NORTHWESTERN UNIVERSITY

Simple Mechanisms and Behavioral Agents: Towards a Theory of Realistic Mechanism Design

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Computer Science

By

Emmanouil Pountourakis

EVANSTON, ILLINOIS

March 2017

ABSTRACT

Simple Mechanisms and Behavioral Agents: Towards a Theory of Realistic Mechanism Design

Emmanouil Pountourakis

Classic mechanism design studies the implementation of *optimal systems or processes* in the presence of multiple selfish and fully rational agents. The theory of mechanism design has been one of the most celebrated advances in the economics discipline, especially with regards to auctions Myerson (1981); Vickrey (1961); Clarke (1971); Groves (1973). Nevertheless, practical applications of mechanism design show a significant disconnect between theory and practice. For instance, anonymous second price auction and pricing is one of the most prevalent forms of online auctions despite what the theory of Myerson (1981) suggests for asymmetric environments, e.g. eBay auctions where bidding history and demographic information provides a very specific information to the seller. On another note, most behavioral experiments indicate that the agent behavior diverges from the rational agent model, especially in dynamic environments.

This thesis addresses this phenomenon by studying simple, widely prevalent auctions like anonymous pricing and developing and analyzing them through the lens of the theory of approximation. The first contribution involves the problem of selling a single to agents with independent but non-identically distributed values. We show that anonymous pricing is a constant approximation to the optimal auction. Our second contributions involves repeated sales of a non-durable good. There, economic theory suggests that in the absence of commitment power the seller generates no revenue. Our contribution was to construct an equilibrium for multiple buyers where the seller is constraint to not discriminate with

constant revenue approximation. Finally, this thesis considers the incorporation of a behavioral model, particularly a model for procrastination: present-bias. We introduce the notion of variable present bias and analyze the performance loss of a procrastination agent versus a rational counterpart.

Acknowledgements

First of all, I would like to thank my advisor Nicole Immorlica for her counsel and guidance throughout all my doctorate studies. Even before joining Northwestern, Nicole has been a great source of inspiration for me. As an advsior she always emphasized the importance of focusing on my weaknesses rather than my strengths and has greatly contributed to me growing holistically as a research. My PhD journey would have been impossible without her continuous support as an advisor, colleague, and a friend.

Secondly, I would like to thank my other committee members, Jason Hartline and Brendan Lucier. Jason has been for me almost like a second advisor. His high expectations and his insistence on excellence helped me numerous times to improve myself and has drastically shaped my perspective on research. Brendan has been an exceptional colleague he was an invaluable mentor during my stay at Microsoft research. Overall, my committee's guidance and productive criticism has been crucial for the quality of my research and the steering the direction of my thesis.

Additionally, I would like to thank all of my co-authors especially the people I co-authored the papers that this thesis is based on: Saeed Alaei, Nick Gravin, Jason, Nicole, Brendan, Rad Niazadeh, Sam Taggart, and Yang Yuan. I am really grateful for our productive collaborations. I would especially like to thank Sam, who apart from our extremely fruitful collaboration has been helpful and supportive.

I would like to express my deepest gratitude to the people at Microsoft Research New England that hosted me for two and a half years. The majority of the work of this thesis took place during my visit at Cambridge lab, where I had great opportunities to be exposed to large number of distinguished researchers.

Apart from Microsoft Research I had the wonderful opportunity to visit several places during my PhD. Particularly, I like to thank Shuchi Chawla for hosting me at University of Wisconsin for the Fall 2016 where I had awesome collaborations with her, Kira Goldner, Ben Miller, and Dimitris Paparas. I would like to thank Guido Schäfer for wonderful collaborating summer 2012 during my internship at CWI, as well as for the numerous times he invited me to continue our research.

Additionally, I would like to to thank my office mates at Northwestern, especially Yorgos Askalidis, Rachel Cummings, Aleck Johnsen, Nima Haghpanah, Darrell Hoy, Greg Stoddard, and Sam. Despite only spending the first three years of my PhD there they helped creating a welcome and encouraging environment. Also, I would like to thank other friend in the algorithmic game theory community that helped me through their friendship and support during this journey, particularly Rediet Abebe, Vasilis Gkatzelis, Vangelis Markakis, and Vasilis Syrgkanis.

Last but not least, I would like to thank my partner Niharika Sravan for the enormous source of help and inspiration, as well as my mother Katerina and my sister Yakinthi for always being supportive and understanding. Finally, I would like to dedicate this work to my father Ioannis Pountourakis, who unfortunately passed away during my studies in 2012 and has been the main source of inspiration for me pursuing my PhD degree.

Dedication

To my father

Table of Contents

ABSTRACT 3

Acknowledgements 5

Dedication 7

List of Figures 10

Chapter 1. Introduction 11

1.1. Beyond Traditional Mechanism Design 13

1.2. Overview of the Results 15

1.3. Organization of the thesis 19

Chapter 2. Preliminaries 21

2.1. Mechanism Design 21

2.2. Revenue maximization 24

2.3. Approximation 28

Chapter 3. Anonymous Pricing: Single-Item Auction 30

3.1. Preliminaries and Notations 34

3.2. Main Result 37

3.3. Upper-Bound Analysis 38

3.4. Lower-Bound Analysis 52

Chapter 4. Anonymous Pricing: Repeated Sales 55

4.1. Introduction 55

4.2. Model 60

4.3. Folk Theorem 63

4.4. Non-Robustness of One-Buyer Learning Equilibria 66

4.5. Two Buyers, One Item 70

4.6. Two Buyers and Digital goods 84

4.7. Discriminatory Pricing for Digital Goods 96

Chapter 5. Behavioral model: Procrastination 100

5.1. Related work 104

5.2. Model 105

5.3. Worst-Case Procrastination 110

5.4. Special Task Graphs 118

5.5. Conclusions and Open problems 126

References 126

Appendix A. Appendix to Chapter 3 130

A.1. Simulations 130

Appendix B. Appendix for Chapter 4 132

B.1. Alternate Sufficient Conditions for Eliminating Learning Equilibria 132

B.2. Multiple buyer equilibrium for a single item 133

B.3. Multiple buyer equilibrium for digital goods 142

List of Figures

3.1 Comparison of our bounds to bounds from the prior literature Where one number is given, the bound is tight; where a range is given, a tight analysis is unknown. The final row labeled "our bounds" are for non-identical distributions and are improvements on the previous row. 33

3.2 Replacing regular distribution F_i with triangular revenue curve distribution $Tri(v_i, q_i)$. 34

3.3 Revenue curve of distribution Tri($\bar{v}, \bar{q}$). 37

3.4 Revenue gap between mechanisms of study; * denotes new bounds. 38

3.5 Replacing regular distribution F_i with triangular revenue curve distribution Tri($\bar{v}_i, \bar{q}_i$). 41

3.6 Intersection of revenue curve and price line p. 41

5.1 Task completion graphs for the examples in Section 5.2. 108

5.2 Task completion graph for the protocol 120

A.1 The ratios of the revenues of various auctions and benchmarks (Table). 131

A.2 The ratios of the revenues of various auctions and benchmarks (Graph). 131

CHAPTER 1

Introduction

The rise of the internet has drastically impacted every aspect of our lives; social media changed the way we communicate and exchange ideas, search engines have become the main source of information retrieval, and e-commerce has been year-by-year a growing percentage of total commerce. Traditionally topics of resource allocation, decision making and market design have been the purview of economic theory. However, in the era of the internet there is a growing need for economic design that is scalable that also incorporates a multitude of naive agents. This has caused a shifting emphasis from mechanism design and economic analysis to simple design and analysis. Algorithmic mechanism design has been an emerging research field for the last two decades that seeks to address these emerging challenges by borrowing analytical tools from theoretical computer science.

A mechanism is an algorithm that collects input by individual agents and aggregates them to a collective decision that effects the whole community. Mechanism design differs from classic algorithm design in that the agents are self-interested; they will behave in a way to maximize their personal utility rather than the goal of the mechanism designer. However, classic mechanism design has primarily focused on optimal mechanisms. This has created a discrepancy between mechanisms that are predicted by theory and implemented in practice. For example, many optimal mechanisms are complicated, lacks robustness, or are simply computationally intractable. While many celebrated results in economic theory demonstrate that optimal mechanism may satisfy additional desired properties, e.g. the revenue maximizing mechanism in single-item auctions is deterministic despite allowing for randomized mechanisms, it lacks the tools to analyze and compare strictly sub-optimal mechanisms generated by a restricted feasibility space.

One of the major contribution of theoretical computer science to mechanism design is the introduction of *approximation*: a systematic analysis of the multiplicative performance loss of a target class of mechanisms versus their optimal counterparts. Algorithmic mechanism design has created a rich theory of approximation in economics, while initially the main goal has been to combat the computational intractability of optimal mechanisms, the focus has shifted in analyzing simple mechanisms even when computational constraints are not an issue. A quintessential example are auctions, where Myerson (1981) provides a robust and elegant characterization of the optimal auction for selling a single item, that is perhaps unexpectedly deterministic, and computationally efficient. Nevertheless, online auction platforms often restrict the users to simpler auctions than the theoretically optimal, e.g. eBay, uses second-price auction with an anonymous reserve.

The phenomenon of widely prevalent mechanisms that, according to economic theory are sub-optimal, has many potential explanations. For example, most realistic environments are asymmetric, and therefore we expect realistic mechanisms to take advantage of this by exhibiting discriminatory behavior. That they don't because implementing a flat out discriminatory mechanism may often be illegal or will be faced with a strong community backlash and may affect future interactions.

This highlights another potential reasoning for the disconnect of theoretical solutions and practical mechanisms: the modeling itself of a mechanism as a single interaction that occurs in a vacuum is unrealistic. In reality any interaction between the mechanism designer and the agents effect potential future interactions, which necessitates the study of more complicated dynamic settings. A rational agent understands that their present behavior impacts future outcomes and may adjust their behavior by taking sub-optimal actions today in order to benefit in the future. This in turn may cause the mechanisms that are optimal for the static setting to become sub-optimal in the dynamic setting.

Another reason optimal mechanisms are scarce in practice is not because of the choices or optimization space of the seller, but rather because of the behavior of the agents. Consider the well studied cognitive bias of anchoring, a human tendency to be influenced by irrelevant numbers. For example a mechanism designer selling multiple items may sacrifice few of them, offering them at outlandish prices in order to increase the revenue generated by the rest. In dynamic settings, behavioral biases are often even more pronounced. For example, economic experiments suggest agents often display present-bias, over-weighting present outcomes with respect to future ones.

What is the right explanation for the disparity between theoretical predictions and practical mechanisms? It is most likely a combination of the reasons outlined above. The purpose of this thesis is to demonstrate cases where approximation can be used to bridge theory and practice, showcasing that the loss induced by accounting for these reasons is, in many cases, not that great.

1.1. Beyond Traditional Mechanism Design

Traditional mechanism design's main question is *What is the optimal mechanism assuming the participating agents are perfectly rational and self-interested?* As we described earlier, this thesis challenges this paradigm. We analyze the interaction of mechanism and agents through the lens of approximation hoping to establishing a more realistic theory of mechanism design.

Beyond Optimal Mechanisms. While it is fair to assume that practical mechanisms strive to optimize their given objective, they are usually complicated and cannot easily be expressed by simplistic models. For example many models fail to account for hidden costs or additional attributes of sub-optimal mechanisms that are not directly effecting the objective. For example, privacy is a very important property in an online environment and agents are becoming increasingly privacy-aware. Pricing mechanisms despite being inferior in terms of revenue extraction to auction-like mechanisms require the agents to reveal much less information. For example, participating in an ascending auction for a single item may require an agent to fully reveal their value. In contrast if the agent is offered a take-it-or-leave it price they would only reveal partial information. What is the revenue trade-off for the mechanism designer if they decide to go with the option that increases privacy to the bidders? While traditional mechanism design fully characterizes the optimal anonymous price and the optimal reserves for the second price auctions individually, it lacks the tools to analyze the performance loss by restricting to an anonymous posted pricing over second price auctions with reserves. In contrast, algorithmic mechanism design address this question using the *approximation ratio*, a well-established tool in algorithmic design theory, that captures the worst case multiplicative performance gap between these objects.

The motivation behind agents seeking privacy is that the outcome of the mechanism with effect future interactions. This highlights an additional challenge of the predictive power of the static single-shot mechanism design analysis. Traditional mechanism design has a vast literature studying dynamic environments. However, it lacks the elegant and succinct results it provides for their static counterparts. Many times the models are analytically intractable, optimal mechanisms are very complicated, and also sometimes the results are counter-intuitive to what happens in practice. Simpler mechanisms may alleviate the intrinsic difficulties of dynamic settings and using the approximation ratio we can argue about the severity of these assumptions. More surprisingly simple mechanisms may out-perform their optimal counterparts in dynamic settings. Consider the example of an agent concerned about their privacy in a repeated interaction with the mechanism. The dynamic environment quantifies the value of privacy in the following way: any information leaked by the agent will be used by the mechanism in the future and potentially reduce the utility of the agent. On one hand, an aggressive mechanism will cause the agent to be strategic about leaking information (e.g. Coase conjecture Coase (1972)). On the other hand, the agent may be less strategic when interacting with a mechanism that is sub-optimal and less aggressive in obtaining information, in a way that may result in better performance than their unrestricted optimal counterpart.

Beyond Rational Agent Behavior. Especially in dynamic settings both the mechanism and the agents are required to solve complicated optimization problems to figure out their optimal strategies. Even if agents had infinite computational power, the assumption of perfect rationality begins to crumble in the face of inter-temporal payoffs: behavioral science shows that humans often exhibit irrational behaviors like time-inconsistency. An optimal mechanism for such agents shall exploit any irrational behavior in order to maximize the given objective. One typical example are credit card contracts where various literature has demonstrated that different plans often are tailored to exploit time-inconsistent behavior Meier and Sprenger (2010) and Ru and Schoar (2016).

Main Question. *Can we use tools of approximation and worst case analysis to provide a theoretical model to identify mechanisms that justifies existing mechanisms; has predictive power for mechanisms that arise in practice; and incorporates into the analysis behavioral models indicated by existing research.*

Approximation. As described earlier, this thesis addresses these questions by studying mechanism design under a restricted space and comparing to a general optimization through the approximation ratio. This approach can serve two purposes. First, we observe a particular property or subclass mechanisms that are implemented in practice and establish a positive correlation between the prevalence of these mechanisms and their approximation ratio. Particularly, we are going focus on anonymous pricing in the static environment of *single-item auction* and the dynamic environment of *repeated sales*. The second purpose of our analysis is to consider cases where agent do not behave perfectly rationally. Here optimal mechanism design often attempts to exploit that irrational behavior for the objective of the mechanisms. Our goal here is instead to design the restricted space for the mechanism so this exploitative power is reduced. Specifically, we consider designing these restrictions when agents display time-inconsistency in their behavior, where an agent overweights present outcomes compared to future ones

1.2. Overview of the Results

Anonoymous Pricing: Single item auctions. Our first contribution considers the problem of revenue maximization in the single-shot single item setting, that is a single item to be auctioned among several agents just once. In this setting the optimal auction is given by Myerson's analysis Myerson (1981). We show that anonymous posted pricing, which is optimal for a single buyer, and second price auction with anonymous reserve, which is optimal for multiple buyers with identical and regular distributions, remain approximately optimal with multiple non-identical buyers.

Instead of comparing to the optimal auction, we use the well known upper-bound benchmark of the *ex ante relaxation* which relaxes the feasibility constraint to bind in expectation rather than point-wise. In this case, this relaxation requires that the expected number of units sold must be at most one. One of the core properties of ex-ante relaxation is that for buyers with non-identical and regular distributions the optimal solution to the ex ante relaxations structurally simple; it is a (discriminatory) posted pricing. This is key to our analysis as comparing anonymous to discriminatory pricing becomes tractable.

We show that for independent and regular distributions posting an anonymous price is an $e \approx 2.718$ approximation to the ex ante relaxation, i.e., the expected revenue of the optimal ex ante relaxation is at most a multiplicative e factor more than the revenue of the optimal anonymous posted pricing. This also implies that second price with optimal anonymous reserve is also an $e \approx 2.718$ approximation. This is the first improvement on the upper bound on the approximation ratio of anonymous pricing (4 to e) in the last half-decade. Our analysis is tight but since ex-ante relaxation is an upper bound on the optimal auction the lower bound on the approximation ratio does not carry over. However, simulation results show that anonymous pricing is at most 2.23 approximation while they don't show any improvement for second-price with anonymous reserves to the previously known lower bound of 2. Finally, we also prove n buyers with values drawn from non-identical and irregular distributions, neither anonymous pricing nor auctions with anonymous reserves are better than an n approximation. This indicates how crucial the regularity assumption is for approximation. These results appear in Chapter 3 and are based on the paper Optimal Auctions vs Anonymous Pricing Alaei et al. (2015).

Anonymous Pricing: Repeated Sales. Our next contribution considers the dynamic problem of repeated sales. There is a single seller, and each day the seller has a single copy of a good to sell, there is a single buyer, who has a private value $v \geq 0$ for obtaining the good each day. The value v is drawn from a distribution known to the seller. We assume that the value does not change from one day to the next; the buyer has the same value for consuming the good on every day. Each day, the seller posts a take-it-or-leave-it price, and the buyer can choose to either accept or reject the offer. The seller is free to set each day's price however she chooses, given the past purchasing behavior of the buyer. On any day that the buyer rejects, the good expires and the seller must discard it. We assume there is an infinite horizon, that is the game is played for infinitely many rounds, and the utilities of the buyer and the seller are discounted by a parameter δ. The goal of the buyer is to maximize the sum of their discounted utility and for the seller to maximize the same of the discounted revenue.

We want to study perfect Bayesian equilibria, which requires that each decision taken by both the seller and the buyer, for any observed history of prices and purchases, is a best response to the anticipated future behavior of the other player, given the seller's belief about the private value. Perfect Bayesian equilibria captures the absence of a commitment device for the seller, in the sense that the buyer cannot for example commit to run the optimal single shot mechanisms (given by Myerson's characterization) in each round and ignore the information the buyer reveals by their previous actions. In fact, this turns out to be an upper bound on the revenue obtainable by any seller-buyer interaction in a perfect Bayesian equilibrium.

Sadly, prior work on this model consists of mostly negative results if the seller lacks commitment power. In particular, there exist a perfect Bayesian equilibrium in which the seller obtains trivial revenue and no knowledge is transferred between the buyer and the seller Devanur et al. (2015); Hart and Tirole (1988); Schmidt (1993). This conclusion is reminiscient of the Coase conjecture; the primary difference is that the Coase conjecture refers to a durable good that a buyer will purchase only once, whereas in our repeated sales version of the problem the good is perishable and can be repurchased each day Coase (1972). Our results make the situation even more dire: if we consider some natural refinements that have been studied previously (see Fudenberg and Tirole (1983) and Hart and Tirole (1988)) this equilibrium is unique.

We next turn to studying a multi-buyer variant of the problem. Suppose that there are two buyers, each buyer's value is drawn i.i.d from a known distribution, and these values are again fixed over all rounds. The seller still has a single copy of the good for sale, and sells that good by posting a single price each day. Each buyer independently chooses whether or not to purchase each day. If both buyers wish to purchase at the offered price on a given day, then one of the accepting buyers is chosen uniformly at random to make the purchase. We manage to construct an equilibrium with a natural explore-exploit form. The seller starts by setting a low price, and slowly raises the price over time as long as at least two buyers purchase in each round. We show that for two buyers, if the distribution over buyer valuations satisfies the standard monotone hazard rate (MHR) condition, our equilibrium yields for the seller a $1/3e^2$ discounted revenue they could achieve from committing to the Myerson optimal mechanism each round.

Finally, we consider the multi-buyer setting where there is no limit on supply: the seller chooses a price each round, each buyer chooses to accept or reject, and all buyers who accept receive a copy of the item. Again we show a similar explore-exploit equilibrium that has similar approximation as the single-item case if we exogenously restrict the seller to use anonymous pricing. More surprisingly we demonstrate that the latter restriction is crucial since if the seller is allowed to use discriminatory pricing in this setting the unique equilibrium provides trivial revenue to the seller. These results are presented in Chapter 4 and are based on the paper Repeated Sales with Multiple Strategic Buyers Immorlica et al. (2017).

Behavioral Agents: Present bias. Our final contribution extends the traditional mechanism design of perfectly rational agent behavior. Particularly, we analyze the time-inconsistent behavior of present-bias. When making a decision, the present-biased agent overestimates her present utility by a multiplicative factor. To study these questions, we adopt a graph-theoretic framework for task completion proposed by Kleinberg and Oren Kleinberg and Oren (2014). An agent starts at vertex s and seeks to reach vertex t in n. On each day i, if the agent's current present bias is b_i and current state is v, then the weight of the first edge along a $v - t$ path is multiplied by an extra factor of b_i, while all the remaining edges are evaluated according to their true weights. The agent chooses a minimal-weight path, under this distortion, and follows the first edge of that path. For a given task graph G and present-bias distribution $\mathcal{F}$, the *procrastination ratio* is the ratio between the expected total weight of the traversed path to the weight of the initial shortest path. We are interested in bounding the procrastination ratio as a function of n, the number of days until the deadline for the task.

Under this model, we can state our question more concretely: are there natural conditions that guarantee a low procrastination ratio? Our first result shows that, for any distribution, the graph with maximal procrastination ratio has exactly this form: on any day, the agent may either complete the task for a growing cost or procrastinate for free. The proof of this result involves a connection to optimal pricing theory: the problem of constructing a worst-case graph can be reduced to the problem of designing a revenue-optimal auction. The fact that on any day there is only one costly edge is a consequence of the fact that the optimal pricing menu for a single-parameter agent has a single deterministic option Myerson (1981).

The connection to pricing highlights the natural application of this model to study the case of contract designer seeking to maximize revenue by exploiting the present-bias behavior of the agent. In contrast to our previous contributions here we are interested in containing the exploitative behavior of the contracts designer by imposing regulations, i.e, restrictions on the solution space. We define two restrictions of the graph structure, the *bounded distance property*: if the weight of the shortest path from any node v to the target t is at most the weight of the shortest path from the initial node s to tm and *monotone distance property*: for any $s - t$ path $(s = v_0, v_1, \dots, v_{k-1}, t = v_k)$, the shortest path from v_i to t is decreasing in i. For the former, we show for any distribution $\mathcal{F}$, the procrastination ratio of a bounded shortest-path graph is at most linear in n, and that this bound is tight even for distributions that have a constant probability of $b_i = 1$. For the latter, we show that if the distribution over present-bias parameters has sufficient mass at biases close to 1, then the procrastination ratio of graphs with the monotone distance property is bounded by a constant. These results are presented in Chapter 5, and are based on the work in Procrastination with variable present bias Gravin et al. (2016).

1.3. Organization of the thesis

Chapter 2. This chapter introduces basic notions in game theory and mechanism design, (Bayes) Nash equilibrium, (Bayes) incentive compatible mechanisms. Section 2.2 presents basics of auction theory, introduces the notion of revenue curves, defines regular distributions and presents Myerson's characterization of optimal auctions, and its corollaries. Finally, Section 2.3 formally defines approximation for revenue maximization and introduces the ex-ante relaxation of the optimal auction, a very useful tool that has been used extensively in approximation results in mechanism design.

The remaining chapters present the contribution of each of the three papers.

Chapter 3. This chapter describes the results of the paper Optimal Alaei et al. (2015) Alaei et al. (2015). Section 3.3 analyzes the worst case distribution for the comparison of the ex-ante relaxation with optimal anonymous pricing by expressing the approximation ratio of anonymous posted pricing as a mathematical program and upper bounds its solution by e. Section 3.4 identifies a sequence of solutions to this mathematical program that approximates e arbitrarily close, implying that the bound is tight. Finally, in Appendix A Section A.1 contains several simulation results that improve the lower bound of anonymous pricing and the option to 2.23.

Chapter 4. This chapter presents the paper Immorlica et al. (2017) Immorlica et al. (2017). Sections 4.3 and 4.4 focus on the single-buyer single-seller interaction and demonstrate the existence of a Folk theorem, and after introducing natural equilibrium refinements, shows that the only remaining equilibrium involves no learning and trivial revenue for the seller. Sections 4.5 and 4.6 study the case of two buyers for a single item and digital goods respectively and shows existence of a natural explore-exploit equilibrium in both cases with non-trivial revenue guarantees if the seller is restricted to anonymous pricing. Finally, Section 4.7 demonstrates that anonymous pricing is crucial for this results as the no-learning equilibrium becomes unique under discriminatory pricing.

Chapter 5. This chapter describes the paper Gravin et al. (2016) Gravin et al. (2016). Section 5.3 analyzes the worst-case task graph for a present-bias agent establishing a connection with optimal pricing. This characterizes the worst-case procrastination ratio given any present-bias distribution. Section 5.4 focuses on two restrictions on the task graph, bounded-distance and monotone distance and showcase the effects of these properties on the procrastination ratio.

CHAPTER 2

Preliminaries

2.1. Mechanism Design

Generally a mechanism design instance is defined as follows. The mechanism designer has to pick an outcome x from a set of alternatives $\mathcal{X}$. There are n agents and each agents $i = 1, \ldots, n$ has a private valuation $v_i(\mathbf{x})$ for every outcome $\mathbf{x} \in \mathcal{X}$, where $v_i \in \mathcal{V}_i$ comes from a known family of valuation functions. The agents are assumed to be rational agents seeking to maximize their personal utility. In addition, the mechanism can transfer utility among the agents via payments $\mathbf{p} = (p_1, \ldots, p_n)$. Given outcome $\mathbf{x}$ and payments $\mathbf{p}$ the utility of player i is given by $v_i(x) - p_i$. The mechanism designer seeks to maximize an objective, e.g., social welfare is defined as $\sum_{i=1}^{n} v_i(x)$.

A typical example is the single item auction. Here $\mathbf{x}$ defines the allocation of the item to the agents and can be represented by a n-dimensional vector $(x_1, \ldots, x_n)$ where x_i represents the probability that agent i gets the item. An allocation is *feasible* if $\sum_{i=1}^{n} x_i \leq 1$. The valuation of an agent is given by $v_i(\mathbf{x}) = v_i \cdot x_i$, where v_i represents the value of agent i for the item. For the rest of this chapter we will focus on this simple setting.

A common set of mechanisms is the set of *direct revelation* mechanisms where agents are assumed to report their valuation function. Let $\mathbf{v} = (v_1, v_2, \ldots, v_n)$, we define as $\mathbf{x}(\mathbf{v}) \in \mathcal{V}$ the allocation function of the mechanism and $\mathbf{p}(\mathbf{v}) \in \mathbb{R}^n$ the payment function. An example of a direct revelation mechanism for the previously described auction setting is the second-price auction, where agents report their values to the mechanisms, and the highest bidder wins and pays the value of the second highest bid.

Typically we assume that the mechanism designer has some prior knowledge about the agent's valuations. Particularly, we assume that the value v_i of agent i is drawn from a publicly known distribution

F_i. In these situations the mechanism designer seeks to maximize a given objective in expectation over the values of the agents according to the prior distributions.

2.1.1. Equilibria of mechanisms

A natural solution concept for the outcome of a mechanism is an *equilibrium*: that is a profile of strategies where the agents cannot improve their personal utility by deviating to a different strategy. In a direct revelation mechanism the strategy σ_i for agent corresponds to a mapping from real valuation v_i to another value $\sigma_i(v_i)$. We say that a set of strategies form an equilibrium if each strategy is a best response to the strategies of the rest of the agents.

The precise definition of the equilibrium may vary depending on our assumptions on the agents' information. *Nash equilibrium* assumes that all agents are fully informed on each others private information.

Definition 2.1.1. A profile of strategies $\sigma_1, \sigma_2, \ldots, \sigma_n$ forms a *Nash equilibrium* if for all $\mathbf{v} \in \mathcal{V}$ it holds that for every $i \in 1, \ldots, n$:

$$v_i \cdot x(\sigma_i(v_i), \sigma_{-i}(v_{-i})) - p_i(\sigma_i(v_i), \sigma_{-i}(v_{-i})) \geq v_i \cdot x(v_i', \sigma_{-i}(v_{-i}))) - p_i(v_i', \sigma_{-i}(v_{-i}))$$

A *Bayes Nash equilibrium* assumes that each strategy σ_i is a best response to strategies σ_{-i} in expectation using the prior F_{-i}

Definition 2.1.2. A profile of strategies $\sigma_1, \sigma_2, \ldots, \sigma_n$ forms a *Bayes Nash equilibrium* if for every $i \in 1, \ldots, n$ and every $v_i \in \mathcal{V}_i$ it holds that:

$$\mathop{\mathbb{E}}_{v_{-i} \sim F_{-i}} [v_i \cdot [x(\sigma_i(v_i), \sigma_{-i}(v_{-i}))] - p_i(\sigma_i(v_i), \sigma_{-i}(v_{-i}))] \geq \mathop{\mathbb{E}}_{v_{-i} \sim F_{-i}} [v_i \cdot x(v_i', \sigma_{-i}(v_{-i}))) - p_i(v_i', \sigma_{-i}(v_{-i}))]$$

Both Nash and Bayes Nash equilibrium require that the strategy chosen is a best response to the strategies of the rest agents. The former requires this to be true even if each agent is aware of each other's valuations while the latter relaxes this assumption where each agent has only knowledge about everyone else's value through the common prior. Since Bayes-Nash equilibrium condition is a relaxation of the Nash equilibrium condition, every Nash-equilibrium is Bayes-Nash equilibrium but no the other way around.

2.1.2. Incentive Compatibility

An alternative method of analyzing mechanisms is to restrict attention to *truthful mechanisms* that is mechanisms where $\sigma_i(v_i) = v_i$ for all i forms an equilibirum. We say that a mechanism is *incentive compatible* if it is in their best interest to reveal their private value. Similarly to equilibria there are various definitions of incentive compatibility depending on the assumptions on agent's information.

Definition 2.1.3. A mechanism $\mathbf{x}, \mathbf{p}$ is *dominant strategy incentive compatible* if for every bid b_i and values $v_{i-1} = v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_n$ it holds that

$$x_i(v_i, v_{-i}) \cdot v_i - p_i(v_i, v_{-i}) \geq x_i(b_i, v_{-i}) \cdot v_i - p_i(b_i, v_{-i}),$$

in other words no matter what the reports of the other agents are it is best for agent i to report their true value v_i. In addition we may also relax the incentive constraints to Bayesian incentive compatibility

Definition 2.1.4. A mechanism $\mathbf{x}, \mathbf{p}$ is *Bayesian incentive compatible* if for every bid b_i to holds that

$$E_{v_{-i} \sim F_{-i}}[x(v_i, v_{-i}) \cdot v_i - p_i(v_i, v_{-i})] \geq E_{v_{-i} \sim F_{-i}}[x(v_i, v_{-i}) \cdot v_i - p_i(v_i, v_{-i})],$$

Since the Bayesian incentive constrained is a relaxed version of the dominant-strategy incentive constraint, every dominant-strategy incentive compatible mechanism is also Bayesian incentive compatible but not necessarily the other way around.

2.1.3. Revelation Principle

Incentive compatible mechanism are a proper subset of the set of feasible mechanisms. This brings up the question of whether restricting to incentive compatible mechanism also reduces the set of implementable outcome. The *revelation principle* establishes that this is not the case.

Proposition 2.1.1. *Let* $\mathbf{x}, \mathbf{p}$ *be a mechanism and* $\sigma_1, \ldots, \sigma_n$ *be a set of strategies that form a (Bayes) Nash equilibrium. There exists a dominant strategy (Bayesian) incentive compatible mechanism* $\mathbf{x}', \mathbf{p}'$ *where truthful reporting implements the same outcome as the equilibrium of mechanism* $\mathbf{x}, \mathbf{p}$ *induced by* $\sigma_1, \ldots, \sigma_n$*:*

$$x_i(\sigma_i(v_i), \sigma_{-i}(v_{-i})) = x'_i(v_i, v_{-i}) \qquad\qquad p_i(\sigma_i(v_i), \sigma_{-i}(v_{-i})) = p'_i(v_i, v_{-i})$$

2.1.4. Individual Rationality

A common assumption in mechanism design is that agents voluntary participate in the mechanism. This can be alternatively be interpreted as if they have the option of leaving the mechanism obtaining a utility of zero.

Definition 2.1.5. A mechanism is *individually rational* if for every agent i and value $\mathbf{v}$,

$$x_i(v_i, v_{-i}) \cdot v_i - p_i(v_i, v_{-i}) \geq 0$$

While we can define the more relaxed Bayesian version of individual rationality we focus on the ex-post version even when studying Bayesian incentive compatible mechanisms.

2.2. Revenue maximization

For the rest of the thesis we focus on Bayes Nash equilibria and Bayesian incentive compatible mechanism. To ease notation we use $x_i(v) = E_{v_{-i} \sim F_{-i}}[x_i(v_i, v_{-i})]$ and $p_i(v) = E_{v_{-i} \sim F_{-i}}[p_i(v_i, v_{-i})]$. There has been an important line of work trying to characterize the optimal mechanisms for a given objective subject to incentive constraints. The revenue of a incentive compatible mechanism with pricing function p is given by.

$$\text{Revenue: } \sum_{i=1}^{n} p_i(v_i).$$

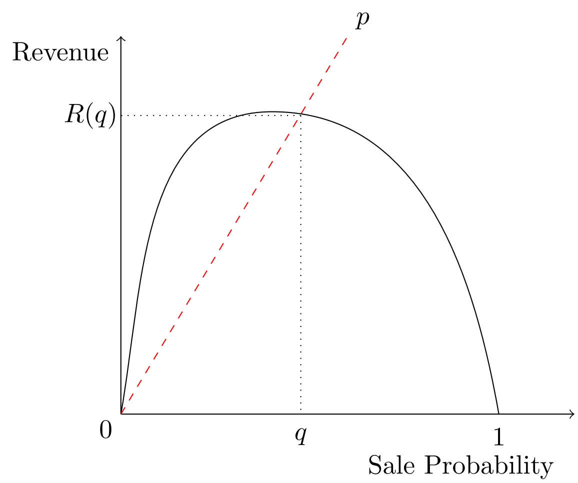

2.2.1. Singe Agent and Revenue Curves

To warm up, consider the problem of a single buyer with distribution F, and a pricing mechanism that offers it at a take-it-or-leave-it price p. Clearly this is a dominant strategy incentive compatible mechanism. The revenue of offering at price p can be simply be computes as $p \cdot (1 - F(p))$. Alternatively, we can analyze the revenue of such a mechanism in *quantile space* : Selling to the agent with probability q is achieved by posting price $p = F^{-1}(1 - q)$. The *revenue curve* of distribution F is defined as $R(q) = q \cdot F^{-1}(1-q)$, captures the revenue generated by posting a price that is accepted with probability q.

Example 1. Consider the case of $F = U[0, 2]$ the uniform distribution with support $[0, 2]$, the revenue of price p is given by $p \cdot (1 - p/2)$ and is maximized at $p = 1$. In quantile space the revenue curve becomes $R(q) = 2 \cdot q \cdot (1 - q)$ which is maximized at $q = 1/2$ (that corresponds to price $p = 1$).

The revenue curve retains all the information for the distribution. Particularly it is easy given a price p to find the probability $1 - F(p)$ by drawing a line with slope p, finding the intersection with the revenue curve and projecting it to the x axis. (see Figure 3.6). Revenue curves play an important role in our understanding of revenue-optimal auctions.

Nevertheless, revenue curves capture the revenue obtained by deterministic mechanisms, that is mechanisms who either allocate with probability one or zero. The natural question that arises is whether there exists a non-deterministic mechanism that allows for randomized allocations obtain more revenue. Myerson's characterization Myerson (1981) shows this is not the case, and even for multiple agents the optimal auction is deterministic.

2.2.2. Myerson characterization

Myerson Myerson (1981) characterized the optimal auction that maximizes expected revenue subject to individual rationality and Bayesian incentive compatibility.

Theorem 2.2.1 (Myerson's characterization Myerson (1981))**.** *Let* $\mathbf{x}, \mathbf{p}$ *be an individually rational and Bayesian incentive compatible mechanism. The following are true*

(1) For each agent i*, interim allocation* $x_i(v_i)$ *is monotone non decreasing in* v_i

(2) For each agent i*, interim payment* $p_i(v_i)$ *is given by*

$$p_i(v_i) = x_i(v_i) \cdot [v - \frac{1 - F_i(v)}{f_i(v)}]$$

One important property of this characterization is that what is referred to as the revenue equivalence, two Bayesian incentive compatible mechanisms with the same allocation, must have the same revenue.

Example 2. Consider the instance of a single agent with $F = U[0, 2]$. The virtual value function is $\phi(v) = 2v - 2$. The mechanism that charges at price p has $x(v) = 1$ for $v \geq p$ and $x(v) = 0$ otherwise. The expected social surplus

$$\int_{y=0}^{2} x(v)\phi(v)f(v)dv = \int_{1}^{2} v - 1 = 1/2$$

which matches the revenue of posting price p as we calculated earlier.

Also,one immediate corollary of this characterization is that for the single agent problem take-it-or-leave it price is optimal.

Corollary 2.2.2. *A deterministic pricing mechanism is optimal for selling one item to single agent.*

In Chapter 5 we use a slight generalization of this corollary, in particular that pricing is still optimal for objectives that are linear in revenue and the allocation probability.

The term of the right hand-side of Myerson's characterization is usually referred to as *virtual value* $\phi_i(v) = v - \frac{1-F_i(v)}{f_i(v)}$. Note that this is equal to the derivative of the revenue curve, that is $R'(F(v)) = \phi_i(v)$. Myerson's characterization can be interpreted as virtual value introduce a different way of counting the total revenue, while the ex-post revenue has been shifting around the different values. There are several important observations of Myerson's characterization.

While Myerson's characterization suggests that the optimal auction should allocate the item to the agent with the highest virtual value agent. However, this may not be feasible if the virtual value functions are not monotone. The distributional property that the corresponding virtual value is monotone we refer to as *regularity* (note that the corresponding revenue curves are concave). An instance is regular if all distributions are regular.[1]

Corollary 2.2.3. *Given a regular instance the mechanism that allocates to the highest virtual valued agent is Bayesian incentive compatible and maximizes revenue*

While pricing was optimal for a single-buyer this is no longer the case. However, consider the second-price with an anonymous reserve action. The highest valued agents wins the item subject to being higher than the fixed reserve and pays the maximum between the second highest value and the reserve.

Example 3. Consider the instance of a two agents with $F_1 = F_2 = U[0, 2]$. The common virtual value function is $\phi(v) = 2v - 2$ In a second price auction the probability that each agent wins the item with value v is exactly $v/2$, that is $x_i(v) = v/2$.

The expected virtual surplus

$$2 \cdot \int_{y=0}^{2} x(v)\phi(v)f(v)dv = \int_{1}^{2} \frac{v}{2} \cdot (v-1) = 10/12.$$

Corollary 2.2.4. *Given a regular and symmetric instance, i.e.* $F_1 = F_2 = \cdots = F_n = F$*, second price with anonymous reserve* $\phi^{-1}(0)$ *maximizes revenue.*

[1] In the case of irregular distributions, the process of *ironing* enables us to construct simple allocations that are incentive-compatible and maximize virtual surplus.

However, if we drop either the regularity or the symmetry assumption is it easy to notice that second price auction with anonymous reserve is no longer optimal. In chapter 2 we investigate the question of how well second price auction with anonymous reserve performs in comparison with the optimal auction via the approximation ratio which we formally define in the subsequent section.

Remark 1. *Another interesting observation is that the optimal auction satisfies dominant strategy incentive compatibility despite the fact that the solution space of the optimization is Bayesian incentive compatible mechanism. In addition, for regular instances the optimal auction is also deterministic.*

2.3. Approximation

Approximation theory studies sub-optimal algorithms/mechanisms and compares them to their optimal counterparts in a worst-case manner. Formally, let A be the target mechanism and let A-REV(I) denote the revenue obtained by mechanism A in instance $I = F_1, F_2, \ldots, F_n$. Let OPT denote the optimal auction and OPT-REV(I) the revenue obtained by the optimal mechanism given instance I.

Definition 2.3.1. Let $\mathcal{I}$ be any subset of instances. The approximation of mechanism A for instances $\mathcal{I}$ is given

$$\max_{I \in \mathcal{I}} \frac{\text{OPT-REV}(I)}{\text{A-REV}(I)}$$

While this is the most canonical definition of approximation one may consider comparing a mechanism A to an arbitrary mechanism B such that A-Rev(I) $\leq$ B-Rev(I) for all $I \in \mathcal{I}$. An important observation is that the approximation of optimal anonymous pricing is always smaller than the approximation of second price auction with optimal anonymous reserve. The revenue of an anonymous pricing mechanisms that sets a price of p is less than the revenue of second price with that uses the same anonymous reserve as the price p; the added competition can only increase revenue.

Proving an approximation result for optimal anonymous pricing provides a lower bound for the approximation ratio of second price auction with an optimal anonymous reserve. This becomes handy as analyzing the revenue of anonymous pricing is quite simple. Despite of Myerson's characterization's elegance the optimal auctions may still be a difficult object to analyze. A frequently used approach is literature is to consider an upper bound on the optimal auction.

2.3.1. Ex-ante relaxation

The set of feasible mechanisms for a single-item auctions consists of mechanisms that sell the item to at most one agent. Consider relaxing this constraints and allowing the seller to sell to more than one person as long as long as the expected number of items sold is less than 1.

Example 4. Consider the instance of two agents with $F_1 = U[0,1]$ and $F_2 = U[0,2]$ and consider the following mechanism: Offer a take-it-or-leave-it price offer $p_1 = 1/2$ and $p_2 = 1$ to first and second agent respectively. The revenue generated by the ex-ante implementation is equal to $3/4$. If instead we tried to implement this pricing scheme in the original setting by offering it first to agent 2 at price $p_2 = 1$ and if agent 2 rejected we offered it to agent 1 at price $p_1 = 1$, then the revenue generated would be equal to $5/8$.

One of the most important upper-bound benchmark in revenue maximization is the ex-ante relaxation, that is the optimal mechanism where the feasibility constraint is satisfied in expectation. For the case of selling a single item this means the seller is allowed to sell the item more than once as long as the sum of the probabilities that each agent receives the item is less than one.

Consider the following optimization problem that uses revenue curves.

$$
\text{(2.1)} \qquad \textsc{ExAnteRev}(\mathcal{I}) \triangleq \quad \begin{aligned} \max \quad & \textstyle\sum_{i=1}^{n} R_i(q_i) \\ \text{subject to} \quad & \textstyle\sum_{i=1}^{n} q_i \leq 1 \\ & q_i \geq 0 \qquad \forall i \in \{1, \dots, n\}. \end{aligned}
$$

Proposition 2.3.1. *The program 2.1 evaluates the revenue of the optimal auctions in the ex-ante relaxation instance.*

CHAPTER 3

Anonymous Pricing: Single-Item Auction

This chapter revisits the quintessential problem in optimal mechanism design: How should a monopolist sell an item to buyers with private and independently distributed values? The most prevalent selling mechanisms are (anonymous) posted pricing and auctions with (anonymous) reserves. These mechanisms are optimal, respectively, when there is a single buyer and when the buyers values are identically distributed from a distribution that satisfies a natural convexity property (henceforth: regular), and non-optimal more generally (Myerson, 1981). This thesis applies the theory of approximation to justify the use of these mechanisms far beyond the settings where they are optimal, e.g., in eBay's auction where bidders are asymmetric as is distinguished by their public bidding history and reputation.

The literature on mechanism design has numerous successful stories where analyses of a principal's optimization problem, of selecting a mechanism that mediates the interaction between strategic agents, yields a simple solution that is both interpretable and observed in practice. Nonetheless, the models for which these sorts of results are obtained tend to be overly idealized. Moreover, embellishment of these ideal models tends to result in solutions to the mechanism design problem that are complex, difficult to interpret, and rarely observed in practice.

Ideal approximation, which quantifies the extent to which the conclusions of an ideal model hold approximately in a richer model, can greatly extend the applicability of optimal mechanism design. In mechanism design, ideal approximation affords two kinds of conclusions. First, it shows that the prediction of the ideal model is robust, not only does it smoothly degrade as the model becomes less ideal, but the degree of the degradation is universally bounded. Second, it allows embellishments of the model to be classified as less relevant details or as salient features. If the optimal mechanism for the ideal model remains approximately optimal when the model is embellished, then that embellishment is a detail. On the other hand, if the optimal mechanism for the ideal model can be far from optimal on the embellished model then the embellishment is a salient feature and must be modeled if the mechanism design problem is to be understood. The reason such an approximation result is important also generates the main difficulty in its proof: the optimal mechanism for an embellished model is complex. A standard approach for overcoming such a difficulty is to identify a relaxation of the optimization problem under which the optimal solution is structurally simpler.

This chapter studies the problem of a revenue-maximizing monopolist in the sale of a single item. We show that posted pricings, which are optimal for a single buyer, and auctions with anonymous reserve prices, which are optimal for multiple buyers with identical and regular distributions, remain approximately optimal with multiple non-identical buyers. The *ex ante relaxation* of a mechanism design problem relaxes the feasibility constraint to bind in expectation rather than point-wise. For a single-item auction, this relaxation requires that the expected number of units sold must be at most one. For buyers with non-identical and regular distributions the optimal solution to the ex ante relaxation is structurally simple; it is a posted pricing. We show that for independent and regular distributions posting an anonymous price is an $e \approx 2.718$ approximation to the ex ante relaxation, i.e., the expected revenue of the optimal ex ante relaxation is at most a multiplicative e factor more than the revenue of the optimal anonymous posted pricing. This bound is tight, specifically, we exhibit a distribution where it holds with equality. Of course an auction that uses the same price as an anonymous reserve obtains only higher revenue than the posted pricing so the auction with the optimal anonymous reserve price is also an e approximation to the ex ante relaxation. This latter bound is, however, not tight.

It is worthwhile to contrast the approximation bounds for that hold for regular distributions (recall: regular distributions satisfy a natural convexity property) to inapproximation bounds that hold for irregular distributions. Specifically, for irregular distributions the assumption of symmetry, i.e., that the bidders' values are identically distributed, is crucial. We prove that, for n buyers with values drawn from non-identical and irregular distributions, neither anonymous pricing nor auctions with anonymous reserves are better than an n approximation. (Moreover, these bounds are tight.) Consequently, as the number of buyers grows the relative performance of these simple mechanisms can degrade linearly.

These approximation and bounds afford the following economic conclusions. First, in comparing posted pricings with the optimal mechanism, there is a tradeoff between asynchronization and competition. The (asynchronous) posted pricing is easier and may thus attract more buyers where as the optimal mechanism seed increased revenues due to competition between buyers. Explicitly modeling a cost for synchronization of would likely give an even more complex environment with even more difficult to interpret conclusions. Our approximation bounds show that there are only modest gains from competition: conclusions from the ideal single-agent model continue to hold with multiple agents. Second, the potential asymmetry between buyers in the model is also of only modest importance: conclusions from the ideal model where buyers' values are identically distributed approximately hold with non-identical distributions. For both of these conclusions the ability to discriminate between high and low types via an anonymous posted price or reserve enables a significant fraction of the revenue to be obtained of an optimal mechanisms which engages in third-degree price discrimination. In contrast, the inapproximation bound for irregular distributions suggests that asymmetry and irregularity together cause challenging revenue optimization problems in auctions and if they are both present then significant revenue can be left on the table if they are not properly treated.

Related Work. See Hartline (2012) for a survey of recent work that applies the theory of approximation to problems in mechanism design; we describe here only a short summary. Chawla et al. (2007) consider a unit-demand multi-dimensional screening problem and identifies a simple item pricing that is approximately optimal. Hartline and Roughgarden (2009) show that discriminatory monopoly reserve pricing, for non-identical but regular agents, is often approximately optimal. For example, for a single-item the second-price auction with monopoly reserve prices is a two approximation to the revenue-optimal mechanism. In the same setting but without the regularity assumption and among other results, Chawla et al. (2010a) show that discriminatory posted pricings are a two approximation. Alaei (2011) makes explicit the concept of the ex ante relaxation and generalizes and improves a number of results for multi-dimensional screening from Chawla et al. (2010a). Yan (2011) gives a very elegant analysis of a number of single-dimensional multi-agent pricing problems and shows that sequential posted pricings (when the agents arrive in a prespecified order) approximate the optimal (simultaneous) mechanism.

	pricings		auctions	
	regular	irregular	regular	irregular
identical	$e/(e-1)$	2	1	2
non-identical	$[2, n]$	$[2, n]$	$[2, 4]$	$[2, n]$
our bounds	$[2.23, e]$	n	$[2, e]$	n

Figure 3.1. Comparison of our bounds to bounds from the prior literature Where one number is given, the bound is tight; where a range is given, a tight analysis is unknown. The final row labeled "our bounds" are for non-identical distributions and are improvements on the previous row.

For anonymous pricing and auctions with anonymous reserves for selling a single item the following approximation bounds are implied by the above literature (see Figure 3.1). The $e/(e-1)$ bound for anonymous pricing with identical and regular distributions is from Yan (2011) and Chawla et al. (2010a); a simple exercise shows that this bound is nearly tight. The upper bound of two for identical and irregular distributions follows from Chawla et al. (2010a); a simple exercise shows that this bound is tight. For non-identical and regular distributions, there is a standard bad example from Hartline and Roughgarden (2009) that gives the lower bounds of two. In the example there are two agents, one with value deterministically one, and the other with value drawn from the so-called *equal revenue distribution* (where every offered price results in an expected revenue of one). The optimal auction has expected revenue two whereas any anonymous pricing or auction with an anonymous reserve has an expected revenue of at most one. For regular distributions, the upper bound of four for auctions with anonymous reserves is a corollary of a main result of Hartline and Roughgarden (2009).

Further Implications. Our result relating anonymous pricing to the optimal auction has implications on mechanism design for agents with multi-dimensional preferences (e.g., for multiple items; cf. Chawla et al., 2007). Understanding these problems, though there has been considerable recent progress, remains an area with fundamental open questions for optimization and approximation. Recently, Haghpanah and Hartline (2015) proved the optimality of uniform pricing for a single unit-demand buyer with values drawn from a large family of item-symmetric distributions. An immediate corollary of our anonymous pricing result is that, for a unit-demand buyer with values drawn from an asymmetric product distribution, uniform pricing is an e approximation (improved from four) to the optimal non-uniform pricing (cf. Cai and Daskalakis, 2011) and, via a result of Chawla et al. (2010b), a $2e$ approximation to the optimal pricing over lotteries (i.e., randomized allocations, improved from eight). Further refinement of this latter bound remains an important open question. These approximation results for a single agent automatically improve the approximation bounds for related multi-agent mechanism design problems based on uniform pricing, e.g., from Alaei et al. (2013). As one example, for selling an object that can be configured on sale in one of m configurations to n agents with independently (but non-identically) distributed values for each configuration (also satisfying a regularity property), the second-price auction with an anonymous reserve that configures the object as the winner most prefers is a $2e^2 \approx 14.8$ approximation to the optimal auction (which is sometimes randomized; improved from 32).

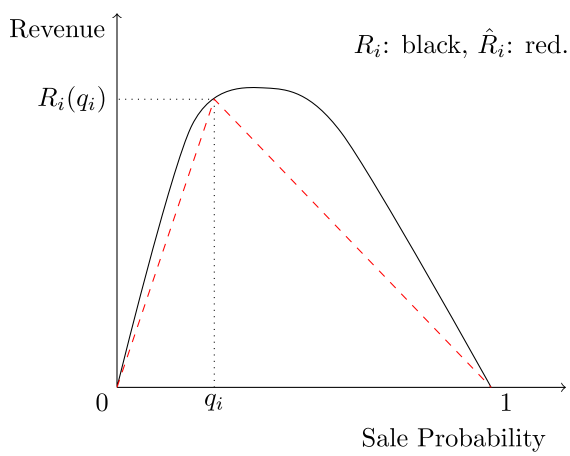

Figure 3.2. Replacing regular distribution F_i with triangular revenue curve distribution $Tri(v_i, q_i)$.

3.1. Preliminaries and Notations

We consider the classical independent private value model from auction theory. A seller has a single good. Agents have values drawn independently, but not necessarily identically from a distribution. We are interested in the revenue of the auction which is the sum of the payments of agents made. [1]

Different auctions of interest in this chapter. Our development of the approximation bound for anonymous pricing and reserves is based on the analysis of four classes of mechanisms:

[1]Most of the definitions here are consistent with Chapter 2 and are repeated for readability

(1) **Ex ante relaxation (a discriminatory pricing):** An ex ante pricing relaxes the feasibility constraint of the auction problem, from selling at most one item ex post, to selling at most one item in expectation over the draws of agents' values, i.e., ex ante. Fixing a probability of serving a given agent the optimal ex ante mechanism offers this agent a posted price irrespective of the outcome of the mechanism for the other agents. This relaxation was identified as a quantity of interest in Chawla et al. (2007) and its study was refined by Alaei (2011) and Yan (2011).

(2) **Auction:** An auction is any mechanism that maps values to outcome and payments subject to incentive and feasibility constraints. The optimal auction was characterized by Myerson (1981) and this characterization, though complex, is the foundation of modern auction theory.

(3) **Anonymous reserve:** An anonymous reserve mechanisms is a variant of the second-price auction where bids below an anonymous reserve are discarded, the winner is the highest of the remaining agents, and the price charged is the maximum of the remaining agents' bids or the reserve if none other remain.

(4) **Anonymous pricing**: An anonymous pricing mechanism posts an anonymous price and the first agent to arrive who is willing to pay this price will buy the item.

For any distribution over agents' values the optimal revenue attainable by each of these classes of mechanisms is non-increasing with respect to the above ordering. The final inequality of optimal anonymous reserve exceeding optimal anonymous pricing follows as with equal reserve and price, the former has only higher revenue as competition drives a higher price. Among the above mechanisms, ex ante relaxation does not generally give a feasible outcome, i.e., it may sell more than one item ex post, and is thus a quantity of interest only in as far as it gives an upper bound on the revenue of the optimal mechanism. While we might prefer to compare the performance of anonymous pricing directly to the optimal auction, such a direct analysis is difficult.

Revenue curves and regular instances. Each agent i has value drawn from a distribution F_i with cumulative distribution function (CDF) denoted by $F_i(\cdot)$. The *revenue curve* $R_i(q) = q \cdot F_i^{-1}(1-q)$ gives the expected revenue obtained by selling an item to agent i with probability exactly q, i.e., by posting price $F^{-1}(1-q)$. The agent is *regular* if its revenue curve $R_i(q)$ is concave in q. An n-agent instance

$\mathcal{I} = \{F_i\}_{i=1}^n$ is regular if each agent's distribution is regular. The family of all regular instances for all $n \geq 1$ is denoted by $\textsc{Reg}$.

Ex ante relaxation and optimal auction. The revenue of the *ex ante relaxation*, which allocates to one agent in expectation, gives an upper bound on the revenue of the optimal auction. For any instance $\mathcal{I}$, it can be easily expressed in terms of the revenue curves of the agents.

$$\textsc{ExAnteRev}(\mathcal{I}) \triangleq \begin{aligned} \max \quad & \textstyle\sum_{i=1}^n R_i(q_i) \\ \text{subject to} \quad & \textstyle\sum_{i=1}^n q_i \leq 1 \\ & q_i \geq 0 \qquad \forall i \in \{1, \ldots, n\}. \end{aligned} \tag{3.1}$$

Anonymous pricings and anonymous reserves. Consider an anonymous pricing that posts a price p. The expected revenue of this anonymous pricing for instance $\mathcal{I} = \{F_i\}_{i=1}^n$ is

$$\textsc{PriceRev}(\mathcal{I}, p) \triangleq p \cdot \Big(1 - \prod\nolimits_i F_i(p)\Big). \tag{3.2}$$

The expected revenue of the optimal anonymous pricing is

$$\textsc{OptPriceRev}(\mathcal{I}) \triangleq \max_{p \in \mathbb{R}_+} \textsc{PriceRev}(\mathcal{I}, p). \tag{3.3}$$

Worst-case approximation ratio. The main task of this chapter is to analyze the worst case ratio of the revenue of the ex ante relaxation to the revenue of the optimal anonymous pricing over all regular instances, that is

$$\rho \triangleq \sup_{\mathcal{I} \in \textsc{Reg}} \frac{\textsc{ExAnteRev}(\mathcal{I})}{\textsc{OptPriceRev}(\mathcal{I})}, \tag{P1}$$

where $\textsc{Reg}$ denotes the space of all regular instances.

Triangular revenue curve instances. We will show that distributions with triangular-shaped revenue curves give worst case instances for program (P1). A *triangular revenue curve* distribution, denoted

$\text{Tri}(\bar{v}, \bar{q})$ with parameters $\bar{v} \in (0, \infty)$ and $\bar{q} \in [0, 1]$, has CDF given by

$$F(p) = \begin{cases} 1 & p \geq \bar{v} \\ \frac{p \cdot (1-\bar{q})}{p \cdot (1-\bar{q}) + \bar{v}\bar{q}} & 0 \leq p < \bar{v} \end{cases} \qquad \forall p \in \mathbb{R}_+. \tag{3.4}$$

The revenue curve corresponding to the above distribution has the form of a triangle with vertices at $(0, 0)$, $(\bar{q}, \bar{v}\bar{q})$, and $(1, 0)$ as illustrated in Fig. 3.3; the revenue curve's concavity implies that the distribution is regular. Note that the CDF is discontinuous at $\bar{v}$ which corresponds to a pointmass of $\bar{q}$ at value $\bar{v}$. A *triangular revenue curve instance* is given by $\mathcal{I} = \{\text{Tri}(\bar{v}_i, \bar{q}_i)\}_{i=1}^n$ with $\sum_{i=1}^n \bar{q}_i \leq 1$; with respect to it the revenue of anonymous pricing p and the ex ante relaxation are given by

$$\textsc{PriceRev}(\mathcal{I}, p) = p \cdot \left(1 - \prod\nolimits_{i: \bar{v}_i \geq p} \left(1 + \frac{\bar{v}_i \bar{q}_i}{p \cdot (1 - \bar{q}_i)}\right)^{-1}\right), \tag{3.5}$$

$$\textsc{ExAnteRev}(\mathcal{I}) = \sum\nolimits_{i=1}^n \bar{v}_i \bar{q}_i. \tag{3.6}$$

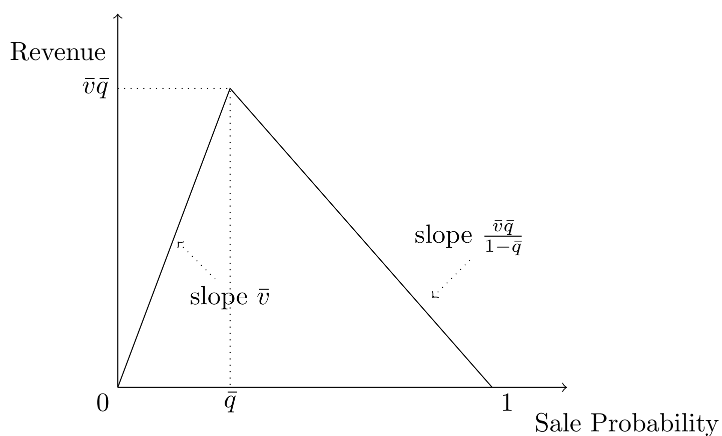

Figure 3.3. Revenue curve of distribution $\text{Tri}(\bar{v}, \bar{q})$.

3.2. Main Result

We obtain the main result of the chapter by deriving and solving a tight-in-the-limit continuous relaxation of program (P1) for triangular shaped revenue curves. This will resolve the open problem of finding the tight approximation gap between ex ante relaxation and anonymous pricing, while it improves the known approximation ratio for anonymous reserve auction and pricing with respect to the optimal auction (see Figure 3.4; the bound of 2.23 derived by simulation in Appendix A.1).

Theorem 3.2.1 (**Anonymous pricing versus ex ante relaxation**)**.** *For a single item environment with agents with independently (but non-identically) distributed values from regular distributions, the worst case approximation factor of anonymous pricing to the ex ante relaxation is* $\left(\mathcal{V}\left(\mathcal{Q}^{-1}(1)\right)+1\right)$ *which evaluates to* $e \approx 2.718$ *where*

$$\mathcal{V}(p) \triangleq p \cdot \ln\left(\frac{p^2}{p^2-1}\right), \qquad \mathcal{Q}(p) \triangleq \int_p^{\infty} -\frac{\mathcal{V}'(v)}{v}\,\mathrm{d}v.$$

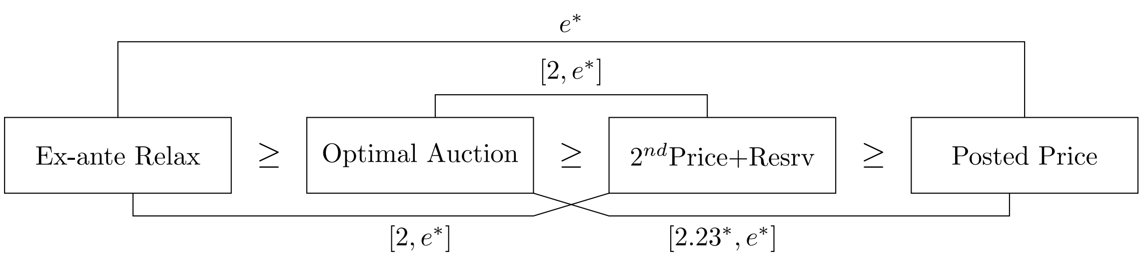

Figure 3.4. Revenue gap between mechanisms of study; * denotes new bounds.

The continuous relaxation of (P1) can be written as follows; maximize the objective $1+\mathcal{V}(p)$ subject to $\mathcal{Q}(p) \leq 1$ which has the following interpretation. The agents are a continuum and each agent value distribution is given by a pointmass at a value with some probability (and then a continuous distribution below the pointmass to minimally satisfy the regularity property). The function $\mathcal{V}(p)$ is the expected pointmass value from agents with pointmass value at least price p;[2] $\mathcal{Q}(p)$ is the expected number of these agents to realize their pointmass value. The optimal p^* meets the constraint with equality, i.e., $\mathcal{Q}(p^*) = 1$.

3.3. Upper-Bound Analysis

Program (P1) defines a tight upper bound on the ratio, denoted by ρ, of the revenue of the ex ante relaxation to the revenue of the optimal anonymous pricing. This program can be thought of

[2] $\mathcal{V}(\cdot)$ excludes the contribution from the "highest valued agent" which is 1; hence the objective $1+\mathcal{V}(p)$.

as a continuous optimization problem over regular distributions with the objective of maximizing the aforementioned ratio. The current section shows the upper bound of Theorem 3.2.1 while Section 3.4 shows tightness.

Overview of the analysis. By normalizing the optimal anonymous pricing revenue to be one, (P1) is equivalent to the following program:

$$\text{(P2)} \qquad \rho = \sup_{\mathcal{I} \in \textsc{Reg}} \quad \textsc{ExAnteRev}(\mathcal{I})$$

$$\text{(P2.1)} \qquad \text{subject to} \quad \textsc{PriceRev}(\mathcal{I}, p) \leq 1 \qquad \forall p \geq 0.$$

We show that for any fixed n the supremum of this program is approached even when restricting to triangular revenue curve instances, i.e., ones of the form $\{\text{Tri}(\bar{v}_i, \bar{q}_i)\}_{i=1}^{n}$ with $\sum_i \bar{q}_i \leq 1$ as defined in Section 3.1. Consequently, the problem is reduced to a discrete optimization problem over variables $\bar{\mathbf{v}} \triangleq (\bar{v}_1, \ldots, \bar{v}_n)$ and $\bar{\mathbf{q}} \triangleq (\bar{q}_1, \ldots, \bar{q}_n)$. An *assignment* for this optimization problem refers to a pair $(\bar{\mathbf{v}}, \bar{\mathbf{q}})$. This optimization problem is still of infinite dimension because n is itself a variable. It also turns out to be highly non-convex. Re-index $\bar{\mathbf{v}}$ such that $\bar{v}_1 \geq \ldots \geq \bar{v}_n$. We will show that, for any fixed n, inequality (P2.1) can be assumed without loss of generality to be tight for all $p \in \{\bar{v}_1, \ldots, \bar{v}_n\}$; otherwise an instance for which at least one of these constraints is not tight could be modified to make all these constraints tight while improving the objective. Thus,

$$\text{(3.7)} \qquad \textsc{PriceRev}(\{\text{Tri}(\bar{v}_i, \bar{q}_i)\}_{i=1}^{n}, \bar{v}_k) = 1 \qquad \forall k \in \{1, \ldots, n\}\,.$$

Observe that for each $k \in \{1, \ldots, n\}$ the left hand side of the above equation only depends on the first k agents, because the the valuations of the rest of the agents are always below $\bar{v}_k$. Consequently once $\bar{v}_1, \ldots, \bar{v}_n$ are fixed, we can compute $\bar{q}_1, \ldots, \bar{q}_n$ by solving equation (3.7) for $k \in \{1, \ldots, n\}$ and using forward substitution. Unfortunately, the resulting formulation of $\bar{q}_k$ in terms of $\bar{v}_1, \ldots, \bar{v}_k$ is analytically intractable for $k \geq 2$. To work around this intractability issue, we relax inequality (P2.1) in such a way that it leads to a tractable formulation of $\bar{\mathbf{q}}$ in terms of $\bar{\mathbf{v}}$. We also show that the relaxed inequality is tight which implies the value of the relaxed program is equal to that of the original program. Finally, we show that the supremum of the relaxed program is attained when $n \to \infty$, and roughly speaking the instance converges to a continuum of infinitesimal agents with triangular revenue curve distributions. For this continuum of agents, ρ is given simply by the optimization of p in the objective $1 + \mathcal{V}(p)$ subject to the constraint $\mathcal{Q}(p) \leq 1$ for the two functions $\mathcal{V}(\cdot)$ and $\mathcal{Q}(\cdot)$ given in the statement of Theorem 3.2.1.

Reduction to triangular revenue curve instances. We begin by showing that without loss of generality we can restrict program (P2) to triangle revenue curve instances.

Lemma 3.3.1. *The supremum of program* (P2) *is approached by triangle revenue curve instances, i.e., of the form* $\hat{\mathcal{I}} = \{\mathrm{Tri}(\bar{v}_i, \bar{q}_i)\}_{i=1}^n$ *with* $\sum_{i=1}^n \bar{q}_i \leq 1$.

Proof. We will show that for any regular instance $\mathcal{I} = \{F_i\}_{i=1}^n$, there exists a corresponding instance $\hat{\mathcal{I}} = \{\mathrm{Tri}(\bar{v}_i, \bar{q}_i)\}_{i=1}^n$ with $\sum_{i=1}^n \bar{q}_i \leq 1$ yielding the same optimal ex ante revenue and (weakly) smaller expected revenue from the optimal anonymous price.

Let $\bar{\mathbf{q}}$ be an optimal assignment for the ex ante relaxation program (3.1) that computes $\textsc{ExAnteRev}(\mathcal{I})$. Set $\bar{v}_i \leftarrow R_i(\bar{q}_i)/\bar{q}_i$ for each $i \in \{1, \dots, n\}$, where R_i is the revenue curve of F_i. We show changing agent i's valuation distribution to $\mathrm{Tri}(\bar{v}_i, \bar{q}_i)$ can only decrease the revenue of any anonymous pricing ($\textsc{PriceRev}$) while preserving the revenue of the ex ante relaxation ($\textsc{ExAnteRev}$), which implies the statement of the lemma.

Let $\hat{R}_i$ be the revenue curve of $\mathrm{Tri}(\bar{v}_i, \bar{q}_i)$. Observe that the change of distributions does not affect $\textsc{ExAnteRev}$ because $\hat{R}_i(\bar{q}_i) = R_i(\bar{q}_i)$ for all $i \in \{1, \dots, n\}$, and $\hat{R}_i$ is a lower bound on R_i elsewhere as R_i is concave (see Fig. 3.5). Therefore the replacement preserves the optimal value of convex program of Eq. (3.1) which implies $\textsc{ExAnteRev}(\hat{\mathcal{I}}) = \textsc{ExAnteRev}(\mathcal{I})$. Next, we show the replacement may only decrease the value of $\textsc{PriceRev}(p)$ at any $p > 0$. Fix a price p, and consider the price line corresponding to p, that is, the line with slope p passing through the origin (see Fig. 3.6). Observe that the probability of agent i's valuation being above p is equal to the q at which $R_i(q)$ intersects price line p. Given that $\hat{R}_i$ is a lower bound on R_i everywhere, the replacement may only decrease the probability of agent i's valuation being above p. Consequently, given that agents' valuations are distributed independently,

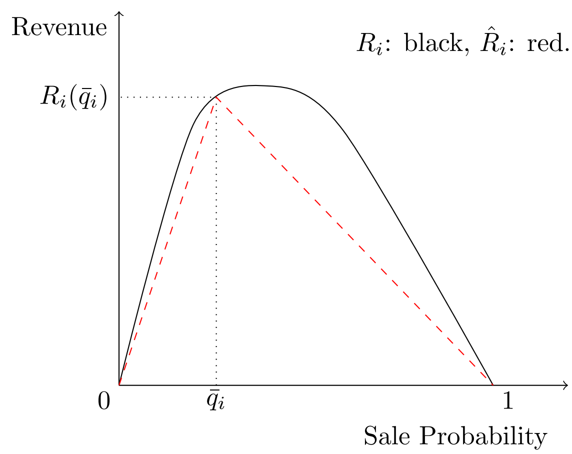

Figure 3.5. Replacing regular distribution F_i with triangular revenue curve distribution $\mathrm{Tri}(\bar{v}_i, \bar{q}_i)$.

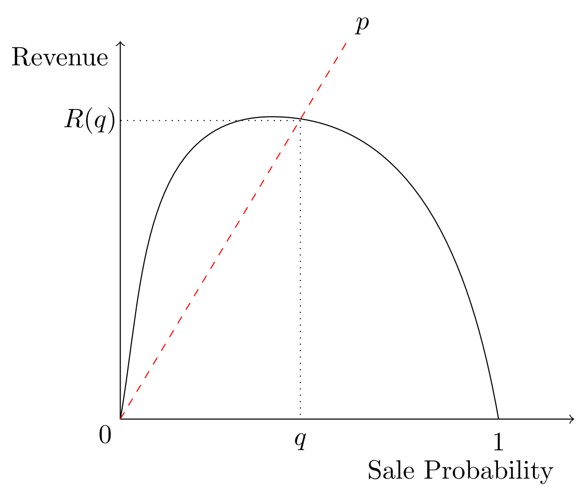

Figure 3.6. Intersection of revenue curve and price line p.

the replacement may only decrease the revenue from sale at any anonymous price p, which implies $\textsc{PriceRev}(\hat{\mathcal{I}}) \leq \textsc{PriceRev}(\mathcal{I})$. □

Combining Lemma 3.3.1 with the formulation of PriceRev and ExAnteRev from Eqs. (3.5) and (3.6) yields the following non-convex program for computing ρ:

$$\text{(P3)} \qquad \rho = \sup_{n\in\mathbb{N},\bar{\mathbf{v}},\bar{\mathbf{q}}} \sum_{i=1}^{n} \bar{v}_i\bar{q}_i$$

$$\text{(P3.1)} \qquad \text{subject to} \quad p \cdot \left(1 - \prod_{i:\bar{v}_i\geq p} \frac{1}{1+\frac{\bar{v}_i\bar{q}_i}{p\cdot(1-\bar{q}_i)}}\right) \leq 1, \qquad \forall p > 0$$

$$\sum_{i=1}^{n} \bar{q}_i \leq 1$$

$$\bar{v}_i \geq 0, \bar{q}_i \geq 0 \qquad \forall i \in \{1,\ldots,n\}\,.$$

Relaxations and canonical assignments. In this section we find a relaxation of program (P3) where the corresponding *pricing revenue constraint* (P3.1) is tight for all $p \in \{\bar{v}_1,\ldots,\bar{v}_n\}$ and can thus be written as a program on variables $\bar{\mathbf{v}}$ alone (i.e., by solving for the appropriate $\bar{\mathbf{q}}$ in terms of $\bar{\mathbf{v}}$). To simplify the solution of $\bar{\mathbf{q}}$ in terms of $\bar{\mathbf{v}}$, we will first make a series of relaxations to the pricing revenue constraint (P3.1). These relaxations will turn out to be tight in the limit with the number of agents n.

Lemma 3.3.2 formalizes these relaxations as sketched below. First, observe that the pricing revenue constraint (P3.1) can be rearranged as

$$\prod_{i:\bar{v}_i\geq p}\left(1+\frac{\bar{v}_i\bar{q}_i}{p\cdot(1-\bar{q}_i)}\right) \leq \left(\frac{p}{p-1}\right) \qquad \forall p > 0.$$

The first relaxation drops the constraint on $p \notin \{\bar{v}_1,\ldots,\bar{v}_n\}$; this is without loss as the optimal anonymous price is always in $\{\bar{v}_1,\ldots,\bar{v}_n\}$. We re-index such that $\bar{v}_1 \geq \cdots \geq \bar{v}_n$ and rephrase the relaxed constraint as

$$\prod_{i=1}^{k}\left(1+\frac{\bar{v}_i\bar{q}_i}{\bar{v}_k\cdot(1-\bar{q}_i)}\right) \leq \left(\frac{\bar{v}_k}{\bar{v}_k-1}\right) \qquad \forall k \in \{1,\ldots,n\}\,.$$

As the second relaxation we drop the term $(1-\bar{q}_i)$ from the denominator of the left hand side and take the logarithm of both sides to get

$$\sum_{i=1}^{k} \ln\left(1+\frac{\bar{v}_i\bar{q}_i}{\bar{v}_k}\right) \leq \ln\left(\frac{\bar{v}_k}{\bar{v}_k-1}\right) \qquad \forall k \in \{1,\ldots,n\}\,.$$

As the third relaxation we upper-bound $\ln\left(1+\frac{\bar{v}_i\bar{q}_i}{\bar{v}_k}\right)$ by $\frac{1}{\bar{v}_k}\ln(1+\bar{v}_i\bar{q}_i)$ for $i \geq 2$. Rearranging gives

$$\sum_{i=2}^{k} \ln\left(1+\bar{v}_i\bar{q}_i\right) \leq \bar{v}_k \ln\left(\frac{\bar{v}_k^2}{(\bar{v}_k-1)(\bar{v}_k+\bar{v}_1\bar{q}_1)}\right) \qquad \forall k \in \{2,\ldots,n\}\,.$$

The previous relaxation uses the fact that $\frac{1}{a}\ln(1+b) \leq \ln(1+\frac{b}{a})$ for all $a \geq 1, b \geq 0$. In the proof, we will also show that $\bar{v}_1\bar{q}_1$ can be replaced with 1 both in the above constraint and in the objective function without loss of generality. Putting everything together, we will obtain the following program as a relaxation of program (P3).

$$\text{(P4)} \qquad \rho' = \sup_{n\in\mathbb{N},\bar{\mathbf{v}},\bar{\mathbf{q}}} \quad 1+\sum_{i=2}^{n}\bar{v}_i\bar{q}_i$$

$$\text{(P4.1)} \qquad \text{subject to} \quad \sum_{i=2}^{k}\ln\left(1+\bar{v}_i\bar{q}_i\right) \leq \mathcal{V}(\bar{v}_k) \qquad \forall k \in \{2,\ldots,n\}$$

$$\sum_{i=2}^{n}\bar{q}_i \leq 1$$

$$\bar{v}_{i+1} \leq \bar{v}_i, \qquad \forall i \in \{2,\ldots,n-1\}$$

$$\bar{v}_i \geq 0, \bar{q}_i \geq 0 \qquad \forall i \in \{1,\ldots,n\}\,.$$

where $\mathcal{V}(\cdot) = p\cdot\ln\left(\frac{p^2}{p^2-1}\right)$ is defined in Theorem 3.2.1.

Lemma 3.3.2. *The value of program* (P4)*, denoted by* ρ'*, is an upper bound on the value of program* (P3) *which is* ρ*.*

Proof. Let $(\bar{\mathbf{v}}',\bar{\mathbf{q}}')$ be any feasible assignment for program (P3) for which $\sum_{i=1}^{n}\bar{v}'_i\bar{q}'_i > 1$. We construct a corresponding assignment $(\bar{\mathbf{v}},\bar{\mathbf{q}})$ which is feasible for program (P4) and yields the same objective value, that is $1+\sum_{i=2}^{n}\bar{v}_i\bar{q}_i = \sum_{i=1}^{n}\bar{v}'_i\bar{q}'_i$ which then implies that $\text{Bound}_I \le \text{Bound}_{II}$. Without loss of generality assume $\bar{v}'_1 \ge \ldots \ge \bar{v}'_n$. Let j be the smallest index for which $\sum_{i=1}^{j}\bar{v}'_i\bar{q}'_i > 1$. Observe that $2 \le j \le n$ because $\bar{v}'_i\bar{q}'_i \le 1$ for all i. Let $\delta = 1 - \sum_{i=1}^{j-1}\bar{v}'_i\bar{q}'_i$. Observe that $0 \le \delta < \bar{v}'_j\bar{q}'_j$. We construct a new optimal assignment $(\bar{\mathbf{v}}, \bar{\mathbf{q}})$ by setting for each $i \in \{2,\ldots,n\}$:

$$\bar{v}_i = \begin{cases} \bar{v}'_{j+i-2} & 2 \le i \le n-j+2 \\ 1 & n-j+3 \le i \le n \end{cases} \qquad \bar{q}_i = \begin{cases} \bar{q}'_j - \frac{\delta}{\bar{v}'_j} & i = 2 \\ \bar{q}'_{j+i-2} & 3 \le i \le n-j+2 \\ 0 & n-j+3 \le i \le n \end{cases}.$$

By the above construction it is easy to see that $1+\sum_{i=2}^{n}\bar{v}_i\bar{q}_i = \sum_{i=1}^{n}\bar{v}'_i\bar{q}'_i$. So we only need to show that $(\bar{\mathbf{v}}, \bar{\mathbf{q}})$ is indeed a feasible assignment. Observe that $\sum_{i=2}^{n}\bar{q}_i \le \sum_{i=1}^{n}\bar{q}'_i$, so the second constraint holds. So we only need to show that $(\bar{\mathbf{v}}, \bar{\mathbf{q}})$ satisfies the constraint (P4.1).

By rearranging (P3.1) we get

$$\prod_{i:\bar{v}'_i \ge p}\left(1+\frac{\bar{v}'_i\bar{q}'_i}{p\cdot(1-\bar{q}'_i)}\right) \le \left(\frac{p}{p-1}\right) \qquad \forall p > 0.$$

We then relax the previous inequality by dropping the term $(1-\bar{q}'_i)$ from the denominator of the left hand side and take the logarithm of both sides to get

$$\sum_{i:\bar{v}'_i \ge p}\ln\left(1+\frac{\bar{v}'_i\bar{q}'_i}{p}\right) \le \ln\left(\frac{p}{p-1}\right) \qquad \forall p > 0. \tag{3.8}$$

On the other hand, by invoking Lemma 3.3.3 we can argue that [3]

$$\ln\left(1+\frac{1}{\bar{v}_k}\right) + \sum_{i=2}^{k}\ln\left(1+\frac{\bar{v}_i\bar{q}_i}{\bar{v}_k}\right) \le \sum_{i=1}^{\min(k+j-2,n)}\ln\left(1+\frac{\bar{v}'_i\bar{q}'_i}{\bar{v}_k}\right) \qquad \forall k \in \{2,\ldots,n\}. \tag{3.9}$$

[3] We invoke Lemma 3.3.3 by setting $b = \frac{1}{\bar{v}_k}$, $a = \frac{\bar{v}_2\bar{q}_2}{\bar{v}_k}$, $m = j$ and $z_i = \frac{\bar{v}'_i\bar{q}'_i}{\bar{v}_k}$.

Observe that for any given k the the right hand side of (3.9) is equal or less than the the left hand side of (3.8) for $p = \bar{v}_k$ which implies

$$\ln\left(1+\frac{1}{\bar{v}_k}\right)+\sum_{i=2}^{k}\ln\left(1+\frac{\bar{v}_i\bar{q}_i}{\bar{v}_k}\right)\leq \ln\left(\frac{\bar{v}_k}{\bar{v}_k-1}\right) \qquad \forall k\in\{2,\dots,n\}\,.$$

We can then further relax the above inequality by replacing the terms $\ln(1+\frac{\bar{v}_i\bar{q}_i}{\bar{v}_k})$ with $\frac{1}{\bar{v}_k}\ln(1+\bar{v}_i\bar{q}_i)$ and rearranging the terms to get the constraint (P4.1) which implies that $(\bar{\mathbf{v}},\bar{\mathbf{q}})$ is feasible with respect to that constraint as well. □

Lemma 3.3.3. *Consider any $a, b, z_1, \dots, z_m \geq 0$ such that $a + b = \sum_{i=1}^m z_i$ and $a \leq z_m \leq b$. Then*

$$\ln(1+b)+\ln(1+a)\leq\sum_{i=1}^{m}\ln(1+z_i).$$

Proof. We can re-write the equation as

$$\ln\big((1+a)(1+b)\big)\leq\ln\prod_{i=1}^{m}(1+z_i) \tag{3.10}$$

Observe that

$$\begin{aligned}\prod_{i=1}^{m}(1+z_i) &\geq 1+\sum_{i=1}^{m}z_i+z_m(\sum_{i=1}^{m-1}z_i)\\ &= (1+\sum_{i=1}^{m-1}z_i)(1+z_m)\\ &\geq (1+a)(1+b),\end{aligned}$$

where the first inequality follows by eliminating some terms from the expansion of $\prod_{i=1}^m(1+z_i)$, and the second inequality from the assumption that $(1+a)\leq(1+z_m)$ and $(1+b)\geq(1+\sum_{i=1}^{m-1}z_i)$. □

of generality the pricing revenue constraint (P4.1) is tight for all $k \in \{1, \dots, n\}$ in program (P4). That will allow us to specify one set of variables (e.g., $\bar{\mathbf{q}}$) in terms of the other set of variables (e.g., $\bar{\mathbf{v}}$), which consequently allows us to eliminate the former variables and drop the pricing revenue constraint

(P4.1). To this end, we first define a *canonical* feasible solution for (P4), restriction to which is without loss given by Lemma 3.3.4.

Definition 3.3.1. A feasible assignment $(\bar{\mathbf{v}}, \bar{\mathbf{q}})$ for (P4) is *canonical* if the pricing constraint (P4.1) is tight for all $k \in \{2, \dots, n\}$.

Lemma 3.3.4. *For any feasible assignment* $(\bar{\mathbf{v}}, \bar{\mathbf{q}})$ *for* (P4)*, there exists an equivalent canonical feasible assignment* $(\bar{\mathbf{v}}', \bar{\mathbf{q}}')$ *obtaining the same objective value, that is* $\sum_i \bar{v}_i \bar{q}_i = \sum_i \bar{v}'_i \bar{q}'_i$.

Proof. Without loss of generality assume $\bar{q}_k > 0$ for all $k \in \{2, \dots, n\}$.[4] The right hand side of the pricing constraint (P4.1) is $\mathcal{V}(\bar{v}_k)$ which is decreasing in $\bar{v}_k$ (see Lemma 3.3.5) and approaches 0 as $\bar{v}_k \to \infty$, so for every $k \in \{2, \dots, n\}$ there exists $\bar{v}'_k \geq \bar{v}_k$ such that

$$\sum_{i=2}^{k} \ln\left(1 + \bar{v}_i \bar{q}_i\right) = \mathcal{V}(\bar{v}'_k) \qquad \forall k \in \{2, \dots, n\}\,.$$

Observe that by the above construction we always have $\bar{v}'_2 \geq \ldots \geq \bar{v}'_n$. We then decrease $\bar{q}_k$ to $\bar{q}'_k = \bar{q}_k \frac{\bar{v}_k}{\bar{v}'_k}$ for each $k \in \{2, \dots, n\}$ to obtain the desired assignment $(\bar{\mathbf{v}}', \bar{\mathbf{q}}')$. □

By Lemma 3.3.4, we can restrict our attention to canonical assignments of (P4) without loss of generality. In particular, we can fully identify such a canonical assignment by specifying only $\bar{\mathbf{v}} = (\bar{v}_1, \dots, \bar{v}_n)$ since the corresponding $\bar{\mathbf{q}}$ is given by

$$\bar{q}_k = \frac{e^{\mathcal{V}(\bar{v}_k) - \mathcal{V}(\bar{v}_{k-1})} - 1}{\bar{v}_k} \qquad \forall k \in \{2, \dots, n\}\,. \tag{3.11}$$

[4]If $\bar{q}_k = 0$, we can drop agent k without affecting feasibility or objective value.

Therefore we can obtain from program (P4) the following program.

$$\text{(P5)} \qquad \rho' = \sup_{n \in \mathbb{N}, \bar{\mathbf{v}}} \quad 1 + \sum_{i=2}^{n} \bar{v}_i \bar{q}_i$$

$$\text{(P5.1)} \qquad \text{subject to} \quad \bar{q}_k = \frac{e^{\mathcal{V}(\bar{v}_k) - \mathcal{V}(\bar{v}_{k-1})} - 1}{\bar{v}_k} \qquad \forall k \in \{2, \ldots, n\}$$

$$\text{(P5.2)} \qquad \sum_{i=2}^{n} \bar{q}_i \leq 1$$

$$\bar{v}_{i+1} \leq \bar{v}_i, \qquad \forall i \in \{2, \ldots, n-1\}$$

$$\bar{v}_i \geq 0, \; \bar{q}_i \geq 0 \qquad \forall i \in \{1, \ldots, n\}\,.$$

Continuum of agents. Given that n itself is a variable, a solution to program (P5) can be practically specified by a finite subset $\bar{\mathbf{v}} \subset \mathbb{R}_+$ where $\bar{v}_i$ is the ith largest value in that subset. We now show that the optimal solution to program (P5) corresponds to $\bar{\mathbf{v}} = [p, \infty)$ (for some $p > 1$) which can be viewed as an instance with infinitely many infinitesimal agents.

For any given $p' > p > 1$, we define a continuum of agents $[p, p')$ by defining for each $m \in \mathbb{N}$ a discrete family of agents of size m spanning $[p, p')$ and by taking the limit of this family as $m \to \infty$. Formally, for each $m \in \mathbb{N}$, we consider the family of agents with distributions $\{\mathrm{Tri}(u_j, (e^{\mathcal{V}(u_j) - \mathcal{V}(u_{j-1})} - 1)/u_j)\}_{j=1}^{m}$ where $u_j = p' + \frac{j}{m}(p - p')$. Observe that the agents in these families satisfy equation (3.11). Furthermore, observe that

$$\lim_{\delta \to 0} \left(\frac{e^{\mathcal{V}(v) - \mathcal{V}(v+\delta)} - 1}{v} \cdot \frac{1}{\delta} \right) = -\frac{\mathcal{V}'(v)}{v}.$$

Therefore in a continuum of agents $[p, p')$ each infinitesimal agent $v \in [p, p')$ has a distribution of $\mathrm{Tri}(v, -\frac{\mathcal{V}'(v)}{v}\,\mathrm{d}v)$, which implies that the contribution of $[p, p')$ to the objective value of (P5) is

$$\text{(3.12)} \qquad \int_{p}^{p'} v \cdot \left(-\tfrac{\mathcal{V}'(v)}{v}\right) \mathrm{d}v = \mathcal{V}(p) - \mathcal{V}(p'),$$

and the contribution of $[p, p')$ to the left hand side of the constraint (P5.2), i.e. $\sum_i \bar{q}_i \leq 1$ which is referred to as the *capacity constraint*, is

(3.13) $$\int_p^{p'} -\frac{\mathcal{V}'(v)}{v}\,\mathrm{d}v = \mathcal{Q}(p) - \mathcal{Q}(p'),$$

where $\mathcal{Q}(v) = \int_p^\infty -\frac{1}{v}\mathcal{V}'(v)\,\mathrm{d}v$ as defined in Theorem 3.2.1.

Via the above derivation of a continuum of agents, program (P5), on the instance corresponding to the continuum $[p, \infty)$, simplifies as:

(P6) $$\begin{aligned} \rho'' = \quad & \max_{p \geq 1} \quad 1 + \mathcal{V}(p) \\ & \text{subject to} \quad \mathcal{Q}(p) \leq 1. \end{aligned}$$

Next we will sketch a construction that demonstrates that any feasible solution $\bar{\mathbf{v}} = (\bar{v}_1, \ldots, \bar{v}_n)$ to program (P5) can be replaced by a continuum of agents that corresponds to an interval $[p, \infty)$ (for some $p > 1$ to be determined) and the objective of (P5) is strictly increased. Note that $\bar{v}_1$ does not appear anywhere in (P5); for notational convenience we redefine it as $\bar{v}_1 = \infty$. Suppose for each $i \in \{2, \ldots, n\}$ we replace the agent $\bar{v}_i$ with the continuum of agents $[\bar{v}_i, \bar{v}_{i-1})$. It follows from Eqs. (3.11) and (3.12) that this replacement changes the object value of (P5) by $\mathcal{V}(\bar{v}_i) - \mathcal{V}(\bar{v}_{i-1}) - \bar{v}_i\bar{q}_i = \ln(1 + \bar{v}_i\bar{q}_i) - \bar{v}_i\bar{q}_i < 0$ which is unfortunately always negative and thus the opposite of what we want to prove. On the other hand, it follows from Eqs. (3.11) and (3.13) that this replacement also changes the left hand side of the capacity constraint (P5.2) by $\mathcal{Q}(\bar{v}_i) - \mathcal{Q}(\bar{v}_{i-1}) - \bar{q}_i$ which is also negative (as we will show later), and thus creates some slack in the capacity constraint (P5.2). Summing over the slack created in the capacity constraint (P5.2) from converting each agent to a continuum, we can add a new continuum of agents $[p, \bar{v}_n)$ where $p < \bar{v}_n$ is chosen to make the capacity constraint (P5.2) tight. As a consequence of the following claims, the net change in the objective value from this transformation is positive.

(i) The amount of slack created in the capacity constraint (P5.2) by replacing $\bar{v}_i$ with $[\bar{v}_i, \bar{v}_{i-1})$ is more than the decrease in the objective value of (P5). Using Eqs. (3.12) and (3.13), we can formally write this claim as $\bar{q}_i - (\mathcal{Q}(\bar{v}_i) - \mathcal{Q}(\bar{v}_{i-1})) > \bar{v}_i\bar{q}_i - (\mathcal{V}(\bar{v}_i) - \mathcal{V}(\bar{v}_{i-1}))$. This is proved below in Lemma 3.3.7.

(ii) If there is a slack of $\Delta > 0$ in the capacity constraint (P5.2), it can be used to extend the last continuum of agents to increase the objective value by more than Δ. Using Eqs. (3.12) and (3.13), we can formalize this claim as follows: if p is chosen such that $\mathcal{Q}(p)-\mathcal{Q}(\bar{v}_n) = \Delta$, then $\mathcal{V}(p)-\mathcal{V}(\bar{v}_n) > \Delta$. This is proved below in Lemma 3.3.5.

The suggestion from the above construction is that from any solution $\bar{\mathbf{v}}$ to program (P5), a price p can be identified such that the continuum of agents on $[p, \infty)$ has higher objective value. In other words, the optimal values of program (P5) and program (P6) are equal. This is proved below in Lemma 3.3.8; though we defer the proof that the solution of program (P6) corresponds to a limit solution of program (P5) to Section 3.4.

Algebraic upper-bound proof. The rest of this section develops a formal but purely algebraic proof that is based on the approach sketched in the previous paragraphs.

Lemma 3.3.5. *The functions $\mathcal{V}(p)$, $\mathcal{Q}(p)$, and $\mathcal{V}(p) - \mathcal{Q}(p)$ are all decreasing in p, for $p > 1$.*

Proof. To prove $\mathcal{V}(p)$ is decreasing, we show its derivative is negative:

$$\mathcal{V}'(p) = \ln\left(1 + \frac{1}{p^2-1}\right) - \frac{2}{p^2-1} < \frac{1}{p^2-1} - \frac{2}{p^2-1} < 0.$$

The first inequality uses the fact that $\ln(1+x) \leq x$. Similarly $\mathcal{Q}(p)$ is also decreasing because $\mathcal{Q}'(p) = \frac{1}{p}\mathcal{V}'(p) < 0$. Finally $\mathcal{V}(p) - \mathcal{Q}(p)$ is also decreasing because

$$\big(\mathcal{V}(p) - \mathcal{Q}(p)\big)' = \mathcal{V}'(p)\left(1 - \frac{1}{p}\right) < 0.$$

□

Lemma 3.3.6. *for any $p' > p > 1$ the following inequality holds: $\mathcal{V}(p) - \mathcal{V}(p') < \ln(\frac{p}{p-1}) - \ln(\frac{p'}{p'-1})$.*

Proof. Define $G(p) = \mathcal{V}(p) - \ln(\frac{p}{p-1})$. Observe that proving the inequality in the statement of the lemma is equivalent to proving $G(p) < G(p')$ which we do by showing $G(p)$ has positive derivative and is therefore increasing.

$$G'(p) = \ln\left(1 + \frac{1}{p^2 - 1}\right) - \frac{1}{p^2 + p}.$$

We will show that $G'(p)$ is decreasing which then implies $G'(p) > 0$ because $\lim_{p\to\infty} G'(p) = 0$. Therefore we only need to show that $G''(p) < 0$.

$$G''(p) = \frac{-(3p+1)}{(p-1)p^2(p+1)^2} < 0.$$

□

Lemma 3.3.7. *For any $p' > p > 1$ and $q = \frac{e^{\mathcal{V}(p)-\mathcal{V}(p')}-1}{p}$ the following inequalities hold:*

$$q - (\mathcal{Q}(p) - \mathcal{Q}(p')) \geq pq - (\mathcal{V}(p) - \mathcal{V}(p')) \geq 0 \tag{3.14}$$

Proof. Define

$$\begin{aligned} W(p,p') &\triangleq \mathcal{V}(p) - \mathcal{Q}(p) - \mathcal{V}(p') + \mathcal{Q}(p') + q - pq \\ &= \mathcal{V}(p) - \mathcal{Q}(p) - \mathcal{V}(p') + \mathcal{Q}(p') - (p-1)\big(e^{\mathcal{V}(p)-\mathcal{V}(p')} - 1\big)/p. \end{aligned}$$

Observe that proving the first inequality in the statement of the lemma is equivalent to proving $W(p,p') > 0$. We instead prove that $W(p,p')$ is increasing in p' which together with the trivial fact that $W(p,p) = 0$ implies $W(p,p') > 0$.

$$\begin{aligned} \tfrac{\partial}{\partial p'} W(p,p') &= -\mathcal{V}'(p') + \mathcal{V}'(p')/p' + (p-1)\,\mathcal{V}'(p')\,e^{\mathcal{V}(p)-\mathcal{V}(p')}/p \\ &= -\mathcal{V}'(p')\left[\frac{p'-1}{p'} - \frac{p-1}{p}\,e^{\mathcal{V}(p)-\mathcal{V}(p')}\right] \\ &> -\mathcal{V}'(p')\left[\frac{p'-1}{p'} - \frac{p-1}{p}\,e^{\ln(\frac{p}{p-1})-\ln(\frac{p'}{p'-1})}\right] = 0. \end{aligned}$$

The finally inequality follows from Lemmas 3.3.5 and 3.3.6: by Lemma 3.3.5, $-\frac{\partial}{\partial p'}\mathcal{V}(p') > 0$; and by Lemma 3.3.6, $e^{\mathcal{V}(p)-\mathcal{V}(p')}$ is less than $e^{\ln(\frac{p}{p-1})-\ln(\frac{p'}{p'-1})}$, so replacing the former with the latter only decreases the value of the expression inside the brackets because its coefficient is $-\frac{p-1}{p}$ which is negative.

The second inequality in the statement of the lemma follows trivially from the fact that $\mathcal{V}(p)-\mathcal{V}(p') = \ln(1+pq)$ thus $pq - (\mathcal{V}(p) - \mathcal{V}(p')) = pq - \ln(1+pq) > 0$. □

Lemma 3.3.8. *The value of program* (P6)*, denoted by* ρ''*, is an upper bound on the value of program* (P5) *which is* ρ'*.*

Proof. Let $(\bar{\mathbf{v}}, \bar{\mathbf{q}})$ be any arbitrary feasible assignment for program (P5). We show there exists a feasible assignment for program (P6) with objective value upper bounding the objective value of $(\bar{\mathbf{v}}, \bar{\mathbf{q}})$ in program (P5). Define $p^* \triangleq \mathcal{Q}^{-1}(1)$, a candidate solution to program (P6) that meets the feasibility constraint with equality. Note that such a p^* exists because $\mathcal{Q}(\infty) = 0$, $\mathcal{Q}(1) = \infty$, and $\mathcal{Q}(\cdot)$ is continuous. Observe that the objective value of (P5) for $(\bar{\mathbf{v}}, \bar{\mathbf{q}})$ satisfies:

$$
\begin{aligned}
1 + \sum_{k=2}^{n} \bar{v}_k \bar{q}_k &\leq 1 + \sum_{k=2}^{n} \big(\mathcal{V}(\bar{v}_k) - \mathcal{V}(\bar{v}_{k-1}) - (\mathcal{Q}(\bar{v}_k) - \mathcal{Q}(\bar{v}_{k-1})) + \bar{q}_k\big) && \text{by Lemma 3.3.7 and } \bar{v}_1 = \infty \\
&= 1 + \mathcal{V}(\bar{v}_n) - \mathcal{Q}(\bar{v}_n) + \sum_{k=2}^{n} \bar{q}_k && \text{as } \mathcal{V}(\infty) = \mathcal{Q}(\infty) = 0 \\
&\leq 1 + \mathcal{V}(\bar{v}_n) - \mathcal{Q}(\bar{v}_n) + 1 && \text{as } \sum_{k=2}^{n} \bar{q}_k \leq 1 \\
(*) \qquad &< 1 + \mathcal{V}(p^*) - \mathcal{Q}(p^*) + 1 && \text{as proved below} \\
&= 1 + \mathcal{V}(p^*) && \text{as } \mathcal{Q}(p^*) = 1.
\end{aligned}
$$

To prove inequality $(*)$ we show that $p^* < \bar{v}_n$ which together with Lemma 3.3.5 implies $\mathcal{V}(p^*) - \mathcal{Q}(p^*) > \mathcal{V}(\bar{v}_n) - \mathcal{Q}(\bar{v}_n)$. To prove $p^* < \bar{v}_n$ observe that Lemma 3.3.7 implies $\sum_{i=2}^{n} \bar{q}_i > \sum_{i=2}^{n} \mathcal{Q}(\bar{v}_i) - \mathcal{Q}(\bar{v}_{i-1}) = \mathcal{Q}(\bar{v}_n)$. On the other hand $\sum_{i=2}^{n} \bar{q}_i \leq 1$. Therefore $\mathcal{Q}(\bar{v}_n) < 1$ which implies $\bar{v}_n > p^*$ because $\mathcal{Q}(p^*) = 1$ and, by Lemma 3.3.5, $\mathcal{Q}(\cdot)$ is decreasing. □

We conclude the section with the proof of the upper-bound of Theorem 3.2.1.

Proof of upper-bound in Theorem 3.2.1. It follows from program (P1), Lemmas 3.3.1, 3.3.2 and 3.3.4, and the rest of the discussion in this section that ρ' which is computed by (P5) is an upper bound on the ratio of the ex ante relaxation to the expected revenue of the optimal anonymous pricing. Following Lemma 3.3.8, ρ' is upper bounded by the objective value of program (P6), i.e. ρ''. As $\mathcal{Q}(\cdot)$ and $\mathcal{V}(\cdot)$ are decreasing (Lemma 3.3.5), the optimal solution to program (P6) is given by $\rho'' = 1 + \mathcal{V}(\mathcal{Q}^{-1}(1))$ which numerically evaluates to $e \approx 2.718$. □

3.4. Lower-Bound Analysis

In this section we show the tightness of our approximation, i.e. the lower-bound in Theorem 3.2.1. As a result of Lemma 3.3.1, it suffices to prove the following lemma.

Lemma 3.4.1. *For any $\epsilon > 0$ there exists a feasible assignment $(n, \bar{\mathbf{v}}, \bar{\mathbf{q}})$ of the program* (P3) *such that $\sum_{i=1}^{n} \bar{v}_i \bar{q}_i \geq 1 + \mathcal{V}(\mathcal{Q}^{-1}(1)) - \epsilon$.*

Proof. Pick $\delta > 0$ s.t. $(1-\delta)^2 \left(1 + \mathcal{V}\left(\mathcal{Q}^{-1}\left(\frac{1}{(1+\delta)^2}\right) + \delta\right)\right) \geq 1 + \mathcal{V}(\mathcal{Q}^{-1}(1)) - \epsilon$. This is always possible as $\mathcal{V}$ and $\mathcal{Q}$ are decreasing. Lets define $\lambda = \mathcal{Q}^{-1}(1)$. The proof is done in two steps:

Step 1: We find $\{v_i, q_i\}_{i=2}^{n}$ such that

$$\sum_{i=2}^{n} q_i \leq 1 \ , \ k = 2, \dots, n : \ \sum_{i=2}^{k} \ln(1 + v_i q_i) = \mathcal{V}(v_k)$$
$$\sum_{i=2}^{n} v_i q_i \geq \mathcal{V}\left(\mathcal{Q}^{-1}\left(\frac{1}{(1+\delta)^2}\right) + \delta\right)$$

In our construction for $\{v_i, q_i\}_{i=2}^{n}$, we use two parameters $\Delta > 0$ and $V_T \geq \lambda$ which we fix later in the proof. Let $v_1 \triangleq \infty$, and for $i \geq 2$ set $v_i = V_T - (i-2)\Delta$ and $q_i = \frac{e^{\mathcal{V}(v_i) - \mathcal{V}(v_{i-1})} - 1}{v_i}$. Now, let $n = \max\{n_0 \in \mathbb{N} : \sum_{i=2}^{n_0} q_i \leq 1\}$. Obviously, $\sum_{i=2}^{n} q_i \leq 1$. Moreover, for any $2 \leq k \leq n$ we have $\sum_{i=2}^{k} \ln(1 + v_i q_i) = \sum_{i=2}^{k} (\mathcal{V}(v_i) - \mathcal{V}(v_{i-1})) = \mathcal{V}(v_k) - \mathcal{V}(v_1) = \mathcal{V}(v_k)$. Now, pick $\delta' > 0$ small enough such that for $x \in [0, \delta']$ we have $\frac{e^x - 1}{x} \leq 1 + \delta$. Moreover, let Δ to be small enough and V_T to be large enough such that $\max\{\mathcal{V}(\lambda) - \mathcal{V}(\lambda + \Delta), \mathcal{V}(V_T), \Delta\} \leq \min\{\delta, \delta'\}$. First observe that due to Lemma 3.3.7 $q_i \geq \mathcal{Q}(v_i) - \mathcal{Q}(v_{i-1})$ which implies $2 \geq \sum_{i=1}^{n} q_i \geq \mathcal{Q}(v_n) - \mathcal{Q}(v_1) = \mathcal{Q}(v_n)$. So, all values v_i are at least equal to λ. As $\mathcal{V}(.)$ is convex over $[1,\infty)$, we have $\mathcal{V}(v_i)-\mathcal{V}(v_{i-1}) \le \mathcal{V}(\lambda)-\mathcal{V}(\lambda+\Delta)\le \delta'$. As a result we have

$$q_i = \frac{e^{\mathcal{V}(v_i)-\mathcal{V}(v_{i-1})}-1}{v_i} \le (1+\delta)\frac{\mathcal{V}(v_i)-\mathcal{V}(v_{i-1})}{v_i} = (1+\delta)\int_{v_i}^{v_{i-1}} \frac{\mathcal{V}'(v)}{v_i}\,\mathrm{d}v.$$

$$= (1+\delta)\int_{v_i}^{v_{i-1}} -\frac{v}{v_i}\mathcal{Q}'(v)\,\mathrm{d}v = (1+\delta)\left(\mathcal{Q}(v_i)-\mathcal{Q}(v_{i-1}) + \int_{v_i}^{v_{i-1}} -\frac{v-v_i}{v_i}\mathcal{Q}'(v)\,\mathrm{d}v\right)$$

$$\le (1+\delta)\left(\mathcal{Q}(v_i)-\mathcal{Q}(v_{i-1}) + \Delta\int_{v_i}^{v_{i-1}} -\mathcal{Q}'(v)\,\mathrm{d}v\right) \le (1+\delta)^2(\mathcal{Q}(v_i)-\mathcal{Q}(v_{i-1})) \tag{3.15}$$

Based on the definition of n (number of distributions in our instance), we have $1 < \sum_{i=2}^{n+1} q_i$. By (3.15), we have $\sum_{i=2}^{n+1} q_i \le (1+\delta)^2 \sum_{i=2}^{n+1}((\mathcal{Q}(v_i)-\mathcal{Q}(v_{i-1})) = (1+\delta)^2\mathcal{Q}(v_{n+1})$. Lets define $\lambda' \triangleq \mathcal{Q}^{-1}\left(\frac{1}{(1+\delta)^2}\right)$. We conclude that $\lambda' \ge v_{n+1}$. Hence, $v_n \le \lambda'+\Delta \le \lambda'+\delta$. Moreover, using Lemma 3.3.1 we have

$$\sum_{i=2}^{n} v_i q_i \ge \sum_{i=2}^{n}(\mathcal{V}(v_i)-\mathcal{V}(v_{i-1})) = \mathcal{V}(v_n) \ge \mathcal{V}(\lambda'+\delta) = \mathcal{V}\left(\mathcal{Q}^{-1}\left(\frac{1}{(1+\delta)^2}\right)+\delta\right) \tag{3.16}$$

where the last inequality is true because $v_n \le \lambda'+\delta$ and $\mathcal{V}$ is decreasing over $[1,\infty)$.

Step 2: Given $\{v_i,q_i\}_{i=2}^n$, we find an instance $\{\bar{v}_i,\bar{q}_i\}_{i=1}^n$ such that is feasible for program (P3) and $\sum_{i=1}^n \bar{v}_i\bar{q}_i \ge 1+\mathcal{V}(\mathcal{Q}^{-1}(1))-\epsilon$. To do so, set $q_1=\delta, v_1=\frac{1}{\delta}-1$. Now, for each $i,k\in[2:n]$ find $\gamma_{i,k}$ such that

$$1+\frac{v_i q_i(1-\gamma_{i,k})}{v_k} = (1+v_i q_i)^{\frac{1}{v_k}} \tag{3.17}$$

and then let $\bar{q}_i = (1-\delta)(1-\max_{k\in[2:n]}\gamma_{i,k})q_i$ and $\bar{v}_i = v_i$, for $i\in[2:n]$. Now we claim $\{\bar{v}_i,\bar{q}_i\}_{i=1}^n$ is a feasible assignment for the program (P3) . We have

$$\sum_{i=1}^n \bar{q}_i = \delta + (1-\delta)\sum_{i=2}^n(1-\max_{k\in[2:n]}\gamma_{i,k})q_i \le \delta+(1-\delta)\sum_{i=2}^n q_i \le 1. \tag{3.18}$$

as $\sum_{i=2}^{n} q_i \leq 1$. Moreover, for $k \in [2:n]$ we have

$$\begin{aligned}
\sum_{i=1}^{k} \ln\left(1+\frac{\bar{v}_i\bar{q}_i}{\bar{v}_k(1-\bar{q}_i)}\right) &\leq \ln\left(1+\frac{\bar{v}_1\bar{q}_1}{\bar{v}_k(1-\bar{q}_1)}\right)+\sum_{i=2}^{k}\ln\left(1+\frac{v_iq_i(1-\gamma_{i,k})}{v_k}\right) \\
&= \ln\left(\frac{v_k+1}{v_k}\right)+\frac{1}{v_k}\sum_{i=2}^{k}\ln(1+v_iq_i) \\
&\leq \ln\left(\frac{v_k+1}{v_k}\right)+\ln\left(\frac{v_k^2}{v_k^2-1}\right) \\
&= \ln\left(\frac{\bar{v}_k}{\bar{v}_k-1}\right)
\end{aligned}$$

By taking exponents from both sides and rearranging the terms it is not hard to see $(n, \bar{\mathbf{v}}, \bar{\mathbf{q}})$ is a feasible assignment of program (P3). Additionally, for a fixed V_T all of the v_i's are bounded, i.e. $1 \leq v_i \leq V_T$. So, as Δ goes to zero we have $v_iq_i \to 0$ as $q_i \to 0$, and the left hand side of (3.17) converges to its right hand side. As a result, for small enough Δ, we can guarantee $\gamma_{i,k} \leq \delta$ for all i, k, and hence $\bar{q}_i \geq (1-\delta)^2 q_i$. So

$$\begin{aligned}
\sum_{i=1}^{n} \bar{v}_i\bar{q}_i &\geq (1-\delta)+(1-\delta)^2\left(\sum_{i=2}^{n} v_iq_i\right) \\
&\geq (1-\delta)^2\left(1+\sum_{i=2}^{n} v_iq_i\right) \\
&\geq (1-\delta^2)\left(1+\mathcal{V}\left(\mathcal{Q}^{-1}\left(\tfrac{1}{(1+\delta)^2}\right)+\delta\right)\right)
\end{aligned}$$

which implies $\sum_{i=1}^{n} \bar{v}_i\bar{q}_i \geq 1+\mathcal{V}(\mathcal{Q}^{-1}(1))-\epsilon$, as desired. $\square$

CHAPTER 4

Anonymous Pricing: Repeated Sales

4.1. Introduction

It is now commonplace for regular, repeated purchases to be made through large online platforms. New parents purchase diapers monthly through Amazon Prime. Firms buy online advertising space millions of times per day through Google, Microsoft and other advertising markets. City-dwellers use delivery services like Foodler and Instacart to purchase their meals and groceries. Each platform is a cornucopia of data, since they can readily observe how pricing decisions affect the purchasing behavior of customers, both in aggregate and individually. It is tempting for a platform to exploit this historical data, by using the past behavior of individual users to tune prices and maximize revenue. However, using revealed preference data in this way runs afoul of game-theoretic considerations. If a regular customer knows that their behavior will impact the prices they will be offered in the future, they will naturally respond by changing their behavior. It is therefore crucial to understand how forward-looking customers will respond to price-learning algorithms, and the implications for how a seller should use historical data to make pricing decisions.

Consider the following simple and fundamental instantiation of the repeated-sales problem, coined the “fishmonger problem” Devanur et al. (2015). There is a single seller, and each day the seller has a single copy of a good to sell. There is a single buyer, who has a private value $v \geq 0$ for obtaining the good each day, drawn from a distribution known to the seller. Crucially, the value does not change from one day to the next; the buyer has the same value for consuming the good on every day. Each day, the seller posts a take-it-or-leave-it price, and the buyer can choose to accept or reject. The seller is free to set each day’s price however she chooses, given the past purchasing behavior of the buyer. On any day that the buyer rejects, the good expires and the seller must discard it. The game is played for infinitely many rounds; the buyer wishes to maximize total time-discounted utility, and the seller wishes to maximize total time-discounted revenue.

How should the seller set her price? If there is only a single round, the well-known solution is to post the Myerson price for the buyer's distribution, which maximizes expected revenue. In the dynamic setting, however, we cannot expect the seller to post the Myerson price each round. After all, if the buyer chose not to purchase on the first day, the seller would naturally want to learn from this information and set a lower price on the following day. It is tempting to guess that the seller can benefit from this oppotunity to learn, by offering a variety of prices to gain information about the value v, then use this knowledge to set an aggressive price just below v. However, a surprising folklore result implies that such techniques can never be beneficial to the seller: the average per-round revenue can never be higher than the one-round Myerson revenue. Intuitively, the issue is that a rational buyer would respond to an explore/exploit strategy by pretending at first to have a low value, passing up some opportunities to buy the item, in order to secure a lower price later on. Indeed, this strategic demand-reduction behavior is the essence of bargaining, and is commonly observed in practice.

So what *can* the seller do? To disentangle the strategic behavior of the buyer and seller, it is necessary to study equilibria. Since ours is a repeated game with private information, the appropriate solution concept is perfect Bayesian equilibrium (PBE). A formal definition is given in Section 5.2, but roughly speaking a PBE requires that the decision taken by each player at each point in time, for any observed history of prices and purchases, is a best response to the anticipated future behavior of the other player, given the seller's belief about the private value (which will depend on the observed behavior of the buyer). Determining how the seller should set prices then reduces to understanding the structure of PBE. Sadly, prior work on equilibria for repeated sales have mostly generated negative results, from the perspective of revenue. In particular, there exist PBE in which the seller posts a price equal to the minimum value in the support of the buyer's distribution, on every round Devanur et al. (2015); Hart and Tirole (1988); Schmidt (1993). For example, if the buyer's value is supported on $[0, 1]$, then there is a PBE with zero revenue for the seller. While it is perhaps unsurprising that there is an efficient equilibrium, the low revenue of this equilibrium is discouraging.[1]

This result is unsatisfying, since the low-revenue equilibrium does not appear to be predictive of real-world outcomes. Why don't we see this behavior in practice? The strongest simplifying assumption in the model is the presence of only a single buyer. Indeed, because there is only one buyer, it is possible for the seller to exploit the buyer's revealed preference in a very targetted way. In contrast, if the seller continues to sell by posting a single price, but that price will be faced by *multiple* buyers, then the opportunity for price-discrimination is diminished. Intuitively, in a market with multiple buyers, each buyer is less worried about being exploited directly, and competition gives an extra incentive to purchase even though this is revealing a signal to the seller. We therefore ask: how does the presence of multiple buyers change the structure of equilibrium?

Our Results. We consider three extensions of the fishmonger problem, in increasing order of complexity: we first re-examine the single-buyer setting of Devanur et al. (2015). We then consider buyers competing for the sale of a single item. Finally, we analyze the setting of multiple buyers with unlimited supply. Our conclusions are as follows:

Single-Buyer. Devanur et al. (2015) took the existence of the efficient, zero-revenue equilibrium in the single-buyer setting as motivation for restricting the seller's strategy space. We offer a more nuanced view that reinforces this motivation. We establish a folk theorem that implies that any fixed posted price, including the Myerson optimal price, as well as prices which are too high for any agent to accept, can be implemented in equilibrium. These equilibria essentially simulate commitment. We also show, however, that this multiplicity of equilibria is pathological - subject to a set of natural refinements, the efficient equilibrium is unique. Among our refinements is the *threshold strategies* refinement of Devanur et al. (2015), which confines buyers to using simple strategies where low types reject and high types buy. This makes more rigorous the conclusion of Devanur et al. (2015) that the seller should not expect non-trivial revenue from equilibria of the single-buyer game.

[1]We note that this conclusion is reminiscient of the Coase conjecture; the primary difference is that the Coase conjecture refers to a durable good that a buyer will purchase only once, whereas in the fishmonger problem the good is perishable and can be repurchased each day Coase (1972).

Multiple Buyers, One Item. We next turn to studying a multi-buyer variant of the Fishmonger problem. Suppose now that there are $n \geq 2$ buyers, each buyer's value is drawn iid from a known distribution, and these values are again fixed over all rounds. The seller still has a single copy of the good for sale, and sells that good by posting a single price each day. Each buyer independently chooses whether or not to purchase each day. If multiple buyers wish to purchase at the offered price on a given day, then one of the accepting buyers is chosen uniformly at random to make the purchase.[2]

We show that the seller can achieve a constant fraction of the benchmark optimal revenue in a PBE that survives the refinements we used to select the efficient equilibrium for one buyer. The equilibrium we construct has a natural form, based upon an explore-exploit structure. The seller starts by setting a low price, and slowly raises the price over time as long as at least two buyers purchase in each round. Once all but one (or zero) buyers have stopped purchasing, the seller switches to an exploitation phase in which she posts the highest price at which she believes an agent is guaranteed to buy. Since an agent is guaranteed to buy, the seller stops learning information about the buyers' values, and will post the same price every round from that point onward.

Note that the ascending price clock structure of our equilibrium asymptotically implements the VCG outcome. This intuitive outcome is notable for two reasons. First, it survives the threshold strategies refinement of Devanur et al. (2015). Hence, efficient equilibrium is achievable with buyers who behave in a simple way. Second, note that the highest-valued agent wins *in the limit*: the seller requires several rounds to raise the price and learn which buyer has the highest type. We show that this learning occurs quickly. In particular, we show that for two buyers, if the distribution over buyer valuations satisfies the standard monotone hazard rate (MHR) condition, our equilibrium yields for the seller a $(1/3e^2)$-fraction of the discounted revenue they could achieve from committing to the Myerson optimal mechanism each round.

Multiple Buyers, Digital Goods. We finally consider a multibuyer setting where there is no limit on supply: the seller chooses a price each round, each buyer chooses to accept or reject, and all buyers who

[2]We choose to model the fishmonger problem as a pricing problem, as this is a common approach taken in practice. We note that one could alternatively model it as a general mechanism design problem, which we leave as a direction for future research.

accept receive a copy of the item. As in the single-item settings, we prove the existence of an efficient equilibrium which survives our refinements. For digital goods, this implies that every agent wins a copy of the item, and the seller makes minimal revenue. Surprisingly, however, this not the only natural outcome: we show that the implement supply restriction, raising the price until just l items sell, for any l. Moreover, for two buyers the seller can raise this price (and thus learn buyers' values) quickly enough to yield total discounted revenue which approximates that of the Myerson optimal mechanism run each round, as in the single-item setting.

We finally consider a seller for the digital goods setting who is able to post discriminatory prices, one for each agent each round, rather than the single, anonymous price of our previous results. We show via a slight strengthening of our refinements that the only natural equilibrium is the efficient one, which yields trivial revenue. In particular, the supply restriction which was possible with an anonymous price no longer surivives. Hence, the exogenous imposition of anonymous pricing was crucial for the seller to obtain high revenue. Intuitively, the information leakage with discriminatory prices resembles that of the single-buyer setting, where the seller could only offer trivial prices. Anonymous prices softens the impact of this information leakage.

4.1.1. Related Work

Hart and Tirole (1988) initiated the study of repeated sales ("rental," in their terms) with a single buyer and a large but finite horizon. They consider a special case with just two possible values. They show that in equilibrium the seller will always post the smaller value for all but for a constant number of final rounds. Schmidt (1993) genereralized their result to general discrete distributions. For a survey of this work and the large body of work on closely related models, see the survey of Fudenberg and Villas-Boas (2006). Some variants include Kennan (2001) and Battaglini (2005) who analyze the setting where the value of the buyer is not constant but evolves according to a Markov process, and Conitzer et al. (2012) who study the case where the buyers are short-lived and given the option to anonymize at a cost.

Closest to our work is Devanur et al. (2015), which was the first attempt by the CS community to attempt to move beyond the strong negative results in the setting of Hart and Tirole (1988) and Schmidt

(1993), and the first to consider continuous distributions. Like us, they analyzed threshold equilibria, proving that no such equilibria exist for large but finite numbers of rounds. They go on to study the case of partial commtiment, where the seller can commit to never increase the price in the future. They prove existence of PBE for power law distributions and provide revenue guarantees for the uniform distribution $U[0, 1]$. Note that our results can be directly compared to Devanur et al. (2015) where instead of relaxing the commitment assumptions we introduce an extra buyer, and provide revenue guarantees for a much larger family of distributions.

4.1.2. Discussion

Our results have several interpretations. First, the folk theorem and subsequent elimination of learning equilibria via refinements can be thought of an extension of the conclusions of Hart and Tirole (1988) and Schmidt (1993) to infinite horizon and continuous distributions. This provides further justification for the modified assuptions Devanur et al. (2015) use to derive their results. Our work is similar in that we show that single-buyer market failure is fragile - we use extra buyers rather than partial commitment to support nontrivial equilibria. Finally, we note the prescriptive flavor of our results - our equilibrium and revenue analysis together provide an approximately optimal solution to the problem of dynamic mechanism design in the presence of distrustful buyers.

4.2. Model

Game Description and Timing: The dynamic pricing game takes place in T rounds, where T may be infinite. Each round, there is one item for sale, which must be allocated using a common price among n buyers. Before the game begins, each buyer i draws their value v_i for the goods independently and identically from some continuous distribution F which is common knowledge. The value for allocation remains unchanged from round to round. Each round k then proceeds in the following way:

(1) The seller chooses a price $p^k \geq 0$, which is posted to the buyers.

(2) Buyers simultaneously decide whether to accept p^k.

(3) The item is allocated uniformly at random among the agents who accept.

Utilities: Agents are risk-neutral expected utility maximizers. Utilities are linear in money, additive across rounds, and discounted by a common discount factor $\delta \in (0,1)$ over time. Formally:

- *Seller:* The seller's utility for an outcome to the above game is $\sum_{k\,:\,p^k \text{ accepted}} \delta^k p^k$.
- *Buyer:* The buyer's utility for an outcome is $\sum_{k\,:\,i \text{ wins}} \delta^k v_i$.

Note that all of our results, except the revenue analyses of Sections 4.5.3 and 4.6.3, hold without modification if the seller holds a different discount factor from the buyers. Moreover, the revenue analyses extend in a natural way.

Information: We assume all information and outcomes are common knowledge, with the exception of buyers' values, which are privately held and unknown by all other agents.

Histories: A *history of play* at round k, denoted h^k is different for buyers and the seller, but generally consists of all past pricing and purchasing decisions. Formally, h^k consists of consists of the vector $\mathbf{p}[k-1] = (p^1, \ldots, p^{k-1})$ of past prices, as well as the purchasing decisions of agents in each past round, denoted $\mathbf{D}[k-1] = (\mathbf{D}^1, \ldots, \mathbf{D}^{k-1})$, where $\mathbf{D}^j = (D_1^j, \ldots, D_n^j) \in \{A, R\}^n$ is the vector of accept/reject decisions for each agent i in round j.

Beliefs: The seller does not know any buyer's values, and buyers only know their own. As mentioned earlier, this uncertainty is modeled with a Bayesian prior. After every round of play, the actions of agents may reveal information about their private values, and hence agents' beliefs must be updated. We consider only outcomes where agents' posteriors after each round are shared, which is possible because all actions are commonly observed. The prior for v_i after history h^k, denoted $\mu_i^k(\cdot \mid h^k)$, is a probability measure over the support of F. The joint posterior at round k is denoted $\boldsymbol{\mu}^k = \times_i \mu_i^k$. After round k, the seller believes values are distributed according to $\boldsymbol{\mu}^k$, and buyer i believes other buyers' values are distributed according to $\boldsymbol{\mu}_{-i}^k$.

Strategies: Generally, strategies are maps from histories and private information to actions in round k:

- A seller strategy $\sigma_S^k(h^k)$ specifies for every history h^k a nonnegative price p^k.
- Buyer i's strategy $\sigma_i^k(h^k, p^k; v_i)$ specifies for every buyer history a response to price p^k for every possible value of buyer i.

Equilibrium: Our solution concept is Perfect Bayesian Equilibrium (PBE). Perfect Bayesian Equilibrium imposes joint requirements on beliefs and strategies: beliefs must be updated accurately given strategies, and given beliefs, strategies must form a subgame-perfect equilibrium. Formally, a profile of strategies $\boldsymbol{\sigma} = (\sigma_S^k(\cdot), \sigma_1^k(\cdot), \dots, \sigma_n^k(\cdot))$ and beliefs $\boldsymbol{\mu}^k(\cdot)$ for $k = 0, \dots, T$ is a PBE if:

- *Bayesian updating:* For every history h^k, if there is some v such that $\mu_i^k(v \mid h^k) > 0$ and $\sigma_i^k(h^k, p^k; v) = D_i^k$, then $\mu_i^k(v \mid h^k)$ is computed according to Bayes' rule. Importantly, for histories which would not occur according to $(\sigma_S^k(\cdot), \sigma_1^k(\cdot), \dots, \sigma_n^k(\cdot))$ under any realization of buyers' values, beliefs may be arbitrary.
- *Subgame perfection:* Let $u_S(\boldsymbol{\sigma} \mid h^k, \boldsymbol{\mu}^k)$ denote the expected utility of the seller from the continuation of the game from stage k according to $\boldsymbol{\sigma}$, given that buyers' values are distributed according to $\boldsymbol{\mu}^k(h^k)$. We require that for every alternate strategy σ_S' of the seller, we have that $u_S(\boldsymbol{\sigma} \mid h^k, \boldsymbol{\mu}^k) \geq u_S(\sigma_S', \boldsymbol{\sigma}_{-S} \mid h^k, \boldsymbol{\mu}^k)$. Similarly if $u_i(\boldsymbol{\sigma} \mid h^k, \boldsymbol{\mu}^k, p^k; v_i)$ is the expected utility of a buyer with value v_i offered price p^k under history h^k under beliefs $\boldsymbol{\mu}_{-i}^k(h^k)$ on other buyers' values, $u_i(\boldsymbol{\sigma} \mid h^k, \boldsymbol{\mu}^k, p^k; v_i) \geq u_i(\sigma_i', \boldsymbol{\sigma}_{-i} \mid h^k, \boldsymbol{\mu}^k; v_i)$ for every alternate strategy σ_i'.

Simple Equilibria: Equilibria may in general be extremely complicated. We focus on equilibria satisfying two refinements:

- *Markovian on path:* An equilibrium is *Markovian on path* if on the equilibrium path, agents condition their actions only on the public beliefs and their private information, rather than the complete history. Formally, for any profile of buyer values $\mathbf{v}$ and strategy profile $\boldsymbol{\sigma}$, let h^k and $h^{k'}$ be the histories generated by $\boldsymbol{\sigma}$ under $\mathbf{v}$. If $\boldsymbol{\mu}^k = \boldsymbol{\mu}^{k'}$, then $p^k = p^{k'}$ and $\mathbf{D}^k = \mathbf{D}^{k'}$.
- *Threshold equilibrium:* If a buyer will buy when their value is v_i, they will also buy with any higher value. Formally, a PBE is a *threshold equilibrium* if for each history h^k and price p^k, there is a threshold $t_i(h^k, p^k)$ such that for each agent i, i accepts p^k if and only if $v_i \geq t_i(h^k, p^k)$. Note that in threshold equilibria, updated beliefs derived from on-path histories will be the value distribution F conditioned to some interval $[a, b]$ for each agent. For such equilibria, we will therefore summarize beliefs over agent i's value with the notation F_a^b to denote F conditioned to the interval $[a, b]$.

We refer to threshold equilibria which are Markovian on path as *simple*. Note that simplicity is a refinement rather than a restriction of the strategy space.

4.3. Folk Theorem

We first explore the space of Markovian on path threshold equilibria with no further refinements. It is well-known from previous work on the subject that there exists an equilibrium for the one-buyer case in which the seller gets no revenue and does not learn anything about the buyer's value. The buyers refuse all positive prices, and deviation is punished by the seller with high prices in the future. We refer to this as the no-learning equilibrium, formally, we have:

Theorem 4.3.1 (Devanur et al. (2015))**.** *For $\delta \geq 1/2$ and any number of buyers there is a simple PBE in which the seller does not learn, and posts a price of* 0 *every round. All buyers accept each round.*

The no-learning equilibrium is considered unnatural and unpredictive. In this and the next section, we offer a more nuanced view. We prove a folk theorem: the no-learning equilibrium can be used to enforce other even less intuitive equilibria, including posting any fixed price every round. In other words, PBE is ineffective at ruling out commitment. We solve this problem in Section 4.4, by offering an additional, intuitive refinement which surprisingly eliminates all equilibria but precisely the no-learning equilibrium. This suggests that such behavior is a reasonable outcome to the game.

Theorem 4.3.2 (Folk theorem)**.** *If $\delta \geq \frac{n}{n+1}$, then for any price p, there is a Markovian on path threshold PBE of the dynamic pricing game with n buyers where the seller offers price p every round on the equilibrium path, regardless of the action of the buyer. This holds regardless of the initial prior over buyers' values.*

Proof. We will explicitly construct an equilibrium where the seller offers price p every round, no matter the buyer's action. We give the buyer's strategy in Algorithm 2, and the seller's strategy in Algorithm 1. Beliefs are simple - on-path, they are updated after the first buying decision and remain constant thereafter. If the seller has caused an off-path history by posting a price other than p, then they expect positive prices to be rejected for the rest of time, as in the zero-revenue equilibrium. As in the latter equilibrium, if a buyer accepts a positive price, then the seller assumes that they have value 1 and posts a price of 1 until the end of time.

To see that the strategies in Algorithms 1 and 2 are a PBE, we first argue that the seller is best responding.

ALGORITHM 1: Folk Theorem Equilibrium - Seller's Strategy

Input : History h^k, Support bounds $(a_1^k, b_1^k), \dots, (a_n^k, b_n^k)$
Output: Price p^{k+1}
if $\mathbf{p}[k-1] = (p, \dots, p)$ **then**
  $p^{k+1} = p$;
else if Any buyer has accepted a price other than p or 0 **then**
  $p^{k+1} = 1$;
else
  $p^{k+1} = 0$

ALGORITHM 2: Folk Theorem Equilibrium - Buyer i's Strategy

Input : History h^k, Support bounds $(a_1^k, b_1^k), \dots, (a_n^k, b_n^k)$, Value v_i, Price p^{k+1}
Output: Purchasing decision for round $k+1$
if $\mathbf{p}[k] = (p, \dots, p)$ **then**
  Accept if and only if $v_i \geq p^{k+1}$;
else if Buyer i has accepted a positive price other than p or 0 **then**
  Accept if and only if $v_i \geq p^{k+1}$;
else if $p^{k+1} \neq 0$ **then**
  Reject;
else
  Accept;

We consider the cases as they are stated in Algorithm 1:

- *All prices offered have been p:* In this case, buyers will behave as price takers if $p^{k+1} = p$, which might yield positive revenue. If the seller offers any other positive price, buyers will reject and demand the item for free for the rest of the game.
- *A price other than p has been offered and accepted by buyer i:* This is off-path. We may therefore set the seller's beliefs to be that $v_i = 1$. Moreover, according to the buyers' strategy, i will be a price-taker from now on. It follows that it is optimal for the seller to set a price of 1.
- *A price other than p has been offered, but no buyer has accepted a positive price other than p:* In this case, buyers will only accept a price of 0, so the seller cannot get any utility with any price; they might as well post 0.

We now argue that an arbitrary buyer i is best responding, using the cases in Algorithm 2.

- *All prices offered have been p:* The seller will continue offering this price no matter what the buyer does. It follows that the buyer should accept if they could get positive utility from doing so.
- *Buyer i has accepted a positive price other than p or* 0*:* In this case, the seller believes that $v_i = 1$, and will post price 1 forever. Buyer i should therefore reject, unless their value is 1, in which case they weakly prefer to accept.
- *A price other than p has been offered, the only positve price accepted has been p, and $p^{k+1} \neq 0$.* Buyer i will get at most utility $v_i - p^{k+1}$ from accepting, as all future prices will be 1. Rejecting, meanwhile, will yield utility of $\frac{\delta}{n(1-\delta)}$, as the seller will offer the item for free for all subsequent rounds, and all buyers will accept. If $\delta \geq \frac{n}{n+1}$, then rejecting will be preferable for any choice of v_i and p^{k+1}.
- *A price other than p has been offered, the only positive price accepted has been p, and $p^{k+1} = 0$:* Rejecting will not change the seller's subsequent prices, and accepting will yield positive expected utility, so accepting is optimal.

□

One way to understand the space of PBE is in terms of pairs of attainable payoffs for the buyers and the seller. Theorem 4.3.2 implies that the Pareto frontier of attainable payoffs under our two simplicity refinements is at least as strong as that attainable from posting the same price each round. A natural question is whether there are PBE which surpass this frontier. The best known bounds on the performance of PBE is a theorem due to Devanur et al. (2015), which we rephrase below.

Theorem 4.3.3 (Devanur et al. (2015))**.** *For any target total expected buyer utility U and revenue R attainable in a PBE, there is a mechanism for the single-round game in which the buyers attain total expected utility $(1-\delta)U$ and the seller attains expected revenue $(1-\delta)R$.*

4.4. Non-Robustness of One-Buyer Learning Equilibria

We now consider the case of one buyer and one seller. In this setting, Theorem 4.3.2 proves that there are Markovian on path threshold equilibria which are totally efficient, totally inefficient, and revenue-maximizing, as well as everything in between. We argue that these equilibria exhibit unnatural seller behavior. In particular, in the equilibria of Theorem 4.3.2, there are continuations in which the seller offers prices at the upper boundary of or outside the support of the current beliefs. We prove in this section that every simple equilibrium of the one-buyer case, except those in which no learning occurs on the part of the seller, requires such odd behavior. This leaves only equilibria in which the seller posts a price at the bottom of the support each round. We first formalize "natural" seller behavior.

Definition 4.4.1. A perfect Bayesian threshold equilbrium $\boldsymbol{\sigma}$ of the single-buyer has *natural prices on-path* (or simply *natural prices*) if for every on-path history h^k with beliefs supported on $[a, b]$, the seller's price $\sigma_S^k(h^k)$ lies in $[a, b)$.

Though this requirement might seem mild, it in fact suffices to eliminate all nontrivial equilibria.

Theorem 4.4.1. *In the single-buyer game, let the value distribution F be supported on $[a, b]$, with $a > 0$. If $\delta > 1/2$, then in any simple PBE with natural prices, the seller posts a every round, which is accepted by all buyers. In other words, no learning will occur.*

Proof. Driving our analysis will be an idea from single-dimensional mechanism design: in equilibrium, allocations are monotone in type. In the repeated pricing setting, agents are maximizing their discounted total utility, which is a function of discounted total allocation and discounted total payments. These quantities satisfy the usual incentive constraints from mechanism design, and hence intuitions from mechanism design carry over. We now define these formally:

Definition 4.4.2. Given an PBE of the single-buyer game with some fixed value distribution, let: $x^k(v)$ be an indicator variable of whether or not the buyer with value v purchases in round k on the equilibrium path, and let $p^k(v)$ be the on-path price offered that round. Then we may define the following for any finite $i \geq 1$ and any $j \in \{1, \ldots, \infty\}$:

- The *total discounted allocation*:

$$X(v) = \sum_{k=1}^{\infty} \delta^k x^k(v)$$

- The *total discounted payments*:

$$P(v) = \sum_{k=1}^{\infty} \delta^k p^k(v) x^k(v)$$

- The *total discounted utility*:

$$U(v) = vX(v) - P(v)$$

.

Lemma 4.4.2. *In any PBE of the one-buyer game, the total discounted allocation, payments, and utility (respectively $X(v)$, $P(v)$, and $U(v)$), are nondecreasing in v.*

To prove this claim, we invoke a theorem of Myerson (1981):

Theorem 4.4.3 (Myerson (1981))**.** *Let $f(\cdot)$, $g(\cdot)$ and be functions from some interval $[a, b]$ ($a > 0$) to the positive reals, and assume the following holds for all v and v' in $[a, b]$:*

$$vf(v) - g(v) \geq vf(v') - g(v'). \tag{4.1}$$

Then the following must be true:

(1) $f(\cdot)$ *is nondecreasing in* v.

(2) $g(v) = vf(v) - \int_a^v f(s)\, ds + g(a)$.

The classic application of this theorem sets $f(\cdot)$ to be the equilibrium allocation probability in a single-item auction and $g(\cdot)$ the equilibrium expected payments. We take a similar approach to prove Lemma 4.4.2.

Proof of Lemma 4.4.2. Consider an buyer with value v, who must choose a strategy. Among their options are to pretend to have a different value, say v', and play the actions that value would play. Doing so would yield total discounted allocation $X(v')$, total discounted payments $P(v')$, and total discounted utility $U(v')$. Since the buyer is best responding, it must be that $vX(v)-P(v) \geq vX(v')-P(v')$. We may now invoke Theorem 4.4.3. Monotonicity of $X(\cdot)$ follows from part (1) of the theorem, and monotonicity of $P(\cdot)$ from part (2). Noting that $U(v) = \int_a^v X(s)\,ds - P(a)$ shows $U(\cdot)$ to be nondecreasing as well. □

We will only use monotonicity of allocations here. In Appendix B.1, we make heavier use of Lemma 4.4.2 to derive alternate sufficient conditions under which the conclusions of Theorem 4.4.1 hold.

We now show that natural prices induces non-monotonicity around any threshold t other than the bottom of the support. In particular, we will show that there must exist a type below t with high cumulative allocation, while the threshold type gets allocated strictly less often. This contradicts Lemma 4.4.2.

Lemma 4.4.4. *For any $\delta > 1/2$, consider any simple PBE of the single-buyer game satisfying natural prices with distribution supported on $[a, b]$ and first-round threshold $t > a$. There exists a type $t' < t$ such that $X(t') > 1$.*

Proof. We argue by contradiction. We will assume that for all $t' < t$, $X(t') \leq 1$ and use natural prices, along with simplicity of equilibrium, to show that there is at least one type less than t who would prefer to deviate from the equilibrium.

We first argue that we may assume the existence of some M such that all types in $[a, t)$ have rejected by round M. Assume this is not the case. Then let k_ϵ be the earliest round such that all agents in $[a, t-\epsilon)$ have rejected at least once. If it is the case that $k_\epsilon \to \infty$ as $\epsilon \to 0$, then because $\delta > 1/2$, it must be the case that there exists some $t' < t$ with $X(t') > 1$, which would prove the lemma. Hence we may assume that the number of rounds before every type in $[a, t)$ would reject at least once on the equilibrium path is finite.

Let M an index such that all agents in $[a, t)$ have rejected before round M. We now claim that there is a round $M^* \leq M$ such that a positive measure of types accept in every one of rounds $1, \ldots, M^*-1$, but all such agents reject in round M^*. If not, then it must be that a positive measure of agents accept in every round up to and including M, a contradiction. Let the interval of such agents be $[a^*, t)$. (The upper bound being t is implied by the threshold property.)

Finally, we show that the existence of M^*, combined with natural prices and the Markovian on path property, implies a profitable deviation for some buyer with type in $[a^*, t)$. First note that the beliefs conditioned on acceptance in rounds $1, \ldots, M^* - 1$ do not change after round M^*, as all agents who accepted in rounds $1, \ldots, M^* - 1$ will reject in round M^*. Because beliefs don't change, the Markovian in path property implies that actions don't change - hence, in this continuation, no agent in $[a^*, t)$ accepts after round $M^* - 1$. On the other hand, the requirement of natural prices on path implies that the seller offers a price $p \in [a^*, t)$ in every round after $M^* - 1$. Some type in (p, t) would clearly prefer to accept at least once rather than reject forever, yielding a contradiction. □

Fix a $\delta > 1/2$, and consider a round of the game in which the beliefs are supported on $[a, b]$ and for which the buyer has a nontrivial threshold t (i.e. above the bottom of the support of the current beliefs). Subgame perfection implies that we may assume this round is the first. We know from Lemma 4.4.4 that there is a value $t' < t$ such that $X(t') > 1$. We will show that we may break ties so that $X(t) = 1$, which contradicts Lemma 4.4.2.

By the definition of threshold equilibrium, the buyer with type t accepts this round. Natural prices implies that upon seeing an acceptance, the seller will never price below t. It follows that the buyer with value t will not get utility from any subsequent round. We may therefore assume they reject in every round without changing their utility. Moreover, such tiebreaking doesn't change the incentives of the seller, as the type t buyer has measure 0. Hence, there is an equilibria with $X(t) = 1$ and $X(t') > 1$ for some $t' < t$, contradicting Lemma 4.4.2. □

Moreover, in Appendix B.1 we give an alternate refinement which similarly eliminates learning equilibria. Theorem 4.4.1 and Appendix B.1 together strongly suggest that with just one buyer, one should not expect the seller to learn from purchasing behavior. This strengthens the conclusions of Hart and Tirole (1988); Schmidt (1993) and extends them to continuous distributions. In Sections 4.5 and B.2, we show that these conclusions are critically dependent on the presence of only a single buyer. With multiple buyers, we give a simple PBE with natural prices in which the seller learns from buyers' actions. Moreover, the seller is able to use this knowledge to obtain revenue comparable with the revenue of the optimal auction run in every round.

4.5. Two Buyers, One Item

In what follows, we describe a simple equilibrium with two buyers whose values are independent and identically distributed, with distribution function F, and discount factor $\delta \geq 2/3$ (In Appendix B.2 we generalize this equilibrium to n buyers). This equilibrium has two desirable properties: first, it survives the refinements proposed in Section 4.4, and can therefore be considered robust. Second, the seller gets nontrivial revenue, which stands in contrast to the robust no-learning equilibrium of the single-buyer case.

The equilibrium consists of two phases: an *exploration phase*, and an *exploitation phase*. In the exploration phase, which starts in the very first round and lasts until one or more agents reject, the seller offers prices which will be rejected with positive probability. Consequently, the buyers' response to the seller's exploration prices cause nontrivial updates to the seller's beliefs.

Once an agent rejects, the equilibrium enters the exploitation phase, which lasts until the end of the game. If a single agent triggered the phase by rejecting, the seller ignores this agent, and posts a price at the bottom of the support of the beliefs for the stronger agents. This price is offered and accepted for the rest of the game. If both agents rejected to trigger exploitation, then the seller posts a price at the bottom of the common support of their beliefs. Below, we informally describe the optimization problems of the seller and the buyers to convey the main ideas of the equilibrium.

We first describe the "explore" phase of the game. Assume buyers are identically distributed according to the distribution F_a^b, with support $[a, b]$, and let p^* be the revenue-optimal anonymous price for two buyers with such a distribution. If $p^* = a$, then the seller posts a for the rest of the game (resulting in a voluntary end to the "explore" phase). Otherwise, the seller selects a price p, and assumes that the buyers (who have all accepted in every previous round, and therefore have the same distribution F_a^b after beliefs are updated) will respond according to a threshold t satisfying the threshold equation:

$$(t-p)\left(F_a^b(t)+\tfrac{1-F_a^b(t)}{2}\right)=\frac{t-a}{2}F_a^b(t)\frac{\delta}{1-\delta}. \tag{4.2}$$

The lefthand side represents the utility of buyer i with value t when they accept. If the other buyer rejects, which occurs with probability $F_a^b(t)$, then buyer i gets the item with certainty, at price p. Otherwise, they get it with probability 1/2 at that same price. The righthand side represents the expected utility from rejection. Rejecting triggers the "exploit" phase of the game. If the other buyer rejects as well, then the seller will post a for the remainder of the game, and both buyers will accept for the remainder of the game, yielding the righthand side of (4.12). If the other buyer accepts, the seller will post t for the rest of the game, yielding no utility for buyer i.

If there are multiple thresholds t satisfying (4.12), we will assume buyers use the highest value. Formally, define $T_a^b(p)$ to be the set of solutions to (4.12). Further let p^* be the monopoly price for the initial value distribution. Define $t_a^b(p)$ to be p^* if $p^* \in T_a^b(p)$, else $t_a^b(p) = \max_{t\in T_a^b(p)} t$ if $T_a^b(p) \neq \emptyset$ and infinity otherwise. Given this threshold function, the seller optimizes their revenue, which is given by the recurrence:

$$\begin{aligned}
R_a^b(p) = &\left(1-F_a^b\left(t_a^b\left(p\right)\right)\right)^2\left(p+\delta R^b_{t_a^b(p)}\right)\\
&+2F_a^b\left(t_a^b(p)\right)\left(1-F_a^b\left(t_a^b(p)\right)\right)\left(p+\frac{t_a^b(p)\delta}{1-\delta}\right)\\
&+F_a^b\left(t_a^b\right)^2\left(\frac{a\delta}{1-\delta}\right),
\end{aligned}$$

where $R^b_{t_a^b(p)}$ is the seller's optimal revenue from the continuation game where the buyers are both distributed acccording to $F^b_{t_a^b(p)}$. This equation comes from a straightforward breakdown of the seller's revenue. If both buyers accept, the seller makes revenue p, and both buyers must have value at least $t_a^b(p)$, so the seller is faced the next round with a fresh game with distributions $F^b_{t_a^b(p)}$. If exactly one rejects, then the seller receives revenue p and prices at $t_a^b(p)$ until the end of the game, and this price is accepted. If neither accepts, then the seller prices at a for the rest of the game. Note that the seller may solve the above recurrence for the optimal price p by value iteration.

We now give a full description of the seller's strategy in Algorithm 3, and the buyers' strategy in Algorithm 4. Note that responses to off-path actions will be dictated by careful updates of the beliefs, which we will explain after presenting the strategies.

ALGORITHM 3: Seller's Strategy

```
Input  : Purchasing history h^k, Support bounds (a_1^k, b_1^k), (a_2^k, b_2^k)
Output: Price p^{k+1}
if h^k == (AA)^k then
    a_1^k = a_2^k = a;
    b_1^k = b_2^k = b;
    if a ≥ p* then
        p^{k+1} = a
    else
        p^{k+1} = arg max_p R_a^b(p)
else
    p^{k+1} = max{a_1^k, a_2^k};
```

ALGORITHM 4: Buyer i Strategy

```
Input  : Purchasing history h_b^k, Support bounds (a_1^k, b_1^k), (a_2^k, b_2^k), value v_i, price p^k
Output: Purchasing decision for round k + 1
if h^k == (AA)^k then
    a = a_1^k = a_2^k;
    b = b_1^k = b_2^k;
    if a ≥ p* then
        Accept if and only if p^k ≤ a
    else
        Accept if and only if v_i ≥ t_a^b(p)
else
    if p^i ≤ max{a_1^k, a_2^k} then
        Accept if and only if v_i ≥ p^{k+1}
    else
        Reject
```

We now describe the belief updates. On-path, beliefs are dictated by standard Bayesian updates. In the "explore" phase, the buyers have the same initial support, $[a, b]$. Any buyers who accept price p^k have their support updated to $[t_a^b(p^k), b]$. Those who reject have new support $[a, t_a^b(p^k)]$. In the "exploit" phase, each buyer either always accepts or always rejects, so updates are trivial.

Off-path, we use belief updates to implement punishments for deviation. In particular, if all buyers are expected to take the same action but one deviates, the seller will update their beliefs to the maximum value of the support. For the rest of the game, the seller will post this value, and with probability 1, all buyers will reject. The details of the belief update algorithm are given in Algorithm 5. Together with Algorithms 3 and 4, this fully specifies equilibrium.

ALGORITHM 5: Belief updates

Input : Purchasing history h^{k+1}, Current support bounds $(a_1^k, b_1^k), (a_2^k, b_2^k)$, First-round common support bounds (a, b).
Output: Updated support bounds $(a_1^{k+1}, b_1^{k+1}), (a_2^{k+1}, b_2^{k+1})$.

```
for i ∈ {1,2} do
    if all types for i should reject p^k then
        if buyer i rejected p^k then
            a_i^{k+1} = a_i^k, b_i^{k+1} = b_i^k
        else
            a_i^{k+1} = b, b_i^{k+1} = b
    else if all types should accept then
        if buyer i accepted then
            a_i^{k+1} = a_i^k, b_i^{k+1} = b_i^k
        else
            a_i^{k+1} = b, b_i^{k+1} = b
    else
        if buyer i accepted then
            a_i^{k+1} = t_{a_i^k}^{b_i^k}(p^k), b_i^{k+1} = b_i^k
        else
            a_i^{k+1} = a_i^k, b_i^{k+1} = t_{a_i^k}^{b_i^k}(p^k)
```

Theorem 4.5.1. *The strategies and beliefs specified by Algorithms 3, 4, and 5 is a PBE for* $\delta \geq 2/3$.

Seller. In the explore phase, the seller's optimization problem is an algorithmic pricing problem. Each round, the seller must jointly choose a threshold and a price satisfying the threshold equation (4.12), for current beliefs F supported on $[a, b]$. They know from the buyers' strategies that such prices will be met with a threshold response. It therefore suffices for the seller to maximize the following value function:

$$R(a, b, p) = (1 - F(t(p)))^2(p + \delta R(t(p), b)) + 2(F(t(p)) - F(t(p))^2)(p + \tfrac{t(p)\delta}{1-\delta}) + F(t(p))^2(\tfrac{a\delta}{1-\delta}), \tag{4.3}$$

where $R(x, y)$ is the optimal continuation revenue from the equilibrium with values distributed according to F conditioned to $[x, y]$ and $t(p)$ is the threshold corresponding to price p. The three terms of this function represent the three possible outcomes to the current round: both buyers accept, exactly one accepts, and both reject.

In our presentation of the equilibrium, we leave the specific price path selected by the seller as the implicit solution to the above optimization, and note that any policy for choosing prices and corresponding thresholds satisfying the threshold equation will support threshold behavior by the buyers. To find a policy arbitrarily close to optimal, one may discretize value space and solve the Markov decision problem associated with the value function (4.3) by value iteration, though we make no claim as to the computational efficiency of this method. For a computationally constrained buyer, we give in the next section a particular threshold-supported pricing policy which obtains provably high revenue.

In Section 4.6.1, we show that the buyers are best-responding, and in Section 4.6.2, we show that the seller is best-responding. Together, these results prove the theorem.

4.5.1. Buyer Incentives

As discussed, the belief updates of Algorithm 5 are designed to punish agents for out-of-equilibrium actions. We will use this as a tool for enforcing equilibrium for the buyers. We must first argue that the punishment is effective.

Lemma 4.5.2. *If any buyer takes an out of an equilibrium action in round k, i.e., an action that has zero probability according to the public beliefs, then that buyer does not obtain any utility from rounds $k + 1$ and on.*

Proof. As described in Algorithm 5, once a buyer i accepts or rejects when all buyers should have behaved in the opposite fashion according to their strategy, the public belief about this buyer's type becomes a pointmass on the top of their original distribution's support, b. As a result, in all future rounds the seller posts a price of at least b. □

We now argue that the buyers are in equilibrium.

Lemma 4.5.3. *For $\delta \geq 2/3$, each buyer is best-responding to the strategies of the seller and the other buyer.*

Proof. We break our analysis into two parts: the exploration phase, which occurs when no buyer has rejected yet, and the exploitation phase, where at least one buyer has rejected. These encompass the on-path incentives, but in fact also apply to any history in which only the seller has deviated from equilibrium. If a buyer has taken an off-path action in the history, then Lemma 4.6.2 implies that incentives are trivial.

Exploration Phase. Assume the seller has offered some sequence of prices, all of which have been accepted by both buyers. Then the beliefs about the buyers' values are distributed IID according to F_a^b, for some support interval $[a, b]$. For notational convenience, we will suppress the subscripts and superscripts for the distribution and simply write F. Now assume that in the current round, the seller has posted price p, and assume that there is at least one threshold t satisfying equation (4.12). We will show that for buyer 1 (without loss of generality), accepting p is a best response if $v_1 \geq t$, and rejecting is a best response if $v_1 < t$.

Assume $v_1 \geq t$. The utility for buyer i from accepting, U_A, satisfies

$$U_A \geq \left(F(t) + \tfrac{1-F(t)}{2}\right)(v_1 - p) + \frac{\delta}{1-\delta}(v_1 - t)F(t). \tag{4.4}$$

The first term comes from the current round: buyer 1 wins at price p with probability 1 if buyer 2 rejects, and with probability 1/2 if buyer 2 accepts. The second term comes from the event that buyer 2 rejects this round, in which case the seller will enter the exploitation phase and offer a price of t for the remainder of the game.

Meanwhile, if buyer 1 rejects this round, the seller will enter the exploitation phase regardless of the action of buyer 2. This leaves us with two cases. If buyer 2 rejects, then the seller will post a price of a for the rest of the game (assuming buyer 1 does not attempt further deviations, which by Lemma 4.6.2 are not profitable). If buyer 2 accepts, then the seller will instead post t in the next round. Note that buyer 1 may now accept t next round. Doing so will be profitable, as $v_1 \geq t$, and because the seller expects rejection from buyer 1, the seller will update their beliefs for buyer 1 to be a pointmass on the highest possible value, yielding no further utility. It follows that if buyer 1 rejects this round, their expected utility U_R satisfies:

$$U_R \leq F(t)\frac{v_1 - a}{2} \cdot \frac{\delta}{1-\delta} + \delta(1 - F(t))\frac{v_1 - t}{2}. \tag{4.5}$$

Write $x = v_1 - t$. By assumption, $x \geq 0$. By subtracting expression (4.13) from expression (4.14) and rearranging, we obtain the following lower bound on the margin by which buyer 1 would prefer to accept:

$$U_A - U_R \geq (t + x - p)\frac{1 + F(t)}{2} - (t + x - a)\frac{F(t)\delta}{2 - 2\delta} + x\left(\frac{\delta F(t)}{1-\delta} - \frac{\delta(1 - F(t))}{2}\right).$$

Next, note that we can rearrange the threshold equation (4.12) to get:

$$(F(t) + 1)(t - p) = F(t)(t - a)\frac{\delta}{1-\delta}. \tag{4.6}$$

Applying equation (4.6) to the first two terms of expression (4.15) allows us to write the as a product of x and another term which is clearly positive:

$$U_A - U_R \geq x\left(\frac{1 - F(t)}{2} - \frac{F(t)}{2}\frac{\delta}{1-\delta} + \frac{\delta F(t)}{1-\delta} - \frac{\delta(1 - F(t))}{2}\right). \tag{4.7}$$

Finally, note that when $p < a$, the corresponding threshold from (4.6) is $t = p$. In this case, the incentive to accept is even stronger than the above analysis would suggest, as rejection yields no utility in the present round, and by Lemma 4.6.2, no utility in the future.

Assume $v_1 < t$. If buyer 1 accepts this round, they will win with probability $F(t) + (1 - F(t))/2 = (1 + F(t))/2$. All subsequent prices will be above t, so they will receive no continuation payoff. Their utility from accepting is therefore:

$$U_A = \frac{1 + F(t)}{2}(v_1 - p). \tag{4.8}$$

Meanwhile, if buyer 1 rejects, they trigger the exploitation phase. If buyer 2 also rejects, then the seller offers a for the remainder of the game. If buyer 2 accepts, then the price for the rest of the game is t, in which case buyer 1 cannot obtain any utility. The utility from rejecting is therefore:

$$U_R = \frac{F(t)}{2}\frac{\delta(v_1 - a)}{1-\delta}. \tag{4.9}$$

Combining equations (4.8) and (4.16) yields:

$$U_R - U_A = \frac{F(t)}{2}\frac{\delta(v_1 - a)}{1-\delta} - \frac{1+F(t)}{2}(v_1 - p)$$

After substituting $x = t - v_1$ and applying equation (4.6), we have:

$$U_R - U_A = x\left(\frac{1+F(t)}{2} - \frac{F(t)}{2}\frac{\delta}{1-\delta}\right) \tag{4.10}$$

Note that $v_1 < t$ only if $p > a$. In this case, $t - p < t - a$, and hence equation (4.6) implies that the term in parentheses in (4.17) is positive. It follows that agents with value below t will reject.

Assume (4.12) has no solution. It is possible for the price p to be such that the threshold equation has no solution. In this case, we require that all buyers reject as a best response.

If buyer 1 accepts, they get utility $U_A = v_1 - p$, as buyer 2 will reject this round, and by Lemma 4.6.2, they will receive no utility in the future. If they reject, then the seller will enter the exploit phase of the equilibrium and post a for the rest of the game. This yields utility $U_R = \frac{\delta}{1-\delta}\frac{v_1-a}{2}$.

When the threshold equation (4.12) has no solution, it must be that the lefthand side of the rearranged threshold equation (4.6) is less than the righthand side. This follows from the fact that for $t < p$, the lefthand side is negative side, while the righthand side is positive. Since both sides are continuous functions of t, it cannot be that the lefthand side grows larger than the right, or else they would cross to yield a solution. Consequently, if (4.6) has no solution, it must be that

$$(1+F(t))(t-p) < \frac{\delta}{1-\delta}(t-a)F(t).$$

Rearranging yields:

$$(t-p) < \frac{\delta}{1-\delta}\frac{F(t)}{F(t)+1}(t-a)$$
$$\leq \frac{1}{2}\frac{\delta}{1-\delta}(t-a).$$

This inequality holds for every value of t - in particular, setting $t = v_1$ yields that $U_R > U_A$, as desired.

Exploitation Phase: In the exploitation phase, the seller prices at the bottom of the belief support of the strongest buyer. If the seller offers this price or lower, the buyers act as price-takers. If the seller offers a higher price, both buyers reject. We must show that each of these behaviors is a best response.

Seller offers $p \leq \max(a_1, a_2)$. Note that if an agent who is expected to accept chooses instead to reject, they get no utility in the current round, and no utility in the future, by Lemma 4.6.2. Accepting the current price is clearly preferable. The only case this does not cover is if $a_1 \neq a_2$ and p lies in between. Assume without loss of generality that $a_1 > a_2$. Buyer 2's purchasing decision does not affect future prices, so they should act as a price taker.

Seller offers $p > \max(a_1, a_2)$. We first argue in the case where $a_1 = a_2 = a$. In this case, the utility of buyer 1 (without loss of generality) from rejecting is $\frac{\delta}{1-\delta}\frac{v_1-a}{2}$, as they will win the item with probability 1/2 for the rest of the game. If they accept, they will receive utility $v_1 - p$, as by Lemma 4.6.2, they will not receive utility in subsequent rounds. Since $\delta \geq 2/3$, we have that $\frac{\delta}{2(1-\delta)} \geq 1$. Since $p > a$ as well, we have that $\frac{\delta}{1-\delta}\frac{v_1-a}{2} > v_1 - p$. Hence, rejection is optimal.

In the case where $a_1 \neq a_2$, the incentive to reject is even greater, due to the fact that the higher-valued agent may receive the item with probability 1 in every subsequent round if they reject. Hence, rejection is optimal here as well.

It follows that in both the exploration and exploitation phases, the buyers are best responding to the strategy of the seller. □

4.5.2. Seller Incentives

Lemma 4.5.4. *The seller is best responding to the actions of the buyers.*

Proof. We break our analysis into three cases: exploration phase, exploitation phase, and off-path analysis.

Exploration Phase. In the exploration phase both buyers have the same support, and the seller has three options: price below the common support, post a price p such that there is a threshold response solving equation (4.12), or post a price p such that no solution to (4.12) exists. We will show that the second option, posting a price within the common support such that there is a threshold response, is optimal. Since the seller's equilibrium strategy is defined implicitly to be the optimal such price, it will follow that they are best responding in the exploration phase.

Let a be the lower bound of the common support. We will first show that pricing below a yields less expected revenue than pricing at a. This implies that the seller prefers to post a price within the common support. To prove this claim, observe that any price $p \leq a$ will be accepted with probability 1 by both buyers and cause the beliefs to remain unchanged. Clearly the seller would prefer to induce this outcome with a higher price.

We now argue that posting a price for which a threshold response does not exist is suboptimal. This follows from the fact that both buyers will reject, yielding no revenue and no update to the beliefs. The next round's decision problem is identical to that of the current round, but with payoffs discounted by δ. The seller clearly does not benefit from skipping a round in this manner.

Exploitation Phase. In the exploitation phase, the seller posts $\max(a_1, a_2)$. We will show that the seller prefers this to their other options, which are posting a lower price, or posting a higher price.

The seller prefers not to post a lower price for the same reason that they would prefer not to post below the common support in the exploration phase. As long as the price p satisfies $p \leq \max(a_1, a_2)$, it will be accepted with probability 1, and will not cause the beliefs about the stronger of the two buyers to change. Posting the largest price which induces this outcome is preferable to other such prices.

If the seller posts a price $p > \max(a_1, a_2)$, then this price is rejected by all agents, and beliefs do not change. The seller effectively skips the round and is faced with the same decision next round, with a discount. This is not optimal. Hence, the optimal price to post in the exploitation phase is $\max(a_1, a_2)$.

Off-Path Analysis. We finally analyze the case where a buyer has taken an action in the current history that was not expected of any type. In this case, the seller updates their beliefs to the highest possible value b, and the buyer strategy dictates that they behave as a price-taker. The optimal response to such a scenario is to post b every round. $\square$

4.5.3. Revenue Guarantees

We now argue that if distributions are well-behaved, the revenue of the equilibrium outlined in the previous section has high revenue. Specifically, we assume that the *hazard rate* $\frac{f(v)}{1-F(v)}$ of the distribution is increasing in v - a standard assumption in mechanism design. As a benchmark, we use the revenue that the seller would obtain if they used the optimal auction in every round. For example, with two $U[0,1]$ buyers (as is considered in Devanur et al. (2015)), the seller obtains $5/12$ every round in expectation, yielding a benchmark of $\frac{5}{12(1-\delta)}$. By Theorem 4.3.3, this revenue is an upper bound on the seller's revenue in any PBE.

Theorem 4.5.5. *Assume the value distribution F of the two buyers has a monotone hazard rate, and assume $\delta \geq 2/3$. In the equilibrium described in Section 4.5, the seller obtains revenue which is at least $\frac{1}{3e^2}$ of the revenue of the optimal auction run each round.*

Proof. To prove the theorem, we give a pricing strategy which is supported by threshold behavior from the buyers and which obtains approximately optimal revenue. Because this is a feasible strategy for the seller, their equilibrium strategy will obtain at least as much revenue. The intuition for this pricing strategy is that the seller will try to offer the monopoly price for the distribution F, i.e. the price p^* maximizing $p(1 - F(p))$, as soon as they are able. If this price is too high to support a threshold response, they will instead try to learn as aggressively as possible from accept decisions, choosing a price which induces a very high threshold. Because the number of buyers is small and the quantile $1 - F(p^*)$ of the monopoly price is not too low, the revenue from offering the monopoly price two two buyers will approximate the revenue of the optimal mechanism for those same buyers. It therefore suffices to show that by aggressively learning, the seller is able to quickly reach point where they are posting the monopoly price. Note that this sequence of prices is easy to compute. To argue the theorem, we first observe that in the exploration phase of the equilibrium, the seller may offer any price which has a threshold response, and the arguments in the previous section ensure that buyers will be incentivized to adhere to threshold responses. It follows that we may analyze any sequence of prices for the exploration phase, and as long as each price has a threshold response, the resulting revenue will be a lower bound on the actual revenue of the seller.

As our upper bound on our benchmark, we use $\frac{2(1-F(p^*))p^*}{1-\delta}$, where p^* is the single-buyer monopoly price. This corresponds to selling two items optimally every round. Because the optimal revenue for the sale of a single item is concave in the number of buyers, this upper bound will always exceed the revenue of the optimal single-item auction every round.

To relate our equilibrium to this upper bound, we imagine the seller choosing a sequence of prices which increases the threshold quickly until it reaches p^*, after which the seller voluntarily enters the exploit phase. Assuming both agents have value above p^*, the seller will receive revenue of p^* in perpetuity starting as soon as the threshold reaches this point. By upperbounding the time it takes for this to occur, we can lowerbound the expected revenue from this sequence of prices, and therefore the revenue from the price sequence actually selected by the seller in equilibrium.

First, consider an arbitrary step of the explore phase, where the current beliefs are over an interval $[a, b]$ with CDF F_a^b. We argue that there is always a way for the seller to induce a threshold t which learns "quickly." Formally:

Lemma 4.5.6. *In the explore phase with beliefs supported on $[a, b]$, there always exists a price $p \geq a$ inducing the threshold t which satisfies $F_a^b(t) = \frac{1-\delta}{\delta}$.*

Proof. Note that the threshold equation for this stage implies:

$$(t-a)F_a^b(t)\frac{\delta}{1-\delta} = (F_a^b(t)+1)(t-p).$$

Substituting in $F_a^b(t) = \frac{1-\delta}{\delta}$ and solving for p yields $p = t - \delta(t-a)$. □

To obtain our bound, we will assume the seller offers the following sequence of prices:

- If there exists some $p \in [a, b]$ inducing threshold p^*, offer p.
- Otherwise, offer a price which induces t satisfying $F_a^b(t) = \frac{1-\delta}{\delta}$.

We now argue that such a sequence of prices will eventually induce a threshold of p^*, if buyers' values are above p^*.

Lemma 4.5.7. *If both sellers have value at least p^*, then the above sequence of prices eventually induces threshold p^*.*

Proof. By Lemma 4.6.6, the seller will eventually reach a stage where the threshold t satisfying $F_a^b(t) = \frac{1-\delta}{\delta}$ is greater than p^*. We show that in this case, there is a price inducing a threshold $t = p^*$.

To see this, assume the current beliefs for buyers who haven't rejected are distributed according to F_a^b with support $[a, b]$. Let t^* be the threshold for which $F_a^b(t^*) = \frac{1-\delta}{\delta}$, and assume $t^* > p^*$. Note that the threshold equation can be rearranged as:

$$\frac{t-p}{t-a} = \frac{F_a^b(t)}{F_a^b(t)+1}\frac{\delta}{1-\delta} \tag{4.11}$$

The righthand side is obviously increasing in t. Since $t^* > p^*$, we therefore have:

$$0 < \frac{F_a^b(p^*)}{F_a^b(p^*)+1}\frac{\delta}{1-\delta} < \frac{F_a^b(t^*)}{F_a^b(t^*)+1}\frac{\delta}{1-\delta} < 1.$$

If we set $t = p^*$, we see that the lefthand side ranges from 0 at $p = p^*$ to 1 at $p = a$, hence, there is a price p that induces the desired threshold. □

We can now argue that under the above sequence of prices, the exploration phase will reach threshold p^* quickly if both agents have values above p^*. Formally:

Lemma 4.5.8. *Let $x = \frac{1-\delta}{\delta}$. If both sellers have value at least p^* and $F(p^*) \leq 1 - 1/e$ then after $1/x + 1$ rounds of the exploration phase using the price sequence defined above, we will have that the lower bound of the support is p^*.*

Proof. Let t_j be the threshold induced in the jth stage of the exploration phase, and assume $t_j < p^*$. We will first lowerbound $F(t_j)$. Note that $F(t_j)$ is exactly the probability that an agent will reject one of the prices in first j stages of the learning phase. This probability can also be written as $F(t_{j-1}) + (1 - F(t_{j-1}))x$, since conditioned on an agent accepting the first $j-1$ prices, the price in round p^j was chosen to be accepted with probability x. This yields the recurrence:

$$F(t_j) = F(t_{j-1}) + (1 - F(t_{j-1}))x,$$

which is solved by $F(t_j) = 1 - (1-x)^j$. If we set $j = 1/x$, we obtain

$$F(t_j) = 1 - (1-x)^{1/x} \geq 1 - 1/e.$$

It therefore must be that after at most $1/x$ rounds, the exploration phase terminates with the threshold reaching p^*. In the subsequent round, the lower bound of the support will be the previous round's threshold, p^*. □

Lemma 4.5.9. *If $F(p^*) \leq 1 - 1/e$ then the equilibrium obtains revenue at least $\frac{1}{1-\delta}\frac{2}{3e}p^*(1-F(p^*))^2$.*

Proof. We lower bound the revenue with the revenue from the sequence of prices described above. The probability that both agents have values above p^*, and therefore that the threshold reaches p^*, is $(1-F(p^*))^2$. By Lemma 4.6.8 the discount factor after reaching p^* is at most $\delta^{1+\frac{\delta}{1-\delta}} \geq \frac{2}{3e}$. After p^* is reached the seller prices the item at p^* for all remaining rounds and the item is accepted with probability one for a total of $\frac{1}{1-d}p^*$ revenue. Overall the revenue obtained by the seller with this price sequence is therefore at least $\frac{1}{1-\delta}\frac{2}{3e}p^*(1-F(p^*))^2$. □

Let OPT denote the total revenue from running the optimal auction for two buyers on F. Our benchmark for revenue is $\mathrm{OPT}/(1-\delta)$. By concavity of the revenue we have that $OPT \leq 2p^*(1-F(p^*))$. By Lemma 4.6.9 the equilibrium gets revenue

$$\frac{1}{1-\delta}\frac{2}{3e}p^*(1-F(p^*))^2 \quad \geq \quad \frac{\mathrm{OPT}}{1-\delta}\frac{1}{3e}(1-F(p^*))$$

For distributions satisfying the monotone hazard rate assumption, it is a standard fact that $F(p^*) \geq 1 - 1/e$. We therefore have that the revenue of our equilibrium is at least $\frac{\text{OPT}}{3(1-\delta)e^2} \geq \frac{\text{OPT}}{3e^2}$, proving the theorem. □

4.6. Two Buyers and Digital goods

We describe an equilibrium for two buyers with natural thresholds in which the seller makes nontrivial revenue and updates their beliefs in a nontrivial way (In Appendix B.3 we extend this equilibrium to n buyers). The equilibrium has two phases: an "explore" phase in which learning occurs, and an "exploit" phase, where the seller ceases to learn from buying behavior and instead posts the bottom of the strongest player's support for the rest of the game. The "explore" phase lasts from the beginning of the game until any buyer rejects a price. The "exploit" phase begins when a buyer rejects and continues for the rest of the game.

We first describe the "explore" phase of the game. Assume buyers' common prior at the start of the game is given by F, and assume the current beliefs are identically distributed according to the distribution F_a^b, with support $[a, b]$. Let p^* be the revenue-optimal anonymous price for selling a single item to two buyers with the intitial distribution F. If $p^* = a$, then the seller posts a for the rest of the game (resulting in a voluntary end to the "explore" phase). Otherwise, the seller selects a price p, and assumes that the buyers (who have all accepted in every previous round, and therefore have the same distribution F_a^b after beliefs are updated) will respond according to a threshold t satisfying the threshold equation:

$$(t - p) = (t - a)F_a^b(t)\frac{\delta}{1 - \delta}. \tag{4.12}$$

The lefthand side represents the utility of buyer i with value t when they accept. They get utility $t - p$ for buying this round. All subsequent prices will be at least t, so such an agent can expect no future utility. The righthand side represents the expected utility of such an agent from rejection. Rejecting triggers the "exploit" phase of the game. If the other buyer rejects as well, then the seller will post a for the remainder of the game, and both buyers will accept for the remainder of the game, yielding the righthand side of (4.12). If the other buyer accepts, the seller will post t for the rest of the game, yielding no utility for buyer i. If at any point of the game, the seller posts a price which is below the minimum value of the highest buyer's support, the buyers punish the seller with no-learning behavior.

If there are multiple thresholds t satisfying (4.12), we will assume buyers use the highest value. Formally, define $T_a^b(p)$ to be the set of solutions to (4.12). Further let p^* be the optimal anonymous price for the initial value distribution F. Define $t_a^b(p)$ to be p^* if $p^* \in T_a^b(p)$, else $t_a^b(p) = \max_{t \in T_a^b(p)} t$ if $T_a^b(p) \neq \emptyset$ and infinity otherwise. Given this threshold function, the seller optimizes their revenue, which is given by the recurrence:

$$\begin{aligned} R_a^b(p) = & \left(1 - F_a^b\left(t_a^b\left(p\right)\right)\right)^2 \left(2p + \delta R^b_{t_a^b(p)}\right) \\ & + 2F_a^b\left(t_a^b(p)\right)\left(1 - F_a^b\left(t_a^b(p)\right)\right)\left(p + \frac{t_a^b(p)\delta}{1-\delta}\right) \\ & + F_a^b\left(t_a^b\right)^2 \left(\frac{2a\delta}{1-\delta}\right), \end{aligned}$$

where $R^b_{t_a^b(p)}$ is the seller's optimal revenue from the continuation game where the buyers are both distributed acccording to $F^b_{t_a^b(p)}$. This equation comes from a straightforward breakdown of the seller's revenue. If both buyers accept, the seller makes revenue $2p$, and both buyers must have value at least $t_a^b(p)$, so the seller is faced the next round with a fresh game with distributions $F^b_{t_a^b(p)}$. If exactly one rejects, then the seller receives revenue p and prices at $t_a^b(p)$ until the end of the game, and this price is accepted by sole buyer who accepted. If neither accepts, then the seller prices at a for the rest of the game and sells to both buyers. Note that the seller may solve the above recurrence for the optimal price p by value iteration.

We now give a full description of the seller's strategy in Algorithm 6, and the buyers' strategy in Algorithm 7. Note that responses to off-path actions will be dictated by careful updates of the beliefs, which we will explain after presenting the strategies.

ALGORITHM 6: Seller's Strategy

Input : Purchasing history h^k, Support bounds $(a_1^k, b_1^k), (a_2^k, b_2^k)$
Output: Price p^{k+1}
if $p^i < \max\{a_1^i, a_2^i\}$ for any $i \leq k$ **then**
  Let i^* be the earliest round where $p^{i^*} < \max\{a_1^{i^*}, a_2^{i^*}\}$.
  if Any buyer has accepted a positive price after round i^* **then**
    $p^{k+1} = b$;
  else if Buyer has ever rejected a price of 0 **then**
    $p^{k+1} = b$;
  else
    $p^{k+1} = 0$;
else if $h^k == (AA)^k$ **then**
  $a_1^k = a_2^k = a$;
  $b_1^k = b_2^k = b$;
  if $a \geq p^*$ **then**
    $p^{k+1} = a$
  else if then
    $p^{k+1} = \arg\max_p R_a^b(p)$
else
  $p^{k+1} = \max\{a_1^k, a_2^k\}$;

We now describe the belief updates. On-path, beliefs are dictated by standard Bayesian updates. In the "explore" phase, the buyers have the same initial support, $[a, b]$. Any buyers who accept price p^k have their support updated to $[t_a^b(p^k), b]$. Those who reject have new support $[a, t_a^b(p^k)]$. In the "exploit" phase, each buyer either always accepts or always rejects, so updates are trivial.

Off-path, we use belief updates to implement punishments for deviation. In particular, if all buyers are expected to take the same action (e.g. reject a price for which no threshold exists, or reject a positive price in a zero-learning continuation) but one deviates, the seller will update their beliefs to the maximum value of the support. For the rest of the game, the seller will post this value, and with probability 1, all buyers will reject. The details of the belief update algorithm are given in Algorithm 5. Together with Algorithms 3 and 4, this fully specifies equilibrium.

Theorem 4.6.1. *The strategies and beliefs specified by Algorithms 6, 7, and 8 is a PBE for* $\delta \geq 1/2$.

In Section 4.6.1, we show that the buyers are best-responding, and in Section 4.6.2, we show that the seller is best-responding. Together, these results prove the theorem.

ALGORITHM 7: Buyer i Strategy

Input : Purchasing history h_b^k, Support bounds $(a_1^k, b_1^k), (a_2^k, b_2^k)$, value v_i, price p^k
Output: Purchasing decision for round $k+1$

```
if p^k < max{a_1^k, a_2^k} then
    Reject;
if p^i < max{a_1^i, a_2^i} for any i ≤ k − 1 then
    if p^k = 0 then
        Accept;
    else if p^k > 0 then
        Let i* be the earliest round where p^{i*} < max{a_1^{i*}, a_2^{i*}}.
        if Buyer has ever accepted a positive price after round i* then
            Accept;
        else if Buyer has ever rejected a price of 0 then
            Accept;
        else
            Reject;
else if h^k == (AA)^k then
    a = a_1^k = a_2^k;
    b = b_1^k = b_2^k;
    if a ≥ p* then
        Accept if and only if p^k ≤ a
    else
        Accept if and only if v_i ≥ t_a^b(p)
else
    if p^i = max{a_1^k, a_2^k} then
        Accept if and only if v_i ≥ p^{k+1}
    else
        Reject
```

4.6.1. Buyer Incentives

As discussed, the belief updates of Algorithm 5 are designed to punish agents for out-of-equilibrium actions. We will use this as a tool for enforcing equilibrium for the buyers. We must first argue that the punishment is effective.

Lemma 4.6.2. *If any buyer takes an out of an equilibrium action in round k, i.e., an action that has zero probability according to the public beliefs, then that buyer does not obtain any utility from rounds $k+1$ and on.*

Proof. As described in Algorithm 5, once a buyer i accepts or rejects when all buyers should have behaved in the opposite fashion according to their strategy, the public belief about this buyer's type

ALGORITHM 8: Belief updates

Input : Purchasing history h^{k+1}, Current support bounds $(a_1^k, b_1^k), (a_2^k, b_2^k)$, First-round common support bounds (a, b).
Output: Updated support bounds $(a_1^{k+1}, b_1^{k+1}), (a_2^{k+1}, b_2^{k+1})$.

```
for i ∈ {1,2} do
    if all types for i should reject p^k then
        if buyer i rejected p^k then
            a_i^{k+1} = a_i^k, b_i^{k+1} = b_i^k
        else
            a_i^{k+1} = b, b_i^{k+1} = b
    else if all types should accept then
        if buyer i accepted then
            a_i^{k+1} = a_i^k, b_i^{k+1} = b_i^k
        else
            a_i^{k+1} = 2 · b, b_i^{k+1} = 2 · b
    else
        if buyer i accepted then
            a_i^{k+1} = t_{a_i^k}^{b_i^k}(p^k), b_i^{k+1} = b_i^k
        else
            a_i^{k+1} = a_i^k, b_i^{k+1} = t_{a_i^k}^{b_i^k}(p^k)
```

becomes a pointmass on the top of their original distribution's support, b. As a result, in all future rounds the seller posts a price of at least b. □

We now argue that the buyers are in equilibrium.

Lemma 4.6.3. *For $\delta \geq 1/2$, each buyer is best-responding to the strategies of the seller and the other buyer.*

Proof. We break our analysis into two parts: the exploration phase, which occurs when no buyer has rejected yet, and the exploitation phase, where at least one buyer has rejected. These encompass the on-path incentives, but in fact also apply to any history in which only the seller has deviated from equilibrium. If a buyer has taken an off-path action in the history, then Lemma 4.6.2 implies that incentives are trivial. If the seller has posted a price below the bottom of the support of the strongest buyer, the incentives are identical to the no-learning equilibrium.

Exploration Phase. Assume the seller has offered some sequence of prices, all of which have been accepted by both buyers. Then the beliefs about the buyers' values are distributed IID according to

F_a^b, for some support interval $[a, b]$. For notational convenience, we will suppress the subscripts and superscripts for the distribution and simply write F. Now assume that in the current round, the seller has posted price p, and assume that there is at least one threshold t satisfying equation (4.12). We will show that for buyer 1 (without loss of generality), accepting p is a best response if $v_1 \geq t$, and rejecting is a best response if $v_1 < t$.

Assume $v_1 \geq t$. The utility for buyer i from accepting, U_A, satisfies

$$U_A \geq (v_1 - p) + \frac{\delta}{1-\delta}(v_1 - t)F(t). \tag{4.13}$$

The first term comes buying at price p in the current round. The second term comes from the event that buyer 2 rejects this round, in which case the seller will enter the exploitation phase and offer a price of t for the remainder of the game.

Meanwhile, if buyer 1 rejects this round, the seller will enter the exploitation phase regardless of the action of buyer 2. This leaves us with two cases. If buyer 2 rejects, then the seller will post a price of a for the rest of the game (assuming buyer 1 does not attempt further deviations, which by Lemma 4.6.2 are not profitable). If buyer 2 accepts, then the seller will instead post t in the next round. Note that buyer 1 may now accept t next round. Doing so will be profitable, as $v_1 \geq t$, and because the seller expects rejection from buyer 1, the seller will update their beliefs for buyer 1 to be a pointmass on the highest possible value, yielding no further utility. It follows that if buyer 1 rejects this round, their expected utility U_R satisfies:

$$U_R \leq F(t)(v_1 - a) \cdot \frac{\delta}{1-\delta} + \delta(1 - F(t))(v_1 - t). \tag{4.14}$$

Write $x = v_1 - t$. By assumption, $x \geq 0$. By subtracting expression (4.13) from expression (4.14) and rearranging, we obtain the following lower bound on the margin by which buyer 1 would prefer to accept:

$$U_A - U_R \geq t + x + p + \frac{\delta}{1-\delta}F(t)x - F(t)(t + x - a)\frac{\delta}{1-\delta} - \delta(1 - F(t))x$$

Rearranging and applying the threshold equation, we obtain the product of x and another term which is clearly positive:

$$U_A - U_R \geq x(1 - \delta(1 - F(t))). \tag{4.15}$$

Assume $v_1 < t$. If buyer 1 accepts this round, all subsequent prices will be above t, so they will receive no continuation payoff. Their utility from accepting is therefore simply $U_A = (v_1 - p)$.

Meanwhile, if buyer 1 rejects, they trigger the exploitation phase. If buyer 2 also rejects, then the seller offers a for the remainder of the game. If buyer 2 accepts, then the price for the rest of the game is t, in which case buyer 1 cannot obtain any utility. The utility from rejecting is therefore:

$$U_R = F(t)\frac{\delta(v_1 - a)}{1 - \delta}. \tag{4.16}$$

We therefore have:

$$U_R - U_A = F(t)\frac{\delta(v_1 - a)}{1 - \delta} - (v_1 - p)$$

After substituting $x = t - v_1$ and applying the threshold equation, we have:

$$U_R - U_A = x\left(\frac{1 + F(t)}{2} - \frac{F(t)}{2}\frac{\delta}{1 - \delta}\right) \tag{4.17}$$

Note that $v_1 < t$ only if $p > a$. In this case, $t - p < t - a$, and hence the threshold equation implies that the term in parentheses in (4.17) is positive. It follows that agents with value below t will reject.

Assume (4.12) has no solution. It is possible for the price p to be such that the threshold equation has no solution. In this case, we require that all buyers reject as a best response.

If buyer 1 accepts, they get utility $U_A = v_1 - p$, as buyer 2 will reject this round, and by Lemma 4.6.2, they will receive no utility in the future. If they reject, then the seller will enter the exploit phase of the equilibrium and post a for the rest of the game. This yields utility $U_R = \frac{\delta}{1-\delta}(v_1 - a)$. The threshold equation implies that $U_R > U_A$, as desired.

Exploitation Phase: In the exploitation phase, the seller prices at the bottom of the belief support of the strongest buyer. If the seller offers this price or lower, the buyers reject, and will only accept prices of zero for the rest of the game.. If the seller offers a higher price, both buyers reject. We must show that each of these behaviors is a best response.

Seller offers $p \leq \max(a_1, a_2)$. The agent can get utility from accepting the current round, but will face the highest possible price for the rest of the game if they do so. They would prefer to reject and receive their item for free for the rest of the game, in a zero-learning continuation.

Seller offers $p > \max(a_1, a_2)$. We first argue in the case where $a_1 = a_2 = a$. In this case, the utility of buyer 1 (without loss of generality) from rejecting is $\frac{\delta}{1-\delta}(v_1 - a)$, as they be allocated every round for the rest of the game. If they accept, they will receive utility $v_1 - p$, as by Lemma 4.6.2, they will not receive utility in subsequent rounds. Since $\delta \geq 1/2$, we have that $\frac{\delta}{1-\delta} \geq 1$. Since $p > a$ as well, we have that $\frac{\delta}{1-\delta} v_1 > v_1 - p$. Hence, rejection is optimal.

In the case where $a_1 \neq a_2$, the incentive to reject is even greater, due to the fact that the higher-valued agent may receive the item with probability 1 in every subsequent round if they reject. Hence, rejection is optimal here as well.

It follows that in both the exploration and exploitation phases, the buyers are best responding to the strategy of the seller. □

4.6.2. Seller Incentives

Lemma 4.6.4. *The seller is best responding to the actions of the buyers.*

Proof. We break our analysis into three cases: exploration phase, exploitation phase, and off-path analysis.

Exploration Phase. In the exploration phase both buyers have the same support, and the seller has three options: price below the common support, post a price p such that there is a threshold response solving equation (4.12), or post a price p such that no solution to (4.12) exists. We will show that the second option, posting a price within the common support such that there is a threshold response, is optimal. Since the seller's equilibrium strategy is defined implicitly to be the optimal such price, it will follow that they are best responding in the exploration phase.

Let a be the lower bound of the common support. Pricing below a will cause both buyers to reject any nonzero price for the remainder of the game. This implies that the seller prefers to post a price within the common support.

We now argue that posting a price for which a threshold response does not exist is suboptimal. This follows from the fact that both buyers will reject, yielding no revenue and no update to the beliefs. The next round's decision problem is identical to that of the current round, but with payoffs discounted by δ. The seller clearly does not benefit from skipping a round in this manner.

Exploitation Phase. In the exploitation phase, the seller posts $\max(a_1, a_2)$. We will show that the seller prefers this to their other options, which are posting a lower price, or posting a higher price.

The seller prefers not to post a lower price for the same reason that they would prefer not to post below the common support in the exploration phase. As long as the price p satisfies $p < \max(a_1, a_2)$, the buyers will reject any nonzero price.

If the seller posts a price $p > \max(a_1, a_2)$, then this price is rejected by all agents, and beliefs do not change. The seller effectively skips the round and is faced with the same decision next round, with a discount. This is not optimal. Hence, the optimal price to post in the exploitation phase is $\max(a_1, a_2)$.

Off-Path Analysis. We finally analyze the case where a buyer has taken an action in the current history that was not expected of any type. In this case, the seller updates their beliefs to the highest possible value b, and the buyer strategy dictates that they behave as a price-taker. The optimal response to such a scenario is to post b every round. □

4.6.3. Revenue guarantees

We now argue that if distributions are well-behaved, the revenue of the equilibrium outlined above has high revenue. Specifically, we assume that the *hazard rate* $\frac{f(v)}{1-F(v)}$ of the distribution is increasing in v - a standard assumption in mechanism design. As a benchmark, we use the revenue that the seller would obtain if they used the optimal auction in every round. For example, with two $U[0, 1]$ buyers (as is considered in Devanur et al. (2015)), the seller obtains $5/12$ every round in expectation, yielding a benchmark of $\frac{5}{12(1-\delta)}$. By Theorem 4.3.3, this revenue is an upper bound on the seller's revenue in any PBE. The result is the following:

Theorem 4.6.5. *Assume the value distribution F of the two buyers has a monotone hazard rate, and assume $\delta \geq 1/2$. In the equilibrium described above, the seller obtains revenue which is at least $\frac{1}{2e^2}$ of the revenue of the optimal auction run each round.*

To argue the theorem, we first observe that in the exploration phase of the equilibrium, the seller may offer any price which has a threshold response, and the arguments in the previous section ensure that buyers will be incentivized to adhere to threshold responses. It follows that we may analyze any sequence of prices for the exploration phase, and as long as each price has a threshold response, the resulting revenue will be a lower bound on the actual revenue of the seller.

As our upper bound on our benchmark, we use $\frac{2(1-F(p^*))p^*}{1-\delta}$, where p^* is the single-buyer monopoly price. This corresponds to selling two items optimally every round.

To relate our equilibrium to this upper bound, we imagine the seller choosing a sequence of prices which increases the threshold quickly until it reaches p^*, after which the seller voluntarily enters the exploit phase. Assuming both agents have value above p^*, the seller will receive revenue of $2p^*$ in perpetuity starting as soon as the threshold reaches this point. By upperbounding the time it takes for this to occur, we can lowerbound the expected revenue from this sequence of prices, and therefore the revenue from the price sequence actually selected by the seller in equilibrium.

First, consider an arbitrary step of the explore phase, where the current beliefs are over an interval $[a, b]$ with CDF F_a^b. We argue that there is always a way for the seller to induce a threshold t which learns "quickly." Formally:

Lemma 4.6.6. *In the explore phase with beliefs supported on $[a, b]$, there always exists a price $p \geq a$ inducing the threshold t which satisfies $F_a^b(t) = \frac{1-\delta}{\delta}$.*

Proof. Note that the threshold equation for this stage implies:

$$(t-a)F_a^b(t)\frac{\delta}{1-\delta} = (t-p).$$

Hence setting $p = a$ satisfies $F_a^b(t(p)) = \frac{1-\delta}{\delta}$. □

To obtain our bound, we will assume the seller offers the following sequence of prices:

- If there exists some $p \in [a, b]$ inducing threshold p^*, offer p.
- Otherwise, offer a price which induces t satisfying $F_a^b(t) = \frac{1-\delta}{\delta}$.

We now argue that such a sequence of prices will eventually induce a threshold of p^*, if buyers' values are above p^*.

Lemma 4.6.7. *If both sellers have value at least p^*, then the above sequence of prices eventually induces threshold p^*.*

Proof. By Lemma 4.6.6, the seller will eventually reach a stage where the threshold t satisfying $F_a^b(t) = \frac{1-\delta}{\delta}$ is greater than p^*. We show that in this case, there is a price inducing a threshold $t = p^*$.

To see this, assume the current beliefs for buyers who haven't rejected are distributed according to F_a^b with support $[a, b]$. Let t^* be the threshold for which $F_a^b(t^*) = \frac{1-\delta}{\delta}$, and assume $t^* > p^*$. Note that the threshold equation can be rearranged as:

$$\frac{t-p}{t-a} = F_a^b(t)\frac{\delta}{1-\delta} \tag{4.18}$$

The righthand side is obviously increasing in t. Since $t^* > p^*$, we therefore have:

$$0 < F_a^b(p^*)\frac{\delta}{1-\delta} < F_a^b(t^*)\frac{\delta}{1-\delta} < 1.$$

If we set $t = p^*$, we see that the lefthand side ranges from 0 at $p = p^*$ to 1 at $p = a$, hence, there is a price p that induces the desired threshold. □

We can now argue that under the above sequence of prices, the exploration phase will reach threshold p^* quickly if both agents have values above p^*. Formally:

Lemma 4.6.8. *Let $x = \frac{1-\delta}{\delta}$. If both sellers have value at least p^* and $F(p^*) \leq 1 - 1/e$ then after $1/x + 1$ rounds of the exploration phase using the price sequence defined above, we will have that the lower bound of the support is p^*.*

Proof. Let t_j be the threshold induced in the jth stage of the exploration phase, and assume $t_j < p^*$. We will first lowerbound $F(t_j)$. Note that $F(t_j)$ is exactly the probability that an agent will reject one of the prices in first j stages of the learning phase. This probability can also be written as $F(t_{j-1}) + (1 - F(t_{j-1}))x$, since conditioned on an agent accepting the first $j-1$ prices, the price in round p^j was chosen to be accepted with probability x. This yields the recurrence:

$$F(t_j) = F(t_{j-1}) + (1 - F(t_{j-1}))x,$$

which is solved by $F(t_j) = 1 - (1-x)^j$. If we set $j = 1/x$, we obtain

$$F(t_j) = 1 - (1-x)^{1/x} \geq 1 - 1/e.$$

It therefore must be that after at most $1/x$ rounds, the exploration phase terminates with the threshold reaching p^*. In the subsequent round, the lower bound of the support will be the previous round's threshold, p^*. □

Lemma 4.6.9. *If $F(p^*) \leq 1 - 1/e$ then the equilibrium obtains revenue at least $\frac{1}{1-\delta}\frac{1}{2e}p^*(1 - F(p^*))^2$.*

Proof. We lower bound the revenue with the revenue from the sequence of prices described above. The probability that both agents have values above p^*, and therefore that the threshold reaches p^*, is $(1 - F(p^*))^2$. By Lemma 4.6.8 the discount factor after reaching p^* is at most $\delta^{1+\frac{\delta}{1-\delta}} \geq \frac{1}{2e}$. After p^* is reached the seller price at p^* for all remaining rounds and the item is accepted with probability one for a total of $\frac{2}{1-\delta}p^*$ revenue. Overall the revenue obtained by the seller with this price sequence is therefore at least $\frac{1}{1-\delta}\frac{1}{e}p^*(1 - F(p^*))^2$. □

Proof of Theorem 4.6.5. Let OPT denote the total revenue from running the optimal auction for two buyers on F. Our benchmark for revenue is $\mathrm{OPT}/(1-\delta)$. By concavity of the revenue we have that $OPT \leq 2p^*(1 - F(p^*))$. By Lemma 4.6.9 the equilibrium gets revenue

$$\frac{1}{1-\delta}\frac{1}{e}p^*(1 - F(p^*))^2 \quad \geq \quad \frac{\mathrm{OPT}}{1-\delta}\frac{1}{2e}(1 - F(p^*))$$

For distributions satisfying the monotone hazard rate assumption, it is a standard fact that $F(p^*) \geq 1 - 1/e$. We therefore have that the revenue of our equilibrium is at least $\frac{\text{OPT}}{2(1-\delta)e^2} \geq \frac{\text{OPT}}{2e^2}$, proving the theorem. □

4.7. Discriminatory Pricing for Digital Goods

The natural question that arises is whether our assumption that the seller is restricted to anonymous pricing is the driving force for the positive results. What if the seller is allowed to discriminate? First we extend the definition of the strategies to include discriminatory pricing.

- A seller strategy $\sigma^k_{i,S}(h^k)$ specifies for every history h^k a nonnegative price p^k_i for buyer i.
- Buyer i's strategy $\sigma^k_i(h^k, p^k_i; v_i)$ specifies for every buyer history a response to price p^k_i for every possible value of buyer i.

We then extend the definition of simple equilibria that considers a slightly different property than Markovian on path.

Definition 4.7.1. An equilibrium for multiple buyers is *simple* if it satisfies the following

- Threshold equilibrium: If for each history h^k and price p^k_i, there is a threshold $t_i(h^k, p^k_i)$ such that for each agent i, i accepts p^k_i if and only if $v_i \geq t_i(h^k, p^k_i)$.
- If given two consecutive on-path histories h^k and h^{k+1} such that $\boldsymbol{\mu}^k_i = \boldsymbol{\mu}^{k+1}_i$, it must be that $\mathbf{D}^k = \mathbf{D}^{k+1}$. In other words if in round k the seller learns nothing for buyer then the next round the response of buyer i should be the same in round $k+1$.

For a single buyer, the second property of Definition 4.7.1 is weaker than the Markovian on path property. Nevertheless, it is sufficient for proving Theorem 4.4.1. For multiple buyers, it actually may be strongerthan the natural generalization of the Markovian on path property: the strategies of both buyer and seller should be a function of the histories. In fact the Markovian on path generelizaiton is very weak if the seller is allowed to use discriminatory pricing: Consider the case of $n + 2$ buyers, where for the first two we implement an explore-exploit equilibrium with a slow pace of learning (In the explore-exploit equilibrium of Strategies 6 and 7, we can disincentive fast learning by appropriately choosing among the multiple solutions of the threshold equation). As learning happens every round with high probability we can approximately implement any arbitrary equilibrium for the remaining n buyers even if it is not Markovian on path.

Similarly to Section 4.4 we need to define a generalization of natural prices for multiple buyers. The intuition behind our initial definition of natural prices is the following. Prices are natural if they expect to occur in the single-shot variant of the game. That if the only thing we know is that F is a continuous distribution supported in $[a, b]$ what is the smallest set of prices that can be optimal? The answer here is $[a, b)$. If the seller allowed to discriminate among buyer we define natural discriminatory prices as follows.

Definition 4.7.2. A perfect Bayesian threshold equilbrium $\boldsymbol{\sigma}$ has *natural discriminatory prices on-path* (or simply *natural discriminatory prices*) if for every on-path history h^k and with beliefs supported on $[a_1, b_1]$,…,$[a_n, b_n]$), the seller's price $\sigma^k_{i,S}(h^k)$ lies in $[a_i, b_i)$ for all $i = 1, \ldots, n$.

Note that this definition does not make sense if the mechanism is constrained to use anonymous pricing. Sticking to the original intuition of the property, if the only information we have is that the distributions are continuous and the value of buyer i is supported in $[a_i, b_i)$ then any price in $p \in \bigcup i = 1^n [a_i, b_i)$ may be the revenue maximizing price. Nevertheless, this version of natural prices would be completely non-binding if discrimination is allowed. Consider the case of $n+1$ buyers where the first buyer is always offered a price equal to the minimum value of the support and always accepts the item. This ensure that the beliefs about this buyer do not change. Therefore, we can implement almost any equilibrium of the Folk theorem.[3].

We show case that simple PBE with natural discriminatory prices are sufficient to obtain a non-learning result as in Theorem 4.4.1. We emphasize again that all these properties are equilibrium refinements and not strategy space restrictions.

Theorem 4.7.1. *In the multiple-buyer game where discriminatory pricing is allowed, let the value distribution F be supported on $[a, b]$. If $\delta > 1/2$, then in any simple PBE with natural discriminatory*

[3]The only equilibrium we cannot eliminate is the totally inefficient equilibrium

prices, the seller posts $p_i^k = a$ for all buyers and every round, which is accepted by all buyers. In other words, no learning will occur.

Proof. The main idea behind our proof is that if the equilibrium is simple and has natural discriminatory prices the problem partially decouples, that is the interaction of each buyer with the seller is almost independent of the rest of the agents. Hence, we can mirror the analysis of Theorem 4.4.1 and show that learning would violate monotonicity.

Similarly to the proof of Theorem 4.4.1 we define $X_i(v), P_i(v)$, and $U_i(v)$ to denote the remaining allocation probability of The following corollary follows directly from the proof of 4.4.2.

Corollary 4.7.2. *In any PBE of the mulitple-buyer game for digital goods, the total discounted allocation, payments, and utility (respectively $X_i(v)$, $P_i(v)$, and $U_i(v)$), are non-decreasing in v.*

Similarly, we show that natural pricing for discriminatory pricing renders it impossible to satisfy monotonicity around threshold t_i. violating Corollary 4.7.2.

Lemma 4.7.3. *For any $\delta > 1/2$, consider any simple PBE of the multiple buyer game satisfying natural discriminatory prices with distribution supported on $[a, b]$ and first-round threshold $t > a$. For any agent i There exists a type $t' < t$ such that $X_i(t') > 1$.*

Proof. We argue by contradiction. We will assume that for all $t_i' < t_i$, $X_i(t_i') \leq 1$ and use natural prices for discriminatory pricing and the simplicity of the equilibrium, to show that there is at least one type less than t who would prefer to deviate from the equilibrium.

We first argue that we may assume the existence of some M such that all types in $[a_i, t_i)$ have rejected by round M. Assume this is not the case. Then let k_ϵ be the earliest round such that all agents in $[a, t-\epsilon)$ have rejected at least once. If it is the case that $k_\epsilon \to \infty$ as $\epsilon \to 0$, then because $\delta > 1/2$, it must be the case that there exists some $t_i' < t_i$ with $X(t_i') > 1$, which would prove the lemma. Hence we may assume that the number of rounds before every type in $[a_i, t_i)$ would reject at least once on the equilibrium path is finite.

Let M an index such that all types in of agent i in $[a_i, t_i)$ have rejected before round M. We now claim that there is a round $M^* \leq M$ such that a positive measure of types accept in every one of rounds $1, \ldots, M^* - 1$, but all such agents reject in round M^*. If not, then it must be that a positive measure of agents accept in every round up to and including M, a contradiction. Let the interval of such agents be $[a_i^*, t_i)$. (The upper bound being t_i is implied by the threshold property.)

Finally, we show that the existence of M^*, combined with natural discriminatory prices and simplicity of the equilibrium, implies a profitable deviation for some buyer with type in $[a_i^*, t_i)$. Consider the history where agent i has accepted all rounds $2, \ldots, M^* - 1$. The current beliefs of the seller for agent one are $[a_i^*, t_i)$. Since the price offered in round M^* is rejected by definition from all types in the interior of the support the second property of Definition 4.7.1 requires that all types of the buyer will reject in the next round. Iteratively applying this property we deduce that buyer i will be reject all future rounds. On the other hand, the requirement of natural discriminatory prices on path implies that the seller offers a price $p_i \in [a_i^*, t_i)$ in every round after $M^* - 1$. Some type in (p_i, t_i) would clearly prefer to accept at least once rather than reject forever, yielding a contradiction. □

Fix a $\delta > 1/2$, and consider a round of the game in which the beliefs are supported on $[a, b]$ and for which the buyer i has a nontrivial threshold t_i (i.e. above the bottom of the support of the current beliefs). Subgame perfection implies that we may assume this round is the first. We know from Lemma 4.4.4 that there is a value $t' < t$ such that $X(t') > 1$. We will show that we may break ties so that $X_i(t) = 1$, which contradicts Lemma 4.4.2.

By the definition of threshold equilibrium, the buyer i with type t_i accepts this round. Natural prices implies that upon seeing an acceptance, the seller will never set the price p_i for buyer i below t_i. It follows that the buyer with value t_i will not get utility from any subsequent round. We may therefore assume they reject in every round without changing their utility. Moreover, such tie-breaking doesn't change the incentives of the seller, as the type t buyer has measure 0. Hence, there is an equilibria with $X_i(t) = 1$ and $X_i(t') > 1$ for some $t' < t$, contradicting Lemma 4.4.2.

□

CHAPTER 5

Behavioral model: Procrastination

Intertemporal tradeoffs – in which an individual incurs short-term costs to achieve long-term goals – have significant implications for health, wealth, and educational outcomes. A student may complete a homework assignment due next week instead of partying with friends. A marathoner may decide to go for a morning training run instead of sleeping in. A skier with a season pass may choose to buy skies instead of renting them each weekend, if it is less costly to do so. Long-term goals like these often involve many sacrifices over time, requiring the individual to form a consistent plan of future action. Anecdotally, however, individuals often have trouble sticking to such plans. This is supported empirically by the observation that experimental subjects discount future payoffs at a non-constant rate. When asked how much money a subject would require in a month/year/decade to offset $15 today, the median response was $20/$50/$100 Thaler (1981). Since these responses are not consistent with any constant discount factor, a subject with these preferences might plan today to invest $20 next month to gain $100 in a decade, but then fail to follow through next month.

These observations have led to an extensive line of work in behavioral economics, proposing behavior models that allow divergence between plan and action. The economic theory of *hyperbolic discounting* is one such model developed to explain these time inconsistencies. A hyperbolic discount function assigns weights to future costs. At each point in time, an individual forms a plan which minimizes total costs calculated according to the hyperbolic discounting function. A special case of hyperbolic discounting, developed by Akerlof (1991), is *present-bias discounting*. In its simplest form, present-bias discounting predicts that individuals will inflate any costs incurred at the present moment by a model parameter $b > 1$, and leave all future costs unweighted.

As an example illustrating how present bias can lead to procrastination, consider a student who must complete a homework assignment.[1] The assignment, which is based on today's lecture, is due next week. The homework takes one night to complete, and there is a cost (in mental effort and missed social opportunities) to complete the assignment immediately after the lecture. Each day that goes by, the student's recollection of lecture becomes dimmer, and the cost for completing the assignment increases. In this scenario, it is clearly optimal, in terms of minimizing induced costs, to complete the homework assignment the day of the lecture. However, if the student inflates present-day costs, he may *perceive* that it is less costly to complete the assignment tomorrow than today. Facing a similar decision tomorrow, he again perceives that procrastinating is less costly than completing the assignment immediately. This procrastination behavior persists, and the student doesn't complete the assignment until the last day, spending substantially more than the optimal plan.

The preceding example contains a very stark prediction: the student will either procrastinate indefinitely, if present-biased, or not at all. As noted in prior work of Kleinberg and Oren (2014), this dichotomy persists even in the face of exponentially increasing costs, causing a present-biased student to spend exponentially more than necessary. In reality, one should expect a more nuanced behavior – any given individual might procrastinate for some amount of time, but eventually even the most lackadaisical students will exert effort to avoid future penalties. Such variation in procrastination behavior can arise from variability in the present-bias parameter. In this work, we assume that the agent has a present-bias parameter b_i drawn on each day i independently at random from a distribution $\mathcal{F}$ supported on $[0, 1)$. For the preceding example, if the student's present bias parameter has a constant probability of taking value 1, then the student will delay for only a constant number of days, in expectation, before completing their assignment. We ask: how does variability in present-bias change procrastination behavior? Are there structured choices, e.g., homework late policies, that guide individuals away from costly procrastination?

To study these questions, we adopt a graph-theoretic framework for task completion proposed by Kleinberg and Oren (2014). In this framework, an agent with a goal n days in the future traverses through a weighted task graph G starting from node s and ending at node t. On each day i, if the agent's

[1] This example is a minor reformulation of an example due to Akerloff.

current present bias is b_i and current state is v, then the weight of the first edge along a $v - t$ path is multiplied by an extra factor of b_i, while all the remaining edges are evaluated according to their true weights. The agent chooses a minimal-weight path, under this distortion, and follows the first edge of that path. For a given task graph G and present-bias distribution $\mathcal{F}$, the *procrastination ratio* is the ratio between the expected total weight of the traversed path to the weight of the initial shortest path. We are interested in bounding the procrastination ratio as a function of n, the number of days until the deadline for the task.[2]

Under this model, we can state our question more concretely: are there natural conditions that guarantee a low procrastination ratio? Recall that in the student example, if there is a constant probability of bias $b_i = 1$, the student will complete the assignment in a constant number of days in expectation. As it turns out, this is not sufficient to guarantee a sub-exponential procrastination ratio (see Section 5.2 for more details). Our first result shows that, for any distribution, the graph with maximal procrastination ratio has exactly this form: on any day, the agent may either complete the task for a growing cost or procrastinate for free. The proof of this result involves a connection to optimal pricing theory: the problem of constructing a worst-case graph can be reduced to the problem of designing a revenue-optimal auction. For example, the fact that on any day there is only one costly edge is a consequence of the fact that the optimal pricing menu for a single-parameter agent has a single deterministic option Myerson (1981).

We can leverage this description of the worst-case task graph to develop bounds on the procrastination ratio as a function of the present-bias distribution. In optimal pricing theory, a special subclass of Pareto distributions, often referred to as *equal-revenue* distributions, are useful for worst-case examples. For a given present-bias distribution, we calculate the "smallest" equal-revenue distribution that stochastically dominates it. As worst-case procrastination ratios grow with stochastic dominance, this gives us an upper bound on the worst-case procrastination ratio of any given present-bias distribution. One can similarly derive lower bounds for any distribution that is *not* dominated by a particular equal-revenue distribution. One implication of our analysis is that present-bias distributions can be roughly divided

[2] Kleinberg and Oren (2014) bounded the procrastination ratio with respect to the total number of nodes in the graph, rather than the path length n, but these are essentially equivalent; see Section 5.2.

into two categories: those with light tails tend to have low (linear or even constant) worst-case procrastination ratios, whereas heavy-tailed distributions tend to have worst-case procrastination ratios that are exponential in n.

These worst-case examples lead us to ask whether there are natural conditions on the task graph, as well as the present-bias distribution, that lead to smaller procrastination ratios. Intuitively, the key feature of the worst-case graph discussed above is that sub-optimal planning can cause an agent to not only incur extra costs today, but also reach a state from which the agent is strictly worse off; i.e., the cost of completing the task can be higher than it was initially. We say that a graph has the *bounded distance property* if the weight of the shortest path from any node v to the target t is at most the weight of the shortest path from the initial node s to t. This condition captures scenarios, like training for a marathon, in which each day of training improves preparedness whereas procrastination diminishes it, but not below the initial base level. We show for any distribution $\mathcal{F}$, the procrastination ratio of a bounded shortest-path graph is at most linear in n, and that this bound is tight even for distributions that have a constant probability of $b_i = 1$.

We next consider a stronger condition under which the procrastination ratio is constant. We say a graph has the *monotone distance property* if, for any $s - t$ path $(s = v_0, v_1, \ldots, v_{k-1}, t = v_k)$, the shortest path from v_i to t is decreasing in i. Roughly speaking, this condition captures scenarios in which progress made by an individual is not lost: regardless of an agent's action in the present round, the total cost required to complete the task has not increased. For example, consider a skier deciding whether to rent or buy skies each weekend of the skiing season: as the cost of buying skies remains constant over time, the total cost required to ski the remainder of the season does not grow. We show that if the distribution over present-bias parameters has sufficient mass at biases close to 1,[3] then the procrastination ratio of graphs with the monotone distance property is bounded by a constant. Moreover, the condition on the

[3]See Section 5.4.2 for a formal statement of the condition on the distribution. For example, the condition is satisfied whenever there is a positive probability that the bias is equal to 1.

present-bias distribution is necessary for this result: if the agent has constant bias $b > 1$, there exist monotone-shortest-path graphs for which the procrastination ratio is $\Omega(n)$.[4]

5.1. Related work

Behavioral Economics. Behavioral economics and game theory study models that can explain anomalies in human behavior which are consistently observed in experimental data, see for example Colin (2003). Our work adopts the model of Akerlof (1991) which aims to explain individual's procrastination that violates classical assumption of utility-maximizing individual's behavior. Closely related to procrastination are issues of abandonment of long-time projects O'Donoghue and Rabin (2008) and benefit of imposing a deadline in the context of task completion Ariely and Wertenbroch (2002). Other examples include models addressing attentiveness issues, e.g., where reduction of choice among the options available to an agent O'Donoghue and Rabin (1999), or delayed notification messages Ely (2015) may improve agent's performance.

Some behavioral models study various levels of agents' rationality. For example Kaur et al. (2010) study sophisticated customers who are aware of their possible future procrastination, and Wright and Leyton-Brown (2010, 2012, 2014) model agents' behavior via quantal cognitive hierarchies in single-round simultaneous-move games where agents are empirically observed not to follow Nash equilibrium strategies.

Hyperbolic Discounting. A significant amount of work in economics literature has been devoted to the study of *quasi-hyperbolic discounting*, see Frederick et al. (2002) for a survey. This form of discounting function generalizes Akerlof's model by modeling agent's behavior with two parameters: present-bias factor b and discounting factor $\delta \leq 1$. In this model the agent at every step scales up the immediate cost by b, and discounts every cost t steps away into the future by a factor of δ^t. This form of discounting has been used in many areas including addiction Beshears et al. (2015), self-control Laibson (2015), and health and pre-retirement savings Choi et al. (2004).

[4]Akerloff's original example of mailing a package is such an example: it satisfies the monotone distance property and has linear procrastination ratio. See Kleinberg and Oren Kleinberg and Oren (2014) for a graph-theoretic formulation of that example.

Graphical Model. Our work uses the model of procrastination proposed by Akerlof (1991) and adapts the graphical framework of Kleinberg and Oren (2014). The latter work studies graphical properties of the network which may cause high procrastination rates, and identifies a graph-minor condition on the task graph that implies a bound on the procrastination ratio. The follow-up work of Tang et al. (2014) provides improved bounds on the procrastination ratio, again as a function of this graph-minor condition. Finally, Kleinberg et al. (2016) considers the case where the behavioral agent is aware of the present-bias and how it effects the procrastination ratio.

Techniques. We model an individual as an agent performing graph traversal under uncertainty. Related models also appear in the literature on route planning under uncertainty Nikolova and Karger (2008), or risk-aversion in routing and congestion gamesNikolova and Moses (2011, 2014, 2015) but under different objectives.

Some of our technical results exploit a connection between agent behavior under hyperbolic discounting and auction theory. In particular, we use tools characterizing the space of optimal single-item auctions under a distribution of agent values, as explored by Myerson Myerson (1981).

5.2. Model

We consider a setting in which an agent must complete a task over a period of n days. We track the progress of the agent towards task completion using states of intermediate progress, which are nodes of a task graph $G = (\cup_{i=1}^{n+1} V_i, E)$ with edge weights $w : E \to \mathbb{R}^+$. The set of nodes V_i are the possible states of the agent on day i. We assume without loss of generality that $V_1 = \{s\}$, the *start* node, and $V_{n+1} = \{t\}$, the *end* node. All edges occur between nodes on consecutive days, i.e., $v_i v_j \in E$ if and only if $v_i \in V_i$ and $v_j \in V_{i+1}$. The weights represent the cost of transitioning between states: i.e., an agent in state v_i wishing to transition to a state v_{i+1} incurs a cost of $w(v_i v_{i+1})$. The agent starts from state $s = v_1$ on day 1, follows a path $(v_1 v_2, \ldots, v_n v_{n+1})$ in G, and ends in state $t = v_{n+1}$ on day $n + 1$, for a total cost of $\sum_{i=1}^{n} w(v_i v_{i+1})$.[5] For convenience we will refer to the shortest path between two nodes

[5]Note we do not assume the agent completes the task; task abandonment can be represented by having a costly edge from a state $v_n \in V_n$ to t.

v_i and v_j in G as the distance, denoted $d(v_i, v_j)$. Thus the minimum cost way to complete the task is $d(s,t)$.

We are interested in the inefficiency of the traversed path due to the variable present-bias of the agent. As such, for each day i we introduce a present-bias factor b_i drawn i.i.d. from a given distribution $\mathcal{F}$ supported on $[1,\infty)$. When calculating costs on day i, the agent inflates costs incurred on that day by b_i while leaving remaining costs unweighted. Thus an agent currently in state v_i on day i will choose a transition $v_i v_{i+1}$ minimizing

$$b_i \cdot w(v_i v_{i+1}) + d(v_{i+1}, t).$$

Note that on the following day $i+1$ the agent again chooses an edge using a (potentially different) present-bias factor b_{i+1}, and may therefore choose a state v_{i+2} different from his planned path on day i.

We consider the ratio of the weight of the chosen path $v_1 v_2, \ldots, v_n v_{n+1}$ with the weight of the shortest $s-t$ path:

$$\frac{\sum_{i=1}^{n} w(v_i v_{i+1})}{d(s,t)}.$$

The *procrastination ratio* is then the expectation (over the choice of the present bias factors) of this ratio.

In summary, an instance of our problem consists of a weighted task graph G and present-bias distribution $\mathcal{F}$. We would like to understand the structure of worst-case instances, i.e., the distributions and tasks graphs that achieve the maximum procrastination ratio. We would also like to describe natural classes of instances with small procrastination ratios. To this end, it is useful to introduce the following restrictions on task graphs.

Definition 5.2.1 (Bounded Distance)**.** A weighted task graph G satisfies *bounded distance*, if $\forall i$, and $\forall v_i \in V_i$, $d(v_i, t) \leq d(s,t)$.

Intuitively, the bounded distance property captures scenarios in which one's current state, in terms of the optimal cost to reach the goal, is never worse than the initial state. For example, this includes scenarios in which "starting over" is always a free and feasible option: i.e., from each vertex there is a 0-cost edge to a vertex whose distance from t is at most $d(v_i, t)$.

Definition 5.2.2 (Monotone Distance)**.** A weighted graph G satisfies *monotone distance* if $\forall i$ and $\forall v_i v_{i+1} \in E$ with $v_i \in V_i, v_{i+1} \in V_{i+1}$, $d(v_i, t) \geq d(v_{i+1}, t)$.

This captures scenarios where the agent always makes progress towards the end state, i.e., no transition causes the agent to lose any of his prior accomplishments. We note that the latter restriction implies the former one, in the sense that any graph with monotone distances also has bounded distances.

Remark. Kleinberg and Oren (2014) studied task planning in generic weighted directed acyclic graphs $G' = (V', E')$ with a designated start node s and end node t, and presented the procrastination ratio as a function of the number of nodes in G', say N. To transform their setup to ours, simply add a self-loop to t of weight 0, create N sets V_i where V_i contains copies of all nodes $v \in V'$ that can be reached from s via a path of exactly $i-1$ edges in G', and create weighted edges between nodes of V_i and V_{i+1} if such edges exist in E'. Note that under this transformation, the number of steps n is at most N, and the total number of vertices in the constructed graph is $O(N^2)$.

To illustrate the model and definitions introduced in Section 5.2, we work through instantiations of the scenarios discussed in the introduction. Throughout the course of this discussion, we will assume an agent with a present-bias factor b_i drawn from distribution $\mathcal{F}$ which takes value 1 with probability 1/3 and value 3 with probability 2/3. Note that this distribution is "close to rational" in the sense that the agent has a constant probability of $b_i = 1$. In each example, we will normalize the shortest $s-t$ path to be 1 and so the procrastination ratio is just the expected cost induced by the agent.

5.2.1. Students Preparing Homework

Our first scenario in the introduction concerned a student assigned homework, where the student's memory of the requisite material diminishes as time passes. Suppose the homework is assigned on day 1 and due on day $n+1$. We construct a task graph $G = (\cup_{i=1}^{n+1} V_i, E)$ with $V_1 = \{s_1\}$ and $V_{n+1} = \{t_{n+1}\}$. On each day $i \in 2, \ldots, n$, there are two possible states s_i and t_i corresponding to whether the student completed the assignment (s_i) or not (t_i). The edges are: $E = \{s_i s_{i+1} | i = 1 \ldots n-1\} \cup \{t_i t_{i+1} | i = 2 \ldots n\} \cup \{s_i t_{i+1} | i = 1 \ldots n\}$. The cost of transitioning from s_i to s_{i+1} or from t_i to t_{i+1} is zero. The cost of transitioning from s_i to t_{i+1} is $w_{s_i t_{i+1}} = 2^{i-1}$. This models the idea that the mental effort to complete

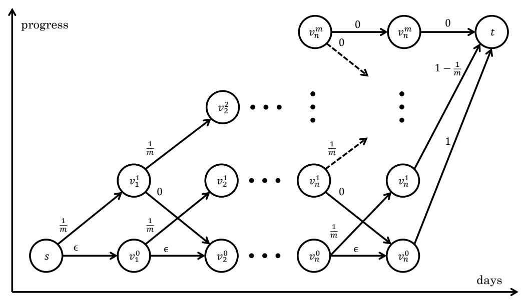

(a) Training for a marathon.

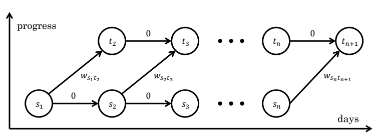

(b) Completing homework.

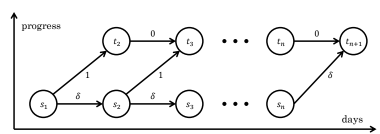

(c) Renting skis.

Figure 5.1. Task completion graphs for the examples in Section 5.2.

the assignment grows by a factor of 2 each day as the student's recollection of the material from lecture becomes dimmer. See Figure 5.1(b) for a representation of this task graph.

Note that this task graph does not have the bounded distance property. Accordingly, the procrastination ratio can be exponential, and indeed is exponential for the distribution $\mathcal{F}$. To see this, note on day i, the perceived cost of completing the assignment that day is $b_i \cdot 2^{i-1}$ while the cost of waiting to complete the homework on the next day is 2^i. Thus, if the instantiation of b_i is greater than 2, then the student will decide to complete the assignment on the following day. For the present-bias distribution $\mathcal{F}$ described above, this event occurs with probability 2/3, and so the total expected cost incurred by the student is $\sum_{i=1}^{n+1}(2/3)^i 2^{i-1}$, which is exponential in n.

5.2.2. Runners Training for a Marathon

Our second scenario in the introduction concerned a runner training for a marathon. The runner's fitness diminishes each day she sleeps in instead of going for a training run. Suppose the runner starts training on day 1 and the marathon is scheduled for day $n+1$. The task graph $G = (\cup_{i=1}^{n+1} V_i, E)$ has $V_1 = \{s\}$ and $V_{n+1} = \{t\}$. The runner has a fitness level in $\{0, \ldots, m\}$ where 0 is the base-level fitness and m is peak fitness. If the runner is at fitness level $j < m$ on day $i < n$, then the runner can pay $\frac{1}{m}$ to train and increase fitness to level $j+1$ on day $i+1$. Alternatively, if the runner is at fitness level $j > 0$ on day $i < n$, then the runner can sleep in, paying 0, and drop to fitness level $j-1$ on day $i+1$. We assume the runner can maintain fitness level m for 0 cost and level 0 at a small cost $\epsilon < \frac{2}{3m}$.[6] On day n, if the runner is at fitness level j, then she pays $1 - \frac{j}{m}$ to run the marathon. See Figure 5.1(a) for a representation of this task graph.

Note that this task graph has the bounded distance property. As we show in Section 5.4.1, such graphs have at most a linear procrastination ratio, for any distribution over present-bias parameters. For $b_i \sim \mathcal{F}$, this task graph nearly achieves this bound. To see this, note that on day $i < n$, if the runner is at fitness level $0 < j < m$, then she perceives training as costing her $b_i \cdot \frac{1}{m} + (1 - \frac{j+1}{m})$ and not training as costing her $1 - \frac{j-1}{m}$. Therefore, she trains when $b_i = 1$ and sleeps in when $b_i = 3$ (ignoring boundary conditions at fitness level 0). This is a bounded random walk on the integers from $\{0, \ldots, m\}$ with an absorbing state at m. The transition probabilities are biased toward lower fitness levels, and so the runner will spend most of her time transitioning between the low states, leading to a total cost of $\Theta(n/m)$ so long as m is sufficiently large that she never hits the absorbing state. Taking $m = \Omega(\log n)$ suffices, and this leads to a bound of $\Theta(n/\log(n))$. In Section 5.4.1, we show how to improve this example to get a linear bound for a broad class of present-bias distributions.

5.2.3. Skiers Renting or Buying Skies

Our third scenario in the introduction concerned a skier debating whether to rent or buy skies. The skier buys a season pass on day 1 and commits to skiing every weekend for the next n weeks. We construct a task graph $G = (\cup_{i=1}^{n+1} V_i, E)$ with $V_1 = \{s\}$ and $V_{n+1} = \{t\}$. On each day $i \in 2, \ldots, n$, there are two possible states, **rent** and **own**, corresponding to whether the skier uses rental equipment or his own equipment. Transitioning from a **rent** state (or s) to an **own** state on the following weekend costs 1, the

[6] At fitness level m, running ceases to become a chore. At fitness level 0, lack of exercise causes remorse and/or disutility from a low fitness level.

cost of buying skies. Transitioning from a **rent** state (or s) to another **rent** state (or t) costs δ, the cost of renting skies. Once the skier is in the **own** state, all future transitions are free. See Figure 5.1(c) for a representation of this task graph.

Note that this task graph has the monotone distance property. As we show in Section 5.4.2, if the present-bias is "close enough" to 1 sufficiently often, such graphs have a constant procrastination ratio. Indeed, for $\delta < 2/3$, we can bound the procrastination ratio by noting that the skier will choose to buy when $b_i = 1$ (except possibly towards the very end of the season) and rents otherwise. The cost is then at most $\sum_{i=1}^{n}(\frac{2}{3})^i(\delta \cdot i + 1)$, which is a constant. Note that if the present-bias was bounded away from 1, e.g., a point mass at 3, then the skier would rent for the entire season and pay a linear cost.

5.3. Worst-Case Procrastination

For a fixed present-bias distribution, we would like to characterize the worst-case procrastination ratio over all task graphs with n days. To this end, we first derive the form of a task graph that achieves the worst-case procrastination ratio for an arbitrary fixed present-bias distribution. We then develop a natural parameterization of distributions and bound this worst-case ratio as a function of the distribution parameters.

5.3.1. Worst-Case Task Graphs

Given present-bias distribution $\mathcal{F}$, we will show that there is a worst-case task graph with a simple form reminiscent of that from the student/homework example. On each day there are only two possible states: the uncompleted state and the completed state. Moving between uncompleted states or completed states has zero cost while moving from an uncompleted state to a completed one has a cost which is a function of the distribution.

The proof uses the following theorem from *optimal pricing theory*: A buyer has a value for an item, drawn from a distribution, and faces a menu of options, each consisting of a probability of allocation and a price.[7] The buyer, acting to maximize his expected utility (defined as value times allocation probability

[7] The menu must contain the option of not buying for a cost of zero, i.e., the buyer is allowed to walk away.

minus price), induces an expected allocation and price of the menu. The theorem states that for any linear function of expected allocation and price, there is a simple *single posted price* menu that optimizes that function.

Theorem 5.3.1 (Myerson (1981)). *Let $M \subseteq [0,1] \times \mathbb{R}$ be such that $(0,0) \in M$. For any such M and $v \in \mathbb{R}^+$, write $(x_M(v), p_M(v)) = \text{argmax}_{(x,p)\in M}[vx - p]$. Then for all distributions $\mathcal{F}$ with support contained in $\mathbb{R}^+$, and for all $\alpha, \beta \in \mathbb{R}$, there exists a p^* and corresponding $M^* = \{(0,0),(1,p^*)\}$ such that*

$$\alpha E_{v\sim\mathcal{F}}[x_M(v)] + \beta E_{v\sim\mathcal{F}}[p_M(v)] \leq \alpha E_{v\sim\mathcal{F}}[x_{M^*}(v)] + \beta E_{v\sim\mathcal{F}}[p_{M^*}(v)].$$

We will construct the worst-case task graph recursively. From a node v_i, the subproblem requires us to construct neighbors v_{i+1}, edge weights $w(v_i v_{i+1})$, and the weights of the shortest paths $d(v_{i+1} t_{n+1})$ from these neighbors to the end state, all subject to the constraint that the weight of the shortest path $d(v_i, t_{n+1})$ from v_i to the end state is preserved. We solve this subproblem by relating the ratio $w(v_i v_{i+1})/d(v_i, t_{n+1})$ to the allocation probability and $d(v_{i+1}, t_{n+1})/d(v_i, t_{n+1})$ to the price of a menu in pricing theory.

Theorem 5.3.2. *Given a present-bias distribution $\mathcal{F}$, there is a task graph over n days, $G^n_{\mathcal{F}} = (\cup_{i=1}^{n+1} V_i, E)$, achieving maximum procrastination ratio with the following form. Let $V_1 = \{s_1\}$, $V_{n+1} = \{t_{n+1}\}$, and $d(s_1, t_{n+1}) = 1$. Then*

- *for $i \in 2, \ldots, n$, $V_i = \{s_i, t_i\}$,*
- *for $i \in 1, \ldots, n-1$, $w(s_i s_{i+1}) = 0$,*
- *for $i \in 2, \ldots, n$, $w(t_i t_{i+1}) = 0$,*
- *and for $i \in 1, \ldots, n$, $w(s_i t_{i+1}) = d(s_i, t_{n+1}) \leq d(s_{i+1}, t_{n+1})$,*

where $d(s_i, t_{n+1})$ is chosen as a function of $\mathcal{F}$.

Proof. First note the assumption $d(s_1, t_{n+1}) = 1$ is without loss of generality as we can always normalize all weights by the initial shortest path without changing the procrastination ratio. We prove the theorem by constructing a worst-case task graph recursively. In the i'th iteration, for $i = 1 \dots n-1$, we will construct:

- the nodes of set V_{i+1},
- the weights of edges from V_i to V_{i+1},
- and the shortest continuation path weights $d(v_{i+1}, t_{n+1})$ for $v_{i+1} \in V_{i+1}$,

to maximize the expected procrastination ratio of the chosen path subject to constraints on the shortest paths $d(v_i, t_{n+1})$ for $v_i \in V_i$.

We first describe the behavior of the agent at node v_i for a fixed set of transition options V_{i+1} consisting of weighted edges $w(v_i v_{i+1})$ and continuation distances $d(v_{i+1}, t_{n+1})$ for $v_{i+1} \in V_{i+1}$. An agent with present-bias parameter b_i selects transition v^*_{i+1} minimizing $b_i w(v_i v_{i+1}) + d(v_{i+1}, t_{n+1})$. This minimization problem is equivalent to the following maximization problem, as can be seen by first changing the sign of the objective, then adding the constant $b_i \cdot d(v_i, t_{n+1})$, and then scaling by the constant $1/d(v_i, t_{n+1})$:

$$\max_{v_{i+1} \in V_{i+1}} \left(b_i \left(1 - \frac{w(v_i v_{i+1})}{d(v_i, t_{n+1})} \right) - \frac{d(v_{i+1}, t_{n+1})}{d(v_i, t_{n+1})} \right).$$

The above problem is the optimization problem faced by a buyer in a pricing menu where: the value of the buyer is b_i, and the allocation and price of menu item v_{i+1} are $x(v_{i+1}) = 1 - \frac{w(v_i v_{i+1})}{d(v_i, t_{n+1})}$ and $p(v_{i+1}) = \frac{d(v_{i+1}, t_{n+1})}{d(v_i, t_{n+1})}$, respectively. We argue this menu satisfies the conditions of Theorem 5.3.1. As $d(v_i, t_{n+1})$ is the weight of the shortest $v_i - t_{n+1}$ path, there must be an option v_{i+1} with $w(v_i v_{i+1}) \leq d(v_i, t_{n+1})$. Furthermore, any option v'_{i+1} with $w(v_i v'_{i+1}) > d(v_i, t_{n+1})$ will never be preferred to v_{i+1} and so we can eliminate such nodes from the graph. Therefore, the allocations $x(v_{i+1})$ are in $[0, 1]$. Furthermore, we can assume without loss of generality that there is a transition $t_{i+1} \in V_{i+1}$ with $w(v_i t_{i+1}) = d(v_i, t_{n+1})$ and $d(t_{i+1}, t_{n+1}) = 0$ as we are required to include a transition option with $w(v_i v_{i+1}) + d(v_{i+1}, t_{n+1}) = d(v_i, t_{n+1})$ and any such option weakly dominates t_{i+1} for any bias factor. This proves that the menu contains an option t_{i+1} with $x(t_{i+1}) = 0$ and $p(t_{i+1}) = 0$.

To apply Theorem 5.3.1, we must argue that the objective of maximizing the procrastination ratio is linear in the expected allocation and pricing of any menu V_{i+1} defined above. Let $\text{OPT}(i, d)$ denote the expected cost of the agent in the worst-case task graph with $n-i$ days and shortest path weight d. We can write $\mathrm{OPT}(i, d(v_i, t_{n+1}))$, the worst-case expected cost conditioned on being at v_i, recursively as:

$$\max_{V_{i+1}} \Big[\mathbf{E}_{b_i \sim \mathcal{F}} \left[w(v_i v^*_{i+1}) \right] + \mathbf{E}_{b_i \sim \mathcal{F}} \left[\mathrm{OPT}(i+1, d(v^*_{i+1}, t_{n+1})) \right] \Big],$$

where v^*_{i+1} is the choice of the agent with bias b_i facing options V_{i+1}, and the base of the recursion $\mathrm{OPT}(n,d) = d$. Using the linearity of $\mathrm{OPT}(i,d)$ in its second argument and the menu pricing notation introduced in the preceding paragraph, we see the optimization problem is equivalent to:

$$\max_{V_{i+1}} \Big(1 - E_{b_i \sim \mathcal{F}}[x(v^*_{i+1})] + E_{b_i \sim \mathcal{F}}[p(v^*_{i+1})] \mathrm{OPT}(i+1, 1) \Big) \cdot d(v_i, t_{n+1}). \tag{5.1}$$

Applying Theorem 5.3.1, we conclude that there is an optimal menu $V_{i+1} = \{s_{i+1}, t_{i+1}\}$ where $x(s_{i+1}) = 1$, $p(s_{i+1}) = p^* \geq 1$ (as the value $b_i \geq 1$), $x(t_{i+1}) = 0$, and $p(t_{i+1}) = 0$. Given a partially constructed task graph with $V_i = \{s_i, t_i\}$, we can therefore optimally extend it by defining $V_{i+1} = \{s_{i+1}, t_{i+1}\}$. The above discussion shows that the transitions from s_i to V_{i+1} should have edge weights $w(s_i s_{i+1}) = 0$ and $w(s_i t_{i+1}) = d(s_i, t_{n+1})$ and continuation shortest path weights $d(s_{i+1}, t_{n+1}) \geq d(s_i, t_{n+1})$ and $d(t_{i+1}, t_{n+1}) = 0$. To satisfy the constraint that $d(t_i, t_{n+1}) = 0$ we also add an edge from t_i to t_{i+1} of weight zero. This completes the i'th iteration. To complete the construction, add edges $s_n t_{n+1}$ and $t_n t_{n+1}$ with weights $d(s_n, t_{n+1})$ and 0, respectively. □

5.3.2. Bounding the Procrastination Ratio for a Bias Distribution

In Section 5.3.1 we characterized the task graph that maximizes the procrastination ratio for any given present-bias distribution. In this section we leverage this characterization to explore how properties of the present-bias distribution impacts the worst-case procrastination ratio. In particular, we are interested in determining features of the bias distribution that imply sub-exponential, or even constant, bounds on the procrastination ratio.

We first note that worst-case procrastination ratio can only be greater for distributions that generate higher present-bias parameters. Recall that distribution $\bar{\mathcal{F}}$ *stochastically dominates* distribution $\mathcal{F}$ if, for all x, $\Pr_{b\sim\bar{\mathcal{F}}}[b > x] \geq \Pr_{b\sim\mathcal{F}}[b > x]$. That is, $\bar{\mathcal{F}}(x) \leq \mathcal{F}(x)$ for all x.

Lemma 5.3.3. *Given two distributions $\mathcal{F}$ and $\bar{\mathcal{F}}$ such that $\bar{\mathcal{F}}$ stochastically dominates $\mathcal{F}$, the worst-case procrastination ratio for $\bar{\mathcal{F}}$ is at least as large as the worst-case procrastination ratio for $\mathcal{F}$.*

Proof. Fix n and let $G^n_{\mathcal{F}}$ be a worst-case task graph for $\mathcal{F}$ of the form specified in Theorem 5.3.2. We use cost_i to denote the expected cost paid by an agent with $b_i \sim \mathcal{F}$ on days $i, \ldots, n$ in the graph $G^n_{\mathcal{F}}$, conditional on being at node s_i on day i. Similarly, we use $\overline{\text{cost}}_i$ to denote the expected cost paid by an agent with $b_i \sim \bar{\mathcal{F}}$ on days $i, \ldots, n$ in the graph $G^n_{\mathcal{F}}$, conditional on being at node s_i on day i. In the following we show using backward induction on i that $\overline{\text{cost}}_i \geq \text{cost}_i$, implying that $\overline{\text{cost}}_0$, a lower bound on the worst-case procrastination ratio for $\bar{\mathcal{F}}$, is at least as large as cost_0, the worst-case procrastination ratio for $\mathcal{F}$.

When $i = n$ we have $\overline{\text{cost}}_n = \text{cost}_n = d(s_n, t_{n+1})$. For $i < n$ note that by definition of $G^n_{\mathcal{F}}$, an agent with threshold bias $b_0 = d(s_{i+1}, t_{n+1})/d(s_i, t_{n+1})$ is indifferent between transitioning to s_{i+1} at a cost of zero, and transitioning to t_{i+1} at a cost of $d(s_i, t_{n+1})$ with zero future costs. Thus

$$(5.2) \qquad \begin{cases} \text{cost}_i &= (1 - \mathcal{F}(b_0)) \cdot \text{cost}_{i+1} + \mathcal{F}(b_0) \cdot d(s_i, t_{n+1}) \\ \overline{\text{cost}}_i &= (1 - \bar{\mathcal{F}}(b_0)) \cdot \overline{\text{cost}}_{i+1} + \bar{\mathcal{F}}(b_0) \cdot d(s_i, t_{n+1}). \end{cases}$$

By induction, $\overline{\text{cost}}_{i+1} \geq \text{cost}_{i+1}$, and so

$$\overline{\text{cost}}_i - \text{cost}_i \geq \left(\mathcal{F}(b_0) - \bar{\mathcal{F}}(b_0)\right) \cdot \left(\text{cost}_{i+1} - d(s_i, t_{n+1})\right).$$

By stochastic dominance, $\mathcal{F}(b_0) \geq \bar{\mathcal{F}}(b_0)$, and by the construction of $G^n_{\mathcal{F}}$, $\text{cost}_{i+1} \geq d(s_{i+1}, t_{n+1}) \geq d(s_i, t_{n+1})$. As a result $\overline{\text{cost}}_i - \text{cost}_i \geq 0$, which completes the induction. □

Lemma 5.3.3 motivates us to categorize distributions into classes, where all distributions in a certain class stochastically dominate a family of distributions. Such a categorization, along with Lemma 5.3.3, will allow us to derive bounds on the procrastination ratio of any given distribution. Our bound will be with respect to the following parametrization of present-bias distributions, which is essentially the largest value z such that the distribution stochastically dominates the equal-revenue distribution with revenue z.[8]

Definition 5.3.1. Let $\mathcal{F}$ be a distribution with support contained in $[1, \infty)$. We write $z(\mathcal{F}) \equiv \max(1 - \mathcal{F}(b))b$. Given $z \in (0, \infty)$, we will write $\mathfrak{F}(z)$ for the family of distributions $\mathcal{F}$ with $z(\mathcal{F}) = z$.

Note that any distribution $\mathcal{F}$ with bounded support belongs to some family $\mathfrak{F}(z)$, potentially with $z = \infty$. For $\mathcal{F} \in \mathfrak{F}(\infty)$, i.e., $\mathcal{F}$ has infinite support and $\mathcal{F} \notin \mathfrak{F}(z)$ for any $z > 0$, then there are task graphs for $\mathcal{F}$ with procrastination ratio growing faster than any exponential function in n. We will largely focus on distributions for which $z(\mathcal{F})$ is finite.

To build some intuition for the classes $\mathfrak{F}(z)$, we give below a few examples.

(1) The uniform distribution over the interval $[1, 3]$ belongs to the family $\mathfrak{F}(1.125)$, with $(1-\mathcal{F}(1.5)) \cdot 1.5 = 1.125$.
(2) The uniform distribution over the interval $[1, 2]$ belongs to the family $\mathfrak{F}(1)$.
(3) Suppose $\mathcal{F}$ is the upper half of a normal distribution with mean 1. That is, $b = \max(\xi, 1)$ where $\xi \sim N(1, 1)$ is a normal random variable with mean and normal deviation equal to 1. In this case, we can maximize $b(1 - \mathcal{F}(b))$ numerically to find that $\mathcal{F}(b) \in \mathcal{F}(z)$ for $z \approx 0.507$.
(4) Suppose $\mathcal{F}(x) = 1 - \frac{1}{2\sqrt{x}}$ for $x \in [1, 100)$, and $\mathcal{F}(100) = 1$. This is a heavy-tailed distribution, which has an atom at $b = 1$ with probability $\frac{1}{2}$. This distribution belongs to $\mathfrak{F}(5)$.

We are now ready to bound the worst-case procrastination ratio for any given distribution $\mathcal{F}$, as a function of $z(\mathcal{F})$. The following two theorems tell us whether a given distribution has an exponential, linear, or constant procrastination ratio, in the worst case over task graphs. Theorem 5.3.4 provides an upper bound on the procrastination ratio, which is tight for the family $\mathfrak{F}(z)$. Theorem 5.3.5 provides a lower bound for any given distribution.

[8]The equal-revenue distribution with revenue z satisfies $F(x) = 1 - \frac{z}{x}$ for all $x \geq z$.

Theorem 5.3.4. *Let $\mathcal{F} \in \mathfrak{F}(z)$ for some $z > 0$. Then the procrastination ratio of any task graph for $\mathcal{F}$ is at most $\sum_{i=0}^{k-1} z^i$ and this bound is tight.*

Proof. To prove an upper bound on the procrastination ratio we analyze the following distribution $\mathcal{F}$ with bounded support on $[1, C]$, where C is some sufficiently large constant:

$$\mathcal{F}: \begin{cases} \Pr[b = 1] & = 1 - z, \qquad \text{if } z < 1 \\ \Pr[b \le x] & = 1 - z/x, \quad x \in [\max(1, z), C] \\ \Pr[b = C] & = z/C \end{cases}$$

We note that $\mathcal{F}$ stochastically dominates every distribution in $\mathfrak{F}(z)$ with support bounded by C.

Recall from Theorem 5.3.2 the form of the worst-case task graph for distribution $\mathcal{F}$. We briefly describe the form of the graph here. Given a current state v_i, the agent is offered two options: an edge with weight $d(v_i, t)$ (and no further costs thereafter), and an edge with weight 0 that transitions to a vertex v_{i+1} with increased distance to t. We claim that, in fact, $d(v_{i+1}, t) = C \cdot d(v_i, t)$ in the worst-case graph. To see why, note that according to (5.1) the worst-case task graph maximizes a linear function of $\mathbf{E}_{b_i \sim \mathcal{F}}[x(v^*_{i+1})]$ and $\mathbf{E}_{b_i \sim \mathcal{F}}[p(v^*_{i+1})]$. The former term does not depend on the set of offered price p, because $\mathbf{E}_{b_i \sim \mathcal{F}}[p(v^*_{i+1})]$ corresponds in (5.1) to the expected revenue of a single item auction for the equal revenue distribution $\mathcal{F}$. Therefore, $\mathbf{E}_{b_i \sim \mathcal{F}}[p(v^*_{i+1})] = z$ for any menu of options, and hence the optimization problem reduces to maximizing the term $1 - E_{b_i \sim \mathcal{F}}[x(v^*_{i+1})]$. It is maximized when the agent's probability of paying a positive cost at step i is maximized. This is maximized when agent can either increase the distance to t to $d(v_{i+1}, t) = C \cdot d(v_i, t)$, or pay the total amount $d(v_i, t)$ right away. Thus we obtain the worst-case task graph.

Now we calculate the procrastination ratio of this task graph. At step i, from a given vertex $v \in \{v_i, u_i\}$, the expected distance to t is $\mathbf{E}[d(v, t)] = z^{i-1}$. The probability that agent takes a non-zero edge is $1 - 1/C$. Thus at the steps $i = 1, \ldots, n-1$ the agents pays a cost, in expectation, of $\mathcal{F}(C) \cdot \mathbf{E}[d(v, t)] = (1 - 1/C) \cdot z^{i-1}$ and $\mathbf{E}[d(v, t)] = z^{n-1}$ at the last step. As a result, the total expected cost paid by the agent is

$$\sum_{i=1}^{n-1}(1-1/C)\cdot z^{i-1}+z^{n-1}$$

As C tends to infinity this quantity converges to $\sum_{i=1}^{n} z^{i-1} = \frac{1-z^n}{1-z}$. This is an upper bound on the procrastination ratio for $\mathcal{F} \in \mathfrak{F}(z)$ by Lemma 5.3.3. Furthermore, this bound is tight for the class $\mathfrak{F}(z)$, since we have exhibited a distribution and a task graph that achieve this bound in the limit as $C \to \infty$. □ □

Theorem 5.3.5. *For arbitrary $\mathcal{F}$, if $\left(1-\mathcal{F}(b_0)\right) b_0 = z$ for some point $b_0 > 1$, then there is a task graph with procrastination ratio $\sum_{i=1}^{n-1}(1-1/b_0)\cdot z^{i-1}+z^{n-1}$.*

Proof. Indeed, we may assume that $\mathcal{F}$ has finite support $\{1, b_0\}$, since every other distribution stochastically dominates this one and by revenue monotonicity (Lemma 5.3.3) the adversary may only get higher procrastination ratio. The task graph from the proof of Theorem 5.3.4 gives us the required bound. □

With these results in hand, we can bound the worst-case competitive ratios for the examples described above:

(1) The uniform distribution over the interval $[1, 3]$ is in $\mathfrak{F}(1.125)$, so Theorem 5.3.4 implies that the worst-case procrastination ratio of any task graph for this distribution is not more than $O(1.125^n)$. Moreover, Theorem 5.3.5 applied at point $b_0 = 1.5$ implies that there are task graphs with procrastination ratio of $\Omega(1.125^n)$.

(2) The uniform distribution over the interval $[1, 2]$ is in $\mathfrak{F}(1)$, and therefore has at most a linear procrastination ratio for any task graph (Theorem 5.3.4). On the other hand, one can obtain sublinear bound of $\Omega(\sqrt{n})$ in Theorem 5.3.5 by taking $b_0 = 1+\varepsilon$ and $z = 1-\varepsilon^2$ for any $\varepsilon > 0$.

(3) For the upper half of a normal distribution with mean 1, since the distribution lies in $\mathfrak{F}(z)$ with $z \approx 0.507$, Theorem 5.3.4 implies that the procrastination ratio is at most 2.03.

(4) For the distribution given by $\mathcal{F}(x) = 1 - \frac{1}{2\sqrt{x}}$ for $x \in [1, 100)$, and $\mathcal{F}(100) = 1$, since the distribution belongs to $\mathfrak{F}(5)$ we have that the procrastination ratio is bounded by $O(5^n)$, and invoking Theorem 5.3.5 with $b_0 = 100$ shows that it can be as large as $\Omega(5^n)$.

5.4. Special Task Graphs

In this section we consider restrictions on the structure of a task graph. We first study the impact of bounding the length of the shortest path to the target node, from any vertex in the network. We then consider a stronger monotonicity property, which roughly states that following an edge in the graph can never increase the length of the shortest path to the target. We will show that the former property implies a linear procrastination ratio, and this is tight, whereas the latter property implies a constant procrastination ratio for appropriate present-bias distributions.

5.4.1. Bounded Distance

Here we consider the class of task graphs such that distance from any given state to t is bounded by the initial distance $d(s, t)$. As it turns out, any such graph has at most a linear procrastination ratio, for any distribution $\mathcal{F}$.

Lemma 5.4.1. *The procrastination ratio of any protocol with bounded distance does not exceed n.*

Proof. Suppose the chosen path is $v_1, v_2, \ldots, v_n, t$. By assumption, $d(v_i, t) \leq d(s, t)$ for all i. Since the agent chooses the edge $v_i v_{i+1}$ to minimize $b_i \cdot w(v_i v_{i+1}) + d(i+1, t)$ where $b_i \geq 1$, and since the first edge of the path realizing distance $d(v_i, t)$ is itself an option, it must be that $w(v_i v_{i+1}) \leq d(v_i, t) \leq d(s, t)$ for each i. The total cost of the chosen path is therefore at most $n \cdot d(s, t)$. □ □

The bound in Lemma 5.4.1 holds for arbitrary procrastination ratios. One might hope that if the present-bias distribution $\mathcal{F}$ is sufficiently well-behaved, the procrastination ratio would improve. However, as we now show, this bound on the procrastination ratio is asymptotically tight for any $\mathcal{F} \in \mathfrak{F}(z)$, for any $z > 1$. In other words, one cannot hope to avoid a linear lower bound unless $\mathcal{F} \in \mathfrak{F}(z)$ for some $z \leq 1$, which we view as a particularly strict condition.

The rough idea for the construction is to simulate a random walk. Whenever the present-bias parameter is low the agent will incur a cost and reduce the length of the shortest path to the goal, but whenever the present-bias parameter is high the length of the shortest path will increase. Over the course of n steps, the agent will have repeatedly made and lost progress sufficiently often to have accrued a total cost that is linear in n, with high probability.

Theorem 5.4.2. *If $\mathcal{F} \in \mathfrak{F}(z)$ for some $z > 1$, then there is a task graph with bounded shortest path that has procrastination ratio $\Omega(n)$.*

Proof. Let $\delta \in (0, 1)$ and integers $\alpha, \beta \geq 1$ be constants that depend on $\mathcal{F}$, to be determined later. We construct a task graph as follows. Each layer V_i contains $n\beta$ vertices; say $V_i = \{v_{i,1}, \ldots, v_{i,n\beta}\}$. Each vertex $v_{i,j}$ has out-degree at most 2. For $i < n$, starting from vertex $v_{i,j}$,

- if $\alpha \leq j \leq (n-1)\beta$, then the agent may go either to $v_{i+1,j-\alpha}$ at no cost, or to $v_{i+1,j+\beta}$ at a cost of $\delta^j - \delta^{j+\beta}$;
- if $j < \alpha$, then the agent must transition to $v_{i+1,\alpha}$ at a cost of $\delta^j - \delta^\alpha$.
- Otherwise, if $j > (n-1)\beta$, the agent must transition to $v_{i+1,j}$ at no cost.

For $i = n$, each vertex $v_{i,j}$ has only a single outgoing edge $(v_{i,j}, t)$, which has cost δ^j.

We claim that, for each vertex $v_{i,j}$, we have $d(v_{i,j}, t) = \delta^j$. This is achieved by the path in which the agent always chooses the option with positive cost, whenever multiple options are available. That is, the agent repeatedly chooses to take the route that increases the second coordinate of its vertex, if possible. If this path is $v_{i,j}, v_{i_2,j_2}, \ldots, v_{i_\ell,j_\ell}, t$, then the total cost is

$$(\delta^j - \delta^{j_2}) + (\delta^{j_2} - \delta^{j_3}) + \ldots + (\delta^{j_{\ell-1}} - \delta^{j_\ell}) + \delta^{j_\ell} = \delta^j,$$

via a telescoping sum. The agent starts at $v_{1,0}$, and $d(v_{1,0}, t) = 1$. We note that the task graph satisfies the bounded distance condition.

Let b^* be such that $(1 - \mathcal{F}(b^*))b^* = z > 1$; we know such a b^* exists by assumption. We will then choose α and β so that

$$\frac{1}{b^*} < \frac{\beta}{\alpha + \beta} < 1 - \mathcal{F}(b^*).$$

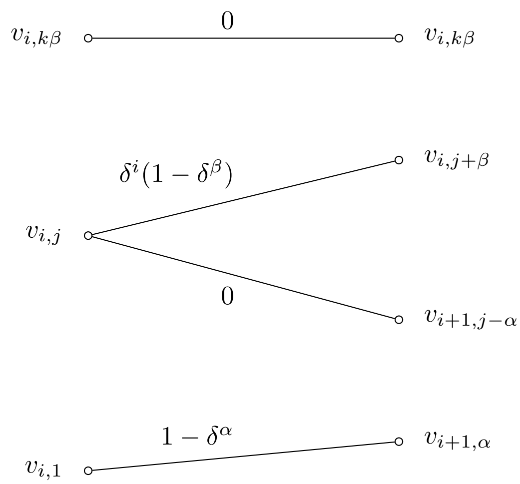

Figure 5.2. Task completion graph for the protocol

Note that the agent will select the 0-cost option, if multiple options are available, when

$$\delta^{j-\alpha} < b_i(\delta^j - \delta^{j+\beta}) + \delta^{j+\beta}.$$

We claim that this will occur for $b_i \geq b^*$, if δ is sufficiently close to 1. To see this, define $H(\delta) \stackrel{\text{def}}{=} b^* \cdot \delta^\alpha - \delta^{\alpha+\beta} \cdot (b^* - 1)$. Note that $H(1) = 1$. Furthermore, $H'(1) = b^* \cdot \alpha - (\alpha+\beta) \cdot (b^* - 1) = \alpha + \beta - \beta \cdot b^*$ is negative if $\frac{1}{b^*} < \frac{\beta}{\alpha+\beta}$. The claim follows, as $H(\delta) = H(1) - H'(1)(1-\delta) + o\,(1-\delta)$ for δ sufficiently close to 1. Choose δ to be sufficiently close to 1 for the claim to hold.

Given $i < n$ and j such that $\alpha \leq j \leq (n-1)\beta$, let γ be the expectation, over the present-bias distribution, of the change in index of the current vertex when an agent makes a single move starting at vertex $v_{i,j}$. Then γ is at most $\beta\mathcal{F}(b^*) - \alpha(1 - \mathcal{F}(b^*))$, which is negative by our choice of α and β.

Note that since the agent started at $v_{1,0}$ it is impossible to reach a state $v_{i,j}$ where $j \geq (n-1)\beta$. We now count the sum of all changes in vertex index, over the course of the agent's selected path, enumerated separately for two cases: (i) when $\alpha \leq j \leq (n-1)\beta$, and when (ii) $j < \alpha$.

Let random variables s_1 and s_2 denote the sum of changes in vertex index for these two cases. We observe that the total sum must be non-negative, and hence

$$E[s_1] + E[s_2] \geq 0. \tag{5.3}$$

Also, by our construction, if $j < \alpha$ then the subsequent step will have $j \geq \alpha$, and hence the second case applies in at least half of the rounds. Therefore,

$$E[s_1] \leq \gamma \cdot n/2 \tag{5.4}$$

where recall that $\gamma < 0$. Equations (5.3) and (5.4) imply that

$$E[s_2] \geq -\gamma \cdot n/2. \tag{5.5}$$

Note that for each index-increasing step, the agent is reducing the shortest path and incurring an immediate cost equal to that reduction. The minimum cost reduction between two consecutive states is $\delta^{\alpha-1} - \delta^{\alpha}$, in the case when $j < \alpha$. We can therefore conclude that

$$\sum_{i=1}^{n} E[w_i] \geq -\gamma \cdot n/2 \cdot (\delta^{a-1} - \delta^{\alpha}) = \Omega(n) \tag{5.6}$$

as required. □

5.4.2. Monotone distance

We next consider the class of task graphs with the monotone distance property. Recall that for such task graphs, the distance from any given state to t is monotonically non-increasing along every edge in the graph.

We first note that if $\mathcal{F}$ is the present-bias distribution, and there exists some $\delta > 0$ such that $\mathcal{F}(1+\delta) = 0$ (that is, $1+\delta$ is a lower bound on the present bias parameter), then there exist task graphs with the monotone distance property with procrastination ratio $\Omega(\delta n)$. Indeed, consider the ski-renting example from Section 5.2: if $\mathcal{F}$ places all of its mass on values greater than $\frac{1}{1-\delta}$, the agent would choose to pay δ at each step. This leads to a total cost of $n\delta$, whereas a cost of 1 was possible.

This example motivates us to consider distributions $\mathcal{F}$ that satisfy a mild condition, which essentially captures the requirement that there be sufficient mass in neighborhoods around $b = 1$. Specifically, we will require that there exist some constants $\beta, \delta > 0$ such that $\mathcal{F}(x) \geq \beta \cdot (x - 1)$ for all $x \in [1, 1 + \delta]$.[9] We show that under this condition, the procrastination ratio is bounded by a constant.

Theorem 5.4.3. *For any $\beta, \delta > 0$, if $\mathcal{F}(x) \geq \beta \cdot (x - 1)$ for all $x \in [1, 1 + \delta]$, then any task graph with monotone distances has procrastination ratio at most* $\max(1 + \frac{1}{\beta}, \frac{1}{\beta\delta})$.

Proof. As in Theorem 5.3.2, we will characterize the task graphs with maximum procrastination ratio by way of an analogy with single-parameter auctions. Following the notation in the proof of Theorem 5.3.2, we have that the problem of maximizing the procrastination ratio can be expressed recursively (up to a constant $d(v_i, t_{n+1})$) as

$$\max_{V_{i+1}} \Big(1 - E_{b_i \sim \mathcal{F}}[x(v^*_{i+1})] + E_{b_i \sim \mathcal{F}}[p(v^*_{i+1})]\mathrm{OPT}(i+1, 1)\Big).$$

The monotone distances property imposes the constraint that $d(v_{i+1}, t) \leq d(v_i, t)$ for all $v_{i+1} \in V_{i+1}$. Since $p(v_{i+1}) = \frac{d(v_{i+1},t)}{d(v_i,t)}$, this translates to a constraint that $p(v^*_{i+1}) \in [0, 1]$. Unlike the proof of Theorem 5.3.2, we cannot employ Theorem 5.3.1, since the price suggested by Theorem 5.3.1 may be greater than 1. Instead, we will use the following variation.

Lemma 5.4.4. *Let $M \subseteq [0, 1] \times [0, 1]$ be a menu of options (x, p) such that $(0, 0) \in M$. For any such M and $v \in \mathbb{R}_{\geq 1}$, write $(x_M(v), p_M(v)) = \mathrm{argmax}_{(x,p) \in M}[vx - p]$. Then for all distributions $\mathcal{F}$ with support on $[1, \infty)$, and for all $\alpha \in \mathbb{R}, \beta \in \mathbb{R}_{\geq 0}$, there exists an $x^* \in [0, 1]$ and corresponding $M^* = \{(0, 0), (x^*, 1)\}$ such that*

$$\alpha E_{v \sim \mathcal{F}}[x_M(v)] + \beta E_{v \sim \mathcal{F}}[p_M(v)] \leq \alpha E_{v \sim \mathcal{F}}[x_{M^*}(v)] + \beta E_{v \sim \mathcal{F}}[p_{M^*}(v)].$$

[9]Sufficient conditions for our requirement including having a positive mass at $b = 1$, or having a constant lower bound on the density function in any neighborhood around $b = 1$.

Proof. For simplicity we will prove the lemma under the assumption that menu M is finite. We note that the result directly extends to general M. Also, in the context in which we apply the lemma (namely, when menu M corresponds to a task graph), M is finite.

Write $M = \{(x_1, p_1), \ldots, (x_k, p_k)\}$. We will say that the item (x_i, p_i) is *chosen* by value v if $(x_M(v), p_M(v)) = (x_i, p_i)$. We will partition the set of possible values v (i.e., $\mathbb{R}_{\geq 1}$) according to their chosen lottery. Note that, for each i, the set of values that choose (x_i, p_i) forms an interval, say $[v_{i-1}, v_i)$. Without loss of generality, we can assume every element of M is chosen by some value. In this case, we can choose to index lotteries so that $1 = v_0 < v_1 < \cdots < v_k = \infty$. In this case, we must have $0 = x_0 \leq x_1 < x_2 < \cdots < x_k \leq 1$ and $p_1 < p_2 < \cdots < p_k \leq 1$. We note that the latter can be derived from well-known payment identities of incentive compatible auctions Myerson (1981). Moreover, the fact that v_{i-1} is the threshold type between choosing (x_{i-1}, p_{i-1}) and (x_i, p_i), and is therefore indifferent between these auctions, implies that

$$v_{j-1} \cdot (x_j - x_{j-1}) = p_j - p_{j-1} \quad \text{if } j \in \{2, \ldots, k\}$$

$$v_0 \cdot x_1 - p_1 > 0 \quad \text{if } j = 1. \tag{5.7}$$

We now consider a collection of two-element menus $M_1, \ldots, M_k$, where $M_j = \{(0, 0), (\frac{1}{v_{j-1}}, 1)\}$. We will construct a distribution over these menus, and consider offering one of them at random: say menu M_j is chosen with probability m_j. We want to choose $\{m_1, \ldots, m_k\}$ such that for every value $v \in [v_{j-1}, v_j]$, the expected allocation chosen by v under this distribution of menus is exactly x_j (the allocation under M), and the expected payment is at least p_j. Note then that this random menu, call it M^r, can only increase our objective of interest:

$$\alpha E_{v \sim \mathcal{F}}[x_M(v)] + \beta E_{v \sim \mathcal{F}}[p_M(v)] \leq \alpha E_{v \sim \mathcal{F}}[x_{M^r}(v)] + \beta E_{v \sim \mathcal{F}}[p_{M^r}(v)].$$

Since M^r is a convex combination of menus M_j, there must exist at least one M_j for which the objective is at least as high as for M^r. This M_j satisfies the requirements of the lemma.

It remains to find appropriate probabilities $m_1, \dots, m_k$. We will choose $m_j = v_{j-1} \cdot (x_j - x_{j-1})$, where $j \in \{1, \dots, k\}$ and $x_0 = 0$.

We observe that each type $v \in (v_{i-1}, v_i)$ chooses the option $(\frac{1}{v_{j-1}}, 1)$ in M_j for every $j \in \{1, \dots, i\}$ and option $(0, 0)$ for the remaining j. Therefore, the allocation probability and expected payment of type v is

$$x(v) = \sum_{j=1}^{i} \frac{1}{v_{j-1}} \cdot m_j = \sum_{j=1}^{i} x_j - x_{j-1} = x_i$$
$$p(v) = \sum_{j=1}^{i} m_j = \sum_{j=1}^{i} v_{j-1} \cdot (x_j - x_{j-1}) = v_0 \cdot x_1 + \sum_{j=2}^{i} (p_j - p_{j-1}) = p_i - p_1 + x_1 \geq p_i,$$

where the last inequality and previous two equalities follow from (5.7) and telescopic summation. Since $p(v_k) \leq 1$ we also get that $\sum_{j=1}^{k} m_j \leq 1$. If $\sum_{j=1}^{k} m_j$ is strictly less than 1, we can extend to a probability distribution by adding a dummy lottery $M_0 = \{(0, 0)\}$ that is offered with probability $1 - \sum_{j=1}^{k} m_j$. $\square$

Lemma 5.4.4 implies that there is an optimal menu $V_{i+1} = \{v_{i+1}, u_{i+1}\}$ where $x(v_{i+1}) \leq 1$, $p(v_{i+1}) \leq 1$, $x(u_{i+1}) = 0$, and $p(u_{i+1}) = 0$. Given a partially constructed task graph with $V_i = \{v_i, u_i\}$, we can therefore optimally extend it by defining $V_{i+1} = \{v_{i+1}, u_{i+1}\}$. Applying the same reasoning as in the proof of Theorem 5.3.2, we can conclude that the transitions from v_i to V_{i+1}, and the continuation shortest path weights, will be as follows:

(1) $w(v_i u_{i+1}) = d(v_i, t)$ and $d(u_{i+1}, t) = 0$, and

(2) $w(v_i v_{i+1}) \leq d(v_i, t)$ and $d(v_{i+1}, t) \leq d(v_i, t)$.

Note that the task graph is now fully specified, except for the values of $w(v_i v_{i+1})$ and $d(v_{i+1}, t)$ for each $i \leq n$. (Recall that we must have $v_{n+1} = u_{n+1} = t$.) For any such graph and distribution $\mathcal{F}$, we can write q_i for the probability that the agent would choose u_{i+1} instead of v_{i+1}, from v_i. That is,

$$q_i = \mathbf{Pr}_{b_i \sim \mathcal{F}} \left[b_i \cdot \frac{w(v_i v_{i+1})}{d(v_i, t)} + \frac{d(v_{i+1}, t)}{d(v_i, t)} \geq b_i \cdot 1 \right].$$

We can write the following expression for the optimum at each state v_i scaled by $d(v_i, t)$, again following the notation from the proof of Theorem 5.3.2:

$$\text{OPT}(i,1) = q_i \cdot 1 + (1-q_i)\left(\frac{w(v_i v_{i+1})}{d(v_i,t)} + \frac{d(v_{i+1},t)}{d(v_i,t)} \cdot \text{OPT}(i+1,1)\right). \tag{5.8}$$

To complete the proof of Theorem 5.4.3, it suffices to show that $\text{OPT}(i,1) \leq \max(1+\frac{1}{\beta}, \frac{1}{\beta\delta})$ for each $1 \leq i \leq n$. We will show this by backward induction on i. The base case $i = n$ follows immediately from the fact that $OPT(n,1) = 1 \leq \max(1+\frac{1}{\beta}, \frac{1}{\beta\delta})$, since the agent has only one option: moving to vertex t.

Consider $i < n$. Let $\Delta \stackrel{\text{def}}{=} \max(1+\frac{1}{\beta}, \frac{1}{\beta\delta})$. By (5.8), it is sufficient to show that $\Delta \geq q_i + (1-q_i)(\frac{w(v_i v_{i+1})}{d(v_i,t)} + \frac{d(v_{i+1},t)}{d(v_i,t)} \cdot \Delta)$, or equivalently that

$$\Delta q_i + \Delta(1-q_i)\left(1 - \frac{d(v_{i+1},t)}{d(v_i,t)}\right) \geq q_i + (1-q_i)\frac{w(v_i v_{i+1})}{d(v_i,t)}. \tag{5.9}$$

We consider two cases based on the value of q_i.

Case 1: $q_i \geq \beta\delta$. Then $\Delta \cdot q_i \geq \frac{q_i}{\beta\delta} \geq 1$, and hence the left hand side of (5.9) is at least 1. On the other hand, since $\frac{w(v_i v_{i+1})}{d(v_i,t)} \leq 1$, the right hand side of (5.9) is at most 1.

Case 2: $q_i < \beta\delta$. We observe that $q_i = \mathcal{F}\left(\frac{d(v_{i+1},t)}{d(v_i,t)-w(v_i v_{i+1})}\right)$ and by the condition of the theorem that $F(x) \geq \beta(x-1)$ for $x \in [1, 1+\delta]$ we therefore have $q_i \geq \beta \cdot \left(\frac{d(v_{i+1},t)}{d(v_i,t)-w(v_i v_{i+1})} - 1\right)$. We therefore have

$$\begin{aligned}
&\Delta q_i + \Delta(1-q_i)\left(1 - \frac{d(v_{i+1},t)}{d(v_i,t)}\right) \\
&\quad \geq q_i + \frac{1}{\beta} q_i + (1-q_i)\left(1 - \frac{d(v_{i+1},t)}{d(v_i,t)}\right) \\
&\quad \geq q_i + (1-q_i)\left(\frac{1}{\beta} q_i + 1 - \frac{d(v_{i+1},t)}{d(v_i,t)}\right) \\
&\quad \geq q_i + (1-q_i)\left(\frac{d(v_{i+1},t)}{d(v_i,t) - w(v_i v_{i+1})} - 1 + 1 - \frac{d(v_{i+1},t)}{d(v_i,t)}\right) \\
&\quad = q_i + (1-q_i)\frac{d(v_{i+1},t)}{d(v_i,t) - w(v_i v_{i+1})}\frac{w(v_i v_{i+1})}{d(v_i,t)} \\
&\quad \geq q_i + (1-q_i)\frac{w(v_i v_{i+1})}{d(v_i,t)},
\end{aligned}$$

where the first inequality follows from $\Delta \geq 1 + \frac{1}{\beta}$, the second inequality we simply reduced the second term by $(1 - q_i)$; in the third inequality we applied our lower bound on q_i. The final inequality follows because $d(v_{i+1}, t) + w(v_i v_{i+1}) \geq d(v_i, t)$, which implies that $\frac{d(v_{i+1},t)}{d(v_i,t) - w(v_i v_{i+1})} \geq 1$. We have therefore established (5.9), as desired. □

5.5. Conclusions and Open problems

We have examined how variability in an individual's decision making process can affect behavior in time-inconsistent planning scenarios. We adopted the graphical framework of Kleinberg and Oren (2014) and characterized worst-case task graphs for the agent. We showed that depending on the distribution of the present-bias parameter, the worst-case procrastination ratio is either bounded by constant, is at most linear, or grows at least exponentially with the number of agent's steps n. We also examined two natural families of tasks: (i) those in which the cost of reaching the goal is never more than at the outset, and (ii) those in which an agent can never lose progress toward the goal. We showed that in the first scenario the worst-case procrastination ratio is at most $O(n)$, and this is tight, and for the second scenario the procrastination ratio is $O(1)$ under some mild assumptions on the present-bias distribution.

Our work leaves open many avenues for future study. The following are a few concrete questions:

(1) We always assume that $b \geq 1$. What could be the worst-case procrastination ratio if b_i can be smaller than 1? For example, an agent might receive a reminder about the task and becomes anxious to complete it as soon as possible.

(2) We examine worst-case procrastination ratio for a fixed distribution of $b \sim \mathcal{F}$. Is there a bound on procrastination ratio for a given pair of the task graph and distribution $\mathcal{F}$? In other words, are certain distributions better suited to certain tasks than others? What is an important intrinsic parameter of a weighted graph that can lead to large procrastination?

(3) We examined the case of a simple, naïve agent who never anticipates their time-inconsistency. What can one say about agents who can reason about their future propensity to procrastinate?

References

Akerlof, G. A. (1991). Procrastination and obedience. *American Economic Review*, 81(2):1–19.

Alaei, S. (2011). Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. In *Proceedings of the 52nd Annual Symposium on Foundations of Computer Science*, pages 512–521. IEEE.

Alaei, S., Fu, H., Haghpanah, N., and Hartline, J. (2013). The simple economics of approximately optimal auctions. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, pages 628–637. IEEE.

Alaei, S., Hartline, J., Niazadeh, R., Pountourakis, E., and Yuan, Y. (2015). Optimal auctions vs. anonymous pricing. In *Proceedings of the 2015 IEEE 56th Annual Symposium on Foundations of Computer Science (FOCS)*, FOCS '15, pages 1446–1463, Washington, DC, USA. IEEE Computer Society.

Ariely, D. and Wertenbroch, K. (2002). Procrastination, Deadlines, and Performance: Self-Control by Precommitment. *Psychological Science*, 13(3):219–224.

Battaglini, M. (2005). Long-term contracting with markovian consumers. *The American economic review*, 95(3):637–658.

Beshears, J., Choi, J. J., Laibson, D., Madrian, B. C., and Milkman, K. L. (2015). The effect of providing peer information on retirement savings decisions. *Journal of Finance*, 70(3):1161–1201.

Cai, Y. and Daskalakis, C. (2011). Extreme-value theorems for optimal multidimensional pricing. In *Proceedings of the 52nd Annual Symposium on Foundations of Computer Science*, pages 522–531. IEEE.

Chawla, S., Hartline, J., and Kleinberg, R. (2007). Algorithmic pricing via virtual valuations. In *Proc. 8th ACM Conf. on Electronic Commerce*.

Chawla, S., Hartline, J., Malec, D., and Sivan, B. (2010a). Sequential posted pricing and multi-parameter mechanism design. In *Proc. 41st ACM Symp. on Theory of Computing*.

Chawla, S., Malec, D., and Sivan, B. (2010b). The power of randomness in bayesian optimal mechanism design. In *Proc. 11th ACM Conf. on Electronic Commerce*, pages 149–158.

Choi, J., Laibson, D., and Madrian, B. C. (2004). Plan design and 401(k) savings outcomes. *National Tax Journal*, 57(2):275–298.

Clarke, E. H. (1971). Multipart pricing of public goods. *Public Choice*, 11(1):17–33.

Coase, R. H. (1972). Durability and monopoly. *The Journal of Law and Economics*, 15(1):143–149.

Colin, C. (2003). *Behavioral game theory: experiments in strategic interaction*, volume The Roundtable series in behavioral economics. Princeton University Press.

Conitzer, V., Taylor, C. R., and Wagman, L. (2012). Hide and seek: Costly consumer privacy in a market with repeat purchases. *Marketing Science*, 31(2):277–292.

Devanur, N. R., Peres, Y., and Sivan, B. (2015). Perfect bayesian equilibria in repeated sales. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 983–1002.

Ely, J. C. (2015). Beeps.

Frederick, S., Loewenstein, G., and O'Donoghue, T. (2002). Time Discounting and Time Preference: A Critical Review. *Journal of Economic Literature*, 40(2):351–401.

Fudenberg, D. and Tirole, J. (1983). Sequential bargaining with incomplete information. *The Review of Economic Studies*, 50(2):221–247.

Fudenberg, D. and Villas-Boas, J. M. (2006). Behavior-based price discrimination and customer recognition. *Handbook on economics and information systems*, 1:377–436.

Gravin, N., Immorlica, N., Lucier, B., and Pountourakis, E. (2016). Procrastination with variable present bias. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, EC '16, pages 361–361, Maastricht, The Netherlands. ACM.

Groves, T. (1973). Incentives in teams. *Econometrica*, 41(4):617–631.

Haghpanah, N. and Hartline, J. (2015). Reverse mechanism design. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 757–758. ACM.

Hart, O. D. and Tirole, J. (1988). Contract renegotiation and coasian dynamics. *The Review of Economic Studies*, 55(4):509–540.

Hartline, J. and Roughgarden, T. (2009). Simple versus optimal mechanisms. In *Proc. 10th ACM Conf. on Electronic Commerce*.

Hartline, J. D. (2012). Approximation in mechanism design. *The American Economic Review*, 102(3):330–336.

Immorlica, N., Lucier, B., Pountourakis, E., and Taggart, S. (2017). Repeated sales with multiple strategic buyers. In *In submission*.

Kaur, S., Kremer, M., and Mullainathan, S. (2010). Self-control and the development of work arrangements. *American Economic Review*, 100(2):624–28.

Kennan, J. (2001). Repeated bargaining with persistent private information. *The Review of Economic Studies*, 68(4):719–755.

Kleinberg, J., Oren, S., and Raghavan, M. (2016). Planning problems for sophisticated agents with present bias. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, EC '16, pages 343–360, Maastricht, The Netherlands. ACM.

Kleinberg, J. M. and Oren, S. (2014). Time-inconsistent planning: a computational problem in behavioral economics. In *ACM Conference on Economics and Computation, EC '14, Stanford , CA, USA, June 8-12, 2014*, pages 547–564.

Laibson, D. (2015). Why don't present-biased agents make commitments? *American Economic Review Papers and Proceedings*, 105(5):267–272.

Meier, S. and Sprenger, C. (2010). Present-biased preferences and credit card borrowing. *American Economic Journal: Applied Economics*, 2(1):193–210.

Myerson, R. (1981). Optimal auction design. *Mathematics of Operations Research*, 6:58–73.

Nikolova, E. and Karger, D. R. (2008). Route planning under uncertainty: The canadian traveller problem. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence, AAAI 2008, Chicago, Illinois, USA, July 13-17, 2008*, pages 969–974.

Nikolova, E. and Moses, N. E. S. (2011). Stochastic selfish routing. In *Algorithmic Game Theory, 4th International Symposium, SAGT 2011, Amalfi, Italy, October 17-19, 2011. Proceedings*, pages 314–325.

Nikolova, E. and Moses, N. E. S. (2014). A mean-risk model for the traffic assignment problem with stochastic travel times. *Operations Research*, 62(2):366–382.

Nikolova, E. and Moses, N. E. S. (2015). The burden of risk aversion in mean-risk selfish routing. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*, pages 489–506.

O'Donoghue, T. and Rabin, M. (1999). Doing it now or later. *American Economic Review*, 89(1):103–124.

O'Donoghue, T. and Rabin, M. (2008). Procrastination on long-term projects. *Journal of Economic Behavior & Organization*, 66(2):161–175.

Ru, H. and Schoar, A. (2016). Do credit card companies screen for behavioral biases? Working Paper 22360, National Bureau of Economic Research.

Schmidt, K. M. (1993). Commitment through incomplete information in a simple repeated bargaining game. *Journal of Economic Theory*, 60:114–114.

Tang, P., Teng, Y., Wang, Z., Xiao, S., and Xu, Y. (2014). Settling the open problems in kleinberg-oren time-inconsistent planning models. *CoRR*, abs/1411.7472.

Thaler, R. (1981). Some empirical evidence on dynamic inconsistency. *Economics Letters*, 8(3):201 – 207.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *J. of Finance*, 16:8–37.

Wright, J. R. and Leyton-Brown, K. (2010). Beyond equilibrium: Predicting human behavior in normal-form games. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2010, Atlanta, Georgia, USA, July 11-15, 2010*.

Wright, J. R. and Leyton-Brown, K. (2012). Behavioral game theoretic models: a bayesian framework for parameter analysis. In *International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2012, Valencia, Spain, June 4-8, 2012 (3 Volumes)*, pages 921–930.

Wright, J. R. and Leyton-Brown, K. (2014). Level-0 meta-models for predicting human behavior in games. In *ACM Conference on Economics and Computation, EC '14, Stanford , CA, USA, June 8-12, 2014*, pages 857–874.

Yan, Q. (2011). Mechanism design via correlation gap. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 710–719. SIAM.

APPENDIX A

Appendix to Chapter 3

A.1. Simulations

In this section, we briefly discuss simulation results for the worst-case instance derived in section Section 3.3. From these simulations we will see how fast, as a function of the number n of agents, the worst-case ratio of the ex ante relaxation to the expected revenue of optimal anonymous pricing converges to e. Moreover, for these worst-case instances we will be able to evaluate the approximation of the optimal auction by anonymous reserves and pricing.

Our worst-case instances are given by a continuum of agents (as given by $\mathcal{V}(\cdot)$ and $\mathcal{Q}(\cdot)$) which we discretize by evaluating $\mathcal{Q}(\cdot)$ on an arithmetic progression, denoted $\{\bar{q}_i\}$. Given the fast convergence that our simulation exhibits, there is little loss in this discretization.

According to the price revenue constraint in (P3), we know that the revenue by posting price $\bar{v}_i$ for every i should be equal to the optimal posting price revenue. Thus, when we have $\{\bar{q}_i\}$, it is simple to get the corresponding $\bar{v}_i$ by binary search and calculating the revenue.

After generating the instances, we also calculate the ratio of the revenue of the optimal mechanism to the anonymous pricing revenue, and the ratio of the revenue of the optimal mechanism to the revenue of the second price with anonymous reserve mechanism for these instances. We use sampling algorithm to calculate the revenue of the second price with anonymous reserve mechanism, while the calculation of the revenue of the optimal mechanism is exact. We report the results of our simulation in Figs. A.1 and A.2 for various numbers n of agents.

n	2	5	10	50	100	500	1000	5000
ExAnteRev/OptPriceRev	2.000	2.507	2.622	2.701	2.710	2.717	2.718	2.718
OptRev/OptPriceRev	2.000	2.138	2.187	2.223	2.227	2.231	2.231	2.232
OptRev/OptReserveRev	2.000	1.794	1.731	1.682	1.676	1.665	1.659	1.607

Figure A.1. The ratios of the revenues of various auctions and benchmarks (Table).
Here ExAnteRev and OptPriceRev are the ex ante relaxation and optimal anonymous pricing revenues (as previously defined). OptRev is the revenue of the optimal auction of Myerson (1981). OptReserveRev is the revenue obtained by the second-price auction with an optimally chosen reserve price.

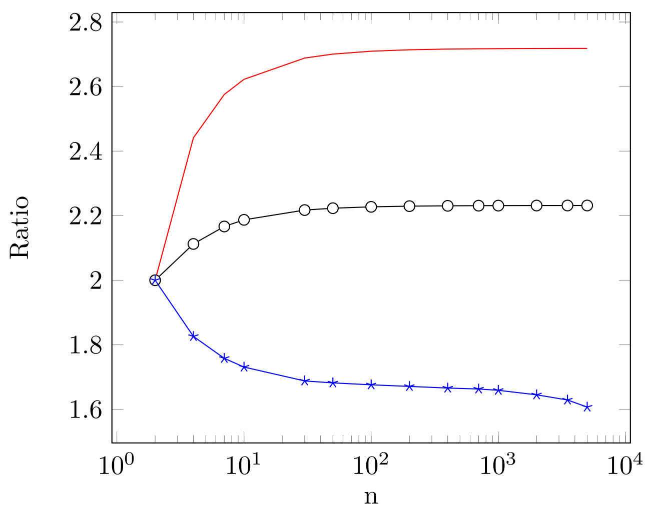

Figure A.2. The ratios of the revenues of various auctions and benchmarks (Graph).
The red sold line represents the ratio of ex-ante benchmark to anonymous pricing revenue. The blue line with star represents the ratio of the optimal revenue to anonymous pricing revenue. The black line with circle represents the ratio of the optimal revenue to the revenue of the second price auction with reserve.

APPENDIX B

Appendix for Chapter 4

B.1. Alternate Sufficient Conditions for Eliminating Learning Equilibria

In this appendix, we give an alternate refinement which selects out all but the equilibrium in which the seller learns nothing about the buyer. As in Section 4.4, we require that the seller posts prices which are above the bottom of the support of beliefs. Rather than upperbounding prices offered by the seller, we require that the buyers accept a price at the bottom of the support, whenever it is offered. Formally:

Definition B.1.1. A threshold PBE of the single-buyer game *respects lower bounds* if the following conditions hold for every history h^k with beliefs lower bounded by a:

- $t(h^k, p) = p$, for all $p \leq a$.
- $\sigma_S^k(h^k) \geq a$.

Theorem B.1.1. *In the single-buyer game, let the value distribution F be supported on $[a, b]$, with $a > 0$. If $\delta > \frac{b}{a+b}$, then in any threshold PBE which respects lower bounds, the seller posts a every round, which is accepted by all buyers. In other words, no learning will occur.*

Note that Theorem B.1.1 does not require that strategies be Markovian on path, though adding this requirement obviously does not change the result.

To prove the theorem, we show that any threshold PBE that respects lowerbounds and must have a non-monotone cumulative allocation function around any threshold other than a or b. To simplify our analysis, we will consider only threshold PBE in which the threshold buyer for each round accepts in the first round. Moreover, since the seller never will offer a subsequent price below t, we may have the threshold type reject for the remainder of the game. For any threshold PBE, one can change the strategies of threshold buyers to an accept followed by nonstop rejection without violating equilibrium, as such a sequence of actions is one of their optimal choices. The change doesn't affect the seller's expected utility because such agents are a measure zero set.

of Theorem B.1.1. Consider a threshold PBE with value distribution supported on $[a, b]$. Assume the threshold buyer in the first round behaves as described above: they accept in the first round, and then reject in every subsequent round. In this equilibrium, we have $X(t) = 1$. We will show that there is a lower-valued agent with with total discounted allocation strictly greater than 1. This violates Lemma 4.4.2 unless $t = a$, proving the theorem.

We will first find a buyer with value less than t with high total discounted payments. To do this, note that after seeing a rejection in the first round, the seller could offer a for the rest of the game, which by the natural threhsolds assumption would yield revenue $a\frac{\delta}{1-\delta}$. Since the seller is best responding, this implies that $\mathbb{E}[P(v) \,|\, v \leq t] \geq a\frac{\delta}{1-\delta}$. Since $P(v)$ is increasing, it must be that there is a set of values with positive measure in $[a, t]$ with total discounted payments at least $a\frac{\delta}{1-\delta}$. Choose some v from this set. We have $P(v) \geq a\frac{\delta}{1-\delta}$.

We now lowerbound $X(v)$. To do this, note that the buyer could choose to reject every round, so $U(v) \geq 0$. This in turn implies that $vX(v) \geq P(v)$, and therefore that $X(v) \geq \frac{a}{v}\frac{\delta}{1-\delta}$. By our assumption that $\delta > \frac{b}{a+b}$, we have:

$$X(v) \geq \frac{a}{v}\frac{\delta}{1-\delta} \geq \frac{a}{b}\frac{\delta}{1-\delta} > 1.$$

This yields the desired non-monotonicity of $X(\cdot)$, contradicting Lemma 4.4.2. It must therefore be that $t = a$. The seller does not learn in such an equilibrium, as the same arguments will hold for every subsequent round. □

B.2. Multiple buyer equilibrium for a single item

In Section 4.5, we showed that the dynamic pricing game with $\delta \geq 2/3$ supports an equilibrium with nontrivial revenue and learning. Moreover, in contrast Appendix B.1. We now generalize these conclusions to games with $n \geq 3$ buyers and $\delta \geq \frac{n}{n+1}$. In particular, we give a recursively-constructed equilibrium, built on the two-buyer equilibrium as a base case, in which the seller obtains non-trivial revenue, learning occurs, and which survives the refinements of Section 4.4 and Appendix B.1.

B.2.1. Equilibrium Description

Much like the two-buyer equilibrium, the multibuyer version has an exploration phase and an exploitation phase. In the exploration phase, the seller posts a price which induces a threshold response among the buyers. Those buyers that reject are priced out of the game, while the k buyers who accept continue in the k-buyer version of the equilibrium. This continues until all or all but one buyer rejects, at which point the seller exploits the buyers who most recently accepted, posting the bottom of their support in perpetuity.

The seller's strategy has two phases: an explore phase and an exploit phase. In the explore phase, the seller offers the price p each round such that there exists a threshold t satisfying the threshold equation

$$P_n(t)(t-p) = \frac{\delta}{1-\delta}\frac{F(t)^{n-1}}{n}(t-a), \tag{B.1}$$

where $P_n(t) = \sum_{j=0}^{n-1} \frac{1}{n-j}\binom{n-1}{j}F(t)^j(1-F(t))^{n-1-j}$ and $F(\cdot)$ is the CDF of the current beliefs of buyers who have not yet rejected, supported on $[a,b]$. Formally, define $T_a^b(p;n)$ to be the set of solutions to (4.12) with n buyers, and let p^* be the price with quantile $1/n$ according to the original distribution of buyers' values. Define $t_a^b(p;n)$ to be p^* if $p^* \in T_a^b(p;n)$, else $t_a^b(p;n) = \max_{t\in T_a^b(p;n)} t$ if $T_a^b(p;n) \neq \emptyset$ and infinity otherwise. Given this threshold function, the seller optimizes their revenue, which is given by the recurrence:

$$\begin{aligned} R(a,b,p,n) = \quad & F(t(p))^n(\tfrac{\delta a}{1-\delta}) + n(1-F(t(p)))F(t(p))^{n-1}\left(p + \tfrac{\delta t(p)}{1-\delta}\right) \\ & + \textstyle\sum_{j=2}^{n}\binom{n}{j}(1-F(t(p)))^j F(t)^{n-j}(p+\delta R(t(p),b,j)), \end{aligned}$$

where $R(t_a^b(p;n),b,j)$ is the seller's optimal revenue from the continuation game with j strong buyers, all distributed acccording to $F^b_{t_a^b(p)}$.

This will result in a price which increases every round until the explore phase ends. Three different events may end the explore phase and trigger the exploit phase: all buyers reject, all but one buyer rejects, or the price reaches the optimal anonymous price p^* for n buyers with the original distribution $F_0(\cdot)$, which is given by $p^* = \max_p p(1-F_0(p)^n)$.

When the exploit phase begins, there will be a set of buyers whose beliefs have a common interval of support $[a, b]$ which is strictly above those of all other buyers. The seller posts the minimum value of this common support for the rest of the game.

On-path, the seller (and buyers) update their beliefs with standard Bayesian updating. If, however, the seller's beliefs are such that a buyer i is expected to always reject or always accept every round, then the seller's belief updates will punish buyer i for not adhering to these expectations. When buyer i deviates to an off-path action, the seller updates their beliefs for that agent to the maximum value in the support of F_0 and then enters the exploit phase, pricing at this maximum value for the rest of the game.

We now describe the buyers strategies. In the explore phase, the seller offers a price which elicits a threshold response, and the buyers who have not yet rejected respond in kind, rejecting below the threshold and accepting above. If the seller offers a price below the support of the beliefs for these buyers, the buyers accept the price, and if the seller offers a price above the support or a price such that there is no threshold response, the buyers reject. If a buyer has rejected in the explore phase, or if they have deviated in the past to an off-path action, then they become price-takers.

In the exploit phase, the seller is targeting the set of buyers who have the strongest support, with lowerbound a. All buyers in this group refuse any price above a, and accept any price below a. All other buyers act as price takers in this phase.

ALGORITHM 9: Seller's Strategy

Input : Purchasing history h^k, Support bounds $\{(a_i^k, b_i^k)\}_{i=1}^n$
Output: Price p^{k+1}
$S^k = \text{argmax}_i \ a_i^k$;
$a = \max_i a_i^k$;
$b = \max_i b_i^k$;
if $\mathbf{D}^k = R^n$ **or** $|S^k| = 1$ **or** $a \geq p^*$ **then**
  $p^{k+1} = a$;
else
  $p^{k+1} = \arg\max_p R(a, b, p, |S^k|)$;

Theorem B.2.1. *The strategies and belief updates laid out in Algorithms 9, 10, and 11 is a PBE as long as* $\delta \geq \frac{n}{n+1}$.

ALGORITHM 10: Buyer i Strategy

Input : Purchasing history h_b^k, Support bounds $\{(a_j^k, b_j^k)\}_{j=1}^n$, value v_i, price p^k
Output: Purchasing decision for round k
$S^k = \text{argmax}_j\ a_j^k$;
$a = \max_j a_j^k$;
$b = \max_j b_j^k$;
if $\mathbf{D}^k = R^n$ **or** $|S^k| = 1$ **or** $a \geq p^*$ **then**
  Accept if and only if $p^k \leq a$ and $v_i \geq p^k$;
else if $i \notin S^k$ **then**
  Accept if and only if $p^k \leq v_i$;
else
  Accept if and only if $v_i \geq t_a^b(p; |S^k|)$;

ALGORITHM 11: Belief updates

Input : Purchasing history h^{k+1}, Current support bounds $\{(a_i^k, b_i^k)\}_{i=1}^n$, First-round common support bounds (a, b).
Output: Updated support bounds $\{(a_i^{k+1}, b_i^{k+1})\}_{i=1}^n$.
$S^k = \text{argmax}_j\ a_j^k$;
for $i \in \{1, \ldots, n\}$ **do**
  if all types for i should reject p^k **then**
    if buyer i rejected p^k **then**
      $a_i^{k+1} = a_i^k$, $b_i^{k+1} = b_i^k$
    else
      $a_i^{k+1} = b$, $b_i^{k+1} = b$
  else if all types should accept **then**
    if buyer i accepted **then**
      $a_i^{k+1} = a_i^k$, $b_i^{k+1} = b_i^k$
    else
      $a_i^{k+1} = b$, $b_i^{k+1} = b$
  else
    if buyer i accepted **then**
      $a_i^{k+1} = t_{a_i^k}^{b_i^k}(p^k, |S^k|)$, $b_i^{k+1} = b_i^k$
    else
      $a_i^{k+1} = a_i^k$, $b_i^{k+1} = t_{a_i^k}^{b_i^k}(p^k, |S^k|)$

We prove the theorem in much the same manner as Theorem 4.6.1, by analyzing the buyer and seller incentives separately, in Sections. Before we proceed, we note that a common term in much of the analysis is the probability a buyer i winning the item when the threshold is t, the common distribution is $F(\cdot)$, and the number of buyers other than i is $n-1$. This value is given by $P_n(t) = \sum_{j=0}^{n-1} \frac{1}{n-j}\binom{n-1}{j} F(t)^j (1-F(t))^{n-1-j}$. This can be simplified by the following lemma:

Lemma B.2.2. *For all $n \geq 2$ and t in the support of $F(\cdot)$,*

$$P_n(t) = \frac{1}{n}\sum_{j=0}^{n-1} F(t)^j$$

Proof. We give a combinatorial proof. The probability $P_n(t)$ was generated as follows: j buyers out of buyer i's $n-1$ competitors will have values below threshold t (and therefore buy) with probability $\binom{n-1}{j}F(t)^j(1-F(t))^{n-1-j}$. Conditioned on this event, the probability of winning the item is $\frac{1}{n-j}$. This same process can be executed in a different way: first permute the n buyers (including i) uniformly at random, and then give the item to the first buyer in the order whose value is over t (or to buyer i, who we have assumed accepts). The probability that buyer i is jth in the order is $1/n$, and conditioned on this event, their probability of winning is $F(t)^{j-1}$. □

B.2.2. Buyer Incentives

As in the two-buyer equilibrium, the seller punishes buyers for out-of-equilibrium actions such as accepting a price with no threshold or rejecting a price with a threshold t and then accepting a price above t in subsequent rounds. We first formalize this:

Lemma B.2.3. *If a buyer takes an out of equilibrium action, then they cannot get positive utility from subsequent rounds.*

Proof. When an agent i takes an out-of-path action in round k, the beliefs on their value are updated to a pointmass the upper bound b of the initial distribution's support. Moreover, since $a_i^{k+1} = b_i^{k+1} = b$ is clearly greater than the target price p^* for the initial distribution, the seller will only post b in subsequent rounds. No buyer can get positive utility from such prices. □

Lemma B.2.4. *For $\delta \geq \frac{n}{n+1}$, all buyers are best responding to the strategies of the seller and other buyers.*

Proof. We argue from the perspective of buyer i in round k, and break our analysis into the exploration and exploitation phases.

Exploration Phase. There are two cases to consider: either buyer i has accepted in all previous rounds, and therefore $i \in S^k$, or buyer i has rejected in a previous round.

First, note that if buyer i has rejected in a previous round buy following the threshold strategy for that round, then they will continue rejecting all subsequent prices. This follows from the fact that all subsequent prices will be above the threshold for the round where i rejected.

If $i \in S^k$, then they are among the agents with a common support, say $[a, b]$, being targeted by the seller. We first argue that buyers in S^k will respond according to the threshold strategy for the $|S^k|$-player version of our equilibrium. All other buyers will reject, by the previous paragraph. We analyze the incentives of an arbitrary buyer in the exploration phase. Without loss of generality, consider the perspective of buyer 1, and consider their decision in the first round of the game. Let $F(\cdot)$ be the CDF of the beliefs, supported on $[a, b]$.

Assume $v_1 < t$. We argue that the buyer would prefer to reject. If they were to accept, they would win at price p with probability $P_n(t)$. They would not receive any utility in the future, as all future prices would be at least t. Hence their utility from accepting is exactly $u_A = P_n(t)(v_1 - p)$.

If they reject, then future prices will be below t only if all other agents also reject (triggering the exploit phase), which occurs with probability $F(t)^{n-1}$. If this occurs, then buyer 1 receives utility $\frac{\delta}{1-\delta}\frac{v_1-a}{n}$. Hence $u_R = \frac{F(t)^{n-1}}{n}\frac{\delta}{1-\delta}(v_1 - a)$.

To show that $u_R \geq u_A$, note that the threshold equation (B.1) implies that

$$P_n(t) \geq \frac{\delta}{1-\delta}\frac{F(t)^{n-1}}{n}. \tag{B.2}$$

Plugging $t = v_1 + x$ for $x \geq 0$ into equation (B.1) and applying inequality (B.2) proves that $u_R \geq u_A$, as desired.

Now assume $v_1 \geq t$. We will show that accepting is preferable to rejecting. To lowerbound buyer 1's utility for accepting, notice that if they accept, they will win the item in the current round at price p with probability $P_n(t)$. They will also likely receive utility in future rounds. As a lower bound on this future utility, they will receive utility $\frac{\delta}{1-\delta}\frac{v_1-t}{n}$ in the event that all other buyers reject the current price, which occurs with probability $F(t)^{n-1}$. Hence,

$$u_A \geq P_n(t)(v_1 - p) + \frac{\delta}{1-\delta}\frac{F(t)^{n-1}}{n}(v_1 - t).$$

If buyer 1 rejects, their utility in future rounds depends on the actions of the other buyers. If all other buyers reject (which occurs with probability $F(t)^{n-1}$, they receive utility $\frac{\delta}{1-\delta}\frac{v_1-a}{n}$ from acccepting for the rest of the game. Otherwise, at least one other buyer accepts, driving future prices above t. The seller believes buyer 1 to have value less than t because of their rejection. If buyer 1 chooses to accept in future rounds, this is seen as an off-path deviation, which will cause the seller to post prohibitively high prices for the rest of the game. Hence, if at least one other buyer accepts, buyer 1 can only get utility from the next round. The expected utility contribution of this event can be shown to be $\delta P_n(t)(v_1 - t)$. Hence:

$$u_R \leq \delta P_n(t)(v_1 - t) + \frac{\delta}{1-\delta}\frac{F(t)^{n-1}}{n}(v_1 - a)$$

Write $v = t + x$. The utility difference between accepting and rejecting satisfies:

$$\begin{aligned} u_A - u_R \quad & \geq (t + x - p)P_n(t) + \tfrac{\delta}{1-\delta}\tfrac{F(t)^{n-1}}{n}(t + x - t) \\ & -\delta P_n(t)(t + x - t) - \tfrac{\delta}{1-\delta}\tfrac{F(t)^{n-1}}{n}(t + x - a) \\ & = P_n(t)(t - p) - \tfrac{\delta}{1-\delta}\tfrac{F(t)^{n-1}}{n}(t - a) \\ & +x\left[P_n(t) + \tfrac{\delta}{1-\delta}\tfrac{F(t)^{n-1}}{n} - \delta P_n(t) - \tfrac{\delta}{1-\delta}\tfrac{F(t)^{n-1}}{n}\right] \\ & = x\left[P_n(t) + \tfrac{\delta}{1-\delta}\tfrac{F(t)^{n-1}}{n} - \delta P_n(t) - \tfrac{\delta}{1-\delta}\tfrac{F(t)^{n-1}}{n}\right] \quad \text{[By (B.1)]} \end{aligned}$$

(B.3)

The last line is obviously nonnegative, so buyer 1 prefers to accept.

The only case left to argue is that sellers will choose to reject a price that has no threshold response. In this case, the buyer strategies dictate that they must all reject. We show that this is a best response. We argue from the perspective of the first round, in which all buyers are in the targeted set S^k. If there are buyers outside of S^k, the argument does not change.

If buyer i accepts the current price p, then they receive utility $U_A = (v_1 - p)$, and by Lemma B.2.3, they will not receive any utility from future rounds. If buyer i rejects, then because all other buyers also reject, the exploitation phase will begin in the next round, yielding utility $U_R = \frac{\delta}{1-\delta}\frac{v_1-a}{n}$.

Becuase there is no threshold response, it must be that the threshold equation (B.1) has no solution. The argument mirrors the two-buyer case: for $t < p$, the lefthand side of (B.1) is negative, while the righthand side is always nonnegative. By continuity, it must be that the lefthand side is always less than the righthand side, or else they would cross, yielding a threshold response. Hence,

$$P_n(t)(t-p) < \frac{\delta}{1-\delta}\frac{F(t)^{n-1}}{n}(t-a).$$

By Lemma B.2.2, $P_n(t) = \frac{1}{n}\sum_{j=0}^{n-1} F(t)^j$. This is an average of n terms which are all larger than $F(t)^{n-1}$. Hence, $P_n(t) \geq F(t)^{n-1}$ for all t. We may therefore conclude that $t - p < \frac{\delta}{1-\delta}(t-a)$ for all values of t, including v_1, proving the claim.

Exploitation Phase. In the exploitation phase, the seller prices at the bottom of the belief support of set of targeted buyers S^k, with common support $[a, b]$. If the seller offers a lower, the buyers act as price-takers. If the seller offers a higher price, all buyers reject. We must show that each of these behaviors is a best response.

Seller offers $p \leq a$. Note that if a buyer who is expected to accept chooses instead to reject, they get no utility in the current round, and no utility in the future, by Lemma 4.6.2. Accepting the current price is clearly preferable. If a buyer is not in S^k, then they cannot affect future prices, and should consequently act as a price taker.

Seller offers $p > a$. We first argue from the perspective of buyer i in S^k. The utility of buyer i from rejecting is $\frac{\delta}{1-\delta}\frac{v_i-a}{|S^k|}$, as they will win the item with probability $1/|S^k|$ for the rest of the game. If they accept, they will receive utility $v_1 - p$, as by Lemma 4.6.2, they will not receive utility in subsequent rounds. Since $\delta \geq \frac{n}{n+1} \geq \frac{|S^k|}{|S^k|+1}$, we have that $\frac{\delta}{|S^k|(1-\delta)} \geq 1$. Since $p > a$ as well, we therefore have that rejection is optimal.

It follows that in both the exploration and exploitation phases, the buyers are best responding to the strategy of the seller. $\square$

B.2.3. Seller Incentives

Lemma B.2.5. *The seller is best responding to the actions of the buyers.*

Proof. As in the two-buyer case, we break our analysis into three cases: exploration phase, exploitation phase, and off-path analysis.

Exploration Phase. In round k of the exploration phase the seller has three options: price below the common support of the targeted set S^k, post a price p such that there is a threshold response solving equation (B.1) for $n = |S^k|$, or post a price p such that no solution to (B.1) exists. We will show that posting a price which induces threshold response is optimal. Since the seller's equilibrium strategy is defined implicitly to be the optimal such price, it will follow that they are best responding in the exploration phase.

Let a be the lower bound of the common support of S^k, which we assume to be $[a, b]$. Pricing below a yields less expected revenue than pricing at a. To see this, note that any price $p \leq a$ will be accepted with probability 1 by all buyers and cause the beliefs to remain unchanged. The seller would prefer to induce this outcome with a higher price.

Posting a price for which a threshold response does not exist is suboptimal. This follows from the fact that all buyers will reject, yielding no revenue and no update to the beliefs. The seller effectively skips a round, which is obviously suboptimal.

Exploitation Phase. In the exploitation phase, the seller posts at the bottom a of the common support of the targeted set S^k. We will show that the seller prefers this to their other options, which are posting a lower price, or posting a higher price.

The seller prefers not to post a lower price for the same reason that they would prefer not to post below the common support in the exploration phase. As long as the price p is at most a, it will be accepted with probability 1, and will not cause the beliefs about the stronger of the two buyers to change. Posting the largest price which induces this outcome is preferable to other such prices.

If the seller posts a price $p > a$, then this price is rejected by all agents, and beliefs do not change. As in the exploitation phase, this is suboptimal.

Off-Path Analysis. We finally analyze the case where a buyer has taken an action in the current history that was not expected of any type. In this case, the seller updates their beliefs to the highest possible value b, and the buyer strategy dictates that they behave as a price-taker. The optimal response to such a scenario is to post b every round. □

B.3. Multiple buyer equilibrium for digital goods

Similar to Appendix B.2 we extend the no-learning equilibrium for digital goods. We show that the dynamic pricing game with $\delta \geq 1/2$ supports an equilibrium with nontrivial revenue and learning.

B.3.1. Equilibrium Description

Likewise, the multi-buyer version has an exploration phase and an exploitation phase. In the exploration phase, the seller posts a price which induces a threshold response among the buyers. Those buyers that reject are priced out of the game, while the k buyers who accept continue in the k-buyer version of the equilibrium. This continues until all or all but l buyers reject, at which point the seller exploits the buyers who most recently accepted, posting the bottom of their support in perpetuity.

The seller's strategy has two phases: an explore phase and an exploit phase. In the explore phase, the seller offers the price p each round such that there exists a threshold t satisfying the threshold equation

$$t - p = \frac{\delta}{1-\delta} F(t)^{n-1}(t-a), \tag{B.4}$$

where $F(\cdot)$ is the CDF of the current beliefs of buyers who have not yet rejected, supported on $[a, b]$. Formally, define $T_a^b(p; n)$ to be the set of solutions to (B.4) with n buyers, and let p^* a target price according to the original distribution of buyers' values. Define $t_a^b(p; n)$ to be p^* if $p^* \in T_a^b(p; n)$, else $t_a^b(p; n) = \max_{t \in T_a^b(p;n)} t$ if $T_a^b(p; n) \neq \emptyset$ and infinity otherwise. Given this threshold function, the seller optimizes their revenue, which is given by the recurrence:

$$R(a,b,p,n) = F(t_a^b(p;n))^n(\tfrac{\delta a}{1-\delta}) + n(1-F(t_a^b(p;n)))F(t_a^b(p;n))^{n-1}\left(p + \tfrac{\delta t_a^b(p;n)}{1-\delta}\right)$$
$$+ \textstyle\sum_{j=2}^{n}\binom{n}{j}(1-F(t_a^b(p;n)))^j F(t_a^b(p;n))^{n-j}(p+\delta R(t_a^b(p;n),b,j)),$$

where $R(t_a^b(p;n),b,j)$ is the seller's optimal revenue from the continuation game with j strong buyers, all distributed according to $F^b_{t_a^b(p)}$.

This will result in a price which increases every round until the explore phase ends. Three different events may end the explore phase and trigger the exploit phase: all buyers reject, at least one but less than or equal to l buyers accept, or the price reaches a target price p^*.

ALGORITHM 12: Seller's Strategy

```
Input  : Purchasing history h^k, Support bounds {(a_i^k, b_i^k)}_{i=1}^n
Output: Price p^{k+1}
S^k = argmax_i a_i^k;
a^k = max_i a_i^k;
b^k = max_i b_i^k;
if p^i < a^i for any i ≤ k then
    Let i* be the earliest round where p^{i*} < a^{i*}.
    if Any buyer has accepted a positive price after round i* then
        p^{k+1} = n · b;
    else if  Buyer has ever rejected a price of 0 then
        p^{k+1} = n · b;
    else
        p^{k+1} = 0;
if D^k = R^n or |S^k| ≤ l or a ≥ p* then
    p^{k+1} = a;
else
    p^{k+1} = arg max_p R(a^k, b^k, p, |S^k|) ;
```

Theorem B.3.1. *The strategies and belief updates laid out in Algorithms 12, 13, and 14 is a PBE as long as $\delta \geq \frac{1}{2}$.*

We prove the theorem in much the same manner as the previous theorems, by analyzing the buyer and seller incentives separately, in Sections.

ALGORITHM 13: Buyer i Strategy

Input : Purchasing history h_b^k, Support bounds $\{(a_j^k, b_j^k)\}_{j=1}^n$, value v_i, price p^k
Output: Purchasing decision for round k

```
S^k = argmax_j a_j^k;
a^k = max_j a_j^k;
b^k = max_j b_j^k;
if p^k < a^k then
    Reject;
if p^i < a^k for any i ≤ k − 1 then
    if p^k = 0 then
        Accept;
    else if p^k > 0 then
        Let i* be the earliest round where p^{i*} < a^{i*}.
        if Buyer has ever accepted a positive price after round i* then
            Accept;
        else if Buyer has ever rejected a price of 0 then
            Accept;
        else
            Reject;
if i ∉ S_i then
    Reject;
else
    if D^k = R^n or |S^k| ≤ l or a ≥ p* then
        Accept if and only if p^k = a;
    else
        Accept if and only if v_i ≥ t_a^b(p; |S^k|);
```

B.3.2. Buyer Incentives

As in the two-buyer equilibrium, the seller punishes buyers for out-of-equilibrium actions such as accepting a price with no threshold or rejecting a price with a threshold t and then accepting a price above t in subsequent rounds. We first formalize this:

Lemma B.3.2. *If a buyer takes an out of equilibrium action, then they cannot get positive utility from subsequent rounds.*

Proof. When an agent i takes an out-of-path action in round k, the beliefs on their value are updated to a pointmass the upper bound b of the initial distribution's support. Moreover, since $a_i^{k+1} = b_i^{k+1} = b$ it is clearly greater than the target price p^* No buyer can get positive utility from such prices. □

Lemma B.3.3. *For $\delta \geq \frac{1}{2}$, all buyers are best responding to the strategies of the seller and other buyers.*

ALGORITHM 14: Belief updates

Input : Purchasing history h^{k+1}, Current support bounds $\{(a_i^k, b_i^k)\}_{i=1}^n$, First-round common support bounds (a, b).

Output: Updated support bounds $\{(a_i^{k+1}, b_i^{k+1})\}_{i=1}^n$.

```
S^k = argmax_j a_j^k;
for i ∈ {1, ..., n} do
    if all types for i should reject p^k then
        if buyer i rejected p^k then
            a_i^{k+1} = a_i^k, b_i^{k+1} = b_i^k
        else
            a_i^{k+1} = b, b_i^{k+1} = b
    else if all types should accept then
        if buyer i accepted then
            a_i^{k+1} = a_i^k, b_i^{k+1} = b_i^k
        else
            a_i^{k+1} = b, b_i^{k+1} = b
    else
        if buyer i accepted then
            a_i^{k+1} = t_{a_i^k}^{b_i^k}(p^k, |S^k|), b_i^{k+1} = b_i^k
        else
            a_i^{k+1} = a_i^k, b_i^{k+1} = t_{a_i^k}^{b_i^k}(p^k, |S^k|)
```

Proof. We argue from the perspective of buyer i in round k, and break our analysis into the exploration and exploitation phases. If a buyer has taken an off-path action in the history, then Lemma B.3.2 implies that incentives are trivial

Exploration Phase. There are two cases to consider: either buyer i has accepted in all previous rounds, and therefore $i \in S^k$, or buyer i has rejected in a previous round. First, note that if buyer i has rejected in a previous round buy following the threshold strategy for that round, then they will continue rejecting all subsequent prices. This follows from the fact that all subsequent prices will be above the threshold for the round where i rejected.

If $i \in S^k$, then they are among the agents with a common support, say $[a, b]$, being targeted by the seller. We first argue that buyers in S^k will respond according to the threshold strategy for the $|S^k|$-player version of our equilibrium. All other buyers will reject, by the previous paragraph. We analyze the incentives of an arbitrary buyer in the exploration phase. Without loss of generality, consider the perspective of buyer 1, and consider their decision in the first round of the game. Let $F(\cdot)$ be the CDF of the beliefs, supported on $[a, b]$.

Assume $v_i < t$. We argue that the buyer would prefer to reject. If they were to accept, they would win at price p. They would not receive any utility in the future, as all future prices would be at least t. Hence their utility from accepting is exactly $u_A = (v_i - p)$.

If they reject, then future prices will be below t only if all other agents also reject (triggering the exploit phase), which occurs with probability $F(t)^{n-1}$. If this occurs, then buyer 1 receives utility $\frac{\delta}{1-\delta}(v_1 - a)$. Hence $u_R = F(t)^{n-1}\frac{\delta}{1-\delta}(v_i - a)$.

To show that $u_R \geq u_A$, note that the threshold equation (B.4) implies that

$$1 \geq \frac{\delta}{1-\delta}\mathcal{F}(t)^{n-1}. \tag{B.5}$$

Plugging $t = v_1 + x$ for $x \geq 0$ into equation (B.4) and applying inequality (B.5) proves that $u_R \geq u_A$, as desired.

Now assume $v_1 \geq t$. We will show that accepting is preferable to rejecting. To lowerbound buyer 1's utility for accepting, notice that if they accept, they will win the item in the current round at price p and will also likely receive utility in future rounds. As a lower bound on this future utility, they will receive utility $\frac{\delta}{1-\delta}v_1 - t$ in the event that all other buyers reject the current price, which occurs with probability $F(t)^{n-1}$. Hence,

$$u_A \geq v_1 - p + \frac{\delta}{1-\delta}F(t)^{n-1}(v_1 - t).$$

If buyer 1 rejects, their utility in future rounds depends on the actions of the other buyers. If all other buyers reject (which occurs with probability $F(t)^{n-1}$, they receive utility $\frac{\delta}{1-\delta}(v_1 - a)$ from accepting for the rest of the game. Otherwise, at least one other buyer accepts, driving future prices above t. The seller believes buyer 1 to have value less than t because of their rejection. If buyer 1 chooses to accept in future rounds, this is seen as an off-path deviation, which will cause the seller to post prohibitively high prices for the rest of the game. Hence, if at least one other buyer accepts, buyer 1 can only get utility from the next round. The expected utility contribution of this event can be shown to be $v_1 - t$. Hence:

$$u_R \leq \delta(v_1 - t) + \frac{\delta}{1-\delta}F(t)^{n-1}(v_1 - a)$$

Write $v = t + x$. The utility difference between accepting and rejecting satisfies:

$$\begin{aligned} u_A - u_R \quad & \geq (t + x - p) + \tfrac{\delta}{1-\delta}F(t)^{n-1}(t + x - t) \\ & -\delta(t + x - t) - \tfrac{\delta}{1-\delta}F(t)^{n-1}(t + x - a) \\ & = (t - p) - \tfrac{\delta}{1-\delta}F(t)^{n-1}(t - a) \\ & +x\left[1 + \tfrac{\delta}{1-\delta}F(t)^{n-1} - \delta - \tfrac{\delta}{1-\delta}F(t)^{n-1}\right] \\ & = x\left[1 + \tfrac{\delta}{1-\delta}F(t)^{n-1} - \delta - \tfrac{\delta}{1-\delta}F(t)^{n-1}\right] \quad \text{[By EquationB.4]} \end{aligned}$$

(B.6)

The last line is obviously non-negative, so buyer 1 prefers to accept.

Assume (B.4) has no solution. It is possible for the price p to be such that the threshold equation has no solution. In this case, we require that all buyers reject as a best response.

If buyer i accepts, they get utility $U_A = v_1 - p$, as rest buyers will reject this round, and by Lemma B.3.2, they will receive no utility in the future. If they reject, then the seller will enter the exploit phase of the equilibrium and post a for the rest of the game. This yields utility $U_R = \frac{\delta}{1-\delta}(v_1 - a)$. Since $\delta \geq 1/2$ we get that $U_R > U_A$, as desired.

Exploitation Phase: In the exploitation phase, the seller prices at the bottom of the belief support of the strongest buyer. If the seller offers this price or lower, the buyers reject, and will only accept prices of zero for the rest of the game. If the seller offers a higher price, both buyers reject. We must show that each of these behaviors is a best response.

Seller offers $p \leq a_i^k$. The agent can get utility from accepting the current round, but will face the highest possible price b for the rest of the game if they do so. They would prefer to reject and receive their item for free for the rest of the game, in a zero-learning continuation.

Seller offers $p > \max(a_i^k)$. . According to the equilibirum all buyers should reject hence if buyer i by Lemma B.3.2, they will not receive utility in subsequent rounds. Since $\delta \geq 1/2$, we have that $\frac{\delta}{1-\delta} \geq 1$. Since $p > 0$ as well, we have that $\frac{\delta}{1-\delta} v_1 > v_1 - p$. Hence, rejection is optimal.

In the case where $a_1 \neq a_2$, the incentive to reject is even greater, due to the fact that the higher-valued agent may receive the item with probability 1 in every subsequent round if they reject. Hence, rejection is optimal here as well. □

B.3.3. Seller Incentives

Lemma B.3.4. *The seller is best responding to the actions of the buyers.*

Proof. We break our analysis into cases: exploration phase and , exploitation phase, and off-path analysis.

Exploration Phase. In the exploration phase all buyers in S^k with $|S^k| > l$. have the same support, and the seller has three options: price below the common support, post a price p such that there is a threshold response solving equation (B.4), or post a price p such that no solution to (4.12) exists. We will show that the second option, posting a price within the common support such that there is a threshold response, is optimal. Since the seller's equilibrium strategy is defined implicitly to be the optimal such price, it will follow that they are best responding in the exploration phase.

Let a^k be the lower bound of the common support in S^k. Pricing below a^k will cause all buyers to reject any nonzero price for the remainder of the game. This implies that the seller prefers to post a price within the common support.

We now argue that posting a price for which a threshold response does not exist is suboptimal. This follows from the fact that all buyers will reject, yielding no revenue and no update to the beliefs. The next round's decision problem is identical to that of the current round, but with payoffs discounted by δ. The seller clearly does not benefit from skipping a round in this manner.

Exploitation Phase. In the exploitation phase, the seller posts a^k. We will show that the seller prefers this to their other options, which are posting a lower price, or posting a higher price.

The seller prefers not to post a lower price for the same reason that they would prefer not to post below the common support in the exploration phase. As long as the price p satisfies $p < a^k$, the buyers will reject any nonzero price.

If the seller posts a price $p > a^k$, then this price is rejected by all agents, and beliefs do not change. The seller effectively skips the round and is faced with the same decision next round, with a discount. This is not optimal. Hence, the optimal price to post in the exploitation phase is a^k.

Off-Path Analysis. We finally analyze the case where a buyer has taken an action in the current history that was not expected of any type. In this case, the seller updates their beliefs to the highest possible value b, and this buyer's strategy dictates that they behave as a price-taker, while the rest of the buyers will not accept non-zero prices. Hence, the optimal response to such a scenario is to post b every round. $\square$