}(\mathbf{v})$ bid $\tau_i(\mathbf{v})$, while winners bid as before. The following inequalities hold, with justifications after.

$$\begin{aligned}
\sum_i s_i(v_i)x_i(\mathbf{v}) &= \sum_i s'_i x_i(\mathbf{v}) \\
&= \sum_i s'_i x_i(s'_i) \\
&\geq \tfrac{1}{\sqrt{m}} \sum_i s'_i x'_i \\
&\geq \tfrac{1}{\sqrt{m}} \sum_i \tau_i(\mathbf{v}_{-i}) x'_i.
\end{aligned}$$

The first line holds because $\mathbf{s}'$ differs from $\mathbf{s}(\mathbf{v})$ only on the bids of losing agents. The second follows from Lemma 3.24, and the third from Lemma 3.23. The last line follows from the fact that $\mathbf{s}'$ doesn't change the bids of winners under $\mathbf{s}(\mathbf{v})$, and for those agents, $s_i(v_i) \geq \tau_i(\mathbf{v}_{-i})$. □

An interpretation of this result is that the pay-your-bid greedy mechanism's welfare loss from incentives is limited to the multiplicative $\frac{e}{e-1}$ factor that appears in value covering. The following hold as corollaries:

Theorem 3.26. *The welfare of the pay-your-bid greedy combinatorial auction is a $\sqrt{m}\frac{e}{e-1}$-approximation to that of any other mechanism.*

Theorem 3.27. *With monopoly reserves and regular bidders, the revenue of the pay-your-bid greedy combinatorial auction is a $(\sqrt{m}+1)\frac{e}{e-1}$-approximation to that of any other mechanism.*

3.5.3. Position Auctions

In first-price position auctions (a.k.a., the generalized first-price auction, GFP), arguments similar to those in the matroid case yield analogous welfare and revenue guarantees.

Formally, a position auction is an auction in which agents can be allocated one of m positions; each of which is valued by an agent at $\alpha_j v_i$. In advertising auctions, these are slots on a web-page to fill where lower slots receive fewer clicks. The positions are ordered such that $\{\alpha_j\}$ is decreasing in j (hence slot 1 is best).

In GFP, agents submit bids b_i, and positions are allocated in order of bid. Each agent pays their bid scaled by the quality of the slot: $\alpha_j b_i$. Equivalently, they pay their bid when they are served, which occurs with probability α_j for position j.

While the correspondence between bids and threshold bids is not as immediate in GFP as in the single-item, first-price auction, GFP satisfies a version of revenue covering where the threshold up to the alternate allocation probability, $T_i(x'_i) = \int_0^{x'_i} \tau(z)\, dz$ is used in place of the expected threshold scaled down by x'_i, which is sufficient for welfare and revenue approximations. The proof is included in Appendix A.4.1.

Theorem 3.28. *The generalized first price (GFP) auction is* 1*-revenue covered.*

3.5.4. Discretized Bids

Often in practice the bid space in an auction is discretized for convenience or feasibility. We note here that discretizing the bidspace of a pay-your-bid auction only results in an additive loss of the discretization amount.

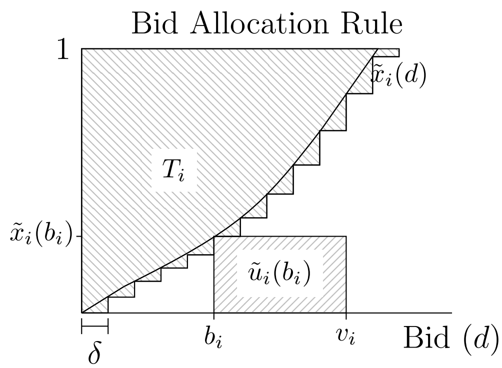

Figure 3.5. When bids are only allowed in multiples of δ, the expected threshold expands, but only by at most an additive amount of δ.

Consider a first-price auction where bids are only allowed in increments of δ. The expected threshold appears larger to the agent with the restriction of bids to discretized amounts than without the restriction. See Figure 3.5 for an illustration.

Consider a pay-your-bid mechanism $\mathcal{M}$ that is μ-revenue covered. Now, let $\mathcal{M}^\delta$ be the discretized version of that mechanism that only allows bids in increments of δ. For any strategy profile in $\mathcal{M}^\delta$, the same strategy profile in $\mathcal{M}$ induces a threshold T that we know is covered by μ copies of the revenue. Consider the expected threshold in $\mathcal{M}^\delta$:

$$T^\delta = \int_0^1 \tau^\delta(z)dz \leq \int_0^1 \tau(z) + \delta dz \leq T + \delta. \tag{3.13}$$

Thus the expected threshold will only be larger by an additive δ, and all of the approximations we have discussed will hold except for that additive δ. If δ is small compared with revenue, this makes effectively no difference: if δ is very large relative to revenue, it could make a large difference.

3.6. Beyond Pay-Your-Bid Auctions

We show that by reducing the problem of bidding in an auction to bidding in a first-price auction, the same value-covering analysis that holds for the first-price auction holds for *any* other auctions as well. This reduction allows us to derive price-of-anarchy results for welfare and revenue in the all-pay auction (Section 3.6.1) and for the simultaneous composition of revenue-covered auctions (Section 3.6.3).

Utility-maximizing agents in any auction must balance two goals: winning the auction and not paying too much to win. In a pay-your-bid auction this relationship is explicit: your bid is exactly your payment if you win, giving an expected utility of $u(b_i) = (v_i - b_i)\tilde{x}_i(b_i)$. For other auctions, this relationship is less clear, but the behavior of agents is still the same: they balance the probability they win with the expected cost of winning, with a utility that we can write for any action as

$$u_i(a_i) = \tilde{x}_i(a_i)v_i - p_i(a_i) = \left(v_i - \frac{p_i(a_i)}{\tilde{x}_i(a_i)}\right)\tilde{x}_i(a_i). \tag{3.14}$$

The term $\frac{p_i(a_i)}{x_i(a_i)}$ plays exactly the same role as the bid in a pay-your-bid auction: it is the price per unit of allocation. We call this the *equivalent bid* of an action; we denote it $\beta_i(a_i) = \frac{p_i(a_i)}{\tilde{x}_i(a_i)}$. We can now define expected thresholds, revenue and value covering using equivalent bids to play exactly the role of first-price bids.

Equivalent Threshold Bids. The pivotal quantity in our pay-your-bid proof framework is the expected threshold bid. First-price auctions have a natural monotonicity property: any bid b is the minimum payment necessary to get the allocation probability $\tilde{x}_i(b)$. For auctions with different payment semantics, we partition actions in agents' choice sets by interim allocation probability, then for each probability consider only the cheapest such action.

For each allocation probability z, let $\alpha_i(z)$ be that cheapest action and let the *equivalent threshold bid* $\tau_i(z) = \beta_i(\alpha_i(z))$ be the equivalent bid of the cheapest action. Formally, $\tau_i(z) = \min_{a_i : \tilde{x}_i(a_i) \geq z} \beta_i(a_i)$, with $\alpha_i(z)$ the arg min. Note that $\tau_i(z)$ depends on $\mathbf{s}$ as it is taken in expectation over the actions of the other agents; for notational convenience, we suppress the strategy profile as an argument. Define the *expected (equivalent) threshold* as $T_i = \int_0^1 \tau_i(z)\, dz$. This quantity will function identically to its counterpart in pay-your-bid auctions, as a representation of how expensive it is for an agent to get allocation.

Covering Conditions and the Price of Anarchy. The use of equivalent bids reduces the optimization problem of a bidder in a general auction to the optimization problem a bidder in a first-price auction faces. As a result, the value covering property of the first-price auction that relates the agents utility and expected threshold still holds:

Lemma 3.29 (Value Covering)**.** *Consider a mechanism $\mathcal{M}$ in BNE with induced allocation and payment rules $(\mathbf{x}, \mathbf{p})$, and an agent i with value v_i. Then*

$$u_i(v_i) + T_i \geq \tfrac{e-1}{e} v_i. \tag{3.15}$$

The proof (included in Appendix A.2) can now be done by reduction to the single-item first-price auction (Lemma 3.4) because the optimization problem is the same in each.

Recall that for the first-price auction, revenue and value covering combined to give approximation results for welfare. The same holds for value covering and μ-revenue covering from Definition 3.1 of general mechanisms.

Theorem 3.30. *If a mechanism is μ-revenue covered, then in any BNE it is a $\mu\frac{e}{e-1}$ approximation to the welfare of the optimal mechanism.*

Proof. The proof proceeds analogously to the proof of Theorem 3.3.

For any agent i with value v_i, value covering (Lemma 3.29) and μ-revenue covering (Definition 3.1) combine to give

$$u_i(v_i) + \mu\,\textsc{Rev}(\mathcal{M}) \geq \tfrac{e-1}{e} v_i.$$

Let $\mathbf{x}^*(\mathbf{v})$ be the welfare-optimal allocation. For any agent, $x_i^* \leq 1$, hence

$$u_i(v_i) + x_i^*(v_i)\mu\textsc{Rev}(\mathcal{M}) \geq x_i^*(v_i)u_i(v_i) + x_i^*(v_i)\mu\textsc{Rev}(\mathcal{M}) \geq x_i^*(v_i)\frac{e-1}{e}v_i. \tag{3.16}$$

Summing over all agents and taking expectation over values gives $\textsc{Util}(\mathcal{M})+\mu\textsc{Rev}(\mathcal{M}) \geq \frac{e-1}{e}\textsc{Welfare}(\textsc{Opt})$. As $\textsc{Util}(\mathcal{M}) + \textsc{Rev}(\mathcal{M}) = \textsc{Welfare}(\mathcal{M})$ and $\mu \geq 1$, we then have our desired result,

$$\mu\frac{e}{e-1}\textsc{Welfare}(\mathcal{M}) \geq \textsc{Welfare}(\textsc{Opt}). \qquad \square$$

3.6.1. All-Pay Auctions

Auctions in which you must pay whether or not you win (*all-pay* auctions) can also be revenue covered. To attain the equivalent bid corresponding to an all-pay bid, divide by the allocation probability: $\beta_i(b_i) = b_i/\tilde{x}_i(b_i)$. The revenue in the all-pay auction is always greater than the expected all-pay bid threshold; and the all-pay threshold is always at least half of the equivalent bid threshold, which gives revenue covering with $\mu = 2$.

Lemma 3.31. *The all-pay matroid auction is 2-revenue covered.*

We include the proof here for the single-item case; the generalization to matroids is included in Appendix A.4.

Proof (single-item). We first translate revenue to threshold bids. In expectation, these thresholds 2-approximate the equivalent threshold bids. Combining the two arguments yields the result.

Let $\mathbf{s}$ be an arbitrary strategy profile and $\mathbf{x}'$ an alternate allocation. First, for any bidder i, let $\tau_i^b(\mathbf{v}_{-i})$ be the threshold (all-pay) bid for i in realized value profile $\mathbf{v}_{-i}$ under $\mathbf{s}$. Since the threshold bid corresponds to some other agent's bid, and agents pay their bids regardless of allocation, $\text{Rev}(\mathcal{M}) \geq \mathbf{E}_{\mathbf{v}_{-i}}[\tau_i^b(\mathbf{v}_{-i})]$.

To relate threshold bids to equivalent thresholds, let $a_i(z)$ be the z-quantile of i's competing bids. That is, $a_i(z) = \arg\min_{a_i} \tilde{p}_i(a_i)$ subject to $\tilde{x}_i(a_i) \geq z$. By the definition of τ_i,

$$\frac{\tilde{p}_i(a_i(z))}{\tilde{x}_i(a_i(z))} \geq \tau_i(z).$$

Rearranging and noting that in an all-pay auction, $\tilde{p}_i(a_i(z)) = a_i(z)$, we obtain

$$a_i(z) \geq \tau_i(z)\tilde{x}_i(a_i(z)) \geq \tau_i(z)z. \tag{3.17}$$

This yields the following sequence of inequalities:

$$\mathbf{E}_{\mathbf{v}_{-i}}\left[\tau_i^b(\mathbf{v}_{-i})\right] = \int_0^1 a_i(z)\,dz \geq \int_0^1 \tau_i(z)z\,dz \geq \frac{1}{2}\int_0^1 \tau_i(z)\,dz = T_i, \tag{3.18}$$

where the first equality follows from noting that expected value can be computed by integrating over quantiles, the first inequality from equation (3.17), and the second inequality from the fact that τ_i is an increasing function and Chebyshev's sum inequality. Finally, since $\mathbf{x}'$ is feasible, $\sum_i x_i' \leq 1$. We can combine this with (3.18) to get

$$2\text{Rev}(\mathcal{M}) \geq \sum\nolimits_i T_i x_i', \tag{3.19}$$

which is the definition of 2-revenue covering. □

Combining revenue covering with Theorem 3.30 gives a welfare bound of $2e/(e-1)$. For revenue, using the techniques of Section 3.4.3 and ensuring that at least two bidders have values drawn from each distribution gives a $4e/(e-1)$-approximation to the revenue of the optimal auction[1].

The bounds can be improved to 2 and 6 respectively by adapting the value-covering condition to the all-pay environment, shown in Appendix A.4.

3.6.2. The Second-Price Auction

Not all mechanisms are revenue covered. In the second-price auction, agents submit sealed bids, the highest bidder wins and is charged the second-highest bid. This auction lacks a direct connection between bidders' threshold bids and the revenue of the auction, which is required for revenue covering. To illustrate, consider a two-agent setting, and assume agents 1 bids 1 and agent 2 bids ϵ, deterministically. The revenue is ϵ, but T_2 is 1, so the second-price auction cannot be revenue covered.

If agents are assumed to never bid above their values, then "welfare-covering" can be used in place of revenue-covering to give approximation bounds. This results in proofs that are structurally very similar to those that assume no-overbidding in Syrgkanis and Tardos [2013]; Caragiannis et al. [2015].

[1]Reserve prices do not work as conveniently in all-pay auctions as they do in first-price auctions, as a reserve price r may eliminate the allocation of agents with values greater than r if they do not attain enough allocation by outbidding the reserve.

3.6.3. Simultaneous Composition

In this section, we prove that μ-revenue covering is closed under simultaneous composition. Consequentially, welfare bounds proved for individual mechanisms via revenue covering extend to collections of such mechanisms run simultaneously. Moreover, we show that revenue covering with the generalized notion of reserves in Section 3.6 is similarly closed under composition. As a corollary, we extend our revenue bound for first-price auctions to the simultaneous composition of such auctions.

Formally, the simultaneous composition of m mechanisms for single-dimensional agents is the following:

Definition 3.2. Let mechanisms $\mathcal{M}_1, \ldots, \mathcal{M}_m$ have allocation and payment rules $(\mathbf{x}^j, \mathbf{p}^j)$ and individual action spaces spaces $A_i^1, \ldots, A_i^m$ for each agent i. The *simultaneous composition* of $\mathcal{M}_1, \ldots, \mathcal{M}_m$ is defined to have:

- Action space $\times_j A_i^j$ for each agent. That is, each agent participates in the global mechanism by participating in each composed mechanism individually.
- Allocation rule $\hat{x}_i(\mathbf{a}) = [\tilde{x}_i^1(\mathbf{a}^1), \ldots, \tilde{x}_i^m(\mathbf{a}^m)]$. In other words, the mechanism gives each agent their allocated bundle from each mechanism.
- Payment rule $\tilde{p}_i(\mathbf{a}) = \sum_j \tilde{p}_i^j(\mathbf{a}^j)$. That is, agents make payments to every composed mechanism.

We assume agent utilities are *unit demand* and *single-valued* over the outcomes of the mechanisms. Agent utilities are therefore of the form $v_i \cdot (\max_{j \in \mathcal{S}_i} \tilde{x}_i^j(\mathbf{a})) - \tilde{p}_i(\mathbf{a})$. The induced single-dimensional allocation rule of the global mechanism is $\tilde{x}_i(\mathbf{a}) = \max_{j \in \mathcal{S}_i} \tilde{x}_i^j(\mathbf{a})$. Using $\tilde{\mathbf{x}}$, define τ_i, $\underline{x}_i(y)$, and T_i^y as discussed in Section 3.6.

We can now state the main theorem of the section.

Theorem 3.32. *Let $\mathcal{M}$ be the simultaneous composition of mechanisms $\mathcal{M}_1, \dots \mathcal{M}_m$ with unit-demand, single-valued agents. If $\mathcal{M}_1, \dots \mathcal{M}_m$ are μ-revenue covered with reserves $\mathbf{r}$ then $\mathcal{M}$ is μ-revenue covered with reserves $\mathbf{r}$.*

The special case with reserves $\mathbf{r} = \mathbf{0}$ yields the following corollary:

Corollary 3.33. *Let $\mathcal{M}$ be the simultaneous composition of μ-revenue covered mechanisms $\mathcal{M}_1, \dots \mathcal{M}_m$ with unit-demand, single-valued agents; then $\mathcal{M}$ is μ-revenue covered.*

The intuition driving the proof of Theorem 3.32 is that getting allocated with probability z in any individual mechanism will always be more costly than getting allocated with that same probability when given all m mechanisms to choose from and bid in. It follows that each agent's expected threshold for the global mechanism is smaller than that for each individual mechanism. Since payments for the global mechanism are just the sum of the individual mechanisms' payments, local revenue covering implies revenue covering for the composed mechanism.

Each agent's strategy in the global mechanism specifies a strategy for each individual mechanism. Therefore, for each individual mechanism j, any global strategy profile $\mathbf{s}$ induces a local strategy profile $\mathbf{s}^j$. The profile $\mathbf{s}^j$ in mechanism j induces local versions of the equivalent bid β_i^j, threshold τ_i^j, reserve allocation $\underline{x}_i^j(y)$, and expected threshold $T_i^{y(j)}$.

We can now formalize the above intuition that individual expected thresholds are larger than the expected thresholds in the global mechanism:

Lemma 3.34. *For any strategy profile $\mathbf{s}$ and allocation level y, $T_i^y \leq T_i^{y(j)}$.*

Proof. Fix an individual mechanism j. By definition, $T_i^y = \int_{\underline{x}_i(y)}^1 \tau_i(z)\,dz$, and $T_i^{y(j)} = \int_{\underline{x}_i^j(y)}^1 \tau_i^j(z)\,dz$. We will show that $\tau_i(z) \leq \tau_i^j(z)$ and $\underline{x}_i(y) \geq \underline{x}_i^j(y)$. These two facts together imply the lemma.

First, the threshold for the global mechanism is defined as $\tau_i(z) = \min_{a_i:\tilde{x}_i(a_i)\geq z} \beta_i(a_i)$, and for the local mechanism as $\tau_i^j(z) = \min_{a_i^j:\tilde{x}_i^j(a_i^j)\geq z} \beta_i^j(a_i^j)$. For every action a_i^j in the local mechanism, there is a corresponding action a_i in the global mechanism where i takes action a_i^j in mechanism j and withdraws from every other mechanism. Since $\beta_i(a_i) = \beta_i^j(a_i^j)$, it follows that $\tau_i(z) \leq \tau_i^j(z)$.

Next, the global reserve allocation is $\underline{x}_i(y) = \max_{a_i:\beta_i(a_i)\leq y} \tilde{x}_i(a_i)$, and the local reserve allocation is $\underline{x}_i^j(y) = \max_{a_i^j:\beta_i^j(a_i^j)\leq y} \tilde{x}_i^j(a_i^j)$. Again, for every a_i^j in the local mechanism, there is an action a_i where i plays a_i^j in j and withdraws everywhere else. Since $\beta_i(a_i) = \beta_i^j(a_i^j)$, it follows that $\underline{x}_i(y) \geq \underline{x}_i^j(y)$. $\square$

Proof of Theorem 3.32. Every feasible allocation $\mathbf{x}'$ in the global mechanism is a matching between agents and items. Define $x'_{i,j} = x'_i$ if i and j are matched, and 0 otherwise. The theorem now follows from the following sequence of inequalities:

$$\begin{aligned}
\mu\textsc{Rev}(\mathcal{M}) &= \sum\nolimits_j \textsc{Rev}(\mathcal{M}_j) \\
&\geq \mu \sum\nolimits_j \sum\nolimits_i T_i^{y(j)} x'_{i,j} \\
&\geq \mu \sum\nolimits_i \sum\nolimits_j T_i^{y(j)} x'_{i,j} \\
&\geq \mu \sum\nolimits_i T_i^y x'_i
\end{aligned}$$

The first equality follows from the definition of simultaneous composition. The second follows revenue covering of mechanism j and the fact that $\mathbf{x'}_{\cdot,j}$ is feasible for j by downward closure. The final inequality follows from Lemma 3.34 and the fact that $\sum_j x'_{i,j} = x'_i$. $\square$

One corollary of Theorem 3.32 is that the simultaneous composition of any number of first-price auctions with the same reserves is 1-revenue covered with those reserves. Using monopoly reserves in a regular environment yields a revenue bound:

Corollary 3.35. *Let $\mathcal{M}$ be the simultaneous composition of m first-price auctions with monopoly reserves $\mathbf{r}^*$, and unit-demand, single-valued agents with regular value distributions. The revenue of $\mathcal{M}$ is a $\frac{2e}{e-1}-$approximation to the revenue of the optimal global mechanism.*

3.7. Conclusion

We have shown a framework for proving price of anarchy results for welfare and revenue in Bayes-Nash Equilibrium. This framework enabled us to prove both welfare and new revenue approximation results for non-truthful auctions in asymmetric settings, including first price and all-pay auctions.

We split this framework in two distinct parts that isolates the analysis of Bayes-Nash Equilibrium from the analysis of the specific mechanism. The first part, value covering, depends only on Bayes-Nash Equilibrium and relates an agents surplus and expected price for additional allocation with her optimal surplus. The second, revenue-covering, depends only on properties of a mechanism over individually rational strategy profiles and feasible allocations. This is especially helpful when equilibria are hard to characterize or understand analytically, as is the case with the first-price auction in asymmetric settings. We expect this framework will aid broadly in understanding properties of equilibria in auctions well beyond the confines of symmetric settings.

We invoked the characterization of Bayes-Nash Equilibrium in a few very specific places in our proofs. For value covering and virtual value covering, it is only important that an agent be best responding to the expected actions of other bidders. For the revenue approximation results, we do rely on the characterization of equilibrium by Myerson [1981] to account for revenue via virtual values. This is the crucial part that allows us to relate the allocation a bidder receives to their contribution to revenue. Extensions beyond single-parameter, risk-neutral, private-valued agents will be challenging without a virtual-value equivalent.

CHAPTER 4

Empirical Price of Anarchy Bounds

This chapter brings data into Price of Anarchy analysis with the *Empirical Price of Anarchy*: the worst case Price of Anarchy of a setting and equilibrium that could have produced a distribution of data. An approach for estimating empirical price of anarchy bounds is presented, based on a near direct refinement of the theoretical framework in Chapter 3 for empirical purposes: estimating the revenue covering bound in place of proving it.

Analysis of welfare in auctions comes traditionally via one of two approaches: precise but fragile inference of the exact details of a setting using data or robust but coarse theoretical price of anarchy bounds that hold in any setting.

In this chapter, we provide tools for analyzing and estimating the *empirical price of anarchy* of an auction. The empirical price of anarchy is the worst case efficiency loss of any auction instance that *could have produced the data*, relative to the optimal.

Our techniques are based on inferring simple properties of auctions from bid data: primarily the expected revenue and the expected payments and allocation probabilities from possible bids. These quantities alone allow us to empirically bound the revenue covering parameter of an auction which allows us to re-purpose the theoretical machinery from Chapter 3 for empirical means. While we focus on the setting of position auctions, and particularly the generalized second price auction, our techniques are applicable far more generally.

Finally, we apply our techniques to a selection of advertising auctions on Microsoft's Bing which are not theoretically revenue covered, and find empirical results that are a significant improvement over the theoretical worst-case bounds.

4.1. Introduction

Evaluation of the revenue and welfare of market mechanisms is a key question in Economics. A typical question of interest is the comparison of a currently deployed mechanism with the best solution implemented by a central planner, taking into account the incentives of participating Economic agents. The *price of anarchy*, first introduced by Koutsoupias and Papadimitriou [1999] for network routing games, provides a worst-case bound on the ratio of the revenue or welfare from the optimal mechanism compared to the currently deployed mechanism.

The worst-case nature of the price of anarchy results in very robust bounds, but this robustness can come at the cost of bounds that are too coarse for an analyst interested in understanding the performance of a currently deployed mechanism. In some types of games this is not a problem because performance can be empirically estimated: processing time or memory usage can be measured; route choice and delay in a network can be tracked, and compared to another benchmark.

However, in auctions and other settings where agents have private information that impacts the objectives of a system, estimating the performance has traditionally required learning that private information. Oftentimes the results from this style of analysis are also very sensitive to the exact decision making of bidders, and for instance are not robust to bidders who choose only the approximately best action, or play learning strategies.

In this chapter, we bridge the robust but coarse theoretical price of anarchy bounds and precise but fragile inference based bounds, by integrating data directly into the price of anarchy style analysis. Instead of quantifying over all settings and uncertainties, we take the worst case over all settings and uncertainties that could induce the observed data. The more we know about the data generated by a mechanism, the higher the potential for an accurate bound.

Our approach benefits from the inherent robustness of worst-case analysis to realistic market features such as differences in details of mechanisms or agents who only approximately best-respond. At the same time, our approach uses the data and effectively informs the price of anarchy bound regarding the "worst case scenario" distributions of uncertainty that are clearly inconsistent with the observed data. That allows us to improve the welfare and revenue bounds given by the theoretical price of anarchy.

4.1.1. Methods

Theoretical. Our theoretical technique for proving empirical price of anarchy bounds is an empirical application of the revenue-covering framework from Chapter 3, targeted to position auctions. First, we analyze the optimization problem of a bidder, comparing actions in the auction based on their expected price-per-click (which is the first-price equivalent bid from Chapter 3).

Second, we relate the revenue of an auction to a threshold quantity, which is based on how expensive allocation is. We call this empirical revenue covering, and differs from theoretical revenue covering of Chapter 3 only in that we measure it for a given instance of an auction instead of proving it for every possible strategy profile. As a result, our empirical revenue covering framework applies even more broadly that theoretical revenue covering: it can be measured for *any Bayes-Nash Equilibrium* of *any* single-parameter auction in the independent, private values model.

Finally, we consider and measure how agents would react to the optimization problem that they are faced with. We measure the value-covering of the auction, which improves on the $1 - \frac{1}{e}$ term shown in Chapter 3. This can be done both with precise knowledge of price-per-click allocation rule, or with rough knowledge of concentration bounds on the price-per-click allocation rule.

Our general approach can also be seen as reducing the empirical analysis of an auction to the econometric question of estimating the revenue of an auction and estimating the allocations and prices-per-click of actions in the auction. Notably, this avoids the need to estimate empirical first order conditions required for traditional econometric analysis.

Robustness. Our results inherit a robustness to changes in the mechanism or the setting by leveraging the theoretical framework of Chapter 3.

- **Beyond Position Auctions.** The analysis in this chapter is focused on position auctions, but the analysis is general and can be applied for Bayes-Nash Equilibria in other setting or mechanism. It is only required that allocation probabilities and expected payments can be estimated.
- **Changes in the Mechanism.** As the thresholds we calculate are based on the price-per-click allocation rule of a bidder, threshold quantities can be compared and computed no matter what the mechanism is as long as these quantities can be estimated.
- **Approximate Equilibrium.** If the agents in an auction only ϵ-best respond to the optimization problem that they are faced with, then our efficiency results only degrade by that ϵ.

  Moreover, if some agents are irrational and some are rational, then our results can be broken out to give efficiency results for individual rational bidders: each bidder's contribution to the optimal welfare is approximated by a combination of their own utility and revenue of the seller if they best respond. Notably, this requires no comparison of utilities between bidders in the auction or assumptions on why the other bidders are playing actions.
- **Learning Quality Scores.** We model the quality score of a bidder as coming from a known distribution. This distribution should be interpreted as the auctioneers knowledge of the quality score of the bidder. Our efficiency results give a comparison to the optimal auction *for the same knowledge of quality scores of the bidders.* This

distribution moreover can have arbitrary correlations, as it only really affects the space of feasible allocations.

4.1.2. Contributions

Our primary contributions are the following:

- **Empirical Price of Anarchy**. We introduce the *empirical Price of Anarchy*(EPoA) benchmark for welfare, representing the worst case efficiency loss of a game consistent with a distribution of data from the game.
- **Empirical Revenue Covering**. We refine the revenue-covering framework of Hartline et al. [2014] for proving robust EPoA bounds, and show that we can empirically estimate the empirical revenue covering of the Generalized Second Price auction.
- **Data**. We apply and bound the empirical price of anarchy from GSP advertising auctions run in Microsoft's Bing, and show that we get EPoA bounds that are significantly stronger than the relevant theoretical bounds.

4.1.3. Related Work

The efficiency of the Generalized Second-Price auction (GSP) was originally modeled and studied in full-information settings in Edelman et al. [2007] and Varian [2009]. Gomes and Sweeney [2014] characterize equilibrium in the Bayesian setting, and give conditions on the existence of efficient equilibria. Athey and Nekipelov [2010] give a structural model of GSP with varying quality scores, which are included in our model. Caragiannis et al. [2015] explores the efficiency of GSP in the Bayesian setting, and finds a theoretical price of anarchy for welfare of 2.927 when the value distributions are independent or correlated, and agents do not overbid. Our results apply for independent distributions of values, but do not need the no-overbidding assumption, and generalizes to the more realistic case that the ranking of the bidders is not exactly equal to the quality score of a bidder. Hence, in principle we could observe higher inefficiency in the data than the theoretical bound above. Despite this fact we find in the data that only better inefficiency bounds are derived, with the exception of one search phrase where we almost exactly match the latter worst-case theoretical bound. The semi-smoothness based approach of Caragiannis et al. [2015] can be seen through our model as using a welfare covering property in place of revenue covering.

4.2. Preliminaries

We consider the position auction setting, with m positions and n bidders. Each bidder i has a private value v_i drawn independently from distribution F_i over the space of possible values V_i. We denote the joint value-space and distribution over values $\mathbf{V} = \Pi_i\ V_i$ and $\mathbf{F} = \Pi_i\ F_i$ respectively. Bidders have a linear utility, so if they pay p_i to receive a probability of service x_i, the utility of the bidder is $u_i = v_i x_i - p_i$.

An outcome o in a position auction is an allocation of positions to bidders. $o(j)$ denotes the bidder who is allocated position j; $o^{-1}(i)$ refers to the position assigned to bidder i. Henceforth we will adopt the terminology of ad auctions and refer to service as a 'click'.

When bidder i is assigned to slot j, the probability of click $c_{i,j}$ is the product of the click-through-rate of the slot α_j and the quality score of the bidder, γ_i, so $c_{i,j} = \alpha_j \gamma_i$. We will generally assume that γ_i is drawn independently from distribution Γ_i, and is observable to the auctioneer, but *not* to the bidder themselves.

Since the auctioneer can use the quality scores in assigning bidders to slots and the quality scores impact the number of clicks that each agent sees, an allocation $\mathbf{x}$ is feasible if and only if there is a quality-score dependent assignment of slots to bidders that gives rise to this allocation.

Denote by $\rho(\gamma, \cdot)$ such an assignment, where $\rho(\gamma, j)$ is the agent who is assigned position j when the quality score profile is γ and $\rho^{-1}(\gamma, i)$ is the position assigned to agent i. Moreover, denote by $\mathcal{M}$ the space of all such quality score dependent assignments. Then an allocation $\mathbf{x}$ is feasible if there exists $\rho \in \mathcal{M}$ such that for each bidder i: $x_i = \mathbf{E}_{\gamma}[\alpha_{\rho^{-1}(\gamma,i)}\gamma_i]$. Call $\mathbf{X}$ the set of all feasible allocations.

A position auction $\mathcal{A}$ consists of a bid space $\mathcal{B}$, allocation rule $\tilde{\mathbf{x}} : \mathcal{B}^n \to \mathbf{X}$ mapping from bid profiles to feasible allocations and payment allocation rule $\tilde{\mathbf{p}} : \mathcal{B}^n \to \mathbb{R}^n$ mapping from bid profiles to payments. A strategy profile $\mathbf{s} : V \to \mathcal{B}^n$ maps values of agents to bids. For a set of values $\mathbf{v}$, the utility generated for each bidder is $U_i(\mathbf{b}; v_i) = v_i\tilde{\mathbf{x}}(\mathbf{b}) - \tilde{p}_i(\mathbf{b})$.

Given a strategy profile $\mathbf{s}$, we will often use the expected allocation and payment an agent expects to receive when playing an bid b_i, taking expectation over other agents values and the quality score γ_i. We call $\tilde{x}_i(b_i) = \mathbf{E}_{\mathbf{v}_{-i}}[\tilde{x}_i(b_i, \mathbf{s}_{-i}(\mathbf{v}_{-i}))]$ the *interim bid allocation rule*. We define $\tilde{p}_i(b_i)$ and $\tilde{u}_i(b_i)$ analogously.

A strategy profile $\mathbf{s}$ is in Bayes-Nash Equilibrium (BNE) if for all agents i, $s_i(v_i)$ maximizes their interim expected utility: e.g., for all bids d, $\tilde{u}_i(s_i(v_i)) \geq \tilde{u}_i(d)$.

The welfare from an allocation $\mathbf{x}$ is the expected utility generated for both the bidders and the auctioneer, $\sum_i x_i v_i$. Thus the expected utility of a strategy profile $\mathbf{s}$ is

$$\textsc{Welfare}(\mathcal{A}(\mathbf{s})) = \mathbf{E}_{\mathbf{v}}\left[\sum_i v_i x_i(s_i(v_i))\right] \tag{4.1}$$

We will break down the welfare of the auction into the revenue paid to the auctioneer, $\textsc{Rev}(\mathcal{A}(\mathbf{s})) = \mathbf{E}_{\mathbf{v}}[\sum_i \tilde{p}_i(s_i(\mathbf{v}))]$ and the utility derived from the agents, $\textsc{Util}(\mathcal{A}(\mathbf{s})) = \mathbf{E}_{\mathbf{v}}[\sum_i \tilde{u}_i(\mathbf{s}(\mathbf{v}))]$, with

$$\textsc{Welfare}(\mathcal{A}(\mathbf{s})) = \textsc{Rev}(\mathcal{A}(\mathbf{s})) + \textsc{Util}(\mathcal{A}(\mathbf{s})).$$

Our benchmark for welfare will be the welfare of the auction that chooses a feasible allocation to maximize the welfare generated, thus $\textsc{Welfare}(\textsc{Opt}) = \mathbf{E}_{\mathbf{v}}[\max_{\mathbf{x}} \sum_i x_i v_i] = \mathbf{E}_{\mathbf{v},\boldsymbol{\gamma}}[\max_o \sum_i \gamma_i \alpha_{o^{-1}(i)} v_i]$. We will denote the resulting optimal value-based allocation rule $\mathbf{x}^*$.

4.2.1. Sponsored Search Auction: model and data

We consider data generated by advertisers repeatedly participating in a sponsored search auction. The mechanism that is being repeated at each stage is an instance of a generalized second price auction triggered by a search query.

The rules of each auction are as follows: Each advertiser i is associated with a click probability γ_i and a scoring coefficient s_i^c and is asked to submit a bid-per-click b_i. Advertisers are ranked by their rank-score $q_i = s_i^c \cdot b_i$ and allocated positions in decreasing order of rank-score as long as they pass a rank-score reserve r. If advertisers also pass a higher mainline reserve m, then they may be allocated in the positions that appear in the mainline part of the page, but at most k advertisers are placed on the mainline.

If advertiser i is allocated position j, then he is clicked with some probability $c_{i,j}$, which we will assume to be separable into a part α_j depending on the position and a part γ_i depending on the advertiser, and that the position related effect is the same in all the participating auctions: $c_{i,j} = \alpha_j \cdot \gamma_i$. We denote with $\boldsymbol{\gamma}$ and $\boldsymbol{\alpha}$ the vectors of quality scores and position coefficients. All the mentioned sets of parameters $\theta = (\mathbf{s^c}, \boldsymbol{\alpha}, \boldsymbol{\gamma}, r, m, k)$ and the bids $\mathbf{b}$ are observable in the data. Moreover, the parameters and bids are known to the auctioneer at the allocation time. We will denote with $\pi_{\mathbf{b},\theta}(j)$ the bidder allocated in slot j under a bid profile $\mathbf{b}$ and parameter profile θ. We denote with $\pi^{-1}_{\mathbf{b},\theta}(i)$ the slot allocated to bidder i.

If advertiser i is allocated position j, then he pays only when he is clicked and his payment, i.e. his cost-per-click (CPC) is the minimal bid he had to place to keep his position, which is:

$$cpc_{ij}(\mathbf{b};\theta) = \max\left\{\frac{s^c_{\pi_{\mathbf{b},\theta}(j+1)} \cdot b_{\pi_{\mathbf{b},\theta}(j+1)}}{s^c_i}, \frac{r}{s^c_i}, \frac{m}{s^c_i} \cdot \mathbf{1}\{j \in M\}\right\} \tag{4.2}$$

where with M we denote the set of mainline positions.

We also assume that each advertiser has a value-per-click (VPC) v_i, which is not observed in the data. If under a bid profile $\mathbf{b}$, advertiser i is allocated slot $\pi^{-1}_{\mathbf{b},\theta}(i)$, his expected utility is:

$$U_i(\mathbf{b}; v_i) = \mathbf{E}_\theta\left[\alpha_{\pi^{-1}_{\mathbf{b},\theta}(i)} \cdot \gamma_i \cdot \left(v_i - cpc_{i\pi^{-1}_{\mathbf{b},\theta}(i)}(\mathbf{b};\theta)\right)\right] \tag{4.3}$$

4.2.2. Price of Anarchy

Definition 4.1. The (Bayesian) price-of-anarchy for welfare of an auction $\mathcal{A}$ is defined as the worst-case ratio of welfare in the optimal auction to the welfare in an equilibrium, taken over all settings, and all equilibrium strategies associated with the setting:

$$PoA(\mathcal{A}) = \max_{\mathbf{F}, \mathbf{s} \in BNE(\mathcal{A},\mathbf{F})} \frac{\text{WELFARE}(\text{OPT}(\mathbf{F}))}{\text{WELFARE}(\mathcal{A}(\mathbf{F}, \mathbf{s}))} \tag{4.4}$$

We will also refer to the price of anarchy with a known restriction on the equilibrium or settings available:

Definition 4.2. The (Bayesian) price-of-anarchy for welfare of an auction and a set $\mathscr{X}$ of value distribution, equilibrium pairs is defined as the worst-case ratio of welfare in the optimal auction to the welfare of an setting-equilibrium pair in $\mathscr{X}$:

$$PoA(\mathcal{A}, \mathscr{X}) = \max_{(\mathbf{F},\mathbf{s})\in\mathscr{X}} \frac{\textsc{Welfare}(\textsc{Opt}(\mathbf{F}))}{\textsc{Welfare}(\mathcal{A}(\mathbf{F},\mathbf{s}))} \tag{4.5}$$

4.3. Price of Anarchy from Data

We begin by defining the empirical price-of-anarchy of an auction, which differs from theoretical price of anarchy in that we assume knowledge of the data generated by the auction, not knowledge about the setting or equilibrium. We use the notation $D(\mathcal{A}, \mathbf{F}, \mathbf{s}) = \mathscr{D}$ to denote that $\mathscr{D}$ is the distribution of data produced by running the mechanism $\mathcal{A}$ with value distributions $\mathbf{F}$ and strategies $\mathbf{s}$.

Definition 4.3 (Empirical Price of Anarchy)**.** The Bayesian empirical price-of-anarchy for welfare of an auction and a distribution of data $\mathscr{D}$ is the Price of Anarchy of the auction restricted to settings and equilibria that could generate $\mathscr{D}$, e.g.,

$$EPoA(\mathcal{A}, \mathscr{D}) = PoA\left(\mathcal{A},\ \{(\mathbf{F},\mathbf{s})\}_{D(\mathcal{A},\mathbf{F},\mathbf{s})=\mathscr{D}\wedge\mathbf{s}\in BNE(\mathcal{A},\mathbf{F})}\right). \tag{4.6}$$

In the case that an equilibrium and setting can be point-identified from the data distribution $\mathscr{D}$, the set $\{(\mathbf{F},\mathbf{s})\}_{D(\mathcal{A},\mathbf{F},\mathbf{s})=\mathscr{D}\wedge\mathbf{s}\in BNE(\mathcal{A},\mathbf{F})}$ has only one element, the point-identified setting and equilibrium.

The distribution of data $\mathscr{D}$ from an equilibrium of an auction is still a theoretical quantity. We will primarily be proving empirical price of anarchy bounds for all distributions that satisfy properties that are easy to infer. For instance, while we will not do this in this chapter, one could prove price of anarchy bounds whenever the average second highest bid is at least a constant fraction of the first highest bid.

4.3.1. Empirical Revenue Covering Framework for Position Auctions

In this section, we refine the revenue covering framework from Chapter 3 for empirical bounds.

Notably, we use a pointwise version of revenue-covering that applies for a given strategy profile and auction rather than taking the worst-case revenue covering over all strategy profiles.

The property of μ-revenue covering is based only on the relationship between the expected revenue of an auction, and a property of the optimization problem that the bidders are solving (the expected threshold).

Both of these quantities are observable in the data, and hence by observing that an instance of an auction is μ-revenue covered, we will get empirical price of anarchy bounds that apply for the auction we are observing.

The rest of this section will proceed in three parts:

(1) Generate PPC Allocation rules: Analyze how to bid in the auction.
(2) Measure μ: Analyze the correspondence between thresholds and revenue in the auction.
(3) Measure λ: Calculate the worst-case tradeoff between utility and thresholds in the auction.

Generate PPC Allocation Rules. We first focus on the optimization problem each bidder faces. When bidding in an auction, each bidder must think about for each possible bid, how

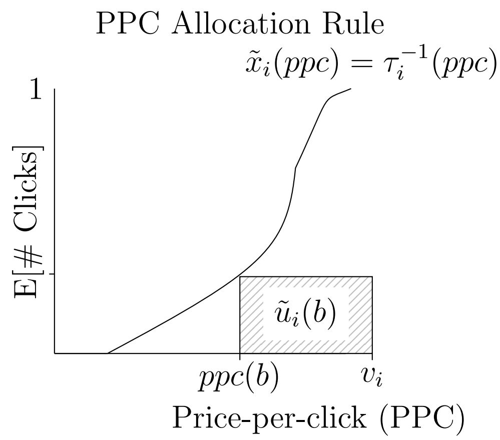

Figure 4.1. For any bid b with PPC $ppc(b)$, the area of a rectangle between $(ppc(b), \tilde{x}_i(ppc(b)))$ and $(v_i, 0)$ on the bid allocation rule is the expected utility $\tilde{u}_i(b)$. The BNE action b^* is chosen to maximize this area.

many clicks she will receive and how much she will have to pay on average for each click. In particular, the utility of an agent can be written to *only* include these terms:

$$\begin{aligned} \tilde{u}_i(b) &= v_i \tilde{x}_i(b) - \tilde{p}_i(b) \\ &= \tilde{x}_i(b) \left(v_i - \frac{\tilde{p}_i(b)}{\tilde{x}_i(b)} \right) \end{aligned} \tag{4.7}$$

The price-per-click term $\frac{\tilde{p}_i(b)}{\tilde{x}_i(b)}$ term now plays exactly the same role in the utility function that the first-price bid does in the first price auction. We call this term $ppc(b) = \frac{\tilde{p}_i(b)}{\tilde{x}_i(b)}$ the price-per-click of the bid in a position auction. Outside of position auctions, it is called the (first-price) equivalent bid in Chapter 3, because it plays the same role as a first-price bid does in a first-price style auction.

Our analysis will be based on the *price-per-click* allocation rule $\tilde{x}(ppc)$, which plots the expected number of clicks of bids against their prices-per-click. See Figure 4.2b for an illustration of the PPC allocation rule.

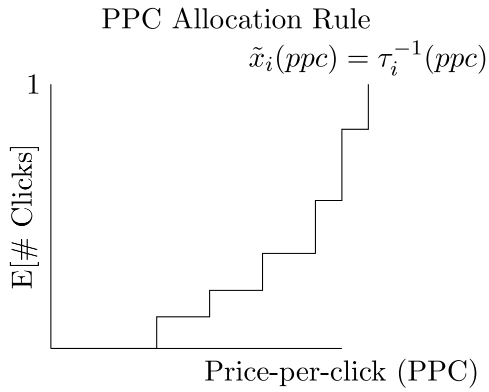

(a) The price-per-click allocation rule in full-information GSP.

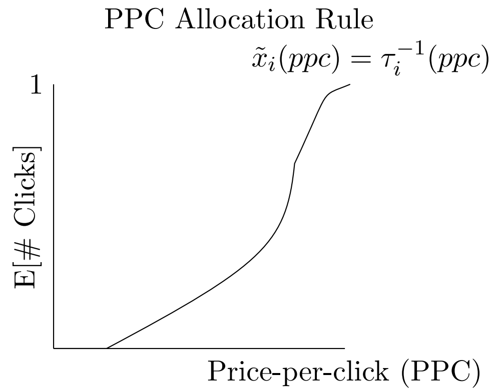

(b) The price-per-click allocation rule in Bayesian GSP.

Figure 4.2. In GSP, the price-per-click curve will be piecewise constant in full-information settings, and smoother in Bayesian settings. In Bayesian settings a bid might result in winning a number of possible slots, and so the expected number of clicks comes from all possible slots.

The utility of a bidder is $\tilde{u}_i(b) = \tilde{x}_i(b)\,(v_i - ppc(b))$, which has a clean visualization on a plot of the PPC allocation rule: is the area of a rectangle between the points $(ppc, \tilde{x}_i(ppc))$ and $(v_i, 0)$. See Figure 4.1 for an illustration.

Thresholds & Revenue Covering. We will most often use the inverse of the PPC allocation rule for our analysis; let $\tau_i(z) = \tilde{x}_i^{-1}(z)$ be the price-per-click of the cheapest bid that achieves allocation at least z. More formally, $\tau_i(z) = \min_{b|\tilde{x}_i(b)\geq z}\{ppc(b)\}$.

The *threshold* for agent i and expected probability of click x_i' is

$$T_i(x_i') = \int_0^{x_i'} \tau_i(z)\ dz \tag{4.8}$$

See Figure 4.3b for an illustration of $T_i(x_i')$ on the plot of the price-per-click allocation rule.

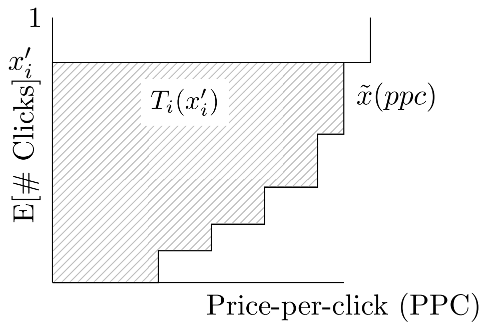

(a) The expected threshold in full-information GSP.

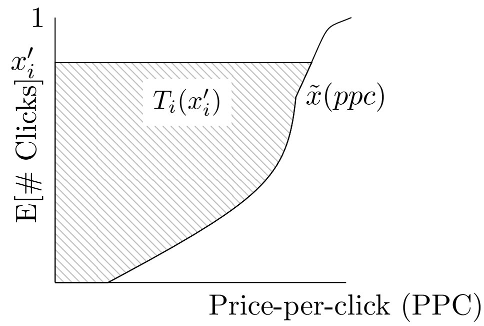

(b) The expected threshold in Bayesian GSP.

Figure 4.3. The expected threshold for each bidder and allocation x_i' is the area above the price-per-click allocation rule, up to x_i'. Note that in Bayesian GSP, the expected thresholds for bidder i in slot j is *not* the expectation of the full-information threshold.

The *total threshold* for the allocation $\mathbf{x}'$ is then the sum of the expected thresholds across all agents, $\sum_i T_i(x_i')$.[1] We now refine the notion of revenue-covering from Chapter 3 to apply for a specific strategy profile.

Definition 4.4 (Revenue Covering)**.** Auction $\mathcal{A}$ with action distribution F^a is μ-revenue covered if for any feasible allocation $\mathbf{x}'$,

$$\mu \textsc{Rev}(\mathcal{A}(F^a)) \geq \sum_i T_i(x_i'). \tag{4.9}$$

If we can prove that for any action distribution, the auction and strategy profile are revenue covered, the auction is theoretically μ-revenue covered in the sense of Definition 3.

We will say a strategy profile and setting are revenue covered for an auction if the action distribution resulting from agents playing in the auction according to the strategy profile is

[1]Note that while this is different than the general definition of expected thresholds in Hartline et al. [2014], it is the same as the definition of thresholds for the generalized-first-price position auction in Hartline et al. [2014]. It is also related to the threshold notion in Syrgkanis and Tardos [2013], which uses $\tau(\mathbf{x}')$ as the threshold quantity rather than $T(\mathbf{x}') = \int_0^z \tau(z)\, dz$.

revenue covered. We will also use $\mathbf{s}(\mathbf{F})$ to refer to the distribution of actions resulting from agents playing from $\mathbf{s}$ when their values are drawn from $\mathbf{F}$.

Value Covering & PoA Results.

Lemma 4.1 (Value Covering). *For any bidder i with value v_i and allocation amount x_i',*

$$u_i(v_i) + \frac{1}{\mu}T_i(x_i') \geq \frac{1-e^{-\mu}}{\mu}x_i'v_i. \tag{4.10}$$

The proof is included in the appendix for completeness: it is a refinement of the proof of value covering from Chapter 3, matching the bound in Syrgkanis and Tardos [2013] for c-threshold approximate auctions.

Combining revenue covering of an action distribution and value covering gives a welfare approximation result for the strategy profile and setting that produces that action distribution:

Theorem 4.2. *The welfare in any μ-revenue covered strategy profile $\mathbf{s}$ and setting of auction $\mathcal{A}$ is at least a $\frac{\mu}{1-e^{-\mu}}$-approximation to the optimal welfare.*

Proof. Let $\mathbf{x}^*(\mathbf{v})$ be the welfare optimal allocation for valuation profile $\mathbf{v}$. Recall that the optimal allocation is also allowed to use the instantiation of the quality scores and is taken in expectation over the quality scores. Applying the value covering inequality of Equation (4.10) with respect to allocation quantity $x_i^*(\mathbf{v})$ gives that for each bidder i with value v_i,

$$u_i(v_i) + \frac{1}{\mu}T_i(x_i^*(\mathbf{v})) \geq \frac{1-e^{-\mu}}{\mu}x_i^*(\mathbf{v})v_i. \tag{4.11}$$

The quantity $x_i^*(\mathbf{v})v_i$ is exactly agent i's expected contribution to the welfare of the optimal auction. Applying the revenue covering inequality (4.9) for $\mathbf{x}' = \mathbf{x}^*(\mathbf{v})$ and taking

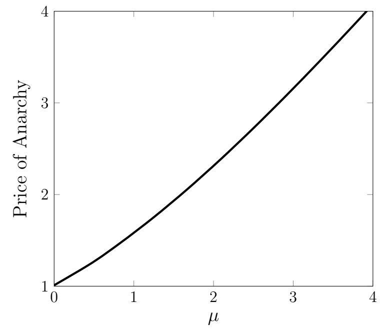

Figure 4.4. The price of anarchy of an auction which is μ-revenue covered, either theoretically or empirically, is $\frac{\mu}{1-e^{-\mu}}$.

expectation over $\mathbf{v}$ yields:

$$\mu \cdot \textsc{Rev}(\mathcal{A}(\mathbf{s}))) \geq \mathbf{E}_{\mathbf{v}}\left[\sum_i T_i(x_i^*(\mathbf{v}))\right]. \tag{4.12}$$

By Equations (4.11) and (4.12) we obtain:

$$\begin{aligned}
\textsc{Util}(\mathcal{A}(\mathbf{s})) + \textsc{Rev}(\mathcal{A}(\mathbf{s})) &\geq \mathbf{E}_{\mathbf{v}}\left[\sum_i u_i(v_i)\right] + \mathbf{E}_{\mathbf{v}}\left[\sum_i \frac{1}{\mu} T_i(x_i^*(\mathbf{v}))\right] \\
&= \sum_i \mathbf{E}_{\mathbf{v}}\left[u_i(v_i) + \frac{1}{\mu} T_i(x_i^*(\mathbf{v}))\right] \\
&\geq \sum_i \mathbf{E}_{\mathbf{v}}\left[\frac{1-e^{-\mu}}{\mu} x_i^*(\mathbf{v}) v_i\right] = \frac{1-e^{-\mu}}{\mu} \textsc{Welfare}(\textsc{Opt})
\end{aligned}$$

Since $\textsc{Welfare}(\mathcal{A}(\mathbf{s})) = \textsc{Rev}(\mathcal{A}(\mathbf{s})) + \textsc{Util}(\mathcal{A}(\mathbf{s}))$, we have our desired result,

$$\textsc{Welfare}(\mathcal{A}(\mathbf{s})) \geq \frac{1-e^{-\mu}}{\mu} \textsc{Welfare}(\textsc{Opt}).$$

$\square$

See Figure 4.4 for a plot of the price of anarchy as a function of μ.

4.3.2. Refining with Observational Data

We now discuss the calculation of μ from the distribution of data $\mathscr{D}$ generated by an auction. The revenue of an auction is observable in $\mathscr{D}$. If we can also upper bound $\sum_i T_i(x_i')$ for any feasible allocation $\mathbf{x}'$, then we have an upper bound on μ. Define $\mathbb{T}$ to be this upperbound, hence $\mathbb{T} = \max_{\mathbf{x}} \sum_i T_i(x_i)$.

Recall that as the auction gets to know the quality scores before deciding the allocation of positions, any feasible allocation corresponds to a quality score dependent assignment of slots to bidders.

If the quality scores γ were deterministic, then we could write

$$\mathbb{T} = \max_{\mathbf{x}} \sum_i T_i(x_i) = \max_{\rho} \sum_i T_i(\mathbf{E}_\gamma \left[\gamma_i \alpha_{\rho^{-1}(\boldsymbol{\gamma},i)}\right]) = \max_{o} \sum_i T_i(\gamma_i \alpha_{o^{-1}(i)}). \tag{4.13}$$

The latter optimization problem would simply be a bipartite weighted matching problem, where the weight of bidder i for position j would be $T_i(\gamma_i \alpha_{o^{-1}(i)})$. However, when the quality scores are random and their distribution has support of size K, then the space of feasible assignments $\mathcal{M}$ has size $(n^m)^K$ and the problem does not have the structure of a matching problem anymore, since the functions $T_i(\cdot)$ are arbitrary convex functions. Solving this complicated maximization problem seems hopeless. In fact it can be shown that the latter problem is NP-hard by a reduction from the maximum hypergraph matching problem, when the size of the support of the correlated distribution of $\boldsymbol{\gamma}$ is not constant. The hardness arises even if each γ_i is either 0 or 1.

However, for the purpose of providing an upper bound on the empirical price of anarchy, it suffices to compute an upper bound on $\mathbb{T}$ and then show that this upper bound is revenue covered. We will use the convexity of functions $T_i(\cdot)$ to provide such an upper bound.

Specifically, let

$$\overline{x_i} = \max_{\mathbf{x}} x_i = \max_{\rho \in \mathcal{M}} \mathbf{E}_\gamma \left[a_{\rho^{-1}(\gamma,i)} \gamma_i \right] = \alpha_1 \, \mathbf{E}_\gamma \left[\gamma_i \right]$$

denote the maximum possible allocation of bidder i. Then observe that by convexity for any feasible x_i: $T_i(x_i) \leq x_i \frac{T_i(\overline{x_i})}{\overline{x_i}}$. Thus we will define:

$$\overline{\mathbb{T}}^1 = \max_{\mathbf{x}} \sum_i x_i \frac{T_i(\overline{x_i})}{\overline{x_i}} \tag{4.14}$$

Then we immediately get the following observation:

Observation 4.1. $\overline{\mathbb{T}}^1 \geq \mathbb{T}$.

Computing $\overline{\mathbb{T}}^1$ is a much easier computational problem than computing $\mathbb{T}$. Specifically, by linearity of expectation:

$$\begin{aligned}
\overline{\mathbb{T}}^1 = \max_{\mathbf{x}} \sum_i x_i \frac{T_i(\overline{x_i})}{\overline{x_i}} &= \max_{\rho \in \mathcal{M}} \sum_i \mathbf{E}_{\boldsymbol{\gamma}} \left[\alpha_{\rho^{-1}(\gamma,i)} \gamma_i \right] \frac{T_i(\overline{x_i})}{\overline{x_i}} \\
&= \max_{\rho \in \mathcal{M}} \mathbf{E}_{\boldsymbol{\gamma}} \left[\sum_i \alpha_{\rho^{-1}(\gamma,i)} \gamma_i \cdot \frac{T_i(\overline{x_i})}{\overline{x_i}} \right] \\
&= \mathbf{E}_{\boldsymbol{\gamma}} \left[\max_{\pi \in \Pi} \sum_i \alpha_{\pi^{-1}(i)} \gamma_i \cdot \frac{T_i(\overline{x_i})}{\overline{x_i}} \right]
\end{aligned}$$

Now observe that the problem inside the expectation is equivalent to a welfare maximization problem where each agent i has a value-per-click of $v_i' = \frac{T_i(\overline{x_i})}{\overline{x_i}}$ and we want to maximize the welfare: $\sum_i \alpha_{\pi^{-1}(i)} \gamma_i \cdot v_i'$. The optimal such allocation is simply the greedy allocation which assigns slots to bidders in decreasing order of $\gamma_i \cdot v_i'$. Thus computing $\overline{\mathbb{T}}^1$ consists of running a greedy allocation algorithm for each quality score profile γ in the support of the distribution of quality scores, which would take time $K \cdot (m + n\log(n))$. When applying it to the data, we will simply compute the optimal greedy allocation for each instance of the quality scores that arrives in each auction (i.e. we compute the latter for the empirical distribution of quality score profiles).

Empirical Revenue Covering. If we can estimate the threshold upper-bound and the revenue, then this is now enough for a revenue-covering result for the strategy profile being played in the auction:

Lemma 4.3. *For any auction $\mathcal{A}$ with action distribution F^a, revenue* $\textsc{Rev}(\mathcal{A})$ *and threshold upper bound $\overline{\mathbb{T}}$, $\mathcal{A}$ is $\frac{\overline{\mathbb{T}}}{\textsc{Rev}(\mathcal{A})}$-revenue covered with action distribution F^a.*

Combining this with Theorem 4.2 directly gives a welfare approximation result:

Corollary 4.4. *For any instance of an auction, with action distribution F^a and (observable) revenue* $\textsc{Rev}(\mathcal{A})$ *and threshold upper bound $\overline{\mathbb{T}}$, the empirical price of anarchy for auction $\mathcal{A}$ is at most*

$$\frac{\overline{\mathbb{T}}}{\textsc{Rev}(\mathcal{A})} \frac{1}{1 - e^{-\overline{\mathbb{T}}/\textsc{Rev}(\mathcal{A})}}. \tag{4.15}$$

Empirical Value Covering. We can also use data to improve the $\frac{1}{1-e^{-\mu}}$ factor in the approximation bound. This term comes from value covering (Lemma 4.1), which analyzes how bidders react to the price-per-click allocation rules they face. In the proof of value covering, it is shown that *no matter what the price-per-click allocation rule* is, it is always the case that $\tilde{u}_i + \frac{1}{\mu}T_i(x_i') \geq \frac{1-e^{-\mu}}{\mu} x_i' v_i$. When we can observe the price-per-click allocation rules, we can simply take the worst case over the price-per-click allocation rules that we observe for each agent, giving an improved price of anarchy result.

Definition 4.5 (Empirical Value Covering)**.** Auction $\mathcal{A}$ and action distribution F^a are empirically λ-value covered if $\mathcal{A}$ is μ-revenue covered, and for any bidder i with value v_i and allocation amount x'_i,

$$u_i(v_i) + \frac{1}{\mu} T_i(x'_i) \geq \frac{\lambda}{\mu} x'_i v_i. \tag{4.16}$$

Lemma 4.5. *If auction $\mathcal{A}$ and action distribution F^a are empirically μ-revenue covered and λ-value covered, then the empirical price of anarchy is at most $\frac{\mu}{\lambda}$.*

Proof. The proof is analogous to the proof of Theorem 4.2, using the value covering parameter λ in place of the general value covering result, Lemma 4.1. □

One approach is to directly look at the threshold curves, and find the worst case ratio of threshold and utility to value over all possible values.

Lemma 4.6. *For a μ-revenue covered auction $\mathcal{A}$ and action distribution F^a, let $\underline{\lambda_i^\mu} = \min_{v_i, x'_i} \frac{\mu u_i(v_i) + T_i(x'_i)}{x'_i v_i}$ and $\underline{\lambda^\mu} = \min_i \underline{\lambda_i^\mu}$.*

Then $\mathcal{A}$ and F^a are empirically $\underline{\lambda^\mu}$-value covered.

In the case that an auction is shown to be μ-revenue covered with respect to the upper bound $\overline{\mathbb{T}}^1$, the maximization can be simplified to only consider the allocation amount $\overline{x_i}$, hence $\underline{\lambda_i^\mu} = \min_{v_i} \frac{\mu \tilde{u}_i + T_i(\overline{x_i})}{v_i}$.

Concentration Bounds. We can also improve on the value covering term even if we only know some properties about the concentration of the price-per-click allocation rule. If the price-per-click allocation rule is highly concentrated, and the minimum feasible price per click is at least a $(1 - \frac{1}{k})$ fraction of the maximum feasible price per click, we can get significantly improved bounds.

μ	$k = 1$	$k = 2$	$k = 4$	$k = 10$	$k = 100$
0.5	1.271	1.	1.	1.	1.
0.75	1.421	1.116	1.	1.	1.
1	1.582	1.302	1.163	1.072	1.009
1.25	1.752	1.506	1.382	1.304	1.256
1.5	1.931	1.717	1.61	1.545	1.505
2	2.313	2.157	2.079	2.032	2.003
4	4.075	4.037	4.019	4.007	4.001
8	8.003	8.001	8.001	8.	8.

Table 4.1. Price of Anarchy bounds when the price-per-click of getting any allocation is at least a $(1 - \frac{1}{k})$ fraction of the price-per-click of getting the maximum allocation, with empirical revenue covering parameter μ.

Lemma 4.7. *For any μ-revenue covered auction $\mathcal{A}$ and action distribution F^a with $\mu \geq 1$, if $\tau(\epsilon) \geq (1 - 1/k)\tau(x_i')$ for any feasible allocation amount x_i' and $\epsilon > 0$, $\mathcal{A}$ and $\mathbf{s}$ are empirically $(1 - 1/k)$-value covered.*

The proof is included in the appendix: see Table 4.1 for better numerical results.

4.4. Data Analysis

We run our analysis on the BingAds auctions. We analyzed eleven phrases from multiple thematic categories. For each phrase we retrieved data of auctions for the phrase for the period of a week. For each phrase and bidder that participated in the auctions for the phrase we computed the allocation curve and by simulating the auctions for the week and computing what would have happened at each auction for each possible bid an advertiser could submit. We discretized the bid space and assumed a hard upper bound on the bid amount.

The left graph in Figure 4.5 shows the allocation curves for a subset of the advertisers for a specific search phrase, the right shows the resulting threshold curves. Most of these keywords have a huge amount of heterogeneity across advertisers as can be seen by the very different bid levels of each advertiser and the very different quality score. For instance, in

Figure 4.6 we depict the average bid, average quality score and average payment of each of the same subset of advertisers for which we depicted the allocation and threshold function in Figure 4.5.

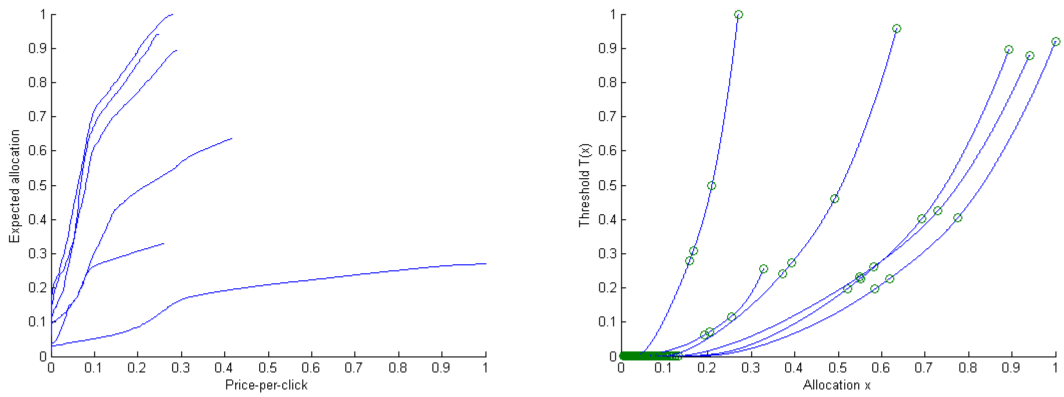

Figure 4.5. Examples of allocation curves (left) and threshold curves (right) for a subset of six advertisers for a specific keyword during the period of a week. All axes are normalized to 1 for privacy reasons. The circles in the left plot correspond to the expected allocation and expected threshold if bidder i was given the j-th slot in all the auctions, i.e. the circle corresponding to the highest allocation and threshold corresponds to the point $(\alpha_1\,\mathbf{E}[\gamma_i], T(\alpha_1\,\mathbf{E}[\gamma_i]))$, the next circle corresponds to $(\alpha_2\,\mathbf{E}[\gamma_i], T(\alpha_2\,\mathbf{E}[\gamma_i]))$, etc.

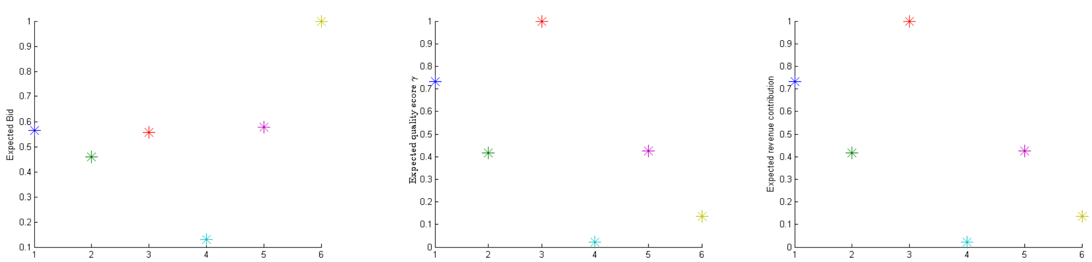

Figure 4.6. Average bid $\mathbf{E}[b_i]$, average quality factor $\mathbf{E}[\gamma_i]$ and average revenue contribution $\mathbf{E}[p_i]$, correspondingly, for the same subset of six advertisers that participated in a specific keyword during the period of a week. Vertical axes are normalized to 1 for privacy reasons.

Subsequently, we applied all the techniques we describe in Section 4.3 for each of the search phrases. We first computed the optimal upper bound on the thresholds $\overline{\mathbb{T}}^1$ and by observing the revenue of the auctions from the data, we can compute an upper bound on the revenue covering of the auction for the phrase, i.e. $\mu^1 = \overline{\mathbb{T}}^1/\textsc{Rev}$. Then for this μ^1 we optimized over λ by using the allocation curves and Lemma 4.6 and assuming some hard upper bound on the valuation of each advertiser and found the optimal such λ, denoted by λ^1. Then an upper bound on the empirical price of anarchy is μ^1/λ^1.

Subsequently we tested the tightness of our analysis by computing the value of the true thresholds on the optimal allocation that was computed under the linear approximations of the thresholds. This is a feasible allocation and hence the true value of $\mathbb{T}$ is at least the value of the thresholds for this allocation. Hence, by looking at the value of the thresholds at this allocation, denoted by $LB - \mathbb{T}$ we can check how good our approximation of $\mathbb{T}$ is $\overline{\mathbb{T}}^1$. Then we also computed the optimal thresholds for any quality score independent allocation rule. Apart from yielding yet another lower bound for $\mathbb{T}$, the latter analysis also yields an empirical price of anarchy with respect to such a handicapped optimal welfare, which can also be used as a welfare benchmark.

We portray our results on these quantities for each of the eleven search phrases in Table 4.2.

	$\frac{1}{EPoA^1}$	$\frac{\overline{\mathbb{T}}^1}{\textsc{Rev}}$	λ^1	$\frac{LB-\mathbb{T}}{\textsc{Rev}}$	$\frac{1}{LB-EPoA}$	$\frac{1}{FA-EPoA}$
phrase1	.567	1.4164	.803	.562	.788	.783
phrase2	.606	1.2848	.779	.553	.792	.784
phrase3	.279	4.182	1.167	3.364	.325	.320
phrase4	.284	3.860	1.098	2.298	.401	.507
phrase5	.673	1.099	.740	.437	.828	.829
phrase6	.628	1.183	.743	.495	.832	.791
phrase7	.672	1.031	.693	.503	.824	.802
phrase8	.645	1.036	.669	.520	.795	.817
phrase9	.622	1.169	.726	.581	.759	.809
phrase10	.597	1.138	.680	.545	.771	.833
phrase11	.573	1.431	.820	.631	.780	.786

Table 4.2. Empirical Price of Anarchy analysis for a set of eleven search phrases on the BingAds system. Phrases are grouped together according to the thematic category of the search phrase. The columns have the following interpretation: 1) $EPoA^1$ is the upper bound on the empirical price of anarchy, i.e. if $1/EPoA^1$ is x it means that the welfare of the auction is at least $x \cdot 100\%$ efficient. This lower bound is computed by using the polynomially computable upper bound $\overline{\mathbb{T}}^1$ of $\mathbb{T}$ and then also optimizing over λ. 2) $\mu^1 = \overline{\mathbb{T}}^1/\textsc{Rev}$ is the ratio of the upper bound on the maximum sum of thresholds over the revenue of the auction. 3) λ^1 is the minimum lambda across advertisers after running the optimization problem presented in Lemma 4.6 for the allocation curve of each advertiser, assuming some upper bound on the value. Then based on Lemma 4.5, $EPoA^1 = \lambda^1/\mu^1$. 4) $LB - \mathbb{T}/\textsc{Rev}$: we use the optimal allocation computed by assuming the linear form of thresholds used for $\overline{\mathbb{T}}^1$. Then we evaluate the true thresholds on this allocation. This is a feasible allocation and hence the value of the thresholds on this allocation, denoted $LB - \mathbb{T}$ is a lower bound on the value of $\mathbb{T}$. Thus this ratio is a lower bound on how well the auction is revenue covered. 5) $LB - EPoA$, this is simply the empirical price of anarchy bound that would have been implied if $\mathbb{T} = LB - \mathbb{T}$ and even if we optimized over λ. Thus $\frac{1}{LB-EPoA}$ is an upper bound on how good our efficiency bound could have been even if we solved the hard problem of computing $\mathbb{T}$. 6) $\underline{\mathbb{T}}$, this corresponds to the optimal thresholds with respect to any quality score independent feasible allocation as defined in Section 4.3. 7) $FA - EPoA$ a bound on the empirical price of anarchy with respect to a quality score independent allocation rule. For this price of anarchy we did not optimize over λ, hence $FA - EPoA = \frac{\mu}{1-\exp(-\mu)}$.

CHAPTER 5

Risk-Averse Bidders

In this chapter, we study simple and approximately optimal auctions for agents with two forms of risk-averse preferences. For capacitated agents, we show that the revenue from the first-price auction is approximately optimal. For CARA agents, we show the revenue from the first-price auction is approximately optimal compared to the first-price auction with a reserve price. Both results rely on adapting aspects of Myerson's [1981] characterization of revenue in risk-neutral settings to risk-averse settings.

5.1. Introduction

We study optimal and approximately optimal auctions for agents with risk-averse preferences. The economics literature on this subject is largely focused on either comparative statics, i.e., is the first-price or second-price auction better when agents are risk averse, or deriving the optimal auction, e.g., using techniques from optimal control, for specific distributions of agent preferences. The former says nothing about optimality but considers realistic prior-independent auctions; the latter says nothing about realistic and prior-independent auctions. Our goal is to study approximately optimal auctions for risk-averse agents that are realistic and not dependent on assumptions on the specific form of the distribution of agent preferences. One of our main conclusions is that, while the second-price auction can be very far from optimal for risk-averse agents, the first-price auction is approximately optimal for an interesting class of risk-averse preferences.

The microeconomic treatment of risk aversion in auction theory suggests that the form of the optimal auction is very dependent on precise modeling details of the preferences of agents, see, e.g., Maskin and Riley [1984] and Matthews [1984]. The resulting auctions are unrealistic because of their reliance on the prior assumption and because they are complex [cf. Wilson, 1987]. Approximation can address both issues. There may be a class of mechanisms that is simple, natural, and much less dependent on exact properties of the distribution. As an example of this agenda for risk neutral agents, Hartline and Roughgarden [2009] showed that for a large class of distributional assumptions the second-price auction with a reserve is a constant approximation to the optimal single-item auction. This implies that the only information about the distribution of preferences that is necessary for a good approximation is a single number, i.e., a good reserve price. Often from this sort of "simple versus optimal" result it is possible to do away with the reserve price entirely. Dhangwatnotai et al. [2010] and Roughgarden et al. [2012] show that simple and natural mechanisms are approximately optimal quite broadly. We extend this agenda to auction theory for risk-averse agents.

The least controversial approach for modeling risk-averse agent preferences is to assume agents are endowed with a concave function that maps their wealth to a utility. This introduces a non-linearity into the incentive constraints of the agents which in most cases makes auction design analytically intractable.

In this work, we focus on two forms of utility functions: *capacitated*, and *constant absolute risk aversion* (CARA) utility functions.

In a capacitated utility function, the utility is risk-neutral until the agent hits her capacity, then flat. Importantly, an agent with such a utility function will not trade off a higher probability of winning for a lower price when the utility from such a lower price is greater than her capacity. While capacitated utility functions are unrealistic, they form a basis for general concave utility functions. In our analyses we will endow the benchmark optimal auction with knowledge of the agents' value distribution and capacity; however, some of the mechanisms we design to approximate this benchmark will be oblivious to them.

We will also attain results comparing the first-price auction to the first-price auction with the optimal reserve price for agents with Constant Absolute Risk Aversion (CARA) preferences. In CARA utility functions, agents' risk attitudes do not change with shifts in wealth. If an agent prefers a constant \$4 to a 50% $\$0, 50\%$ \$10 lottery, she will also prefer a constant \$1004 to 50% $\$1000, 50\%$ \$1010 lottery.

As an illustrative example, consider the problem of maximizing welfare by a single-item auction when agents have known capacitated utility functions (but unknown values). Recall that for risk-neutral agents the second-price auction is welfare-optimal as the payments are

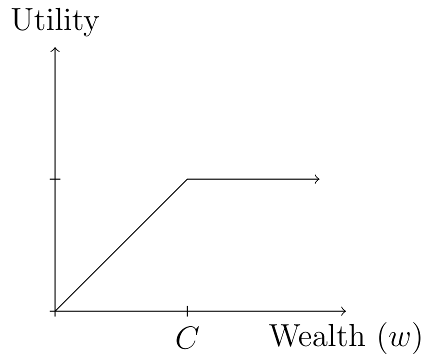

Figure 5.1. The capacitated utility function, $U_C(w) = \min(w, C)$.

transfers from the agents to the mechanism and cancel from the objective welfare which is thus equal to value of the winner. (The auctioneer is assumed to have linear utility.) For agents with capacitated utility, the second-price auction can be far from optimal. For instance, when the difference between the highest and second highest bid is much larger than the capacity then the excess value (beyond the capacity) that is received by the winner does not translate to extra utility because it is truncated at the capacity. Instead, a variant of the second-price auction, where the highest bidder wins and is charged the maximum of the second highest bid and her bid less her capacity, obtains the optimal welfare. Unfortunately, this auction is parameterized by the form of the utility function of the agents. There is, however, an auction, not dependent on specific knowledge of the utility functions or prior distribution, that is also welfare optimal: If the agents values are drawn i.i.d. from a common prior distribution then the first-price auction is welfare-optimal. To see this: (a) standard analyses show that at equilibrium the highest-valued agent wins, and (b) no agent will shade her bid more than her capacity as she receives no increased utility from such a lower payment but her probability of winning strictly decreases.

Our main goal is to duplicate the above observation for the objective of revenue. It is easy to see that the gap between the optimal revenues for risk-neutral and capacitated agents can be of the same order as the gap between the optimal welfare and the optimal revenue (which can be unbounded). When the capacities are small the revenue of the welfare-optimal auction for capacitated utilities is close to its welfare (the winners utility is at most her capacity). When capacities are infinite or very large then the risk-neutral optimal revenue is close to the capacitated optimal revenue (the capacities are not binding). One of our main technical results shows that even for mid-range capacities one of these two mechanisms that are optimal at the extremes is close to optimal.

As a first step towards understanding profit maximization for capacitated agents, we characterize the optimal auction for agents with capacitated utility functions. We then give a “simple versus optimal” result showing that either the revenue-optimal auction for risk-neutral agents or the above welfare-optimal auction for capacitated agents is a good approximation to the revenue-optimal auction for capacitated agents. In symmetric settings, the Bulow-Klemperer [1996] Theorem implies that with enough competition (and mild distributional assumptions) welfare-optimal auctions are approximately revenue-optimal. The first-price auction is welfare-optimal and prior-independent; therefore we conclude that it is approximately revenue-optimal for capacitated agents in symmetric settings, and leveraging the approach of Chapter 3, approximately optimal in asymmetric settings.

Our “simple versus optimal” result (which holds for both symmetric and asymmetric settings) comes from an upper bound on the expected payment of an agent in terms of her allocation rule [cf. Myerson, 1981]. This upper bound is the most technical result in the paper; the difficulties that must be overcome by our analysis are exemplified by the following observations. First, unlike in risk-neutral mechanism design, Bayes-Nash equilibrium does not imply monotonicity of allocation rules. There are mechanisms where an agent with a high value would prefer less overall probability of service than she would have obtained if she had a lower value (Example 5.1 in Section 5.3). Second, even in the case where the capacity is higher than the maximum possible value of any agent, the optimal mechanism for risk-averse agents can generally obtain more revenue than the optimal mechanism for risk-neutral agents (Example 5.3 in Section 5.4). This may be surprising because, in such a case, the revenue-optimal mechanism for risk-neutral agents would give any agent a wealth that is within the linear part of her utility function. Finally, while our upper bound on risk-averse payments implies that this relative improvement is bounded by a factor of two for large capacities, it can be arbitrarily large for small capacities (Example 5.2 in Section 5.4).

It is natural to conjecture that the first-price auction will continue to perform nearly optimally well beyond capacitated and CARA preferences. It is a relatively straightforward calculation to see that for a large class of risk-averse utility functions from the literature [e.g., Matthews, 1984] the first-price auction is approximately optimal at extremal risk parameters (risk-neutral or extremely risk-averse). We leave to future work the extension of our analysis to mid-range risk parameters for these other families of risk-averse utility functions.

It is significant and deliberate that our main theorem is about the first-price auction which is well known to not have a truthtelling equilibrium. Our goal is a prior-independent mechanism. In particular, we would like our mechanism to be parameterized neither by the distribution on agent preference nor by the capacity that governs the agents utility function. While it is standard in mechanism design and analysis to invoke the *revelation principle* [cf. Myerson, 1981] and restrict attention to auctions with truthtelling as equilibrium, this principle cannot be applied in prior-independent auction design. An auction with good equilibrium can be implemented by one with truthtelling as an equilibrium if the agent strategies can be simulated by the auction. In a Bayesian environment, agent strategies are parameterized by the prior distribution and therefore the suggested revelation mechanism is not generally prior independent.

Risk Aversion, Universal Truthfulness, and Truthfulness in Expectation. Our results have an important implication on a prevailing and questionable perspective that is explicit and implicit broadly in the field of algorithmic mechanism design. Two standard solution concepts from algorithmic mechanism design are "universal truthfulness" and "truthfulness in expectation." A mechanism is universally truthful if an agent's optimal (and dominant) strategy is to reveal her values for the various outcomes of the mechanism regardless of the reports of other agents or random coins flipped by the mechanism. In contrast, in a truthful-in-expectation mechanism, revealing truthfully her values only maximizes the agent's utility in expectation over the random coins tossed by the mechanism. Therefore, a risk-averse agent modeled by a non-linear utility function may not bid truthfully in a truthful-in-expectation mechanism designed for risk-neutral agents, whereas in a universally truthful mechanism an agent behaves the same regardless of her risk attitude. For this reason, the above-mentioned perspective sees universally truthful mechanisms superior because the performance guarantees shown for risk-neutral agents seem to apply to risk-averse agents as well.

This perspective is incorrect because the optimal performance possible by a mechanism is different for risk-neutral and risk-averse agents. In some cases, a mechanism may exploit the risk attitude of the agents to achieve objectives better than the optimal possible for risk-neutral agents; in other cases, the objective itself relies on the utility functions (e.g. social welfare maximization), and therefore the same outcome has a different objective value. In all these situations, the performance guarantee of universally truthful mechanisms measured by the risk-neutral optimality loses its meaning. We have already discussed above two examples for capacitated agents that illustrate this point: for welfare maximization the second-price auction is not optimal, for revenue maximization the risk-neutral revenue-optimal auction can be far from optimal.

The conclusion of the discussion above is that the universally truthful mechanisms from the literature are not generally good when agents are risk averse; therefore, the solution concept of universal truthfulness buys no additional guarantees over truthfulness in expectation. Nonetheless, our results suggest that it may be possible to develop a general theory for prior-independent mechanisms for risk-averse agents. By necessity, though, this theory will look different from the existing theory of algorithmic mechanism design.

Summary of Results. Our main theorem is that the first-price auction is a prior-independent 5-approximation for revenue for two or more agents with i.i.d. values and risk-averse preferences (given by a common capacity). The technical results that enable this theorem are as follows:

- The optimal auction for agents with capacitated utilities is a two-priced mechanism where a winning agent either pays her full value or her value less her capacity.
- The expected revenue of an agent with capacitated utility and regular value distribution can be bounded in terms of an expected (risk-averse) virtual surplus, where the (risk-averse) virtual value is twice the risk-neutral virtual value plus the value minus capacity (if positive).
- Either the mechanism that optimizes value minus capacity (and charges the Clarke payments or value minus capacity, whichever is higher) or the risk-neutral revenue optimal mechanism is a 3-approximation to the revenue optimal auction for capacitated utilities.

- We characterize the Bayes-Nash equilibria of auctions with capacitated agents where each bidder's payment when served is a deterministic function of her value. An example of this is the first-price auction. The BNE strategies of the capacitated agents can be calculated formulaically from the BNE strategies of risk-neutral agents.

Some of these results extend beyond single-item auctions. In particular, the characterization of equilibrium in the first-price auction holds for position auction environments (i.e., where agents are assigned to positions greedily by bid with decreasing probabilities of service and charged their bid if served). If valuations are symmetric in a position auction, then our prior-independent 5-approximation (Theorem 5.15) holds. Our simple-versus-optimal 3-approximation (Theorem 5.17) holds generally for downward-closed environments, non-identical distributions, and non-identical capacities.

For CARA agents in symmetric settings that satisfy a regularity condition, the revenue of the first-price auction is an $\frac{n-1}{n}$ approximation to the revenue of the first-price auction with the optimal reserve. Note that this exactly matches the risk-neutral bound from Theorem 2.9.

Related Work. The comparative performance of first- and second-price auctions in the presence of risk aversion has been well studied in the Economics literature. From a revenue perspective, first-price auctions are shown to outperform second-price auctions very broadly. Riley and Samuelson [1981] and Holt [1980] show this for symmetric settings where bidders have the same concave utility function. Maskin and Riley [1984] show this for more general preferences.

Matthews [1987] shows that in addition to the revenue dominance, bidders whose risk attitudes exhibit *constant absolute risk aversion (CARA)* are indifferent between first- and second-price auctions, even though they pay more in expectation in the first-price auction. Hu et al. [2010] considers the optimal reserve prices to set in each, and shows that the optimal reserve in the first price auction is less than that in the second price auction. Interestingly, under light conditions on the utility functions, as risk aversion increases, the optimal first-price reserve price decreases.

Matthews [1983] and Maskin and Riley [1984] have considered optimal mechanisms for a single item, with symmetric bidders (i.i.d. values and identical utility function), for CARA and more general preferences. Both approaches apply only when the optimal auction involves a deterministic price upon winning, which is not satisfied in our setting.

Dughmi and Peres [2012] have shown that by insuring bidders against uncertainty, any truthful-in-expectation mechanism for risk-neutral agents can be converted into a dominant-strategy incentive compatible mechanism for risk-averse buyers with no loss of revenue. However, there is potentially much to gain—mechanisms for risk-averse buyers can achieve unboundedly more welfare and revenue than mechanisms for risk-neutral bidders, as we show in Example 5.2 of Section 5.4.

5.2. Preliminaries

Risk-Averse Agents. Consider selling an item to an agent who has a private valuation v drawn from a known distribution F. Denote the outcome by (x, p), where $x \in \{0, 1\}$ indicates whether the agent gets the item, and p is the payment made. The agent obtains a *wealth* of $vx - p$ for such an outcome and the agent's utility is given by a concave utility function $U(\cdot)$ that maps her wealth to utility, i.e., her utility for outcome (x, p) is $U(vx - p)$. Concave utility functions are a standard approach for modeling risk-aversion.[1]

A *capacitated* utility function is $U_C(z) = \min(z, C)$ for a given C which we refer to as the *capacity*. Intuitively, small C relative to value corresponds to severe risk aversion; large C

[1]There are other definitions of risk aversion; this one is the least controversial. See Mas-Colell et al. [1995] for a thorough exposition of expected utility theory.

corresponds to mild risk aversion; and $C = \infty$ corresponds to risk neutrality. An agent views an auction as a deterministic rule that maps a random source and the (possibly random) reports of other agents which we summarize by π, and the report b of the agent, to an allocation and payment. We denote these coupled allocation and payment rules as $x^\pi(b)$ and $p^\pi(b)$, respectively. The agent wishes to maximize her expected utility which is given by $\mathbf{E}_\pi[U_C(vx^\pi(b) - p^\pi(b))]$, i.e., she is a von Neumann-Morgenstern utility maximizer.

Incentives. A strategy profile of agents is $\mathbf{s} = (s_1, \ldots, s_n)$ mapping values to reports. Such a strategy profile is in *Bayes-Nash equilibrium* (BNE) if each agent i maximizes her utility by reporting $s_i(v_i)$. I.e., for all i, v_i, and z:

$$\mathbf{E}_\pi \left[U(v_i x_i^\pi(s_i(v_i)) - p_i^\pi(s_i(v_i)))\right] \geq \mathbf{E}_\pi \left[U(v_i x_i^\pi(z) - p_i^\pi(z))\right]$$

where π denotes the random bits accessed by the mechanism as well as the random inputs $s_j(v_j)$ for $j \neq i$ and $v_j \sim F_j$. A mechanism is *Bayesian incentive compatible* (BIC) if truthtelling is a Bayes-Nash equilibrium: for all i, v_i, and z

$$\mathbf{E}_\pi \left[U(v_i x_i^\pi(v_i) - p_i^\pi(v_i))\right] \geq \mathbf{E}_\pi \left[U(v_i x_i^\pi(z) - p_i^\pi(z))\right] \tag{IC}$$

where π denotes the random bits accessed by the mechanism as well as the random inputs $v_j \sim F_j$ for $j \neq i$.

We will consider only mechanisms where losers have no payments, and winners pay at most their bids. These constraints imply ex post *individual rationality* (IR). Formulaically, for all i, v_i, and π, $p_i^\pi(v_i) \leq v_i$ when $x_i^\pi(v_i) = 1$ and $p_i^\pi(v_i) = 0$ when $x_i^\pi(v_i) = 0$.

Auctions and Objectives. The revenue of an auction $\mathcal{M}$ is the total payment of all agents; its expected revenue for implicit distribution F and Bayes-Nash equilibrium is denoted

$\textsc{Rev}(\mathcal{M}) = \mathbf{E}_{\pi,\mathbf{v}}[\sum_i p_i^\pi(v_i)]$. The welfare of an auction $\mathcal{M}$ is the total utility of all participants including the auctioneer; its expected welfare is denoted $\textsc{Welfare}(\mathcal{M}) = \textsc{Rev}(\mathcal{M}) + \mathbf{E}_{\pi,\mathbf{v}}[\sum_i U(v_i x_i^\pi(v_i) - p_i^\pi(v_i))]$.

Some examples of auctions are: the *first-price auction* (FPA) serves the agent with the highest bid and charges her her bid; the *second-price auction* (SPA) serves the agent with the highest bid and charges her the second-highest bid. The second price auction is incentive compatible regardless of agents' risk attitudes. The *capacitated second-price auction* (CSP) serves the agent with the highest bid and charges her the maximum of her value less her capacity and the second highest bid. The second-price auction for capacitated agents is incentive compatible for capacitated agents because, relative to the second-price auction, the utility an agent receives for truthtelling is unaffected and the utility she receives for any misreport is only (weakly) lower.

Two-Priced Auctions. The following class of auctions will be relevant for agents with capacitated utility functions.

Definition 5.1. A mechanism $\mathcal{M}$ is *two-priced* if, whenever $\mathcal{M}$ serves an agent with capacity C and value v, the agent's payment is either v or $v - C$; and otherwise (when not served) her payment is zero. Denote by $x_{\text{val}}(v)$ and $x_C(v)$ probability of paying v and $v - C$, respectively.

Note that from an agent's perspective the outcome of a two-priced mechanism is fully described by a x_C and x_{val}.

Auction Theory for Risk-Neutral Agents. For risk-neutral agents, i.e., with $U(\cdot)$ equal to the identity function, only the probability of winning and expected payment are relevant. The *interim allocation rule* and *interim payment rule* are given by the expectation of x^π and p^π

over π and denoted as $x(b) = \mathbf{E}_\pi[x^\pi(b)]$ and $p(b) = \mathbf{E}_\pi[p^\pi(b)]$, respectively (recall that π encodes the randomization of the mechanism and the reports of other agents).

For risk-neutral agents, Myerson [1981] characterized interim allocation and payment rules that arise in BNE and solved for the revenue optimal auction. These results are summarized in the following theorem.

Theorem 5.1 (Myerson, 1981)**.** *For risk-neutral bidders with valuations drawn independently and identically from F,*

(1) *(monotonicity) The allocation rule $x(v)$ for each agent is monotone non-decreasing in v.*

(2) *(payment identity) The payment rule satisfies $p(v) = vx(v) - \int_0^v x(z)\mathrm{d}z$.*

(3) *(virtual value) The ex ante expected payment of an agent is $\mathbf{E}_v[p(v)] = \mathbf{E}_v[\varphi(v)x(v)]$ where $\varphi(v) = v - \frac{1-F(v)}{F(v)}$ is the* virtual value *for value v.*

(4) *(optimality) When the distribution F is* regular*, i.e., $\varphi(v)$ is monotone, the second-price auction with reserve $\varphi^{-1}(0)$ is revenue-optimal.*

The payment identity in part 2 implies the *revenue equivalence* between any two auctions with the same BNE allocation rule.

A well-known result by Bulow and Klemperer shows that, in part 4 of Theorem 5.1, instead of having a reserve price to make the second-price auction optimal, one may as well add in another identical bidder to get at least as much revenue.

Theorem 5.2 (Bulow and Klemperer, 1996)**.** *For risk-neutral bidders with valuations drawn i.i.d. from a regular distribution, the revenue from the second-price auction with $n+1$ bidders is at least that of the optimal auction for n bidders.*

5.3. The Optimal Auctions

In this section we study the form of optimal mechanisms for capacitated agents. In Section 5.3.1, we show that it is without loss of generality to consider two-priced auctions, and in Section 5.3.2 we characterize the incentive constraints of two-priced auctions. In Section 5.3.3 we use this characterization to show that the optimal auction (in discrete type spaces) can be computed in polynomial time in the number of types.

5.3.1. Two-priced Auctions Are Optimal

Recall a two-priced auction is one where when any agent is served she is either charged her value or her value minus her capacity. We show below that restricting our attention to two-priced auctions is without loss for the objective of revenue.

Theorem 5.3. *For any BIC auction with capacitated agents (with heterogeneous capacities) there is a two-priced auction with no lower revenue.*

Proof. We prove this theorem in two steps. In the first step we show, quite simply, that if an agent with a particular value received more wealth than C then we can truncate her wealth to C (by charging her more). With her given value she is indifferent to this change, and for all other values this change makes misreporting this value (weakly) less desirable. Therefore, such a change would not induce misreporting and only (weakly) increases revenue. This first step gives a mechanism wherein every agent's wealth is in the linear part of her utility function. The second step is to show that we can transform the distribution of wealth into a two point distribution. Whenever an agent with value v is offered a price that results in a wealth $w \in [0, C]$, we instead offer her a price of $v - C$ with probability w/C, and a price of v with the remaining probability. Both the expected revenue and the utility of a truthful bidder is unchanged. The expected utility of other types to misreport v, however, weakly decreases by the concavity of U_C, because mixing over endpoints of an interval on a concave function gives less value than mixing over internal points with the same expectation. $\square$

5.3.2. Characterization of Two-Priced Auctions

In this section we characterize the incentive constraints of two-priced auctions. We focus on the induced two-priced mechanism for a single agent given the randomization π of other agent values and the mechanism. The interim two-priced allocation rule of this agent is denoted by $x(v) = x_{\text{val}}(v) + x_C(v)$.

Lemma 5.4. *A mechanism with two-price allocation rule $x = x_{\text{val}} + x_C$ is BIC if and only if for all v and v^+ such that $v < v^+ \leq v + C$,*

$$\frac{x_{\text{val}}(v)}{C} \leq \frac{x_C(v^+) - x_C(v)}{v^+ - v} \leq \frac{x(v^+)}{C}. \tag{5.1}$$

Equation (5.1) can be equivalently written as the following two linear constraints on x_C, for all $v^- \leq v \leq v^+ \in [v - C, v + C]$:

$$x_C(v^+) \geq x_C(v) + \frac{v^+ - v}{C} \cdot x_{\text{val}}(v), \tag{5.2}$$

$$x_C(v^-) \geq x_C(v) - \frac{v - v^-}{C} \cdot x(v). \tag{5.3}$$

Equations (5.2) and (5.3) are illustrated in Figure 5.2. For a fixed v, (5.2) with $v^+ = v + C$ yields a lower bounding line segment from $(v, x_C(v))$ to $(v + C, x_C(v) + x_{\text{val}}(v))$, and (5.3) with $v^- = v - C$ gives a lower bounding line segment from $(v, x_C(v))$ to $(v - C, x_C(v) - x(v))$. Note that (5.2) implies that x_C is monotone.

In the special case when x_C is differentiable, by taking v^+ approaching v in (5.1), we have $\frac{x_{\text{val}}(v)}{C} \leq x'_C(v) \leq \frac{x(v)}{C}$ for all v. In general, we have the following condition in the integral form (see Section D.1 for a proof).

Corollary 5.5. *The allocation rule $x = x_{\text{val}} + x_C$ of a BIC two-priced mechanism for all $v < v^+$ satisfies:*

$$\int_v^{v^+} \frac{x_{\text{val}}(z)}{C}\, \mathrm{d}z \leq x_C(v^+) - x_C(v) \leq \int_v^{v^+} \frac{x(z)}{C}\, \mathrm{d}z. \tag{5.4}$$

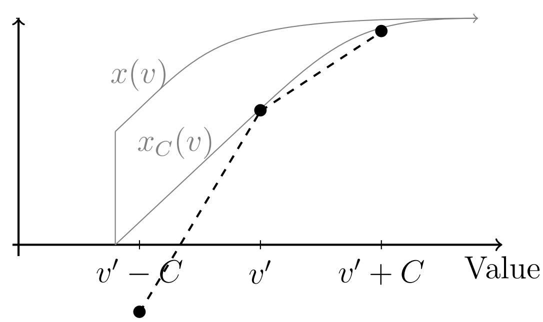

Figure 5.2. Fixing $x(v') = x_{\text{val}}(v') + x_C(v')$, the dashed line between points $(v' - C, x_C(v') - x(v'))$, $(v', x_C(v'))$, and $(v' + C, x_C(v') + x_{\text{val}}(v'))$ (denoted by "•") depicts the lower bounds from (5.2) and (5.3) on x_C for values in $[v' - C, v' + C]$.

Importantly, the equilibrium characterization of two-priced mechanisms does not imply monotonicity of the allocation rule $\mathbf{x}$. This is in contrast with mechanisms for risk-neutral agents, where incentive compatibility requires a monotone allocation rule (Theorem 5.1, part 1). This non-monotonicity is exhibited in the following example.

Example 5.1. *There is a single-agent two-priced mechanism with a non-monotone allocation rule.* Our agent has two possible values $v = 3$ and $v = 4$, and capacity C of 2. We give a two price mechanism. Recall that $x_C(v)$ is the probability with which the mechanism sells the item and charges $v - C$; $x_{\text{val}}(v)$ is the probability with which the mechanism sells the item and charges v; and $x(v) = x_C(v) + x_{\text{val}}(v)$. The mechanism and its outcome are summarized in the following table.

v	x	x_C	x_{val}	utility from truthful reporting	utility from misreporting
3	5/6	1/2	1/3	1	2/3
4	2/3	2/3	0	4/3	4/3

5.3.3. Optimal Auction Computation

Solving for the optimal mechanism is computationally tractable for any discrete (explicitly given) type space D. Given a discrete valuation distribution on support D, one can use $2\,|D|$ variables to represent the allocation rule of any two-priced mechanism, and the expected revenue is a linear sum of these variables. Lemma 5.4 shows that one can use $O(|D|^2)$ linear constraints to express all BIC allocations, and hence the revenue optimization for a single bidder can be solved by a $O(|D|^2)$-sized linear program. Furthermore, using techniques developed by Cai et al. [2012] and Alaei et al. [2012], in particular the "token-passing" characterization of single-item auctions by Alaei et al. [2012], we obtain:

Theorem 5.6. *For n bidders with independent valuations with type spaces $D_1, \cdots, D_n$ and capacities $C_1, \cdots, C_n$, one can solve for the optimal single-item auction with a linear program of size $O\big((\sum_i |D_i|)^2\big)$.*

5.4. An Upper Bound on Two-Priced Expected Payment

In this section we will prove an upper bound on the expected payment from any capacitated agent in a two-priced mechanism. This upper bound is analogous in purpose to the identity between expected risk-neutral payments and expected virtual surplus of Myerson

[1981] from which optimal auctions for risk-neutral agents are derived. We use this bound in Section 5.5.2 and Section 5.5.3 to derive approximately optimal mechanisms.

As before, we focus on the induced two-priced mechanism for a single agent given the randomization π of other agent values and the mechanism. The expected payment of a bidder of value v under allocation rule $x(v) = x_C(v) + x_{\text{val}}(v)$ is $p(v) = v \cdot x_{\text{val}}(v) + (v - C) \cdot x_C(v) = v \cdot x(v) - C \cdot x_C(v)$.

Recall from Theorem 5.1 that the (risk-neutral) virtual value for an agent with value drawn from distribution F is $\varphi(v) = v - \frac{1-F(v)}{F(v)}$ and that the expected risk-neutral payment for allocation rule $\mathbf{x}(\cdot)$ is $\mathbf{E}_v[\varphi(v)\mathbf{x}(v)]$. Denote $\max(0, \varphi(v))$ by $\varphi^+(v)$ and $\max(v - C, 0)$ by $(v - C)^+$.

Theorem 5.7. *For any agent with value $v \sim F$, capacity C, and two-priced allocation rule $x(v) = x_{\text{val}}(v) + x_C(v)$,*

$$\mathbf{E}_v\left[p(v)\right] \leq \mathbf{E}_v\left[\varphi^+(v) \cdot x(v)\right] + \mathbf{E}_v\left[\varphi^+(v) \cdot x_C(v)\right]$$
$$+ \mathbf{E}_v\left[(v - C)^+ \cdot x_C(v)\right].$$

Corollary 5.8. *When bidders have regular distributions and a common capacity, either the risk-neutral optimal auction or the capacitated second price auction (whichever has higher revenue) gives a 3-approximation to the optimal revenue for capacitated agents.*

Proof. For each of the three parts of the revenue upper bound of Theorem 5.7, there is a simple auction that optimizes the expectation of the part across all agents. For the first two parts, the allocation rules across agents (both for $x(\cdot)$ and $x_C(\cdot)$) are feasible. When the distributions of agent values are regular (i.e., the virtual value functions are monotone), the risk-neutral revenue-optimal auction optimizes virtual surplus across all feasible allocations (i.e., expected virtual value of the agent served); therefore, its expected revenue upper bounds the first and second parts of the bound in Theorem 5.7. The revenue of the third part is again the expectation of a monotone function (in this case $(v - C)^+$) times the service probability. The auction that serves the agent with the highest (positive) "value minus capacity" (and charges the winner the maximum of her "minimum winning bid," i.e., the second-price payment rule, and her "value minus capacity") optimizes such an expression over all feasible allocations; therefore, its revenue upper bounds this third part of the bound in Theorem 5.7. When capacities are identical, this auction is the capacitated second price auction. $\square$

Before proving Theorem 5.7, we give two examples. The first shows that the gap between the revenue of the capacitated second-price auction and the risk-neutral revenue-optimal auction (i.e., the two auctions from Corollary 5.8) can be arbitrarily large. This means that there is no hope that an auction for risk-neutral agents always obtains good revenue for risk-averse agents. The second example shows that even when all values are bounded from above by the capacity (and therefore, capacities are never binding in a risk-neutral auction) an auction for risk-averse agents can still take advantage of risk aversion to generate higher revenue. Consequently, the fact that we have two risk-neutral revenue terms in the bound of Theorem 5.7 is necessary (as the "value minus capacity" term is zero in this case).

Example 5.2. *Auctions with capacitated agents can achieve unboundedly more revenue than with risk-neutral agents.* The *equal revenue* distribution on interval $[1, h]$ has distribution function $F(z) = 1 - 1/z$ (with a point mass at h). The distribution gets its name because such an agent would accept any offer price of p with probability $1/p$ and generate an expected revenue of one. With one such agent the optimal risk-neutral revenue is one.

Of course, an agent with capacity $C = 1$ would happily pay her value minus her capacity to win all the time (i.e., $x(v) = x_C(v) = 1$). The revenue of this auction is $\mathbf{E}[v] - 1 = \ln h$. For large h, this is unboundedly larger than the revenue we can obtain from a risk-neutral agent with the same distribution.

Example 5.3. *The revenue from a two-priced mechanism can be better than the optimal risk-neutral revenue even when all values are no more than the capacity.* Consider selling to an agent with capacity of $C = 1000$ and value drawn from the equal revenue distribution from Example Example 5.2 with $h = 1000$.

The following two-priced rule is BIC and generates revenue of approximately 1.55 when selling to such a bidder. Let $x_C(v) = \frac{0.6}{1000}(v-1)$, $x(v) = \min(x_C(v) + 0.6, 1)$, and $x_{\text{val}}(v) = x(v) - x_C(v)$ (shown in Figure 5.3). Recall that the expected payment from an agent with value v can be written as $vx(v) - Cx_C(v)$; for small values, this will be approximately 0.6; for large values this will increase to 400. The expected revenue is

$$\int_1^{1000} (z \cdot x(z) - 1000\; x_C(z))\, F(z)\mathrm{d}z + \frac{1}{1000}(1000 \cdot x_{\text{val}}(1000)) \approx 1.55.$$

The optimal risk-neutral revenue is 1, giving our desired bound.

In the remainder of this section we instantiate the following outline for the proof of Theorem 5.7. First, we transform any given two-priced allocation rule $x = x_{\text{val}} + x_C$ into a new two-priced rule $\bar{x}(v) = \bar{x}_C(v) + \bar{x}_{\text{val}}(v)$ (for which the expected payment is $\bar{p}(v) = v\bar{x}(v) - C\bar{x}_C(v)$). While this transformation may violate some incentive constraints (from Lemma 5.4), it enforces convexity of $\bar{x}_C(v)$ on $v \in [0, C]$ and (weakly) improves revenue. Second, we derive a simple upper bound on the payment rule $\bar{p}(\cdot)$. Finally, we use the enforced convexity property of $\bar{x}_C(\cdot)$ and the revenue upper bound to partition the expected payment $\mathbf{E}_v[\bar{p}(v)]$ by the three terms that can each be attained by simple mechanisms.

5.4.1. Two-Priced Allocation Construction

We now construct a two-priced allocation rule $\bar{x} = \bar{x}_{\text{val}} + \bar{x}_C$ from $x = x_{\text{val}} + x_C$ for which (a) revenue is improved, i.e., $\bar{p}(v) \geq p(v)$, and (b) the probability the agent pays her value minus capacity, $\bar{x}_C(v)$, is convex for $v \in [0, C]$. In fact, given x_{val}, $\bar{x}_C$ is the smallest function for which IC constraint (5.2) holds; and in the special case when x_{val} is monotone, the left-hand side of (5.4) is tight for $\bar{x}_C$ on $[0, C]$. Other incentive constraints may be violated by $\bar{x}$, but we use it only as an upper bound for revenue.

Definition 5.2 ($\bar{x}$)**.** We define $\bar{x} = \bar{x}_C + \bar{x}_{\text{val}}$ as follows:

(1) $\bar{x}_{\text{val}}(v) = x_{\text{val}}(v)$;

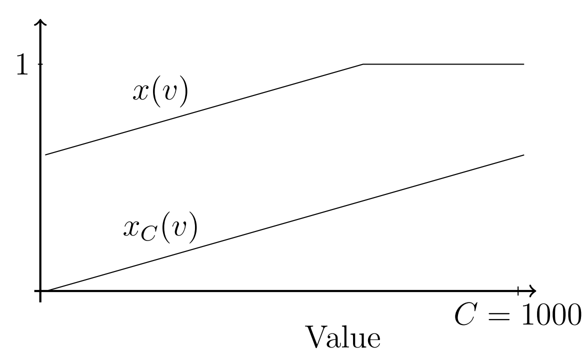

Figure 5.3. With $C = 1000$ and values drawn from the equal revenue distribution on $[1, 1000]$, this two-priced mechanism is BIC and achieves 1.55 times the revenue of the optimal risk-neutral mechanism.

(2) Let $r(v)$ be $\frac{1}{C}\sup_{z\le v} x_{\text{val}}(z)$, and let

$$\bar{x}_C(v) = \begin{cases} \int_0^v r(y)\,\mathrm{d}y, & v \in [0, C]; \\ x_C(v), & v > C. \end{cases} \tag{5.5}$$

Lemma 5.9 (Properties of $\bar{x}$)**.**

(1) *On $v \in [0, C]$, $\bar{x}_C(\cdot)$ is a convex, monotone increasing function.*

(2) *On all v, $\bar{x}_C(v) \le x_C(v)$.*

(3) *The incentive constraint from the left-hand side of* (5.4) *holds for $\bar{x}_C$:*

$$\frac{1}{C}\int_v^{v^+} \bar{x}_{\text{val}}(z)\,\mathrm{d}z \le \bar{x}_C(v^+) - \bar{x}_C(v) \text{ for all } v < v^+.$$

(4) *On all v, $\bar{x}_C(v) \le x_C(v)$, $\bar{x}(v) \le x(v)$, and $\bar{p}(v) \ge p(v)$.*

The proof of part 2 is technical, and we give a sketch here. Recall that, for each v, the IC constraint (5.2) gives a linear constraint lower bounding $x_C(v^+)$ for every $v^+ > v$. If one decreases $x_C(v)$, the lower bound it imposes on $x_C(v^+)$ is simply "pulled down" and is less binding. The definition of $\bar{x}_C$ simply lands $\bar{x}_C(v)$ on the most binding lower bound, and therefore not only makes $\bar{x}_C(v)$ at most $x_C(v)$, but also lowers the linear constraint that v imposes on larger values. If the number of values is countable or if x_{val} is piecewise constant, the lemma is easy to see by induction. A full proof for the general case of part 2, along with the proofs of the other more direct parts of Lemma 5.9, is given in Section D.2.

5.4.2. Payment Upper Bound

Recall that $\bar{p}(v)$ is the expected payment corresponding with two-priced allocation rule $\bar{x}(v)$. We now give an upper bound on $\bar{p}(v)$.

Lemma 5.10. *The payment $\bar{p}(v)$ for v and two-priced rule $\bar{x}(v)$ satisfies*

$$\bar{p}(v) \leq v\bar{x}(v) - \int_0^v \bar{x}(z)\,\mathrm{d}z + \int_0^v \bar{x}_C(z)\,\mathrm{d}z. \tag{5.6}$$

Proof. View a two-priced mechanism $\bar{x} = \bar{x}_{\mathrm{val}} + \bar{x}_C$ as charging v with probability $\bar{x}(v)$ and giving a rebate of C with probability $\bar{x}_C(v)$. We bound this rebate as follows (which proves the lemma):

$$\begin{aligned} C \cdot \bar{x}_C(v) &\geq C \cdot \bar{x}_C(0) + \int_0^v \bar{x}_{\mathrm{val}}(z)\,\mathrm{d}z \\ &\geq \int_0^v \bar{x}(z)\,\mathrm{d}z - \int_0^v \bar{x}_C(z)\,\mathrm{d}z. \end{aligned}$$

The first inequality is from part 3 of Lemma 5.9. The second inequality is from the definition of $\bar{x}_C(0) = 0$ in (5.5) and $\bar{x}_{\mathrm{val}}(v) = \bar{x}(v) - \bar{x}_C(v)$. See Figure 5.4 for an illustration. □

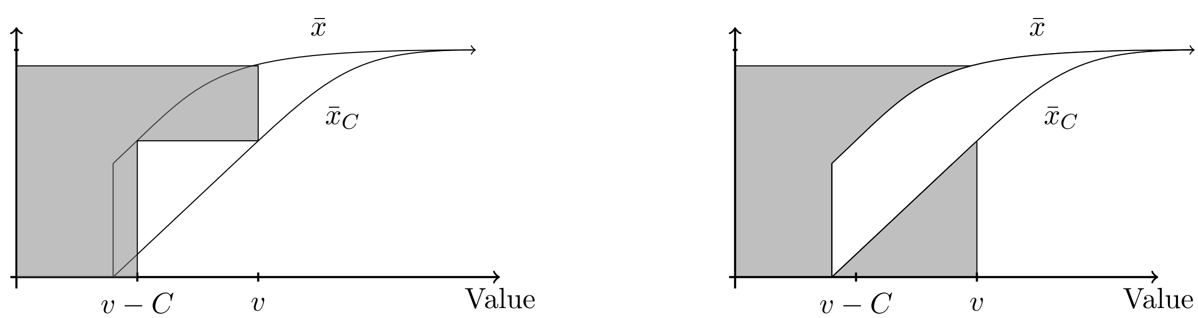

(a) The shaded region is the expected payment from an agent of value v.

(b) The shaded region upper bounds expected payment from an agent with value v, shown in Lemma 5.10.

Figure 5.4. Depictions and bounds on the payments in two-priced auctions.

5.4.3. Three-part Payment Decomposition

Below, we bound $\bar{p}(\cdot)$ (and hence $p(\cdot)$) in terms of the expected payment of three natural mechanisms. As seen geometrically in Figure 5.5, the bound given in Lemma 5.10 can be broken into two parts: the area above $\bar{x}(\cdot)$, and the area below $\bar{x}_C(\cdot)$. We refer to the former as $\bar{p}^{\mathrm{I}}(\cdot)$; we further split the latter quantity into two parts: $\bar{p}^{\mathrm{II}}(\cdot)$, the area corresponding to $v \in [0, C]$, and $\bar{p}^{\mathrm{III}}(\cdot)$, that corresponding to $v \in [C, v]$. We define these quantities formally below:

$$\bar{p}^{\mathrm{I}}(v) = \bar{x}(v)v - \int_0^v \bar{x}(z)\,\mathrm{d}z, \tag{5.7}$$

$$\bar{p}^{\mathrm{II}}(v) = \int_0^{\min\{v,C\}} \bar{x}_C(z)\,\mathrm{d}z \tag{5.8}$$

$$\bar{p}^{\mathrm{III}}(v) = \begin{cases} 0, & v \leq C; \\ \int_C^v \bar{x}_C(z)\,\mathrm{d}z, & v > C. \end{cases} \tag{5.9}$$

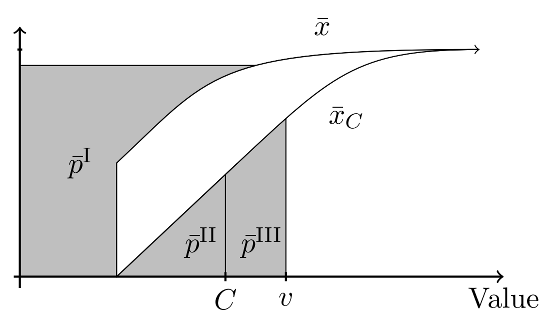

Figure 5.5. Breakdown of the expected payment upper bound in a two-priced auction.

Proof of Theorem 5.7. We now bound the revenue from each of the three parts of the payment decomposition. These bounds, combined with part 4 of Lemma 5.9 and Lemma 5.10, immediately give Theorem 5.7.

Part 1.: $\mathbf{E}_v[\bar{p}^{\mathrm{I}}(v)] = \mathbf{E}_v[\varphi(v) \cdot \bar{x}(v)] \leq \mathbf{E}_v[\varphi^+(v) \cdot x(v)]$.

Formulaically, $\bar{p}^{\mathrm{I}}(\cdot)$ corresponds to the risk-neutral payment identity for $\bar{x}(\cdot)$ as specified by part 2 of Theorem 5.1; by part 3 of Theorem 5.1, in expectation over v, this payment is equal to the expected virtual surplus $\mathbf{E}_v[\varphi(v) \cdot \bar{x}(v)]$.[2] The inequality follows as terms $\varphi(v)$ and $\bar{x}(v)$ in this expectation are point-wise upper bounded by $\varphi^+(v) = \max(\varphi(v), 0)$ and $x(v)$, respectively, the latter by part 4 of Lemma 5.9.

Part 2.: $\mathbf{E}_v[\bar{p}^{\mathrm{II}}(v)] \leq \mathbf{E}_v[\varphi(v) \cdot \bar{x}_C(v)] \leq \mathbf{E}_v[\varphi^+(v) \cdot x_C(v)]$.

By definition of $\bar{p}^{\mathrm{II}}(\cdot)$ in (5.8), if the statement holds for $v = C$ it holds for $v > C$; so we argue it only for $v \in [0, C]$. Formulaically, with respect to a risk-neutral agent with allocation rule $\bar{x}_C(\cdot)$, the risk-neutral payment is $v \cdot \bar{x}_C(v) - \int_0^v \bar{x}_C(z)\,\mathrm{d}z$, the surplus is $v \cdot \bar{x}_C(v)$, and the risk-neutral agent's utility (the difference between the surplus and payment) is $\int_0^v \bar{x}_C(z)\,\mathrm{d}z = \bar{p}^{\mathrm{II}}(v)$. Convexity of $\bar{x}_C(\cdot)$, from part 1 of Lemma 5.9, implies that the risk-neutral payment is at least half the surplus, and so is at least the risk-neutral utility. The lemma follows, then, by the same argument as in the previous part.

Part 3.: $\mathbf{E}_v[\bar{p}^{\mathrm{III}}(v)] \leq \mathbf{E}_v[(v - C)^+ \cdot \bar{x}_C(v)] = \mathbf{E}_v[(v - C)^+ \cdot x_C(v)]$.

The statement is trivial for $v \leq C$ so assume $v \geq C$. By definition $\bar{x}_C(v) = x_C(v)$ for $v > C$. By (5.2), $x_C(\cdot)$ is monotone non-decreasing. Hence, for $v > C$, $\bar{p}^{\mathrm{III}}(v) =$

[2] Note: This equality does not require monotonicity of the allocation rule $\bar{x}(\cdot)$; as long as part 2 of Theorem 5.1 formulaically holds, part 3 follows from integration by parts.

$\int_C^v x_C(z)\,\mathrm{d}z \leq \int_C^v x_C(v)\,\mathrm{d}z = (v - C) \cdot x_C(v)$. Plugging in $(v - C)^+ = \max(v - C, 0)$ and taking expectation over v, we obtain the bound.

□

5.5. Approximation Mechanisms and a Payment Identity

In this section we first give a payment identity for Bayes-Nash equilibria in mechanisms that charge agents a deterministic amount upon winning (and zero upon losing). Such one-priced payment schemes are not optimal for capacitated agents; however, we will show that they are approximately optimal. In both symmetric and asymmetric cases, we use this payment identity to prove that the first-price auction is approximately optimal. We also give a simple direct-revelation one-priced mechanism that is BIC and approximately optimal and improves on the approximation guarantees of the first-price auction.

5.5.1. A One-price Payment Identity

For risk-neutral agents, the Bayes-Nash equilibrium conditions entail a payment identity: given an interim allocation rule, the payment rule is fixed (Theorem 5.1, part 2). For risk-averse agents there is no such payment identity: there are mechanisms with identical BNE allocation rules but distinct BNE payment rules. We restrict attention to auctions wherein an agent's payment is a deterministic function of her value (if she wins) and zero if she loses. We call these *one-priced* mechanisms; for these mechanisms there is a (partial) payment identity.

We consider a single agent and the induced allocation rule she faces from a Bayesian incentive compatible auction (or, by the revelation principle, any BNE of any mechanism). This allocation rule internalizes randomization in the environment and the auction, and specifies the agents' probability of winning, $\mathbf{x}(v)$, as a function of her value. Given allocation rule $\mathbf{x}(v)$, the risk-neutral expected payment is $p^{\mathrm{RN}}(v) = v \cdot \mathbf{x}(v) - \int_0^v \mathbf{x}(z)\,\mathrm{d}z$ (Theorem 5.1, part 2). Given an allocation rule $\mathbf{x}(v)$, a one-priced mechanism with payment rule $p(v)$ would charge the agent $p(v)/\mathbf{x}(v)$ upon winning and zero otherwise (for an expected payment of $p(v)$). Define $p^{\mathrm{VC}}(v) = (v - C) \cdot \mathbf{x}(v)$ which, intuitively, gives a lower bound on a capacitated agent's willingness to trade-off decreased probability of winning for a cheaper price.

Theorem 5.11. *An allocation rule* $\mathbf{x}$ *and payment rule* p *are the BNE of a one-priced mechanism if and only if (a)* $\mathbf{x}$ *is monotone non-decreasing and (b) if* $p(v) \geq p^{\mathrm{VC}}(v)$ *for all* v *then* $p = p^C$ *is defined as*

$$p^C(0) = 0, \tag{5.10}$$

$$p^C(v) = \max\left(p^{\mathrm{VC}}(v),\ \sup_{v^- < v} \left\{p^C(v^-) + (p^{\mathrm{RN}}(v) - p^{\mathrm{RN}}(v^-))\right\}\right). \tag{5.11}$$

Moreover, if $\mathbf{x}$ *is strictly increasing then* $p(v) \geq p^{\mathrm{VC}}(v)$ *for all* v *and* $p = p^C$ *is the unique equilibrium payment rule.*

The payment rule should be thought of in terms of two "regimes": when $p^C = p^{\mathrm{VC}}$, and when $p^C > p^{\mathrm{VC}}$, corresponding to the first and second terms in the max argument of (5.11) respectively. In the latter regime, (5.11) necessitates that $\frac{\mathrm{d}}{\mathrm{d}v}p^C(v) = \frac{\mathrm{d}}{\mathrm{d}v}p^{\mathrm{RN}}(v)$; for nearby such points v and $v + \epsilon$, the v^- involved in the supremum will be the same, and thus $p^C(v+\epsilon) - p^C(v) = p^{\mathrm{RN}}(v+\epsilon) - p^{\mathrm{RN}}(v)$.

The proof is included in Section D.3. The main intuition for this characterization is that risk-neutral payments are "memoryless" in the following sense. Suppose we fix $p^{\mathrm{RN}}(v)$ for a v and ignore the incentive of an agent with value $v^+ > v$ to prefer reporting $v^- < v$, then the risk-neutral payment for all $v^+ > v$ is $p^{\text{RN}}(v^+) = p(v) + \int_v^{v^+} (\mathbf{x}(v^+) - \mathbf{x}(z))\,\mathrm{d}z$. This memorylessness is simply the manifestation of the fact that the risk-neutral payment identity imposes local constraints on the derivatives of the payment, i.e., $\frac{\mathrm{d}}{\mathrm{d}v} p^{\text{RN}}(v) = v \cdot \frac{\mathrm{d}}{\mathrm{d}v}\mathbf{x}(v)$.

There is a simple algorithm for constructing the risk-averse payment rule p^C from the risk-neutral payment rule p^{RN} (for the same allocation rule $\mathbf{x}$).

(0) For $v < C$, $p^C(v) = p^{\text{RN}}(v)$.

(1) The $p^C(v) = p^{\text{RN}}(v)$ identity continues until the value v' where $p^C(v') = p^{\text{VC}}(v')$, and $p^C(v)$ switches to follow $p^{\text{VC}}(v)$.

(2) When v increases to the value v'' where $\frac{\mathrm{d}}{\mathrm{d}v} p^{\text{RN}}(v'') = \frac{\mathrm{d}p^{\text{VC}}}{\mathrm{d}v}(v'')$ then $p^C(v)$ switches to follow $p^{\text{RN}}(v)$ shifted up by the difference $p^{\text{VC}}(v'') - p^{\text{RN}}(v'')$ (i.e., its derivative $\frac{\mathrm{d}}{\mathrm{d}v} p^C(v)$ follows $\frac{\mathrm{d}}{\mathrm{d}v} p^{\text{RN}}(v)$).

(3) Repeat this process from Step 1.

Lemma 5.12. *The one-priced BIC allocation rule $\mathbf{x}$ and payment rule p^C satisfy the following*

(1) *For all v, $p^C(v) \geq \max(p^{\text{RN}}(v), p^{\text{VC}}(v))$.*

(2) *Both $p^C(v)$ and $p^C(v)/\mathbf{x}(v)$ are monotone non-decreasing.*

The proof of part 2 is contained in the proof of Lemma D.3 in Section D.3, and part 1 follows directly from equations (5.10) and (5.11).

5.5.2. Approximate Optimality of First-price Auction in Symmetric Settings

We show herein that for agents with a common capacity and values drawn i.i.d. from a continuous, regular distribution F with strictly positive density the first-price auction is approximately optimal.

It is easy to solve for a symmetric equilibrium in the first-price auction with identical agents. First, guess that in BNE the agent with the highest value wins. When the agents are i.i.d. draws from distribution F, the implied allocation rule is $\mathbf{x}(v) = F^{n-1}(v)$. Theorem 5.11 then gives the necessary equilibrium payment rule $p^C(v)$ from which the bid function $b^C(v) = p^C(v)/\mathbf{x}(v)$ can be calculated. We verify that the initial guess is correct as Lemma 5.12 implies that the bid function is symmetric and monotone. There is no other symmetric equilibrium.[3] Moreover, Chawla and Hartline [2013] show that there are no asymmetric equilibria of the first-price auction for a class of environments including ours.

Proposition 5.13. *The first-price auction for identical (capacity and value distribution) agents has a unique BNE wherein the highest valued agent wins.*

The expected revenue at this equilibrium is $n\,\mathbf{E}_v[p^C(v)]$. Lemma 5.12 implies that p^C is at least p^{RN} and p^{VC}.

Corollary 5.14. *The expected revenue of the first-price auction for identical (capacity and value distribution) agents is at least that of the capacitated second-price auction and at least that of the second-price auction.*

In conjunction with Theorem 2.9 by Bulow and Klemperer [1996], this now allows us to use the capacitated first-price auction to approximate each of the three terms in the revenue bound of Theorem 5.7, giving us our main result:

[3]Any other symmetric equilibrum must have an allocation rule that is increasing but not always strictly so. For this to occur the bid function must not be strictly increasing implying a point mass in the distribution of bids. Of course, a point mass in a symmetric equilibrium bid function implies that a tie is not a measure zero event. Any agent has a best response to such an equilibrium of bidding just higher than this pointmass so at essentially the same payment, she always "wins" the tie.

Theorem 5.15. *For symmetric, regular environments with $n \geq 2$ agents having common capacities, the revenue in the BNE of the first price auction (FPA) is a 5-approximation to the optimal revenue.*

Proof. An immediate consequence of Theorem 2.9 is that for $n \geq 2$ risk-neutral, regular, i.i.d. bidders, the second-price auction extracts a revenue that is at least half the optimal revenue; hence, by Corollary 5.8, the optimal revenue for capacitated bidders by any BIC mechanism is at most four times the second-price revenue plus the capacitated second-price revenue. Since the first-price auction revenue in BNE for capacitated agents is at least the capacitated second-price revenue and the second-price revenue, the first-price revenue is a 5-approximation to the optimal revenue.[4] $\square$

5.5.3. Approximate Optimality of First-Price Auctions in Asymmetric Settings

We now consider the case of asymmetric value distributions and capacities. For risk-neutral settings, Chapter 3 showed an approach for analyzing the revenue of the first-price auction in asymmetric settings, achieving approximation bounds of $3e/(e-1)$ for the first-price auction with at least two bidders from each distribution.

We can also extend this result to asymmetric settings by using Theorem 3.12 from Chapter 3 to compare the revenue of the risk-neutral first-price auction to the risk-neutral optimal auction.

Theorem 5.16. *For asymmetric, regular environments, where $k \geq 2$ agents have values drawn from each distribution F_i and a common capacity C_i, the revenue in any BNE of the First-Price Auction is a $\frac{7e-1}{e-1} \approx 10.5$ approximation to the revenue of the optimal auction.*

[4]In fact, the Bulow and Klemperer [1996] result shows that the second-price auction is asymptotically optimal; using this, our result can be tightened to a $(3 + \frac{2}{n-1})$-approximation.

If we knew that the revenue of the capacitated first-price auction is greater than the revenue of the risk-neutral first-price auction, then we could immediately attain the result by using Theorem 3.12 in place of Theorem 2.9 in the proof of Theorem 5.15. However, the allocation rule of the first-price auction is not efficient in asymmetric settings, and will potentially be different between the normal and the capacitated settings, rather than the first-price auction.

However, nowhere in the proof of Theorem 3.12 depended on the exact allocation rule between agents of different types. Rather, it depended on how agents react to distributions of bids, as well as risk-neutral virtual values serving as an amortization of first-price revenue. Capacitated agents will bid higher than risk-neutral agents, and risk-neutral virtual values will serve as a lower bound on first-price revenue, and so Theorem 3.12 is easily extended to hold for capacitated agents, to give $\textsc{Rev}(\text{FPA}) \geq \frac{3e}{e-1}\textsc{Rev}(\textsc{Opt})$. Integrating into the proof of Theorem 5.15 in place of the bound of 2 from Theorem 2.9 gives our desired result, a bound of $\frac{7e-1}{e-1} \approx 10.5$ as compared to the revenue optimal auction.

Much of the loss in the approximation factor of Theorem 5.16 comes from not knowing the allocation rule. The best one-priced auction can do significantly better: we show here it is at least a three approximation to the optimal.

Theorem 5.17. *For n (non-identical) agents with capacities $C_1, \ldots, C_n$, and regular value distributions $F_1, \ldots, F_n$, there is a one-priced BIC mechanism whose revenue is at least one third of the optimal (two-priced) revenue.*

Proof. Recall from Theorem 5.7 that either the risk-neutral optimal revenue or $\mathbf{E}_{v_1,\ldots,v_n}[\max\{(v_i - C_i)_+\}]$ is at least one third of the optimal revenue. We apply Theorem 5.11 to two monotone allocation rules for each agent:

(1) the interim allocation rule of the risk-neutral optimal auction, and

(2) the interim allocation rule specified by: serve agent i that maximizes $v_i - C_i$, if positive; otherwise, serve nobody.

As both allocations are monotone, we apply Theorem 5.11 to obtain two single-priced BIC mechanisms. By Lemma 5.12, the expected revenue of the first mechanism is at least the risk neutral optimal revenue, and the expected revenue of the second mechanism is at least $\mathbf{E}_{v_1,\ldots,v_n}[\max\{(v_i - C_i)_+\}]$. The theorem immediately follows for the auction with the higher expected revenue. $\square$

Although Theorem 5.17 is stated as an existential result, the two one-priced mechanisms in the proof can be described analytically using the algorithm following Theorem 5.11 for calculating the capacitated BIC payment rule. The interim allocation rules are straightforward (the first: $\mathbf{x}_i(v_i) = \prod_{j\neq i} F_j(\varphi_j^{-1}(\varphi_i(v_i)))$, and the second: $\mathbf{x}_i(v_i) = \prod_{j\neq i} F_j(v_i - C_i + C_j))$, and from these we can solve for $p_i^{C_i}(v)$.

5.6. First Price Auctions with CARA Agents

In this section, we generalize the Bulow-Klemperer result that competition can approximately replace a reserve price (Theorem 2.9) to the case of symmetric, risk-averse bidders with constant absolute risk-aversion (CARA) utility functions. A CARA utility function has the form

$$U(z) = \frac{1 - e^{-Rz}}{R},$$

for some parameter of risk-aversion R. As $R \to 0$, $U(z) \to z$ and agents behave as if risk-neutral; as $R \to \infty$, the agent approaches being so risk-averse that she is bidding to maximize her chance of winning, subject to paying less than her value.

5.6.1. Bidding Behavior

Hu et al. [2010] show that the only equilibrium of a first-price auction with symmetric CARA bidders will be symmetric, and each bidders bid function will be:

$$b_i(v_i) = \frac{1}{R}\log\left(e^{Rv_i} - R\int_0^{v_i} \frac{F(z)}{F(v_i)} e^{Rz}\, dz\right). \tag{5.12}$$

We can generalize this beyond the first-price auction to first-price style auctions. Consider a general allocation rule $x_i(v_i)$. The payment rule $\mathbf{p}_i$ combined with x_i gives a BNE if and only if

$$b_i(v_i) = \frac{1}{R}\log\left(e^{Rv_i} - R\int_0^{v_i} \frac{x_i(z)}{x_i(v_i)} e^{Rz}\, dz\right) \tag{5.13}$$

$$= v_i + \frac{1}{R}\log\left[1 - R\int_0^{v_i} \frac{x_i(z)}{x_i(v_i)} e^{-R(v_i-z)}\,\mathrm{d}z\right] \tag{5.14}$$

5.6.2. Characterizing Revenue in Equilibrium

We begin by presenting a characterization of the revenue, in the line of Myerson's [1981] virtual value based characterization of revenue in equilibrium for risk-neutral agents.

Lemma 5.18. *For symmetric, single-item CARA environments, the revenue of a first-price-style auction A with CARA bidders satisfies*

$$\textsc{Rev}(A) = \sum_i \int_0^1 \phi_i^x(v_i)\, x_i(v_i)\, f(v_i)\,\mathrm{d}v_i, \tag{5.15}$$

where $\phi_i^x(v_i)$ is the risk-averse virtual value,

$$\phi_i^x(v_i) = v_i - \frac{1}{f(v_i)}\int_{v_i}^1 \frac{f(z)}{e^{-R(v_i-z)} + R\int_0^{v_i} \frac{x_i(y)}{x_i(z)} e^{-R(v_i-y)}\,\mathrm{d}y}\,\mathrm{d}z. \tag{5.16}$$

The risk-averse virtual value represents an amortization of revenue across the bidders in the auction, just as the risk-neutral virtual value $\phi(v) = v - \frac{1-F(v)}{f(v)}$ does for risk-neutral environments. Note that the risk-averse virtual value *does* depend on the allocation rule, as opposed to the risk-neutral virtual value. The risk-averse virtual value thus can only be used to account for the revenue, it cannot as directly be used for determining the optimal auction. Changing the allocation rule to serve more higher virtual valued agents will not necessarily result in more revenue, as it may change the virtual values of all the bidders.

Proof of Lemma 5.18. By the CARA first-price payment rule, the bid of the agents satisfies

$$b_i(v_i) = v_i + \frac{1}{R}\log\left[1 - R\int_0^{v_i} \frac{x_i(z)}{x_i(v_i)} e^{-R(v_i-z)}\,\mathrm{d}z\right] \tag{5.17}$$

Rewriting log as $\int \frac{1}{z}\ dz$ and then substitution with $w(z) = R\int_0^z \frac{x_i(y)}{x_i(v_i)} e^{-R(v_i-y)}\ dy$, and $w'(z) = R\frac{x_i(z)}{x_i(v_i)} e^{-R(v_i-y)}\ dy$ gives:

$$\begin{aligned}
b_i(v_i) &= v_i + \frac{1}{R}\log\left[1 - R\int_0^{v_i} \frac{x_i(z)}{x_i(v_i)} e^{-R(v_i-z)}\,\mathrm{d}z\right] \\
&= v_i + \frac{1}{R}\int_1^{1-R\int_0^{v_i}\frac{x_i(y)}{x_i(v_i)}e^{-R(v_i-y)}\ dy} \frac{1}{z}\ dx \\
&= v_i - \frac{1}{R}\int_0^{R\int_0^{v_i}\frac{x_i(y)}{x_i(v_i)}e^{-R(v_i-y)}\ dy} \frac{1}{1+z}\ dz \\
&= v_i - \int_0^{v_i} \frac{\frac{x_i(z)}{x_i(v_i)}e^{-R(v_i-z)}}{1 + R\int_0^z \frac{x_i(y)}{x_i(v_i)}e^{-R(v_i-y)}dy}\ dz
\end{aligned} \tag{5.18}$$

We can then integrate over expected values to get the expected revenue from agent i:

$$\begin{aligned}
\textsc{Rev}_i &= \int_0^1 b_i(v_i)x_i(v_i)f_i(v_i)\ dv_i \\
&= \int_0^1 \left(v_i - \int_0^{v_i} \frac{\frac{x_i(z)}{x_i(v_i)} e^{-R(v_i - z)}}{1 + R\int_0^z \frac{x_i(y)}{x_i(v_i)} e^{-R(v_i-y)} dy}\ dz \right) x_i(v_i) f_i(v_i)\, dv_i. \\
&= \int_0^1 v_i x_i(v_i) f_i(v_i) dv_i - \int_0^1 \int_0^{v_i} \frac{\frac{x_i(z)}{x_i(v_i)} e^{-R(v_i - z)}}{1 + R\int_0^z \frac{x_i(y)}{x_i(v_i)} e^{-R(v_i-y)} dy}\, dz\ x_i(v_i)\, f_i(v_i)\, dv_i. \\
&= \int_0^1 v_i x_i(v_i) f_i(v_i) dv_i - \int_0^1 \left(\int_z^1 \frac{x_i(v_i) f_i(v_i) e^{-R(v_i - z)}}{f_i(z)(x_i(v_i) + R\int_0^z x_i(y) e^{-R(v_i-y)} dy)}\, dv_i \right) x_i(z)\, f_i(z)\, dz.
\end{aligned}$$

The last step followed by swapping the order of integration. Swapping z and v_i in the second integral and recombining the integrals gives

$$\begin{aligned}
\textsc{Rev}_i &= \int_0^1 v_i x_i(v_i) f_i(v_i) dv_i - \int_0^1 \left(\int_{v_i}^1 \frac{x_i(v_i) f_i(z) e^{-R(z - v_i)}}{f_i(v_i)(x_i(z) + R\int_0^{v_i} x_i(y) e^{-R(z-y)} dy)}\ dz \right) x_i(v_i)\, f_i(v_i)\, dv_i \\
&= \int_0^1 \left(v_i - \int_{v_i}^1 \frac{x_i(z) f_i(z) e^{-R(z - v_i)}}{f_i(v_i)(x_i(z) + R\int_0^{v_i} x_i(y) e^{-R(z-y)} dy)}\ dz \right) x_i(v_i)\, f_i(v_i)\, dv_i \qquad (5.19) \\
&= \int_0^1 \phi_i^x(v_i) x_i(v_i)\, f_i(v_i)\, dv_i. \qquad (5.20)
\end{aligned}$$

Summing over all agents gives our desired result. □

5.6.3. Extending Bulow Klemperer Results

We first define a regularity condition analogous to risk-neutral regularity of a distribution: the virtual-value must be non-decreasing.

Definition 5.3. A distribution is CARA-regular for risk parameter R and allocation rule x if $\phi_i^x(v)$ is monotone non-decreasing in v.

Theorem 5.19. *For symmetric, single-item, CARA-regular, CARA environments with n agents, the revenue of the first-price auction is at least an $(1 - 1/n)$ fraction of the first-price auction with an optimal reserve. Moreover, as agents become more risk averse, the bound approaches 1. Specifically,*

$$\textsc{Rev}(\text{FPA}) \geq \left(1 - \frac{1}{n}\right) \textsc{Rev}(\text{FPAR}). \tag{5.21}$$

The proof now proceeds in two steps: first, we bound the revenue that the first-price auction loses to serving negative virtual values using the same techniques as in the risk-neutral case; second, we compare the virtual values in the first-price auction to the virtual values in the first-price auction with reserve.

Proof. To align with the risk-neutral proof, we will refer to agents by their quantile, $q = 1 - F_i(v)$, and use the inverse value function $v(q) = F_i^{-1}(1 - q)$. Let $\phi_i^x(q) = \phi_i^x(v(q))$.

Let $R_i(q) = \int_0^1 \phi_i^x(q)\ dq$ be the "risk-averse revenue curve" for a given quantile. Note that this revenue curve is *not* generated by considering the revenue from selling to an agent with probability q: it is rather only playing the role of the risk-neutral revenue curve in the calculation of revenue, as a change of variables and integration by parts gives

$$\textsc{Rev}_i = \int_0^1 \phi_i^x(v_i) x_i(v) f(v)\, \mathrm{d}v_i \tag{5.22}$$

$$= \int_0^1 \phi_i^x(q) x_i(q)\, \mathrm{d}q \tag{5.23}$$

$$= \int_0^1 -x_i'(q) R_i(q)\, \mathrm{d}q. \tag{5.24}$$

As $R_i(q)$ is concave in q and the allocation rule is identical to the risk-neutral setting, we can reduce directly to the risk-neutral setting. Theorem 2.9 gives

$$\begin{aligned} \textsc{Rev}_i &= \int_0^1 -x_i'(q) R_i(q)\,\mathrm{d}q. \\ &\geq \frac{n-1}{n} \int_0^1 -x_i'(q) R^+(q)\,\mathrm{d}q \\ &\geq \frac{n-1}{n} \textsc{Rev}_i^+. \end{aligned} \tag{5.25}$$

We now compare $\textsc{Rev}_i^+$ to the revenue from the first-price auction with reserve.

Let $\mathbf{x}^*$ be the allocation rule for the first-price auction with optimal reserves; thus the virtual values are:

$$\phi_i^{x^*}(v_i) = v_i - \frac{1}{f_i(v_i)} \int_{v_i}^1 \frac{f_i(z)}{e^{-R(v_i-z)} + R\int_0^{v_i} \frac{x_i^*(y)}{x_i^*(z)} e^{-R(v_i-y)}\,\mathrm{d}y}\,\mathrm{d}z. \tag{5.26}$$

As $\mathbf{x}^*$ is the allocation rule for the first-price auction with a reserve, we know that for any $y < z$, $\frac{x_i^*(y)}{x_i^*(z)} \leq \frac{x_i(y)}{x_i(z)}$: either y is above the reserve and $\frac{x_i^*(y)}{x_i^*(z)} = \frac{x_i(y)}{x_i(z)}$ or below the reserve and $\frac{x_i^*(y)}{x_i^*(z)} = 0$.[5]

$$\phi_i^{x_i^*}(v_i) = v_i - \frac{1}{f_i(v_i)} \int_{v_i}^1 \frac{f_i(z)}{e^{-R(v_i-z)} + R\int_0^{v_i} \frac{x_i^*(y)}{x_i^*(z)} e^{-R(v_i-y)}\,\mathrm{d}y}\,\mathrm{d}z \tag{5.27}$$

$$\leq v_i - \frac{1}{f_i(z)} \int_{v_i}^1 \frac{f_i(z)}{e^{-R(v_i-z)} + R\int_0^{v_i} \frac{x_i(y)}{x_i(z)} e^{-R(v_i-y)}\,\mathrm{d}y}\,\mathrm{d}z \tag{5.28}$$

$$= \phi_i^x(v_i) \tag{5.29}$$

[5]We need only consider the case that $x_i^*(z) > 0$.

We can then upper bound the optimal revenue with the revenue from agents with positive virtual values in the first price auction:

$$\textsc{Rev}^+(A) = \sum_i \int_0^1 \max(0, \phi_i^x(v_i))\, x_i(v_i)\, f_i(v_i)\, \mathrm{d}v_i \tag{5.30}$$

$$\geq \sum_i \int_0^1 \max(0, \phi_i^{x^*}(v_i))\, x_i(v_i)\, f_i(v_i)\, \mathrm{d}v_i \tag{5.31}$$

$$\geq \sum_i \int_0^1 \phi_i^{x^*}(v_i)\, x_i(v_i)\, f_i(v_i)\, \mathrm{d}v_i \tag{5.32}$$

$$= \textsc{Rev}(\text{FPAR}) \tag{5.33}$$

Combining Equations (5.25) and (5.33) and summing over agents gives our desired result, $\textsc{Rev}(A) \geq \frac{n-1}{n}\textsc{Rev}(\text{FPAR})$. □

5.7. Conclusions

For the purpose of keeping the exposition simple, we have applied our analysis only to single-item auctions. Our techniques, however, as they focus on analyzing and bounding revenue of a single agent for a given allocation rule, generalize easily to structurally rich environments. Notice that the main theorems of Sections 5.3, 5.4, and the first part of Section 5.5 do not rely on any assumptions on the feasibility constraint except for downward closure, i.e., that it is feasible to change an allocation by withholding service to an agent who was formerly being served.

For example, our prior-independent 5-approximation result (Theorem 5.15) generalizes to symmetric feasibility constraints such as position auctions. A position auction environment is given by a decreasing sequence of weights $\alpha_1, \ldots, \alpha_n$ and the first-price position auction assigns the agents to these positions greedily by bid. With probability α_i the agent in position i receives an item and is charged her bid; otherwise she is not charged. (These position auctions have been used to model pay-per-click auctions for selling advertisements on search engines where α_i is the probability that an advertiser gets clicked on when her ad is shown in the ith position on the search results page.) For agents with identical capacities and value distributions, the first-price position auction where the bottom half of the agents are always rejected is a 5-approximation to the revenue-optimal position auction (that may potentially match all the agents to slots).

Our one- versus two-price result (Theorem 5.17) generalizes to asymmetric capacities, asymmetric distributions, and asymmetric downward-closed feasibility constraints. A downward-closed feasibility constraint is given by a set system which specifies which subset of agents can be simultaneously served. Downward-closure requires that any subset of a feasible set is feasible. A simple one-priced mechanism is a 3-approximation to the optimal mechanism in such an environment. The mechanism is whichever has higher revenue of the standard (risk neutral) revenue-optimal mechanism (which serves the subset of agents with the highest virtual surplus, i.e., sum of virtual values) and the one-priced revelation mechanism that serves the set of agents S that maximizes $\sum_{i\in S}(v_i - C_i)^+$ subject to feasibility.

A main direction for future work is to relax some of the assumptions of our model. Our approach to optimizing over mechanisms for risk-averse agents relies on (a) the simple model of risk aversion given by capacitated utilities and (b) that losers neither make (i.e., ex post individual rationality) nor receive payments (i.e., no bribes). These restrictions are fundamental for obtaining linear incentive compatibility constraints. Of great interest in future study is relaxation of these assumptions.

Our analytical (and computational) solution to the optimal auction problem for agents with capacitated utilities requires an *ex post individual rationality* constraint on the mechanism that is standard in algorithmic mechanism design. This constraint requires that an agent who loses the auction cannot be charged. While such a constraint is natural in many settings, it is with loss and, in fact, ill motivated for settings with risk-averse agents. One of the most standard mechanisms for agents with risk-averse preferences is the "insurance mechanism" where an agent who may face some large liability with small probability will prefer to pay a constant insurance fee so that the insurance agency will cover the large liability in the event that it is incurred. This mechanism is not ex post individually rational. Does the first-price auction (which is ex post individual rational) approximate the optimal interim individually rational mechanism?

For CARA agents, our results provide bounds relative to the first-price auction with optimal reserve prices. We do not yet have bounds comparing either of these quantities to the optimal mechanism from Matthews [1983]. We would also like to understand a simpler set of conditions for regularity.

CHAPTER 6

Simplifying Strategic Behavior: The Utility-Target Auction

In this chapter, we explore the nature of strategic behavior from a different perspective than the earlier chapters: we focus on a revelation mechanism, the *utility-target auction*. Rather than bidding for how much to pay, in the utility target auction bidders report their valuations and then bid for how much utility they want to receive from the mechanism. This simplifies the strategic behavior of agents in complicated settings while inheriting good performance guarantees from profit-target bidding strategies in first-price auctions.

This chapter also differs in that it focuses on full-information settings, and dynamics in full information settings where bidders learn of the behavior of others only through their own participation in the mechanism.

6.1. Introduction

In 1961, Vickrey initiated the formal study of auctions. He first considered common auctions of the day — including the first-price auction, the Dutch auction, and the English auction — and studied their equilibria. Vickrey observed that the English auction was, in theory, more robust because each agent had a strategy that dominated all others regardless of other agents' bids. As a solution, he proposed[1] the second-price auction as a means to achieve the same robustness in a sealed-bid format. The subsequent development of auction theory largely followed Vickrey's paradigm: existing auctions were evaluated in terms of their equilibria, meanwhile the field of mechanism design emerged with dominant strategy incentive compatibility as a *sine qua non*.

Fifty years later, it is apparent that Vickrey's analysis does not always give best guide to implementing a real auction. In mechanisms without dominant strategies, Vickrey's original concern still stands — equilibrium is a highly questionable predictor of outcome due (at least in part) to agents' informational limitations [Vickrey, 1961; Harrison, 1989]. Neither is dominant strategy incentive compatibility a magic solution: incentive compatible mechanisms have sufficiently many drawbacks that their real attractiveness rarely matches theory — the simple and elegant second-price auction has earned the title "Lovely but Lonely" [Ausubel and Milgrom, 2006] for its sparse use. Even the supposition that bidders will play strategies that are theoretically dominant is discredited by a wide variety of practical issues [Klemperer, 2002].

Dynamic analysis offers a powerful complement to Vickrey's static approach. For example, certain behavior will be clearly irrational when an auction is repeated. Such reasoning

[1]While Vickrey was the first to discover the second-price auction in the economics literature, it has been used in practice as early as 1893 [Lucking-Reiley, 2000].

was used by Edelman and Schwarz [2010] in the context of the generalized second-price (GSP) ad auction — they analyzed a dynamic game to derive bounds on reasonable outcomes of the auction, then studied the static game under the assumption that these bounds were satisfied. Dynamic settings also introduce new pitfalls: Edelman and Ostrovsky [2007] showed that the instability of Overture's generalized first-price (GFP) ad auction could be attributed to its lack of a pure-strategy equilibrium.

We study repeated first-price auctions and show that they offer powerful performance guarantees. We begin with a static perspective and observe that the equilibrium properties of the auction depend significantly on the types of bids that bidders can express. We propose a generalization of the first-price auction called the *utility-target auction* which effectively encapsulates truncation or profit-target bidding strategies which have been shown to perform well in first-price and ascending proxy package auctions [Milgrom, 2004; Day and Milgrom, 2007].

We show that the utility-target auction effectively inherits the static performance guarantees of profit-target bidding strategies in general settings, with many advantages over incentive compatible mechanisms in a static equilibrium analysis, including revenue, simplicity, and transparency.

More significantly, we show that *the same performance guarantees may be derived using only a few simple behavioral axioms and limited information* in a repeated setting. These dynamic results are particularly powerful because they do not require an a priori assumption that the auction will reach equilibrium — for example, assuming only that losers will not wait too long to raise their bids *the auctioneer's revenue satisfies a natural lower bound regardless of whether bidders' behavior converges to equilibrium.* Moreover, bidders need only know if they are winning or losing to implement the dynamics. We build on these axioms to demonstrate behavior that offers progressively stronger performance guarantees, culminating with a set of axioms that together imply convergence to the egalitarian equilibrium.

First-Price Auctions. The virtues of a first-price auction — and other auctions in the pay-your-bid family — arise from its simplicity. From the bidders' perspective, the pay-your-bid property offers transparency, credibility, and privacy: not only is the auction easy to understand, but it ensures that the auctioneer cannot cheat (say, by unreasonably inflating the reserve price in a repeated auction) and allows a bidder to participate without expressing her true willingness to pay.

The auctioneer can also benefit from this simplicity because agents' bids represent guaranteed revenue. By comparison, the revenue from a dominant strategy incentive compatible auction is almost always less than the bids and, in the most general settings, may even be zero [Ausubel and Milgrom, 2006; Roberts, 1979]. Supposing a first-price auction reaches equilibrium, a variety of work presents settings where they generate more revenue for the seller than their incentive compatible brethren [Milgrom, 2004; Leme et al., 2012; Hoy et al., 2012] (though they may also generate less revenue [Maskin and Riley, 2000]).

Yet, running a first-price auction is risky. While first-price auctions have been quite successful in settings like procurement auctions, Overture's generalized first-price (GFP) auction for sponsored search advertising was erratic: bids rapidly rose and fell in a sawtooth pattern, rendering the auction unpredictable and depressing revenue [Edelman and Ostrovsky, 2007]. As a result, the sponsored search industry has moved to a generalized second-price (GSP) auction that leverages the intuition of the second-price auction to dis-incentivize small adjustments to a agent's bid.

The challenges of a first-price auction are many and complex. Vickrey identified a major source of risk in the first-price single-item auction: since a rational bidder's optimal bid depends on other agents' bids, actual behavior will depend on beliefs about others' strategies. A first-price auction also requires bidders to strategize, a task that is may be difficult and expensive. At best, agents will be in a Bayes-Nash equilibrium, and, at worst, they will be completely unpredictable. Indeed, predicting the outcome of a first-price auction lies at the center of a lively debate between experimental and theoretical economists [Harrison, 1989].

Experience with GFP highlights another potential pitfall of first-price auctions: when generalized beyond the single-item setting, a first-price auction may not have a pure-strategy equilibrium. Edelman and Ostrovsky [2007] showed that this was the case with GFP and demonstrated how it generated the rapid sawtooth behavior seen in practice. Our goal is to demonstrate that how proper design coupled with dynamic arguments can support strong performance guarantees.

The Utility-Target Auction. The equilibria and performance of a pay-your-bid auction depend on its implementation. Within the pay-your-bid constraint, the auctioneer chooses the form of agents' bids, potentially restricting or broadening the bids that agents may express.

The historical performance of the GFP ad auction exemplifies the importance of choosing a good bidding language. In the GFP auction, advertisers placed a single bid and paid the bid price for each click regardless of where their ads were shown. In retrospect, the rapid sawtooth motion observed in bids is not surprising because the auction had no pure-strategy equilibrium [Edelman et al., 2007; Edelman and Ostrovsky, 2007]; however, as we show and as Dütting et al. [2014] have shown, a pure-strategy equilibrium would have existed if bidders could have placed more expressive bids, such as bidding different prices for each slot.

A natural question arises: what are good bidding languages and how complicated must a language be to offer good performance? In GFP, the bidding language is precisely sufficient to represent any possible valuation function; hence, it is possible that bids may need to be more expressive than the space of valuation functions.

We show that the overhead required for a good bidding language is at most a single value: it is sufficient to ask bidders for their valuation function and their final desired utility. We call such an auction a *utility-target auction*: a agent's bid includes a specification of her value for every outcome and a single number representing the utility-target that she requests regardless of the outcome. Her payment is her claimed value for the final outcome minus the utility-target that she requested, and the auctioneer chooses the outcome that maximizes the total payment. In essence, the utility-target auction isolates the single dimension (utility) along which a bidder truly wishes to strategize

We begin with a static analysis of the utility-target auction's equilibria. We first show that the utility-target auction is quasi-incentive compatible: a bidder never has an incentive to misreport her valuation function — it is always sufficient for her to manipulate the utility-target she requests. We show that a pure-strategy equilibrium always exists and that the egalitarian equilibrium is efficiently computable. These results are really driven by the fact that the utility target auction allows bidders to communicate a profit-target bidding strategy with less strategic communication, and profit-target bidding strategies always exist.

Next, we show that the utility-target auction equilibria have good performance, effectively inheriting the performance guarantees of profit-target/truncation bidding strategies in package auctions [Milgrom, 2004; Day and Milgrom, 2007]. Similar to the approach of Edelman et al. [2007] on the generalized second-price (GSP) auction, we show that all equilibria satisfying a natural envy-free criterion have good performance. First, such equilibria are efficient and generate at least as much revenue as the incentive compatible Vickrey-Clarke-Groves (VCG) mechanism. Moreover, they generate revenue even when the incentive compatible mechanisms fail — the revenue of the envy-free equilibria of a utility-target auction all meet an intuitive benchmark we call the second-price threat, even settings where the VCG mechanism may make little or no revenue.

Dynamic Analysis through Behavioral Axioms. A significant novelty of our work is our use of simple behavioral axioms to prove guarantees on the performance of utility-target auctions.

Dynamic arguments are generally fraught with peril: in addition to being difficult to prove, more complex auctions (or markets, or games) require more complex bidding behavior to converge to an equilibrium and therefore sacrifice robustness. For example, Walrasian tâtonnement[2] [Walras, 1954] is perhaps the earliest concrete dynamic procedure proposed in economics — it converges in general markets when modeled as a particular continuous process [Samuelson, 1941; Arrow et al., 1959] but may or may not converge as a discrete process [Bala and Majumdar, 1992; Cole and Fleischer, 2008]. More recent results have sought stronger guarantees, e.g. by showing that agents' behavior will converge to equilibrium in repeated games as long as their learning strategies are "adaptive and sophisticated" [Milgrom and Roberts, 1991] or no-regret [Hart and Mas-Colell, 2000; Even-dar et al., 2009]. However, these properties are sufficiently complicated that it is difficult to evaluate whether agents' strategies indeed satisfy them in practice.

In contrast, we build simple behavioral axioms and use them to prove performance guarantees. Our first axioms are that (a) a bidder who is losing will raise her bid to try to win, and (b) a bidder who is losing is more impatient than a bidder who is winning. After formalizing these axioms in the context of utility-target auctions, we show that the auction will eventually reach an outcome that satisfies a natural notion of envy-freeness and, by

[2]To justify market equilibrium as a predictor of actual market behavior, Leon Walras described a dynamic procedure called tâtonnement that might converge to it.

extension, a natural second-price type bound on revenue. Significantly, this result neither implies nor requires that agents' bids converge to a steady-state. Moreover, bidder behavior requires only knowing whether one is winning or losing, not the precise bids of other agents.

Next, we show that bidders will not overpay if two more axioms are also satisfied, namely that (c) bidders who are winning will try to lower their bid to save money. Axioms (a)-(c) guarantee that bids will ultimately remain close to the boundary between envy-free and non-envy-free outcomes, a boundary which contains the envy-free equilibria. These axioms offer a degree of robustness, since bids will seek this boundary even as bids and ads change.

Finally, we show that bids will converge to the egalitarian equilibrium — the equilibrium that distributes utility most evenly — if a fourth axiom is satisfied. The fourth axiom concerns the timing of raised bids: (d) the bidder who has the most value at risk is the least patient and therefore raises her bid first. When bidder behavior satisfies all five axioms (a)-(d), we show that bids will converge to the egalitarian equilibrium. Together, these results offer powerful guarantees about the performance of a utility-target auction in a repeated setting.

Related Work. Our utility-target auction is most closely related to first-price package auctions [Bernheim and Whinston, 1986] and the ascending proxy auction [Milgrom, 2004]. Profit-target bidding in these auctions is closely related to quasi-truthful bidding in utility-target auctions, and the static properties we prove in Section 6.4 all have direct analogues. our dynamic analysis is entirely new, a more general confirmation of Milgrom's postulate that profit-target equilibria "may describe a central tendency for some kinds of environments" [Milgrom, 2004].

Auctions in which a agent's bid directly specifies her payment are known as *pay-your-bid* auctions. The first-price auction, as well as the Dutch an English auctions, are members of this family. Our utility-target auction is closely related to first-price and ascending proxy package auctions [Milgrom, 2004]. Engelbrecht-Wiggans and Kahn [1998] explored multi-unit, sealed-bid pay-your-bid auctions and found their equilibria to be substantially different from the standard first-price auction — the core issue they encounter is the same one arising in GFP.

A key reason repeated auctions may admit more robust performance guarantees is that bidders can learn about others' valuations. A similar informational exchange is present in and a motivation for classic ascending auctions. In addition to his discussion of ascending proxy auctions [Milgrom, 2004], Milgrom [2000] offers a broad discussion of this literature. Some recent work studies ascending auctions for position auctions like sponsored search [Edelman et al., 2007; Ashlagi et al., 2010].

Our work can also be seen through the lens of *simple versus optimal* mechanisms [Hartline and Roughgarden, 2009]. The general goal of this line of research is to design a mechanism that is simple and transparent while (possibly) sacrificing efficiency or revenue. For example, Hart and Nisan analyze the tradeoff between the number of different bundles offered to a buyer and an auction's performance [Hart and Nisan, 2013]. By comparison, our results show that a first-price auction can guarantee good performance when the bidding complexity is only slightly larger than that of the valuation functions.

6.2. Definitions and Preliminaries

The *utility-target auction* is a generalization of the first-price auction. Its key feature is an extra utility-target parameter in the bid — this parameter highlights the key dimension along which bidders care to compete. It gives bidders sufficient flexibility to guarantee the existence of pure-strategy equilibria while minimizing the communication required between the bidders and the auctioneer.

6.2.1. First-Price and Pay-Your-Bid Auctions.

An auction is a protocol through which agents bid to select an outcome. A standard sealed-bid auction can be decomposed into three stages: (1) each agent i submits a bid b_i, (2) the auctioneer uses agents' bids to pick an outcome o from a set $\mathbf{O}$, and finally (3) each agent i pays a price p_i. The final utility of agent i is given by $v_i(o) - p_i$, where $v_i(o) \geq 0$ denotes i's value for the outcome o, i.e. $v_i \in V_i$ is i's valuation function (drawn from a publicly known set V_i).

From this perspective, the standard first-price auction is described as follows: (1) each agent submits a single number $b_i \in \Re$, (2) the auctioneer chooses to give the item to the agent i^* who submits the largest bid b_i, and (3) the winner i^* pays b_{i^*} and everyone else pays zero. For comparison, the second-price auction is identical to the first-price auction except that the price paid is equal to the second-highest value of b_i.

When the outcomes are few, we will use v_i and b_i to denote the profile of values and bids across outcomes, e.g., $v_i = (1, 1.5, 0)$.

When considering settings beyond the single-item auction there are many ways to generalize the first-price auction. Even within the single-item setting, the auctioneer could choose an arbitrary encoding for agents' bids. Moreover, the auctioneer might choose an encoding that changes the space of possible bids, e.g by forcing bidders to place integer bids when values are actually real numbers. In such cases, the principle feature that we wish to preserve is that the winner "pays what she bid," or alternatively that a agent's bid precisely specifies her payment. Formally, we say that such an auction has the pay-your-bid property:

Definition 6.1. An auction has the *pay-your-bid property* if the payment p_i depends only on the outcome o and i's bid b_i (it does not directly depend on others' bids).

The first-price auction as described above clearly satisfies this property while a second-price auction does not.

Not all sealed-bid pay-your-bid auctions are equivalent. Edelman et al. [2007] showed that GFP, where the set of possible bids is precisely V_i, did not have a pure-strategy equilibrium:

Observation 6.1. The pay-your-bid property does not guarantee the existence of a pure-strategy equilibrium in a sealed-bid auction when the space of bids is the same as the space of valuation functions.

Moreover, as we discuss in Section 6.5, any pay-your-bid ad auction where bids are restricted to a subset of V_i must suffer in terms of its welfare and revenue guarantees. Thus, it is import to consider auctions that allows bids $b_i \notin V_i$. This motivates us to introduce the utility-target auction, a sealed-bid pay-your-bid auction that allows such bids and always has pure-strategy equilibria with strong performance guarantees.

6.2.2. Utility-Target Auctions

A utility-target auction is a sealed-bid pay-your-bid auction with a special bidding language. A agent's bid specifies payments using two pieces of information: her valuation function and the amount of utility she requests (a single real number). Her payment for an outcome is her (claimed) valuation for that outcome minus the utility that she specified in her bid. Formally:

Definition 6.2. A *utility-target auction* for a finite outcome space $\mathbf{O}$ is defined as follows:

ALGORITHM 6.1: A generic utility-target auction.

input : Players' bids $b_i = (x_i, \omega_i)$
output: An outcome o^* and first-price payments p_i.

Let $b_i(o) = \max(0, x_i(o) - \omega_i)$; `// ` $b_i(o)$ ` is ` i`'s effective bid for outcome ` o.
Compute $o^* = \arg\max_o \sum_{i \in [n]} b_i(o)$; `// Choose the outcome with the highest total bid.`
For all i, set $p_i = b_i(o^*)$; `// Each agent pays what she bid.`

- A bid is a tuple $b_i = (x_i, \omega_i)$ where $x \in V_i$ is a function mapping outcomes $o \in \mathbf{O}$ to nonnegative values and ω is a real number. We call the parameters x_i and ω_i the *value bid* and *utility-target bid* respectively.
- A bidder's effective bid for outcome o is

$$b_i(o) = \max(x_i(o) - \omega_i, 0) \ .$$

  Note this may generate $b_i \notin V_i$ when the set V_i is sufficiently restricted.
- The auctioneer chooses the outcome $o^* \in \mathbf{O}$ that maximizes $\sum_{i \in [n]} b_i(o)$. Ties are broken in favor of the most-recent winning outcome when applicable.
- When the outcome is o, bidder i pays $p_i(o) = b_i(o)$ and derives utility $u_i(o) = v_i(o) - b_i(o)$. Note that if a bidder reports $x_i = v_i$, then $u_i(o) = \omega_i$ whenever $v_i(o) \geq \omega_i$.

A generic utility-target auction is illustrated in Algorithm 6.1.

6.3. Quasi-Truthful Bidding

An idealist's intuition for the utility-target auction is that agents truthfully reveal their valuation function through their value bids (i.e. they bid bid $x_i = v_i$) and then use the utility-target bid ω_i to strategize. Clearly, bidders need not follow this ideal; however, it turns out that they have no incentive to do otherwise — the utility-target auction is quasi-truthful in the sense that for any bid a agent might consider, there is another bid in which she reveals v_i truthfully and obtains at least as much utility:

Lemma 6.1 (Quasi-Truthfulness)**.** *Fix the total bid of agents $j \neq i$ for all outcomes, i.e. fix $\sum_{j\in[n]\setminus\{i\}} b_j(o)$ for all o, and suppose ties are broken according to a fixed total-ordering on outcomes. If bidder i gets u_i^I by bidding (x_i^I, ω_i^I), then she gets the same utility u_i^I by bidding (v_i, u_i^I).*

Significantly, this implies bidder i always has a quasi-truthful best-response.

Proof. Since ties are broken according to a fixed total ordering, the outcome is fully specified by the total bids for each outcome (i.e. by $\sum_{i\in[n]} b_i(o)$ for all o). Thus, given $\sum_{j\in[n]\setminus\{i\}} b_j(o)$ and a bid $b_i^I = (x_i^I, \omega_i^I)$ for i, the outcome o^I is uniquely defined. Let ω_i^I be the utility i gets by bidding (x_i^I, ω_i^I), i.e.

$$u_i^I = v_i(o^I) - b_i(o^I) = v_i(o^I) - \max(x_i^I(o^I) - \omega_i^I, 0) \ .$$

Now suppose i bids $b_i^Q = (v_i, u_i^I)$ instead of (x_i^I, ω_i^I). There are two possible results of this change:

- *The outcome doesn't change.* If the outcome doesn't change, then i gets the same utility by construction.
- *The outcome changes to $o^Q \neq o^I$.* Notice that i did not change the amount bid for outcome o^I, so the total bid for o^I did not change. Given this and the tie-breaking rule, the only way the outcome can switch from o^I to o^Q is if the total bid for o^Q

strictly increased. Given that $\sum_{j\in[n]\setminus\{i\}} b_j(o^Q)$ is fixed, this implies i's bid for o^Q increased, i.e. $b_i^Q(o^Q) > b_i^I(o^Q) \geq 0$.

Next, by definition of a utility-target auction, $b_i(o) > 0$ implies $x_i(o) > \omega_i(o)$. Since $b_i^Q(o^Q) \geq 0$, this implies $v_i(o^Q) > \omega_i^Q$, from which it immediately follows that i's final utility in o^Q will be ω_i^Q.

In either case, i's final utility is precisely u_i^I, so i is indifferent between bidding (x_i^I, ω_i^I) and (v_i, u_i^I). □

6.4. Static Equilibrium Analysis

We begin by studying the utility-target auction from a static perspective and show that they offer strong revenue and welfare guarantees. First, we show that pure-strategy equilibria always exist:

Theorem 6.2. *A utility-target auction with n outcomes always has a pure-strategy co-operatively envy-free (defined below) equilibrium that is computable in time poly(n).*

Specifically, the egalitarian equilibrium exists and is efficiently computable by Algorithm 6.2 (proof omitted).

Next, we show that such cooperatively envy-free equilibria not only maximize welfare but offer as much revenue as the VCG mechanism as well as a new revenue benchmark we call the second-price threat (defined below):

Theorem 6.3. *Any cooperatively envy-free equilibrium of a utility-target auction*

(1) *maximizes social welfare,*

(2) *dominates the revenue of the VCG mechanism,*

(3) *and has revenue lower-bounded by the second-price threat.*

We formalize and prove the theorem below.

6.4.1. Cooperatively Envy-Free Equilibria

While a typical utility-target auction may have many equilibria, some of them are unrealistic in repeated auctions. In particular, it is possible to have an equilibrium in which a group of "losers" envy the "winners" — the losers would be happy to collectively raise their bids to make an alternate outcome win, but the outcome is an equilibrium because no single bidder is willing to raise her bid high enough. In a repeated setting, one would expect all the losers to eventually raise their bids.

For example, consider a setting with three bidders (A, B, C) and three outcomes $(1, 2, 3)$, in which the first two bidders are symmetric and value the first two outcomes and the third bidder values only the third outcome. Let the specific values, indexed by outcome, be

$$v_A = (1, 1.5, 0),$$

$$v_B = (1, 1.5, 0),$$

$$v_D = (0, 0, 2).$$

Now, let A and B bid for the first outcome, and C bid for the third outcome, with bids: $b_A = (1, 0, 0)$, $b_B = (1, 0, 0)$ and $b_C = (0, 0, 2)$.

Bidders A and B would prefer the second outcome to the first as they see a value of 1.5 instead of 1. Moreover, they would be happy to make the second outcome win by cooperating and each bidding $1 + \epsilon$. However, since both are bidding 0 for the second outcome, neither can unilaterally cause the second outcome to win, making this outcome an equilibrium. The problem in this example is that, at a total price of 2, bidders A and B would prefer that the second outcome wins. In a sense, bidders A and B in the second outcome envy the deal they received in the first outcome.

Hence, we are interested in bids such that agents have no incentive to cooperatively deviate to get a better outcome. We will call such a set of bids *cooperatively-envy free.*

To define such a notion, we must also consider bidders who are happy with the winning outcome. Consider a four bidder setting with three possible outcomes, with the following values:

$$v_A = (1, 1.5, 0),$$

$$v_B = (1, 1.5, 0),$$

$$v_C = (1, 0.5, 0),$$

$$v_D = (0, 0, 2).$$

In this case, C cannot get a better deal from the second outcome, so she will not cooperate with A and B. In order to win, A and B must collectively bid 1.75 in the second outcome to make up the deficit between it and the winning outcome (which they are willing to do).

Definition 6.3. The set of bids $\{b_i\}_{i\in[n]}$ are *cooperatively envy-free (CEF)* if there is no subset of bidders $\mathbf{J} \subseteq [n]$ who would prefer to cooperatively pay the extra money required to make an alternate outcome o win over the current winner o^*.

Formally, a set of bids is cooperatively envy-free if

$$\sum_{i\in[n]} \max\left((v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)), 0\right) \leq \sum_{i\in[n]} b_i(o^*) - b_i(o)$$

for all outcomes o.

The CEF constraints are similar to the core property described by Milgrom [2004] in package auctions (as well as notions like group-strategyproofness); however, the notion of a CEF outcome is weaker. For example, it does not require that bidders are playing equilibrium strategies.

Equilibria that are CEF have nice properties analogous to those of core equilibria in package auctions. The following claims are straightforward and proven in the appendix of the full version of the paper:

Claim 6.4. *CEF bids maximize welfare.*

Claim 6.5. *The revenue from CEF bids dominate that of the VCG mechanism: every agent pays at least as much in the CEF equilibrium as she would in the VCG mechanism.*

6.4.2. The "Second-Price Threat"

The revenue of a CEF equilibrium also meets or exceeds a benchmark we call the second-price threat. The revenue of the second-price auction has a convenient intuition: the price paid by the winner should be at least as large as the maximum willingness to pay of any other bidder. We can ask the same question in more general settings: how much would "losers" be willing to pay to get an outcome o instead of the socially optimal outcome o^*? In general, agent i should be willing to pay up to $v_i(o) - v_i(o^*)$ to help o beat o^*, hence we can generalize the intuition of the second-price auction to give a natural lower bound on the revenue the auctioneer might hope to earn:

Definition 6.4. The *second-price threat* for outcome o^* is given by

$$\max_{o \in \mathbf{O}} \sum_{i \in [n]} \max(v_i(o) - v_i(o^*), 0) \ .$$

ALGORITHM 6.2: An algorithm for computing the egalitarian equilibrium in a utility-target auction.

input : A utility-target auction problem.
output: The egalitarian equilibrium bids $b_i^* = (v_i, \omega_i^*)$.

1 Set all bids to $(v_i, 0)$. Call the socially optimal outcome o^*.
2 Increase ω_i for all bidders uniformly until some bidder i reaches $\omega_i = v_i(o^*)$ or a CEF constraint would be violated for some outcome o.
3 Fix the bids of the newly-constrained advertisers.
4 Repeat (2) and (3), lowering only unfixed bids until all bidders are fixed.

This bound is particularly powerful in cases where bidders share value for an outcome (cases where VCG would make little or no revenue). For example, consider the following 4-bidder, 2-outcome setting:

$$v_A = (1, 0)$$

$$v_B = (1, 0)$$

$$v_C = (1, 0)$$

$$v_D = (0, 2)$$

In a VCG auction, nobody pays anything. However, a naïve auctioneer might expect the first outcome to win, with A, B, and C paying a total of \$2 (the second-price threat) since they are beating D.

The CEF constraints quickly imply that a CEF outcome generates at least as much revenue as the second-price threat:

Claim 6.6. *The revenue in any CEF outcome is lower-bounded by the second-price threat.*

The proof is given in Appendix C.

6.5. Utility-Target Auctions for Sponsored Search

Sponsored search advertising demonstrates the benefits of a utility-target auction. The standard auction in this setting is the generalized second-price (GSP) auction; however, it (and incentive-compatible VCG mechanisms) lack transparency: payments are complicated to compute and bidders must trust the auctioneer not to abuse their knowledge when an auction is repeated. Moreover, its performance may degrade when using more accurate models of user behavior [Roughgarden and Tardos, 2012] and advertiser value [Hoy et al., 2012]. It can have misaligned incentives when parameters are estimated incorrectly [Wilkens and Sivan, 2012]. Some of these problems would be solved by a first-price auction; however, Overture's implementation of GFP demonstrated that such schemes might be highly unstable. A utility-target auction offers the benefits of a pay-your-bid auction without the instability of GFP.

6.5.1. The Utility-Target Ad Auction

We illustrate a utility-target auction in the standard model of sponsored search: n advertisers compete for $m \leq n$ slots associated with a fixed keyword. An advertiser's value depends on the likelihood of a click, called the click-through-rate (CTR) $\alpha_{i,j}$, and the value v to the advertiser of a user who clicks. The CTR $\alpha_{i,j}$ is separable into a parameters β_i that depends on the advertiser and α_j that depends on the slot, so the expected value to advertiser i for having her ad shown in slot j is $\alpha_{i,j}v_i = \alpha_j\beta_i v_i$. As is standard, we assume that slots are naturally ordered from best ($j = 1$) to worst ($j = m$), i.e. $\alpha_j \geq \alpha_{j'}$ for all $j < j'$. Without loss of generality, we assume bidders are ordered in decreasing order of bid, i.e. $b_1 \geq b_2 \geq \cdots \geq b_n$.

The auctioneer chooses a matching of advertisements to slots and charges an advertiser a per-click price ppc_i. For example, in the GFP auction, advertisers submitted bids b_i representing their per-click payment and paid $ppc_i = b_i$ whenever a their ads were clicked. Similarly, in the standard GSP auction, bidder i is charged according bid of the next highest bidder.[3] To account for differences in CTRs, this quantity is normalized by β so that bidder i pays a per-click price of $ppc_i = \frac{\beta_i}{\beta_{i+1}} b_{i+1}$.

In a *utility-target auction*, bidders submit both their per-click value x_i and the utility-target bid ω_i (the utility that they request). The auctioneer picks the assignment $j(i)$ maximizing

$$\sum_{i \in [n]} \max(0, \alpha_{j(i)} \beta_i x_i - \omega_i)$$

and charges i so that her expected payment is

$$\mathbf{E}[p_i] = \max(0, \alpha_{j(i)} \beta_i x_i - \omega_i) \ .$$

There are at least two interesting ways the utility-target auction can be implemented. The first implementation charges

$$ppc_i = \max\left(0, x_i - \frac{\omega_i}{\alpha_{j(i)} \beta_i}\right)$$

to achieve the desired expected payment. In effect, it uses the utility request ω_i to compute a different per-click bid for each slot. A practical downside to this implementation is that the payments are still somewhat complicated from the bidders' perspectives; however, the

[3]The designers of the GSP auction intended it to inherit the incentive compatibility of the second-price auction. It does not; however, it has the nice property that bidder i pays the minimum amount required to win the slot that she received.

auctioneer could mitigate this problem by publishing CTRs and displaying the per-click payments in the bidding interface.

An alternative implementation of the utility-target auction pays a rebate of ω_i regardless of whether a click occurred and charges precisely $ppc_i = x_i$ when a click occurs. This auction is even simpler from the bidders' perspective; however, when a click does not occur the auctioneer will be paying the bidder (in expectation the bidder still pays the auctioneer). This implementation of the utility-target auction is illustrated in Algorithm 6.3.

Such a utility-target auction offers many benefits over existing auction designs like GSP and VCG. As noted earlier, a first-price auction directly increases transparency and simplicity from the bidders' perspective. Even if bidders reveal their true valuation functions v_i, the pay-your-bid property ensures that increasing a reserve price will not increase payments unless bidders subsequently raise their bids.

The auction also easily generalizes to more complicated bidding languages. Whereas the welfare and revenue performance of GSP degrades (albeit gracefully) Roughgarden and Tardos [2012] when considering externalities imposed by the presence of competing ads, the reasonable (CEF) equilibria of the utility-target auction guarantee good performance. An open question of Hoy et al. [2012] is to find an auction that performs well in multi-slot settings when multiple bidders can benefit from clicks on the same ad (e.g. Microsoft and Samsung both benefit from an ad for a Samsung laptop running Windows) — the utility-target auction offers good revenue guarantees in these 'coopetitive' ad auctions with multiple slots. The utility-target auction is also less sensitive to estimation errors in the CTRs. As shown in [Wilkens and Sivan, 2012], incentive-compatibility can be broken because the auctioneer only knows estimates of the α and β parameters. Informally, the utility-target auction is much less sensitive to such errors because the payments need not explicitly depend on the auctioneer's estimates.

6.5.2. Utility-Target vs. GFP

Juxtaposing GFP with the utility-target auction illustrates the benefits of a more complex bidding language. GFP is identical to the utility-target ad auction except that bids contain only the per-click payment x_i and not the utility-target bid ω_i. Consequently, a agent's bid necessarily offers the same per-click payment regardless of the slot won by the bidder. By comparison, the utility-target auction permits bids that encode a different per-click payment depending on the slot in which an ad is shown.

In retrospect, it is easy to see that different per-click bids are important for a good pure-strategy equilibrium. In GFP, all advertisers who are shown must bid so that $\beta_i x_i$ is the same, otherwise some bidder can lower her value of x_i without changing her assignment; however, if this is true, then some bidder can move up to the top slot by bidding $x_i + \epsilon$. In fact, any bidding language that requires the same per-click payment for all slots could not have a pure-strategy equilibrium unless the potential benefit of being in the top slot was less than the effective bid increment required to get there. This necessarily weakens any revenue guarantees and, worse, implies that the auction cannot differentiate between the winning bidders to pick the best ordering of ads.

As noted earlier, the existence of a pure-strategy equilibrium is directly related to the dynamic performance of the auction. Edelman and Ostrovsky [2007] discuss how the lack of such an equilibrium naturally leads to sawtooth cycling behavior in GFP, as bidders alternate between increasing their bids to compete for higher slots and decreasing their bids to avoid overpaying for the slots they have. They also show that this cyclic behavior potentially

---

ALGORITHM 6.3: A utility-target auction for search advertising.

---

input : Bids $b_i = (x_i, \omega_i)$.
output: An assignment of advertisements to slots, per-click payments ppc_i, and unconditional payments r_i.

For each bidder i and slot j, compute $\mathbf{E}[p_{i,j}] = \max(\alpha_j \beta_i x_i - \omega_i, 0)$;
Compute assignment $j(i)$ of advertisements to slots that maximizes $\sum_{i\in[n]} \mathbf{E}[p_{i,j(i)}]$;
if $\mathbf{E}[p_{i,j(i)}] > 0$ **then**
  Always pay i the rebate $r_i = \omega_i$;
  Whenever i's ad is clicked, charge $ppc_i = x_i$;
else
  Do not charge/pay anything to i;

---

reduced revenue below that of the VCG mechanism. In contrast, Theorem 6.2 shows that utility-target auctions have pure-strategy equilibria, and Theorem 6.3 shows that revenue at equilibrium dominates the VCG mechanism; moreover, our dynamic results show that bids will naturally approach this equilibrium (or the set of such equilibria) as bidders adjust their utility targets.

6.6. Dynamic Analysis

In this section, we consider the behavior of utility-target auctions in a very simple dynamic setting, and under very simple assumptions. We show that a few rules and simple knowledge of whether one is winning or losing are enough to guarantee the revenue and welfare bounds from Theorem 6.3.

Following Lemma 6.1, we assume that bidders are quasi-truthful and report bids of the form (v_i, ω_i). Bidders compete using the utility-target terms ω_i and employ strategies to optimize their utility.

Winners and Losers. Our dynamic axioms are based on a natural decomposition of bidders into winners and losers. In a standard first-price auction, the winner is the bidder who gets what he wants — the item — and the losers are those who do not get what they want. In a utility-target auction, a bidder effectively reports her valuation v_i and requests that the auctioneer give her a certain utility ω_i. This suggests partitioning bidders into winners and losers based on whether a bidder gets the utility she requests, giving us the following formal definition:

Definition 6.5 (Winners and Losers)**.** A *winner* is a bidder who gets the utility she requests, i.e. if o^b is the outcome of the auction, then i is a winner if and only if

$$u_i(o^b) = v_i(o^b) - b_i(o^b) = \omega_i \ .$$

Any bidder who is not a winner is a *loser*.

Observation 6.2. Bidder i is a winner if and only if $v_i(o^b) \geq \omega_i$ and is always a winner when $\omega_i = 0$. When bidder i is a loser, $u_i(o^b) < \omega_i$.

Note that this definition does not coincide with the standard definition of winners and losers in a single item auction because a bidder who does not get the item is still a winner if $\omega_i = 0$.

Raising and Lowering Bids. In a dynamic setting, we want to think about how winners and losers manipulate ω_i. In the utility-target auction, the effective bid b_i (what bidder i is actually offering to pay) and the utility-target term ω_i move in opposite directions, so when we talk about raising i's bid we are talking about decreasing the utility-target term ω_i:

Definition 6.6 (Raising and Lowering Bids)**.** We say that bidder i *raises her bid* from (v_i, ω_i) if she chooses a new bid (v_i, ω_i') where $\omega_i' < \omega_i$, i.e. she raises her bid if she *decreases her utility-target bid.*

Similarly, a bidder who *lowers her bid* correspondingly *increases her utility-target bid* from ω_i to $\omega_i' > \omega_i$.

Importantly, our definition of winners and losers shares a natural property with the standard definition: winners cannot benefit by offering to pay more, and losers cannot benefit by offering to pay less:

Claim 6.7. *Fixing other agents' bids, a loser cannot increase her utility by raising her bid. Likewise, a winner cannot increase her utility by lowering her bid.*

The claim is straightforward to prove.

Our definition of winners and losers also has a new property that is important:

Claim 6.8. *A loser can always raise her bid in a way that weakly increases her utility.*

Proof. Suppose i is a loser bidding (v_i, ω_i) and receiving utility $u_i < p_i$. If she raises her bid to (v_i, u_i), Lemma 6.1 says that she will receive utility of precisely u_i, making her a winner. $\square$

In our model, bidders locally adjust their bids by ϵ. To mimic settings where auctions happen frequently and no two bidders move simultaneously, bid changes are modeled as asynchronous events. As noted earlier, our model assumes agents bid quasi-truthfully, that is, they always submit their true valuation functions in their bids. As a result, the history of the auction is characterized by a sequence of utility-target vectors $\omega^0, \ldots$.

We assume that $0 \leq \inf_o v_i(o)$ and $\sup v_i(o) < \infty$, so utility-targets will always lie in the finite interval $[0, \sup v_i(o)]$. Unless a agent's utility-target hits the boundary of this interval, all bid changes are made in increments of ϵ.

Notions of Convergence. We will show that progressively stronger assumptions imply progressively stronger convergence guarantees. Our first results show that bids will eventually be close to the set of CEF (or non-CEF) bids. As noted earlier, the utility-targets ω are sufficient to characterize bidders' strategies, so we define $\mathcal{C}$ to be the set of all such utility-target:

Definition 6.7 (The CEF Set)**.** $\mathcal{C}$ is the set of all utility-target vectors ω where the quasi-truthful bids (v_i, ω_i) produce a cooperatively envy-free outcome.[4]

The set $\overline{\mathcal{C}}$ is the set of all utility-target vectors which are *not* CEF, i.e. $\overline{\mathcal{C}} = \Re^n_+ \setminus \mathcal{C}$.

Significantly, $\mathcal{C}$ is never empty. In particular, it always contains the 0 vector ($0^n \in \mathcal{C}$).

Since bidders are continually experimenting with their bids, it is not realistic to expect bids to explicitly converge to $\mathcal{C}$; rather, they will remain close. For a set of bids ω, let ω_ϵ denote the set of bids that are close to some vector in ω, i.e.

Definition 6.8. Let S_ϵ be the set of all utility-targets ω which are close to some vector in ω coordinate-wise. Formally,

$$S_\epsilon = \{\omega \mid \exists \omega' \in S \text{ s.t. } ||\omega - \omega'||_\infty \leq \epsilon\} \ .$$

In particular, we will care about the sets $\mathcal{C}_\epsilon$ and $\overline{\mathcal{C}}_\epsilon$, the sets representing bids close to being CEF and close to being not CEF, respectively.

Next we define the convergence of an auction to utility-target bids ω:

[4]Note that membership in $\mathcal{C}$ depends on both the vector ω and the outcome chosen by the auction. This is because certain utility-target vectors ω will be in $\mathcal{C}$ if ties are broken in favor of o^* but not if ties are broken in favor of a suboptimal outcome.

Definition 6.9. An auction *converges to a set of utility-targets S* if, for any $\delta > 0$, there exists a sufficiently small bid adjustment parameter ϵ for which the auction always reaches a utility-target ω such that all future bids are in S_δ.

Our strongest result wull show that bids converge to the egalitarian equilibrium:

Definition 6.10. The egalitarian equilibrium is the CEF equilibrium which distributes utility as evenly as possible. Formally, for each equilibrium let $u_\uparrow$ be the vector of bidders' utilities with its coordinates sorted in increasing order. The *egalitarian equilibrium* is the one for which $u_\uparrow$ is lexicographically maximized.

An auction converges to the egalitarian equilibrium ω^E if it converges to $\{\omega^E\}$.

6.6.1. Axioms and Results

Our convergence theorems show that progressively stronger assumptions about bidder behavior lead to progressively stronger convergence results.

Our first axiom of bidder behavior captures some intuition about how winners and losers behave. Following Claim 6.7, a winner cannot benefit by raising her bid and a loser cannot benefit by lowering it, so we suppose that they never do this. Additionally, a loser who is actively engaged in the auction should raise her bid if it is beneficial. By Claim 6.8 we know that a loser can always raise her bid in a way that is weakly beneficial, so we suppose that a loser will always try to raise her bid.

(A1). *A losing bidder will raise her bid in an effort to win; a loser will not lower her bid and a winner will not raise her bid.* Formally, if the current utility-target is ω and i is a loser, then i must raise her bid at some point in the future unless she becomes a winner through the actions of other bidders.

Anecdotal evidence suggests that advertisers bidding in an ad auction generally expend substantial effort to launch advertising campaigns but are much slower to change things once they appear to work. Our second axiom generalizes this idea by supposing that winners (who, by definition, get the utility-target they request) view the outcome of the auction as a success while losers are unhappy with the results:

(A2). *A bidder who is losing is more impatient than a bidder who is winning.* Formally, if the current utility-target is ω and a set of bidders $L \subseteq [n]$ are losers, then the next time bids change it will necessarily be because some loser $i \in L$ raised her bid.

Our third axiom is analogous to (A1) but for winners — a winner who is actively engaged should lower her bid from time to time to see if she can win at a lower bid.

(A3). *A winner will try lowering her effective bid to win at a lower price. Specifically, if a bidder is currently a winner, then she must lower her bid at some point in the future unless she becomes a loser through the action of another agent.* Formally, if the current utility-targets are ω and i is a winner, then i must lower her bid at some point in the future unless she becomes a loser through the actions of other bidders.

Our final axiom concerns the relative timing of events. Intuition and anecdotal evidence suggests that larger bidders who have more at stake tend to invest more heavily in active bidding strategies. This axiom roughly represents that intuition:

(A4). *Between two losers, the bidder with the higher utility-target is more impatient.* Formally, if the current utility-targets are ω and bidders i and j are both losers, then i will raise her bid before j if $\omega_i > \omega_j$.

These simple properties of bidder behavior imply the following convergence results. Proofs follow in Section 6.6.2 and Appendix C.1.

Theorem 6.9. *If losing bidders will only raise their effective bids (A1) and are more impatient than winning bidders (A2), the auction converges to the set of bids that are cooperatively envy-free (i.e. bids will be in $\mathcal{C}_\epsilon$).*

Theorem 6.10. *If winners try to lower their effective bids (A3) and losers try to raise but not lower their effective bids (A1), the auction converges to the set of bids that are non-cooperatively envy-free (i.e. bids will be in $\overline{\mathcal{C}}_\epsilon$).*

Combining Theorems 6.9 and 6.10 shows that bids will converge to the frontier of the CEF set. The strict Pareto frontier of this set is the set of CEF equilibrium bids.

Corollary 6.11. *If losing bidders will try raising their bids (A1), losers are less patient than winners (A2), and winners try lowering their bids (A3), the auction converges to the boundary between CEF and non-CEF bids (bids will be in the set $\mathcal{C}_\epsilon \cup \overline{\mathcal{C}}_\epsilon$).*

Finally, adding A4 induces convergence to a particular equilibrium:

Theorem 6.12. *If losing bidders will raise their effective bids (A1), winning bidders will try lowering their effective bids (A3), and the most impatient bidder is the losing bidder bidding for the highest utility-target (A2, A4), then bids will converge to the Egalitarian envy-free equilibrium.*

6.6.2. Convergence Proofs

In this section we give proofs of Theorems 6.9 and 6.10. Theorem 6.12 and omitted proofs may be found in Appendix C.1. Throughout this section, we assume that there is a single welfare optimal outcome for clarity of presentation.

Observation 6.3. Under assumptions A1 and A2, a bidder will only lower her bid if all bidders are winners.

Lemma 6.13. *If all bidders are winners under utility-targets ω, then ω is in the CEF set $\mathcal{C}$.*

Proof. If all bidders are winners, then we know that they are receiving precisely the utility-target they request when they bid ω. Intuitively, this means that raising bids necessarily implies receiving less utility.

Formally, if bids are $b_i = (v_i, \omega_i)$ and the outcome of the auction is o^b, then we want to show that the CEF condition holds for any outcome o. Since all bidders are winners, we know $v_i(o^b) - b_i(o^b) = \omega_i$. Moreover, $v_i(o) - b_i(o) \leq \omega_i$ by definition, so $v_i(o) - b_i(o) \leq v_i(o^b) - b_i(o^b)$. Thus

$$\max\big((v_i(o) - b_i(o)) - (v_i(o^b) - b_i(o^b)), 0\big) = 0 \ .$$

Since o^b is the outcome of the auction, we know $\sum_{i\in[n]} b_i(o^b) \geq \sum_{i\in[n]} b_i(o)$ for any outcome o. Thus, $0 \leq \sum_{i\in[n]} b_i(o^b) - b_i(o)$ and therefore

$$\sum_{i\in[n]} \max\big((v_i(o) - b_i(o)) - (v_i(o^b) - b_i(o^b)), 0\big) \leq \sum_{i\in[n]} b_i(o^b) - b_i(o)$$

as desired. □

Since A1 and A2 imply that a agent will only lower her bid from ω if all bidders are winners, an important corollary is that a bidder will only lower her bid if the current utility-target vector is in the CEF set $\mathcal{C}$:

Corollary 6.14. *Under assumptions A1 and A2, if a agent lowers her bid from ω, then ω is in the CEF set $\mathcal{C}$.*

A corollary of Claim 6.4 is that any set of CEF bids maximizes welfare, hence this implies that a agent will only lower her bid if the welfare-optimal outcome is winning:

Corollary 6.15. *Under assumptions A1 and A2, a bidder will only lower her bid if a welfare-optimal outcome o^* is winning.*

Another useful fact about $\mathcal{C}$ is that it is leftward-closed (the proof is in the appendix) and the natural corollary that $\overline{\mathcal{C}}$ is rightward-closed:

Lemma 6.16. *If ω is in the CEF set $\mathcal{C}$, then $\omega - \Delta$ is in the CEF set $\mathcal{C}$ for any $\omega \geq \Delta \geq 0$.*

Corollary 6.17. *If ω is in the not-CEF set $\overline{\mathcal{C}}$, then $\omega + \Delta$ is in the not-CEF set $\overline{\mathcal{C}}$ for $\Delta \geq 0$.*

To prove that bids will converge, we first show that bids will not be stuck at arbitrarily low values:

Lemma 6.18. *Suppose the initial vector of utility-targets is ω^0. Under properties A1 and A2, the auction will always reach a configuration in which all bidders are winners and will do so within $\left|\left\lceil \frac{1}{\epsilon}\omega^0 \right\rceil\right|_1$ steps.*

We can now prove the our first theorem, that bids will be close to $\mathcal{C}$ when A1 and A2 are satisfied.

Proof of Theorem 6.9: Lemma 6.18 implies that all bidders will be winners within a finite time. Once all bidders are winners, the only way bids will change is if someone lowers her bid. Thus, after a finite amount of time, we can conclude that either all bidders are winners or some bidder has lowered her bid.

Let ω be the vector of utility-targets at any point after the first time all bidders are winners. If all bidders are still winners, then $\omega \in \mathcal{C}$ by Lemma 6.13. Otherwise, let i be the most recent agent to lower her bid, increasing the utility-target vector from ω' to $\omega'' = \omega' + \epsilon e_i$. We show that if i raises her bid again then the resulting utility-targets must be CEF regardless of how bids have changed since i's raise.

By construction, agents have only raised their bids since i lowered hers, so we can define $\Delta = \omega'' - \omega$ where $\Delta \geq 0$. Corollary 6.14 tells us that $\omega' \in \mathcal{C}$. If i raised her bid between ω'' and ω, then $\omega \leq \omega'$ and Lemma 6.16 tells us that $\omega \in \mathcal{C}$, so were done. Otherwise, we know $\omega'' \geq \Delta \geq 0$ and Lemma 6.16 tell us that $\omega' - \Delta \in \mathcal{C}$. Therefore $\omega = \omega' - \Delta + \epsilon e_j \in \mathcal{C}_\epsilon$. □

To prove Theorem 6.10, we need a lemma similar to Lemma 6.18 showing that the auction will reach a bid vector that is CEF:

Lemma 6.19. *Under properties A1 and A3, as long as there is some outcome o and bidder j such that $v_j(o) > v_j(o^*)$, the auction will always reach a configuration that is not CEF when ϵ is sufficiently small. If there is no such outcome o and bidder j, then the auction may converge to $\omega_j = v_j(o^*)$ instead.*

Proof Sketch of Theorem 6.10: By Lemma 6.19, the auction will typically reach a utility-target vector in $\overline{\mathcal{C}}$. Our primary goal is to show that the auction will be in $\overline{\mathcal{C}}_\epsilon$ from that point onwards.

Suppose $\omega \in \mathcal{C}$. Let ω' be the most recent utility-targets that were in $\overline{\mathcal{C}}$ and let $\omega'' \in \mathcal{C}$ be the utility-targets immediately after ω'. Let i be the bidder who changed her bid between ω' and ω''. Let o be an outcome that violated the CEF constraints at ω'.

Consider a bidder j whose bid is higher at ω than at ω''. Notice that j must be a loser to raise her bid, roughly $v_j(o^*) < \omega_j$. It then follows from the definition of the utility-target auction that $b_j(o^*)$ essentially did not change from ω' to ω. Moreover, $b_j(o)$ can only have increased from ω' to ω, so $b_j(o)$ increased more than $b_j(o^*)$ did. Similar logic leads to an analogous conclusion for bidders whose bids were lower at ω'' than ω, roughly giving

$$b_j(o) - b_j(o^*) \geq b'_j(o) - b'_j(o^*)$$

for each bidder $j \neq i$. For bidder i, a similar inequality holds accounting for the fact that i raised her bid from ω' to ω'':

$$b_i(o) - b_i(o^*) \geq b'_i(o) - b'_i(o^*) - \epsilon \ .$$

Finally, given that the CEF constraint for o was violated at ω', these inequalities imply it must be nearly violated at ω. □

We have now shown that when losing bidders raise their effective bids and winning bidders lower their effective bids, bids remain close to the frontier of the CEF set $\mathcal{C}$. Adding in the specific behavior that the first agent to raise their bid will be the losing bidder with the highest utility-target results in convergence to one specific equilibrium: the egalitarian equilibrium (Theorem 6.12). The full proof is included in the appendix; we provide a sketch of it here.

Proof Sketch of Theorem 6.12:

Arrange bidders into levels $L_1, \ldots, L_k$ in increasing order of the utility each bidder gets at the egalitarian equilibrium.

For each level L_{i+1}, bids from all bidders in the level will converge close to the egalitarian equilibrium once the bids of lower level bidders are sufficiently close to their egalitarian bids.

Thus, beginning with the bidders who get the least utility in equilibrium, and working on up to the lucky bidders with the most utility, bids will converge close to the egalitarian outcome. □

6.7. Conclusion and Open Questions

Pay-your-bid auctions — and utility-target auctions in particular — offer many advantages over incentive compatible mechanisms in terms of transparency and simplicity. Moreover, in many complex settings they even appear to generate more revenue.

Our work first shows that the bidding language is important in first-price auction design. In particular, it is both important and sufficient that bidders can compete in terms of their final utility. Also, a key feature in a repeated first-price auction is a pure-strategy equilibrium, something that GFP does not have Edelman et al. [2007]. This is a question of design: the existence of pure-strategy equilibria may be guaranteed through a carefully crafted bidding language (e.g. the utility-target auction) that can encode different per-click payments for different ad slots.

More significantly, when agents compete on utility, our results show that robust performance guarantees may be derived using only simple axioms of bidder behavior that merely require knowledge of whether one is winning or losing. These results are powerful because they do not require an a priori assumption that the auction is in equilibrium or full information about others' bids.

Yet, reflection raises a concern about utility-target auctions: *why should bidders reveal their true valuation functions in a repeated auction*? We claimed that first-price auctions were better because the auctioneer could not cheat, but it would seem that quasi-truthfulness is just as dangerous. In fact, a quasi-truthful pay-your-bid auction is still strongly preferable to a standard second-price auction: even if the auctioneer knows a bidder's true valuation function, it cannot immediately increase the amount of money the bidder pays. By comparison, the auctioneer in a second-price auction might force a bidder to pay her full value in the second round by increasing the reserve price. The auctioneer is welcome to engage in a game of chicken or a "negotiation" with the bidder to see if she is willing to raise her bid, but the pay-your-bid property ensures that final approval still rests with the bidder.

In practice, systems may also be designed to encourage competition on the utility-target term and thereby recover stability. For example, Overture exacerbated the instability of the GFP auction by offering an API automating the sawtooth behavior. If an API were offered to compete on the utility-target term, bidders would likely use the API and stability would be restored, regardless of whether they were reporting their true valuation functions.

Issues of quasi-truthfulness aside, our work also raises questions about dynamic axioms of bidder behavior. Our axioms may be simple and natural, but strict adherence to them is clearly unrealistic. In this vein, many interesting questions are open:

(1) *How does the behavior of the auction change with small modifications to the axioms?* For example, we showed that bids would converge to the egalitarian equilibrium when the bidder with the most to gain raised first. Can we prove convergence to a different equilibrium by modifying agents' delays?

(2) *Do the performance guarantees still hold if axioms only hold probabilistically or on average?* It seems unlikely that bidder behavior always satisfies any particular set

of axioms. How do the dynamic guarantees change when axioms only hold most of the time?

(3) *What dynamic axioms do bidders actually obey?* An interesting experimental question is to determine what axioms are actually satisfied by bidder behavior. For example, could one experimentally measure bidders' delays and combine this with an answer to (1) to predict a particular equilibrium outcome?

References

Alaei, S., Fu, H., Haghpanah, N., Hartline, J., and Malekian, A. (2012). Bayesian optimal auctions via multi- to single-agent reduction. In *ACM Conference on Electronic Commerce*.

Arrow, K. J. (1971). *Essays in the theory of risk-bearing*. Markham.

Arrow, K. J., Block, H. D., and Hurwicz, L. (1959). On the stability of the competitive equilibrium, ii. *Econometrica*, 27(1):pp. 82–109.

Ashlagi, I., Braverman, M., Hassidim, A., Lavi, R., and Tennenholtz, M. (2010). Position auctions with budgets: existence and uniqueness. *The B.E. Journal of Theoretical Economics*, 10(1).

Athey, S. and Nekipelov, D. (2010). A structural model of sponsored search advertising auctions. In *Sixth ad auctions workshop*.

Ausubel, L. M. and Milgrom, P. (2006). The lovely but lonely vickrey auction. *Combinatorial auctions*, pages 17–40.

Bala, V. and Majumdar, M. (1992). Chaotic tatonnement. *Economic Theory*, 2(4):437–45.

Bernheim, B. D. and Whinston, M. D. (1986). Menu auctions, resource allocation, and economic influence. *The Quarterly Journal of Economics*, CI(1).

Bulow, J. and Klemperer, P. (1996). Auctions versus negotiations. *The American Economic Review*, 86(1):180–194.

Cai, Y., Daskalakis, C., and Weinberg, S. M. (2012). An algorithmic characterization of multi-dimensional mechanisms. In *ACM Symposium on Theory of Computing*, pages 459–478.

Caragiannis, I., Kaklamanis, C., Kanellopoulos, P., Kyropoulou, M., Lucier, B., Paes Leme, R., and Tardos, E. (2015). Bounding the inefficiency of outcomes in generalized second price auctions. *Journal of Economic Theory*, 156(C):343–388.

Chawla, S. and Hartline, J. D. (2013). Auctions with unique equilibria. In *ACM Conference on Electronic Commerce*, pages 181–196.

Christodoulou, G., Kovács, A., Sgouritsa, A., and Tang, B. (2013). Tight bounds for the price of anarchy of simultaneous first price auctions. *arXiv preprint arXiv:1312.2371*.

Clarke, E. H. (1971). Multipart pricing of public goods. *Public choice*, 11(1):17–33.

Cole, R. and Fleischer, L. (2008). Fast-converging tatonnement algorithms for one-time and ongoing market problems. In *STOC '08: Proceedings of the 40th annual ACM symposium on Theory of computing*, pages 315–324, New York, NY, USA. ACM.

Day, R. and Milgrom, P. (2007). Core-selecting package auctions. *International Journal of Game Theory*, 36(3-4):393–407.

Dhangwatnotai, P., Roughgarden, T., and Yan, Q. (2010). Revenue maximization with a single sample. In *ACM Conference on Electronic Commerce*, pages 129–138.

Dughmi, S. and Peres, Y. (2012). Mechanisms for risk averse agents, without loss. *arXiv preprint arXiv:1206.2957*, pages 1–9.

Dütting, P., Fischer, F., and Parkes, D. C. (2014). Expressiveness and robustness of first-price position auctions. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 57–74. ACM.

Edelman, B. and Ostrovsky, M. (2007). Strategic bidder behavior in sponsored search auctions. *Decision Support Systems*, 43(1):192–198.

Edelman, B., Ostrovsky, M., and Schwarz, M. (2007). Internet advertising and the generalized second-price auction: selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259.

Edelman, B. and Schwarz, M. (2010). Optimal auction design and equilibrium selection in sponsored search auctions.

Engelbrecht-Wiggans, R. and Kahn, C. M. (1998). Multi-unit pay-your-bid auctions with variable awards. *Games and Economic Behavior*, 23(1):25 – 42.

Even-dar, E., Mansour, Y., and Nadav, U. (2009). On the convergence of regret minimization dynamics in concave games. In *STOC '09: Proceedings of the 41st annual ACM symposium on Theory of computing*, pages 523–532, New York, NY, USA. ACM.

Fiat, A. and Papadimitriou, C. (2010). When the players are not expectation maximizers. In *Algorithmic Game Theory*, pages 1–14. Springer.

Gomes, R. and Sweeney, K. (2014). Bayes–nash equilibria of the generalized second-price auction. *Games and Economic Behavior*, 86:421–437.

Groves, T. (1973). Incentives in teams. *Econometrica: Journal of the Econometric Society*, pages 617–631.

Harrison, G. W. (1989). Theory and misbehavior of first-price auctions. *The American Economic Review*, 79(4):pp. 749–762.

Hart, S. and Mas-Colell, A. (2000). A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150.

Hart, S. and Nisan, N. (2013). The menu-size complexity of auctions.

Hartline, J., Hoy, D., and Taggart, S. (2014). Price of anarchy for auction revenue. In *ACM EC*, pages 693–710, New York, New York, USA. ACM Press.

Hartline, J. D. (2014). *Mechanism Design and Approximation*.

Hartline, J. D. and Roughgarden, T. (2009). Simple versus optimal mechanisms. In *ACM Conference on Electronic Commerce*, pages 225–234.

Holt, C. A. (1980). Competitive bidding for contracts under alternative auction procedures. *The Journal of Political Economy*, 88(3):433–445.

Hoy, D., Jain, K., and Wilkens, C. A. (2012). Coopetitive ad auctions. *arXiv*, abs/1209.0832.

Hu, A., Matthews, S. A., and Zou, L. (2010). Risk aversion and optimal reserve prices in first- and second-price auctions. *Journal of Economic Theory*, 145(3):1188–1202.

Kaplan, T. R. and Zamir, S. (2012). Asymmetric first-price auctions with uniform distributions: analytic solutions to the general case. *Economic Theory*, 50(2):269–302.

Kirkegaard, R. (2009). Asymmetric first price auctions. *Journal of Economic Theory*, 144(4):1617–1635.

Kirkegaard, R. (2012). A mechanism design approach to ranking asymmetric auctions. *Econometrica*, 80(5):2349–2364.

Klemperer, P. (2002). What really matters in auction design. *The Journal of Economic Perspectives*, 16(1):pp. 169–189.

Koutsoupias, E. and Papadimitriou, C. (1999). Worst-case equilibria. In *STACS*, pages 404–413.

Krishna, V. (2009). *Auction Theory*. Academic Press/Elsevier.

Lebrun, B. (2006). Uniqueness of the equilibrium in first-price auctions. *Games and Economic Behavior*, 55(1):131–151.

Lehmann, D., Ita O'Callaghan, L., and Shoham, Y. (2002). Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM*, 49(5):577–602.

Leme, R. P., Syrgkanis, V., and Tardos, E. (2012). Sequential auctions and externalities. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '12, pages 869–886. SIAM.

Lucier, B. and Borodin, A. (2010). Price of anarchy for greedy auctions. In *ACM-SIAM Symposium on Discrete Algorithms*, pages 537–553.

Lucking-Reiley, D. (2000). Vickrey auctions in practice: from nineteenth-century philately to twenty-first-century e-commerce. *Journal of Economic Perspectives*, 14(3):183–192.

Mas-Colell, A., Whinston, M. D., and Green, J. R. (1995). *Microeconomic theory*. Oxford University Press.

Maskin, E. and Riley, J. (1984). Optimal auctions with risk averse buyers. *Econometrica*, 52(6):1473–1518.

Maskin, E. and Riley, J. (2000). Asymmetric auctions. *The Review of Economic Studies*, 67(3):pp. 413–438.

Maskin, E. and Riley, J. (2003). Uniqueness of equilibrium in sealed high-bid auctions. *Games and Economic Behavior*, 45(2):395 – 409.

Matthews, S. (1983). Selling to risk averse buyers with unobservable tastes. *Journal of Economic Theory*, 30(2):370–400.

Matthews, S. A. (1984). On the implementability of reduced form auctions. *Econometrica*, 52(6):pp. 1519–1522.

Matthews, S. A. (1987). Comparing auctions for risk averse buyers: a buyer's point of view. *Econometrica*, 55(3):633–646.

Milgrom, P. (2000). Putting auction theory to work: the simultaneous ascending auction. *Journal of Political Economy*, 108(2):245–272.

Milgrom, P. (2004). *Putting Auction Theory to Work*. Churchill Lectures in Economics. Cambridge University Press.

Milgrom, P. and Roberts, J. (1991). Adaptive and sophisticated learning in normal form games. *Games and Economic Behavior*, 3(1):82 – 100.

Myerson, R. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73.

Nash, J. (1951). Non-cooperative games. *The Annals of Mathematics*, 54(2):286–295.

Pratt, J. W. (1964). Risk aversion in the small and in the large. *Econometrica: Journal of the Econometric Society*, pages 122–136.

Riley, J. and Samuelson, W. (1981). Optimal auctions. *The American Economic Review*, 71(3):381–392.

Roberts, K. (1979). The characterization of implementable social choice rules. In *Aggretaion and Revelation of Preferences*, J-J.Laffont (ed.), North Holland Publishing Company.

Roughgarden, T. (2009). Intrinsic robustness of the price of anarchy. In *ACM Symposium on Theory of Computing*, pages 513–522.

Roughgarden, T., Talgam-Cohen, I., and Yan, Q. (2012). Supply-limiting mechanisms. In *ACM Conference on Electronic Commerce*, pages 844–861.

Roughgarden, T. and Tardos, E. (2012). Do externalities degrade gsp's efficiency? In *The Eighth Ad Auctions Workshop*.

Samuelson, P. A. (1941). The stability of equilibrium: comparative statics and dynamics. *Econometrica*, 9(2):97–120.

Syrgkanis, V. (2012). Bayesian games and the smoothness framework. *CoRR*, abs/1203.5155.

Syrgkanis, V. and Tardos, E. (2013). Composable and efficient mechanisms. In *ACM Symposium on Theory of Computing*, pages 211–220.

Talwar, K. (2003). The price of truth: Frugality in truthful mechanisms. In *STACS 2003*, pages 608–619. Springer.

Varian, H. R. (2009). Online ad auctions. *The American Economic Review*, pages 430–434.

Varian, H. R. and Harris, C. (2014). The VCG auction in theory and practice. *American Economic Review*, 104(5):442–445.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37.

Von Neumann, J. and Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton University Press.

Walras, L. (1954). *Elements of pure economics*. American Economic Association.

Wilkens, C. A. and Sivan, B. (2012). Single-call mechanisms. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, EC '12, pages 946–963, New York, NY, USA. ACM.

Wilson, R. (1987). Game-theoretic analyses of trading processes. In Bewley, T. F., editor, *Advances in Economic Theory*, pages 33–70. Cambridge University Press. Cambridge Books Online.

APPENDIX A

Revenue Covering

A.1. Single-item First-Price Auction Lowerbounds

We present a lower bound of 1.15 on the single-item, first-price auction Price of Anarchy.

Consider a setting with one high bidder with a fixed value of 10, and n small bidders with values drawn from some distribution with value always less than 10. The welfare-optimal allocation always serves the high bidder. We parameterize the expected utility of the high bidder as u_H. Assume the low bidders will bid such that the highest of their bids is distributed according to the CDF $F_L(b) = u_H/(10-b)$, with a point mass of probability $u_H/10$ at 0. With this distribution, agent H achieves utility u_h for any bid in the range $[0, 10-u_h]$.

The high agent plays a mixed strategy according to the bid CDF $B_H(b) = \sqrt{b/(10-u_H)}$. The competing bid CDF for each low bidder is $F_c(b) = B_H(b) \cdot B_L(b)^{(n-1/n)}$.

With $u_H = 5.7$, solving for the first order conditions in the first price auction tells us that for any low agent bidding b, $v = b + F_c(b)/F_c'(b)$; solved numerically it is approximately $v(b) = \frac{15b-0.5b^2}{5+0.5b}$. Solving numerically gives welfare of 8.69; thus the price of anarchy for welfare is approximately 1.15.

This example is almost tight against the expected cumulative threshold lowerbound $\underline{T}$ used in the proof of the value covering lemma (Lemma 3.4). However, the $\frac{e}{e-1}$ price of anarchy proof ignores the bid from the agent allocated in the optimal allocation and the utility of the agents allocated in FPA but not Opt. Both of these quantities are non-zero, which leads to the 1.15 figure being reasonably far from the $\frac{e}{e-1}$. Bounding these quantities is a likely required step for improving the $\frac{e}{e-1}$ bound for single-item settings.

A.2. Framework Proofs

Lemma 3.29 (Restatement). *Consider a mechanism $\mathcal{M}$ in BNE with induced allocation and payment rules $(\mathbf{x}, \mathbf{p})$, and an agent i with value v_i. For any $x' \in [0,1]$,*

$$v_i x_i(v_i) + T_i \geq \tfrac{e-1}{e} v_i x'. \tag{3.15}$$

Proof of Lemma 3.29. If $x_i(v_i) > x'$, $T_i = 0$ and the result follows. Otherwise, note that by the definition of BNE, i chooses an action which maximizes utility. It follows that

$$u_i(v_i) \geq v_i x_i(\alpha_i(z)) - p_i(\alpha_i(z)) = \left(v_i - \frac{p_i(\alpha_i(z))}{x_i(\alpha_i(z))}\right) x_i(\alpha_i(z)) \geq \left(v_i - \frac{p_i(\alpha_i(z))}{x_i(\alpha_i(z))}\right) z. \tag{A.1}$$

Rearranging (A.1) yields

$$v_i - \frac{u_i(v_i)}{z} \leq \frac{p_i(\alpha_i(z))}{x_i(\alpha_i(z))} = \tau_i(z). \tag{A.2}$$

This bound is meaningful as long as $v_i - \frac{u_i(v_i)}{z} \geq 0$, or alternatively $z \geq u_i(v_i)/v_i$. It follows that

$$\begin{aligned} v_i x_i(v_i) + T_i &\geq v_i x_i(v_i) + \int_{x_i(v_i)}^{x'} \max\left(0, v_i - \frac{u_i(v_i)}{z}\right) dz \\ &\geq v_i x_i(v_i) + \int_{u_i(v_i)/v_i}^{x'} v_i - \frac{u_i(v_i)}{z}\, dz \\ &= v_i x' + u_i(v_i) \ln \frac{u_i(v_i)}{x' v_i}. \end{aligned} \tag{A.3}$$

Holding $x'v_i$ fixed and minimizing this quantity as a function of $u_i(v_i)$ yields a minimum at $u_i(v_i) = \frac{x'v_i}{e}$, and at that point assumes value $(1-1/e)x'v_i$. This is precisely the righthand side of (3.15), implying the lemma. $\square$

A.3. Revenue Extension Proofs

A.3.1. FPA with duplicate bidders

For convenience, let us define shorthand to refer to the virtual surplus provided by agents with positive and negative virtual values respectively.

Definition A.1. For any BNE $\mathbf{s}$ of mechanism $\mathcal{M}$, let $\textsc{Rev}^+(\mathcal{M})$ and $\textsc{Rev}^-(\mathcal{M})$ be the virtual surplus from agents with positive and negative virtual values, respectively. Thus,

$$\textsc{Rev}^+(\mathcal{M}) = \sum_i \mathbf{E}_{v_i}\left[\max(0, \phi_i(v_i))x_i(v_i)\right],$$
$$\textsc{Rev}^-(\mathcal{M}) = \sum_i \mathbf{E}_{v_i}\left[\min(0, \phi_i(v_i))x_i(v_i)\right].$$

Recall that our proof of revenue approximation results for the first price auction with monopoly reserves (Theorem 3.11) only used monopoly reserves to eliminate the impact of agents with negative virtual values. Thus by the same proof we can show that for any first-price auction, the revenue from only positive-virtual valued agents approximates the optimal revenue:

Corollary A.1 (of Theorem 3.11)**.** *For any first price auction,*

$$\frac{2e}{e-1}\textsc{Rev}^+(\text{FPA}) \geq \textsc{Rev}(\textsc{Opt}). \tag{A.4}$$

Now let us prove our desired result.

Theorem 3.12 (Restatement). *The revenue in any BNE of the first price auction with k-duplicates* (FPA_k) *and agents with regularly distribution values is at least a* $\frac{k}{k-1}\frac{2e}{e-1}$ *approximation to the revenue of the optimal auction.*

For shorthand, we will consider a partition of the agents into groups $B_1, B_2, \ldots B_p$ such that each group has size at least k and all agents in each group B_i have values drawn from the distribution F_i.

Next, in any BNE of a first-price auction with k-duplicates, agents of the same group will play symmetric strategies. This follows from Theorem 3.1 of Chawla and Hartline [2013], which gives that any two agents when competing against a reserve distribution will behave identically.

Corollary A.2 (of Theorem 3.1, Chawla and Hartline [2013])**.** *In any BNE of a first-price or all-pay auctions with k-duplicates, for any group* B_j *of agents who have identically distributed values, all agents in the group play by identical strategies everywhere except on a measure zero set of values.*

We now relate the revenue from each group of bidders to the revenue from a symmetric second price auction with reserves among only the bidders within the group of duplicates, allowing us to use the symmetric auction approximation results of Bulow and Klemperer [1996].

Let $\mathrm{SPA}_R(B)$ be a second price auction run among agents in group B with a random reserve drawn according to the distribution R.

Lemma A.3. *There exist reserve value distributions $R_1, R_2 \ldots R_p$ such that in any first price auction* FPA_k *with k-duplicates,*

$$\textsc{Rev}(\mathrm{FPA}_k) = \sum\nolimits_j \textsc{Rev}(\mathrm{SPA}_{R_j}(B_j)), \tag{A.5}$$

$$\textsc{Rev}^+(\mathrm{FPA}_k) = \sum\nolimits_j \textsc{Rev}^+(\mathrm{SPA}_{R_j}(B_j)). \tag{A.6}$$

Proof. Fixing the values and actions of bidders outside a group j results in a threshold value for the highest bidding member of a group, above which he bids high enough to be allocated and below which he bids less than the amount necessary to be allocated. Let the distribution of such thresholds be R_j; then a second price auction among the group members with reserve drawn precisely from R_j will induce exactly the same allocation rule for all members of the group. By revenue equivalence (Theorem 3.1), the revenue from members of group j in FPA_k will be the same as $\textsc{Rev}(\mathrm{SPA}_{R_j}(B_j))$. The same argument holds for $\textsc{Rev}^+(\mathrm{FPA}_k)$ and $\textsc{Rev}^+(\mathrm{SPA}_{R_j}(B_j))$. $\square$

A second-price auction within a group is now a symmetric setting, and thus we can now use the work of Bulow and Klemperer [1996] to relate (A.5) and (A.6). By Bulow and Klemperer [1996], if $k \geq 2$, $\textsc{Rev}(\mathrm{SPA}_{R_j}(B_j)) \geq \frac{k-1}{k}\textsc{Rev}^+(\mathrm{SPA}_{R_j}(B_j))$ and hence:

$$\begin{aligned}
\textsc{Rev}(\mathrm{FPA}) &= \sum\nolimits_j \textsc{Rev}(\mathrm{SPA}_{R_j}(B_j)) \\
&\geq \sum\nolimits_j \frac{k-1}{k}\textsc{Rev}^+(\mathrm{SPA}_{R_j}(B_j)) \\
&= \frac{k-1}{k}\textsc{Rev}^+(\mathrm{FPA}).
\end{aligned}$$

Combining with Corollary A.1 gives our desired result,

$$\frac{k}{k-1}\frac{2e}{e-1}\textsc{Rev}(\text{FPA}) \geq \frac{2e}{e-1}\textsc{Rev}^{+}(\text{FPA}) \geq \textsc{Rev}(\textsc{Opt}). \tag{A.7}$$

A.4. Revenue Covering Proofs

A.4.1. GFP

In a randomized mechanism like a position auction, fixing the actions of other agents results not in a single threshold for winning but a number of thresholds for different allocation amounts corresponding with the value of each slot. See Figure A.1a for an illustration.

As a result, the generalized first price position auction will not satisfy the same "complete" revenue covering that the normal first-price auction satisfies, but rather it will satisfy a different version which will be sufficient to prove the same approximation results[1]. Let $T_i(x'_i)$ be the expected threshold up to x'_i.

Notably, for any alternate allocation $\mathbf{x}'$, the revenue will cover for each agent the threshold up to their alternate allocation amount $\mathbf{x}'$. See Figure A.1b for an illustration.

Definition A.2. Position auction $\mathcal{A}$ is μ-revenue covered if for any feasible allocation $\mathbf{x}'$,

$$\mu\textsc{Rev}(\mathcal{M}) \geq \sum_i T(x'_i) \tag{A.8}$$

Lemma A.4. *The generalized first-price position auction is* 1*-revenue covered.*

The proof comes in two steps: first we show that GFP is revenue covered when bidders play deterministic strategies; next, that deterministic revenue-covering implies general revenue-covering.

[1]This version of revenue covering can also be used for all of the results for pay-your-bid mechanisms in the rest of Section 3.5

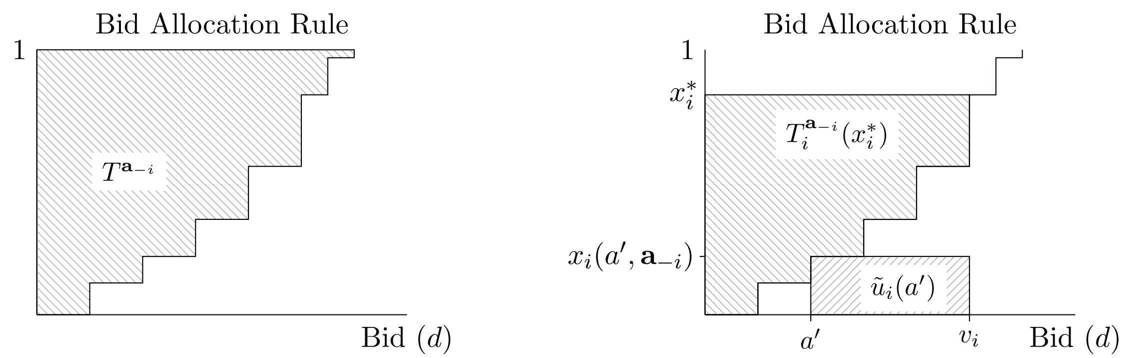

(a) In GFP, when other agents play actions a, there is not just one threshold for allocation, but a threshold for allocation in each slot.

(b) For revenue covering in position auctions, $T_i^{\mathbf{a}_{-i}}(x_i')$ is used in place of $x_i' T_i^{\mathbf{a}_{-i}}$. Note that by the convexity of $T_i^{\mathbf{a}_{-i}}(x_i')$ in x_i', $T_i^{\mathbf{a}_{-i}}(x_i') \leq x_i' T_i^{\mathbf{a}_{-i}}$.

Figure A.1. The threshold in the GFP auction.

Proposition A.5. *The generalized first-price position auction is* 1*-revenue covered in deterministic strategies.*

Proof. Consider the bid-based allocation rule of an agent in GFP, $\tilde{x}_i(b_i, \mathbf{b}_{-i})$. For any bid b_i, if b_i would be the jth highest bid, then $\tilde{x}_i(b_i, \mathbf{b}_{-i})$ is the position weight of slot j. So, $\tilde{x}_i(b_i, \mathbf{b}_{-i})$ is a stair function, with a stair corresponding to each position.

Let $T_i^{\mathbf{a}}(x')$ denote the expected threshold up to x' when agents play actions $\mathbf{a}$. Denote by b^j and b_{-i}^j the jth highest bid from all bidders including and excluding i, respectively; then

$$T_i^{\mathbf{a}}(\alpha_j) = \sum_{i=j}^{m} (\alpha_i - \alpha_{i+1}) b_{-i}^j. \tag{A.9}$$

The revenue in GFP is then $\textsc{Rev}(\mathcal{M}(\mathbf{a})) = \sum_j \alpha_j b^j \geq \sum_j \alpha_j b_{-i}^j$. For any slot j, the threshold amount for the bidder allocated j in the alternate allocation is less than payment of the bidder who won the slot j: $\alpha_j b^j \geq \sum_{i=j}^{m}(\alpha_i - \alpha_{i+1}) b^j_{-i}$. Summing over all bidders gives $\textsc{Rev}(\mathcal{M}(\mathbf{a})) \geq \sum_i T_i^{\mathbf{a}}(x'_i)$, our desired result. $\square$

We now show that GFP is revenue-covered in the general Bayesian setting.

Proof of Lemma A.4. First, by linearity of expectations we can view the expected threshold up to $\tilde{x}_i(b_i)$ of an agent as the expectation of their full-information thresholds, hence

$$\begin{aligned}
T_i(\tilde{x}_i(b_i)) &= b_i \tilde{x}_i(b_i) - \int_0^{b_i} \tilde{x}_i(d)\ dd \\
&= \mathbf{E}_{\mathbf{v}_{-i}}\left[b_i \tilde{x}_i(b_i, s_i(\mathbf{v}_{-i})) - \int_0^{b_i} \tilde{x}_i(d, s_i(\mathbf{v}_{-i}))\ dd \right] \\
&= \mathbf{E}_{\mathbf{v}_{-i}}\left[T_i^{\mathbf{s}_{-i}(\mathbf{v}_{-i})}(\tilde{x}_i(b_i, \mathbf{s}_{-i}(\mathbf{v}_{-i}))\right]. \qquad \text{(A.10)}
\end{aligned}$$

Note that in Equation A.10 the bids are fixed, not the allocation probabilities. Keeping bids fixed is actually the expected-threshold minimizing way to get allocation $\tilde{x}_i(b_i)$ even if the bidder is allowed to change her bid in reaction to the actions of other agents. Formally,

$$\begin{aligned}
&\mathbf{E}_{\mathbf{v}_{-i}}\left[T_i^{\mathbf{s}_{-i}(\mathbf{v}_{-i})}(\tilde{x}_i(b_i, \mathbf{s}_{-i}(\mathbf{v}_{-i})))\right] \\
&\qquad = \min_{\pi_i(\cdot)\ \text{s.t.}\ \tilde{x}_i(\pi_i) = \tilde{x}_i(b_i)} \mathbf{E}_{\mathbf{v}_{-i}}\left[T_i^{\mathbf{s}_{-i}(\mathbf{v}_{-i})}(\tilde{x}_i(\pi_i(\mathbf{s}_{-i}(\mathbf{v}_{-i})), \mathbf{s}_{-i}(\mathbf{v}_{-i})))\right]. \qquad \text{(A.11)}
\end{aligned}$$

By first-order conditions, the threshold-minimizing bidding policy to get allocation $\tilde{x}_i(b_i)$ comes from equating the marginals $\frac{\partial}{\partial z} T_i^{\mathbf{s}_{-i}(\mathbf{v}_{-i})}(z) = \tau(z)$, which is the bid required to get allocation z. Thus, bidding b_i is the expected-threshold minimizing way to get allocation $\tilde{x}_i(b_i)$.

One alternative bidding policy is to bid to get a fixed probability amount. By the optimality of fixed bids, we have

$$\mathbf{E}_{\mathbf{v}_{-i}}\left[T_i^{\mathbf{s}_{-i}(\mathbf{v}_{-i})}(\tilde{x}_i(b_i, \mathbf{s}_{-i}(\mathbf{v}_{-i})))\right] \le \mathbf{E}_{\mathbf{v}_{-i}}\left[T_i^{\mathbf{s}_{-i}(\mathbf{v}_{-i})}(\tilde{x}_i(b_i))\right]. \tag{A.12}$$

Combining Equations A.10 and A.12 with Proposition A.5 then gives our desired result,

$$\begin{aligned} T_i(\tilde{x}_i(b_i)) &\le \mathbf{E}_{\mathbf{v}_{-i}}\left[T_i^{\mathbf{s}_{-i}(\mathbf{v}_{-i})}(\tilde{x}_i(b_i))\right] \\ &\le \mathbf{E}_{\mathbf{v}_{-i}}\left[\mu \textsc{Rev}(\mathcal{M})\right] \\ &\le \mu \textsc{Rev}(\mathcal{M}). \end{aligned}$$

□

The position auction variant of revenue-covering is weaker than the general revenue covering (Definition 3.1) condition, and so we also need a stronger value-covering condition.

Lemma A.6 (Value Covering (Position Auctions)). *In any BNE of* Gfp*, for any bidder* i *with value* v_i*,*

$$u_i(v_i) + T_i \ge \tfrac{e-1}{e} x_i' v_i. \tag{A.13}$$

The proof is omitted; it proceeds almost identically to Lemma 3.13 but with T_i in place of T_i and $x_i' v_i$ in place of v_i, and the worst case u_i is found at $x_i' v_i / e$, not v_i/e. Combining revenue and value covering then gives a welfare approximation:

Theorem A.7. *The welfare of any BNE of* Gfp *is an* $\frac{e}{e-1}$*-approximation to the welfare of the optimal auction.*

Using an extension from values to virtual values (as in Lemma 3.6) and using reserve amounts to cover the additional threshold generated by the reserve just gives a revenue approximation result with monopoly reserves:

Theorem A.8. *The revenue in any BNE of the* GFP *with regular agents and monopoly reserves is a $\frac{2e}{e-1}$-approximation to the revenue of the optimal auction.*

A.5. All-Pay Auctions

A.5.1. All-Pay Matroid Auction

Lemma 3.31 (Restatement). *The all-pay matroid auction is 2-revenue covered.*

Proof. The proof for matroid environments differs from the single-item proof only in the use of Lemma 3.20 to relate revenue to threshold bids. Note that the all-pay matroid auction selects a basis maximizing the sum of the bids, so Lemma 3.20 still holds. Let $\mathbf{x}'$ be a feasible allocation and $\mathbf{a}$ be an action profile. By the payment semantics of the mechanism,

$$\text{REV}(\mathcal{M}) = \mathbf{E}_v\left[\sum\nolimits_i s_i(v_i)\right] \geq \mathbf{E}_v\left[\sum\nolimits_i s_i(v_i)x_i(v)\right].$$

Now let $\tau_i^b(\mathbf{v}_{-i})$ be the threshold bid for i in realized value profile $\mathbf{v}_{-i}$ under strategy profile $\mathbf{s}$ (without index i). By Lemma 3.20 implies that

$$\mathbf{E}_v\left[\sum\nolimits_i s_i(v_i)x_i(v)\right] \geq \mathbf{E}_v\left[\sum\nolimits_i \tau_i^b(\mathbf{v}_{-i})x_i'\right] = \sum\nolimits_i \mathbf{E}_v\left[\tau_i^b(\mathbf{v}_{-i})\right]x_i'. \tag{A.14}$$

The rest of the proof proceeds exactly the same as the proof for the single-item case: translating between threshold bids and equivalent threshold bids (losing the factor of 2), and finally summing over all agents to achieve the desired result. ∎

A.5.2. All-Pay Native Framework

In the proof of Lemma 3.31, we lost a factor of 2 because we needed to translate all-pay bids into their first-price equivalents. By switching from a first-price centric version of the value covering and revenue covering framework to one which works directly in terms of all-pay bids, we can match the welfare results of [Syrgkanis and Tardos, 2013]. We can also derive a tighter revenue result with duplicates than was possible in the old framework.

First, for each allocation probability z, let $\tilde{\tau}_i(z)$ be the lowest bid i needs to place to get allocated with probability at least z. Formally, $\tilde{\tau}_i(z) = \min\{b \,|\, \tilde{x}_i(b) \geq z\}$. As in the first-price auction, we can compute the expected value of this threshold bid as $\tilde{T}_i = \int_0^1 \tilde{\tau}_i(z)\, dz$. For the first-price auction, we used the payment semantics to derive a distribution of threshold bids for which i would be indifferent between all bids less than v_i. We can do the same thing for the all-pay auction and get the following result:

Lemma A.9 (All-Pay Value Covering)**.** *For any BNE of an all-pay auction and agent i with value v_i,*

$$u_i(v_i) + \tilde{T}_i \geq \frac{v_i}{2}.$$

Proof. The proof parallels that of Lemma 3.4 - we lower bound $\tilde{T}_i$ using the payment semantics of the all-pay auction, then minimize the lower bound. As with Lemma 3.4, the deviation-based approach of Syrgkanis and Tardos [2013] also suffices.

Lowerbounding $\tilde{T}_i$. Bidder i chooses a best response bid b_i which maximizes her utility, $\tilde{u}_i(b_i) = v_i\tilde{x}_i(b_i) - b_i$. It follows that for any other deviation bid d, $\tilde{u}_i(b_i) \geq v_i\tilde{x}_i(d) - d$. Rearranging, we get $\tilde{x}_i(d) \leq \frac{\tilde{u}_i(b_i)+d}{v_i}$. Since $\tilde{x}_i(d)$ is the CDF of i's threshold bid, we can lower bound $\tilde{T}_i$ by integrating above the curve $\frac{\tilde{u}_i(b_i)+d}{v_i}$. In other words: $\tilde{T}_i \geq \int_0^1 \max(0, v_i z - \tilde{u}_i(b_i))\, dz$. Call the latter quantity $\underline{\tilde{T}_i}$.

Optimizing $\underline{\tilde{T}_i}$. Evaluating the integral for $\underline{\tilde{T}_i}$ gives $\underline{\tilde{T}_i} = v_i/2 - \tilde{u}_i(b_i) + \tilde{u}_i(b_i)^2/2v_i$, hence $u_i + \tilde{u}_i(b_i) = v_i/2 + \tilde{u}_i(b_i)^2/2v_i$. Holding v_i fixed and minimizing with respect to $\tilde{u}_i(b_i)$ yields a minimum at $u_i(b_i) = 0$, hence $\tilde{u}_i(b_i) + \underline{\tilde{T}_i} \geq v_i/2$. Using the facts that $\tilde{u}_i(b_i) = u_i(v_i)$ and $\tilde{T}_i \geq \underline{\tilde{T}_i}$ yields the result. $\square$

As in the original framework, value covering characterizes the tradeoff between an agent's utility and the difficulty they face getting allocated. Now, however, the latter quantity is represented by $\tilde{T}_i$, which comes from all-pay rather than equivalent first-price bids. In proving revenue covering, we can therefore skip the translation from all-pay bids to equivalent first-price bids, yielding revenue covering with $\mu = 1$:

Lemma A.10 (All-Pay Revenue Covering)**.** *For any BNE of the all-pay action and any agent i, the expected revenue is at least $\tilde{T}_i$.*

Proof. The revenue of the all-pay auction is expected sum of all bids. This is at least the expected highest bid from all agents except i, which is exactly $\tilde{T}_i$. $\square$

Combining Lemmas A.9 and A.10, summing over all agents, and taking expectations in the manner used to prove Theorem 3.3 yields the welfare bound of Syrgkanis and Tardos [2013]:

Theorem A.11. *The welfare in any BNE of the all-pay auction is at least a 2-approximation to the welfare of the welfare optimal mechanism.*

Furthermore, from Lemma A.9, we can derive a virtual value covering result:

Lemma A.12 (All-Pay Virtual Value Covering)**.** *For any BNE of the all-pay auction, any agent i with value v_i such that $\phi_i(v_i) \geq 0$,*

$$\phi_i(v_i)x_i(v_i) + \tilde{T}_i \geq \frac{\phi_i(v_i)}{2}.$$

For revenue, combining this with the duplicates results in Section 3.4.3 yields:

Theorem A.13. *The revenue in any BNE of the all-pay auction with at least k bidders from each distribution is at least a 6-aproximation to the revenue of the optimal mechanism.*

Finally, note that all of the above can be extended to matroids using Lemma 3.20.

APPENDIX B

Price of Anarchy from Data

Lemma 4.1 (Restatement). *For any bidder i with value v_i and allocation amount x_i',*

$$u_i(v_i) + \frac{1}{\mu} T_i(x_i') \geq \frac{1 - e^{-\mu}}{\mu} x_i' v_i.$$

(Proof sketch). The proof proceeds analogously to the proof of value covering in Chapter 3 with a small optimization to improve the bound. We first defining a lower bound $\underline{T}(x) = \int_0^x \underline{\tau}(z)\ dz$ s.t. $\underline{\tau}(z) \leq \tau(z)$ and hence $\underline{T}(x) \leq T_i(x)$.

$$\begin{aligned}
\underline{T}(x_i') &= \int_0^{x_i'} \underline{\tau}(z)\ dz \\
&= \int_0^{x_i'} \max(0, v - u_i(v_i)/z)\ dz
\end{aligned}$$

Evaluating the integral gives $\underline{T}(x_i') = (v_i x_i' - u_i(v_i)) - u_i(v_i)\left(\log x_i' - \log \frac{u_i(v_i)}{v_i}\right)$, thus

$$u_i(v_i) + \frac{1}{\mu}\underline{T}(x_i') = u_i(v_i) + \frac{1}{\mu}\left(v_i x_i' - u_i(v_i)\left(1 + \log x_i' - \log \frac{u_i(v_i)}{v_i}\right)\right)$$

and

$$\frac{u_i(v_i) + \frac{1}{\mu}\underline{T}(x_i')}{v_i} = \frac{u_i(v_i)}{v_i} + \frac{1}{\mu}\left(x_i' - \frac{u_i(v_i)}{v_i}\left(1 + \log x_i' - \log \frac{u_i(v_i)}{v_i}\right)\right) \tag{B.1}$$

The right side of Equation (B.1) is convex in $\frac{u_i(v_i)}{v_i}$, so we can minimize it by taking first-order conditions in $\frac{u_i}{v_i}$, giving

$$0 = 1 - \frac{1}{\mu}\left(\log x_i' - \log\frac{u_i(v_i)}{v_i}\right).$$

Thus the right side of Equation (B.1) is minimized with $u_i(v_i)/v_i = x_i' e^{-\mu}$, giving our desired result,

$$\frac{u_i(v_i) + \frac{1}{\mu}\underline{T}(x_i')}{v_i} \geq \frac{1-e^{-\mu}}{\mu} x_i'.$$

$\square$

Lemma 4.7 (Restatement). *For any μ-revenue covered mechanism $\mathcal{M}$ and strategy profile* $\mathbf{s}$ *with $\mu \geq 1$, if $\tau(\epsilon) \geq (1-1/k)\tau(\mathbf{x}')$ for any feasible allocation amount $\mathbf{x}'$ and $\epsilon > 0$, $\mathcal{M}$ and* $\mathbf{s}$ *are empirically $(1-1/k)$-value covered.*

(Proof sketch). First, for bidders with values $v_i < \tau(1)$, the bound holds even without the u_i term, as

$$T_i(x_i') = \int_0^{x_i'} \tau(z)\ dz \tag{B.2}$$

$$\geq \int_0^{x_i'} \tau(0) \tag{B.3}$$

$$\geq x_i'(1-1/k)\tau(1) \tag{B.4}$$

$$\geq x_i'(1-1/k)v_i \tag{B.5}$$

Consider bidders with values $v_i \geq \tau(1)$. As such a bidder can always choose the bid with price-per-click $\tau(1)$ and get utility $v_i - \tau(1)$, we know $u_i(v_i) \geq v_i - \tau(1)$. For any allocation they choose, we then have

$$u_i(v_i) + \frac{1}{\mu}T_i(x'_i) \geq v_i - \tau(1) + \frac{1}{\mu}\int_0^{x'_i} \tau(z)\ dz \tag{B.6}$$
$$\geq v_i - \tau(1) + \frac{1}{\mu}x'_i(1 - 1/k)\tau(1) \tag{B.7}$$
$$\geq v_i - \tau(1) + \frac{x'_i(1 - 1/k)\tau(1)}{\max(1, \mu)} \tag{B.8}$$
$$\geq (v_i - \tau)\left(1 - \frac{x'_i(1 - 1/k)}{\max(1, \mu)}\right) + \frac{x'_i(1 - 1/k)v_i}{\max(1, \mu)} \tag{B.9}$$
$$\geq \frac{x'_i(1 - 1/k)v_i}{\max(1, \mu)} \tag{B.10}$$

We can improve on the bound by considering the worst-case price-per-click allocation rule that satisfies $\tau(1) = 1$ and $\tau(0) = 1 - \frac{1}{k}$, much like in the proof of value covering.

The worst case price-per-click allocation rule $\tilde{x}$, for agents with value $v = u + 1$ is

$$\tilde{x}(z) = \begin{cases} 1 & \text{if } 1 \leq z \\ \frac{u}{v-z} & \text{if } 1 - \frac{1}{k} \leq z \leq 1 \\ 0 & \text{if } z \leq 1 - \frac{1}{k} \end{cases} . \tag{B.11}$$

Note that this is exactly the price-per-click allocation rule that results in the bidder being indifferent over all bids in $[1 - \frac{1}{k}, 1]$, as opposed to the indifference over $[0, 1]$ for the normal value covering proof (with a little more normalization).

We can again define $\underline{T}(x_i')$ to be the threshold based on $\tilde{x}$. We will solve numerically for the case that $x_i' = 1$ as every other case is strictly worse. So,

$$\underline{T}(1) = \int_0^1 \tau(z)\ dz \tag{B.12}$$

$$= 1 - \int_{1-1/k}^1 \tilde{x}(y)\ d\,y \tag{B.13}$$

$$= 1 - \int_{1-1/k}^1 \frac{u}{v-y}\ d\,y \tag{B.14}$$

$$= 1 + u\left(\log(v-1) - \log(v - (1 - \frac{1}{k}))\right) \tag{B.15}$$

$$= 1 + u \log \frac{v-1}{v - (1-\frac{1}{k})} \tag{B.16}$$

Thus, $u_i + \frac{1}{\mu}\underline{T}(1) = u + \frac{1}{\mu}\left(1 + u \log \frac{v-1}{v-(1-\frac{1}{k})}\right)$, and

$$\frac{v}{u_i + \frac{1}{\mu}\underline{T}(1)} = \frac{v}{u + \frac{1}{\mu}\left(1 + u\log\frac{v-1}{v-(1-\frac{1}{k})}\right)} \tag{B.17}$$

In the worst case, $u = v - 1$, so

$$\frac{v}{u_i + \frac{1}{\mu}\underline{T}(1)} = \frac{v}{v - 1 + \frac{1}{\mu}\left(1 + (v-1)\log\frac{v-1}{v-(1-\frac{1}{k})}\right)} \tag{B.18}$$

Numerically minimizing for a variety of μ and k values give the results in Table 4.1.

$\square$

APPENDIX C

The Utility Target Auction

C.1. Selected Proofs

This appendix contains selected proofs that were previously omitted.

C.1.1. CEF Claims

Proof of Claim 6.4: We want to show that $\sum_{i\in[n]} v_i(o^*) \geq \sum_{i\in[n]} v_i(o)$ for any outcome o and CEF equilibrium o^*. The envy-freeness constraints give

$$\sum_{i\in[n]} \max\left((v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)), 0\right) \leq \sum_{i\in[n]} b_i(o^*) - b_i(o)$$

$$\sum_{i\in[n]} (v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)) \leq \sum_{i\in[n]} b_i(o^*) - b_i(o)$$

$$\sum_{i\in[n]} v_i(o) - v_i(o^*) \leq 0$$

as desired. □

Proof of Claim 6.5: Define o^i as the outcome that maximizes the welfare of bidders except i:

$$o^i = \arg\max_o \sum_{j\neq i} v_j(o) \quad .$$

Thus, the VCG price of agent i is $\sum_{j\neq i} v_j(o^i) - v_j(o^*)$.

Now, the envy-freeness constraints give

$$
\begin{aligned}
b_i(o^*) \geq & \sum_j \max\big((v_j(o^i) - b_j(o^i)) - (v_j(o^*) - b_j(o^*)), 0\big) \\
& + b_i(o^*) - \sum_j (b_j(o^*) - b_j(o^i)) \\
b_i(o^*) \geq & \sum_{j \neq i} (v_j(o^i) - b_j(o^i)) - (v_j(o^*) - b_j(o^*)) \\
& - \sum_{j \neq i} (b_j(o^*) - b_j(o^i)) + b_i(o^i) \\
& + \max\big((v_i(o^i) - b_i(o^i)) - (v_i(o^*) - b_i(o^*)), 0\big) \\
b_i(o^*) \geq & \sum_{j \neq i} (v_j(o^i) - v_j(o^*)) + b_i(o^i) \\
& + \max\big((v_i(o^i) - b_i(o^i)) - (v_i(o^*) - b_i(o^*)), 0\big) \\
b_i(o^*) \geq & \sum_{j \neq i} (v_j(o^i) - v_j(o^*)) \ .
\end{aligned}
$$

□

Proof of Claim 6.6: We want to show that

$$
\sum_{i \in [n]} b_i(o^*) \geq \max_o \sum_{i \in [n]} \max(v_i(o) - v_i(o^*), 0) \ .
$$

For any outcome o, the envy-freeness constraints give

$$\sum_{i\in[n]} \max\left((v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)), 0\right) \leq \sum_{i\in[n]} b_i(o^*) - b_i(o)$$

$$\sum_{i\in[n]} \max\left(v_i(o) - v_i(o^*) + b_i(o^*), b_i(o)\right) \leq \sum_{i\in[n]} b_i(o^*)$$

$$\sum_{i\in[n]} \max\left(v_i(o) - v_i(o^*), 0\right) \leq \sum_{i\in[n]} \sum_{i\in[n]} b_i(o^*)$$

as desired. $\square$

C.1.2. Convergence Lemmas

Proof of Lemma 6.16: Let b be the bids at ω and b^δ be the bids at $\omega - \delta$. Note that Claim 6.4 implies a welfare-optimal outcome o^* is winning at ω.

First, suppose that all bidders for whom $\delta_i > 0$ are winners at ω. In this case, $v_i(o^*) \geq \omega_i$ and so $v_i(o^*) \geq \omega_i - \delta_i$ and for any outcome o we get

$$\begin{aligned}\sum_{i\in[n]} b_i^\delta(o^*) &= \sum_{i\in[n]} b_i(o^*) + \delta_i \\ &\geq \sum_{i\in[n]} b_i(o) + \delta_i \\ &\geq \sum_{i\in[n]} b_i^\delta(o) \ ,\end{aligned}$$

implying o^* is still winning at b_i^δ. Since $v_i(o^*) \geq \omega_i - \delta_i$, we can conclude that all bidders are winners, ergo $\omega - \delta \in \mathcal{C}$ by Lemma 6.13.

Now, suppose some bidders in ω may be losers, but that the vector δ has the following property:

$$\delta_i \leq \max(\omega_i - v_i(o*), 0) \ .$$

This condition says that only losers will raise their bids, and they will not raise them enough to affect $b_i(o^*)$.

Our goal is to show

$$\sum_{i \in [n]} b_i^\delta(o^*) - b_i^\delta(o) \geq \sum_{i \in [n]} \max\left((v_i(o) - b_i^\delta(o)) - (v_i(o^*) - b_i^\delta(o^*)), 0\right) \ .$$

First, we see that $b_i^\delta(o) = b_i(o)$ as long as $v_i(o) \leq v_i(o^*)$. For any bidder i we have

$$b_i^\delta(o) = \max(v_i(o) - \omega_i - \delta_i, 0)$$

which can only be nonzero if $v_i(o) > \omega_i$. However, $b_i(o)$ can only change if $\delta_i > 0$, which requires $\omega_i > v_i(o^*)$ and thus $v_i(o) > \omega_i > v_i(o^*)$. By construction, this also holds for $\omega_i - \delta_i$:

$$v_i(o) > \omega_i - \delta_i \geq v_i(o^*) \ .$$

Now, when $v_i(o) > \omega_i - \delta_i \geq v_i(o^*)$, we have

$$\begin{aligned}
(v_i(o) - b_i^\delta(o)) - (v_i(o^*) - b_i^\delta(o^*)) &= \min(v_i(o), \omega_i - \delta_i) - \min(v_i(o^*), \omega_i - \delta_i) \\
&= \omega_i - \delta_i - v_i(o^*) \\
&\geq 0 \ .
\end{aligned}$$

Importantly, if $\Delta(o)$ is the set of bidders for which $b_i^\delta(o) \neq b_i(o)$, we may conclude that

$$\sum_{i\in[n]} \max\big((v_i(o) - b_i^\delta(o)) - (v_i(o^*) - b_i^\delta(o^*)), 0\big) =$$

$$= \sum_{i\notin\Delta(o)} \max\big((v_i(o) - b_i^\delta(o)) - (v_i(o^*) - b_i^\delta(o^*)), 0\big)$$
$$+ \sum_{i\in\Delta(o)} (v_i(o) - b_i^\delta(o)) - (v_i(o^*) - b_i^\delta(o^*))$$

and likewise

$$\sum_{i\in[n]} \max((v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)), 0) =$$

$$= \sum_{i\notin\Delta(o)} \max((v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)), 0)$$
$$+ \sum_{i\in\Delta(o)} (v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)) \ .$$

The desired CEF condition quickly follows, using the fact that bidders $i \notin \Delta(o)$ did not change their bids:

$$\sum_{i\in[n]} \max\big((v_i(o) - b_i^\delta(o)) - (v_i(o^*) - b_i^\delta(o^*)), 0\big) =$$

$$= \sum_{i \notin \Delta(o)} \max\left((v_i(o) - b_i^\delta(o)) - (v_i(o^*) - b_i^\delta(o^*)), 0\right) + \sum_{i \in \Delta(o)} (v_i(o) - b_i^\delta(o)) - (v_i(o^*) - b_i^\delta(o^*))$$

$$= \sum_{i \notin \Delta(o)} \max((v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)), 0) + \sum_{i \in \Delta(o)} (v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*))$$

$$+ \sum_{i \in \Delta(o)} (b_i(o) - b_i^\delta(o)) - (b_i(o^*) - b_i^\delta(o^*))$$

$$= \sum_{i \in [n]} \max((v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)), 0) + \sum_{i \in [n]} (b_i(o) - b_i^\delta(o)) - (b_i(o^*) - b_i^\delta(o^*))$$

$$\leq \sum_{i \in [n]} (b_i(o^*) - b_i(o) + \sum_{i \in [n]} (b_i(o) - b_i^\delta(o)) - (b_i(o^*) - b_i^\delta(o^*))$$

$$= \sum_{i \in [n]} b_i^\delta(o^*) - b_i^\delta(o)$$

as desired.

Finally, for general δ, split it as $\delta = \delta^1 + \delta^2$ where

$$\delta_i^1 = \min(\delta_i, \omega_i - v_i(o^*)) \ .$$

The vector δ^1 satisfies the condition $\delta_i \leq \max(\omega_i - v_i(o*), 0)$, so $\omega - \delta^1 \in \mathcal{C}$. Moreover, all bidders are winners in $\omega - \delta^1$, so

$$s - \delta^1 - \delta^2 = s - \delta \in \mathcal{C}$$

as desired. □

Proof of Lemma 6.18: Properties A1 and A3 imply that a bid will only be lowered if there are no losers. Thus, bids will only be raised (utility-targets decreased) until all bidders are simultaneously winners. Since any bidder i is always a winner when bidding $\omega_i = 0$ and bidders never decrease their utility-targets when they are winners (A1), utility-target can be decreased at most $\left|\left\lceil \frac{1}{\epsilon} s \right\rceil\right|_1$ times before all bidders are winners. Moreover, since losers will always try to decrease their utility-targets (A1), the auction will never stall in a configuration where some bidder is a loser. $\square$

Proof of Lemma 6.19: If the auction reaches a vector ω that induces an outcome $o \neq o^*$, then $\omega \in \overline{\mathcal{C}}$ and we are done. Thus, it remains to show that an auction will reach a vector $\omega \in \overline{\mathcal{C}}$ even if the outcome is always o^*.

Consider a bidder j. By A1 we know that j will only decrease ω_j if she is a loser and increase ω_j if she is a winner. By A3 we can conclude that j will eventually decrease her bid until $\omega_j \geq v_j(o^*) - \epsilon$, implying $b_j(o^*) \leq \epsilon$. Thus,

$$\sum_{i \in [n]} b_i(o^*) - b_i(o) \leq n\epsilon \ .$$

Now, as long as there is some outcome o and bidder j such that $v_j(o) > v_j(o^*)$, when ϵ is sufficiently small it will be the case that

$$n\epsilon < \sum_{i \in [n]} \max\left((v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)), 0\right)$$

Unfortunately, this implies

$$\sum_{i \in [n]} \max\left((v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)), 0\right) > \sum_{i \in [n]} b_i(o^*) - b_i(o) \ ,$$

and therefore $\omega \in \overline{\mathcal{C}}$.

If there is no outcome o and bidder j such that $v_j(o) > v_j(o^*)$, then by similar logic bidders will increase their utility-targets until precisely $\omega_j = v_j(o^*)$ (as a result of our restriction that bidders always bid $\omega_j \leq \sup_o v_j(o^*)$). $\square$

Theorem 6.10 (Restatement). *If winners try to lower their effective bids (A3) and losers try raising but not lowering (A1), bids will eventually remain close to the boundary of the set of CEF bids or entirely outside it.*

Proof of Theorem 6.10: By Lemma 6.19, the auction will eventually reach a utility-target vector in $\overline{\mathcal{C}}$ or the degenerate case where nobody is paying anything and o^* is winning. In the degenerate case, bids converge to a point on the boundary of $\mathcal{C}$, so the theorem is true. For the standard case, we show that $\omega \in \overline{\mathcal{C}}_\epsilon$ from the first time a bid in $\overline{\mathcal{C}}$ is reached.

If $\omega \in \overline{\mathcal{C}}$, we are done, so suppose $\omega \in \mathcal{C}$. Let ω' be the most recent utility-targets that were in $\overline{\mathcal{C}}$ and let $\omega'' \in \mathcal{C}$ be the utility-targets immediately after ω'. Let i be the bidder who changed her bid between ω' and ω''. Corollary 6.17 implies that i must have raised her bid between ω' and ω''.

First, suppose that the outcome changed from o' to o'' when i raised her bid. Since o^* must be the outcome of any CEF bid, we know that $o'' = o^*$ and that the outcome does not change again before bids reach ω. Define the utility-target vector $\tilde{s}$ with associated bids $\tilde{b}$

as follows:

$$\tilde{\omega}_j = \begin{cases} \min(\omega_i, \omega'_i) & j = i \\ \omega_j - \epsilon & \omega_j > \omega''_j \\ \omega_j + \epsilon & \omega_j < \omega'_j \\ \omega_j & \text{otherwise.} \end{cases}$$

Let $\delta_j = \tilde{\omega}_j - \omega''_j$. We argue later that i will not increase her utility-target from $\omega''_i = \omega'_i + \epsilon$, so $|\tilde{\omega}_j - \omega_j| \leq \epsilon$ for all j. Thus, it is sufficient to show that $\tilde{s} \in \overline{\mathcal{C}}$.

Consider a bidder $j \neq i$ and suppose $\omega_j > \omega''_j$. By definition, we get a simple bound on j's bid for o':

$$\tilde{b}_j(o') \geq b'_j(o') - \delta_j \ \ .$$

We also know that j lowered her bid at some point between ω'' and ω. Since j would only increase her utility-target if she were a winner, she must have been a winner at some value $\geq \omega_j - \epsilon = \tilde{\omega}_j$. Thus, $\tilde{\omega}_j = \omega_j - \epsilon \leq v_j(o^*)$. We can thus upper-bound her bid for o^*:

$$\tilde{b}_j(o^*) \leq b'_j(o^*) - \delta_j \ \ .$$

Combining these two bounds and noting that $b''_j = b'_j$ for $j \neq i$ gives

$$\tilde{b}_j(o') - \tilde{b}_j(o^*) \geq b'_j(o') - b'_j(o^*) \ \ .$$

For bidders $j \neq i$ with $\omega_j < \omega''_j$, analogous reasoning based the fact that j must have been a loser to decrease her utility-target gives

$$\tilde{b}_j(o') - \tilde{b}_j(o^*) \geq b'_j(o') - b'_j(o^*) \ \ .$$

For bidders $j \neq i$ with $\omega_j = \omega''_j$, we trivially have

$$\tilde{b}_j(o') - \tilde{b}_j(o^*) = b'_j(o') - b'_j(o^*) \ ,$$

so it remains to consider bidder i.

For bidder i, we know that decreasing her utility-target from ω'_i to ω''_i increased her bid for o^* more than it increased her bid for o'. This implies $v_i(o') < v_i(o^*)$ and $\omega''_i < v_i(o*)$. Consequently, i is a winner with ω''_i at o^* and will not decrease her utility-target further. Firstly, this implies that $|\tilde{\omega}_i - \omega_i| \leq \epsilon$. First, suppose $\omega_i > \omega''_i$. In this case, $\omega_i \geq \omega'_i$, and since $v_i(o^*) > v_i(o')$ we have

$$\tilde{b}_i(o') - \tilde{b}_i(o^*) \geq b'_i(o') - b'_i(o^*) \ .$$

Otherwise, i does not change her bid from ω'' to ω, so $\tilde{\omega}_i = \omega'_i$ and therefore $\tilde{b}_i = b'_i$

Thus, for any bidder j we have

$$\tilde{b}_j(o') - \tilde{b}_j(o^*) \geq b'_j(o') - b'_j(o^*) \ ,$$

and thus

$$\sum_{j \in [n]} \tilde{b}_j(o') - \sum_{j \in [n]} \tilde{b}_j(o^*) \geq \sum_{j \in [n]} b'_j(o') - \sum_{j \in [n]} b'_j(o^*) \ .$$

Since o' was winning at b'_j, this implies o^* cannot be winning under $\tilde{\omega}$, and therefore $\tilde{\omega} \in \overline{\mathcal{C}}$. By construction, $|\tilde{\omega}_j - \omega_j| \leq \epsilon$, so this implies $\omega \in \overline{\mathcal{C}}_\epsilon$.

So far, we showed that $\omega \in \overline{\mathcal{C}}_\epsilon$ as long as the outcome changed when i raised her bid. In the case where the outcome was already $o' = o^*$, we want to analyze the CEF constraints directly. Since $\omega' \in \overline{\mathcal{C}}$, there is some outcome o^w for which the CEF constraints are violated, i.e.

$$\sum_{i\in[n]} b'_i(o^*) - b'_i(o^w) < \sum_{i\in[n]} \max\left((v_i(o^w) - b'_i(o^w)) - (v_i(o^*) - b'_i(o^*)), 0\right) \ .$$

Observing that the outcome does not change from ω' to ω, the logic from the case where $o' \neq o^*$ gives

$$\tilde{b}_j(o') - \tilde{b}_j(o^*) \geq b'_j(o') - b'_j(o^*)$$

for any bidder j. It immediately follows that

$$\sum_{i\in[n]} \tilde{b}_i(o^*) - \tilde{b}_i(o^w) < \sum_{i\in[n]} \max\Big((v_i(o^w) - \tilde{b}_i(o^w)) - (v_i(o^*) - \tilde{b}_i(o^*)), 0\Big) \ ,$$

and so $\tilde{\omega} \in \overline{\mathcal{C}}$ and $\omega \in \overline{\mathcal{C}}_\epsilon$. □

C.1.3. Convergence to the Egalitarian Equilibrium

Theorem 6.12 (Restatement). *If losing bidders will raise their effective bids (A1), winning bidders will try lowering their effective bids (A3), and the most impatient bidder is the losing bidder bidding for the highest utility (A2, A4), then bids will converge to the Egalitarian envy-free equilibrium.*

Proof of Theorem 6.12: The proof will proceed as follows. We first categorize bidders into levels based on their utility in the egalitarian outcome. We define upper and lower bounds on utility-targets as multiples of ϵ, the amount by which agents change their bids. Then, we show that if for a given bidder j, the bid of every lower-utility bidder has converged to within their bounds, the bid of j will also converge to within her bounds - first to at least her lower bound (Lemma C.4), and then to at most her upper bound (Lemma C.6). Combining these via induction gives our final result that the bids of all agents converge near their egalitarian outcome.

Let o^* be the egalitarian outcome; let ω_i^* be the corresponding utility-target of bidder i. Let $B_b(o) = \sum_{i \in [n]} b_i(o)$ be the total bid for a given outcome o. Let $B_X(o) = \sum_{j \in L_X} b_j(o)$, and $B_X^*(o)$ be similarly defined.

First, consider all utility-targets in the egalitarian equlibrium; let z_i be the ith smallest (distinct) utility-target. Let L_i be the set of all agents with a utility-target of z_i in the egalitarian equilibrium. We will use $L(j)$ to denote the level of a bidder j.

We will show convergence by showing that there exist functions $b^-(i)$ and $b^+(i)$ s.t. for any $j \in L_i$, utility-targets converge into and remain in the interval $[\omega_j^* - \epsilon b^-(i), \omega_j^* + \epsilon b^+(i)]$.

Bidding Bounds. We now precisely define the bounding functions $b^-(\cdot)$ and $b^+(\cdot)$.

Definition C.1.

$$b^-(i) = 2^{2|L_{<i}|} \tag{C.1}$$

$$b^+(i) = 2^{2|L_{<i}|+|L_i|} \tag{C.2}$$

These bounds are given specifically so that for any level k, the sum over upper bounds in lower levels is at most the lower bound in level k, and the sum over all lower bounds for lower (or equal) levels is at most the upper bound for level k. Intuitively, we are saying that lower-level bidders cannot over or under bid enough to make up for bidders in level k.

Claim C.1.

$$b^-(k) > \sum_{i=0}^{k-1} |L_i| b^+(i) \tag{C.3}$$

$$b^+(k) > \sum_{i=0}^{k} |L_i| b^-(i) \tag{C.4}$$

We omit the proof; it follows from manipulation of exponential sums.

Witness outcomes. Recall from Algorithm 6.2 that utility-targets for any given agent are raised until the CEF constraint for some outcome o is violated. These outcomes have an important role to play in the egalitarian equilibrium — they are the reason that a bidder cannot achieve any more utility. We will call them *witness* outcomes.

Three properties of these witness outcomes are important for us. First, bidder i values the witness at less than her egalitarian bid, hence she would be 'losing' if it was chosen above the egalitarian winning ad; that all bidders with higher utility value it at at least their utility-target; and that with the final winning bids, the total bid of each is tied. We define *witness* outcomes precisely as follows:

Definition C.2. Outcome o^w is a *witness* outcome for bidder i at the egalitarian utility-targets ω^* if its total bid is tied with the egalitarian outcome, i asking for more utility at the egalitarian equilibrium results in a higher total bid for o^w than for the optimal egalitarian ad and i is the highest-utility bidder to lose if o^w wins over o^*.

Recall the intuition behind these outcomes: they are the reason that a bidder cannot achieve more utility at the egalitarian equilibrium. If there is no witness for a bidder who must pay something, then the bidder could ask for more utility, and higher utility bidders could effectively 'pick up the slack', resulting in a more egalitarian outcome.

Claim C.2. *At the egalitarian equilibrium, every bidder i s.t. $\omega_i^* < vi(o^*)$ has at least one witness.*

Proof. We will prove via contradiction. Assume at the CEF egalitarian outcome o^* bidder i has no witness. Now, let bidder i increase her utility-target by a small enough $\epsilon > 0$, that only outcomes that were previously tied with o^* win over o^*. For each of these outcomes, there must be a higher utility bidder than i who does not win with the outcome; otherwise it would be a witness for i. Decrease the utility-targets of the highest utility bidder not in each of these outcomes by ϵ. At this point, all outcomes will be tied again — and we can have the optimal outcome win the tiebreaker via having a higher utility, or assume that one agent will decrease, then raise their utility-target to ensure that it was the previous outcome to win. These bids will be CEF, and will be more egalitarian than o^*, as bidder i achieved more utility, and only higher utility bidders achieved less utility. □

Another important property will be that each outcome is only a witness for bidders of a single level:

Claim C.3. *An outcome is only a witness outcome for bidders of a single level.*

This really follows from the definition — bidders in different levels cannot both be the highest utility bidder to not win with an outcome. More intuitively though, if agents of different levels were both not in an outcome, and the lower utility bidder had no other witness outcome, then a more egalitarian outcome would involve increasing his utility-target, and decreasing the utility-target of the higher utility bidder.

Bidding convergence. We now present the core of our convergence result. This convergence is a two step process for bidders in a given level; after the utility bids of all lower level bidders have converged within their bounds, convergence in the given level to at least the lower bound takes place first, and then bids in the given level will converge to below their upper bound.

Lemma C.4. *Under assumptions A1, A2, A3 and A4, the utility-target of each bidder i in level L_i will converge to at least their lower bounds, $\omega_i^* - \epsilon \cdot b^-(i)$ if for every bidder j' in level L_{i^-} s.t. $i^- < i$, $\omega_{j'} \leq \omega_{j'}^* + \epsilon \cdot b^+(i^-)$.*

Proof. Our argument consists of two parts: first, that if a bidder is bidding for utility at or below her lower bound then she will never reduce her utility-target further. Second, she will eventually try raising her bid (by A3). These two combined will lead to her eventually raising her bid to at least the lower bound.

Claim C.5. *Under the assumptions of Lemma C.4, no bidder i in level L_i with a utility-target of $\omega_i \leq \omega_i^* - \epsilon \cdot b^-(i)$ will lower her utility-target.*

We will prove via contradiction. Assume for bidder i that with a utility-target of $\omega_i \leq \omega_i^* - \epsilon \cdot b^-(i)$, she wishes to lower her utility-target further. Let o be the winning outcome with bids ω. As i will only lower her bid if she is losing (A1), $\omega_i > vi(o)$. We will now try to derive the contradiction that the total bid for the optimal outcome is at least the total effective bid for o ($B(o^*) > B(o)$), hence she must win and would not care to lower her utility-target.

For i to decrease her utility-target, by A4 she must be the highest utility bidder who is losing such that $\omega_i > vi(o)$. By our bound, we know that for every lower utility bidder j' in level L_{i^-}, $\omega_{j'} \leq \omega_{j'}^* + \epsilon b^+(i^-)$. Since o^* is the optimal winning outcome and o the currently winning outcome, $B^*(o^*) \geq B^*(o)$ and $B(o^*) \leq B(o)$.

In the egalitarian outcome, every bidder i receives the utility she bids for; hence $b^*_i(o^*) = vi(o^*) - \omega^*_i$. By our assumption on the utility-target bounds, for all bidders $j' \in L_{<i+1}$, $b^*_{j'}(o^*) - b_{j'}(o^*) \leq \epsilon b^+(L(j'))$.

Consider a bidder j' in a lower level than i and first, is requesting more utility relative to the egalitarian outcome, specifically that $\omega_{j'} > \omega^*_{j'}$. Hence, we will have $0 \geq b_{j'}(o^*) - b^*_{j'}(o^*) \geq -(\omega_{j'} - \omega^*_{j'})$ and $0 \geq b_{j'}(o) - b^*_{j'}(o) \geq -(\omega_{j'} - \omega^*_{j'})$. Hence,

$$(b_{j'}(o^*) - b^*_{j'}(o^*)) - (b_{j'}(o) - b^*_{j'}(o)) \geq -(\omega_{j'} - \omega^*_{j'}) \tag{C.5}$$

$$\geq -b^+(L(j')). \tag{C.6}$$

Consider the case that $\omega_{j'} \leq \omega^*_{j'}$, that j' is requesting less utility than in the egalitarian outcome. Then $b_{j'}(o^*) - b^*_{j'}(o^*) = -(\omega_{j'} - \omega^*_{j'}) \geq 0$, and $b_{j'}(o) - b^*_{j'}(o) \leq -(\omega_{j'} - \omega^*_{j'})$. Hence,

$$(b_{j'}(o^*) - b^*_{j'}(o^*)) - (b_{j'}(o) - b^*_{j'}(o)) \geq 0. \tag{C.7}$$

Summing over all lower-level bidders via Equations (C.6) and (C.7) gives $(B_{<i}(o^*) - B^*_{<i}(o^*)) - (B_{<i}(o) - B^*_{<i}(o)) \geq -\sum_{i'<L(i)} b^+(i')$ and hence by Claim C.1,

$$(B_{<i}(o^*) - B^*_{<i}(o^*)) - (B_{<i}(o) - B^*_{<i}(o)) > -b^-(i). \tag{C.8}$$

Now, consider a bidder j' in the same or a higher level than i. If j' is overbidding and not winning in outcome o with bids b, then she would have decreased her utility-target faster than i. She could however be overbidding and winning in o; in which case the decrease in bids for o^* must be bounded by the decrease for o, hence: $(b_{j'}(o^*) - b^*_{j'}(o^*)) - (b_{j'}(o) - b^*_{j'}(o)) \geq 0$. If she is requesting less utility, o^* will see the full increase in bid while o may not. Denote the total bid of all bidders aside from i in the same or higher level as i as $B_{\geq i \setminus i}(o)$. Then, summing over all such bidders gives

$$B_{\geq i \setminus i}(o^*) - B^*_{\geq i \setminus i}(o^*)) - (B_{\geq i \setminus i}(o) - B^*_{\geq i \setminus i}(o)) \geq 0. \tag{C.9}$$

Our original assumption on i gives $(b_i(o^*) - b_i^*(o^*)) - (b_i(o) - b_i^*(o)) \leq b^-(i)$. Now, taking the sum over this and equations (C.8) and (C.9) gives $(B(o^*) - B^*(o^*)) - (B(o) - B^*(o)) > -b^+(L(j)) + b^+(L(j)) = 0$. By our assumption that o^* is the egalitarian winning outcome, we have $B^*(o^*) - B^*(o) \geq 0$. Adding these yields

$$B(o^*) - B(o) > 0. \tag{C.10}$$

This is in violation of our assumption that o wins with bids b. Hence, no such bidder i can ever wish to lower her utility-target past the lower bound when all lower-level agents have bids within their upper bounds. By Assumption A3, she will eventually try and lower her bid when winning, hence her bid will converge above her lower bound. □

Lemma C.6. *Under assumptions A1, A2, A3 and A4, the utility-target ω_i of each bidder i in level L_i will converge to at most the upper bound, $\omega_i^* + \epsilon \cdot b^+(i)$ if for every bidder j' in level L_{i^-} s.t. $i^- \leq i$, $\omega_{j'} \geq s^*_{j'} - \epsilon \cdot b^-(i^-)$.*

Proof. By Assumptions A1, A2 and Observation 6.3, a bidder will only request more utility from a set of bids b with winning outcome o if all other bidders are winning with bids b, and by Lemma 6.13, b must be CEF.

Our proof will proceed by showing that in any such o, $\omega_i < \omega_i^* + \epsilon \cdot b^+(i)$, and hence her utility-target must stay below $\omega_i^* + \epsilon \cdot b^+(i)$ in winning outcomes. Furthermore, by Theorem 6.9 bids will become CEF; hence i will be forced to decrease her utility-target.

By Claim C.2, there is a witness outcome o^w which includes every bidder j' in a strictly higher level i^+ than i. We will now show that if all other agents are winning with the egalitarian winning outcome, then i's utility-target must be below her upper bound, otherwise the witness outcome o^w would win over o^*.

By Definition C.2, $B^*(o^w) = B^*(o^*)$. Consider the quantity $B(o^*) - B^*(o^*)$, and break it into sums over bidders in levels at or below bidder i, i and bidders in levels above i:

$$\begin{aligned}(B(o^*) - B^*(o^*)) =& (B_{\leq i \setminus i}(o^*) - B^*_{\leq i \setminus i}(o^*)) \\ &+ (b_i(o^*) - b_i^*(o^*)) \\ &+ (B_{>i}(o^*) - B^*_{>i}(o^*))\end{aligned}$$

We will now proceed by separately considering bidders in higher and lower levels than bidder i. We will bound the change in bids from each, and see that there is no way for bidder i to ask for utility above her upper bound and still be in the winning outcome.

Higher-level bidders. By properties of witness sets, any such bidder j' must be winning in the witness outcome at both the egalitarian bids and the current bids, hence $b_{j'}(o^w) - b^*_{j'}(o^w) = -(\omega_{j'} - \omega^*_{j'})$. Since we know that at the egalitarian bids, such a bidder must be winning in the egalitarian outcome, $b_{j'}(o^*) - b^*_{j'}(o^*) = -(\omega_{j'} - \omega^*_{j'}) \leq b_{j'}(o^w) - b^*_{j'}(o^w)$. Summing over all such bidders yields

$$(B_{>i}(o^*) - B^*_{>i}(o^*)) - (B_{>i}(o^w) - B^*_{>i}(o^w)) \leq 0. \tag{C.11}$$

Lower-level bidders. By our initial assumption that bidding has converged above lower bounds for these bidders, for any bidder j' in $L_{\leq i}$, $\omega_{j'} \geq \omega^*_{j'} - \epsilon b^-(L(j'))$, and hence $b_{j'}(o^*) \leq b^*_{j'}(o^*) + \epsilon b^-(L(j'))$ and $b_{j'}(o^w) \leq b^*_{j'}(o^w) + \epsilon b^-(L(j'))$.

Recall that all bids $b^*_{j'}(o^*)$ and $b_{j'}(o^*)$ are winning by assumption — since o^* is the egalitarian outcome, and no agent wishes to decrease their utility-target in the current bids. If for some bidder j', $v_{j'}(o^w) \geq v_{j'}(o^*)$, then $(b_{j'}(o^w) - b^*_{j'}(o^w)) = (b_{j'}(o^*) - b^*_{j'}(o^*)) = -(\omega_{j'} - \omega^*_{j'})$.

Consider then the case that $vj'(o^w) < vj'(o^*)$; that is, that j' values the witness o less than the egalitarian outcome. We will consider two cases: that her utility-target is lower or higher than her egalitarian utility-target respectively.

$(\omega_{\mathbf{j'}} < \omega^*_{\mathbf{j'}})$ If the bidder j' bids for less utility than in the egalitarian outcome, then that increase in effective bid will be bounded by the increase in the bid for the egalitarian outcome. That is, we have $b_{j'}(o^w) - b^*_{j'}(o^w) = \max(v_{j'}(o^w), \omega_{j'}) - \omega_{j'} - \max(v_{j'}(o^w), \omega^*_{j'}) + \omega^*_{j'}$, and hence $b_{j'}(o^w) - b^*_{j'}(o^w) = -(\omega_{j'} - \omega^*_{j'}) + (\max(v_{j'}(o^w), \omega_{j'}) - \max(v_{j'}(o^w), \omega^*_{j'}))$. As $b_{j'}(o^*) - b^*_{j'}(o^*) = -(\omega_{j'} - \omega^*_{j'})$, we then have:

$$0 \leq b_{j'}(o^w) - b^*_{j'}(o^w) \leq b_{j'}(o^*) - b^*_{j'}(o^*) = -(\omega_{j'} - \omega^*_{j'}). \tag{C.12}$$

Furthermore, since $\omega_{j'} \geq \omega^*_{j'} - \epsilon b^-(L(j'))$ by assumption, we have:

$$0 \leq b_{j'}(o^w) - b^*_{j'}(o^w) \leq b_{j'}(o^*) - b^*_{j'}(o^*) \leq \epsilon b^-(L(j')) \tag{C.13}$$

and

$$0 \leq (b_{j'}(o^*) - b^*_{j'}(o^*)) - (b_{j'}(o^w) - b^*_{j'}(o^w)) \leq \epsilon b^-(L(j')). \tag{C.14}$$

$(\omega_{\mathbf{j'}} \geq \omega^*_{\mathbf{j'}})$ If bidder j' instead is bidding for at least as much utility as in the egalitarian outcome, the decrease in total bid is bounded by the change in bids for the egalitarian outcome, hence the change in utility-targets will be between $b_{j'}(o^*) - b^*_{j'}(o^*) = \omega^*_{j'} - \omega_{j'}$ and 0. Hence,

$$b_{j'}(o^*) - b^*_{j'}(o^*) \leq b_{j'}(o^w) - b^*_{j'}(o^w) \leq 0 \tag{C.15}$$

and

$$-(\omega_{j'} - \omega^*_{j'}) \leq (b_{j'}(o^*) - b^*_{j'}(o^*)) - (b_{j'}(o^w) - b^*_{j}(o^w)) \leq 0. \tag{C.16}$$

We now have upper bounds on $-(\omega_{j'} - \omega^*_{j'}) \leq (b_{j'}(o^*) - b^*_{j'}(o^*)) - (b_{j'}(o^w) - b^*_{j}(o^w))$ for all lower-level bidders. Taking the sum across all members of $L_{\leq i}$ via Equations (C.14) and (C.16) gives:

$$\sum_{j' \in L_{\leq i}} (b_{j'}(o^*) - b^*_{j'}(o^*)) - (b_{j'}(o^w) - b^*_{j'}(o^w)) \leq \sum_{j' \in L_{\leq i}} \epsilon b^-(L(j')) \tag{C.17}$$

Rearranging and noting that $\sum_{j' \in L_{\leq i}} \epsilon b^-(L(j')) < b^+(i)$ by Claim C.1 gives

$$(B_{\leq i \setminus i}(o^*) - B^*_{\leq i \setminus i}(o^*)) - (B_{\leq i \setminus i}(o^w) - B^*_{\leq i \setminus i}(o^w)) < \epsilon b^+(i). \tag{C.18}$$

Summing over equations (C.18) and (C.11) gives us:

$$\begin{aligned}&(B_{\leq i \setminus i}(o^*) - B^*_{\leq i \setminus i}(o^*)) - (B_{\leq i \setminus i}(o^w) - B^*_{\leq i \setminus i}(o^w)) \\ &\quad + (B_{>i}(o^*) - B^*_{>i}(o^*)) - (B_{>i}(o^w) - B^*_{>i}(o^w)) < \epsilon b^+(i).\end{aligned} \tag{C.19}$$

By assumption, $(b_i(o^*) - b_i^*(o^*)) = -(\omega_i - \omega_i^*) \leq -\epsilon b^+(i)$ and $b_i(o^w) = b_i^*(o^w)) = 0$. Thus, $(b_i(o^*) - b_i^*(o^*)) - (b_i(o^w) - b_i i^*(o^w)) \leq -\epsilon b^+(i)$. Adding this to (C.19) gives:

$$(B(o^*) - B^*(o^*)) - (B(o^w) - B^*(o^w)) < \epsilon b^+(i) - \epsilon b^+(i) = 0 \tag{C.20}$$

By our initial assumption that o^w is a witness outcome, $B^*(o^*) - B^*(o^w) = 0$. Adding this to the above equation yields

$$B(o^*) < B(o^w) \tag{C.21}$$

This contradicts our assumption that o^* is a winning set with bids $b(\cdot)$. Hence, i will be forced to decrease her utility-target to at most $\omega_i^* + \epsilon b^+(i)$ before the egalitarian winning set o^* is winning again. □

Combining Lemma C.4 and Lemma C.6 gives us convergence of each bidder in each level i to within their bounds as soon as lower level bidders have all converged. It follows then from straightforward induction on levels that all bids converge to within their bounds. □

APPENDIX D

Risk Averse bidders

D.1. Proofs from Section 5.3

Lemma 5.4 (Restatement). *A mechanism with two-price allocation rule $x = x_{\text{val}} - x_C$ is BIC if and only if for all v and v^+ such that $v < v^+ \leq v + C$,*

$$\frac{x_{\text{val}}(v)}{C} \leq \frac{x_C(v^+) - x_C(v)}{v^+ - v} \leq \frac{x(v^+)}{C}. \tag{5.1}$$

Proof of Lemma 5.4. Consider an agent and fix two possible values of the agent $v \leq v^+ \leq v + C$. The utility for truthtelling with value v is $C \cdot x_C(v)$ in a two-price auction. The utility for misreporting v^+ from value v is $x_{\text{val}}(v^+) \cdot (v - v^+) + x_C(v^+) \cdot (C + v - v^+)$: when the mechanism sells and charges v^+, the agent's utility is $v - v^+$; when the mechanism sells and charges $v^+ - C$, her utility is $U_C(C + v - v^+) = C + v - v^+$ (since $v < v^+$). Likewise, the utility for misreporting v from true value v^+ is $x_{\text{val}}(v) \cdot (v^+ - v) + x_C(v) \cdot C$. Note that here when the mechanism charges $v - C$, the utility of the agent is C because the wealth $C - v + v^+$ is more than C; when the mechanism charges v, her utility is $v^+ - v$ because we assumed $v^+ \leq v + C$.

An agent with valuation v (or v^+) would not misreport v^+ (or v) if and only if

$$x_C(v) \cdot C \geq x_{\text{val}}(v^+) \cdot (v - v^+) + x_C(v^+) \cdot (C + v - v^+); \tag{D.1}$$

$$x_C(v^+) \cdot C \geq x_{\text{val}}(v) \cdot (v^+ - v) + x_C(v) \cdot C. \tag{D.2}$$

Now the right side of (5.1) follows from (D.1) and the left side follows from (D.2).

When $v^+ > v+C$, the agent with value v certainly has no incentive to misreport v^+, since any outcome results in non-positive utility. Alternatively, the agent with value v^+ will derive utility $C \cdot x(v)$ from misreporting v and thus will misreport if and only if $x(v) > x_C(v^+)$. Substituting $v + C$ for v^+ in equation (5.1) gives $x(v) \leq x_C(v + C)$, and taking this for intermediate points between $v + C$ and v^+ gives monotonicity of $x_C(v)$ over $[v + C, v^+]$. Combining these gives $x(v) \leq x_C(v + C) \leq x_C(v)$ and hence v^+ will not misreport v. □

Corollary 5.5 (Restatement): *The allocation rules x_C and x_{val} of a BIC two-priced mechanism satisfies that for all $v < v^+$,*

$$\int_v^{v^+} \frac{x_{\text{val}}(z)}{C} \, \mathrm{d}z \leq x_C(v^+) - x_C(v) \leq \int_v^{v^+} \frac{x(z)}{C} \, \mathrm{d}z. \tag{5.4}$$

Proof of Corollary 5.5. Without loss of generality, suppose $v^+ \leq v + C$ (the statement then follows for higher v^+ by induction). Define function

$$\bar{x}_C(z) = x_C(v) + \int_v^z \frac{\sup_{y' \in [v,y]} x_{\text{val}}(y')}{C} \, \mathrm{d}y, \quad \forall z \in [v, v^+],$$

then $\bar{x}_C(z) \geq x_C(v) + \int_v^z \frac{x_{\text{val}}(y)}{C} \, \mathrm{d}y$ and hence

$$\int_v^z \frac{x_{\text{val}}(y)}{C}\,\mathrm{d}y \le \bar{x}_C(z) - x_C(v).$$

By the argument in the proof of Lemma 5.9,part 2, we have $\bar{x}_C(z) \le x_C(z)$, for all z. This gives the left side of (5.4). The other side is proven similarly. □

D.2. Proofs from Section 5.4

Definition 5.2 (Restatement). We define $\bar{x} = \bar{x}_C + \bar{x}_{\text{val}}$ as follows:

(1) $\bar{x}_{\text{val}}(v) = x_{\text{val}}(v)$;

(2) Let $r(v)$ be $\frac{1}{C}\sup_{z\le v} x_{\text{val}}(z)$, and let

$$\bar{x}_C(v) = \begin{cases} \int_0^v r(y)\,\mathrm{d}y, & v \in [0, C]; \\ x_C(v), & v > C. \end{cases} \tag{5.5}$$

Lemma 5.9 (Restatement).

(1) *On $v \in [0, C]$, $\bar{x}_C(\cdot)$ is a convex, monotone increasing function.*

(2) *On all v, $\bar{x}_C(v) \le x_C(v)$.*

(3) *The incentive constraint from the left-hand side of* (5.4) *holds for $\bar{x}_C$: $\int_v^{v^+} \bar{x}_{\text{val}}(z)\mathrm{d}z \le \bar{x}_C(v^+) - \bar{x}_C(v)$ for all $v < v^+$.*

(4) *On all v, $\bar{x}_C(v) \le x_C(v)$, $\bar{x}(v) \le x(v)$, and $\bar{p}(v) \ge p(v)$.*

Proof of Lemma 5.9.

(1) On $[0, C]$, $\bar{x}_C(v)$ is the integral of a monotone, non-negative function.

(2) The statement holds directly from the definition for $v > C$; therefore, fix $v \le C$ in the argument below.

Since $r(v)$ is an increasing function of v, it is Riemann integrable (and not only Lebesgue integrable).

Fixing v, we show that, given any $\epsilon \leq 0$, $\bar{x}_C(v) \leq x_C(v) + \epsilon$. Fix an integer $N > v/\epsilon$, and let Δ be $v/N < \epsilon$. Consider Riemann sum $S = \sum_{j=1}^{N} \Delta \cdot r(\xi_j)$, where each ξ_j is an arbitrary point in $[(j-1)\Delta, j\Delta]$.[1] We will also denote by $S(k) = \sum_{j=1}^{k} \Delta \cdot r(\xi_j)$, $k \leq N$, the partial sum of the first k terms. Since $\bar{x}_C(v) = \lim_{\Delta \to 0} S$, it suffices to show that for all $\Delta < \epsilon$, $S \leq x_C(v) + \epsilon$. In order to show this, we define a piecewise linear function y. On $[0, \Delta]$, y is 0, and then on interval $[j\Delta, (j+1)\Delta]$, y grows at a rate $r((j-1)\Delta)$. Intuitively, y "lags behind" x_C by an interval Δ and we will show it lower bounds x_C and upper bounds $S + \epsilon$. Note that since r is an increasing function, y is convex.

We first show $y(v) \leq x_C(v)$. We will show by induction on j that $y(z) \leq x_C(z)$ for all $z \in [0, j\Delta]$. Since y is 0 on $[0, \Delta]$, the base case $j = 1$ is trivial. Suppose we have shown $y(z) \leq x_C(z)$ for all $z \in [0, (j-1)\Delta]$, let us consider the interval $[(j-1)\Delta, j\Delta]$. Let z^* be $\arg\max_{z \leq (j-1)\Delta} x_{\text{val}}(z)$.[2] By the induction hypothesis, $y(z^*) \leq x_C(z^*)$. Recall that $z^* \leq z \leq C$. By the BIC condition (5.2), for all $z \geq z^*$,

$$x_C(z) \geq x_C(z^*) + \frac{x_{\text{val}}(z^*)}{C}(z - z^*).$$

[1]Obviously S depends both on Δ and the choice of ξ_j's. For cleanness of notation we omit this dependence and do not write $S_{\Delta,\xi}$.

[2]Here we assumed that $\sup_{z<(j-1)\Delta} x_{\text{val}}(z)$ can be attained by z^*, which is certainly the case when x_{val} is continuous. It is straightforward to see though that we do not need such an assumption. It suffices to choose z^* such that $x_{\text{val}}(z^*)$ is close enough to $r((j-1)\Delta)$. The proof goes almost without change, except with an even smaller choice of Δ.

On the other hand, by definition, r is constant on $[z^*, z]$, and the derivative of y is no larger than $r(z^*)$ on $[z^*, z]$. Hence for all $z \le j\Delta$,

$$\begin{aligned} y(z) &\le y(z^*) + \frac{x_{\text{val}}(z^*)}{C}(z - z^*) \\ &\le x_C(z^*) + \frac{x_{\text{val}}(z^*)}{C}(z - z^*) \le x_C(z). \end{aligned}$$

This completes the induction and shows $y(z) \le x_C(v)$ for all $z \in [0, v]$.

Now we show $S \le y(v) + \epsilon$. Note that since $r(z) \le 1$ for all z, $S \le S(N - 1) + \Delta < S(N-1) + \epsilon$. We will show by induction that $S(N-1) \le y(v)$. Our induction hypothesis is $S(j-1) \le y(j\Delta)$. The base case for $j = 1$ is obvious as $S(0) = y(\Delta) = 0$.

$$\begin{aligned} S(j) &= S(j-1) + \Delta \cdot r(\xi_j) \\ &\le y(j) + \Delta \cdot r(j\Delta) \\ &= y(j+1). \end{aligned}$$

In the inequality we used the induction hypothesis and the monotonicity of r. The last equality is by definition of y.

This completes the proof of part 2.

(3) For $v \le v^+ \le C$, by definition of $\bar{x}_C$,

$$\bar{x}_C(v^+) - \bar{x}_C(v) = \int_v^{v^+} r(z)\,\mathrm{d}z \ge \int_v^{v^+} \frac{x_{\text{val}}(z)}{C}\,\mathrm{d}z.$$

For $C \le v \le v^+$, $\bar{x}_C$ and $\bar{x}_{\text{val}}$ are equal to $\bar{x}_C$ and $\bar{x}_{\text{val}}$ on $[v, v^+]$, and the inequality follows from Corollary 5.5. For $v \le C$ and $v^+ \ge C$, we have

$$\begin{aligned}
\bar{x}_C(v^+) - \bar{x}_C(v) &= [\bar{x}_C(v^+) - \bar{x}_C(C)] + [\bar{x}_C(C) - \bar{x}_C(v)] \\
&\ge \int_C^{v^+} \frac{x_{\text{val}}(z)}{C}\,\mathrm{d}z + \int_v^C \frac{x_{\text{val}}(z)}{C}\,\mathrm{d}z \\
&= \int_v^{v^+} \frac{x_{\text{val}}(z)}{C}\,\mathrm{d}z.
\end{aligned}$$

(4) The first part, $\bar{x}_C(v) \le x_C(v)$, is from part 2 of the lemma and the definition of $\bar{x}_C(v) = x_C(v)$ on $v > C$. The second part, $\bar{x}(v) \le x(v)$, follows from the definition of $\bar{x}_{\text{val}}(v) = x_{\text{val}}(v)$, the first part, and the definition of $x(v) = x_{\text{val}}(v) + x_C(v)$. The third part, $\bar{p}(v) \ge p(v)$, follows because lowering $x_C(v)$ to $\bar{x}_C(v)$ on $v \in [0, C]$ foregoes payment of $v - C$ which is non-positive (for $v \in [0, C]$). □

D.3. Proofs from Section 5.5

Theorem 5.11 (Restatement). *An allocation rule* $\mathbf{x}$ *and payment rule* p *are the BNE of a one-priced mechanism if and only if (a)* $\mathbf{x}$ *is monotone non-decreasing and (b) if* $p(v) \ge p^{\text{VC}}(v)$ *for all* v *then* $p = p^C$ *is defined as*

$$p^C(0) = 0, \tag{5.10}$$

$$p^C(v) = \max\left(p^{\text{VC}}(v),\ \sup_{v^- < v} \left\{p^C(v^-) + (p^{\text{RN}}(v) - p^{\text{RN}}(v^-))\right\}\right). \tag{5.11}$$

Moreover, if $\mathbf{x}$ *is strictly increasing then* $p(v) \ge p^{\text{VC}}(v)$ *for all* v *and* $p = p^C$ *is the unique equilibrium payment rule.*

The proof follows from a few basic conditions. First, with strictly monotone allocation rule $\mathbf{x}$, the payment upon winning must be at least $v - C$; otherwise, a bidder would wish to overbid and see a higher chance of winning, with no decrease in utility on winning. Second, when the payment on winning is strictly greater than $v - C$, the bidder is effectively risk-neutral and the risk-neutral payment identity must hold locally. Third, when an agent is paying exactly $v - C$ on winning, they are capacitated when considering underbidding, but risk-neutral when considering overbidding. As a result, at such a point, p^C must be at least as steep as p^{RN}, i.e., if $\frac{\mathrm{d}}{\mathrm{d}v}p^{\mathrm{RN}}(v) > \frac{\mathrm{d}p^{\mathrm{VC}}}{\mathrm{d}v}(v)$, p^C will increase above p^{VC}, at which point it must follow the behavior of p^{RN}.

Theorem 5.11 follows from the following three lemmas which show the necessity of monotonicity, the (partial) necessity of the payment identity, and then the sufficiency of monotonicity and the payment identity.

Lemma D.1. *If $\mathbf{x}$ and p are the BNE of a one-priced mechanism, then $\mathbf{x}$ is monotone non-decreasing.*

Lemma D.1 shows that monotonicity of the allocation rule is necessary for BNE in a one-priced mechanism. Compare this to Example 5.1 where we exhibited a non-one-priced mechanisms that was not monotone. Because the utilities may be capacitated, the standard risk-neutral monotonicity argument; which involves writing the IC constraints for a high-valued agent reporting low and a low-valued agent reporting high, adding, and canceling payments; does not work.

Lemma D.2. *If $\mathbf{x}$ and p are the BNE of a one-priced mechanism and $p(v) \geq p^{\mathrm{VC}}(v)$ for all v, then $p = p^C$ (as defined in Theorem 5.11); moreover, if $\mathbf{x}$ is strictly monotone then $p(v) \geq p^{\mathrm{VC}}(v)$ for all v.*

From Lemma D.2 we see that one-priced mechanisms almost have a payment identity. It is obvious that a payment identity does not generically hold as a capacitated agent with value v is indifferent between payments less than $v - C$; therefore, the agent's incentives does not pin down the payment rule if the payment rule ever results in a wealth for the agent of more than C. Nonetheless, the lemma shows that this is the only thing that could lead to a multiplicity of payment rules. Additionally, the lemma shows that if $\mathbf{x}$ is strictly monotone, then these sorts of payment rules cannot arise.

Lemma D.3. *If allocation rule $\mathbf{x}$ is monotone non-decreasing and payment rule $p = p^C$ (as defined in Theorem 5.11), then they are the Bayes-Nash equilibrium of a one-priced mechanism.*

The following claim and notational definition will be used throughout the proofs below.

Claim D.4. *Compared to the wealth of type v on truthtelling, when type $v^+ > v$ misreports v she obtains strictly more wealth (and is more capacity constrained) and when type $v^- < v$ misreports v she obtains strictly less wealth (and is less capacity constrained) and if $p(v) \geq p^{\mathrm{VC}}(v)$ then type v^- is strictly risk neutral on reporting v.*

Definition D.1. Denote the utility for type v misreporting v' for the same implicit allocation and payment rules by $u^C(v, v')$ and $u^{\mathrm{RN}}(v, v')$ for risk-averse and risk-neutral agents, respectively.

Proof of Lemma D.1. We prove via contradiction. Assume that $\mathbf{x}$ is not monotone, and hence there is a pair of values, $v^- < v^+$, for which $\mathbf{x}(v^-) > \mathbf{x}(v^+)$. We will consider this in three cases: (1) when a type of v^- is capacitated upon truthfully reporting and winning, and when a type of v^- is strictly in the risk-neutral section of her utility upon winning and a type of v^+ is either in the (2) capacitated or (3) strictly risk-neutral section of her utility upon winning.

(1) (v^- capacitated). If v^- is capacitated upon winning, then v^+ will also be capacitated upon winning and misreporting v^- (Claim D.4). A capacitated agent is already receiving the highest utility possible upon winning. Therefore, v^+ strictly prefers misreporting v^- as such a report (strictly) increases probability of winning and (weakly) increases utility from winning.

(2) (v^- risk-neutral, v^+ capacitated). We split this case into two subcases depending on whether the agent with type v^- is capacitated with misreport v^+.

   (a) (v^- capacitated when misreporting v^+). As the truthtelling v^+ type is also capacitated (by assumption of this case), the utilities of these two scenarios are the same, i.e.,

$$u^C(v^-, v^+) = u^C(v^+, v^+). \tag{D.3}$$

   Since type v^- truthfully reporting v^- is strictly uncapacitated, if her value was increased she would feel a change in utility (for the same report); therefore, type v^+ reporting v^- has strictly more utility (Claim D.4), i.e.,

$$u^C(v^+, v^-) > u^C(v^-, v^-). \tag{D.4}$$

   Combining (D.3) and (D.4) we arrive at the contradiction that type v^+ strictly prefers to report v^-, i.e.,

$$u^C(v^+, v^-) > u^C(v^+, v^+).$$

(b) (v^- risk-neutral when misreporting v^+). First, it cannot be that the bidder of type v^+ is capacitated for both reports v^+ and v^- as, otherwise, misreporting v^- gives the same utility upon winning but strictly higher probability of winning. Therefore, both types are risk neutral when reporting v^-. Type v^- is risk-neutral for both reports so she feels the discount in payment from reporting v^+ instead of v^- linearly; type v^+ feels the discount less as she is capacitated at v^+. On the other hand, v^+ has a higher value for service and therefore feels the higher service probability from reporting v^- over v^+ more than v^-. Consequently, if v^- prefers reporting v^- to v^+, then so must v^+ (strictly).

(3) (v^- risk-neutral, v^+ risk-neutral). First, note that the price upon winning must be higher when reporting v^- than v^+, i.e., $p(v^-)/\mathbf{x}(v^-) > p(v^+)/\mathbf{x}(v^+)$; otherwise a bidder of type v^+ would always prefer to report v^- for the higher utility upon winning and higher chance of winning. Thus, a bidder of type v^+ must be risk-neutral upon underreporting v^- and winning; furthermore, risk-neutrality of v^+ for reporting v^+ implies the risk-neutrality of v^- for reporting v^+ (Claim D.4). As both v^+ and v^- are risk-neutral for reporting either of v^- or v^+, the standard monotonicity argument for risk-neutral agents applies.

Thus, for $\mathbf{x}$ to be in BNE it must be monotone non-decreasing. □

Proof of Lemma D.2. First we show that if $\mathbf{x}$ is strictly monotone then $p(v) \geq p^{\mathrm{VC}}(v)$ for all v. If $p(v) < p^{\mathrm{VC}}(v)$ then type v on truthtelling obtains a wealth w strictly larger than C. Type $v^- = v - \epsilon$, for $\epsilon \in (0, w - C)$, would also be capacitated when reporting v; therefore, by strict monotonicity of $\mathbf{x}$ such a over-report strictly increases her utility and BIC is violated.

The following two claims give the necessary condition.

$$p^C(v) \geq p^C(v^-) + (p^{\mathrm{RN}}(v) - p^{\mathrm{RN}}(v^-)), \quad \forall v^- < v \tag{D.5}$$

$$p^C(v) \leq \sup_{v^- < v} \left\{ p^C(v^-) + (p^{\mathrm{RN}}(v) - p^{\mathrm{RN}}(v^-)) \right\}, \quad \forall v \text{ } s.t. \text{ } p^C(v) > p^{\mathrm{VC}}(v). \tag{D.6}$$

Equation (D.5) is easy to show. Since $p^C(v) \geq p^{\mathrm{VC}}(v)$, the wealth of any type v^- when winning is at most C, and strictly smaller than C if overbidding. In other words, when overbidding, a bidder only uses the linear part of her utility function and therefore can be seen as risk neutral. Equation (D.5) then follows directly from the standard argument for risk neutral agents.[3]

Equation (D.6) would be easy to show if p^C is continuous: for all v where $p^C(v) > p^{\mathrm{VC}}(v)$, there is a neighborhood $(v - \epsilon, v]$ such that deviating on this interval only incurs the linear part of the utility function and the agent is effectively risk neutral. We give the following general proof that deals with discontinuity and includes continuous cases as well.

To show (D.6), it suffices to show that, for each v where $p^C(v) > p^{\mathrm{VC}}(v)$, for any $\epsilon > 0$, $p^C(v) < p^C(v^-) + (p^{\mathrm{RN}}(v) - p^{\mathrm{RN}}(v^-) + \epsilon$ for some $v^- < v$. Consider any $v^- > v - \frac{\epsilon}{2}$. Since $p^C(v^-) \geq p^{\mathrm{VC}}(v^-) = (v^- - C)\mathbf{x}(v^-) > (v - \frac{\epsilon}{2} - C)\mathbf{x}(v^-)$, the utility for v to misreport v^-, i.e., $u^C(v, v^-)$ is not much smaller than if the agent is risk neutral:

$$u^{\mathrm{RN}}(v, v^-) - u^C(v, v^-) < \frac{\epsilon}{2}\mathbf{x}(v^-).$$

[3]For a risk neutral agent, the risk neutral payment maintains the least difference in payment to prevent all types from overbidding.

The following derivation, starting with the BIC condition, gives the desired bound:

$$
\begin{aligned}
0 \leq u^C(v,v) - U_C(v,v^-) < U_C(v,v) &- u^{\mathrm{RN}}(v,v^-) + \frac{\epsilon}{2}\mathbf{x}(v^-) \\
&= (\mathbf{x}(v)v - p^C(v)) - (\mathbf{x}(v^-)v - p^C(v^-)) + \frac{\epsilon}{2}\mathbf{x}(v^-) \\
&= (\mathbf{x}(v) - \mathbf{x}(v^-))v - (p^C(v) - p^C(v^-)) + \frac{\epsilon}{2}\mathbf{x}(v^-) \\
&\leq p^{\mathrm{RN}}(v) - p^{\mathrm{RN}}(v^-) + (v - v^-)\mathbf{x}(v) - (p^C(v) - p^C(v^-)) + \frac{\epsilon}{2}\mathbf{x}(v^-) \\
&\leq p^{\mathrm{RN}}(v) - p^{\mathrm{RN}}(v^-) - (p^C(v) - p^C(v^-)) + \epsilon.
\end{aligned}
$$

The first equality holds because $p^C(v) > p^{\mathrm{VC}}(v)$; the second to last inequality uses the definition of risk neutral payments (Theorem 5.1, part 2), and the last holds because $\mathbf{x}(v^-) < \mathbf{x}(v) \leq 1$. □

Proof of Lemma D.3. The proof proceeds in three steps. First, we show that an agent with value v does not want to misreport a higher value v^+. Second, we show that the expected payment on winning, i.e., $p^C(v)/\mathbf{x}(v)$ is monotone in v. Finally, we show that the agent with value v does not want to misreport a lower value v^-. Recall in the subsequent discussion that p^{RN} is the risk-neutral expected payment for allocation rule $\mathbf{x}$ (from Theorem 5.1, part 2).

(1) (Type v misreporting v^+.) This argument pieces together two simple observations. First, Claim D.4 and the fact that $p^C \geq p^{\mathrm{VC}}$ imply that v is risk-neutral upon reporting v^+. Second, by definition of p^C, the difference in a capacitated agent's payments given by $p^C(v^+) - p^C(v)$ is at least that for a risk neutral agent given by $p^{\mathrm{RN}}(v^+) - p^{\mathrm{RN}}(v)$. The risk-neutral agent's utility is linear and she prefers reporting v to v^+. As the risk-averse agent's utility is also linear for payments in the given

range and because the difference in payments is only increased, then the risk-averse agent must also prefer reporting v to v^+.

(2) (Monotonicity of $p^C/\mathbf{x}$.) The monotonicity of $\frac{p^C}{\mathbf{x}}$, which is part 2 of Lemma 5.12, will be used in the next case (and some applications of Theorem 5.11). We consider v and v^+ and argue that $\frac{p^C(v)}{\mathbf{x}(v)} \leq \frac{p^C(v^+)}{\mathbf{x}(v^+)}$. First, suppose that the wealth upon winning of an agent with value v is C, i.e., $p^C(v) = p^{\mathrm{VC}}(v)$. If $p^C(v^+) = p^{\mathrm{VC}}(v^+)$ as well, then by definition of p^{VC} (by $\frac{p^{\mathrm{VC}}(v)}{\mathbf{x}(v)} = v - C$) monotonicity of $p^C/\mathbf{x}$ holds for these points. If p^C is higher than p^{VC} at v^+ then this only improves $p^C/\mathbf{x}$ at v^+. Second, suppose that the wealth of an agent with value v is strictly larger than C, meaning this agent's utility increases with wealth. The allocation rule $\mathbf{x}(\cdot)$ is weakly monotone (Lemma D.1), suppose for a contradiction that $\frac{p^C(v)}{\mathbf{x}(v)} > \frac{p^C(v^+)}{\mathbf{x}(v^+)}$ on $v < v^+$. Then the agent with value v can pretend to have value v^+, obtain at least the same probability of winning, and obtain strictly lower payment. This increase in wealth is strictly desired, and therefore, this agent strictly prefers misreport v^+. Combined with part 1, above, which argued that a low valued agent would not prefer to pretend to have a higher value, this is a contradiction.

(3) (Type v misreporting v^-.) If $p^C(v) = p^{\mathrm{VC}}(v)$, then paying less on winning does not translate into extra utility, and hence by the monotonicity of $p^C/\mathbf{x}$, the agent would never misreport.

We thus focus then on the case that $p^C(v) > p^{\mathrm{VC}}(v)$. By the monotonicity of $p^C/\mathbf{x}$, there is a point $v_0 < v$ such that for every value v^- between v_0 and v, if an agent with value v reported v^-, she would still be in the risk-neutral section of her utility function. Specifically, this entails that $\forall v^-$ such that $v_0 < v^- < v$, $p^C(v^-)/\mathbf{x}(v^-) \geq v - C$. Consider such a v_0 and any such v^-. For any such point,

$p^C(v^-)/\mathbf{x}(v^-) > v^- - C$, and hence a bidder with value v^- would also be strictly in the risk-neutral part of her utility function upon winning.

For every such point, by our formulation in (5.11), $p^C(v) - p^C(v^-) = p^{\text{RN}}(v) - p^{\text{RN}}(v)$. As a result, since she is effectively risk-neutral in this situation, she cannot wish to misreport v^-; otherwise, the combination of $\mathbf{x}$ and p^{RN} would not be BIC for risk-neutral agents.

For any $v^- \leq v_0$, the wealth on winning for a bidder with value v would increase, but only into the capacitated section of her utility function, hence gaining no utility on winning, but losing out on a chance of winning thanks to the weak monotonicity of $\mathbf{x}$. Hence, she would never prefer to bid v^- over bidding v_0. Combining this argument with the above argument, our agent with value v does not prefer to misreport any $v^- < v$. $\square$