

Information Economics in the Age of e-Commerce:
Models and Mechanisms for Information-Rich Markets

L. Elisa Celis

A dissertation submitted in partial fulfillment
of the requirements for the degree of

Doctor of Philosophy

University of Washington

2012

Program Authorized to Offer Degree: Computer Science and Engineering

In presenting this dissertation in partial fulfillment of the requirements for the doctoral degree at the University of Washington, I agree that the Library shall make its copies freely available for inspection. I further agree that extensive copying of this dissertation is allowable only for scholarly purposes, consistent with "fair use" as prescribed in the U.S. Copyright Law. Requests for copying or reproduction of this dissertation may be referred to Proquest Information and Learning, 300 North Zeeb Road, Ann Arbor, MI 48106-1346, 1-800-521-0600, to whom the author has granted "the right to reproduce and sell (a) copies of the manuscript in microform and/or (b) printed copies of the manuscript made from microform."

Signature_____

Date_____

University of Washington

Abstract

Information Economics in the Age of e-Commerce:
Models and Mechanisms for Information-Rich Markets

L. Elisa Celis

Co-Chairs of the Supervisory Committee:
Professor Anna R. Karlin
Computer Science and Engineering

Senior Researcher Yuval Peres
Microsoft Research

The internet has dramatically changed the landscape of both markets and computation with the advent of electronic commerce (e-commerce). It has accelerated transactions, informed buyers, and allowed interactions to be computerized, enabling unprecedented sophistication and complexity. Since the environment consists of multiple owners of a wide variety of resources, it is a distributed problem where participants behave strategically and selfishly. This led to the birth of algorithmic game theory, whose goal is to understand equilibria arising in these strategic environments, study their computational complexity, and design mechanisms accordingly.

More recently, the prevalence and access to information about consumers and products has increased dramatically and changed the landscape accordingly. In this thesis I focus on two simple yet fundamental observations which require further investigation as the field of algorithmic game theory progresses in the context of information economics. Specifically:

1. Access to information widens an agent's strategy space, and
2. the generation and exchange of information between agents is itself a game.

There are three specific problems we address in this thesis.

Informed Valuations: Increasingly sophisticated consumer tracking technology gives advertisers a wealth of information which they use to reach narrowly targeted consumer demographics. With targeting, advertisers are able to bid differently depending on the age, location, computer, or even

browsing history of the person viewing a website. This is preferable to bidding a fixed value since they can choose to only bid on consumers who are more likely to be interested in their product. This results in an increase in revenue to the advertiser. Notice, however, that this implies information has changed the distribution of bids. With this change, common assumptions, which are reasonable in the absence of information, no longer hold. Thus, the mechanisms currently in place are no longer optimal.

Using historical bidding data from a large premium advertising exchange, we show that the bidding distributions are unfavorable to the standard mechanisms. This motivates our new BIN-TAC mechanism, which is simple and effective in extracting revenue in the presence of targeting information. Bidders can *buy-it-now*, or alternatively *take-a-chance* in an auction, where the top $d > 1$ bidders are equally likely to win. The randomized take-a-chance allocation incentivizes high valuation bidders to buy-it-now. We show that for a large class of distributions, this mechanism outperforms the second-price auction, and achieves revenue performance close to Myerson’s optimal mechanism. We apply structural methods to our data to estimate counterfactual revenues, and find that our BIN-TAC mechanism improves revenue by 4.5% relative to a second-price auction with optimal reserve.

Information Acquisition: A prevalent assumption in traditional mechanism design is that the buyers know their precise value for an item; however, this assumption is rarely accurate in practice. Judging the value of a good is difficult since it may depend on many unknown parameters such as the intrinsic quality of the good, the saving it will yield in the future, or the buyer’s emotional state. Additionally, the estimated value for a good is not static; buyers can *deliberate*, i.e. spend money or time, in order to refine their estimates by acquiring additional information. This information, however, comes at some cost, either financial, computational or emotional. It is known that when deliberative agents participate in traditional auctions, surprising and often undesirable outcomes can occur; see [12, 70].

We consider optimal dominant strategy mechanisms for one-step deliberative setting where each user can determine their exact value for a fixed cost. We show that for single-item auctions under mild assumptions these are equivalent to a sequential posted price mechanism. Additionally, we propose a new approach that allows us to leverage classical revenue-maximization results in deliberative environments. In particular, we use Myerson (1981) to construct the first non-trivial (i.e., dependent on deliberation costs) upper bound on revenue in deliberative auctions. This bound allows us to

apply existing results in the classical environment to a deliberative environment; specifically, for single-parameter matroid settings, sequential posted pricing is a 2-approximation or better.

Exchange Networks: Information is constantly being exchanged in many forms; i.e. communication among friends, company mergers and partnerships, and more recently, selling of user information by companies such as BlueKai. Exchange markets capture the trade of information between agents for profit, and we wish to understand how these trades are agreed upon. Understanding information markets helps us determine the power and influence structure of the network. To do this, we consider a very general network model where nodes are people or companies, and weighted edges represent profitable *potential* exchanges of information, or any other good. Each node is capable of finalizing an exactly one of its possible transactions; this models the situation where some form of exclusivity is involved in the transaction. This model is in fact very general, and can capture everything from targeting information exchange to the marriage market.

A balanced outcome [39, 94] in an exchange network is an equilibrium concept that combines notions of stability and fairness. In a recent paper, Kleinberg and Tardos [62] introduced balanced outcomes to the computer science community and provided a polynomial time algorithm to compute the set of such outcomes. Their work left open a pertinent question: are there natural, local dynamics that converge to a balanced outcome? We answer this question in the affirmative by describing such a process on general weighted graphs, and show that it converges to a balanced outcome whenever one exists. In addition, we present a new technique for analyzing the rate of convergence of local dynamics in bargaining networks. The technique reduces balancing in a bargaining network to optimal play in a random-turn game, and allows a tight polynomial bound on the rate of convergence for a nontrivial class of unweighted graphs.

Table of Contents

	Page
List of Figures	iv
List of Tables	vi
Chapter 1: Introduction	1
1.1 Informed Valuations	3
1.2 Information Acquisition	4
1.3 Information Exchange	5
Chapter 2: Background	7
2.1 Environments, Mechanisms and Utilities	7
2.1.1 Set Systems	8
2.2 Truthfulness and Equilibria	9
2.2.1 Individual Rationality	9
2.2.2 Dominant Strategy Equilibria	9
2.2.3 Nash Equilibria	10
2.2.4 Revelation Principle	10
2.2.5 Monotonicity	11
2.3 Revenue Maximization	11
2.3.1 Regular distributions	11
2.3.2 Irregular distributions	12
2.4 Approximation Mechanisms	12

Chapter 3:	Informed Valuations: Price Discrimination in Ad Auctions	13
3.0.1	Related Work:	15
	Organization:	16
3.1	The Display Advertising Market	17
3.2	Model and Analysis	19
3.2.1	The Environment	19
	Discussion:	20
3.2.2	Pricing Mechanisms	20
3.2.3	A Motivating Example	21
3.2.4	Equilibrium Analysis	22
	Equilibrium Characterization	23
	Optimizing BIN-TAC Parameters	24
3.2.5	Performance Comparisons	26
	Myerson Mechanism in Our Model	29
3.2.6	Monte Carlo Simulations	31
3.3	Empirical Application	33
3.3.1	Descriptive Evidence	36
3.3.2	Estimation and Counterfactual Simulations	39
	Mechanisms with Targeting:	41
	Bundling Mechanisms:	42
	Robustness to Informational Assumptions:	42
3.3.3	Results	43
3.4	Conclusion and Future Work	44
Chapter 4:	Information Acquisition: Mechanism Design for Deliberative Agents	47
4.1	Deliberative Model	48
4.1.1	Special Types of Deliberative Mechanisms	50
4.1.2	Information Models	51
4.1.3	Related Work:	51
4.2	Revelation Principle for DDMs	52
4.3	Characterization	53
4.3.1	Preliminary Lemmas	55
4.3.2	Main Result	57
4.3.3	Discussion and Examples	61
	Non-Optimal DDMs are not SPPs	62
	Minimal-Information DDMs can outperform SPPs	62
4.4	Approximate Revenue Maximization	64
4.4.1	Upper Bound in Classical Environment	65
4.4.2	Approximation	69
4.5	Future Work	70

Chapter 5:	Information Exchange: Convergence to Balanced Outcomes on Networks . .	71
5.1	Preliminary Definitions	75
5.2	Convergence for General Graphs	76
5.2.1	Supplementary Definitions	76
5.2.2	Convergence	78
5.2.3	Fixed Point Characterization	79
5.3	Rates of Convergence for Classes of Graphs	92
5.3.1	From a Bargaining Game to a Random-Turn Game	92
5.3.2	Rate of Convergence	95
5.3.3	Extensions	100
5.4	Conclusion and Future Work	103
Bibliography	105

List of Figures

Figure Number	Page
3.1 The display advertising market.	18
3.2 Comparison of Allocations and Payments: Allocation probabilities (top panel) and expected payments (bottom panel) for the OPT, VA-T and BIN-TAC mechanisms when the distributions F_L and F_H are uniform. The x -axis corresponds to the bid.	30
3.3 Revenue Performance vs Number of bidders: Simulated expected revenues for different mechanisms as the number of bidders n varies, in an environment where F_L has mean 1 and standard deviation 0.5, the match probability is 0.05 and the match increment is 5.	32
3.4 Revenue Performance vs Match Probability: Simulated expected revenues for different mechanisms as the probability of a match α varies, where F_L has mean 1 and standard deviation 0.5, the number of bidders is 10 and the match increment is 5.	33
3.5 Revenue Performance vs Match Increment: Simulated expected revenues for different mechanisms as the match increment Δ varies, where F_L has mean 1 and standard deviation 0.5, the match probability is 0.05 and the number of bidders is 10.	34
3.6 Bids over Time: The (rescaled) bids of five advertisers in our data, selected at random from the top 50 advertisers (ranked by purchases) on 50 randomly chosen successive impressions of the most popular product. Note that the set of impressions differs across bidders (there are no impressions on which all 5 participate).	36
3.7 Bidding Gap and Virtual Valuations: The left panel shows a kernel density estimate of the pdf of the (normalized) gap between the highest and second highest (rescaled) bids in auctions for the product with the highest sales volume in our dataset. The right panel shows the estimated virtual valuations as a function of bids.	39

4.1	A “proof by picture” of the deliberative revelation principle. The truthful DDM \mathcal{N} simulates both the dominant strategy mechanism \mathcal{M} and the dominant strategies of the agents.	54
4.2	Two main cases in the proof of Lemma 31.	60
4.3	A non-optimal deliberative mechanism that is not an SPP.	63
4.4	A mechanism that is dominant-strategy truthful in the minimal-information setting, and Bayes-Nash truthful in any information setting, but not dominant-strategy truthful in the full-information setting. This mechanism is not an SPP, and in fact attains more revenue than the optimal SPP.	64
4.5	The relationships of u_i^a, u_i^d and u_i^r in Lemma 30. (a) The case where $t_1 > t_2$. (b) The case where $t_1 \leq t_2$	67
5.1	Examples of fixed points of the balancing process that are not, in fact, balanced. In Figure 5.1(a), the fixed point is not quasi-balanced; however, no stable outcome exists. In Figure 5.1(b), the fixed point is quasi-balanced but not stable; however, the matching is not maximum.	77
5.2	The subgraph H found by the exploration algorithm takes one of the above forms. Matched edges are depicted in bold.	81
5.3	An unweighted bipartite graph G and its corresponding digraph D with balanced allocations.	93
5.4	A graph G and its corresponding digraph D using the general reduction. Note that we omit some (sv) and (vr) edges.	101
5.5	A weighted bipartite graph G and its corresponding digraph D with balanced allocations.	103

List of Tables

Table Number	Page
3.1 Revenue Comparison: Uniform Environment: Expected revenue and welfare under different mechanisms, for the uniform environment with $\Delta = 3$, $\alpha = 0.05$ and the number of bidders $n = 5$	31
3.2 Summary Statistics for AppNexus Display Ad Auctions: The full dataset is a 0.1 percent sample of a week's worth of auction data sampled sometime within the last two years. An observation is a bid in the top panel; an auction in the middle panel; and an advertiser in the last panel. Bids have been normalized so that their average is 1 for confidentiality reasons. The bid correlation is measured by selecting a pair of bids at random in every auction with at least two bidders, and computing the correlation coefficient.	35
3.3 Matching on Region: Results from OLS Regressions. In the first two columns, the dependent variable is a dummy for participation. The sample used in the regressions consists of all auction-bidder pairs, limited to the 10% of bidders who participate most often. In the last two columns, the dependent variable is the bid. The sample used in the regressions only includes bids from the 10% of bidders who bid most often. The independent variable is the population-weighted fraction of pageviews of the advertiser's website that come from the region the user is in. Time-of-day fixed effects refer to a dummy for each quarter of the day, starting at midnight. Product-region fixed effects are dummies for the page-group advertised on, and the state the user is located in. Standard errors are robust. Significance levels are denoted by asterisks (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).	38
3.4 Optimal Parameter Choices	44

3.5	Revenue-maximizing parameter choices. For each of the above mechanisms, we find these by maximizing the revenue functions defined in the main text over the available parameters numerically using a grid search. The two modified BIN-TAC mechanisms are robustness checks, varying the informational assumptions made for BIN-TAC. In the complete information case, bidders know the valuations of the other participants, and made BIN decisions accordingly. In the rationalizable worst-case model, bidders assume they will only have to pay the reserve price in TAC auction, and therefore take the BIN option more rarely.	44
3.6	Counterfactual Revenues and Welfare	45
3.7	Counterfactual simulations of average advertiser revenues, social welfare and total welfare (sum of social welfare and revenue). All statistics reported outside parentheses are averages across impressions; those in parentheses are standard errors computed by bootstrapping the full dataset (i.e. they reflect uncertainty over the true DGP). Six different simulations are run. First, a Vickrey auction without reserve. Then, a Vickrey auction with optimal (revenue-maximizing) reserve. Third, the BIN-TAC mechanism under incomplete information – the primary model considered in the text. Fourth, a Vickrey auction where the impressions are bundled according to the product and user region. The last two simulations are robustness checks, varying the informational assumptions made for BIN-TAC. In the complete information case, bidders know the valuations of the other participants, and made BIN decisions accordingly. In the rationalizable worst-case model, bidders assume they will only have to pay the reserve price in TAC auction, and therefore take the BIN option more rarely. Where applicable, the parameters used are the optimal parameters from Table 3.4.	45

Chapter 1

Introduction

In classic algorithm design, a standard unspoken assumption is that the algorithm, before processing, is given its input. However, this assumption does not hold for many settings, especially for the buying and selling of goods, where self-interested participants control the input; namely, how much they are willing to pay for an item. Thus, classical algorithmic game theory assumes agents provide input with the intent to skew the algorithm's output in their favor. However, this assumption also may not hold since the agent may not have all the necessary information to make an informed decision. In this thesis we push algorithmic game theory to this new level, incorporating ideas from information economics into mechanism design and analysis.

The objective of algorithmic game theory is to design algorithms in strategic environments; instead of assuming we will know the input to the algorithm, the input is provided by a collection of self-interested *agents*. These agents benefit or lose depending on the algorithm's *outcome*, so they may have *incentive* to manipulate their piece of the input in order to achieve a more desirable outcome for themselves. The usual considerations from classical algorithm design remain; e.g. polynomial running times, or good approximation ratio. However, the designer must also consider incentives in order to ensure the input is useful and their goals can be achieved. Algorithmic game theory has two primary components: 1) Analysis, where we look at the current algorithms or real-world scenarios and study their properties such as *equilibria* or best-response dynamics, and 2) Design: where we define the algorithm (in this case called a *mechanism*) to have both good game-theoretic

and algorithmic properties.

Information economics is a branch of microeconomics that studies how information affects economic decisions. As agents gain access to more information, they are able to condition their behavior accordingly. This expands their strategy set, and gives them more power overall. Sometimes information is freely available before an agent enters the marketplace, but other times, as with try-before-you-buy software, information can be acquired at some cost before a final decision is made. In both cases we must design stronger mechanisms that still incentivize agents appropriately.

To demonstrate the interplay between algorithmic game theory and information economics, assume you are selling your used car. How will you decide who to sell to? How much will you charge? You could announce a price and sell your car a first-come first-served basis. But, assuming your goal is to make the most money, can you do better? One option is to solicit *bids*, i.e. have all interested buyers state a price, and sell to the highest bidder.

Now, as a bidder, how will I bid? This is not clear. My bid will vary based not only on my *value* for the car, but also on my *information* about other bidders. I can lower my bid to try to save some money, but then also risk someone else raising their bid to beat mine. This introduces the study agents' *equilibrium behavior*; e.g. if I know that no one else is willing to pay more than $\$d$, I can simply offer to pay $\$d + 1$ and win, even if my value is much higher.

You, as a seller, could also decide to simplify matters by selling your car in a *truthful* manner; i.e. set up a situation such that potential buyers accurately report their values because this is, in fact, acting in their own best interest. The Vickery auction is one such model: sell to the highest bidder, at the second highest price. In this case, no one will overbid because they may have to pay more than their value, and no one will underbid because it does not provide any advantage. However, this assumes everyone has the same information about the product.

Consider, instead, the following situation: there is one highly-informed bidder (e.g. a mechanic) who knows if the value of the car is $\$2d$, or whether it is $\$0$. To me, either seems equally likely. Then, even though my expected value for the car is $\$d$, I should bid $\$0$. Why? Well, clearly I won't bid more than $\$d$. Thus, if the mechanic knows the car is valuable, they can simply outbid me and I will lose. Thus, I will only win if the mechanic knows the car is worthless, which means I got a very bad deal for $\$d$. Thus, I have no incentive to bid higher than $\$0$. Taking this one step farther, if the mechanic knows that I know, then they can buy the car for $\$1$ when it is valuable, since this will be higher than my bid! Of course, if I know that they know that I know, I may want to bid higher again.

Alternately, I may choose to hire my own mechanic to examine the car. This, however, requires

me to spend some money before even buying the car, and thus affects the final price I am willing to pay.

Clearly, information, its use and acquisition add a new twist to the study of algorithmic game theory. We now introduce the three specific problems addressed in this thesis.

1.1 Informed Valuations

The online display advertising market is an example of such a market that is highly dependent on the wealth of information we have about users. On one side of the market are the *publishers*: these are websites who have desirable content and therefore attract Internet users to browse their sites. These publishers earn revenue by selling advertising slots on these sites, small pieces of webpage *real estate* in standardized sizes. The other side of the market consists of advertisers. They would like to display their advertisements to users browsing the publisher's websites, and are in essence buying user attention. Each instance of displaying an advertisement to a user is called an *impression*. An advertiser's utility for each impression is determined by which user they are reaching, and what the user's current desires or intent are. For example, a Ferrari dealer might value high income users located close to the dealership. A mortgage company might value people that are reading an article on "how to refinance your mortgage" more than those who are reading an article on "ways to survive your midlife crisis", while the dealership might prefer the reverse.

Often, content is sold by auction through a centralized platform called an advertising exchange. Examples of leading advertising exchanges include the Microsoft Advertising Exchange, Google's DoubleClick, and Yahoo's RightMedia. Advertising exchanges work in real-time. When a user loads a participating publisher's webpage, a *request-for-content* is sent to the advertising exchange. This request will specify the type and size of advertisement to be displayed on the page, as well as information about the webpage itself (potentially including information about its content), and information about the user browsing the page such as age, gender, location, computer type, or even browsing history.

The advertising exchange will then either allocate the impression to an advertiser at a previously negotiated price, or hold an auction between participating advertisers. If an auction is held, all or some of the information about the webpage and user is passed along to ad brokers who bid on behalf of the advertisers. These ad brokers can be thought of as proprietary algorithms that take as input an advertiser's budget and preferences, and output decisions on whether to participate in an auction and how much to bid. The winning bidder's ad is then served by the ad exchange, and shown on the publisher's webpage.

The bids placed in the auction are jointly determined by the preferences advertisers have, the ad broker interface and the disclosure policies of the ad exchanges or the publishers they represent. The ad brokers can only condition the bids they place on the information provided to them: if the user’s past browsing history is not made available to them, they can’t use it in determining their bid, even if their valuation would be influenced by this information. Similarly, the advertisers are constrained in expressing their preferences by the technology of the ad broker: if the algorithm doesn’t allow the advertiser to specify a different willingness to pay based on some particular user characteristic, then this won’t show up in their bids.

Providing information, in general, holds many advantages for the advertisers. They can specify their desired audience, and thus make more money from each ad placed. Additionally, they have less competition since other advertisers will be interested in different users, further reducing the cost. However, while increasing overall surplus, the revenue of the mechanism may actually decrease [14, 70]. In Chapter 3 we propose a new second degree price discrimination mechanism for advertising platform markets. We improve revenue over current mechanisms, while retaining the benefit of targeting for the advertiser.¹

1.2 Information Acquisition

There are many settings where agents may not be given more information, but can acquire it at some cost. This is the case when one hires a mechanic to examine a used car, or spend time comparison shopping online. As an agent gains information, they refine their beliefs about the product in question and thus their valuation for the item is changing. In fact, with the advent of data exchange markets in display advertising, we can no longer ignore information acquisition as part of the agent’s strategy space. In an advertising market as described above, it is not only the ad exchange that has information about the users. New companies, such as BlueKai, are emerging that collect user data with the purpose of selling it to advertisers

Traditional mechanism design does not allow for changing valuations, leaving open the question of optimal mechanisms for this richer strategy setting.

Suppose an agent is considering buying a used car for \$8,000. The value of the car *to her* depends on her needs and preferences. She initially believes the value is uniformly between \$5,000 and \$10,000. However, she can *deliberate*, that is, she can act to reduce her uncertainty about this value. For example, she can hire a mechanic to examine the car, or take it for a test-drive. Each deliberation has a different cost (in money or time), and reveals different information. As a rational

¹Chapter 3 is based on joint work with Greg Lewis, Markus Mobius and Hamid Nazerzadeh [26, 27].

agent, she evaluates the cost and value of information for each deliberation, and chooses the best one. She then decides whether or not to purchase the car, based on what she learned.

This example introduces a deliberative agent, who is uncertain about her preferences, but can take actions to reduce the uncertainty. Judging the value of a good is difficult since it depends on many parameters, some of which are unknown, such as the intrinsic quality of the good, the saving it will yield in the future, or personal biases. Additionally, there may be computational constraints that prevent an agent from achieving certainty about her valuation.

In this Chapter 4, we study the design of revenue-maximizing auctions with dominant strategies in deliberative settings, and make two contributions. First, we show that posted-price-based auctions characterize the space of dominant strategy auctions in significantly more general deliberative settings. Second, we show how to design auctions that obtain revenue that is within a small constant factor of the maximum possible revenue in these settings. ²

1.3 Information Exchange

Finally, we consider the exchange of information in a network. Often, one party cannot monetize information on their own, they need a partner; companies such as BlueKai who sell cookie information to advertisers is a prime example. Together, the information holder and the information user can make some profit. However, this raises the question of how they should split their profit. There are a few key aspects of these information exchange models:

- information loses its potency the more it spreads, confidentiality contracts are common – i.e. each piece of information is sold to one person;
- there are many information sources available, if i do not get it from one source, i can get it (or some other useful information) from somewhere else;
- all transactions are local, partnerships are formed without regard to the rest of the network.

We model these interactions using general exchange models and network bargaining games which have been studied previously in economics and the social sciences.

In a network bargaining game, nodes in a graph are involved in pairwise transactions with their neighbors. This type of game was introduced by Cook and Yamagishi [39] to capture the *power* of a node derived from its position in a network, and has also been used in economics to model

²Chapter 4 is based on joint work with Anna Karlin, Dimitrios Gklezakos, Kevin Leighton-Brown, Thach Nguyen, and David Thompson [24, 25].

two-sided markets [90, 94]. Recently these games have been analyzed from a computational point of view. First, balanced outcomes were computed in a centralized model [62]. Analyzing simple, local dynamics that converge quickly to an equilibrium in such games was an important open problem that attracted much interest [49, 60, 61].

In Chapter 5 we first focus on a specific local bargaining procedure defined on a matching, and show that it converges to a balanced outcome whenever a balanced outcome exists. The convergence proof uses a potential function argument. Given the potential function, the proof of convergence is fairly easy. However, the potential function we use is somewhat peculiar, owing to the fact that many natural potential functions do not work. This is because balancing one matched edge might lead to an increase in the imbalance of many other edges. The only special case where an alternate (easier) potential function is known is for a path.

We then focus on the rate of convergence and obtain a tight polynomial bound on the rate of convergence for a variety of natural dynamics on a certain class of unweighted graphs. The potential function used in our initial convergence proof is exponential, so does not provide any immediate bounds; instead, we draw a connection between network bargaining games and random-turn games. Random-turn games are a well-studied class of two-player combinatorial games in which the outcome of a coin flip determines which player moves next [66, 67]. Combinatorial games can be represented as a game on a directed graph where players move a token along edges until one reaches their goal state. We transform the network bargaining game into an equivalent random-turn game which we can analyze using martingale techniques to obtain bounds on the rate of convergence. In particular, the convergence rate for the dynamics is related to the *absorption time* of the corresponding random-turn game.³

³Chapter 5 is based on joint work with Yossi Azar, Ben Birnbaum, Nikhil Devanur and Yuval Peres [8, 22].

Chapter 2

Background

In this chapter we define important concepts that appear throughout this thesis. We also introduce a number of classical results which we will build upon or use. Specifically, this includes the Vickrey auction, equilibrium concepts, information sets in games, truthfulness, the Myerson mechanism, and previous approximation mechanisms and characterization results.

2.1 Environments, Mechanisms and Utilities

We consider mechanisms for the following general setting.

Definition 1 (Single-Parameter Environment). *A single-parameter environment is a set of n agents and a collection of feasible sets of agents \mathcal{S} , which represent the subsets of agents that can be served simultaneously. Each bidder i has a valuation $v_i \sim \mathcal{F}_i$ that parameterizes their benefit from being served by the mechanism.*

A common and useful case is a *single-item* auction, where $\mathcal{S} = \{\{1\}, \{2\}, \dots, \{n\}\}$. In a k -unit auction, any subset of k agents can be served. Alternately, there could be combinatorial constraints given by, say, a graph $G = ([n], E)$. If $uv \in E$ this means both u and v want the same item, so the mechanism cannot serve both simultaneously. Thus, $\mathcal{S} \subseteq \{S \subseteq [n] : \forall uv \in E, \{u, v\} \not\subseteq S\}$. In a general single-minded combinatorial auction, any subset of agents whose desired item sets don't overlap can be served. See [56] for a further discussion of single-parameter settings.

Valuations are assumed to be *private*, i.e., unknown to the mechanism and to other agents. We denote the vector of valuations as $v = (v_1, \dots, v_n)$, and let $v_{-i} = \{v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_n\}$. We let $\mathcal{F} = \mathcal{F}_1 \times \mathcal{F}_2 \times \dots \times \mathcal{F}_n$ be the product distribution from which valuations are drawn. These distributions are public knowledge. This definition is quite general, and can capture most classical settings.

We now wish to define how these agents are served.

Definition 2 (Mechanism). *A mechanism $\mathcal{M} = (x, p)$ is an allocation rule and a payment rule which take bids $b = (b_1, \dots, b_n)$ as input.*

- An allocation rule $x : b \rightarrow \{0, 1\}^n$, indicates the winners ($x_i = 1$) and losers ($x_i = 0$).
- A payment rule $p : b \rightarrow \mathbb{R}_+^n$ indicates the amount each agent is asked to pay.

We assume mechanisms only charge agents they serve.

Each agent i provides the mechanism with the bid b_i .

Definition 3 (Strategy). *A strategy of an agent is a function $\sigma : (v_i, \mathcal{F}, \mathcal{M}) \rightarrow b_i$ that maps her information about her valuation, other valuation profiles and the mechanism at hand, to a bid.*

To define how an agent chooses a bid, we must first understand their objective.

Definition 4 (Utility). *Given a mechanism the utility for an agent i is $u_i(b) = b_i x_i(b) - p_i(b)$.*

Note that if the mechanism is *randomized*, i.e. uses random bits in x or p , the utility is in fact defined to be the *expected* utility. We rarely state this distinction explicitly. Agents are assumed to be selfish, and choose the strategy σ which maximizes their utility.

2.1.1 Set Systems

Definition 5. *A set system \mathcal{C} is downward-closed if $A \subseteq B$ and $B \in \mathcal{C}$ imply $A \in \mathcal{C}$. In particular $\emptyset \in \mathcal{C}$ unless $\mathcal{C} = \emptyset$. An environment whose feasible set system is downward-closed is called an downward-closed environment.*

Matroids are a downward-closed environments.

Definition 6. *A matroid \mathcal{C} is a downward-closed set system that have the following property: if $A, B \in \mathcal{C}$ and $|A| < |B|$ then there exists $x \in B \setminus A$ such that $A \cup \{x\} \in \mathcal{C}$. The sets in \mathcal{C} are called independent sets. An environment whose feasible set system is a matroid is called a matroid environment.*

Some special matroids are usually considered in approximation mechanisms.

Definition 7. • A uniform matroid is characterized by a ground set X and a number k . The matroid contains all subsets of size at most k of X .

- A partition matroid is characterized by a partition of a ground set X into t partitions X_1, X_2, \dots, X_t and t numbers k_1, k_2, \dots, k_t . The matroid contains all subsets whose intersections with each X_i contain at most k_i elements.

In particular, single-item auctions and k -unit auctions are uniform matroid environments.

2.2 Truthfulness and Equilibria

A mechanism is truthful if for every bidder i and every fixed set v_{-i} of bids by the other agents, the bidder maximizes her utility by bidding her true valuation (i.e. $b_i = v_i$). This is also known as *incentive compatibility*; since the *incentive* of the user is aligned with that of the mechanism, the bidder reports the truth. We often study truthful mechanisms and, in such cases, do not formally distinguish between agents valuations and the bids they submit to the mechanism.

2.2.1 Individual Rationality

The following common notion guarantees that agents, at the end of their interaction with the mechanism, suffer no loss.

Definition 8 (Individual Rationality). A mechanism is individually rational if truthful agents are guaranteed non-negative utility by the mechanism; i.e. $p_i(v) \leq v_i x_i(v)$ for all i and v .

Depending on the situation, this may or may not be a reasonable condition. For example, in the deliberative setting of Chapter 4, gathering information comes at some cost, and does not guarantee service from the mechanism. Hence, restricting a mechanism to be individually rational severely constrains our possibilities. On the other hand, as in Chapter 5, we consider partnerships where it does not make sense that a person would accept a deal where they lose money; in this case we assume individual rationality as a constraint.

2.2.2 Dominant Strategy Equilibria

We first introduce a strong notion of equilibria. If we have two agents, Alice and Bob, then they are in a dominant-strategy equilibrium if Alice's is making the best decision taking into account *all possible* decisions Bob could make, and Bob does the same considering all possible decisions Alice could make.

Definition 9 (Dominant Strategy (DS)). *A strategy σ_i is a dominant strategy for an agent i if for all values v_i , all potential alternate strategies σ'_i , and all strategies σ_{-i} of the other agents, the utility for i using σ_i is at least the utility using σ'_i .*

A dominant strategy equilibrium is a set of strategies σ such that each σ_i is a dominant strategy. A dominant strategy equilibrium is said to be *truthful* if the agent's strategies at this equilibrium are to reports their value truthfully. Note that a such equilibria need not exists. In Chapter 4, we focus primarily on dominant strategy truthful equilibria.

2.2.3 Nash Equilibria

We can now consider a weaker notion of an equilibrium; namely, a stable set of strategies. For example, Alice and Bob are in a *Nash equilibrium* if Alice is making the best decision she can, taking into account Bob's decision, and Bob is making the best decision he can, taking into account Alice's decision. Likewise, a group of players are in a Nash equilibrium if each one is making the best decision that he or she can, taking into account the decisions of the others. The formal definition is as follows:

Definition 10 (Nash Equilibria). *A set of strategies $\sigma = (\sigma_1, \dots, \sigma_n)$ is a Nash equilibria if for all i , assuming σ_{-i} remains fixed, the utility for i using σ_i is at least the utility using σ'_i .*

We consider a generalization of Nash equilibria in Chapter 5 to network settings.

A Bayes-Nash Equilibria is one where, given a known prior \mathcal{F}_i for each agent, the strategy where agents report their values truthfully is a nash equilibria; i.e. agent i 's expected utility over \mathcal{F} for reporting truthfully is at least as much as any other strategy σ'_i assuming the other agents report their values truthfully.

2.2.4 Revelation Principle

The classical revelation principle asserts that any dominant strategy mechanism is equivalent to a single-stage truthful mechanism:

Definition 11 (Revelation Principle [53]). *For any dominant strategy \mathcal{M} and any of its dominant strategy equilibria, there is a truthful mechanism \mathcal{N} such that the outcome of \mathcal{N} when agents bid truthfully is exactly the outcome of \mathcal{M} at this dominant strategy equilibrium.*

The proof of this result is simple: \mathcal{N} simply simulates both \mathcal{M} and the dominant strategies of the agent. Since \mathcal{N} simulates both \mathcal{M} and the strategy of each agent, it simply needs the agents' values

to proceed, thus the mechanism is “flattened” to a single-stage mechanism – the multiple stages are simulated internally.

The importance of this result lies in the fact that we can now restrict ourselves to study truthful mechanisms. If we wish to implement some outcome, we first restrict to truthful single-stage mechanisms; if no such truthful mechanism exists, then no mechanism, even those that are multi-stage or *not* truthful, can implement this outcome. By narrowing our search space, designing mechanisms and proving optimality becomes more tractable.

As written, this statement of the revelation principle holds for dominant strategy mechanisms. The revelation principle also holds in other settings, such as for Bayes-Nash equilibria [44, 58, 80]. For deliberative settings, variants of the revelation principle also exist [65, 97], and we state and prove a general version in Chapter 4.

2.2.5 Monotonicity

Lemma 12 (Monotonicity ([79])). *In any truthful mechanism, for any agent i and bids b_{-i} from other agents, there exists a critical value $t_i(b_{-i})$ where if $b_i > t_i(b_{-i})$ then i is served and pays exactly $t_i(b_{-i})$, otherwise she is not served and pays nothing.*

We will extend this lemma naturally to deliberative settings in Chapter 4.

2.3 Revenue Maximization

The mechanism designer, in general will also have a goal. In this thesis, we focus solely on the goal of *revenue maximization*.

Definition 13 (Revenue). *The revenue of a mechanism $M = (x, p)$ on an input b is $M(b) = \sum_{i=1}^n p_i(v)$, namely the sum of payments collected.*

Different mechanisms may earn drastically different revenue on a specific input b . However, for any given \mathcal{F} , the expected revenues of different mechanisms are comparable. As is traditional in optimal mechanism design, we assume that the mechanism designer knows \mathcal{F} but not the instantiation of v .

2.3.1 Regular distributions

Definition 14 (Regular distribution). *Consider a distribution F with density function f . A virtual valuation of F at v is $\phi(v) = v - \frac{1-F(v)}{f(v)}$. The distribution F is regular if $\phi(v)$ is increasing.*

Definition 15 (Vickrey auction with reserve (VA)). *labeldef:Vickrey* The Vickrey auction with reserve r is as follows: Let $b^* = \max_{i \in [n]} \{b_i\}$ and $i^* = \operatorname{argmax}_{i \in [n]} \{b_i\}$. Additionally, let $b_{-i^*}^* = \max_{i \in [n] \setminus \{i^*\}} \{b_i\}$, namely, the second highest price. If $b^* < r$, no allocation is made. Otherwise, the mechanism allocates to the bidder i^* , and charges $\max\{b_{-i^*}^*, r\}$.

The Vickrey auction with reserve is also commonly referred to as a second-price auction for obvious reasons. In general, we may omit stating a reserve explicitly.

Lemma 16. *The Vickrey auction with reserve is a truthful auction.*

Theorem. *For single-item auctions where agents valuations are drawn i.i.d. from a regular distribution, the Vickrey auction with a reserve price maximizes the seller's expected revenue over all truthful individually rational auctions.*

2.3.2 Irregular distributions

We turn our attention to irregular distributions; namely, distributions that do not satisfy the regularity condition in Definition 14. For these distributions, the mechanism in Definition ?? is not optimal in general. We refer the reader to [57] for an exposition of Myerson's mechanism. The details of Myerson's construction are not essential for understanding the results in this thesis, though a special case is considered and explained in Chapter 3. The original result was stated for continuous distributions which are strictly positive on an interval. See [77] and [84] for optimal mechanisms when distributions have discrete, disjoint or mixed support.

Theorem (Optimal Mechanism [79]). *The Myerson mechanism given in [79] is the optimal profit-maximizing truthful mechanism.*

2.4 Approximation Mechanisms

Often we either cannot prove a mechanism is optimal, or we wish to use a mechanism that we know is not optimal because it has some other desirable properties. However, we can still say something useful about how *close to optimality* the mechanism is; e.g., in single-item auctions, the Vickrey auction with optimal reserve is a 2-approximation; namely, the Myerson mechanism will never make more than twice what Vickrey auction makes.

For our approximation results, it is often useful to restrict the environment. One option would be to regulate the distributions we consider (as in when considering only regular distributions). Alternately, we can constrain the environments' feasible sets.

Chapter 3

Informed Valuations: Price Discrimination in Ad Auctions

Advertising technology is changing fast. Consumers can now be reached while browsing the internet, playing games on their phone or watching videos on YouTube. The large companies that control these new media — household names like Google, Facebook and Yahoo! — generate a substantial part of their revenue by selling advertisements. They also know increasing amounts of information about their users. This allows them to match advertisers to potential buyers with ever greater efficiency. While this matching technology increases advertisers' expected utility, it also tends to create thin markets where perhaps only a single advertiser has a high willingness to pay. These environments pose special challenges for the predominant auction mechanisms that are used to sell online ads because they reduce competition among bidders, making it difficult for the platform to increase revenue despite the fact agent's utility is increased by targeting.

For example, a sportswear firm advertising on the New York Times website may be willing to pay much more for an advertisement placed next to a sports article than one next to a movie review. It might pay an additional premium for a local consumer who lives in New York City and an even higher premium if the consumer is known to browse websites selling sportswear. Each layer of targeting increases the sportswear firm's valuation for the consumer but also dramatically narrows the set of participating bidders, to fellow sportswear firms in New York City. Without competition, revenue performance may be adversely affected [14, 70].

Consider a simple model: When advertisers “match” with users, they have high valuations; other-

wise they have low valuations. Assume that match probabilities are independent across bidders, and sufficiently low that the probability that *any* bidder matches is relatively small. Then a second-price auction will typically get low revenue, since the probability of two “matches” occurring in the same auction is small. On the other hand, setting a high fixed price is not effective since the probability of zero “matches” occurring is relatively large and many impressions would go unallocated. Hence, allowing targeting creates asymmetries in valuations that can increase an advertiser’s utility, but decrease the mechanism’s revenue. In fact, because of this phenomenon, some have suggested that it is better to create market with fewer types (and hence more bidders per type) by not disclosing information; i.e., “bundling” many different impressions together [48, 52, 75]. The question of how to optimally bundle is a subject of ongoing research [15].

Since targeting increases agent’s utility, platforms would like to allow targeting while still increasing their revenue. This chapter outlines a new and simple mechanism for doing so. We call it *buy-it-now or take-a-chance* (BIN-TAC), and it can be thought of as a two stage mechanism where in the first stage agents self-select into either a high buy-it-now (BIN) auction or a low take-a-chance (TAC) auction. In the second stage, we run one of the two auctions. The BIN auction is a Vickrey auction with a high buy-it-now reserve p . The TAC auction is such that the top d bidders are eligible to receive the item, and the item is awarded to one of them with equal probability at the $(d + 1)$ st price. If at least one agent self-selects to buy-it-now we run the BIN auction, otherwise we run the TAC auction.

When matches occur, advertisers can self-select into the buy-it-now option, allowing the mechanism to profit from the match. Advertisers are incentivized to “buy-it-now” because in the event that they “take-a-chance”, there is a significant probability they will not win the impression, even if their bid is the highest. On the other hand, when no matches occur, the auction mechanism ensures the impression is still allocated, thereby earning revenue. This is Bayes-Nash incentive compatible since both BIN and TAC are dominant-strategy incentive compatible, and the choice of which option can be optimized based on other agent’s expected behavior.¹

BIN-TAC is simple, both in that it is easy to explain to advertisers and in that it requires relatively little input from the mechanism designer: a choice of buy-it-now price, randomization parameter d and optionally a reserve in the take-a-chance auction. As we show both analytically and through Monte Carlo simulation, BIN-TAC generally outperforms the two leading alternatives: a Vickrey auction with reserve, or the “bundling” solution in which the platform withholds targeting information. At least in principle one could do better still by using the revenue-optimal mechanism

¹For any equilibrium, this mechanism can be reinterpreted in the usual form where a bidder simply gives the mechanism a value and the mechanism acts in the agent’s best interest.

suggested in [81], which is considerably more complicated. We demonstrate that in our context BIN-TAC closely approximates the allocations and payments of the optimal mechanism, achieving similar performance.

To analyze its performance in a real-world setting, we turn to historical data from the Microsoft Advertising Exchange. By estimating the distribution of advertiser valuations, we can simulate the effect of introducing the BIN-TAC mechanism. We also consider a bundling strategy in which all impressions on a given webpage browsed by a user located in a particular geographic region are sold as identical products. We find that the optimal BIN-TAC mechanism generates 4.5% more revenue than the optimal second-price auction, while at the same time improving consumer utility by 11%. This is possible because in order to benefit from the high bidders the optimal second-price auction must use a high reserve, hence excluding low bidders. The BIN-TAC mechanism benefits from high bidders through a high buy-it-now price, and avoids excluding low valuation bidders by allowing the TAC auction. Both these mechanisms outperform the bundling strategy we consider, although performance may be improved using a different bundling scheme.

We view the main contribution of this chapter as introducing and analyzing a new and simple price discrimination mechanism that makes use of randomized auctions, and then testing its performance in a realistic environment. While our focus is on the display advertising market, we note that there are other markets in which randomized allocations are used as a screening tool. For example, Priceline offers users the choice between a hotel of their choice at a fixed high price, or the opportunity to bid for a random hotel room of certain guaranteed characteristics (e.g. location, star rating).

A secondary contribution of the chapter is to document participation and bidding behavior in the display advertising market. While there has been theoretical work on this market [73, 78], and some empirical work on the search advertising market ([83], [6]), there has been little empirical work on display advertising. We document the large gap between the highest and second highest valuations in these auctions. We also show that advertisers do, in fact, condition their bids on the geographical location of the viewers, taking advantage of user demographics provided by the platform to achieve better matches. Overall this work supports the assumptions typically made in the theoretical chapters cited above.²

3.0.1 *Related Work:*

Our work is related to the literature on price discrimination and screening. Here we consider a mechanism that treats all bidders symmetrically, and proceeds sequentially. Other chapters have

²This chapter is based on joint work with Greg Lewis, Markus Mobius and Hamid Nazerzadeh [26, 27].

suggested sequential screening approaches. [42] consider a setting where the buyers themselves learn their type dynamically, in two stages. In this case, offering contracts after the first type revelation but before the second may be optimal; see [17] for a survey on dynamic mechanisms. In the static setting, sequential screening and posted-price mechanisms can be used to design optimal (or near-optimal) mechanisms when the bidders have multi-dimensional private information (see for example [89] and [34]).

More generally, the question of whether mechanisms should provide information that allows buyers to target their bids arises in the analysis of optimal mechanism disclosure (see for example [16]). The idea of bundling items together to take advantage of negative correlation in valuations — in this case the negative correlation in the valuations from “match” or “no match” — dates back to [4]; see also [74]. Our chapter is similar in style to [36], who combine theory, simulations and empirics to argue that bundle-size pricing is a good approximation to the more complicated (but theoretically superior) mixed bundling pricing scheme for a monopolist selling multiple items.

Finally, our model considers only the private values setting. [1] consider a problem that arises in a pure common value setting when some bidders are privately informed. This is motivated by the case when some advertisers are better able to utilize the user information provided by the platform. They show that asymmetry of information can sometimes lead to low revenue in this market.

From an empirical perspective, this chapter contributes to the growing literature on online advertising and optimal pricing. Much of the work here is experimental in nature — for example, [71] ran a randomized experiment to test advertising effectiveness, while [83] used an experimental design to test the impact of reserve prices on revenues. There has also been recent work on privacy and targeting in online advertising [54, 55].

Organization:

The chapter proceeds in three parts. First, we give an overview of the market for display advertising. In the second part we introduce a stylized environment, and prove existence and characterization results for the BIN-TAC mechanism. We also provide analytic results concerning the revenue maximizing parameter choices, and compare our mechanism to others using both theory and Monte Carlo simulation. Finally, in the third part we provide an empirical analysis of a display advertising marketplace, including counterfactual simulations of our mechanism’s performance. .

3.1 *The Display Advertising Market*

This chapter proposes a new mechanism for advertising platforms such as Microsoft, Google and Facebook. In these markets, advertisers care about the characteristics of the users they advertise to, but it is up to the platform to choose whether or not to disclose what they know about their users. The online display advertising market is an example of such a market. Its organization is depicted in Figure 3.1. On one side of the market are the “publishers”: these are websites who have desirable content and therefore attract Internet users to browse their sites. These publishers earn revenue by selling advertising slots on these sites.

The other side of the market consists of advertisers. They would like to display their advertisements to users browsing the publisher’s websites. They are buying user attention. Each instance of showing an advertisement to a user is called an “impression”. Advertiser demand for each impression is determined by which user they are reaching, and what the user’s current desires or intent are. For example, a Ferrari dealer might value high income users located close to the dealership. A mortgage company might value people that are reading an article on “how to refinance your mortgage” more than those who are reading an article on “ways to survive your midlife crisis”, while the dealership might prefer the reverse.

Some large publishers, primarily AOL, Microsoft and Yahoo!, sell directly to advertisers. Since the number of users browsing such publishers is extremely large (e.g. 1.5% of total worldwide Internet pageviews are on Yahoo), they can predict with high accuracy their user demographics. Consequently, they can contract to sell a large fixed number of impressions in pre-defined categories to a particular advertiser. For example, they could agree to sell 1 million impressions of male 15-24 year olds living in New York City viewing the Yahoo! homepage. Transactions of this kind are generally negotiated between the publisher and the advertiser.

Alternatively, content is sold by auction through a centralized platform called an advertising exchange. Examples of leading advertising exchanges include the Microsoft Advertising Exchange (a subset of which we examine in this chapter), Google’s DoubleClick, and Yahoo’s RightMedia [78]. Advertising exchanges are a minor technological wonder. They work in real-time. When a user loads a participating publisher’s webpage, a “request-for-content” is sent to the advertising exchange. This request will specify the type and size of advertisement to be displayed on the page, as well as information about the webpage itself (potentially including information about its content), and information about the user browsing the page.³

The advertising exchange will then either allocate the impression to an advertiser at a previously

³For example, it may include their IP address and cookies that indicate their past browsing behavior.

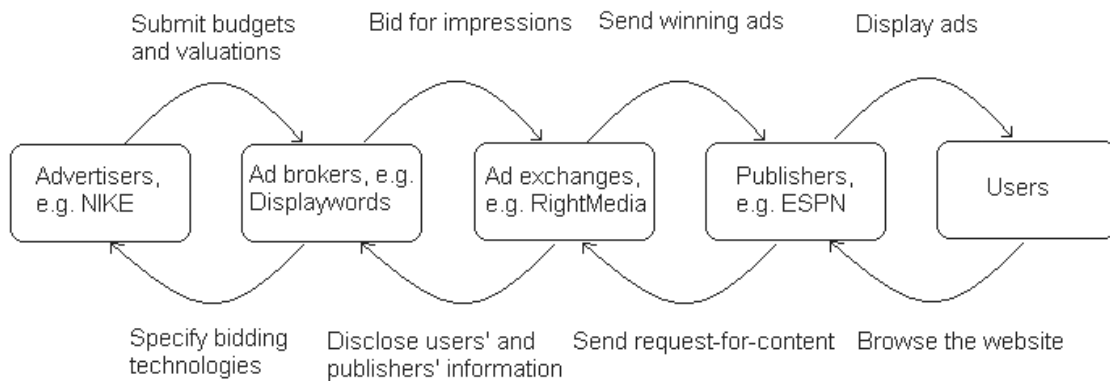


Figure 3.1: The display advertising market.

negotiated price, or hold a second-price auction between participating advertisers. If an auction is held, all or some of the information about the webpage and user is passed along to ad brokers who bid on behalf of the advertisers. These ad brokers can be thought of as proprietary algorithms that take as input an advertiser’s budget and preferences, and output decisions on whether to participate in an auction and how much to bid. The winning bidder’s ad is then served by the ad exchange, and shown on the publisher’s webpage.⁴

The bids placed in the auction are jointly determined by the preferences advertisers have, the ad broker interface and the disclosure policies of the ad exchanges or the publishers they represent. The ad brokers can only condition the bids they place on the information provided to them: if the user’s past browsing history is not made available to them, they can’t use it in determining their bid, even if their valuation would be influenced by this information. Similarly, the advertisers are constrained in expressing their preferences by the technology of the ad broker: if the algorithm doesn’t allow the advertiser to specify a different willingness to pay based on some particular user characteristic, then this won’t show up in their bids.

Ad exchanges have two main advantages over direct negotiation. First, they economize on transaction costs, by creating a centralized market for selling ad space. Second, they allow for very

⁴To make things yet more complicated, in some ad exchanges — though not Microsoft Advertising Exchange — two different pricing models coexist. The first is pay-per-impression, which is what we analyze in the current chapter; the second is pay-per-click, where the payment depends on whether or not the user clicks on the advertisement. Ad exchanges use expected click through rates to compare these different bids through a single expected revenue number.

detailed products to be sold, such as the attention of a male 15-24 year old living in New York City viewing an article about hockey that has previously browsed articles about sports and theater. There is no technological reason why the products need to be sold in “categories”, as publishers tend to do when guaranteeing sales in advance. This “real-time” sales technology is often touted as the future of this industry, as it potentially improves the match between the advertiser and their target audience. We will focus on developing a real-time pricing mechanism for display advertising exchanges.

3.2 Model and Analysis

3.2.1 The Environment

A publisher has an impression to sell in real time, and they have information about the user viewing the webpage, summarized in a cookie. The mechanism is considering one of two policies: either disclosing the cookie content to the advertiser (the “targeting” policy), or withholding it (the “bundling” policy). When they allow targeting, bidders know whether the user is a “match” for them or not. When a match occurs, the bidder has a high valuation. But the probability of a match is low and matches are assumed independent, so it is likely that everyone in the auction has a low valuation. Additionally, allowing targeting may make the market “thin”, meaning that there are few bidders of each type; this is natural, targeting differentiates between types more precisely, so the same average number of bidders per type will decrease.

Instead, the mechanism may choose to withhold the cookie, so that bidders are uncertain about whether the user is a match for them or not. The mechanism thus bundles item impressions with bad ones, so that bidders have intermediate valuations. This keeps the market “thick” since we can no longer distinguish between types with such granularity. However, it also reduces agent’s utility since they can no longer form matches.

The formal model is as follows. There are n symmetric bidders who participate in an auction for a single item which is valued at zero by the mechanism. Bidders are risk neutral. They have value V_H for the item when a match occurs, and value V_L for the item if no match occurs, where $V_L \sim F_L$ and $V_H \sim F_H$. We assume that F_L has support $[\underline{\omega}_L, \bar{\omega}_L]$ and F_H has support $[\underline{\omega}_H, \bar{\omega}_H]$, and that these supports are disjoint (so $\bar{\omega}_L < \underline{\omega}_H$). We assume both F_L and F_H have continuous densities f_L and f_H . The Bernoulli random variable X indicates whether a match has occurred, and the event $X = 1$ occurs with probability $\alpha \in (0, 1)$.

The bidder type is a triple (X, V_L, V_H) is drawn identically and independently across bidders, so that a user who is a match for one advertiser need not be a match for the others. With targeting,

each advertiser’s realized valuation $V = (1 - X)V_L + XV_H$ is private information, known only to the advertiser. Instead, if the mechanism bundles all impressions, the advertiser knows V_L and V_H but does not know the realization of X , implying their expected valuation is $E[V] = (1 - \alpha)V_L + \alpha V_H$.

For simplicity of presentation, we also make some technical assumptions on the *virtual valuations* $\psi(v) = v - \frac{1-F(v)}{f(v)}$. We assume that $\psi(v)$ is continuous and increasing over the regions $[\underline{\omega}_L, \bar{\omega}_L]$ and $[\underline{\omega}_H, \bar{\omega}_H]$. We additionally assume that $\psi(v)$ single-crosses zero, that this intersection occurs in the low valuation region $[\underline{\omega}_L, \bar{\omega}_L]$, and that $\psi(\bar{\omega}_L) \leq \psi(\underline{\omega}_H)$.⁵ Overall, our environment is fully characterized by the tuple (n, α, F_L, F_H) .

Discussion:

We assume that the match random variables X and the valuations V_L and V_H are independent across bidders. We focus on independence for two reasons. First, it is an assumption that is often made in the screening and mechanism design literatures, and so is a natural starting point. Second, in the log data examined in this chapter we observe little correlation in bids.⁶

We also will focus on environments where α is small, since this implies that the probability of either a single match or no match is high. This is the interesting case, reflecting the industry concern that providing “too much” targeting information reduces competition and hurts revenue. In our data we often observe a large gap between the highest and second highest bid, which provides support for this focus.

3.2.2 Pricing Mechanisms

Our BIN-TAC mechanism works as follows. A *buy-it-now price* p is posted. Buyers simultaneously indicate whether they wish to *buy-it-now* (BIN). In the event that exactly one bidder elects to buy-it-now, that bidder wins the auction and pays p . If two or more bidders elect to BIN, a second-price sealed bid auction with reserve p is held between those bidders. Bidders who chose to BIN are obliged to participate in this auction. Finally, if no-one elects to BIN, a sealed bid *take-a-chance* (TAC) auction is held between all bidders, with a reserve r . In that auction, one of the top d bidders is chosen uniformly at random, and if that bidder’s bid exceeds the reserve, they win the auction and pay the maximum of the reserve and the $(d + 1)$ -th bid. Ties among d -th highest bidders are broken

⁵Without these assumptions we would have to analyze multiple cases, which is straightforward but tedious.

⁶In practice, the information that platforms may choose to disclose is multidimensional, and some user characteristics may be “vertical” (e.g. income) and therefore induce positive correlation in match probabilities; while others may be “horizontal” and have correlation implications that depend on the population of advertisers (e.g. age).

randomly prior to the random allocation. We call r the TAC-reserve, and d the randomization parameter.

To analyze the performance of BIN-TAC, it will be useful to have some benchmarks for comparison. A natural benchmark is the pricing mechanism that is most commonly used in practice, the Vickrey auction (VA). We distinguish between when a VA is used and targeting is allowed (VA-T), and when it is used without targeting (VA-B). Throughout this chapter we assume, in general, that targeting is used unless specified otherwise.

A third benchmark is the revenue-optimal mechanism within the class of those that allow targeting (i.e. those that commit to reveal the cookie to all bidders for free).⁷ Usually this mechanism is the second-price auction with an optimally chosen reserve price. However in this case the virtual valuations $\psi(v)$ are not increasing over the whole support of F — indeed they are (infinitely) negative over the region $(\bar{\omega}_L, \underline{\omega}_H)$. In this case, the optimal mechanism requires ironing [81].

In plain terms, ironing implies that sometimes the allocation will be randomized among bidders with different valuations. Just as in our TAC auction, the winner of the auction need not have the highest valuation. The difference is that in the optimal mechanism, the randomization only takes place when two or more bidders — including the highest valuation bidder — have valuations in a given “ironing” region. By contrast, in BIN-TAC this randomization occurs whenever no-one takes the BIN option. The differences will be clearer later when we compare the performance of the mechanisms. For now, we would like to present a simple example to illustrate why our BIN-TAC mechanism may be good at increasing both social welfare and revenue, as a motivation for the detailed equilibrium analysis that follows.

3.2.3 A Motivating Example

Consider a special case of our environment with just two bidders, and a match probability of 10%. Bidders have fixed symmetric valuations, equal to 10 if they match, and 1 otherwise.

Now consider the expected outcomes of the two second-price auction mechanisms. First, consider running the Vickrey auction in a setting where we allow targeting (VA-T). The allocation will be fully efficient – the highest valued bidder will be allocated the item. The probability that at least one bidder has a high valuation is $1 - (0.9)^2 = 0.19$, and so the expected value of the winning bidder for the item is $10(0.19) + 1(0.81) = 2.71$. On the other hand, the probability that both bidders match is only 1%, so expected revenue is only $(0.01)10 + (0.99)1 = 1.09$.

⁷A mechanism may potentially do better by withholding match information altogether (bundling), or by selling the rights to the match information — see [16].

Now, consider the Vickrey auction in a setting where targeting is not allowed (VA-B). This is effectively bundling, the bidders don't know if the impression is a match, and thus bid their expected value $(0.1)10 + (0.9)1 = 1.9$. They bid identically, yielding a revenue of 1.9. This is a significant improvement. But now the allocation may be inefficient with the lower valuation bidder getting the impression. In this case, the expected value of the winning bidder for the item is $(0.01)10 + (0.18)5.5 + (0.81)1 = 1.9$, much lower than before. Notice that the bundling strategy has decreased the agent's utilities, but increased revenue.

Next, consider the BIN-TAC mechanism with a BIN price of 5.5, TAC reserve of 1, and randomization parameter 2. Without loss of generality, assume the first bidder matches. If the second bidder takes the BIN option, then the first bidder will (weakly) prefer to BIN since her expected utility is 0 regardless of her choice. Additionally, if the second bidder takes the TAC option, then once again the first bidder (weakly) prefers the BIN price, since their utility when doing so is $10 - 5.5 = 4.5$, whereas if they take-a-chance, her expected utility is $(0.5)(10 - 1) = 4.5$. So if both buyers match, there will be an auction with revenue 10; if one buyer matches, the revenue will be 5.5; and if none match it will be 1. Adding this up gets an expected revenue of 1.9, as in the VA-B case. But notice that the BIN-TAC allocation is efficient, and thus agents get the same utility as in VA-T. So the BIN-TAC mechanism may improve revenue relative to the VA-T, and welfare relative to the VA-B.

3.2.4 Equilibrium Analysis

Moving back to the general environment, we characterize equilibrium strategies under BIN-TAC. We proceed by backward induction. The auctions that follow the initial BIN decision admit simple strategies. If multiple players choose to BIN, the allocation mechanism reduces to a second-price auction with reserve p . Thus, it is weakly dominant for players to bid their valuations.⁸

Truth-telling is also weakly dominant in the TAC auction. The logic is standard: if a bidder with valuation v bids $b' > v$, it can only change the allocation when the maximum of the d -th highest rival bid and the reserve price is in $[v, b']$. But whenever this occurs, the resulting price of the object is above the bidder's valuation and if she wins she will regret her decision. Alternatively, if she bids $b' < v$, when she wins the price is not affected, and her probability of winning will decrease.

Taking these strategies as given, we turn to the buy-it-now decision. Intuitively, the BIN option should be more attractive to higher types: they have the most to lose from either random allocation (they may not get the item even if they are willing to pay the most) or from rivals taking the BIN option (they certainly do not get the item). This suggests that in a symmetric Bayesian equilibrium,

⁸Since participation is obligatory at this stage, the minimum allowable bid is p ; but no bidder would take the BIN option unless they had a valuation of at least p .

the BIN decision takes a threshold form: $\exists \bar{v}$ such that types with $v \geq \bar{v}$ elect to BIN, and the rest do not. This is in fact the case.

Prior to stating a formal theorem, we introduce the following notation. Let the random variable Y^j be the j -th highest draw in an iid sample of size $n-1$ from F (i.e. the j -th highest rival valuation) and let Y^* be the maximum of Y^d and the TAC reserve r .

Equilibrium Characterization

Which type is indifferent between the BIN and TAC options? If strategies are increasing, the only time the choice is relevant is when there are no higher valuation bidders (since otherwise those bidders would BIN and win the resulting auction). So if a bidder has the highest value and chooses to BIN, they get a utility of $v - p$. Choosing to TAC gives $\frac{1}{d}E[v - Y^* | Y^1 < v]$, since they only win with probability $\frac{1}{d}$, although their payment of Y^* is on average much lower. Equating these two finds the indifferent type \bar{v} . We state this formally in the following theorem.

Theorem 17. *Assume $d > 1$ and $p \leq \frac{d-1}{d}\bar{\omega}_H + \frac{1}{d}E[Y^*]$.⁹ Then there exists a unique symmetric pure strategy Bayes-Nash equilibrium of the game, characterized by a threshold \bar{v} satisfying:*

$$\bar{v} = p + \frac{1}{d}E[\bar{v} - Y^* | Y^1 < \bar{v}] \quad (3.1)$$

Types with $v \geq \bar{v}$ take the BIN option; and all types bid their valuation in any auction that may occur.

Proof. Let a be a binary choice variable equal to 1 if the agent takes BIN and zero if TAC. Fix a player i , and fix arbitrary measurable BIN strategies $a_j(v)$ for the other players. Let q be the probability that no other agent takes the BIN option, equal to $\prod_{j \neq i} (\int 1(a_j(v) = 0) dF(v))$.

Let $\pi(a, v)$ be the expected payoff to action a for type v given that the agent bids their valuation in any auction that follows. Then we have that $\frac{\partial}{\partial v}\pi(1, v) \geq q$, as a marginal increase in type increases the payoff by the probability of winning, which is lower bounded by q when taking the BIN option. Similarly we have that $\frac{\partial}{\partial v}\pi(0, v) \leq \frac{q}{d}$, as the probability of winning when taking-a-chance is bounded above by q/d . Then $\pi(a, v)$ satisfies the strict single crossing property in (a, v) ; it follows by Theorem 4 of [76], the best response function must be strictly increasing in v , which in this case implies a threshold rule. It follows that any symmetric equilibrium must be in symmetric threshold strategies. So fix an equilibrium of the form in the theorem, and let the payoffs to taking BIN be

⁹The assumption that $p \leq \frac{d-1}{d}\bar{\omega}_H + \frac{1}{d}E[Y^*]$ rules out uninteresting cases where the BIN price is so high that no-one ever chooses BIN.

$\pi_B(v)$ and to TAC be $\pi_T(v)$. They are given by:

$$\begin{aligned}\pi_B(v) &= E [1(v > Y^1 > \bar{v})(v - Y^1)] + E [1(Y^1 < \bar{v})(v - p)] \\ \pi_T(v) &= E \left[1(Y^1 < \bar{v})1(Y^* < v) \frac{1}{d}(v - Y^*) \right]\end{aligned}$$

The threshold type \bar{v} must be indifferent, so

$$\begin{aligned}\pi_B(\bar{v}) &= E [1(Y^1 < \bar{v})(\bar{v} - p)] \\ &= E \left[1(Y^1 < \bar{v}) \frac{1}{d}(\bar{v} - Y^*) \right] = \pi_T(\bar{v}).\end{aligned}\tag{3.2}$$

Next, we show a \bar{v} satisfying Eq. (3.1) exists and is unique. The right hand side of Eq. (3.1) is a function of \bar{v} with first derivative $\frac{1}{d}(1 - \frac{\partial}{\partial \bar{v}} E[Y^* | Y^1 < \bar{v}]) < 1$. Since at $\bar{v} = 0$ it has value $p > 0$ and globally has slope less than 1, it must cross the 45° line exactly once. Thus there is exactly one solution to the implicit Eq. (3.1). \square

Optimizing BIN-TAC Parameters

Now we consider the revenue-maximizing choices of the design parameters: the BIN price p , the TAC reserve r and the randomization parameter d . It is hard to characterize the optimal d , as it is an integer programming problem which doesn't admit standard optimization approaches. However for a given d , the optimal BIN price and TAC reserve are given by some familiar looking equations. Again, we must introduce some notation. Let $p(\bar{v}, r) = \bar{v} - \frac{1}{d} E[\bar{v} - Y^* | Y^1 < \bar{v}]$ be the solution of Equation (3.1), expressing the BIN price as a function of the threshold and the TAC reserve. Let $R(\bar{v}, r)$ be the conditional expected revenue from a TAC auction when the highest valuation is exactly equal to \bar{v} and the reserve is r . Then we have the following theorem.

Theorem 18. *For any (p, d) , the revenue-maximizing TAC reserve r^* satisfies:*

$$r^* = \frac{1 - F(r^*)}{f(r^*)}\tag{3.3}$$

The optimal BIN price is given by $p(\bar{v}^, r^*)$ where \bar{v}^* is the solution of the equation below:*

$$F(\bar{v}) (p(\bar{v}, r^*) - R(\bar{v}, r^*)) + (n - 1)(1 - F(\bar{v})) (\bar{v} - p(\bar{v}, r^*)) = \frac{(1 - F(\bar{v}))F(\bar{v})}{f(\bar{v})} \frac{\partial p(\bar{v}, r^*)}{\partial \bar{v}}\tag{3.4}$$

If no such solution exists in $[\underline{\omega}_H, \bar{\omega}_H]$, then the optimal BIN price is $p(\underline{\omega}_H, r^)$.*

Equation (3.3) is somewhat surprising; the optimal TAC reserve is exactly the standard reserve in [81], ensuring that no types with negative virtual valuation are ever awarded the object. The key insight is that the TAC reserve is relevant for the BIN choice. Raising the TAC reserve lowers the

utility from participating in the TAC auction, and so the mechanism can also raise the BIN price while keeping the indifferent type \bar{v} constant. So the trade-off is exactly the usual one: raising the TAC reserve extracts revenue from types above r^* — even those above \bar{v} — at the cost of losing revenue from types below the new reserve.

Notice that the BIN price in some sense sets a reserve at \bar{v} . If two bidders meet the reserve, the mechanism gets the second highest bid; if only one, the BIN price; and if none, he gets the TAC revenue. So a marginal increase in the threshold has three effects: 1) If the highest bidder has valuation exactly equal to the threshold, following an increase she will shift from BIN to TAC, decreasing the mechanism's revenue. 2) If the second highest bidder has valuation equal to the threshold, an increase will knock her out of the BIN auction, decreasing the mechanism's revenue. 3) if the highest bidder is above the reserve and the second highest is below, the mechanism's revenue increases by a small amount. Working out the probabilities of these various events, and equating expected costs and benefits, we get the result.

Proof. By assumption, $\psi(v)$ single-crosses zero exactly once from below on $[\underline{\omega}_L, \bar{\omega}_L]$, so the implicit equation for r^* has exactly one solution. We next show that this is necessary. Fix d and $\bar{v} > p \geq r$ and define $p(r)$ implicitly as the BIN price that holds \bar{v} constant as r changes. Then there are two effects of increasing the reserve r slightly: first, you can raise the BIN price without changing \bar{v} ; second, if all bidders TAC, increasing the reserve raises the expected payment of some types, while decreasing the probability of sale. The marginal increase in revenue from BIN auctions is:

$$nF(\bar{v})^{n-1}(1 - F(\bar{v}))\frac{1}{d}\Pr[(Y^d \leq r)$$

With probability $F(\bar{v})^n$ there are no BIN bidders. Writing $F_{\bar{v}}$ for $F(v|v < \bar{v})$, revenue from the TAC auction is given by:

$$F(\bar{v})^n \frac{1}{d} \sum_{k=1}^d \left[\sum_{j=k}^d \binom{n}{j} (1 - F_{\bar{v}}(r))^j F_{\bar{v}}(r)^{n-j} r \right. \\ \left. + \int_r^{\bar{v}} \frac{n!}{d!(n-1-d)!} f_{\bar{v}}(s) F_{\bar{v}}(s)^{n-d-1} (1 - F_{\bar{v}}(s))^d ds \right]$$

Taking a first order condition in r , canceling telescoping terms and simplifying:

$$F(\bar{v})^n \frac{1}{d} \sum_{k=1}^d \binom{n}{k} k (1 - F_{\bar{v}}(r))^{k-1} F_{\bar{v}}(r)^{n-k} (1 - F_{\bar{v}}(r) - r f_{\bar{v}}(r))$$

Summing both marginal effects and expanding $P(Y^d \leq r)$:

$$n(1 - F(\bar{v})) \left(\sum_{k=0}^{d-1} \binom{n-1}{k} (1 - F_{\bar{v}}(r))^k F_{\bar{v}}(r)^{n-1-k} \right) + \\ F(\bar{v}) \sum_{k=1}^d \binom{n}{k} k (1 - F_{\bar{v}}(r))^{k-1} F_{\bar{v}}(r)^{n-k} (1 - F_{\bar{v}}(r) - r f_{\bar{v}}(r))$$

Changing summation limits, factorizing, eliminating constants and setting the equation to 0:

$$(1 - F(\bar{v})) + (1 - F_{\bar{v}}(r) - r f_{\bar{v}}(r)) F(\bar{v}) = 0$$

Now since $F_{\bar{v}} = F(v|v < \bar{v}) = F(v)/F(\bar{v})$, we can simplify and solve to get $r^* = \frac{1-F(r^*)}{f(r^*)}$.

Next, the optimal BIN price $p > r^*$ must be such that $\bar{v} \geq \underline{\omega}_H$ (only high types take the BIN option). There are three effects of a marginal increase in \bar{v} . First, the second highest bidder may have valuation \bar{v} and choose not to take BIN, which decreases revenue by $\bar{v} - p(\bar{v}, r)$. The probability of $V^2 = \bar{v}$ is given by $n(n-1)f(\bar{v})(1-F(\bar{v}))F(\bar{v})^{n-2}$. The second is that that highest bidder may have valuation \bar{v} and choose not to take BIN, reducing revenue by $p(\bar{v}, r) - R(\bar{v}, r)$. This happens with probability $n f(\bar{v}) F(\bar{v})^{n-1}$. Finally, the highest bidder may have valuation above \bar{v} and the second highest below it, which raises revenue by $\frac{\partial p(\bar{v}, r)}{\partial \bar{v}}$. This happens with probability $n(1-F(\bar{v}))F(\bar{v})^{n-1}$. Setting the sum of these effects equal to zero, evaluating the expression at r^* and eliminating common factors we get:

$$f(\bar{v}) (((n-1)(1-F(\bar{v}))(\bar{v} - p(\bar{v}, r^*)) + F(\bar{v})(p(\bar{v}, r^*) - R(\bar{v}, r^*))) = (1-F(\bar{v}))F(\bar{v}) \frac{\partial p(\bar{v}, r^*)}{\partial \bar{v}}$$

□

3.2.5 Performance Comparisons

We are interested in comparing the BIN-TAC mechanism to the benchmarks in terms of both revenue and total welfare. For any mechanism M with parameters θ , define a payoff function $\pi(M, \theta, \beta)$ as follows (suppressing the dependence on the environment):

$$\pi(M, \theta, \beta) = ER(M, \theta) + \beta ECS(M, \theta) \tag{3.5}$$

where ER denotes expected revenue and ECS expected consumer utility (i.e., the expected sum of the advertiser's utilities). Notice that when $\beta = 0$, the platform objective is just to maximize revenue as in the usual optimal mechanism design problem. Similarly, when $\beta = 1$ the objective aligns with the social planner problem of maximizing welfare. We say that mechanism M dominates M' over the interval $[a, b] \subseteq [0, 1]$ if $\max_{\theta} \pi(M, \theta, \beta) \geq \max_{\theta} \pi(M', \theta, \beta)$ for all $\beta \in [a, b]$ and for all environments

(n, α, F_L, F_H) . If M dominates M' over the whole interval $[0, 1]$ we say that M dominates M' . We say such dominance is strict if for some environment and some β the inequality holds strictly. Strict dominance means that regardless of whether the platform is maximizing revenue, joint welfare, or some combination of the two, and regardless of the environment, mechanism M is better able to achieve that objective than M' .

Theorem 19. 1. *BIN-TAC strictly dominates VA-T.*

2. *For any environment $\exists \underline{\beta} < 1$ such that BIN-TAC dominates VA-B on $(\underline{\beta}, 1]$.*

3. *In the special case when F_H and F_L are degenerate with atoms at V_H and V_L respectively, BIN-TAC strictly dominates VA-B.*

We first provide some intuition. The first result follows by showing that VA-T is just a special case of BIN-TAC, and therefore any performance achievable by VA-T is also achievable by BIN-TAC. The idea is to turn the TAC auction into a VA-T, by setting the randomization parameter $d = 1$ and the BIN-price p so high that the BIN option is never taken. The second result follows by showing that BIN-TAC is always better at achieving an efficient allocation (since the bundling solution suppresses the information needed to ensure good match outcomes), and therefore as long as the weight on consumer utility is sufficiently high, there is some interval of weights for which BIN-TAC is better.

The next question is whether BIN-TAC could be a good solution even when the platform is only interested in revenue maximization. Our last statement shows that this is possible: when the only source of private information is the match variable X , disclosing that information and then using BIN-TAC is better than withholding it and running a second-price auction.

Proof. We prove each statement in turn.

(i) Construct a BIN-TAC mechanism that achieves exactly the same outcomes as VA-T for any type realization. Let the VA-T have optimal reserve r^* , and let the BIN-TAC mechanism have TAC reserve r^* , randomization parameter $d = 1$ and BIN price $p = \bar{\omega}_H$. Then no type will take the BIN option in equilibrium (it is strictly dominated), and so the TAC auction will always occur. Since $d = 1$, this is just a VA-T with reserve r^* . Since for this particular choice of parameters BIN-TAC does as well as VA-T, in general BIN-TAC dominates VA-T. Strict dominance follows from the uniform example in the text.

- (ii) We will argue that there is an open interval $(\underline{\beta}, 1]$ on which VA-T dominates VA-B, and therefore by part (i), so does BIN-TAC. First we argue that when $\beta = 1$, VA-T does better than VA-B. Since the objective is surplus maximization, the optimal VA-T and VA-B share a common reserve of zero. Total surplus is just the valuation of the winning bidder. Under VA-T this is obviously maximized, since in the VA-T bidders bid their valuation, and the highest bidder wins. Under VA-B this needn't be maximized, since bidders bid their expected valuation, and the highest bidder wins. So whenever the bidder with the highest realized valuation did not have the highest expected valuation, the highest valuation bidder does not win. Since these events happen with positive probability, VA-T achieves higher expected payoff than VA-B for $\beta = 1$. Now both $\pi(\text{VA-T}, r, \beta)$ and $\pi(\text{VA-B}, r, \beta)$ are continuous in r and β . This holds by inspection for β , and is true for r because the type distributions are atomless. Then a triangle inequality argument suffices to extend the result at $\beta = 1$ by continuity to an interval $(\underline{\beta}, 1]$ for the functions $\max_r \pi(\text{VA-T}, r, \beta)$ and $\max_r \pi(\text{VA-B}, r, \beta)$.
- (iii) Consider a BIN-TAC mechanism with $d = n$, and a reserve of v_L . The highest BIN price that makes electing BIN optimal for high types is $p = v_H - \frac{(v_H - v_L)}{n} = \frac{n-1}{n}v_H + \frac{1}{n}v_L$. Then for any $\beta \in [0, 1]$, we have:

$$\begin{aligned}
& \max_{\theta} \pi(\text{BIN-TAC}, \theta, \beta) \\
& \geq (1 - (1 - \alpha)^n - n\alpha(1 - \alpha)^{n-1}) v_H \\
& \quad + n\alpha(1 - \alpha)^{n-1} (\beta v_H + (1 - \beta)p^{BIN}) + (1 - \alpha)^n v_L \\
& = (1 - (1 - \alpha)^n - n\alpha(1 - \alpha)^{n-1}) v_H \\
& \quad + n\alpha(1 - \alpha)^{n-1} \left(\left(1 - \frac{1 - \beta}{n}\right) v_H + \frac{1 - \beta}{n} v_L \right) + (1 - \alpha)^n v_L \\
& \geq (1 - (1 - \alpha)^n - \alpha(1 - \alpha)^{n-1}) v_H + (\alpha(1 - \alpha)^{n-1} + (1 - \alpha)^n) v_L
\end{aligned}$$

where the last step uses the fact that the minimum of the function occurs at $\beta = 0$. Under the VA-B, all types bid the same, there is no consumer surplus, and revenue is equal to:

$$\max_{\theta} \pi(\text{VA-B}, \theta, \beta) = \alpha v_H + (1 - \alpha) v_L$$

The final line of both expressions is a weighted average of v_L and v_H ; it suffices to show that the mass on v_L is lower under BIN-TAC. This requires $(\alpha(1 - \alpha)^{n-1} + (1 - \alpha)^n) < (1 - \alpha)$. After a bit of simple algebra, this is equivalent to showing $(1 - \alpha)(1 - \alpha(1 - \alpha)^{n-2} - (1 - \alpha)^{n-1}) \geq 0$, which holds by binomial expansion of 1 with equality for $n = 2$ and strictly for $n > 2$. This

proves dominance; strict dominance follows by noting that the BIN-TAC payoff is strictly higher for $\beta > 0$.

□

By definition, BIN-TAC will have (weakly) worse revenue performance than the revenue-optimal mechanism, which we refer to here as OPT. Note that by OPT we mean the optimal mechanism when targeting is allowed – it is possible (as we see below) to in fact improve revenue through bundling in certain cases. The question here is how close BIN-TAC gets to OPT in the targeting setting. We will show informally, using graphs and simulations, that it gets very close indeed. To begin, we must first describe the optimal mechanism.

Myerson Mechanism in Our Model

Consider the case when $\alpha \underline{\omega}_H < r^*(1 - F(r^*))$, where r^* is the optimal reserve of equation (3.3).¹⁰ We will define the optimal mechanism for this setting using Myerson’s construction [81]. Let $\psi(v) = v - \frac{1-F(v)}{f(v)}$ be the virtual values for our distribution. Note that, in our setting, the distribution is irregular (see Definition 14). The *ironed virtual values* (defined in [81]) are defined with respect to the virtual values as follows:

$$\phi(v) = \begin{cases} 0 & v \in [\underline{\omega}_L, r^*) \\ \psi(v) & v \in [r^*, \tilde{v}] \\ \psi(\tilde{v}) & v \in (\tilde{v}, \underline{\omega}_H) \\ \psi(v) & v \in [\underline{\omega}_H, \bar{\omega}_H], \end{cases} \quad (3.6)$$

for some \tilde{v} where $r^* \leq \tilde{v} \leq \bar{\omega}_L$. Myerson’s allocation procedure is the following: award the item to the bidder with the highest ironed virtual valuation, breaking ties at random, provided the virtual valuation is positive. Notice that all types between \tilde{v} and $\underline{\omega}_H$ get the same ironed virtual valuations; therefore, if they tie, the winner is selected at random within this range. Like BIN-TAC, it may seem that the randomization could give a suboptimal allocation, but, in fact, randomization allows the mechanism to maintain incentive compatibility and profit from higher valued bidders.

Having obtained this characterization, we can compare BIN-TAC with OPT. For now, let us focus on a simple environment, where F_L is uniform over $[0, 1]$ and F_H is uniform over $[\Delta, \Delta + 1]$. Figure 3.2 shows the interim allocation probabilities (top panel) and expected payments by type

¹⁰When this condition fails, ironing is not required and the optimal mechanism is just a Vickrey auction with reserve, which can be implemented as a BIN-TAC auction with $d = 1$.




Figure 3.2: Comparison of Allocations and Payments: Allocation probabilities (top panel) and expected payments (bottom panel) for the OPT, VA-T and BIN-TAC mechanisms when the distributions F_L and F_H are uniform. The x -axis corresponds to the bid.

(bottom panel) as a function of bidder type, in the case where $\Delta = 3$, $\alpha = 0.05$ and $n = 5$ (with optimal parameter choices). The optimal mechanism has a discontinuous jump in the allocation probability at $\tilde{v} = 0.676$, and then irons until the high valuation region on $[3, 4]$. As you can see, BIN-TAC is able to approximate the discontinuous increase in allocation probability at \tilde{v} with a smooth curve, by randomizing the allocation in that region using the TAC auction. By contrast, the slope of the VA-T allocation schedule is steep on this region and so the VA-T cannot extract revenue from the high types (who could easily pretend to be a lower type while barely changing their probability of winning). This is clear from the bottom panel.

Table 3.1 compares the expected revenue and welfare obtained by all the mechanisms. The performance of BIN-TAC is close to the optimal mechanism (about 96% of OPT), much better than the optimal VA-T (85%). The table also shows that VA-B performs less well than both BIN-TAC and OPT, especially in terms of expected consumer utility. This is because it often fails to match advertisers and users correctly.

Mechanism	OPT	SPA	BIN-TAC (d=2)	BIN-TAC (d=3)	Bundling
Expected Revenue	0.89	0.76	0.85	0.83	0.81
Expected Consumer Surplus	0.51	0.67	0.48	0.40	0.16

Table 3.1: **Revenue Comparison: Uniform Environment:** Expected revenue and welfare under different mechanisms, for the uniform environment with $\Delta = 3$, $\alpha = 0.05$ and the number of bidders $n = 5$.

3.2.6 Monte Carlo Simulations

We would like to test our mechanism against the benchmarks in a variety of other settings. We drop the assumption that F_L and F_H have disjoint support. The optimal BIN-TAC mechanism remains easy to calculate. Nothing in the proof of Theorem 3.3 required the disjoint supports for determining r^* and p^* , and so these can be solved for numerically for each d . Thus the optimization problem reduces to a one dimensional discrete optimization problem in the randomization parameter d , which can be quickly solved. Finding the optimal mechanism is more challenging, but can be done using standard optimization techniques.

For our simulations, we restrict simple distributions where $F_H(\cdot) = F_L(\cdot - \Delta)$ for some shift-parameter Δ , as in the uniform case above. This Δ is the difference in mean valuation between the high and low groups, which we call the “match increment”. We consider two cases: one where V_L (and hence H_L) is normally and another where V_L (and hence H_L) is log-normally distributed. In both cases V_L has mean 1 and standard deviation 0.5. We allow Δ , n and α to vary across experiments, and compute r^* , p^* and d^* as discussed. The default parameters we consider are $n = 10$, $\Delta = 5$, and $\alpha = .05$, and we vary one parameter at a time. Each experiment is repeated for 100000 impressions, and we calculate the average revenues.

The results are presented in Figures 3.3, 3.4 and 3.5. In all cases, on the y-axis we plot the revenue as a fraction of the revenue from the optimal mechanism. Recall that BIN-TAC generalizes VA-T, so its performance is always at least as good, and often significantly better. In all cases, the BIN-TAC extracts at least 90% of the optimal revenue, compared to a worst-case performance of around 82% for the VA-T. Consistent with Theorem 19, VA-B in some cases does even better than OPT (recall that OPT is defined in the targeting setting, so it is possible that disallowing targeting improves revenue). However, this only occurs when there are very few bidders, and the performance of VA-B degrades sharply as the probability of a match gets large, and also as the gap Δ increases.

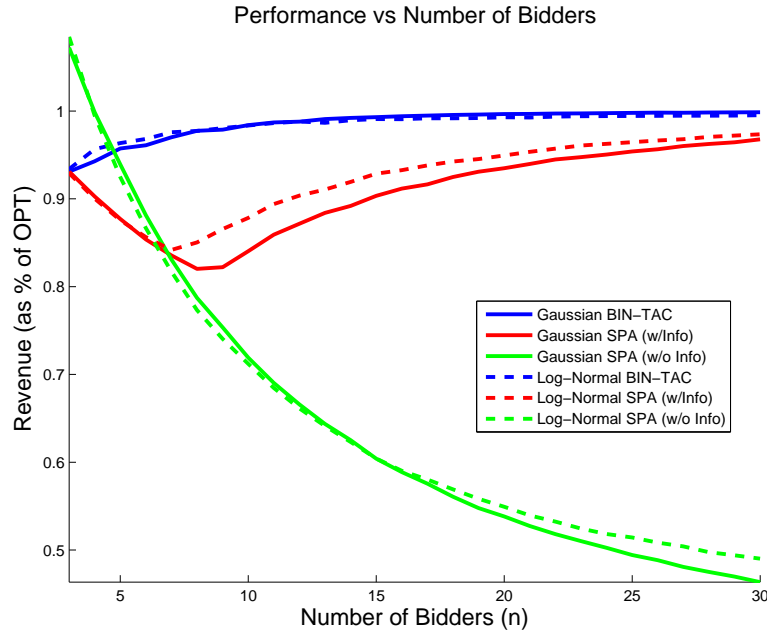


Figure 3.3: Revenue Performance vs Number of bidders: Simulated expected revenues for different mechanisms as the number of bidders n varies, in an environment where F_L has mean 1 and standard deviation 0.5, the match probability is 0.05 and the match increment is 5.

We see this in Figures 3.3 and 3.4. The expected number of matches is αn , and so as either α or n increases, the performance of the mechanisms that allow targeting improves relative to the VA-B. Over some range, BIN-TAC also significantly outperforms the VA-T, but as the number of bidders or the probability of match get sufficiently high, both converge to the OPT mechanism (which is itself a VA-T with high reserve).

Figure 3.5 shows the dependence on the gap Δ . As expected, the performance of BIN-TAC increases while that of VA-T falls as Δ gets larger, over some range. Since there is more revenue to be gained from high-valued bidders, BIN-TAC can only perform better with a large Δ . For sufficiently high Δ though, both BIN-TAC and VA-T set high reserves, “throwing away” low-valued impressions and extracting all their revenue from matches, with equal revenue performance.

Overall, the performance of BIN-TAC is very good, at least for the distributions and parameters chosen. The main caveats are that it doesn’t perform well with very few bidders (when bundling is preferable), and has little to recommend it when matches are highly probable or very valuable (a second-price auction would do as well). Its niche is in markets with relatively large numbers of bidders but low match probabilities, so that markets are “thin” in the sense of having relatively low matches in expectation.




Figure 3.4: Revenue Performance vs Match Probability: Simulated expected revenues for different mechanisms as the probability of a match α varies, where F_L has mean 1 and standard deviation 0.5, the number of bidders is 10 and the match increment is 5.

3.3 Empirical Application

Our theoretical analysis has shown that there are cases in which BIN-TAC performs well. We now test our mechanism’s performance in a real-world setting. We have historical data from Microsoft Advertising Exchange, one of the world’s leading ad exchanges. Our data comes from a single large publisher’s auctions on this exchange and consists of a 0.1% random sample of a week’s worth of auction data from this publisher, sampled within the last two years. This publisher sells multiple “products”, where a product is a URL-ad size combination (e.g. a large banner ad on the sports landing page of the New York Times). This data includes information on both the publisher and advertiser side. On the publisher side, we see the url of the webpage the ad will be posted on, the size of the advertising space and the IP address of the user browsing the website. We form a unique identifier for the url-size pair, and call that a product. We determine which US state the user IP originates from, and call that a region. We use controls for product and region throughout the descriptive regressions. Unfortunately, we don’t have more detailed information on the product or the user, as the tags and cookies passed by the publisher to the ad exchange were not stored.

On the advertiser side, we see the bid they placed, the company name, the ad broker they

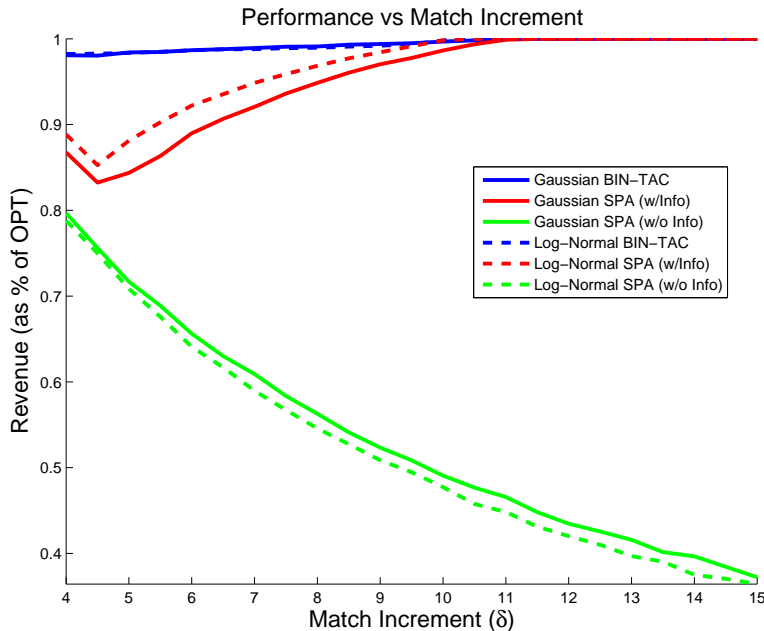


Figure 3.5: Revenue Performance vs Match Increment: Simulated expected revenues for different mechanisms as the match increment Δ varies, where F_L has mean 1 and standard deviation 0.5, the match probability is 0.05 and the number of bidders is 10.

employed, and a variable indicating the ad they intend to show. In the overwhelming majority of cases there is a single ad for each company, but some larger firms have multiple ad campaigns simultaneously. We treat these as being a single ad campaign in what follows because each firm should have the same per impression valuation across campaigns. We observe who won the auction and the final price.

We drop auctions in which the eventual allocation was determined by biased bids and modifiers.¹¹ We also restrict attention to impressions that originate in the US, and where the publisher content is in English. Finally, we restrict only to reasonably frequently sold products, those with at least 100 sales in the dataset. This leaves us with a sample of 83515 impressions.

The dataset is summarized in Table 3.2. For confidentiality reasons, bids have been rescaled so that the average bid across all observations is equal to 1 unit. Bids are very skew, with the median bid being only 0.57 units. Perhaps as a consequence of this skewness, the winning bid — which is more heavily sampled from the right tail of the bid distribution — is much higher at 2.96 units.

¹¹When the advertiser has a technologically complex kind of ad to display, their bid is modified. When the advertiser has a previously negotiated contract with the platform, their bid may be biased. This can affect the allocation and payments.

There are on average 6 bidders per auction, but there is considerable variation in participation, with a standard deviation of nearly 3. Bids are not strongly correlated: as the table shows, the correlation between a randomly selected pair of bids from each auction is only 0.01. This is not statistically significant at 5% (p-value 0.116, $N = 15827$).¹²

	Mean	Median	Std. Dev.	Min	Max
Bid-Level Data					
Average bid	1.000	0.565	2.507	0.0000157	130.7
Number of bids	508036				
Auction-Level Data					
Winning bid	2.957	1.614	5.543	0.00144	130.7
Second highest bid	1.066	0.784	1.285	0.00132	39.22
Number of bidders	6.083	6	2.970	1	15
Bid correlation	0.01				
Number of auctions	83515				
Advertiser-Level Data					
% of auctions participated in (p1)	0.697	0.0251	4.641	0.00120	88.28
% of auctions won if participated (p2)	38.90	29.59	35.50	0	100
Correlation of (p1,p2)	-0.09				

Table 3.2: Summary Statistics for AppNexus Display Ad Auctions: The full dataset is a 0.1 percent sample of a week’s worth of auction data sampled sometime within the last two years. An observation is a bid in the top panel; an auction in the middle panel; and an advertiser in the last panel. Bids have been normalized so that their average is 1 for confidentiality reasons. The bid correlation is measured by selecting a pair of bids at random in every auction with at least two bidders, and computing the correlation coefficient.

The advertisers are themselves quite active in the market. On average they bid on 0.7% of all impressions, and win nearly 40% of those they bid on. These averages are somewhat misleading though. The median advertiser is far less active, bidding on only 0.02% of impressions, while the most active advertiser participates in nearly 90% of auctions. Our hypothesis is that some advertisers choose to participate in relatively few auctions, but tend to bid quite high and therefore win with relatively high probability. Others bid lower amounts in many auctions, and win with lower

¹²That bids are not positively correlated should not be taken to mean that underlying valuations are not positively correlated; it could just be that informational and technological constraints prevent advertisers from fully expressing their preferences.




Figure 3.6: Bids over Time: The (rescaled) bids of five advertisers in our data, selected at random from the top 50 advertisers (ranked by purchases) on 50 randomly chosen successive impressions of the most popular product. Note that the set of impressions differs across bidders (there are no impressions on which all 5 participate).

probability. The first strategy is followed by companies who want to place their advertisements only on webpages with specific content or to target specific demographics, while the latter strategy is followed by companies whose main aim is brand visibility.

3.3.1 Descriptive Evidence

Before proceeding to the main estimation and simulations, we provide some evidence that advertisers bid differently on different users (i.e. there is matching on user demographics). We also show that the platform is doing poorly in extracting the benefit from matches as revenue.

Leading advertisers do vary their bids on the same product over short periods of time. Figure 3.6 shows re-scaled bids in 50 auctions by five large advertisers for the most popular webpage slot sold by this publisher. The advertisers were chosen at random from the top 50 advertisers in our dataset (ranked by purchases). The 50 auctions are chosen to be consecutive for each bidder.¹³ The bids exhibit considerable variation, even though all of these impressions were auctioned within a 3-hour period. While this could in principle be driven by decreases in the advertisers' available budget, since the bids go both up and down it seems more likely that this variation arises from matching on

¹³Since the same set of bidders don't participate in every auction, the impression number on the x-axis corresponds to different impressions for different bidders.

user demographics.

One direct test of advertiser-user matching is to look for the significance of advertiser-user fixed effects in explaining bids. Specifically, we estimate an unrestricted model where the dependent variable is bids and the controls are advertiser-user dummies, versus a restricted model with just advertiser and user fixed effects, but not their interaction. The restricted model is overwhelmingly rejected by the data: the relevant F-statistic is over 15, while the 99% critical value is just over 1. This points towards matching on demographics.

Proving that this matching is motivated by economic considerations is a little more difficult. The only user demographic we observe is the user region, and it is hard to know a priori what the advertisers' preferences over regions are. To get a handle on this, we turn to another proprietary dataset that indicates how often an advertiser's webpage was viewed by internet users in different regions of the country during the calendar month prior to the auction.¹⁴ Our intention is to proxy for the advertisers' geographic preferences (insofar as these exist) using this pageview data. The idea is that firms who operate in only a few regions probably attract all their pageviews from those regions, and also only want to advertise in those regions. If this is right, advertisers who attract a large fraction of their pageviews from a particular region should participate more frequently and bid higher on users from those regions.¹⁵ We normalize the pageviews from a particular state by the state population to get a per capita pageview measure, and then construct the fraction of normalized pageviews each region receives, calling this the "pageview ratio".

In Table 3.3, we present results from regressions of auction participation (a dummy equal to one if the advertiser participated), and bid (conditional on participation) on the pageview ratio, as well as a number of fixed effects. Because the sheer size of our dataset makes it difficult to run the fixed effect regressions, we run this on a subsample consisting of the top 10% of advertisers.¹⁶ The first column shows participation as a function of the pageview ratio, as well as product-region fixed effects, and time-of-day fixed effects (since participation and bids may vary with the user's local time). We find a positive but insignificant effect. But when we include advertiser fixed effects to control for different participation frequencies across advertisers, we find a much bigger and now highly significant effect. All else equal, an advertiser is 3.3% more likely to bid on a user from a state that contributes 10% of the population-weighted pageviews for their site. This is a large increase,

¹⁴For example, if these auctions were in May, the pageview data would be taken from April.

¹⁵Because the pageview data dates from a period before our exchange data we are not worried about reverse causality (i.e. advertisers who win more impressions from region X later get more views from region X).

¹⁶Fortunately since participation is highly skewed, these advertisers account for 90% of the bids. With only bidder fixed effects we could use a within transformation to reduce the computational burden; but unfortunately this is not possible with multiple non-interacting fixed effects.

as the average probability of participation is only around 1%.

	Participation		Bids	
Advertiser Website Pageview Ratio	0.029 (0.022)	0.329*** (0.015)	0.264*** (0.052)	0.286*** (0.053)
Time-of-Day Fixed Effects	yes	yes	yes	yes
Product-Region Fixed Effects	yes	yes	yes	yes
Advertiser Fixed Effects	no	yes	no	yes
N	5581749	5581749	417557	417557
R^2	0.02	0.34	0.04	0.26

Table 3.3: Matching on Region: Results from OLS Regressions. In the first two columns, the dependent variable is a dummy for participation. The sample used in the regressions consists of all auction-bidder pairs, limited to the 10% of bidders who participate most often. In the last two columns, the dependent variable is the bid. The sample used in the regressions only includes bids from the 10% of bidders who bid most often. The independent variable is the population-weighted fraction of pageviews of the advertiser’s website that come from the region the user is in. Time-of-day fixed effects refer to a dummy for each quarter of the day, starting at midnight. Product-region fixed effects are dummies for the page-group advertised on, and the state the user is located in. Standard errors are robust. Significance levels are denoted by asterisks (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

Turning to the bids, we find similar estimates and significance levels from the specifications with and without advertiser fixed effects. We find that firms bid higher on users from more relevant regions, although this effect is relatively modest in economic terms. Given that our proxy for advertiser preferences is relatively crude, it is notable that we find these effects. This provides some evidence that advertisers are able to target regions where their most valuable customers are.

A second stylized fact is that there is often a substantial gap between the highest and second highest bid in the auction. To facilitate bid comparisons, we look at the product with the highest sales volume in the data (over 38% of all impressions). The left panel of Figure 3.7 shows a kernel density estimate of this gap. The average bid in an auction is 0.88, while the mean gap is much larger at 1.89, indicating that there is a lot of money left on the table by a second-price mechanism (see Table 3.2 for other summary statistics). That gap itself is extremely skewed.

Assuming bids are equal to valuations — an assumption we will motivate in the next section — the right panel shows the virtual valuations $\psi(v)$ as a function of the bids. Although the virtual valuations are never infinitely negative, as in our stylized model, they are certainly non-monotone. This implies that BIN-TAC may be able to extract more revenue than a Vickrey auction. We test




Figure 3.7: Bidding Gap and Virtual Valuations: The left panel shows a kernel density estimate of the pdf of the (normalized) gap between the highest and second highest (rescaled) bids in auctions for the product with the highest sales volume in our dataset. The right panel shows the estimated virtual valuations as a function of bids.

this in the next section.

3.3.2 Estimation and Counterfactual Simulations

Our theoretical model is of a single auction with a particular valuation structure, rather than a whole market with a general valuation structure, and so in order to provide micro-foundations for our simulation approach, we need to enrich the model.

We make the following assumptions for the estimation and counterfactual simulations. There is a fixed set of N bidders who are always present in the market.¹⁷ As in the text, the model is symmetric

¹⁷The assumption that bidders are continuously present in the market is in principle relatively innocuous since bidding is done by ad broker algorithms. Yet some bidding algorithms ignore certain auctions in order to respect

independent private values. Each bidder draws their valuations for each impression identically, independently and privately according to some distribution F_j supported on $[0, \infty)$ (where j indexes products). So bidder valuations are independent both across bidders and within a bidder over time. This is a strong assumption, as it rules out common preferences for particular user demographics. For example, it rules out the possibility that all bidders prefer high income bidders, in which case we would observe positive correlation in bids. Some partial support for this assumption comes from the lack of bid correlation reported in Table 3.2. The symmetry assumption is also strong — and probably rejected by the data given the significance of the advertiser fixed effects in the reduced form regressions — but helps to keep the problem computationally tractable. To address the concern that the symmetry and independence assumptions are driving our results, we will do some robustness tests based on different informational assumptions in a later subsection.

From the summary statistics we also know that participation varies across advertisers. We assume that participation costs are zero, and thus we can infer from non-participation that an advertiser has zero valuation for the impression. This may seem like a strong assumption, but given that the 5th percentile of bids in our data is equal to 0.013 — tiny in real terms, with an almost zero probability of winning, and even lower utility — it is hard to believe that participation costs are substantially different from zero. One reason for this may be that bidding is automated.

Given these assumptions, we are able to make the following inference from the second-price auction data. Let $i = 1 \dots I$ index bidders and $t = 1 \dots T$ index auctions, if bidder i makes a bid of $b_{i,t}$ in auction t , their valuation is $b_{i,t}$, since it is weakly dominant for them to bid their valuations. Moreover, if bidder i did not participate in auction t , their valuation for that particular impression must have been zero. Since there is a one-to-one mapping from the distribution of bids and participation to the valuations, F_j is non-parametrically identified. We could therefore estimate the valuation density for each product using non-parametric methods. But, as we will show below, the counterfactual simulations will never require estimating more than some conditional moments of order statistics (e.g. the expected value of the d -th highest valuation when the highest valuation is less than \bar{v}). So instead we estimate these moments by the corresponding sample average.

We are interested in comparing the “optimal” BIN-TAC mechanism to other leading mechanisms. For simplicity, we restrict attention throughout to the class of mechanisms that make the same parameter choices for all products (e.g. we rule out different reserves or randomization parameters by product or user-region). In each case we find these optimal parameters by maximizing the revenue

advertiser budget constraints. We will not model this “inattention”, especially because it is hard to rationalize such behavior as optimal: bidding close to zero has almost no effect on the budget constraint since the maximum possible payment in a second-price auction is bounded above by the bid.

functions defined in equations (3.7) and (3.8) below, using standard optimization methods.¹⁸ To get standard errors on our revenue and consumer utility estimates, we bootstrap the estimation sample and re-run the simulation procedure, holding the parameter choices fixed.¹⁹

Mechanisms with Targeting:

The policies we want to compare here are the Vickrey auction with targeting, and BIN-TAC. We will consider two different VA-T mechanisms for comparison. The first will be the mechanism with no reserve (i.e., with reserve 0), and the second with an optimal reserve. With reserve of r , the expected revenue of VA-T depends on the joint distribution of the top two valuations: since bidders bid their valuations, the item sells if the highest valuation exceeds r , and then the revenue is the maximum of the second highest bid and r . Let the k -th highest bid in an auction t be $b_t^{(k)}$, our estimate is then given by the sample average across the T auctions:

$$\text{Revenue}^{\text{VA-T}}(r) = \frac{1}{T} \sum_{t=1}^T \mathbf{1}(b_t^{(1)} > r) \max\{b_t^{(2)}, r\} \quad (3.7)$$

BIN-TAC is harder, as an agent's equilibrium decision to take the BIN option depends on their beliefs about the distribution of rival valuations. From the model, advertiser behavior is characterized by a threshold value $\bar{v}_j = \bar{v}_j(p, d, r)$ for each product, above which they will take the BIN option, and below which they will TAC. From Theorem 17, this threshold solves the implicit equation $\bar{v}_j - p = \frac{1}{d} E[\bar{v} - Y^* | Y^1 < \bar{v}_j]$, where $Y^* = \max\{Y^d, r\}$ and Y^1 and Y^d are the 1st and d -th order statistics of rival bids on product j . To solve this equation for fixed (p, d, r) , we need to estimate the expected TAC payment $E[Y^* | Y^1 < s]$ for varying s .

Under symmetry, the joint distribution of valuations is exchangeable, and so the joint distribution of rival bids is exactly the same as the joint distribution of $N - 1$ randomly selected bids. So our estimate of the TAC payment conditional on winning on product j is given by:

$$\text{TAC Payment}(s, r) = \frac{1}{T} \sum_{t=1}^T \frac{\sum_k \mathbf{1}(b_t^{(1)} < s) \max\{b_t^{(d)}, r\}}{\sum_k \mathbf{1}(b_t^{(1)} < s)}$$

where k indexes the N choices of $N - 1$ -length bid vectors for each auction, including zeros for bidders that didn't participate and restricting the sample only to product j .²⁰ We can then solve for

¹⁸This raises an over-fitting concern, in that the parameters are optimized for this specific realization of the data generating process. However given our sample size, the bias this introduces is likely to be small.

¹⁹We use 100 bootstrap samples (i.e. samples of T impressions drawn randomly with replacement).

²⁰It is correct to include the non-participating bidders, as in principle all N bidders are present in every period and so the distribution of rival bids drops only one of them — probably a bidder who would not have participated in any case.

the equilibrium $\bar{v}(p, d, r)$ for each set of BIN-randomization parameters (p, d, r) , and get a revenue estimate as follows:

$$\begin{aligned} \text{Revenue}^{\text{BIN-TAC}}(p, d, r) &= \frac{1}{T} \sum_{t=1}^T \mathbb{1}(b_t^{(2)} \geq \bar{v}(p, d, r)) b_t^{(2)} + \frac{1}{T} \sum_{t=1}^T \mathbb{1}(b_t^{(1)} \geq \bar{v}(p, d, r) > b_t^{(2)}) p \\ &+ \frac{1}{T} \sum_{t=1}^T \mathbb{1}(b_t^{(1)} < \bar{v}(p, d, r)) \sum_{j=1}^d \mathbb{1}(b_t^{(j)} \geq r) \max\{b^{(d+1)}, r\} \end{aligned} \quad (3.8)$$

Bundling Mechanisms:

As we do not observe all the impression characteristics provided to advertisers in this market, we cannot consider the optimal bundling strategy. But we can consider bundling by product and user region, where the platform strips away all other user characteristics except for the region, so that advertisers are buying a random impression of a given size, on a given website, being viewed by a user from a particular US state. This is unlikely to be optimal, but provides a lower bound on the revenues from the bundling strategy.

For this analysis, we allow for bidder valuations to be asymmetric and vary by both product and region. Our estimate of a bidder’s willingness to pay for this “generic impression” is just their average bid across all auctions of this product-region combination, taking their implicit bids when they didn’t participate as equal to zero. Given that participation costs are zero and all bidders have strictly positive mean valuations, in the counterfactual world all bidders will participate in all auctions. We assume that these impressions are sold by second-price auction without reserve (since the bundling creates thick markets, a reserve isn’t necessary).

Robustness to Informational Assumptions:

The above theory and structural estimation follows the empirical auctions literature in treating bidder’s valuations as private information.²¹ A different modeling approach was suggested in an influential chapter by [47]. They proposed a complete information model of sponsored search auctions. Their logic was that since these players compete with high frequency and can potentially learn each others’ valuations, a complete information model may be a better approximation to reality than an incomplete information model.

Following this intuition, we also consider counterfactual simulations under complete information. The only model this affects is the BIN-TAC model, as under weak refinements the VA equilibria under

²¹See for example [63]. See also [6] for a model of sponsored search models in this tradition.

incomplete and complete information coincide. However in the BIN-TAC model we unfortunately now have multiple equilibria.²²

To see this, consider a case where the bidder with the highest valuation is going to take the BIN option. Then the remaining bidders are indifferent between BIN and TAC, since in either case they will lose the auction and get payoff 0. We employ a trembling hand perfection refinement (see [?]) to eliminate this multiplicity. Specifically, for any probability $\epsilon > 0$ that the highest bidder will take the TAC option instead, the second-highest bidder faces a non-trivial choice between BIN and TAC. Applying this logic restores a generically unique equilibrium prediction. We can therefore solve for the unique trembling hand perfect equilibrium of each auction, and estimate the expected revenues from the average sample revenues at any parameter vector.

We also perform a worst-case analysis over beliefs about rival strategies. From the point of view of revenue, the worst-case for BIN-TAC occurs when agents are least inclined to take the BIN option: specifically, when they believe that all other agents will choose to TAC and then bid zero. This implies that incentives to take the BIN option must be provided directly by the design, through the randomization parameter d and the reserve price r in the TAC auction. Since these beliefs are identical across all auctions, we can compute the indifference threshold $\bar{v}(p, d, r)$ implied by these beliefs, and then calculate revenue in exactly the same way as in the incomplete information case.

3.3.3 Results

The results are in Tables 3.4 and 3.6. We find that the optimal reserve when running a Vickrey auction is high: nearly twice as high as the second highest bid. By contrast, BIN-TAC always uses relatively low reserves (all well below the average bid), and instead threatens to randomize among 3-4 bidders in order to get agents to take the high buy price (which is close in magnitude to the optimal VA-T reserve). Interestingly, in the worst-case scenario the platform has to threaten randomization among 4 agents to get bidders to take BIN, since bidder beliefs are such that the TAC auction looks relatively attractive.

The welfare performance of these mechanisms is detailed in Table 3.6. The VA-T without reserve earns revenue of 0.98 per auction, and does not capitalize on the increase in social welfare — on average 1.97 per auction.²³ Adding the large optimal reserve improves revenue slightly (to 1.03 per auction), but hurts social welfare substantially (it falls to 1.44).

²²This arises also in the generalized Vickrey auction — see [47] and [99].

²³The per auction revenue of 0.98 is lower than the average second highest bid of 1.07 in Table 3.2 because of a small fraction (2.3%) of auctions with only a single bidder, which will realize zero revenue in a VA-T without reserve.

Table 3.4: Optimal Parameter Choices

Policy	p	d	r
VA-T	-	-	1.96
BIN-TAC	2.60	3	0.43
BIN-TAC (complete information)	1.95	3	0.65
BIN-TAC (rationalizable worst case)	2.10	4	0.65

Table 3.5: Revenue-maximizing parameter choices. For each of the above mechanisms, we find these by maximizing the revenue functions defined in the main text over the available parameters numerically using a grid search. The two modified BIN-TAC mechanisms are robustness checks, varying the informational assumptions made for BIN-TAC. In the complete information case, bidders know the valuations of the other participants, and made BIN decisions accordingly. In the rationalizable worst-case model, bidders assume they will only have to pay the reserve price in TAC auction, and therefore take the BIN option more rarely.

BIN-TAC does better than both of these mechanisms in terms of revenue. Interestingly, the social welfare is higher than under VA-T and reserve, implying BIN-TAC dominates VA-T in terms of both revenue and social welfare. This happens because the optimal VA-T reserve price is very high — to extract revenue from the long right tail — and so many impressions are not sold, resulting in inefficiency and lower total welfare. By contrast, BIN-TAC has the BIN price to extract this revenue, and so the reserve is much lower, and more impressions are sold. Even accounting for distortions owing to the TAC auction, this is a welfare improvement.

By contrast, the bundling strategy underperforms. Revenues are much lower in VA-B, and social welfare falls even more dramatically. This is because there is considerable variation in the utility of matched impressions even after conditioning on product and region, and so bundling along only these two dimensions destroys a lot of social welfare.

Finally, the BIN-TAC results in the bottom part of the table show that the revenue estimates are relatively robust to how we model the information structure. However in models where the bidders are more informed, or dubious about the BIN option, social welfare is lower. In those cases the BIN decision is taken less often, thereby increasing the distortion from TAC auctions.

3.4 Conclusion and Future Work

We have introduced the BIN-TAC mechanism, designed to allow publishers to capture the increase in social welfare created by providing match information. This mechanism outperforms the second-

Table 3.6: Counterfactual Revenues and Welfare

Policy	Revenue	Social Welfare	Total Welfare
VA-T (no reserve)	0.983 (0.004)	1.974 (0.019)	2.957 (0.019)
VA-T (optimal reserve)	1.028 (0.005)	1.471 (0.018)	2.499 (0.020)
BIN-TAC	1.075 (0.005)	1.633 (0.018)	2.708 (0.020)
VA-B (bundling by product-region)	0.644 (0.006)	0.730 (0.016)	1.374 (0.015)
Robustness to Informational Assumptions			
BIN-TAC (complete information)	1.072 (0.005)	1.589 (0.018)	2.661 (0.020)
BIN-TAC (rationalizable worst case)	1.066 (0.005)	1.530 (0.018)	2.596 (0.020)

Table 3.7: Counterfactual simulations of average advertiser revenues, social welfare and total welfare (sum of social welfare and revenue). All statistics reported outside parentheses are averages across impressions; those in parentheses are standard errors computed by bootstrapping the full dataset (i.e. they reflect uncertainty over the true DGP). Six different simulations are run. First, a Vickrey auction without reserve. Then, a Vickrey auction with optimal (revenue-maximizing) reserve. Third, the BIN-TAC mechanism under incomplete information – the primary model considered in the text. Fourth, a Vickrey auction where the impressions are bundled according to the product and user region. The last two simulations are robustness checks, varying the informational assumptions made for BIN-TAC. In the complete information case, bidders know the valuations of the other participants, and made BIN decisions accordingly. In the rationalizable worst-case model, bidders assume they will only have to pay the reserve price in TAC auction, and therefore take the BIN option more rarely. Where applicable, the parameters used are the optimal parameters from Table 3.4.

price auction mechanism in this setting, and is preferable to bundling items together by withholding information, at least when there is a reasonable size population of potential bidders. Moreover, we demonstrated that the mechanism can closely approximate Myerson's optimal mechanism with ironing, despite its relative simplicity.

Our analysis of the exchange marketplace revealed that it has many features that make it a good place to apply our mechanism: large differences between the highest and second highest bid, and evidence of matching on user characteristics that the platform has chosen to make available to advertisers. Although the market does not fit our stylized model, we found that the BIN-TAC mechanism would nonetheless improve revenues and social welfare relative to the existing mechanism, a Vickrey auction with reserve.

Due to data limitations we were not able to compare our mechanism to an optimal bundling strategy. Instead, we looked at what would happen if the platform only provided advertisers with product and user location information, rather than more detailed demographics. This bundling strategy performed poorly, but it is an interesting and open research question as to whether switching mechanisms to BIN-TAC is in fact better than retaining VA-B with a more clever bundling strategy.

Chapter 4

Information Acquisition: Mechanism Design for Deliberative Agents

The previous chapter showed that due to the influx of information created in modern e-commerce, we must be willing to adjust commonly held assumptions about agents in order to develop good mechanisms. In this chapter we focus on a different aspect; what if the information is not held by the agent *or* the mechanism, but can be acquired from some third party. In the recent past, new companies, whose goal is to aggregate information about consumers for the purpose of selling it to advertisers, have sprung up.¹ This opens up a whole new strategy space for bidders where their valuation for an item can be influenced based on the information they receive. The added complication is that this information always comes at some cost.

Other settings, outside of the ad world, also lend themselves to a model where agents acquire information. It is reasonable to assume that agents do not have perfect information about their values, but can pay some cost (either in money, time or effort) to attain more certainty. For example, suppose you are making a decision about buying a house. You will do a lot of research about the area and schools, pay for a real estate agent who can provide more information about the house's value, and possibly pay for an inspections of the wiring, plumbing or grounds. Each such action has some cost, but also helps better evaluate the worth of the house. In some cases, i.e. if you find out the school district is the worst in the country, you may simply walk away. However, if the school district is reasonable, you may choose to proceed further, taking more actions (at more cost)

¹Well-known companies that offer this kind of service are BlueKai, eXelate and OwnerIQ.

and gaining even more information until you have refined your prior to the point where you are comfortable making an offer. We introduce a general model for agents that have information-buying capabilities, and call them *deliberative* agents.

In particular, we consider the seller’s perspective; i.e., how should the house be sold given that agents are deliberative? In this chapter we focus solely on deterministic dominant strategy mechanisms. Our results are of two kinds. First, we show that in several settings, revenue can be maximized with a sequential posted price mechanism, where one agent after another is offered a take-it-or-leave-it price. Additionally, we show that for a simple class of deliberations, approximately revenue-optimal SPPs can be constructed efficiently. Our approximation result follows from a reduction to the classical setting.

We consider *single-parameter* environments (see Definition 1) where an auctioneer is offering goods or a service, and, based on the outcome of a mechanism, chooses a subset of the agents to be served (we call them *winners*) and a price p_i for each winning agent i (see Definition 2). Depending on the particular setting, only certain subsets of agents can be feasibly served. Throughout, we will assume that the environment is common knowledge to the agents and to the auctioneer.²

4.1 Deliberative Model

In general deliberative environments, the mechanism is an iterative process and agents may acquire information throughout their interaction with the mechanism. We consider single-parameter environments, but now, an agent *may not know her own value* v_i ! Rather, they have the opportunity to *learn* about their valuation at some *cost*.

Definition 20 (Deliberative Agent). *A deliberative agent is represented by a tuple $(\mathcal{F}, \mathcal{D}, c)$, where*

- $v \sim \mathcal{F}$ is the agent’s value. This value is unknown to the agent.
- \mathcal{D} is a set of deliberations the agent can perform. This set includes the trivial deliberation $\emptyset(\mathcal{F}) = \mathcal{F}$.
- $c : \mathcal{D} \rightarrow \mathbb{R}_+$, where $c(d)$ represents the cost to perform the deliberation d , the value $c(\emptyset) = 0$, and $c(d) > 0$ for all $d \neq \emptyset$.

Each deliberation is a (potentially random) function between agent types and gives a new instantiation $(\mathcal{F}', \mathcal{D}', c')$. A deliberation $d \in \mathcal{D}$ yields a new prior for the agent’s value $d(\mathcal{F}) = \mathcal{F}'$ (and similarly

²This chapter is based on joint work with Dimitrios Gklezakos, Anna Karlin, Kevin Leighton-Brown, Thach Nguyen, and David Thompson [24, 25].

a new deliberation set \mathcal{D}' and costs c'). The distribution over new priors is such that the marginals agree with \mathcal{F} , i.e. $v \sim d(\mathcal{F})$ is identically distributed as $v \sim \mathcal{F}$ (though $v \sim \mathcal{F}'$ may be different); one can think of d as providing some information which the agent uses to update her prior.

The agent after deliberations d_1, \dots, d_t can be represented by $(\mathcal{F}^{(t)}, \mathcal{D}^{(t)}, c^{(t)})$.

We can now formally define a deliberative environment, which is an analogue of the original single-parameter setting.

Definition 21 (Deliberative Environment). A deliberative environment is a tuple $(A; \mathcal{S})$ where

- A is a set of deliberative agents,
- a set \mathcal{S} of feasible subsets of A , which represent the subsets of agents that can be served simultaneously.

For notational ease, we let $\mathcal{F} = \mathcal{F}_1 \times \dots \times \mathcal{F}_n$ be the joint distribution of agents' values. We think of each agent as having a single value $v_i \sim \mathcal{F}_i$ which parameterizes her benefit from being served by the mechanism, however, this value is unknown, and the agent's estimate of it may change over time (with deliberation).

Definition 22 (Mechanisms in Deliberative Environments). A deliberative mechanism $\mathcal{M} = (x, p)$ is a multi-stage process in which at each step the mechanism

- solicits input from a subset of agents,
- reveals information to a subset of agents, and
- outputs a partial allocation and payments x_i, p_i on a subset of agents.

Note that the subsets may differ within a time step. We assume that mechanisms only charge agents they serve, though agents may end up with negative utility since they may pay for a deliberation and then not be served.

Upon receiving a request from the mechanism, the agent may decide to perform some deliberations and respond. To define the utility of a deliberative agent, we make two important observations. First, the cost of deliberations obviously factor into the agent's utility. Second, since the agents may not know their own valuations even after participating in a mechanism, they may not know their real utility. Therefore, a more relevant notion is the agents' *expected* utility given their information. Thus, in this setting, rational agents act to maximize their expected utility.

Definition 23 (Expected Utility for Deliberative Agents). *Suppose an agent performs deliberations d_1, d_2, \dots, d_t during the course of her interaction with the mechanism. If the mechanism serves her at a price p then her expected utility is*

$$\mathbb{E}_{v \sim \mathcal{F}^{(t)}}[v] - p - \sum_{i=1}^t c(d_i).$$

If the mechanism does not serve the agent then her perceived utility is

$$- \sum_{i=1}^t c(d_i)$$

4.1.1 Special Types of Deliberative Mechanisms

We will consider two specific kinds of mechanism in this chapter, namely, dynamically direct and sequential posted price mechanisms. Dynamically direct mechanisms are a special case of general deliberative mechanisms and are introduced, as we will see later, because they allow us to state a deliberative form of the revelation principle.

Definition 24 (Dynamically Direct Mechanism (DDM)). *A dynamically direct mechanism is a multi-stage mechanism where at each stage either*

- *a single agent is asked to perform and report the result of a specified deliberation,*
- *reveals information to a subset of agents (as discussed in Section 4.1.2),*
- *outputs an allocation and payment rule.*

Thus, a DDM interacts with a single agent at a time, and restricts interactions to soliciting and receiving the results of their deliberations.

Sequential posted price mechanisms, where agents are offered take-it-or-leave-it prices in turn, are well studied in the classical setting.

Definition 25 (Sequential Posted Price (SPP)). *A sequential posted price mechanism \mathcal{M} is a multi-stage process which at each step,*

- *a single agent is offered a take-it-or-leave-it price p .*

The mechanism is committed to sell to the agent at the offered price if she accepts, and will not serve the agent now or at any future point if she rejects.

SPPs are much simpler than DDMs because an agents' decision is made in isolation; the agent has a concrete offer she can accept or reject without concern as to how other agents may affect her price or service in the future. Additionally, the mechanism never tells her how to act, she decides which, if any, deliberations to make based on the offered price. Hence, SPPs are inherently truthful.

Note that an SPP is not a DDM as defined above. However, we will show that SPPs are robust enough to capture the complexity of DDMs.

4.1.2 Information Models

Recall that in the definition of a DDM we give above, we allow the mechanism to reveal information to the agents. Consider the case where the mechanism reveals no information. In this case, agents do not know the history of the mechanism and hence their strategies are independent of the history. We call this the *minimal-information model*. The *full-information model* is when *all* interactions are observable (or revealed), and hence agents can condition their strategy accordingly.

One can think of the minimal-information model as corresponding to mechanisms that employ *private* communication; i.e., other agents are not privy to the interaction between an agent and the mechanism. However, it is possible that the mechanism cannot control this information; perhaps agents can observe the mechanism's interactions with other agents directly. The full-information model corresponds to public-communication; i.e., every agent observes the interaction between any other agent and the mechanism.

There is, naturally, a spectrum of information models that lie in between these two extremes, and many of our results do not depend on the particular information model. However, the environments for which SPPs maximize revenue differ depending on the information model. This is natural; the full-information model allows the agents to have more complex strategies. This, in turn, restricts the mechanism since it must afford a dominant strategy equilibria. Hence, the class of mechanisms in the minimal-information setting is larger (and includes) the mechanisms in the full-information setting. Thus, it is not surprising that our characterization results are stronger in the full-information setting.³

4.1.3 Related Work:

Previous work shows that auction design for deliberative agents is fundamentally different from classical auction design due to the greater flexibility in the agents' strategies. In classical mechanism

³In fact, any single-parameter environment is revenue-maximized by an SPP in the full-information setting, while we can only show that single-item auctions are revenue-maximized by SPPs in the minimal-information setting. This is stated formally in Theorem 28.

design, an agent only has to decide how much information to reveal. In deliberative-agent mechanism design, an agent first has to decide how much information to acquire and then how much to reveal. This affects the equilibrium behavior. For example, in second-price auctions, deliberative agents do not have dominant strategies [65] and must coordinate their information gathering [98]. Furthermore, the standard revelation principle, which asserts that every multi-stage auction is equivalent to some sealed-bid auction, no longer holds.

There has been considerable interest in designing novel auctions for deliberative agents. This research has mostly focused on maximizing social welfare subject to various constraints [13, 21, 64], with some research on revenue maximization in Bayes-Nash equilibrium [19, 43]. More recently, [97] demonstrated that dominant strategy revenue optimization is possible via SPPs. However, their result only holds for single-item auctions in the minimal-information model with binary-valued agents who have a single deliberative action. Our contributions here generalize a) the types of agents, b) the types of environments, and c) the types of information models.

4.2 *Revelation Principle for DDMs*

The traditional revelation principle does not hold for deliberative environments. Since the deliberations performed by an agent at each stage may depend on the information they receive, and the mechanism cannot simulate this information gathering on behalf of the agent, strategies cannot be “flattened” to work with a single-stage mechanism. For example, in a Japanese auction, bidders can condition their information gathering on information revealed at earlier stages, coordinating in ways that are not possible in sealed-bid auctions [37].

However, we can recover a version of the revelation principle that states we can consider only truthful dynamically direct mechanisms (DDMs). When interacting with a DDM, a *strategy* of a deliberative agent consists of either

1. not deliberating and reporting a result \hat{r} , or
2. performing the requested deliberation and reporting a result \hat{r} , which may depend on the real result r of the deliberation,
3. or performing other or additional deliberation(s) and reporting a result \hat{r} , which may depend on the real result(s) r of the deliberation(s).

A *truthful strategy* is to do (2) and report $\hat{r} = r$, the real result of the deliberation. A *truthful DDM* is one in which the truthful strategy is a dominant strategy for every agent.

Versions of the revelation principle have been shown for DDMs in the past (see [65,97]), however we state and prove it in full generality here with no restriction on the agents, information model, or use of randomness as was the case in prior work.

Theorem 26 (Revelation principle). *For any deliberative mechanism \mathcal{M} and some equilibrium σ of \mathcal{M} , there exists a truthful DDM \mathcal{N} which implements the same outcome as \mathcal{M} .*

Note that we do not specify which type of equilibria we require (though we focus exclusively on dominant-strategy equilibria in this chapter); truthfulness in \mathcal{N} is an equilibrium the same way that σ is an equilibrium in \mathcal{M} (i.e. if \mathcal{M} has a Bayes-Nash equilibria, then in \mathcal{N} truthfulness is a Bayes-Nash equilibria). Additionally, we do not specify the information model; the information model in \mathcal{N} is the same as in \mathcal{M} . These facts are implicit and obvious from the proof below.

The proof is similar to that of the classical revelation principle and prior deliberative models; \mathcal{N} simulates both \mathcal{M} and the agents' dominant strategies.

Proof. Let σ be the desired equilibrium for \mathcal{M} where σ_i is the strategy of bidder i . Note that \mathcal{N} can simulate the behavior of a bidder *except* for her actual deliberation. Let \mathcal{N} proceed from time step to time step as does \mathcal{M} simulating both \mathcal{M} and σ 's behavior. Specifically, at each time step, we first “solicit” input from the corresponding *simulated* agents. We can simulate perfectly until the point where the strategy demands a deliberation. If this occurs, \mathcal{N} asks the appropriate bidder to perform the deliberation.⁴ Once we have completed the solicitation step from \mathcal{M} we “reveal” information to our *simulated* agents, and output a partial allocation as in \mathcal{M} .

Clearly we have constructed a DDM. Recall that the agents only control their response to the deliberation. However, by either not deliberating or reporting untruthfully, they are using a different strategy σ'_i overall. Since σ is an equilibrium in \mathcal{M} , reporting truthfully is an equilibrium in \mathcal{N} .

Figure 4.2 illustrates this argument. □

As with the original revelation principle, this allows us to restrict our search to truthful dominant-strategy DDMs.

4.3 Characterization

In this section, we prove that SPPs are revenue-optimal dominant strategy mechanisms for bounded deliberative agents.

⁴If this occurs synchronously for several bidders in \mathcal{M} , we can simply ask sequentially in \mathcal{N} with no effect.

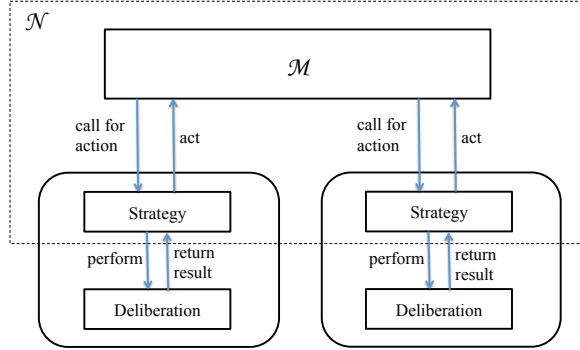


Figure 4.1: A “proof by picture” of the deliberative revelation principle. The truthful DDM \mathcal{N} simulates both the dominant strategy mechanism \mathcal{M} and the dominant strategies of the agents.

Definition 27 (Bounded Deliberative Agent). *A bounded deliberative agent i is represented by a tuple $(\mathcal{F}_i, \mathcal{D}_i, c_i)$ as in Definition 20, with the additional restrictions:*

- *Every prior has bounded expectation, i.e. $\mathbb{E}[\mathcal{F}_i^{(t)}] < \infty$ for all i, t .*
- *Every set of deliberative actions is finite, i.e. $|D_i^{(t)}| < \infty$ for all i, t .*
- *There is some finite T such that $D_i^{(t)} = \{\emptyset\}$ for all $t \geq T$.*

In other words, there are finitely many deliberation paths, and along any such path the current prior always has bounded expectation. For the remainder of this chapter we assume all agents are bounded.

The fact that SPPs have dominant strategies is straightforward and was given in [97]. Intuitively, the fact that all SPPs have a dominant strategy equilibrium follows from the fact that an agent is indifferent to what happened before she received an offer (values are independent and private) and indifferent to what happens after she rejects (she has no externalities). The argument that SPPs characterize dominant strategy mechanisms is far more involved, and the precise statement is given in the following theorem.

Theorem 28 (Characterization). *If agents are bounded, any dominant strategy mechanism \mathcal{M} in single-parameter deliberative environments is revenue-dominated by a deterministic SPP \mathcal{N} when in*

- *single-item environment, or a*

- *full-information setting.*

Specifically, for these two settings, any optimal mechanism can be transformed into a revenue-equivalent SPP.

Proof. The proof of this theorem follows from a sequence of lemmas, the outline is as follows:

1. By the revelation principle, any mechanism and corresponding equilibrium is outcome-equivalent to a truthful DDM (Theorem 26).
2. Any deterministic dominant-strategy truthful DDM is revenue-dominated by a DDM where the price an agent is charged, if served, is determined at the time she is *first* consulted for
 - single-item environments, and
 - full-information environments

(Lemma 31).

3. Any such DDM is outcome-equivalent to an SPP (Lemma 31).

□

We proved Lemma 26 above. We now complete the remaining parts of the proof, which first require some supporting lemmas.

4.3.1 Preliminary Lemmas

As in classical settings, any truthful mechanism must satisfy a monotonicity condition. Notice, however, that in a deliberative setting we do not know the *exact* values \mathbf{v}_{-i} of the other agents. However, after a given *instance* (i.e., after a run) of the mechanism, we know a set of *effective values* $\hat{\mathbf{v}}$, namely the expected values of agents at the time the mechanism ends. We use effective values in the statement of our monotonicity result.

Lemma 29 (Monotonicity (Deliberative)). *In any truthful DDM, for any agent i there exists a “critical value” $q_i(\hat{\mathbf{v}}_{-i})$ where if $\hat{v}_i > q_i(\hat{\mathbf{v}}_{-i})$ then i is served and pays exactly $q_i(\hat{\mathbf{v}}_{-i})$, otherwise she is not served and pays nothing.*

The proof follows as in the classical setting [79]. For characterization via SPPs, we need a stronger result. Namely, that the price of an agent depends only on the values of the agents which deliberated before her, i.e., the information the mechanism has when she is asked to deliberate. [97] proved such a result for binary-valued agents using an ‘‘Influence Lemma’’, which states that if an agent with two possible values deliberates, then she must be served when she reports the higher value and not served when she reports the lower. We generalize that lemma to show that such high and low values exist.

Lemma 30 (Generalized Influence Lemma). *Consider a truthful DDM \mathcal{M} , and a bounded agent i who has performed $t - 1$ deliberations. Let $S = \text{support}(\mathcal{F}_i^{(t)})$. If i is asked to perform a nontrivial deliberation d , then there exist thresholds $L_d, H_d \in S$ such that*

- *If an agent deliberates and reports value $\leq L$ she will not be served,*
- *If an agent deliberates and reports value $\geq H$ she will be served.*
- *It is possible that the deliberation results in a value $\leq L$ or $\geq H$.*

Recall that the value we refer to here is the effective value $\mathbb{E}[\mathcal{F}_i^{(t+1)}]$ where $\mathcal{F}_i^{(t+1)} \sim d(\mathcal{F}_i^{(t)})$.

Proof. Recall, from Lemma 29, that for an agent i and set of effective values \hat{v}_{-i} , there exists a threshold price $q(\hat{v}_{-i})$ such that i is served if she is above this threshold, and not served if she is below. Let $Q = \{q(\hat{v}_{-i}) : i \text{ is asked to take deliberation } d\}$, i.e., Q is the set of all such thresholds conditioned on the fact that i is asked to take deliberation d . We show that $L = \inf(Q)$ and $H = \sup(Q)$ satisfy the lemma’s conditions. With some abuse of notation, we will use v to denote i ’s effective value *right before* she is asked to take deliberation d , and \hat{v} to denote i ’s effective value *right after* she takes the deliberation.

First, we ensure L and H are well defined. Since a mechanism only chooses nonnegative prices, L is well-defined. However, we must prove this for H . Assume for sake of contradiction, that Q is not bounded above, i.e. $q \in Q$ can be arbitrarily large. Note that when the other agents are such that i ’s threshold price is q , the agent’s utility if she deliberates is

$$u_d(q) = \mathbb{E}[\hat{v} - q | \hat{v} \geq q] \Pr[\hat{v} \geq q] - c.$$

Since $v < \infty$ and the marginals of \hat{v} must align with the probability distribution of v , for large enough q we have

$$\mathbb{E}[\hat{v} - q | \hat{v} > q] \Pr[\hat{v} > q] < c.$$

Thus, for large enough q , the agent's utility $u_d(q) < 0$ (since $c > 0$), which contradicts truthfulness. Hence Q must be bounded above and $H < \infty$.

Now, let $R = \{\hat{v}\}$; i.e., R is the set of possible *effective values* after deliberation. We first show that there is some effective value $\hat{v} \geq H$; i.e., some effective value which is above all thresholds and hence, for which we are always served. If R is unbounded, then there is an outcome of our deliberation such that our effective value \hat{v} is greater than H . Now, let us assume that R is bounded and let $r = \sup(R)$. We will now show that $H \leq r - c$, and thus, since $c > 0$, there is some effective value above H . Assume for sake of contradiction that there is some threshold price $q > r - c$. Then,

$$\begin{aligned} u_d(q) &= (\mathbb{E}[\hat{v}|\hat{v} \geq q] - q)\Pr(\hat{v} \geq q) - c \\ &< \mathbb{E}[\hat{v}|\hat{v} \geq q] - q - c \\ &< r - q - c < 0 \end{aligned}$$

since $\hat{v} \leq r$ by definition. This contradicts truthfulness, so there exists some $\hat{v} > H$.

We now show that there is some effective value $\hat{v} \leq L$. To begin, we show that $L \geq s = \inf(R)$. Assume, for sake of contradiction, that there is some effective value $q < s$. Note that $P[\hat{v} < q] = 0$ since there is no $\hat{v} < s$ by definition. Thus,

$$\begin{aligned} v - q &= (\mathbb{E}[\hat{v}|\hat{v} \geq q] - q)\Pr[\hat{v} \geq q] + (\mathbb{E}[\hat{v}|\hat{v} < q] - q)\Pr[\hat{v} < q] \\ &= (\mathbb{E}[\hat{v}|\hat{v} \geq q] - q)\Pr[\hat{v} \geq q] \\ &> (\mathbb{E}[\hat{v}|\hat{v} \geq q] - q)\Pr[\hat{v} \geq q] - c. \end{aligned}$$

Note that the left hand side of this equation is the agent's utility if she *does not* deliberate, yet reports a value above H (hence forcing that mechanism to serve her at price q), while the right hand side is the agent's utility if she deliberates. This contradicts truthfulness. Thus, $L \geq s$. If $s \in R$ (i.e. $\min(R)$ is well defined), then $s \leq L$ is an effective value. If $s \notin R$, then the proof above holds when $q = s$ (not just $q < s$), and proves that $L > s$ (not just $L \geq s$). Therefore, there is some effective value $s < v \leq L$. \square

4.3.2 Main Result

We can now prove our main characterization lemma which states that the price an agent is charged, if served, is completely determined upon her first interaction with the mechanism. Specifically, this means the price does not depend on the future potential actions of any agent, including any future interaction she may have with the mechanism.

Lemma 31. *Let \mathcal{M} be a dominant-strategy truthful DDM for bounded deliberative agents. If \mathcal{M} asks an agent i to deliberate, we can modify \mathcal{M} such that the price it charges i if served only depends on the history before \mathcal{M} 's first interaction with i . This modification preserves truthfulness and can only improve revenue whenever we are in a*

- *single-item environment, or a*
- *full-information setting.*

These conditions are in fact necessary for our result to hold, and we show some examples to this effect in Section 4.3.3.

Proof. Let \mathcal{M} be a DDM as above. We can think of \mathcal{M} as being given by a tree where at each node the mechanism asks an agent to deliberate, and each child corresponds to a responses from the agent. With some abuse of notation, we say an agent reports a value in H if her value is above H_d , reports a value in L if her value is below L_d , and reports a value in M otherwise. The mechanism, of course, is told the precise outcome of her deliberation, however considering the above categories suffices for our proof.

Consider some instance of \mathcal{M} in which agent i is asked to deliberate. Let h denotes the *path* of the instance. Let h_i^b represent the part of the path before the *last time* agent i is asked to deliberate, and h_i^a the path after. Let d be the deliberative action she is asked to take. We will show that we can modify \mathcal{M} without any loss to revenue such that her price, if served, does not depend on h_i^a . We will then apply this result inductively to the next-to last time i is asked to deliberate, and so on until we reach our desired result.

Base Case (last time i deliberates):

$i \in H$. Consider the case where, after deliberation, i 's effective value is in H . Hence, by Lemma 30, i must be served at some price, say p_i . Note that since \mathcal{M} is truthful, these prices cannot depend on what i reports. Now, consider our two cases:

- **single item:** The mechanism can ask no one else to deliberate since it must allocate to i , and would not be able to allocate to another j reporting value H_j . Additionally, the mechanism is deterministic. Hence, there is nothing on which the mechanism can condition for which it charges i different prices in this case. Thus, when i reports a value in H , she is always charged $\hat{p}_i = p_i$.

- **full-information:** We will assume, for sake of contradiction, an alternate price p'_i may be reached by some path $h_i^{\mathbf{a}'}$ (see Figure 4.2(a)). There exist strategies σ_{-i} and σ'_{-i} such that paths $h_i^{\mathbf{a}}$ and $h_i^{\mathbf{a}'}$ respectively are always followed. More generally, there is a strategy $\hat{\sigma}_{-i}$ such that conditioned on this history $h_i^{\mathbf{b}}$ if i reports H , the mechanism always charges i $\hat{p}_i = \max_{h_i^{\mathbf{a}}} \{p_i\}$. Modify \mathcal{M} such that i is charged \hat{p}_i for all $h_i^{\mathbf{a}}$.

This change can only increase \mathcal{M} 's revenue. Additionally, this change does not affect any agent other than i and hence does not affect their strategies. Furthermore, since i 's price was increased when she reports a value in H , the mechanism is still truthful when i is in L or M . We now need only ensure i is still truthful when her value is in H .

Recall that in the full-information setting, agents may condition their strategy on the observed history at the time they are asked to deliberate. Hence, the full-information setting decouples strategies when i reports H from strategies when i reports another value, say M . Namely, for every possible outcome $h_i^{\mathbf{a}}$ when $i \in M$, there is a strategy $\hat{\sigma}_{-i}$ which implements outcome \hat{p}_i when $i \in H$. Hence, since \mathcal{M} is a dominant-strategy truthful mechanism, i will not have incentive to lie against $\hat{\sigma}_{-i}$. Therefore, i will also not have incentive to lie in the modified mechanism when $i \in H$.⁵


$i \in M$. Now, consider the case where i 's effective value is in M . If i is never served, there is nothing to show. Assume there is at least one path $h_i^{\mathbf{a}}$ such that i is served at some price p_i (see Figure 4.2(b)). There exists a strategy σ_{-i} which implements outcome $h_i^{\mathbf{a}}$ when $i \in M$.

From Lemma 30, we know there is at least one possible effective value above H , hence, \hat{p}_i is defined by \mathcal{M} as above for any σ_{-i} . Clearly, $p_i \leq \hat{p}_i$, otherwise even if $i \in M$, she prefers to report H and attain a better price with service guaranteed. Additionally, if $p_i < \hat{p}_i$, then against σ_{-i} , agent i prefers to report M when her value is in H , contradicting dominant-strategy truthfulness. Therefore, $p_i = \hat{p}_H$ for all such p_i .

Therefore, at the *last* time i is asked to deliberate, there is a fixed price \hat{p}_i which, if served, she will be charged.

Inductive Step: Consider a node u where i is asked to deliberate. Assume that for any node v in the subtree of \mathcal{M} rooted at u where i is asked to deliberate, if served, i is charged \hat{p}_v . This is our inductive hypothesis.

⁵This difference in strategies according to information setting is what allows this result. In the minimal information setting, such a change may not preserve truthfulness when $i \in H$.



(a) If i has an effective value $\geq H$, she is always charged the same price.

(b) If i has an effective value in M and is served, she is charged the same price as if she had an effective value $\geq H$.

Figure 4.2: Two main cases in the proof of Lemma 31.

Note that we will only ask i to deliberate again if she reports some value in M (otherwise \mathcal{M} has already committed to either serve or not serve i). Hence, by the same proof as in the base case, for $i \in H$ we can modify \mathcal{M} such that $i \in H$ is always charged \hat{p} .

Now, consider the case where $i \in M$. Consider some h_i^a where i is not asked to deliberate again and i is served at some price p' . By the same proof as in the base case, $p' = \hat{p}$. Finally, consider the some h_i^a such that i is asked to deliberate again, and let σ_{-i} be a set of strategies that reach this point. By induction, we know that the price i is charged, if served, is fixed. Call this price \hat{p}' . Thus, if $\hat{p} < \hat{p}'$, i prefers to report H and be served at price \hat{p} . If instead $\hat{p}' < \hat{p}$, then against σ_{-i} , when $i \in H$, she in fact prefers to first report $i \in M$ the first time she is asked to deliberate and then report $i \in H'$ the second time she is asked to deliberate – this guarantees her the item at the smaller price \hat{p}' . Therefore, since \mathcal{M} is dominant-strategy truthful, it must be the case that $\hat{p}' = \hat{p}$.

Thus, by induction, the price i is charged, if served, is fixed before the *first* time she is asked to deliberate. \square

Lemma 32. *Let \mathcal{M} be a DDM such that the price it charges i , assuming i is served, only depends on the history before \mathcal{M} 's first interaction with i . Then, \mathcal{M} is revenue-equivalent to an SPP \mathcal{N} .*

Proof. Given \mathcal{M} , we construct \mathcal{N} . Consider any deliberation node in the DMM \mathcal{M} , and let h_i^b be the history leading up to this node. Without loss of generality, let us refer to the agent asked to deliberate at this node as i . Recall from Lemma 31 that from this point forward, whenever agent i

is served, she is charged this same price. Call this price p_i . Let \mathcal{N} , under the same h_i^b , offer price p_i to agent i .

Since \mathcal{M} is truthful, each deliberation she asks i to perform is in her best interest. Additionally, there is no deliberation i would like to perform which the mechanism did not ask her to perform. Let this final effective value be \hat{v}_i . Clearly, by truthfulness, i gets the item if and only if $p_i \leq \hat{v}_i$. Similarly, when \mathcal{N} offers her price p_i it will be in her interest to take the same sequence of deliberations, otherwise \mathcal{M} was not truthful to begin with. She accepts if and only if $p_i \leq \hat{v}_i$, matching the scenario under which she accepts \mathcal{M} 's offer.

Thus, for a fixed history, the outcome of \mathcal{M} and \mathcal{N} is the same for this agent. This is true of any deliberative node. Hence, \mathcal{M} and \mathcal{N} have the same expected revenue. \square

4.3.3 Discussion and Examples

The above completes our characterization results for revenue-optimal dominant-strategy mechanisms. Note that the conditions in both cases are in fact necessary, as examples below will show. That there would be a difference in what is characterizable in the minimal- and full-information settings may not be obvious up front. However, note that the strategy space for an agent is much wider for the full-information setting – an agent can behave differently depending on the history. Thus, dominant strategies are much stronger since they must protect against more options. This, in turn, restricts the set of dominant strategy mechanisms, making the characterization easier in a way.

An important fact to come back to at this point is that deliberative settings are fundamentally different than classical settings – they are not a generalization. Additionally, one is not necessarily better than the other. Consider the following examples:

Example 33 (Revenue Higher in Deliberative Environment). *Consider a single-item auction with a single agent whose value is uniformly distributed on $[0, 1]$. She can deliberate to learn her exact value at cost 1. The optimal mechanism in the classical environment sells at price 0.5 and the offer is accepted with probability 0.5; hence the expected profit of the mechanism is 0.25. On the other hand, in the deliberative setting the agent would take a posted price of 0.5, the dominant strategy for the agent in the deliberative setting is to accept it without deliberating. Hence, the revenue of the optimal DDM in the deliberative setting is 0.5.*

Example 34 (Revenue Higher in Classical Environment). *Consider a single-item auction with two agents whose values are uniformly distributed on $[0, 1]$ where either can deliberate to learn their exact value at a small cost $\epsilon > 0$. The optimal mechanism in the classical environment is a second-price*

auction with reserve 0.5 and the expected profit is $5/12$. On the other hand, in the deliberative setting, a simple optimization exercise shows that the optimal SPP offers one agent the price of $3/4$ (which is accepted with probability $1/4$), and if rejected offers the next agent a price of $1/2$ (which, if asked, is accepted with probability $1/2$) since for ϵ small enough, agents will deliberate at these prices. Hence, the revenue of the expected optimal revenue of this optimal DDM in the deliberative setting is $3/8$, lower than in the classical setting.

The differences comes down to the amount of information available and the volatility created by the information asymmetries.

The following two examples show where the characterization result breaks down if a condition is not met, and in doing so give some intuition about the consequences of the different strategy spaces.

Non-Optimal DDMs are not SPPs

The following example is a non-optimal but truthful DDM in the full-information model. It is easy to see this is not a SPP since the price A pays depends on B 's answer. However, this DDM is also not optimal; A 's price when $B = 4$ can be set to $4 - 2c$ in all cases, allowing more revenue.

Example 35. *The deliberative environment in this example contains 2 items and 3 agents A, B and C . Each agent's value is drawn from a distribution over three values $\{1, 2, 4\}$ such that $\Pr[1] = \Pr[2] = \Pr[4] = 1/3$. All of the agents have deliberative costs equal to a small constant c . Then the DDM \mathcal{M} depicted in Figure 4.3 and described below is not an SPP, but also not optimal.*

1. Ask A to deliberate

(a) If $v_A = 4$ then ask B to deliberate

i. If $v_B = 4$ then sell the two items to A and B at price $4 - 2c$.

ii. If $v_B < 4$ then sell one item to A at price 2, and offer the other item to C at price 2 (without deliberation).

(b) If $v_A < 4$ then offer one item to each of the other two agents at price 2 (without deliberation).

Minimal-Information DDMs can outperform SPPs

The mechanism shown here is dominant-strategy truthful in the minimum-information setting, but not dominant-strategy truthful in the full-information setting. Specifically, in the full-information




Figure 4.3: A non-optimal deliberative mechanism that is not an SPP.

setting, if B has the strategy that she will report 4 whenever A reports 2 and report 2 whenever A reports 4, then A has incentive to always report 4 without deliberation. However, B 's strategy is not admissible in the full-information case – she must take a decision independently of A 's report. Thus, there are examples that distinguish these two information models. More significantly, this mechanism outperforms the optimal SPP for this environment, showing that minimal-information DDMs can strictly outperform full-information DDMs (when not in the single-item case).⁶

Example 36. *The deliberative environment in this example contains 2 items and 3 agents A , B and C . Each agent's value is drawn from a distribution over three values $\{1, 2, 4\}$ such that $\Pr[1] = \Pr[2] = \Pr[4] = 1/3$. All of the agents have deliberative costs equal to a small constant c . The following DDM \mathcal{M} , also depicted in Figure 4.4, does not have a dominant strategy for A . This mechanism outperforms the optimal SPP.*

1. Ask A to deliberate

(a) If $v_A = 4$ then ask B to deliberate

i. If $v_B = 4$ then sell the two items to A and B at price $4 - 3c$.

ii. If $v_B < 4$ then sell one item to A at price 2, and offer the other item to C at price 2.

(b) If $v_A = 2$ then ask B to deliberate

i. If $v_B = 4$ then sell one item to B at price $4 - 3c$ and offer the other item to C at price 2.

ii. If $v_B < 4$ then sell one item to A at price 2 and offer the other item to C at price 2.

⁶Additionally, this mechanism is Bayes-Nash truthful in either information setting. Thus, the fact that SPPs maximize revenue is dependent on the fact that we only consider dominant strategy mechanisms.




Figure 4.4: A mechanism that is dominant-strategy truthful in the minimal-information setting, and Bayes-Nash truthful in any information setting, but not dominant-strategy truthful in the full-information setting. This mechanism is not an SPP, and in fact attains more revenue than the optimal SPP.

(c) If $v_A = 1$ then offer one item to each of the other two agents at price 2.

4.4 Approximate Revenue Maximization

Our characterization result shows that SPPs optimize revenue for either full-information or single-item mechanisms. Additionally, as shown in Example 36, our result relies on the fact that we consider only dominant strategy mechanisms. In this chapter we make no restriction on the kind of mechanisms we allow. Our goal is to show that in general settings, even though SPPs may not optimize revenue, they still approximate the optimal deliberative mechanism. We succeed in showing this for simple deliberative agents and certain matroid environments with no restriction to the information model or equilibrium type.

Additionally, a significant advantage this approximation result is that it is constructive. While the optimal SPP can be found via dynamic programming, the size of the program is often exponential. The fact that this can be computationally difficult is not surprising since an SPP must define an ordering over agents. The approximation algorithm, on the other hand, uses a reduction to the classical environment where ordering is not a concern. Additionally, this reduction could allow us to apply existing results in the classical setting. In particular, here we focus on approximation via SPPs, however this reduction may be useful in a more general sense.

Definition 37 (Simple Deliberative Agent). *A simple deliberative agent i is a bounded deliberative agent that has exactly one nontrivial deliberative option; i.e., $D_i = \{\emptyset, d_i\}$ for all i and $D'_i = \{\emptyset\}$.*

Thus, after a non-trivial deliberation d , an agent has an effective value $\mathbb{E}_{d(\mathcal{F}_i)}[v_i]$ (often denoted \widehat{v}_i), and can gather no additional information. A natural example is a deliberation that reveals an agent's value exactly, as considered in related literature [13,97]. Hence, an agent wants an item priced at less than \widehat{v}_i , and does not want an item priced any higher. Thus, after a non-trivial deliberation, we can represent an agent by her expected value. Similarly, if an agent never deliberates, her value is $\mathbb{E}_{\mathcal{F}_i}[v_i]$.

Our approximation result is shown via a transformation from a deliberative environment to a related classical environment. We show that the expected revenue of the optimal mechanism in the classical environment upper bounds the optimal revenue in the deliberative setting. This allows us to use near-optimal SPPs in the classical environments to obtain near-optimal SPPs in the deliberative setting.

For ease of presentation, for the remainder of this chapter we assume the agents' distributions are continuous, positive on some interval, and that for each distribution F_i , the probability density function f_i exists.

4.4.1 Upper Bound in Classical Environment

Again, by the revelation principle, we can restrict our search to truthful DDMs. We first show that for agent i there exists $[\ell_i, h_i]$ such that if i is asked to deliberate, the price she is charged lies within this interval.

Lemma 38. *For any deliberative agent i there are two threshold $\ell_i \leq h_i$ such that*

- *if offered a price $p \leq \ell_i$, i prefers to take it without deliberation,*
- *if offered a price $p \geq h_i$, i prefers to reject the offer without deliberation,*
- *if offered a price $p \in (\ell_i, h_i)$, i prefers to deliberate and accept the offer only if $v_i \geq p$.*

When the above inequalities hold with equality, the preferences are weak. We name ℓ_i and h_i the low and high deliberation thresholds of i respectively.⁷

Proof. By monotonicity, agent i is effectively faced with a posted price offer of p , where $p = t_i(\widehat{v}_{-i})$. Faced with this offer, she has three possible strategies: she can accept the offer without deliberation, reject it without deliberation, deliberate and accept if her expected value is greater than p . Denote

⁷Note that ℓ_i and h_i are not explicitly related to L_i and H_i ; the former specify an agent's preference, the latter a mechanism's.

the expected utility of these three strategies $u_i^a(p)$, $u_i^r(p)$ and $u_i^d(p)$. (See Figure 4.5.) With some abuse of notation, we let v_i be the effective value *after* deliberation (in this case, we can think of it as an explicit value since no further deliberations are possible). Clearly,

- $u_i^a(p) = \mathbb{E}[v_i] - p$.
- $u_i^d(p) = \mathbb{E}[v_i - p | v_i \geq p] \Pr[v_i \geq p] - c_i$.
- $u_i^r(p) = 0$.

Let ℓ_i denote the price where at which i is indifferent between accepting and deliberating, and let h_i denote the price where she is indifferent between deliberating and rejecting. We show that $\ell_i \leq h_i$ satisfy the lemma's conditions.

Observe that if $\mathbb{E}[v_i] \leq c_i$, then $u_i^d(p) \leq \mathbb{E}[v_i | v_i \geq p] \Pr[v_i \geq p] + \mathbb{E}[v_i | v_i < p] \Pr[v_i < p] - c_i \leq 0$ for any p . Thus, i never deliberates, and $\ell_i = h_i = v_i$. Now, let us assume $\mathbb{E}[v_i] > c_i$.

First, consider the function $\alpha(p) = u_i^a(p) - u_i^d(p)$. Note that $\alpha(0) = c > 0$. Additionally, since we assumed continuous distributions (denoted here by f), we can rewrite α as

$$\begin{aligned} \alpha(p) &= \int_0^\infty (v_i - p) f(v_i) dv_i - \int_p^\infty (v_i - p) f(v_i) dv_i + c_i \\ &= \int_0^p (v_i - p) f(v_i) dv_i + c_i. \end{aligned}$$

Thus, it is clear that the derivative of $\alpha(p)$ is negative, so $\alpha(p)$ is strictly decreasing. Additionally, let $\beta(p) = u_i^r(p) - u_i^d(p)$. Note that $\beta(0) = -\mathbb{E}[v_i] + c < 0$. Additionally, as above, we can rewrite

$$\beta(p) = - \int_p^\infty (v_i - p) f(v_i) dv_i + c.$$

Thus, the derivative of $\beta(p)$ is positive so β is strictly increasing.

Hence, since α is decreasing and starts above 0 and β is increasing and starts above 0, there are unique points t_1 and t_2 such that $u_i^a(p) - u_i^d(p) \geq 0$ if and only if $p \leq t_1$, and $u_i^r(p) - u_i^d(p) \geq 0$ if and only if $p \geq t_2$.

If $t_1 \leq t_2$ then it is easy to check that $\ell_i = t_1$ and $h_i = t_2$ satisfy the lemma's conditions. Otherwise, we must have $t_2 \leq \mathbb{E}[v_i]$ since $u_i^a(t_2) \geq u_i^d(t_2) = u_i^r(t_2)$, and $t_1 \geq \mathbb{E}[v_i]$ since $u_i^a(t_1) = u_i^d(t_1) \leq u_i^r(t_1)$. In this case, $\ell_i = h_i = \mathbb{E}[v_i]$ satisfy the lemma's conditions.

Hence, these values satisfy the lemma's conditions: below ℓ_i we always prefer to accept without deliberation, above h_i we always prefer to reject without deliberation, and in-between we prefer to deliberate. \square




Figure 4.5: The relationships of u_i^a , u_i^d and u_i^r in Lemma 30. (a) The case where $t_1 > t_2$. (b) The case where $t_1 \leq t_2$.

Definition 39. Consider an agent i with value v_i and low and high deliberation thresholds ℓ_i and h_i respectively. The representative value of i is v'_i , defined by

$$v'_i = \begin{cases} \ell_i & \text{if } v_i \leq \ell_i \\ v_i & \text{if } v_i \in (\ell_i, h_i) \\ h_i & \text{if } v_i \geq h_i. \end{cases}$$

The classical agent with value v'_i is the representative of agent i . We will use i' to denote the representative of i , and use v'_i to denote the representative value of i , i.e., the value of i' .

Definition 40. Let i be a deliberative agent with values drawn from a distribution F_i with low and high deliberation thresholds ℓ_i and h_i . Then the representative distribution of i is defined over the interval $[\ell_i, h_i]$ by $G_i(x) = F_i(x)$ for all $\ell_i \leq x < h_i$ and $G_i(h_i) = 1$. Hence, G_i has a point mass of $F_i(\ell_i)$ at ℓ_i and a point mass of $1 - F_i(h_i)$ at h_i . Furthermore, if $\ell_i = h_i$ then the support of G_i contains a single value.

We say that $E' = (\mathcal{A}', \mathcal{S}')$ is the representative environment of $E = (\mathcal{A}, \mathcal{S})$ if E' is obtained from E by replacing each agent i with value v_i drawn from F_i by an agent i' with value v'_i drawn from G_i (which induces a corresponding set of agents \mathcal{A}' and feasible subsets \mathcal{S}'). The following theorem relates the expected revenue of truthful mechanisms in the two environments.

Theorem 41. For any truthful DDM \mathcal{M} in E , there is mechanism \mathcal{N} where

1. \mathcal{N} is a truthful (in expectation) mechanism in E' and
2. the expected revenue of \mathcal{M} is at most the expected revenue of \mathcal{N} .

Before proving Theorem 41, we introduce some notation and definitions. For a deliberative agent i whose value comes from a distribution F_i and whose low and high deliberation thresholds are ℓ_i and h_i respectively, we denote by F_i^ℓ the distribution over $[0, \ell]$ defined by $F_i^\ell(0) = 0$ and $F_i^\ell(x) = F_i(x)/F_i(\ell_i)$ for all $x \in (0, \ell]$. Similarly, F_i^h denote the distribution over $[h, \infty)$ defined by $F_i^h(h) = 0$ and $F_i^h(x) = (F_i(x) - F_i(h))/(1 - F_i(h))$.

Definition 42. Let $\mathbf{v}' \in \text{support } \mathcal{G}$ be a valuation profile of the representative agents in E' . We say that \mathbf{v} is an originator of \mathbf{v}' if

- $v'_i = \ell_i$ implies $v_i \leq \ell_i$,
- $v'_i = h_i$ implies $v_i \geq h_i$, and
- $v'_i \in (\ell_i, h_i)$ implies $v_i = v'_i$.

Given a valuation profile \mathbf{v}' of the representative agents, we can construct a random originator \mathbf{v} of \mathbf{v}' by setting (i) $v_i = v'_i$ if $v'_i \in (\ell_i, h_i)$, (ii) v_i equal to a random number drawn from F_i^ℓ if $v'_i = \ell_i$, and (iii) v_i equal to a random number drawn from F_i^h if $v'_i = h_i$. A random originator constructed this way is called a *sampled originator* of \mathbf{v}' .

Proof of Theorem 41. Given a truthful DDM \mathcal{M} on E , we construct the following mechanism \mathcal{N} on E' :

1. Solicit a bid vector \mathbf{b} from the agents.
2. Construct a sampled originator \mathbf{u} of \mathbf{b} .
3. Run \mathcal{M} on (E, \mathbf{u}) .

First, we show that \mathcal{N} is truthful in expectation. Consider a representative $i' \in N'$ with value $v_{i'}$. We show that submitting $b_{i'} = v_{i'}$ yields the best expected utility for i' . To this end, we show that this is the case for any fixed \mathbf{u}_{-i} . Once \mathbf{u}_{-i} is fixed, there are three cases regarding whether i is served.

1. \mathcal{M} does not ask i to deliberate and does not serve her. In this case, i' is not served by \mathcal{N} and her bid does not matter, therefore she is truthful.

2. \mathcal{M} asks i to deliberate. Then by Lemma 29, there is a threshold $t(\mathbf{u}_{-i})$ that does not depend on b_i such that i is served if and only if $u_i \geq t(\mathbf{u}_{-i})$. Moreover, $t(\mathbf{u}_{-i})$ is the price of i , hence i' , if she is served. By Lemma 38, we have $t(\mathbf{u}_{-i}) \in (\ell_i, h_i)$. If $v_{i'} \geq t(\mathbf{u}_{-i})$ then i' prefers to buy at this price, and bidding truthfully makes sure that this happens. On the other hand, if $v_{i'} < t(\mathbf{u}_{-i})$ then i' prefer to not buy, and bidding truthfully also makes sure this outcome is chosen.
3. \mathcal{M} does not ask i to deliberate but serves i and charges her some price p_i , which is independent of i 's value. Since \mathcal{M} is a truthful DDM, by Lemma 38, we must have $p_i \leq \ell_i \leq v_{i'}$, therefore i' prefers buying to withdrawing from the mechanism, hence bidding truthfully.

The second part of the theorem follows from the fact that if \mathbf{b} is randomly drawn from \mathcal{G} then \mathbf{u} is a random draw from \mathcal{F} . □

As an immediate corollary of Theorem 41, we get an upper bound on the revenue of all truthful DDMs on E .

Corollary 43. *For any truthful DDMs \mathcal{M} , the expected revenue of \mathcal{M} in E is the expected revenue of the optimal revenue-maximizing auction in E' .*

4.4.2 Approximation

Corollary 43 suggests that in order to approximate the expected revenue of the optimal DDM in a deliberative environment, we can design a truthful DDM that approximates the optimal auction in their representative environment.

For this we apply the following theorem of Chawla et. al. [33], which shows how, in classical single parameter environments, to obtain a constant factor approximation of the optimal auction [79] with an SPP.

Theorem 44 ([33]). *There are SPPs that approximate the expected revenue of the optimal mechanism in various classical environments. In particular,*

- *For any general matroid environment, there is an SPP whose expected revenue is at least $1/2$ of the optimal expected revenue.*
- *For any uniform matroid or partition matroid environment, there is a SPP whose expected revenue is at least $(e - 1)/e$ times the optimal expected revenue.*

- *For any environment whose feasible set system is the intersection of two matroids, there is a SPP whose expected revenue is at least $1/3$ times the optimal expected revenue.*

Moreover, it is immediate that any SPP outputs the same outcome on a instance (E, \mathbf{v}) of a deliberative environment E and the representative instance (E', \mathbf{v}') . This gives us our approximation result:

Corollary 45. *There are SPPs that approximate the expected revenue of the optimal truthful DDMs in the single-parameter settings of Theorem 44. These settings include multi-unit auctions, single-minded combinatorial auctions, and many other natural settings.*

4.5 Future Work

We view these results as a first step towards understanding deliberative environments. While the characterization of optimal deliberative mechanisms by SPPs shows that deliberative settings are fundamentally different than their classical counterparts, our approximation results show that they are related, at least in this simple deliberative setting. The future is open for constructive optimal or approximation mechanisms in for general deliberative agents.

There are still numerous open problems concerning revenue maximization in dominant strategies. As the examples show, (non-single-item) minimal-information and Bayes-Nash mechanisms are not revenue-dominated by SPPs. Characterizing optimal mechanisms in this setting remains open.

Additionally, although there are social-welfare optimizing mechanisms [13, 21] they rely on restricted environments (e.g., single-good auctions or models where all deliberations must happen simultaneously) and are only Bayes-Nash incentive compatible, not dominant-strategy truthful. It may be possible to use existing SPP results in classical settings [20] to get approximate SPPs for deliberative environments, but not using the proof techniques used here.

Lastly, the general deliberative model with no restriction on the number of non-trivial deliberations performed is still untouched. Here, the question of defining a meaningful class of deliberations is an important one. One potential model allows for agents to choose from a variety of noisy deliberations, trading off accuracy against cost. Another model allows for agents to do multiple stages of deliberation, for example, getting tighter and tighter bounds on their true value. Still another allows for the possibility of one agent deliberating about another agent's value (so called "strategic-deliberation"). Nothing is known about dominant-strategy mechanism design in these settings.

Information Exchange: Convergence to Balanced Outcomes on Networks

In a network bargaining game, nodes in a graph are involved in pairwise transactions with their neighbors. Each node in the graph is a player, and the weight on an edge represents the dollar amount available to be shared between the two adjacent players *if* they choose to form a partnership. A partnership guarantees the players this income, however, they must reach an agreement as to how to split their profit. Additionally, each player is constrained to make at most *one* partnership.

We consider two notions of equilibrium in this game. An outcome is *stable* if no two adjacent players have incentive to deviate from their current matches. The surplus is the amount the players make from forming an agreement minus the amount they would each make by partnering with someone else. An outcome is *balanced* if it is stable *and* all matched players divide the surplus equally between themselves. This can be seen as a generalization of Nash's bargaining solution for two players [82] where the outside options are given by the network itself.

In this paper, we consider the dynamics via which players may bargain in order to reach agreement. Specifically, consider a situation where at each time step, matched players *balance* their agreement by splitting their surplus. Since the surplus is a function of a player's neighbors, this process is dynamic and not guaranteed to converge.

We consider a specific type of edge balancing dynamics and show (exponential time) convergence on all graphs. The convergence point is a balanced outcome whenever a balanced outcome exists. We then give the first polynomial time bound on local dynamics converging to a balanced outcome

for a class of graphs which includes unweighted trees. The only polynomial time bound known perviously was for paths. We show how this proof can be generalized to other graphs and balancing dynamics, and show certain cases where the bound is tight.¹

History of Bargaining Games

Network bargaining games have a long history in two communities: sociology and game theory. In sociology, they are studied under the name *network exchange theory*, where the goal is to understand the power of a node as a function of its position in the network (see the overview by Willer [100]). Network bargaining games as we define here were introduced by Cook and Yamagishi [39], who also introduced the notion of balanced outcomes. In fact, they also introduced local dynamics similar to what we consider in this paper, but without a theoretical analysis of the convergence of their dynamics. There have also been experimental results [28,38] which validate the relevance and applicability of this work.

In game theory, the study of bargaining can be traced back to Nash's bargaining solution [82]. Many results in this field focus on two-sided markets, which naturally give rise to the bipartite version of the network bargaining game as was introduced by Shapley and Shubik [94]. This version, known as the *assignment game*, can also be viewed as the classic Gale-Shapley stable marriage problem [50] with the addition of transferable utilities. Rochford [90] defined balanced outcomes for assignment games under the name symmetrically pairwise-bargained allocations. She also showed that they are the intersection of the core and the kernel, two common solution concepts in cooperative game theory. Other solution concepts such as the nucleolus [72] have also been considered. In fact, the computability of these solution concepts has been much studied [85,96]. Other related models consider price setting as a result of a bargaining process [40]. Faigle, Kern and Kuipers [49] considered similar dynamics for a more general class of games, but do not show bounds on the rate of convergence.

Balanced outcomes can be computed by centralized polynomial time algorithms [62], but the game is by nature distributed; individual players working on individual deals. An important open problem was to show there exist simple and natural *local dynamics* that converge quickly to a balanced outcome.

¹This chapter is based on joint work with Yossi Azar, Ben Birnbaum, Nikhil Devanur and Yuval Peres [8,23].

Our Results

In this paper we first focus on a specific local bargaining procedure defined on a matching, and show that it converges to a fixed point. We then show that this fixed point is a balanced outcome whenever a balanced outcome exists. The convergence proof uses a potential function argument. Given the potential function, the proof of convergence is straightforward. However, the potential function we use is somewhat peculiar, owing to the fact that many natural potential functions do not work. This is because balancing one matched edge might lead to an increase in the imbalance of many other edges. The only special case where an alternate easier potential function is known is for a path.

We then focus on the rate of convergence and obtain a tight polynomial bound for a variety of natural dynamics on a certain classes of simple graphs including trees. The potential function used in our initial convergence proof is exponential, so does not provide any immediate bounds; instead, we draw a connection between network bargaining games and random-turn games.

Comparison with Contemporaneous Work

Network bargaining games were introduced to the theoretical computer science community by Kleinberg and Tardos [62]. They gave a polynomial time algorithm to compute the set of balanced outcomes. Since then, there has been an intensive study of various aspects of network bargaining [9, 10, 29, 30, 60]. Independently and concurrently with our work, Kanoria, et al. [61] considered the same problem and showed convergence of a different dynamics to a balanced outcome. We discuss some of the differences here.

The work by Kanoria et. al. uses local dynamics which synchronously update with a dampening factor; i.e. the update rule takes a weighted average of the current and balanced allocations. Additionally, follow up work [60] shows similar results for dynamics where nodes have unequal *bargaining power*; i.e. the surplus is not split evenly, rather it favors the node with more inherent skill. Our general convergence result (Section 5.2) holds for asynchronous balancing with no dampening factor. However, we extend this to several other dynamics (including damped or unequal bargaining) for certain classes of graphs (Section 5.3).

The dynamics Kanoria, et al. consider has the advantage that it does not need a matching to be known and fixed; rather, the dynamics also finds a matching. In our case, on the other hand, the matching M is fixed throughout. This may seem counterintuitive since the premise for a balanced outcome is that players may switching partners. However, once such a threat is acknowledged, the players do not need to switch in order to bargain; in fact, our convergence result shows that for

maximum M , we do reach stability. A different approach would be to consider known distributed dynamics that find matchings [11, 93]. These algorithms also have a bargaining flavor in their dynamics. One can imagine a two phase approach, where in the first phase the players find a matching and in the second find a balanced outcome with the matching fixed.

We say that an allocation f is ε -close to balanced if there exists a balanced outcome B such that $|B(v) - f(v)| \leq \varepsilon$ for all v ; i.e. all vertices are no more than ε -away from their value in a balanced outcome. Note that this is stronger than an ε -balanced outcome which only guarantees that each edge close to balanced; i.e. if an edge splits its surplus evenly the endpoint's values do not change by more than ε . Our work gives rates of convergence for the stronger global notion of ε -close to balanced as opposed to the local notion of ε -balanced considered in other work.

Other Related Work

In general, analysis of the convergence of local dynamics to an equilibrium of a game is a common theme. Examples include analysis of random best response dynamics for the Gale-Shapley stable matching game [2, 50]. In fact, a major philosophical hypothesis of algorithmic game theory [35, 45, 59] is that the existence of such dynamics is crucial to validate a solution concept.

Random-turn games are natural, and many variants have been analyzed [66–68]. Most interestingly, a variant called the tug-of-war game has been found to be related to partial differential equations such as the infinity Laplacian and the p -Laplacian [87], due to which these games have received considerable attention [31, 102]. Some variants have also been related to percolation [86].

Papers in a similar vein that consider price setting through a bargaining processes include the work by Corominas-Bosch [41]. At a meta-level, the convergence of local dynamics to a global equilibrium is a common theme. Ackerman et. al [3] showed an exponential lower bound for random best-response dynamics for the Gale-Shapley stable matching game [51]. Several papers [7, 32, 92] have studied the convergence of best-response dynamics to Nash equilibria in congestion games. In terms of structural results, Driessen [46] shows that the *kernel* is included in the *core* of an assignment game. This is in a similar spirit to one of the structural results we prove (Proposition 48).

Organization

We begin by giving a formal definition of a bargaining game and bargaining dynamics in Section 5.1. In Section 5.2 we show a general proof of convergence of a local bargaining process, and prove that it converges to a balanced outcome whenever one exists. Section 5.3 contains a tight bound on the rate of convergence for unweighted bipartite graphs with unique balanced outcomes, and shows how

some of these techniques can be extended to other types of bargaining processes. We conclude and suggest future work in Section 5.4.

5.1 Preliminary Definitions

We begin by formally defining the network bargaining game, balanced outcomes, and local dynamics in this context.

A network bargaining game is defined on a weighted undirected graph $G = (V, E)$ with $w : E \rightarrow \mathbb{R}_+$. Every node in the graph is a player, and the weight on an edge represents the dollar amount available to be shared between the two adjacent players. However, each player is constrained to make at most one such sharing agreement. An *outcome* of this game is a *matching* in the graph $M \subseteq E$ and an *allocation* describing each player's profit, $f_0 : V \rightarrow \mathbb{R}_+$ where for all $(uv) \in M$, we have $f_0(u) + f_0(v) = w(uv)$, and for all unmatched $u \in V$, $f_0(u) = 0$. Note that at times, it is useful to think of f_0 as a function as described here; however, especially in Section 5.2, it is also useful to think of f_0 as a point in \mathbb{R}_+^n where $n = |V|$.

We now formally define *stable* and *balanced* outcomes. Let the *best alternate* of a node u be

$$\alpha_{f_0}(u) = \max\{0, \max_{v:uv \in E \setminus M} \{w(uv) - f_0(v)\}\},$$

i.e. the maximum profit a player could get from a neighbor she is not currently matched to. For every matched edge (uv) define the surplus as

$$s_{f_0}(uv) = w(uv) - (\alpha_{f_0}(u) + \alpha_{f_0}(v)).$$

Note that the surplus may be negative. An outcome is balanced if it is stable and for all matched edges (uv) , $f_0(u) = \alpha(u) + s_{f_0}(uv)/2$ and $f_0(v) = \alpha(v) + s_{f_0}(uv)/2$, or equivalently, $f_0(u) - \alpha(u) = f_0(v) - \alpha(v)$. This can be seen as a generalization of Nash's bargaining solution for two players [82]. It is known that the following are equivalent: (1) a balanced outcome exists, (2) a stable outcome exists and (3) the matching polytope has no integrality gap [62].

The allocation in round t is denoted by $B_{f_0}(u, t)$, the best alternatives by $\alpha_{f_0}(u, t)$ and the surpluses by $s_{f_0}(uv, t)$.² The initial allocation is $B_{f_0}(v, 0) = f_0(v)$, and for each time $t > 0$ that allocation will be updated for some subset of edges $M_t \subseteq M$. How the subset M_t is chosen depends on the specific process. For example, it may just be a single edge. EDGE BALANCING is defined by the following update rule: for all $u \in V$, $(uv) \in M_t$ and $t > 0$,

$$B_{f_0}(u, t) = \alpha_{f_0}(u, t-1) + s_{f_0}(uv, t-1)/2.$$

²The subscript f_0 may be dropped when it is clear from context.

Note that since the surplus may be negative, B may also be negative. Thus, the allocation for the next round is determined by *balancing* each matched edge in M_t using the allocation in the current round.

5.2 Convergence for General Graphs

We begin by introducing our definitions and results for convergence on general graphs. In this section we will show that there are local dynamics that converge on general graphs; and specifically, whenever a balanced outcome exists, the dynamics will converge to such an outcome.

5.2.1 Supplementary Definitions


Recall that M is a matching on a graph $G = (V, E)$ with weights $w : E \rightarrow \mathbb{R}_+$, and let f be the current allocation on V . We have already defined *stable* and *balanced outcomes* and *edges* in Section 5.1. We now define two other types of matched edges for clarity in describing situations in which the surplus is negative. We say that an edge $uv \in M$ is *unhappy* if $\alpha_f(u) + s(uv)/2 < 0$ or $\alpha_f(v) + s(uv)/2 < 0$. That is, an edge is unhappy whenever the balancing process suggests a negative value for an endpoint; prohibiting negative values is equivalent to saying that the players at the nodes are *Individually Rational (IR)* at every time step. We say that vertex v is *saturated* if the unhappy edge $uv \in M$ is as close to being balanced as it can get without causing $f(u)$ to be negative, namely, $\alpha_f(u) + s(uv)/2 < 0$, $f(v) = w(uv)$ and $f(u) = 0$. Finally, we introduce the term *quasi-balanced* to denote an edge $uv \in M$ such that $f(u) = \alpha_f(u) + s(uv)/2$ and $f(v) = \alpha_f(v) + s(uv)/2$, regardless of whether or not $s(uv) \geq 0$. In a *quasi-balanced outcome* every matched edge is quasi-balanced.

It is not hard to see that a quasi-balanced outcome is stable if and only if no matched edge has negative surplus. Therefore, a balanced outcome is equivalent to a stable quasi-balanced outcome; this is the characterization we will use. Since a stable outcome must be on a maximum weight matching [62,95], a balanced outcome must therefore also be on a maximum weight matching. It is also known that if M is maximum, then a balanced outcome with M exists if and only if a stable outcome on G exists [62,91]. Hence, we have the following fact.


Fact 46. *A balanced outcome exists on a matching $M \subseteq G$ if and only if M is a maximum matching and there exists a stable outcome on G .*

In this section we consider a specific balancing process that maintains the individual rationality constraint. Additionally, we consider asynchronous dynamics where only one edge is updated at a time; otherwise it is easy to show that cyclic behavior can occur.³ This is a modification of the

³For example, consider an unweighted even cycle with $f(i) = \mathbb{1}_{\{i \text{ is even}\}}$.



(a) Fixed point that is not quasi-balanced.



(b) Fixed point that is quasi-balanced, but not stable.

Figure 5.1: Examples of fixed points of the balancing process that are not, in fact, balanced. In Figure 5.1(a), the fixed point is not quasi-balanced; however, no stable outcome exists. In Figure 5.1(b), the fixed point is quasi-balanced but not stable; however, the matching is not maximum.

simple dynamics defined in the introduction. For any v, t, f let $b(u) = \alpha_f(u) + s_f(uv)/2$. We define the *IR balancing function* below. The update rule we consider in this chapter is as follows: For each time $t > 0$ fix some u with $uv \in M$ and set

$$B_{f_0}(z, t+1) = \begin{cases} B_{f_0}(z, t) & \text{if } z \notin \{u, v\}, \text{ otherwise} \\ 0 & \text{if } b(z) < 0, \\ w(vu) & \text{if } b(z) > w(vu) \\ b(z) & \text{otherwise.} \end{cases}$$

We assume the natural non-starvation condition that each edge be considered an infinite number of times.

Clearly we maintain the invariant that $f(u) + f(v) = w(uv)$ for all $uv \in M$, and if $z \notin M$ then $f(z)$ does not change. In a fixed point of this process, every matched edge is either quasi-balanced or unhappy with a saturated endpoint, and unmatched edges may or may not be stable. There are examples of fixed points that are not quasi-balanced or stable (see Figure 5.1). However, we show that in any such example either M is not maximum or no stable outcome exists in G . More formally stated, our result is as follows:

Theorem. *If there exists a balanced outcome on M , then any fixed point of the balancing process is balanced.*

This theorem is proved in Section 5.2.3, and we prove the process converges to a fixed point in Section 5.2.2.

Theorem. *For any initial allocation \mathbf{f} and any matching M , the balancing process converges to a fixed point.*

5.2.2 Convergence

We begin by proving Theorem 5.2.1. Note the statement of the theorem holds for any G and M , although as previously stated, if M is not a maximum fractional matching then the point we converge to is not balanced.

Proof of Theorem 5.2.1. Let $W = \max\{w(uv) : uv \in E\}$, so a valid allocation f can be thought of as a point $f \in [0, W]^n$. Consider the multiset $S = \{f(v) : v \in V\} \cup \{f(u) + f(v) - w(uv) : uv \in E\}$. We call the values in S the *slacks* of f . Let $\mathbf{s} = \mathbf{s}(f) \in [-W, 2W]^{n+m}$ be the vector obtained by sorting the values in S in non-decreasing order. Define the potential function $\Phi : [-W, 2W]^{n+m} \rightarrow \mathbb{R}$ by

$$\Phi(\mathbf{s}) = \sum_{i=1}^{n+m} 2^{-i} s_i$$

where s_i is the i th index of \mathbf{s} . We first show the value of Φ converges as we run the balancing process, and then prove this implies the convergence of f . We know the entries of \mathbf{s} are at most $2W$, so $\Phi(\mathbf{s}) \leq 2W$ for any such slack vector. We now show that Φ is monotonically increasing as we bargain, and thus converges.

Assume f is not fixed, and let f' be the allocation obtained from from one step of the balancing process. Let $uv \in M$ be the edge selected for this step, and suppose that $x'_u = f(u) + \epsilon$ (and thus $x'_v = f(v) - \epsilon$) for some $\epsilon > 0$. Note that the slacks that change are those corresponding to edges adjacent to u or v (except for uv), and the vertices u and v . Additionally, each such slack changes by exactly ϵ , and thus $\|f - f'\|_\infty = \|\mathbf{s} - \mathbf{s}'\|_\infty = \epsilon$.

For any $au \in E$ where $a \neq v$ we have $f(u) - \alpha_f(u) \leq f(u) - (w(au) - f(a)) = f(a) + f(u) - w(au)$ and $f(u) - \alpha_f(u) \leq f(u)$. Similarly, for any $bv \in E$ where $b \neq u$, we have $f(v) - \alpha_f(v) \leq f(v) - (w(bv) - f(b)) = f(b) + f(v) - w(bv)$ and $f(v) - \alpha_f(v) \leq f(v)$. Note that since $f(u) < f(u)'$ we know that $f(u)' \neq 0$. Thus, either $|\alpha_f(u) - \alpha_f(v)| \leq w(uv)$, or $\alpha_f(u) - \alpha_f(v) > w(uv)$. In either case, $f(u) < x'_u \leq \frac{1}{2}(w(uv) + \alpha_f(u) - \alpha_f(v))$. Since $f(u) = w(uv) - f(v)$, we have $f(u) + (w(uv) - f(v)) < w(uv) + \alpha_f(u) - \alpha_f(v)$, or equivalently $f(u) - \alpha_f(u) < f(v) - \alpha_f(v)$. Hence, the minimum slack before the balancing step is $f(u) - \alpha_f(u)$. (Note that when u has degree 1 this is simply $f(u)$.) Similarly, after the balancing step, we can show that $f(u)' - \alpha_f(u)$ is the minimum slack. Hence, $\Phi(\mathbf{s}') - \Phi(\mathbf{s}) \geq (f(u)' - f(u))2^{-(m+n)} = \epsilon 2^{-(m+n)}$. Hence Φ is increasing, and thus convergent.

However, Φ is simply a function of the vector f . It is not clear that this implies the allocation vector f also converges. Despite this fact, it is not hard to show that the convergence of Φ implies the convergence of f . First, since Φ converges, the quantity $2^{n+m}\Phi_i$ must also converge and is therefore

Cauchy. From this, we have

$$2^{m+n}\Phi(\mathbf{s}') - 2^{m+n}\Phi(\mathbf{s}) \geq \epsilon = \|f' - f\|_\infty.$$

Hence f is Cauchy under the ℓ_∞ norm. Therefore, as the balancing process proceeds, f converges to a point in $[0, W]^{n+m}$.

Thus the balancing process converges regardless of the sequence of edges chosen. However, without a non-starvation condition, it is possible that we will not converge to a fixed point of the balancing process.⁴ Despite this fact, it suffices for every edge to be considered an infinite number of times. Assume for sake of contradiction that we converge to a point \mathbf{z} that is not a fixed point of the process. Since Φ is increasing, we know $\Phi(\mathbf{f}) \leq \Phi(\mathbf{z})$ for all intermediate f . Additionally, since \mathbf{z} is not a fixed point, there exists some edge uv whose endpoints converge to values that are not quasi-balanced and not saturated. Let ϵ be such that uv is ϵ -quasi-balanced in allocation \mathbf{z} . Since the process converges, we can proceed until uv is always at least $\epsilon/2$ unbalanced and $\Phi(\mathbf{z}) - \Phi(\mathbf{f}) < \epsilon 2^{-(n+m)}/4$. Since the non-starvation condition ensures that uv eventually be balanced eventually, at some point $\Phi(\mathbf{f})$ must increase by at least $\epsilon 2^{-(n+m)}/2$, and $\Phi(\mathbf{f}) > \Phi(\mathbf{z})$, which gives a contradiction. Hence we converge to a fixed point of the process. \square

Recall that while the process always converges, it may not reach a fixed point in a finite amount of time.¹ However, if we fix some ϵ and modify the balancing process so that edges only balance if they are not ϵ -quasi-balanced or ϵ -close to saturated, then Φ must increase by at least $\epsilon 2^{-(n+m)}$ in each step. In addition to preventing starvation, this shows a convergence time of $2W\epsilon \cdot 2^{n+m}$ for this process. We do not know if this bound is tight.

5.2.3 Fixed Point Characterization

We will prove Theorem 5.2.1 in two steps. The first step is to show that when a balanced outcome exists, a fixed point of the balancing process is quasi-balanced.

Proposition 47. *If there exists a stable outcome on G , M is a maximum matching, and \mathbf{f} is a fixed point of the balancing process, then \mathbf{f} has no unhappy edges (and is therefore quasi-balanced).*

The second step is to show that, given these same conditions, a quasi-balanced allocation is stable.

Proposition 48. *If there exists a stable outcome on G , M is a maximum matching, and \mathbf{f} is quasi-balanced, then \mathbf{f} contains no unstable edges (and is therefore balanced).*

⁴This is easy to show, even for a path of 6 nodes where we only choose to balance two edges.

Combined with Fact 46, these two propositions prove Theorem 5.2.1.

Proposition 48 is of independent interest because it shows that although quasi-balanced is a weaker notion than balanced, the two are equivalent on matchings that allow balanced outcomes. In fact, we can prove a generalization that is motivated by practical concerns as follows. Recall that although the balancing process converges, it may not reach a fixed point in finite time. However, once parties are sufficiently satisfied they may not want to negotiate further. Thus we consider a notion of sufficiently quasi-balanced or sufficiently stable.

We say that allocation \mathbf{f} is ε -quasi-balanced if the surplus is split evenly, to within an additive constant of ε ; namely every edge $uv \in M$ satisfies $|(f(u) - \alpha_f(u)) - (f(v) - \alpha_f(v))| \leq \varepsilon$. Similarly, an allocation is δ -stable if no two unmatched players have an unrealized exchange that is more than δ better for them; namely $f(u) + f(v) \geq w(uv) - \delta$ for each edge $uv \notin M$. In light of Proposition 48, one might ask if there is a similar relationship between ε -quasi-balanced and δ -stable. We settle this question in the affirmative in the following proposition.

Proposition 49. *If there exists a stable outcome on G , M is a maximum matching, and \mathbf{f} is an ε -quasi-balanced allocation on M , then \mathbf{f} is $(n\varepsilon)$ -stable, where $n = |V(G)|$.*

The proof of Proposition 49 follows similar ideas as the proof of Proposition 48.

In the proofs of Propositions 47 and 48, we use the following fact about the existence of stable outcomes. It can be proved by a straightforward duality argument [62].


Fact 50. *There exists a stable outcome on G if and only if a maximum fractional matching on G is integral.*

Thus, the conditions that M be a maximum integral matching and that there exists a stable outcome are equivalent to the condition that integral matching M be a maximum fractional matching.

The proofs of Propositions 47 and 48 share a common technique for proving the contrapositive. We assume there exists an edge uv that is unhappy (in the case of Proposition 47) or an edge uv that is unstable (in the case of Proposition 48). Starting from this edge, we explore the graph along matched edges and best alternatives and show that this exploration must terminate by finding a structure with properties that imply M is not a maximum fractional matching. This exploration algorithm is shown below. To ease notation, for the remainder of the section we will usually write $w(u_i u_j)$, $f(u_i)$ and $\alpha_f(u_i)$ as simply $w_{i,j}$, \mathbf{f}_i and α_i respectively.




(a) *Individually Rational (IR)*: at least one endpoint is matched. The other endpoint can take any form.




(b) *Lollipop*: at least one endpoint forms an alternating (even length) cycle. The other endpoint can take any form.




(c) *Augmenting path*: both endpoints are unmatched.



(d) *Flower*: one endpoint forms a blossom (odd cycle), and the other endpoint is unmatched.



(e) *Bicycle*: both endpoints form disjoint blossoms.



(f) *Pretzel*: both endpoints form blossoms which are not disjoint.

Figure 5.2: The subgraph H found by the exploration algorithm takes one of the above forms. Matched edges are depicted in bold.

Exploration Algorithm

Choose $u_0 \in V$.

Let $S = \emptyset$ be the set of explored vertices.

For $i = 0, 1, 2, 3, \dots$

If i is even:

If $u_i \notin M$ or $u_i \in S$, **break**.

Else, let $u_{i+1} = M(u_i)$.

Else (i is odd):

If $\alpha_i = 0$, **break**.

Else, choose $u_{i+1} \sim u_i$ such that $\alpha_i = w_{i,i+1} - \mathbf{f}_{i+1}$.

$S \leftarrow S \cup \{u_i\}$.

For $i = 0, -1, -2, -3, \dots$

If i is odd:

If $u_i \notin M$ or $u_i \in S$, **break**.

Else, let $u_{i-1} = M(u_i)$.

Else (i is even):

If $\alpha_i = 0$, **break**.

Else, choose $u_{i-1} \sim u_i$ such that $\alpha_i = w_{i,i-1} - \mathbf{f}_{i-1}$.

$S \leftarrow S \cup \{u_i\}$.

Note that by the definition of best alternate α_i , when $\alpha_i > 0$ a vertex $u_{i'} \sim u_i$ exists such that $\alpha_i = w_{ii'} - \mathbf{f}_{i'}$ as desired. Thus, this process is well defined. Let H be the subgraph formed by the set of vertices S and the edges travelled in their discovery. Because both directions of the exploration terminate if a vertex in S is rediscovered, the subgraph H can be classified in one of the six ways shown in Figure 5.2. In both proofs, we will proceed by case analysis on the type of H , making heavy use of the following lemma from [5, 62].

Lemma 51. *Let G be a graph, M a matching on G , and x an allocation on G with respect to M . If there is subgraph $H \subseteq G$ that is an augmenting path, alternating cycle, flower, or bicycle (see Figure 5.2) such that*

- for each $uv \in H \cap M$, $f(u) + f(v) = w(uv)$, and
- for each $uv \in H \setminus M$, $f(u) + f(v) < w(uv)$,

then there exists a (possibly fractional) matching on G with weight strictly greater than M .

Because the proofs of Proposition 48 and 49 are simpler than the proof of Proposition 47 and more clearly illustrate our exploration technique, we present them first.

Quasi-Balanced Outcomes are Stable

Proof of Proposition 48. We can rephrase the contrapositive using Fact 50 as follows: *If \mathbf{f} is a quasi-balanced allocation on M that contains some unstable edge, then M is not a maximum fractional matching.*

Assume there exists an unstable edge $uv \in M$. If neither u nor v are matched, then M is not maximum. Without loss of generality, assume u is matched and let $u_0 = u$. Let $u_\ell, \dots, u_0, \dots, u_r$ be the set of vertices discovered by the exploration algorithm, and let $H \subseteq G$ be the subgraph of vertices and edges traversed. We first prove that every edge $uv \in H \setminus M$ is unstable.

Specifically, we prove by induction that for $0 \leq i \leq \lfloor r/2 \rfloor$, we have $\mathbf{f}_{2i-1} + \mathbf{f}_{2i} < w_{2i-1,2i}$, and an equivalent argument holds for $0 \geq i \geq \lceil (\ell+1)/2 \rceil$. To prove the base case we show $u_{-1}u_0$ is unstable. Note $\alpha_0 \geq w_{v0} - f(v)$, which implies $\mathbf{f}_0 - \alpha_0 \leq \mathbf{f}_0 + \mathbf{f}_v - w_{v0} < 0$. Specifically, $0 < \alpha_0$, so u_{-1} exists and $\alpha_0 = w_{-1,0} - \mathbf{f}_{-1}$. Thus, $\mathbf{f}_0 - (w_{-1,0} - \mathbf{f}_{-1}) = \mathbf{f}_0 - \alpha_0 < 0$, or equivalently $\mathbf{f}_0 + \mathbf{f}_{-1} < w_{-1,0}$. Hence $u_{-1}u_0$ is unstable.

Now let us assume $u_{2i-3}u_{2i-2}$ is an unmatched edge that is unstable and let us show that $u_{2i-1}u_{2i}$ is unstable. Since $u_{2i-2}u_{2i-1}$ is a matched edge, it is by assumption quasi-balanced. Also, since $u_{2i-3} \sim u_{2i-2}$, it follows that $\alpha_{2i-2} \geq w_{2i-3,2i-2} - \mathbf{f}_{2i-3}$. Hence, $\mathbf{f}_{2i-1} + \mathbf{f}_{2i} = \mathbf{f}_{2i-1} + w_{2i-1,2i} - \alpha_{2i-1} = \mathbf{f}_{2i-2} - \alpha_{2i-2} + w_{2i-1,2i} \leq \mathbf{f}_{2i-2} - w_{2i-3,2i-2} + \mathbf{f}_{2i-3} + w_{2i-1,2i} < w_{2i-1,2i}$, and thus $u_{2i-1}u_{2i}$ is unstable. We conclude that every edge $uv \in H \setminus M$ is unstable.

Now consider H , and recall it takes the form of one of six structures in Figure 5.2. We consider each case, and show our desired conclusion. Since every unmatched edge in H is unstable, we know $\alpha_i > \mathbf{f}_i \geq 0$ for all i . Thus, H cannot be IR (Figure 5.2(a)). Now suppose that H is a lollipop (Figure 5.2(b)) or a pretzel (Figure 5.2(f)). If H is a lollipop then by definition it contains an alternating cycle. Alternately, if H is a pretzel then a simple parity argument shows that one of its three cycles must have even length, and therefore be alternating.

Hence, H is either an augmenting path (Figures 5.2(c)), a flower (Figures 5.2(d)), a bicycle (Figure 5.2(e)), or contains an alternating cycle. Since each unmatched edge is unstable, we apply Lemma 51 to conclude that M is not a maximum fractional matching, thus proving the contrapositive. \square

We prove Proposition 49 using similar ideas for this more general setting.

Proof of Proposition 49. Let w be the edge that maximizes $w_{u'v'} - \mathbf{f}_{u'} - \mathbf{f}_{v'}$ for all unmatched edges $u'v'$. Let $\delta = w(uv) - f(u) - f(v)$. To prove that \mathbf{f} is in the $(n\varepsilon)$ -core, it suffices to show that $\delta \leq n\varepsilon$. Since M is a maximum matching, at least one of u and v is matched. Without loss of generality, assume $u_0 = u$ is matched. Consider the exploration algorithm, and let H be the vertices $u_\ell, \dots, u_0, \dots, u_r$ and edges traversed by the algorithm.

We first prove the following lemma.

Lemma 52. *Any edge $ab \in H \setminus M$ satisfies $f(a) + f(b) \leq w_{ab} - \delta + (n-1)\varepsilon$.*

Proof. We prove a stronger claim: for $0 \leq i \leq \lfloor r/2 \rfloor$, $\mathbf{f}_{2i-1} + \mathbf{f}_{2i} \leq w_{2i-1,2i} - \delta + i\varepsilon$. The proof of this claim is by induction on i .

The base case is when $i = 0$. Since $u_0 \sim v$, we have $\alpha_0 = w_{-1,0} - \mathbf{f}_{-1} \geq w_{0v} - f(v) = \mathbf{f}_0 + \delta$. Hence, $\mathbf{f}_{-1} + \mathbf{f}_0 \leq w_{-1,0} - \delta$. (In fact, by the definition of $u = u_0$ and v , this implies that $\mathbf{f}_{-1} + \mathbf{f}_0 = w_{-1,0} - \delta$.)

Now suppose that the claim holds for $i-1$. That is, suppose that $\mathbf{f}_{2i-3} + \mathbf{f}_{2i-2} \leq w_{2i-3,2i-2} - \delta + (i-1)\varepsilon$. Since $u_{2i-3} \sim u_{2i-2}$, it follows that $\alpha_{2i-2} \geq w_{2i-3,2i-2} - \mathbf{f}_{2i-3}$. Thus, we have

$$\begin{aligned} \mathbf{f}_{2i-1} + \mathbf{f}_{2i} &= \mathbf{f}_{2i-1} + w_{2i-1,2i} - \alpha_{2i-1} && \text{(by choice of } u_{2i}\text{)} \\ &\leq w_{2i-1,2i} + \mathbf{f}_{2i-2} - \alpha_{2i-2} + \varepsilon && \text{(\varepsilon-quasi-balanced)} \\ &\leq w_{2i-1,2i} + \mathbf{f}_{2i-2} - w_{2i-3,2i-2} + \mathbf{f}_{2i-3} + \varepsilon && \text{(from above)} \\ &\leq w_{2i-1,2i} - \delta + i\varepsilon && \text{(induction hypothesis),} \end{aligned}$$

as desired.

By an analogous argument, we can show that for $0 \geq i \geq \lceil (\ell+1)/2 \rceil$, we have $\mathbf{f}_{2i-1} + \mathbf{f}_{2i} \leq w_{2i-1,2i} - \delta + i\varepsilon$. Hence, every edge $ab \in H \setminus M$ satisfies $f(a) + f(b) \leq w_{ab} - \delta + \max\{\lceil (\ell+1)/2 \rceil, \lfloor r/2 \rfloor\}\varepsilon \leq w_{ab} - \delta + (n-1)\varepsilon$. \square

With the lemma proved, we consider the possible types of structures H (see Figure 5.2). Since M is a maximum matching, we can use the standard matching results in [5, 62] to show that $\delta \leq n\varepsilon$ in every case as follows:

- Suppose that H is IR (Figure 5.2(a)). Without loss of generality, assume $u_{r-1}u_r$ is matched. Hence, $\alpha_r = 0$, and $\mathbf{f}_r - \alpha_r = \mathbf{f}_r \geq 0$. Because $u_{r-1}u_r$ is ε -balanced, $\mathbf{f}_{r-1} - \alpha_{r-1} \geq -\varepsilon$. Since $u_{r-1} \sim u_{r-2}$, we have $\alpha_{r-1} \geq w_{r-2,r-1} - \mathbf{f}_{r-2}$. From this and Lemma 52, we have

$$0 \leq \mathbf{f}_{r-1} - \alpha_{r-1} + \varepsilon \leq \mathbf{f}_{r-1} - w_{r-2,r-1} + \mathbf{f}_{r-2} + \varepsilon \leq -\delta + n\varepsilon ,$$

and hence $\delta \leq n\varepsilon$ as desired.

- Suppose that H is a lollypop (Figure 5.2(b)) or a pretzel (Figure 5.2(f)). If the latter, then a simple parity argument on the three intersecting cycles of H show that at least one must be even. Thus in either case, H contains an even alternating cycle. Relabel the vertices of this cycle v_0, v_1, \dots, v_{2k} , where $v_{2k} = v_0$ and $v_0 v_1 \in M$. By Lemma 52 and because M is a maximum matching, we have

$$\sum_{i=1}^k (\mathbf{f}_{2i-1} + \mathbf{f}_{2i}) \leq \sum_{i=1}^k w_{2i-1, 2i} - k\delta + kn\varepsilon \leq \sum_{i=0}^{k-1} w_{2i, 2i+1} - k\delta + kn\varepsilon .$$

On the other hand,

$$\sum_{i=1}^k (\mathbf{f}_{2i-1} + \mathbf{f}_{2i}) = \sum_{i=0}^{k-1} (\mathbf{f}_{2i} + \mathbf{f}_{2i+1}) = \sum_{i=0}^{k-1} w_{2i, 2i+1} .$$

Combining the above and dividing by k yields $\delta \leq n\varepsilon$ as desired.

- Suppose that H is an augmenting path (Figure 5.2(c)). Relabel the vertices of the augmenting path $v_0, v_1, \dots, v_{2k+1}$. By Lemma 52 and because M is a maximum matching, we have

$$\begin{aligned} \sum_{i=0}^k (\mathbf{f}_{2i} + \mathbf{f}_{2i+1}) &\leq \sum_{i=0}^k w_{2i, 2i+1} - (k+1)\delta + (k+1)n\varepsilon \\ &\leq \sum_{i=1}^k w_{2i-1, 2i} - (k+1)\delta + (k+1)n\varepsilon . \end{aligned}$$

On the other hand, since $\mathbf{f}_0 = \mathbf{f}_{2k+1} = 0$, we have

$$\sum_{i=0}^k (\mathbf{f}_{2i} + \mathbf{f}_{2i+1}) = \sum_{i=1}^k (\mathbf{f}_{2i-1} + \mathbf{f}_{2i}) = \sum_{i=1}^k w_{2i-1, 2i} .$$

Combining the above and dividing by $k+1$ yields $\delta \leq n\varepsilon$ as desired.

- Suppose that H is a flower (Figure 5.2(d)). Relabel the vertices of the flower starting at the bottom of the stem by $v_0, v_1, \dots, v_{2k}, \dots, v_{2\ell+1}$, where $v_{2k} = v_{2\ell+1}$. By Lemma 52 and because M is a maximum matching, we know

$$\begin{aligned} &2 \sum_{i=0}^{k-1} (\mathbf{f}_{2i} + \mathbf{f}_{2i+1}) + \sum_{i=k}^{\ell} (\mathbf{f}_{2i} + \mathbf{f}_{2i+1}) \\ &\leq 2 \sum_{i=0}^{k-1} w_{2i, 2i+1} + \sum_{i=k}^{\ell} w_{2i, 2i+1} - (k+\ell+1)\delta + (k+\ell+1)n\varepsilon \\ &\leq 2 \sum_{i=1}^k w_{2i-1, 2i} + \sum_{i=k+1}^{\ell} w_{2i-1, 2i} - (k+\ell+1)\delta + (k+\ell+1)n\varepsilon . \end{aligned}$$

On the other hand, since $\mathbf{f}_0 = 0$ and $\mathbf{f}_{2k} = \mathbf{f}_{2\ell+1}$, we have

$$\begin{aligned} 2 \sum_{i=0}^{k-1} (\mathbf{f}_{2i} + \mathbf{f}_{2i+1}) + \sum_{i=k}^{\ell} (\mathbf{f}_{2i} + \mathbf{f}_{2i+1}) &= 2 \sum_{i=1}^k (\mathbf{f}_{2i-1} + \mathbf{f}_{2i}) + \sum_{i=k+1}^{\ell} (\mathbf{f}_{2i-1} + \mathbf{f}_{2i}) \\ &= 2 \sum_{i=1}^k w_{2i-1,2i} + \sum_{i=k+1}^{\ell} w_{2i-1,2i} . \end{aligned}$$

Combining the above and dividing by $(k + \ell + 1)$ yields $\delta \leq n\varepsilon$ as desired.

- Finally, suppose that H is a bicycle. Relabel the vertices by

$$v_0, \dots, v_{2j+1}, \dots, v_{2k}, \dots, v_{2\ell+1} ,$$

such that $v_0 = v_{2j+1}$ and $v_{2k} = v_{2\ell+1}$. By Lemma 52 and because M is a maximum matching, we have

$$\begin{aligned} &2 \sum_{i=j+1}^{k-1} (\mathbf{f}_{2i} + \mathbf{f}_{2i+1}) + \sum_{i=0}^j (\mathbf{f}_{2i} + \mathbf{f}_{2i+1}) + \sum_{i=k}^{\ell} (\mathbf{f}_{2i} + \mathbf{f}_{2i+1}) \\ &\leq 2 \sum_{i=j+1}^{k-1} w_{2i,2i+1} + \sum_{i=0}^j w_{2i,2i+1} + \sum_{i=k}^{\ell} w_{2i,2i+1} - (\ell + k - j)\delta + (\ell + k - j)n\varepsilon \\ &\leq 2 \sum_{i=j+1}^k w_{2i-1,2i} + \sum_{i=1}^j w_{2i-1,2i} + \sum_{i=k+1}^{\ell} w_{2i-1,2i} - (\ell + k - j)\delta + (\ell + k - j)n\varepsilon . \end{aligned}$$

On the other hand,

$$\begin{aligned} &2 \sum_{i=j+1}^{k-1} (\mathbf{f}_{2i} + \mathbf{f}_{2i+1}) + \sum_{i=0}^j (\mathbf{f}_{2i} + \mathbf{f}_{2i+1}) + \sum_{i=k}^{\ell} (\mathbf{f}_{2i} + \mathbf{f}_{2i+1}) \\ &= 2 \sum_{i=j+1}^k (\mathbf{f}_{2i-1} + \mathbf{f}_{2i}) + \sum_{i=1}^j (\mathbf{f}_{2i-1} + \mathbf{f}_{2i}) + \sum_{i=k+1}^{\ell} (\mathbf{f}_{2i-1} + \mathbf{f}_{2i}) \\ &= 2 \sum_{i=j+1}^k w_{2i-1,2i} + \sum_{i=1}^j w_{2i-1,2i} + \sum_{i=k+1}^{\ell} w_{2i-1,2i} . \end{aligned}$$

Combining the above and dividing by $(\ell + k - j)$ yields $\delta \leq n\varepsilon$ as desired.

Thus, in all cases, we have shown that $\delta \leq n\varepsilon$, and hence \mathbf{f} is $(n\varepsilon)$ -stable. \square

Note that this result is tight to within a constant factor, even for unweighted graphs. Specifically, consider an unweighted lollipop (Figure 5.2(b)) consisting of $2k+1$ vertices $v_0, u_1, v_1, u_2, v_2, \dots, u_k, v_k$ where $u_1, v_1, \dots, u_k, v_k$ form an alternating cycle, and v_0 is adjacent only to u_1 . (Hence $u_j v_j$ is matched for $1 \leq j \leq k$.) Let $\varepsilon = 4/(k^2 + 4)$, and consider the allocation \mathbf{f} given by $\mathbf{f}_{u_j} =$

$1 - j(k + 2 - j)\varepsilon$ and $\mathbf{f}_{v_j} = j(k + 2 - j)\varepsilon$ for $1 \leq j \leq k$. It can be checked that every vertex gets an amount between 0 and 1, so this is a valid allocation. It can also be checked that this allocation is (2ε) -quasi-balanced. However, $\mathbf{f}_{v_0} + \mathbf{f}_{u_1} = \mathbf{f}_{u_1} = 1 - (k + 1)\varepsilon$, and therefore this allocation is not δ -stable for any $\delta = o(n\varepsilon)$.

Fixed Points are Quasi-Balanced

To prove that fixed points are quasi-balanced, we begin with the following lemma.

Lemma 53. *Let f be a fixed point of the balancing process on a maximum matching M , and let $u_0 u_1 \in M$ be an unhappy edge with u_1 saturated. Consider the vertices u_0, u_1, \dots, u_r discovered by the first part of the exploration algorithm. If $\ell \leq r$, then for any odd ℓ :*

1. For any j such that $0 \leq j < \lfloor \ell/2 \rfloor$, we have

$$\mathbf{f}_\ell + \sum_{i=j+1}^{\lfloor \ell/2 \rfloor} w_{2i-1, 2i} > \mathbf{f}_{2j+1} + \sum_{i=j+1}^{\lfloor \ell/2 \rfloor} w_{2i, 2i+1} .$$

2. $\mathbf{f}_\ell > 0$.
3. $\mathbf{f}_\ell - \alpha_\ell \leq \mathbf{f}_{\ell-1} - \alpha_{\ell-1}$.
4. $\alpha_\ell > \mathbf{f}_\ell$.
5. For any odd $\ell' < \ell$, we have $\alpha_\ell - \mathbf{f}_\ell \geq \alpha_{\ell'} - \mathbf{f}_{\ell'}$.

And, for any even ℓ :

6. For $\ell \geq 2$, $w_{\ell-1, \ell} > \mathbf{f}_{\ell-1} + \mathbf{f}_\ell$.
7. For $\ell \geq 2$, we have $\alpha_\ell > \mathbf{f}_\ell$.
8. For any j such that $0 \leq j < \ell/2$, we have

$$\mathbf{f}_\ell + \sum_{i=j}^{\ell/2-1} w_{2i, 2i+1} < \sum_{i=j+1}^{\ell/2} w_{2i-1, 2i} .$$

Proof. We prove this lemma by joint induction on ℓ . When $\ell = 0$, Lemma 53(6) and Lemma 53(7) do not apply (note that their proofs when $\ell = 2$ only rely on $\ell = 1$), and Lemma 53(8) holds vacuously. Additionally, when $\ell = 1$, Lemma 53(1) holds vacuously; Lemma 53(2) holds because $\mathbf{f}_1 = w_{01} > 0$; Lemma 53(3) holds because $\mathbf{f}_0 = 0$ and $\alpha_1 - \alpha_0 > w_{01} = \mathbf{f}_1$; Lemma 53(4) holds because $\alpha_1 > \mathbf{f}_1 + \alpha_0$ and $\alpha_0 \geq 0$; and finally, Lemma 53(5) holds vacuously. Let us assume the lemma holds for for all $0, \dots, \ell - 1$ where $\ell - 1 < r$. We now show that the claims hold for ℓ , proving the inductive step.

First consider the case where ℓ is odd, and denote $\ell = 2k + 1$:

Lemma 53(1): Let $0 \leq j < k$. We have

$$\begin{aligned}
\mathbf{f}_\ell + \sum_{i=j+1}^k w_{2i-1,2i} &= \mathbf{f}_\ell + w_{\ell-2,\ell-1} + \sum_{i=j+1}^{k-1} w_{2i-1,2i} \\
&= (w_{\ell-1,\ell} - \mathbf{f}_{\ell-1}) + w_{\ell-2,\ell-1} + \sum_{i=j+1}^{k-1} w_{2i-1,2i} \quad (u_{\ell-1}u_\ell \in M) \\
&> w_{\ell-1,\ell} + \mathbf{f}_{\ell-2} + \sum_{i=j+1}^{k-1} w_{2i-1,2i} \quad (\text{Lemma 53(6), induction hypothesis}) \\
&\geq w_{\ell-1,\ell} + \mathbf{f}_{2j+1} + \sum_{i=j+1}^{k-1} w_{2i,2i+1} \quad (\text{Lemma 53(1), induction hypothesis}) \\
&= \mathbf{f}_{2j+1} + \sum_{i=j+1}^k w_{2i,2i+1} .
\end{aligned}$$

Note that while we do use Lemma 53(6) for $\ell - 1$, we know that $\ell \geq 3$, so this does not present a problem.

Lemma 53(2): We have

$$\begin{aligned}
\mathbf{f}_\ell &> \mathbf{f}_1 + \sum_{i=1}^k w_{2i,2i+1} - \sum_{i=1}^k w_{2i-1,2i} \quad (\text{Lemma 53(1), with } j = 0) \\
&= \sum_{i=0}^k w_{2i,2i+1} - \sum_{i=1}^k w_{2i-1,2i} \quad (\mathbf{f}_1 = w_{01}) \\
&\geq 0 . \quad (\text{see below})
\end{aligned}$$

Note that the edges u_0u_1 and $u_{\ell-1}u_\ell$ are both matched. Thus the last inequality follows since M is a maximum matching.

Lemma 53(3): Because \mathbf{f} is a fixed point, $u_{\ell-1}u_\ell$ is either balanced or unhappy. By Lemma 53(2), it cannot be unhappy with $u_{\ell-1}$ saturated. Thus, it is either balanced and $\mathbf{f}_\ell - \alpha_\ell = \mathbf{f}_{\ell-1} - \alpha_{\ell-1}$

or it is unhappy with u_ℓ saturated. In the latter case, note that $\mathbf{f}_\ell + \mathbf{f}_{\ell-1} = w_{\ell, \ell-1} < a_\ell - a_{\ell-1}$ by definition. Since $\mathbf{f}_{\ell-1} = 0$, we have $\mathbf{f}_\ell - \alpha_\ell < \mathbf{f}_{\ell-1} - \alpha_{\ell-1}$, as desired.

Lemma 53(4): From Lemma 53(3), we have $\alpha_\ell \geq \mathbf{f}_\ell - \mathbf{f}_{\ell-1} + \alpha_{\ell-1}$. By Lemma 53(7) for $\ell - 1$ we know $\alpha_{\ell-1} > \mathbf{f}_{\ell-1}$. Hence $\alpha_\ell > \mathbf{f}_\ell$.

Lemma 53(5): First consider $\ell' = \ell - 2$. By Lemma 53(3), we have $\alpha_\ell - \mathbf{f}_\ell \geq \alpha_{\ell-1} - \mathbf{f}_{\ell-1}$. Since $u_{\ell-2} \sim u_{\ell-1}$, we have $\alpha_{\ell-1} \geq w_{\ell-1, \ell-2} - \mathbf{f}_{\ell-2}$. Hence $\alpha_\ell - \mathbf{f}_\ell \geq w_{\ell-1, \ell-2} - \mathbf{f}_{\ell-2} - \mathbf{f}_{\ell-1} = \alpha_{\ell-2} - \mathbf{f}_{\ell-2}$ by our choice of $u_{\ell-1}$. Thus, $\alpha_\ell - \mathbf{f}_\ell \geq \alpha_{\ell-2} - \mathbf{f}_{\ell-2}$. By the induction hypothesis, this holds for all odd $\ell' < \ell$ as desired.

Now consider the case where ℓ is even, and denote $\ell = 2k$:

Lemma 53(6): By our choice of u_ℓ , we know $w_{\ell-1, \ell} = \alpha_{\ell-1} + \mathbf{f}_\ell$. By Lemma 53(4), we have $\alpha_{\ell-1} > \mathbf{f}_{\ell-1}$. Thus $w_{\ell-1, \ell} > \mathbf{f}_{\ell-1} + \mathbf{f}_\ell$ as desired.

Lemma 53(7): Because $u_{\ell-1} \sim u_\ell$ and $u_{\ell-1}u_\ell \notin M$, we have $\alpha_\ell \geq w_{\ell-1, \ell} - \mathbf{f}_{\ell-1}$. By Lemma 53(6) we have $\alpha_\ell > \mathbf{f}_\ell$ as desired.

Lemma 53(8): We have

$$\begin{aligned}
\mathbf{f}_\ell + \sum_{i=j}^{k-1} w_{2i, 2i+1} &= \mathbf{f}_\ell + \mathbf{f}_{\ell-1} + \mathbf{f}_{\ell-2} + \sum_{i=j}^{k-2} w_{2i, 2i+1} \\
&\leq \sum_{i=j+1}^{k-1} w_{2i-1, 2i} + \mathbf{f}_{\ell-1} + \mathbf{f}_\ell && \text{(see below)} \\
&< w_{\ell-1, \ell} + \sum_{i=j+1}^{k-1} w_{2i-1, 2i} && \text{(Lemma 53(6))} \\
&= \sum_{i=j+1}^k w_{2i-1, 2i} .
\end{aligned}$$

The first inequality follows by the inductive hypothesis of Lemma 53(8) if $\ell = 2k \geq 4$. If $\ell = 2k = 2$, it follows because $\mathbf{f}_{\ell-2} = \mathbf{f}_0 = 0$.

Note that each case of the above lemma relies only on previous cases, or on the inductive hypothesis. Thus this concludes the proof. \square

We can now proceed with the proof of the proposition.

Proof of Proposition 47. We once again rephrase the contrapositive using Fact 50 as follows: *If \mathbf{f} is a fixed point of the balancing process on M that contains some unhappy edge, then M is not a maximum fractional matching.*

Recall that the only unhappy edges that appear in a fixed point must be saturated. Suppose $u_0u_1 \in M$ is unhappy with u_1 saturated. Consider the structure formed by the vertices u_0, u_1, \dots, u_r of the exploration process started at u_0 . Note that if the process ends because u_r is a previously labeled vertex, then for some $s \geq 0$ the sequence u_s, \dots, u_r forms an even alternating cycle or a blossom. Alternately, the process ends with u_r either IR or unmatched. We examine each possible case:

- If u_r is IR (Figure 5.2(a)): From Lemma 53(4) we know that $\alpha_i > \mathbf{f}_i \geq 0$ for all odd $i \geq 0$, hence by the definition of the exploration process this cannot occur.
- If u_r is unmatched (Figure 5.2(c)): By Lemma 53(8) with $j = 0$ we know

$$\sum_{i=0}^{r/2-1} w_{2i,2i+1} < \sum_{i=1}^{r/2} w_{2i-1,2i} .$$

Since u_0 is matched to u_1 , and u_r is not matched at all, M is not a maximum matching.

- If u_s, \dots, u_r forms an even alternating cycle for some $0 \leq s < r$ (Figure 5.2(b)):

First suppose $u_s = u_0$, so u_0, \dots, u_r forms a single alternating cycle, and r is even. By Lemma 53(5), $\alpha_{r-1} - \mathbf{f}_{r-1} \geq \alpha_1 - \mathbf{f}_1$. Since $\mathbf{f}_0 = 0$, and $u_r = u_0$, we know $\alpha_{r-1} = w_{r-1,0}$. Because u_{r-1} is adjacent to u_0 , we have $\alpha_0 \geq w_{r-1,0} - \mathbf{f}_{r-1}$. Hence $\alpha_0 \geq \alpha_1 - \mathbf{f}_1$, and $\mathbf{f}_1 = w_{01}$ implies $\alpha_1 - \alpha_0 \leq w_{01}$, contradicting the fact that u_0u_1 is unhappy with u_1 saturated.

Otherwise, we have an alternating cycle where $0 < s$. Note that r is even since any preceding u_s is already matched. If r is even, then s is also even, and by Lemma 53(8) with $j = s/2$, we know

$$\sum_{i=s/2}^{r/2-1} w_{2i,2i+1} < \sum_{i=s/2+1}^{r/2} w_{2i-1,2i} .$$

Thus, M is not a maximum matching.

- If u_s, \dots, u_r form a blossom (odd cycle, see Figure 5.2(d)): Thus, $r = 2k$ must be even and $s = 2k' - 1$ must be odd. Note that although $u_s \dots u_r$ forms a blossom and $u_0u_1 \dots u_s$ forms a stem, the bottommost edge is matched. Hence this is not equivalent to a flower, and Lemma 51

does not apply. Nevertheless, we can argue in a similar way to its proof by constructing a fractional \mathbf{y} matching of weight higher than M . This is done as follows.

$$y_e = \begin{cases} 1 & \text{if } e = u_i u_{i+1} \notin M \text{ for } 1 \leq i \leq 2k' - 2 \\ \frac{1}{2} & \text{if } e = u_i u_{i+1} \text{ for } 2k' - 1 \leq i \leq 2k - 1 \\ 1 & \text{if } e \in M \text{ but } e \text{ not in explored structure} \\ 0 & \text{otherwise} \end{cases} .$$

This places weight $1/2$ on the edges in the blossom, weight 1 the edges of the stem which are not in M . We wish to show that

$$\sum_{i=1}^{k'-1} w_{2i-1,2i} + \frac{1}{2} \sum_{i=s}^{r-1} w_{i,i+1} > \sum_{i=0}^{k-1} w_{2i,2i+1} ,$$

or equivalently

$$2 \sum_{i=1}^{k'-1} w_{2i-1,2i} + \sum_{i=k'}^k w_{2i-1,2i} > 2 \sum_{i=0}^{k'-1} w_{2i,2i+1} + \sum_{i=k'}^{k-1} w_{2i,2i+1} .$$

We have

$$\begin{aligned} & 2 \sum_{i=1}^{k'-1} w_{2i-1,2i} + \sum_{i=k'}^k w_{2i-1,2i} \\ & > 2 \sum_{i=1}^{k'-1} w_{2i-1,2i} + \left(\sum_{i=k'}^{k-1} w_{2i-1,2i} + \mathbf{f}_{2k-1} \right) + \mathbf{f}_{2k} && \text{(Lemma 53(6))} \\ & > 2 \sum_{i=1}^{k'-1} w_{2i-1,2i} + \left(\sum_{i=k'}^{k-1} w_{2i,2i+1} + \mathbf{f}_{2k'-1} \right) + \mathbf{f}_{2k} && \text{(Lemma 53(1), } \ell = 2k - 1) \\ & = 2 \left(\sum_{i=1}^{k'-1} w_{2i-1,2i} + \mathbf{f}_{2k'-1} \right) + \sum_{i=k'}^{k-1} w_{2i,2i+1} && (u_{2k} = u_{2k'-1}) \\ & = 2 \left(\sum_{i=1}^{k'-1} w_{2i-1,2i} - \mathbf{f}_{2k'-2} + w_{2k'-2,2k'-1} \right) + \sum_{i=k'}^{k-1} w_{2i,2i+1} && (u_{2k'-2,2k'-1} \text{ matched}) \\ & \geq 2 \left(\sum_{i=0}^{k'-2} w_{2i,2i+1} + w_{2k'-2,2k'-1} \right) + \sum_{i=k'}^{k-1} w_{2i,2i+1} && \text{(see below)} \\ & = 2 \sum_{i=0}^{k'-1} w_{2i,2i+1} + \sum_{i=k'}^{k-1} w_{2i,2i+1} . \end{aligned}$$

The last inequality follows from Lemma 53(8) if $k' \geq 2$, and if $k' = 1$, it follows because $\mathbf{f}_{2k'-2} = \mathbf{f}_0 = 0$. Hence, we conclude that M is not a maximum weighted fractional matching.

Thus, in all cases M is not a maximum fractional matching, and we have proved the contrapositive. \square

5.3 Rates of Convergence for Classes of Graphs

Now we focus on rates of convergence. We first consider a class of unweighted bipartite graphs with unique balanced outcomes. Our results generalize to a variety of settings as described in Section 5.3.3. Our results in this section rely crucially on transforming a local bargaining process into a random-turn game.

Random-turn games are a well-studied class of two-player combinatorial games in which the outcome of a coin flip determines which player moves next [66–68]. Combinatorial games can be represented as a game on a directed graph where players move a token along edges until one reaches their goal state. We transform the network bargaining game into an equivalent random-turn game which we can analyze using martingale techniques to obtain bounds on the rate of convergence. In particular, the convergence rate for the dynamics is related to the *absorption time* of the corresponding random-turn game.

Random-Turn Games

Every two-player game from Tic-Tac-Toe to Chess can be formalized as a combinatorial game on a directed graph where each turn consists of moving a token from one vertex to another along an edge [18]. Random-turn games are combinatorial games where the turns are determined by a coin flip.

In this section, a RANDOM-TURN GAME consists of a directed graph $D = (V, E)$, payoff function $f : V \rightarrow [0, 1]$, initial vertex v_0 , and horizon $T \in \mathbb{N}$. The set V of game states contains two terminal states s and r and all payoff functions set $f(s) = 0$ and $f(r) = 1$. The game is played by *Max* and *Mini* where *Max*'s goal is to maximize the value of the end state, and *Mini*'s goal is to minimize it. Game play for horizon T is as follows: a token is initially placed at v_0 and at every step a fair coin is tossed to determine who gets to move the token. *Max* must always move to a predecessor of v and *Mini* to a successor (as determined by the edge set E). We repeat until either T moves have been made, or we reach an absorbing state $\{s, r\}$. At the end of the game, *Mini* pays *Max* $\$f(v)$ if the game terminates at node v . Since this is a full-information game, for any finite horizon, one can compute the optimal strategies for the two players. This defines a *value* of the game, which is the expected payoff for *Max* under optimal play.

5.3.1 From a Bargaining Game to a Random-Turn Game

We now give a reduction from a network bargaining game to a random-turn game, the concept that lies at the heart of our results. We first restrict ourselves to unweighted bipartite graphs for clarity.




Figure 5.3: An unweighted bipartite graph G and its corresponding digraph D with balanced allocations.

Consider a graph $G = (V', E')$ where $w'(uv) = 1$ for all $(uv) \in E'$ and V' is bipartitioned as $\{L, R\}$. Create a directed graph D as follows: let $D = (V, E)$ where V is the subset of matched vertices in L along with two special vertices, s and r . Let the set of vertices other than s and r be denoted by \dot{V} . Add an edge $(uv) \in E$ if $(M(u)v) \in E'$. Additionally, place an edge from s to all vertices in \dot{V} and an edge from all vertices in \dot{V} to r . Finally, add an rv edge in E if there exists an edge $(vu) \in E'$ where $u \notin M$. Similarly, add a (vs) edge if there is a $(M(v)u)$ edge with $u \notin M$. We also give an allocation $f : V \rightarrow [0, 1]$ on D , given the allocation f' on G . Define $f(v) = f'(v)$ if $v \in \dot{V}$, $f(s) = 0$ and $f(r) = 1$. See Figure 5.3 for an example of this reduction. Note that an allocation f' on G takes values between 0 and 1 since the edge weights all have weight 1. Thus, the definition of an allocation allows us to reconstruct f' from f , since $f'(M(v)) = 1 - f(v)$ and $f(u) = 0$ when $u \notin M$.

The concepts (from the bargaining game described earlier) translate as follows.

- An allocation is *stable* if for all edges $(uv) \in E$, $f(u) \leq f(v)$.
- Let the *best predecessor* and *successor* of a node v be $v_f^+ = \arg \max_{u:uv \in E} \{f(u)\}$ and $v_f^- = \arg \min_{u:vu \in E} \{f(u)\}$ respectively. An allocation is *balanced* if it is stable, and for all vertices $v \in \dot{V}$, $f(v) = \frac{1}{2}(f(v^+) + f(v^-))$.
- Let the allocation in round t of EDGE BALANCING be $B_f(v, t)$ where $B_f(v, 0) = f(v)$. Then, balancing is equivalent to $B(v, t + 1) = \frac{1}{2}(B(v^+, t) + B(v^-, t))$.

An interesting aspect of this reduction is the *time reversal*. By that we mean that if one considers a T -horizon RANDOM-TURN GAME and T steps of EDGE BALANCING, then the first step of EDGE

BALANCING actually corresponds to the last step in the RANDOM-TURN GAME. In general, the t^{th} balancing step corresponds to t steps remaining in the game.

Throughout this paper, we say a graph D is *weakly acyclic* if the only directed cycles it contains go through s or r . If a graph G reduces to a digraph D that is weakly acyclic then the balanced outcome on G is *unique*. The converse also holds for unweighted bipartite graphs.

Consider the RANDOM-TURN GAME defined by the digraph $D = (V, E)$ and the payoff function f as above. The following theorem relates the value of the RANDOM-TURN GAME to EDGE BALANCING, and shows it is sufficient to analyze the convergence of the RANDOM-TURN GAME.

Theorem. *The value of a RANDOM-TURN GAME with starting vertex v and horizon T is exactly $B(v, T)$ when the directed graph is weakly acyclic.*

Let the balanced outcome be denoted by $B(v)$. For such games, we give the *optimal* rate of convergence, which is as follows. Let h be the maximum length of a path from s to r in D .

Theorem. *There exists a $T \in O(h^2 \log(1/\varepsilon))$ such that for all $t \geq T$ the value of the RANDOM-TURN GAME starting at vertex v with horizon t is within ε of $B(v)$, given that D is weakly acyclic.*

The proof of this theorem is the most technical part of this section, and uses techniques from the theory of martingales. Recall that an allocation f' is ε -close to balanced if there exists a balanced outcome B' such that $|B'(v) - f'(v)| \leq \varepsilon$ for all v . We can now restate the result and the corresponding rate of convergence in EDGE BALANCING. The proofs are the focus of Section 5.3.2.

Theorem. *EDGE BALANCING on unweighted bipartite graphs with a unique balanced outcome results in an allocation that is ε -close to a balanced outcome after at most $O(|M|^2 \log(1/\varepsilon))$ rounds of the balancing process.*

This result follows directly from Theorem 5.3.1 and the fact that $h \leq |V| = |M| + 2$. Lastly, we show our result is tight.

Theorem. *There exist graphs G with matchings M and initial allocations such that the balancing process requires $\Omega(|M|^2 \log(1/\varepsilon))$ time to be ε -close to a balanced outcome.*

Sketch of Convergence Proof

We now give a brief sketch of the proof of Theorem 2 for the case when M is a *perfect matching* in G . Observe that if a game with finite horizon ends in an absorbing state, then the vertex payoffs don't matter. Thus one approach is to show that with high probability, a RANDOM-TURN GAME with a sufficiently large horizon ends in an absorbing state. To be precise, let $\{X_t\}$ be a sequence

of vertices in a run of the RANDOM-TURN GAME under optimal play. We wish to show that for a game with sufficiently large horizon T , $X_T \in \{s, r\}$ with high probability. However, it is unclear how to analyze the behavior of X_t . Instead we show it is sufficient to analyze the related sequence of vertices $\{Y_t\}$ obtained when *Max* plays optimally, but *Mini* plays as if the payoff function was B . We show $B(Y_t)$, the value of the balanced outcome of vertex Y_t , is a supermartingale. Moreover, we know that it is bounded in $[0, 1]$ and show that its conditional variance is at least $1/h^2$. These suffice to prove the desired bound on the absorption time.

5.3.2 Rate of Convergence

We begin with the proof of Theorem 5.3.1. We recall some notation: given an allocation f , $v_f^+ = \arg \max_{u:uv \in E} \{f(u)\}$ and $v_f^- = \arg \min_{u:vu \in E} \{f(u)\}$. The allocation in round t of EDGE BALANCING is $B_f(v, t)$ (we now drop the subscript f for convenience). The updates are, $B(v, t+1) = \frac{1}{2}(B(v^+, t) + B(v^-, t))$ where v^+ and v^- are defined with respect to $B(v, t)$. Theorem 5.3.1 says that $B(v, T)$ is the value of the RANDOM-TURN GAME starting at vertex v with horizon T . The proof is by induction on T . We first strengthen the inductive hypothesis to assume the *optimal strategies* for *Max* and *Mini* are to choose v^+ and v^- respectively. We refer to this strategy as the *balancing strategy*.

Theorem. *Given a RANDOM-TURN GAME with horizon T , the optimal strategy for either player is the balancing strategy.*

Proof of Theorems 5.3.1 and 5.3.2. The proof is by a joint induction on the horizon t to prove (a) $B(v, t)$ is the value of the game and (b) the optimal strategy when $t+1$ moves remain is the balancing strategy.

In the base case, $t = 0$. To show (a), note that the expected payoff of the game for *Max* at node v is exactly $B_f(v, 0) = f(v)$ since there are no moves to be made. To show (b), consider the horizon $t + 1 = 1$ at a given node v . In this case, optimal moves for *Max* and *Mini* are clearly v_f^+ and v_f^- respectively, since the payoff at the end of this turn will be the terminal payoff of the game.

For the inductive step, let us assume that for all v and some $t \in \mathbb{N}$, the value of the game of horizon $t-1$ is $B_f(v, t-1)$, and in the t horizon game the bargaining strategy is optimal. To prove (a) we note that the latter statement implies *Max* will move to $v_{B(v, t-1)}^+$ if he wins the coin toss and *Mini* will move to $v_{B(v, t-1)}^-$ if she wins the coin toss. From the first part of the inductive hypothesis we know $B(v, t-1)$ is the expected payoff for *Max* in the $t-1$ horizon game. Thus, the expected payoff of the game for *Max* under optimal play in the t horizon game is $\frac{1}{2}(B(v^+, t-1) + B(v^-, t-1)) = B(v, t)$. To prove (b), consider the $t+1$ horizon game. Under optimal play, *Max* wishes to maximize his

expected payoff, and *Mini* wishes to minimize the expected amount she has to pay. Assume we are at vertex u , and recall that *Max* must move to a predecessor of u and *Mini* to a successor. Since there are t steps remaining after the initial step, an optimal strategy for *Max* (*Mini*) will maximize (minimize) the expected payoff $B_f(v, t)$. Thus, if *Max* wins the toss he will move to $v_{B(v,t)}^+$ and if *Mini* wins it she will move to $v_{B(v,t)}^-$, which is precisely the balancing strategy. \square

We now give the proof of Theorem 5.3.1 for the case where we have a perfect matching. Note that with the assumptions of the theorem, this implies D is *strongly acyclic*; i.e. it does not contain cycles of any kind. We briefly explain the technical extension for non-perfect matchings at the end of this section. The main idea behind the proof is to first reduce the analysis to showing that a particular sequence $\{Y_t\}$ (of vertices in V) gets absorbed at $\{s, r\}$ with high probability, and then show this happens in polynomial time using techniques from the theory of martingales.

Proof of Theorem 5.3.1. Consider two allocations, f and g such that $f(v) \leq g(v)$ for all v . We show in Lemma 54 that $B_f(v, t) \leq B_g(v, t)$ for all v, t . Hence, if we consider the initial allocations

$$\mathbf{0}(v) = \begin{cases} 0 & \text{if } v \neq r; \\ 1 & \text{otherwise.} \end{cases} \quad \text{and} \quad \mathbf{1}(v) = \begin{cases} 1 & \text{if } v \neq s; \\ 0 & \text{otherwise.} \end{cases},$$

we have $B_{\mathbf{0}}(v, t) \leq B_f(v, t) \leq B_{\mathbf{1}}(v, t)$ for all v, t , and f . Thus, it suffices to prove that $B_{\mathbf{0}}(v, T) \geq B(v) - \varepsilon$ and $B_{\mathbf{1}}(v, T) \leq B(v) + \varepsilon$ for $T \in O(h^2 \log(1/\varepsilon))$. We will prove the latter, and the proof for the former follows exactly with the roles of *Mini* and *Max* reversed and the payoff function $\mathbf{0}$ instead of $\mathbf{1}$.

Consider the game with payoff function $\mathbf{1}$ and horizon T where $T \in O(h^2 \log(1/\varepsilon))$. Consider the sequence of vertices $\{X_t\}$ with $X_0 = v$ that occurs if *Mini* and *Max* play optimally. From Theorem 5.3.1,

$$B_{\mathbf{1}}(v, t) = \mathbb{E}\mathbf{1}(X_t). \tag{5.1}$$

Now consider the half-optimal sequence $\{Y_t\}$ with $Y_0 = v$, where *Max* plays optimally for the payoff function $\mathbf{1}$ and *Mini* plays optimally for the payoff function B . For the game with payoffs $\mathbf{1}$ *Max*'s expected payoff is only higher. That is

$$\mathbb{E}\mathbf{1}(X_t) \leq \mathbb{E}\mathbf{1}(Y_t). \tag{5.2}$$

Our key result in Lemma 56 shows that for any function f , $\mathbb{E}_v[|f(Y_T) - B(Y_T)|] \leq \varepsilon$. (The proof of this lemma follows by showing convergence of the sequence $\{Y_t\}$.) If we take $f = \mathbf{1}$ and note that $\mathbf{1}(Y_t) \geq B(Y_t)$, we get

$$\mathbb{E}\mathbf{1}(Y_T) \leq \mathbb{E}B(Y_T) + \varepsilon. \tag{5.3}$$

Now consider the sequence $\{Z_t\}$ with $Z_0 = v$ that occurs when *Mini* and *Max* play optimally for the payoff function B . The expected payoff for *Max* with payoff function B is higher in $\{Z_t\}$ than in $\{Y_t\}$. Thus

$$\mathbb{E}B(Y_T) \leq \mathbb{E}B(Z_T). \quad (5.4)$$

Finally, we show in Lemma 55 that

$$\mathbb{E}B(Z_T) = B(v). \quad (5.5)$$

From (5.1) – (5.5), it follows that $B_{\mathbf{1}}(v, T) \leq B(v) + \varepsilon$ as desired. \square

Lemma 54. *The balancing process is monotonic, namely if $f(v) \leq g(v)$ for all $v \in V$, then $B_f(v, t) \leq B_g(v, t)$ for all v, t .*

Proof. Assume $f(v) \leq g(v)$ for all $v \in V$. Let us prove the claim by induction on t . For the base case, note that $B_f(v, 0) = f(v) \leq g(v) = B_g(v, 0)$. Now, assume that $B_f(v, t) \leq B_g(v, t)$ for some $t \in \mathbb{N}$. Then by the definition of the balancing process and the induction hypothesis,

$$\begin{aligned} B_f(v, t+1) &= \frac{1}{2}(B_f(v^+, t) + B_f(v^-, t)) \\ &\leq \frac{1}{2}(B_g(v^+, t) + B_g(v^-, t)) \\ &= B_g(v, t+1) \end{aligned}$$

as desired. \square

Lemma 55. *The value of a RANDOM-TURN GAME with function $f = B$ is equal to B for all horizons $T \in \mathbb{N}$.*

This Lemma follows from Theorem 5.3.1 and the observation that B is a fixed point of EDGE BALANCING.

Proof of Lemma 55. From Theorem 5.3.1 we know that $\mathbb{E}_v B(X_t) = B_B(v, t)$. We will use the latter formulation to prove this statement by induction on t . In the base case, we know $B_B(v, 0) = B(v)$ since no balancing has occurred. Now, let us assume that $B_B(v, t) = B(v)$ for some $t \in \mathbb{N}$. By definition, $B_B(v, t+1) = \frac{1}{2}(B_B(v^+, t) + B_B(v^-, t))$ and by induction this is equivalent to $\frac{1}{2}(B(v^+) + B(v^-)) = B(v)$. Hence, $B_B(v, t) = B(v)$ for all t . \square

Lemma 56. *Consider the expected payoff for *Max* in the half-optimal chain $\{Y_t\}$ defined in the proof of Theorem 5.3.1. For sufficiently large t , the expected payoff for *Max* with payoff function f is close to the balanced outcome B . Specifically, $\mathbb{E}_v[|f(Y_T) - B(Y_T)|] \leq \varepsilon$ when $T \geq 4h^2 \log(1/\varepsilon)$.*

Proof. Clearly if $Y_t \in \{s, r\}$, then the game has ended and $f(Y_t) - B(Y_t) = 0$. Additionally, the difference $|f(Y_t) - B(Y_t)|$ is at most 1 since $f(v), B(v) \in [0, 1]$ for all $v \in V$. Thus, the expected difference $\mathbb{E}_v[|f(Y_t) - B(Y_t)|]$ is at most the probability that Y_t has not been absorbed.

Let us now show this probability is bounded, namely $\mathbb{P}r_v[Y_t \notin \{s, r\} \text{ for } t \geq 4h^2 \log(1/\varepsilon)] \leq \varepsilon$ for all $v \in V$. The main convergence is shown in Lemma 57 which says that the probability that $Y_t \notin \{s, r\}$ for $t = 4h^2$ is at most $\frac{1}{4}$. Since the statement holds for all $v \in V$, if we are not at s or r after $4h^2$ time steps we can simply apply the lemma again. Thus, after $4h^2 \log(1/\varepsilon)$ time steps, the probability that we are not at s or r is $(\frac{1}{4})^{\log(1/\varepsilon)} = 4^{\log \varepsilon} \leq \varepsilon$. \square \square

Lemma 57. $\mathbb{P}r_v[Y_t \notin \{s, r\} \text{ for } t \geq 4h^2] \leq \frac{1}{4}$ for all $v \in V$ where h is the height of D^5 and $\{Y_t\}$ is the half-optimal chain defined above.

Proof. Let the absorption time be $\tau = \min\{t : Y_t \in \{s, r\}\}$. Note that $\mathbb{P}r[Y_t \notin \{s, r\} \text{ for some } t \geq 4h^2] = \mathbb{P}r[\tau \geq 4h^2]$. We show that $\mathbb{E}[\tau] \leq h^2$. Then by Markov's inequality, $\mathbb{P}r[\tau \geq 4h^2] \leq \frac{1}{4}$ as desired.

Consider the sequence $\{\Psi_t\} = \{B(Y_t)\}$. In the half-optimal chain $\{Y_t\}$, *Max* plays suboptimally and *Mini* plays optimally according to payoff function B (see the proof of Theorem 5.3.1). Hence $B(Y_t)$ is an upper bound on the expected payoff for *Max* at time t , and therefore $\{\Psi_t\}$ is a supermartingale.⁶

Now consider the quadratic chain $\Phi_t = 2\Psi_t - \Psi_t^2 + t\sigma^2$ where σ^2 is a lower bound on the conditional variance of Ψ_t . We show that Φ_t is also a supermartingale (Lemma 58). Therefore, since $\Phi_t \geq 0$, the optional stopping theorem⁷ gives $\mathbb{E}[\Phi_\tau] \leq \Phi_0 \leq 1$. The bounds on Ψ_t also imply that $2\Psi_t - \Psi_t^2 \geq 0$, and hence we get $\mathbb{E}[\Phi_\tau] \geq \mathbb{E}[\tau]\sigma^2$, or equivalently $\mathbb{E}[\tau] \leq \frac{1}{\sigma^2}$. By Lemma 59, we know that we can take $\sigma^2 = \frac{1}{h^2}$, so $\mathbb{E}[\tau] \leq h^2$ as required. \square \square

Lemma 58. *Given a supermartingale $0 \leq \Psi_t \leq 1$ with conditional variance at least σ^2 , the quadratic chain $\Phi_t = 2\Psi_t - \Psi_t^2 + t\sigma^2$ is a supermartingale.*

Proof. To prove Φ_t is a supermartingale, we show that the expected step size of Φ_t is not positive, i.e. $\mathbb{E}[\Phi_{t+1} - \Phi_t | Y_t] \leq 0$.

⁵The *height* is the length of the longest path from s to r .

⁶Recall that a supermartingale is a sequence $\{a_t\}$ in which $a_t \geq \mathbb{E}[a_{t+1} | a_t]$.

⁷See Theorem 10.10 (d) in [101].

$$\begin{aligned}
\mathbb{E}[\Phi_{t+1} - \Phi_t | Y_t] &\leq 2\mathbb{E}[\Psi_{t+1} | Y_t] - \mathbb{E}[\Psi_{t+1}^2 | Y_t] + (t+1)\sigma^2 - 2\Psi_t + \Psi_t^2 - t\sigma^2 \\
&\leq -1 + 2\mathbb{E}[\Psi_{t+1} | Y_t] - \mathbb{E}[\Psi_{t+1} | Y_t]^2 + 1 - 2\Psi_t + \Psi_t^2 \\
&= -(1 - \mathbb{E}[\Psi_{t+1} | Y_t])^2 + (1 - \Psi_t)^2 \leq 0.
\end{aligned}$$

Since Ψ_t is a supermartingale and $\Psi_t \leq 1$. □

Lemma 59. *The variance of a step in $\{\Psi_t\}$ is at least $\sigma^2 = 1/h^2$.*

Proof. Let B be the balanced outcome for a graph D with no cycles of any kind, and let h be the height of the digraph. Consider a path from s to r of length h . The stability condition for B says the values along this path are increasing. Additionally, the balanced condition states that each vertex takes the value of the average between its maximum predecessor and minimum successor. Since $B(s) = 0$ and $B(r) = 1$ and the length of the path is maximal, it is clear that the i th vertex will have value i/h in B .

Thus, in a graph D with no cycles, the minimum difference between a vertex and its predecessors and successors is at least $1/h$. Additionally, recall that $\{\Psi_t\}$ is a supermartingale, so $\Psi_t \geq \mathbb{E}[\Psi_{t+1} | \Psi_t]$. Therefore,

$$\begin{aligned}
\sigma^2 &= \mathbb{E}[\Psi_{t+1}^2 | \Psi_t] - \mathbb{E}[\Psi_{t+1} | \Psi_t]^2 \\
&\geq \mathbb{E}[\Psi_{t+1}^2 | \Psi_t] - \Psi_t^2 \\
&\geq \frac{1}{2} ((\Psi_t + 1/h)^2 + (\Psi_t - 1/h)^2) - \Psi_t^2 \\
&= (\Psi_t^2 + 1/h^2) - \Psi_t^2 = \frac{1}{h^2},
\end{aligned}$$

giving the desired bound. □

For completeness, we now prove the matching lower bound.

Proof of Theorem 5.3.1. Consider the unweighted graph G which corresponds to an alternating path which starts and ends on matched edges. This corresponds to a digraph D which is a directed path of length $h = |M| + 1$ from s to r . This graph is easy to analyze since *Max* and *Mini* always have exactly one choice. Therefore the sequence X_t is a simple random walk on a path, regardless of the payoff function. Now consider the payoff function $\mathbf{1}$. Then $\mathbb{E}[\mathbf{1}(X_t)]$ (for any starting vertex) is the probability that $X_t \neq s$. Similarly if the payoff function is $\mathbf{0}$, then $\mathbb{E}[\mathbf{0}(X_t)]$ is the probability that $X_t = r$. Therefore the difference $\mathbb{E}[\mathbf{1}(X_t)] - \mathbb{E}[\mathbf{0}(X_t)]$ is the probability that $X_t \notin \{s, r\}$, that is

the probability of non-absorption. Hence if the probability of non-absorption is large, then either $\mathbb{E}[\mathbf{1}(X_t)]$ or $\mathbb{E}[\mathbf{0}(X_t)]$ is far from the balanced outcome.

We now show a lower bound on the absorption probability. Consider a simple random walk along the path of best alternates, and the mixing time τ_{mix} where we identify s with r to form a single state. The stationary distribution is 1 at the $s-r$ vertex, and 0 elsewhere. Thus, $\tau_{mix}(\varepsilon)$ corresponds to the time at which we are absorbed with probability $1 - \varepsilon$. Lower bounding τ_{mix} will give a lower bound on the rate of convergence. From standard techniques developed to analyze Markov chain mixing times (e.g., see [69]) we know $\tau_{mix}(\varepsilon) \geq (1/\gamma_* - 1) \log(1/2\varepsilon)$ where γ_* is related to the spectral gap. Again, from the literature, we know that $\gamma_* \in \Theta(h^2)$, so $\tau_{mix}(\varepsilon) \in \Omega(h^2 \log(1/\varepsilon))$. \square

When M is not a perfect matching, D is a weakly acyclic digraph with a cycles through s and/or r . Any vertex in such a cycle must take value exactly 0 or 1 in the balanced outcome, thus these cycles can be treated as absorbing states. Hence, we can first analyze the mixing time of the cycle using spectral techniques⁸, and then apply the theorems above to get the same time bound.

5.3.3 Extensions

To summarize, the approach outlined to prove convergence of EDGE BALANCING is as follows: reduce it to convergence of a RANDOM-TURN GAME (Theorem 5.3.1) and show bounds on this game (Theorem 5.3.1). The first part of this approach can be extended naturally to show convergence (but not rates) for many variants of the dynamics and general graphs. For *non-bipartite graphs* we maintain a vertex in D for each matched vertex in G . If the graph is *weighted* we use running payoffs in the random-turn game. *Damped dynamics* correspond to lazy random-turn games. And if we wish vertices to be *individual rational*, then the corresponding *IR dynamics* are captured by a random-turn game where the players are allowed to quit. This list is far from exhaustive, but illustrates the flexibility and robustness of our technique.

Asynchronous dynamics: In the asynchronous process, matched edges update their allocation one at a time according to an arbitrary predetermined order. In other words, for every time step t there is a fixed edge that updates during that time step. This is more natural than forcing a synchronous update to a local process.

The reduction uses a modified RANDOM-TURN GAME where in each step a move can be made at only one vertex; more precisely, we allow a move if and only if the current vertex matches the vertex updated in the corresponding step of EDGE BALANCING. Thus, it is defined for a sequence $\{v_1, \dots, v_T\}$ of vertices in \dot{V} which determine the balancing order. We say a sequence has *starvation*

⁸See Chapter 12 in Peres, et al. [69] for an exposition on spectral techniques.




Figure 5.4: A graph G and its corresponding digraph D using the general reduction. Note that we omit some (sv) and (vr) edges.

$N = \max_{v \in \dot{V}, i} \{\min\{j - i : v_j = v\} : v_i = v \text{ or } v_k \neq v \forall k < i\}$; i.e., this is the maximum number of time steps an edge must wait before being balanced. This leads to the asynchronous version of Theorems 5.3.1 and 5.3.1.

Theorem. *For starvation N , there exists $T \in \Theta(Nh^2 \log(1/\varepsilon))$ such that for all $t \geq T$ the allocation of ASYNCHRONOUS EDGE BALANCING after t steps is ε -close to a balanced outcome when D is weakly acyclic.*

Unweighted non-bipartite graphs: For a bipartite graph, the reduction maintained one side of the matched vertices, which was sufficient to determine the entire allocation. In the non-bipartite case, we cannot partition the vertices in this manner. Hence, we create a vertex in D for every matched vertex in G . Specifically, let \dot{V} be the matched vertices in V' , and let $V = \dot{V} \cup \{s, r\}$. The edges are defined exactly as before with the new \dot{V} , and an allocation f on D is defined starting with an allocation f' on G as before. See Figure 5.4 for an example of this reduction.

This leads to the unweighted non-bipartite graph version of Theorems 5.3.1 and 5.3.1.

Theorem. *There exists a $T \in \Theta(h^2 \log(1/\varepsilon))$ such that for all $t \geq T$ the allocation of NON-BIPARTITE EDGE BALANCING after t steps is ε -close to a balanced outcome when D is weakly acyclic.*

Note that in this case, D weakly acyclic no longer corresponds to graphs with unique balanced outcomes.

Damped dynamics: We can similarly consider a *damped* process where the update is a fixed convex combination of the current allocation and the allocation given by EDGE BALANCING. Specif-

ically, the update rule becomes

$$B_f(u, t) = \kappa(\alpha_f(u, t-1) + s_f(uv, t-1)/2) + (1 - \kappa)B_f(u, t-1)$$

for some $0 \leq \kappa \leq 1$. This corresponds to a lazy RANDOM-TURN GAME in which at every step with probability κ you move as before, but with probability $(1 - \kappa)$ you stay at the same vertex.

This leads to the damped version of Theorems 5.3.1 and 5.3.1.

Theorem. *For dampening factor $0 \leq \kappa \leq 1$, there exists a $T \in \Theta(h^2 \log(1/\varepsilon)/\kappa)$ such that for all $t \geq T$, the value of the DAMPED EDGE BALANCING after t steps is ε -close to a balanced outcome when D is weakly acyclic.*

Damped updates could be useful in analyzing cases where D has an even-length cycle, by breaking the periodicity of updates.

Unequal Bargaining Powers: We can similarly consider a process where node u has skill δ_u in bargaining as considered in [60]. The update rule then splits the surplus in favor of the node with higher bargaining power. Specifically, edge uv would give u a $\delta_u/(\delta_u + \delta_v)$ proportion of the surplus. Specifically, the update rule becomes

$$B_f(u, t) = \alpha_f(u, t-1) + \frac{\delta_u s_f(uv, t-1)/2}{\delta_u + \delta_v}.$$

This corresponds to a lazy RANDOM-TURN GAME in which the ammount of laziness varies from vertex to vertex depending on the values of δ .

This leads to the damped version of Theorems 5.3.1 and 5.3.1.

Theorem. *Let $\hat{\delta} = \min_{u \in V, uv \in M} \{\delta_u/(\delta_u + \delta_v)\}$. Then, there exists a $T \in \Theta(h^2 \log(1/\varepsilon)/\hat{\delta})$ such that for all $t \geq T$, the value of the unequal bargaining power process after t steps is ε -close to a balanced outcome when D is weakly acyclic.*

Weighted Graphs: Recall that the original process we defined was for weighted graphs. In a weighted graph G , the directed graph will have edge weights w . The reduction is as before, except that the random-turn game is modified to include running payoffs. Each time *Max* moves (backward) along an edge from u to v , he pays $\$w(uv)$ to *Mini*. Similarly, if *Mini* moves along an edge from u to v , she pays $\$w(uv)$ to *Max*. These additional payments are called *running payoffs*, and have been studied previously in combinatorial game literature [18, 66, 88].

Our technique for weighted graphs shows convergence, but we can no longer attain a bound on the rate of convergence since the running payoffs are unbounded, so our proof cannot be easily generalized.




Figure 5.5: A weighted bipartite graph G and its corresponding digraph D with balanced allocations.

Theorem. WEIGHTED EDGE BALANCING converges for weakly acyclic D .

Individually Rational dynamics: In weighted graphs, the updates as defined could lead to negative allocations for some vertices as considered in Section 5.2. Note that this is not an issue for unweighted graphs, which is why it was not a concern in the bulk of Section 5.3. If we wish vertices to be individually rational, we must modify the process as in Section 5.2. This leads to a modification of the RANDOM-TURN GAME where if at any stage the expected future payoff for *Max* under optimal play is less than zero, then we stop the game. Similarly, there is a modification for *Mini* (that is not entirely symmetric). The convergence rate is the same as for weighted graphs. The IR version presents few extra difficulties, and the theorems for weighted graphs generalize as they are.

Theorem. WEIGHTED EDGE BALANCING converges for weakly acyclic D .

5.4 Conclusion and Future Work

In this paper, we first focused on a specific local bargaining procedure defined on a matching, and showed that it converges to a balanced outcome whenever a balanced outcome exists. Next, we reduced the problem of analyzing the convergence of local dynamics for a network bargaining game to that of a random-turn game. With this reduction we bring all the machinery from the analysis of random processes, especially the theory of Markov chains and martingales, to the analysis of local dynamics. We used these techniques to give the optimal bound on unweighted graphs with a unique balanced outcome.

Our work opens up a promising line of approach to analyze many variants of local dynamics on general graphs. The most immediate is perhaps to bound the convergence rate for weighted graphs.

The difficulty with our current analysis is that the supermartingale Ψ_t we used in the unweighted case is unbounded when there are weights. We believe a different supermartingale that does not suffer from this drawback could give the appropriate bound.

The most significant technical hurdle arises when D is cyclic. In this case, the game may never end, since it might get stuck in a *stalemate*, where the players travel in a cycle indefinitely. Thus, a bound on the absorption time of the game does not suffice – we must analyze the behavior on the cycle separately by internally considering its *mixing time*, and externally treating it as an absorbing state.⁹ However, the details of such an analysis remain unclear.

A final important direction is to obtain tight polynomial bounds for dynamics which find both the matching and the balanced outcome simultaneously. One approach would be to combine the dynamics by Kanoria et. al. [61] with our techniques to attain a tight polynomial rate of convergence.

⁹This proposed approach is a generalization of our analysis of non-perfect matchings in unweighted graphs.

Bibliography

- [1] Ittai Abraham, Susan Athey, Moshe Babioff, and Michael Grubb. Peaches, lemons and cookies: The impact of information structure on revenue in common value auctions. working paper, Microsoft Research, 2010.
- [2] Heiner Ackermann, Paul W. Goldberg, Vahab S. Mirrokni, Heiko Röglin, and Berthold Vöcking. Uncoordinated two-sided matching markets. In *ACM Conference on Electronic Commerce*, pages 256–263, 2008.
- [3] Heiner Ackermann, Paul W. Goldberg, Vahab S. Mirrokni, Heiko Röglin, and Berthold Vöcking. Uncoordinated two-sided matching markets. In *ACM Conference on Electronic Commerce*, pages 256–263, 2008.
- [4] William J. Adams and Janet Yellen. Commodity bundling and the burden of monopoly. *Quarterly Journal of Economics*, 90(3):475–498, 1976.
- [5] Bengt Aspvall and Yossi Shiloach. A polynomial time algorithm for solving systems of linear inequalities with two variables per inequality. *SIAM Journal on Computing*, 9(4):827–845, 1980.
- [6] Susan Athey and Denis Nekipelov. A structural model of sponsored search advertising auctions. Working Paper, Harvard University, 2010.
- [7] Baruch Awerbuch, Yossi Azar, Amir Epstein, Vahab S. Mirrokni, and Alexander Skopalik. Fast convergence to nearly optimal solutions in potential games. In *ACM Conference on Electronic Commerce*, pages 264–273, 2008.
- [8] Yossi Azar, Benjamin Birnbaum, L. Elisa Celis, Nikhil R. Devanur, and Yuval Peres. Convergence of local dynamics to balanced outcomes in exchange networks. In *IEEE Symp. on Foundations of Computer Science*, 2009.

- [9] Yossi Azar, Nikhil Devanur, Kamal Jain, and Yuval Rabani. Monotonicity in bargaining networks. In *ACM Symp. on Discrete Algorithms*, 2010.
- [10] Mohammad Hossein Bateni, Mohammad Taghi Hajiaghayi, Nicole Immorlica, and Hamid Mahini. Network bargaining with capacity constraints. In *Intl. Colloq. on Automata, Languages and Programming*, 2010.
- [11] Mohsen Bayati, Christian Borgs, Jennifer T. Chayes, and Riccardo Zecchina. Belief-propagation for weighted b-matchings on arbitrary graphs and its relation to linear programs with integer solutions. *SIAM J. in Discrete Math*, 2010.
- [12] D. Bergemann and J. Valimaki. Information acquisition and efficient mechanism design. *Econometrica*, 70(3), 2002.
- [13] D. Bergemann and J. Valimaki. Information acquisition and efficient mechanism design. *Econometrica*, 70(3), 2002.
- [14] Dirk Bergemann and Alessandro Bonatti. Targeting in advertising markets: Implications for offline vs. online media. Cowles Foundation Discussion Paper 1758R, 2010.
- [15] Dirk Bergemann, Alessandro Bonatti, and Maher Said. Product design in ad exchanges. Working paper, MIT Sloan, 2011.
- [16] Dirk Bergemann and Martin Pesendorfer. Information structures in optimal auctions. Cowles Foundation Discussion Paper 1323, 2001.
- [17] Dirk Bergemann and Maher Said. Dynamic auctions: A survey. *Wiley Encyclopedia of Operations Research and Management Science*, 2010.
- [18] Elwyn R. Berlekamp, John H. Conway, and Richard K. Guy. *Winning ways for your mathematical Plays*. Academic Press, London, 1982.
- [19] Sushil Bikhchandani. Information acquisition and full surplus extraction. *Journal of Economic Theory*, 2009.
- [20] Liad Blumrosen and Thomas Holenstein. Posted prices vs. negotiations: an asymptotic analysis. In *ACM Conference on Electronic Commerce*, page 49, 2008.
- [21] Ruggiero Cavallo and David C. Parkes. Efficient metadeliberation auctions. In *AAAI*, pages 50–56, 2008.
- [22] L. Elisa Celis, Nikhil R. Devanur, and Yuval Peres. Local dynamics in exchange networks via random turn games. In *Workshop in Network Economics*, 2010.
- [23] L. Elisa Celis, Nikhil R. Devanur, and Yuval Peres. Local dynamics in exchange networks via random turn games. In *Workshop in Network Economics*, 2010.

- [24] L. Elisa Celis, Dimitrios Gklezakos, and Anna Karlin. Revenue maximization for information acquisition via spps. Working Paper, 2012.
- [25] L. Elisa Celis, Anna Karlin, Kevin Leyton-Brown, Thach Nguyen, and David Thompson. Approximately revenue-maximizing mechanisms for deliberative agents. In *International World Wide Web Conference*, 2011.
- [26] L. Elisa Celis, Gregory Lewis, Markus, Mobius, and Hamid Nazerzadeh. Buy-it-now or take-a-chance: A simple sequential screening mechanism. In *International World Wide Web Conference*, 2011.
- [27] L. Elisa Celis, Gregory Lewis, Markus, Mobius, and Hamid Nazerzadeh. Buy-it-now or take-a-chance: A mechanism for real-time price discrimination. Working Paper, 2012.
- [28] Tanmoy Chakraborty, Stephen Judd, Michael Kearns, and Jinsong Tan. A behavioral study of bargaining in social networks. In *ACM Conference on Electronic Commerce*, 2010.
- [29] Tanmoy Chakraborty and Michael Kearns. Bargaining solutions social networks. In *Workshop on Internet and Network Economics*, 2008.
- [30] Tanmoy Chakraborty, Michael Kearns, and Sanjeev Khanna. Nash bargaining: Algorithms and structural results. In *ACM Conference on Electronic Commerce*, 2009.
- [31] Fernando Charro, Jesus García Azorero, and Julio D. Rossi. A mixed problem for the infinity Laplacian via tug-of-war games. *Var. & Par. Differential Eq.*, 34(3):307–320, 2009.
- [32] K. Chatterjee and W. Samuelson. Bargaining under incomplete information. *Operations Research*, 31:835–851, 1983.
- [33] S. Chawla, J. Hartline, D. Malec, and B. Sivan. Sequential posted pricing and multi-parameter mechanism design. In *Proc. 41st ACM Symp. on Theory of Computing*, 2010.
- [34] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd ACM Symp. on Theory of Computing (STOC)*, pages 311–320, 2010.
- [35] Xi Chen and Xiaotie Deng. Settling the complexity of two-player Nash equilibrium. In *IEEE Symp. on Foundations of Computer Science*, pages 261–272, 2006.
- [36] Sean Chu, Phillip Leslie, and Alan Sorensen. Bundle-size pricing as an approximation to mixed bundling. *American Economic Review*, 101(1), 2011.
- [37] Olivier Compte and Philippe Jehiel. Auctions and information acquisition: Sealed-bid or dynamic formats? Levine’s Bibliography 78482800000000495, UCLA Department of Economics, October 2005.
- [38] K. S. Cook, R. M. Emerson, M. R. Gillmore, and T. Yamagishi. The distribution of power in exchange networks: Theory and experimental results. *American Journal of Sociology*, 89:275–305, 1983.

- [39] Karen S. Cook and Toshio Yamagishi. Power in exchange networks: a power-dependence formulation. *Social Networks*, 14(3-4):245–265, 1992.
- [40] Margarida Corominas-Bosch. Bargaining in a network of buyers and sellers. *Journal of Economic Theory*, 115(1):35 – 77, 2004.
- [41] Margarida Corominas-Bosch. Bargaining in a network of buyers and sellers. *Journal of Economic Theory*, 115(1):35 – 77, 2004.
- [42] Pascal Courty and Hao Li. Sequential screening. *The Review of Economic Studies*, 67(4):pp. 697–717, 2000.
- [43] Jacques Cramer, Yossi Spiegel, and Charles Zheng. Optimal selling mechanisms with costly information acquisition. Technical report, 2003.
- [44] Dasgupta, Hammond, and Maskin. The implementation of social choice rules: some results on incentive compatibility. *Review of Economic Studies*, 46, 1979.
- [45] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a Nash equilibrium. In *ACM Symp. on Theory of Computing*, pages 71–78, 2006.
- [46] Theo S.H. Driessen. A note on the inclusion of the kernel in the core of the bilateral assignment game. *International Journal of Game Theory*, 27, 1998.
- [47] Benjamin Edelman, Michael Ostrovsky, and Michael Schwartz. Internet advertising and the generalized second price auction: Selling billions of dollars of keywords. *American Economic Review*, 97(1):242—259, 2007.
- [48] Eyal Even-Dar, Michael J. Kearns, and Jennifer Wortman. Sponsored search with contexts. In *Internet and Network Economics, Third International Workshop (WINE)*, pages 312–317, 2007.
- [49] Ulrich Faigle, Walter Kern, and Jeroen Kuipers. *An efficient algorithm for nucleolus and prekernel computation in some classes of TU-games*. PhD thesis, Memorandum No. 1464, University of Twente, 1998.
- [50] D. Gale and L. S. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.
- [51] D. Gale and L. S. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.
- [52] Arpita Ghosh, Hamid Nazerzadeh, and Mukund Sundararajan. Computing optimal bundles for sponsored search. In *Internet and Network Economics, Third International Workshop (WINE)*, pages 576–583, 2007.
- [53] A. Gibbard. Manipulation of voting schemes: a general result. *Econometrica*, 41:211–215, 1973.

- [54] Avi Goldfarb and Catherine Tucker. Online display advertising: Targeting and intrusiveness. *Marketing Science*, 30(3):389–404, 2011.
- [55] Avi Goldfarb and Catherine Tucker. Privacy regulation and online advertising. *Management Science*, 57(1):57–71, 2011.
- [56] J. Hartline and A. Karlin. Profit maximization in mechanism design. In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 13, pages 331–362. Cambridge University Press, 2007.
- [57] Jason D. Hartline. Lectures on optimal mechanism design. Lecture notes, 2006.
- [58] B. Holmstrom. *On incentives and control in organizations*. PhD thesis, Stanford University, 1977.
- [59] Kamal Jain. A polynomial time algorithm for computing an Arrow-Debreu market equilibrium for linear utilities. *SIAM J. Comput.*, 37(1):303–318, 2007.
- [60] Yashodhan Kanoria. An FPTAS for bargaining networks with unequal bargaining powers. In *Workshop on Internet and Network Economics*, 2010.
- [61] Yashodhan Kanoria, Mohsen Bayati, Christian Borgs, Jennifer T. Chayes, and Andrea Montanari. Fast convergence of natural bargaining dynamics on exchange networks. In *ACM Symp. on Discrete Algorithms*, 2011.
- [62] Jon Kleinberg and Eva Tardos. Balanced outcomes in social exchange networks. In *ACM Symp. on Theory of Computing*, 2008.
- [63] Jean-Jacques Laffont and Quang Vuong. Structural analysis of auction data. *The American Economic Review*, 86(2):414–420, may 1996.
- [64] Kate Larson. Reducing costly information acquisition in auctions. In *AAMAS*, pages 1167–1174, 2006.
- [65] Kate Larson and Tuomas Sandholm. Strategic deliberation and truthful revelation: an impossibility result. In *ACM Conference on Electronic Commerce*, pages 264–265, 2004.
- [66] Andrew Lazarus, Daniel E. Loeb, James G. Propp, Walter R. Stromquist, and Daniel H. Ullman. Combinatorial games under auction play. *Games and Economic Behavior*, 27, 1999.
- [67] Andrew Lazarus, Daniel E. Loeb, James G. Propp, and Daniel H. Ullman. Richman games. *Games of No Chance*, 29:439–449, 1996.
- [68] David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov Chain Mixing Times*. Academic Mathematical Society, London, 2009.
- [69] David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov Chain Mixing Times*. Academic Mathematical Society, London, 2009.

- [70] Johathan Levin and Paul Milgrom. Online advertising: Heterogeneity and conflation in market design. *American Economic Review, Papers and Proceedings*, 100(2), 2010.
- [71] Randall A Lewis and David H Reiley. Does retail advertising work? measuring the effects of advertising on sales via a controlled experiment on yahoo! Yahoo! Research Working Paper, 2011.
- [72] Michael Maschler. The bargaining set, kernel, and nucleolus. In R.J. Aumann and S. Hart, editors, *Handbook of Game Theory with Economic Applications*, volume 1, chapter 18, pages 591–667. Elsevier, July 1992.
- [73] Preston McAfee. The design of advertising exchanges. *Review of Industrial Organization*, 39(3):169–185, 2011.
- [74] Preston McAfee, John McMillan, and Michael D. Whinston. Multiproduct monopoly, commodity bundling, and correlation of values. *Quarterly Journal of Economics*, 104(2):371–384, 1989.
- [75] R. Preston McAfee, Kishore Papineni, and Sergei Vassilvitskii. Maximally representative allocations for guaranteed delivery advertising campaigns. Yahoo! Research Working Paper, 2010.
- [76] Paul Milgrom and Chris Shannon. Monotone comparative statics. *Econometrica*, 62(1):157–180, 1994.
- [77] Paulo Klinger Monteiro and Benar Fux Svaiter. Optimal auction with a general distribution: Virtual valuation without densities.
- [78] S. Muthukrishnan. Ad exchanges: Research issues. manuscript, 2010.
- [79] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.
- [80] Roger B. Myerson. Incentive-compatibility and the bargaining problem. *Econometrica*, 47, 1979.
- [81] Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- [82] John Nash. The bargaining problem. *Econometrica*, 18:155–162, 1950.
- [83] Michael Ostrovsky and Michael Schwarz. Reserve prices in internet advertising auctions: A field experiment. Working Paper, 2009.
- [84] Mallesh Pai and Rakesh Vohra. Optimal auctions with financially constrained bidders. Working paper, 2009.
- [85] Daniel Paulusma. *Complexity Aspects of Cooperative Games*. PhD thesis, University of Twente, Enschede, the Netherlands, 2001.

- [86] Yuval Peres, Oded Schramm, Scott Sheffield, and David B. Wilson. Random-turn hex and other selection games. *Amer. Math. Monthly*, 114(5):373–387, 2007.
- [87] Yuval Peres, Oded Schramm, Scott Sheffield, and David B. Wilson. Tug-of-war and the infinity laplacian. *Journal of the AMS*, 22:167–210, 2009.
- [88] Yuval Peres and Scott Sheffield. Tug-of-war with noise: A game-theoretic view of the p -laplacian. *Duke Math Journal*, 145(1):91–120, 2008.
- [89] Jean-Charles Rochet and Philippe Chone. Ironing, sweeping, and multidimensional screening. *Econometrica*, 66(6):788–826, July 1998.
- [90] Sharon C. Rochford. Symmetrically pairwise-bargained allocations in an assignment market. *Journal of Economic Theory*, 34(2):262–281, 1984.
- [91] Sharon C. Rochford. Symmetrically pairwise-bargained allocations in an assignment market. *Journal of Economic Theory*, 34(2):262–281, 1984.
- [92] R.W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *Int. Journal of Game Theory*, 2:65–67, 1973.
- [93] Sujay Sanghavi, Dmitry M. Malioutov, and Alan S. Willsky. Linear programming analysis of loopy belief propagation for weighted matching.
- [94] Lloyd S. Shapley and Martin Shubik. The assignment game I: The core. *Int. Journal of Game Theory*, 1(2):111–130, 1972.
- [95] Lloyd S. Shapley and Martin Shubik. The assignment game I: The core. *International Journal of Game Theory*, 1(2):111–130, 1972.
- [96] Tamás Solymosi and Tirukkannamangai E. S. Raghavan. An algorithm for finding the nucleolus of assignment games. *Int. Journal of Game Theory*, 23(2):119–143, 1994.
- [97] David R.M. Thompson and Kevin Leyton-Brown. Dominant-strategy auction design for agents with uncertain, private values. In *Twenty-Fifth Conference of the Association for the Advancement of Artificial Intelligence (AAAI-11)*, 2011.
- [98] David Robert Martin Thompson and Kevin Leyton-Brown. Valuation uncertainty and imperfect introspection in second-price auctions. In *AAAI*, pages 148–153, 2007.
- [99] Hal Varian. Position auctions. *International Journal of Industrial Organization*, 25(7):1163–1178, 2007.
- [100] D. Willer, editor. *Network Exchange Theory*. Praeger, 1999.
- [101] David Williams. *Probability with Martingales*. Cambridge University Press, 1991.

- [102] Yifeng Yu. Maximal and minimal solutions of an Aronsson equation: L^∞ variational problems versus the game theory. *Var. & Par. Differential Eq.*, 37:63–74, 2010.