

Simple Pricing Schemes For Consumers With Evolving Values

Shuchi Chawla^{*} Nikhil R. Devanur[†] Anna R. Karlin[‡] Balasubramanian Sivan[§]

Abstract

We consider a pricing problem where a buyer is interested in purchasing/using a good, such as an app or music or software, repeatedly over time. The consumer discovers his value for the good only as he uses it, and the value evolves with each use. Optimizing for the seller’s revenue in such dynamic settings is a complex problem and requires assumptions about how the buyer behaves before learning his future value(s), and in particular, how he reacts to risk. We explore the performance of a class of pricing mechanisms that are extremely simple for both the buyer and the seller to use: the buyer reacts to prices myopically without worrying about how his value evolves in the future; the seller needs to optimize for revenue over a space of only two parameters, and can do so without knowing the buyer’s risk profile or fine details of the value evolution process. We present simple-versus-optimal type results, namely that under certain assumptions, simple pricing mechanisms of the above form are approximately optimal *regardless of the buyer’s risk profile*.

Our results assume that the buyer’s value per usage evolves as a martingale. For our main result, we consider pricing mechanisms in which the seller offers the product for free for a certain number of uses, and then charges an appropriate fixed price per usage. We assume that the buyer responds by buying the product for as long as his value exceeds the fixed price. Importantly, the buyer does not need to know anything about how his future value will evolve, only how much he wants to use the product *right now*. Regardless of the buyers’ initial value, our pricing captures as revenue a constant fraction of the total value that the buyers accumulate in expectation over time.

1 Introduction

A common assumption in auction theory and mechanism design is that a consumer knows how much he values an item at the moment he is considering a purchase and that he has an accurate prior over how much he will value that item in the future (if not precise knowledge). The consumer is then modeled as a utility maximizer, where the utility is typically quasi-linear — equal to the consumer’s valuation minus payments made by him.

In this paper, we explore scenarios where the consumer discovers his valuation for a good as he *uses* it. For example, consider a consumer buying a video game. The consumer’s valuation of the game depends on how much pleasure he derives from playing and how many times he ends up wanting to play it. The consumer does not know these quantities ahead of time, and only discovers them as he repeatedly plays the game. Indeed the consumer’s enjoyment of the game, rather than being constant over time, may evolve as he plays more and more. At some point of time, the consumer may tire of playing the game altogether and stop using it. The principle of value evolution applies to a wide range of products, including apps, songs, and cable TV.

We model such scenarios by having a value per usage. We use V_t to denote the value to the consumer for the $t + 1^{\text{st}}$ usage. The consumer discovers V_t only when he is considering using the good for $t + 1^{\text{st}}$ time. Note that t denotes usage and not real time: if the consumer does not use the product his valuation for the next usage doesn’t change. The initial value V_0 is drawn from an arbitrary distribution with bounded support.

How should such a good be priced? The prevalent mechanism for selling songs, apps, software, and other digital goods is to offer a one-time “contract” to the consumer for unlimited usage of the product. We call such a mechanism a “Buy-It-Now” (BIN) scheme. BIN pricing is straightforward in our value evolution model, as long as the seller knows the distribution of initial values and the stochastic process governing the value evolution. Specifically, if the buyer is risk-neutral, the pricing problem can be reduced to the

^{*}University of Wisconsin-Madison. shuchi@cs.wisc.edu. Supported in part by NSF grants CCF-1101429 and CCF-1320854.

[†]Microsoft Research. nikdev@microsoft.com.

[‡]University of Washington. karlin@cs.washington.edu. Supported by NSF CCF-1420381.

[§]Microsoft Research. balu2901@gmail.com.

standard problem of pricing an item in a one-shot game: the seller computes the buyer’s *cumulative value*,¹ as a function of V_0 , defined as $\mathcal{C}(v) := \mathbf{E}[\sum_{t \geq 0} V_t | V_0 = v]$, and then chooses the price p that maximizes²

$$p \mathbf{P}[\mathcal{C}(v) \geq p].$$

Alternatively, one could consider a dynamic mechanism, that interacts with the buyer and *adjusts prices with each usage*. Optimizing for the seller’s revenue in the dynamic setting can be quite complex, and depends on the fine details of the value evolution process, and the buyer’s risk profile. For certain types of value evolution processes, optimal mechanisms have been derived, e.g. by Kakade et al. [13] and Pavan et al. [16]. (See Section 2 for a more detailed discussion of related work.) However, these results are somewhat complex, assume that the buyer is risk-neutral, depend on the assumption that both the buyer and seller know the precise stochastic process governing the buyer’s value evolution, and in some cases require the buyer to solve complex MDPs. (The buyer needs to know the evolution process because these mechanisms guarantee only interim individual rationality.)

In this paper, we take a different point of view. We focus on the performance of only the simplest mechanisms: simple for the seller to find, simple for the seller to implement, simple for the buyer to understand, and simple for the buyer to optimize over. In particular, we consider schemes in which the seller sets a price p_t for the $(t + 1)$ -st usage, and the buyer knows the sequence $\{p_t\}_{t=0}^\infty$ up front. We call such a mechanism a “Pay-Per-Play” (PPP) scheme. We restrict our attention to a particularly simple class of PPP mechanisms where the seller offers a free trial to the buyer for the first T uses, and then charges a constant price of c per usage thereafter, that is, $p_t = 0$ for $t < T$, and $p_t = c$ for $t \geq T$.

Faced with a PPP scheme, a buyer buys the item for at least as long as his value exceeds the current price: $V_t \geq p_t$; Let $T^* = (\min\{t : V_t < p_t\} - 1)$ denote the length of this time period. During this period the buyer’s utility steadily increases. When his value dips below the current price, the buyer’s purchase behavior depends on his beliefs about his future value for the item and his risk profile. In order to obtain a revenue guarantee that is robust against buyer behavior, in accounting for the revenue of the seller,

we disregard any revenue accrued beyond time T^* . Specifically, we use the term **risk-robust revenue**, defined as $\sum_{t=0}^{T^*} p_t$, to denote a lower bound on the seller’s revenue regardless of the buyer’s risk profile.

Can a PPP scheme obtain good risk-robust revenue? To get a feel for the problem, consider the simplest possible PPP scheme: a constant price $p_t = p$ per usage for all t . Suppose, for example, that the buyer’s initial value is $v \in (0, 1)$, and evolves according to a simple symmetric random walk³ on the values $\{i\delta\}_{i=0}^{1/\delta}$ with reflection at 1 and absorption at 0. Simple calculations⁴ show that $\mathcal{C}(v) = \Omega(v/\delta^2)$ whereas no constant price PPP scheme can get risk-robust revenue more than $O(v^2/\delta^2)$. The latter can be much smaller than $\mathcal{C}(v)$ if v is small. (Recall that $v \in (0, 1)$.)

Of course, the astute reader has observed that this is an apples and oranges comparison: we are comparing the revenue of a seller facing, essentially, an infinitely risk-averse buyer in the PPP case to the expected cumulative value of the buyer. What hope is there for PPP to be constant-competitive for this type of comparison?

1.1 Formal model and results

We consider the following general model of value evolution. Our main result shows that a very simple pricing scheme guarantees the seller an expected risk-robust revenue equal to a constant fraction of the buyer’s expected cumulative value for any sequence V_t under this model.

The bounded martingale model:

We assume that the evolution of V_t satisfies the following properties.

1. $V_0 \in (0, 1)$ and $V_t \in [0, 1]$ for $t > 0$.
2. 0 is an absorbing state: $V_t = 0$ implies $V_{t+1} = 0$. In other words, the buyer loses interest in the product and does not want to continue purchasing it at any price.
3. **(Martingale.)** If $V_t \in (0, 1)$, then the value evolution satisfies the martingale property, i.e., $\mathbf{E}[V_{t+1} | V_0, V_1, \dots, V_t] = V_t$. We use Δ_t to denote the difference $V_{t+1} - V_t$.

¹One could also consider a discounted model where $\mathcal{C}(v) := \mathbf{E}[\sum_{t \geq 0} \delta^t V_t | V_0 = v]$ for some discount parameter $\delta > 0$.

²This assumes that $\mathcal{C}(v)$ is weakly increasing with the initial value v .

³For $0 < i < 1/\delta$, $V_{t+1} = V_t + \delta$ with probability 1/2 and $V_t - \delta$ with probability 1/2.

⁴See Section 3

4. **(Bounded step size.)** There is a sufficiently small constant $\epsilon > 0$ which is an upper bound on $|\Delta_t|$ for every $t \geq 0$.
5. **(Minimum variance.)** The second moment of Δ_t is bounded from below everywhere except when the value reaches 0. Thus, there is $\delta \in (0, \epsilon)$ such that $V_t \in (0, 1]$ implies

$$(1.1) \quad \mathbf{E} \left[\Delta_t^2 \mid V_0, \Delta_0, \dots, \Delta_{t-1} \right] \geq \delta^2.$$

The simple symmetric random walk mentioned earlier is an example of a bounded martingale process. As another example, consider a buyer repeatedly playing a video game. With each usage, the buyer has the potential to discover new features, improve his skill level, or be exposed to new competitors. Suppose that each of these adds a small random shock of mean zero to his value, but this shock could depend on the entire history of game play up to that time including his current valuation. Such a process fits into our model.

Let $\mathcal{R}_{T,c}$ denote the expected risk-robust revenue of the PPP mechanism with a free trial period of T uses and a per-play price of c thereafter, faced with a buyer whose value evolves according to the bounded martingale model. Let $\mathcal{R}_{T,c|V_0=v}$ denote the seller's expected risk-robust revenue when the buyer's initial value is v . Our main result is as follows.

THEOREM 1.1. *Suppose that ϵ^2/δ is a sufficiently small constant. Then there exist constants T and c , depending only on ϵ and δ but independent of any further details of the value evolution process or its initial value $v = V_0$, such that*

$$\forall v \in [0, 1], \quad \mathcal{R}_{T,c|V_0=v} = \Omega \left((\delta/\epsilon)^3 \mathcal{C}(v) \right).$$

Thus if $\delta = \Omega(\epsilon)$, the free trial scheme obtains revenue which is a constant fraction of $\mathcal{C}(v)$.

For the special case of value evolution processes with independent increments Δ_t , the constant in the approximation ratio to $\mathcal{C}(v)$ can be improved, and the requirement that ϵ^2/δ is a sufficiently small constant can be removed. However, this comes at the cost of requiring the seller to know the variance lower bounds δ_t^2 for each t . For the details, see Theorem B.1 in the appendix.

1.2 Some remarks on the model

As in any single buyer mechanism design setting, it is instructive to think of the buyer as being drawn

randomly from a heterogenous population. In this context, our model possesses some nice properties and limitations that we now discuss.

- While we require every buyer in the population to satisfy the stated properties, our model allows each buyer in the population to have a different value evolution process as well as a different starting value. Indeed, the change in a buyer's value over time can depend on his entire history of value evolution.
- Our model requires every buyer to have the same upper bound on value (normalized to 1), however, our results are robust to variations in this upper bound. If the seller prices the product with respect to the smallest value upper bound, our approximation factor worsens by the square of the ratio of the largest possible upper bound to the smallest possible upper bound.
- A limitation of our model is that the Assumptions 3-5 along with the value upper bound at 1 impose significant restrictions on what the value evolution process can look like when the value is very close to 1. (For values at most $1 - 2\epsilon$, Assumptions 3-5 are not affected by the upper bound at all.)

Relaxing these assumptions, particularly when the value is close to the upper bound, is an interesting direction for future work. For digital goods it is quite reasonable to assume that per-usage values are bounded. However, another interesting direction is to remove the value upper bound altogether and to analyze the time-discounted revenue⁵ of the seller. Finally, a very compelling direction for future work is to allow for super- or sub-martingale processes for value evolution.

1.3 Some remarks on the main theorem

- We view our main theorem as a “simple-versus-optimal” type result, namely that in our model, simple pricing mechanisms of the above form are approximately optimal regardless of the buyer's risk profile.
- As is apparent from the theorem statement, the buyer need not know anything about the process governing the evolution of his own value V_t .

⁵Time-discounting accounts for the fact that in the absence of an upper bound the buyer's cumulative value is unbounded.

Thus, our revenue bounds hold even under the strongest definition of risk aversion for the buyer.

- Most positive results in mechanism design rely on knowledge of the distribution of buyers’ types. In contrast, our model is almost *detail-free*, that is, it requires very little distributional information; It suffices for the seller to possess (conservative) estimates of the parameters δ and ϵ satisfied by the value evolution processes of each buyer in the population. Moreover, in order to obtain the revenue guarantee, the seller needs to optimize over a space of only two parameters (T and c). The downside of detail-freeness is that the seller does not exploit information that might lead to considerable improvement in revenue.
- In this value evolution model, even against a risk-neutral buyer, the optimal BIN scheme cannot always get an expected revenue which is a constant fraction of $\mathbf{E}[\mathcal{C}(V_0)]$. Intuitively, one reason that our free-trial pricing scheme is able to do so well in comparison is that it price discriminates between buyers: those that retain their interest in the product for longer pay more than those who lose interest quickly. It is worth noting, however, that this price discrimination is envy-free⁶ and therefore, fair.
- Another reason that our free-trial PPP can beat out BIN in terms of revenue is that the mitigation of buyer risk yields a larger pool of buyers. Indeed, both the seller *and the buyers* benefit from PPP in the following sense: while a BIN pricing immediately excludes low-initial-value buyers from using the product, in the free-trial PPP scheme, every buyer, regardless of his initial value, gets to explore the product for some amount of time; in the event that the buyer’s value climbs up during the free-trial, both the buyer and seller obtain greater utility from continued usage. This is particularly important in competitive markets where many sellers sell identical/similar products and marketshare becomes an important consideration. We formalize this intuition in the context of another natural model of value evolution that we call the binary value model (see appendix Section A).
- In the special case where the buyer’s value follows a simple random walk, coming up with a

good pricing is not too difficult: depending on the buyer’s starting value, it is possible to obtain fairly tight bounds on how long the random walk stays alive, and how much value it accumulates. Our model, however, allows for much more general processes: the change in value at a step t may depend in a complex way on the entire history of value evolution. Consequently, all the random variables of interest, such as how large the value is at any given point of time, the amount of time that the value remains positive, or the total amount of value accumulated by the buyer, can be poorly behaved, especially when subject to conditioning of any kind.

Moreover, our goal is to obtain a *single* PPP pricing that obtains a *point-wise* guarantee for revenue regardless of the buyer’s initial value: for every possible starting value of a buyer, we obtain a constant fraction of the total value he will accumulate in the future in expectation. This is especially challenging when the buyer’s starting value is very low⁷. In that case, there is a high probability that the value will quickly decrease to zero. However, there is also a small probability that the value becomes very large, and our PPP scheme must exploit this event in order to remain competitive. We need to simultaneously avoid having the buyer drop out in the initial “ramp-up” phase, as well as charge a high enough price in the following “sustained-high-value” phase, without sufficient information on how long each of these phases last.

- Finally, in Section 3 we show that if we relax any of the modeling Assumptions 3-5, then there is no PPP scheme (with *any* sequence $\{p_t\}$) that yields the seller an expected revenue of $\Omega(\mathbf{E}[\mathcal{C}(V_0)])$.

Simulations. Owing to the generality of our model, our approximation factors are quite large. In Appendix C, we present the results of simulations, experimentally evaluating the gap between PPP revenue and cumulative value for two models of value evolution: (a) a simple random walk model; (b) the binary value model, namely, where the value remains equal to the initial value for a random number of steps, and then drops to zero. In the former setting, PPP obtains at least 17% of the buyer’s cumulative

⁶meaning that no buyer prefers the price and allocation received by another buyer to their own

⁷For example, even if we knew the initial value, in this case, charging a constant price equal to initial value does not guarantee a constant fraction of cumulative value.

value regardless of the initial value in all of our experiments. In the latter setting, we show that PPP obtains more revenue for the seller as well as generates more utility for the buyer, compared to the optimal BIN scheme.

2 Related Work

Our paper is closely related to the literature on dynamic mechanism design. Initiated by the work of Baron and Besanko [4], there have been a number of papers [7, 11, 5, 6] that consider mechanism design where the private information of the agents evolves over time. Eso and Szentes [12] consider a setting where the seller controls information that helps the buyer determine his valuation, and the problem is how to optimally use this power to maximize revenue. Bergemann and Strack [6] and Athey and Segal [1] consider very general value evolution models where the evolution could depend on the action of the mechanism. They characterize the *welfare* maximizing mechanism, a generalization of the VCG mechanism to dynamic settings. These papers mostly focus on characterizing the revenue/welfare maximizing mechanism, typically under an interim-IR type condition. (The condition is ex-post over the history up to any time but ex-ante over the future.) This could lead to elaborate contracts, or require the buyer and/or seller to solve computationally expensive optimization problems.

The work that is most closely related to our model is that of Kakade et al. [13] and Pavan et al. [16]. They consider a revenue-optimal dynamic mechanism design problem where the buyer valuation changes over time, based on the signals that the buyers get with each time period. Both works assume that buyers are risk neutral, and construct mechanisms that are interim IR.

Kakade et al. [13] characterize the optimal Bayesian incentive compatible (BIC) and interim IR⁸ mechanism for many bidders, for value evolution processes that satisfy one of two separability conditions. Roughly, their additive separability condition is that $V_t = V_0 + M_t$ where M_t is some random process, independent of V_0 . The process we consider does not satisfy this condition because of boundary condi-

⁸They call their mechanism ex-post IR, but what this means is that for every period, the expected utility conditioned on the history up to that point and the expectation taken over all the signals in the future is non-negative. Ex-post IR (at the end) would mean that the utility at the end of the mechanism is non-negative and per period ex-post IR would mean that the utility in every period is non-negative.

tions.⁹ Their multiplicative separability condition is that $V_t = V_0 M_t$.

Pavan et al. [16] identify a set of sufficient conditions on the value evolution process under which envelope theorems can be used to characterize incentive-compatible allocation and payment rules. They then characterize perfect Bayesian equilibrium-implementable allocation rules in Markov environments. The mechanisms from Kakade et al. [13] and Pavan et al. [16] apply in many settings to which our model does not apply, but are computationally more complex, both for the seller and the buyers. They also require detailed distributional information, and would not work for risk-averse buyers.

Another major focus in dynamic mechanism design is on revenue/welfare maximization when buyers arrive sequentially, with either adversarial value distributions, or values drawn from a known/unknown distribution [15, 3, 10, 14, 9, 2]. A special case is dynamic pricing where the seller posts a price for each buyer (as opposed to more general schemes like auctions). There is a huge body of work on dynamic pricing and revenue management problems in the Operations Research literature. See [8] and references there in. The book by Vohra and Krishnamurthi [17] compares various pricing strategies like posted-prices, auctions, and haggling in the presence/absence of competition. This line of work is mostly unrelated to the problem we consider.

3 The necessity of assumptions in the bounded martingale model

We now present an example showing that a constant price PPP (without a free trial) cannot obtain a constant approximation against the buyer's cumulative value, as well as some examples demonstrating that the assumptions we make about the value evolution model are necessary.

Free trial is necessary: Consider a setting where the value evolves according to a simple random walk with steps of size δ , with reflection at 1 and absorption at 0, where $1/\delta \in \mathbb{Z}$. In this model we have $\epsilon = \delta$, and thus our main theorem states that free trial + PPP should get a constant factor approximation. However, we show that without free trial, the approximation factor for a starting value of $V_0 = \delta$ is $\Theta(\frac{1}{\delta})$. First we note that starting at $V_0 = \delta$, the expected time to absorption at 0 is

⁹The following is an easy way to see this is. With additive separability, the processes with different initial V_0 s have different maxima and minima, whereas the maxima/minima for our process are the same for all initial values.

$\delta(2-\delta)/\delta^2 < 2/\delta$ (see Lemma 4.6). In order to obtain non-zero revenue, any constant price PPP scheme must charge a price of at most δ , and therefore gets revenue no more than δ times $2/\delta$, which is $O(1)$. On the other hand, we can compute the cumulative value for an initial value of v by solving the following system of equations: $C(v) = v + \frac{C(v+\delta)+C(v-\delta)}{2}$ for $v = i\delta$, $i \in \{1, \dots, 1/\delta - 1\}$, $C(0) = 0$, and $C(1) = 1 + C(1 - \delta)$. Solving this system, we obtain $C(v) \geq \frac{2v}{3\delta^2}$,¹⁰ and therefore, $C(\delta) \geq \frac{2}{3\delta}$.

Martingale assumption (3) is necessary:

Let $\delta, \eta > 0$ be some sufficiently small numbers. Let V_t be either η or $\eta + \delta$, uniformly at random and independent of all previous values. Let the process have a fixed horizon of T (after which, say it is deterministically 0). It is easy to see that this satisfies all assumptions except that it is not a martingale. The expected value at any time $t \leq T$ is $\mathbf{E}[V_t] = \eta + \delta/2$, and so the cumulative value is $T(\eta + \delta/2)$. We claim that no PPP scheme, even one with time-varying prices, can obtain a revenue more than $T\eta + \delta$. For $\eta \ll \delta$, this gives an approximation ratio of $\Omega(T)$, which is unbounded. To prove the claim, we note that at every step, the optimal PPP scheme offers a price of either η or $\eta + \delta$. The total revenue from time periods in which the price is $\eta + \delta$ is at most $\eta + \delta$, since in each such time period the buyer stops buying with probability $1/2$. The revenue from each time period in which the price is η is exactly η . Therefore, the total revenue obtained by PPP is at most $T\eta + \delta$.

Minimum variance assumption (5) is necessary: Let μ be a random variable drawn from the “equal revenue” distribution over the set $[n]$, i.e., $\Pr[\mu \geq i] = \frac{1}{i}$. Consider a bin with μ red balls and $n - \mu$ blue balls, and draw these balls one by one uniformly at random, without replacement. For all $t \in [n]$, let $X_t := 1$ if the t^{th} ball drawn is red, and 0 otherwise. Note that $\sum_{t=1}^n X_t = \mu$ by definition.

Suppose that there is only one type of buyer whose value evolves as follows. Let for all $t \in [n]$,

$$V_t := \mathbf{E}[\mu | X_1, \dots, X_{t-1}] / n.$$

For all $t \in \{n+1, n+2, \dots, n+T\}$, $V_t := V_{t-1}$. (The process has a finite horizon.) This is a Doob martingale with bounded differences, and hence it satisfies the martingale and step size assumptions. It is, by definition, in the interval $[0, 1]$, and 0 is trivially an absorbing state (since it is never reached). Hence

¹⁰The unique solution to this system of equations is $C(v) = \frac{v}{3\delta^2}(3 + \delta^2 - v^2)$.

it satisfies all the properties except for the minimum variance property (5).

The cumulative value of the buyer is $(n + T)\mathbf{E}[\mu]/n = H_n(1 + T/n) = O(T \ln n/n)$, for $T \gg n$. We will now argue that no PPP scheme (including ones that can have a different price for every step) can get a revenue more than $O(T/n)$. Note that V_n (and hence V_t for any $t > n$) is exactly equal to μ , by definition. For any fixed time $t > n$, and any price $p = i/n$, $\Pr[V_t \geq i/n] = \Pr[\mu \geq i] = \frac{1}{i}$. The expected revenue in the t^{th} step is hence equal to $1/n$ for all such prices. Further, the continuation probability is maximized when $p = 1/n$, therefore the optimal PPP scheme in rounds $t > n$ is to set a price of $1/n$ for all time steps and get a revenue of T/n . For $T \gg n$, the total revenue of the PPP scheme is at most $n + T/n = O(T/n)$.

Dependence on the ratio of ϵ to δ (assumptions (4) and (5)): We now show that the approximation factor achieved by free trial + PPP must depend on the ratio of the upper bound on the step size (ϵ), and the square root of the lower bound on the variance (δ). We construct an example where $\epsilon = 1$, and show that any free trial + fixed price PPP obtains an approximation factor of $\Omega(\log^2(\frac{1}{\delta}))$.¹¹

Let the initial value be distributed according to an equal revenue distribution as follows: let $\delta = 2^{-N}$; for $k = 0, \dots, N$, the initial value is $2^k \delta$ with probability $\frac{1}{2^{k+1}}$; The remaining mass of $\frac{1}{2^{N+1}}$ is at 0. The value evolves as follows: for $k = 0, \dots, N-1$, when the current value is $2^k \delta$, the next value is $2^{k+1} \delta$ with probability $\frac{1}{2}$ and 0 with probability $\frac{1}{2}$. When the current value is 1, the next value is 0 with probability 1. Given this value evolution, the cumulative value starting at an initial value of $2^k \delta$ can be computed to be $(N - k + 1)/2^{N-k}$. Therefore, the expected cumulative value C is $\sum_{k=0}^N \frac{1}{2^{k+1}} \cdot \frac{N-k+1}{2^{N-k}} = \Theta(N^2 \delta)$.

We will now show an upper bound on the revenue of any free trial + fixed price PPP scheme. First, we argue that it is suboptimal for the PPP scheme to offer any free trial period. This follows from the following observation: the probability distribution of V_t , given the above initial distribution over V_0 and the evolution process, stochastically dominates the distribution of V_{t+1} , even when conditioned on V_t being above a certain price. Consequently, for any T, p , the revenue from offering a free trial for T

¹¹We remark that for the example we construct, variable price PPP can obtain an approximation factor of $\Theta(\log(\frac{1}{\delta}))$, but no better; This is essentially the familiar revenue versus social welfare gap for the equal revenue value distribution.

rounds, followed by a per-round price of p is no larger than that from offering no free trial and a fixed price of p . Finally, let $2^k \delta$ be the price charged by the fixed-price PPP scheme. The revenue of such a scheme is

$$\begin{aligned} \sum_{r=1}^{N+1} 2^k \delta \sum_{z=k+r}^{N+1} \frac{1}{2^z} &= \sum_{r=1}^{N+1} \frac{\delta}{2^{r-1}} \left(1 - \frac{1}{2^{N+2-k-r}} \right) \\ &\leq 2\delta. \end{aligned}$$

Thus, the gap from cumulative value is $\Theta(N^2) = \Theta(\log^2(\frac{1}{\delta}))$.

4 Proof of the main theorem

In this section we will present the proof of Theorem 1.1 which claims that there is a free trial period pricing scheme (i.e. T and c) that yields the seller a constant fraction of the expected cumulative value, for $\delta/\epsilon = \Omega(1)$. We begin with some notation:

- Throughout, we condition on $V_0 = v$.
- Let c denote the value in $[0, 1]$ that will end up being the PPP price that we charge once the free trial period ends.
- Define $w = \frac{\delta}{\epsilon}$, and $c = c' \frac{\delta}{\epsilon}$ for some constant $c' \in (0, 1]$.
- We will consider the following stopping times.

- τ : the first time t such that $V_t \geq w$ or $V_t = 0$;
- τ_1 : the first time after τ that $V_t = 1$ or $V_t < c$, conditioned on $V_\tau \geq w$;
- τ_2 : the first time after τ_1 that the value reaches c , conditioned on $V_\tau \geq w$ and $V_{\tau_1} = 1$.

The lemmas in the next subsection present analysis of the martingale process that we need. The key tool here is the optional stopping theorem, applied to a number of different martingales derived from V_t .

4.1 Analysis of the value evolution process

The first lemma determines the probability that the process reaches w before dying out and being absorbed at 0. Note that w is a constant that depends on the ratio between δ and ϵ . The second and third lemma show that conditioned on reaching w before 0, the expected time to reach w is relatively small. It is this last lemma that enables us to set the free trial period length T . The goal is to ensure that there is a reasonable chance that τ occurs before the

free trial period ends. The fourth lemma shows that conditioned on reaching w before 0, the value has sufficient probability of reaching 1 before dropping below the PPP price. The fifth lemma shows that conditioning on these latter events, the expected time between hitting 1 and dropping below c (i.e., the length of the period during which the PPP scheme is making money), is sufficiently large.

LEMMA 4.1. *The probability of crossing w before 0 starting at v , i.e. $\Pr[V_\tau \geq w | V_0 = v]$, is at least $v/(w + \epsilon)$ and at most v/w .*

Proof. τ is a stopping time and V_t is a martingale in $[0, \tau - 1]$, so by the optional stopping theorem

$$\begin{aligned} \mathbf{E}[V_\tau] &= \mathbf{E}[V_0] \\ &= v \\ &= \mathbf{P}[V_\tau \geq w] \mathbf{E}[V_\tau | V_\tau \geq w] \\ &\quad + (1 - \mathbf{P}[V_\tau \geq w]) \mathbf{E}[V_\tau | V_\tau = 0] \\ &\in [w, w + \epsilon] \mathbf{P}[V_\tau \geq w], \end{aligned}$$

which yields this fact. \square

LEMMA 4.2. *Let τ be the first time V_t is either 0 or crosses w . Then*

$$\mathbf{E}[\sum_{i=0}^{\tau-1} \Delta_i^2 | V_0 = v] \leq v((w + \epsilon)^2/w - v).$$

Proof. $X_t = V_t^2 - \sum_{j=0}^{t-1} \Delta_j^2$ is a martingale. Applying optional stopping at τ yields

$$\begin{aligned} v^2 &= \mathbf{E}[V_\tau^2] - \mathbf{E}\left[\sum_{j=0}^{\tau-1} \Delta_j^2\right] \\ &\leq (w + \epsilon)^2 v/w - \mathbf{E}\left[\sum_{j=0}^{\tau-1} \Delta_j^2\right], \end{aligned}$$

which implies the lemma. \square

LEMMA 4.3. *Let τ be the first time V_t is either 0 or crosses w . Conditioned on the fact that V_t crosses w before 0 (i.e. $V_\tau \geq w$), the expected time to cross w is “small”:*

$$\mathbf{E}[\tau | V_\tau \geq w] \leq 1/3\epsilon^2 (1 + \epsilon^2/\delta)^3 (1 + 3\epsilon^2/\delta).$$

Proof. Define M_t as follows:

$$M_t := V_t^3 - 3V_t \sum_{j=0}^{t-1} \Delta_j^2 + 2 \sum_{j=0}^{t-1} \Delta_j^3.$$

Then M_t is a martingale in $[0, \tau - 1]$. To see this, observe that

$$\begin{aligned} M_{t+1} - M_t &= (V_t + \Delta_t)^3 - V_t^3 - 3(V_t + \Delta_t) \sum_{j=0}^t \Delta_j^2 \\ &\quad + 3V_t \sum_{j=0}^{t-1} \Delta_j^2 + 2\Delta_t^3 \\ &= (3V_t \Delta_t^2 + 3V_t^2 \Delta_t + \Delta_t^3) - 3V_t \Delta_t^2 \\ &\quad - 3\Delta_t \sum_{j=0}^t \Delta_j^2 + 2\Delta_t^3 \\ &= 3\Delta_t \left(V_t^2 - \sum_{j=0}^{t-1} \Delta_j^2 \right). \end{aligned}$$

Since $\mathbf{E}[\Delta_t | V_0, \dots, V_t] = 0$, we have that $\mathbf{E}[M_{t+1} | V_0, \dots, V_t] = M_t$. Applying the optional stopping theorem, we obtain

$$\begin{aligned} \mathbf{E}[M_0] &= v^3 = \mathbf{E}[M_\tau] \\ &\leq (w + \epsilon)^3 \mathbf{P}[V_\tau \geq w] - 3 \mathbf{E} \left[V_\tau \sum_{j=0}^{\tau-1} \Delta_j^2 \right] \\ &\quad + 2 \mathbf{E} \left[\sum_{j=0}^{\tau-1} \Delta_j^3 \right]. \end{aligned}$$

Next we recall that for all j , $\Delta_j \leq \epsilon$, and therefore, $\Delta_j^3 \leq \epsilon \Delta_j^2$. Thus, using $V_\tau \in \{0\} \cup [w, w + \epsilon]$ and $v/(w + \epsilon) \leq \mathbf{P}[V_\tau \geq w] \leq v/w$, and applying Lemma 4.2,

$$\begin{aligned} v^3 &\leq (w + \epsilon)^3 \mathbf{P}[V_\tau \geq w] \\ &\quad - 3w \mathbf{E} \left[\sum_{j=0}^{\tau-1} \Delta_j^2 \mid V_\tau \geq w \right] \mathbf{P}[V_\tau \geq w] \\ &\quad + 2\epsilon \mathbf{E} \left[\sum_{j=0}^{\tau-1} \Delta_j^2 \right] \\ &\leq (w + \epsilon)^3 \frac{v}{w} - 3w \frac{v}{w + \epsilon} \mathbf{E} \left[\sum_{j=0}^{\tau-1} \Delta_j^2 \mid V_\tau \geq w \right] \\ &\quad + 2\epsilon v \left(\frac{(w + \epsilon)^2}{w} - v \right). \end{aligned}$$

Dividing by $3vw/(w + \epsilon)$ we obtain

$$\begin{aligned} &\mathbf{E} \left[\sum_{j=0}^{\tau-1} \Delta_j^2 \mid V_\tau \geq w \right] \\ &\leq \frac{(w + \epsilon)^2}{3w^2} \left((w + \epsilon)^2 + 2\epsilon(w + \epsilon) - v(v + 2\epsilon) \frac{w}{w + \epsilon} \right) \\ &\leq \frac{1}{3w^2} (w + \epsilon)^3 (w + 3\epsilon). \end{aligned}$$

The lemma then follows by noting that $\mathbf{E}[\Delta_j^2 | \Delta_0, \dots, \Delta_{j-1}] \geq \delta^2$ for all j , and substituting the value of $w = \delta/\epsilon$. \square

We use the following corollary of this lemma.

COROLLARY 4.1. *Suppose that $w = \delta/\epsilon$ and $\epsilon^2/\delta < 1/9$. Then,*

$$\mathbf{E}[\tau \mid V_\tau \geq w] \leq 5/8\epsilon^2.$$

The following lemma follows from an argument similar to the proof of Lemma 4.1.

LEMMA 4.4. *Conditioned on $V_\tau \geq w$, and any history up to time τ , the value reaches 1 before c with sufficient probability:*

$$\Pr[V_{\tau_1} = 1 \mid V_\tau \geq w] \geq \frac{w-c}{1-c} \geq \frac{\delta}{\epsilon} (1 - c').$$

LEMMA 4.5. *Conditioned on $V_{\tau_1} = 1$ and any history up to time τ_1 , let τ_2 be the first time $V_t < c$ for $t \geq \tau_1$. Then the time between τ_1 and τ_2 is large:*

$$\mathbf{E}[\tau_2 - \tau_1 \mid V_{\tau_1} = 1] \geq \frac{(1-c)^2 - c\epsilon}{\epsilon^2}.$$

Proof. Suppose that $V_{\tau_1} = 1$. We apply a version of the reflection principle for simple random walks: Consider the martingale X_t that is coupled with V_t . $X_0 = V_{\tau_1}$ (conditioned on the history). For $t > 0$, when $V_{\tau_1+t} < 1$, $X_{t+1} = V_{\tau_1+t+1}$. When $V_{\tau_1+t} = 1$, with probability 1/2, we set $X_{t+1} = V_{\tau_1+t+1}$, and with probability 1/2, we set it to $2 - V_{\tau_1+t+1}$. When $X_t > 1$, we set $X_{t+1} = 2 - V_{\tau_1+t+1}$. One can check that X_t is a martingale with absorbing states at 0 and 2.

Now, let us suppose that $X_0 = 1$, and determine the expected time it takes to reach either a value less than c or a value greater than $2 - c$ for some $c \in [0, 1]$. This time, $\tau_2 - \tau_1$, is a stopping time. Note that each of these possibilities happen with probability 1/2. Consider the martingale $Z_t = X_t^2 - \sum_{i=\tau_1}^{i=\tau_1+t-1} \Delta_i^2$. Then, we have

$$\begin{aligned} 1 &= \mathbf{E}[Z_0] = \mathbf{E}[Z_{\tau_2 - \tau_1}] \\ &\geq (c - \epsilon)^2/2 + (2 - c)^2/2 - \mathbf{E} \left[\sum_{i=\tau_1}^{i=\tau_2-1} \Delta_i^2 \right]. \end{aligned}$$

Rearranging this expression, and using $\Delta_i \leq \epsilon$ for all i , implies the lemma. \square

4.2 Upper bound on cumulative value

The following lemma bounds the buyer's cumulative value: we use the variance in the value evolution process to argue that the value gets absorbed at zero before too long.

LEMMA 4.6. *The cumulative value starting at v is not too large: $\mathcal{C}(v) \leq \frac{v(2-v)}{\delta^2}$.*

Proof. Let $\tilde{\tau}$ be the first time that V_t is 0. We will prove that starting at value v , the expected value of $\tilde{\tau}$ is at most $v(2-v)/\delta^2$. This implies the theorem because at every step, while the value is non-zero, the cumulative value accrues an amount no larger than 1.

To prove the claim, as in the proof of Lemma B.4, let us define a martingale X_t that is coupled with V_t . $X_0 = V_0$. For $t > 0$, when $V_t < 1$, $X_{t+1} = V_{t+1}$. When $V_t = 1$, with probability $1/2$, we set $X_{t+1} = V_{t+1}$, and with probability $1/2$, we set it to $2 - V_{t+1}$. When $X_t > 1$, we set $X_{t+1} = 2 - V_{t+1}$. One can check that X_t is a martingale with absorbing states at 0 and 2.

Now, let us compute the expected time that X_t takes to reach 0 or 2 (which is $\mathbf{E}[\tilde{\tau}]$). Applying the optional stopping theorem, the probability that we hit 2 before 0 is $v/2$. Consider the martingale $Z_t = X_t^2 - \sum_{i=0}^{t-1} \delta_i^2$. Then, we have

$$v^2 = \mathbf{E}[Z_0] = \mathbf{E}[Z_{\tilde{\tau}}] = 4\frac{v}{2} - \mathbf{E}\left[\sum_{i=0}^{i=\tilde{\tau}-1} \delta_i^2\right].$$

Rearranging this expression and using $\delta_j^2 \geq \delta^2$ for all j implies the claim. \square

4.3 Proof of Theorem 1.1

We will assume that $\epsilon^2/\delta < 1/9$. Set the free trial period length to be

$$T := 3/4\epsilon^2 \geq 6/5 \mathbf{E}[\tau \mid V_\tau \geq w];$$

(The above inequality follows from Corollary 4.1.) Thus, by Markov's inequality,

$$(4.2) \quad \mathbf{P}_v[\tau < T \mid V_\tau \geq w] \geq 1/6.$$

Let T' be the length of time after the free trial period ends for which V_t is at least c : $T' = \min\{t - T : t > T, \text{ and, } V_{t+1} < c\}$. We will condition on the event

$$E = \{V_\tau \geq w \text{ and } \tau \leq T \text{ and } V_{\tau_1} = 1\}.$$

Note that conditioned on E , we have $T' \geq \tau_2 - \max\{\tau_1, T\} \geq \tau_2 - \tau_1 - T$.

Thus, for a buyer with starting value $V_0 = v$ we obtain

$$\begin{aligned} \mathcal{R}_{T,c} &\geq \mathbf{E}[c(\tau_2 - \tau_1 - T) \mid E] \cdot \mathbf{P}_v[E] \\ &\geq c \mathbf{E}[\tau_2 - \tau_1 - T \mid E] \cdot \Pr[V_{\tau_1} = 1 \mid \tau < T, V_\tau \geq w] \\ &\quad \cdot \Pr[\tau < T \mid V_\tau \geq w] \cdot \Pr[V_\tau \geq w \mid V_0 = v] \end{aligned}$$

For $v < w$, we simplify this using the following facts:

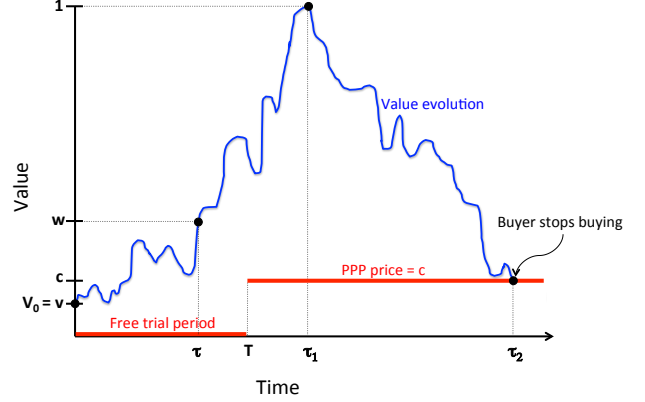


Figure 1: The figure illustrates the event E that we are conditioning on.

- By Lemma 4.1, $\Pr[V_\tau > w \mid V_0 = v] \geq v\frac{\epsilon}{\delta}$.
- By inequality (4.2): $\mathbf{P}_v[\tau < T \mid V_\tau \geq w] \geq \frac{1}{6}$.
- By Lemma 4.4, $\Pr[V_{\tau_1} = 1 \mid \tau < T, V_\tau > w] \geq \frac{\delta}{\epsilon}(1 - c')$. (Notice that the extra conditioning on $\tau < T$ does not affect anything except the history: since T is a constant, this event is determined at time τ .)

Together, these imply that the probability of E , i.e. the value reaches 1 before the buyer stops buying the product, is at least $v\frac{\epsilon}{\delta} \cdot \frac{1}{6} \cdot \frac{\delta}{\epsilon}(1 - c') = \frac{1}{6}v(1 - c')$.

Finally, we apply Lemma 4.5. Again, we observe that, since T is a constant, the event $\tau < T$ is determined at time τ , and thus affects only the history before time τ_1 in Lemma 4.5. Therefore, we have

$$\begin{aligned} \mathcal{R}_{T,c} \mid_{V_0=v} &\geq c \mathbf{E}[\tau_2 - \tau_1 - T \mid E] \cdot \frac{1}{6}v(1 - c') \\ &\geq c \frac{((1 - c)^2 - c\epsilon - 3/4)}{\epsilon^2} \frac{1}{6}v(1 - c') \\ &= \Omega\left(\frac{\delta}{\epsilon^3}v\right). \end{aligned}$$

The final equality follows by substituting $c = c'\delta/\epsilon$, picking a c' small enough and simplifying.

For $v \geq w$, define $\tau = 0$, and condition on the event

$V_{\tau_1} = 1$. Using a similar analysis, we obtain

$$\begin{aligned} \mathcal{R}_{T,c}|_{v \geq w} &\geq c \mathbf{E} [\tau_2 - \tau_1 - T | V_{\tau_1} = 1] \\ &\quad \cdot \Pr[V_{\tau_1} = 1 | V_0 = v] \\ &\geq c \frac{((1-c)^2 - c\epsilon - 3/4)}{\epsilon^2} \cdot \frac{v-c}{1-c} \\ &= \Omega\left(\frac{\delta}{\epsilon^3} v\right). \end{aligned}$$

Finally, using Lemma 4.6, for any value of V_0 , we obtain the theorem. ■

References

- [1] Susan Athey and Ilya Segal. An Efficient Dynamic Mechanism. Levine’s Bibliography 122247000000001134, UCLA Department of Economics, March 2007. URL <http://ideas.repec.org/p/cla/levrem/122247000000001134.html>.
- [2] Moshe Babaioff, Shaddin Dughmi, Robert Kleinberg, and Aleksandrs Slivkins. Dynamic pricing with limited supply. In *ACM Conference on Electronic Commerce, EC ’12, Valencia, Spain, June 4-8, 2012*, pages 74–91, 2012.
- [3] Ziv Bar-Yossef, Kirsten Hildrum, and Felix Wu. Incentive-compatible online auctions for digital goods. In *Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA ’02*, pages 964–970, 2002.
- [4] David P. Baron and David Besanko. Regulation and information in a continuing relationship. *Information Economics and Policy*, 1(3):267–302, 1984.
- [5] Marco Battaglini. Long-term contracting with markovian consumers. *American Economic Review*, 95(3):637–658, 2005.
- [6] Dirk Bergemann and Philipp Strack. Dynamic revenue maximization: A continuous time approach. 2014.
- [7] David Besanko. Multi-period contracts between principal and agent with adverse selection. *Economics Letters*, 17(12):33–37, 1985.
- [8] Omar Besbes and Assaf Zeevi. Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Operations Research*, 57(6):1407–1420, November 2009.
- [9] Avrim Blum and Jason D. Hartline. Near-optimal online auctions. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 1156–1163, 2005.
- [10] Avrim Blum, Vijay Kumar, Atri Rudra, and Felix Wu. Online learning in online auctions. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms, January 12-14, 2003, Baltimore, Maryland, USA.*, pages 202–204, 2003.
- [11] Pascal Courty and Li Hao. Sequential screening. *Review of Economic Studies*, 67(4):697–717, 2000.
- [12] Peter Eso and Balzs Szentes. Optimal information disclosure in auctions and the handicap auction. *Review of Economic Studies*, 74(3):705–731, 2007.
- [13] Sham M Kakade, Ilan Lobel, and Hamid Nazarzadeh. Optimal dynamic mechanism design and the virtual-pivot mechanism. *Operations Research*, 61(4):837–854, 2013.
- [14] Robert D. Kleinberg and Frank Thomson Leighton. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In *44th Symposium on Foundations of Computer Science (FOCS)*, pages 594–605, 2003.
- [15] Ron Lavi and Noam Nisan. Competitive analysis of incentive compatible on-line auctions. In *Proceedings of the 2Nd ACM Conference on Electronic Commerce, EC ’00*, pages 233–241, 2000.
- [16] Alessandro Pavan, Ilya Segal, and Juuso Toikka. Dynamic mechanism design: A myersonian approach. *Econometrica*, 82(2):601–653, 2014.
- [17] Rakesh V. Vohra and Lakshman Krishnamurthi. *Principles of Pricing - An Analytical Approach*. Cambridge University Press, 2012. ISBN 978-1-107-01065-9.

A Binary value model

In this section we consider an alternate model of value evolution, which we call the **binary value** model. In this model, $V_t = V_0$ for all $t \leq T(V_0)$ and is 0 otherwise, where $T(V_0)$ is a random variable that depends on V_0 . We assume that $\mathbf{E}[T(V_0)|V_0]$

is monotone non decreasing in V_0 . This captures scenarios where the buyer's value per usage doesn't change with usage, but the number of times he uses the product is uncertain. Buyers with higher value per usage tend to use the product more.

For this model, we compare the revenue of a simple constant-price PPP scheme with that of an optimal BIN scheme with a risk neutral buyer. In the BIN scheme, a price p_{BIN} is offered for unlimited usage of the product and the buyer accepts this price if his expected total value $\mathbf{E}[V_0 T(V_0)]$ is at least p_{BIN} . In the constant price PPP scheme with price p , each buyer with $V_0 \geq p$ keeps buying the product upto time $T(V_0)$. We show that there exists a price p for which, PPP gets more revenue as well as more buyer utility as compared to the optimal BIN.

THEOREM A.1. *In the binary value model there is a constant price PPP scheme that gets at least as much risk-robust revenue, as much social welfare, and as much buyer utility as the revenue optimal BIN scheme with a risk-neutral buyer.*

Proof. In the binary value model the buyer's expected cumulative value is $\mathcal{C}(V_0) = V_0 \mathbf{E}[T(V_0)|V_0]$. The optimal BIN price is the monopoly price for the distribution of this random variable where V_0 is drawn from some distribution F . Since $\mathbf{E}[T(V_0)|V_0]$ is non-decreasing in V_0 , there is a threshold initial value v^{BIN} such that the risk neutral buyer purchases in this BIN scheme if and only if $V_0 \geq v^{\text{BIN}}$. The price of this BIN scheme is then $p_{\text{BIN}} = \mathcal{C}(v^{\text{BIN}}) = v^{\text{BIN}} \mathbf{E}[T(v^{\text{BIN}})]$.

Now consider a constant price PPP scheme with price p equal to v^{BIN} . At this price, any buyer with $V_0 \geq v^{\text{BIN}}$ purchases the product for $T(V_0)$ periods, i.e., the same set of buyers purchase the product in the optimal BIN scheme as well as this one. However, a buyer with initial value V_0 pays a total of $v^{\text{BIN}} \mathbf{E}[T(V_0)]$, which is at least as large as $v^{\text{BIN}} \mathbf{E}[T(v^{\text{BIN}})] = p_{\text{BIN}}$. Thus, the revenue of PPP is at least that of BIN. Further, the two schemes generate the same total social welfare.

Now, consider gradually decreasing the per-play price in the PPP scheme to below v^{BIN} . Then, the social welfare generated by PPP increases, while its revenue may or may not decrease. We can continue decreasing the price as long as the revenue of PPP stays above that of BIN, and at some point, both the revenue and the buyer utility (which is social welfare – revenue) of PPP exceed the corresponding quantities for BIN, since revenue eventually becomes 0 at a price of 0. \square

B Martingales with independent increments

In this section, we consider a special case of the martingale model where the increments $\Delta_t = V_{t+1} - V_t$ are independent of the current value V_t and of each other. In this case, we can get a sharper bound.

B.1 Preliminary lemmas

We develop the analogues of the lemmas used in Section 4. The proofs are very similar, but we are able to control the random quantities more easily because of the independence across steps. This allows us to replace certain random variables by their expectations. Let τ denote the first time that the V_t hits 0 or 1. Conditioned on V_τ being 1, we will use $\tau' > \tau$ to denote the first time after τ that the value becomes less than or equal to c (the PPP price). Also let

$$\delta_t^2 := \mathbf{E}[\Delta_t^2].$$

LEMMA B.1. *The probability of hitting 1 before 0 starting at v , i.e. $\Pr[V_\tau = 1|V_0 = v]$, is v .*

In the analysis of this section, we use $\mathbf{E}[\sum_{i=0}^{i=\tau-1} \delta_i^2]$ as a proxy of sorts for the time it takes for the random walk to reach 0 or 1. The following lemma bounds this expectation.

LEMMA B.2. *Let τ be the first time V_t is either 0 or 1. Then*

$$\mathbf{E}\left[\sum_{i=0}^{i=\tau-1} \delta_i^2 \mid V_0 = v\right] = v(1-v).$$

LEMMA B.3. *Let τ be the first time V_t is either 0 or 1. Conditioned on the fact that V_t hits 1 before 0 (i.e. $V_\tau = 1$), the expected time to hit 1 is “small”:*

$$\mathbf{E}\left[\sum_{i=0}^{i=\tau-1} \delta_i^2 \mid V_\tau = 1\right] \leq \frac{1+\epsilon}{3}.$$

LEMMA B.4. *Let τ be the first time V_t is either 0 or 1. Let τ' be the first time after τ that V_t drops below c conditioned on $V_\tau = 1$. The time between τ and τ' is large:*

$$\mathbf{E}\left[\sum_{i=\tau}^{i=\tau'-1} \delta_i^2 \mid V_\tau = 1\right] \geq (1-c)^2 - c\epsilon.$$

B.2 Proofs of key lemmas

Proof of Lemma B.1. τ is a stopping time and V_t is a martingale in $[0, \tau - 1]$, so applying the optional

stopping theorem, and noting $V_\tau \in \{0, 1\}$,

$$\begin{aligned} \mathbf{E}[V_\tau] &= \mathbf{E}[V_0] \\ &= v \\ &= \mathbf{P}[V_\tau = 1] \mathbf{E}[V_\tau | V_\tau = 1] \\ &\quad + (1 - \mathbf{P}[V_\tau = 1]) \mathbf{E}[V_\tau | V_\tau = 0]. \end{aligned}$$

which yields this fact. \blacksquare

Proof of Lemma B.2. Consider the random variable

$$X_t = V_t^2 - \sum_{j=0}^{t-1} \delta_j^2.$$

X_t is a martingale. Applying optional stopping at τ yields

$$\begin{aligned} v^2 &= \mathbf{E}[X_0] = \mathbf{E}[X_\tau] = \mathbf{E}[V_\tau^2] - \mathbf{E}\left[\sum_{j=0}^{\tau-1} \delta_j^2\right] \\ &= v - \mathbf{E}\left[\sum_{j=0}^{\tau-1} \delta_j^2\right]. \end{aligned}$$

Here the last inequality follows by noting $V_\tau^2 \in \{0, 1\}$, and applying Lemma B.1. Rearranging the equation implies the lemma. \blacksquare

Proof of Lemma B.3. Define M_t as follows:

$$M_t := V_t^3 - 3V_t \sum_{j=0}^{t-1} \delta_j^2 - \sum_{j=0}^{t-1} \gamma_j^3,$$

where $\gamma_j = \mathbf{E}[\Delta_t^3]$. Then M_t is a martingale in $[0, \tau - 1]$:

$$\begin{aligned} M_{t+1} - M_t &= (V_t + \Delta_t)^3 - V_t^3 - 3(V_t + \Delta_t) \sum_{j=0}^t \delta_j^2 \\ &\quad + 3V_t \sum_{j=0}^{t-1} \delta_j^2 - \gamma_t^3 \\ &= (3V_t \Delta_t^2 + 3V_t^2 \Delta_t + \Delta_t^3) - 3V_t \delta_t^2 \\ &\quad - 3\Delta_t \sum_{j=0}^t \delta_j^2 - \gamma_t^3 \\ &= 3\Delta_t (V_t^2 - \sum_{j=0}^t \delta_j^2) + 3V_t (\Delta_t^2 - \delta_t^2) \\ &\quad + (\Delta_t^3 - \gamma_t^3). \end{aligned}$$

Thus,

$$\mathbf{E}[M_{t+1} - M_t | V_0, \dots, V_t] = 0.$$

Applying the optional stopping theorem, we obtain

$$\begin{aligned} \mathbf{E}[M_0] &= v^3 \\ &= \mathbf{E}[M_\tau] \\ &= \mathbf{P}[V_\tau = 1] - 3 \mathbf{E}\left[V_\tau \sum_{j=0}^{\tau-1} \delta_j^2\right] - \mathbf{E}\left[\sum_{j=0}^{\tau-1} \gamma_j^3\right]. \end{aligned}$$

Next we recall that for all j , $\Delta_j \geq -\epsilon$, and therefore, $\gamma_j^3 \geq (-\epsilon)\delta_j^2$. Thus, using $V_\tau \in \{0, 1\}$ and $\mathbf{P}[V_\tau = 1] = v$,

$$\begin{aligned} v^3 &\leq \mathbf{P}[V_\tau = 1] - 3 \mathbf{E}\left[\sum_{j=0}^{\tau-1} \delta_j^2 \mid V_\tau = 1\right] \mathbf{P}[V_\tau = 1] \\ &\quad + \epsilon \mathbf{E}\left[\sum_{j=0}^{\tau-1} \delta_j^2\right] \\ &= v - 3v \mathbf{E}\left[\sum_{j=0}^{\tau-1} \delta_j^2 \mid V_\tau = 1\right] + \epsilon \mathbf{E}\left[\sum_{j=0}^{\tau-1} \delta_j^2\right] \end{aligned}$$

Dividing by $3v$ and applying Lemma B.2, we obtain

$$\mathbf{E}\left[\sum_{j=0}^{\tau-1} \delta_j^2 \mid V_\tau = 1\right] \leq \frac{1}{3} (1 - v^2 + \epsilon(1 - v)) \leq \frac{1 + \epsilon}{3}.$$

\blacksquare

Proof of Lemma B.4. Let $V_\tau = 1$. We apply the reflection principle by considering the following martingale X_t that is coupled with V_t . $X_0 = V_\tau$. For $t > 0$, when $V_{\tau+t} < 1$, $X_{t+1} = V_{\tau+t+1}$. When $V_{\tau+t} = 1$, with probability $1/2$, we set $X_{t+1} = V_{\tau+t+1}$, and with probability $1/2$, we set it to $2 - V_{\tau+t+1}$. When $X_t > 1$, we set $X_{t+1} = 2 - V_{\tau+t+1}$. One can check that X_t is a martingale with absorbing states at 0 and 2.

Now, let us suppose that $X_0 = 1$, and determine the expected time it takes to reach either a value less than c or a value greater than $2 - c$ for some $c \in [0, 1)$. This time, $\tau' - \tau$, is a stopping time. Note that each of these possibilities happen with probability $1/2$. Consider the martingale $Z_t = X_t^2 - \sum_{i=\tau}^{i=\tau+t-1} \delta_i^2$.

Then, we have

$$\begin{aligned} 1 &= \mathbf{E}[Z_0] = \mathbf{E}[Z_{\tau'-\tau}] \\ &\geq \frac{(c-\epsilon)^2}{2} + \frac{(2-c)^2}{2} - \mathbf{E}\left[\sum_{i=\tau}^{i=\tau'-1} \delta_i^2\right]. \end{aligned}$$

Rearranging this expression implies the lemma. \blacksquare

B.3 Main Theorem

We can now prove the main theorem in this setting.

THEOREM B.1. *There exist constants T and c independent of $v = V_0$, such that the expected seller revenue from a PPP pricing scheme with a free trial period of length T and PPP price c for a risk averse buyer is a constant fraction of the buyer's expected cumulative value:*

$$\forall v \in [0, 1], \mathcal{R}_{T,c}|_{V_0=v} = \Omega\left(\left(\frac{\delta}{\epsilon}\right)^2 \mathcal{C}(v)\right).$$

Proof. Recall that τ is the time at which the value first becomes 0 or 1, and conditioned on $V_\tau = 1$, τ' is the first time the value drops below c . Let T be defined to be the smallest time such that

$$\sum_{i=0}^{i=T-1} \delta_i^2 \geq \frac{2}{3}(1+\epsilon).$$

Note that the right hand side of this inequality is at least $2\mathbf{E}[\sum_{i=0}^{i=\tau-1} \delta_i^2 | V_\tau = 1]$ by Lemma B.3, so by Markov's inequality,

$$(B.1) \quad \mathbf{P}_v\left[\sum_{i=0}^{i=\tau-1} \delta_i^2 \leq \sum_{i=0}^{i=T-1} \delta_i^2 \mid V_\tau = 1\right] \geq \frac{1}{2}.$$

Let T' be the length of time after the free trial period ends for which V_i is at least c : $T' = \min\{t-T : t > T, \text{ and, } V_{t+1} < c\}$. Note that conditioned on $\tau < T$ and $V_\tau = 1$, we have $T' \geq \tau' - T$. Then, $\mathcal{R}_{T,c} = \mathbf{E}[cT']$.

For a buyer with starting value $V_0 = v$ we obtain

$$\begin{aligned} (B.2) \quad \mathcal{R}_{T,c} &\geq \mathbf{E}_v[cT' \mid V_\tau = 1 \text{ and } \tau \leq T] \\ &\quad \cdot \mathbf{P}_v[\tau \leq T \mid V_\tau = 1] \mathbf{P}_v[V_\tau = 1] \\ &\geq c \mathbf{E}[\tau' - T \mid V_\tau = 1] \\ &\quad \cdot \mathbf{P}_v\left[\sum_{i=0}^{i=\tau-1} \delta_i^2 \leq \sum_{i=0}^{i=T-1} \delta_i^2 \mid V_\tau = 1\right] v \\ &\geq c \mathbf{E}[\tau' - T \mid V_\tau = 1] \frac{1}{2} v \\ &\geq c \cdot \frac{(1-c)^2 - c\epsilon - 2/3(1+\epsilon)}{\epsilon^2} \cdot \frac{v}{2} \end{aligned}$$

The second inequality here follows from Lemma B.1. The third is (B.1). The last inequality applies Lemma B.4 to get that

$$\mathbf{E}\left[\sum_{i=T}^{i=\tau'-1} \delta_i^2\right] \geq (1-c)^2 - c\epsilon - \frac{2}{3}(1+\epsilon),$$

and then uses $\delta_i^2 \leq \epsilon^2$ for all i .

Finally, we get that for sufficiently small c and ϵ , the RHS of (B.2) is $\Omega\left(\frac{v}{\epsilon^2}\right)$. Lemma 4.6 then implies the lemma. \square

C Simulations

In this section, we report the results of computer simulations we performed to empirically measure the performance of our PPP schemes. While our theorems already establish the benefits of such pricing schemes, the results of simulations are stronger, and favor PPP even more. In particular, for the binary value model, our theorem shows PPP schemes that simultaneously get more revenue and higher utility for buyers (in expectation over buyer types) than a Buy-It-Now scheme (BIN). Our simulations show that even for distributions very close to worst case, there are PPP schemes that simultaneously get more revenue and higher utility for *each* buyer type (not just in expectation over buyer types) than BIN. For the martingale model, we focus on the simple random walk as a special case and our simulations show that the constant factors are significantly smaller than what we get in our theorems, and we also identify distributions where PPP gets more revenue than BIN.

C.1 Binary value model

In the binary value model, we performed simulations for the initial value distributions being normal and uniform. When the initial value distribution is a point mass, it is impossible to obtain more revenue than a Buy-It-Now scheme because BIN revenue is the entire cumulative value. However, simulations show that even if there is a tiny spread in the initial value, say a Normal distribution with a small standard deviation, PPP already performs better in all dimensions. Thus, we perform simulations for Normal distributions with various means and standard deviations, and observe good performance even with tiny standard deviations. The improvement in revenue gets even better when the spread in initial value is more. To demonstrate this, we simulate the Uniform in $[0, 1]$ distribution and show the markedly higher improvement in revenue. Apart from revenue, we also measure the improvement in the number of

buyers that PPP brings in. The percentage improvement in the number of buyers is even higher than that of revenue.

Simulation Details. Since we want our PPP schemes to get higher utility (than BIN) for each buyer type, our PPP schemes stop charging a buyer once he has paid the BIN optimal price in total. We call this PPP-CAP. This immediately gives us the required pointwise guarantee on buyer utility. To measure revenue and buyer pool improvement, we consider 10000 samples from Normal distribution truncated between 0 and 1, with values discretized to multiples of 0.05. The three parameters of interest are the mean, standard deviation and the expected time alive (the expected time for which the initial value lasts at where it is, before it hits 0). As for expected time alive, our theorem just required that it is increasing in initial value. For our simulations, we fix time alive to be a polynomial in initial value, and vary the exponent in the polynomial. We begin with $\text{time alive} = \text{initial value}^{0.5}$.

1. **Normal with varying standard deviation:**

In Figure 2, we consider a Normal distribution with mean 0.2, and vary standard deviation from 0.2 to 0.02 (from mean to one-tenth of mean). If the standard deviation were 0, we get a point mass distribution and it is impossible to get any revenue improvement in this deterministic case. However, even with a small standard deviation like 0.02, PPP already gets a close to 5% revenue improvement, and more than a 5% increase in number of buyers.

2. **Normal with varying mean:** In Figure 3 we vary the mean from 0 to 1, fixing a small standard deviation of 0.02. As expected, the increase in revenue and number of buyers drops as mean increases. This is because, for a fixed standard deviation, as mean increases, the distribution looks closer and closer to a point-mass distribution.

3. **Normal with varying time alive:** Finally, in Figure 4, we vary the exponent q in expected time alive from 0 to 1, and observe that the increase in revenue and number of buyers increases with q . This is expected because as q increases, the expected time alive drops quicker as the initial value decreases, making BIN obtain revenue from a small fraction of high valued buyers. On the other hand, PPP-CAP will charge a low per-round price that greatly increases the pool of buyers, and still manages to get higher revenue

because the high valued buyers, though they pay a small per-round price, pay it till their value hits the BIN price, and the revenue from large pool of smaller valued buyers is additional.

4. **Uniform:** We also simulate the $U[0, 1]$ distribution, with values discretized to multiples of 0.05, and observed a 13% increase in revenue and 25% increase in number of buyers.

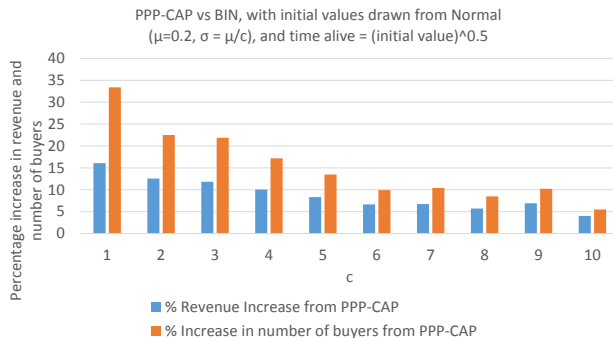


Figure 2: Binary value model, with initial values Normally distributed, and varying standard deviation.




Figure 3: Binary value model, with initial values Normally distributed, and varying mean.

C.2 Random walk model

For our main results in this paper, we studied a general martingale evolution of values. Here, we fix a special case of this, namely, a simple random walk between 0 and 1, with step size of $\delta = 0.05$. The walk is reflected at 1 and is absorbed at 0. This corresponds to the case of $\epsilon = \delta = 0.05$ in our model. We simulate this random walk, starting with different initial value distributions, and measure the revenue of free-trial+PPP scheme as a fraction of BIN revenue

| Initial value distribution | Rev(PPP) as Percentage of Rev(BIN)) |
|---|-------------------------------------|
| Point mass anywhere in $[0.05, 1]$ | 17% |
| Normal, with $\mu \in [0.05, 1]$, $\sigma = 0.1$ | 21% |
| Uniform $[0, 1]$ | 32% |

Table 1: Revenue of free trial + PPP scheme as a fraction of BIN revenue




Figure 4: Binary value model, with initial values Normally distributed, and varying time alive.

(and sometimes cumulative value). For every fixed initial value, we perform our simulations over 10000 samples of random walk trajectories and average over them. We summarize our results in Table 1.

1. **Point mass:** The worst-case is when the initial values are just point masses, in which case

BIN will extract the entire cumulative value as revenue. Even in this extreme case, PPP gets more than 17% of the cumulative value, with a free trial period that is agnostic to the initial value, and a per-round price that is also agnostic to the initial value. This corresponds to at most a factor 6 approximation, that is stronger than what the theorems suggest.

2. **Normal:** When we move slightly away from the worst-case, to have initial values drawn from Normal distributions, with means ranging from 0 to 1, and standard deviation of 0.1, PPP already gets more than 21% of BIN revenue, using the same initial-value-agnostic price and free trial period as in the point mass case above.
3. **Uniform:** This percentage increases further to 32% as we move to the Uniform $[0, 1]$ distribution, again using the same price and free trial period as for the point mass case.