

Intertemporal Price Discrimination with Time-Varying Valuations

Victor F. Araman

Olayan School of Business, American University of Beirut, Beirut, Lebanon, va03@aub.edu.lb

Bassam Fayad

Institut Mathématique de Jussieu, Paris, France, bassam.fayad@imj-prj.fr

October 11, 2021

Abstract

A firm that sells a non perishable product considers intertemporal price discrimination in the objective of maximizing its long-run average revenue. We consider a general model of patient customers with changing valuations. Arriving customers wait for a random but bounded amount of time and purchase the product whenever its price falls below their valuation, which varies following a stochastic process. We show the optimality of a class of cyclic strategies and obtain an algorithm that yields them. When the pace of intertemporal pricing is constrained to be comparable to the upper bound on customers patience levels, we have a good control on the cycle length and on the structure of the optimizing cyclic policies. The results hold also for forward looking strategic customers as well as under an alternative objective of maximizing infinite horizon discounted revenue.

We cast our results in a general framework of optimizing the long-run average revenue for a class of revenue functions that we call *weakly coupled*, in which the revenue per period depends on a finite number of neighboring prices. We analyze in detail the case of Markovian valuations where we can obtain closed form formulations of the revenue function and some refinements of our main results. We also extend our results to the case of patience levels that are bounded only in expectation.

1 Introduction

Typically, customers try to estimate the value they will generate from a product and accordingly decide whether to purchase or not. This valuation exercise is quite complex as it could depend on many external factors and often done based on incomplete information on the product. As a result, it is not uncommon to see customers changing repeatedly their willingness-to-pay, sometimes drastically, before eventually making the purchase or deciding to drop it completely.

When customers become interested in a product, they often go through an *evaluation* phase, where they gather information about the product and reach a first estimate of their willingness-to-pay or valuation. They move next to the *purchasing decision* phase. Some customers would instantly purchase

the product if they find the current price adequate. Others simply delay their decision aiming for a better match in the future. Those customers will then be exposed to additional information, be it about the product or their personal state, that could affect their valuation and thus their purchasing decision. It is this latter phase that we are interested in modeling and analyzing in this work.

Consider for instance a customer who is looking into buying a new camera. The customer has an estimate of her willingness-to-pay based on a previous experience she had with that brand, some information gathered online, and the urgency to get the product. Given this valuation, she can decide to buy the camera or delay the purchase depending on how attractive the price is and whether she is expecting, say, a better deal in the near future. If she delays the purchase, she might indeed witness changes in the prices but possibly also in her own valuation of the product. These changes in the valuation could be due to a number of factors including, new product/customers-reviews highlighting a specific feature or a flaw, marketing campaigns (e.g., targeted online advertising), availability and prices of other competing products, and so on. The valuation can also be affected by the customer's personal information: e.g., a customer decides to go on an unplanned vacation and for whom getting the camera becomes more valuable; another customer gets a bonus at work and can now allocate a larger budget for the camera. All this information could be received sequentially feeding the decision process and resulting in a dynamic update of the valuation.

Such behavior has been generally overlooked in the pricing literature where most often customers are assumed to have a reservation price (possibly drawn from some given distribution) constant throughout the purchasing process. In this work, we consider customers interested in purchasing a product who see their valuations change through time. We try to shed some light on the impact that such varying valuations have on the seller's optimal pricing policy.

Intertemporal pricing is experienced in practice and often argued for as a way to price discriminate, and to take advantage of customers heterogeneity, scarcity of capacity, and demand stochasticity. In our case, we analyze this phenomenon in a specific context characterized by the following modeling choices. We consider a setting where a non perishable product is being sold throughout an infinite horizon during which the seller has *committed* to a sequence of prices. A continuous influx of heterogeneous customers get interested by the product and approach the seller at different times. They could purchase on arrival or decide to remain in the system for a (random) period of time during which their valuations *change* stochastically. We disregard any inventory and cost related issues and assume the product is available throughout. The fact that the firm commits to future prices is in line with the recent literature on intertemporal pricing with patient and/or heterogeneous customers (see, ? ?, ?, ?, ? and ?). Customers' heterogeneity is modeled through the dynamics of the stochastic valuation process governing each arrival.

The type of questions we are interested in are whether, in this context, there is an opportunity for price discrimination; and if there is one, how to go about characterizing it formally. We specifically examine the optimality of cyclic policies and how numerically efficient these policies are, especially since such cyclic policies have been proven to be optimal in similar settings in the literature.

Our main contributions are the following.

1. We introduce a very general and innovative model that incorporates time-varying valuations in a simple setting of intertemporal pricing for which we obtain structural results for the optimal pricing strategies.
2. In the case where we bound customers lifespan in the system, we cast our results in a general framework of optimizing the long-run average revenues for a class of revenue functions that we call *weakly coupled*, in which the revenue per period depends on a finite number of neighboring prices. Within this framework, we show the optimality of cyclic policies. We characterize such solutions and present an algorithm to obtain them.
3. We introduce a control parameter M for the firm that measures the number of price-changes that any customer may witness. Our algorithm yields the optimal pricing policies in polynomial time in the number of prices available, but exponential in M . By constructing an example where the optimal policy is cyclic, *increasing* and involves all but one of the available prices, we argue that, in general, in such changing valuation setting, the exponential complexity is unlikely to be reduced.
4. Our algorithm becomes fully relevant for small values of M . When customers can only see at most two different prices (i.e. $M = 1$), we show that the pricing policy is either a fixed price policy or cyclic and simple i.e. each price in the cycle is set only once for exactly a specified amount of time.
5. We analyze in detail the case of Markovian valuations where we can obtain closed form formulations of the revenue function and some refinements of our main results. Such Markovian model paired with a simple intertemporal pricing setting allows us to argue that the complexity of the problem can be traced primarily to the sole property of changing valuations. In realistic Markovian models, we confirm numerically that taking into account the consumer's changing valuation behavior may have a major impact on the pricing policy that the firm should adopt. In particular, overlooking such behavior can cause a sizable opportunity cost.
6. We establish the robustness of the optimality of cyclic policies by extending our results in three important directions : *i*) patience levels bounded only in expectation; *ii*) forward looking or strategic customers; *iii*) infinite horizon discounted revenue instead of long-run average revenue.

1.1 Literature Review

The literature on intertemporal pricing is quite extensive starting with the early work on durable goods (e.g., ?, ?, ?, ?, ? and ?). The work of ?, ? and ? are among the first to argue that price commitment policies in the presence of strategic customers are an effective way to exercise market power. In their model, the firm faces a population of customers that is present in the system from the

start. In contrast, others (e.g., ? , ? and ?) consider a firm that sets dynamically its price and face a continuous flux of strategic customers. The work of ? is the closest in this literature to our work. Customers in this paper have two possible valuations of the product, Low and High, and they remain in the system indefinitely until they purchase. Interestingly, ? show that the optimal pricing policy is both cyclic and decreasing.

More recently, the operations management literature have looked at different variants of intertemporal pricing models (see, the recent reviews of ? and ?). Often in this literature, customers are myopic and impatient. They arrive through time endowed each with a valuation drawn from a given distribution. Customers purchase on arrival if the price is higher than their valuation, otherwise they leave the system. Some of this literature incorporates customers' strategic behavior (see, the review of ?). ? and ? consider respectively a monopolistic and a set of firms selling to a set of differentiated customers segments and assume that all customers are in the system from the start. ? consider forward looking customers with declining valuations who arrive following a Poisson process and where the listed price changes only once. More closely related to our work is the paper of ? that considers a deterministic model where strategic customers are differentiated not only through their valuation, but also, through their patience level. Optimal prices are shown to be monotone. The very recent work of ? considers an intertemporal pricing problem under minimax regret where both strategic and myopic patient-customers are considered. Finally, two recent notes (? and ?) tackle the problem of intertemporal pricing in the presence of reference price effects in an infinite horizon setting with discounted revenue.

Besides that of ?, the papers that are the most related to our model are the recent works of ?, ? and ?. Their models are variants of ?, but, like us, consider a seller who commits to the pricing policy at the start of the horizon and maximizes the long-run average revenue. In the case of ?, heterogeneous customers, through their valuations and their patience level, are strategic. They prove that a cyclic pricing policy is optimal but, does not need to be decreasing. In an almost exactly similar setting, ? consider patient customers, who purchase as soon as the price drops below their willingness-to-pay. In their case, when the patience levels are deterministic and fixed, cyclic and decreasing policies are optimal. Finally, ? consider the exact setting of ? and suggest a dynamic programming algorithm that runs in polynomial time for finding optimal pricing policies under bounded but arbitrary distribution of the patience levels.

None of the above considers valuations that change through time. It is only recently that a handful of papers have tackled such behavior (see, ?, ? and ?). ? considers, similarly to us, an infinite horizon setting where customers arrive through time with valuations that stochastically change while facing a committed pricing path. There are though a number of differences with our work. First, the model is set in continuous time where both buyers and sellers have a common discount rate and where customers are strategic and homogeneous with a valuation governed by a Markov process that continuously switches between only two values, Low and High. The seller, is optimizing on the set of pricing paths (on the entire positive line). The optimal pricing path is shown to be cyclic decreasing

(à la ?). As opposed to ?, we consider in this paper a very general valuation process that is only required to be stationary. We do consider the case where the process is Markovian but we do not restrict it to take two values (except for the simplified model of Appendix B). Moreover, we assume that prices belong only to a finite and discrete set. Another recent paper that incorporates changing valuations is that of ?. The model considers one unit of a good and one buyer who is present in the system from the start and has a valuation drawn from some distribution. At some random time, a stochastic shock occurs and causes the buyer’s valuation to be drawn again (from the same initial distribution). Finally, we mention the paper of ? that models changing valuations in the context of revenue management. They consider a firm selling a finite number of units of a product that can be purchased anytime during a finite horizon and get consumed at the end of it. Customers are present at the beginning of the horizon and see their expected valuations changing through time. The firm commits to the prices and the customers need to decide when to purchase.

1.2 Outline of the Paper

In Section 2 we briefly introduce a Markovian model of changing valuations. The objective of this section is to highlight, through a simple setting of an intertemporal pricing problem, the complexity that the moving valuations property is bringing and compare that to the literature that focused on constant valuations. In Section 3, we introduce a very general model of changing valuations. For that we formulate an open loop optimization problem. We state the main results with respect to the optimality of cyclic policies, their structure, and how efficiently these can be obtained. In Section 4, we introduce the notion of *weakly coupled* revenue functions and show that the revenue function of a pricing policy under a general but bounded valuation model belongs to the class of weakly coupled functions. We next characterize the solutions of the optimization problem, defined in the previous section, when applied to this class of revenue functions. We also present an algorithm to obtain them. The proofs of our general results on optimal policies follow in a straightforward manner. In Section 5, we focus on the special case where the valuations processes follow discrete time Markov chains. In this context, we are able to obtain a closed form of the revenue function. We also show in this setting that the revenue function satisfies an affine structure which allows us to refine some of our main results. In Section 6, we discuss the complexity of the optimal policy. To support our claim that, in general, the optimal policy lacks any simplifying structure, we give two examples, one for general weakly coupled functions and one for the specific setting of Markovian valuations. Consequently, we argue that our results can prove to be valuable in obtaining efficient approximations of the optimal policy while preserving the changing valuations feature. In Section 7, we analyze some important extensions. First, we consider the case where the patience levels are bounded only in expectation and show, in the Markovian setting, that all the results hold when optimality is replaced with ε -optimality. Next, we discuss how the main results can be adjusted to the case where the firm is maximizing an infinite horizon discounted revenue function. Finally, we argue how our results extend to the case where customers are forward looking. We devote Section 8 for an extensive numerical analysis that

takes full advantage of the algorithm suggested earlier, and shows the opportunity cost of discarding the changing valuations property. We conclude in Section 9. In Appendix A, we relate the proofs of the Markovian model, in particular, when the revenue function is affine. In Appendix B, we analyze the setting where the firm can set only two prices and the valuation process is Markovian with unbounded patience levels. We specifically discuss in this setting the existence of a reset time.

2 A simple time varying valuation model

The objective of this section is to introduce a simple model of intertemporal pricing in an infinite horizon setting and highlight the complexity that the changing valuation feature of our model brings in. For this reason, we do not at this point precisely define some of the concepts (e.g., cyclic policy) but delay that to the next section where, we carefully introduce a general model, as well as the different definitions needed for the analysis, and state rigorously our main results.

Consider a stream of customers interested each in buying one unit of a non perishable product. At each period $t \in \mathbb{N}$, one customer arrives to the system. If customers don't purchase on arrival, we assume in this simple model that they all have a lifespan of τ periods, beyond the period of their arrival, to purchase the product. This window of time is known as the patience level. If customers don't purchase by the end of their lifespan, they leave the system (if $\tau = 0$, all customers are impatient and they either purchase on arrival or leave without purchasing).

Customers arriving to the system have a valuation $v \in \Omega = \{v_1, v_2, \dots, v_K\}$ with $K < \infty$ which is drawn following some distribution γ . During the "lifespan" in the system, each customer's valuation of the product changes following a simple Markov chain with a given transition matrix \bar{Q} on Ω .

The seller adopts a committed pricing policy, whereby at time 0, the prices are set for the entire horizon, $\pi = (p_1, p_2, \dots)$, in a way to maximize the long-run average revenue. As mentioned earlier, the fact that the firm commits to future prices is in line with the recent literature on intertemporal pricing and has been also shown to be the optimal strategy to follow (see, ? ?, ?, ?, ? and ?). We denote by $L_t(p_1, \dots, p_t)$ the revenue function, which is the expected revenues generated during the first t periods, having set the first t prices to be p_1, \dots, p_t . Let then $\mathcal{R}(\pi) = \limsup_{t \rightarrow \infty} \frac{1}{t} L_t(p_1, \dots, p_t)$ be the long-run average revenue generated by a policy π . WLOG we assume that the p_i 's belong to Ω . The seller is looking to solve for

$$\sup_{\pi \in \mathcal{P}} \mathcal{R}(\pi). \tag{1}$$

This simple model is a special case of the model considered in ? and ? but for these papers, customers keep the same valuation (drawn on arrival) during their lifespan, which is equivalent to setting \bar{Q} equal to the identity matrix, $\bar{Q} = \mathcal{I}$. In the case of ?, customers are strategic while in the case of ? customers are patient in the sense they are not forward looking but purchase as soon as their current valuation is larger than the current price. Whether customers are patient or forward looking the results happen to be of similar nature. We focus mainly on patient customers and tackle later in the extensions the case of forward looking customers.

2.1 Optimality of cyclic policies

As an example of the types of results obtained in the literature in the case of non varying valuations, we now recall, using our notations, the main result of ?.

Theorem 1 (?) *If $\bar{Q} = \mathcal{I}$ then there exists a decreasing cyclic policy of cycle size smaller or equal to $K + \tau - 1$ which can be obtained in polynomial in K and τ elementary computations.*

A similar result has also been obtained in ? except that they showed in their strategic setting that optimal policies are not necessarily decreasing but satisfy the so-called *refelction principle*. With our simple Markovian model for the valuation process, we obtain that

Theorem 2 *For a given transition matrix \bar{Q} , there exists a cyclic policy that optimizes \mathcal{R} . The cycle size of the policy is in the order of $\mathcal{O}(K^\tau)$. The policy can be obtained through an algorithm that requires $\mathcal{O}(K^{4\tau})$ elementary computations.*

First, we observe that our result confirms that cyclic policies remain optimal even under changing valuations. As for the cycle size and the complexity of finding the optimal policy, we note that for $\tau = 1$, our results and those of ? are comparable. However, for larger τ , Theorem 2 cannot guarantee a good performance of the algorithm that generates the pricing policy while in ?, the optimal pricing policy remains polynomial in τ . Next, we discuss these points in some detail.

2.2 The structure of optimal cyclic policies and the complexity of the optimization

The recent literature, as reflected in Theorem 1, has recognized that cyclic policies are optimal in some variants of our current setting as long as $\bar{Q} = \mathcal{I}$, (see also, ? and ?). Our main results and specifically Theorem 2 are confirming the optimality of cyclic policies even when $\bar{Q} \neq \mathcal{I}$. However in this context of time-varying valuations, the proof of our result requires a more complex approach.

In order to prove that optimal policies are cyclic and tractable, all the aforementioned papers relied on a regenerative point argument. Indeed, it is observed that for any pricing policy, the essential minimum price (the lowest price that appears infinitely many times in the policy) acts as a regenerative point. The optimization problem of choosing the prices in between each two appearances of the minimum price is the same, which yields the optimality of cyclic policies.

This regenerative point argument does not apply when valuations change through time, even in the simple Markovian setting of Theorem 2. Given a pricing policy, customers with lower valuations can possibly have higher valuations in the future and hence the regenerative points have no reason to exist *even* for the optimal policy.

The regenerative point argument also helps in showing some very interesting structure results on the optimizing cyclic policies in the constant valuations models. For instance, in Theorem 1 the optimal cyclic policies were shown to be decreasing. As a matter of fact, it is the decreasing property of the policy that allows ? to show that the optimal policy can be obtained in polynomial time. Recently, ? generalized this latter result to customers having constant valuations and different patience levels.

He observed that in this more general setting, the optimal cyclic policy may fail to be decreasing; it remains however made of subcycles, each of which is *decreasing*. In ? it is the reflection principle property that gives the tractability of the optimal cyclic solutions.

The next proposition shows how, even in the simplest Markovian setting, a changing valuation model does not allow optimal cyclic policies to inherit any of the *nice* structures of optimal policies obtained in the constant valuation case.

Proposition 1 *Let $\tau := 1$. For any $K \geq 5$, there exists $\mathbf{v} = (v_1, \dots, v_K)$, and $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_K)$ and a transition matrix \bar{Q} , such that the optimal solution is cyclic, increasing in its prices, and with a cycle size equal to $K - 1$.*

We will give the explicit example that proves Proposition 1 in section 6. The proposition shows that despite the fact that customers spend at most two periods in the system ($\tau = 1$), the optimal cyclic policy can have a cycle as large as $K - 1$. We also recall that in the case where both valuations are constant and all customers have one patience level, ? observed numerically and conjectured that under some broad conditions the cycle size is independent of K and depends only on the patience level. In the case of varying valuations, it seems that the optimal cyclic policy can be of any size.

The other surprising feature of the example of Proposition 1, which is in full contrast with the results in the literature pertaining to constant valuations, is that the prices in a cycle of an optimal policy are in fact increasing.

In conclusion, in the presence of stochastic valuations (even in the basic Markovian model), the optimal policies may fail to have any of the nice structures of optimal policies in the constant valuations models, that allowed to reduce the complexity of the optimization problem. With the lack of simplifying structure, it is hard to imagine how one can reduce the complexity of obtaining the cycle of the policy and thus, the curse of dimensionality seems inherent to this type of intertemporal pricing problems.

2.3 Weakly coupled revenues

An influential work on intertemporal pricing is ? that proved among others the optimality of cyclic policies for durable goods. The approach followed by ? is based on customers' accumulation type argument. The accumulation argument relies on the fact that customers with the lowest valuation who don't purchase remain in the system indefinitely, making it eventually worth resetting the system by setting for one period the price at the lowest valuation. This approach again does not apply in our case where the expected number of customers in the system remains uniformly bounded, preventing any accumulation of customers in the system. In Appendix B, in the case of $K = 2$ under a Markovian setting, we obtain necessary and sufficient conditions under which the seller is better off setting a reset price at the end of the cycle, showing therefore that cyclic policies are strictly optimal. When $K > 2$, it does not hold in general that optimal (or near optimal) cyclic policies contain a reset price.

To handle that, we introduce in Section 4 the concept of weakly coupled revenue functions in which the revenue per period depends on a finite number of neighboring prices. This property generalizes to some extent the concept of regenerative point, and will allow us to establish the optimality of cyclic policies, as well as to get an algorithm that detects them.

We now move to introduce a general model of time varying valuations and state the corresponding results.

3 The General Bounded Patience Level Model, or BP Model.

3.1 Model Description

We assume that at each period $t \in \mathbb{N}$, a random number $\xi_t^i \in \mathbb{N}$ of customers arrive to the system where $i \in \mathcal{I}$ indicates the customer's class, and \mathcal{I} is a finite set of integers. Customers either decide to purchase on arrival or else remain in the system. During the "lifespan" in the system, each customer's valuation of the product changes following a stochastic process. All customers have a *patience level*, which is a maximum window of time during which they are willing to wait to purchase, before they lose interest in the product. We assume that the patience level of all customers is bounded by some $\tau \in \mathbb{N} \setminus \{0\}$. Customers may also lose interest in the product before time τ and this is modeled by the existence of an absorbing state 0 for the valuation processes.

More precisely, our assumptions on the arrival process are the following:

- For any $i \in \mathcal{I}$, $\xi = (\xi_t^i : t \geq 0)$ is a stationary process in t , such that $\sup_{i \in \mathcal{I}} \mathbb{E} \xi_0^i < \infty$.
- For any $i \in \mathcal{I}$, $t \geq 0$, $1 \leq k \leq \xi_t^i$, the process $(\mathcal{V}_t^{i,k}(s) : s \geq 0)$ is the valuation process that takes values on the non-negative real line and is independently and identically generated for each customer k of class i that arrives at time t . Here s is the time after the instant t , with $\mathcal{V}_t^{i,k}(0)$ being the valuation on arrival. We define the sequence of processes $\mathcal{V}_t^i = (\mathcal{V}_t^{i,k}(\cdot) : 1 \leq k \leq \xi_t^i)$ and $\mathcal{V}^i \equiv (\mathcal{V}_t^i : t \geq 0)$ is assumed to be a stationary process (in t).
- There exists a positive integer preset constant τ such that for any $i \in \mathcal{I}$, $t \geq 0$, and $1 \leq k \leq \xi_t^i$, $\mathcal{V}_t^{i,k}(s) = 0$, for all $s \geq \tau$. Moreover, if $\mathcal{V}_t^{i,k}(s) = 0$, then for all $s' \geq s$, $\mathcal{V}_t^{i,k}(s') = 0$.

This general model will be denoted hereafter as the (BP) model, highlighting the Bounded Patience feature. In Section 5, we discuss in more detail the Markovian valuation case under bounded patience levels, denoted (M-BP) which was initially introduced in Section 2; in Section 7, we discuss the Markovian valuation case under *unbounded* patience levels, denoted (M-UP).

We assume for the main exposition that customers are patient and will purchase the product as soon as their valuation reaches a value higher than the current price. If that does not happen on arrival, or during their patience level, they do not purchase and leave the system. In the extension section (see, Section 7), we show that most of our results generalize to the case where customers are *strategic*.

On the supply end, we recall that the seller commits at the beginning of the horizon to an infinite sequence of prices that belong to a finite set of values Ω , where $\Omega := \{v_j : 1 \leq j \leq K\}$, with $0 < v_1 < v_2 < \dots < v_K$. To ease the notations we define the infinite price sequences $\pi = (p_t : t \geq 0)$ with $p_t \in [1, K]$ denoting the index of the price at time t . Hence, $p_t = k$ means that at time t the value of the price is set at v_k . We denote by \mathcal{P} the set of all possible pricing policies.

Observe that the primitives of the model are given by the pair $(\xi^i, \mathcal{V}^i)_{i \in \mathcal{I}}$. Now, given such a pair, and a policy $\pi = (p_1, p_2, \dots) \in \mathcal{P}$, we are again looking to solve for (1). We say that π^* is optimal if it is a solution to equation (1).

The results of this paper can also be extended beyond the long-run average revenue type-objective that the related literature has been focusing on. For example, we show in the extensions, Section 7, that most of our main results hold for the case of an infinite horizon discounted revenue function, where $\mathcal{R}(\pi)$ is of the form $\sum_{t=0}^{\infty} e^{-rt} L_t(p_1, \dots, p_t)$.

3.2 Main Results

We state in this section our main results on the optimality of a class of cyclic policies in solving problem (1), together with an algorithm that finds these optimal policies with a complexity that is logarithmically small compared to the *a priori* number of cyclic solutions in our class. Moreover, with an additional constraint on what we will define as the firm's *pricing pace*, these optimal solutions become effectively tractable.

We start with a few definitions and notations about cyclic policies and pricing pace, then we state our results in a unified way for the unconstrained and the pace-constrained cases.

Definition 1 (Cyclic policies)

A policy $\pi = (k_1, k_2, \dots)$ is called *cyclic* if there exists $n \geq 1$ such that $k_{j+n} = k_j$ for all $j \in \mathbb{N}$. When n is the minimal integer with the latter property, we say that (k_1, k_2, \dots, k_n) is the *cycle* of π , n its *size*, and denote π by its cycle $\pi = (k_1, \dots, k_n)$.

A fixed-price policy, one where the same price is set indefinitely, can also be viewed as cyclic with $n = 1$.

At this point, we introduce a way to express pricing policies that will be convenient for our exposition. Any pricing policy $\pi \in \mathcal{P}$ can be viewed as a sequence of *phases* of the form $\pi = ((k_j, \tau_j) : j \geq 1)$, where each $(k_j, \tau_j) \in \Omega \times \mathbb{N} \cup \{\infty\}$ is called a phase (of the policy) with k_j being the price of the j^{th} phase and τ_j its duration, which is the time during which the price is continuously set at k_j . For all $j \geq 1$, set $T_j := \tau_1 + \tau_2 + \dots + \tau_j$. Note that this expression of the policy is not unique unless if we impose that $k_j \neq k_{j+1}$. [Check \(AE\) Here it would be better to check and confirm that we can assume \$k_j \neq k_{j+1}\$ everywhere, besides the fixed price policy that we can present as an exception for this rule.](#)

We next introduce the notion of a pricing pace. This is a constraint on the prices that firms can set in order to regulate how frequently a price will change through time. The pace will be regulated by a constant $\sigma \geq 1$, defined as the minimal duration that the firm allows each price to be set for.

From a practical point view, the existence of such pricing pace seems quite natural due to the direct and indirect costs incurred by frequently changing prices, whether from the logistical end, or, as a reflection of customers' dissatisfaction. As a result of such pace, the set of admissible policies will be reduced.

Definition 2 (σ -policies) For any $\sigma \geq 1$, we denote by \mathcal{A}_σ the set of policies where each phase has a duration that is a multiple of σ . We can always write such policy as a sequence of σ -phases of the form $\pi = ((k_j, \sigma) : j \geq 1)$, where k_{j+1} can be equal to k_j . Policies in \mathcal{A}_σ will be called σ -policies.

When the pricing pace is set to σ , and if we restrict ourselves to policies where prices are set in multiples of σ , the firm's problem becomes a constrained optimization over \mathcal{A}_σ

$$\sup_{\pi \in \mathcal{A}_\sigma} \mathcal{R}(\pi). \quad (2)$$

Note that the special case of $\sigma = 1$ corresponds to the unconstrained optimization on \mathcal{P} , as initially defined in (1).

We define the quantity $M := \lceil \tau/\sigma \rceil$. By setting a minimum pace σ , the firm can control the maximum number of price-changes, M , customers could witness during their lifespan in the system. For example, if σ is set so that $M = 2$, customers will see one price on arrival and utmost two *additional* prices during their patience level. Up to enlarging τ if necessary, we will assume all through the paper that

$$\tau = M\sigma.$$

In the unconstrained case $\sigma = 1$, and hence $M = \tau$.

Before we state our results we introduce the notion of M -simple cyclic policies.

Definition 3 (M -simple) For $M \in \mathbb{N} \setminus \{0\}$, a cyclic policy $\pi = ((k_j, \sigma) : j \geq 1) \in \mathcal{A}_\sigma$ is said to be M -simple if any string of M prices $(k_{j+1}, \dots, k_{j+M})$ corresponding to M consecutive σ -phases of the policy appears at most once during a cycle. For $M = 1$, we just say the policy is simple (i.e., all possible values of the prices appear at most one time in a cycle during exactly σ consecutive periods).

Theorem 3 There exists an M -simple cyclic σ -policy that optimizes \mathcal{R} on \mathcal{A}_σ . In particular, there exists a τ -simple cyclic policy that optimizes \mathcal{R} on \mathcal{P} .

We move now to discuss how computationally meaningful this result is.

Observe that an M -simple policy in \mathcal{A}_σ has a cycle of length at most K^M σ -phases. For $i \in \{1, \dots, K^M\}$, the number of M -simple policies in \mathcal{A}_σ with i σ -phases is of the order of $K^M(K^M - 1) \dots (K^M - i + 1)$.¹

¹I think it is $((K^{M-1})!)^{\frac{1}{M}}$. Indeed, we have K^M choices to start with an M string call it W_1 . Then we have to pick among the remaining $K^M - 1$ M -strings a string W_2 that is compatible with our choice. The number of such strings should be larger than K^{M-1} (to be checked). Hence the choice of W_1W_2 gives $K^M K^{M-1}$ possibilities. The choice of W_3

Summing over all $i \in \{1, \dots, K^M\}$, we get that the total number of M -simple strategies in \mathcal{A}_σ is of the order of $(K^M)!$, and by Stirling formula is actually roughly of the order of K^{MK^M} . The next result shows that this *a priori* double-exponential complexity can be logarithmically reduced by an adequate algorithm that computes an optimal strategy of Theorem 3.

Theorem 4 *The optimal cyclic policies of Theorem 3 can be obtained, through an algorithm that requires $\mathcal{O}(K^{4M})$ elementary computations.*

Theorems 3 and 4 will be proved in the next two sections by relying on Proposition 2 and respectively, Proposition 4 and Corollary 1.

Theorem 4 shows that the computational complexity of obtaining the optimal policy is affected polynomially in the number of possible prices K . Hence, for small values of M the problem is tractable, while for M large, despite a major reduction, it remains exponentially complex. We believe that in the context of changing valuations one cannot reduce much further the order of complexity. We revisit this claim in Section 6.

Next, we cast our main results in a general framework of optimizing the long-run average revenue for a specific class of revenue functions.

4 Weakly Coupled Revenue Functions

In this section, we show that the revenue function corresponding to a (BP) model belongs to a general class of revenue functions that we call weakly coupled. In the latter framework, it is shown that the optimization problems formulated in (1) and (2) admit cyclic policies. We also provide an algorithm that detects the optimal cyclic policies. With respect to the latter, finding the optimal cyclic policy is similar to calculating minimum mean cycles on graphs and efficient algorithms of calculating such minima were obtained in several works (e.g., ?). However, we stay in the context of optimizing long-run average revenue and propose our own algorithm for detecting the optimal cyclic policies that essentially has [the minimal possible complexity](#).

We first introduce a general class of revenue functions that we call weakly coupled revenue functions in which the revenue per period depends on a finite number of neighboring prices.

4.1 Definition of Weakly coupled functions

For a pricing strategy $\pi \in \mathcal{P}$ we define a class of long-run average revenue called weakly coupled.

: must be compatible with W_2 , that gives K^{M-1} choices, but must also avoid all the M strings contained in W_1W_2 that is M . Hence the choice of $W_1W_2W_3$ gives at least $K^M K^{M-1}(K^{M-1} - M)$. Continuing like this we get for $W_1W_2 \dots W_i$: $K^M K^{M-1}(K^{M-1} - M) \dots (K^{M-1} - iM)$ that we can divide by i to eliminate the permutations. Taking $i = K^{M-1}/M$ we get the maximal contribution that is of the order of $((K^{M-1})!)^{\frac{1}{M}}$, that is K^{K^M}

Definition 4 (Weakly coupled revenue functions) We say that a function $\mathcal{R} : \mathcal{P} \rightarrow \mathbb{R}_+$ is a τ -weakly coupled revenue function for some positive τ , if there exists a function $f : \Omega^\tau \times \Omega^\tau \rightarrow \mathbb{R}^+$, such that for any policy of the form, $\pi = ((w_n : n \geq 1) : w_n \in \Omega^\tau)$, the following holds

$$\mathcal{R}(\pi) = \frac{1}{\tau} \varphi(\pi)$$

$$\varphi(\pi) := \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n f(w_i, w_{i+1}).$$

4.2 The Weakly Coupled property of Revenue Functions in (BP) Models

We show next that the revenue functions that appear in the context of (BP) models introduced in Section 3.1 are weakly coupled.

Proposition 2 Under a (BP) model with patience level τ , the revenue function is τ -weakly coupled.

Proof. WLOG we can restrict ourselves to valuation processes that take values in Ω . Let us fix a pricing policy that we write as $\pi = ((w_j : j \geq 1) : w_n \in \Omega^\tau)$. The prices given by the vector w_j are set during the j^{th} time phase $[(j-1)\tau + 1, j\tau]$. Given a realization of the processes $(\xi_t^i, \mathcal{V}_t^i)_{i \in \mathcal{I}}, t \geq 0$, let us denote by \hat{L}_j the revenues generated during the phase w_j . For $j \geq 2$, the revenues \hat{L}_j is the result of purchases that occur from either customers who arrive to the system during the j^{th} phase and who buy during this same phase, or from customers who arrived to the system earlier and who are still in the system because they did not buy and did not reach the absorbing exit state, v_0 . Due to the fact that the patience level is bounded by τ , the latter customers all come from the $j-1^{\text{th}}$ phase (during which prices w_{j-1} are set). Hence \hat{L}_j is a function of (w_{j-1}, w_j) and the realizations of $(\xi_t^i, \mathcal{V}_t^i)_{i \in \mathcal{I}}$, for time $t \in [(j-2)\tau + 1, j\tau]$.

Since the arrival process $(\xi_t^i, \mathcal{V}_t^i)_{i \in \mathcal{I}}$ is assumed to be stationary in t , it follows that the expectation L_j of \hat{L}_j over all realizations of $(\xi_t^i, \mathcal{V}_t^i)_{i \in \mathcal{I}}$ is a deterministic function $f(w_{j-1}, w_j)$.

It then follows that

$$\mathcal{R}(\pi) = \limsup_{n \rightarrow \infty} \frac{1}{n\tau} \sum_{j=1}^n f(w_j, w_{j-1})$$

is a τ weakly-coupled function of the pricing policies. \square

4.3 Optimizing Weakly Coupled Revenue Functions

We consider a τ -weakly coupled revenue function following Definition 4. We are interested in finding the policy that maximizes φ . We first extend the definition of φ to finite strings $W = (w_1, \dots, w_n)$ as follows:

$$\varphi(W) = \frac{1}{n} \sum_{i=1}^n f(w_i, w_{i+1}), \quad (3)$$

with the notation $w_{n+1} = w_1$. Note that $\varphi(W) = \varphi(\pi)$, where $\pi = (W, W, \dots)$. When $W = (w_1, \dots, w_P)$ is the shortest period of π we write $\pi = (w_1, \dots, w_P)$.

Proposition 3 (*General optimization*). *There exists a cyclic $\pi^* \in \mathcal{P}$ such that*

$$\varphi(\pi^*) = \sup_{\pi \in \mathcal{P}} \varphi(\pi)$$

Moreover, $\pi^* = (w_1, w_2, \dots, w_P)$, $w_i \in \Omega^\tau$ and all the w_i 's are distinct, with $P \leq K^\tau$.

Proof. Denote by \mathcal{W}_n the set of strings of length less than n , that is $W \in \mathcal{W}_n$ if $W = (\omega_1, \dots, \omega_l)$ with $\omega_i \in \Omega^\tau$ for all $i \leq l$ and $l \leq n$. We have that $\varphi(W) = \varphi(\pi(W))$, where $\pi(W) = (W, W, \dots)$. Let $\bar{W}_n = \operatorname{argmax}_{W \in \mathcal{W}_n} \varphi(W)$. From the definition of φ , it follows that for any $\pi \in \mathcal{P}$,

$$\varphi(\pi) \leq \varphi(\bar{W}_n) + \frac{\max f}{n}.$$

Hence it is sufficient to prove that for any fixed n , the maximizer \bar{W}_n can be taken to be simple. Suppose ω is such that $\bar{W}_n = (W, W')$ with W and W' starting with some $\omega \in \Omega^\tau$, and let N and N' be the sizes of W and W' . Then since

$$\varphi(\bar{W}_n) = \frac{1}{N + N'}(N\varphi(W) + N'\varphi(W'))$$

we get that $\varphi(W) = \varphi(W') = \varphi(\bar{W}_n)$. Continuing this procedure we reach a simple maximizer of φ on \mathcal{W}_n . From that we also conclude that

$$\varphi(\pi) \leq \varphi(\bar{W}),$$

where \bar{W} is the maximizer of φ on all cyclic policies where the elements of the cycle are all distinct. □

Since we are also interested in the constrained optimization on \mathcal{A}_σ , we state the following proposition, which is immediately obtained from the proof of Proposition 3 by restricting the pricing policies to those in \mathcal{A}_σ .

Proposition 4 (*Optimization on \mathcal{A}_σ*). *For any $\tau \geq 1$, for any function $f : \Omega^\tau \times \Omega^\tau \rightarrow \mathbb{R}^+$ we have the following. There exists an M -simple cyclic $\pi^* \in \mathcal{A}_\sigma$ such that*

$$\varphi(\pi^*) = \sup_{\pi \in \mathcal{A}_\sigma} \varphi(\pi).$$

By recalling Proposition 2, the previous proposition completes the proof of Theorem 3.

We observe here that the proof of the optimality of cyclic policies does not rely on the notion of regenerative points (compared to ? or ?), but takes full advantage (see, proof of Proposition 3) of the defining expression of weakly coupled functions.

4.4 Optimization Algorithm

The objective of this section is to present an algorithm to find an optimal cyclic policy π^* as in Proposition 3.

The algorithm is the result of a greedy-type construction that generates an optimal solution. We first fix two consecutive strings of prices w_1 and w_2 , and set the third one in a way to insure that the second string w_2 is the maximizer of the revenue generated by these three strings in the first two periods: $f(w_1, w_2) + f(w_2, w_3)$. Once w_3 (not necessarily unique) is identified, we move to find w_4 , given w_2 and w_3 . The complexity of the algorithm is reduced in particular thanks to the “simple” property of the pricing policy.

For a finite string $W = (w_1, \dots, w_n)$, we denote by

$$\tilde{\varphi}(W) = \frac{1}{n-1} \sum_{i=1}^{n-1} f(w_i, w_{i+1}).$$

We define the following map from $\Omega^\tau \times \Omega^\tau$ to the subsets of Ω^τ ,

$$\psi(w, w') = S; \quad S = \{\bar{w} \in \Omega^\tau : \varphi(w, w', \bar{w}) = \max_{\tilde{w} \in \Omega^\tau} \varphi(w, \tilde{w}, \bar{w})\}.$$

A collection \mathcal{W}_n of finite simple strings of the same size n is called *admissible* if for all $W \in \mathcal{W}_n$, it is of the form (w_1, \dots, w_n) with the same $w_1 := a$; moreover, if $(W, W') \in \mathcal{W}_n^2$ such that $w_n = w'_n$, then $W = W'$. Set $\tilde{K} = K^\tau$.

Algorithm.

STEP 1. INITIALIZATION. Fix $(w_1, w_2) \in \Omega^\tau \times \Omega^\tau$, with $w_1 \neq w_2$, let $\mathcal{W}_2(w_1, w_2) = \{(w_1, w_2)\}$, $\mathcal{S}_2(w_1, w_2) = \emptyset$ and set $n = 2$

STEP 2. COMPUTATION. Given an admissible simple collection \mathcal{W}_n with $n < \tilde{K}$ and $\#\mathcal{W}_n \leq \tilde{K}$, define an admissible collection $\mathcal{W}_{n+1} = \Psi(\mathcal{W}_n)$ in the following way

1. for each $W^i \in \mathcal{W}_n := (w_1^i, \dots, w_n^i)$, consider $S^i = \psi(w_{n-1}^i, w_n^i) := \{s_1^i, \dots, s_{J_i}^i\}$ where, $J_i \leq \tilde{K}$ depends on the pair (w_{n-1}^i, w_n^i) . Consider all the strings

$$\{W_j^i = (w_1^i, \dots, w_n^i, s_j^i) : 1 \leq j \leq J_i, 1 \leq i \leq \#\mathcal{W}_n\}.$$

2. Select from these strings those such that $s_j^i = w_1$. This set of strings (possibly empty) is denoted by \mathcal{S}_{n+1}
3. Eliminate from the remaining all the strings W_j^i 's that are not simple
4. From those remaining, consider any two strings, such that $s_j^i = s_{j'}^{i'}$ for some (i, j) and (i', j') . If $\tilde{\varphi}(W_j^i) < \tilde{\varphi}(W_{j'}^{i'})$ or $\tilde{\varphi}(W_j^i) > \tilde{\varphi}(W_{j'}^{i'})$, then eliminate the string that has the smaller value of $\tilde{\varphi}$. If $\tilde{\varphi}(W_j^i) = \tilde{\varphi}(W_{j'}^{i'})$, we eliminate (randomly) one of them. We do that for all such strings. The remaining strings must define an admissible collection of simple strings which we denote by \mathcal{W}_{n+1} and note that the cardinality $\#\mathcal{W}_{n+1} \leq \tilde{K}$

STEP 3. ITERATION. If $n = \tilde{K}$, STOP. Otherwise, set $n = n + 1$ and GO TO Step 2.

We then make the following elementary observation.

Lemma 1 *If the cycle of an optimal strategy of Proposition 3 is $W := (w_1, w_2, \dots, w_P)$, for some integer $P \geq 2$, then $W \in \mathcal{S}_P(w_1, w_2)$.*

Given that each computation step generates simple strings, and by noticing that the first part of Step 2 can be done offline, we need at most $\mathcal{O}(\tilde{K}^2)$ computations to obtain $\mathcal{W}_n(w_1, w_2)$ and $\mathcal{S}_n(w_1, w_2)$, $n \leq \tilde{K}$. This algorithm takes as an input, w_1 and w_2 and hence, needs to be repeated for all possible values of w_1 and w_2 . As a consequence of Lemma 1, if we consider

$$\mathcal{S} := \bigcup_{(w_1, w_2), w_1 \neq w_2} \bigcup_{n \leq \tilde{K}} \mathcal{S}_n(w_1, w_2)$$

and the set of single value strategies

$$\mathcal{S}_0 := \bigcup_{w \in \Omega^\tau} \{(w)\}$$

then an optimal strategy W of Proposition 3 satisfies $W \in \mathcal{S}$. Note that the computation of \mathcal{S} needs at most $\mathcal{O}(\tilde{K}^4)$ calculations of the type $\varphi(w, w', w'')$. We have thus proved the following.

Proposition 5 *If a revenue function is τ -weakly coupled, then an optimal cyclic policy of Proposition 3 can be obtained with $\mathcal{O}(K^{4\tau})$ elementary calculations.*

2

As for the constrained optimization, the immediate corollary of Proposition 5 is the following.

Corollary 1 *If a revenue function is τ -weakly coupled, then any optimal M -simple cyclic σ -policy of Proposition 4 can be obtained with $\mathcal{O}(K^{4M})$ elementary calculations.*

Given Proposition 2, the previous corollary completes the proof of Theorem 4.

5 Markovian Models of Time-Varying Valuations

In this section, we revisit the model introduced in Section 2, where we specifically restrict ourselves to the class of (BP) valuations processes that follow a Markov chain. This special and important class allows us to achieve a number of things. First, we obtain an explicit expression of the revenue function. This is particularly helpful in the numerical study of Section 8. Besides retaining the weakly coupled

²We can put the following remark here : Observe that in step 2, when finding all admissible strings that start with (ω_1, ω_2) we get actually the informations on how other couples such as (ω_2, ω_3^1) can be completed in admissible strings. If one takes into account this information, the algorithm here presented can be slightly altered to give a complexity of the order of \tilde{K}^3 which corresponds to the minimal complexity known for such problems.

property, the revenue function is shown to be affine in the duration of a phase and in the expected number of customers available initially. As a result, we obtain some refinements of the main results of Section 3.2. Finally, this Markovian setting allows us also to extend our analysis to the case where the patience levels are bounded only in expectation (see, Section 7).

5.1 Model Ingredients

Similarly to Section 2, and in order to simplify the presentation, we consider a constant number of customers arriving every period. WLOG we assume that this number is one. We restrict our analysis to a single class of customers since the revenue function of a multiple class model is just a linear combination of revenue functions of single class models, and since the additional properties that we want to show for the revenue function of the Markovian model are stable under linear combination.

We recall that every arrival is endowed with an initial valuation $\omega \in \Omega \equiv \{v_j : 1 \leq j \leq K\}$ with a distribution given by a probability vector $\gamma = (\gamma_1, \dots, \gamma_K)$.

We also assume that the valuation process of the customer arriving at time t , $(\mathcal{V}_t(s) : s \geq 0)$ (where we dropped the superscripts i and k), follows a Markov chain on Ω , determined by his arrival valuation and by a transition matrix that we denote by \bar{Q} . Each entry q_{ij} of \bar{Q} , is the probability that a customer with current valuation v_i sees his valuation changing to v_j in the next period. Recall that $v_0 := 0$, is an absorbing state that once reached indicates that the client has lost interest in purchasing. Hence, we take the first line of \bar{Q} to be all zero except for its first entry that is equal to 1. From now on we will refer to the transition matrix by its minor Q obtained by removing from \bar{Q} the first line and the first column. This model allows customers to have different life spans, but we recall that there exists a maximum patience level τ so that $\mathcal{V}_t(s) = v_0 = 0$ for every $s \geq \tau$.

In this Markovian setting, besides τ , the primitives of the problem are summarized by (γ, Q) . With some abuse of notations, the pair (γ, Q) will also be used to represent the process \mathcal{V}_t .

When $\tau < \infty$, we call our Markovian model (M-BP), for *Markovian, Bounded Patience*. The (M-BP) model satisfies the assumptions of the general setting discussed in Section 3, that is: it is a (BP) model. As a consequence, its revenue function is weakly coupled and all the results that we obtained for (BP) models hold for these Markovian models.

In contrast, we will refer to the case where $\tau = \infty$, as the (M-UP) model for *Markovian, Unbounded Patience*. We will see in Section 7 how the analysis of the (M-BP) yields approximate solutions to the (M-UP) case.

5.2 Revenue Pairs

We start this section by defining a revenue pair that will be helpful in calculating the revenue function associated to the (M-BP) model.

We say that a vector θ denotes the expected number of customers in the system at some given time, when the coordinates of θ , θ_m , $1 \leq m \leq K$, are equal to the expected number of customers in

the system with valuation v_m at this time. Note that the number of customers in the system at any point in time is bounded by τ .

Let $\mathcal{M} = \Omega \times \mathbb{N} \times [0, \tau]^K$.

Definition 5 (Revenue pair) Let $\boldsymbol{\theta}$ be a vector determining the expected number of customers in the system at the beginning of a phase (k, t) , we define

$$\begin{aligned} \Theta : \quad \mathcal{M} &\rightarrow [0, \tau]^K \\ (k, t, \boldsymbol{\theta}) &\mapsto \boldsymbol{\theta}' = \Theta(k, t | \boldsymbol{\theta}), \end{aligned}$$

where $\boldsymbol{\theta}'$ is the vector measuring the expected number of customers in the system at the end of the phase (k, t) .

Let

$$\begin{aligned} L : \quad \mathcal{M} &\rightarrow \mathbb{R}_+ \\ (k, t, \boldsymbol{\theta}) &\mapsto L(k, t | \boldsymbol{\theta}), \end{aligned}$$

where $L(k, t | \boldsymbol{\theta})$ is the total expected revenues generated during phase (k, t) .

It is worth stressing that $\boldsymbol{\theta}'$ accounts (in expected value) for both those customers included in $\boldsymbol{\theta}$ and who did not exit by time t , as well as those who arrived during the phase (k, t) and did not exit yet. Similarly, $L(k, t | \boldsymbol{\theta})$ are the expected revenues generated from either customers included in $\boldsymbol{\theta}$ or, from those who arrived to the system during the phase (k, t) . Either way, these revenues must have been also generated during time t .

The fact that the expected number of customers in the system, with valuation v_m , at the end of the phase (k, t) , depends only on (k, t) and $\boldsymbol{\theta}$, is a consequence of : *i.*) the Markovian feature of the valuation process that guarantees that the evolution of the customers included in $\boldsymbol{\theta}$ does not depend on the past, and *ii.*) the stationarity of the processes (γ, Q) , (i.e., being independent of t). For these same reasons we can also define the expected revenue as a function of (k, t) and $\boldsymbol{\theta}$.

We move now to express the revenue function *via* the pair (Θ, L) . Since we are interested in the optimization of the long-run average revenue, we may assume WLOG that the system starts empty at $t = 0$.

Proposition 6 Let $\pi = ((k_j, t_j) : j \geq 1)$. The revenue function is given by

$$\begin{aligned} \mathcal{R}(\pi) &= \limsup_{n \rightarrow \infty} \frac{1}{T_n} \sum_{j=1}^n L(k_j, t_j | \boldsymbol{\theta}^{j-1}), \\ \boldsymbol{\theta}^j &= \Theta(k_j, t_j | \boldsymbol{\theta}^{j-1}). \end{aligned}$$

with $T_n = \sum_{i=1}^n t_i$.

Proof. Fix a pricing policy $\pi = ((k_j, t_j) : j \geq 1)$. For a realization through time of the arrival and the transition processes $(\mathcal{V}_t(s) : s \geq 0)$, we define $\hat{\theta}^j$ to be the realized number of customers in the system at time T_j . As in Definition 5 we have that the $\hat{\theta}^j$ and the realized revenue during the j^{th} phase are functions of $\hat{\theta}^{j-1}$ and the realization of the arrival and the transition processes of customers during this phase. Considering all the possible realizations of (γ, Q) for $0 \leq t \leq T_j$, we find that the expected value of $\hat{\theta}^j$ is exactly the sequence $\theta^j = \Theta(k_j, t_j | \theta^{j-1})$ and that the expected value of the revenue generated during the j^{th} phase is $L(k_j, t_j | \theta^{j-1})$, from which Proposition 6 immediately follows. \square

5.3 The (M-BP) Model. Main Results

Our first result shows that under the (M-BP) model the revenue function satisfies the so-called τ -affine property. We use the notation $\langle \cdot, \cdot \rangle$ to denote the scalar product of two vectors.

Proposition 7 (Affine property) *Under (M-BP) and given (γ, Q, τ) , the following holds. For any price $1 \leq k \leq K$, there exist $\bar{\Theta}_k \in [0, \tau]^K$, $A_k, B_k \in \mathbb{R}^+$, and $\mathbf{C}_k \in \mathbb{R}^K$ such that for any $t \geq \tau$, and for any $\theta \in [0, \tau]^K$, we have that*

$$(A1) \quad \Theta(k, t | \theta) = \bar{\Theta}_k,$$

$$(A2) \quad L(k, t | \theta) = A_k + B_k(t - \tau) + \langle \mathbf{C}_k, \theta \rangle.$$

In this case, the pair (Θ, L) is said to be τ -affine. The parameters A_k, B_k, \mathbf{C}_k and $\bar{\Theta}_k$ can be given in closed form.

The proof of the affine property is postponed to Appendix A.

This previous property not only allows one to obtain an explicit expression of the revenue function but also leads to some refinements of the main results of Section 3.2.

While \mathcal{A}_σ is the set of policies where each phase is a multiple of σ , we introduce \mathcal{P}_σ^+ as the set of policies where each phase has a duration larger than σ . Clearly, $\mathcal{A}_\sigma \subseteq \mathcal{P}_\sigma^+$.

Proposition 8 *Under (M-BP):*

- i.) One can find an optimizing strategy of (1) that is either a fixed price strategy or during which no price is set for more than τ periods.*
- ii.) If $\sigma \geq \tau$, there exists a simple cyclic σ -policy that optimizes \mathcal{R} on \mathcal{P}_σ^+ . Moreover, there exists a threshold σ_0 such that, if $\sigma \geq \sigma_0$, then a fixed-price policy is optimal on \mathcal{P}_σ^+ .*

Proof.

- i.) Suppose that the optimum cannot be reached by a fixed price policy. Let $\pi^* = ((k_1, \tau_1), \dots, (k_n : \tau_n))$ be a cyclic optimal strategy. We will show by contradiction that $\tau_1 \leq \tau$. Suppose by*

contradiction : repetition that $\tau_1 > \tau$. Replace π^* by $\pi_s = ((k_1, s), (k_2, \tau_2), \dots, (k_n, \tau_n))$ with $s \in [\tau, \infty)$. Observe that

$$\mathcal{R}(\pi_s) = \frac{1}{s + \sum_{i=2}^n \tau_i} (U + V(s - \tau)),$$

where U and V do not depend on the value of $s \geq \tau$. It follows that the derivative of $\mathcal{R}(\pi_s)$ has a sign that does not depend on s . The sign must be negative otherwise the optimum could be reached by the fixed price policy with price k_1 . Hence $\mathcal{R}(\pi_s)$ for $s \in [\tau, \infty)$ reaches its optimum at $s = \tau$ a contradiction with $\tau_1 > \tau$. Since the choice of which phase is the first one in a cyclic strategy is arbitrary it means that all the τ_i are less than τ .

ii.) The same argument as that used in the proof of Proposition 8 *i.)*, implies that for $\sigma \geq \tau$ the optimum of \mathcal{R} on \mathcal{P}_σ^+ is reached on a fixed price policy or on a σ -policy $\pi = ((k_j, \tau_j) : j \geq 1)$ with $\tau_j = \sigma$ for all $j \in \mathbb{N}$. By Proposition 7 4, an optimal σ -policy can be taken to be simple cyclic (**in particular $k_{j+1} \neq k_j$ unless it is a fixed price policy**). Let $\pi_c = (k_1, \dots, k_n)$ (where each k_i is set for σ periods) be an optimal policy.

From the definition of affine revenue functions, we have that

$$\mathcal{R}(\pi_c) = \frac{1}{n\sigma} \sum_{i=1}^n A_i + B_i(\sigma - \tau) + T_{i,i+1}$$

where A_i and B_i are functions of k_i and $T_{i,i+1} = \langle \mathbf{C}_i, \bar{\Theta}(i+1) \rangle$ is a function of (k_i, k_{i+1}) (and not of σ) and where $k_{n+1} := k_1$. Observe also that B_i is the long-run average revenue of the fixed price policy with price k_i .

As $\sigma \rightarrow \infty$, it is clear that the fixed price policy with price k such that $B_k = \max_{k'} B_{k'}$ becomes an increasingly better approximation of an optimal policy in \mathcal{A}_σ . In the generic case where $\max_{k'} B_{k'}$ is a strict maximum for $k' = k$, then as soon as $\sigma \geq \sigma_0$, where σ_0 is such that for any simple cyclic vector (k_1, \dots, k_n) , we have

$$B_k(n\sigma_0 + \tau - \sum_{i=1}^n \frac{B_i}{B_k}(\sigma_0 - \tau)) > \sum_{i=1}^n A_i + T_{i,i+1},$$

then the fixed-price policy with price k is strictly optimal on \mathcal{P}_σ^+ . □

The previous proposition brings indeed some refinements to Theorem 3. In the unconstrained case, since an optimal strategy is τ -simple and cyclic, we already know that no single price is set during more than 2τ periods. However, the affine structure of the revenue reduces this bound to τ periods. Moreover, recall that Theorem 3 characterized an optimal policy in \mathcal{A}_σ . When $\tau \geq \sigma$ part *ii.)* of the previous proposition generalizes the statement of Theorem 3 to \mathcal{P}_σ^+ . Finally, note that *ii.)* of this proposition claims also that if the pricing pace is too slow (σ large), then a fixed-price policy is optimal.

6 Absence of Structure and Complexity

We devote this section to complement our discussion on the complexity of the solutions obtained in Theorem 3 and Proposition 3. We give respectively two examples, one for the general setting of weakly coupled revenue function and one for the specific model of changing valuation, that show some unexpected behavior of the optimal policy. These examples confirm that one cannot expect to reduce the optimization (feasible) set through some simplifying structure of the optimal policy.

6.1 Example of an optimum along an arbitrary cyclic strategy for weakly coupled revenue functions.

We revisit here the general setting of weakly coupled revenue functions. Despite the results obtained in Section 4, the question remains of whether one could reduce the complexity of such cyclic policy even further. The next very simple example shows that in general this is not possible, in particular, any cycle of any size can be an optimal solution of an adequately chosen weakly coupled revenue function.

Example 1 Let $\bar{\pi} = (w_1, \dots, w_n)$ be a finite simple cyclic strategy. Fix $U > 0$, and define f such that

$$f(w, w') = \begin{cases} U & \text{if } (w, w') = (w_i, w_{i+1}) \text{ for some } i \in [1, n] \\ 0 & \text{if not,} \end{cases}$$

where we used the notation $(w_n, w_{n+1}) := (w_n, w_1)$.

The following is straightforward from the definition of f . We have

$$\varphi(\bar{\pi}) = U,$$

and for any strategy $\pi \neq \bar{\pi}$, we have that

$$\varphi(\bar{\pi}) - \varphi(\pi) \geq \frac{1}{K}U.$$

6.2 Example of an M-BP model with an optimizing strategy that is increasing with length $K - 1$.

In this section we construct the example that proves Proposition 1. This proposition confirms that due to the changing valuations feature of the model, one can construct an optimal cyclic policy that is increasing and involves all but one of the available prices.

The idea behind this example is the following. First, we standardize the return of any constant price strategy and set it equal to 1 by carefully choosing the parameters γ and \mathbf{v} . Next we choose the transition probabilities $q_{i,j}$ as follows. We say that a customer is in state i if her current valuation is

v_i . We set $q_{i,j} = 0$, from any state $i \in [1, K-2]$ to any state $j \leq K$, except for $j = i+2$. We denote by $(i(i+1))$ the expected revenue generated by all customers with valuation strictly lower than v_i , who enter at a date where the price is i and see the price $i+1$ at the subsequent date. We set $q_{i,i+2}$ so that $(i(i+1))$ is a positive quantity $U \ll 1$ that does not depend on $i \in [2, K-1]$. Since the transition probabilities from $i-1$ to states that are strictly higher than $i+1$ are zero, one observes that whenever a price i appears in the strategy, it is optimal to follow it with the price $i+1$, otherwise, there is a shortfall of U in the revenue. Finally, the transition probability, $q_{K-1,j}$, is set to zero for all states j , except for state $j = 2$; while the transition $q_{K-1,2}$ is chosen sufficiently large so that the expected revenue $(K, 2)$ is of order $\frac{11}{10}U$. The latter forces an optimal strategy to contain the string $(K, 2)$. Then, as discussed above, 2 must be followed by 3, then 4, etc., until K , which determines the optimal strategy as $(2, 3, \dots, K-1, K)$.

In terms of how to determine the values of \mathbf{v}, γ and Q , these are designed so that the key properties (E0)–(E2), defined below, hold. These properties are sufficient to prove the optimality of the strategy $\bar{\pi} = (2, \dots, K)$. We first set \mathbf{v} and γ such that (E0) and (E1) hold. Then the first $K-2$ rows of Q are chosen so that the first identity of (E2) holds. Then the row $K-1$ is chosen to guarantee the second identity of (E2). Finally the last row of Q , that plays no role in the optimization, is chosen to be $(0, 0, \dots, 0, 1)$.

Example 2 Fix $K \geq 5$. Let $\varepsilon = \frac{1}{200K}$. Take \mathbf{v}, γ and Q as follows.

$$\begin{aligned} \forall i \in [1, K] : \gamma_i &= (1 - \varepsilon)\varepsilon^{-K+i} \\ \forall i \in [1, K] : v_i &= \varepsilon^{K-i}/(1 - \varepsilon^{K+1-i}) \\ \forall i \in [2, K-1], \forall j \neq i+1 : q_{i-1,i+1} &= \frac{10}{11}\varepsilon^{K+1}(1 - \varepsilon^{K-i}), \quad q_{i-1,j} = 0 \\ \forall j \neq 2 : q_{K-1,2} &= \varepsilon^2(1 - \varepsilon^{K-1})/(1 - \varepsilon), \quad q_{K-1,j} = 0 \\ \forall j \in [1, K-1] : q_{K,j} &= 0, \quad q_{K,K} = 1 \end{aligned}$$

Let $\Gamma_i = \sum_{l=i}^K \gamma_l$. With our choice of parameters, let $\alpha = 1$ and $U = (1 - \varepsilon)\frac{10}{11}\varepsilon^{K-1} \leq \alpha$. Then the properties (E0) – (E2) hold true, where we denote by $\bar{q}_{i,j} := \gamma_i q_{i,j}$, and where $\xi \in [\frac{1}{10}, \frac{1}{10} + 2\varepsilon]$.

$$\forall i \in [1, K] : \Gamma_i v_i = \alpha \tag{E0}$$

$$\forall i \in [1, K-1] : v_i \leq \frac{\varepsilon}{1 - \varepsilon} v_{i+1}. \tag{E1}$$

$$\begin{cases} \forall i \in [2, K-1] : \bar{q}_{i-1,i+1} v_{i+1} = U, \quad \forall j \neq i+1 : q_{i-1,j} = 0 \\ (\bar{q}_{1,3} + \bar{q}_{2,4} + \dots + \bar{q}_{K-2,K} + \bar{q}_{K-1,2}) v_2 = (1 + \xi)U, \quad \forall j \neq 2 : q_{K-1,j} = 0 \\ \forall j \in [1, K-1] : q_{K,j} = 0, \quad q_{K,K} = 1 \end{cases} \tag{E2}$$

The following claim shows that **Example 2** satisfies Proposition 1.

CLAIM. Let $\bar{\pi} = (2, \dots, K)$. We have that $\mathcal{R}(\bar{\pi}) = \alpha + U + \frac{\xi}{K-1}U$ and for any $\pi \neq \bar{\pi}$

$$\mathcal{R}(\pi) \leq \mathcal{R}(\bar{\pi}) - \frac{1}{20(K-1)}U.$$

Remark 1 Observe that for (E0) – (E2) to hold, while $K \geq 5$, we must have $U = \mathcal{O}(\varepsilon^{K-4}\alpha)$ and thus the loss if one adopts a single price strategy with any price v_2, \dots, v_K , compared to an optimal strategy is small. Note also the importance of the term $\bar{q}_{K-1,2}v_2$ in our assumptions since $(\bar{q}_{1,3} + \bar{q}_{2,4} + \dots + \bar{q}_{K-2,K})v_2 \leq \varepsilon U$.

PROOF OF THE CLAIM.

We denote by (ij) the expected revenue function generated by the customers with valuation strictly lower than v_i , who enter at a date where the price is i and see the price j at the subsequent date. Thus, for $\pi = (w_1, \dots, w_J)$

$$\mathcal{R}(\pi) = \alpha + \frac{1}{J} \left(\sum_{i=1}^{J-1} (w_i w_{i+1}) + (w_J w_1) \right).$$

Lemma 2 Under (E1)–(E2) we have that

$$\begin{aligned} (i(i+1)) &= U : \forall i \in [1, K-1] \\ (i(i+\ell)) &= 0 : \forall 1 \leq i \leq i+\ell \leq K, \ell \geq 2 \\ (K2) &= \frac{11}{10}U \\ (ij) &\leq \left(1 + \frac{1}{100K}\right) U : \forall i, j \in [1, K], (ij) \neq (K2). \end{aligned}$$

Proof. The first three equalities follow directly from the assumption (E2).

Consider $i \leq K-1$ and $2 \leq j \leq i$. We have that

$$(ij) = (\bar{q}_{j-2,j} + \bar{q}_{j-1,j+1} + \dots + \bar{q}_{i-1,i+1})v_j \leq (1 + 2\varepsilon)U$$

due to assumptions (E1) and (E2) (when $j = 2$ we used the notation $\bar{q}_{0,2} = 0$).

Now, for $j > 2$

$$(Kj) = (\bar{q}_{j-2,j} + \bar{q}_{j-1,j+1} + \dots + \bar{q}_{K-2,K})v_j \leq (1 + 2\varepsilon)U.$$

and finally,

$$(K1) = (\bar{q}_{1,3} + \bar{q}_{2,4} + \dots + \bar{q}_{K-2,K})v_1 + \bar{q}_{K-1,2}v_1 \leq \varepsilon \frac{11}{10}U.$$

□

It follows from the first two equalities of the lemma that $\mathcal{R}(\bar{\pi}) = \alpha + U + \frac{1}{10(K-1)}U$. Take π to be another simple cyclic strategy. If π does not contain the string $(K, 2)$ it follows from the third equation of the lemma that $\mathcal{R}(\pi) \leq \alpha + \left(1 + \frac{1}{100K}\right)U \leq \mathcal{R}(\bar{\pi}) - \frac{1}{20(K-1)}U$. If, to the contrary, π contains the string $(K, 2)$ but is distinct from $\bar{\pi}$ we have that π contains another string (i, j) with $j \neq i + 1$, hence using the first, second and third inequalities of the lemma we get if L is the length of the cycle of π :

$$\mathcal{R}(\pi) \leq \alpha + \frac{1}{L} \left(\frac{11}{10}U + (L-2) \left(1 + \frac{1}{100K} \right) U \right) \leq \mathcal{R}(\bar{\pi}) - \frac{1}{20(K-1)}U$$

and the claim is proved. □

6.3 Approximating the Solution of Problem (1)

The lack of simplifying structure of the optimal policies confirms the curse of dimensionality, that seems inherent to the changing valuations feature of such type of intertemporal pricing. This has been argued for in Section 2 by relying specifically on example 2. In such context, the exponential complexity makes our results even more valuable in suggesting efficient approximations for problem (1). Indeed, as a result of the untractability of problem (1), one would naturally look for approximations. For instance, one could consider optimizing over fixed price policies. Such policies have been shown to be asymptotically optimal in various settings (e.g., the recent work of ? who prove, the asymptotic optimality of fixed price policies under price commitment). Another approach would be to approximate a customer's stochastic valuation by a carefully selected constant and accordingly, use the results of ?. Our results can be viewed as a third approach. By introducing the notion of pricing pace, we formulated (2) which is a constrained version of (1) (and of independent interest). Hence, the optimal solutions of (2) defined and obtained through Theorem 3 and 4 are therefore natural approximations of (1). Moreover, such approximations have the advantage of preserving the nature of the problem with respect to changing valuations. We show by many numerical examples (see, Section 8, Finding III and V), that such solutions even for small values of M , outperform sensitively the two other approaches of fixed price policies and constant valuations. This in turns shows that the algorithm does give an added value to the pricing optimization even if it is operated in a constrained setting.

7 Extensions

7.1 Unbounded Patience Setting. The (M-UP) model

One main assumption throughout this paper was the boundedness of the maximum patience level. In this section, we look at relaxing this assumption and setting $\tau = \infty$ in the Markovian model. This model was denoted by (M-UP) and is only determined by the pair (γ, Q) . Under this setting, the time at which the valuation process eventually reaches the value $v_0 = 0$ is endogenous, governed only by the matrix Q . We still assume that its expected value is finite. A sufficient condition for this to hold

is

$$\|Q\| := 1 - \nu < 1, \quad (\mathcal{C})$$

for some $\nu \in (0, 1)$, where $\|\cdot\|$ is the infinity norm for matrices, that is the maximum absolute row sum of the matrix. This same condition will guarantee that there will be no accumulation of customers in the system, equivalently, at any time t the expected number of customers in the system is uniformly bounded, (see, Appendix A for a proof of this claim).

We now analyze the (M-UP) case by approximating it by an (M-BP) model. For this, fix any $\varepsilon > 0$ and introduce the quantity

$$\tau_\varepsilon = \lceil \ln \varepsilon / \ln(1 - \nu) \rceil. \quad (4)$$

By definition, $1 - \nu$ is an upper bound on the probability of a customer - no matter the state - to remain for an additional period in the system. It then follows from (C) and the expression of τ_ε that the probability that an arriving customer spends more than τ_ε in the system is lower than $(1 - \nu)^{\tau_\varepsilon} = \varepsilon$. We therefore introduce the (M-BP) model with patience level τ_ε and the same primitives (γ, Q) as the (M-UP) we are studying.

To proceed, we say that a pricing policy π^* is ε -optimal on a subset \mathcal{A} of \mathcal{P} for some $\varepsilon > 0$, if

$$\sup_{\pi \in \mathcal{A}} \mathcal{R}(\pi) \leq \mathcal{R}(\pi^*) + \varepsilon.$$

We have the following

Proposition 9 *Under condition (C), if π^* is an optimal strategy on any subset \mathcal{A} of \mathcal{P} , for the (M-BP) model given by τ_ε and (γ, Q) , then π^* is $(v_K \varepsilon / \nu)$ -optimal on \mathcal{A} for the (U-BP) model given by (γ, Q) .*

Proof. Consider one customer that arrives to the system at some time t with a valuation v_m governed by a (U-BP) model to which we apply some policy $\pi \in \mathcal{A}$. Recall that an (M-BP) model disregards any revenue generated by any customer who spends more than τ_ε in the system. Therefore, the difference between the expected revenues generated from respectively applying the same policy π to both an (M-UP) model and an (M-BP) model with the same primitives, is bounded by $\eta^\pi(m, \tau_\varepsilon) v_K$, where $\eta^\pi(m, \tau_\varepsilon)$ is the probability in the (M-UP) model, that under π , the customer remains in the system for a period greater than τ_ε . By definition of τ_ε and condition (C), $\eta^\pi(m, \tau_\varepsilon) \leq \varepsilon / \nu$. By considering all customers that arrived to the system by time t , the difference in expected revenues from these t customers throughout the horizon will be upper bounded by $t v_K \varepsilon / \nu$. Let us denote by $\mathcal{R}^U(\pi)$ (resp., $\mathcal{R}^B(\pi)$) the long-run average revenue generated by applying policy π under an (M-BP) (resp., (M-UP)) model. From the above we conclude that for any policy $\pi \in \mathcal{A}$ we have that

$$\mathcal{R}^U(\pi) \leq \mathcal{R}^B(\pi) + v_K \varepsilon / \nu \leq \mathcal{R}^B(\pi^*) + v_K \varepsilon / \nu \leq \mathcal{R}^U(\pi^*) + v_K \varepsilon / \nu.$$

This completes our proof. \square

As a corollary we obtain the following

Proposition 10 *Under (M-UP), Theorems 3 and 4 as well as Propositions 7 and 8 hold after replacing the optimality concept by ε -optimality, and τ by τ_ε .*

7.2 Infinite Horizon Discounted Revenue Function

In this extension, we show how the main results can be easily extended to the context of infinite-horizon discounted revenue functions. For that it is enough to revisit the weakly coupled framework of Section 4 and extend Proposition 3 accordingly. The rest of the analysis on weakly coupled functions and its implications on problems (1) and (2) will also extend in a similar way.

We start by a definition of weakly coupled revenue function in the discounted context. Let

$$\mathcal{R}(\pi) = \sum_{n=0}^{\infty} e^{-rn} f(w_n, w_{n+1}).$$

The definition of φ given in Equation (3) is also modified in a similar way. Finally, we say that a policy π is *pre-cyclic*, if it is of the form (W_0, W_1, W_1, \dots) for some “period” W_1 and for some initial finite string of prices W_0 .

Proposition 11 *For any $\tau \geq 1$, for any function $f : \Omega^\tau \times \Omega^\tau \rightarrow \mathbb{R}^+$, we have the following. There exists a pre-cyclic $\pi^* \in \mathcal{P}$ such that,*

$$\mathcal{R}(\pi^*) = \sup_{\pi \in \mathcal{P}} \mathcal{R}(\pi).$$

Moreover, $\pi^ = (W_0, W_1, W_1, \dots)$, with $W_0 = (w_1^0, \dots, w_{P_0}^0)$ and $W_1 = (w_1^1, \dots, w_{P_1}^1)$, such that all the coordinates of (W_0, W_1) are distinct, with $P_0 + P_1 \leq K^\tau$.*

Proof. Decompose any $\pi \in \mathcal{P}$ as $\pi = (W_0, W_1, \bar{\pi})$ where $W_0 = (w_1^0, \dots, w_{P_0}^0)$ and $W_1 = (w_1^1, \dots, w_{P_1}^1)$ and $\bar{\pi} = (\bar{w}, \dots) \in \mathcal{P}$ where $\bar{w} \in \Omega^\tau$ is the first string that appears twice in the expression of π . By definition of \bar{w} we have that all the coordinates of (W_0, W_1) are distinct.

Define $\tilde{\pi} = (W_1, \bar{\pi})$. Observe that for some functions C and C' one has

$$\mathcal{R}(\pi) = C(W_0, \bar{w}) + e^{-rP_0} \mathcal{R}(\tilde{\pi})$$

as well as

$$\mathcal{R}(\pi) = C'(W_0, W_1, \bar{w}) + e^{-r(P_0+P_1)} \mathcal{R}(\bar{\pi}).$$

Applying the first equality to an optimal strategy π^* , we see that $\tilde{\pi}^*$ should maximize \mathcal{R} among all strategies that start with \bar{w}^* . Hence, one can take $\bar{\pi}^* = \tilde{\pi}^*$ and still get an optimal strategy $\pi^{**} = (W_0^*, W_1^*, W_1^*, \dots)$. \square

7.3 Extension to Strategic Consumers

We end this section by arguing how the weakly coupled property of the revenue function is also satisfied when customers are strategic in the sense of forward looking.

The decision for a patient customer to purchase the product occurs as soon as the current valuation becomes larger than the current price. In the case of a strategic customer, this decision is more subtle and requires a specific model of strategic behavior. Strategic customers make their purchasing decision not only based on the current valuation and price, but also based on future committed prices and possibly on the distribution of their future valuation (e.g., a customer is expecting shortly a bonus and thus the willingness-to-pay will be affected by how large the bonus will be). Independently of the specifics of the strategic behavior model, we can show, in the (BP) model that Theorems 3 and 4 remain valid for strategic customers. In fact, since every arriving customer has a bounded patience level τ , if we consider a policy of the form, $\pi = ((w_j : j \geq 1) : w_j \in \Omega^\tau)$, then the customers who arrive between time $j\tau + 1$ and $(j+1)\tau$ will generate revenues that are only function of w_j and w_{j+1} no matter how their strategic behavior is modeled (as long as the stationarity assumption holds). Hence the weakly-coupled property holds as in the case of strategic customers, and the results of Theorem 3 and 4 apply.

We end this discussion with a couple of examples on possible strategic models that could govern forward looking customers.

Example 3. Customers present in the system purchase in a specific period if the current surplus is larger than all possible surplus in the future calculated using the current valuation. Suppose a customer is currently present in the system with valuation v_k and who could remain in the system for at most another u periods. The customer observes the current and future committed prices in his remaining window say, $(k_s, k_{s+1}, \dots, k_{s+u})$ and would decide to purchase the product at the current price, k_s , if $k - k_s \geq 0$ and $k_s \leq \min\{k_{s+1}, \dots, k_{s+u}\}$. The customer would delay a purchase if the product is offered at a cheaper price in the “future”. In this example, the customer believes that his current valuation is the best prediction of her future valuation.

Example 4. In this example, the customers of every given class have a common model for their stochastic valuation evolution process, call it $\tilde{\mathcal{V}}_t(\cdot)$ (it does not need to be the same as the firm’s model). Customers purchase in a specific period if the current surplus is positive and larger than the expected value of the maximum surplus that can be generated in the remaining periods of the patience level, that is if $k - k_s \geq 0$ and

$$k - k_s \geq \mathbb{E} \max_{s+1 \leq j \leq s+u} \{\tilde{\mathcal{V}}_t(j) - k_j\}.$$

8 Numerical Analysis

Our objective in this section is to use the algorithm of Theorem 4 to shed some light on the behavior of optimal pricing strategies. We confine our analysis to the (M-BP) and (M-UP) models where we

have a closed form of the revenue function. We use the algorithm to obtain and analyze an optimal cyclic policy for many choices of the matrix Q , while also varying τ , the valuation vector \mathbf{v} and its size K . Without much loss of generality we set $\gamma = [1/K, \dots, 1/K]$. The parameters that we use in drawing different realizations of Q and \mathbf{v} are the following:

- i.) The parameter $\nu > 0$ represents the probability in any given period that a customer leaves the system by reaching the absorbing state v_0 . Without a great loss of generality we often restrict, for simplicity of notations, our numerical analysis to the case where ν is the same for any given state of the customer. The coefficients of the matrix Q can thus be drawn randomly in $(0, 1)$ such that the sum on each row is equal to $1 - \nu$.
- ii.) The parameter $e > 0$ is introduced to represent the average gap between consecutive valuations. That is, the vector $\mathbf{v} = [v_1, \dots, v_K]$ is such that $v_1 = 1$ and $v_{i+1} = v_i(1 + e\zeta_i)$ where ζ_i is uniformly drawn in $(0, 1)$.

When the findings with respect to the (M-BP) and (M-UP) cases are the same, we present these under the (M-BP) setting. We later discuss the findings that are specific to the (M-UP) case. For now, we set $M = 1$, or equivalently, $\sigma = \tau$. For each set of parameters (τ, ν, e, K) we make $n \gg 1$ drawings of the different parameters to which we apply the analytic algorithm. We collected and summarized the statistics of the outcomes in the table below.

We call k -cyclic policy a cyclic and simple τ -policy of size $k\tau$ (i.e., with k different prices).

We say that a cyclic policy π outperforms π' by $\delta\%$ if $100(\mathcal{R}(\pi) - \mathcal{R}(\pi'))/\mathcal{R}(\pi') = \delta$.

Set \mathfrak{f} the percentage of times where the maximal policy performs better than all the fixed price policies, and \mathfrak{d} denotes the average ratio (in percent) of how much the maximal policy outperforms the fixed price policy (this average is computed conditional on the fact that the maximal policy is not a fixed price one).

Set \mathfrak{f}' the percentage of times where the maximal policy performs better than all the fixed price and two-cyclic policies, and \mathfrak{d}' denotes the average ratio (in percent) of how much the maximal policy outperforms the maximum over the fixed price and two-cyclic policies (this average is computed conditional to the fact that the maximal policy has period larger or equal to three).

Finding I. Cyclic policies with cycle at most three and decreasing are almost always optimal.

Under (M-BP), each customer sees at most two prices and this is true with probability larger than $1 - \varepsilon$ in the (M-UP) case. It is hence reasonable to expect that the optimal cyclic policies given by Theorem 3 actually have a short cycle size, even if K is large. We obviously know based on Proposition 1 that this cannot be always true. However, our numerical analysis actually yielded, in all the cases we considered, optimal policies that are monotone decreasing and k -cyclic with k smaller or equal than 3. Hence, our first finding is that for all values of the parameters and for all the instances we considered (as described above), the optimal policies were decreasing and k -cyclic with $k \leq 3$.

Finding II. Fixed price and two-cyclic policies are practically optimal.

Our second finding is that the fixed price and two-cyclic policies were nearly optimal, with just very few instances where a three-cyclic policy outperformed them; and in those cases the performance ratio was less than 1%. This is reflected in the two last columns of Table 1.

Finding III. As σ increases fixed price policies perform increasingly better compared to two-cyclic ones.

In line with our result of Propositions 8 we find that, as long as the probability to reach the state v_0 or v_K from all the other valuations is supposed to be strictly positive, then as σ gets very large the fixed price policies become the best. This is because customers tend to see just one price as σ goes to infinity. As shown by the fourth column of Table 1, the frequency at which two-cyclic policies outperform the fixed price policies is decreasing with σ . As shown by the fifth column of Table 1 we also see the outperforming rate of the two-cyclic policies \mathfrak{d} is also a decreasing function of σ . Finally, the last line of these two columns confirm our result of Proposition 8, namely that as σ goes to infinity the optimal policy is a fixed price one.

Finding IV. As ν increases fixed price policies perform increasingly better compared to two-cyclic ones.

We also find that as the probability to reach the states v_0 or v_K , from all the other valuations, increases, the fixed price policies become the best. This is natural since when the latter probability grows, the customers spend less time in the system and it is clear for example that if each customer stays in the system for just one period, then the price that optimizes the revenue of a single period (without consideration of Θ) gives rise to an optimal policy that is the fixed price policy with this same price. As shown by the fourth column of Table 1, the frequency at which two-cyclic policies outperform the fixed price policies is decreasing with ν . As shown by the fifth column of the Table 1, we also see the outperforming rate of the two-cyclic policies \mathfrak{d} is also a decreasing function of ν .

τ	ν	e	\mathfrak{f}	\mathfrak{d}	\mathfrak{f}'	\mathfrak{d}'
1	0.1	0.1	76	2.6	5	0.1
5	0.1	0.1	43	0.6	2	0
9	0.1	0.1	21	0.3	0	0
13	0.1	0.1	13	0	0	0
5	0.	0.1	67	3	5	0.3
5	0.1	0.1	54	0.6	4	0.1
5	0.2	0.1	16	0.1	0	0
5	0.3	0.1	1	0	0	0
5	0.1	0.05	45	0.3	2	0.1
5	0.1	0.15	43	0.7	3	0.1
5	0.1	0.25	29	1.2	2	0.2
5	0.1	0.35	34	1.7	2	0.2

Table 1. (BPL), $\sigma = \tau$, $K = 4$.

Finding V. Impact of disregarding changing valuations on revenues.

We illustrate these findings with a couple of specific examples. Consider first the following. Set $K = 4$ and $\tau = \sigma = 10$ under the (M-BP) case and set the transition matrix Q close to the Identity,

$$Q = \begin{bmatrix} 0.9 & 0.05 & 0.03 & 0.02 \\ 0.05 & 0.8 & 0.1 & 0.05 \\ 0 & 0.05 & 0.95 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

Note that $\nu(Q) = 1$ as is the case for the Identity matrix. Let $\mathbf{v} = [1.0, 1.3, 1.45, 1.6]$. For this valuations' vector and this transition matrix Q , an optimal policy is the two-cyclic policy $[3, 2]$. If however the matrix is the Identity (i.e., the valuations are not changing), an optimal policy is given by $[2, 1]$. Of course, if the seller disregards the impact of the changing valuations, and thus solves for an optimal policy using $Q' = \text{Id}$, she will end up setting the two-cyclic policy $[2, 1]$ instead of $[3, 2]$. The loss in profit of doing so is found numerically to be 4%. This result is quite robust. Indeed, the policy $[2, 1]$ is in fact numerically found to be optimal for all diagonal matrices that are close to the Id. Thus, the seller would incur the same loss of 4% if he replaces Q by any diagonal matrix close to the Identity.

Now, let $\mathbf{v} = [1.0, 1.1, 1.25, 1.4]$ and

$$Q = \begin{bmatrix} 0.7 & 0.1 & 0.05 & 0.07 \\ 0.05 & 0.7 & 0.07 & 0.07 \\ 0.07 & 0.07 & 0.7 & 0 \\ 0 & 0 & 0.05 & 0.8 \end{bmatrix}.$$

For any transition matrix that is diagonal, $Q' = \text{Diag}(d_1, d_2, d_3, d_4)$ where $d_i \in [0.6, 1]$, we find that the fixed price policy $[1]$ is optimal. On the other hand, when we consider Q to be the transition matrix, the policy $[3, 2, 1]$ is optimal and outperforms $[1]$ by 5%.

Finding VI. An example of three-cyclic optimal policy.

In the case of $\tau = M\sigma$, $M \geq 2$, one expects that optimal policies with longer cycles appear. However each customer sees $M + 1$ prices in this context and it is reasonable to conjecture that the length of the cycles will be small if M is taken to be small. Because of the complexity K^{4M} we are bound to let M be small if we want to apply our algorithm.

In all the examples we have tested ($M \leq 4$) the optimal policies were decreasing and had a cyclic length of at most 3. However, as opposed to the case $M = 1$, we often got cases where a three-cyclic policy outperforms both two-cyclic and the fixed price policies by more than 5%. We illustrate this finding with the following example.

Consider the following (M-BP) example with $K = 4$, $\tau = 6$ and $\sigma = 2$, that is $M = 3$. We set $\mathbf{v} = [1.0, 1.2, 1.4, 2]$ and $\boldsymbol{\gamma} = [0.25, 0.25, 0.25, 0.25]$. We let

$$Q = \begin{bmatrix} 0.8 & 0.05 & 0.0 & 0 \\ 0.2 & 0.6 & 0.0 & 0.2 \\ 0.1 & 0.2 & 0.6 & 0 \\ 0.1 & 0.2 & 0.6 & 0 \end{bmatrix}.$$

The maximum fixed price policy is given by [1] with $\mathcal{R}([1]) = 1$. The optimal two-cyclic policy is given by [3, 1] with $\mathcal{R}([3, 1]) = 1.09$. Finally, an optimal policy is a three-cyclic policy and is given by [4, 3, 1] with $\mathcal{R}([4, 3, 1]) = 1.14$. It is an example where an optimal policy outperforms the fixed price policies by 14% and the two-cyclic ones by 4%.

Finding VII. Tuning ε in the (M-UP) case.

Consider the (M-UP) case. We fix K, Q, \mathbf{v} , and $\boldsymbol{\gamma}$. One could aim at selecting ε in a way to obtain the best possible solution of (1). For every $\varepsilon > 0$ we define $\tau_\varepsilon = -\ln \varepsilon / \nu$ and we set $\tau = \tau_\varepsilon$ and take $\sigma = \tau$ or equivalently $M = 1$ (a similar reasoning below can be made for any M).

We denote by π_ε an optimal policy given by the algorithm that uses the analytic approximation of the revenue by a τ -affine revenue function. We denote by $\bar{\mathcal{R}}(\pi_\varepsilon)$ the approximate return of an optimal policy as given by the algorithm. Given that the error of the affine approximation is bounded by ε we may introduce the quantity $\bar{\mathcal{R}}(\pi_\varepsilon) - \varepsilon$ and focus on the choices of ε that make it positive.

For example, to compare the performance of π_ε versus the fixed price policies, we can introduce the notation $\Delta_\varepsilon = \bar{\mathcal{R}}(\pi_\varepsilon) - \varepsilon - \mathcal{R}_f^*$ where \mathcal{R}_f^* denotes the revenue of the best fixed price policy. Typically, as $\varepsilon \rightarrow 0$, fixed price policies are optimal as is confirmed by the (M-UP) version of Proposition 8 (included in Proposition 10) and the algorithm (see Table 2). As ε becomes large, we often find two-cyclic optimal policies π_ε . However as ε becomes too large, Δ_ε becomes negative which implies that any optimal policy we find is not guaranteed to perform better than the fixed price ones, once we account for the error ε of our approximation. In conclusion, as we vary $\varepsilon > 0$, we aim to maximize $\bar{\mathcal{R}}(\pi_\varepsilon) - \varepsilon$ and track Δ_ε as a way to compare the approximation with fixed price policies.

We illustrate this approach using the following example. We let $\mathbf{v} = [1.0, 1.2, 1.5, 2]$ and

$$Q = \begin{bmatrix} 0.65 & 0.05 & 0 & 0.05 \\ 0.05 & 0.65 & 0 & 0.05 \\ 0.05 & 0.05 & 0.6 & 0 \\ 0 & 0 & 0 & 0.6 \end{bmatrix}.$$

Note that $\nu(Q) = 0.3$. From Table 2, we observe that there is a broad range of values of ε , where a two-cyclic policy is strictly better than fixed-price policies. In particular, by setting $\varepsilon = 0.01$ the algorithm generates a policy that outperforms the best fixed price policy by more than 3%, and represents our best approximation of the unconstrained optimization.

ε	τ_ε	π_ε	$\bar{\mathcal{R}}(\pi_\varepsilon)$	Δ_ε
10^{-13}	83	[4]	1.02	0
10^{-5}	32	[4,1]	1.03	0.02
10^{-3}	19	[4,1]	1.04	0.04
10^{-2}	12	[4,1]	1.05	0.05
10^{-1}	6	[4,1]	0.99	0
0.2	4	[4,1]	0.91	-0.1
0.4	2	[4,1]	0.73	-0.27

Table 2. An (M-UP) case, with $M = 1$.

9 Conclusion

We tackled in this paper an intertemporal pricing problem that incorporates what we argued to be a natural consumer’s behavior, namely, customers with time-varying valuations. We considered a continuous flux of patient customers with random but bounded patience levels during which their valuation processes stochastically change through time. The model depicting the valuation process is extremely general and requires only time stationarity. We show, with some explicit numerical examples, that disregarding stochastic valuations can make the firm miss a sizable profit. As for the analytical analysis, and despite the generality of the model, the additive revenue function is shown in this context to preserve an interesting property, denoted by weakly coupled, which guarantees that cyclic policies are optimal. We thus established that the optimality of cyclic policies in the context of intertemporal pricing is quite robust.

The time-varying valuations feature of the model brings however a major complexity with respect to the structure and the tractability of these policies. Concerning the latter, one can obtain well behaved and numerically tractable policies by limiting the pricing pace, i.e., by forcing the firm to keep each price for a minimum amount of time. As for the structure, previous work (see, ?) have highlighted the complexity of finding the optimal policy even for constant valuations as soon as one considers multiple patience levels. In the latter setting, ? shows that it can remain tractable. However with general time-varying valuations, we argue here that it will not be possible to reduce the exponential complexity of the optimal policy. For that, we constructed an example showing that even when customers spend no more than two periods in the system ($M = 1$), an optimal policy can be cyclic increasing and the size of the cycle can be as large as $K - 1$. Interestingly however, we observed numerically that in the same setting of the example ($M = 1$), when we randomly select the primitives of the problem, it is quite rare that an optimal cyclic policy has a size strictly larger than two (and hence, independent of K) or that an optimal policy is not monotone *decreasing*.

In terms of potential next steps, we believe that the last observation is worth investigating carefully by looking for conditions under which the cycle size is independent of the number of prices. The problem is difficult but a better understanding of this possible independence in K would represent

a valuable breakthrough. It will be also valuable to look more carefully into the strategic customer case and study the best way to model such behavior. It has been argued (see ? and ?) that the strategic case is easier to handle than the patient case. We note that under changing valuations, it seems impossible to get better than the weak coupling dissociation on the revenue function. Finally, most of the closely related literature have been looking at a single product. It could be an interesting challenge to look at substitutable products and see whether one could identify the structure of optimal policies under such a more complex demand model.

Appendix A.

1 Markovian Valuations - Formulation of the Revenue Pair

This section is devoted to some preliminary analysis and calculations when the valuation process follows a discrete time Markov chain on Ω with an absorbing state at $v_0 = 0$; hence its transition matrix can be written as follows

$$\bar{Q} = \begin{pmatrix} 1 & 0 \\ H & Q \end{pmatrix},$$

where H is a column vector with dimension K , which entries are the probabilities to go from any state $k > 0$ to v_0 . We also assume, WLOG, that one customer arrives every period.

Our objective is to characterize how the expected number of customers that are in the system evolve through time and we also obtain a closed form formulation of the revenue function (L, Θ) .

We start by introducing new notations. We set the following

- $P_{i,j^+} = \sum_{l=j}^K q_{il}$. The probability to transfer from state i to any possible state $l \geq j$.
- $P_{i^-,j^+} = (\sum_{l=j}^K q_{1l}, \sum_{l=j}^K q_{2l}, \dots, \sum_{l=j}^K q_{(i-1)l})^T$. The column vector whose entries are the probabilities to go from a state strictly lower than i to all possible states higher than j .
- $U_{i,j^-} = (q_{i1}, \dots, q_{i(j-1)})$. The row vector whose entries are the probabilities to go from state i to a specific state strictly lower than j
- $Q_{i^-,j^-} = \{q_{kl} : 1 \leq k < i, 1 \leq l < j\}$ is the matrix of transfer probabilities from states lower than i to states lower than j . It is a minor of the matrix \bar{Q} . We denote by Q_m the square matrix Q_{m^-,m^-}

Note first that “ $-$ ” involves states strictly lower while “ $+$ ” involves states larger or equal; secondly, all the above quantities involve transition probabilities between states in Ω and do not involve the absorbing state v_0 .

Expected Number of Customers in the System.

We consider $\theta = (\theta_1, \dots, \theta_K) \in [0, \tau]^K$, where θ_m is again the expected number of individuals with valuation v_m who are currently in the system. Given a phase (k, t) , an expected state of the system, θ , we recall that $\Theta(k, t | \theta) = (\Theta_1(k, t | \theta), \dots, \Theta_K(k, t | \theta))$, is the expected state of the system one phase later. Therefore, if $m \geq k$ then $\Theta_m(k, \tau | \theta) \equiv 0$, otherwise,

$$\Theta_m(k, t | \theta) = \left(\sum_{l=1}^{k-1} \theta_l U_{l,k^-} Q_k^t + \sum_{i=0}^{t \wedge \tau - 1} \sum_{l=1}^{k-1} \gamma_l U_{l,k^-} Q_k^i \right)_m + \gamma_m \quad (5)$$

We recognize that for $t \geq \tau$ and for $m < k$, we have that

$$\Theta_m(k, t) = \bar{\Theta}_m(k) := \left(\sum_{i=0}^{\tau-1} \sum_{l=1}^{k-1} \gamma^l U_{l,k^-} Q_k^i \right)_m + \gamma_m, \quad (6)$$

and $\Theta_m(k, t) = \bar{\Theta}_m(k) \equiv 0$, when $m \geq k$.

Expected Revenues.

We denote by $\mathcal{L}_{(k,t)}(m)$ the expected revenue generated in the next t periods by one individual currently present in the system with valuation v_m , facing in the next t periods the price v_k . It is possible that the individual does not purchase during these t periods; any revenues generated after t are not included in $\mathcal{L}_{(k,t)}(m)$. Therefore,

$$\begin{aligned} \mathcal{L}_{(k,t)}(m) &= [P_{m,k^+} + U_{m,k^-} P_{k^-,k^+} + U_{m,k^-} Q_k P_{k^-,k^+} + U_{m,k^-} Q_k Q_k P_{k^-,k^+} + \dots] \cdot v_k \\ &= [P_{m,k^+} + \sum_{l=0}^{t \wedge \tau - 2} U_{m,k^-} Q_k^l P_{k^-,k^+}] \cdot v_k \end{aligned} \quad (7)$$

We also denote by $\bar{\mathcal{L}}_{(k,t)}(m)$ the expected revenue generated from one individual who just arrived into the system with valuation v_m , and who will be facing the phase (k, t) . Therefore,

$$\bar{\mathcal{L}}_{(k,t)}(m) = \delta_{m \geq k} v_k + (1 - \delta_{m \geq k}) \mathcal{L}_{(k,t)}(m).$$

The first term depicts the expected revenue generated if at arrival the customer has a valuation larger or equal than the listed price k . Otherwise, the customer now is in the system and the expected revenue generated from the remaining period are given by \mathcal{L} defined above.

Putting together some of these formulations we obtain a closed formulation of the expected revenue L .

Proposition 12 *Given a transition matrix, Q , a vector γ of initial proportions and a set Ω of valuations, the expected revenues generated from a set of consecutive phases $((k_j, t_j) : 1 \leq j \leq n)$, is given by*

$$\begin{aligned} L(k_1, \dots, k_n : t_1, \dots, t_n) &= \sum_{j=1}^n L(k_j, t_j | \theta^{j-1}) \\ &= \sum_{j=1}^n \sum_{m=1}^K [\gamma_m \sum_{t=1}^{t_j} \bar{\mathcal{L}}_{(k_j,t)}(m) + \theta_m^{j-1} \mathcal{L}_{(k_j,t_j)}(m)] \end{aligned} \quad (8)$$

with $\theta_m^j = \Theta_m(k_j, t_j | \theta^{j-1})$ is given in closed form in (5) and $\theta^0 = \mathbf{0}$.

We end this section by a result that is relevant primarily in the unbounded case.

Proposition 13 *Suppose that the system starts empty. Then, for any policy $\pi \in \mathcal{P}$, and any $1 \leq m \leq K$, the expected number of customers in the system with valuation v_m is bounded by $\rho = \min\{\tau, 1/\nu\}$.*

The result is obvious in the (M-BP) case. For the (M-UP), it is enough to prove that for any policy $\pi = ((k_j, \tau_j) : j \geq 1) \in \mathcal{P}$, we have that for all $j \geq 1$ and $1 \leq m \leq K$,

$$\theta_m^j = \Theta_m(k_j, \tau_j | \theta^{j-1}) \leq \rho < \infty.$$

We start with two observations. First, the probability that a customer with valuation v_m jumps next to v_0 is upper bounded by $\Lambda := 1 - \nu$. Second, the expected number of customers with valuation v_m at time t , is always less than the expected number of customers with valuation v_m at time t , that would have been in the system if the pricing policy was set at K throughout. Recall, that $\Theta_K = 0$ and $\theta^0 = \mathbf{0}$, and thus from Equation 5 we have that

$$\begin{aligned} \|\Theta(K, t)\| &\leq \left\| \sum_{i=0}^{t-1} \sum_{l=1}^{K-1} \gamma_l U_{l,K-} Q_K^i \right\| + \|(\gamma_1, \gamma_2, \dots, \gamma_{K-1}, 0)\| \\ &\leq \sum_{i=0}^{t-1} \sum_{l=1}^{K-1} \gamma_l \|U_{l,K-}\| \|Q_K\|^i + 1 \\ &\leq \Lambda \sum_{i=0}^{t-1} \Lambda^i + 1 \leq \frac{1}{1 - \Lambda}. \end{aligned}$$

The third inequality is due to the fact that $\sum_{l=1}^{K-1} \gamma_l \|U_{l,K-}\| \leq \Lambda \sum_l \gamma_l \leq \Lambda$. The bound above shows that each component of Θ is bounded by $1/(1 - \Lambda)$. Hence, the expected number of customers in the system is bounded by $K/(1 - \Lambda)$. \square

2 Proof of Propositions 7

2.1 Additional Notations

We introduce some additional notations. Recall that $\mathcal{L}_{(k,t)}(m)$ (resp. $\bar{\mathcal{L}}_{(k,t)}(m)$) is the expected revenues generated in the next t periods from one individual who is already in the system (resp. just arrived to the system) facing the price v_k for t consecutive periods. Given a certain duration τ , we denote by

- $C_k(m) = \mathcal{L}_{(k,\tau)}(m)$, is the expected revenues generated during the phase (k, τ) by an individual initially in state m ,
- $B_k(m) = \bar{\mathcal{L}}_{(k,\tau)}(m)$, is the expected revenues generated during the phase (k, τ) by an individual arriving with valuation v_m ,
- $A_k(m) = \sum_{t=1}^{\tau} \bar{\mathcal{L}}_{(k,t)}(m)$
- $A_k := \sum_m \gamma_m A_k(m)$, is the aggregate expected revenues generated during the phase, (k, τ) , by all the customers that arrived during this same phase,

- $B_k := \sum_m \gamma_m B_k(m)$, is the aggregate expected revenues generated during the phase, (k, τ) , by all the customers that newly arrived at the beginning of this phase,
- $T_{k,k'} := \sum_m \bar{\Theta}_m(k) C_{k'}(m)$, is the expected revenues generated by $\bar{\Theta}(k)$ during a phase (k', τ) , where, $\bar{\Theta}$ is given by (6)
- $p_{k,t,m}$ is the probability that a customer who faces a phase (k, t) remains in the system at the end of the phase.

2.2 The proof

Next we detail the proof of Proposition 7. Property (A1) holds as a direct result of $\tau < \infty$. As for (A2), we start with the quantity $L(k, t | \boldsymbol{\theta})$ for $t \geq \tau$ and that can be written as the sum of three terms. First, the expected revenues generated during the phase from new customers in the last τ periods of the phase. That is exactly A_k . Secondly, the expected revenues generated from new comers that arrived in the first $t - \tau$ periods of the phase and that is exactly, $\sum_m \sum_{s=\tau+1}^t \gamma_m \bar{\mathcal{L}}_{(k,s)}(m)$. Finally, the revenues generated from customers that were in the system at the beginning of the phase, $\boldsymbol{\theta}$, and that is $\sum_m \theta_m \mathcal{L}_{k,t}(m)$. We write that

$$\begin{aligned}
L(k, t | \boldsymbol{\theta}) &= A_k + \sum_m \sum_{s=\tau+1}^t \gamma_m \bar{\mathcal{L}}_{(k,s)}(m) + \sum_m \theta_m \mathcal{L}_{k,t}(m) \\
&= A_k + \sum_m \sum_{s=\tau+1}^t \gamma_m (B_k(m) + p_{k,\tau,m} v_0) + \sum_m \theta_m C_k(m) + \sum_m \theta_m p_{k,\tau,m} v_0 \\
&= A_k + B_k (t - \tau) + \sum_m \theta_m C_k(m).
\end{aligned}$$

By setting \mathbf{C}_k to be the vector with entry $C_k(m)$, we complete our proof.

Appendix B

A Simpler Model. The case of $K = 2$

In this section, we consider an (M-UP) model with $K = 2$. Let $\Omega = \{v_1, v_2\}$, with $v_0 < v_1 \leq v_2$. A seller facing a patient customer would only consider a pricing policy \mathbf{p} where at any time t , $p_t \in \{v_1, v_2\}$. We start by observing that the price v_1 is a reset price and that is, once set the system empties. Having this in mind, it is easy to see that the optimal policy is either a fixed price policy where for all t , $p_t = v_i$ $i = 1, 2$, or is cyclical of the form $\pi_c = ((v_2, t), (v_1, 1))$; the cycle starts with v_2 that is set for $0 < t < \infty$ consecutive periods and then v_1 is set once. Note that the extreme cases of $t \in \{0, \infty\}$ cover the cases where the prices are constant.

The arriving customers has valuation v_2 with probability γ , and have valuation v_1 with probability $\bar{\gamma} = 1 - \gamma$. The average revenue per period collected by the seller is $R(t)$, with

$$\mathcal{R}(t) = \frac{1}{t+1} \left(\gamma v_2 t + v_1 + v_1 \frac{\bar{\gamma} q_{12}}{1 - q_{11}} \left(t - \frac{1 - q_{11}^t}{1 - q_{11}} \right) + v_1 \bar{\gamma} \frac{1 - q_{11}^t}{1 - q_{11}} (q_{11} + q_{12}) \right)$$

The first two terms in parenthesis represent the revenues from customers that arrive and buy right away: those with valuation v_2 at price v_2 during t periods, and everyone at price v_1 in the last period. The last two terms in parenthesis account for the customers that accumulate over time and end up buying at either v_1 or v_2 . The expected number of customers buying in a particular period l , $1 \leq l \leq t$, at price v_2 is given by the expression

$$\bar{\gamma} q_{12} \sum_{n=0}^{l-2} q_{11}^n = \bar{\gamma} q_{12} \frac{1 - q_{11}^{l-1}}{1 - q_{11}},$$

i.e., these are all customers that arrive earlier with low valuation v_2 , and for whom the valuation remains at v_2 and jumps to v_1 from at period $l - 1$. Summing over all l , $1 \leq l \leq t$, we get:

$$\sum_{l=1}^t \left(\bar{\gamma} q_{12} \frac{1 - q_{11}^{l-1}}{1 - q_{11}} \right) = \frac{\bar{\gamma} q_{12}}{1 - q_{11}} \left(t - \sum_{l=1}^t q_{11}^{l-1} \right) = \frac{\bar{\gamma} q_{12}}{1 - q_{11}} \left(t - \sum_{l=0}^{t-1} q_{11}^l \right) = \frac{\bar{\gamma} q_{12}}{1 - q_{11}} \left(t - \frac{1 - q_{11}^t}{1 - q_{11}} \right)$$

The expected number of customers buying in period $t + 1$ at price v_1 is given by those who arrived earlier with valuation v_1 and remained at v_1 until period t , and then either stayed at v_1 or jumped to v_2 , i.e.,

$$\bar{\gamma} \sum_{l=0}^{t-1} q_{11}^l (q_{11} + q_{12}) = \bar{\gamma} (q_{11} + q_{12}) \frac{1 - q_{11}^t}{1 - q_{11}}.$$

We next state our main result here where we show that under some conditions involving the primitives of the problem, the seller is better off implementing a strictly cyclic policy of the form $\pi_c = ((v_2, t), (v_1, 1))$ with t finite.

Proposition 14 *We denote by $t^* = \arg \max_{t \in \mathbb{N}} \{\mathcal{R}(t)\}$. We have the following*

$$0 < t^* < \infty \quad \text{iff} \quad C \ln q_{11} < B - v_1 < C,$$

where

$$B = \gamma v_2 + v_2 \bar{\gamma} \frac{q_{12}}{1 - q_{11}} = v_2 \frac{\gamma q_{20} + q_{12}}{1 - q_{11}} > 0,$$

and

$$C = -v_2 \bar{\gamma} \frac{q_{12}}{(1 - q_{11})^2} + v_2 \bar{\gamma} \frac{q_{11} + q_{12}}{1 - q_{11}}.$$

If these conditions do not hold, then $t^* = 0$ iff $B - v_1 < C \ln q_{11}$, and $t^* = \infty$ otherwise.

Proof. Given the expressions of B and C , we rewrite the revenue function as follows

$$(t + 1) \mathcal{R}(t) = v_1 + B t + C (1 - q_{11}^t).$$

Note that $C > 0$ iff $v_2 < v_1 (1 + \frac{q_{11}}{q_{12}}) (1 - q_{11})$ where the term $(1 + \frac{q_{11}}{q_{12}}) (1 - q_{11}) \geq 1$. In order to prove that it is optimal for the seller to set p_t at v_1 and reset the system, it is enough to prove that t^* the maximizer of \mathcal{R} is finite i.e. $0 \leq t^* < \infty$. Simple calculations show that

$$\mathcal{R}'(t) = \frac{B - v_1 - C + C q_{11}^t (1 - (t + 1) \ln(q_{11}))}{(t + 1)^2}.$$

We denote by $\mathcal{N}(t)$ the numerator of $\mathcal{R}'(t)$. By taking the derivative with respect to t of $g(t) = q_{11}^t (1 - (t + 1) \ln(q_{11}))$, we observe that this quantity is non-increasing in t , from which we conclude that $\mathcal{N}(t)$ is itself monotone in t . Therefore, the equation $\mathcal{R}'(t) = 0$ admits at most one solution. In light of that, we denote by \hat{t}^* the supremum of \mathcal{R} on the positive real line and conclude that a necessary and sufficient condition for $0 < \hat{t}^* < \infty$ is that

$$i.) \mathcal{N}(0) > 0 \quad \text{and} \quad ii.) \mathcal{N}(\infty) < 0.$$

Hence, by putting these two inequalities together we get that $-C(1 - \ln q_{11}) < B - v_1 - C$ and $B - v_1 - C < 0$. We finally get that

$$0 < \hat{t}^* < \infty \quad \text{iff} \quad C \ln q_{11} < B - v_1 < C.$$

Note that under these condition $\mathcal{R}(t) > \mathcal{R}(\infty) \equiv B$ for any $t \geq \hat{t}^*$ and $\mathcal{R}(t) > \mathcal{R}(0)$ for any $t \leq \hat{t}^*$ and hence $t^* := \arg \max_n \mathcal{R}(n) \in (0, \infty)$, which completes our proof.