

# Repeated Sales with Multiple Strategic Buyers

NICOLE IMMORLICA, Microsoft Research

BRENDAN LUCIER, Microsoft Research

EMMANOUIL POUNTOURAKIS, Northwestern University

SAMUEL TAGGART, Northwestern University

---

In a market with repeated sales of a single item to a single buyer, prior work has established the existence of a zero revenue perfect Bayesian equilibrium in the absence of a commitment device for the seller. This counter-intuitive outcome is the result of strategic purchasing decisions, where the buyer worries that the seller will update future prices in response to past purchasing behavior. We first show that in fact almost any revenue can be achieved in equilibrium, but the zero revenue equilibrium uniquely survives natural refinements. This establishes that single buyer markets without commitment are subject to market failure. However, our main result shows that this market failure depends crucially on the assumption of a single buyer. If there are multiple buyers, the seller can approximate the revenue that is possible with commitment. We construct an intuitive equilibrium for multiple buyers that survives our refinements, in which the seller learns from past purchasing behavior and obtains a constant factor of the per-round Myerson optimal revenue. Moreover, we describe a simple and computationally tractable pricing algorithm for the seller that achieves this approximation when buyers best-respond.

---

## 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 Fodler 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 opportunity 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 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. 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. This extreme and counterintuitive equilibrium is driven by a self-fulfilling prophecy: the buyer never accepts any positive price out of fear that doing so will lead the seller to charge very high prices in the future; as a result, the seller infers that only a buyer with very high type would ever accept a positive price, so the seller would indeed charge very high prices in response. The formal details of the equilibrium are described in Section 3. This construction illustrates that in the absence of commitment power, a seller might suffer extremely low revenue in long-term interaction with a buyer. We note that this conclusion is reminiscent 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].

This result is quite negative, but also 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? One 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 targeted 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: would the presence of multiple buyers change the structure of equilibrium?

*Our Results.* The existence of a zero-revenue equilibrium is discouraging, but we begin by showing that the single-buyer situation is even more dire than that. One might wonder whether the low-revenue equilibrium is simply an edge case, and that better and more plausible equilibria exist. Indeed, we establish a folk theorem that implies that any amount of revenue between the trivial lower bound (that of posting the minimum-supported value every round) and that of Myerson pricing every round can be realized at a PBE of the game. However, despite the rich space of equilibria, we prove that the zero-revenue equilibrium is the unique equilibrium that survives a natural refinement of the set of PBE. Specifically, it is the unique equilibrium in which the buyer uses threshold strategies (i.e., on each round and for any offered price, a buyer purchases if and only if their value is sufficiently high), strategies are Markovian on-path (meaning that on the equilibrium path, the players' strategies depend only on their beliefs and the current price, and not the full history of past play), and the seller offers prices in the support of buyers' distributions. These refinements have been studied previously in the context of repeated sales (see [Fudenberg and Tirole, 1983] and [Hart and Tirole, 1988]), and are natural conditions for "simple" strategies. We interpret this as strong evidence that the zero-revenue equilibrium, and the market failure it implies, is actually a plausible and natural outcome of the single-buyer repeated game.

*Main Result: Multiple Buyers.* 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.<sup>1</sup>

In contrast to the single-buyer variant, we show that the seller can achieve a constant fraction of the benchmark optimal revenue in a PBE that employs threshold strategies and is Markovian on-path, surviving our refinements. 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.

This equilibrium structure sets up a natural optimization problem for the seller: how quickly should prices be increased, given the way that rational buyers will respond at equilibrium? Typical of explore/exploit algorithms, the seller must balance the rate of learning with the revenue ultimately generated in the exploitation phase. Different pricing policies will correspond to different equilibria, with potentially different amounts of revenue. We provide an approximation result: for two buyers, if the distribution over buyer valuations satisfies the standard monotone hazard rate (MHR) condition, then we show how to compute prices (and the corresponding equilibrium thresholds for the buyers) that generate a constant fraction ( $1/3e^2$ ) of the Myerson optimal revenue. For  $n \geq 3$  bidders, we obtain a stronger guarantee (.12-approximation) for a broader class of distributions (regular with monopoly

---

<sup>1</sup>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.

quantile at least  $1/n$ ). The revenue of the prices we analyze are a lower bound on the seller’s revenue in our equilibrium.

### 1.1 Related Work

Hart and Tirole [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 a constant number of final rounds. Schmidt [1993] generalized 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 commitment, 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.

### 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 assumptions 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.

## 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 5.2 and 6.2, 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 the vector  $\mathbf{p}[k-1] = (p_1, \dots, p_{k-1})$  of past prices, as well as the purchasing decisions of agents in each past round, denoted  $\mathbf{D}[k-1] = (\mathbf{D}^1, \dots, \mathbf{D}^{k-1})$ , where  $\mathbf{D}^j = (D_1^j, \dots, D_n^j) \in \{A, R\}^n$  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  $\mu_i^k(\cdot | h^k)$ , 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  $\mu_i^k(v | h^k) > 0$  and  $\sigma_i^k(h^k, p_k; v) = D_i^k$ , then  $\mu_i^k(v | h^k)$  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  $u_S(\boldsymbol{\sigma} | h^k, \boldsymbol{\mu}^k)$  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  $\sigma'_S$  of the seller, we have that  $u_S(\boldsymbol{\sigma} | h^k, \boldsymbol{\mu}^k) \geq u_S(\sigma'_S, \boldsymbol{\sigma}_{-S} | h^k, \boldsymbol{\mu}^k)$ . Similarly if  $u_i(\boldsymbol{\sigma} | h^k, \boldsymbol{\mu}^k, p_k; v_i)$  is the expected utility of a buyer with value  $v_i$  offered price  $p_k$

under history  $h^k$  under beliefs  $\mu_{-i}^k(h^k)$  on other buyers' values,  $u_i(\sigma | h^k, \mu^k, p_k; v_i) \geq u_i(\sigma'_i, \sigma_{-i} | h^k, \mu^k; v_i)$  for every alternate strategy  $\sigma'_i$ .

*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  $\sigma$ , let  $h^k$  and  $h^{k'}$  be the histories generated by  $\sigma$  under  $\mathbf{v}$ . If  $\mu^k = \mu^{k'}$ , 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.

### 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, and for completeness present the equilibrium in Appendix A. Formally, we have:

**THEOREM 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 unproductive. 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, 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 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.*

We prove the theorem in Appendix B. Intuitively, we use the no-learning equilibrium to commit the seller to a strategy. One way to understand the space of PBE is in terms of pairs of attainable payoffs for the buyers and the seller. Theorem 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 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  $(1 - \delta)R$ .*

The proof is constructive: given the PBE attaining  $R$  and  $U$ , the mechanism designer may in essence simulate the PBE on the reported values of the sellers. In other words, PBE resemble single-shot mechanisms with stronger incentive constraints. Theorem 3.3 implies that the utility-revenue Pareto frontier for PBE cannot generally exceed that of the single-shot mechanism design problem. For one buyer, Theorem 3.3 implies that the folk theorem is tight - the utility and revenue guarantees are the best possible. Theorem 3.2 implies a troubling multiplicity of equilibria, all with very different outcomes for both the seller and the buyers. It implies that further study of PBE is not worthwhile without a manner of refining away equilibria. We provide such a selection tool in the next section.

#### 4 NON-ROBUSTNESS OF ONE-BUYER LEARNING EQUILIBRIA

We now consider the case of one buyer and one seller. In this setting, Theorem 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 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 4.1.* A perfect Bayesian threshold equilibrium  $\sigma$  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 4.2.** *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.*

Driving the proof of Theorem 4.2 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 4.3.* 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  $j \in \{1, \dots, \infty\}$ :

- 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. *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 4.5 ([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  $[a, b]$ :*

$$vf(v) - g(v) \geq vf(v') - g(v'). \quad (1)$$

*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.

PROOF OF LEMMA 4.4. 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 4.5. 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.  $\square$

We will only use monotonicity of allocations here. In Appendix C, we make heavier use of Lemma 4.4 to derive alternate sufficient conditions under which the conclusions of Theorem 4.2 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.

LEMMA 4.6. *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  $k_\epsilon$  be the earliest round such that all agents in  $[a, t - \epsilon)$  have rejected at least once. If it is the case that  $k_\epsilon \rightarrow \infty$  as  $\epsilon \rightarrow 0$ , 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  $1, \dots, M^* - 1$ , 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  $1, \dots, M^* - 1$  do not change after round  $M^*$ , as all agents who accepted in rounds  $1, \dots, M^* - 1$  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.  $\square$

PROOF OF THEOREM 4.2. 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 4.6 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.

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 equilibrium with  $X(t) = 1$  and  $X(t') > 1$  for some  $t' < t$ , contradicting Lemma 4.4.  $\square$

In Appendix C we give an alternate refinement which similarly eliminates learning equilibria. Theorem 4.2 and Appendix C 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 5.1 and 6, 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.

## 5 EQUILIBRIUM WITH TWO BUYERS

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  $\delta \geq 2/3$ . This equilibrium has two desirable properties: first, it survives the refinements proposed in Section 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. We describe the main ideas of the equilibrium in section 5.1 and leave the full formal description to Appendix D. In Section 5.2, we derive revenue guarantees for our equilibrium.

### 5.1 Equilibrium Description

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 agent. 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.

*Buyers.* In any given round of the exploration phase, neither buyer has rejected a price yet. The seller offers a new price, say,  $p$ , and the buyers, whose beliefs are distributed i.i.d. according to some posterior  $F$  supported on  $[a, b]$ , behave according to a threshold  $t$  solving the equation:

$$\left(\frac{1-F(t)}{2} + F(t)\right)(t-p) = \frac{\delta}{2(1-\delta)}(t-a)F(t). \quad (2)$$

The lefthand side represents the utility of a buyer with type  $t$  who accepts the price  $p$ , which is  $t-p$  times the probability of winning in the current round, with a continuation utility of zero. The righthand side represents the threshold buyer's utility from rejection - if the other buyer accepts, then the seller will only post prices above  $t$  in the subsequent game yielding zero continuation utility for the threshold buyer. If both buyers reject, then they share the item for the rest of the game at price  $a$ . In Appendix D, we show that as expected, higher-valued buyers will prefer to accept, and lower-valued buyers will reject.

In the exploitation phase, the seller targets the buyers with the strongest value distribution conditioned on past behavior, and prices at the bottom of their support. The buyer incentives in this phase of the game are similar to those of the no-learning equilibrium.

*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 (2), 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:

$$R(a, b, p) = (1 - F(t(p)))^2(p + \delta R(t(p), b)) + 2(F(t(p)) - F(t(p))^2)(p + \frac{t(p)\delta}{1-\delta}) + F(t(p))^2(\frac{a\delta}{1-\delta}), \quad (3)$$

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 (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.

## 5.2 Revenue Guarantees

We now argue that if distributions are well-behaved, the revenue of the equilibrium outlined in the previous section (and discussed in full detail in Appendix D) 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 3.3, this revenue is an upper bound on the seller’s revenue in any PBE. The result is the following:

**THEOREM 5.1.** *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 5.1 and Appendix D, the seller obtains revenue which is at least  $\frac{1}{3e^2}$  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. 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 5.2. *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  $p = t - \delta(t - a)$ .  $\square$

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

PROOF. By Lemma 5.2, 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:

$$\frac{t - p}{t - a} = \frac{F_a^b(t)}{F_a^b(t) + 1} \frac{\delta}{1 - \delta} \quad (4)$$

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.  $\square$

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 5.4. *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  $j$ th 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,  $p^*$ .  $\square$

LEMMA 5.5. *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 5.4 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  $\frac{1}{1-\delta} p^*$  revenue. Overall the revenue obtained by the seller with this price sequence is therefore at least  $\frac{1}{1-\delta} \frac{2}{3e} p^* (1 - F(p^*))^2$ .  $\square$

PROOF OF THEOREM 5.1. Let  $OPT$  denote the total revenue from running the optimal auction for two buyers on  $F$ . Our benchmark for revenue is  $OPT/(1 - \delta)$ . By concavity of the revenue we have that  $OPT \leq 2p^*(1 - F(p^*))$ . By Lemma 5.5 the equilibrium gets revenue

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

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  $\frac{OPT}{3(1-\delta)e^2} \geq \frac{OPT}{3e^2}$ , proving the theorem.  $\square$

## 6 EQUILIBRIUM WITH MANY BUYERS

In Section 5, we showed that the dynamic pricing game with  $\delta \geq 2/3$  supports an equilibrium with nontrivial revenue and learning. Moreover, in contrast to the one-buyer case, this equilibrium is robust to the refinements laid out in Section 4 and Appendix C. 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 and Appendix C. We derive revenue guarantees in Section 6.2 and Appendix F.

### 6.1 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. We give the full description of the equilibrium in Appendix E.

*Buyers.* In the exploration phase, the seller targets the set of buyers  $S$  who have not yet rejected a price. The not in  $S$ , i.e. those who have rejected, are ignored by the seller and priced out of the market. We may therefore continue our discussion with  $|S| = n$ . For the

sellers in  $S$ , the price offered by the seller induces a threshold response. For  $n$  buyers, the threshold equation is

$$P_n(t)(t - p) = \frac{\delta}{1 - \delta} \frac{F(t)^{n-1}}{n} (t - a), \quad (5)$$

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  $F$  is the common distribution of buyers in  $S$ , supported on  $[a, b]$ . The lefthand side again corresponds to the utility of a buyer with type  $t$  from accepting, which is made more complicated now by the fact that anywhere from 0 to  $n - 1$  other buyers could accept. Meanwhile, the righthand side is again the utility of the threshold agent from rejecting, which is only positive from the case where all other buyers also reject, in which case all buyers split the item in perpetuity. In the exploitation phase, the seller continues targeting either the set of agents who have never rejected, or the agents who only rejected in the previous round, if everyone rejected that round. As in the two-buyer equilibrium, the buyer incentives are similar to the no-learning equilibrium.

*Seller.* As in the two buyer equilibrium, the seller is faced with an algorithmic pricing problem, where each round they are restricted to prices which support threshold responses. To do so, they optimize the value function

$$\begin{aligned} R(a, b, p, n) = & F(t(p))^n \left( \frac{\delta a}{1 - \delta} \right) + n(1 - F(t(p))) F(t(p))^{n-1} \left( p + \frac{\delta t(p)}{1 - \delta} \right) \\ & + \sum_{j=2}^n \binom{n}{j} (1 - F(t(p)))^j F(t)^{n-j} (p + \delta R(t(p), b, j)), \end{aligned}$$

where  $t(p)$  is the corresponding threshold price solving (5) for  $p$  and  $R(t(p), b, j)$  is the expected continuation revenue from the  $j$ -buyer equilibrium with the belief distribution conditioned to  $[t(p), b]$ . Note that for  $n \geq 3$  buyers, the recurrence has an additional argument, which is the number of players who have not yet rejected. As before, this problem can be solved to arbitrary precision using discretization and value iteration. In Section 6.2, we extend the ideas of Section 5.2 to give a computationally efficient and simple strategy which secures for the seller a constant fraction of the maximum revenue possible while lowerbounding the revenue of the seller's best response.

## 6.2 Approximate Optimality of Multibuyer Equilibrium

In this section, we state our revenue guarantees, assuming buyers' value distributions are sufficiently well-behaved. As with  $n \geq 2$  buyers, the proof demonstrates a squence of prices which the seller could select in equilibrium, for which buyers would behave according to their threshold responses. This lowerbounds the revenue of the seller if they best respond. Intuitively, our sequence of prices learns from "accept" decisions of the buyers as aggressively as possible, seeking high thresholds until it is possible to offer the price with quantile  $1/n$  according to the original distribution. Our analysis shows that this sequence of prices quickly reaches this price, and that offering this price every round approximates the optimal revenue. Formally:

*Definition 6.1.* A distribution is *regular* if the Myerson virtual value function  $\phi(v) = v - \frac{1 - F(t)}{f(t)}$  is increasing in  $v$ .

**THEOREM 6.2.** *For all regular value distributions  $F$  with monopoly quantile at least  $1/n$ , the equilibrium described in Section 6 and Appendix E earns the seller at least a  $\frac{7}{27} (2/3)^{1 - \frac{\ln 3}{\ln(1-1/\sqrt{2})}} \approx .1202$ -fraction of the revenue of the revenue-optimal auction for  $F$  run each round.*

We reiterate that all MHR distributions are regular with monopoly quantile at least  $1/e$ . Hence, the requirements for Theorem 6.2 are weaker than for two buyers. The proof can be found in Appendix F.

## REFERENCES

- Saeed Alaei. 2011. Bayesian Combinatorial Auctions: Expanding Single Buyer Mechanisms to Many Buyers. In *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*. 512–521. DOI:<http://dx.doi.org/10.1109/FOCS.2011.90>
- Marco Battaglini. 2005. Long-term contracting with Markovian consumers. *The American economic review* 95, 3 (2005), 637–658.
- Yang Cai and Constantinos Daskalakis. 2011. Extreme-Value Theorems for Optimal Multidimensional Pricing. In *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*. 522–531. DOI:<http://dx.doi.org/10.1109/FOCS.2011.76>
- R. H. Coase. 1972. Durability and Monopoly. *The Journal of Law and Economics* 15, 1 (1972), 143–149. DOI:<http://dx.doi.org/10.1086/466731>
- Vincent Conitzer, Curtis R Taylor, and Liad Wagman. 2012. Hide and seek: Costly consumer privacy in a market with repeat purchases. *Marketing Science* 31, 2 (2012), 277–292.
- Nikhil R. Devanur, Yuval Peres, and Balasubramanian Sivan. 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*. 983–1002. DOI:<http://dx.doi.org/10.1137/1.9781611973730.67>
- Drew Fudenberg and Jean Tirole. 1983. Sequential bargaining with incomplete information. *The Review of Economic Studies* 50, 2 (1983), 221–247.
- Drew Fudenberg and J Miguel Villas-Boas. 2006. Behavior-based price discrimination and customer recognition. *Handbook on economics and information systems* 1 (2006), 377–436.
- Oliver D Hart and Jean Tirole. 1988. Contract renegotiation and Coasian dynamics. *The Review of Economic Studies* 55, 4 (1988), 509–540.
- John Kennan. 2001. Repeated bargaining with persistent private information. *The Review of Economic Studies* 68, 4 (2001), 719–755.
- Roger Myerson. 1981. Optimal Auction Design. *Mathematics of Operations Research* 6, 1 (1981), 58–73.
- Klaus M Schmidt. 1993. Commitment through incomplete information in a simple repeated bargaining game. *Journal of Economic Theory* 60 (1993), 114–114.

## A NO-LEARNING EQUILIBRIUM

In this appendix, we give the full description of the no-learning equilibrium. The seller’s strategy can be found in Algorithm 1, and the buyer’s strategy in Algorithm 2. The beliefs which support this strategy profile are simple: if a buyer has ever accepted a positive price or rejected 0 (neither of which is on-path), the seller believes the buyer’s value is 1. Otherwise,

the seller learns nothing about the buyer's value and offers the item for free every round. Moreover, it is clear from inspection that this equilibrium survives the simplicity refinement.

---

**ALGORITHM 1:** Zero-Revenue Equilibrium - Seller's Strategy
 

---

**Input** : Purchasing history  $h^k$ , belief support  $[a, b]$ .  
**Output** : Price  $p_{k+1}$

```

if Buyer has ever accepted a positive price then
  |  $p_{k+1} = 1$ ;
else if Buyer has ever rejected a price of 0 then
  |  $p_{k+1} = 1$ ;
else
  |  $p_{k+1} = 0$ ;

```

---



---

**ALGORITHM 2:** Zero-Revenue Equilibrium - Buyer's Strategy
 

---

**Input** : Purchasing history  $h^k$ , belief support  $[a, b]$ , value  $v$ , price  $p_{k+1}$   
**Output** : Purchasing decision for round  $k + 1$

```

if  $p_{k+1} = 0$  then
  | Accept;
else if  $p_{k+1} > 0$  then
  | if Buyer has ever accepted a positive price then
  | | Accept;
  | else if Buyer has ever rejected a price of 0 then
  | | Accept;
  | else
  | | Reject;

```

---

## B PROOF OF THEOREM 3.2

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 4, and the seller's strategy in Algorithm 3. 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.

---

**ALGORITHM 3:** Folk Theorem Equilibrium - Seller's Strategy
 

---

**Input** : History  $h^k$ , Support bounds  $(a_k^1, b_k^1), \dots, (a_k^n, b_k^n)$   
**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 4:** Folk Theorem Equilibrium - Buyer  $i$ 's Strategy
 

---

**Input** : History  $h^k$ , Support bounds  $(a_k^1, b_k^1), \dots, (a_k^n, b_k^n)$ , 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;

---

To see that the strategies in Algorithms 3 and 4 are a PBE, we first argue that the seller is best responding. We consider the cases as they are stated in Algorithm 3:

- *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 4.

- *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 positive 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.

## C 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, 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 C.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 C.2.** *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 C.2 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.

**PROOF OF THEOREM C.2.** 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 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 thresholds 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. 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.  $\square$

## D FULL DESCRIPTION OF TWO-BUYER EQUILIBRIUM

In this appendix, 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. 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 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:

$$(t - p) \left( F_a^b(t) + \frac{1 - F_a^b(t)}{2} \right) = \frac{t - a}{2} F_a^b(t) \frac{\delta}{1 - \delta}. \quad (6)$$

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 (6). 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 (6), we will assume buyers use the highest value. Formally, define  $T_a^b(p)$  to be the set of solutions to (6). 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:

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

where  $R_{t_a^b(p)}^b$  is the seller’s optimal revenue from the continuation game where the buyers are both distributed according to  $F_{t_a^b(p)}^b$ . 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  $F_{t_a^b}^b$ . 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 5, and the buyers’ strategy in Algorithm 6. 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 5:** 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**  $a \geq p^*$  **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 6:** 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**  $a \geq p^*$  **then**  
        Accept if and only if  $p_k \leq a$   
    **else**  
        Accept if and only if  $v_i \geq t_a^b(p)$   
**else**  
    **if**  $p_i \leq \max\{a_1^k, a_2^k\}$  **then**  
        Accept if and only if  $v_i \geq p_{k+1}$   
    **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 7. Together with Algorithms 5 and 6, this fully specifies equilibrium.

---

**ALGORITHM 7:** 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  $i \in \{1, 2\}$  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 D.1.** *The strategies and beliefs specified by Algorithms 5, 6, and 7 is a PBE for  $\delta \geq 2/3$ .*

In Section D.1, we show that the buyers are best-responding, and in Section D.2, we show that the seller is best-responding. Together, these results prove the theorem.

### D.1 Buyer Incentives

As discussed, the belief updates of Algorithm 7 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 D.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 7, 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  $b$ .  $\square$

We now argue that the buyers are in equilibrium.

**LEMMA D.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 D.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 (6). 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

$$U_A \geq \left(F(t) + \frac{1-F(t)}{2}\right)(v_1 - p) + \frac{\delta}{1-\delta}(v_1 - t)F(t). \quad (7)$$

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 D.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:

$$U_R \leq F(t)\frac{v_1 - a}{2} \cdot \frac{\delta}{1-\delta} + \delta(1 - F(t))\frac{v_1 - t}{2}. \quad (8)$$

Write  $x = v_1 - t$ . By assumption,  $x \geq 0$ . By subtracting expression (7) from expression (8) 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 (6) to get:

$$(F(t) + 1)(t - p) = F(t)(t - a)\frac{\delta}{1 - \delta}. \quad (9)$$

Applying equation (9) to the first two terms of expression (9) allows us to write the as a product of  $x$  and another term which is clearly positive:

$$U_A - U_R \geq x\left(\frac{1 - F(t)}{2} - \frac{F(t)}{2}\frac{\delta}{1 - \delta} + \frac{\delta F(t)}{1 - \delta} - \frac{\delta(1 - F(t))}{2}\right). \quad (10)$$

Finally, note that when  $p < a$ , the corresponding threshold from (9) 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 D.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:

$$U_A = \frac{1 + F(t)}{2}(v_1 - p). \quad (11)$$

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:

$$U_R = \frac{F(t)}{2} \frac{\delta(v_1 - a)}{1 - \delta}. \quad (12)$$

Combining equations (11) and (12) 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 (9), we have:

$$U_R - U_A = x \left( \frac{1 + F(t)}{2} - \frac{F(t)}{2} \frac{\delta}{1 - \delta} \right) \quad (13)$$

Note that  $v_1 < t$  only if  $p > a$ . In this case,  $t - p < t - a$ , and hence equation (9) implies that the term in parentheses in (13) is positive. It follows that agents with value below  $t$  will reject.

Assume (6) 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 D.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 (6) has no solution, it must be that the lefthand side of the rearranged threshold equation (9) 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 (9) has no solution, it must be that

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

Rearranging yields:

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

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 D.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 D.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.  $\square$

## D.2 Seller Incentives

LEMMA D.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 (6), or post a price  $p$  such that no solution to (6) 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  $\delta$ . 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.  $\square$

## E EQUILIBRIUM FOR THREE OR MORE BUYERS

In this appendix, we provide a full description of our equilibrium for three or more buyers. The full formal descriptions of the seller's strategy, the buyers' strategies, and the belief updates are given in Algorithms 8, 9, and 10.

The seller's strategy has two phases: an explore phase and an exploit phase. In the explore phase, the seller offers the revenue  $p$  each round such that there exists a threshold  $t$  satisfying the threshold equation

$$P_n(t)(t - p) = \frac{\delta}{1 - \delta} \frac{F(t)^{n-1}}{n} (t - a), \quad (14)$$

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 (6) 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:

$$R(a, b, p, n) = F(t_a^b(p; n))^n \left( \frac{\delta a}{1 - \delta} \right) + n(1 - F(t_a^b(p; n))) F(t_a^b(p; n))^{n-1} \left( p + \frac{\delta t_a^b(p; n)}{1 - \delta} \right) \\ + \sum_{j=2}^n \binom{n}{j} (1 - F(t_a^b(p; n)))^j F(t_a^b(p; n))^{n-j} (p + \delta R(t_a^b(p; n), b, j)),$$

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  $F_{t_a^b(p)}^b$ .

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 8:** Seller's Strategy
 

---

**Input** : Purchasing history  $h^k$ , Support bounds  $\{(a_i^k, b_i^k)\}_{i=1}^n$   
**Output** : Price  $p_{k+1}$   
 $S_k = \operatorname{argmax}_i a_i^k$ ;  
 $a = \max_i a_i^k$ ;  
 $b = \max_i b_i^k$ ;  
**if**  $D^k = R^n$  **or**  $|S_k| = 1$  **or**  $a \geq p^*$  **then**  
 |  $p_{k+1} = a$ ;  
**else**  
 |  $p_{k+1} = \operatorname{arg max}_p R(a, b, p, |S_k|)$  ;

---



---

**ALGORITHM 9:** 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  $k$   
 $S_k = \operatorname{argmax}_j a_j^k$ ;  
 $a = \max_j a_j^k$ ;  
 $b = \max_j b_j^k$ ;  
**if**  $D^k = R^n$  **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 10:** 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  $\{(a_i^{k+1}, b_i^{k+1})\}_{i=1}^n$ .  
 $S_k = \operatorname{argmax}_j a_j^k$ ;  
**for**  $i \in \{1, \dots, n\}$  **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**  
 | | |  $a_i^{k+1} = a_i^k, b_i^{k+1} = t_{a_i^k}^{b_i^k}(p_k, |S_k|)$

---

**THEOREM E.1.** *The strategies and belief updates laid out in Algorithms 8, 9, and 10 is a PBE as long as  $\delta \geq \frac{n}{n+1}$ .*

We prove the theorem in much the same manner as Theorem D.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 E.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  $j$ th in the order is  $1/n$ , and conditioned on this event, their probability of winning is  $F(t)^{j-1}$ .  $\square$

## E.1 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 E.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 optimal anonymous price  $p^*$  for the initial distribution, the seller will only post  $b$  in subsequent rounds. No buyer can get positive utility from such prices.  $\square$

**LEMMA E.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 (14) implies that

$$P_n(t) \geq \frac{\delta}{1-\delta} \frac{F(t)^{n-1}}{n}. \quad (15)$$

Plugging  $t = v_1 + x$  for  $x \geq 0$  into equation (14) and applying inequality (15) 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 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  $\delta P_n(t)(v_1 - t)$ . 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:

$$\begin{aligned}
u_A - u_R &\geq (t + x - p)P_n(t) + \frac{\delta}{1-\delta} \frac{F(t)^{n-1}}{n} (t + x - t) \\
&\quad - \delta P_n(t)(t + x - t) - \frac{\delta}{1-\delta} \frac{F(t)^{n-1}}{n} (t + x - a) \\
&\quad = P_n(t)(t - p) - \frac{\delta}{1-\delta} \frac{F(t)^{n-1}}{n} (t - a) \\
&\quad + x \left[ P_n(t) + \frac{\delta}{1-\delta} \frac{F(t)^{n-1}}{n} - \delta P_n(t) - \frac{\delta}{1-\delta} \frac{F(t)^{n-1}}{n} \right] \\
&= x \left[ P_n(t) + \frac{\delta}{1-\delta} \frac{F(t)^{n-1}}{n} - \delta P_n(t) - \frac{\delta}{1-\delta} \frac{F(t)^{n-1}}{n} \right] \quad [\text{By (14)}]
\end{aligned} \tag{16}$$

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 E.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  $U_R = \frac{\delta}{1-\delta} \frac{v_1 - a}{n}$ .

Because there is no threshold response, it must be that the threshold equation (14) has no solution. The argument mirrors the two-buyer case: for  $t < p$ , the lefthand side of (14) 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 E.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 D.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 D.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.  $\square$

## E.2 Seller Incentives

LEMMA E.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 (14) for  $n = S_k$ , or post a price  $p$  such that no solution to (14) 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.  $\square$

## F PROOF OF THEOREM 6.2

We now give our pricing algorithm: let  $F$  be the common distribution of all sellers who have not yet rejected. Further let  $p^*$  be the price such that  $F(p^*) = 1 - 1/n$ .

- If there is a price  $p$  with threshold  $t(p) = p^*$ , post  $p$ .
- Otherwise, post a price  $p$  which induces threshold  $t(p) = n^{-1}\sqrt{\frac{1-\delta}{\delta}}$ .

LEMMA F.1. *For any continuation where at least two buyers have accepted every round, let  $F$  be the distribution of the common beliefs of the buyers who have not rejected. Let  $t^*$  be the unique value satisfying  $F(t^*) = n^{-1}\sqrt{\frac{1-\delta}{\delta}}$ . There is a price  $p$  solving the threshold equation for  $t = t^*$ .*

PROOF. We may rearrange the threshold equation as

$$nP_n(t) \frac{t-p}{t-a} = F(t)^{n-1} \frac{\delta}{1-\delta}. \quad (17)$$

For  $t = t^*$ , the righthand side has value 1. We must simply show that there is a value of  $p \in [a, b]$  satisfying  $nP(t^*) \frac{t^*-p}{t^*-a} = 1$ . Note that as a function of  $p$ ,  $nP(t^*) \frac{t^*-p}{t^*-a}$  ranges continuously from  $nP(t^*)$  at  $p = a$  to 0 at  $p = t^*$ . From Lemma E.2, it is obvious that  $nP(t^*) > 1$ . Hence, there must be a price satisfying (17).  $\square$

LEMMA F.2. *If at least two buyers have value at least  $p^*$ , then the above sequence of prices eventually induces threshold  $p^*$ . In particular, let  $\Delta = n^{-1} \sqrt{\frac{1-\delta}{\delta}}$ , and let  $k$  solve  $1 - (1 - \Delta)^k = F(p^*)$ . Then at most  $k + 1$  steps are required before the above price sequence begins posting price  $p^*$ .*

PROOF. Assume there are at least two buyers with value at least  $p^*$ . Lemma F.1 guarantees that until the seller offers a price with threshold  $p^*$ , round  $j$ 's threshold  $t_j$  will satisfy  $F(t_j) = 1 - (1 - \Delta)^j$ , where  $F$  is the initial distribution of values. We show that as soon as  $1 - (1 - \Delta)^j \geq F(p^*)$ , i.e.  $t_j \geq p^*$ , the seller will post  $p^*$  in the subsequent round.

Consider a round where the common beliefs of the buyers who have not yet rejected are distributed according to  $F$ . Let  $t^*$  be the threshold satisfying  $F(t^*) = \Delta$ , and assume that  $t^* > p^*$ . We show that in this round, there will be a price  $p$  and corresponding threshold  $t$  such that  $t = p^*$ . Note that the threshold equation can be rewritten as

$$\frac{t-p}{t-a} = \frac{F(t)^{n-1}}{nP_n(t)} \frac{\delta}{1-\delta}. \quad (18)$$

Note that by applying Lemma E.2, we see that the righthand side is equal to

$$\frac{F(t)^{n-1}}{\sum_{j=0}^{n-1} F(t)^j} \frac{\delta}{1-\delta},$$

which is increasing in  $t$ . Since  $t^* > p^*$ , we have

$$0 < \frac{F(p^*)^{n-1}}{\sum_{j=0}^{n-1} F(p^*)^j} \frac{\delta}{1-\delta} < \frac{F(t^*)^{n-1}}{\sum_{j=0}^{n-1} F(t^*)^j} \frac{\delta}{1-\delta} < 1.$$

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.  $\square$

LEMMA F.3. *For  $\delta \geq \frac{n}{n+1}$ , as long as there are two or more buyers with value at least  $p^*$ , the discount factor in the round where the seller first posts  $p^*$  will be at most  $(2/3)^{1 - \frac{\ln 3}{\ln(1-1/\sqrt{2})}} \approx .4638$ .*

PROOF. By Lemma F.2, it will take at most  $k + 1$  stages for the seller to post  $p^*$  for the first time, where  $k$  solves  $1 - (1 - \Delta)^k = F(p^*) = 1 - 1/n$ . In other words, the discount factor in this round will be at least  $\gamma(\delta, n) = \delta^{1 + \log_{1-\Delta} \frac{1}{n}}$ .

We first show that  $\gamma(\delta, n)$  is increasing in  $\delta$  for all  $n$ . The derivative of  $\gamma(\delta, n)$  with respect to  $\delta$  is

$$\delta^{1 + \frac{\ln \frac{1}{n}}{\ln \left(1 - \left(\frac{1-\delta}{\delta}\right)^{\frac{1}{n-1}}\right)}} \left( \frac{1 + \frac{\ln \frac{1}{n}}{\ln \left(1 - \left(\frac{1-\delta}{\delta}\right)^{\frac{1}{n-1}}\right)}}{\delta} + \frac{-\left(\frac{1}{\delta} + \frac{1-\delta}{\delta^2}\right) \left(\frac{1-\delta}{\delta}\right)^{\frac{1}{n-1}-1} \ln \delta \ln \frac{1}{n}}{\left(1 - \left(\frac{1-\delta}{\delta}\right)^{\frac{1}{n-1}}\right) (n-1) \ln \left(1 - \left(\frac{1-\delta}{\delta}\right)^{\frac{1}{n-1}}\right)} \right).$$



First, since

$$\delta^{1 + \frac{\ln \frac{1}{n}}{\ln \left(1 - \left(\frac{1-\delta}{\delta}\right)^{\frac{1}{n-1}}\right)}} > 0,$$

we need only analyze the quantity in the outermost parentheses. We first show that

$$\frac{1 + \frac{\ln \frac{1}{n}}{\ln \left(1 - \left(\frac{1-\delta}{\delta}\right)^{\frac{1}{n-1}}\right)}}{\delta} \geq 0.$$

This follows from noting that for  $\delta \geq 2/3$  and  $n \geq 3$ , the arguments of both logarithms are less than 1, and hence both logarithms are negative, causing their signs to cancel. What remains is obviously positive. We next show

$$\frac{-\left(\frac{1}{\delta} + \frac{1-\delta}{\delta^2}\right) \left(\frac{1-\delta}{\delta}\right)^{\frac{1}{n-1}-1} \ln \delta \ln \frac{1}{n}}{\left(1 - \left(\frac{1-\delta}{\delta}\right)^{\frac{1}{n-1}}\right) (n-1) \ln \left(1 - \left(\frac{1-\delta}{\delta}\right)^{\frac{1}{n-1}}\right)} \geq 0.$$

To do so, we note the sign of each of the factors in the numerator and denominator:

- $-\left(\frac{1}{\delta} + \frac{1-\delta}{\delta^2}\right)$ : negative
- $\left(\frac{1-\delta}{\delta}\right)^{\frac{1}{n-1}-1}$ : positive
- $\ln \delta$ : negative
- $\ln \frac{1}{n}$ : negative
- $1 - \left(\frac{1-\delta}{\delta}\right)^{\frac{1}{n-1}}$ : positive
- $n-1$ : positive
- $\ln \left(1 - \left(\frac{1-\delta}{\delta}\right)^{\frac{1}{n-1}}\right)$ : negative.

Taking these signs into account yields that the term in question is positive, proving that  $\gamma(\delta, n)$  is increasing in  $\delta$ . It therefore suffices to analyze  $\gamma\left(\frac{n}{n+1}, n\right)$ .

We next show that  $\gamma\left(\frac{n}{n+1}, n\right)$  is increasing in  $n$ . Hence, the worst case will be when  $n = 3$ . To do so, treat  $n$  as a continuous variable. We have

$$\frac{d}{dn} \gamma\left(\frac{n}{n+1}, n\right) = \frac{1}{n} \left(\frac{n}{n+1}\right)^{1 + \frac{\ln \frac{1}{n}}{\ln \left(1 - \left(\frac{1}{n}\right)^{\frac{1}{n-1}}\right)}} \times \left( \frac{1 + \frac{\ln \frac{1}{n}}{\ln \left(1 - \left(\frac{1}{n}\right)^{\frac{1}{n-1}}\right)}}{n+1} + \frac{\alpha_n}{\left(\left(\frac{1}{n}\right)^{\frac{1}{n-1}} - 1\right) (n-1)^2 \ln^2 \left(\left(\frac{1}{n}\right)^{\frac{1}{n-1}} - 1\right)} \right),$$

Where

$$\alpha_n = \left(1 - \left(\frac{1}{n}\right)^{\frac{1}{n-1}}\right) (n-1)^2 \ln \left(1 - \left(\frac{1}{n}\right)^{\frac{1}{n-1}}\right) + \left(\frac{1}{n}\right)^{\frac{1}{n-1}} \ln \frac{1}{n} \left(n \ln \frac{1}{n} + n - 1\right) \ln \frac{n}{n+1}.$$

As with the derivative of  $\gamma(\delta, n)$  with respect to  $\delta$ , one may check term-by-term that the quantities above come out positive. We omit the analysis for brevity.

It follows that  $\gamma(\delta, n) \geq (2/3)^{1 - \frac{\ln 3}{\ln(1-1/\sqrt{2})}} \approx .4638$ .  $\square$

LEMMA F.4. For  $n \geq 3$ , the probability that at least two buyers have value  $p^*$  is at least  $7/27$ .

PROOF. The probability that two or more buyers have value at least  $p^*$  is

$$1 - \left[ \left(1 - \frac{1}{n}\right)^n + \left(1 - \frac{1}{n}\right)^{n-1} \right]. \quad (19)$$

This function is increasing in  $n$ , and hence we may take the worst case, where  $n = 3$ , with value  $7/27$ .  $\square$

PROOF OF THEOREM 6.2. A standard upper bound on the revenue of a single-item mechanism is the ex ante relaxation benchmark, the revenue of the optimal mechanism for the unlimited supply setting which sells with probability at most 1 (see [Alaei, 2011, Cai and Daskalakis, 2011]). For regular distributions with monopoly quantile at least  $1/n$ , this benchmark has value exactly  $p^*$  for each round, or  $\frac{p^*}{1-\delta}$  in total. Combining Lemmas F.2 and F.4 yields that the expected revenue from the exploit phase of the price sequence we give is at least  $\frac{7}{27}(2/3)^{1-\frac{\ln 3}{\ln(1-1/\sqrt{2})}} p^*$ .  $\square$