

Economics and Computation: Ad Auctions and Other Stories

Christopher Wilkens



Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2013-66

<http://www.eecs.berkeley.edu/Pubs/TechRpts/2013/EECS-2013-66.html>

May 15, 2013

Copyright © 2013, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Acknowledgement

I want to thank those who have helped me reach this point --- my adviser Christos Papadimitriou for his incomparable wisdom and guidance; Kamal Jain and Aranyak Mehta for fantastic internships and advice; and many others at Berkeley, Microsoft, Google, ebay, MIT, and Sanford who have educated me. I also want to thank Darrell Hoy and Balasubramanian Sivan without whom many works in this thesis would have been shadows of their current forms. Last but not least, I want to thank my family and friends for setting a high standard, for making the educational process enjoyable, and for encouraging me to follow my passions even when my work was "too simple" or when "nobody would hire me." Both proved false, but only through the many named and unnamed people who made this dissertation

possible.

**Economics and Computation:
Ad Auctions and Other Stories**

by

Christopher A. Wilkens

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Computer Science

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Christos H. Papadimitriou, Chair
Professor Shachar Kariv
Professor Elchanan Mossel

Spring 2013

**Economics and Computation:
Ad Auctions and Other Stories**

Copyright 2013
by
Christopher A. Wilkens

Abstract

Economics and Computation:
Ad Auctions and Other Stories

by

Christopher A. Wilkens

Doctor of Philosophy in Computer Science

University of California, Berkeley

Professor Christos H. Papadimitriou, Chair

There is a growing research tradition in the interface between Economics and Computer Science: Economic insights and questions about incentives inform the design of systems, while concepts from the theory of computation help illuminate classical Economics problems. This dissertation presents results in both directions of the intellectual exchange.

Originally designed by industry engineers, the sponsored search auction has raised many interesting questions and spurred much research in auction design. For example, early auctions were based on a first-price payment model and proved to be highly unstable — this dissertation explores how improvements in the bidding language could restore stability. We also show that a first-price auction offers substantially better performance guarantees when a single advertiser may benefit from multiple ads. Another interesting problem arises because sponsored search auctions must operate with limited information about a user’s behavior — we show how sampling can maintain incentive compatibility even when the auctioneer incorrectly predicts the user’s behavior.

Computational tools also offer novel ways to understand the limits of complex economic systems. For example, a fundamental observation in this intellectual exchange is that people cannot be expected to solve computationally intractable problems. We show that this insight engenders a new form of stability we call *complexity equilibria*: when production has economies of scale, markets may be stable because finding a good deviation is computationally intractable. We also use techniques from communication complexity to show that equilibrium prices, even when they exist, may need to encode an impractical amount of information to guarantee that a market clears.

To Mom, Dad, and Molly.

Contents

Contents	ii
List of Figures	iv
1 Introduction	1
1.1 The Sponsored Search Auction	2
1.2 Complexity Theory and Markets	3
I Mechanism Design and Sponsored Search	6
2 Dynamic Axioms and First-Price Auctions	7
2.1 Introduction	8
2.2 Definitions and Preliminaries	12
2.3 Quasi-Truthful Bidding	14
2.4 Static Equilibrium Analysis	15
2.5 Utility-Target Auctions for Sponsored Search	20
2.6 Dynamic Analysis	22
2.7 Conclusion and Open Questions	33
3 Coopetitive Ad Auctions	35
3.1 Introduction	35
3.2 The Coopetitive Ad Auction	39
3.3 Pitfalls of Standard Mechanisms	40
3.4 Equilibria of the First-Price Auction	41
3.5 The Egalitarian Solution	44
3.6 Conclusion and Open Problems	46
4 Single-Call Mechanisms	47
4.1 Introduction	48
4.2 Preliminaries	51
4.3 Single-Call Mechanisms and Single-Call Reductions	53
4.4 Maximal-in-Distributional-Range Reductions	59

4.5	A Single-Call Application — PPC AdAuctions	68
4.6	Single-Parameter Reductions	71
4.7	Characterizing Reductions for Single-Parameter Domains	76
II Markets		84
5	Non-Convex Production and Complexity Equilibria	85
5.1	Introduction	85
5.2	Foundations and Models	87
5.3	Computing Pareto Optima	90
5.4	Computing Core Allocations	94
5.5	Computing Equilibria in a Sequential Model	100
5.6	Complexity Equilibria	104
5.7	Discussion and Open Problems	107
6	Market Communication in Production Economies	108
6.1	Introduction	108
6.2	Markets and Market Communication	109
6.3	A Lower Bound for the Arrow-Debreu Model	112
6.4	Conclusion	117
Bibliography		118
A	Convergence in Utility-Target Auctions	127
A.1	Convergence to the Egalitarian Equilibrium	127
B	Optimality in Single-Call Mechanisms	134
B.1	Optimality Proofs for Generalized BKS	134
B.2	Analysis Definitions, Facts, and Lemmas	150

List of Figures

- 3.1 A Google search for “samsung intel laptop” illustrates the pitfalls of ignoring the mutual benefits of an ad — the top ad for “Intel Laptops” is competing against the ad below it for “Samsung...w/ Intel.” Google is charging Intel a premium to show its ad on top; however, Intel should be happy if a user clicks on the Samsung ad (and possibly even the Newegg ad) and thus should be unwilling to pay this premium. Even worse, the inclusion of “w/ Intel” in Samsung’s ad suggests that it is subsidized by Intel through the Intel Inside program, further driving up the price of the top slot. 36

Acknowledgments

I want to thank the people who have helped me reach this point — my adviser Christos Papadimitriou for his incomparable wisdom and guidance; Kamal Jain and Aranyak Mehta for fantastic internship experiences and advice; and many others at Berkeley, Microsoft, Google, ebay, MIT, and Sanford who have educated me. I also want to thank coauthors Darrell Hoy and Balasubramanian Sivan without whom many works in this thesis would have been mere shadows of their current forms. Last but not least, I want to thank my family and friends for setting a high standard, for making the educational process enjoyable, and for encouraging me to follow my passions even when my work was “too simple” or when “nobody would hire me.” Both predictions proved false, but only through the many named and unnamed people who made this dissertation possible.

Chapter 1

Introduction

In 1936, an aspiring young mathematician named Alan Turing realized that understanding the true limits of Logic required more than just syntax and semantics: it required a machine. This is how the Turing machine — and the field of Computer Science — was born.

Today, theoretical economics has arrived at a similar crossroads. Its mathematical formalization has been wildly successful, yet, it still fails to give us a complete understanding of the way markets, social networks, and other complex economic systems function. Indeed, the behavior of an economic system is fundamentally a computation; thus, understanding such systems requires that we make the same leap Turing did many years ago.

Over the past twenty years, researchers at the intersection of Economics and Computer Science have explored many aspects of this fundamental relationship. One of the first connections between the two fields came from the recognition that complex economic systems are fundamentally computational. Economists have long sought to understand basic questions about economic systems — what does a market do? What strategies do bidders use in a repeated auction? Such questions occasionally have simple answers — the market price of oil goes up when a refinery burns down, or a bidder who feels cheated will lower her bid. However, just as often the answers are beyond the grasp of modern economics — we do not know how to accurately predict stock market crashes or the aggregate behavior of advertisers in a complex ad auction. Viewing a market, auction, or other economic system as a computation invites us to apply techniques from computer science to understand its behavior.

Computational tools are also useful for designing economic systems. In 1981, Myerson [71] asked a straightforward question: how should a seller, who wishes to maximize revenue, auction an item to a group of bidders? While simple, his question was groundbreaking because it was a question of engineering — alongside the work of Maskin and Hurwicz, Myerson's seminal study launched the engineering side of game theory now known as mechanism design. Computer science, as a discipline largely focused on engineering, has much to offer mechanism design — it has both techniques for measuring and designing good systems as well as systems in need of economic engineering. This dissertation explores a variety of descriptive and prescriptive questions inspired by ad auctions and markets.

1.1 The Sponsored Search Auction

In 1996, Goto.com introduced a revolutionary model of paid search. Instead of building a complicated ranking algorithm like its contemporaries, Goto.com (later Overture.com) asked websites to bid on search keywords and displayed the highest bidders on top. In effect, the engineers had built a real-time auction for selling internet advertising. Today, such sponsored search auctions dominate the multi-billion-dollar search advertising world (ads shown alongside search results) and raise a wide variety of research questions.

Dynamic Auctions

In the standard microeconomic approach, an economic system (auction, market, etc.) is constructed as an aggregate of individual agents (bidders, buyers, sellers, etc.). Each agent is defined by her behavior, i.e., by the algorithm that describes her choices.

An algorithmic perspective is especially natural in dynamic (repeated) auctions: because agents' behavior evolves over time, it can be likened to a learning algorithm. Economists and computer scientists alike have pursued this connection, for example, by borrowing the notion of regret from machine learning and showing that a system will converge to a correlated equilibrium if all agents' behavior satisfies a no-regret property [37, 32].

In Chapter 2 (work with Darrell Hoy and Kamal Jain [45]), we study repeated first-price auctions and show that simple, natural properties of bidder behavior can imply robust guarantees about an auction's outcome. First-price auctions, where the highest bidder wins and pays what she bid, are popular in practice for their simplicity and transparency; unfortunately, they are poorly understood in theory because the strategy of a rational bidder depends substantially on her beliefs about other bidders. While many implementations have been quite successful, including procurement auctions and treasury bill auctions, others have not. Our poor understanding of first-price auctions was exemplified by the failure of Overture's (Goto's) generalization of the first-price ad auction — the company's groundbreaking generalized first-price auction for sponsored search advertising was so unstable that history records it merely as a stepping stone on the road to Google's generalized second-price auction.

Our results show that when a repeated first-price auction is generalized carefully — in a way that allows bidders to express the final utility they desire — simple and natural properties of bidder behavior imply robust guarantees about the auction's outcome. As an illustration, consider a standard first-price single-item auction and suppose that a loser will raise her bid to try to win. Similar to the behavior of an ascending auction, it naturally follows that the final sale price will eventually reach the maximum willingness to pay of the second highest bidder. In more general settings, we show that analogous axioms of bidder behavior can imply similarly robust guarantees for the outcome of the auction. Importantly, while we do demonstrate bidder behavior that explicitly converges, such results show that important properties of an auction's outcome (revenue, efficiency, etc.) may be characterized without any guarantee that bids will converge to an equilibrium.

Coopetitive Ad Auctions. Another issue arising in advertising is that a single ad will often benefit many advertisers. For example, an ad for a Samsung laptop with an Intel processor benefits Samsung, Intel, and Microsoft. In Chapter 3 (work with Darrell Hoy and Kamal Jain [46]), we study the performance of ad auctions when accounting for such effects. We first show that the status quo — advertisers can only cooperate through external contracts — can and currently does create outcomes that are undesirable either for the advertisers or for the auctioneer. We thus argue that an auctioneer should run a *coopetitive ad auction*, that is, the auctioneer should account for the broad benefit of a single ad. Unfortunately, we also show that the standard VCG auction is unpalatable for the auctioneer because it may not generate any revenue. Instead, we show that a first-price auction can offer desirable guarantees in both efficiency and revenue, using techniques from Chapter 2 to characterize its performance.

Single-Call Mechanisms. A third problem that arises in sponsored search auctions is a lack of information: Google and other companies that sell advertising do not know the likelihood that a user will click on any particular advertisement. Instead, they build complicated systems to learn and predict these likelihoods. Unfortunately, standard auction theory assumes the auctioneer has complete information. As a result, properties like incentive compatibility fall apart. In fact, this is an instance of a more general problem: any time bidders' values depends on some unknown action of nature, standard auctions may lose incentive compatibility.

In Chapter 4 (work with Balasubramanian Sivan [91]), we explore randomized mechanisms called *single-call mechanisms* that recover incentive compatibility in expectation with very limited information. The key idea of a single-call mechanism, first demonstrated by Babaioff et al. [9], is that the allocation procedure (e.g., the computer system that estimates the likelihood of a click, assigns advertisements to slots, and measures the result) can be viewed as a black box which the mechanism can only call once.

The following question then arises: what mechanisms can be implemented under the single-call restriction, and what are their trade-offs? To answer this question, we first characterize all general single-call constructions for two important domains (single-parameter and MIDR) and show that they have very simple structures. Next, we use our characterizations to show that the auctioneer faces a substantial trade-off between the expected quality of the auction and the variance of its payments. While our lower-bounds suggest that the most general single-call mechanisms are an impractical solution for real-world auctions, they leave open the possibility of better single-call mechanisms for specific domains. Our characterizations also offer useful insight into the structure of certain randomized mechanisms.

1.2 Complexity Theory and Markets

The study of markets often focuses on a straightforward question: what do they do? A dominant perspective in theoretical microeconomics is that markets reach a market equilibrium

— a configuration in which each agent is selfishly optimizing her own profit or utility subject to market prices. This so-called theory of general equilibrium has developed over the past few centuries, beginning with classical economic ideas like the “invisible hand of the market” from Adam Smith’s *The Wealth of Nations* and the “fictitious auctioneer” of Leon Walras’s *Elements of Pure Economics*, and climaxing with Arrow and Debreu’s seminal proof that equilibria always exist in a large class of markets [5]. Implicit in much of this theory is an assumption that a market can reach an equilibrium whenever it exists.

Adopting a computational perspective raises a new question: even if a market equilibrium exists, is it reasonable to assume that a market will reach it? Unfortunately, finding an equilibrium may be computationally intractable, and therefore neither a large computer nor a market — which is effectively a large, distributed computer — can be expected to compute an equilibrium. Twenty years ago, Christos Papadimitriou [78] recognized that complexity theory offered rigorous tools to understand when an equilibrium might be reachable and inspired a long line of research on the complexity of equilibria.

More generally, complexity theory offers rigorous techniques to understand what a computation can and cannot do and thus provides a non-constructive framework to understand the limitations of complex economic systems.

Markets with Economies of Scale and Complexity Equilibria. In Chapter 5 (work with Christos Papadimitriou [80]), we ask how the computational tractability of equilibrium is affected by economies of scale in production. Economies of scale — e.g. the fact that the millionth car is easier to produce than the first — occur quite naturally, yet they are explicitly disallowed in Arrow and Debreu’s model. The economic theorist’s approach is to relax the requirements of equilibrium and hope that a market might still satisfy its first-order conditions [33], reaching a so-called marginal-price equilibrium. We show that even satisfying these first-order conditions in the presence of economies of scale is harder than computing Arrow and Debreu’s equilibrium — NP-hard¹ versus PPAD-complete — suggesting that marginal-price equilibrium is not a priori a reasonable predictor of market behavior. The key idea in our work is the observation that up-front investment costs (fixed costs) can force markets to make discrete choices — e.g. a factory can produce microchips or cars, but not a mixture of the two — and therefore encode a discrete problem like SAT.

In [80], we also identify a new phenomenon which we call *complexity equilibria*: in markets with economies of scale in production, some configurations can masquerade as equilibria because it is computationally intractable (NP-hard) to discover that the configuration is not a market equilibrium. We offer examples of complexity equilibria where the notion of NP-hardness is both worst-case and average-case.

Market Communication Complexity. Communication complexity offers another set of tools with which to study the limitations of market equilibrium. Since players make independent decisions, the market must reveal enough information that players can correctly

¹Specifically, it is $F\Delta_2^P$ -complete.

identify their equilibrium behaviors. Classical economic results on the dimensionality of message spaces show that, when information is communicated through real numbers, normalized prices (one real number per good) are both necessary [48, 70] and sufficient [5]. However, these results ignore the fact that prices are rarely represented as real numbers; instead, they are often fixed-precision values rounded to the nearest whole number, hundred, or even thousand. Deng et al. [23] introduced a model of market communication complexity and offered a lower-bound on the number of bits of information that prices must contain to communicate equilibrium behavior to the players.

In Chapter 6 (based on [90]), we give a substantially tighter bound than [23] and extend their results to markets with production. We show that the number of bits of information that must be contained in prices is polynomial not only in the number of goods, but also in the number of consumers and firms in the market.

Part I

Mechanism Design and Sponsored Search

Chapter 2

Dynamic Axioms and First-Price Auctions

The first-price auction is popular in practice for its simplicity and transparency. Moreover, its potential virtues grow in complex settings where incentive compatible auctions may generate little or no revenue. Unfortunately, the first-price auction is poorly understood in theory because equilibrium is not *a priori* a credible predictor of bidder behavior.

In this chapter, we take a dynamic approach to studying first-price auctions: rather than basing performance guarantees solely on static equilibria, we study the repeated setting and show that robust performance guarantees may be derived from simple axioms of bidder behavior. For example, as long as a loser raises her bid quickly, a standard first-price auction will generate at least as much revenue as a second-price auction.

We generalize this dynamic technique to complex pay-your-bid auction settings: as long as losers do not wait too long to raise bids, a first-price auction will reach an envy-free state that implies a strong lower-bound on revenue; as long as winners occasionally experiment by lowering their bids, the outcome will near the boundary of this envy-free set so bidders do not overpay; and when players with the largest payoffs are the least patient, bids converge to the egalitarian equilibrium. Significantly, bidders need only know whether they are winning or losing in order to implement such behavior.

Along the way, we find that the auctioneer's choice of bidding language is critical when generalizing beyond the single-item setting, and we propose a specific construction called the *utility-target auction* that performs well. The utility-target auction includes a bidder's final utility as an additional parameter, identifying the single dimension along which she wishes to compete. This auction is closely related to profit-target bidding in first-price and ascending proxy package auctions and gives strong revenue guarantees for a variety of complex auction environments. Of particular interest, the guaranteed existence of a pure-strategy equilibrium in the utility-target auction shows how Overture might have eliminated the cyclic behavior in their generalized first-price sponsored search auction if bidders could have placed more sophisticated bids.

2.1 Introduction

In 1961, Vickrey [87] 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 player had a strategy that dominated all others regardless of other players' bids. As a solution, he proposed¹ 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 players' informational limitations [87, 35]. 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” [7] 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 [57].

Dynamic analysis offers a powerful complement to Vickrey's static approach. For example, certain behavior will be unsustainable when an auction is repeated. Such reasoning was used by Edelman and Schwarz [30] 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 [28] 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* that is closely related to profit-target bidding in first-price and ascending proxy package auctions [67, 22]. Like these package auctions, we show that the utility-target auction possess 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

¹While Vickrey was the first to discover the second-price auction in the economics literature, it has been used in practice as early as 1893 [64].

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 players' 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 [7, 82]. 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 [67, 63] (though they may also generate less revenue [66]), and we discuss a specific example from ad auctions in Chapter 3.

Yet, running a first-price auction is risky. While first-price auctions have been quite successful in settings like treasury bill and 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 [28]. 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 disincentivize small adjustments to a player'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 players' 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, players will be in a Bayesian 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 [35].

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 [28] 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 players’ bids, potentially restricting or broadening the bids that players 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 [29, 28]; however, we show that 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 player’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. Moreover, we show that a pure-strategy equilibrium always exists and that the egalitarian equilibrium is efficiently computable. These results are closely related to profit-target equilibria in package auctions [67].

Next, we show that the utility-target auction offers good equilibrium performance. Similar to the approach of Edelman, Ostrovsky, and Schwarz [29] 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. Again, this bound is related to the core property of profit-target equilibria in package auctions [67, 22].

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 tâtonnement² [88] is perhaps the earliest concrete dynamic procedure proposed in economics — it converges in general markets when modeled as a particular continuous process [84, 4] but may or may not converge as a discrete process [12, 20]. More recent results have sought stronger guarantees, e.g. by showing that players’ behavior will converge to equilibrium in repeated games as long as their learning strategies are “adaptive and sophisticated” [69] or no-regret [37, 32]. However, these properties are sufficiently complicated that it is difficult to evaluate whether players’ 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 extension, a natural second-price type bound on revenue. Significantly, this result neither implies nor requires that players’ bids converge to a steady-state. Moreover, bidder behavior requires only knowing whether one is winning or losing, not the precise bids of other players.

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 [14] and the ascending proxy auction [67]. 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 2.4 all have direct analogues. In contrast, the utility-target auction can be applied beyond the package auction setting (e.g. to ad auctions), and 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” [67].

²To justify market equilibrium as a predictor of actual market behavior, Leon Walras described a dynamic procedure called tâtonnement that might converge to it.

Auctions in which a player’s bid directly specifies her payment are known as *pay-your-bid* auctions. The first-price auction, as well as the Dutch and English auctions, are members of this family. Our utility-target auction is closely related to first-price and ascending proxy package auctions [67]. Engelbrecht-Wiggans and Kahn [31] 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 [67], Milgrom offers a broad discussion of this literature in [68]. Some recent work studies ascending auctions for position auctions like sponsored search [29, 6].

Our work can also be seen through the lens of *simple versus optimal* mechanisms [40]. 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 [38]. 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.

2.2 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.

First-Price and Pay-Your-Bid Auctions.

An auction is a protocol through which players bid to select an outcome. A standard sealed-bid auction can be decomposed into three stages: (1) each player i submits a bid b_i , (2) the auctioneer uses players’ bids to pick an outcome o from a set \mathcal{O} , and finally (3) each player i pays a price p_i . The final utility of player 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 player submits a single number $b_i \in \mathbb{R}$, (2) the auctioneer chooses to give the item to the player 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 players' 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 player's bid precisely specifies her payment. Formally, we say that such an auction has the pay-your-bid property:

Definition 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. [29] showed that GFP, where the set of possible bids is precisely V_i , did not have a pure-strategy equilibrium:

Observation 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 2.5, 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 important 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.

Utility-Target Auctions

A utility-target auction is a sealed-bid pay-your-bid auction with a special bidding language. A player'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 2 *A utility-target auction for a finite outcome space \mathcal{O} is defined as follows:*

- *A bid is a tuple $b_i = (x_i, \pi_i)$ where $x \in V_i$ is a function mapping outcomes $o \in \mathcal{O}$ to nonnegative values and π is a real number. We call the parameters x_i and π_i the value bid and utility-target bid respectively.*

ALGORITHM 1: A generic utility-target auction.

input : Players' bids $b_i = (x_i, \pi_i)$
output: An outcome o^* and first-price payments p_i .

- 1 Let $b_i(o) = \max(0, x_i(o) - \pi_i)$; // $b_i(o)$ is i 's effective bid for outcome o .
 - 2 Compute $o^* = \operatorname{argmax}_o \sum_{i \in [n]} b_i(o)$; // Choose the outcome with the highest total bid.
 - 3 For all i , set $p_i = b_i(o^*)$; // Each player pays what she bid.
-

- A bidder's effective bid for outcome o is

$$b_i(o) = \max(x_i(o) - \pi_i, 0) .$$

Note this may generate $b_i \notin V_i$ when the set V_i is sufficiently restricted.

- The auctioneer chooses the outcome $o^* \in \mathcal{O}$ 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 $u_i(o) = \pi_i$ whenever $v_i(o) \geq \pi_i$.

A generic utility-target auction is illustrated in Algorithm 1.

2.3 Quasi-Truthful Bidding

An idealist's intuition for the utility-target auction is that players truthfully reveal their valuation function through their value bids (i.e. they bid $x_i = v_i$) and then use the utility-target bid π_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 player might consider, there is another bid in which she reveals v_i truthfully and obtains at least as much utility:

Lemma 1 (Quasi-Truthfulness) *Fix the total bid of players $j \neq i$ 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, π_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 $b_i^I = (x_i^I, \pi_i^I)$ for i , the outcome o^I is uniquely defined. Let u_i^I be the utility i gets by bidding (x_i^I, π_i^I) , i.e.

$$u_i^I = v_i(o^I) - b_i(o^I) = v_i(o^I) - \max(x_i^I(o^I) - \pi_i^I, 0) .$$

Now suppose i bids $b_i^Q = (v_i, u_i^I)$ instead of (x_i^I, π_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 $x_i(o) > \pi_i(o)$. Since $b_i^Q(o^Q) \geq 0$, this implies $v_i(o^Q) > \pi_i^Q$, from which it immediately follows that i 's final utility in o^Q will be π_i^Q .

In either case, i 's final utility is precisely u_i^I , so i is indifferent between bidding (x_i^I, π_i^I) and (v_i, u_i^I) . ■

2.4 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 2 *A utility-target auction with n outcomes always has a pure-strategy cooperatively envy-free (defined below) equilibrium that is computable in time $\text{poly}(n)$.*

Specifically, the egalitarian equilibrium exists and is efficiently computable by Algorithm 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 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.

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 value the first two outcomes and the third bidder values only the third outcome. Let the specific values, indexed by outcome, be

$$\begin{aligned} v_A &= (1, 1.5, 0), \\ v_B &= (1, 1.5, 0), \\ v_D &= (0, 0, 2). \end{aligned}$$

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 players 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:

$$\begin{aligned} v_A &= (1, 1.5, 0), \\ v_B &= (1, 1.5, 0), \\ v_C &= (1, 0.5, 0), \\ v_D &= (0, 0, 2). \end{aligned}$$

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 3 *The set of bids $\{b_i\}_{i \in [n]}$ are cooperatively envy-free (CEF) if there is no subset of bidders $\mathcal{B} \subseteq [n]$ 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

$$\sum_{i \in [n]} \max((v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)), 0) \leq \sum_{i \in [n]} b_i(o^*) - b_i(o)$$

for all outcomes o .

The CEF constraints are similar to the core property described by Milgrom [67] 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:

Claim 1 *CEF bids maximize welfare.*

Proof: 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

$$\begin{aligned} \sum_{i \in [n]} \max((v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)), 0) &\leq \sum_{i \in [n]} b_i(o^*) - b_i(o) \\ \sum_{i \in [n]} (v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)) &\leq \sum_{i \in [n]} b_i(o^*) - b_i(o) \\ \sum_{i \in [n]} v_i(o) - v_i(o^*) &\leq 0 \end{aligned}$$

as desired. ■

Claim 2 *The revenue from CEF bids dominates that of the VCG mechanism: every player pays at least as much in the CEF equilibrium as she would in the VCG mechanism.*

Proof: Define o^i as the outcome that maximizes the welfare of bidders except i :

$$o^i = \operatorname{argmax}_o \sum_{j \neq i} v_j(o) .$$

Thus, the VCG price of player i is $\sum_{j \neq i} v_j(o^i) - v_j(o^*)$.

Now, the envy-freeness constraints give

$$\begin{aligned}
b_i(o^*) &\geq \sum_j \max((v_j(o^i) - b_j(o^i)) - (v_j(o^*) - b_j(o^*)), 0) \\
&\quad + b_i(o^*) - \sum_j (b_j(o^*) - b_j(o^i)) \\
b_i(o^*) &\geq \sum_{j \neq i} (v_j(o^i) - b_j(o^i)) - (v_j(o^*) - b_j(o^*)) \\
&\quad - \sum_{j \neq i} (b_j(o^*) - b_j(o^i)) + b_i(o^i) \\
&\quad + \max((v_i(o^i) - b_i(o^i)) - (v_i(o^*) - b_i(o^*)), 0) \\
b_i(o^*) &\geq \sum_{j \neq i} (v_j(o^i) - v_j(o^*)) + b_i(o^i) \\
&\quad + \max((v_i(o^i) - b_i(o^i)) - (v_i(o^*) - b_i(o^*)), 0) \\
b_i(o^*) &\geq \sum_{j \neq i} (v_j(o^i) - v_j(o^*)) .
\end{aligned}$$

■

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, player 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 4 *The second-price threat for outcome o^* is given by*

$$\max_{o \in \mathcal{O}} \sum_{i \in [n]} \max(v_i(o) - v_i(o^*), 0) .$$

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,

ALGORITHM 2: An algorithm for computing the egalitarian equilibrium in a utility-target auction.

input : A utility-target auction problem.

output: The egalitarian equilibrium bids $b_i^* = (v_i, \pi_i^*)$.

- 1 Set all bids to $(v_i, 0)$. Call the socially optimal outcome o^* .
 - 2 Increase π_i for all bidders uniformly until some bidder i reaches $\pi_i = v_i(o^*)$ 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.
-

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 3 *The revenue in any CEF outcome is lower-bounded by the second-price threat.*

Proof: 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

$$\begin{aligned} \sum_{i \in [n]} \max((v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)), 0) &\leq \sum_{i \in [n]} b_i(o^*) - b_i(o) \\ \sum_{i \in [n]} \max(v_i(o) - v_i(o^*) + b_i(o^*), b_i(o)) &\leq \sum_{i \in [n]} b_i(o^*) \\ \sum_{i \in [n]} \max(v_i(o) - v_i(o^*), 0) &\leq \sum_{i \in [n]} \sum_{i \in [n]} b_i(o^*) \end{aligned}$$

as desired. ■

2.5 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 [83] and advertiser value (see Chapter 3). It can have misaligned incentives when parameters are estimated incorrectly [91]. 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.

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) c , and the value v to the advertiser of a user who clicks. The CTR c 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 $c_{i,j}v_i = \alpha_j\beta_iv_i$. 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. $b_1 \geq b_2 \geq \dots \geq b_n$.

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 π_i (the utility that they request). The auctioneer picks the assignment $j(i)$ maximizing

$$\sum_{i \in [n]} \max(0, \alpha_{j(i)}\beta_ix_i - \pi_i)$$

and charges i so that her expected payment is

$$\mathbf{E}[p_i] = \max(0, \alpha_{j(i)}\beta_ix_i - \pi_i) .$$

³The 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.

There are at least two interesting ways the utility-target auction can be implemented. The first implementation charges

$$ppc_i = \max\left(0, x_i - \frac{\pi_i}{\alpha_{j(i)}\beta_i}\right)$$

to achieve the desired expected payment. In effect, it uses the utility request π_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 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 π_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 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) [83] when considering externalities imposed by the presence of competing ads, the reasonable (CEF) equilibria of the utility-target auction guarantee good performance. For example, in Chapter 3 we study “coopetitive” ad auctions where 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 is necessary to extend these results to settings with multiple slots. The utility-target auction is also less sensitive to estimation errors in the CTRs. As shown in [91], 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.

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 π_i . Consequently, a player'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

ALGORITHM 3: Autility-target auction for search advertising.

input : Bids $b_i = (x_i, \pi_i)$.

output: An assignment of advertisements to slots, per-click payments ppc_i , and unconditional payments r_i .

- 1 For each bidder i and slot j , compute $\mathbf{E}[p_{i,j}] = \max(\alpha_j \beta_i x_i - \pi_i, 0)$;
 - 2 Compute assignment $j(i)$ of advertisements to slots that maximizes $\sum_{i \in [n]} \mathbf{E}[p_{i,j(i)}]$;
 - 3 **if** $\mathbf{E}[p_{i,j(i)}] > 0$ **then**
 - └ Always pay i the rebate $r_i = \pi_i$;
 - └ Whenever i 's ad is clicked, charge $ppc_i = x_i$;
 - else**
 - └ Do not charge/pay anything to i ;
-

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 [28] 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 reduced revenue below that of the VCG mechanism. In contrast, Theorem 2 shows that utility-target auctions have pure-strategy equilibria, and Theorem 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.

2.6 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 3.

Following Lemma 1, we assume that bidders are quasi-truthful and report bids of the form (v_i, π_i) . Bidders compete using the utility-target terms π_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 π_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 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*

$$u_i(o^b) = v_i(o^b) - b_i(o^b) = \pi_i .$$

Any bidder who is not a winner is a loser.

Observation 2 *Bidder i is a winner if and only if $v_i(o^b) \geq \pi_i$ and is always a winner when $\pi_i = 0$. When bidder i is a loser, $u_i(o^b) < \pi_i$.*

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 $\pi_i = 0$.

Raising and Lowering Bids. In a dynamic setting, we want to think about how winners and losers manipulate π_i . In the utility-target auction, the effective bid b_i (what bidder i is actually offering to pay) and the utility-target term π_i move in opposite directions, so when we talk about raising i 's bid we are talking about decreasing the utility-target term π_i :

Definition 6 (Raising and Lowering Bids) *We say that bidder i raises her bid from (v_i, π_i) if she chooses a new bid (v_i, π'_i) where $\pi'_i < \pi_i$, 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 π_i to $\pi'_i > \pi_i$.

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 4 *Fixing other players' 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 5 *A loser can always raise her bid in a way that weakly increases her utility.*

Proof: Suppose i is a loser bidding (v_i, π_i) and receiving utility $u_i < p_i$. If she raises her bid to (v_i, u_i) , Lemma 1 says that she will receive utility of precisely u_i , making her a winner.

■

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 players 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 π^0, \dots .

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 player'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 π are sufficient to characterize bidders' strategies, so we define \mathcal{C} to be the set of all such utility-target:

Definition 7 (The CEF Set) \mathcal{C} is the set of all utility-target vectors π where the quasi-truthful bids (v_i, π_i) produce a cooperatively envy-free outcome.⁴

The set $\bar{\mathcal{C}}$ is the set of all utility-target vectors which are not CEF, i.e. $\bar{\mathcal{C}} = \mathbb{R}_+^n \setminus \mathcal{C}$.

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 π , let π_ϵ denote the set of bids that are close to some vector in π , i.e.

Definition 8 Let S_ϵ be the set of all utility-targets π which are close to some vector in π coordinate-wise. Formally,

$$S_\epsilon = \{\pi \mid \exists \pi' \in S \text{ s.t. } \|\pi - \pi'\|_\infty \leq \epsilon\} .$$

In particular, we will care about the sets \mathcal{C}_ϵ and $\bar{\mathcal{C}}_\epsilon$, 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 π :

Definition 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 π such that all future bids are in S_δ .

⁴Note that membership in \mathcal{C} depends on both the vector π and the outcome chosen by the auction. This is because certain utility-target vectors π 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.

Our strongest result will show that bids converge to the egalitarian equilibrium:

Definition 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 π^E if it converges to $\{\pi^E\}$.

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 4, 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 5 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 π 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 π 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 player.* Formally, if the current utility-targets are π 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 π and bidders i and j are both losers, then i will raise her bid before j if $\pi_i > \pi_j$.

These simple properties of bidder behavior imply the following convergence results. Proofs follow in Section 2.6 and Appendix A.1.

Theorem 4 *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}_ϵ).*

Theorem 5 *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 $\bar{\mathcal{C}}_\epsilon$).*

Combining Theorems 4 and 5 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 *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 $\mathcal{C}_\epsilon \cup \bar{\mathcal{C}}_\epsilon$).*

Finally, adding A4 induces convergence to a particular equilibrium:

Theorem 7 *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.*

Convergence Proofs

In this section we give proofs of Theorems 4 and 5. Theorem 7 is sketched, and a full proof may be found in Appendix A.1. Throughout this section, we assume that there is a single welfare optimal outcome for clarity of presentation.

Observation 3 *Under assumptions A1 and A2, a bidder will only lower her bid if all bidders are winners.*

Lemma 8 *If all bidders are winners under utility-targets π , then π 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 π . Intuitively, this means that raising bids necessarily implies receiving less utility.

Formally, if bids are $b_i = (v_i, \pi_i)$ 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 $v_i(o^b) - b_i(o^b) = \pi_i$. Moreover, $v_i(o) - b_i(o) \leq \pi_i$ by definition, so $v_i(o) - b_i(o) \leq v_i(o^b) - b_i(o^b)$. Thus

$$\max((v_i(o) - b_i(o)) - (v_i(o^b) - b_i(o^b)), 0) = 0 .$$

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

$$\sum_{i \in [n]} \max((v_i(o) - b_i(o)) - (v_i(o^b) - b_i(o^b)), 0) \leq \sum_{i \in [n]} b_i(o^b) - b_i(o)$$

as desired. ■

Since A1 and A2 imply that a player will only lower her bid from π 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 9 *Under assumptions A1 and A2, if a player lowers her bid from π , then π is in the CEF set \mathcal{C} .*

A corollary of Claim 1 is that any set of CEF bids maximizes welfare, hence this implies that a player will only lower her bid if the welfare-optimal outcome is winning:

Corollary 10 *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 $\bar{\mathcal{C}}$ is rightward-closed:

Lemma 11 *If π is in the CEF set \mathcal{C} , then $\pi - \Delta$ is in the CEF set \mathcal{C} for any $\pi \geq \Delta \geq 0$.*

Corollary 12 *If π is in the not-CEF set $\bar{\mathcal{C}}$, then $\pi + \Delta$ is in the not-CEF set $\bar{\mathcal{C}}$ for $\Delta \geq 0$.*

Proof of Lemma 11. Let b be the bids at π and b^δ be the bids at $\pi - \delta$. Note that Claim 1 implies a welfare-optimal outcome o^* is winning at π .

First, suppose that all bidders for whom $\delta_i > 0$ are winners at π . In this case, $v_i(o^*) \geq \pi_i$ and so $v_i(o^*) \geq \pi_i - \delta_i$ 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 $v_i(o^*) \geq \pi_i - \delta_i$, we can conclude that all bidders are winners, ergo $\pi - \delta \in \mathcal{C}$ by Lemma 8.

Now, suppose some bidders in π may be losers, but that the vector δ has the following property:

$$\delta_i \leq \max(\pi_i - v_i(o^*), 0) .$$

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

$$\sum_{i \in [n]} b_i^\delta(o^*) - b_i^\delta(o) \geq \sum_{i \in [n]} \max((v_i(o) - b_i^\delta(o)) - (v_i(o^*) - b_i^\delta(o^*)), 0) .$$

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

$$b_i^\delta(o) = \max(v_i(o) - \pi_i - \delta_i, 0)$$

which can only be nonzero if $v_i(o) > \pi_i$. However, $b_i(o)$ can only change if $\delta_i > 0$, which requires $\pi_i > v_i(o^*)$ and thus $v_i(o) > \pi_i > v_i(o^*)$. By construction, this also holds for $\pi_i - \delta_i$:

$$v_i(o) > \pi_i - \delta_i \geq v_i(o^*) .$$

Now, when $v_i(o) > \pi_i - \delta_i \geq v_i(o^*)$, we have

$$\begin{aligned} (v_i(o) - b_i^\delta(o)) - (v_i(o^*) - b_i^\delta(o^*)) &= \min(v_i(o), \pi_i - \delta_i) - \min(v_i(o^*), \pi_i - \delta_i) \\ &= \pi_i - \delta_i - v_i(o^*) \\ &\geq 0 . \end{aligned}$$

Importantly, if $\Delta(o)$ is the set of bidders for which $b_i^\delta(o) \neq b_i(o)$, we may conclude that

$$\sum_{i \in [n]} \max((v_i(o) - b_i^\delta(o)) - (v_i(o^*) - b_i^\delta(o^*)), 0) =$$

$$\begin{aligned}
&= \sum_{i \notin \Delta(o)} \max((v_i(o) - b_i^\delta(o)) - (v_i(o^*) - b_i^\delta(o^*)), 0) \\
&\quad + \sum_{i \in \Delta(o)} (v_i(o) - b_i^\delta(o)) - (v_i(o^*) - b_i^\delta(o^*))
\end{aligned}$$

and likewise

$$\begin{aligned}
&\sum_{i \in [n]} \max((v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)), 0) = \\
&= \sum_{i \notin \Delta(o)} \max((v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)), 0) \\
&\quad + \sum_{i \in \Delta(o)} (v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)) .
\end{aligned}$$

The desired CEF condition quickly follows, using the fact that bidders $i \notin \Delta(o)$ did not change their bids:

$$\begin{aligned}
&\sum_{i \in [n]} \max((v_i(o) - b_i^\delta(o)) - (v_i(o^*) - b_i^\delta(o^*)), 0) = \\
&= \sum_{i \notin \Delta(o)} \max((v_i(o) - b_i^\delta(o)) - (v_i(o^*) - b_i^\delta(o^*)), 0) + \sum_{i \in \Delta(o)} (v_i(o) - b_i^\delta(o)) - (v_i(o^*) - b_i^\delta(o^*)) \\
&= \sum_{i \notin \Delta(o)} \max((v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)), 0) + \sum_{i \in \Delta(o)} (v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)) \\
&\quad + \sum_{i \in \Delta(o)} (b_i(o) - b_i^\delta(o)) - (b_i(o^*) - b_i^\delta(o^*)) \\
&= \sum_{i \in [n]} \max((v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)), 0) + \sum_{i \in [n]} (b_i(o) - b_i^\delta(o)) - (b_i(o^*) - b_i^\delta(o^*)) \\
&\leq \sum_{i \in [n]} (b_i(o^*) - b_i(o)) + \sum_{i \in [n]} (b_i(o) - b_i^\delta(o)) - (b_i(o^*) - b_i^\delta(o^*)) \\
&= \sum_{i \in [n]} b_i^\delta(o^*) - b_i^\delta(o)
\end{aligned}$$

as desired.

Finally, for general δ , split it as $\delta = \delta^1 + \delta^2$ where

$$\delta_i^1 = \min(\delta_i, \pi_i - v_i(o^*)) .$$

The vector δ^1 satisfies the condition $\delta_i \leq \max(\pi_i - v_i(o^*), 0)$, so $\pi - \delta^1 \in \mathcal{C}$. Moreover, all bidders are winners in $\pi - \delta^1$, so

$$s - \delta^1 - \delta^2 = s - \delta \in \mathcal{C}$$

as desired. ■

To prove that bids will converge, we first show that bids will not be stuck at arbitrarily low values:

Lemma 13 *Suppose the initial vector of utility-targets is π^0 . Under properties A1 and A2, the auction will always reach a configuration in which all bidders are winners and will do so within $\lceil \frac{1}{\epsilon} \pi^0 \rceil_1$ steps.*

Proof: 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 $\pi_i = 0$ and bidders never decrease their utility-targets when they are winners (A1), utility-target can be decreased at most $\lceil \frac{1}{\epsilon} s \rceil_1$ 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. ■

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 4. Lemma 13 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 π be the vector of utility-targets at any point after the first time all bidders are winners. If all bidders are still winners, then $\pi \in \mathcal{C}$ by Lemma 8. Otherwise, let i be the most recent player to lower her bid, increasing the utility-target vector from π' to $\pi'' = \pi' + \epsilon e_i$. 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, players have only raised their bids since i lowered hers, so we can define $\Delta = \pi'' - \pi$ where $\Delta \geq 0$. Corollary 9 tells us that $\pi' \in \mathcal{C}$. If i raised her bid between π'' and π , then $\pi \leq \pi'$ and Lemma 11 tells us that $\pi \in \mathcal{C}$, so were done. Otherwise, we know $\pi'' \geq \Delta \geq 0$ and Lemma 11 tell us that $\pi' - \Delta \in \mathcal{C}$. Therefore $\pi = \pi' - \Delta + \epsilon e_j \in \mathcal{C}_\epsilon$. ■

To prove Theorem 5, we need a lemma similar to Lemma 13 showing that the auction will reach a bid vector that is CEF:

Lemma 14 *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 $\pi_j = v_j(o^*)$ instead.*

Proof: If the auction reaches a vector π that induces an outcome $o \neq o^*$, then $\pi \in \bar{\mathcal{C}}$ and we are done. Thus, it remains to show that an auction will reach a vector $\pi \in \bar{\mathcal{C}}$ even if the outcome is always o^* .

Consider a bidder j . By A1 we know that j will only decrease π_j if she is a loser and increase π_j if she is a winner. By A3 we can conclude that j will eventually decrease her bid until $\pi_j \geq v_j(o^*) - \epsilon$, 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

$$n\epsilon < \sum_{i \in [n]} \max((v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)), 0)$$

Unfortunately, this implies

$$\sum_{i \in [n]} \max((v_i(o) - b_i(o)) - (v_i(o^*) - b_i(o^*)), 0) > \sum_{i \in [n]} b_i(o^*) - b_i(o) ,$$

and therefore $\pi \in \bar{\mathcal{C}}$.

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 $\pi_j = v_j(o^*)$ (as a result of our restriction that bidders always bid $\pi_j \leq \sup_o v_j(o)$). ■

Proof of Theorem 5. By Lemma 14, the auction will eventually reach a utility-target vector in $\bar{\mathcal{C}}$ 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 $\pi \in \bar{\mathcal{C}}_\epsilon$ from the first time a bid in $\bar{\mathcal{C}}$ is reached.

If $\pi \in \bar{\mathcal{C}}$, we are done, so suppose $\pi \in \mathcal{C}$. Let π' be the most recent utility-targets that were in $\bar{\mathcal{C}}$ and let $\pi'' \in \mathcal{C}$ be the utility-targets immediately after π' . Let i be the bidder who changed her bid between π' and π'' . Corollary 12 implies that i must have raised her bid between π' and π'' .

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 π . Define the utility-target vector $\tilde{\pi}$ with associated bids \tilde{b} as follows:

$$\tilde{\pi}_j = \begin{cases} \min(\pi_i, \pi'_i) & j = i \\ \pi_j - \epsilon & \pi_j > \pi''_j \\ \pi_j + \epsilon & \pi_j < \pi'_j \\ \pi_j & \text{otherwise.} \end{cases}$$

Let $\delta_j = \tilde{\pi}_j - \pi_j''$. We argue later that i will not increase her utility-target from $\pi_i'' = \pi_i' + \epsilon$, so $|\tilde{\pi}_j - \pi_j| \leq \epsilon$ for all j . Thus, it is sufficient to show that $\tilde{\pi} \in \bar{\mathcal{C}}$.

Consider a bidder $j \neq i$ and suppose $\pi_j > \pi_j''$. 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 π'' and π . Since j would only increase her utility-target if she were a winner, she must have been a winner at some value $\geq \pi_j - \epsilon = \tilde{\pi}_j$. Thus, $\tilde{\pi}_j = \pi_j - \epsilon \leq v_j(o^*)$. 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 $\pi_j < \pi_j''$, 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 $\pi_j = \pi_j''$, 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 π_i' to π_i'' increased her bid for o^* more than it increased her bid for o' . This implies $v_i(o') < v_i(o^*)$ and $\pi_i'' < v_i(o^*)$. Consequently, i is a winner with π_i'' at o^* and will not decrease her utility-target further. Firstly, this implies that $|\tilde{\pi}_i - \pi_i| \leq \epsilon$. First, suppose $\pi_i > \pi_i''$. In this case, $\pi_i \geq \pi_i'$, 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 π'' to π , so $\tilde{\pi}_i = \pi_i'$ 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 $\tilde{\pi}$, and therefore $\tilde{\pi} \in \bar{\mathcal{C}}$. By construction, $|\tilde{\pi}_j - \pi_j| \leq \epsilon$, so this implies $\pi \in \bar{\mathcal{C}}_\epsilon$.

So far, we showed that $\pi \in \bar{\mathcal{C}}_\epsilon$ 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 $\pi' \in \bar{\mathcal{C}}$, there is some outcome o^w for which the CEF constraints are violated, i.e.

$$\sum_{i \in [n]} b'_i(o^*) - \tilde{b}'_i(o^w) < \sum_{i \in [n]} \max((v_i(o^w) - \tilde{b}'_i(o^w)) - (v_i(o^*) - \tilde{b}'_i(o^*)), 0) .$$

Observing that the outcome does not change from π' to π , 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

$$\sum_{i \in [n]} \tilde{b}_i(o^*) - \tilde{b}_i(o^w) < \sum_{i \in [n]} \max((v_i(o^w) - \tilde{b}_i(o^w)) - (v_i(o^*) - \tilde{b}_i(o^*)), 0) ,$$

and so $\tilde{\pi} \in \bar{\mathcal{C}}$ and $\pi \in \bar{\mathcal{C}}_\epsilon$. ■

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 player 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 7). The full proof is included in the appendix; we provide a sketch of it here.

Proof Sketch of Theorem 7:

Arrange bidders into levels L_1, \dots, 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. ■

2.7 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 [29]. 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 players 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 players' 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?

Chapter 3

Coopetitive Ad Auctions

A single advertisement often benefits many parties, for example, an ad for a Samsung laptop benefits Microsoft. In this chapter, we study this phenomenon in search advertising auctions and show that standard solutions, including the status quo ignorance of mutual benefit and a benefit-aware Vickrey-Clarke-Groves mechanism, perform poorly. In contrast, we show that a first-price auction (using the utility-target auction of Chapter 2 for multi-slot settings) has well-behaved equilibria in a single-slot ad auction: all equilibria that satisfy a natural cooperative envy-freeness condition select the welfare-maximizing ad and satisfy an intuitive lower-bound on revenue.

3.1 Introduction

In 1991, Intel launched its “Intel Inside” advertising campaign and forever changed the way people buy computers. Previously, buyers only considered hardware insofar as it affected the software that would run on their new machine. The “Intel Inside” campaign aimed to change that behavior — Intel coordinated with PC vendors to advertise not just the processor’s capabilities but the Intel brand. Twenty years later, the “Intel Inside” mark has become one of the most recognized in the tech industry, their signature five-note chime is known worldwide, and, most importantly, buyers think about the brand of processor inside their computers [49].

Intel’s benefit from the “Intel Inside” campaign is an obvious example of a general phenomenon. Intel clearly has a vested interest in the sale of computers containing its products and, in an amortized sense, derives a specific benefit from every sale. This exemplifies a fundamental aspect of marketing: *a single advertisement often benefits many different companies*. Companies commonly recognize this benefit and team up with partners in so-called *cooperative advertising* agreements similar to the “Intel Inside” campaign — in 2000, an estimated \$15 billion was spent on cooperative advertising in the United States alone [72]. This phenomenon is also recognized in the operations research and marketing fields where it has been modeled using a variety of Stackelberg and dynamic games [13, 41]. However, one

key question seems to have gone unasked in both the practical and theoretical realms: *how can the companies who sell advertising space exploit the broad benefit of a single ad?* We study this question in the context of online ad auctions.

Our main results show that an auctioneer may improve both his own revenue and consumers’ welfare by using an auction that allows and encourages cooperation among the advertisers bidding on a single ad but maintains competition between ads — we call this a cooperative¹ ad auction. We first show that conventional cooperative advertising contracts and the Vickrey-Clarke-Groves (VCG) mechanism may perform poorly — Figure 3.1 shows a real query in which Google’s current ad auction produces an unreasonable and unsustainable outcome. In contrast, we show that equilibria of the first-price auction which satisfy a cooperative envy-freeness condition have a natural performance guarantee similar to that of a second-price auction.

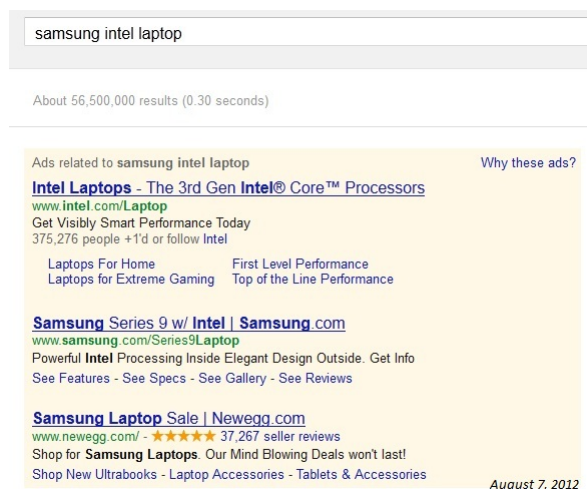


Figure 3.1: A Google search for “samsung intel laptop” illustrates the pitfalls of ignoring the mutual benefits of an ad — the top ad for “Intel Laptops” is competing against the ad below it for “Samsung...w/ Intel.” Google is charging Intel a premium to show its ad on top; however, Intel should be happy if a user clicks on the Samsung ad (and possibly even the Newegg ad) and thus should be unwilling to pay this premium. Even worse, the inclusion of “w/ Intel” in Samsung’s ad suggests that it is subsidized by Intel through the Intel Inside program, further driving up the price of the top slot.

Mutual Benefit and Cooperative Advertising. Most advertisements (directly or indirectly) benefit many parties. For example, an iPhone ad benefits cell phone providers, a Samsung ad for a Windows laptop benefits Microsoft, and an ad for the Boston Red Sox

¹Coopetition is a business term describing an environment where the same parties are simultaneously cooperating in some areas and competing in others [89].

benefits a bar across the street from Fenway Park. Moreover, these secondary benefits are often significant — when Best Buy sells a Samsung laptop running Windows, Microsoft makes more money from the laptop’s sale than either Best Buy or Samsung. As a result, companies have strong incentives to share advertising costs, particularly when competing against a more integrated adversary as when Microsoft and Samsung compete against Apple.

The status quo technique for pooling advertising dollars is an ad-hoc system of external contracts. In such external contracts, one company agrees to pay a portion of another’s advertising costs when its branding is included in the ads. In the Intel example, Intel agrees to pay a percentage of the advertising costs when Samsung or Dell include the “Intel Inside” branding in their ad.

In the research community, cooperative advertising has mostly been studied in marketing and operations research. The setting is typically modeled as a Stackelberg game in which an upstream manufacturer makes an offer to downstream manufacturers or retailers [13, 55], sometimes incorporating dynamic components [41]. Other recent research has relaxed the assumption that the upstream manufacturer is the first mover and considers external contracts based on alternative bargaining solutions [47].

Pay-Per-Click Ad Auctions. Search advertising today is sold through a pay-per-click (PPC) ad auction. In the standard setting, each bidder comes to the auction with its own ad and places a bid in terms of its willingness to pay for each click. The auctioneer subsequently assigns ads to slots on the web page — bidders effectively compete for slots — but only charges a bidder when his ad is clicked.

In a standard PPC ad auction, the only way for advertisers to share costs is through external contracts. Unfortunately, this creates undesirable results in an auction format. First, a manufacturer may compete with itself. For example, Figure 3.1 shows a Google query for “samsung intel laptop” in which an ad from Intel is shown above a Samsung ad that explicitly advertises Intel-based laptops. In general, Intel will be happy if the buyer visits Samsung’s site with the intent to buy an Intel-based laptop. Thus, even if Intel would prefer the user to click on its own ad instead of Samsung’s, Intel should be unhappy paying the premium required to beat Samsung in the auction. Even worse, Intel may be subsidizing Samsung’s ad, further increasing the price Intel pays in a GSP auction.

A second downside to external cooperative advertising agreements is that downstream producers may face a moral hazard. In some situations, the downstream producer may be incentivised to overspend on advertising and waste the upstream manufacturer’s money. In the worst case, the upstream manufacturer will refuse to participate and cooperation will collapse entirely. We give examples of both phenomena in Section 3.3.

The theory of mechanism design suggests another approach: bids should express complex preferences over ads. In particular, as it is possible for multiple bidders to value a single ad and for a single bidder to value multiple ads, bidders should be able to express these preferences. This however is not enough to guarantee good behavior: we show in Section 3.3 that a Vickrey-Clarke-Groves (VCG) auction may generate little or no revenue — the implicit

cooperation among advertisers (which is desired) mimics the kind of strategic collusion that is known to reduce revenue in VCG mechanisms.

The First-Price Coopetitive Ad Auction. In contrast to the pitfalls of external contracts and VCG mechanisms, we use techniques from Chapter 2 to show that a first-price auction that consider bidders' complex preferences will have nice equilibria. (We focus on the special case of a single-slot ad auction in this chapter.) First, we generalize the intuition of a second-price auction to give a natural lower-bound on the revenue that the auctioneer should expect. Second, analogous to the results of Edelman et al. [29] for the commonly used generalized second-price (GSP) auction, we show that all equilibria satisfying a natural cooperative envy-freeness condition maximize welfare while satisfying this revenue lower-bound. We also show that the cooperative envy-free equilibria dominate VCG in terms of revenue — each bidder is individually paying more in the first-price auction than in the VCG mechanism.

Next, we show that such cooperative envy-free equilibria can be found easily. The envy-freeness constraints define a polytope of which the equilibria form the Pareto frontier, suggesting a family of convex programs for computing equilibria. Finally, we specifically identify the egalitarian equilibrium and give an efficient algorithm for computing it.

As discussed in Chapter 2, first-price auctions are often preferable to VCG and GSP auctions. Most significantly, they are more transparent — since prices precisely correspond to bids, bidders do not face uncertainty in their payments and there is no opportunity for the auctioneer to manipulate the auction, particularly by learning in a repeated setting. For example, eBay historically offered a VCG-like auction for selling multiple identical items. The auction was sufficiently disliked by the bidders that eBay no longer offers it as an option. Today, eBay sells single items through a transparent ascending price auction. In the specific case of ad auctions, Overture's experience has discredited the first-price auction; however, our work in Chapter 2 showed how Overture's auction might have been fixed.

Related Work. In addition to the afore-mentioned work on cooperative advertising and online ad auctions, our work is related to auctions with externalities. Whereas simple auction models assume players are indifferent to the bundle received by another player, in reality there may be externalities, i.e. players may care about the bundles received by other agents. Incorporating externalities has been studied in both the economics and computer science literatures (e.g. [61, 53]), typically producing a mechanism in which bidders can express a different value for every possible outcome. Our ad auctions may not always have nice interpretations in terms of these externalities; however, like mechanisms with externalities they cannot be expressed in the standard bidding language and, in the extreme, degenerate to a mechanism which requires bidders to express a value for every possible outcome.

3.2 The Coopetitive Ad Auction

The coopetitive advertising model generalizes the standard pay-per-click (PPC) advertising auction. As in the standard auction, the auctioneer must choose which of m competing ads to show in s slots. Each advertiser derives a value of v from a click on one of its ads and has a utility that is quasilinear in money, i.e. $u(p) = v - p$ when the advertiser gets a click and pays p . The likelihood that a user clicks on ad j in slot k , called the click-through-rate (CTR), is given by $c_{j,k}$. Hence, the expected utility of bidder whose ad is shown in slot with CTR c is $\mathbf{E}[u_i(p_i)] = c(v - p)$. In this paper, we focus on the special case with one slot ($s = 1$) where CTRs are independent of the ad. In this case, the CTR can be taken to be 1 without loss of generality, so henceforth we ignore them. All the results in this chapter can be extended to the multi-slot setting using utility-target auctions introduced in Chapter 2.

The new feature of the coopetitive model is that an advertiser can derive value from clicks on multiple ads. In general, advertiser i 's value for a click $v_{i,j}$ depends on the particular ad j ; however, for the sake of presentation we consider a simpler model. An advertisement $j \in \{1, \dots, m\}$ is defined by a publicly-known set of advertisers $S_j \subseteq [n]$ who all derive value from a click on advertisement j . Advertiser i derives the same value v_i from a click on any ad j where $i \in S_j$ (an advertiser does not benefit if $i \notin S_j$). We use T to denote the set of bidders in the ad with the maximum total value, i.e. $T = \operatorname{argmax}_{S_j} \sum_{i \in S_j} v_i$ (in case of a tie, T denotes the particular winning ad chosen by the auction). All our results generalize to the more complicated $v_{i,j}$ setting through the utility-target auction, where bids effectively encode the same utility $v_{i,j} - b_{i,j}$ for each ad.

We will often use shorthand for examples. Our notation is itself best described by an example:

$$\{(A^2, B^1, C^3), (A^2, B^1), (C^3)\}$$

This denotes an auction with three advertisements and three advertisers (A , B , and C). Each advertiser derives value when the first ad (A^2, B^1, C^3) is shown, while only advertisers A and B benefit from the second ad and only advertiser C benefits from the third. In this example, advertisers are indifferent between the ads they are involved with, with values of 2, 1, and 3 respectively when an ad of theirs is shown.

The External Contracts Mechanism. External contracts are the way advertisers currently cooperate: each ad is “owned” by a single bidder, and any party wishing to increase the bid on an ad must negotiate an external contract with the ad’s owner.

We model this as a standard VCG PPC ad auction in which advertiser i can (before the auction is run) commit to pay an α_i fraction of the cost each time one of its ads is clicked, up to a maximum β_i . This payment goes directly to the owner $o(j)$ of the clicked ad j . Thus, the utility of advertiser i would be $u_i = v_i - \min(\alpha_i p_{o(i)}, \beta_i)$, while the utility of the ads owner $o(j)$ would be $u_{o(j)} = v_{o(j)} - p_{o(j)} + \min(\alpha_i p_{o(j)}, \beta_i)$.

The VCG Mechanism. In our setting with a single slot, the Vickrey-Clarke-Groves (VCG) mechanism chooses the ad j maximizing $\sum_{i \in S_j} v_i$ and charges bidder i his externality. When the bidder values all ads the same, this will be the minimum value he needed to have reported to be in the winning ad. (Bidders submit their true values v_i because VCG is incentive compatible.)

The First-Price Auction. In a first-price auction, each advertiser submits a bid b_i . The auctioneer displays the ad j maximizing $\sum_{i \in S_j} b_i$ and charges each bidder in the winning ad $p_i = b_i$ when the ad is clicked.

3.3 Pitfalls of Standard Mechanisms

Many standard mechanisms behave poorly with respect to advertisers' cooperative valuations. For example, Figure 3.1 shows how Google's current system caused Intel to compete against itself. In this section, we will give further examples detailing the poor behavior of the external contracts mechanism (the status quo) and the Vickrey-Clarke-Groves mechanism.

External Contracts. A few pitfalls specifically arise in the current system of external contracts. First, if contracts are made with insufficient granularity, an advertiser might easily compete with itself:

Example 1 Consider the following two single-slot ad auctions:

$$\{(S^3, M^{10}), (D^2, M^{10}), (A^{11})\} \quad \text{and} \quad \{(S^3, M^{10}), (D^2, M^{10})\} .$$

In the first ad auction, M will happily contribute advertising funds to help beat A ; however, in the second auction M wins regardless. The only effect of M 's dollars in the second auction is to fund a useless bidding war between the SM and DM ads, so it should not offer cooperative advertising contracts to S and D .

Ideally, M would only contribute advertising funds in the first auction. However, since the granularity of real cooperative advertising contracts is somewhat limited, this example demonstrates a legitimate concern.

Additionally, the advertiser receiving the external contract faces a moral hazard: he is often incentivized to overspend the money of his advertising partner. In equilibrium, the result is that cooperation collapses:

Example 2 Consider the following three ads with four interested parties:

$$\{(S^3, M^{10}), (D^2, M^{10}), (A^{11})\} .$$

Suppose there are three ad slots with CTR's 0.1, 0.08, and 0.05 respectively (if a bidder appears in multiple ads, their likelihood of a click is the sum of the likelihoods for those ads).

In the external contracts model described in Section 3.2, M will not offer a cooperative advertising contract in equilibrium. As a result, the auction degenerates to $\{(S^3), (D^2), (A^{11})\}$. Not only will revenue decrease substantially, but the auction will be inefficient because (A^{11}) wins the top slot.

We omit the calculations.

The VCG Mechanism. The VCG mechanism charges a player based on the externality it imposes on other users, i.e. the welfare that others lose because of its presence. A downside of the VCG mechanism is that it may not generate any revenue. It is well-known that collusion has a negative effect on revenue in the VCG mechanism [7]. While such collusion can be illegal in other settings, in cooperative ad auctions we specifically want advertisers to cooperate on ads that are of mutual benefit — hence we cannot assume that they will not collude.

For example, two players can make their payments zero by simultaneously claim sufficiently large values for the winning outcome o . If the players' bids are sufficiently large that o is still the welfare-maximizing outcome even if one of the pair were removed, the externality that each player imposes is zero, and nobody pays anything. This happens easily in cooperative ad auctions. For example:

Example 3 *In the single slot ad auction $\{(A^1, B^1, C^1, D^1), (E^{2-9})\}$, the VCG mechanism will show the first ad in the slot, but nobody pays anything.*

In this case, the $ABCD$ ad remains the best ad even if a single bidder is removed. As a result, nobody pays anything. Such a scenario can occur naturally if the winning ad is valued by many small bidders.

The weak revenue of the VCG mechanism is not limited to extremes like the above example. In general, payments will be lower than an auctioneer might hope:

Example 4 *In the single slot ad auction $\{(A^2, B^2), (E^3)\}$, the VCG mechanism will show the first ad in the slot and players A and B will each pay 1.*

Intuition says that the auctioneer should hope to make A and B pay a total of 3, since that is the total bid required to beat E . However, the total VCG payment for the first ad is only 2. In contrast, we show that a cooperative first-price auction will indeed generate a revenue of 3.

3.4 Equilibria of the First-Price Auction

In this section, we show that the equilibria of first-price cooperative ad auctions have desirable revenue and welfare properties. The straightforward results we prove are special-case adaptations of more general theorems in Chapter 2, where we show how the same properties

can always be guaranteed when a first-price auction is implemented with an appropriate bidding language. We focus on the single-slot setting to illustrate intuition; however, analogous results can be shown for multi-slot settings using the utility-target auction framework.

First-price auctions may have many equilibria — we will focus on equilibria that satisfy *cooperative envy-freeness* as introduced in Chapter 2. The intuition for expecting an envy-free outcome in a repeated auction is simple: if a group of losing bidders (bidders who do not get value from the winning ad) can unilaterally increase their utilities by increasing their bids, then over time it is reasonable to expect that all bidders in the losing group will eventually raise their bids appropriately even if they do not explicitly collude.

More formally, cooperatively envy-free equilibria are ones in which cooperating partners cannot jointly raise their bids to beat out the presently winning ad without overpaying (much like the core in cooperative game theory). The interpretation of CEF equilibria in this setting is the following:

Definition 11 *The bids $(b_i)_{i \in T}$ for the winning ad T are cooperatively envy-free (CEF) if and only if for all alternate ads S_j ,*

$$\sum_{i \in T \setminus S_j} b_i \geq \sum_{i \in S_j \setminus T} v_i \quad (3.1)$$

We also insist that the agents bidding are individually rational (IR) and bids are non-negative, that for each advertiser i , $0 \leq b_i \leq v_i$.

Combining IR and CEF yields efficiency - if the bids of agents in $T \setminus S_j$ are at least the values of agents in $S_j \setminus T$, then so too are the values.

Lemma 15 *If the bids $(b_i)_{i \in T}$ for the winning ad T are IR and CEF, then T is the efficient winning ad.*

These CEF, IR and non-negativity conditions form a polytope of possible payments associated with the correct winning ad. Not all of these are equilibria - those will instead form the Pareto frontier of the polytope. On this frontier, there can still be many possible equilibria. For instance, consider the following ad auction with three advertisers and three outcomes: $\{(A^{100}, B^{100}), (C^{99})\}$. Every set of bids $b_A = x, b_B = 99 - x$ for $0 \leq x \leq 99$ constitute an equilibrium.

Naturally, the constraints defining this polytope come from other ads — for any bidder (dimension), there must be an alternative ad that would win if she lowered her bid. We will use this structural result later.

Lemma 16 *The IR, CEF, non-negative bids (b_i) form an equilibrium for the winning ad T if and only if for each bidder k , $b_k = 0$ or there exists an ad $S_j \not\ni k$ s.t.*

$$\sum_{i \in T} b_i = \sum_{i \in S_j} b_i \quad (3.2)$$

Proof: First, the ‘if’ direction. Assume such an S_j exists for every k . Then, were k to lower his bid, S_j would win and k would no longer be in the winning set.

Consider the other direction. Assume a set of bids equilibrium bids are CEF, IR and non-negative. For every winning advertiser k , they must not be able to lower their bids and still win - otherwise they would, and we would not be in equilibrium. Thus, for any k s.t. $b_k > 0$, there must be such a set S_j . ■

Revenue

In this section, we consider the revenue behavior of equilibrium points in the polytope. First, note that the revenue is not the same for all equilibrium points - this is not simply a matter of dividing a fixed payment up. As an example, consider the following three ad, five interested party setting: $\{(A^1, B^1, C^1), (A^1, D^1), (B^1, E^1)\}$. Clearly the first ad should win, but what should the payments be?

Our CEF, IR and non-negativity conditions give us the following polytope: $b_A + b_C \geq 1$, $b_B + b_C \geq 1$, $0 \leq b_A, b_B, b_C \leq 1$. The set of equilibrium points includes $(1, 0, 1)$, $(0, 1, 0)$ and every convex combination of the two. Thus, the revenue of the equilibrium points can range from 1 to 2.

How would other mechanisms do? VCG will charge nothing, as no advertiser is integral to the ad being displayed. Were A and B to lie and say they are not affiliated with C or D , any second or first price mechanism would insist on a payment of 1 from the three of them. In this example then, the revenue of our first price equilibria are lower bounded by VCG, and by simpler first and second price auctions.

Lemma 17 *In the first price cooperative ad auction, the bid of each advertiser is at least their VCG payment.*

Proof: We provide a simple proof here. Consider advertiser i . Let T be the winning ad, and let S_j be the winning ad without i . If $S_j = T$, then i 's VCG payment is 0, and hence we need only worry about the case that $S_j \neq T$. By the CEF constraints, we have $b_i \geq \sum_{k \in S_j \setminus T} v_k - \sum_{k \in S_j \setminus T - \{i\}} b_k \geq \sum_{k \in S_j \setminus T} v_k - \sum_{k \in S_j \setminus T - \{i\}} v_k$. The latter quantity is exactly i 's VCG payment, and hence every advertiser's bid is at least their VCG payment. ■

Another revenue benchmark is the “second-price threat”. That is, we could imagine being in a second-price auction in which the winning bidders collude to only profess their desire for the winning outcome, and pretend to be uninterested in the other outcomes. In this case, a second price-esque auction would have them pay the maximum of the values of non-winning players in non-winning ads.

Any CEF equilibria will get at least the revenue of this outcome.

Definition 12 *The second-price threat is defined as the maximum sum of non-winning values in a non-winning ad. That is, given a winning ad T , the second price threat is:*

$$\max_{S \neq T} \sum_{i \in S \setminus T} v_i.$$

Lemma 18 *The revenue of any CEF equilibrium is at least the second price threat.*

This follows directly from the CEF conditions. Note that this is the same as the revenue a second price auction would get if treated advertisements as single agents and removed interests of winning advertisers in losing ads.

Thus, any first price equilibria that satisfies CEF (eg., no losing advertisers can collaborate to increase their bids and win) has good revenue - revenue that beats both VCG and a natural analogue to a second price auction.

3.5 The Egalitarian Solution

As discussed earlier, there are many ways the winners can split payments while still satisfying our equilibrium and cooperative envy-freeness constraints. In a first price auction, the exact split that we expect bidders to reach will depend on the preferences and bidding dynamics of the advertisers.

In this section, we consider one such split in particular — the *egalitarian bargaining solution*. In the egalitarian solution, the utility of the worst-off player is maximized, and on up the line. This can be defined as the equilibrium with the lexicographically maximum utility:

Definition 13 *The egalitarian solution in the cooperative first-price auction is an equilibrium that displays the highest surplus ad and charges advertisers so that the utility vector is lexicographically maximal when bidders are ordered in terms of increasing utility.*

In many normal settings, this will result in all players sharing equally the surplus generated.

Algorithm

The egalitarian equilibrium can be computed by repeatedly lowering bids uniformly. The algorithm is described in Algorithm 4.

Lemma 19 *Algorithm 4 computes the egalitarian equilibrium point.*

Proof:

First, we show that the resulting point is efficient and cooperatively envy-free — then we'll show that the resulting equilibrium must be the egalitarian one.

Begin with the following claim:

ALGORITHM 4: An algorithm for computing the egalitarian equilibrium in the single-slot first-price auction.

input : A coopetitive ad auction problem.

output: The egalitarian equilibrium bids b_i .

- 1 Set the bids of all advertisers to their values. Call T the winning ad.
 - 2 Lower bids of advertisers in the winning ad T uniformly until some bidder i reaches $b_i = 0$ or a constraint $\sum_{i \in T} b_i \geq \sum_{i \in S_j} b_i$ would be violated for some ad S_j .
 - 3 Fix the bids of advertisers in $T \setminus S_j$ (or fix the bid of i if b_i reached 0).
 - 4 Repeat (2) and (3), lowering only unfixed bids until all bidders in T are fixed.
-

Claim 6 *The total bid for the winning ad T never drops below the total bids of non-winning ads.*

At any point, the algorithm lowers only the bids of players in *every* ad tied for the highest value. As a result, every ad tied for the highest value is decreasing by the same amount. Thus, the winning ad at the beginning of the algorithm remains the winning ad at the end and hence the final winning ad is the ad with the most surplus.

At the end of the auction, as T remains the ad with the highest value and no non-winning advertisers see their utilities decrease, the CEF constraints will be satisfied for every alternate ad S_j .

By Lemma 16 we know that our point will be in equilibrium if and only if there is a set for every agent tied with the winning set that he is not in. Our algorithm will only stop when for every bidder there is such a set, or they are bidding 0 — hence it must be such an equilibrium.

Thus, we've argued that the final point is an efficient, CEF point. Now, we discuss whether or not it is in fact the egalitarian solution. We'll prove this with induction. Assume that the algorithm gives the egalitarian bargaining solution for the first $i - 1$ lowest utility advertisers. Consider the algorithm after those advertisers are fixed - in particular, the next time an advertiser has their bid fixed, with a utility of z . At this point, there will be a set S_j s.t.

$$\sum_{i \in T \setminus S_j} \max(v_i - z, b'_i) = \sum_{i \in T \setminus S_j} v_i \quad (3.3)$$

First, note that the algorithm cannot reduce the bid of i further than the egalitarian solution — otherwise that would be the egalitarian solution. Assume now that the algorithm results in a lower utility for player i — hence, i 's bid is fixed before being lowered to his egalitarian utility. By our assumption, then, $\sum_{i \in T \setminus S_j} \max(v_i - z, b'_i) > \sum_{i \in T \setminus S_j} b'_i \geq \sum_{i \in T \setminus S_j} b_i$. By our CEF constraints in the egalitarian solution, we have that $\sum_{i \in T \setminus S_j} b_i \geq \sum_{i \in T \setminus S_j} v_i$. Then $\sum_{i \in T \setminus S_j} \max(v_i - z, b'_i) < \sum_{i \in T \setminus S_j} b_i$ and hence $\sum_{i \in T \setminus S_j} b'_i = \sum_{i \in T \setminus S_j} b_i$. Since all advertisers with utility less than z have the egalitarian utility, and all advertisers with more utility are in S_j , $b'_i = b_i$. ■

3.6 Conclusion and Open Problems

As we have discussed, a wide variety of ads provide value to more than one party — ads for computers, cell phones, and even baseball teams to name a few. As evidenced by the “samsung intel laptop” query shown in Figure 3.1, current ad auctions completely fail to account for such shared benefits. While Google may enjoy a little extra revenue at Intel’s expense in the interim, Figure 3.1 does not represent a sustainable equilibrium — our theoretical results show that failure to account for the shared benefit of an advertisement can have a substantial negative effect on both welfare and revenue.

As a possible solution, we showed that the coopetitive first-price auction behaves well in equilibrium. In particular, all equilibria that satisfy cooperative envy-freeness are efficient and have good revenue. Moreover, such equilibria can be computed efficiently. Finally, as noted in Section 3.2, these results generalize to a model where bidders’ have different values for different ads and multiple slots through the utility-target auction.

The main open question is *what is the best way for auctioneers to accommodate the shared value of an ad?* While the first-price auction may have nice equilibria and offer transparency to bidders, it is not clear if it is the best solution. On the other hand, the VCG auction may perform poorly, and naïvely building a GSP-style auction will either inherit the downsides of VCG or allow players to affect the distribution of payments by tweaking their bids. Thus, while the properties of first-price equilibria presented herein are desirable, the question of how auctioneers should design real auctions remains unresolved.

Chapter 4

Single-Call Mechanisms

Truthfulness is fragile and demanding. It is oftentimes harder to guarantee truthfulness when solving a problem than it is to solve the problem itself. Even worse, truthfulness can be utterly destroyed by small uncertainties in a mechanism's outcome. One obstacle is that *truthful payments depend on outcomes other than the one realized*, such as the lengths of non-shortest-paths in a shortest-path auction. Single-call mechanisms are a powerful tool that circumvents this obstacle — they implicitly charge truthful payments, guaranteeing truthfulness in expectation using only the outcome realized by the mechanism. The cost of such truthfulness is a trade-off between the expected quality of the outcome and the risk of large payments.

In this chapter, we study two of the most general domains for truthful mechanisms and largely settle when and to what extent single-call mechanisms are possible. The first single-call construction was discovered by Babaioff, Kleinberg, and Slivkins [9] in single-parameter domains. They give a transformation that turns any monotone, single-parameter allocation rule into a truthful-in-expectation single-call mechanism. Our first result is a natural complement to [9]: we give a new transformation that produces a single-call VCG mechanism from any allocation rule for which VCG payments are truthful. Second, in both the single-parameter and VCG settings, we precisely characterize the possible transformations, showing that a wide variety of transformations are possible but that all take a very simple form. Finally, we study the inherent trade-off between the expected quality of the outcome and the risk of large payments. We show that our construction and that of [9] simultaneously optimize a variety of metrics in their respective domains.

Our study is motivated by settings where uncertainty in a mechanism renders other known techniques untruthful. As an example, we analyze pay-per-click advertising auctions, where the truthfulness of the standard VCG-based auction is easily broken when the auctioneer's estimated click-through-rates are imprecise.

4.1 Introduction

In their seminal work that sparked the field of Algorithmic Mechanism Design, Nisan and Ronen [73] made a striking observation: naively computing VCG payments for shortest-path auctions requires computing “ n versions of the original problem.” In their case, it requires solving $n + 1$ different shortest path problems in a network. Over the next decade, as researchers studied computation in mechanisms, they repeatedly noticed that computing payments is harder than solving the original problem. Babaioff et al. [11] exhibited a problem for which deterministic truthfulness is precisely $(n + 1)$ -times harder than the original problem. In the case of Nisan and Ronen’s own path auction, Hershberger et al. [44] showed that computing VCG prices for a directed graph requires time equivalent to \sqrt{n} shortest path computations.¹

Surprisingly, Babaioff, Kleinberg, and Slivkins [9] recently showed that randomization eliminated this difficulty for a large class of problems. They showed that, if in a single-parameter domain payments need only be *truthful in expectation*, then they may be computed by solving the original problem only once. They apply their result to Nisan and Ronen’s path auctions to get a truthful-in-expectation mechanism that uses precisely one shortest-path computation and chooses the shortest path with probability arbitrarily close to 1. We call this a *single-call mechanism*.

The usefulness of Babaioff, Kleinberg, and Slivkins’ result goes far beyond speeding up computation: *Their construction enables truthfulness in cases in which computing “ n versions of the original problem” is informationally impossible.* To use again the Nisan-Ronen path auction, suppose that the graph represents a packet network with existing traffic. In this case, the actual transit times (i.e. costs to edges) may be increased by congestion. While it is possible to estimate congestion ex ante, it is generally impossible to precisely know its effect without transmitting a packet and explicitly measuring its transit time. Unfortunately, since VCG prices depend on the transit times for many different paths, naively computing them will inherit any estimation errors. Even worse, when bidders have conflicting beliefs about such errors, *naively computing “VCG” prices with bad estimates may not guarantee truthfulness* even if the errors are small enough that they not affect the path chosen by the mechanism. In such a case, truthfulness may be regained using a mechanism that only requires measurements along a single path, that is, a mechanism that only requires measurements returned by a single call to the shortest-path algorithm. We will concretely demonstrate this phenomenon later using an example based on pay-per-click advertising auctions.

An important question arises then: *In which mechanism design problems, and to what extent, are single-call mechanisms possible?* In this paper we study, and largely settle, this question for two large domains of truthful mechanisms. First, we show that this it is possible

¹ Interestingly, the undirected case is easier. Hershberger and Suri [43, 42] show that it only requires time equivalent to a single shortest-path computation. Their work is orthogonal to our own — single-call mechanisms achieve truthfulness in a limited-information setting using only one shortest-path computation, while [44, 43, 42] assume complete information and study an algorithmic problem.

to transform any mechanism that charges VCG prices in expectation into a roughly equivalent single-call mechanism. While similar in spirit to [9], our reduction charges prices that are fundamentally different from the mechanism in that paper — they do not coincide even when applied to the same allocation rule. Second, we give characterization theorems, delineating precisely the single-call mechanisms that are possible, for both the VCG and single-parameter settings. Finally, single-call constructions offer a tradeoff between expectation and risk. Our characterization theorems allow us to derive lower bounds on this tradeoff, establishing that our VCG construction and the construction of [9] are optimal in a general sense.

Mechanisms, Allocations, and Payments One cornerstone of mechanism design is the decomposition of a mechanism into two distinct parts: an allocation function and a payment function. This approach has borne much fruit — it first revealed fundamental relationships between allocation functions and their nearly unique truthful prices, and it subsequently allowed researchers to study the two problems in isolation. Like [9], we leverage this decomposition to study payment techniques that apply to large classes of allocation functions — naturally, our primary requirement is that the allocation function may only be evaluated once.

We will focus on single-call mechanisms for two classes of allocation functions that, together, comprise most allocation functions for which truthful payments are known: *monotone single-parameter functions* and *maximal in distributional range (MIDR) functions*.

An allocation function is said to be *single-parameter* if an agent’s bid can be expressed as a single number. This setting was first studied by Myerson [71] in the context of single-item auctions. Subsequent generalizations showed that truthful prices existed if and only if a single-parameter allocation is monotone and provided an explicit characterization of truthful payments. We will use one such characterization developed by Archer and Tardos [2].

An allocation function is said to be *maximal in distributional range (MIDR)* if, for some fixed set of distributions over outcomes, the allocation always chooses one that maximizes the social welfare of the bidders. MIDR allocation functions are important because they are precisely the ones for which VCG payments are truthful [26].

Truthfulness Under Uncertainty Our motivation for developing and optimizing single-call mechanisms comes from scenarios where nature prohibits computing an allocation more than once, most often due to parameter uncertainty. We give a few examples here; more generally, we conjecture that most mechanism design problems have similar variants.

In the uncertain shortest-path auction described earlier, truthful prices will depend on the incremental effect of transit times adjusted for congestion. If the auctioneer generates the network traffic, he may be able to predict the congestion in an edge better than the edge itself and use this prediction when computing the shortest path. However, each edge may individually disagree with the auctioneer’s estimate, and these beliefs are generally unknown to the auctioneer. If the auctioneer were to simply compute VCG payments by combining his estimates with players’ bids, the prices would likely not be truthful. On the other hand, we

can require that payments are computed using measured transit times instead of estimates; however, it is informationally impossible to know the precise delay along edges that were not actually traversed. A single-call mechanism sidesteps this hurdle by using only the delays along traversed edges for which the delay had been precisely known.

Machine scheduling offers another application for single-call mechanisms. In some applications (e.g. cloud services), it is common for machines to bid in terms of cost per unit time (or other resource). It is then the responsibility of the scheduler to estimate the time required for the job on that machine. If the scheduler’s estimates differ from a machine’s belief about a job’s runtime, then we find ourselves in the same situation as the path auction — the standard truthful prices for this single-parameter setting will depend on machines’ beliefs about the runtimes of jobs under alternate schedules. A single-call mechanism sidesteps this problem because it requires only the runtimes of jobs under the schedule chosen by the mechanism, which may be measured.

Another interesting example arises in the application of learning procedures such as multi-arm-bandits (MABs). In recurring mechanisms, it is natural for the auctioneer to run a learning algorithm across multiple auctions. For example, when an online advertising auction is repeated, the auctioneer tries to learn the likelihood that a particular ad will get clicked. Computing truthful prices requires knowing what would have happened if the learner had been initialized with a different set of bids. This setting was the original motivation of [9], where they showed that their single-call construction allowed a MAB to be implemented truthfully with $O(\sqrt{T})$ regret. This contrasts with results of Babaioff, Sharma, and Slivkins [10] and Devanur and Kakade [24] who showed that any universally truthful mechanism must have regret at least $\Omega(T^{\frac{2}{3}})$ for different measurements of regret.

Finally, in Section 4.5 we analyze *single-shot pay-per-click (PPC) advertising auctions*. A PPC advertising auction ranks bidders using their pay-per-click bid (i.e. they only pay when they receive a click) and an estimate of the probability of a click (the click-through rate, or CTR). If the bidders’ estimates of their own CTRs are different from the auctioneer’s, truthful prices necessarily depend on bidders’ beliefs about the CTRs, which are unknown.

Single-Call Mechanisms and Reductions The informational limitations of single-call mechanisms are formalized through the concept of a single-call reduction, the main object of study in this paper. A *single-call reduction* is a transformation that takes an allocation function as a black-box and produces a truthful-in-expectation mechanism that calls the allocation function once. Since the expected payment is equal to the truthful payment for the resulting mechanism, the payments are dubbed *implicit*.

Babaioff, Kleinberg, and Slivkins [9] discovered such a reduction for single-parameter domains. Using only the guarantee that the black-box allocation rule is monotone, their reduction produces a truthful-in-expectation mechanism that implements the same outcome as the original allocation rule with probability arbitrarily close to 1.²

²The authors of [9] have observed that their construction may be extended to any domain where the bid space is convex.

VCG is a mechanism design framework much broader than single-parameter. *Can we construct similar single-call mechanisms that charge VCG prices?* We answer this in the affirmative by giving a reduction producing, for any MIDR allocation function, a single-call mechanism that charges VCG prices in expectation. Analogous to [9], our reduction transforms any MIDR allocation rule into a truthful-in-expectation mechanism that implements the same outcome as the original allocation rule with probability arbitrarily close to 1. However, our construction is fundamentally different in that the distribution of payments does not coincide with [9] when an allocation is both MIDR and single-parameter. This reduction can guarantee truthfulness in multi-parameter mechanisms with uncertainty, as described above, and can also be used to speed up payment computation in MIDR settings like Dughmi and Roughgarden’s [27] truthful FPTAS for welfare-maximization packing problems.

We next ask *what single-call reductions are possible?* Babaioff et al. generalize to a class of self-resampling procedures. Subsequent research [39] generalized further (and simplified substantially), but concisely characterizing single-call reductions remained an open question. We give tight characterization theorems, showing that a wide variety of reductions are possible and that payments have a very simple characterization in both scenarios. The key technical idea is a simple proof equating a reduction’s expected payments with those required for truthfulness, giving a sharp characterization of the parameters in the reduction. Our technique is a very simple alternative to the contraction mapping argument in [9].

Finally, we ask *what are the best single-call reductions?* As noted above, known single-call reductions choose an outcome different from the original allocation rule with some small probability δ . The penalty for making δ small is that the payments may occasionally be very large — we study this tradeoff. Our study is not unprecedented: [9] asked, as an open question, if their reduction optimized payments with respect to the welfare loss, and Lahaie [62] show a similar tradeoff between the size and complexity of kernel-based payments achieving ϵ -incentive compatibility in single-call combinatorial auctions.

We study the tradeoff inherent to single-call mechanisms with respect to three measures of expectation — welfare, revenue, and a technical (but natural) precision metric — and two measures of risk — variance and worst-case payments. We show that our VCG reduction and the single-parameter reduction of [9] *simultaneously* optimize the tradeoff between expectation and risk for all these criteria.

Subsequent Work In recent follow-up work, Babaioff, Kleinberg, and Slivkins [8] showed how their reduction for single-parameter domains could be extended to arbitrary convex type-spaces. They use their reduction to design a truthful multi-armed bandit for a multi-parameter ad auction setting.

4.2 Preliminaries

A mechanism is a protocol among n rational agents that implements a social choice function over a set of outcomes \mathcal{O} . Agent i has preferences over outcomes $o \in \mathcal{O}$ given by a *valuation*

function $v_i : \mathcal{O} \rightarrow \mathbf{R}$. The function v_i is private but is drawn from a publicly known set $V_i \subseteq \mathbf{R}^{\mathcal{O}}$.

A *deterministic direct revelation mechanism* \mathcal{M} is a social choice function $A : V_1 \times \dots \times V_n \rightarrow \mathcal{O}$, also known as an *allocation rule*, and a vector of payment functions p_1, \dots, p_n where $p_i : V_1 \times \dots \times V_n \rightarrow \mathbf{R}$ is the amount that agent i pays to the mechanism designer. When a direct revelation³ mechanism is instantiated, each agent reports a *bid* $b_i \in V_i$. The mechanism uses bids $b = (b_1, \dots, b_n)$ to choose an outcome $A(b) \in \mathcal{O}$ and to compute payments $p_i(b)$. The utility $u_i(v_i, o)$ that agent i receives is $u_i(v_i, o) = v_i(o) - p_i$.

A mechanism is *truthful* (or incentive compatible) if bidding truthfully (i.e. $b_i = v_i$) is a dominant strategy. Formally, for each i , each $v_{-i} \in V_{-i}$, and every $v_i, v_i' \in V_i$, we have $u_i(v_i, A(v)) \geq u_i(v_i', A(v_i', v_{-i}))$, where v_{-i} denotes the vector of valuations for all agents except agent i .

A mechanism is *ex-post individually rational* (IR) if agents always get non-negative utility, and mechanism has *no positive transfers* (NPT) if for each agent i and each $v \in V$, $p_i(v) \geq 0$, i.e., the mechanism never pays a player money.

A randomized mechanism is a distribution over deterministic mechanisms. Thus, $A(b)$ and $p_i(b)$ are random variables. For randomized mechanisms, properties like truthfulness may be said to hold universally or in expectation. A randomized mechanism is *universally truthful* if it is truthful for every deterministic mechanism in its support. It is *truthful in expectation* if, in expectation over the randomization of the mechanism, truthful bidding is a dominant strategy. Henceforth, we use truthful, IR, and NPT to mean truthful in expectation unless otherwise noted.

MIDR Allocation Rules MIDR mechanisms are variants of *VCG mechanisms*, mechanisms that maximize social welfare and charge “VCG payments”. Formally, a VCG mechanism’s social choice rule satisfies $A(v) \in \operatorname{argmax}_{o \in \mathcal{O}} \sum_j v_j(o)$, and its payments are $p_i(v) = h_i(v_{-i}) - \sum_{j \neq i} v_j(A(v))$ for some function $h_i : V_{-i} \rightarrow \mathbb{R}$. VCG payments are the only universal technique known to induce truthful bidding. The most common implementation of VCG payments uses the Clarke-Pivot payment rule: set $h_i(v_{-i}) = \max_{o \in \mathcal{O}} (\sum_{j \neq i} v_j(o))$, which gives the only payments that simultaneously satisfy truthfulness, IR, and NPT.

More generally, any allocation rule that maximizes an affine function of agents’ valuations can be truthfully implemented with VCG payments. Moreover, Roberts’ theorem [82] implies that in a general setting (when $V_i = \mathbf{R}^{\mathcal{O}}$), if A is onto (every outcome can be realized), then A has truthful payments if and only if it is an affine maximizer. If the “onto” restriction is relaxed, a social choice function is truthfully implementable with VCG payments if and only if it is (weighted) maximal-in-range (MIR) [74] or, for randomized mechanisms, maximal-in-distributional-range (MIDR) [26]:

³“Direct revelation” means that an agent’s bid b_i is an element of V_i . In general this need not be the case; however, by the revelation principle, any social choice rule that may be truthfully implemented may be implemented as a direct revelation mechanism that charges the same payments in equilibrium.

Definition 14 An allocation rule A is MIDR if there is a set \mathcal{D} of probability distributions over outcomes such that A outputs a random sample from the distribution $D \in \mathcal{D}$ that maximizes expected welfare. Formally, for each $v \in V$, $A(v) = o \sim D^*$ where $D^* \in \operatorname{argmax}_{D \in \mathcal{D}} \mathbf{E}_{o \sim D}[\sum_i v_i(o)]$.

A weighted MIDR allocation rule maximizes the weighted social welfare $\sum_i w_i v_i(o)$ for $w_i \geq 0$.

Single-Parameter Domains A larger class of social choice rules can be implemented when V_i is single dimensional. We say that a social choice rule has a single-parameter domain if $v_i(o) = t_i f_i(o)$ for some publicly known function $f_i : \mathcal{O} \rightarrow \mathbb{R}_+$. The value $t_i \in T_i$ is an agent's type (T_i is her type-space, and $T = T_1 \times \dots \times T_n$), and submitting i 's bid precisely requires stating $b_i = t_i$. When $T = \mathbb{R}_+^n$, we say that bidders have *positive types*. We also use $A_i(b) = f_i(A(b))$ as shorthand, and we say A is bounded if the functions A_i are bounded functions.

A single-parameter social choice rule may be implemented if and only if it is *monotone*, where $A : T \rightarrow \mathcal{O}$ is said to be monotone if for each agent i , for all $b_{-i} \in T_{-i}$ and for every two bids $b_i \geq b'_i$, we have $A_i(b_i, b_{-i}) \geq A_i(b'_i, b_{-i})$. This was first shown for a single item auction by Myerson [71]; Archer and Tardos [2] gave the current generalization:

Theorem 20 [Myerson + Archer-Tardos] For a single parameter domain, an allocation rule A has truthful payments (p_1, \dots, p_n) if and only if A is monotone. These payments take the form

$$p_i(b) = h_i(b_{-i}) + b_i A_i(b_i, b_{-i}) - \int_0^{b_i} A_i(u, b_{-i}) du,$$

where $h_i(b_{-i})$ is independent of b_i .

These payments simultaneously satisfy IR and NPT if and only if $p_i^0(b_{-i}) = 0$. Such a mechanism is said to be normalized.

4.3 Single-Call Mechanisms and Single-Call Reductions

Single-call mechanisms capture the idea that the outcome of an auction is generated by a process that is not fully known and cannot be repeated. In this section, we will see that single-call mechanisms are naturally formalized as the the output of a black-box reduction that we call a single-call reduction. We also formalize some of the trade-offs inherent in single-call mechanisms.

In each of our motivating examples, the mechanism designer has access to a procedure $A(\cdot)$ that he wishes to implement truthfully. For example, A might be a procedure that uses

links' reported costs to route a packet along the shortest path in a network. Alternatively, A could be a procedure that ranks ads then displays them to the user.

Moreover, A can only be called once by the mechanism because it produces an irreversible outcome. For example, the purpose of an ad auction is to decide which ads to display to the user. Thus, if A displays ads, it can only be called once per ad auction. In effect, the outcome generated by calling A is necessarily the outcome of the mechanism.

This intuition gives us the following formal definition:

Definition 15 *A single-call mechanism \mathcal{M} for an allocation rule A is a truthful mechanism that computes both the allocation and payments with a single oracle call to A . The output of \mathcal{M} is the outcome of the single call to A .*

The definition of a single-call mechanism captures some important ideas, but it is still missing a key ingredient of the settings we presented earlier: *limited information*. The critical obstacle faced by the mechanism designer in our examples is that he does not know exactly what A will do. For example, when A routed packets, the designer could not know the precise transit time of a packet without actually calling A and measuring the result.

In a worst-case setting, the mechanism designer knows whether he *might* be given a particular A , but otherwise does not know anything about the A he is actually given until he calls it. Formally, we let Ω be a set of allocation rules. The mechanism designer knows that A was drawn from Ω , but otherwise does not know anything about which $A \in \Omega$ besides the result of his one call to A .

Our goal now is to define a *family* of mechanisms $\mathcal{F}(\cdot)$ such that $\mathcal{F}(A)$ is a single-call mechanism for every $A \in \Omega$. In effect, \mathcal{F} will be a procedure that reduces the problem of implementing a truthful mechanism to the problem of calling A once, hence it is a black box reduction:

Definition 16 *A single-call reduction is a procedure that takes any allocation function A from a known set Ω (as a black box) and returns a single-call mechanism.*

For example, the procedure of [9] is a single-call reduction whose input allocation function A is drawn from the set of all monotone, bounded, single-parameter allocation rules and whose output is a single-call mechanism. Similarly, our construction for VCG prices is a single-call reduction that takes any A that is MIDR and returns a single-call mechanism.

The following example illustrates single-call mechanisms and some of their subtleties:

Example 5 *An auctioneer has a single item to sell as well as a monotone mystery procedure $A(\cdot)$ that selects a winner. The auctioneer takes bids $b_i \in \mathbb{R}_+$ and calls $A(b)$ to give the item to the winner.*

Having used A to run the mystery auction, the auctioneer must now compute payments. Since this is a single-parameter setting, we know [71, 2] that payments are given by $p_i(b) = b_i A_i(b_i, b_{-i}) - \int_0^{b_i} A_i(u, b_{-i}) du$. Here are some strategies the auctioneer could consider:

1. Numerically compute $p_i(b)$ by calling A whenever the computation requires $A(x)$ for $x \neq b$.

Unfortunately, this requires evaluating A not just at b , but also at many points of the form (b_{-i}, u) . Thus, it would not be a single-call mechanism.

2. Numerically compute $p_i(b)$ by simulating $A(x)$ for all $x \neq b$.

This would be a single-call mechanism, but the auctioneer does not have enough information to simulate $A(x)$. (On the other hand, if he could simulate $A(x)$, then this strategy would work and the problem of computing payments is easy.)

3. Guess a function \tilde{A} and numerically compute $p_i(b)$ by computing $\tilde{A}(x)$ whenever $A(x)$ is required.

This solves the informational problem that we cannot simulate A and it only requires calling A once; however, it will only be truthful if the auctioneer guessed \tilde{A} correctly, so it does not necessarily produce a single-call mechanism.

4. Use the single-call reduction of [9]. (As will be clear later, the auctioneer must choose this option before calling A .)

Since A is a monotone, single-parameter allocation procedure, these payments give a single-call mechanism.

Remark 1 The “single-call” property is fundamentally computational. As noted earlier, the definition of a single-call mechanism does not limit the information available to the mechanism designer — he might have no knowledge of A or he might have some information about what A does. In the extreme, the designer might even know precisely what A does and already know how to compute truthful payments $p_i(b)$. However, according to our definition, such mechanism is a single-call mechanism as long as it only calls A once. Thus, calling something a “single-call mechanism” is not a statement about the allocation and payments but rather a statement about how they are implemented. Informational limitations are formalized through the idea of a single-call reduction.

To parameterize single-call reductions, we first note the following requirements:

- A reduction must take a bid vector b and a black-box allocation function A as input.
- A reduction must evaluate A on at most one bid vector \hat{b} , causing the outcome $A(\hat{b})$ to be realized.⁴
- A reduction must charge payments λ_i that are a function of b , \hat{b} , and $A(\hat{b})$ (and possibly its own randomness).

These requirements suggest the following generic definition of a single-call reduction to turn an allocation function A into a truthful-in-expectation single-call mechanism $\mathcal{M} = (\mathcal{A}, \{\mathcal{P}_i\})$:

1. Solicit the bid vector b from agents.
2. Use b to compute the modified bid vector \hat{b} . This implicitly defines a probability measure $\mu_b(B)$ denoting the probability of choosing $\hat{b} \in B \subseteq V_1 \times \cdots \times V_n$ as the modified (resampled) bid vector when b is the actual bid vector. When $\hat{b}_i \neq b_i$, we say that i 's bid was resampled.
3. Declare the outcome to be $A(\hat{b})$, i.e. evaluate A at the modified bid vector \hat{b} . This implicitly defines the allocation function $\mathcal{A}(b)$ which samples $\hat{b} \sim \mu_b$ and chooses the outcome $A(\hat{b})$. The resampling procedure must ensure that truthful payments $\mathcal{P}(b)$ exist for $\mathcal{A}(b)$; Note that $\mathcal{A}(b)$ and $\mathcal{P}(b)$ are random variables that depend on the randomly resampled bid vector \hat{b} . Also, $\mathcal{A}(b)$ and $\mathcal{P}(b)$ are randomized even if $A(b)$ and $p(b)$ are deterministic;
4. Use b , \hat{b} , and $A(\hat{b})$ to compute payments $\lambda_i(A(\hat{b}), \hat{b}, b)$ that satisfy truthfulness in expectation, that is, charge player i a payment $\lambda_i(A(\hat{b}), \hat{b}, b)$ such that $\mathbf{E}_{\hat{b}}[\lambda_i(A(\hat{b}), \hat{b}, b)] = \mathbf{E}_{\hat{b}}[\mathcal{P}_i(b)]$.

This general procedure is illustrated in Algorithm 5.

We describe a single-call reduction in the above framework by the tuple $(\mu, \{\lambda_i\})$, where μ implies specifying the resampling measure μ_b for all $b \in V_1 \times \cdots \times V_n$. Since payments should be finite, we require that λ_i be finite everywhere, and we also require that it be integrable. For the rest of this paper, we assume that λ_i 's are deterministic. For randomized λ_i 's, the characterization theorems still hold with λ_i 's replaced by their expectations over the randomness used.

We say that a reduction is *normalized* if $b_i(A(b)) = 0$ for all i implies $\lambda_i(A(\hat{b}), \hat{b}, b) = 0$, i.e. when every agent receives zero value, all payments are zero.

⁴Strictly speaking, there may be settings where a single-call reduction could realize an outcome other than $A(\hat{b})$. However, our restriction follows naturally in scenarios where “computing $A(b)$ ” means realizing $A(b)$ and making measurements. It is also required for complete generality because there is no reason to believe that the designer knows how to realize any outcome other than $A(\hat{b})$.

ALGORITHM 5: Generic Single-Call Reduction $(\mu, \{\lambda_i\})$

input : Black box access to an allocation function A , which is drawn from a known set.

output: Truthful-in-expectation mechanism $\mathcal{M} = (\mathcal{A}, \{\mathcal{P}_i\})$.

- 1 Solicit bid vector b from agents;
- 2 Sample $\hat{b} \sim \mu_b$;
- 3 Realize the outcome $A(\hat{b})$; // $\mathcal{A}(b)$ is the random function $A(\hat{b})$ where $\hat{b} \sim \mu_b$
- 4 Charge payments $\lambda(A(\hat{b}), \hat{b}, b)$; // $\mathcal{P}_i(b)$ is the random function $\lambda_i(A(\hat{b}), \hat{b}, b)$ where $\hat{b} \sim \mu_b$

Optimal Reductions — Expectation vs. Risk

There are two downsides to the mechanisms produced by single-call reductions. First, there is a penalty in *expectation*, i.e., the expected outcome $\mathbf{E}_{\hat{b}}[A(\hat{b})]$ produced by the reduction is not identical to the desired outcome, $A(b)$. This modified outcome may reduce the expected welfare or revenue of the mechanism, or it may simply cause it to do the “wrong” thing.

Second, there is a penalty in *risk* because the payments λ may vary significantly, i.e. for a fixed b the payments at different resampled bids \hat{b} could be very different. In particular, the magnitude of the payment charged by the single-call mechanism may be much larger than the payments in the original mechanism, i.e. it may be that $|\lambda_i| \gg |p_i|$ for certain outcomes.

Our characterization theorems reveal that there is a fundamental trade-off between expectation and risk. Thus, we call a reduction optimal if it minimizes risk with respect to a lower bound on the expectation.

Expectation

We study three criteria for measuring the expectation of a reduction: $\Pr(\hat{b} = b|b)$, social welfare, and revenue.

The first criterion, $\Pr(\hat{b} = b|b)$ (the precision), measures the likelihood that the reduction modifies players’ bids. This criterion is natural when modifying bids is inherently undesirable:

Definition 17 *The precision of a reduction α_P is the probability that the reduction does not alter any player’s bid:*

$$\alpha_P \equiv \min_b \Pr(\hat{b} = b|b) .$$

The other criteria measure standard quantities in mechanism design:

Definition 18 *The welfare approximation α_W of a single-call reduction is given by the worst-case ratio between the welfare of the single-call mechanism and the welfare of the*

original allocation function:

$$\alpha_W = \min_{A,b} \frac{\mathbf{E}_{\hat{b}} [\sum_i b_i(\mathcal{A}_i(b))]}{\sum_i b_i(A_i(b))} .$$

When the welfare of A is zero, $\alpha_W = 1$ if the welfare of \mathcal{A} is also zero and unbounded otherwise.

Definition 19 The revenue approximation α_R of a single-call reduction is given by the worst-case ratio between the revenue of the single-call mechanism and the revenue of the original allocation function:

$$\alpha_R = \min_{A,b} \frac{\mathbf{E}_{\hat{b}} [\sum_i \mathcal{P}_i(b)]}{\sum_i p_i(b)} .$$

When the revenue of A is zero, then $\alpha_R = 1$ when the revenue of \mathcal{A} is also zero and unbounded otherwise.

In the case of continuous spaces we replace min/max with inf/sup as appropriate for infinite domains.

Risk

We measure risk through both the variance of payments and their worst-case magnitude.⁵ In order to make a meaningful comparison across different allocation functions and bids, we normalize by players' bids:⁶

Definition 20 Decompose λ_i into terms which depend only on the payoff to a single bidder j (i.e. on $b_j(A(\hat{b}))$ instead of $A(\hat{b})$):

$$\lambda_i(A(\hat{b}), \hat{b}, b) = \sum_j \lambda_{ij}(b_j(A(\hat{b})), \hat{b}, b)$$

(our characterizations in Sections 4.4 and 4.6 show that this is possible for our settings). Then the bid-normalized payments of the reduction are given by

$$\sum_j \frac{\lambda_{ij}(b_j(A(\hat{b})), \hat{b}, b)}{b_j(A(\hat{b}))} .$$

⁵Intuition suggests optimizing with respect to a high-probability bound. Unfortunately, this is problematic because ignoring low-probability events can dramatically change the expected payment. Thus, in general it is not reasonable to conclude a priori that low-probability events can be ignored.

⁶Intuition also suggests normalizing by the truthful prices for A (i.e. by p_i), but constant allocation functions such as $A_i(b) = 1$ have $p_i = 0$, making this impossible. Bid-normalized payments are a next logical choice.

We can thus write the variance of bid-normalized payments as

$$\max_{A,i} \text{Var}_{\hat{b} \sim \mu_b} \left(\sum_j \frac{\lambda_{ij}(b_j(A(\hat{b})), \hat{b}, b)}{b_j(A(\hat{b}))} \right)$$

and the worst-case magnitude as

$$\max_{A,i,\hat{b}} \left| \sum_j \frac{\lambda_{ij}(b_j(A(\hat{b})), \hat{b}, b)}{b_j(A(\hat{b}))} \right|$$

where we replace min/max with inf/sup as appropriate for infinite domains.

Optimality

We define an optimal reduction as one that *simultaneously* optimizes the six-way trade-off between expectation and risk:

Definition 21 *A single-call reduction optimizes the variance of/worst-case payments with respect to precision/welfare/revenue for a set of allocation functions if for every bid b , it minimizes the variance of/worst-case normalized payments over all possible reductions that achieve a precision of α_P / welfare approximation of α_W / revenue approximation of α_R .*

4.4 Maximal-in-Distributional-Range Reductions

In this section, we show how to construct a single-call reduction for MIDR allocation rules, i.e. we show how to construct a randomized, truthful mechanism from an arbitrary MIDR allocation rule A using only a single black-box call to A . The main results are Theorem 21, a characterization of all reductions that use VCG payments for an arbitrary MIDR allocation rule, and an explicit construction that optimizes the expectation-risk tradeoff.

Truthful payments for MIDR allocation rules are given by VCG payments with the Clarke-Pivot rule:⁷

$$\mathbf{E}[p_i] = \mathbf{E}[\text{total welfare of bidders without } i] - \mathbf{E}[\text{total welfare of bidders } j \neq i \text{ with } (4.1)]$$

(where the expectation is over the randomization in the given MIDR allocation rule). The reduction comes from this formula for $\mathbf{E}[p_i]$: we need to measure the welfare without agent i (the first term in the RHS), so, with some probability, we ignore agent i and maximize the welfare of the remaining agents. Intuitively, this is equivalent to evaluating the allocation function where i 's bid is changed to a “zero” bid while other bids remain the same.

⁷If we relax the no positive transfers requirement, a trivial way to construct a single-call mechanism is to ignore the first term in (4.1). However, the resulting mechanism would make a huge loss because no agent would ever pay the mechanism.

Unfortunately, having removed agent i , even with a small probability, means that computing truthful payments for agent $j \neq i$ requires knowing the allocation where both i and j are ignored. By induction, a single-call mechanism must generate all sets of agents $M \subseteq [n]$ with some probability. Thus, we get an intuitive picture of the reduction's behavior: it will randomly pick a set of bidders $M \subseteq [n]$ and zero the bids of agents not in M .

Characterizing Truthfulness

We consider reductions in which i 's resampled bid \hat{b}_i is always b_i or zero,⁸ where “zero” means that the agent has a valuation of zero for all outcomes. That is, the resampling measure $\mu_b(B)$ represents a discrete distribution over the bids $\{\hat{b}^M\}$ where $M \subseteq [n]$ is a set of agents and

$$\hat{b}_i^M = \begin{cases} b_i & i \in M \\ 0 & i \notin M \end{cases}$$

Resampling to \hat{b}^M is equivalent to ignoring the welfare of agents outside M and evaluating A at b .

In the most general setting, our restriction to zeroing reductions is without loss of generality because b and zero are the only bids that are guaranteed to be valid inputs to A for all MIDR allocation functions A . That said, even if a multi-parameter bid structure were known, VCG payments do not depend on the outcome at any other bid. Thus, intuition suggests that resampling to other bids will not be helpful even if it is possible. This intuition can be formalized, but we do not do it here.

Let $\pi(M)$ be a distribution over sets $M \subseteq [n]$. We define the *associated coefficients* $c_i^\pi(M)$ as:

$$c_i^\pi(M) = \begin{cases} -1, & i \in M \\ \frac{\pi(M \cup \{i\})}{\pi(M)}, & i \notin M \end{cases}$$

Intuitively, c_i^π is the weighting that ensures $-\pi(M \cup \{i\})c_i^\pi(M \cup \{i\}) = \pi(M)c_i^\pi(M)$ (where $i \notin M$) to match the terms in (4.1).

We prove the following characterization of all truthful MIDR reductions $(\pi, \{\lambda_i\})$ that work for all MIDR A :

Theorem 21 *A normalized single-call reduction, with VCG payments, for the set of all MIDR allocation rules satisfies truthfulness, individual rationality, and no positive transfers in an ex-post sense if and only if it takes the form $(\pi, \{\lambda_i\})$ where $\pi(M)$ is a distribution over sets $M \subseteq [n]$, the coefficients $c_i^{\pi(M)}$ are finite, and payments take the form*

$$\lambda_i(A(\hat{b}^M), \hat{b}^M, b) = c_i^\pi(M) \sum_{j \neq i} b_j(A(\hat{b}^M)) .$$

⁸Even if explicit “zero” bids are not known to the reduction, we assume that the reduction can induce A to optimize the utility of an arbitrary subset of agents. Note that a black-box allocation function can only be turned into a truthful mechanism (even if multiple calls to A are allowed) if it can ignore at least one bidder at a time, so our assumption is not unreasonable.

Proof: Recall that in general, a multi-parameter allocation function that can be rendered truthful by VCG payments must be MIDR. Thus, our reduction must ensure that \mathcal{A} is MIDR, and we first derive the implications of this requirement on the single-call reduction. We have already assumed that $\mu_b(B)$ is a distribution over bids $\{\hat{b}^M\}$. Let $\pi_b(M)$ be the probability of selecting \hat{b}^M given b .

First, we show that \mathcal{A} is always MIDR if and only if $\pi_b(M)$ does not depend on b . For the *if* direction, if $\pi_b(M)$ is independent of b then \mathcal{A} is a distribution over MIDR allocation rules, and by [27], such an allocation rule is MIDR.

For the *only if* direction, we use contradiction. Assume that there are some bids x and y such that $\pi_x(M) \neq \pi_y(M)$ for some M . Then there exists a set $S \subseteq [n]$ such that $\Pr_\pi(M \subseteq S|x) \neq \Pr_\pi(M \subseteq S|y)$ (by contradiction and induction, start with $S = \emptyset$). Consider an allocation function that has welfare $\sum_i b_i(A(\hat{b}^M)) = 0$ for $M \subseteq S$ and $\sum_i b_i(A(\hat{b}^M)) = 1$ otherwise. The welfare of \mathcal{A} will be precisely $1 - \Pr_\pi(M \subseteq S)$, implying that for either x or y , \mathcal{A} did not chose the distribution that maximized social welfare and is therefore not MIDR. Thus, the allocation rule \mathcal{A} is MIDR for all MIDR A if and only if $\mu_b(B)$ is a discrete distribution $\pi(M)$ independent of b .

Next, we write VCG payments for \mathcal{A} that satisfy individual rationality and no positive transfers using the Clarke-Pivot payment rule:

$$\begin{aligned} \mathbf{E}[\mathcal{P}_i] &= \sum_{j \neq i} \sum_{M \subseteq [n]} \pi(M) b_j(A(\hat{b}^{M \setminus \{i\}})) - \sum_{j \neq i} \sum_{M \subseteq [n]} \pi(M) b_j(A(\hat{b}^M)) \\ &= \sum_{M|i \notin M} \pi(M \cup \{i\}) \sum_{j \neq i} b_j(A(\hat{b}^M)) - \sum_{M|i \in M} \pi(M) \sum_{j \neq i} b_j(A(\hat{b}^M)) . \end{aligned} \quad (4.2)$$

By definition of $\lambda_i(A(\hat{b}^M), \hat{b}^M, b)$, we know that the expected payment made by i will be

$$\mathbf{E}[\mathcal{P}_i] = \sum_{M \subseteq [n]} \pi(M) \lambda_i(A(\hat{b}^M), \hat{b}^M, b) . \quad (4.3)$$

The two formulas for payments in (4.2) and (4.3) must be equal:

$$\sum_{M \subseteq [n]} \pi(M) \lambda_i(A(\hat{b}^M), \hat{b}^M, b) = \sum_{M|i \notin M} \pi(M \cup \{i\}) \sum_{j \neq i} b_j(A(\hat{b}^M)) - \sum_{M|i \in M} \pi(M) \sum_{j \neq i} b_j(A(\hat{b}^M)) .$$

Since A may be any MIDR allocation function, the only way this can hold is when terms corresponding to each M are equal, i.e., for all i, M

$$\pi(M) \lambda_i(A(\hat{b}^M), \hat{b}^M, b) = \begin{cases} \pi(M \cup \{i\}) \sum_{j \neq i} b_j(A(\hat{b}^M)), & i \notin M \\ -\pi(M) \sum_{j \neq i} b_j(A(\hat{b}^M)) & i \in M . \end{cases} \quad (4.4)$$

To see that this is necessary, construct two allocation functions A and A' such that $b_j(A(\hat{b}^M)) = b_j(A'(\hat{b}^M))$ for all $M \neq \bar{M}$ and $b_j(A(\hat{b}^{\bar{M}})) = 0$. It immediately follows that if the reduction

works for both A and A' , then (4.4) must hold for \bar{M} under A . Since \bar{M} is arbitrary, it follows that (4.4) must hold for all M .

The theorem immediately follows from the above equality. \blacksquare

Remark 2 *Note that this theorem forbids some distributions $\pi(M)$ from being used to construct a single-call reduction — in particular, it requires that $\pi(M) > 0$ for all $M \subseteq [n]$, otherwise some payment $\lambda_i(\cdot)$ will be infinite for nontrivial allocation rules. For example, an obviously forbidden distribution is the one that never changes bids, i.e. the one with $\pi([n]) = 1$. This matches the intuition that a single-call mechanism must occasionally modify bids.*

A Single-Call MIDR Reduction

We now give an explicit single-call reduction for MIDR allocation functions. Our reduction $\text{MIDRTOMECH}(A, \gamma)$ (illustrated in Algorithm 6) is defined by the following resampling distribution $\bar{\pi}$ parameterized by a constant $\gamma \in (0, 1)$:

$$\bar{\pi}(M) = \gamma^{n-|M|}(1-\gamma)^{|M|} \quad (4.5)$$

That is, each agent i is independently dropped from M with probability γ . Thus sampling from the distribution $\bar{\pi}$ is computationally easy. Following Theorem 21, we charge payments $\lambda_i(A(\hat{b}^M), \hat{b}^M, b) = c_i^{\bar{\pi}}(M) \sum_{j \neq i} b_j(A(\hat{b}^M))$ where

$$c_i^{\bar{\pi}}(M) = \begin{cases} -1, & i \in M \\ \frac{1-\gamma}{\gamma}, & i \notin M \end{cases}$$

Corollary 22 (of Theorem 21) *The mechanism*

$$\mathcal{M} = (\mathcal{A}, \{\mathcal{P}_i\}) = \text{MIDRTOMECH}(A, \gamma)$$

calls A once and it satisfies truthfulness, individual rationality, and no positive transfers in an ex-post sense for all MIDR A .

Optimal Single-Call MIDR Reductions

We now prove that the construction $\text{MIDRTOMECH}(A, \gamma)$ is optimal for the definitions of optimality given in Section 4.3. Theorem 21 implies that the bid-normalized payments will be

$$\sum_j \frac{\lambda_{ij}(b_j(A(\hat{b})), \hat{b}, b)}{b_j(A(\hat{b}))} = (n-1)c_i^{\bar{\pi}}(M)$$

Thus, it is sufficient to optimize the variance as $\max_i \text{Var}_{M \sim \bar{\pi}} c_i^{\bar{\pi}}(M)$ and the worst-case as $\max_{i, M} |c_i^{\bar{\pi}}(M)|$.

ALGORITHM 6: MIDRTOMECH(A, γ) — A single-call reduction for MIDR allocation functions

input : MIDR allocation function A .

output: Truthful-in-expectation mechanism $\mathcal{M} = (\mathcal{A}, \{\mathcal{P}_i\})$.

1 Solicit bids b from agents;

2 **for** $i \in [n]$ **do**

with probability $1 - \gamma$
 | Add agent i to set M ;
otherwise
 | Drop agent i from M ;
 └──────────────────────────────────┘

3 Realize the outcome $A(\hat{b}^M)$;

4 Charge payments

$$\lambda_i(A(\hat{b}^M), \hat{b}^M, b) = \left(\sum_{j \neq i} b_j(A(\hat{b}^M)) \right) \times \begin{cases} -1, & i \in M \\ \frac{1-\gamma}{\gamma}, & i \notin M \end{cases}$$

Optimizing Risk vs. Precision

Theorem 23 *The reduction MIDRTOMECH(A, γ) uniquely minimizes both the payment variance and the worst-case payment among all reductions that achieve a precision of at least $\alpha_P = (1 - \gamma)^n$.*

That is, for any other distribution π with precision $\pi([n]) \geq (1 - \gamma)^n$, the payment variance is larger, i.e.

$$\max_i \text{Var}_{M \sim \pi} c_i^\pi(M) > \max_i \text{Var}_{M \sim \bar{\pi}} c_i^{\bar{\pi}}(M) ,$$

and the worst-case payment is larger, i.e.

$$\max_{i, M} |c_i^\pi(M)| > \max_{i, M} |c_i^{\bar{\pi}}(M)| .$$

Proof: First we prove optimality for the worst-case payment $\max_{i, M} |c_i^\pi(M)|$ by contradiction. Assume that some distribution $\pi(M)$ does as well as $\bar{\pi}(M)$. Then it must be that $\max_{i, M} c_i^\pi(M) \leq \max_{i, M} c_i^{\bar{\pi}}(M)$ (the largest coefficient is not bigger), and $\pi([n]) \geq \bar{\pi}([n]) = \alpha_P$ (it respects the lower bound on precision). Since $\max c_i^{\bar{\pi}}(M) = \frac{1-\gamma}{\gamma}$, it must be that for all M and $i \notin M$,

$$\frac{\pi(M \cup \{i\})}{\pi(M)} \leq \max_{i, M} c_i^{\bar{\pi}}(M) = \frac{1 - \gamma}{\gamma} = \frac{\bar{\pi}(M \cup \{i\})}{\bar{\pi}(M)} .$$

Therefore, for any bidder i , it must be that

$$\frac{\pi([n])}{\pi([n] \setminus \{i\})} \leq \frac{\bar{\pi}([n])}{\bar{\pi}([n] \setminus \{i\})} .$$

Since $\pi([n]) \geq \bar{\pi}([n])$, it follows that $\pi([n] \setminus \{i\}) \geq \bar{\pi}([n] \setminus \{i\})$. Repeating this argument, it follows by induction that $\pi(M) \geq \bar{\pi}(M)$ for any set M .

However, we also know that both $\pi(M)$ and $\bar{\pi}(M)$ are distributions so both have to sum to one over all M . Given that $\pi(M) \geq \bar{\pi}(M)$ for all M , this implies $\pi(M) = \bar{\pi}(M)$. Thus, $\bar{\pi}(M)$ is uniquely optimal.

Second, we argue that $\bar{\pi}$ optimizes the payment variance. The variance of bidder i 's payments is

$$\begin{aligned} \text{Var}_{M \sim \pi} c_i^\pi(M) &= \sum_{M \subseteq [n]} \pi(M) (c_i^\pi(M))^2 - \left(\sum_{M \subseteq [n]} \pi(M) c_i^\pi(M) \right)^2 \\ &= \sum_{M \subseteq [n]} \pi(M) (c_i^\pi(M))^2 - 0 \\ &= \sum_{M \subseteq [n] \setminus \{i\}} (\pi(M) + \pi(M \cup \{i\})) \frac{\pi(M \cup \{i\})}{\pi(M)} \end{aligned}$$

This is minimized when $\Pr(i \in M)$ is independent of other bidders (Lemma 66), i.e. $\frac{\pi(M \cup \{i\})}{\pi(M)} = \frac{1 - \gamma_i}{\gamma_i}$ for some constant γ_i . For such a distribution, the precision will be

$$\pi([n]) = \prod_i (1 - \gamma_i) .$$

It follows that the maximum variance is $\max_i \frac{1 - \gamma_i}{\gamma_i}$, and it will only be minimized when $\gamma_i = \gamma_j$ for all $i \neq j$, which corresponds precisely to the distribution $\bar{\pi}$. ■

Optimizing Risk vs. Welfare

A natural optimization metric is the social welfare of \mathcal{A} (indeed, this was an open question from [9] in the single-parameter setting).

Unfortunately, since MIDR allocation rules may generate negative utilities and remain MIDR under additive shifts of the valuation function, one can make the welfare approximation arbitrarily bad (indeed, even undefined) by subtracting a constant from each player's valuation. Thus, if valuation functions may be negative, we cannot meaningfully optimize the loss in social welfare.

However, when valuation functions are known to be nonnegative, the following lemma shows that the worst-case welfare approximation is bounded:

Lemma 24 *The reduction $\text{MIDR}_{\text{TO MECH}}(A, \gamma)$ obtains an $\alpha_W = \min_i \Pr_\pi(i \in M) = 1 - \gamma$ approximation to the social welfare, and there is an allocation function A and bid b such that this bound is tight.*

The idea for the lower bound is that the sum of welfare of bidders in M cannot be lower at $A(\hat{b}^M)$ than at $A(\hat{b}^{[n]})$ because that would imply A did not maximize the social welfare of bidders in M at \hat{b}^M . The worst case scenario occurs when one player receives all the welfare. The proof is below.

Using this lemma, we can show that $\text{MIDRTOMECH}(A, \gamma)$ is optimal:

Theorem 25 *The reduction $\text{MIDRTOMECH}(A, \gamma)$ minimizes payment variance and worst-case payments among all reductions that achieve a welfare approximation of at least $\alpha_W = 1 - \gamma$.*

Proof of Lemma 24. The expected social welfare of the single-call mechanism, where the expectation is over the randomness in the resampling function is given by $\mathbf{E} \left[\sum_{j \in [n]} b_j(\mathcal{A}(b)) \right]$. We now prove the required lower bound on this quantity.

$$\begin{aligned}
\mathbf{E} \left[\sum_{j \in [n]} b_j(\mathcal{A}(b)) \right] &= \sum_{j \in [n]} \sum_{M \subseteq [n]} \pi(M) b_j(A(\hat{b}^M)) \\
&= \sum_{M \subseteq [n]} \pi(M) \sum_{j \in [n]} b_j(A(\hat{b}^M)) \\
&\geq \sum_{M \subseteq [n]} \pi(M) \sum_{j \in M} b_j(A(\hat{b}^M)) \\
&\geq \sum_{M \subseteq [n]} \pi(M) \sum_{j \in M} b_j(A(\hat{b}^{[n]})) \\
&= \sum_{j \in [n]} \Pr_{\pi}(j \in M) b_j(A(\hat{b}^{[n]})) \\
&\geq \left(\min_{j \in [n]} \Pr_{\pi}(j \in M) \right) \sum_{j \in [n]} b_j(A(\hat{b}^{[n]})) \\
&= \left(1 - \max_{j \in [n]} \Pr_{\pi}(j \notin M) \right) \sum_{j \in [n]} b_j(A(\hat{b}^{[n]})) .
\end{aligned}$$

Finally, we observe that this is tight. Consider a valuation and allocation function pair for which, every agent other than some agent j has a zero value for every outcome, and agent j has a non-zero value only for those outcomes that were chosen taking j into consideration, i.e.:

$$b_k(A(\hat{b}^M)) = \begin{cases} 0, & k \neq j \\ 0, & j \notin M \\ 1, & \text{otherwise} \end{cases}$$

When $j = \operatorname{argmax}_{k \in [n]} \Pr_{\pi}(k \notin M)$, the preceding bound is tight. ■

Lemma 26 *Let π be a distribution such that $\max_{i,M} c_i^\pi(M) < \max_{i,M} c_i^{\bar{\pi}}(M)$. Then*

$$\max_i \Pr_\pi(i \notin M) > \max_i \Pr_{\bar{\pi}}(i \notin M) .$$

Proof: Let $\bar{c} = \max_{i,M} c_i^{\bar{\pi}}(M)$. Note that for all $M|i \notin M$, $c_i^{\bar{\pi}}(M) = \frac{\bar{\pi}(M \cup \{i\})}{\bar{\pi}(M)} = \bar{c}$. It follows by algebra that $\frac{\sum_{M|i \notin M} \bar{\pi}(M \cup \{i\})}{\sum_{M|i \notin M} \bar{\pi}(M)} = \bar{c}$ and therefore by the conditions of the lemma

$$\max_{i,M} c_i^\pi(M) < \frac{\sum_{M|i \notin M} \bar{\pi}(M \cup \{i\})}{\sum_{M|i \notin M} \bar{\pi}(M)} . \quad (4.6)$$

Next we have,

$$\max_{i,M} c_i^\pi(M) \geq \max_M c_i^\pi(M) \geq \max_{M|i \notin M} \frac{\pi(M \cup \{i\})}{\pi(M)} \geq \frac{\sum_{M|i \notin M} \pi(M \cup \{i\})}{\sum_{M|i \notin M} \pi(M)} \quad (4.7)$$

Combining (4.6) and (4.7) gives

$$\frac{\sum_{M|i \notin M} \pi(M \cup \{i\})}{\sum_{M|i \notin M} \pi(M)} < \frac{\sum_{M|i \notin M} \bar{\pi}(M \cup \{i\})}{\sum_{M|i \notin M} \bar{\pi}(M)} \quad (4.8)$$

Note that since π and $\bar{\pi}$ are probability distributions, the sum of the numerator and denominator of both the LHS and the RHS of (4.8) equals 1. Thus, it immediately follows that the denominator of the LHS is larger than the denominator of the RHS, i.e.,

$$\sum_{M|i \notin M} \pi(M) > \sum_{M|i \notin M} \bar{\pi}(M) \quad (4.9)$$

Inequality (4.9) when restated, reads as

$$\Pr_\pi(i \notin M) > \Pr_{\bar{\pi}}(i \notin M) .$$

But since the above inequality is true for all i , and the RHS of the above inequality is the same for all i (namely the parameter μ by which the reduction is parametrized), the statement of the lemma follows. \blacksquare

Proof of Theorem 25. By Lemma 24, the worst case loss in social welfare of a distribution π is given by

$$1 - \alpha_\pi = \max_i \Pr_\pi(i \notin M) .$$

For worst-case payments, the contrapositive of Lemma 26 precisely says that if $1 - \alpha_\pi \leq 1 - \alpha_{\bar{\pi}}$, then the largest payment $\max_{M,i} c_i^\pi(M) \geq \max_{M,i} c_i^{\bar{\pi}}(M)$, thus proving that any other reduction will be worse.

For payment variance, arguing along the lines of Theorem 21 again says that variance will be minimized when π is an independent distribution and $\Pr(i \in M)$ is the same for all i . Since $\bar{\pi}$ is precisely the distribution that does this, it follows that it is optimal. \blacksquare

Optimizing Risk vs. Revenue

Our next lemma implies that a lower bound on the factor of approximation to revenue is equivalent to a lower bound on precision.

Lemma 27 *The reduction $\text{MIDRTOMECH}(A, \gamma)$ obtains an $\alpha_R = \alpha_P = \pi([n]) = (1 - \gamma)^n$ approximation to the revenue, and this is tight.*

Lemma 27 shows that the revenue approximation α_R is the same as the precision α_P for the reduction $\text{MIDRTOMECH}(A, \gamma)$. Since Theorem 23 says that $\text{MIDRTOMECH}(A, \gamma)$ optimizes payments with respect to precision, it similarly follows that it optimizes payments with respect to revenue:

Theorem 28 *The reduction $\text{MIDRTOMECH}(A, \gamma)$ minimizes payment variance and the worst-case payment among all reductions that guarantee an $\alpha_R = (1 - \gamma)^n$ approximation to revenue.*

Proof of Lemma 27. For any b with non-negative valuations, the revenue under a single call reduction will be

$$\begin{aligned} \sum_{i \in [n]} \mathbf{E}[\mathcal{P}_i] &= \sum_{i \in [n]} \sum_{M \subseteq [n]} \pi(M) \left(\sum_{k \neq i} b_k(A(\hat{b}^{M \setminus \{i\}})) - \sum_{k \neq i} b_k(A(\hat{b}^M)) \right) \\ &\geq \pi([n]) \sum_{i \in [n]} \left(\sum_{k \neq i} b_k(A(\hat{b}^{[n] \setminus \{i\}})) - \sum_{k \neq i} b_k(A(\hat{b}^{[n]})) \right) \end{aligned}$$

where $\sum_{i \in [n]} \left(\sum_{k \neq i} b_k(A(\hat{b}^{[n] \setminus \{i\}})) - \sum_{k \neq i} b_k(A(\hat{b}^{[n]})) \right)$ is the revenue generated by A under VCG prices. Thus, any distribution $\pi(M)$ gives an $\alpha = \pi([n])$ approximation to the revenue.

To see that this is tight, consider the following allocation function:

$$b_i(A(\hat{b}^M)) = \begin{cases} \frac{1}{n} & M = [n] \\ \frac{1}{n-1} & i \in M \text{ but } M \neq [n] \\ 0, & \text{otherwise.} \end{cases}$$

The revenue under VCG prices is $\sum_{i \in [n]} \left(\sum_{k \neq i} b_k(A(\hat{b}^{[n] \setminus \{i\}})) - \sum_{k \neq i} b_k(A(\hat{b}^{[n]})) \right)$, which is $n \left(\frac{n-1}{n-1} - \frac{n-1}{n} \right) = 1$.

Under any single-call reduction, the revenue will be given by

$$\begin{aligned}
\sum_{i \in [n]} \mathbf{E}[\mathcal{P}_i] &= \sum_{i \in [n]} \sum_{M \subseteq [n]} \pi(M) \left(\sum_{k \neq i} b_k(A(\hat{b}^{M \setminus \{i\}})) - \sum_{k \neq i} b_k(A(\hat{b}^M)) \right) \\
&= \sum_{i \in [n]} \pi([n]) \left(\sum_{k \neq i} b_k(A(\hat{b}^{[n] \setminus \{i\}})) - \sum_{k \neq i} b_k(A(\hat{b}^{[n]})) \right) \\
&= \sum_{i \in [n]} \pi([n]) \left(1 - \frac{n-1}{n} \right) \\
&= \pi([n]) .
\end{aligned}$$

■

4.5 A Single-Call Application — PPC AdAuctions

Pay-per-click (PPC) AdAuctions are a prime example of mechanisms in which uncertainty can destroy truthfulness. There is a deep literature on truthful ad auctions, much of which makes a powerful assumption: the likelihood that a user clicks in any given setting is a commonly-held belief. In reality, this simply is not true. Auctioneers make their best effort to estimate the likelihood of a click; however, anecdotal evidence [50] suggests that advertisers manipulate their bids according to the perceived accuracy of the auctioneer’s estimates. As we will illustrate in this section, even if the auctioneer’s estimates are good enough to (say) maximize welfare given the current bids, they are not sufficient to compute truthful prices. We show that single-call mechanisms can recover truthfulness in PPC ad auctions in spite of these conflicting beliefs.

In a standard PPC ad auction, n advertisers compete for $m \ll n$ slots. The value to an advertiser depends on the likelihood of a click, called the click-through-rate (CTR) c , and the value to the advertiser once the user has clicked, the value-per-click v . The expected value to an advertiser is thus cv . The auctioneer’s job is to assign advertisers to slots and compute per-click payments — bidders are only charged when a click occurs. Both tasks require knowing the CTRs for common objectives like welfare or revenue maximization, so the auctioneer must also maintain estimates of the CTRs, which we denote by c' .

Researchers generally acknowledge that, in reality, both c and v may depend arbitrarily on the outcome — they certainly depend on the quality and relevance of the particular ad being shown, but they also depend on where the ad is shown and on which other ads are shown nearby. However, for analytical tractability, the parameters c and v are often assumed to have a very restricted structure. We discuss two different structures to illustrate the pervasiveness of the problem caused by estimation error and to show how different single-call reductions may be applied.

Outcome-Independent Values and Separable CTRs In the ad auction literature, it is common to assume that a bidder's value-per-click v_i is independent of the assignment and that the CTR is separable, that is, it takes the form $c = \alpha_j \beta_i$, where β_i depends only on the ad and α_j depends only on the slot $j \in [m]$ where the ad is shown. Unfortunately, even in this restricted setting, estimation errors may break the truthfulness of VCG prices. The following example shows that even if the auctioneer's estimates correctly identify the welfare-maximizing allocation, they may not yield truthful prices, even in the special case where $\beta_i = 1$:

Example 6 Consider a 2-slot, 2-advertiser setting with CTRs c_j and bids b_i . Assume that $b_1 > b_2$ and $c_1 > c_2$, so that the welfare-optimizing assignment is to assign ad-1 to slot-1 and ad-2 to slot-2, i.e.,

$$c_1 b_1 + c_2 b_2 \geq c_1 b_2 + c_2 b_1 . \quad (4.10)$$

The auctioneer wishes to optimize welfare, so he uses c'_j to implement the VCG allocation. It is quite plausible that maximizing welfare w.r.t c'_j results in the same welfare maximizing allocation, namely given (4.10), it is not unreasonable to assume that the following is true if the auctioneer's estimates are good enough:

$$c'_1 b_1 + c'_2 b_2 \geq c'_1 b_2 + c'_2 b_1 .$$

However, we will show that this is not enough to guarantee truthfulness.

We show that advertiser-1 may have an incentive to lie. According to the estimates c'_j , The expected VCG payment should be $c'_1 b_2 - c'_2 b_2$. Since advertiser 1 will only be charged when he actually receives a click, the price-per-click charged will be

$$p_1 = \frac{1}{c'_1} [c'_1 b_2 - c'_2 b_2] .$$

and the expected utility to bidder i will be

$$u_1 = c_1 \left(b_1 - \frac{1}{c'_1} [c'_1 b_2 - c'_2 b_2] \right) ,$$

where the extra c_1 gets multiplied because the utility is non-zero only upon a click, which happens with probability c_1 .

Now, for example, let the inaccurate c'_j be as follows: $c'_1 = \alpha c_1$, $c'_2 = c_2$ where $\alpha > 1$. Notice that in this example we always have

$$\alpha c_1 b_1 + c_2 b_2 \geq \alpha c_1 b_2 + c_2 b_1$$

and thus the mechanism will always maximize welfare in spite of the estimation errors.

The utility of advertiser-1 will be

$$u_1 = c_1 \left(b_1 - \frac{1}{\alpha c_1} [\alpha c_1 b_2 - c_2 b_2] \right) .$$

Now, suppose advertiser-1 decides to lie and bid zero, he gets the second slot, pays zero, and gets utility of c_2b_1 . Lying is clearly profitable if

$$c_2b_1 > c_1 \left(b_1 - \frac{1}{\alpha c_1} [\alpha c_1 b_2 - c_2 b_2] \right).$$

Rearranging, lying is profitable if

$$\alpha c_1 b_1 + c_2 b_2 < \alpha c_1 b_2 + \alpha c_2 b_1 \quad (4.11)$$

It is quite possible that lying might be profitable, that is inequality (4.11) holds true. For example, if $c_1 = 0.1$, $c_2 = 0.09$, $b_1 = 1.1$, $b_2 = 1$, and $\alpha = 1.1$, payments computed using c'_j are nontruthful, even though the mechanism always picks the welfare-maximizing assignment for any $\alpha > 1$.

In the language of allocations and payments, truthfulness is broken because the auctioneer only knows an estimate of A and thus does not have enough information to compute true VCG prices. However, once ads are shown, clicks may be measured, giving an unbiased estimate of bidders' values. Unfortunately, this can only be done once — since the auctioneer only has one opportunity to show ads to the user, these unbiased estimates can only be measured under a single advertiser-slot assignment. Fortunately, *these unbiased estimates are exactly the information required to compute truthful payments using a single-call mechanism.*

Since a player's bid b_i is merely its value-per-click v_i , this version of a PPC ad auction is a single-parameter domain and we can apply the result of [9]. Their result says that we can turn any monotone allocation rule into a truthful-in-expectation mechanism — maximizing welfare subject to estimates α'_j and β'_i is a monotone allocation rule as long as the estimates α'_j have the same order as α_j (i.e. $\alpha'_{j_1} \geq \alpha'_{j_2}$ if $\alpha_{j_1} \geq \alpha_{j_2}$). Thus the mechanism $\text{SPTOMECHBKS}(A^{PPC}, \gamma)$ of [9] (we define the mechanism formally in Section 4.6) gives a truthful mechanism for all properly ordered estimates of α 's.

Theorem 29 *Consider a single-parameter PPC auction with separable CTRs and let A^{PPC} be the allocation rule that maximizes welfare using estimated CTR parameters α'_j and β'_i , where the estimates α'_j are properly ordered. Then $\text{SPTOMECHBKS}(A^{PPC}, \gamma)$, the single-call reduction of [9], gives a mechanism that is truthful in expectation and has expected welfare within a factor of $(1 - \gamma)^n$ of A^{PPC} .*

Outcome-Dependent Values and CTRs While most research uses single-parameter models for analytical tractability, an advertiser's value-per-click v really depends on the advertiser-slot assignment chosen by the auctioneer as noted earlier. As in the preceding single-parameter setting, estimated CTRs are insufficient to guarantee truthfulness; however, the reduction of [9] no-longer applies in such a multi-parameter domain — we show how our MIDR single-call reduction can be used to recover truthfulness.

To capture the dependence on the advertiser-slot assignment, we assume that a bidder's CTR $c_{i,j}$ and value-per-click $v_{i,j}$ depend arbitrarily on both the bidder i and the slot j .

Since the only allocation rules that have truthful prices in general multi-parameter domains are MIDR, we assume that the auctioneer can generate a MIDR allocation, specifically we assume the auctioneer can query an oracle to determine the allocation that maximizes the welfare of any set of bidders under the actual bid b (but not necessarily for an arbitrary bid b) and apply our MIDR reduction in Section 4.4 to obtain the following:

Theorem 30 *Consider a multi-parameter PPC auction where a bidder’s value-per-click $v_{i,j}$ depends on the bidder and the slot. Let A^{PPC} be an allocation rule that chooses the advertiser-slot assignment returned by the welfare-maximizing oracle described above. Then the mechanism $\text{MIDRTOMECH}(A^{PPC}, \gamma)$ is truthful in expectation and approximates the welfare of A^{PPC} to within a factor of $(1 - \gamma)$.*

4.6 Single-Parameter Reductions

In this section, we characterize truthful reductions for single-parameter domains and show that the construction of [9] is optimal. Theorem 31 characterizes all reductions that are truthful for an arbitrary monotone, bounded, single-parameter allocation function A . Our characterization is more general than the self resampling procedures described by Babaioff et al. and shows that a wide variety of probability measures may be used to construct a truthful reduction. Theorem 33 shows that the construction given in Babaioff et al. is optimal among such reductions for a fixed bound on the precision, welfare approximation, or revenue approximation of the reduction.

Generalization. Both Theorem 31 and Theorem 33 assume that the resampling measures are “nice” for simplicity in exposition. See Theorem 37 in Section 4.7 for the fully general version of Theorem 31, and Theorem 57 in Appendix B.1 for the fully general version of Theorem 33.

As in the MIDR setting, truthful payments give intuition for the structure of a single-call reduction. As noted in Section 4.2, payments are truthful if and only if they are given by the Archer-Tardos characterization:

$$p_i(b) = b_i A_i(b) - \int_0^{b_i} A_i(u, b_{-i}) du . \quad (4.12)$$

Loosely speaking, this says “charge i the value she receives minus what she would expect if she lowered her bid.” Thus, a single call reduction should, with some probability, lower agents’ bids to compute the value of allocation function at (u, b_{-i}) for $u \leq b_i$.

Characterizing Single-Call Reductions

For the sake of intuition, we start with the special case that the resampling measure μ_b has a nicely behaved density representation $f_b(\hat{b})$ (the resampling density) that is continuous

in \hat{b} and b . The proof for arbitrary measures μ_b requires significant measure theory and is deferred until Theorem 37 of Section 4.7.

Define the coefficients $c_i^f(\hat{b}, b)$ as $c_i^f(\hat{b}, b) = 1 - \frac{1}{b_i} \int_0^{b_i} \frac{f_{u, b-i}(\hat{b})}{f_b(\hat{b})} du$ when $b_i \neq 0$, and to be 0 when $b_i = 0$. We characterize truthful reductions as follows:

Theorem 31 *A normalized single-parameter reduction $(f, \{\lambda_i\})$ for the set of all monotone bounded single-parameter allocation functions satisfies truthfulness, individual rationality and no positive transfers in an ex-post sense if and only if the following conditions are met:*

1. *The resampling density f_b is such that the single-call mechanism's randomized allocation procedure $\mathcal{A}_i(b)$ is monotone in expectation, i.e., for all agents i , for all b , and $b'_i \geq b_i$, $\mathbf{E}_{\hat{b} \sim f_b}[\mathcal{A}_i(b'_i, b_{-i})] \geq \mathbf{E}_{\hat{b} \sim f_b}[\mathcal{A}_i(b)]$. (See below.)*
2. *The resampling density f_b is such that $f_b(\hat{b}) \neq 0$ if $\int_0^{b_i} f_{u, b-i}(\hat{b}) du \neq 0$.*⁹
3. *The payment functions $\lambda_i(A(\hat{b}), \hat{b}, b)$ satisfy: $\lambda_i(A(\hat{b}), \hat{b}, b) = b_i c_i^f(\hat{b}, b) A_i(\hat{b})$ almost surely, i.e. for all \hat{b} except possibly a set with probability zero under f_b .*

Proof: (See Section 4.7 for the proof when μ_b is an arbitrary measure.)

Necessity. The first condition, that \mathcal{A} must be monotone in expectation, follows directly from Archer-Tardos characterization of truthful allocation functions. The second and third conditions, as we prove below, are necessary for the expected payment to take the form required by the Archer-Tardos characterization.

The allocation function \mathcal{A} is a single-parameter allocation function, so the Archer-Tardos characterization gives truthful prices if they exist:

$$\begin{aligned} \mathbf{E}[\mathcal{P}_i] &= b_i \mathbf{E}_{\hat{b} \sim f_b}[\mathcal{A}_i(b)] - \int_0^{b_i} \mathbf{E}_{\hat{b} \sim f_{u, b-i}}[\mathcal{A}_i(u, b_{-i})] du \\ &= b_i \mathbf{E}_{\hat{b} \sim f_b}[A_i(\hat{b})] - \int_0^{b_i} \mathbf{E}_{\hat{b} \sim f_{u, b-i}}[A_i(\hat{b})] du \\ &= b_i \int_{\hat{b} \in \mathbb{R}^n} A_i(\hat{b}) f_b(\hat{b}) d\hat{b} - \int_0^{b_i} \int_{\hat{b} \in \mathbb{R}^n} A_i(\hat{b}) f_{u, b-i}(\hat{b}) d\hat{b} du . \end{aligned}$$

Rearranging, where changing the order of integration may be justified by Tonelli's theorem, gives

$$\mathbf{E}[\mathcal{P}_i] = \int_{\hat{b} \in \mathbb{R}^n} f_b(\hat{b}) b_i A_i(\hat{b}) \left(1 - \frac{1}{b_i} \int_0^{b_i} \frac{f_{u, b-i}(\hat{b})}{f_b(\hat{b})} du \right) d\hat{b} .$$

By construction, we can express the expected price as

$$\mathbf{E}[\mathcal{P}_i] = \int_{\hat{b} \in \mathbb{R}^n} f_b(\hat{b}) \lambda_i(A(\hat{b}), \hat{b}, b) d\hat{b} .$$

⁹This condition effectively requires $c_i^f(\hat{b}, b)$ to be finite.

Thus truthfulness in expectation necessarily implies

$$\int_{\hat{b} \in \mathbb{R}^n} f_b(\hat{b}) \lambda_i(A(\hat{b}), \hat{b}, b) d\hat{b} = \int_{\hat{b} \in \mathbb{R}^n} f_b(\hat{b}) b_i A_i(\hat{b}) \left(1 - \frac{1}{b_i} \int_0^{b_i} \frac{f_{u, b_{-i}}(\hat{b})}{f_b(\hat{b})} du \right) d\hat{b} . \quad (4.13)$$

Note that proving the necessity of condition three in the theorem is equivalent to proving that the integrands in the LHS and the RHS of (4.13) are equal almost everywhere. That is, we have to show that the only way for Equation (4.13) to hold for all monotone bounded A is when the integrands are equal almost everywhere. To show this, it is sufficient to show that Equation (4.13) must still hold if we restrict the range of integration to an arbitrary rectangular parallelepiped (hence forth called as rectangle) $\mathbb{S} \subseteq \mathbb{R}^n$ (see why this is enough in Section 4.7 for a more general setting), that is, it is sufficient to show that for all rectangles $\mathbb{S} \subseteq \mathbb{R}^n$

$$\int_{\hat{b} \in \mathbb{S}} f_b(\hat{b}) \lambda_i(A(\hat{b}), \hat{b}, b) d\hat{b} = \int_{\hat{b} \in \mathbb{S}} f_b(\hat{b}) b_i A_i(\hat{b}) \left(1 - \frac{1}{b_i} \int_0^{b_i} \frac{f_{u, b_{-i}}(\hat{b})}{f_b(\hat{b})} du \right) d\hat{b} . \quad (4.14)$$

Showing (4.14) would be straight-forward if we are given that (4.13) holds for *all* A — we could take any A and make it zero for all points not in \mathbb{S} , and then (4.13) immediately implies (4.14). However (4.13) is guaranteed to be true only for monotone bounded A , since those are the allocation functions that could possibly be input to our reduction. To see that it is still true when (4.13) is only guaranteed for monotone bounded A , define the function $1_{\mathbb{S}}(\hat{b})$ as

$$1_{\mathbb{S}}(\hat{b}) = \begin{cases} 1, & \hat{b} \in \mathbb{S} \\ 0, & \text{otherwise.} \end{cases}$$

Observe that $1_{\mathbb{S}}$ can be written as $1_{\mathbb{S}}(\hat{b}) = 1_{\mathbb{S}}^+(\hat{b}) - 1_{\mathbb{S}}^-(\hat{b})$ where $1_{\mathbb{S}}^+$ and $1_{\mathbb{S}}^-$ are both $\{0, 1\}$, monotone functions. Moreover, the functions $A^+(b) = 1_{\mathbb{S}}^+(b)A(b)$ and $A^-(b) = 1_{\mathbb{S}}^-(b)A(b)$ are also monotone, and they agree with A on \mathbb{S} . If we plug A^+ and A^- into (4.13) and subtract the results, we get precisely (4.14). Thus condition three is necessary.

For the necessity of condition two, note that if it were not to hold, the coefficients c_i^f will become $-\infty$, and hence the payments as defined in condition three will not be finite. Clearly finiteness of payments is a requirement.

This proves that all three conditions in the theorem are necessary for truthfulness.

Sufficiency. We now show that the three stated conditions are sufficient. In a single-parameter setting, for a mechanism to be truthful, we need the allocation function to be monotone in expectation and the payment function to satisfy the Archer-Tardos payment functions. Condition one guarantees that the allocation function output by the single-call reduction is a monotone in expectation allocation function. It remains to show that the second and third conditions result in payments that agree with Archer-Tardos payments. Given condition two, finiteness of payments as defined in condition three is satisfied. All we

need to show is that under the formula of $\lambda_i(A(\hat{b}), \hat{b}, b)$ described in condition three, the single-call payments match in expectation with Archer-Tardos payments, i.e., (4.13) holds.

Since $c_i^f(\hat{b}, b) = 1 - \frac{1}{b_i} \int_0^{b_i} \frac{f_{u, b-i}(\hat{b})}{f_b(\hat{b})} du$, taking

$$\lambda_i(A(\hat{b}), \hat{b}, b) = b_i c_i^f(\hat{b}, b) A_i(\hat{b}) \quad a.s.$$

trivially satisfies (4.13), implying that the reduction is truthful. \blacksquare

Unfortunately, our assumption that μ_b has a density representation is unreasonable. Most significantly, one would expect $\hat{b} = b$ with some nonzero probability, implying that μ_b would have at least one atom for most interesting distributions. In particular, the distribution used in the BKS transformation has such an atom, so it cannot be analyzed in this fashion.

To handle general measures μ_b we apply the same ideas in Section 4.7 using tools from measure theory.

The BKS Reduction for Positive Types

The central construction of Babaioff, Kleinberg, and Slivkins [9] is a reduction for scenarios where bidders have positive types.¹⁰

Their resampling procedure (implicitly defining μ_b) is described Algorithm 7. In the language of our characterization, the coefficients c_i^{BKS} are

$$c_i^{BKS}(\hat{b}, b) = \begin{cases} 1, & \hat{b}_i = b_i \\ 1 - \frac{1}{\gamma} & \text{otherwise.} \end{cases}$$

They proved that $\text{SPTOMECHBKS}(A, \gamma)$ is truthful. This fact can be easily derived from Theorem 31:

Theorem 32 (Babaioff, Kleinberg, and Slivkins 2010.) *For all monotone, bounded, single-parameter allocation rules A , the single-call mechanism given by $\text{SPTOMECHBKS}(A, \gamma)$ satisfies truthfulness and no positive transfers in an ex-post sense and is ex-post universally individually rational.*

Optimal Single-Call Reductions

Analogous to our MIDR construction, we show that, the BKS construction for positive types is optimal with respect to precision, welfare, and revenue as defined in Section 4.3. Using our characterization from Theorem 31, the bid-normalized payments we wish to optimize will be

$$\sum_j \frac{\lambda_{ij}(b_j(A(\hat{b})), \hat{b}, b)}{b_j(A(\hat{b}))} = \frac{c_i^\mu(\hat{b}, b) b_i A_i(\hat{b})}{b_i A_i(\hat{b})} = c_i^\mu(\hat{b}, b) .$$

¹⁰They also give a reduction that applies to more general type spaces, but we do not state it here.

ALGORITHM 7: $\text{SPTOMECHBKS}(A, \gamma)$ — The BKS reduction for single-parameter domains

input : Bounded, monotone allocation function A .

output: Truthful-in-expectation mechanism $\mathcal{M} = (\mathcal{A}, \{\mathcal{P}_i\})$.

1 Solicit bids b from agents;

2 **for** $i \in [n]$ **do**

| **with probability** $1 - \gamma$

| Set $\hat{b}_i = b_i$;

| **otherwise**

| Sample x_i uniformly at random from $[0, \hat{b}_i]$;

| Set $\hat{b}_i = b_i x_i^{\frac{1}{1-\gamma}}$;

3 Realize the outcome $A(\hat{b})$;

4 Charge payments

$$\lambda_i(A(\hat{b}^M), \hat{b}^M, b) = b_i A_i(\hat{b}) \times \begin{cases} 1, & \hat{b}_i = b_i; \\ \frac{1-\gamma}{\gamma}, & \hat{b}_i < b_i; \end{cases}$$

Thus, optimizing variance of normalized payments is equivalent to optimizing $\max_i \text{Var}_{\hat{b}_i \sim \mu_b} c_i^\mu(\hat{b}, b)$, and optimizing the worst-case normalized payment is equivalent to optimizing $\sup_{i, \hat{b}} |c_i^\mu(\hat{b}, b)|$.

For this section, we make a “nice distribution” assumption that for any $u \neq b_i$, $\Pr(\hat{b}_i = u|b) = 0$. That is, if we compute the marginal distribution of \hat{b}_i , the only bid \hat{b}_i that has an atom is b_i (other bids only have positive density). We handle the general case with full proofs in Appendix B.1. The generalized versions of Theorem 33, Lemma 34 and Theorem 35 are Theorem 57, Lemma 58 and Theorem 59 in Section B.1.

Our main result is that the BKS transformation is optimal:

Theorem 33 *The single-call reduction $\text{SPTOMECHBKS}(A, \gamma)$ optimizes the variance of bid-normalized payments and the worst-case bid-normalized payment for every b subject to a lower bound $\alpha = (1 - \gamma)^n \in (\frac{1}{e}, 1)$ on the precision, the welfare approximation, or the revenue approximation.*

We prove Theorem 33 in two steps: Lemma 34 and Theorem 35. Lemma 34 shows that the three metrics we study are equivalent for interesting reductions in the single parameter setting:

Lemma 34 *For $\alpha > \frac{1}{e}$ and $n \geq 2$, a reduction that optimizes the variance of normalized payments or the maximum normalized payment subject to a precision constraint of $\Pr(\hat{b} = b|b) \geq \alpha$ also optimizes the maximum payment subject to a welfare or revenue approximation of α .*

Proof: (Sketch. The full proof is in Appendix B.1.) Consider the following allocation function:

$$A_i(b) = \begin{cases} 1, & b \geq \bar{b} \\ 0, & \text{otherwise.} \end{cases}$$

Intuitively, a reduction should not resample to higher bids because Archer-Tardos payments do not depend on higher bids, and hence no useful information is obtained through raising bids. However, if a reduction never raises bids (i.e. $\Pr(\hat{b} \leq b|b) = 1$), then the welfare and revenue of a single-call reduction will both be precisely $\Pr(\hat{b} = b|b)$ if we consider the above mentioned A at a bid of \bar{b} . ■

Thus, to prove Theorem 33, it is sufficient to prove that the BKS reduction optimizes precision.

Theorem 35 *The single-call reduction $\text{SPTOMECHBKS}(A, \gamma)$ optimizes the variance of normalized payments and the worst-case normalized payment among reductions with a precision of at least $\alpha_P = (1 - \gamma)^n > \frac{1}{e}$.*

Proof: (Sketch. The full proof is in Appendix B.1.) When $\Pr(\hat{b} = b|b)$ is large, the mechanism extracts a modest payment from i when $\hat{b}_i = b_i$ and pays a large rebate otherwise. Thus, we bound $\inf_{\hat{b}_i} c_i^\mu(\hat{b}, b)$. Let $\pi^\mu(M, b)$ be the probability (given b) that $\hat{b}_i = b_i$ for all $i \in M$ and $\hat{b}_i < b_i$ for all $i \notin M$. Then the key step is to prove the following lower bound on $\inf_{\hat{b}_i} c_i^\mu$:

$$\inf_{\hat{b}_i} c_i^\mu(\hat{b}, b) \leq -\frac{\pi^\mu(M \cup \{i\}, b)}{\pi^\mu(M, b)}.$$

Notably, this bound takes the same form as the truthful payment coefficients for MIDR reductions. Applying the same logic as Theorem 23 shows that the BKS transformation is optimal. ■

4.7 Characterizing Reductions for Single-Parameter Domains

In this section, we generalize the single-parameter characterization theorem from Section 4.6 to reductions using arbitrary measures μ_b . We refer the reader to Appendix B.2 for some background and definitions from measure theory.

Before we begin, we must formalize some properties of the functions A and the measures μ_b . The following assumptions would typically be implicit in Algorithmic Mechanism Design; however, it is necessary that they be formalized for some of the tools in our proof. We assume the following:

1. Any allocation function A that the reduction receives as input (as a black box) is a Borel measurable function, i.e., each of the A_i 's as a function from $\mathbb{R}^n \rightarrow \mathbb{R}_+$ is a bounded Borel measurable function.
2. For every b , the resampling measure $\mu_b(\cdot)$ is a Borel probability measure.
3. The function mapping the bid b to the resampling measure $\mu_b(\cdot)$ is measurable w.r.t to the Borel σ -algebra on the space of Borel probability measures over \mathbb{R}^n .

First, we use the measure $\mu_b(\cdot)$ to define a signed measure $\nu_{b,i}(B) = b_i\mu_b(B) - \int_0^{b_i} \mu_{u,b-i}(B)du$ which has the property:

$$\int_{\hat{b} \in \mathbb{R}^n} A_i(\hat{b}) d\nu_{b,i} = b_i \mathbf{E}_{\hat{b} \sim \mu_b} [A_i(\hat{b})] - \int_0^{b_i} \mathbf{E}_{\hat{b} \sim \mu_{u,b-i}} [A_i(\hat{b})] du ,$$

that is, integrating A_i with respect to $\nu_{b,i}$ is equivalent to computing the Archer-Tardos prices.

Lemma 36 *The function $\nu_{b,i}(B) = b_i\mu_b(B) - \int_0^{b_i} \mu_{u,b-i}(B)du$ is a finite signed measure satisfying*

$$\int_{\hat{b} \in \mathbb{R}^n} A_i(\hat{b}) d\nu_{b,i} = b_i \mathbf{E}_{\hat{b} \sim \mu_b} [A_i(\hat{b})] - \int_0^{b_i} \mathbf{E}_{\hat{b} \sim \mu_{u,b-i}} [A_i(\hat{b})] du$$

for any bounded A_i

Proof: First, we show that $\nu_{b,i}(B)$ is a finite signed measure. Since μ_b is a probability measure, we have $\mu_b(B) \leq 1$ for all B . Thus, $\nu_{b,i}(B)$ is well-defined and finite for all Borel sets B (note that the integral is well defined by our assumptions on the measurability of μ_b). From this it is easy to see that $\nu_{b,i}(\emptyset) = 0$ because $\mu_b(\emptyset) = 0$. It remains to show countable additivity, i.e. $\sum_{k=1}^{\infty} \nu_{b,i}(B_k) = \nu_{b,i}(\cup_k B_k)$, which follows because integrals obey countable additivity for nonnegative functions (see Fact 77):

$$\begin{aligned} \sum_{k=1}^{\infty} \nu_{b,i}(B_k) &= \sum_{k=1}^{\infty} \left(b_i \mu_b(B_k) - \int_0^{b_i} \mu_{u,b-i}(B_k) du \right) = \sum_{k=1}^{\infty} b_i \mu_b(B_k) - \int_0^{b_i} \sum_{k=1}^{\infty} \mu_{u,b-i}(B_k) du \\ &= b_i \mu_b(\cup_k B_k) - \int_0^{b_i} \mu_{u,b-i}(\cup_k B_k) du = \nu_{b,i}(\cup_k B_k) . \end{aligned}$$

Second, we show from first-principles that integrating A_i with respect to $\nu_{b,i}$ is equivalent to calculating the Archer-Tardos prices for A_i . We begin by showing this equality for characteristic functions over Borel measurable sets. The proof for more general functions (in our case A_i) can be built-up from characteristic functions precisely as in the definition of an integral, so we omit it (see Definition 46). Let 1_B be the characteristic function of a Borel

measurable set. By definition of an integral, $\int 1_X d\nu = \nu(X)$, and plugging in we observe the desired equality:

$$\begin{aligned} b_i \mathbf{E}_{\hat{b} \sim \mu_b} [1_B(\hat{b})] - \int_0^{b_i} \mathbf{E}_{\hat{b} \sim \mu_{u, b-i}} [1_B(\hat{b})] du &= b_i \mu_b(B) - \int_0^{b_i} \mu_{u, b-i}(B) du \\ &= \int_{\hat{b} \in \mathbb{R}^n} 1_B(\hat{b}) d\nu_{b,i} . \end{aligned}$$

■

The general version of the characterization theorem shows that the payment functions precisely correspond to the density function $\rho_{b,i}^\mu(\hat{b})$ relating $\nu_{b,i}$ to μ_b (i.e. the Radon-Nikodym derivative of $\nu_{b,i}$ with respect to μ_b — its existence is guaranteed by the absolute continuity that figures in the characterization theorem 37 below). In this setting, we can equivalently define the associated coefficients $c_i^\mu(\hat{b}, b)$ as the function that satisfies

$$b_i c_i^\mu(\hat{b}, b) = \rho_{b,i}^\mu(\hat{b}) .$$

Theorem 37 (Characterizing single-call reductions) (Generalization of Theorem 31.)

A single-call single-parameter reduction $(\mu, \{\lambda_i\})$ for the set of all monotone bounded single-parameter allocation functions satisfies truthfulness, individual rationality, and no positive transfers in expectation if and only if the following conditions are met:

1. *The distribution μ is such that for all monotone, locally bounded A , the randomized allocation procedure $\mathcal{A}_i(b)$ is monotone in expectation, i.e., for all agents i , for all b , and $b'_i \geq b_i$, $\mathbf{E}[\mathcal{A}_i(b)] \leq \mathbf{E}[\mathcal{A}_i(b', b_{-i})]$ (see Lemma 39 for further discussion).*
2. *For all i , and for all Borel measurable sets B , the measure $\mu_b(B) \neq 0$ if $\int_0^{b_i} \mu_{u, b-i}(B) du \neq 0$, or equivalently, the signed measure $\nu_{b,i}$ is absolutely continuous w.r.t. measure μ_b .*
3. *The payment functions $\lambda_i(A(\hat{b}), \hat{b}, b)$ satisfy*

$$\lambda_i(A(\hat{b}), \hat{b}, b) = \rho_{b,i}^\mu(\hat{b}) A_i(\hat{b}) + \lambda_i^0(\hat{b}, b) \text{ a.s.}$$

where $\mathbf{E}_{\hat{b} \sim \mu_b} [\lambda_i^0(\hat{b}, b)] = 0$ and $\rho_{b,i}^\mu(\hat{b})$ is the density function relating $\nu_{b,i}$ to μ_b .

(Almost surely, or a.s., means that it holds everywhere except for a set with measure zero under $\mu_b(\cdot)$.)

Proof:

Necessity We first prove the necessity of the three conditions above. The first condition, that \mathcal{A} is monotone in expectation, follows directly from Archer-Tardos characterization of truthful allocation functions. The second and third conditions, as we prove below, are necessary for the expected payment to take the form required by the Archer-Tardos characterization.

We now write down the truthful payments give by the Archer-Tardos characterization, and rewrite it using the signed measure $\nu_{b,i}$.

$$\begin{aligned} \mathbf{E}[\mathcal{P}_i] &= b_i \mathbf{E}[\mathcal{A}_i(b)] - \int_0^{b_i} \mathbf{E}[\mathcal{A}_i(u, b_{-i})] du \\ &= b_i \mathbf{E}_{\hat{b} \sim \mu_b} [A_i(\hat{b})] - \int_0^{b_i} \mathbf{E}_{\hat{b} \sim \mu_{u, b_{-i}}} [A_i(\hat{b})] du \\ &= \int_{\hat{b} \in \mathbb{R}^n} A_i(\hat{b}) d\nu_{b,i} . \end{aligned}$$

where the last equality follows from the definition of the signed measure $\nu_{b,i}$, and Lemma 36.

By definition of the reduction, we can write the expected payment as:

$$\mathbf{E}[\mathcal{P}_i] = \int_{\hat{b} \in \mathbb{R}^n} \lambda_i(A(\hat{b}), \hat{b}, b) d\mu_b .$$

Equating these two gives

$$\int_{\hat{b} \in \mathbb{R}^n} \lambda_i(A(\hat{b}), \hat{b}, b) d\mu_b = \mathbf{E}[\mathcal{P}_i] = \int_{\hat{b} \in \mathbb{R}^n} A_i(\hat{b}) d\nu_{b,i} . \quad (4.15)$$

Next, we define the normalized payment function $\tilde{\lambda}$ as

$$\tilde{\lambda}_i(A(\hat{b}), \hat{b}, b) = \lambda_i(A(\hat{b}), \hat{b}, b) - \lambda_i(0^n, \hat{b}, b) .$$

By (4.15), $\int_{\hat{b} \in \mathbb{R}^n} \lambda_i(0^n, \hat{b}, b) d\mu_b(B) = 0$, and therefore we may write

$$\int_{\hat{b} \in \mathbb{R}^n} \tilde{\lambda}_i(A(\hat{b}), \hat{b}, b) d\mu_b = \int_{\hat{b} \in \mathbb{R}^n} A_i(\hat{b}) d\nu_{b,i} .$$

If the above equality were to hold for all bounded, monotone, measurable allocation functions A , then by Lemma 38, this implies for all Borel measurable sets $X \subseteq \mathbb{R}^n$:

$$\int_{\hat{b} \in X} \tilde{\lambda}_i(A(\hat{b}), \hat{b}, b) d\mu_b = \int_{\hat{b} \in X} A_i(\hat{b}) d\nu_{b,i} . \quad (4.16)$$

This statement would be intuitive if we allowed A_i to be any function — we could pick the function $A'_i(b) = 1_X(b)A_i(b)$, i.e. we could zero A_i except on X , and plug back into the previous equality. Unfortunately, this A'_i is not monotone. The work of Lemma 38 is to show

that the space of bounded, monotone functions is still sufficiently general as to guarantee equality for any Borel measurable set X .

Having derived Equation (4.16), we now show how it makes conditions two and three in theorem necessary. If we substitute the constant function $A_i(b) = 1$ into (4.16), we see that for all measurable X

$$\int_{\hat{b} \in X} \tilde{\lambda}_i(1^n, \hat{b}, b) d\mu_b = \int_{\hat{b} \in X} d\nu_{b,i} ,$$

that is, $\tilde{\lambda}_i(1^n, \hat{b}, b)$ satisfies the definition of the derivative of $\nu_{b,i}$ w.r.t μ_b , and therefore $\rho_{b,i}^\mu(\hat{b}) = \tilde{\lambda}_i(1^n, \hat{b}, b)$. Thus, given that finite payments λ exist it follows that the density relating $\nu_{b,i}$ to μ_b , namely $\rho_{b,i}^\mu(\hat{b})$, also exists and is finite. But given that both μ_b and $\nu_{b,i}$ are finite measures, this also means that $\nu_{b,i}$ is absolutely continuous w.r.t. μ_b . If not, then there exists a Borel measurable set V such that $\nu_{b,i}(V) \neq 0$ but $\mu_b(V) = 0$. We run into an immediate contradiction as follows:

$$0 = \int_{\hat{b} \in V} \rho_{b,i}^\mu(\hat{b}) d\mu_b = \int_{\hat{b} \in V} d\nu_{b,i} = \nu_{b,i}(V) \neq 0.$$

Thus we have proved that condition two, absolute continuity of $\nu_{b,i}$ w.r.t. μ_b , is necessary. Returning to (4.16), by the definition of $\rho_{b,i}^\mu(\hat{b})$ we can write

$$\begin{aligned} \int_{\hat{b} \in X} \tilde{\lambda}_i(A(\hat{b}), \hat{b}, b) d\mu_b &= \int_{\hat{b} \in X} A_i(\hat{b}) \rho_{b,i}^\mu(\hat{b}) d\mu_b \\ \int_{\hat{b} \in X} \left(\tilde{\lambda}_i(A(\hat{b}), \hat{b}, b) - A_i(\hat{b}) \rho_{b,i}^\mu(\hat{b}) \right) d\mu_b &= 0 \end{aligned}$$

for all Borel measurable sets $X \subseteq \mathbb{R}^n$. By a standard argument (Fact 82), this implies

$$\tilde{\lambda}_i(A(\hat{b}), \hat{b}, b) - A_i(\hat{b}) \rho_{b,i}^\mu(\hat{b}) = 0$$

almost surely with respect to $\mu_b(B)$, the third condition. Thus we have shown that all the three conditions are necessary.

Sufficiency We now show that the three stated conditions are sufficient. In a single-parameter setting, for a mechanism to be truthful, we simply need the allocation function to be monotone in expectation, and the payment function must satisfy the Archer-Tardos payment functions. Condition one guarantees that the allocation function output by the single-call reduction is a monotone in expectation allocation function. It remains to show that the second and third conditions result in payments that agree with Archer-Tardos payments. Given condition two, we see that $\nu_{b,i}$ is absolutely continuous w.r.t the resampling measure μ_b , and thus by Radon Nikodym theorem, the density function $\rho_{b,i}^\mu(\cdot)$ is finite and

exists. All we need to show is that under the formula of $\lambda_i(A(\hat{b}), \hat{b}, b)$ described in condition three, we have

$$\int_{\hat{b} \in \mathbb{R}^n} \lambda_i(A(\hat{b}), \hat{b}, b) d\mu_b = b_i \mathbf{E}[\mathcal{A}_i(b)] - \int_0^{b_i} \mathbf{E}[\mathcal{A}_i(u, b_{-i})] du.$$

Once we substitute the formula for $\lambda_i(A(\hat{b}), \hat{b}, b)$ from condition three, this equality follows from the definition of $\rho_{b,i}^\mu(\cdot)$ and $\nu_{b,i}$. \blacksquare

Lemma 38 *Let μ and ν be finite measures (possibly signed), and let $g : \mathbb{R}_+^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ be a function with $g(0, \hat{b}) = 0$ satisfying*

$$\int_{\hat{b} \in \mathbb{R}^n} g(A(\hat{b}), \hat{b}) d\mu = \int_{\hat{b} \in \mathbb{R}^n} A_i(\hat{b}) d\nu$$

for all Borel measurable functions $A : \mathbb{R}^n \rightarrow \mathbb{R}_+^n$ where A is bounded and monotone in the sense that $b' \geq b \Rightarrow A(b') \geq A(b)$.

Then for any such A and all Borel measurable sets $X \subseteq \mathbb{R}^n$,

$$\int_{\hat{b} \in X} g(A(\hat{b}), \hat{b}) d\mu = \int_{\hat{b} \in X} A_i(\hat{b}) d\nu .$$

Proof: First, assume that the characteristic function of X can be written as the difference of two $\{0, 1\}$ monotone functions, that is, $1_X(b) = f^+(b) - f^-(b)$ where f^+ and f^- are monotone functions mapping \mathbb{R}^n to $\{0, 1\}$. Note that this includes all rectangular parallelepipeds (a product of open, closed, or half-open intervals).

Define as $A_i^+(b) = A_i(b) \cdot f^+(b)$ and $A_i^-(b) = A_i(b) \cdot f^-(b)$. Note that for any bounded, monotone, measurable A , the functions A^+ and A^- are similarly bounded and monotone. Therefore the conditions of the lemma imply

$$\int_{\hat{b} \in \mathbb{R}^n} g(A^+(\hat{b}), \hat{b}) d\mu = \int_{\hat{b} \in \mathbb{R}^n} A_i^+(\hat{b}) d\nu$$

and

$$\int_{\hat{b} \in \mathbb{R}^n} g(A^-(\hat{b}), \hat{b}) d\mu = \int_{\hat{b} \in \mathbb{R}^n} A_i^-(\hat{b}) d\nu$$

Taking the difference, we get

$$\int_{\hat{b} \in \mathbb{R}^n} \left(g(A^+(\hat{b}), \hat{b}) - g(A^-(\hat{b}), \hat{b}) \right) d\mu = \int_{\hat{b} \in \mathbb{R}^n} \left(A_i^+(\hat{b}) - A_i^-(\hat{b}) \right) d\nu .$$

Note that $A^+ = A^-$ everywhere except on the set X , so the integrands are only nonzero on X , thus we can replace \mathbb{R}^n with X in the integrals:

$$\int_{\hat{b} \in X} \left(g(A^+(\hat{b}), \hat{b}) - g(A^-(\hat{b}), \hat{b}) \right) d\mu = \int_{\hat{b} \in X} \left(A_i^+(\hat{b}) - A_i^-(\hat{b}) \right) d\nu .$$

Now, note that on X , $A^+ = A$ and $A^- = 0$. Thus, also using the fact $g(A^-(\hat{b}), \hat{b}) = 0$, we have

$$\int_{\hat{b} \in X} g(A(\hat{b}), \hat{b}) d\mu = \int_{\hat{b} \in X} A_i(\hat{b}) d\nu \ ,$$

as desired.

To show that the lemma holds for all Borel measurable sets X , we observe that it holds for all rectangular parallelepipeds (a product of open, closed, or half-open intervals) by the above argument. Since the set of rectangular parallelepipeds is closed under finite intersections, the lemma applies to all finite intersections of rectangular parallelepipeds, which is the π -system that generates the Borel σ -algebra of \mathbb{R}^n .

Additionally, if the lemma holds for a countable sequence of disjoint sets X_k , then it clearly holds for their union as well, implying that the sets for which the lemma is true must be a λ -system.

Therefore, by Dynkin's π - λ theorem, the λ -system (the sets satisfying the lemma) must contain all sets in the σ -algebra generated by the π -system (the set of rectangular parallelepipeds) — namely, it must contain all sets in the Borel σ -algebra of \mathbb{R}^n . Thus, the lemma must hold for all Borel measurable sets X . ■

Monotonicity and μ_b

Theorem 37 requires μ_b to be such that $\mathcal{A}_i(b)$ is monotone in expectation. The following lemma gives a necessary condition:

Lemma 39 *Let B be a set of bids that is leftward closed with respect to b_i , i.e. if $\hat{b} \in B$, then $(u, \hat{b}_{-i}) \in B$ for all $u \in (-\infty, b_i] \cap T_i$. If $\mu_b(B)$ satisfies the monotonicity condition*

$$\Pr(\hat{b} \in B | b) = \mu_b(B)$$

is weakly decreasing in b_i . Similarly, if B is rightward closed with respect to b_i (i.e. $\hat{b} \in B$ implies $\hat{b}_{-i}u \in B$ for $u \in [b_i, \infty)$), then $\Pr(\hat{b} \in B | b)$ is weakly increasing in b_i , and if B is both rightward and leftward closed with respect to b_i then $\Pr(\hat{b} \in B | b)$ is constant in b_i .

Proof: First, we prove the case where B is rightward closed. For contradiction, let B be a rightward closed set on which f violates the statement of the lemma for some b and $b'_i > b_i$, i.e.

$$\Pr(\hat{b} \in B | b) = \mu_b(B) > \mu_{b'_i, b_{-i}}(B) = \Pr(\hat{b} \in B | b'_i, b_{-i}) \ .$$

Consider the monotone allocation function

$$A_i(b) = \begin{cases} 1, & b \in B \\ 0, & \text{otherwise.} \end{cases}$$

Noting that the $\mathbf{E}[\mathcal{A}(b)] = \mu_b(B)$, we have

$$E[\mathcal{A}_i(b'_i, b_{-i})] = \mu_{b'_i, b_{-i}}(B) < \mu_b(B) = E[\mathcal{A}_i(b)] .$$

Thus, under this allocation function, bidder i lowers her expected utility by raising her bid to b'_i , contradicting the monotonicity condition.

Finally, any leftward closed set B is the complement (probabilistically) of a rightward closed set, therefore $\Pr(\hat{b} \in B|b)$ must be weakly decreasing. For a set B that is both leftward and rightward closed, the theorem follows because $\Pr(\hat{b} \in B|b)$ must be both weakly increasing and weakly decreasing. ■

Part II

Markets

Chapter 5

Non-Convex Production and Complexity Equilibria

The convexity assumptions required for the Arrow-Debreu theorem are reasonable and realistic for preferences; however, they are highly problematic for production because they rule out economies of scale. In this chapter, we take a complexity-theoretic look at economies with non-convex production. It is known that in such markets equilibrium prices may not exist; we show that it is an intractable problem to achieve Pareto efficiency, the fundamental objective achieved by equilibrium prices. The same is true for core efficiency or any one of an array of concepts of stability, with the degree of intractability ranging from $F\Delta_2^P$ -completeness to PSPACE-hardness. We also identify a novel phenomenon that we call *complexity equilibrium* in which agents quiesce, not because there is no way for any one of group of them to improve their situation, but because discovering the changes necessary for (individual or group) improvement is intractable. In fact, we exhibit a somewhat natural distribution of economies that has an average-case hard complexity equilibrium.

5.1 Introduction

General Equilibrium Theory studies stable outcomes in markets — outcomes where each agent is doing as well as he can given the actions of others [65]. In the standard model, a market consists of consumers with initial endowments (vectors of goods) and preferences (utility functions), and firms with production sets (specifying what vector(s) of goods can be produced with each combination of raw materials); consumers own shares in firms.

By far the most studied kind of stable market outcome is the price equilibrium: Each firm optimizes its profit at market prices, and each consumer optimizes her utility at the same prices, selling her endowment and purchasing her preferred bundle of goods. Magically, this uncoordinated activity *clears the market*: no goods are left unsold, and all demand is satisfied. Most importantly, the resulting allocation of goods is *efficient* in the sense of Pareto: there is no allocation that is better in the sense that it dominates, in terms of

utility, the allocation achieved via the price mechanism (this is known as the First Theorem of Welfare Economics).

Price equilibria had been studied by economists since the mid 19th century, but it was not until 1954 that Arrow and Debreu [5] made the idea irresistibly powerful and attractive by proving that (under assumptions) an equilibrium price vector is guaranteed to always exist. This result promises a kind of *Arrow-Debreu paradise*, where equilibrium is both beneficent (it achieves Pareto efficiency) and universally guaranteed¹. The theory has spawned an entire area of Economics, and of course more recently a variety of results in Algorithmic Game Theory, including many algorithms for special cases (see [76], Chapters 5 and 6 for a survey, as well as [56, 51, 52] for production-specific algorithms).

There are of course wrinkles in General Equilibrium Theory. The existence proof in [5] is non-constructive, and this has been shown to imply some form of intractability, weaker than NP-completeness [79, 19]. The basic theorem holds for a very simplified model; in more realistic models parameters may be stochastic, time-varying, and generation-specific, among many other complications, and much work has been done addressing such difficulties. The model also hides tricky externalities (for example, production, or consumption, by one can harm the environment for all). Many other objections (e.g., that goods are available at different places and times) can be absorbed in the model by enlarging it. The focus of this paper is one of the most fundamentally objectionable assumptions of the theory, namely *the convexity assumption for production*.²

Convexity in utilities is quite natural: it states that you may draw less pleasure from your tenth evening dress than you did from your first. In contrast, convexity in production is very questionable because it *rules out economies of scale*. In other words, producing the hundredth airplane cannot, in any way, be easier than producing the first one. In the absence of this utterly unrealistic assumption — that is to say, in realistic economies — a price equilibrium may not exist, and thus the First Theorem cannot guarantee Pareto efficiency: Paradise lost.

Market Equilibrium Theory without Convexity of Production Since price equilibria may not exist in the absence of convexity in production, economists have studied the set of Pareto optima (which do generally exist). The first work in this line was by Guesnerie [33], whose stated goal was “to characterize precisely Pareto-optimal states and to examine the possibility of achieving them in a decentralized economy” (a task which is, as we point out, unattainable for reasons of complexity). Since then, a vast literature has developed; for two excellent surveys, see [17, 1] — the first one actually contains a discussion of computation.

¹In an excellent Microeconomics textbook [60] one reads after the statement of the First Theorem: “You should now be hearing choirs of angels and choruses of trumpets. The invisible hand of the price mechanism produces equilibria that cannot be improved upon.” The author goes on to expose and discuss the many problems of the theory.

²Convexity in production refers to convexity of the set of net production vectors. For example, if a firm can produce according to net production (input/output) vectors y_1 and y_2 , then it can also produce according to any production vector $\alpha y_1 + (1 - \alpha)y_2$ for $\alpha \in [0, 1]$.

A standard approach is to assume that firms price goods at marginal cost, in other words, postulate that prices depend on production decisions; this assumption is quite strong and rather artificial and unrealistic, but it often yields an allocation that is Pareto efficient — not always, of course. The state of the art in this direction (e.g. the marginal pricing rule of [16]) still seeks a decentralized model of agent behavior that is guaranteed to achieve Pareto efficiency. Our results, outlined next, suggest that there are huge computational impediments in the way of this ambition.

Our Results: Computational Complexity in Markets with Non-Convex Production We study markets with non-convex production from the perspective of computational complexity. To the best of our knowledge, the only other work with a computational flavor is [86], which employs dimensional communication complexity to differentiate this case from the convex one: it is shown in [86] that $\approx m \cdot n$ real numbers are needed to achieve Pareto efficiency in this case, where m is the number of goods and n the number of agents and firms, as opposed to only m in the convex case.

We show that the theory of markets with non-convex production is plagued with very bleak negative complexity results, as many natural concepts of rationality are hard for various levels of the polynomial hierarchy. We start by showing that computing a Pareto efficient outcome in a market with non-convex production is $F\Delta_2^P$ -complete. Economists regard Pareto efficiency as a sine qua non for any concept of stability or rationality in markets. Hence, our negative result for the complexity of finding Pareto efficient outcomes is a lower-bound for any “reasonable” equilibrium concept. Finally, in sections 5.4 and 5.5, we give similar results for two concepts of stability more sophisticated than Pareto efficiency: It is $F\Sigma_2^P$ -complete to tell if an allocation is in the core (no coalition of agents has an incentive to defect and create its own economy). And for a natural models of sequential production, we show that computing equilibria is $F\Delta_3^P$ -complete and PSPACE-hard, respectively.

Perhaps most significantly, we show in the process that such economies can have a novel kind of “equilibrium,” from which deviation may yield tremendous improvement for any and all agents, but the agents are stuck at a suboptimal solution of a particular instance of an NP-hard optimization problem. We call such a situation a *complexity equilibrium* (Definition 27). When agents are at such an equilibrium, standard complexity-theoretic assumptions imply that no computationally efficient procedure would generally allow them to improve — indeed, it is even intractable to recognize that improvement is possible. We also present a somewhat natural *average-case* NP-hard construction of a complexity equilibrium.

Apart from deliberate complexity-theoretic studies in game theory (e.g. [85]), the only other natural economic setting we know where computational complexity begets stability is voting, where outcome manipulation may be computationally intractable (e.g. [81, 36]).

5.2 Foundations and Models

In this section we introduce the standard economic model and relevant complexity classes.

The Economic Model

We will employ a slightly simplified version of the standard private ownership economy used in general equilibrium theory [65]. We define an economy E as follows:

Agents: An economy has n agents.

Goods: An economy has m divisible, tradable goods.

Utilities: Each agent i has a utility function $u_i : \mathbb{R}^m \rightarrow \mathbb{R}$ mapping bundles of goods to amounts of utility. An agent's consumption in an economy is specified by a vector of goods $x_i \in \mathbb{R}^m$, and her utility is $u_i(x_i)$. In general, we assume that $u_i(x)$ satisfies standard assumptions and is efficiently computable. Our hardness results use *Leontief* utilities:

$$u(x) = \min_j \frac{x_j}{\alpha_j},$$

i.e. goods are demanded in constant proportions specified by the parameters $\{\alpha_j\}$ (possibly 0).

Endowments: Each agent i is endowed with a quantity of each good, i.e. a vector $e_i \in \mathbb{R}^m$.

Production: Each firm in the economy is specified by a production function $f_k : \mathbb{R}_+^m \rightarrow \mathbb{R}_+^m$ where f_k maps a bundle of input goods to a bundle of output goods³. The behavior of a production firm may be specified by either the vector of inputs consumed by the firm x_k (in which case the output produced is $f_k(x_k)$), or by a net production vector $y_k = f_k(x_k) - x_k$. Note that the vector x_i denotes consumption by agent i and x_k denotes consumption by firm k .

Each agent in an economy owns⁴ a set of production firms F_i , meaning that agent i provides the inputs and receives the outputs of production. This is a slight restriction of the natural extension of the private ownership economy to this domain.

Our reductions use one very simple form of non-convex production function, namely *Leontief production functions with fixed costs*. Such a function f takes the form

$$f(x) = z \cdot \max \left(\min_j \left(\frac{x_j - \beta_j}{\alpha_j} \right), 0 \right)$$

where z is a bundle of goods and β_j is a fixed cost of each good required to have positive production. Interestingly, the addition of fixed costs is sufficient to force the agent to solve a discretized problem. This will be a key technique in our reductions.

³Standard General Equilibrium Theory specifies a production firm by a set of possible net production vectors Y_k instead of a function f_k . However, when modeling economies it is common to use a production function instead of a production set. While there are scenarios that differentiate the two representations, e.g. where firms would like to have free disposal, we will not encounter them here.

⁴In the standard private ownership economy, agents are said to own *shares* in production firms and receive the appropriate fraction of the profit. Since we avoid discussion of prices, the total ownership restriction avoids the issue of cooperative production for which there is already a literature, e.g. [21].

We make the following assumptions about production in the economy:

1. The production function $f(x)$ is efficiently computable for all x .
2. The total set of production possibilities is closed and bounded. The total set of production possibilities contains any net production vector that may be achieved by the economy. I.e. it contains a vector $y \geq 0$ if and only if there is a set of vectors $\{y_k \in Y_k\}$ such that $\sum_i e_i + \sum_k y_k = y$.
3. There is no x such that $f(x) - x \geq 0$ other than $x = 0$ (no free lunch).
4. For all $f(x)$, a bounded input x implies a bounded output $f(x)$.

With the exception of the computability assumption, these are standard or weakened versions of standard assumptions in the economic literature. The efficient computability assumption is nonstandard insofar as the issue has not been considered.

Finally, we recognize that smoothness is a common assumption in economics. While the functional forms we use for f (and u) are not smooth, they may be made smooth without affecting the results in this paper.

We also use the following standard economic vocabulary (see [65]):

Definition 22 *An economic allocation (hereafter an allocation) is an assignment (x, y) such that x_i is the vector of goods consumed by agent i and y_k is the vector of net outputs for firm k .*

An allocation is feasible if the amount consumed is less than the amount available in the economy, i.e.

$$\sum_i x_i \leq \sum_i e_i + \sum_k y_k$$

where $y_k = f_k(x_k) - x_k$ for the vector of goods x_k used as input to f_k .

Definition 23 *An economic allocation (x_1, y_1) is said to be strictly Pareto preferred to another allocation (x_2, y_2) if some agent receives more utility in allocation 1 than in allocation 2 and no agent receives less utility.*

An allocation (x^, y^*) is Pareto efficient or Pareto optimal if no feasible allocation is strictly Pareto preferred to it.*

Definition 24 *The social welfare W of an allocation (x, y) is the sum of the utilities obtained by consumers in the economy, i.e.*

$$W = \sum_i u_i(x_i).$$

The particular form of Leontief utilities gives the following:

Proposition 40 (*Free disposal under Leontief utilities.*) *When utility functions are Leontief, then $x' \geq x$ implies $u_i(x') \geq u_i(x)$.*

This allows us to make a key assumption: when an agent has a good \hat{s} and only one possible use for that good, no agent may be harmed by assuming that \hat{s} is applied to that use.

The Polynomial Hierarchy

Our computational complexity results will locate variations on the General Equilibrium problem in different classes of the polynomial hierarchy. The relevant portions of the polynomial hierarchy, Σ_k^P and Δ_k^P , are defined recursively as

$$\Sigma_0^P = \Delta_0^P = P$$

$$\Sigma_k^P = NP^{\Sigma_{k-1}^P}$$

$$\Delta_k^P = P^{\Sigma_{k-1}^P}$$

in other words, Σ_k^P is equal to NP with an oracle for Σ_{k-1}^P . A prefixed “F” denotes the corresponding class of functional problems, e.g. $F\Sigma_k^P$.

Krentel [59] defines the related complexity class OptP as the class of problems that may be expressed as the maximum (or minimum) value along any branch of a nondeterministic Turing machine. Relevant to our work, he shows that any OptP-complete problem is complete for the class $F\Delta_2^P$ and shows a similar generalization to $F\Delta_k^P$ [58]. For our purposes, we use the fact that any OptP problem is in $F\Delta_2^P$.

5.3 Computing Pareto Optima

In Economics, Pareto efficiency (the requirement that no allocation is preferred to the current one by all) is essentially a prerequisite for any reasonable solution concept or prediction. By showing negative complexity results for finding a Pareto efficient allocation, we lower bound the complexity of any equilibrium concept that achieves Pareto efficiency. Our first result classifies the hardness of computing a Pareto efficient allocation:

Theorem 41 *Computing a Pareto efficient allocation in an economy with non-convex production functions and polynomial-time computable utilities is $F\Delta_2^P$ -complete.*

We prove this theorem after introducing two gadgets.

Gadgets

The proof for Theorem 41 constructs an economy out of the following gadgets:

The Choice Gadget The choice gadget $Choice(\alpha_1 \cdot \hat{j}_1, \dots, \alpha_c \cdot \hat{j}_c)$ enforces indivisibility: given a set of production options, agent i must choose to produce α_k units of good \hat{j}_k for some $k \in \{1, \dots, c\}$ (she cannot produce a convex combination). Specifically, when agent i has a choice gadget $Choice(\dots)$, the economy includes the following:

Goods: $\hat{s}, \hat{j}_1, \dots, \hat{j}_c$.

Production: agent i owns firms with the following production functions:

$$\forall \hat{j}_k : x_{\hat{j}_k} = f_{j_k}(x_{\hat{s}}) = \alpha_k \cdot \max(x_{\hat{s}} - 1, 0)$$

In words, for each good \hat{j}_k , agent i has a production function to turn $x_{\hat{s}}$ units of good \hat{s} into $\alpha_k \cdot (x_{\hat{s}} - 1)$ units of good \hat{j}_k . where the values α_k are scalar parameters specified by $Choice(\alpha_1 \cdot \hat{j}_1, \dots, \alpha_c \cdot \hat{j}_c)$.

Endowment: $e_{i,\hat{s}} = 2$, i.e. agent i is endowed with 2 units of \hat{s} .

This construction ensures that only one good \hat{j}_k may be produced in positive quantities. If $x_{\hat{s}}^{\hat{j}_k}$ is the amount of \hat{s} used to produce \hat{j}_k , then because of fixed costs $x_{\hat{j}_k} > 0$ implies $x_{\hat{s}}^{\hat{j}_k} > 1$. Thus, if two or more goods \hat{j}_k is present in positive quantities, $x_{\hat{s}} \geq \sum_k x_{\hat{s}}^{\hat{j}_k} > 2$. However, agent i has only 2 units of good \hat{s} , so only one good may be present.

In order to ensure that all of good \hat{s} is consumed, we stipulate that no agent has any other use for good \hat{s} , either as a source of utility in consumption or as an input to production. Thus, by Proposition 40 (free disposal), we may assume that agent i will use all 2 units of good \hat{s} and, therefore, produce exactly α_k units of the chosen good \hat{j}_k .

The Limit Gadget The limit gadget enables the economy to limit the production of a specific good \hat{j} to α units. An instance of $Limit(\hat{j}, \alpha)$ consists of

Goods: \hat{j}, \hat{r} , and \hat{j} .

Production: (owned by agent i):

$$x_{\hat{j}} = f(x_{\hat{j}}, x_{\hat{r}}) = \min(x_{\hat{j}}, x_{\hat{r}}) .$$

Endowment: $e_{i,\hat{r}} = \alpha$.

The good \hat{r} acts as a limiting reagent in the production function $f(x_{\hat{j}}, x_{\hat{r}})$ — since the endowment of \hat{r} is fixed at α , agent i may produce as much \hat{j} as desired up to α units.

To enforce this limit, all production functions that produce \hat{j} are modified to produce \hat{j} , thereby forcing all of good \hat{j} to come from $f(x_{\hat{j}}, x_{\hat{r}})$ or an endowment.

Proof

First, we must note that Pareto optima always exist in our economies:

Proposition 42 *Under the general assumptions of Section 5.2, a Pareto efficient allocation always exists.*

Proof: Observe that if an allocation (x', y') is strictly Pareto preferred to (x, y) , then (x', y') must have a higher social welfare than (x, y) . Since the set of net production possibilities are closed and bounded and utilities are continuous (a standard assumption), it follows that the maximum social welfare W^* is well defined, and any allocation that achieves W^* is a Pareto optimum. ■

We now prove Theorem 41.

Proof: (Proof of Theorem 41.) First, we show that an efficient allocation may be computed in $F\Delta_2^P$.

Consider the following problem: *compute a feasible allocation (x, y) with social welfare at least W .* It may be solved in NP by guessing the goods x_i consumed by each agent and the goods x_k used as inputs by each firm. (The assumptions of computability imply that we may efficiently compute (x, y) and W from the x_i and x_k vectors.) As in Proposition 42, an allocation with optimal social welfare W^* must be Pareto efficient. Thus, a Pareto efficient allocation may be expressed as the optimum of an NP problem, so it is in OptP and therefore $F\Delta_2^P$.

To show that computing a Pareto efficient allocation is $F\Delta_2^P$ -hard, we reduce the $F\Delta_2^P$ -complete problem Weighted MAX-SAT [59] an economy. Let $(\Phi = \bigwedge_j \phi_j, \{\alpha_j\})$ be a Weighted MAX-SAT instance with n_Φ variables and m_Φ clauses, i.e. we desire a boolean assignment χ to the CNF formula Φ that maximizes $\sum_j \alpha_j \phi_j(\chi)$. The economy is:

Agents: One agent i for each variable χ_i and one agent j for each clause ϕ_j .

Goods: A utility good $\hat{\gamma}$.

For each SAT variable χ_i : two goods $\hat{\chi}_i$ and $\hat{\bar{\chi}}_i$.

For each clause ϕ_j : one good $\hat{\phi}_j$.

Utilities: All agents have $u(x) = x_{\hat{\gamma}}$, i.e. they only want $\hat{\gamma}$.

Production: Each variable agent i owns:

$$Choice(\hat{\chi}_i, \hat{\bar{\chi}}_i) ,$$

and each clause agent j owns:

$$\forall \chi_i \in \phi_j : x_{\hat{\phi}_j} = f_{\phi_j, \chi_i}(x_{\hat{\chi}_i}) = m_\Phi \cdot x_{\hat{\chi}_i}$$

$$\forall \bar{\chi}_i \in \phi_j : x_{\hat{\phi}_j} = f_{\phi_j, \bar{\chi}_i}(x_{\hat{\chi}_i}) = m_{\Phi} \cdot x_{\hat{\chi}_i}$$

$$\text{Limit} \left(\hat{\phi}_j, 1 \right)$$

$$x_{\hat{\gamma}} = f_{\gamma, \phi_j}(x_{\hat{\phi}_j}) = \alpha_j \cdot x_{\hat{\phi}_j} - \frac{1}{2} .$$

In this economy, a the goods $\hat{\chi}_i$ and $\hat{\bar{\chi}}_i$ represent an assignment to the variables χ_i . As argued previously, the choice gadget ensures that the economy will either produce precisely 1 unit of $\hat{\chi}_i$ or 1 unit of $\hat{\bar{\chi}}_i$. Thus, we say that $\chi_i = 1$ if $\hat{\chi}_i$ is present and $\chi_i = 0$ if $\hat{\bar{\chi}}_i$ is present.

Given the goods $\hat{\chi}_i$ and $\hat{\bar{\chi}}_i$, the clause agents will use them to produce clause goods $\hat{\phi}_j$. By construction the only way to use $\hat{\chi}_i$ (resp. $\hat{\bar{\chi}}_i$) is to create the clause good $\hat{\phi}_j$ for a clause j which is satisfied by $\chi_i = 1$ (resp. $\chi_i = 0$). Conversely, if ϕ_j is not satisfied by the assignment, there is no way to produce $\hat{\phi}_j$. Thus, the economy can produce $\hat{\phi}_j$ if and only if ϕ_j is satisfied by the assignment χ .

The limit gadgets ensure that exactly 1 unit of $\hat{\phi}_j$ is created for each satisfied clause. Consider the following production plan: $\frac{1}{m_{\Phi}}$ units of each χ_i good are ‘‘allotted’’ to each clause ϕ_j , and if the assignment to χ_i satisfies ϕ_j , then $m_{\Phi} \cdot \frac{1}{m_{\Phi}} = 1$ units of $\hat{\phi}_j$ are produced (otherwise the $\frac{1}{m_{\Phi}}$ units of the χ_i good are not used). This produces ≥ 1 units of $\hat{\phi}_j$ for each satisfied clause ϕ_j and, since each clause good $\hat{\phi}_j$ is limited to 1 unit, it therefore yields the maximum production of clause goods possible given the assignment χ . Since the variable goods $\hat{\chi}_i$ and $\hat{\bar{\chi}}_i$ have no use apart from creating clause goods, we may assume that 1 unit of $\hat{\phi}_j$ is present in the economy if and only if ϕ_j is satisfied by χ .

Finally, the clause goods are turned into the utility good $\hat{\gamma}$. Given that 1 unit of a clause good will be present if and only if ϕ is satisfied, the form of $f_{\gamma, \phi_j}(x_{\hat{\phi}_j})$ ensures that the amount of $\hat{\gamma}$ is precisely the weight of the assignment.

Since agents only desire $\hat{\gamma}$, it follows that any Pareto efficient allocation must create the maximum amount of $\hat{\gamma}$, i.e. solve the weighted MAX-SAT instance. Thus, computing a Pareto efficient allocation is $\text{F}\Delta_2^P$ -complete. ■

Marginal Price Equilibria

Of course some non-convex economies may have price equilibria, or, more generally, have equilibria in marginal prices; is the previous complexity result irrelevant in such favorable special cases? Alas, it is hard to recognize economies with equilibria:

Theorem 43 *It is NP-hard to distinguish an economy that has no efficient marginal price equilibria from one that has a price equilibrium.*

Proof: Given a SAT formula Φ , we construct an economy with the following properties: if Φ is unsatisfiable, the economy has a trivial price equilibrium. If Φ is satisfiable, then we get the economy in Section 4 of [17] that has no efficient price equilibrium (marginal or regular). The main trick is to manipulate the set of production possibilities for the two goods \hat{a} and \hat{b} : when Φ is satisfiable, we want it to be $T = \{x_{\hat{a}}, x_{\hat{b}} | x_{\hat{a}} \leq 2 \text{ and } x_{\hat{b}} \leq 2 \text{ and } (x_{\hat{a}} \leq 1 \text{ or } x_{\hat{b}} \leq 1)\}$ (see the picture in [17]), and when it is unsatisfiable, we want it to be $F = \{x_{\hat{a}}, x_{\hat{b}} | x_{\hat{a}} \leq 1 \text{ and } x_{\hat{b}} \leq 1\}$ (this is a convex set, so there will be a price equilibrium).

To accomplish this, we construct a SAT gadget, as above. When Φ is false, we have a false good. We give agents the technology to turn this good into any quantity of goods \hat{a} and \hat{b} from set F . Since this set is convex (and the other goods, i.e. those used in the SAT formula are not traded), the economy has a price equilibrium. Now, when Φ is true, we allow the economy to produce from T as follows: a choice gadget produces one of two intermediate goods. The first can be used in production of any vector of goods up to $\hat{a} = 1$ and $\hat{b} = 2$, while the other can be used in production of any vector of goods up to $\hat{a} = 2$ and $\hat{b} = 1$. It is straightforward to integrate the utilities and endowments to make the example work. ■

5.4 Computing Core Allocations

Pareto efficient allocations are stable in a cooperative sense, that is, with respect to a concept of deviation that requires all agents to cooperate and change their production and consumption. In this section, we consider allocations that are also stable with respect to certain selfish defections. A standard concept of rationality in economics is the *core*. An allocation is said to be in the core if no coalition would prefer to defect, i.e. no subset of agents can achieve strictly higher utility among themselves by creating a separate economy in which they are the only agents [65]. Price equilibria are always in the core; however, since price equilibria do not exist as such, we potentially lose this property when we lose convexity. A complexity equilibrium therefore is an allocation from which it is intractable to find another allocation where some agents do strictly better using only their own endowment and production technologies.

We find that computing core allocations is harder than computing Pareto optima:

Theorem 44 *Computing an allocation in the core is $F\Sigma_2^P$ -complete (and such an allocation may not exist).*

Additionally, if we relax our rationality requirements to include only single-agent deviations, then rational allocations are guaranteed to exist and are as easy to find as Pareto optima:

Theorem 45 *Computing an allocations that is rational with respect to single-agent deviations and Pareto improvement is $F\Delta_2^P$ -complete.*

The proofs for Theorems 44 and 45 follow after we introduce a few more gadgets.

Gadgets

The SAT Gadget The $SAT_\Phi(\{\hat{\chi}_i\}, \{\hat{\bar{\chi}}_i\}, \hat{\Phi}_{TRUE}, \hat{\Phi}_{FALSE})$ gadget allows an agent i to evaluate an arbitrary boolean formula Φ to produce exactly one unit of either $\hat{\Phi}_{TRUE}$ or $\hat{\Phi}_{FALSE}$. Let Φ be expressed as a tree T on which each literal in Φ is a leaf and each internal node represents the AND or OR of its children. Then the SAT gadget is described by the following subset of an economy:

Goods: For all variables χ_i in Φ : goods $\hat{\chi}_i$ and $\hat{\bar{\chi}}_i$.

For each internal node t in tree T : goods \hat{t} and $\hat{\bar{t}}$.

Let r refer to the root of the tree. Then \hat{r} and $\hat{\bar{r}}$ are synonyms for $\hat{\Phi}_{TRUE}$ and $\hat{\Phi}_{FALSE}$.

Production: For all AND nodes t in T with children c_j : the functions

$$x_{\hat{t}} = f_t(x) = \min_j x_{c_j}$$

$$\forall \hat{c}_j : x_{\hat{\bar{t}}} = f_{\bar{t}}(x) = x_{\hat{c}_j}$$

and for all OR nodes t with children c_1, \dots : two functions

$$\forall c_j : x_{\hat{t}} = f_t(x) = x_{c_j}$$

$$x_{\hat{\bar{t}}} = f_{\bar{t}}(x) = \min_j x_{\hat{c}_j}$$

Finally, we ensure that at most one unit of the true and false good exists with the following limit gadgets: $Limit(\hat{\Phi}_{TRUE}, 1)$ and $Limit(\hat{\Phi}_{FALSE}, 1)$. (When necessary, constant scalars may be added to ensure that at least one unit of $\hat{\Phi}_{TRUE}$ or $\hat{\Phi}_{FALSE}$ is created.)

In the manner of previous SAT reductions, this gadget allows direct evaluation of Φ given sufficient quantities of the goods setting each variable χ_i . The main differences from our previous reductions are that this gadget evaluates arbitrary formulas and that it explicitly signals false as well as true.

The Circle-of-Death Gadget The circle-of-death constructs a group of three agents who will continually defect unless they are given a particular good. The gadget $COD(\hat{d})$ includes the following:

Agents: Three agents 1, 2, and 3.

Goods: For each agent i , there is one source good \hat{s}_i and one product good \hat{p}_i . There is also a deactivator good \hat{d} .

Endowments: Agent i is endowed with 2 units of \hat{s}_i and nothing else.

Utilities: Agent i has utility function $u_i(x) = x_{\hat{p}_i}$.

Production: Agent i has a production function for producing the *bundle*

$$\begin{aligned} [x_{\hat{p}_{i-1}}, x_{\hat{p}_i}] &= f_i(x) \\ &= [2, 1] \cdot \max(\min(x_{\hat{s}_{i-1}} - 1, x_{\hat{s}_i} - 1), 0) \end{aligned}$$

and a “deactivated” production function

$$x_{\hat{p}_i} = f_{\hat{d}, i}(x_{\hat{s}_i}, x_{\hat{d}}) = \min(x_{\hat{s}_i}, 6x_{\hat{d}})$$

The f_i are designed with three properties:

1. Because of fixed costs, only one function f_i may be used at a time. Thus, if f_i is used, agent $i + 1$ will get 0 units of utility, agent i will get 1 unit, and agent $i - 1$ will get 2.
2. For any choice of f_i to use, the function f_{i+1} gives 2 units of utility to agent i , 1 unit to agent $i + 1$ and 0 units to agent $i - 1$. Moreover, agents i and $i + 1$, both of whom strictly prefer using f_{i+1} , have both the endowment and production technology to achieve this result in isolation. Thus, in the absence of \hat{d} , there is always a defecting coalition.
3. Agents may opt out of the circle-of-death and use the deactivator good to produce utility if present, but the “losing” agent in the cycle must be able to generate at least 1 unit of utility if the cycle is to be broken, requiring $\geq \frac{1}{2}$ units of \hat{d} .

Consequently, if $< \frac{1}{2}$ units of \hat{d} are available in the economy, then the core is empty, i.e. some pair of agents would always benefit from defecting.

Binary Counting Gadget The gadget $BCG(\hat{\gamma}, \hat{j}_0, \dots, \hat{j}_c)$ treats $x_{\hat{j}_0}, \dots, x_{\hat{j}_c}$ as the binary representation of a c -bit number x and produces x units of $\hat{\gamma}$. It includes the following:

Goods: The “counting” good $\hat{\gamma}$ and the c input bit goods \hat{j}_k .

Production: For each good \hat{j}_k :

$$x_{\hat{\gamma}} = f_k(x_{\hat{j}_k}) = 2^k \cdot x_{\hat{j}_k} .$$

Generalized Choice Gadget The gadget $GChoice(\hat{s}, x_1, \dots, x_c)$ is identical to the choice gadget (see Section 5.3) except that the input good \hat{s} is provided by the economy and the choice is over bundles x_k instead of individual goods \hat{j}_k :

Goods: A source good \hat{s} and bundles x_1, \dots, x_c over the space of all other goods \mathbb{R}^{m-1} .

Production: Agent i owns firms with the following production functions: $Limit(\hat{s}, 2)$ and

$$\forall j_k : z = f_{j_k}(x_{\hat{s}}) = x_k \max(x_{\hat{s}} - 1, 0) .$$

Proofs

Proof: (Proof of Theorem 44.) It is easy to certify that an allocation (x, y) is not in the core (i.e. testing core membership is in coNP): demonstrate a coalition and an allocation (x', y') such that the coalition strictly prefers (x', y') to (x, y) . Thus, a core allocation (x^c, y^c) may be found in $F\Sigma_2^P = FNP^{NP}$ as follows: guess the core allocation and use an oracle call to check that it is in the core.

Next, we give the reduction from the $F\Sigma_2^P$ -complete problem $\Sigma_2\text{SAT}$ [3]: find x_1 such that $\forall x_2 \Phi(x_1, x_2) = 1$. Let n_Φ be the number of variables in Φ .

We will describe the economy in stages. The globally relevant parts of the economy include:

Agents: Two directors A and B .

For each variable χ in x_1 : two agents χ and $\bar{\chi}$.

Goods: One utility good $\hat{\gamma}_i$ for each agent i .

One deactivator good for a circle-of-death, \hat{d} .

Utilities: Agent i desires his utility good, i.e. $u_i(x) = x_{\hat{\gamma}_i}$.

Production: A circle of death $COD(\hat{d})$.

Director A wants Φ to be true and director B wants it to be false. The circle of death will ensure that Φ can never actually be falsified in a core allocation. (Note that the χ agents are defined only for the variables in x_1 .)

In the first stage, director B decides whether to pick x_1 himself or defer to director A :

Goods: Authorization goods \hat{a}_A and \hat{a}_B .

A choice source good $\hat{s}_{D,\chi}$ for each director D and variable χ in Φ .

Production: Director B has

$$Choice((2n_\Phi + 2)\hat{a}_A, (2n_\Phi + 1)\hat{a}_B)$$

Director A has $Limit(\hat{d}, 1)$ and

$$x_{\hat{d}} = f_{\hat{d}}(x) = x_{a_A} .$$

For each variable χ , each director D has $Limit(\hat{s}_{D,\chi}, 2)$ and

$$x_{s_{D,\chi}} = f_{s_{D,\chi}}(x) = x_{a_D} .$$

The director D whose good \hat{a}_D is chosen will produce exactly two units of $\hat{s}_{D,\chi}$ for each variable χ . If A is chosen, he will also produce one unit of \hat{d} . (Because of the limit gadgets, the directors have no other way to consume all the authorization goods, so, following Proposition 40, we may assume that they do. This logic carries through the remainder of the construction.)

Next, the chosen director picks x_1 . Both directors D have the following infrastructure (note that we use generalized choice gadgets).

Goods: For each variable χ in x_1 , two choice authorization goods $\hat{a}_{D,\chi}$ and $\hat{a}_{D,\bar{\chi}}$, and two assignment goods $\hat{\chi}$ and $\hat{\bar{\chi}}$.

For each variable χ in x_2 , a choice source good $\hat{s}_{D,\chi}$ and two assignment goods $\hat{\chi}$ and $\hat{\bar{\chi}}$.

Production: Each director D owns, for each variable $\chi \in x_1$:

$$GChoice(\hat{s}_{D,\chi}, 3\hat{a}_{D,\chi}, 3\hat{a}_{D,\bar{\chi}})$$

and for each variable $\chi \in x_2$:

$$GChoice(\hat{s}_{D,\chi}, \hat{\chi}, \hat{\bar{\chi}}) .$$

Each χ agent has (for both directors): $Limit(x_{\hat{\chi}}, 1)$ and

$$x_{\hat{\chi}} = f_{\hat{\chi}}(x) = x_{\hat{a}_{D,\chi}}$$

(the $\bar{\chi}$ agents have similar functions).

In essence, the director produces x_2 himself; however, though he chooses x_1 , he must delegate the production of the $\hat{\chi}$ goods for x_1 to the χ and $\bar{\chi}$ agents. The extra authorization good \hat{a} will later be used to generate utility.

Next Φ is evaluated:

Goods: True and false goods $\hat{\Phi}_{TRUE}$ and $\hat{\Phi}_{FALSE}$ to represent the value of Φ .

Production: Director B has a

$$SAT_{\Phi}(\{\hat{\chi}_i\}, \{\hat{\bar{\chi}}_i\}, \hat{\Phi}_{TRUE}, \hat{\Phi}_{FALSE})$$

gadget to evaluate Φ given the $\hat{\chi}$ and $\hat{\bar{\chi}}$ goods.

Note that a coalition of B and the χ agents may evaluate Φ without the help of any other agents. Moreover, those same agents are required to compute Φ even if A is involved. This will restrict the possible defecting coalitions.

Finally, agents receive their payoffs:

Goods: For each χ (resp. $\bar{\chi}$) agent and each director D , an intermediate utility good $\hat{a}_{D,\chi}$ (resp. $\hat{a}_{D,\bar{\chi}}$).

Production: Director A has $Limit(\hat{\gamma}_A, 1)$ and

$$x_{D, \hat{\gamma}_A} = f_{\gamma_A}(x) = \min((2n + 2)x_{\Phi_{TRUE}}, x_{\hat{a}_A}) .$$

Similarly, director B has $Limit(\hat{\gamma}_B, 1)$ and

$$x_{\hat{\gamma}_B} = f_{\gamma_B}(x) = \min((2n + 2)x_{\Phi_{FALSE}}, x_{\hat{a}_B}) .$$

Each χ agent has, for each director D : $Limit(\hat{\alpha}_{D, \chi}, 1)$ and

$$x_{\hat{\alpha}_{D, \chi}} = f_{D, \alpha_\chi}(x) = x_{\hat{a}_{D, \chi}} ,$$

(and $\bar{\chi}$ have similar functions).

Finally, each χ agent has

$$\begin{aligned} x_{\hat{\gamma}_\chi} &= f_{A, \gamma_\chi, \chi}(x) = \min(x_{\hat{\alpha}_{A, \chi}}, (2n + 2)x_{\Phi_{TRUE}}) \\ x_{\hat{\gamma}_\chi} &= f_{A, \gamma_\chi, \bar{\chi}}(x) = 2 \min(x_{\hat{\alpha}_{A, \bar{\chi}}}, (2n + 2)x_{\Phi_{TRUE}}) \\ x_{\hat{\gamma}_\chi} &= f_{B, \gamma_\chi, \chi}(x) = \frac{3}{2} \min(x_{\hat{\alpha}_{B, \chi}}, (2n + 2)x_{\Phi_{FALSE}}) \\ x_{\hat{\gamma}_\chi} &= f_{B, \gamma_\chi, \bar{\chi}}(x) = \frac{3}{2} \min(x_{\hat{\alpha}_{B, \bar{\chi}}}, (2n + 2)x_{\Phi_{FALSE}}) \end{aligned}$$

(similar for the $\bar{\chi}$ agents).

Payoffs are summarized in the following table (D represents the chosen director):

Agent	$\Phi = 1,$ $D = A$	$\Phi = 0,$ $D = A$	$\Phi = 1,$ $D = B$	$\Phi = 0,$ $D = B$
A	1	0	0	0
B	0	0	0	1
χ produces x_1	1	0	0	$\frac{3}{2}$
χ does not produce x_1	2	0	0	$\frac{3}{2}$

Agents A and B fight over whether Φ is true or false. In allocations where A makes Φ true, the χ agents who don't participate in producing x_1 will get 2 units of utility and be happy, but the agents who do produce x_1 will only get 1 unit. Thus, if they can falsify Φ with B , the coalition would rather defect and get $\frac{3}{2}$. Note that if they defect, the only setting of x_1 that they can produce is the one originally chosen by A (because any other setting requires the production firms of the χ agents who got 2 and, therefore, would not want to join the coalition).

Since nobody gets any utility if Φ is not evaluated, we may assume that it is evaluated in any core allocation. Moreover, for the same reason, A must be the director if Φ is evaluated

to true, and B must be the director if Φ is false. (A coalition of agents would clearly like to defect from an allocation in which nobody receives any utility.)

For the moment, let us assume that the $\Sigma_2\text{SAT}$ instance is true, i.e. a “solution” x_1 exists. In this case, a core allocation is the one in which A picks an x_1 that solves the $\Sigma_2\text{SAT}$ problem. In this case, the only agents who could reasonably defect are B and the χ agents who produce x_1 . However, as noted, they are bound by A ’s choice of x_1 , therefore they will not be able to falsify Φ and would receive no utility if they defect. Thus, the allocation is in the core. In contrast, if A picks the wrong x_1 and still tries to satisfy Φ , that same coalition will defect and falsify Φ on its own. Finally, any allocation in which A is not in charge is precluded by the circle-of-death. The case that no solution x_1 exists is merely a subset of the cases mentioned above.

Thus, any core allocation must correspond to an x_1 that “solves” the $\Sigma_2\text{SAT}$ instance. It follows that core allocations are $\text{F}\Sigma_2^P$ -complete to compute. ■

Proof Sketch of Theorem 45: To show $\text{F}\Delta_2^P$ -hardness, first, assume a single agent i to leaves the economy. The defecting agent will maximize his utility using only his endowment e_i and his own production functions F_i . We will call his resulting utility his batna,⁵ b_i . Note that we can compute b_i in $\text{F}\Delta_2^P$, and, given the b_i ’s, optimizing production subject to individual rationality is also in $\text{F}\Delta_2^P$. ■

5.5 Computing Equilibria in a Sequential Model

Thus far, we have considered the most common models of rationality in markets — individual defection and core rationality. However, since production often involves intermediate goods, it seems natural to consider a model that makes this sequential property explicit.

For example, consider a production sequence in which \hat{a} is transformed into \hat{b} , \hat{b} is traded and transformed into \hat{c} , and \hat{c} is again traded. Moreover, imagine that the agent who transforms \hat{b} into \hat{c} would rather just keep the \hat{b} that he gets. It seems natural that this agent should be able to wait to defect until she receives \hat{b} , but core rationality ignores this possibility.

Since we are not aware of a standard, general model of sequential production, we adopt what we believe is a natural model. Specifically, we augment the production model to specify time:

Definition 25 *A sequential production plan is a specification $(\{x_i\}, \{x_{k,t}\})$ including the vector of goods x_i consumed by agent i and the vector of production inputs $x_{k,t}$ used by firm k at time t . A feasible production plan is one in which all production inputs required at time*

⁵The acronym BATNA, meaning “Best Alternative To a Negotiated Agreement,” is commonly used in negotiation theory.

t exist in the economy, i.e.

$$\sum_k x_{k,t} \leq \sum_{\tau=0}^{t-1} \sum_k (f_k(x_{k,\tau}) - x_{k,t}) + \sum_i e_i .$$

To ensure that all sequential production plans have polynomial size, we require that each function f_k may be used at most once and that some production function f_k is used at each timestep.

First, we assume that defectors are isolated for the remainder of the production plan, i.e. an agent will participate until time t_d and then choose to deviate, after which he cannot trade and can only use his own production technologies F_i . In this setting, we classify the complexity and show that complexity equilibria exist under the natural generalization:

Theorem 46 *In a model in which defecting individuals will face subsequent isolation, it is $F\Delta_3^P$ -complete to find the optimal sequential production plan in which no individual wishes to defect.*

Proof: Membership in $F\Delta_3^P$ is the easy direction. As with core defections, given a plan that represents a defection, it is easy (efficiently computable) to check. It follows that an individually-rational production plan may be found in NP with an NP oracle (guess the plan and use the oracle to verify its rationality). To find the optimal such plan, we maximize social welfare. Thus, the problem lies in a generalization of OptP to the second level of the polynomial hierarchy and, therefore, in $F\Delta_3^P$.

To prove that this problem is $F\Delta_3^P$ -hard, we reduce from the $F\Delta_3^P$ -complete problem lexicographically maximum Σ_2 SAT [58], i.e. find the lexicographically maximum x_1 such that for all x_2 , $\Phi(x_1, x_2) = 1$ (where Φ has n_Φ variables). We define the following economy:

Agents: Two agents: agent A and B .

Goods: Two utility goods $\hat{\gamma}_A$ and $\hat{\gamma}_B$.

An initial seed good \hat{s}_A .

Approval goods \hat{a} and \hat{a}' .

For each $\chi \in \Phi$, a source good \hat{s}_χ .

For each $\chi \in \Phi$, assignment goods $\hat{\chi}$ and $\hat{\bar{\chi}}$.

True and false goods $\hat{\Phi}_{TRUE}$ and $\hat{\Phi}_{FALSE}$.

For each agent and value of Φ , intermediate utility goods $\hat{\alpha}$, e.g. $\hat{\alpha}_{A,TRUE}$.

For each $\chi \in x_1$, an “intermediate counting good” $\hat{\beta}_\chi$.

Utilities: $u_i(x) = x_{\hat{\gamma}_i}$.

Endowments: For each $\chi \in x_1$, $e_{A,\hat{s}_\chi} = 2$.

Production: For each $\chi \in x_1$, agent A has

$$GChoice(\hat{s}_\chi, [2\hat{\chi}, \hat{a}], [2\hat{\chi}, \hat{a}], \frac{1}{n_\Phi} \hat{\gamma}_A) .$$

Agent B has

$$x_{\hat{a}'} = f_{a'}(x) = |x_2| \cdot \max(x_{\hat{a}} - (|x_1| - 1), 0)$$

and for each $\chi \in x_2$

$$\begin{aligned} x_{\hat{s}_\chi} &= f_{s_\chi}(x) = 2 \cdot x_{\hat{a}'} \\ &GChoice(\hat{s}_\chi, \hat{\chi}, \hat{\chi}) . \end{aligned}$$

Agent B also has $SAT_\Phi(\{\hat{\chi}_i\}, \{\hat{\chi}_i\}, \hat{\Phi}_{TRUE}, \hat{\Phi}_{FALSE})$ to evaluate Φ .

For payoffs, agent A has

$$\begin{aligned} x_{\hat{\alpha}_{A,TRUE}} &= f_{\alpha_{A,T}}(x) = 4 \cdot x_{\hat{\Phi}_{TRUE}} \\ &Limit(\hat{\alpha}_{A,TRUE}, 1) \\ x_{\hat{\gamma}_B} &= f_{\gamma_b}(x) = 2 \cdot x_{\hat{\alpha}_{A,TRUE}} \end{aligned}$$

and agent B has

$$\begin{aligned} x_{\hat{\alpha}_{B,FALSE}} &= f_{\alpha_{B,F}}(x) = 4 \cdot x_{\hat{\Phi}_{FALSE}} \\ &Limit(\hat{\alpha}_{B,FALSE}, 1) \\ x_{\hat{\gamma}_B} &= f_{\gamma_b}(x) = 2 \cdot x_{\hat{\alpha}_{B,FALSE}} . \end{aligned}$$

Meanwhile, agent A has the following for each variable $\chi \in x_1$:

$$\begin{aligned} x_{\hat{\beta}_\chi} &= f_{\beta_\chi}(x) = \min(x_{\hat{\chi}}, 2n_\Phi \cdot \Phi_{TRUE}) \\ &Limit(\hat{\beta}_\chi, 1) . \end{aligned}$$

and a binary counting gadget

$$BCG(\hat{\gamma}_A, \hat{\beta}_{\chi_0}, \hat{\beta}_{\chi_1}, \dots) .$$

The economy functions in three stages. First, A picks x_1 . Once x_1 is chosen, B picks x_2 (the structure of approval goods \hat{a} ensures that no variable in x_2 is chosen before all variables in x_1 have been fixed). Finally, payoffs are computed based on the results of evaluating Φ . Agent B receives 2 units if it is false, and agent A receives $2 + \sigma$ units if Φ is true, where σ is the value found by taking x_1 as the binary representation of a number. Agent A also has a default option to refuse to produce x_1 , thereby generating a small amount of utility $0 < u_A < 1$. (The complicated structure of intermediate goods merely ensures the correct discretization and distribution of goods.)

Production plans come in three flavors: satisfy Φ , falsify Φ , or neither. If Φ is satisfied, then agent A will be happy. However, agent B would rather falsify Φ . Thus, if B can pick x_2 so that Φ is false, he will defect, since he does not need to interact with any other players once he has the goods specifying x_1 . The only individually rational plan that satisfies Φ will include an x_1 such that $\Phi(x_1, x_2)$ is always true.

In plans that falsify Φ , agent A gets nothing and, therefore, will defect at the beginning and choose the default option. Thus, no such plan may be individually rational.

Finally, in plans that do not compute Φ , agent B does not receive any utility, and A receives at most 1 unit (from the default option). This is strictly dominated by any plan in which Φ is made to true, so it can only be the optimal individually rational plan if for any x_1 , there exists an x_2 such that $\Phi(x_1, x_2)$ is false.

Thus, if there is an x_1 such that $\Phi(x_1, x_2)$ is always true, the optimal individually rational production plan is the one in which A picks the x_1 that maximizes σ , i.e. the lexicographically maximum x_1 . In other cases, the only individually rational option is for A to take the default option, signaling that no such x_1 exists. It follows that computing an optimal individually rational sequential production plan is equivalent to lexicographically maximum Σ_2 SAT and, therefore, is $\text{F}\Delta_3^P$ -complete. ■

Second, we observe that isolation is not always a credible threat. We would like a “subgame perfect” allocation; however, a subgame perfect equilibrium may require an exponentially large specification. Thus, we ask for a production plan that is consistent with the realization of some subgame perfect equilibrium. We show that this problem is computationally harder:

Theorem 47 *In a model in which defectors are not isolated, it is PSPACE-hard to find a sequential production plan that is consistent with a subgame perfect equilibrium.*

Proof: (Sketch.) We reduce from a TQBF instance:

$$\exists x_1 \forall x_2 \dots x_{n_T} \Phi(x_1, x_2, \dots x_{n_T}) .$$

Consider a generalized choice gadget $G\text{Choice}(\hat{s}, x_1, \dots x_c)$ in which the source good \hat{s} may be produced in the economy (instead of given as an endowment), and the agent’s choice is over bundles of goods x_k instead of individual goods \hat{j}_k .

As before, agents will use choice gadgets to set the vectors x_i , leveraging the generalizations to enforce order. In particular, instead of choosing only a good $\hat{\chi}_i$, an agent chooses a bundle of goods $\hat{\chi}_i$ and \hat{s}_{i+1} . This is the only source of \hat{s}_{i+1} , so obtaining the quantity of \hat{s}_{i+1} necessary to choose x_{i+1} necessarily requires choosing x_i .

The choices are split among two agents A and B . Agent A picks all x_i related to \exists quantifiers (i.e. all x_{2i-1}), and agent B picks x_i related to \forall quantifiers (i.e. all x_{2i}). If Φ is evaluated to true, agent A gets 2 units of utility and agent B gets 1 unit. If Φ is evaluated to

false, B gets 2 units and A gets 1. Finally, if Φ is not evaluated at all, neither agent receives any utility.

Under this payoff structure, the only credible threat an agent can make is to change his assignment of some x_i . (Refusal to assign some x_i or compute Φ is not rational.) Thus, from the perspective of subgame perfect equilibria, this economy is equivalent to a TQBF game in which players alternate picking x_i , and player A wins if Φ is true in the end. In this case, a subgame perfect equilibrium is equivalent to a winning strategy, and the production plan we seek is equivalent to the result of playing a winning strategy in this game. By definition, A will have a winning strategy if and only if the TQBF instance is true. Thus, the value of Φ in a production plan consistent with a subgame perfect equilibrium is equal to the value of the TQBF instance.

It follows that finding such a plan is PSPACE-hard. (The details of the construction follow naturally from the sequential choice mechanism of Theorem 46.)

■

5.6 Complexity Equilibria

Perhaps the most interesting complexity-theoretic phenomenon in non-convex economies is the existence of allocations that are stable because profitable deviation is computationally intractable. To formalize complexity equilibria, one must recognize that stable outcomes in Economics are defined in terms of an appropriate concept of *deviation*.

Definition 26 *A deviation scheme is a mapping \mathcal{D} assigning each feasible allocation and subset of agents a set of feasible allocations. Intuitively, if (x, y) is a feasible allocation and $S \subseteq [n]$, then $\mathcal{D}((x, y), S)$ is the set of all allocations to which the agents in S can drive the economy in one step called a deviation by S . Deviation $(x', y') \in \mathcal{D}((x, y), S)$ is a profitable deviation by S if each $i \in S$ has at least as good utility in (x', y') than in (x, y) , and at least one agent $i \in S$ has strictly better utility. A \mathcal{D} -equilibrium is an allocation (x, y) such that for all S $\mathcal{D}((x, y), S)$ contains no profitable deviations.*

Suppose that, for all allocations (x, y) $\mathcal{D}((x, y), S)$ is the set of all feasible allocations whenever $S = [n]$, and is the empty set otherwise; then \mathcal{D} -equilibria are precisely the Pareto optimal allocations. To define the core, $\mathcal{D}((x, y), S)$ contains all feasible allocations that are also feasible if the endowments, consumption, and production by agents not in S is set to zero. And for the sequential production model, $\mathcal{D}((x, y), \{i\})$ contains all allocations that can be achieved by having agent i unilaterally change her production decisions; all other values of \mathcal{D} are empty.

We can now define complexity equilibria:

Definition 27 *We say that a family of economies \mathcal{E} has complexity equilibria with respect to deviation scheme \mathcal{D} if the following problem is NP-complete: Given an allocation $(x^+, y^+)_E$ in an economy $E \in \mathcal{E}$, find a profitable \mathcal{D} -deviation.*

Theorem 48 *Economies with non-convex production possibilities have complexity equilibria with respect to Pareto improvement, core rationality, and sequential rationality. Moreover, the inefficiency of the complexity equilibrium is unbounded.*

Proof: A nonconstructive existence proof follows from the economy of Theorem 41 and the fact that it is NP-hard to satisfy $k + 1$ clauses of a CNF SAT formula given an assignment satisfying k clauses. To give a constructive example and demonstrate unbounded inefficiency, we modify the economy to satisfy a standard CNF SAT instance.

Let $\Phi = \bigwedge \phi_j$ be a CNF SAT instance with n_Φ variables χ_i and m_Φ clauses ϕ_j . We use the economy of Theorem 41 with the following modified production functions:

Production: The functions for producing $\hat{\gamma}$ are replaced by

$$\begin{aligned} x_{\hat{\gamma}} &= f_{\gamma, \Phi}(x_{\hat{\phi}_1}, \dots, x_{\hat{\phi}_{m_\Phi}}) = \min_j x_{\hat{\phi}_j} \\ x_{\hat{\gamma}} &= f_{\gamma, \hat{s}}(x_{\hat{s}_1}, \dots, x_{\hat{s}_{n_\Phi}}) = \frac{1}{4} \min_i x_{\hat{s}_i} \end{aligned}$$

where \hat{s}_i is the source good from i 's choice gadget.

As in Theorem 41, the agents may pick an assignment ϕ and produce 1 unit of the clause good for each satisfied clause. The welfare W will be the amount of $\hat{\gamma}$ produced. The production function $f_{\gamma, \Phi}$ only produces $\hat{\gamma}$ if all clause goods are present, so agents want to satisfy Φ . When Φ is satisfied, the maximum amount of $\hat{\gamma}$ is 1.

The function $f_{\gamma, \hat{s}}$ provides a back door to generate $\hat{\gamma}$. However, since there are only 2 units of each \hat{s}_i good, at most $\frac{1}{2}$ unit of $\hat{\gamma}$ will be produced. Since $f_{\gamma, \hat{s}}$ is the only way to produce $\hat{\gamma}$ apart from $f_{\gamma, \Phi}$, no allocation can achieve welfare higher than $W = \frac{1}{2}$ without satisfying Φ .

Thus, an allocation in which the agents achieve a social welfare of $\frac{1}{2}$ by using all their \hat{s}_i in $f_{\gamma, \hat{s}}$ is a complexity equilibrium — it is certainly not Pareto efficient if Φ has a solution; however, finding an allocation that is Pareto preferred requires achieving $W > \frac{1}{2}$ and therefore satisfying Φ , which is NP-hard. Moreover, if we eliminate $f_{\gamma, \hat{s}}$, the relative inefficiency (relative social welfare of the complexity equilibrium compared to a true Pareto optimum) is unbounded because the social welfare at the complexity equilibrium is 0 and the social welfare at a Pareto optimum is 1. ■

Through standard complexity theory arguments involving so-called *complexity cores*, this result implies that, unless $P = NP$, there are families of allocations on which any group of polynomial-time agents would almost always be stuck. Specifically, a result of Orponen and Schöning [77] implies the following corollary:

Corollary 49 *Unless $P = NP$, there exists an infinite set of economies \mathcal{E} with the following properties:*

1. Any set of polynomial-time agents will be stuck at complexity equilibria in almost all economies $E \in \mathcal{E}$ (in all but a finite number of economies).
2. The set \mathcal{E} is not too small — if $f(n)$ is the number of economies $E \in \mathcal{E}$ with size $|E| = n$, then $f(n)$ is superpolynomial in n .

Moreover, we show families of economies exist with complexity equilibria relative to average-case NP-hardness. We refer the reader to Bogdanov and Trevisan’s survey [15] for background on average-case complexity.

Theorem 50 *There exists a distribution of economies $(\mathcal{D}, \mathcal{E})$ with complexity equilibria from which Pareto improvement is average-case NP-hard.*

Proof: (Rough sketch.) Consider modeling an economy as an undirected graph $G = (V, E)$. Each vertex $v \in V$ corresponds to a location, and each edge corresponds to a route along which goods may be transferred. This is repeated for each time $t \in \{1, \dots, T\}$, and a vertex (location) may choose to save goods or borrow against the future (like an edge from v_t to v_{t+1} in a product graph).

A feasible configuration of the economy is a specification of the net production at each location v , the net transfer along each edge, and the net savings at v from time t to $t+1$. The production choices at v will dictate how goods are exchanged, i.e. a production process will inherently require that goods be saved or that they be exchanged with a particular neighbor u . Using non-convexity as in the choice gadget, a location is forced to make a discrete choice among possible sets of net-transfer vectors.

In the Arrow-Debreu framework, we model this economy by creating separate copies of the goods and production functions for each time t and location and edge. (Note that Arrow and Debreu [5] suggest modeling an economy over time and space by copying goods, so this is not unprecedented.) Arbitrarily, we assume one agent per vertex v .

Hardness will follow because this economic model is a superset of an edge tiling problem defined by Gurevich [34]. An edge tiling problem consists of a set of tiles T , an $n \times n$ square, and some initial conditions. The goal is to place one tile at each location in the square such that adjacent labels match and all initial constraints are satisfied.

Gurevich [34] shows that when the first row is randomly filled according to a certain “uniform” distribution (the initial conditions) and all possible sets of tiles T occur with positive probability, it is average-case NP-complete to decide if the $n \times n$ square may be tiled.

The edge tiling problem corresponds to an economy where agents are organized on a line. A production task (a tile) is chosen for each vertex at each time such that the net transfers to neighbors (the left and right labels) and the net savings (the top and bottom labels) match accordingly. Similar to the construction in Theorem 48, players only receive a payoff if the square is completely tiled.

Reducing from Gurevich shows that when production at time $t = 0$ (i.e. filling the first row) is done according to the proper distribution, *improving from payoff 0 in such an*

economy is average-case NP-hard for any distribution over possible production tasks such that every set of tiles occurs with positive probability. ■

The economy in this proof is somewhat reasonable. It is powerful because the precise distribution does not matter, provided it satisfies the very general (though questionable) condition that all sets of tiles are possible. The linear topology is an artifact of this particular proof rather than a fundamental requirement.

The main drawback of this construction is that it requires randomization over discrete production tasks. In contrast, a more natural reduction would fix the discrete choices and randomize over continuous parameters of those choices. We do not know of a natural construction that does this.

5.7 Discussion and Open Problems

We showed that economies with nonconvexities — in other words, real economies — can be theaters of extreme complexity phenomena, including a novel kind of equilibrium in which agents quiesce because of the intractability of the task of finding a better allocation. One remark here is in order: economists often respond to complexity results such as the PPAD-completeness of Nash equilibria by questioning the relevance, and plausibility in real life, of the complex games with specialized structure that arise in those reductions. In the present situation, however, the intractability is, intuitively, more “generic.” Nonconvex optimization is a hard problem, and in hard optimization problems “gaps” between optima and defaults are common. As a result, the present complexity results may be a little more compelling to economists.

One could hope for a proof that, in a well-defined sense to be determined, nonconvex economies are “often,” or even “almost always,” computationally hard. Our average-case hard construction takes a step in this direction, and, we believe, gives hope that stronger results are possible.

Chapter 6

Market Communication in Production Economies

In this chapter, we study the information content of equilibrium prices using the market communication model of Deng, Papadimitriou, and Safra [23]. We show that, in the worst case, communicating an exact equilibrium in a production economy requires a number of bits that is a quadratic polynomial in the number of goods, the number of agents, the number of firms, and the number of bits used to represent an endowment.

6.1 Introduction

In the European Union, prices are typically expressed in whole-Euro amounts (or as “nice” decimals when they are small). In contrast, buyers and sellers in the United States cling to every penny and advertise prices to the $\frac{1}{100}$ -th of a dollar. Does such accuracy serve a computational purpose? We study this question in the case of market equilibrium: how many bits of information must prices express in order to ensure that the economy achieves equilibrium?

The *market communication* model of Deng et al. [23] highlights the unusual properties of communication in markets. In standard market models, communication often comes from central authority, such as a market maker or Walras’s fictitious auctioneer [88]. This omniscient authority must broadcast enough information (e.g. prices) for each agent to decide his own behavior without further communication — because each agent has private information (e.g. an endowment), it may be that agents are ignorant of others’ equilibrium allocations. By comparison, in Yao’s basic two-party model [92], two players follow a protocol (where both may send information) to communicate enough information that both players know the answer to the problem. Here, we study the communication requirements of reaching equilibrium in the market communication model.

Classical economic treatment of communication costs studies the *dimensionality of the message space* required to communicate a Pareto-efficient outcome. In standard convex

economies, the seminal work of Arrow and Debreu [5] may be interpreted as a proof that $(m - 1)$ real numbers — i.e. normalized prices — are sufficient. Subsequent work [48, 70] shows that normalized prices are optimal. A priori, the results for convex economies are powerful because the amount of communication is independent of the number of agents and firms. Many subsequent works have sharpened and extended these results [54, 18, 86]. Of particular relevance, Calsamiglia’s introduction of parametric communication [18] precisely captures the notion that communication may leverage private information to reduce communication.

Our work focuses on the bit-wise communication requirements for reaching equilibrium — while $(m - 1)$ real numbers may be dimensionally optimal, they may hide many bits. Since most real-world applications communicate a price with fixed precision, we follow Deng et al. [23] in believing that bit-wise communication bounds are important. Related communication complexity results [75, 76] consider the problem of communicating preferences or complete allocations, while most research on market equilibria has focused on developing efficient algorithms (e.g. [25, 56, 20]). To the best of our knowledge, Deng et al. give the only result specifically applicable to this model.

Our main result gives a lower bound on the number of bits of information that must be communicated in an Arrow-Debreu market with production. We show that the number of bits depends polynomially on the number of agents, the number of firms, and the amount of private information they hold.

Our bound is significantly stronger than the bound of Deng et al. [23]. First, Deng et al. need $\Theta(\frac{n}{m})$ -bit numbers to show a $poly(n)$ lower bound, i.e. they give each agent polynomially many bits of private information. We achieve the same lower bound with a logarithmic number of such bits. Second, Deng et al. must relax the standard non-satiation requirement on utility functions;¹ we do not. Thirdly, our bound is more general because it considers a production economy.

The main shortcoming of our bound is that it critically exploits the fact that real numbers rarely sum to the same value, even if they are very close. Thus, it is unlikely to extend to approximate equilibria.

6.2 Markets and Market Communication

Market communication complexity aims to study the amount of information that prices must encode to induce equilibrium in an Arrow-Debreu economy [5].

Arrow-Debreu Markets

An Arrow-Debreu market with production consists of n agents, m goods, and l production firms (indexed by i , j , and k respectively). A bundle of goods is a vector $x \in \mathbb{R}^m$ where x_j

¹They call it “strict concavity.” Nonsatiation is required for Arrow and Debreu’s proof of the existence of equilibrium [5].

represents a quantity of good j .

Each agent has a utility function and an endowment. The utility function $u_i(x_i) : \mathbb{R}^M \rightarrow \mathbb{R}$ maps bundles of goods to utilities, and the endowment $e_i \in \mathbb{R}^m$ is a bundle of goods. In order to guarantee the existence of an equilibrium, it is sufficient to assume that u_i is strictly concave in x .

A production firm is specified by a set of net production possibilities $Y_k \subset \mathbb{R}^M$. A vector $y_k \in Y_k$ represents the net quantities of goods produced: a positive value $y_{j,k}$ represents an output of good j , and a negative value $y_{j,k}$ represents an input. Notice that at prices π , the profit of firm k may be written as $\pi \cdot y_k$. Again, the sets Y_k must satisfy convexity requirements. In particular, it is sufficient to assume the following: Y_k is closed, convex, and contains the 0 vector, and if $y \in \bigcup_k Y_k$, then $-y \in \bigcup_k Y_k$ if and only if $y = 0$.

To link production to consumption, a firm is owned by agents. Agent i may own a share $\alpha_{i,k} \in [0, 1]$ of the profits of firm k , i.e. at prices π , agent i 's budget will be the value of his endowment plus the profit derived from firms he owns, i.e.

$$M_i = \pi \cdot e_i + \sum_{k \in [l]} \sigma_{i,k} \pi \cdot y_k . \quad (6.1)$$

Since $\sigma_{i,k}$ denotes a share of firm k , it must be that $\sum_{i \in [n]} \sigma_{i,k} = 1$. We omit the precise restrictions on production sets and utility functions for brevity.

The following economic definitions are standard [65]:

Definition 28 *An economic allocation is a tuple $(\{x_i\}, \{y_k\})$ specifying the bundle x_i consumed by each agent and the production vector y_k chosen by each firm.*

Definition 29 *An economic allocation is feasible if $x_i \geq 0$, $y_k \in Y_k$, and the total demand is less than or equal to the total supply, i.e.*

$$\sum_{i \in [n]} x_i \leq \sum_{i \in [n]} e_i + \sum_{j \in [m]} y_k \quad (6.2)$$

Definition 30 *A competitive equilibrium (hereafter equilibrium) in an Arrow-Debreu market is a set of prices $\pi \in \mathbb{R}^M$ and a feasible allocation $(\{x_i\}, \{y_k\})$ such that agents maximize their utilities and firms maximize their profits at current prices, i.e.*

$$x_i \in \operatorname{argmax}_{x \in \{x | x \cdot \pi \leq e_i \cdot \pi\}} u_i(x) \quad (6.3)$$

$$y_k \in \operatorname{argmax}_{y \in Y_k} \pi \cdot y . \quad (6.4)$$

Market Communication

Deng et al. [23] define the market communication model as follows:

Definition 31 Market Communication: n agents $[n]$ have private information $x_i \in X_i$ (the sets X_i are common knowledge). Agent i wishes to compute the function $f_i(x_1, \dots, x_n)$. Another agent, agent 0 (the “invisible hand”), knows (x_1, \dots, x_n) .

A protocol is a set of functions $(g_0(\cdot), g_1(\cdot), \dots, g_n(\cdot))$ where $g_0 : X_1 \times \dots \times X_n \rightarrow X_0$, $g_{i \in [n]} : X_0 \times X_i \rightarrow \mathbb{R}$, and $g_i(g_0(x_1, \dots, x_n), x_i) = f_i(x_1, \dots, x_n)$. The amount of market communication is the number of bits in $x_0 = g_0(x_1, \dots, x_n)$.

In essence, the omniscient agent 0 computes $x_0 = g_0(x_1, \dots, x_n)$ and broadcasts x_0 to agents $i \in [n]$. Next, each agent privately uses x_i to compute $g_i(x_0, x_i) = f_i(x_1, \dots, x_n)$.

The Power of Market Communication

The addition of an omniscient agent substantially increases the model’s power: it is as powerful as standard nondeterministic communication.

Theorem 51 Assume communication costs are measured in bits. Then any problem $f(x_1, \dots, x_n)$ in NP^{CC} has an efficient market communication protocol.

Proof: By assumption, there is a communication sequence σ of poly-logarithmic length that solves the problem. Let $T = \{(i_t, \sigma_t)\}$ be a transcript of the communication, i.e. agent i_t sent σ_t at time t .

Note that agent 0 may compute T because she is omniscient. Thus, in the market communication protocol, agent 0 computes T and broadcasts it to the agents. Each agent then simulates his behavior based on T to solve the problem. The size of i_t is $\log n$, so $|T| = \Theta(|\sigma| \log n)$, thus giving an efficient market communication protocol. ■

Market Communication in Arrow-Debreu Markets

We wish to discuss the number of bits of private information an agent or firm receives; however, such private information is often given in terms of real numbers or functions. To generate a meaningful measure of each agent’s private information, we assume that endowments, utility functions and production sets are drawn from finite sets.

Specifically, an agent’s utility function u_i is drawn from a finite set \mathbb{U}_i , and an agent’s endowment is an m -dimensional vector in which each coordinate is represented in β bits. Similarly, a firm’s production set Y_k is drawn from a finite set \mathbb{Y}_k . Our bound will be a function of β , the number of possible utility functions $|\mathbb{U}_i|$ and the number of possible production sets $|\mathbb{Y}_k|$.

The goal of an agent or firm is to compute its consumption vector x_i or production vector y_k . Thus, if E represents all private information in the economy, we have $g_i = x_i(E)$ for the

agents and $g_k = y_k(E)$ for the firms. (While the definition of an equilibrium includes prices, we take the position that prices are merely a communication tool and that, at the end of the day, we only care if each agent chooses the correct allocation. Thus, we do not explicitly require agents to compute prices as part of g_i .)

For example, a trivial protocol might broadcast everyone's endowments, utility functions, and production sets; each agent individually computes the equilibrium. This naïve protocol requires $nm\beta + \sum_{i \in [n]} |\mathbb{U}_i| + \sum_{k \in [l]} \lg |\mathbb{Y}_k|$ bits of communication. In the next section, we will see that this amount of communication is nearly necessary.

6.3 A Lower Bound for the Arrow-Debreu Model

Our main result shows that the number of bits required to communicate an equilibrium in the worst case is polynomial in the number of goods, agents, firms, and bits of private information. In particular, it requires communicating the total amount of each good available in the economy as well as the utility functions of all agents and the production sets of all firms.

Theorem 52 *In the worst case, communicating a market equilibrium in the market communication model requires at least*

$$\frac{m}{2} (\beta + \lg(n - 1)) + \sum_{i \in [n] \setminus \{1\}} \lg |\mathbb{U}_i| + \sum_{k \in [l] \setminus \{1\}} \lg |\mathbb{Y}_k| \tag{6.5}$$

bits of communication to reach equilibrium.

The terms in this bound have natural interpretations. The $\frac{m}{2} (\beta + \lg(n - 1))$ term corresponds to communicating the total global endowment of resources, since the total endowment of each resource is, in general, a $\Theta(\beta + \lg n)$ -bit number. The remaining terms correspond to communicating everyone's utility functions ($\sum \lg |\mathbb{U}_i|$) and production sets ($\sum \lg |\mathbb{Y}_k|$).

Significantly, the main difference between this lower-bound and a naive protocol that broadcasts everyone's private information is that it only needs to communicate the total amount of each good i available in the economy instead of the actual endowments of each agent ($nm\beta$ bits).

This theorem has a couple of interesting special cases. In a setting where each agent/firm has only two possible utility/production functions, the number of bits of communication required is already linear in the number of agents n and firms l :

Corollary 53 *Communicating a market equilibrium in the market communication model where $|\mathbb{U}_i| = |\mathbb{Y}_k| = 2$ requires at least*

$$\frac{m}{2} (\beta + \lg(n - 1)) + n + l - O(1) \tag{6.6}$$

bits of communication.

Moreover, we can substantially improve the lower bound of Deng et al. [23]. Deng et al. achieve an $\Omega(n \log(m+n))$ lower bound using an exponentially large set of utility functions that require $\text{poly}(m, n)$ -bit numbers, i.e. $|\mathbb{U}_i| = \Omega(2^{p(m,n)})$ and $\beta = q(m, n)$ for polynomials p and q . By comparison, our construction achieves an equivalent lower bound with constant β and $|\mathbb{U}_i| = \text{poly}(m, n)$. Theorem 52 gives a stronger bound for these parameters:

Corollary 54 *Communicating a market equilibrium in the market communication model where $|\mathbb{U}_i| = \Omega(2^{p(m,n)})$ and $\beta = q(m, n)$ requires at least*

$$\Omega(m(q(m, n) + \lg(n-1)) + n \cdot p(m, n)) \tag{6.7}$$

bits of communication.

We now prove the lower bound theorem.

Proof:(of Theorem 52) We construct an economy with m goods, n agents, and l firms. The main trick is to make each combination of utility functions (or production functions) correspond to a unique prime factorization. Thus, no two combinations of utility functions (or production functions) will have the same optimal allocation.

The second trick is to leave one agent (and firm) without any private information, so the number of communication sequences is trivially lower-bounded by the number of possible equilibrium choices she may make.

The Economy. Assume m is divisible by 4, $n \geq 2$, and $l \geq 2$.

Partition the goods into four groups modulo 4, i.e.

$$M_a = \{j | j \in [m] \text{ and } j \equiv a \pmod{4}\} \tag{6.8}$$

The sets will serve the following purposes:

- Goods in M_0 are production inputs and goods in M_1 are outputs. Nobody wants goods in M_0 . Consequently, the entire supply of goods M_0 is converted to goods in M_1 . Goods in M_1 are indistinguishable to the agents, so Pareto-optimality will imply that producers maximize the total output of goods in M_1 .
- Goods in M_2 and M_3 are traded among agents. Goods are paired such that an agent balances the quantity of a good $m_2 \in M_2$ with some good $m_3 \in M_3$ to match marginal utilities.

Agents $i > 1$ have utility functions of the form

$$u_i(x_i) = \sum_{j \in M_3} \left(2\sqrt{x_{i,j} \lg c_{i,j}} + x_{i,j-1} \right) + \sum_{j \in M_1} x_{1,j} \tag{6.9}$$

where $c_{i,j} \in C_{i,j}$, and the sets $C_{i,j}$ will be determined later. Agents $i > 1$ are endowed with goods from M_0 , M_2 and M_3 only, i.e.

$$e_{i,j} = \begin{cases} e_{i,j}, & j \in M_0 \cup M_3 \\ \bar{e}, & j \in M_2 \\ 0 & \text{otherwise.} \end{cases} \quad (6.10)$$

For endowed goods, $e_{i,j} \in [2^\beta]$ is a β -bit integer and $\bar{e} = n \cdot 2^\beta$ is a large number (large enough that, in equilibrium, an agent will always keep a positive quantity of each good in M_2).

Agent 1 has the utility function

$$u_1(x_1) = \sum_{j \in M_3} (2\sqrt{x_{1,j}} + x_{1,j-1}) \quad (6.11)$$

(the first term is equivalent to setting $c_{1,j} = 2$). She is endowed with 1 unit each of goods in M_2 , and M_3 , i.e.

$$e_{1,j} = \begin{cases} 1, & j \in M_2 \cup M_3 \\ 0 & \text{otherwise.} \end{cases} \quad (6.12)$$

Note that agent 1 has no private information.

The firms have technology to convert goods in M_0 to goods in M_1 . Like agents' utilities, the production functions are parameterized by coefficients $c_{j,k} \in C_{j,k}$. We define the production of firm k in terms of a production function, i.e. firm k may transform $y_{j-1,k}$ units of good $(j-1)$ into $y_{j,k}$ units of good j according to the following function $f_{j,k}$:

$$y_{j,k} = f_{j,k}(y_{j-1,k}) = 2\sqrt{y_{j-1,k} \cdot \lg c_{j,k}} \quad (6.13)$$

This function is translated to a set of vectors to match the model.² In order to create a firm with no private information, we require that $c_{j,1} = 2$. For simplicity, we also specify that all firms are owned by agents $i > 1$.

Analysis. First, we show that agent 1 must be able to select

$$((n-1)2^\beta)^{\frac{m}{4}} \prod_{i \in [n]} |\mathbb{U}_i| \quad (6.14)$$

distinct consumption vectors. A similar proof gives a lower bound for the production side of the economy.

²The only trick to converting $f_{j,k}$ to a set is to allow firm k to produce any amount of good j between 0 and $f_{j,k}$. In equilibrium, production will always occur on the boundary defined by $f_{j,k}$, so this change is inconsequential.

Consider a good $j \in M_3$. (Note that good $(j-1)$ is in M_2 .) Let $p_j = \frac{\pi_j}{\pi_{j-1}}$ be the relative price of good j compared to good $(j-1)$. Note that each agent has a term of the form $2\sqrt{x_{i,j} \lg c_{i,j}} + x_{i,j-1}$ in u_i . In equilibrium, we know that agent i does not wish to sell good $j-1$ to get good j (or vice-versa). Thus, agent i must balance her marginal utilities from the $2\sqrt{x_{i,j} \lg c_{i,j}}$ and $x_{i,j-1}$ terms. This gives the relation

$$\frac{\partial (2\sqrt{x_{i,j} \lg c_{i,j}})}{\partial x_{i,j}} = \frac{\partial (p_j x_{i,j-1})}{\partial x_{i,j-1}} \quad (6.15)$$

$$\sqrt{\frac{\lg c_{i,j}}{x_{i,j}}} = p_j \quad (6.16)$$

$$x_{i,j} = \frac{\lg c_{i,j}}{(p_j)^2} \quad (6.17)$$

(Note that by construction, i.e. by choice of \bar{e} , this is always possible.) Since goods in M_2 and M_3 do not involve production, we know that

$$\sum_{i \in [n]} x_{i,j} = \sum_{i \in [n]} e_{i,j} . \quad (6.18)$$

Let $\alpha_j = \sum_{i \in [n]} e_{i,j}$. Using this constraint and the equations $x_{i,j} = \frac{\lg c_{i,j}}{(p_j)^2}$, it follows that

$$p_j^2 = \frac{1}{\alpha_j} \sum_{i \in [n]} \lg c_{i,j} = \frac{1}{\alpha_j} \lg \left(\prod_{i \in [n]} c_{i,j} \right) , \quad (6.19)$$

and thus

$$x_{i,j} = \frac{\alpha_j \lg c_{i,j}}{\lg \left(\prod_{i \in [n]} c_{i,j} \right)} . \quad (6.20)$$

For agent 1, we get

$$x_{1,j} = \frac{1}{\lg \left(\left(\prod_{i \in [n] \setminus \{1\}} c_{i,j} \right)^{\frac{1}{\alpha_j}} \right)} . \quad (6.21)$$

To show a lower bound, we want to show that we can choose the sets $C_{i,j}$ such that the number of possible values for $x_{1,j}$ is large. We take each set $C_{i,j}$ to contain only prime numbers. To count the number of possible values for $x_{1,j}$, consider the value

$$\prod_{i \in [n] \setminus \{1\}} (c_{i,j})^{\frac{1}{\alpha_j}} = \prod_{c \in \cup_i C_{i,j}} (c_{i,j})^{\frac{k_{c,j}}{\alpha_j}} \quad (6.22)$$

where $k_{c,j}$ is the number of times that $c_{i,j} = c$ (over i). We want to count the number of possible values for this product.

Suppose the exponents $\left\{ \frac{k_{c,j}}{\alpha_j} \right\}$ and $\left\{ \frac{k'_{c,j}}{\alpha'_j} \right\}$ yield the same value, i.e.

$$\prod_{c \in \cup_i C_{i,j}} (c_{i,j})^{\frac{k_{c,j}}{\alpha_j}} = \prod_{c \in \cup_i C_{i,j}} (c_{i,j})^{\frac{k'_{c,j}}{\alpha'_j}} . \quad (6.23)$$

Then we naturally get that

$$\prod_{c \in \cup_i C_{i,j}} (c_{i,j})^{k_{c,j} \alpha'_j} = \prod_{c \in \cup_i C_{i,j}} (c_{i,j})^{k'_{c,j} \alpha_j} . \quad (6.24)$$

Since the terms $k_{c,j} \alpha'_j$ and $k'_{c,j} \alpha_j$ are all integers by construction, each side represents a prime factorization of the same number. Since prime factorizations are unique, we get that $k_{c,j} \alpha'_j = k'_{c,j} \alpha_j$ for all c . Summing over all c and observing that $\sum_i k_{c,j} = \sum_i k'_{c,j} = n - 1$ we see that $\alpha_j = \alpha'_j$ and therefore $k_{c,j} = k'_{c,j}$.

Thus, every possible combination of $k_{c,j}$ and α_j gives a different value for the product. Taking the sets $C_{i,j}$ to be disjoint (across i) and using the fact that there are $(n - 1)2^\beta$ possible values for α_j , we get

$$(n - 1)2^\beta \prod_{i \in [n] \setminus \{1\}} |C_{i,j}| \quad (6.25)$$

possible values for the product, and, therefore, the same number of possible values for $x_{1,j}$.

Aggregating over all goods in M_3 , we see that the total number of possible vectors x_1 is

$$\prod_{j \in M_3} (n - 1)2^\beta \prod_{i \in [n] \setminus \{1\}} |C_{i,j}| = ((n - 1)2^\beta)^{\frac{m}{4}} \prod_{i \in [n]} |\mathbb{U}_i| . \quad (6.26)$$

The analysis for the firms is similar: we sketch the argument that firm 1 must be able to select

$$((n - 1)2^\beta)^{\frac{m}{4}} \prod_{k \in [l]} |\mathbb{Y}_k| \quad (6.27)$$

distinct production vectors. First, we characterize optimal production. Consider a single good $j \in M_1$ and observe that

$$\sum_{k \in [l]} y_{j-1,k} = \sum_{i \in [n]} e_{i,j-1} = \alpha_j . \quad (6.28)$$

Let $p_{j-1} = \frac{\pi_{j-1}}{\pi_j}$ be the equilibrium price of good $(j - 1)$ relative to good j . Then we know that firm k maximizes

$$y_{j,k} - p_{j-1} y_{j-1,k} = 2\sqrt{y_{j-1,k} \cdot \lg c_{j,k}} - p_{j-1} y_{j-1,k} \quad (6.29)$$

Taking the first derivative with respect to $y_{j-1,k}$ implies that $y_{j,k} = \frac{\lg c_{j,k}}{(p_{j-1})^2}$, so we repeat the analysis used for agent 1.

Because the choices of agent 1 and firm 1 are independent, all combinations of choices are possible. Thus, the total number of communication sequences must be at least

$$((n - 1)2^\beta)^{\frac{m}{2}} \prod_{i \in [n]} |\mathbb{U}_i| \prod_{k \in [l]} |\mathbb{Y}_k| \tag{6.30}$$

and the total number of bits of communication is at least

$$\frac{m}{2} (\beta + \lg(n - 1)) + \sum_{i \in [n]} \lg |\mathbb{U}_i| + \sum_{k \in [l]} \lg |\mathbb{Y}_k| . \tag{6.31}$$

■

6.4 Conclusion

Our main theorem tarnishes the power of prices, with the caveat that we demand an exact equilibrium. While $(m - 1)$ prices are sufficient, the amount of information they contain may be highly dependent on the parameters of the market.

Most significantly, they must communicate much of agents' private information, including agents' utility functions and firms' production sets. Thus, the number of bits of information they must communicate is linear in the number of agents and firms in the worst case. Consequently, even though a price is supposed to be "universal," prices must contain unique bits of information for every agent in the economy. In the context of decimal prices, this roughly translates to at least one digit for every four buyers of a good. When the number of goods is small, this is quite impractical.

It remains an open problem to give tight bounds. For example, we currently do not have any nontrivial upper bounds. Also, there are a few reasons why our lower bound may not be tight. First, instead of a multiplicative factor of $\frac{m}{2}$, one might expect a multiplicative factor of $(m - 1)$ since that is the number of prices that must be communicated. Second, the multiplicative $\log(m + n)$ factor shown by Deng et al. [23] arises from an effect not present in our construction.

A more significant open problem is to give lower bounds for communicating approximate equilibria. Since our construction is highly dependent on the fact that two sets of irrational numbers rarely sum to the same value, it is unlikely to survive when an approximate equilibrium is sufficient. Furthermore, lower bounds for approximate equilibria would be more realistic. Because market clearing is also measured to finite precision, a lower bound approximate equilibria would give a stronger result on the amount of precision required in prices.

Bibliography

- [1] C. D. Aliprantis, B. Cornet, and R. Tourky. “Economic Equilibrium: Optimality and Price Decentralization”. In: *Positivity* 6.3 (2002), pp. 205–241.
- [2] A. Archer and É. Tardos. “Truthful Mechanisms for One-Parameter Agents”. In: *FOCS '01: Proceedings of the 42nd IEEE symposium on Foundations of Computer Science*. Washington, DC, USA: IEEE Computer Society, 2001, p. 482.
- [3] Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach*. New York, NY, USA: Cambridge University Press, 2009. ISBN: 0521424267, 9780521424264.
- [4] Kenneth J. Arrow, H. D. Block, and Leonid Hurwicz. “On the Stability of the Competitive Equilibrium, II”. In: *Econometrica* 27.1 (1959), pp. 82–109. ISSN: 00129682. URL: <http://www.jstor.org/stable/1907779>.
- [5] Kenneth J. Arrow and Gerard Debreu. “Existence of an Equilibrium for a Competitive Economy”. In: *Econometrica* 22.3 (1954), pp. 265–290. ISSN: 00129682. URL: <http://www.jstor.org/stable/1907353>.
- [6] Itai Ashlagi et al. *Position auctions with budgets: Existence and uniqueness. Working paper*. 2010.
- [7] Lawrence M. Ausubel and Paul Milgrom. “The Lovely but Lonely Vickrey Auction”. In: *Combinatorial Auctions, chapter 1*. MIT Press, 2006.
- [8] Moshe Babaioff, Robert D. Kleinberg, and Aleksandrs Slivkins. “Multi-parameter Mechanisms with Implicit Payment Computation”. In: EC '13 (2013).
- [9] Moshe Babaioff, Robert D. Kleinberg, and Aleksandrs Slivkins. “Truthful mechanisms with implicit payment computation”. In: *EC '10: Proceedings of the 11th ACM conference on Electronic commerce*. Cambridge, Massachusetts, USA: ACM, 2010, pp. 43–52. ISBN: 978-1-60558-822-3. DOI: <http://doi.acm.org/10.1145/1807342.1807349>.
- [10] Moshe Babaioff, Yogeshwer Sharma, and Aleksandrs Slivkins. “Characterizing truthful multi-armed bandit mechanisms: extended abstract”. In: *EC '09: Proceedings of the tenth ACM conference on Electronic commerce*. Stanford, California, USA: ACM, 2009, pp. 79–88. ISBN: 978-1-60558-458-4.
- [11] Moshe Babaioff et al. “Informational overhead of incentive compatibility”. In: *EC '08: Proceedings of the 9th ACM conference on Electronic commerce*. Chicago, IL, USA: ACM, 2008, pp. 88–97. ISBN: 978-1-60558-169-9.

- [12] Venkatesh Bala and Mukul Majumdar. “Chaotic Tatonnement”. In: *Economic Theory* 2.4 (1992), pp. 437–45. URL: <http://ideas.repec.org/a/spr/joecth/v2y1992i4p437-45.html>.
- [13] Paul D. Berger. “Vertical Cooperative Advertising Ventures”. English. In: *Journal of Marketing Research* 9.3 (1972), pp. 309–312. ISSN: 00222437. URL: <http://www.jstor.org/stable/3149542>.
- [14] B. Douglas Bernheim and Michael D. Whinston. “Menu Auctions, Resource Allocation, and Economic Influence”. In: *The Quarterly Journal of Economics* CI.1 (1986).
- [15] Andrej Bogdanov and Luca Trevisan. “Average-Case Complexity”. In: *CoRR* abs/cs/0606037 (2008).
- [16] J.-M. Bonnisseau and B. Cornet. “Existence of equilibria with a tight marginal pricing rule”. In: *Journal of Mathematical Economics* 44.7-8 (2008). Special Issue in Economic Theory in honor of Charalambos D. Aliprantis, pp. 613–624. ISSN: 0304-4068. DOI: DOI:10.1016/j.jmateco.2007.03.001. URL: <http://www.sciencedirect.com/science/article/B6VBY-4NN1TT3-1/2/9618bea2fff2a37e22990d151e06a9df>.
- [17] Donald J. Brown. “Chapter 36 Equilibrium analysis with non-convex technologies”. In: ed. by Werner Hildenbrand and Hugo Sonnenschein. Vol. 4. *Handbook of Mathematical Economics*. Elsevier, 1991, pp. 1963–1995. DOI: DOI:10.1016/S1573-4382(05)80011-6. URL: <http://www.sciencedirect.com/science/article/B7P5Y-4FKY4C6-C/2/0bd1a73374b0b690f702691e1f7fe671>.
- [18] X. Calsamiglia. “Informational Requirements of Parametric Resource Allocation Processes”. In: *Information, Incentives, and Economic Mechanisms*. Ed. by T. Groves, R. Radner, and S. Reiter. Basil Blackwell, 1988.
- [19] Bruno Codenotti et al. “Leontief economies encode nonzero sum two-player games”. In: *SODA '06: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*. Miami, Florida: ACM, 2006, pp. 659–667. ISBN: 0-89871-605-5. DOI: <http://doi.acm.org/10.1145/1109557.1109629>.
- [20] Richard Cole and Lisa Fleischer. “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*. Victoria, British Columbia, Canada: ACM, 2008, pp. 315–324. ISBN: 978-1-60558-047-0. DOI: <http://doi.acm.org/10.1145/1374376.1374422>.
- [21] Luis C. Corchn and M. Socorro Puy. “Individual rationality and voting in cooperative production”. In: *Economics Letters* 59.1 (1998), pp. 83–90. ISSN: 0165-1765. DOI: DOI:10.1016/S0165-1765(98)00027-5. URL: <http://www.sciencedirect.com/science/article/B6V84-3T51RH8-4C/2/d54277f4bade8131d06141b8d15f55d7>.

- [22] Robert Day and Paul Milgrom. “Core-selecting package auctions”. In: *International Journal of Game Theory* 36.3-4 (July 2007), pp. 393–407. ISSN: 0020-7276. DOI: 10.1007/s00182-007-0100-7. URL: <http://link.springer.com/10.1007/s00182-007-0100-7>.
- [23] Xiaotie Deng, Christos H. Papadimitriou, and Shmuel Safra. “On the complexity of equilibria”. In: *STOC '02: Proceedings of the thirty-fourth annual ACM symposium on Theory of computing*. Montreal, Quebec, Canada: ACM, 2002, pp. 67–71. ISBN: 1-58113-495-9. DOI: <http://doi.acm.org/10.1145/509907.509920>.
- [24] Nikhil R. Devanur and Sham M. Kakade. “The price of truthfulness for pay-per-click auctions”. In: *EC '09: Proceedings of the tenth ACM conference on Electronic commerce*. Stanford, California, USA: ACM, 2009, pp. 99–106. ISBN: 978-1-60558-458-4. DOI: <http://doi.acm.org/10.1145/1566374.1566388>.
- [25] Nikhil R. Devanur et al. “Market Equilibrium via a Primal-Dual-Type Algorithm”. In: *FOCS '02: Proceedings of the 43rd Symposium on Foundations of Computer Science*. Washington, DC, USA: IEEE Computer Society, 2002, pp. 389–395. ISBN: 0-7695-1822-2.
- [26] Shahar Dobzinski and Shaddin Dughmi. “On the Power of Randomization in Algorithmic Mechanism Design”. In: *FOCS*. 2009, pp. 505–514.
- [27] Shaddin Dughmi and Tim Roughgarden. “Black-Box Randomized Reductions in Algorithmic Mechanism Design”. In: *Proceedings of the 51st IEEE symposium on Foundations of Computer Science*. 2010, pp. 775–784.
- [28] Benjamin Edelman and Michael Ostrovsky. “Strategic bidder behavior in sponsored search auctions”. In: *Decision Support Systems* 43.1 (2007), pp. 192–198.
- [29] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. “Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords”. In: *American Economic Review* 97.1 (2007), pp. 242–259. URL: <http://ideas.repec.org/a/aea/aecrev/v97y2007i1p242-259.html>.
- [30] Benjamin Edelman and Michael Schwarz. *Optimal Auction Design and Equilibrium Selection in Sponsored Search Auctions*. 2010.
- [31] Richard Engelbrecht-Wiggans and Charles M. Kahn. “Multi-Unit Pay-Your-Bid Auctions with Variable Awards”. In: *Games and Economic Behavior* 23.1 (1998), pp. 25 – 42. ISSN: 0899-8256. DOI: 10.1006/game.1997.0599. URL: <http://www.sciencedirect.com/science/article/pii/S0899825697905996>.
- [32] Eyal Even-dar, Yishay Mansour, and Uri Nadav. “On the convergence of regret minimization dynamics in concave games”. In: *STOC '09: Proceedings of the 41st annual ACM symposium on Theory of computing*. Bethesda, MD, USA: ACM, 2009, pp. 523–532. ISBN: 978-1-60558-506-2. DOI: <http://doi.acm.org/10.1145/1536414.1536486>.
- [33] Roger Guesnerie. “Pareto Optimality in Non-Convex Economies”. In: *Econometrica* 43.1 (1975), pp. 1–29. ISSN: 00129682. URL: <http://www.jstor.org/stable/1913410>.

- [34] Yuri Gurevich. “Average Case Completeness”. In: *Journal of Computer and System Sciences* 42 (1991), pp. 346–398.
- [35] Glenn W. Harrison. “Theory and Misbehavior of First-Price Auctions”. English. In: *The American Economic Review* 79.4 (1989), pp. 749–762. ISSN: 00028282. URL: <http://www.jstor.org/stable/1827930>.
- [36] Glenn W. Harrison and Tanga McDaniel. “Voting games and computational complexity”. In: *Oxford Economic Papers* 60.3 (2008), pp. 546–565. DOI: 10.1093/oep/gpm045. eprint: <http://oep.oxfordjournals.org/content/60/3/546.full.pdf+html>. URL: <http://oep.oxfordjournals.org/content/60/3/546.abstract>.
- [37] Sergiu Hart and Andreu Mas-Colell. “A Simple Adaptive Procedure Leading to Correlated Equilibrium”. In: *Econometrica* 68.5 (2000), pp. 1127–1150. ISSN: 1468-0262. DOI: 10.1111/1468-0262.00153. URL: <http://dx.doi.org/10.1111/1468-0262.00153>.
- [38] Sergiu Hart and Noam Nisan. *The Menu-Size Complexity of Auctions*. 2012.
- [39] Jason D. Hartline. *Approximation in Economic Design*. Draft. 2011. URL: <http://users.eecs.northwestern.edu/~hartline/amd.pdf>.
- [40] Jason D. Hartline and Tim Roughgarden. “Simple versus optimal mechanisms”. In: *Proceedings of the 10th ACM conference on Electronic commerce*. EC ’09. Stanford, California, USA: ACM, 2009, pp. 225–234. ISBN: 978-1-60558-458-4. DOI: 10.1145/1566374.1566407. URL: <http://doi.acm.org/10.1145/1566374.1566407>.
- [41] Xiuli He, Ashutosh Prasad, and Suresh P. Sethi. “Cooperative Advertising and Pricing in a Dynamic Stochastic Supply Chain: Feedback Stackelberg Strategies”. In: *Production and Operations Management* 18.1 (2009), pp. 78–94. ISSN: 1937-5956. DOI: 10.1111/j.1937-5956.2009.01006.x. URL: <http://dx.doi.org/10.1111/j.1937-5956.2009.01006.x>.
- [42] John Hershberger and Subhash Suri. “Erratum to “Vickrey Pricing and Shortest Paths: What is an Edge Worth?””. In: *Proceedings of the 43rd Symposium on Foundations of Computer Science*. FOCS ’02. Washington, DC, USA: IEEE Computer Society, 2002.
- [43] John Hershberger and Subhash Suri. “Vickrey Prices and Shortest Paths: What is an Edge Worth?” In: *Proceedings of the 42nd IEEE symposium on Foundations of Computer Science*. FOCS ’01. Washington, DC, USA: IEEE Computer Society, 2001. ISBN: 0-7695-1390-5.
- [44] John Hershberger, Subhash Suri, and Amit Bhosle. “On the difficulty of some shortest path problems”. In: *ACM Trans. Algorithms* 3 (1 2007), 5:1–5:15.
- [45] Darrell Hoy, Kamal Jain, and Christopher A. Wilkens. “A Dynamic Axiomatic Approach to First-Price Auctions”. In: *Proceedings of the 14th ACM Conference on Electronic Commerce*. EC ’13. 2013.
- [46] Darrell Hoy, Kamal Jain, and Christopher A. Wilkens. “Coopetitive Ad Auctions”. In: *arXiv abs/1209.0832* (2012).

- [47] Zhimin Huang, Susan X. Li, and Vijay Mahajan. “An Analysis of Manufacturer-Retailer Supply Chain Coordination in Cooperative Advertising*”. In: *Decision Sciences* 33.3 (2002), pp. 469–494. ISSN: 1540-5915. DOI: 10.1111/j.1540-5915.2002.tb01652.x. URL: <http://dx.doi.org/10.1111/j.1540-5915.2002.tb01652.x>.
- [48] Leonard Hurwicz. “Optimality and Informational Efficiency in Resource Allocation Processes”. In: *Mathematical Methods in the Social Sciences*. Ed. by K.J. Arrow, S. Karlin, and P. Suppes. Stanford, CA: Stanford University Press, 1960, pp. 27–46.
- [49] Intel Corporation. *Intel Inside Program - Anatomy of a Brand Campaign*. http://www.intel.com/pressroom/intel_inside.htm.
- [50] David Jabas. Private communication. 2010.
- [51] Kamal Jain and Kasturi Varadarajan. “Equilibria for economies with production: constant-returns technologies and production planning constraints”. In: *SODA '06: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*. Miami, Florida: ACM, 2006, pp. 688–697. ISBN: 0-89871-605-5. DOI: <http://doi.acm.org/10.1145/1109557.1109632>.
- [52] Kamal Jain, Vijay V. Vazirani, and Yinyu Ye. “Market equilibria for homothetic, quasi-concave utilities and economies of scale in production”. In: *SODA '05: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*. Vancouver, British Columbia: Society for Industrial and Applied Mathematics, 2005, pp. 63–71. ISBN: 0-89871-585-7.
- [53] Philippe Jehiel, Benny Moldovanu, and Ennio Stacchetti. “Multidimensional Mechanism Design for Auctions with Externalities”. In: *Journal of Economic Theory* 85.2 (1999), pp. 258–293. ISSN: 0022-0531. DOI: 10.1006/jeth.1998.2501. URL: <http://www.sciencedirect.com/science/article/pii/S0022053198925017>.
- [54] J. S. Jordan. “The competitive allocation process is informationally efficient uniquely”. In: *Journal of Economic Theory* 28.1 (1982), pp. 1–18. ISSN: 0022-0531. DOI: 10.1016/0022-0531(82)90088-6. URL: <http://www.sciencedirect.com/science/article/B6WJ3-4CYG8P6-J/2/b8d8f9397e31e1edf18f286e61da1fdf>.
- [55] Raja Kali. “Minimum Advertised Price”. In: *Journal of Economics & Management Strategy* 7.4 (1998), pp. 647–668. ISSN: 1530-9134. DOI: 10.1111/j.1430-9134.1998.00647.x. URL: <http://dx.doi.org/10.1111/j.1430-9134.1998.00647.x>.
- [56] Sanjiv Kapoor, Aranyak Mehta, and Vijay V. Vazirani. “An auction-based market equilibrium algorithm for a production model”. In: *Theor. Comput. Sci.* 378.2 (2007), pp. 153–164. ISSN: 0304-3975. DOI: <http://dx.doi.org/10.1016/j.tcs.2007.02.018>.
- [57] Paul Klemperer. “What Really Matters in Auction Design”. English. In: *The Journal of Economic Perspectives* 16.1 (2002), pp. 169–189. ISSN: 08953309. URL: <http://www.jstor.org/stable/2696581>.

- [58] Mark W. Krentel. “Generalizations of Opt P to the polynomial hierarchy”. In: *Theor. Comput. Sci.* 97.2 (1992), pp. 183–198. ISSN: 0304-3975. DOI: [http://dx.doi.org/10.1016/0304-3975\(92\)90073-0](http://dx.doi.org/10.1016/0304-3975(92)90073-0).
- [59] Mark W. Krentel. “The complexity of optimization problems”. In: *STOC '86: Proceedings of the eighteenth annual ACM symposium on Theory of computing*. Berkeley, California, United States: ACM, 1986, pp. 69–76. ISBN: 0-89791-193-8. DOI: <http://doi.acm.org/10.1145/12130.12138>.
- [60] David M. Kreps. *A Course in Microeconomic Theory*. Princeton University Press, 1990. ISBN: 0691042640. URL: <http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20&path=ASIN/0691042640>.
- [61] Piotr Krysta et al. “Combinatorial auctions with externalities”. In: *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1 - Volume 1*. AAMAS '10. Toronto, Canada: International Foundation for Autonomous Agents and Multiagent Systems, 2010, pp. 1471–1472. ISBN: 978-0-9826571-1-9. URL: <http://dl.acm.org/citation.cfm?id=1838206.1838437>.
- [62] Sébastien Lahaie. “Stability and Incentive Compatibility in a Kernel-Based Combinatorial Auction”. In: *AAAI*. 2010.
- [63] Renato Paes Leme, Vasilis Syrgkanis, and Éva Tardos. “Sequential auctions and externalities”. In: *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '12. Kyoto, Japan: SIAM, 2012, pp. 869–886. URL: <http://dl.acm.org/citation.cfm?id=2095116.2095186>.
- [64] David Lucking-Reiley. “Vickrey Auctions in Practice: From Nineteenth-Century Philately to Twenty-First-Century E-Commerce”. In: *Journal of Economic Perspectives* 14.3 (2000), pp. 183–192. DOI: 10.1257/jep.14.3.183. URL: <http://www.aeaweb.org/articles.php?doi=10.1257/jep.14.3.183>.
- [65] Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, 1995. ISBN: 0195073401. URL: <http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20&path=ASIN/0195073401>.
- [66] Eric Maskin and John Riley. “Asymmetric Auctions”. English. In: *The Review of Economic Studies* 67.3 (2000), pp. 413–438. ISSN: 00346527. URL: <http://www.jstor.org/stable/2566960>.
- [67] Paul Milgrom. *Putting Auction Theory to Work*. Churchill Lectures in Economics. Cambridge University Press, 2004. ISBN: 9780521536721. URL: <http://books.google.com/books?id=AkeHTU7XW4kC>.
- [68] Paul Milgrom. “Putting Auction Theory to Work: The Simultaneous Ascending Auction”. English. In: *Journal of Political Economy* 108.2 (2000), pp. 245–272. ISSN: 00223808. URL: <http://www.jstor.org/stable/10.1086/262118>.

- [69] Paul Milgrom and John Roberts. “Adaptive and sophisticated learning in normal form games”. In: *Games and Economic Behavior* 3.1 (1991), pp. 82–100. ISSN: 0899-8256. DOI: 10.1016/0899-8256(91)90006-Z. URL: <http://www.sciencedirect.com/science/article/pii/089982569190006Z>.
- [70] Kenneth Mount and Stanley Reiter. “The informational size of message spaces”. In: *Journal of Economic Theory* 8.2 (1974), pp. 161–192.
- [71] Roger B. Myerson. “Optimal Auction Design”. In: *Mathematics of Operations Research* 6.1 (1981), pp. 58–73.
- [72] Matthew Nagler. “An exploratory analysis of the determinants of cooperative advertising participation rates”. In: *Marketing Letters* 17.2 (2006), pp. 91–102. URL: <http://ideas.repec.org/a/kap/mktlet/v17y2006i2p91-102.html>.
- [73] Noam Nisan and Amir Ronen. “Algorithmic Mechanism Design”. In: *Games and Economic Behavior* 35.1-2 (2001), pp. 166–196.
- [74] Noam Nisan and Amir Ronen. “Computationally Feasible VCG Mechanisms”. In: *J. Artif. Intell. Res. (JAIR)* 29 (2007), pp. 19–47.
- [75] Noam Nisan and Ilya Segal. “The communication requirements of efficient allocations and supporting prices”. In: *Journal of Economic Theory* 129.1 (2006), pp. 192–224.
- [76] Noam Nisan et al. *Algorithmic Game Theory*. New York, NY, USA: Cambridge University Press, 2007. ISBN: 0521872820.
- [77] Pekka Orponen and Uwe Schöningh. “The density and complexity of polynomial cores for intractable sets”. In: *Information and Control* 70.1 (1986), pp. 54–68. ISSN: 0019-9958. DOI: DOI:10.1016/S0019-9958(86)80024-9. URL: <http://www.sciencedirect.com/science/article/B7MFM-4G3DGHT-2T/2/938db249da575aa0557c0c5ca2cd781c>.
- [78] Christos H. Papadimitriou. *Computational complexity*. Reading, Massachusetts: Addison-Wesley, 1994. ISBN: 0201530821.
- [79] Christos H. Papadimitriou. “On the complexity of the parity argument and other inefficient proofs of existence”. In: *J. Comput. Syst. Sci.* 48.3 (1994), pp. 498–532. ISSN: 0022-0000. DOI: [http://dx.doi.org/10.1016/S0022-0000\(05\)80063-7](http://dx.doi.org/10.1016/S0022-0000(05)80063-7).
- [80] Christos H. Papadimitriou and Christopher A. Wilkens. “Economies with non-convex production and complexity equilibria”. In: *Proceedings of the 12th ACM conference on Electronic commerce. EC '11*. San Jose, California, USA: ACM, 2011, pp. 137–146. ISBN: 978-1-4503-0261-6. DOI: 10.1145/1993574.1993595. URL: <http://doi.acm.org/10.1145/1993574.1993595>.

- [81] Ariel D. Procaccia and Jeffrey S. Rosenschein. “Junta distributions and the average-case complexity of manipulating elections”. In: *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*. AAMAS '06. Hakodate, Japan: ACM, 2006, pp. 497–504. ISBN: 1-59593-303-4. DOI: <http://doi.acm.org/10.1145/1160633.1160726>. URL: <http://doi.acm.org/10.1145/1160633.1160726>.
- [82] Kevin Roberts. *The characterization of implementable social choice rules*. In *Aggregation and Revelation of Preferences*, J-J.Laffont (ed.), North Holland Publishing Company, 1979.
- [83] Tim Roughgarden and Éva Tardos. “Do Externalities Degrade GSP’s Efficiency?” In: *The Eighth Ad Auctions Workshop*. 2012.
- [84] Paul A. Samuelson. “The Stability of Equilibrium: Comparative Statics and Dynamics”. In: *Econometrica* 9.2 (1941), pp. 97–120. ISSN: 00129682. URL: <http://www.jstor.org/stable/1906872>.
- [85] Grant Schoenebeck and Salil Vadhan. “The computational complexity of nash equilibria in concisely represented games”. In: *EC '06: Proceedings of the 7th ACM conference on Electronic commerce*. Ann Arbor, Michigan, USA: ACM, 2006, pp. 270–279. ISBN: 1-59593-236-4. DOI: <http://doi.acm.org/10.1145/1134707.1134737>.
- [86] Guoqiang Tian. “On the Informational Requirements of Decentralized Pareto-Satisfactory Mechanisms in Economies with Increasing Returns”. Oct. 2006. URL: <http://econweb.tamu.edu/tian/mechl35.pdf>.
- [87] William Vickrey. “Counterspeculation, Auctions, and Competitive Sealed Tenders”. In: *The Journal of Finance* 16.1 (1961), pp. 8–37. ISSN: 1540-6261. DOI: 10.1111/j.1540-6261.1961.tb02789.x. URL: <http://dx.doi.org/10.1111/j.1540-6261.1961.tb02789.x>.
- [88] Leon Walras. *Elements of pure economics, or, The theory of social wealth / Leon Walras ; translated by William Jaffe*. English. Published for the American Economic Association, the Royal Economic Society by Allen, and Unwin, London : 1954, 620 p. :
- [89] Wikipedia. *Coopetition*. <http://en.wikipedia.org/wiki/Coopetition>.
- [90] Christopher A. Wilkens. “Market communication in production economies”. In: *Proceedings of the 6th international conference on Internet and network economics*. WINE'10. Stanford, CA, USA: Springer-Verlag, 2010, pp. 391–401. ISBN: 3-642-17571-6, 978-3-642-17571-8. URL: <http://dl.acm.org/citation.cfm?id=1940179.1940212>.
- [91] Christopher A. Wilkens and Balasubramanian Sivan. “Single-call mechanisms”. In: *Proceedings of the 13th ACM Conference on Electronic Commerce*. EC '12. Valencia, Spain: ACM, 2012, pp. 946–963. ISBN: 978-1-4503-1415-2. DOI: 10.1145/2229012.2229084. URL: <http://doi.acm.org/10.1145/2229012.2229084>.

- [92] Andrew Chi-Chih Yao. “Some complexity questions related to distributive computing(Preliminary Report)”. In: *STOC '79: Proceedings of the eleventh annual ACM symposium on Theory of computing*. Atlanta, Georgia, United States: ACM, 1979, pp. 209–213. DOI: <http://doi.acm.org/10.1145/800135.804414>.

Appendix A

Convergence in Utility-Target Auctions

A.1 Convergence to the Egalitarian Equilibrium

Theorem 7 (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 7. 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 players 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 55), and then to at most her upper bound (Lemma 56). Combining these via induction gives our final result that the bids of all players converge near their egalitarian outcome.

Let o^* be the egalitarian outcome; let π_i^* be the corresponding utility-target of bidder i . Let $\mathbf{B}_b(o) = \sum_{i \in [n]} b_i(o)$ be the total bid for a given outcome o . Let $\mathbf{B}_X(o) = \sum_{j \in L_X} b_j(o)$, and $\mathbf{B}_X^*(o)$ be similarly defined.

First, consider all utility-targets in the egalitarian equilibrium; let z_i be the i th smallest (distinct) utility-target. Let L_i be the set of all players 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 $[\pi_j^* - \epsilon b^-(i), \pi_j^* + \epsilon b^+(i)]$.

Bidding Bounds. We now precisely define the bounding functions $b^-(\cdot)$ and $b^+(\cdot)$.

Definition 32

$$b^-(i) = 2^{2|L_{<i}|} \tag{A.1}$$

$$b^+(i) = 2^{2|L_{<i}|+|L_i|} \tag{A.2}$$

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 7

$$b^-(k) > \sum_{i=0}^{k-1} |L_i| b^+(i) \tag{A.3}$$

$$b^+(k) > \sum_{i=0}^k |L_i| b^-(i) \tag{A.4}$$

We omit the proof; it follows from manipulation of exponential sums.

Witness outcomes. Recall from Algorithm 2 that utility-targets for any given player 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 33 *Outcome o^w is a witness outcome for bidder j at the egalitarian utility-targets π^* if its total bid is tied with the egalitarian outcome, j asking for more utility at the egalitarian equilibrium results in a higher total bid for o^w than for the optimal egalitarian ad and j 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 8 *At the egalitarian equilibrium, every bidder j s.t. $\pi_i^* < vi(o^*)$ has at least one witness.*

Proof: We will prove via contradiction. Assume at the CEF egalitarian outcome o^* bidder j has no witness. Now, let bidder j 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 j who does not win with the outcome; otherwise it would be a witness for j . 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 player 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 j 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 9 *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 players 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 55 *Under assumptions A1, A2, A3 and A4, the utility-target of each bidder j in level L_i will converge to at least their lower bounds, $\pi_j^* - \epsilon \cdot b^-(i)$ if for every bidder j' in level L_{i^-} s.t. $i^- < i$, $\pi_{j'} \leq \pi_{j'}^* + \epsilon \cdot b^+(i^-)$.*

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 10 *Under the assumptions of Lemma 55, no bidder j in level L_i with a utility-target of $\pi_j \leq \pi_j^* - \epsilon \cdot b^-(i)$ will lower her utility-target.*

We will prove via contradiction. Assume for bidder j that with a utility-target of $\pi_j \leq \pi_j^* - \epsilon \cdot b^-(i)$, she wishes to lower her utility-target further. Let o be the winning outcome with bids π . As i will only lower her bid if she is losing (A1), $\pi_j > v_j(o)$. 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 ($\mathbf{B}(o^*) > \mathbf{B}(o)$), hence she must win and would not care to lower her utility-target.

For j to decrease her utility-target, by A4 she must be the highest utility bidder who is losing such that $\pi_j > v_j(o)$. By our bound, we know that for every lower utility bidder j' in level L_{i-} , $\pi_{j'} \leq \pi_{j'}^* + \epsilon b^+(i^-)$. Since o^* is the optimal winning outcome and o the currently winning outcome, $\mathbf{B}^*(o^*) \geq \mathbf{B}^*(o)$ and $\mathbf{B}(o^*) \leq \mathbf{B}(o)$.

In the egalitarian outcome, every bidder j receives the utility she bids for; hence $b_j^*(o^*) = v_j(o^*) - \pi_j^*$. 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 j and first, is requesting more utility relative to the egalitarian outcome, specifically that $\pi_{j'} > \pi_{j'}^*$. Hence, we will have $0 \geq b_{j'}(o^*) - b_{j'}^*(o^*) \geq -(\pi_{j'} - \pi_{j'}^*)$ and $0 \geq b_{j'}(o) - b_{j'}^*(o) \geq -(\pi_{j'} - \pi_{j'}^*)$. Hence,

$$(b_{j'}(o^*) - b_{j'}^*(o^*)) - (b_{j'}(o) - b_{j'}^*(o)) \geq -(\pi_{j'} - \pi_{j'}^*) \quad (\text{A.5})$$

$$\geq -b^+(L(j')). \quad (\text{A.6})$$

Consider the case that $\pi_{j'} \leq \pi_{j'}^*$, that j' is requesting less utility than in the egalitarian outcome. Then $b_{j'}(o^*) - b_{j'}^*(o^*) = -(\pi_{j'} - \pi_{j'}^*) \geq 0$, and $b_{j'}(o) - b_{j'}^*(o) \leq -(\pi_{j'} - \pi_{j'}^*)$. Hence,

$$(b_{j'}(o^*) - b_{j'}^*(o^*)) - (b_{j'}(o) - b_{j'}^*(o)) \geq 0. \quad (\text{A.7})$$

Summing over all lower-level bidders via Equations (A.6) and (A.7) gives $(\mathbf{B}_{<i}(o^*) - \mathbf{B}_{<i}^*(o^*)) - (\mathbf{B}_{<i}(o) - \mathbf{B}_{<i}^*(o)) \geq -\sum_{i' < L(j)} b^+(i')$ and hence by Claim 7,

$$(\mathbf{B}_{<i}(o^*) - \mathbf{B}_{<i}^*(o^*)) - (\mathbf{B}_{<i}(o) - \mathbf{B}_{<i}^*(o)) > -b^-(i). \quad (\text{A.8})$$

Now, consider a bidder j' in the same or a higher level than j . If j' is overbidding and not winning in outcome o with bids b , then she would have decreased her utility-target faster than j . 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 j in the same or higher level as j as $\mathbf{B}_{\geq i \setminus j}(o)$. Then, summing over all such bidders gives

$$\mathbf{B}_{\geq i \setminus j}(o^*) - \mathbf{B}_{\geq i \setminus j}^*(o^*) - (\mathbf{B}_{\geq i \setminus j}(o) - \mathbf{B}_{\geq i \setminus j}^*(o)) \geq 0. \quad (\text{A.9})$$

Our original assumption on j gives $(b_j(o^*) - b_j^*(o^*)) - (b_j(o) - b_j^*(o)) \leq b^-(i)$. Now, taking the sum over this and equations (A.8) and (A.9) gives $(\mathbf{B}(o^*) - \mathbf{B}^*(o^*)) - (\mathbf{B}(o) - \mathbf{B}^*(o)) > -b^+(L(j)) + b^+(L(j)) = 0$. By our assumption that o^* is the egalitarian winning outcome, we have $\mathbf{B}^*(o^*) - \mathbf{B}^*(o) \geq 0$. Adding these yields

$$\mathbf{B}(o^*) - \mathbf{B}(o) > 0. \quad (\text{A.10})$$

This is in violation of our assumption that o wins with bids b . Hence, no such bidder j 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 56 *Under assumptions A1, A2, A3 and A4, the utility-target π_j of each bidder j in level L_i will converge to at most the upper bound, $\pi_j^* + \epsilon \cdot b^+(i)$ if for every bidder j' in level L_{i^-} s.t. $i^- \leq i$, $\pi_{j'} \geq s_{j'}^* - \epsilon \cdot b^-(i^-)$.*

Proof:

By Assumptions A1, A2 and Observation 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 8, b must be CEF.

Our proof will proceed by showing that in any such o , $\pi_j < \pi_j^* + \epsilon \cdot b^+(i)$, and hence her utility-target must stay below $\pi_j^* + \epsilon \cdot b^+(i)$ in winning outcomes. Furthermore, by Theorem 4 bids will become CEF; hence i will be forced to decrease her utility-target.

By Claim 8, there is a witness outcome o^w which includes every bidder j' in a strictly higher level i^+ than j . We will now show that if all other players are winning with the egalitarian winning outcome, then j 's utility-target must be below her upper bound, otherwise the witness outcome o^w would win over o^* .

By Definition 33, $\mathbf{B}(o^w) = \mathbf{B}^*(o^*)$. Consider the quantity $\mathbf{B}(o^*) - \mathbf{B}^*(o^*)$, and break it into sums over bidders in levels at or below bidder j , j and bidders in levels above j :

$$\begin{aligned} (\mathbf{B}(o^*) - \mathbf{B}^*(o^*)) &= (\mathbf{B}_{\leq i \setminus j}(o^*) - \mathbf{B}_{\leq i \setminus j}^*(o^*)) \\ &\quad + (b_j(o^*) - b_j^*(o^*)) \\ &\quad + (\mathbf{B}_{> i}(o^*) - \mathbf{B}_{> i}^*(o^*)) \end{aligned}$$

We will now proceed by separately considering bidders in higher and lower levels than bidder j . We will bound the change in bids from each, and see that there is no way for bidder j 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 $b_{j'}(o^w) - b_{j'}^*(o^w) = -(\pi_{j'} - \pi_{j'}^*)$. Since we know that at the egalitarian bids, such a bidder must be winning in the egalitarian outcome, $b_{j'}(o^*) - b_{j'}^*(o^*) = -(\pi_{j'} - \pi_{j'}^*) \leq b_{j'}(o^w) - b_{j'}^*(o^w)$. Summing over all such bidders yields

$$(\mathbf{B}_{> i}(o^*) - \mathbf{B}_{> i}^*(o^*)) - (\mathbf{B}_{> i}(o^w) - \mathbf{B}_{> i}^*(o^w)) \leq 0 \tag{A.11}$$

Lower-level bidders. By our initial assumption that bidding has converged above lower bounds for these bidders, for any bidder j' in $L_{\leq i}$, $\pi_{j'} \geq \pi_{j'}^* - \epsilon b^-(L(j'))$, 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 player wishes to decrease their utility-target in the current bids. If for some bidder j' , $v_{j'}(o^w) \geq v_{j'}(o^*)$, then $(b_{j'}(o^w) - b_{j'}^*(o^w)) = (b_{j'}(o^*) - b_{j'}^*(o^*)) = -(\pi_{j'} - \pi_{j'}^*)$.

Consider then the case that $v_{j'}(o^w) < v_{j'}(o^*)$; 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.

($\pi_{j'} < \pi_{j'}^*$) 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 $b_{j'}(o^w) - b_{j'}^*(o^w) = \max(v_{j'}(o^w), \pi_{j'}) - \pi_{j'} - \max(v_{j'}(o^w), \pi_{j'}^*) + \pi_{j'}^*$, and hence $b_{j'}(o^w) - b_{j'}^*(o^w) = -(\pi_{j'} - \pi_{j'}^*) + (\max(v_{j'}(o^w), \pi_{j'}) - \max(v_{j'}(o^w), \pi_{j'}^*))$. As $b_{j'}(o^*) - b_{j'}^*(o^*) = -(\pi_{j'} - \pi_{j'}^*)$, we then have:

$$0 \leq b_{j'}(o^w) - b_{j'}^*(o^w) \leq b_{j'}(o^*) - b_{j'}^*(o^*) = -(\pi_{j'} - \pi_{j'}^*). \quad (\text{A.12})$$

Furthermore, since $\pi_{j'} \geq \pi_{j'}^* - \epsilon b^-(L(j'))$ by assumption, we have:

$$0 \leq b_{j'}(o^w) - b_{j'}^*(o^w) \leq b_{j'}(o^*) - b_{j'}^*(o^*) \leq \epsilon b^-(L(j')) \quad (\text{A.13})$$

and

$$0 \leq (b_{j'}(o^*) - b_{j'}^*(o^*)) - (b_{j'}(o^w) - b_{j'}^*(o^w)) \leq \epsilon b^-(L(j')). \quad (\text{A.14})$$

($\pi_{j'} \geq \pi_{j'}^*$) 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 $b_{j'}(o^*) - b_{j'}^*(o^*) = \pi_{j'}^* - \pi_{j'}$ and 0. Hence,

$$b_{j'}(o^*) - b_{j'}^*(o^*) \leq b_{j'}(o^w) - b_{j'}^*(o^w) \leq 0 \quad (\text{A.15})$$

and

$$-(\pi_{j'} - \pi_{j'}^*) \leq (b_{j'}(o^*) - b_{j'}^*(o^*)) - (b_{j'}(o^w) - b_{j'}^*(o^w)) \leq 0. \quad (\text{A.16})$$

We now have upper bounds on $-(\pi_{j'} - \pi_{j'}^*) \leq (b_{j'}(o^*) - b_{j'}^*(o^*)) - (b_{j'}(o^w) - b_{j'}^*(o^w))$ for all lower-level bidders. Taking the sum across all members of $L_{\leq i}$ via Equations (A.14) and (A.16) gives:

$$\sum_{j' \in L_{\leq i}} (b_{j'}(o^*) - b_{j'}^*(o^*)) - (b_{j'}(o^w) - b_{j'}^*(o^w)) \leq \sum_{j' \in L_{\leq i}} \epsilon b^-(L(j')) \quad (\text{A.17})$$

Rearranging and noting that $\sum_{j' \in L_{\leq i}} \epsilon b^-(L(j')) < b^+(i)$ by Claim 7 gives

$$(\mathbf{B}_{\leq i \setminus j}(o^*) - \mathbf{B}_{\leq i \setminus j}^*(o^*)) - (\mathbf{B}_{\leq i \setminus j}(o^w) - \mathbf{B}_{\leq i \setminus j}^*(o^w)) < \epsilon b^+(i). \quad (\text{A.18})$$

Summing over equations (A.18) and (A.11) gives us:

$$\begin{aligned}
& (\mathbf{B}_{\leq i \setminus j}(o^*) - \mathbf{B}_{\leq i \setminus j}^*(o^*)) - (\mathbf{B}_{\leq i \setminus j}(o^w) - \mathbf{B}_{\leq i \setminus j}^*(o^w)) \\
& + (\mathbf{B}_{> i}(o^*) - \mathbf{B}_{> i}^*(o^*)) - (\mathbf{B}_{> i}(o^w) - \mathbf{B}_{> i}^*(o^w)) < \epsilon b^+(i). \tag{A.19}
\end{aligned}$$

By assumption, $(b_j(o^*) - b_j^*(o^*)) = -(\pi_j - \pi_j^*) \leq -\epsilon b^+(i)$ and $b_j(o^w) = b_j^*(o^w) = 0$. Thus, $(b_j(o^*) - b_j^*(o^*)) - (b_j(o^w) - b_j^*(o^w)) \leq -\epsilon b^+(i)$. Adding this to (A.19) gives:

$$(\mathbf{B}(o^*) - \mathbf{B}^*(o^*)) - (\mathbf{B}(o^w) - \mathbf{B}^*(o^w)) < \epsilon b^+(i) - \epsilon b^+(i) = 0 \tag{A.20}$$

By our initial assumption that o^w is a witness outcome, $\mathbf{B}^*(o^*) - \mathbf{B}^*(o^w) = 0$. Adding this to the above equation yields

$$\mathbf{B}(o^*) < \mathbf{B}(o^w) \tag{A.21}$$

This contradicts our assumption that o^* is a winning set with bids $b(\cdot)$. Hence, j will be forced to decrease her utility-target to at most $\pi_j^* + \epsilon b^+(i)$ before the egalitarian winning set o^* is winning again. ■

Combining Lemma 55 and Lemma 56 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 B

Optimality in Single-Call Mechanisms

B.1 Optimality Proofs for Generalized BKS

In this section, we generalize our optimality result of Section 4 to arbitrary probability measures and give a complete proof. Theorem 37 shows that truthful payments take the form

$$\lambda_i(A(\hat{b}), \hat{b}, b) = \rho_b^\mu(\hat{b})A_i(\hat{b}) + \lambda_i^0(\hat{b}, b) \text{ a.s.}$$

and thus optimizing the bid-normalized payments means optimizing the following quantity:

$$\sum_j \frac{\lambda_{ij}(b_j(A(\hat{b})), \hat{b}, b)}{b_j(A(\hat{b}))} = \frac{\rho_b^\mu(\hat{b})A_i(\hat{b})}{b_i A_i(\hat{b})} = \frac{\rho_b^\mu(\hat{b})}{b_i}.$$

This means that for worst-case payments we will optimize $\sup_{i, \hat{b}} \left| \frac{\rho_b^\mu(\hat{b})}{b_i} \right|$, and for payment variance we will optimize $\max_i \text{Var}_{\hat{b} \sim \mu_b} \left(\frac{\rho_b^\mu(\hat{b})}{b_i} \right)$. We show that the BKS transformation is optimal for both, subject to an almost everywhere caveat:

Theorem 57 (Optimality of the BKS Transformation) (*Generalization of Theorem 33*)
The BKS reduction $\text{SPTOMECHBKS}(A, \gamma)$ optimizes the payment variance and worst-case normalized payment subject to a lower bound of $\alpha = (1 - \gamma)^n \in (\frac{1}{e}, 1)$ on the precision, the welfare approximation ($n \geq 2$), or the revenue approximation ($n \geq 2$). That is, for any other truthful reduction $(\mu, \{\lambda_i\})$ that achieves a precision, welfare approximation, or revenue approximation of α , the worst-case normalized payments are at least as large almost everywhere over b :

$$\sup_{A, i} \text{Var}_{\hat{b} \sim \mu_b} \left(\sum_j \frac{\lambda_{ij}(b_j(A(\hat{b})), \hat{b}, b)}{b_j(A(\hat{b}))} \right) = \max_i \text{Var}_{\hat{b} \sim \mu_b} \left(\frac{\rho_b^\mu(\hat{b})}{b_i} \right) \geq \max_i \text{Var}_{\hat{b} \sim \mu_b} \left(\frac{\rho_b^{\text{BKS}}(\hat{b})}{b_i} \right) \text{ a.e.}$$

and

$$\sup_{A, i, \hat{b}} \left| \sum_j \frac{\lambda_{ij}(b_j(A(\hat{b})), \hat{b}, b)}{b_j(A(\hat{b}))} \right| = \sup_{i, \hat{b}} \left| \frac{\rho_b^\mu(\hat{b})}{b_i} \right| \geq \sup_{i, \hat{b}} \left| \frac{\rho_b^{\text{BKS}}(\hat{b})}{b_i} \right| \text{ a.e.}$$

Under the nice distribution assumption, this holds for every b .

The theorem is proven in two steps. First, we show in Lemma 58 that a distribution which optimizes precision also optimizes the welfare and revenue approximations. Second, we prove in Theorem 59 that the BKS transform optimizes precision.

Lemma 58 (Generalization of Lemma 34) *For $\alpha > \frac{1}{e}$ and $n \geq 2$, a probability measure that optimizes the variance of normalized payments or the maximum normalized payment subject to a precision constraint of $\Pr(\hat{b} = b|b) \geq \alpha$ also optimizes the maximum normalized payment almost everywhere subject to a welfare or revenue approximation of α .*

Lemma 58 is proven in Section B.1, building on technical lemmas from Section B.1.

Theorem 59 (Precision Optimality of the BKS Transformation) (Generalization of Theorem 35) *The BKS reduction $\text{SPTOMECHBKS}(A, \gamma)$ optimizes the variance of normalized payments and the worst-case normalized payment subject to a lower bound of $\alpha_P = (1 - \gamma)^n \in (\frac{1}{e}, 1)$ on the precision almost everywhere over b . Under the nice distribution assumption, it is optimal for every b .*

Definitions

To prove Theorem 59, we give names to certain probabilities. As in the MIDR setting, we use a set $M \subseteq [n]$ to denote the set of bidders with $\hat{b}_i = b_i$. Bidders $i \notin M$ have their bids lowered, that is $\hat{b}_i < b_i$. We define the probability $\pi^\mu(M, b)$ to be the probability that such an event occurs, that is, $\pi^\mu(M, b)$ is the probability when b is bid that $\hat{b}_i = b_i$ if $i \in M$, and $\hat{b}_i < b_i$ if $i \notin M$:

$$\pi^\mu(M, b) \equiv \Pr \left((\hat{b}_i = b_i \text{ for } i \in M) \text{ and } (\hat{b}_i < b_i \text{ for } i \notin M) \middle| b \right) .$$

Note that for the BKS transformation, $\pi^\mu(M, b) = (1 - \gamma)^{|M|} \gamma^{n-|M|}$ so $\frac{\pi^\mu(M \cup \{i\}, b)}{\pi^\mu(M, b)} = \frac{1-\gamma}{\gamma}$.

The second probability quantifies the behavior of μ_b near b as follows. Fix a bid \bar{b} and assume player i actually bids $b_i - \delta$. Does the distribution $\mu_{b_i - \delta, b_{-i}}$ cause the reduction to select $\hat{b} = b$ with positive probability in spite of the fact that i said $b_i - \delta$? In particular, we care about the average behavior for $\delta \in [0, b_i]$, which we represent by $z^\mu(M, i, \bar{b})$. Formally, we define

$$\zeta^\mu(M, i, b, z) \equiv \Pr \left(\hat{b}_i = z \text{ and } (\hat{b}_j = b_j \text{ for } j \in M \setminus \{i\}) \text{ and } (\hat{b}_j < b_j \text{ for } j \notin M \cup \{i\}) \middle| b \right)$$

and

$$z^\mu(M, i, b) \equiv \frac{1}{b_i} \int_0^{b_i} \zeta^\mu(M, i, (u, b_{-i}), b_i) .$$

Of particular importance, we will show $z^\mu(M, i, b) = 0$ almost everywhere in general and everywhere under the nice distribution assumption.

Precision Optimality of the BKS Transformation

The optimality proof for the BKS transformation

The first result follows as a corollary of Lemma 64:

Corollary 60 (of Lemma 64) *If a resampling distribution μ satisfies the monotonicity condition, then for all M , $i \notin M$:*

$$\sup_{\hat{b}} \left| \frac{\rho_b^\mu(\hat{b})}{b_i} \right| \geq \frac{\pi^\mu(M \cup \{i\}, b) - z^\mu(M, i, b)}{\pi^\mu(M, b)}$$

and

$$\begin{aligned} & \int_{\hat{b}_i \leq b_i \wedge (j \in M \Rightarrow \hat{b}_j = b_j) \wedge (j \notin M \cup \{i\} \Rightarrow \hat{b}_j < b_j)} \left(\frac{\rho_b^\mu(\hat{b})}{b_i} \right)^2 d\mu_b \\ & \geq (\pi^\mu(M, b) + \pi^\mu(M \cup \{i\}, b)) \frac{\pi^\mu(M \cup \{i\}, b)}{\pi^\mu(M, b)} \left(1 - \frac{z^\mu(M, i, b)}{\pi^\mu(M \cup \{i\}, b)} \right)^2. \end{aligned}$$

Proof: Apply Lemma 64 where B_{-i} is the set of \hat{b}_{-i} where $\hat{b}_j = b_j$ if $j \in M$ and $\hat{b}_j < b_j$ for $j \notin M$. \blacksquare

If we ignore the $z^\mu(M, i, b)$ terms, this looks precisely like the normalized payments from the MIDR setting. Fortunately, $z^\mu(M, i, b)$ is almost always zero:

Corollary 61 (of Lemma 67) *For any resampling distribution μ and a fixed M and i ,*

$$z^\mu(M, i, b) = 0 \text{ a.e.}$$

(i.e. for all but a set of b with zero measure).

Under the nice distribution assumption, $z^\mu(M, i, b) = 0$ for all b .

Proof: Note that $\zeta^\mu(M, i, (u, b_{-i}), b_i) \leq \Pr_\mu(\hat{b}_i = b_i | u, b_{-i})$, so by Lemma 67

$$z^\mu(M, i, b) = \frac{1}{b_i} \int_0^{b_i} \zeta^\mu(M, i, (u, b_{-i}), b_i) \leq \int_0^{b_i} \Pr_\mu(\hat{b}_i = b_i | u, b_{-i}) = 0 \text{ a.e.}$$

\blacksquare

Thus, Corollaries 60 and 61 together imply the following bound:

Lemma 62 *If a resampling distribution μ with precision $\alpha \geq (1 - \gamma)^n$ satisfies the monotonicity condition, then*

$$\sup_{i, \hat{b}} \left| \frac{\rho_b^\mu(\hat{b})}{b_i} \right| \geq \frac{1 - \gamma}{\gamma} \text{ a.e.}$$

that is, for all b but a set with measure zero. This holds everywhere if $z^\mu(M, i, b) = 0$ everywhere.

Proof: We first prove the bound on the worst-case normalized payment. By assumption on the precision of μ , we have $\pi^\mu([n], b) \geq (1 - \gamma)^n$ for some γ and all b . By Corollary 60, we know that

$$\sup_{\hat{b}} \left| \frac{\rho_{\hat{b}}^\mu(\hat{b})}{b_i} \right| \geq \frac{\pi^\mu(M \cup \{i\}, b) - z^\mu(M, i, b)}{\pi^\mu(M, b)} .$$

Applying Lemma 65 with $\eta(S) = \pi^\mu(S, b)$, $\alpha = (1 - \gamma)^n$, and $\beta = 1$ we get that

$$\max_{M, i \notin M} \frac{\pi^\mu(M \cup \{i\}, b)}{\pi^\mu(M, b)} \geq \frac{1 - \phi}{\phi}$$

where

$$\phi = 1 - \left(\frac{(1 - \gamma)^n}{1} \right)^{\frac{1}{n}} = \gamma .$$

Thus,

$$\max_{M, i \notin M} \frac{\pi^\mu(M \cup \{i\}, b)}{\pi^\mu(M, b)} \geq \frac{1 - \gamma}{\gamma} .$$

Aggregating Corollary 60 over all M and $i \notin M$, we have

$$\begin{aligned} \sup_{i, \hat{b}} \left| \frac{\rho_{\hat{b}}^\mu(\hat{b})}{b_i} \right| &\geq \max_{M, i \notin M} \frac{\pi^\mu(M \cup \{i\}, b)}{\pi^\mu(M, b)} - \max_{M, i \notin M} \frac{z^\mu(M, i, b)}{\pi^\mu(M, b)} \\ &\geq \frac{1 - \gamma}{\gamma} - \max_{M, i \notin M} \frac{z^\mu(M, i, b)}{\pi^\mu(M, b)} . \end{aligned}$$

If we assume $z^\mu(M, i, b) = 0$ everywhere (e.g. by the nice distribution assumption), then we get

$$\sup_{i, \hat{b}} \left| \frac{\rho_{\hat{b}}^\mu(\hat{b})}{b_i} \right| \geq \frac{1 - \gamma}{\gamma} .$$

Otherwise, Corollary 61 says that $z^\mu(M, i, b) = 0$ almost everywhere, giving the more general bound. \blacksquare

Lemma 63 *If a resampling distribution μ with precision $\alpha \geq (1 - \gamma)^n \geq \frac{1}{e}$ satisfies the monotonicity condition, then*

$$\max_i \text{Var}_{\hat{b}} \left(\frac{\rho_{\hat{b}}^\mu(\hat{b})}{b_i} \right) \geq \frac{1 - \gamma}{\gamma} \text{ a.e.}$$

that is, for all b but a set with measure zero. This holds everywhere if $z^\mu(M, i, b) = 0$ everywhere.

Proof: The proof for variance is similar to Lemma 62, but we apply Lemma 66 instead of Lemma 65. First, note that since μ_b is a probability measure, $\mu_b(\mathbb{R}^n) = 1$ and thus

$$\int_{\hat{b} \in \mathbb{R}^n} \frac{\rho_b^\mu(\hat{b})}{b_i} d\mu_b = \frac{1}{b_i} \nu_{b,i}(\mathbb{R}^n) = \mu_b(\mathbb{R}^n) - \frac{1}{b_i} \int_0^{b_i} \mu_{u,b-i}(\mathbb{R}^n) du = 1 - \frac{1}{b_i} \int_0^{b_i} 1 du = 0 .$$

We begin with the variance for player i , applying Corollaries 60 and 61:

$$\begin{aligned} \text{Var}_{\hat{b}} \left(\frac{\rho_b^\mu(\hat{b})}{b_i} \right) &= \int_{\hat{b} \in \mathbb{R}^n} \left(\frac{\rho_b^\mu(\hat{b})}{b_i} \right)^2 d\mu_b - \left(\int_{\hat{b} \in \mathbb{R}^n} \frac{\rho_b^\mu(\hat{b})}{b_i} d\mu_b \right)^2 \\ &= \int_{\hat{b} \in \mathbb{R}^n} \left(\frac{\rho_b^\mu(\hat{b})}{b_i} \right)^2 d\mu_b \\ &\geq \sum_{M|i \notin M} \int_{\hat{b}_i \leq b_i \wedge (j \in M \Rightarrow \hat{b}_j = b_j) \wedge (j \notin M \cup \{i\} \Rightarrow \hat{b}_j < b_j)} \left(\frac{\rho_b^\mu(\hat{b})}{b_i} \right)^2 d\mu_b \\ &\geq \sum_{M|i \notin M} (\pi^\mu(M, b) + \pi^\mu(M \cup \{i\}, b)) \frac{\pi^\mu(M \cup \{i\}, b)}{\pi^\mu(M, b)} \text{ a.e.} \end{aligned}$$

Applying Lemma 66 with $\eta(S) = \pi^\mu(S, b)$, $\alpha = (1 - \gamma)^n$ and $\beta = \Pr(\hat{b} \leq b | b)$ immediately implies

$$\max_i \text{Var}_{\hat{b}} \left(\frac{\rho_b^\mu(\hat{b})}{b_i} \right) \geq \Pr(\hat{b} \leq b | b) \frac{1 - \phi}{\phi} \text{ a.e.}$$

where

$$\phi = 1 - \left(\frac{(1 - \gamma)^n}{\Pr(\hat{b} \leq b | b)} \right)^{\frac{1}{n}} .$$

One can check that when $\frac{(1-\gamma)^n}{\Pr(\hat{b} \leq b | b)} \geq \frac{1}{e}$, the quantity $\Pr(\hat{b} \leq b | b) \frac{1-\phi}{\phi}$ is decreasing in $\Pr(\hat{b} \leq b | b)$. Taking the worst case $\Pr(\hat{b} \leq b | b) = 1$ implies the desired result:

$$\max_i \text{Var}_{\hat{b}} \left(\frac{\rho_b^\mu(\hat{b})}{b_i} \right) \geq \max_i \geq \frac{1 - \gamma}{\gamma} \text{ a.e.}$$

■

Theorem 59 – optimality of the BKS transformation with respect to a precision bound – follows from the two previous lemmas:

Proof:[of Theorem 59] For worst-case payments, we show that for any measure μ , with precision at least 2^{-n} ,

$$\sup_{i, \hat{b}} \left| \frac{\rho_b^{BKS}(\hat{b})}{b_i} \right| \leq \sup_{i, \hat{b}} \left| \frac{\rho_b^\mu(\hat{b})}{b_i} \right| \quad a.e.$$

For $\Pr(\hat{b} = b|b) = (1 - \gamma)^n$, the BKS transform achieves $\sup_{i, \hat{b}} \left| \frac{\rho_b^{BKS}(\hat{b})}{b_i} \right| = \max\left(1, \frac{1-\gamma}{\gamma}\right)$ for all b . Provided $\gamma > \frac{1}{2}$, the dominant term is $\frac{1-\gamma}{\gamma}$ and Lemma 62 shows that this is a lower bound for any such μ almost everywhere. When $\alpha > 2^{-n}$ we get $\gamma > \frac{1}{2}$, and thus *BKS* is optimal.

Moreover, under the nice distribution assumption (implying $z^\mu(M, i, b) = 0$), Lemma 62 says that this holds everywhere.

For the variance of normalized payments, we need to show that for any measure μ with precision at least $\frac{1}{e}$:

$$\text{Var}_{\hat{b} \sim \mu_b} \left(\frac{\rho_b^{BKS}(\hat{b})}{b_i} \right) \leq \text{Var}_{\hat{b} \sim \mu_b} \left(\frac{\rho_b^\mu(\hat{b})}{b_i} \right) \quad a.e.$$

Again, for $\Pr(\hat{b} = b|b) = (1 - \gamma)^n$, the BKS transform achieves $\text{Var}_{\hat{b} \sim \mu_b} \left| \frac{\rho_b^{BKS}(\hat{b})}{b_i} \right| = \frac{1-\gamma}{\gamma}$ for all b . Lemma 63 shows that this is a lower bound for any such μ almost everywhere. ■

Technical Lemmas

The next lemma gives our main lower bound on the worst coefficient:

Lemma 64 *If a measure μ satisfies the monotonicity condition, then for any player i , bid b , and set of bids $B_{-i} \subseteq \mathbb{R}_+^{n-1}$:*

$$\sup_{\hat{b}} \left| \frac{\rho_b^\mu(\hat{b})}{b_i} \right| \geq \frac{\Pr\left(\hat{b}_i = b_i \wedge \hat{b}_{-i} \in B_{-i} \mid b\right) - \frac{1}{b_i} \int_0^{b_i} \Pr\left(\hat{b}_i = u \wedge \hat{b}_{-i} \in B_{-i} \mid u, b_{-i}\right) du}{\Pr\left(\hat{b}_i < b_i \wedge \hat{b}_{-i} \in B_{-i} \mid b\right)},$$

and

$$\begin{aligned} \int_{\hat{b}_i \leq b_i \wedge \hat{b}_{-i} \in B_{-i}} \left(\frac{\rho_b^\mu(\hat{b})}{b_i} \right)^2 d\mu_b &\geq \Pr\left(\hat{b}_i \leq b_i \wedge \hat{b}_{-i} \in B_{-i} \mid b\right) \frac{\Pr\left(\hat{b}_i = b_i \wedge \hat{b}_{-i} \in B_{-i} \mid b\right)}{\Pr\left(\hat{b}_i < b_i \wedge \hat{b}_{-i} \in B_{-i} \mid b\right)} \\ &\quad \times \left(1 - \frac{\frac{1}{b_i} \int_0^{b_i} \Pr\left(\hat{b}_i = u \wedge \hat{b}_{-i} \in B_{-i} \mid u, b_{-i}\right) du}{\Pr\left(\hat{b}_i = b_i \wedge \hat{b}_{-i} \in B_{-i} \mid b\right)} \right)^2, \end{aligned}$$

where the integral terms are zero almost everywhere in b by Lemma 67.

Proof: Define the sets

$$B^{(=)} = \{b_i\} \times B_{-i} \quad \text{and} \quad B^{(<)} = [0, b_i) \times B_{-i} ,$$

i.e. the set $B^{(=)}$ contains bids \hat{b} where $\hat{b}_i = b_i$ and $\hat{b}_{-i} \in B_{-i}$, and the set $B^{(<)}$ contains bids \hat{b} where $\hat{b}_i < b_i$ and $\hat{b}_{-i} \in B_{-i}$. The main work of the lemma is to bound the following term:

$$\begin{aligned} \int_{\hat{b} \in B^{(<)}} \frac{\rho_b^\mu(\hat{b})}{b_i} d\mu_b &= \frac{\nu_{b,i}(B^{(<)})}{b_i} \\ &= \frac{b_i \mu_b(B^{(<)}) - \int_0^{b_i} \mu_{u,b_{-i}}(B^{(<)}) du}{b_i} \\ &= \mu_b(B^{(<)}) - \frac{1}{b_i} \int_0^{b_i} \mu_{u,b_{-i}}(B^{(<)}) du \\ &= \Pr\left(\hat{b} \in B^{(<)} \mid b\right) - \frac{1}{b_i} \int_0^{b_i} \Pr\left(\hat{b} \in B^{(<)} \mid u, b_{-i}\right) du . \end{aligned}$$

By monotonicity, $\Pr\left(\hat{b} \in B^{(<) \cup B^{(=)}} \mid u, b_{-i}\right)$ is weakly decreasing in u (Lemma 39). This implies

$$\Pr\left(\hat{b} \in B^{(=)} \mid b\right) + \Pr\left(\hat{b} \in B^{(<)} \mid b\right) \leq \frac{1}{b_i} \int_0^{b_i} \left(\Pr\left(\hat{b} \in B^{(=)} \mid u, b_{-i}\right) + \Pr\left(\hat{b} \in B^{(<)} \mid u, b_{-i}\right)\right) du$$

and thus

$$\begin{aligned} \int_{\hat{b} \in B^{(<)}} \frac{\rho_b^\mu(\hat{b})}{b_i} d\mu_b &= \Pr\left(\hat{b} \in B^{(<)} \mid b\right) - \frac{1}{b_i} \int_0^{b_i} \Pr\left(\hat{b} \in B^{(<)} \mid u, b_{-i}\right) du \\ &\leq - \left(\Pr\left(\hat{b} \in B^{(=)} \mid b\right) - \frac{1}{b_i} \int_0^{b_i} \Pr\left(\hat{b} \in B^{(=)} \mid u, b_{-i}\right) du \right) . \end{aligned}$$

To bound $\sup_{\hat{b} \in B^{(<)}} \left| \frac{\rho_b^\mu(\hat{b})}{b_i} \right|$, we have

$$\sup_{\hat{b} \in B^{(<)}} \left| \frac{\rho_b^\mu(\hat{b})}{b_i} \right| \geq \left| \frac{\int_{\hat{b} \in B^{(<)}} \frac{\rho_b^\mu(\hat{b})}{b_i} d\mu_b}{\mu_b(B^{(<)})} \right| \geq \frac{\Pr\left(\hat{b} \in B^{(=)} \mid b\right) - \frac{1}{b_i} \int_0^{b_i} \Pr\left(\hat{b} \in B^{(=)} \mid u, b_{-i}\right) du}{\Pr\left(\hat{b} \in B^{(<)} \mid b\right)}$$

Lemma 67 implies that the limit term is zero almost everywhere in b .

For our partial bound on the second moment, we write

$$\begin{aligned}
\int_{\hat{b} \in B^{(<)} \cup B^{(=)}} \left(\frac{\rho_b^\mu(\hat{b})}{b_i} \right)^2 d\mu_b &\geq \int_{\hat{b} \in B^{(=)}} \left(\frac{\rho_b^\mu(\hat{b})}{b_i} \right)^2 d\mu_b + \int_{\hat{b} \in B^{(<)}} \left(\frac{\rho_b^\mu(\hat{b})}{b_i} \right)^2 d\mu_b \\
&\geq \mu_b(B^{(=)}) \left(\frac{\int_{\hat{b} \in B^{(=)}} \frac{\rho_b^\mu(\hat{b})}{b_i} d\mu_b}{\mu_b(B^{(=)})} \right)^2 + \mu_b(B^{(<)}) \left(\frac{\int_{\hat{b} \in B^{(<)}} \frac{\rho_b^\mu(\hat{b})}{b_i} d\mu_b}{\mu_b(B^{(<)})} \right)^2 \\
&\geq \mu_b(B^{(=)}) \left(\frac{\Pr(\hat{b} \in B^{(=)} | b) - \frac{1}{b_i} \int_0^{b_i} \Pr(\hat{b} \in B^{(=)} | u, b_{-i}) du}{\mu_b(B^{(=)})} \right)^2 \\
&\quad + \mu_b(B^{(<)}) \left(\frac{\Pr(\hat{b} \in B^{(=)} | b) - \frac{1}{b_i} \int_0^{b_i} \Pr(\hat{b} \in B^{(=)} | u, b_{-i}) du}{\mu_b(B^{(<)})} \right)^2 \\
&\geq (\mu_b(B^{(<)}) + \mu_b(B^{(=)})) \frac{\mu_b(B^{(=)})}{\mu_b(B^{(<)})} \\
&\quad \times \left(1 - \frac{\frac{1}{b_i} \int_0^{b_i} \Pr(\hat{b} \in B^{(=)} | u, b_{-i}) du}{\Pr(\hat{b} \in B^{(=)} | b)} \right)^2
\end{aligned}$$

Which is the desired bound. ■

Lemma 65 *Let $\eta : \{0, 1\}^n$ be a function over subsets $S \subseteq [n]$ with $\eta([n]) \geq \alpha \in [0, 1]$ and $\sum_{S \subseteq [n]} \eta(S) \leq \beta \in [0, 1]$. Then*

$$\max_{S, i \in [n] \setminus S} \frac{\eta(S \cup \{i\})}{\eta(S)} \geq \frac{1 - \phi}{\phi}$$

where $\phi = 1 - \left(\frac{\alpha}{\beta}\right)^{\frac{1}{n}}$.

Proof: By contradiction. Assume that for every S and $i \notin S$,

$$\frac{\eta(S \cup \{i\})}{\eta(S)} < \frac{1 - \phi}{\phi}$$

where $\phi = 1 - \left(\frac{\alpha}{\beta}\right)^{\frac{1}{n}}$.

Then by multiplying $\frac{\eta(S)}{\eta(S \cup \{i\})}$ terms together we get

$$\eta(S) \geq \eta([n]) \left(\frac{\phi}{1 - \phi} \right)^{n - |S|}.$$

Summing over all $S \subseteq [n]$, substituting for α and β , and algebra gives

$$\begin{aligned}
 \sum_{S \subseteq [n]} \eta(S) &> \eta([n]) \sum_{S \subseteq [n]} \left(\frac{\phi}{1-\phi} \right)^{n-|S|} \\
 \beta &> \alpha \sum_{S \subseteq [n]} \left(\frac{\phi}{1-\phi} \right)^{n-|S|} \\
 \beta(1-\phi)^n &> \alpha \sum_{S \subseteq [n]} (1-\phi)^{|S|} \phi^{n-|S|} \\
 \beta(1-\phi)^n &> \alpha \\
 \beta \left(1 - \left(1 - \left(\frac{\alpha}{\beta} \right)^{\frac{1}{n}} \right) \right)^n &> \alpha \\
 \alpha &> \alpha .
 \end{aligned}$$

Which is a contradiction. ■

Lemma 66 *Let $\eta : \{0, 1\}^n$ be a function over subsets $S \subseteq [n]$ with $\eta([n]) \geq \alpha \in [0, 1]$ and $\sum_{S \subseteq [n]} \eta(S) = \beta \in [0, 1]$. Then*

$$\max_i \sum_{S | i \notin S} (\eta(S) + \eta(S \cup \{i\})) \frac{\eta(S \cup \{i\})}{\eta(S)} \geq \beta \frac{1-\phi}{\phi}$$

where $\phi = 1 - \left(\frac{\alpha}{\beta} \right)^{\frac{1}{n}}$.

Proof: We lower-bound the sum. Fix i and differentiate the sum:

$$\frac{\partial}{\partial \eta(S)} \left(\sum_{T | i \notin T} (\eta(T) + \eta(T \cup \{i\})) \frac{\eta(T \cup \{i\})}{\eta(T)} \right) = \begin{cases} 2 \frac{\eta(S)}{\eta(S \setminus \{i\})} + 1, & i \in S \\ - \left(\frac{\eta(S \cup \{i\})}{\eta(S)} \right)^2, & i \notin S \end{cases} .$$

The conditions of the lemma bound $\sum_S \eta(S)$ and $\eta([n])$, otherwise the values of η are only constrained to be in $[0, 1]$. The derivative tells us that in an optimal assignment, for all sets S that do not contain i , the ratio $\frac{\eta(S \cup \{i\})}{\eta(S)}$ is constant. Construct such an optimal assignment and define ϕ_i as satisfying

$$\frac{\eta(S \cup \{i\})}{\eta(S)} = \frac{1-\phi_i}{\phi_i}$$

for all S that do not contain i . Note that this implies

$$\sum_{S | i \notin S} (\eta(S) + \eta(S \cup \{i\})) \frac{\eta(S \cup \{i\})}{\eta(S)} \geq \beta \frac{1-\phi_i}{\phi_i}$$

For any set S it follows that

$$\begin{aligned} \eta(S) &= \eta([n]) \prod_{i \notin S} \frac{\phi_i}{1 - \phi_i} \\ \sum_{S \subseteq [n]} \eta(S) &= \eta([n]) \sum_{S \subseteq [n]} \prod_{i \notin S} \frac{\phi_i}{1 - \phi_i} \\ \beta \prod_{i \in [n]} (1 - \phi_i) &\geq \alpha \sum_{S \subseteq [n]} \prod_{i \in S} (1 - \phi_i) \prod_{i \notin S} \phi_i \\ \prod_{i \in [n]} (1 - \phi_i) &\geq \frac{\alpha}{\beta} . \end{aligned}$$

This implies there is some i such that $\phi_i \leq 1 - \left(\frac{\alpha}{\beta}\right)^{\frac{1}{n}}$, which implies the lemma. \blacksquare

The next lemma is our main analysis lemma. We will ultimately use it to claim that our lower bound must hold almost everywhere for any μ :

Lemma 67 *For any resampling distribution μ that satisfies the monotonicity condition, any bid b , and any bidder i ,*

$$\int_0^{b_i} \Pr_{\mu}(\hat{b}_i = b_i | u, b_{-i}) = 0 \text{ a.e.}$$

(i.e. for all but a set of b with zero measure).

Proof: Define the marginalized measure μ_b^i for a set of bids $B \subseteq \mathbb{R}$ as

$$\mu_b^i(B) \equiv \mu_b(\{b \in \mathbb{R}^n | b_i \in B\}) .$$

Note that

$$\mu_{u, b_{-i}}^i(\{b_i\}) = \Pr_{\mu}(\hat{b}_i = b_i | u, b_{-i})$$

and therefore our task is to show that

$$\lim_{u \rightarrow -b_i} \mu_{u, b_{-i}}^i(\{b_i\}) = 0 \text{ a.e.}$$

Next we show that for any b we can prove the desired limit is zero by proving that a related integral is zero. Assume that for some b we have

$$\lim_{u \rightarrow -b_i} \mu_{u, b_{-i}}^i(\{b_i\}) > 0 .$$

Then there exists a δ_b such that

$$\forall u \in (b_i - \delta_b, b_i) : \mu_{u, b_{-i}}^i(\{b_i\}) > 0 .$$

Since $\mu_{u,b-i}^i(\{b_i\})$ is nonnegative, this implies

$$\int_{u \in \mathbb{R}} \mu_{u,b-i}^i(\{b_i\}) du \geq \int_{u \in (b_i - \delta_b, b_i)} \mu_{u,b-i}^i(\{b_i\}) du > 0 .$$

Taking the contrapositive, it follows that if the integral is zero at a bid b then the limit is also zero:

$$\int_{u \in \mathbb{R}} \mu_{u,b-i}^i(\{b_i\}) du = 0 \Rightarrow \lim_{u \rightarrow^- b_i} \mu_{u,b-i}^i(\{b_i\}) = 0 . \quad (\text{B.1})$$

Henceforth, we will prove that $\int_{u \in \mathbb{R}} \mu_{u,b-i}^i(\{b_i\}) du = 0$ almost everywhere.

We start with the integral

$$\int_{b \in \mathbb{R}} \int_{u \in \mathbb{R}} \mu_{u,b-i}^i(\{b_i\}) du db .$$

Manipulating the integral and noting that $\int_{u \in \mathbb{R}} 1_{\{u\}}(\hat{b}_i) du = 0$, we get

$$\begin{aligned} \int_{b \in \mathbb{R}} \int_{u \in \mathbb{R}} \mu_{u,b-i}^i(\{b_i\}) du db &= \int_{b \in \mathbb{R}} \int_{u \in \mathbb{R}} \mu_{b_i,b-i}^i(\{u\}) du db \\ &= \int_{b \in \mathbb{R}^n} \int_{u \in \mathbb{R}} \int_{\hat{b}_i \in \mathbb{R}} 1_{\{u\}}(\hat{b}_i) d\mu_b^i du db \\ &= \int_{b \in \mathbb{R}^n} \int_{\hat{b}_i \in \mathbb{R}} \int_{u \in \mathbb{R}} 1_{\{u\}}(\hat{b}_i) du d\mu_b^i db \\ &= \int_{b \in \mathbb{R}^n} \int_{\hat{b}_i \in \mathbb{R}} 0 d\mu_b^i db \\ &= 0 \end{aligned}$$

(where integral rearrangements may be justified by Tonelli's Theorem). By Fact 81, this implies

$$\int_{u \in \mathbb{R}} \mu_{u,b-i}^i(\{b_i\}) du = 0 \text{ almost everywhere over } b,$$

which implies the desired result. ■

Welfare and Revenue Optimality

Under mild assumptions, one can show that optimizing precision is equivalent to optimizing the social welfare approximation or the revenue approximation. We include only the worst-case optimality proofs; the variance proof is similar, applying ideas from Lemma 63.

The optimality proof is divided into two steps:

1. *Lemmas 68 and 69:* Show that the welfare/revenue approximation of a resampling distribution μ is essentially

$$\inf_b \min_{i \in [n]} \Pr \left(\hat{b}_i \geq b_i \text{ and } \hat{b}_{-i} = b_{-i} \mid b \right) .$$

The welfare and revenue lemmas use different techniques to give a lower bound on the approximation; however, they use the same “bad” allocation function.

2. *Lemma 70 and finally Lemma 58:* Show that a distribution that optimizes the worst-case normalized payment with respect to

$$\min_{i \in [n]} \Pr \left(\hat{b}_i \geq b_i \text{ and } \hat{b}_{-i} = b_{-i} \mid b \right) \geq \alpha$$

must take $\Pr(\hat{b} \not\leq b \mid b) = 0$ and, therefore

$$\min_{i \in [n]} \Pr \left(\hat{b}_i \geq b_i \text{ and } \hat{b}_{-i} = b_{-i} \mid b \right) = \Pr \left(\hat{b} = b \mid b \right)$$

implying that it is sufficient to optimize with respect to $\Pr(\hat{b} = b \mid b) \geq (1 - \gamma)^n = \alpha$.

The following lemmas characterize the welfare and revenue approximations of the reduction generated by a resampling distribution μ :

Lemma 68 *The welfare approximation of a resampling distribution μ for a bid b is*

$$\alpha = \min_{i \in [n]} \Pr \left(\hat{b}_i \geq b_i \text{ and } \hat{b}_{-i} = b_{-i} \mid b \right) .$$

Proof: For a bid b , define the set $B^i \subset \mathbb{R}_+^n$ as

$$B^i = \{ \hat{b} \mid \hat{b}_i \geq b_i \text{ and } \hat{b}_{-i} = b_{-i} \} .$$

Monotonicity of A requires that for all $u \geq b_i$,

$$A_i(u, b_{-i}) \geq A_i(b) .$$

Thus, the allocation received by player i under \mathcal{A} is at least

$$\Pr \left(\hat{b}_i \geq b_i \text{ and } \hat{b}_{-i} = b_{-i} \mid b \right) A_i(b) = \Pr \left(\hat{b} \in B^i \mid b \right) A_i(b)$$

and thus the social welfare is at least

$$\begin{aligned} \sum_{i \in [n]} b_i \mathcal{A}_i(b) &\geq \sum_{i \in [n]} b_i \Pr \left(\hat{b} \in B^i \mid b \right) \mathcal{A}_i(b) \\ &\geq \min_{i \in [n]} \left(\Pr \left(\hat{b} \in B^i \mid b \right) \right) \sum_{i \in [n]} b_i \mathcal{A}_i(b) . \end{aligned}$$

This lower bound is tight in the following allocation rule

$$A_i(\hat{b}) = \begin{cases} 1 & i = j \text{ and } \hat{b} \in B^i \\ 0 & \text{otherwise} \end{cases}$$

when $j = \operatorname{argmin}_{i \in [n]} b_i \Pr(\hat{b} \in B^i | b)$. ■

Lemma 69 *The revenue approximation α_R of a reduction given by a resampling distribution μ is bounded from below by the precision*

$$\alpha_P = \inf_b \Pr(\hat{b} = b | b) \leq \alpha_R$$

and above by

$$\alpha_R \leq \inf_b \min_{i \in [n]} \Pr(\hat{b}_i \geq b_i \wedge \hat{b}_{-i} = b_{-i} | b) .$$

Proof: To see that the precision $\alpha_P = \inf_b \Pr(\hat{b} = b | b)$ is a lower bound on the revenue approximation, consider decomposing the mechanism produced by the reduction as follows: with probability α_P , the mechanism uses the original allocation function, and with probability $1 - \alpha_P$ it chooses an allocation function A^{rs} that resamples bids more frequently. Since prices are linear, the final expected price will be the weighted sum of the truthful prices for A and the truthful prices for A^{rs} .

For positive types, revenue from both A and A^{rs} will be nonnegative, and the revenue of the resulting mechanism will be the weighted sum of the revenues from A and A^{rs} . Thus, since A is chosen with probability α_P , the revenue of their combination will be at least α_P times the revenue from A .

Next we use the allocation function from Lemma 68 to give an upper bound. For clarity, we assume that the infimum in the bound of α is attained by some b . (The proof when the infimum is not attained is messier but fundamentally the same.) Let b be a bid such that

$$\min_{i \in [n]} \Pr(\hat{b}_i \geq b_i \wedge \hat{b}_{-i} = b_{-i} | b) = \alpha .$$

Again, let $B^i \subset \mathbb{R}_+^n$ be the set

$$B^i = \{\hat{b} | \hat{b}_i \geq b_i \text{ and } \hat{b}_{-i} = b_{-i}\} ,$$

and consider following allocation function, where $j = \operatorname{argmin}_{i \in [n]} b_i \Pr(\hat{b} \in B^i | b)$:

$$A_i(\hat{b}) = \begin{cases} 1 & i = j \text{ and } \hat{b} \in B^i \\ 0 & \text{otherwise.} \end{cases}$$

When this allocation function is implemented directly with the Archer-Tardos pricing rule, the revenue when bidders say b will be

$$\sum_{i \in [n]} b_i A_i(b) - \int_{-\infty}^{b_i} A_i(u, b_{-i}) du = b_j .$$

Now, for any single call reduction, the expected revenue will be

$$\begin{aligned} \sum_{i \in [n]} b_i \mathbf{E}[A_i^{sc}(b)] - \int_{-\infty}^{b_i} \mathbf{E}[A_i^{sc}(u, b_{-i})] du &\leq b_j \mathbf{E}[A_j^{sc}(b)] \\ &= b_j \Pr(\hat{b} \in B^j | b) . \end{aligned}$$

Thus, the revenue approximation when players bid b is at most $\Pr(\hat{b} \in B^j | b)$. \blacksquare

Lemma 70 *The worst-case bid-normalized payment for a resampling distribution μ is at least*

$$\sup_{\hat{b}} \left| \frac{\rho_b^\mu(\hat{b})}{b_i} \right| \geq \max \left(\frac{1 - \gamma^{(=)}}{\gamma^{(=)}}, \frac{1 - \gamma^{(>)}}{\gamma^{(>)}} \right) \text{ a.e.}$$

where

$$\gamma^{(=)} = 1 - \left(\frac{\Pr(\hat{b} = b | b)}{\Pr(\hat{b} \leq b | b)} \right)^{\frac{1}{n}}$$

and

$$\gamma^{(>)} = 1 - \left(\frac{\min_{i \in [n]} \Pr(\hat{b}_i > b_i \wedge \hat{b}_{-i} = b_{-i} | b)}{\frac{1}{n} \Pr(\hat{b} \not\leq b | b)} \right)^{\frac{1}{n-1}} .$$

The bound holds everywhere under the nice distribution assumption.

Proof: For the sake of clarity, we assume the nice distribution assumption. The general case follows naturally by carrying extra terms through the analysis.

Corollary 60 says that for any $M \subset [n]$ and $i \notin M$,

$$\sup_{\hat{b}} \left| \frac{\rho_b^\mu(\hat{b})}{b_i} \right| \geq \frac{\pi^\mu(M \cup \{i\}, b)}{\pi^\mu(M, b)} .$$

Since $\sum_{M \subseteq [n]} \pi(M, \bar{b}) = \Pr(\hat{b} \leq \bar{b} | \bar{b})$, applying Lemma 65 with $\eta(S) = \pi^\mu(M, b)$ implies that

$$\max_{M \subseteq [n]} \frac{\pi^\mu(M \cup \{i\}, b)}{\pi^\mu(M, b)} \geq \frac{1 - \gamma^{(=)}}{\gamma^{(=)}}$$

where $\gamma^{(=)}$ is

$$\gamma^{(=)} = 1 - \left(\frac{\Pr(\hat{b} = b|b)}{\Pr(\hat{b} \leq \bar{b}|\bar{b})} \right)^{\frac{1}{n}}.$$

Thus,

$$\sup_{\hat{b}} \left| \frac{\rho_b^\mu(\hat{b})}{b_i} \right| \geq \frac{1 - \gamma^{(=)}}{\gamma^{(=)}}.$$

Next, define $\nu^\mu(M, j, b)$ as the probability that $\hat{b}_j > b_j$ while bids $i \neq j$ obey M (that is, $\hat{b}_i = b_i$ for $i \in M$ and $\hat{b}_i < b_i$ if $i \notin M$). Lemma 64 implies that for all j , $M \subseteq [n] \setminus \{j\}$ and $i \notin M \cup \{j\}$,

$$\sup_{\hat{b}} \left| \frac{\rho_b^\mu(\hat{b})}{b_i} \right| \geq \frac{\nu^\mu(M \cup \{i\}, j, b)}{\nu^\mu(M, j, b)} \quad a.e.$$

For any particular j , applying Lemma 65 with $\eta(S) = \nu^\mu(S, j, b)$ as above implies that

$$\max_{M \subseteq [n] \setminus \{j\}} \frac{\nu^\mu(M \cup \{i\}, j, b)}{\nu^\mu(M, j, b)} \geq \frac{1 - \gamma^{(j)}}{\gamma^{(j)}}$$

where $\gamma^{(j)}$ is

$$\gamma^{(j)} = 1 - \left(\frac{\Pr(\hat{b}_j > b_j \wedge \hat{b}_{-j} = b_{-j}|b)}{\Pr(\hat{b}_j > b_j \wedge \hat{b}_{-j} \leq b_{-j}|b)} \right)^{\frac{1}{n-1}}.$$

Since the probabilities $\Pr(\hat{b}_j > b_j \wedge \hat{b}_{-j} \leq b_{-j}|b)$ are disjoint, there must be some j such that

$$(1 - \gamma^{(j)})^{n-1} \geq \frac{\min_{i \in [n]} \Pr(\hat{b}_i > b_i \wedge \hat{b}_{-i} = b_{-i}|b)}{\frac{1}{n} \Pr(\hat{b} \not\leq b|b)}.$$

$$\frac{\Pr(\hat{b}_j > b_j \wedge \hat{b}_{-j} = b_{-j}|b)}{\Pr(\hat{b}_j > b_j \wedge \hat{b}_{-j} \leq b_{-j}|b)} \geq \frac{\min_{i \in [n]} \Pr(\hat{b}_i > b_i \wedge \hat{b}_{-i} = b_{-i}|b)}{\frac{1}{n} \Pr(\hat{b} \not\leq b|b)}.$$

Thus, it must be that

$$\max_{j, M \subseteq [n] \setminus \{j\}, i \notin M \cup \{j\}} \frac{\nu^\mu(M \cup \{i\}, j, b)}{\nu^\mu(M, j, b)} \geq \frac{1 - \gamma^{(>)}}{\gamma^{(>)}}$$

where $\gamma^{(>)}$ satisfies

$$\gamma^{(>)} = 1 - \left(\frac{\min_{i \in [n]} \Pr(\hat{b}_i > b_i \wedge \hat{b}_{-i} = b_{-i}|b)}{\frac{1}{n} \Pr(\hat{b} \not\leq b|b)} \right)^{\frac{1}{n-1}}.$$

Consequently,

$$\sup_{\hat{b}} \left| \frac{\rho_b^\mu(\hat{b})}{b_i} \right| \geq \frac{1 - \gamma^{(>)}}{\gamma^{(>)}}$$

as desired. ■

We now have the tools to prove that a resampling distribution that optimizes payments subject to a precision bound also optimizes them subject to a welfare approximation or revenue approximation bound: *Proof:*[of Lemma 58] For clarity, we argue under the nice distribution assumption. Subject to $\min_{i \in [n]} \Pr_{\mu}(\hat{b}_i > b_i \wedge \hat{b}_{-i} = b_{-i} | b) \geq \alpha > 2^{-n}$, the BKS transformation achieves

$$\sup_{\hat{b}} \left| \frac{\rho_b^{BKS}(\hat{b})}{b_i} \right| = \frac{\alpha^{\frac{1}{n}}}{1 - \alpha^{\frac{1}{n}}} ,$$

so any optimal distribution must do at least as well.

Let μ be some resampling distribution. If $\Pr_{\mu}(\hat{b} \not\leq b | b) \neq 0$, either

$$\frac{\Pr_{\mu}(\hat{b} = b | b)}{\Pr_{\mu}(\hat{b} \leq \bar{b} | \bar{b})} > \alpha ,$$

or

$$\frac{\min_{i \in [n]} \Pr_{\mu}(\hat{b}_i > b_i \wedge \hat{b}_{-i} = b_{-i} | b)}{\Pr_{\mu}(\hat{b} \not\leq b | b)} \geq \alpha .$$

In the first case, applying Lemma 70 gives

$$\sup_{\hat{b}} \left| \frac{\rho_b^{\mu}(\hat{b})}{b_i} \right| \geq \frac{1 - \gamma^{(=)}}{\gamma^{(=)}} > \frac{\alpha^{\frac{1}{n}}}{1 - \alpha^{\frac{1}{n}}} = \sup_{\hat{b}} \left| \frac{\rho_b^{BKS}(\hat{b})}{b_i} \right|$$

and therefore μ cannot be optimal.

In the second case, Lemma 70 and the assumption that $\alpha > 2^{-n} \geq \frac{1}{n^n}$ gives

$$\begin{aligned} \gamma^{(>)} &\leq 1 - (n\alpha)^{\frac{1}{n-1}} \\ &< 1 - (\alpha^{-\frac{1}{n}}\alpha)^{\frac{1}{n-1}} \\ &= 1 - \alpha^{\frac{1}{n}} . \end{aligned}$$

Thus, $\gamma^{(>)} < 1 - \alpha^{\frac{1}{n}}$, so

$$\sup_{\hat{b}} \left| \frac{\rho_b^{\mu}(\hat{b})}{b_i} \right| \geq \frac{1 - \gamma^{(>)}}{\gamma^{(>)}} > \frac{\alpha^{\frac{1}{n}}}{1 - \alpha^{\frac{1}{n}}} = \sup_{\hat{b}} \left| \frac{\rho_b^{BKS}(\hat{b})}{b_i} \right|$$

so again μ cannot be optimal.

It follows that any optimal distribution μ must have $\Pr(\hat{b} \not\leq b | b) = 0$ and, therefore

$$\min_{i \in [n]} \Pr \left(\hat{b}_i > b_i \wedge \hat{b}_{-i} = b_{-i} \mid b \right) = \Pr \left(\hat{b} = b \mid b \right) .$$

Thus, a distribution which wishes to optimize the worst-case normalized payment subject to $\Pr(\hat{b} = b | b) \geq \alpha$ will also optimize payments subject to $\min_{i \in [n]} \Pr(\hat{b}_i > b_i \wedge \hat{b}_{-i} = b_{-i} | b) \geq \alpha$, and will have $\Pr(\hat{b} = b | b) = \min_{i \in [n]} \Pr(\hat{b}_i > b_i \wedge \hat{b}_{-i} = b_{-i} | b)$. ■

B.2 Analysis Definitions, Facts, and Lemmas

This section provides a limited background on analysis concepts.

Measures and Integrals

We begin with various possible set of axioms a collection of sets may satisfy, and their technical names.

Definition 34 (σ -algebra) *The σ -algebra over a set U is a non-empty collection Σ of subsets of U that is closed under complementation and countable union of its members. The pair (U, Σ) is called a measurable space.*

Definition 35 (Generated σ -algebra) *Given a set U and a collection of subsets F of U , there is a unique smallest σ -algebra over U containing all the elements of F . This σ -algebra is denoted by $\sigma(F)$ and is called as the σ -algebra generated by F .*

Definition 36 (Borel σ -algebra) *The Borel σ -algebra $\mathcal{B}(U)$ of a metric space U is the σ -algebra generated by the collection of all open sets of U .*

Definition 37 (Measurable sets) *Once we fix a measurable space (U, Σ) , the sets $X \in \Sigma$ are called measurable sets.*

Definition 38 (Measurable functions) *Given two measurable spaces (U, Σ) and (U', Σ') , a function $f : U \rightarrow U'$ is measurable if for each $X' \in \Sigma'$, $f^{-1}(X') \in \Sigma$.*

We are now ready for the definition of a measure.

Definition 39 (Measure) *Given a measurable space (U, Σ) , we equip it with a measure ν , which is function $\nu : \Sigma \rightarrow [0, \infty]$ that satisfies*

1. $\nu(\emptyset) = 0$
2. *Countable additivity, i.e. for all countable sequences $\{X_i\}_{i \in \mathbb{Z}}$ of pairwise-disjoint sets in Σ , $\nu(\cup_{i \in \mathbb{Z}} X_i) = \sum_{i \in \mathbb{Z}} \nu(X_i)$.*

A measure ν is said to be finite if $\nu(U)$ is finite.

Definition 40 (Probability measure) *A measure is a probability measure if $\nu(U) = 1$.*

Definition 41 (Signed measure) *A signed measure is a function $\nu : \Sigma \rightarrow [-\infty, \infty]$ that satisfies $\nu(\emptyset) = 0$ and countable additivity.*

Fact 71 *If ν_1 and ν_2 are finite (signed) measures, then $\nu_3(X) = \nu_1(X) - \nu_2(X)$ is a finite signed measure.*

Convention According to standard convention, a measure is not signed unless explicitly stated. For the purposes of this paper, the set U will always be \mathbb{R}^n .

Apart from the set collections defined via σ -algebras, we also need some weaker set collections, which we define below.

Definition 42 (π -system) *The π -system over a set U is a non-empty collection P of subsets of U that is closed under finite intersection of its members, i.e., $X_1 \cap X_2 \in P$ whenever X_1 and $X_2 \in P$.*

Definition 43 (λ -system, or Dynkin system) *The λ -system over a set U is a non-empty collection L of subsets of U that is closed under complementation and countable disjoint union of its members.*

Fact 72 (Dynkin's theorem) *If P is a π -system and L is a λ -system over the same set U , and $P \subseteq L$, then $\sigma(P) \subseteq L$, i.e., the σ -algebra generated by P is contained in L .*

The Hahn and Jordan decompositions decompose a signed measure into two measures. They will be useful when we discuss the integral with respect to a signed measure.

Fact 73 (Hahn decomposition theorem) *The Hahn decomposition of a signed measure ν over a measurable space (U, Σ) consists of two sets $P, N \in \Sigma$ such that $P \cup N = U$, $P \cap N = \emptyset$, and for all measurable sets $X \subseteq P$, $\nu(X) \geq 0$ and for all measurable sets $X \subseteq N$, $\nu(X) \leq 0$. The Hahn decomposition is guaranteed to exist and be unique (up to a set of measure 0)*

Fact 74 (Jordan decomposition theorem) *This theorem is a consequence of Hahn decomposition theorem, and states that every signed measure ν can be decomposed as two (non-negative) measures $\nu^+(X) = \nu(X \cap P)$ and $\nu^-(X) = -\nu(X \cap N)$, where P and N are the Hahn decomposition of ν . The measures satisfy $\nu(X) = \nu^+(X) - \nu^-(X)$. The Jordan decomposition is guaranteed to exist and to be unique, and at least one of ν^+ and ν^- is guaranteed to be a finite measure. If ν is finite, then both ν^+ and ν^- are finite.*

Definition 44 (Characteristic Function) *The characteristic function $1_S(x)$ of a set S is the function that is 1 if $x \in S$ and zero elsewhere, i.e.*

$$1_S(x) = \begin{cases} 1, & x \in S \\ 0, & \text{otherwise.} \end{cases}$$

Definition 45 (Simple Function) *Given a measurable space (U, Σ) , a function $s : U \rightarrow \mathbb{R}$ is a simple function if it can be written as a finite linear combination of indicator function of measurable sets. That is,*

$$s(x) = \sum_{k=1}^n a_k 1_{S_k}(x)$$

for finite sequences of measurable sets $\{S_k\} \in \Sigma$ and coefficients $\{a_k\} \in \mathbb{R}$.

Fact 75 For any non-negative, measurable function f , there is a monotonic increasing sequence of non-negative simple functions $\{s_k\}$ such that

$$f = \lim_{k \rightarrow \infty} s_k .$$

Definition 46 (Integral) Given a measurable space (U, Σ) , the integral of a function $f : U \rightarrow \mathbb{R}$ with respect to a measure ν is defined incrementally. For any measurable set X , the integral of 1_X is

$$\int_U 1_X d\nu = \nu(X) .$$

For any simple function $s : U \rightarrow \mathbb{R}$,

$$\int_U s d\nu = \sum_{k=1}^n a_k \nu(X_k) .$$

For a general non-negative function $f : U \rightarrow \mathbb{R}$,

$$\int_U f d\nu = \sup \left\{ \int_U s d\nu \quad : \quad 0 \leq s \leq f \text{ and } s \text{ is simple} \right\} .$$

For general f , let $f^+(x) = \max(f(x), 0)$ and $f^-(x) = \max(-f(x), 0)$, i.e. f^+ and f^- are the positive and negative parts of f respectively. Then

$$\int_U f d\nu = \int_U f^+ d\nu - \int_U f^- d\nu .$$

Finally, for some measurable set Y ,

$$\int_Y f d\nu = \int_U f d\nu_Y$$

where $\nu_Y(X) = \nu(U \cap Y)$.

Fact 76 (Monotone Convergence Theorem) For any countable, monotone sequence of measurable functions $\{f_k\}$ (that is, sequence where $f_k \geq f_{k-1}$ pointwise),

$$\lim_{k \rightarrow \infty} \int f_k d\nu = \int \lim_{k \rightarrow \infty} f_k d\nu .$$

The following fact follows because $g_k = \sum_{i=1}^k f_i$ satisfies the monotone convergence theorem:

Fact 77 For any countable sequence of nonnegative measurable functions $\{f_k\}$

$$\sum_{k=1}^{\infty} \int f_k d\nu = \int \sum_{k=1}^{\infty} f_k d\nu .$$

Fact 78 Let $\{X_k\}$ be a countable sequence of disjoint sets. Then

$$\sum_k \int_{X_k} f d\nu = \int_{\cup_k X_k} f d\nu .$$

Definition 47 (Integral with respect to a Signed Measure) The integral of a function f with respect to a signed measure ν is

$$\int_U f d\nu = \int_U f d\nu^+ - \int_U f d\nu^- ,$$

where ν^+ and ν^- are the Jordan decomposition of ν .

Densities and Derivatives

Definition 48 (Absolute continuity) Given a signed measure ν and a measure μ on the same measurable space, ν is absolutely continuous w.r.t. μ , if for every measurable set V where $\mu(V) = 0$, we have $\nu(V) = 0$.

We now state below the Radon-Nikodym theorem the way we use it, though the theorem itself is more general.

Fact 79 (Radon-Nikodym Theorem) The Radon-Nikodym theorem states that given a finite signed measure ν and a finite measure μ on the same measurable space such that ν is absolutely continuous w.r.t. μ , the measure ν has a density, or “Radon-Nikodym derivative”, with respect to μ , i.e., there exists a μ -measurable function ρ taking values in $[0, \infty]$, such that for any μ -measurable set X we have

$$\nu(X) = \int_X \rho d\mu .$$

Fact 80 If ρ is a Radon-Nikodym derivative of measure ν w.r.t. measure μ , then

$$\int_X f(x) d\nu = \int_X \rho(x) f(x) d\mu$$

wherever $\int_X f(x) d\nu$ is well defined.

Almost Everywhere

Definition 49 (Almost Everywhere) A property $P(s)$ is said to hold almost everywhere on a set S if the subset of S on which $P(s)$ is false has measure zero (or is contained in a set that has measure 0). It is abbreviated a.e.. The exact measure used will become clear from the context.

Definition 50 (Almost Surely) *If a property $P(s)$ is false with probability 0 with respect to some distribution over s , then it is said to hold almost surely. This is equivalent to saying $P(s)$ is true almost everywhere with respect to the probability measure associated with the distribution.*

Fact 81 *For a non-negative measurable function f and measure μ , $\int f d\mu = 0$ if and only if $f(x) = 0$ almost everywhere.*

Fact 82 *For any measurable function f and signed measure ν , if $\int_X f d\nu = 0$ for all measurable X , then $f = 0$ almost everywhere.*

The second fact follows from the first by a standard argument — decompose f into its positive and negative parts and decompose ν according to its Hahn decomposition. This partitions the space into four sets over which the integral may be written as a non-negative function with respect to a non-negative measure. Apply Fact 81 to each of the four sets.

Extrema

Definition 51 (Supremum/Infimum) *For a set S , the supremum of S , denoted $\sup S$, is the smallest value x such that $x \geq s$ for all $s \in S$. Similarly, the infimum of S is the largest value x such that $x \leq s$ for all $s \in S$.*

Definition 52 (Limit Superior/Inferior) *For a real-valued function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, the limit superior, denoted $\limsup_{u \rightarrow b} f(u)$, may be defined as follows:*

$$\limsup_{u \rightarrow b} f(u) = \lim_{\epsilon \rightarrow 0} \left(\sup_{u \in \text{BALL}(b, \epsilon)} f(u) \right)$$

where $\text{BALL}(b, \epsilon)$ is the open ball of radius ϵ centered at b . It is an upper bound on the limit of $f(u_i)$ for any sequence of values $\{u_i\}$ that converges to b . The \liminf is defined similarly. Note that while the limit may not exist as $u \rightarrow b$, the \limsup and \liminf are always well defined for real-valued functions.

It is natural to generalize \sup and \limsup to almost everywhere:

Definition 53 (Essential Supremum/Infimum) *The essential supremum of a set S , denoted $\text{ess sup } S$, is the smallest value x such that $x \geq s$ almost everywhere, i.e. the set of values $T = \{s \mid s \in S \text{ and } s > x\}$ has measure zero. The essential infimum ess inf is defined similarly.*

Definition 54 (limesssup/limessinf) *For a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, the $\lim \text{ess sup}_{u \rightarrow b} f(u)$ can be defined as follows:*

$$\lim_{u \rightarrow b} \text{ess sup } f(u) = \lim_{\epsilon \rightarrow 0} \left(\text{ess sup}_{u \in \text{BALL}(b, \epsilon)} f(u) \right) .$$

It can be understood as a version of the \limsup that will ignore values that $f(x)$ only attains on sets with measure zero. The $\lim\text{ess inf}$ is defined similarly. Like the \limsup and \liminf , the $\lim\text{ess sup}$ and $\lim\text{ess inf}$ are always well defined for real-valued functions.