WORKING PAPERS

Dynamic Pricing with Search Frictions

Daniel Garcia

Mai 2017

Working Paper No: 1703

DEPARTMENT OF ECONOMICS

UNIVERSITY OF VIENNA

Dynamic Pricing with Search Frictions*

Daniel Garcia†

May 24, 2017

Abstract

This paper studies dynamic pricing in markets with search frictions. Sellers have a single unit of a good and post prices in every trading period. Buyers have to incur a search cost to match with a new seller and upon matching they observe the price and the realization of some idiosyncratic match value. There is no discounting but trade ends at an exogenously given deadline. We show that equilibrium involves trading in finitely many trading periods and the volume of trade increases over time. Under mild conditions on the buyer-to-seller ratio and the distribution of valuations, prices decrease at increasing rates as the deadline approaches. We derive the gains from trade in equilibrium and their distribution between buyers and sellers. For the case in which the measures of buyers and sellers coincide, we provide a full characterization of the (unique) equilibrium for a class of distribution functions. We finally discuss implications for market design, including the use of platform fees and cancellation policies.

JEL Classification: D11, D83, L13

Keywords: Consumer Search; Dynamic Pricing; Sharing Economy.

*I thank Marc Goñi, Philipp Schmidt-Dengler, Karl Schlag, Juha Tolvanen and Matan Tsur for helpful discussions and Jose Luis Moraga for his comments on an earlier draft. I also thank various audiences at University of Vienna, JKU Linz and MACCI Annual Conference for their suggestions. The usual disclaimer applies.

†Department of Economics, University of Vienna. Oskar-Morgenstern-Platz 1, 1090 Vienna, Austria. Email: daniel.garcia@univie.ac.at.

1 Introduction

In the *sharing economy*, individual buyers and sellers perform market transactions in fundamentally novel ways. Using a simple mobile application, a user can schedule a car ride with a private driver; book a night at someone's place; or hire someone else to clean her house, assemble her furniture or take care of her pet. Conversely, owners of empty houses, drivers with free time or people willing and able to perform home chores for someone else can use these same digital platforms to provide goods and services for profit. The numbers are staggering: according to a recent report by McKinsey and Co. 162 million people have provided goods or services through online marketplaces in the last 12 months in the US and the EU alone.[1]

While these markets differ in many dimensions, many share some specific features.[2] First, transactions are specific to a certain time frame. Buyers obtain the right to use a good or hire someone to perform a certain task at a particular date. Second, providers have very limited capacity (often a single unit) and they post prices which may vary over time. Third, because goods are heterogeneous, interested buyers have to devote time and effort to search among the different providers. Bound by their limited capacity, sellers adjust their prices over time in order to maximize revenue, while buyers search repeatedly trying to obtain the best possible deal. It is the goal of this paper to provide a simple framework in which to analyze large dynamic markets with search frictions.

More precisely, we set up a very simple model of a large market with search frictions and unit capacity. Sellers are homogeneous and hold one unit of a good that can be rented at some future date, which we refer to as the *deadline*. In each period they set a price and may get randomly matched with an interested buyer. If the buyer accepts the offer, both parties commit to perform the transaction at the deadline. If the buyer rejects, the seller carries her unit of the good to the next period. Similarly, we assume that buyers are homogeneous and demand a single unit of the good. In each period they may decide to remain idle (and wait) or actively search for a seller. In the latter case, they get to observe the price and their idiosyncratic valuation for the good offered but have to pay a search cost. An equilibrium is characterized by a sequence of prices, reservation values and searching decisions such that firms maximize profits given the behavior of buyers and these maximize their expected utility given the sequence of prices.

In order to understand the incentives at play, first note that, from the perspec-

[1]See http://www.mckinsey.com/global-themes/employment-and-growth/independent-work-choice-necessity-and-the-gig-economy

[2]Einav et al. (2016). See also the FTC Report 'The Sharing Economy'.

tive of a seller, the opportunity cost of selling the good in a given period is the continuation profit she expects to obtain in the future. Therefore, the price must be equal to the expected future profit plus a markup that depends on the elasticity of demand. Since the elasticity of demand depends on buyers' reservation values, it evolves over time as they become more pessimistic about their likelihood of finding a suitable match in the ever-shortening future. Similarly, sellers' continuation profits decrease over time, as they run out of future opportunities to sell. The evolution of prices and reservation values, therefore, determines the equilibrium in this market.

Given a sequence of prices, the typical buyer's decision problem has two components. First, he has to decide which offers to accept if he decides to search actively. A natural indifference condition can be obtained as follows: the net surplus he demands in order to accept an offer in a given period must equal the net surplus he will demand in the following period plus the expected net gains from an additional search. The reservation value can thereby be computed as the surplus demanded plus the expected equilibrium price. Second, he has to decide whether to actively search or not. As it turns out, this is a static decision that only compares the expected benefit of search, which depends on the likelihood of obtaining a better draw than the current reservation value, and his intrinsic search cost. In general, there is a unique cutoff value such that consumers search if and only if their reservation value is below this cutoff. Importantly, this cutoff is independent of prices and is equal to the constrained-efficient gains from trade in the market.

These simple observations lead to some striking results. Under a mild restriction on the distribution of match values, we can show that an (un-dominated) equilibrium exists and all equilibria are approximately outcome-equivalent. In equilibrium, buyers' reservation values are decreasing and convex as time approaches the deadline, so that trading probabilities increase over time. As a consequence, net gains from search increase as the market approaches the deadline, which yields a unique time period such that buyers participate in every subsequent period but never before. Given the induced dynamics on demand, the evolution of prices depends on whether buyers' price elasticity increases or decreases over time (which depends on the shape of the distribution of valuations) and on the initial ratio of buyers and sellers in the market. In the particular case in which this ratio equals one, prices are decreasing over time. At the deadline, sellers become effective monopolists with no marginal cost so that the price equals the inverse elasticity of demand.

Our first main result concerns the welfare properties of this equilibrium. If there are equal numbers of buyers and sellers, total welfare is given by the difference between the reservation value of buyers in the first active trading period and the sellers' markup in this period. As mentioned above, this reservation value equals the gains from trade in the constrained-efficient allocation and, therefore, is independent of the sequence of prices, while the markup depends on the inverse of the hazard ratio of the distribution of valuation. As search costs vanish, the reservation value converges to the upper bound of the support of the distribution, so that the equilibrium welfare depends on its shape. For some distributions with increasing hazard rates (like the uniform), the markup converges to zero and, therefore, the equilibrium outcome is approximately efficient. For a large class of distributions, including the exponential or the Pareto, however, the hazard rate does not vanish at the upper bound and the equilibrium remains inefficient as search frictions vanish.[3]

More generally, if the number of buyers and sellers do not coincide, the gains from trade depend on the whole sequence of prices. In this case, it follows that the rents of the long side of the market are bounded by a function of the ratio of buyers and sellers and the markup in the initial period. Interestingly, if this markup becomes sufficiently small as frictions vanish (which ensures efficiency), the long side will only appropriate a tiny fraction of the whole pie.

A final observation from a normative point of view is that trade starts too late. This is because sellers appropriate some of the gains that buyers generate by searching an additional period and so buyers decide to wait even if it is socially inefficient. Therefore, any policy that fosters early search has the potential to increase the total gains from trade.

A very attractive feature of this model is that it allows us to solve in closed form the equilibrium prices and reservation values for certain distributions in the case of equal measures of buyers and sellers. Using the continuous time limit we can also derive the number of trading periods and perform comparative statics on prices and reservation values as a function of the search cost. Not surprisingly, in any given period, increasing search costs leads to higher prices and lower reservation values. Strikingly, however, prices in the first trading period *decrease* in search costs. We also show that the shape of the distribution crucially affects the elasticity of total search as a function of search costs, so that distributions with a long tail induce buyers to increase their total search effort when search costs decrease.

We then use the model to answer some classical questions in Industrial Organization. First, we establish a relation between the pass-through rate in a static environment (which depends only on the curvature of demand) and the elasticity of the initial price with respect to search costs. This relation can be used to derive

[3]For instance, for the Pareto distribution, which induces an iso-elastic demand function, the equilibrium gains from trade are only a fraction of the available surplus.

estimates of the incidence of taxes or merchant fees with basic knowledge of the demand function. Second, we perform the classical comparison between fixed and variable as a source of revenue for the platform. We show that neither policy has implications on the gains from trade but they have different implications for the distribution of the burden between buyers and sellers. In particular, while variable fees have no impact in market outcomes, fixed fees lead to a steeper profile of prices and, therefore, to a lower time horizon, which benefits buyers because they incur lower search expenditures. We finally study the impact on equilibrium outcomes of a very popular feature in many dynamic markets: cancellation policies. We show that cancellation policies have no impact on total gains from trade because they are never taken up in the earlier periods. We then show that if they are actually taken up, they may benefit sellers, since they increase the buyer-to-seller ratio in the market and induce buyers to become less patient.

1.1 Related Literature

The literature on dynamic pricing is lengthy and rich, but focuses mostly on the case of a monopolistic seller with and without commitment who faces a sequence of buyers who may or may not wait for better offers and whose valuation may be independently distributed or may include some correlation.[4] A smaller strand of papers have studied the case of duopoly with myopic buyers (c.f. Dudey (1992) and more recently Martínez-de Albéniz and Talluri (2011)). In general, these models have not been very fruitful because they are plagued with non-existence issues, and even if it exists, equilibrium requires complicated mixed strategies and may not be unique. In many real-world applications, however, a large number of sellers actively compete with one another and buyers, while able to wait for better deals need to incur time and effort in order to discover prices. Meisner (2016) studies an oligopolistic version of Hörner and Samuelson (2011) but focuses on the case in which a monopolist supplier would implement the efficient allocation, in which case an oligopolistic supply side induces the same equilibrium outcome. On the other hand, there is a large literature on dynamic matching models (Lauermann, 2013) where sellers and buyers engage in bargaining in an stationary environment and heterogeneity is on a *vertical* dimension. Our focus is different since we are concerned with *horizontally differentiated* products (e.g. products differing in their location) and we explicitly model a deterministic deadline (as opposed to exponential discounting).

[4]Prominent recent contributions include, Hörner and Samuelson (2011) and Board and Skrzypacz (2016) and references therein

A different strand of the literature has studied dynamic demand uncertainty. Deneckere and Peck (2012) provides a model in which firms compete over a finite horizon with fixed capacities, buyers arrive at random times but may wait to buy in future periods and aggregate demand is highly uncertain. In their model, price competition does not induce marginal cost pricing because firms' may 'bet' on different demand states, so that high price firms will sell in high demand states. Demand uncertainty is an alternative rationale for dynamic pricing and therefore we view our papers as complementary. In addition, our setup is much simpler and we can derive a much sharper equilibrium characterization.

The need for a better understanding of dynamic pricing in competitive setups has been already identified in empirical work. Sweeting (2015) provides a computational model of an oligopolistic markets in which sellers face myopic buyers who arrive over time and who can sample all available products at no cost and applies it to the secondary market for sports tickets. Buyers and sellers are heterogeneous in many dimensions, which translates in a large state space so that computing the equilibrium of the game is not feasible. We assume monopolistic competition without ex-ante heterogeneity and search frictions but model buyers as long-run players and obtain a thorough characterization of equilibrium outcomes.

Finally, the model contributes to the literature on consumer search by introducing dynamic pricing and capacity constraints. To the best of our knowledge only Moraga-González and Watanabe (2016) and Garcia et al. (2015) introduce capacity constraints in markets in which consumers search sequentially over different producers offering heterogeneous products but they abstract from dynamic considerations. Similarly, Coey et al. (2016) study a random search model in which buyers search over time for better prices and sellers sort themselves into a faster auction market and a slower posted-price market, but do not face capacity constraints. [5]

2 The Model

We now present the benchmark model of monopolistic competition and search frictions with capacity constraints and a deadline. We consider a market with a continuum of measure m of buyers and a continuum of measure 1 of sellers. Sellers are ex-ante identical and have one unit of the good that they can supply at no

[5]Dynamic considerations play a role in a handful of papers in the consumer search literature. For instance, in Armstrong and Zhou (2015) oligopolistic sellers price discriminate over time between different buyers. Similarly, in Garcia and Shelegia (2015) firms compete over time in an environment in which consumers learn from each other.

additional cost, while buyers have unit demand.[6] Time is discrete and indexed by $t = 1, 2, ..., T$, so that $t = 1$ corresponds to the deadline.

Buyers search sequentially for sellers without recall. Upon paying a search cost s, buyer i may get matched with a randomly chosen seller j. Upon matching, buyer i learns her valuation for the good of seller j, v_{ij} and its price p_j, from which he derives surplus $v_{ij} - p_j$. The buyer's outside option is normalized to 0. The random utility draw v_{ij} is specific to the match between i and j and has a distribution $F(v)$, with density $f(v)$ and $F(0) = 0$. We shall consider both distributions with bounded support (like uniform) and distributions with semi-infinite support (like exponential). Let $H(v) = \frac{1-F(v)}{f(v)}$ be the inverse hazard rate of $F(v)$. We assume throughout that the *virtual valuation*, defined as $v - H(v)$ is strictly increasing.

The measure of matches in each period equals a fixed proportion of the total potential matches so that if m_t buyers search among n_t sellers, the total number of matches will be $q \min\{m_t, n_t\}$ with $0 < q \leq 1$.[7] One can think of q as the reciprocal of the mean number of periods that a given buyer needs to inspect a seller. For most of the paper, we assume that $m \leq 1$, so that at every period $q \min\{m_t, n_t\} = qm_t$. We study the case of $m > 1$ in Subsection 7.2.

The timing is as follows. At the beginning of each period t, sellers post prices and commit to them for a single period,[8] and buyers decide whether to search or not based on their expectation of prices and match values. Each buyer is either matched with a seller or idle. Upon matching, a buyer observes his realization and the price and decides whether he searches or not. We concentrate throughout on (Perfect-Bayesian) Symmetric Equilibrium, in which sellers use un-dominated strategies and buyers have passive beliefs. The restriction to un-dominated equilibrium rules out no-trade equilibria originating in coordination failures by restricting sellers to choose optimal prices even in those periods in which they expect no demand. Passive beliefs imply that buyers' expectations regarding future prices and trading probabilities do not change following a deviation of an individual seller.[9]

[6] The assumption of zero marginal costs is made for expositional reasons and together with the assumptions on the distribution of valuations implies that trade is always efficient in the last period. For an analysis with positive marginal costs see Section 5

[7] Formally, we assume that a gatekeeper ensures a frictionless matching subject to delay which is distributed a la Poisson. For the case of $m = 1$, as we will see an urn-ball process would deliver the same results, and even for the case of $m \neq 1$, the qualitative results are unaffected.

[8] Posted prices are common in most online markets (FTC 2016). Here we restrict sellers' commitment possibilities to one-period only, but in Appendix C we provide some analysis for the case of partial commitment, so that sellers choose a fixed price for the next $1 \leq k \leq T$ periods.

[9] Passive beliefs is common in the consumer search literature and are natural in the case of infinitesimal sellers. For a recent discussion see Janssen et al. (2014).

2.1 Optimal Search

We begin the analysis describing the problem of a consumer in period $t \leq T$, who faces a sequence of future prices $\{p_\tau\}_{\tau \leq t}$. Although different buyers may have had different histories up to period t, provided that sellers' prices are independent of aggregate variables (see below), these histories play no role in their decisions. It then follows from standard arguments that if a consumer searches in period t, observes a utility draw $v_{i,j}$ and a price p_j, he will accept the offer if and only if $v_{i,j} - p_j \geq w_t - p_t$, which can be defined recursively as follows

$$w_t - p_t = \big(1 - q(1 - F(w_{t-1})\big)(w_{t-1} - p_{t-1}) + q \int_{w_{t-1}}^{\infty} (v - p_{t-1}) dF(v) - s$$

with $w_1 = p_1$. This expression can be easily understood. Upon paying a cost s, the buyer faces a lottery in the following period. With probability $1 - q(1 - F(w_{t-1}))$ he fails to draw a suitable match and move forward to the next period, which yields an expected surplus of $w_{t-1} - p_{t-1}$. With complementary probability, however, he draws a desirable match and buys the good, obtaining $v - p_{t-1}$. Let $V_t = w_t - p_t$ be the continuation value in period t. We can then rewrite this expression as

$$V_t = V_{t-1} + q \int_{w_{t-1}}^{\infty} (v - w_{t-1}) dF(v) - s.$$

It follows that the continuation value in period t equals the continuation value in period $t-1$ plus the net expected gains from search in that period. Since $V_1 = 0$, we have

$$w_t - p_t = V_t = q \sum_{\tau=1}^{t-1} \int_{w_\tau}^{\infty} (v - w_\tau) dF(v) - (t-1)s, \tag{1}$$

which is a dynamic version of the standard stationary threshold rule in sequential search models. Since the gains from search are strictly decreasing in w_t, it follows that $V_{t+1} - V_t \geq 0$ if and only if $w_t \leq w^*$ defined as

$$q \int_{w^*}^{\infty} (v - w^*) dF(v) = s. \tag{2}$$

As a result, consumers will search if and only if $w_t < w^*$. Let T^* be the earliest time period t, such that $w_t \leq w^*$. That is, $w_{T^*} \leq w^* \leq w_{T^*+1}$. Similarly, define $\bar{s}$ as the lowest search cost such that this market exists, i.e.,

$$q \int_{p_1}^{\infty} (v - p_1) dF(v) = \bar{s}. \tag{3}$$

Finally, in what follows we assume that the date at which the market opens formally, T is earlier than T^* so that buyers always start their search at T^*.[10] Therefore, the reservation rule that a buyer uses in the first period in which he searches is (up to the discreteness of time) the same as the one she would use in a stationary setup with the same search cost and distribution of valuations. The intuition from this result stems from the fact that buyers must be (approximately) indifferent between searching or not at T^* so that the change in continuation value from periods T^* and T^*+1 is (approximately) zero.

2.2 Dynamic Pricing

We turn now our attention to sellers. Sellers have a single unit to supply at zero marginal cost and can adjust their prices in every period t. Upon trading, a seller leaves the market immediately. Therefore, we need only consider the problem of seller j holding a unit of the good in period t. She anticipates a future sequence of prices by her competitors $\{p_\tau\}$ and a future sequence of reservation utilities by buyers $\{w_\tau\}$. She also knows that in every future period $\tau < t$ she will match a buyer with probability q_τ whose distribution can be computed recursively from the initial probability $q_{T^*} = qm$ and the distribution of quantities traded in the market at every date $\tau = T, T-1, .., t-1$. Since the measure of traders who leave the market in every period is super-atomless (Podczeck, 2010), a strong Law of Large Numbers applies, ensuring that the probability that a seller matches a buyer in period τ can be computed as,

$$q_\tau = q\frac{m_{\tau+1}F(w_{\tau+1})}{1 - qm_{\tau+1}(1 - F(w_{\tau+1}))}, \tag{4}$$

with $m_{T^*} = m$. Equation (4) can be understood as follows. In period $\tau+1$, there are $m_{\tau+1} \leq 1$ buyers for each seller. The measure of matches in period $\tau+1$ was, therefore, $q_{\tau+1}(1-F(w_{\tau+1}))$ since a match in period $\tau+1$ was successful if $v \geq w_{\tau+1}$. Thus, out of a measure of $m_{\tau+1}$ buyers in period $\tau+1$, only $m_{\tau+1}F(w_{\tau+1})$ remain in period t. Similarly, a proportion of $qm_{\tau+1}(1-F(w_{\tau+1}))$ sellers matched in period $t+1$ so that the remaining sellers constitute a proportion $1 - qm_{\tau+1}(1 - F(w_{\tau+1}))$ of the original sellers. Notice that $q_\tau \leq q_{\tau+1}$ so that the matching rate of sellers decreases over time.

Two extreme cases are particularly interesting. If the initial ratio of buyers-to-sellers equals one, we have that $m_t = m_{t-1} = 1$ for all t, and trading probabilities are stationary. On the other hand, if there are infinitely more sellers than buyers, $m = 0$ and so $m_t = m_{t-1} = 0$ for all t. The case of $m = 0$ corresponds to a model in

[10]A sufficient condition in terms of underlying parameters is $q\int_{Ts}^{\infty}(v - Ts)dF(v) < s$.

which sellers' capacity constraint is not binding, since the likelihood of selling the good in the future is negligible. In this case, each seller will use the static pricing rule,

$$p_t = \frac{1 - F(w_t)}{f(w_t)} = H(w_t),$$

where $H(w_t)$ is the inverse hazard rate and determines the price elasticity of demand if consumers use w_t as a reservation value in period t. For $m > 0$, however, the capacity constraint binds with positive probability and so the seller has to take into account the option value of holding his unit in period $t-1$ when deciding which price to post in period t. The value function representation satisfies:

$$\Pi_t = \max_{p_j}(1 - q_t(1 - F(w_t - p_t + p_j))\Pi_{t-1} + q_t(1 - F(w_t - p_t + p_j))p_j.$$

In a symmetric equilibrium, this yields the necessary condition for p_t.

$$p_t = \Pi_{t-1} + H(w_t). \tag{5}$$

This condition is also sufficient since $v - H(v)$ is strictly increasing as we show formally in the Appendix. Notice that $p_1 = H(p_1) = p^m$ gives the static monopoly price, if it lies in the interior; otherwise, p_1 is such that $F(p_1) = 0$. In every previous period, the expected continuation profits of a seller can also be computed recursively. In particular,

$$\begin{aligned}
\Pi_{t-1} &= \big(1 - q_{t-1}(1 - F(w_{t-1}))\big)\Pi_{t-2} + q_{t-1}(1 - F(w_{t-1}))p_{t-1} \\
&= \Pi_{t-2} + q_{t-1}(1 - F(w_{t-1}))H(w_{t-1}).
\end{aligned}$$

Notice then that the evolution of prices across periods is driven by changes in the elasticity of demand driven by changes in the reservation value of consumers (w_t) and by the changes in the future continuation profits of sellers, as captured by the value of the foregone trading opportunities in earlier periods.

3 Equilibrium

Equation (5) together with Equation (1) implicitly characterize a unique sequence of prices and reservation values for each T^*. If $m = 1$ or $m = 0$, these sequences are independent of T^* and so the equilibrium number of periods is the unique maximizer of V_t. As we saw in Section 2.1 this corresponds to the earliest time period in which the reservation value drops below w^*. On the other hand, if $m \in (0, 1)$, for each candidate initial period l and an associated sequence of prices and reservation values there is a (potentially different) maximizer of V_τ. Let $\tau(l)$ be its maximizer, assuming that all buyers participate in every period. A pure-strategy equilibrium is a fixed point of $\tau(l)$ so that all buyers enter the market in period l and their reservation values satisfy $w_{l+1} > w^* > w_l$. In order to prove existence, however, we shall use the fact that if $w_l = w^*$, buyers are indifferent between entering the market in period l or $l-1$ and we can adjust the participation rate in order to adjust the reservation value.

3.1 Equilibrium Existence

In order to show that such an equilibrium exists, we first show that, for a given candidate initial period l, the sequence $\{w_t^l\}_1^{\tau(l)}$ of reservation values is strictly increasing. This follows directly from the assumption of monotone hazard rates.

Lemma 1. *Suppose that virtual valuations are monotone, i.e. $v - H(v)$ is strictly increasing. Then, $w_t > w_{t-1}$ for all t such that $w_t < w^*$. Therefore, $\tau(l)$ is well-defined.*

Proof. In the Appendix. □

Lemma 1 delivers a simple testable implication for dynamic markets in which buyers are not the long side of the market. Namely, it predicts that the trading probability of a randomly drawn buyer is increasing over time. Conversely, the average match quality is decreasing over time.

Using this monotonicity property, we can construct an un-dominated equilibrium. For every l, let $\{w_t^l\}_1^l$ be the sequence of reservation values assuming full participation in every period $t \leq l$. Namely,

$$w_t^l = w_{t-1}^l + \int_{w_{t-1}^l}^{\infty} (v - w_{t-1}^l) dF(v) - s + \Delta(w_{t-1}, t, l)$$

where $\Delta(w, t, l)$ is the difference in prices between periods t and $t-1$ given some matching rate for sellers q_t^l and some reservation value of consumers w. In particular,

$$\Delta(w, t, l) = H(w) - H(w)(1 - q_{t-1}^l F(w)).$$

The crucial part of the argument shows that $w_t^l > w_t^{l+1}$ if $\tau(l) \geq l$ and $\tau(l+1) \geq l+1$. This follows by induction noticing that the difference between both sequences is difference in the buyer-to-seller ratio in period l. Since the difference equation is monotone with respect to m, it follows that $w_t^l > w_t^{l+1}$. This monotonicity carries over to the function $\tau(l)$; that is, $\tau(l) \geq \tau(l+1)$. Since $\tau(\infty) = T_1 < \infty$. Denote by l^* the lowest initial period such that $\tau(l^*) \leq l^*$. In the proof of Proposition 2 we show that either $l^* = \tau(l^*)$ or we can construct an equilibrium with partial participation in period l^* so that $w_{l^*} = w^*$. Hence,

Proposition 1. *If $v - H(v)$ is non-increasing, an un-dominated equilibrium exists and all un-dominated equilibria yield approximately the same surplus. Further, if either $m = 0$, $m = 1$ or $F(v)$ is exponential, the equilibrium is unique.*

3.2 Welfare

We now turn our attention to the welfare properties of these equilibria, as represented by the total gains from trade (net of search costs). This is the main object of interest both for market designers and policy makers.

Fix an equilibrium $\{w_t\}, \{p_t\}$ with a corresponding starting period T^* so that $w_{T^*} \leq w^* \leq w_{T^*+1}$ and $w_1 = p_1 = H(p_1)$. The gains from trade can be computed as the sum of the consumers' continuation value in the initial period and the expected profit of an entrant seller. More precisely,

$$\begin{aligned} W &= mV_{T^*} + \Pi_{T^*} \qquad (6) \\ &= m\big(w_{T^*} - p_{T^*}\big) + \big(p_{T^*} - H(w_{T^*})\big) + qmqm\big(1 - F(w_{T^*})\big)H(w_{T^*}) \\ &= m\big(w_{T^*} - H(w_{T^*})\big) + (1-m)\big(p_{T^*} - H(w_{T^*}) + qm\big(1 - F(w_{T^*})\big)H(w_{T^*}) \\ &\approx m\big(w^* - H(w^*)\big) + (1-m)\big(p_{T^*} - H(w^*)\big). \end{aligned}$$

The first equality follows by definition of the gains from trade. The second equality uses the fact that $\Pi_t = \Pi_{t-1} + qm(1 - F(w_{t-1}))H(w_{t-1})$. The third equality is simply the product of a reorganization of terms. The last condition follows when taking the continuous time limit as the per-period matching rate converges to zero so that $w_{T^*-1} \to w^*$.[11]

Notice first that if $m = 1$, then $W \approx w^* - H(w^*)$, which is independent of the sequence of prices and only depends on the first period's markup, the distribution of valuations and the search costs. Indeed, gains from trade depend only on the equilibrium outcomes in the initial period, because buyers and sellers internalize all future transactions in their value and buyers adjust their search behavior so that, regardless of the future sequence of prices, their reservation value at the initial

[11]Formally, we use the symbol $\approx$ to represent the following relation between two quantities,

$$x(q,s) \approx y \Longleftrightarrow \lim_{\tilde{q}\to 0} x(\tilde{q}, s\tilde{q}) = y.$$

period in which they search is the same. In this case, the role of prices is simply to redistribute rents between buyers and sellers but do not have consequences for the allocation, except those of the first trading period. It follows that if $H(v)$ is non-increasing, the (static) monopoly price is an upper bound for the difference in the gains from trade between the constrained-efficient allocation and the equilibrium.[12]

The following proposition summarizes the implications of this result for different distributions of valuations.

Proposition 2. *Suppose that $s > 0$. In any equilibrium, gains from trade are approximately given by (6). Moreover, as frictions vanish, if $m = 1$.*

1. *if $H(v) \to 0$ as $F(v) \to 1$ (e.g. $F(v)$ is uniform), equilibrium is approximately efficient.*

2. *if $H(v) \to B$ as $F(v) \to 1$, and $H(v)$ is non-increasing ($F(v)$ is exponential or logistic), equilibrium welfare converges to $w^* - B$.*

3. *if $F(v)$ is Pareto so that $H(v) = av$ for some $a < 1$, then, the equilibrium welfare converges to $(1-a)w^*$.*

Figure 1 depicts the relative welfare as a function of s for the uniform, Pareto and exponential distributions.

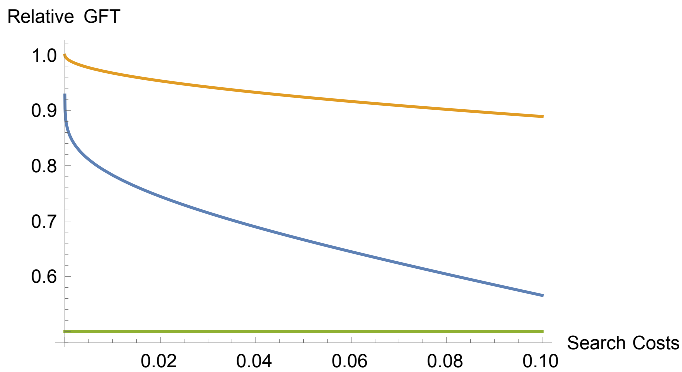

Figure 1: Relative Gains from Trade for different Distributions.

Notice that the invariance of gains from trade with respect to the sequence of prices crucially relies on the endogenous adjustment of the time-span of the market. For instance, if $T < T^*$ so that trade would start immediately, different sequences of prices would lead to different realized welfare gains.

We now consider the case of a buyer's market so that $m < 1$. If buyers are the short side of the market, aggregate welfare depends on the distribution of rents

[12]If $H(v)$ is non-decreasing, then the static monopoly price is a lower bound for this difference.

between buyers and sellers, and, thereby, also on the whole sequence of prices. In equilibrium, at most m units will be sold so that $1 - m$ will always remain idle. However, at the outset, sellers do not which of those units will be sold and so the expected profit per seller decrease with m, reducing their rents and increasing the share of the pie accruing to buyers. Since in equilibrium there is inefficient delay because buyers do not fully appropriate all their surplus, per-buyer welfare is higher if $m < 1$.

To further elaborate on this argument, suppose that there is a gatekeeper that blocks the entry of these idle $1 - m$ sellers so that the market has an equal measure of active buyers and sellers. This exclusion has a direct effect on per-seller revenue, leading to an increase in the value of holding a unit of the good for those sellers who remain, which is reflected in higher prices in earlier periods. Since the price in the last period is always the (static) monopoly price, the profile of prices necessarily becomes steeper. This induces buyers to wait more before purchasing, inducing more distortions and lower welfare. The extent to which this is true depends on the relative magnitude of m. If $m \to 0$ and $H(v)$ is non-decreasing, industry profits also tend to zero so that restricting entry does not change prices, while if $m \to 1$, blocking entry leads to little redistribution and, therefore, small effects on prices. Therefore, there is some intermediate m^* such that per-buyer gains from trade are maximal for m^*.

Flipping this argument, we can obtain a bound for the rents accruing to sellers when they are the long side of the market. Notice that it must be the case that $W \leq mw^*$ since in any equilibrium a buyer is willing to accept $w^* - p_{T^*}$ in the first period and sellers optimality condition implies that $\Pi_{T^*-1} \approx \Pi_{T^*} \leq p_{T^*}$. As a result,

$$\Pi_{T^*} \leq H(w^*)\frac{m}{1-m}\big(1 - qm(1 - F(w^*))\big) \leq \frac{m}{1-m}H(w^*). \tag{7}$$

Further since $p_{T^*} \approx \Pi_{T^*} + H(w^*)$, we have $p_{T^*} \leq \frac{1}{1-m}H(w^*)$ and, therefore, $V_{T^*} \geq w^* - \frac{1}{1-m}H(w^*)$. Equation (7) gives us an idea of the effect of competition on seller's profits. If there are many sellers for each buyer in the market or if, at the outset, buyer' demand is very elastic, sellers will appropriate a very small fraction of the surplus. While the first part is rather obvious, the second is not so. The idea is as follows. Buyers arrive at the market with a reservation value of w^* which induces a price elasticity of $-H(w^*)^{-1}$. Over time, this elasticity may be changing and eventually they will become so desperate that prices will tend to monopoly. If $m < 1$, however, sellers are more *impatient* than buyers because as time goes by their matching rate drops and, therefore, they cannot hold for those higher prices in the future.

Proposition 3. *In equilibrium,*

1. *if $H(v)$ is decreasing and converges to zero as v approaches its upper bound, then for every $m < 1$, $\Pi_t < \epsilon$ for all $s < s(m, \epsilon)$.*

2. *if $H(v)$ is non-increasing and converges to a constant as $v \to \infty$, then for every $m < 1$, the share of the pie accruing to buyers converges to one as $s \to 0$.*

For distributions in the Generalized Pareto Class Bulow and Pfleiderer (1983), which includes the Pareto, the exponential and the uniform as special cases, we have that $H(v)$ is linear in in v. If the slope of $H(v)$ is negative, then the first case applies and the rents of sellers are uniformly bounded by a function of s for any $m < 1$. If, on the other hand, $H(v)$ is constant, so that $F(v)$ is exponential, the share of the total rents that sellers appropriate vanish as frictions disappear. Finally, if the slope is increasing, profits represent at most $m/(1-m)$ of the total surplus.

A corollary of Proposition 4 is that sellers need not benefit from improvements in the matching technology if they are the long-side of the market.

Corollary 1. *Suppose that $1 - F(v)$ is strictly log-concave. Then, there exists some $0 < m^* < 1$, such that for all $m < m^*$, $\lim_{s \to \bar{s}} \Pi_t > \lim_{s \to 0} \Pi_t$.*

This is very intuitive. Higher search costs increase the market power of sellers and, therefore, they can extract a bigger fraction of the total pie. For the case of log-concave distributions and small enough m, this gain more than compensates the loss in terms of a lower quantity sold. Obviously, competition and lack of commitment are crucial for this result, for a single monopolist who can commit to a single price would never lose from increases in search costs. As we will see in Section 4, however, sellers do benefit from reductions in search costs if the market is balanced ($m = 1$).

These formulae also can be used to obtain some insight on price dynamics. To see this first notice that as search costs vanish, gains from trade converge to first-best and so the probability that a buyer exits the market at the deadline without buying must converge to zero. Therefore, each seller sells with probability approaching m. Rewriting the expression for profits we have.

$$\bar{p} = \frac{1}{m}\Pi_{T^*} = \frac{1}{1-m}H(w^*)$$

where $\bar{p}$ is the average transaction price of a seller, and since $p_{T^*} = \Pi_{T^*} + H(w^*) = \frac{1}{1-m}H(w^*)$, if $H(w^*) \to 0$, we have that $\bar{p} = p_T$. Notice also that $p_1 = H(p_1) > p_{T^*}$ and so prices cannot be monotonically decreasing. This is in stark contrast with models of experimentation where sellers learn about demand and prices decrease as they become more pessimistic, regardless of market conditions.

Delay

Consider now an intervention that would force consumers to start searching one period before. The gains from trade in period $T^* + 1$

$$\begin{aligned} W_{T^*+1} - W_{T^*} &= \int_{w_{T^*+1}}^{\infty} (v - w_{T^*+1}) dF(v) + qm(1 - F(w_{T^*+1}))H(w_{T^*+1}) - s \\ &\approx qm(1 - F(w_{T^*+1}))H(w_{T^*+1}) > 0 \end{aligned}$$

since $w_{T^*+1} \approx w^*$. Therefore, if $F(w^*) < 1$, $W_{T^*+1} - W_{T^*} > 0$. Because consumers do not fully appropriate all the rents they create when searching, there is inefficient delay in that the search starts too late.

Proposition 4. *Suppose that $m > 0$, then in any equilibrium there is inefficient delay.*

Inefficient delay is a direct result of the wedge between the social value of an additional search and the private value accruing to the consumer. This wedge depends on the share of rents that sellers can appropriate. A redistributive policy may then lead not only to lower initial prices but also to an increase in the number of trading periods.

This result also establishes a connection between initial prices and equilibrium welfare. An estimate of the initial markup can be obtained by using the standard demand elasticity formulas if the researcher has access to some cost shifters (e.g. shipping costs or merchant fees). This estimate can then be used to obtain a direct estimate of the inefficiency in equilibrium.

From the perspective of the design of optimal platforms, an interesting observation then concerns the issue of time-varying fees. Most transactions occur in the final periods but welfare can be enhanced if sellers reduced their markups in the initial periods. Interestingly, in a recent work Sweeting (2015) proposes a counterfactual commission system for Stubhub involving a discount for trade one week before the deadline. Our model suggests that indeed such reforms are likely to induce more trade and higher profits.

3.3 Prices and Quantities

As we discussed in the Introduction, the bulk of the literature on dynamic pricing has neglected the issue of competition. This may be because of a variety of reasons including tractability (both in theory (Martínez-de Albéniz and Talluri, 2011) and in empirical applications (Sweeting, 2015)), a particular focus in certain industries with few competitors (e.g. airlines), and the striking success of the monopolistic model in replicating salient features of the data (Sweeting, 2012). Therefore, it is important to provide a set of clear and empirically testable predictions that could offer guidance to researchers considering introducing competition in their dynamic pricing models.

A first observation is that if the measures of buyers and sellers are the same ($m = 1$), the capacity constraint of a single seller is just as tight as the market-wise capacity constraint so that competition yields exactly the same profile of prices and quantities as the case in which a single manufacturer supplies the same number of units would.[13] Thing are, however, very different if $m < 1$, for in such a case the empirical co-evolution of prices and quantities is rather different. A monopolistic seller facing limited capacity will always *increase* the price following a period of high sales and will let the price drop if the demand has been sluggish in the previous period. On the other hand, a competitive seller's behavior will always *decrease* the price following a period of higher sales because her own capacity constraint has not been softened but the total future demand drops. Thus, whether a market is better modelled through competition or monopoly may be tested directly using the correlation between lagged demand and current prices. Moreover, our model predicts that the price change brought about by earlier sales is bigger the higher is the ratio of buyers to sellers. For instance, if $m = 0$, profits are zero in every period and so prices are independent of earlier sales.

4 Comparative Statics

We now focus on the case in which $m = 1$, so that matching probabilities are constant over time.[14] As a result the sequence of reservation values and prices are independent of T^* and, therefore there is always a unique equilibrium. Furthermore, if valuations are distributed according to the Generalized Pareto Distribution, the inverse hazard rate of the distribution of valuations is linear (Ausubel et al., 2014), which allows us to derive the difference equation that describes the evolution of reservation values over time and derive the associated differential equation in the continuous time model. For some distributions within this class, including the exponential, the uniform (linear demand) and the Pareto (with scale $v = 2$, representing a constant

[13] Meisner (2016) provides a similar result, but in his case the allocation must also be efficient.

[14] While we do not model entry or exit, the rents of the long side of the market are low enough that a balanced market is probably a good approximation. Furthermore, by Propositions 2 and 3, the case of $m = 1$ provides a lower-bound in terms of welfare. See also Subsection 6.2. for the case of a sellers' market.

elasticity of 1/2), we can then solve for the time-span of the market and derive the initial prices, seller expected profits and buyers' surplus as a function of the fundamental parameters only.

4.1 Evolution of Reservation Values

Suppose then that both sides of the market are evenly matched, $m = 1$, and the distribution of valuations belongs to the Generalized Pareto Distribution class, so that $F(v) = 1-(1+\xi\frac{w-\mu}{\sigma})^{-1/\xi}$.[15] In this case, $H(w) = \sigma-\xi\mu+\xi w$, with $\xi < 1$. This class includes the uniform, the exponential, the Pareto as well as some particular cases of other distributions like the $Beta(1, b)$. Notice that as long as $\xi < 1$, virtual valuations are increasing so that a unique (un-dominated) equilibrium exists.

For this class of distributions, we can express the gains from search for a consumer using a reservation value of w as,

$$\int_{w}^{\infty}(u-w)dF(u) = \frac{1}{1-\xi}H(w)(1-F(w)).$$

In terms of demand theory, this implies that consumer surplus and monopoly profits are proportional to each other, so that increases in consumer surplus always coincide with consumer surplus. In the context of a model with search, this implies that the private returns from search are a constant fraction of the social surplus they generate. Since the price in period t can be expressed as,

$$p_t = \Pi_{t-1} + H(w_t) = q\sum_{\tau=1}^{t-1}(1-F(w_\tau))H(w_\tau) + H(w_t),$$

we can obtain the reservation value in period t as,

$$w_t = q\frac{2-\xi}{1-\xi}\sum_{\tau=1}^{t-1}(1-F(w_\tau))H(w_\tau) + H(w_t) - (t-1)s.$$

Further, because $H(w_t) = \sigma - \xi\mu + \xi w_t$, we can rewrite the previous expression as

$$w_t = q\frac{2-\xi}{(1-\xi)^2}\sum_{\tau=1}^{t-1}(1-F(w_\tau))H(w_\tau) - \frac{1}{1-\xi}(t-1)s + \frac{\sigma-\xi\mu}{1-\xi}. \tag{8}$$

Therefore, the reservation value in period t is a linear function of the sum of the

[15] In demand theory, if buyers' valuations are drawn from a GPD, their demand function satisfies the constant-curvature assumption so that in a monopolistic market with constant marginal costs the pass-through rate is constant.

expected consumer and producer surpluses in every future period net of search costs plus a term that depends on the hazard rate. As it happens, this term coincides with the static monopoly price if it is interior (i.e. if $\sigma > \mu$). Notice further that w_t is increasing in t and concave since $(1 - F(w))H(w)$ is decreasing in w as long as $H'(w) = \xi < 1$. This property of the evolution of the reservation values directly translates to prices and, as a result, as the deadline approaches, prices decrease at increasing rates.

Recall that w^* provides a good approximation of w_{T^*} when (q, s) are small enough. In the case of the Generalized Pareto Distribution, we can also obtain a closed-form solution for w^* as a function of the parameters of the distribution and the search cost only as,

$$w^* = \frac{1}{\xi}\left(\frac{s(1-\xi)}{q\sigma}\right)^{\frac{\xi}{\xi-1}}\sigma - \frac{\sigma - \xi\mu}{\xi}.$$

If $\xi < 0$, the support of the distribution is $[\mu, -\frac{\sigma-\xi\mu}{\xi}]$ so as $s \to 0$, w^* approaches the upper bound at a rate $s^{\frac{\xi}{\xi-1}}$. On the other hand, if $\xi > 0$, the support is $[\mu, \infty)$ and as $s \to 0$, the second term vanishes so that w^* tends to $ks^{\xi}\xi - 1$ for some $k > 0$. Notice also that w^* is homogeneous of degree zero in (q, s) so that only their ratio matters and if $\xi \geq 0$, as $s \to 0$, w^* is independent of μ and decreasing in σ for $\xi \leq 0$.

We can now use Eq. (8) to obtain an expression of the price that a firm would charge in the first period in which she expects consumers to be active, T^*,

$$\begin{aligned} p_{T^*} &= q\sum_{\tau=1}^{T^*-1}(1 - F(w_\tau))H(w_\tau) + H(w_{T^*}) \qquad (9)\\ &\approx \frac{1}{2-\xi}w^* + \frac{1}{2-\xi}\frac{\sigma-\xi\mu}{1-\xi} + \frac{1-\xi}{2-\xi}(T^*-2)s. \qquad (10)\end{aligned}$$

According to (9) the initial price is a linear combination of the (second-best) gains from trade and the initial price (if interior) plus a share of the total search expenditure. Using (9) we can derive the profits that a seller expects to make when entering the market. To see this notice that $\Pi_{T^*-1} \approx p_{T^*} - H(w^*)$, which yields,

$$\begin{aligned} \Pi_{T^*} &= q\sum_{\tau=1}^{T^*}(1 - F(w_\tau))H(w_\tau)\\ &= \frac{1-\xi}{2-\xi}(w^* - H(w^*)) + \frac{1-\xi}{2-\xi}(T^*-2).\end{aligned}$$

As a result, sellers obtain a fraction $\frac{1-\xi}{2-\xi}$ of the equilibrium gains from trade and the buyers' search effort. The fraction of the social surplus accruing to sellers is, therefore, bounded below by $\frac{1-\xi}{2-\xi}$, their static share. Similarly, buyers' surplus is

$$V_{T^*} \approx w^* - H(w^*) - \Pi_{T^*} = \frac{1}{2-\xi}(w^* - H(w^*)) - \frac{1-\xi}{2-\xi}(T^* - 2).$$

So that buyers extract at most their static fraction of the gains from trade. We now summarize these results in the following Proposition.

Proposition 5. *Suppose that the distribution of valuations belongs to the GPD and $m = 1$. Then, there exists a unique equilibrium. Further,*

1. *Prices and reservation values decrease at increasing rates over time.*
2. *Buyers appropriate at most a fraction $\frac{1}{2-\xi}$ of the social surplus.*
3. *If search costs increase,*
   - *The equilibrium price in any active period (after the change), decreases.*
   - *The quantity traded in any period in any active period increases.*

A typical equilibrium outcome is depicted in Figure 2. The blue line depicts the reservation value as a function of the number of periods before the deadline for consumers whose valuations are exponentially distributed with unitary mean ($\xi = 0, \sigma = 0$). For $t > T^*$, w_t is constant since consumers are too picky to find search profitable. Eventually, w_t decreases as consumers run out of options and face decreasing prices (red line). The green line depicts the expected continuation profit of a seller in period t and equals the price minus the (constant) markup. The total welfare in the market can be obtained by subtracting the difference between the red and the green lines from the value of the blue line.

4.2 Examples

For some distributions, the differential equation describing the evolution of reservation values can be solved analytically and used to obtain a closed-form characterization of T^*. In particular, the differential equation

$$\dot{w} = q\frac{2-\xi}{(1-\xi)^2}(\sigma + \xi(w-\mu))(1+\xi\frac{w-\mu}{\sigma})^{-1/\xi} - \frac{1}{1-\xi}s,$$

with boundary conditions $w(0) = \max\{\frac{\sigma-\xi\mu}{1-\xi}, \mu\}$ and $w(T^*) = w^*$ can be solved if $F(v)$ is uniform, exponential or Pareto with shape parameter 2.

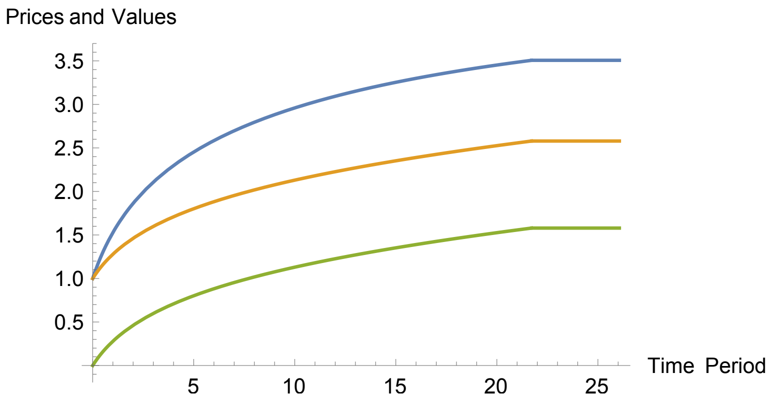

Figure 2: Prices (red), Profits (green) and Reservation Values (blue) for exponentially distributed valuations $\sigma = 1$ and $s = 0.004$.

Exponential Distribution

If valuations are distributed exponentially according to $F(v) = 1 - e^{-\lambda v}$, the Hazard Rate is constant and equal to λ, so that $\xi = 0$ and $\sigma = \lambda^{-1}$ and the equilibrium gains from trade, as shown in Proposition 3, equal $w^* - H(w^*) = w^* - \sigma$, with

$$w^* = \sigma \ln\left(\frac{\sigma q}{s}\right)$$

It follows that the differential equation describing the evolution of reservation values is

$$\dot{w} = 2q\sigma e^{-\frac{w}{\sigma}} - s \tag{11}$$

with initial conditions $w(0) = \sigma$ and $w(T^*) = w^*$. This equation can be solved in closed-form to obtain

$$T^* \approx \frac{\sigma}{s} \ln\left(2 - \frac{es}{q\sigma}\right). \tag{12}$$

Therefore $T^* s$ is decreasing in s but the elasticity of the total expenditure with respect to s is proportional to s and therefore very small in magnitude. Since the exponential distribution has a constant hazard rate, the gains from search decrease (almost) proportionally as w increases so that, if s is small enough, a one percent increase in the search cost leads to a one percent increase in the number of active trading periods. As a result, initial prices, profits and buyers' surplus decrease in search costs and the share of rents accruing to buyers also decreases in search costs and is bounded above by $1/2$. In the static monopoly model, total surplus is split evenly among buyers and sellers, while in this dynamic market sellers obtain some extra surplus due to buyers' search effort.

Perhaps surprisingly, as search frictions vanish prices in the initial period increase. More trading periods increase the opportunities to trade, pushing prices upwards. On the other hand, if we compare prices in a *given* period, with higher search costs, consumers are less picky and, therefore, the probability of future trade is higher so that prices decrease as search frictions decrease. This result is depicted in Figure 3.

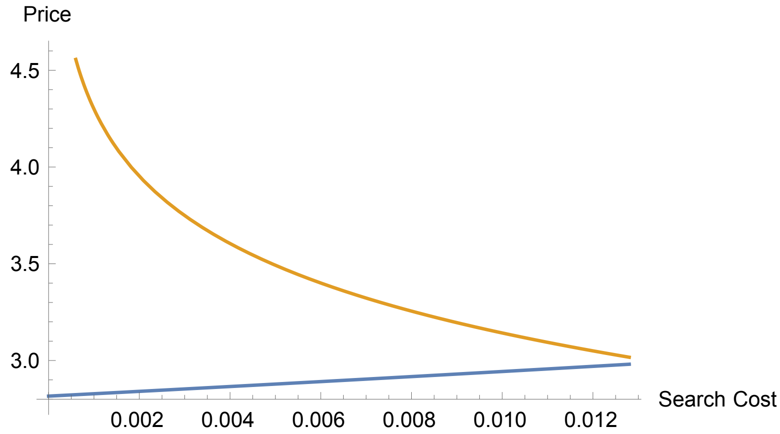

Figure 3: Prices in periods $T = T^*$ and $T = 50$ as a function of search costs.

Uniform Distribution

If valuations are uniformly distributed, so that $F(v) = \frac{v-\mu}{\sigma}$, the inverse hazard rate is decreasing and equal to $\sigma + \mu - v$, so that in equilibrium sellers charge a decreasing markup. The price at the deadline (p_1) satisfies $p_1 = \sigma + \mu - p_1$, and the reservation value in the initial period is $w^* = \mu + \sigma - \sqrt{\frac{2s}{q}}$. As shown in Proposition 3, equilibrium gains from trade are $\mu + \sigma - 2\sqrt{2s/q}$, which converges to w^* as $s \to 0$, so the market becomes approximately efficient as search frictions vanish. The differential equation is

$$\dot{w} = \frac{3}{2\sigma} q(\mu + \sigma - w)^2 - \frac{1}{2} s$$

with boundary conditions $w(0) = p_1$ and $w(T^*) = w^*$. This gives $w(t)$ as the solution of Riccati equation with constant coefficients (Egorov (2007)). Normalizing $q = 1$ and for s small enough, the total search expenditure can be well approximated by,

$$C(s) = T^* s \approx \sqrt{\frac{8\sigma s}{3}} A(s, \sigma, \mu) \tag{13}$$

This function is increasing in s for s sufficiently small so that the number of periods can be well-approximated for a continuum (i.e. $s < (10\sigma)^{-3}$). Further, $C(0) = 0$ and $C(s)$ is increasing in σ and μ. Using these observations we obtain the following expression for the value of a buyer who begins her search in period T^*,

$$V_{T^*} \approx \frac{1}{3} - \frac{2}{3}\sqrt{\frac{2s}{q}} - \frac{2}{3}C(s),$$

which is decreasing in s, and for the final price,

$$p_{T^*} \approx \frac{2}{3} - \frac{1}{3}\sqrt{\frac{2s}{q}} + \frac{2}{3}C(s),$$

which can also be shown to be decreasing in s since the marginal effect of a reduction in search costs on total effort is an order of magnitude lower than its effect on total surplus. Naturally, then, the share of total surplus accruing to buyers decreases in search costs and converges to 1/3 as $s \to 0$. Notice that 1/3 is also the static share of rents for buyers.

Pareto Distribution

If valuations are distributed according to a Pareto with shape parameter $a = 2$ and lower bound $x_m = 1$, so that $F(v) = 1 - v^{-2}$, the inverse hazard rate is simply $v/2$. Notice then that the static monopoly price is given by the corner solution $p_1 = w_1 = 1$ and the threshold reservation value is $w^* = q/s$. As shown in Proposition 3, equilibrium gains from trade as $w^*/2$. The associated differential equation is,

$$\dot{w} = \frac{-3q}{w} + 2s,$$

which can be solved in closed-form to obtain

$$T^* \approx \frac{3q\Big(\ln\big(3q - 2s\big) - \ln(q)\Big) - 2(q - s)}{4s^2},$$

so that $T^* s \approx E(q, s)/s$ which explodes as $s \to 0$. In this case the gains from search decrease faster than the cost as w increases because of the long tail of the Pareto. The value of a buyer who enters the market in the first trading period can be written as

$$V_{T^*} = \frac{1}{3}\frac{q - E(q, s)}{s} - \frac{1}{3},$$

decreasing in s. Similarly,

$$p_{T^*} = \frac{1}{3}\frac{2q + E(q,s)}{s} + \frac{1}{3}$$

which is also decreasing in s. The share of the gains from trade accruing to buyers is less than 45% and decreases in search costs.

4.3 Summary

Remarkably, these different distributions yield a consistent picture. First, prices and reservation values decrease at increasing rates as the deadline approaches. This is indeed consistent both with casual observations and the available empirical evidence. For instance, there are a number of commercial software providers that help *hosts* in Airbnb who want to implement simple dynamic pricing schemes. By and large, they suggest that hosts should keep the base-rate price unchanged until the last week or so and then start decreasing it at increasing rates. In the last period, they suggest some 50% markup over the reported marginal cost of the seller, while the average markup is approximately 100%.[16] Sweeting (2012) and Sweeting (2015) provide substantial evidence that in the secondary market for sports tickets, prices decrease at increasing rates as the event approaches and that trading probabilities are increasing and convex.

In Figure 4 we depict the trading probabilities (red line) and the unsold quantities (blue line) at each period t for the exponential case with the search cost calibrated so that the total quantity sold is approximately 75% of the available units ($s = 0.011$).

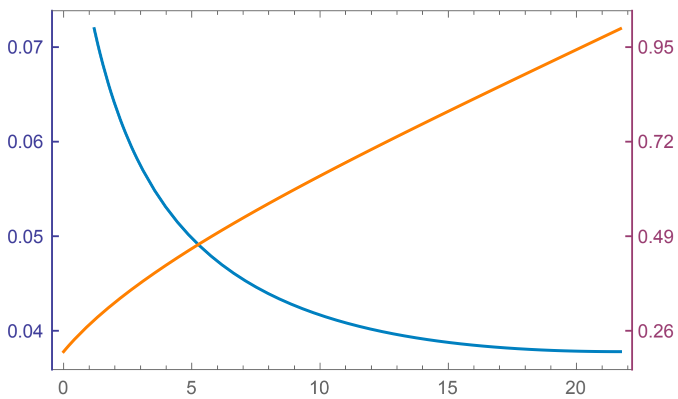

Figure 4: Unsold Quantities and Trading Probabilities over time.

Second, in markets with lower trading frictions (as captured by search costs),

[16]See, for instance, BeyondPricing.com

prices and trading probabilities, *in a given time period* are lower. For instance, the model predicts that the establishment of an online marketplace in a decentralized market leads to a decrease in prices in those periods that are sufficiently close to the deadline. Equilibrium profits and consumer surplus in the market, however, do increase as search costs vanish. The intuition is that lower search costs have a direct effect on the number of periods in which buyers actively search, increasing the likelihood that a given unit is sold and, thereby, increasing their option value. In a way, the positive effect of search intensity on prices stems on the activation of an extensive margin (the number of periods) rather than an intensive margin (a change in the reservation values) and, therefore, is related to the results in Moraga-González et al. (2015).

Third, the elasticity of the number of trading periods with respect to search costs depends on the shape of the distribution of valuations. Distributions with decreasing hazard rates have fat tails and induce buyers to increase their search effort when search costs drop while distributions with increasing hazard rates allow buyers to slack-off. The expenditure in search effort has an important redistributive effect, so that higher search costs lead to lower rents for buyers as search costs increase.

5 Pass-Through

The welfare and redistributive effects of many policies depend on the response of prices to changes in marginal costs. While in most static models of search with differentiated products, cost is passed one-to-one to consumers, in our dynamic setup the degree of pass-through is determined by the elasticity of demand and the response of search effort to changes in prices. In order to characterize the pass-through rates of the model, therefore, assume that sellers face a cost c to deliver their unit. This cost may be intrinsic to the service (e.g. cleaning of an apartment), represent a tax or a fixed fee charged by a platform.[17]

In every period t, a seller carrying a unit will set up a markup equal to her continuation profit, so that

$$p_t - c = \Pi_{t-1} + H(w_t).$$

[17] Variable taxes or fees do not have effects on equilibrium outcomes since we do not explicitly model entry and buyers have unit demand.

Solving for w_{T^*} we obtain,

$$w_{T^*} = \frac{\sigma - \mu\xi + c}{(1-\xi)} + \frac{2-\xi}{(1-\xi)^2} q \sum_{\tau=1}^{T^*-1} (1 - F(w_\tau))H(w_\tau) - \frac{1}{1-\xi} s(T^* - 2) = w^*,$$

because w^* is independent of c. Hence, we can solve for p_{T^*} as

$$p_{T^*} = \frac{1}{2-\xi}(w^* + \sigma - \mu\xi + c) + \frac{1-\xi}{2-\xi} s(T^* - 2).$$

Recall that if $\sigma + c \geq \mu$, this expression can be rewritten as,

$$p_{T^*} = \frac{1}{2-\xi} w^* + \frac{1-\xi}{2-\xi} p_1(c) + \frac{1-\xi}{2-\xi} s(T^* - 2).$$

where we use the fact that $p_1(c)$ is the static monopoly price. Notice then that T^* now depends on c only through the initial price. Absent any effects on the total search expenditure, therefore, the dynamic pass-through rate of a unit increase in the marginal cost is $(2-\xi)^{-1} < (1-\xi)^{-1}$, which is the pass-through rate at the deadline. In addition, higher marginal costs lead to an increase in the price at the deadline which represents the initial condition for the difference equation describing the evolution of w_t. Therefore, the total number of periods, T^*, and the associated search effort decrease as c increases. As a result, the pass-through rate of the initial price, which determines the change in the distribution of rents across buyers and sellers, is at most $(2-\xi)^{-1}$.[18]

For the exponential case, we can solve for the price as a function of t in the continuous time limit as well as the equilibrium quantities in every period so that we can go further and study the effect of increases in marginal costs on the sequence of prices. Since prices are homogeneous of degree zero in (q, s), throughout we assume that $q = 1$. The price at time t equals,

$$p(t) = \sigma + c + \frac{\sigma}{2} \ln \left(\frac{e^{\frac{\sigma+c}{\sigma}} s + 2\sigma \left(e^{\frac{st}{\sigma}} - 1 \right)}{e^{\frac{c+\sigma}{\sigma}} s} \right).$$

[18]In models of imperfect competition, under GPD, pass-through rates can be written as $(1 - \xi\theta)^{-1}$, where θ measures the degree of competition (see, e.g. Weyl and Fabinger (2013)). Typically, $0 \leq \theta \leq 1$, so that pass-through rates are substantially higher in static frameworks than in this dynamic market.

As a result, the pass-through rate in time t, $\rho(t)$ equals

$$\rho(t) = \frac{p(t)}{\partial c} = 1 - \left(\frac{e^{\frac{\sigma+c}{\sigma}} s}{\sigma - \sigma e^{\frac{st}{\sigma}} - 2} \right)$$

Naturally, at the deadline, $\rho(0) = 1$, while

$$\rho(T^*) = \frac{\sigma}{2\sigma - e^{\frac{\sigma+c}{\sigma}} s}.$$

As $s \to 0$, $\rho(T^*)$ approaches $(2 - \xi)^{-1}$ while as frictions increase so that the market collapses, $\rho(T^*) \to \rho(0)$. In general, $\rho(t)$ is decreasing and convex in t.

Pass-through rates are also useful as a method for identification of the shape of the distribution of valuations and the search costs. For this purpose, suppose that we have a good estimate of marginal costs and we have the sequence of prices and quantities. Provided that we have enough observations and we are able to assume exponentially distributed valuations, we can consistently estimate σ from the transactions sufficiently close to the deadline. In principle, using σ we could directly estimate s from $\rho(T^*)$. However, by definition, the number of transactions at T^* will be very limited and subject to outliers. A more practical approach is to use data on quantities and prices over time and infer from those the search cost. As it turns out, as we show in Figure 5, one can infer search costs from either the (un)weighted average over time of the pass-through rates (i.e. $\int \rho(t)q(t)dt$) but not from the pass-through rate on the average transaction price (i.e. $\rho(\int p(t)q(t)dt)$). This is because higher marginal costs lead to a reduction of total quantity and a shift from earlier to later periods so that the average transaction price becomes non-monotone on search costs. Hence, the model suggests that an effective way to recover search costs from dynamic markets is to compute the average change in the time-specific price following a change in marginal costs.

6 Cancellation Policy

In many of these markets, buyers hold the right to renege on their promise in a later period with or without some compensation to the seller (Xu (2005)). Many platforms such as Uber and Airbnb establish cancellation policies with the stated goal of reducing uncertainty and promoting fairness. Airbnb allows for a sophisticated menu of options that the seller may offer, depending on market characteristics. In markets with low demand, the typical seller will keep the option to cancel without penalty up to 24 hours before the arranged date, while in markets with extremely

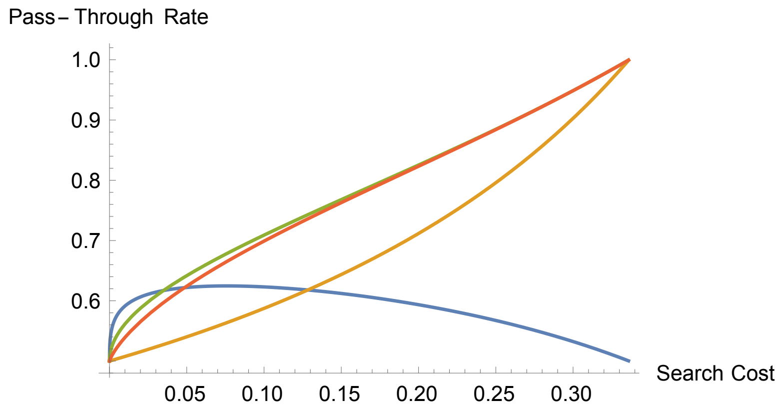

Figure 5: Pass-Through on the Average Price (Blue), Weighted (Green) and Unweighted (Red) Average Pass-Through, Pass-Through on Inital Price (Orange)

high demand, Airbnb offers options that allow for partial refund only for very early cancellations.

From the perspective of the buyer, the right of cancellation is valuable since it decouples purchase from search, allowing a buyer to gather more information before committing to a final decision. From the perspective of the seller, the cancellation policy is a blessing and a curse. On the one hand, since prices decrease over time, locking in early buyers is profitable. On the other hand, given his capacity constraint, a seller who offers a cancellation policy is excluded from the market and faces a risk of not matching with a buyer. This risk is mitigated if the buyer has to pay a penalty to those sellers whose option she gives up.[19].

In order to compare the results with the benchmark model and ensure a tractable characterization we simply extend the model to incorporate a renegotiation clause so that if a buyer purchased a good in period t but found a better good in period $t' < t$, the buyer can make a take-it-or-leave-it offer to the seller to return her the good in exchange of a refund $f_{t'}$. This arrangement is consistent with the original model since it always gives the current owner of the good the right to post a price.

Since the seller has no private information, the buyer will naturally make an offer that gives him no additional surplus so that the continuation surplus of the seller is $\Pi_{t'}$. Therefore, upon purchasing the good, the buyer becomes the residual claimant on the joint continuation surplus of the match and the optimal offer is simply to compensate the seller from the incurred losses if he executes the offer in any future period (i.e. her *externality* on the buyer).[20] Since the continuation refund

[19]The literature on consumer search has studied the role of costly recall extensively, both because of exogenous costs and through an optimal seller policy (Armstrong and Zhou (2015), Janssen and Parakhonyak (2014))

[20]Since there is no discounting, the resulting allocation can be implemented in multiple ways.

is independent of the initial price p or the period of purchase t, the problem of the buyer is fully characterized by the pair period s and continuation utility u, while the seller is perfectly insured against the risk of cancellation.[21]

Therefore, consider a buyer who holds an offer with utility u in period $t'+1$ and meets a seller with an offer (u', p'). Taking up the option gives a expected surplus of $V_{t'}(u') - p'$ while keeping up the previous offer delivers $V_{t'}(u) - \Pi_{t'}$. Thus, taking up the option is profitable if $V_{t'}(u') - V_{t'}(u) \geq p' - \Pi_{t'}$, which, in equilibrium, equals the period-t' markup. In order to further simplify the problem, we shall assume that $F(u)$ is exponential with parameter μ^{-1}, so the markup is constant and equal to μ.[22] Hence, the condition simply translates to $V_{t'}(u') \geq V_{t'}(u) + \mu$. If the buyer does not currently hold an option, she obtains a value of $V_{t'} = w_{t'} - p_{t'}$. Let $w_{t'}(u)$ be such that $V_{t'}(w_{t'}(u)) = V_{t'}(u) - \mu$.

Given this ex-post payoffs, the returns from search of a buyer in period t, when a measure α_t of buyers have no option and a measure β_t of buyers have one option but participate in the market are

$$q\frac{\alpha_t}{\alpha_t + \beta_t}\int_{w_t(u)}^{\infty}(V_{t-1}(v) - V_{t-1}(u) - \mu)dF(v) - s.$$

if he holds a draw u and

$$q\frac{\alpha_t}{\alpha_t + \beta_t}\int_{w_t}^{\infty}(V_{t-1}(u) - V_{t-1})dF(u) - s$$

if he does not currently hold any option. A buyer stops her search if this returns are negative and continues as long as they are positive. This induces a cutoff value v_t as a function of α_t, such that if the current alternative is $u \geq v_t$, the buyer leaves the market. Since v_t is decreasing in α_t and $\alpha_t = F(v_t) - F(w_t)$, is increasing in v_t, there is a unique solution for the system. Using the second expression, we can derive the reservation value of a consumer who joins the market for the first time in

For instance, the buyer may pay the markup μ upon visiting the seller and then a rental price of $\Pi_{t'} - \Pi_{t'-1}$ for each period in which he holds the good. Alternatively, the buyer may pay upfront the whole price and then enter the market as a seller to recover $\Pi_{t'-1}$.

[21]It should be noticed that this contract is also locally optimal if the seller had the right to sell the good. This follows because of the special nature of the screening problem at hand. High valuation buyers are much less likely to execute the policy, and, therefore, price reductions that attract more buyers lead to higher costs and become comparatively less profitable. On the other hand, a price increase that is compensated with a more lenient policy is comparatively unattractive because the infra-marginal buyers are unlikely to execute it.

[22]For general distributions, the markup would be a function of the distribution of values of buyers in the market.

period T^*.

$$q \int_{w^*}^{\infty} (V_{T^*}(u) - V_{T^*-1}) dF(u) = s$$

since $\beta_{T^*} = 0$. Further notice that if a buyer purchases the right to buy a good in period T^* he will never search in period $T^* - 1$ because the gains from search jump down by μ while the difference between periods is negligible (if needed take $(s, q) \to 0$). In other words, cancellation policies are only useful if the buyer expects to search further but his incentives to do so are sharply reduced by the very fact that he holds a much improved outside option. This implies that $V_{T^*-1}(w_t) = w_t$ and the previous formula collapses to the condition of the benchmark model. Namely,

$$q \int_{w^*}^{\infty} (u - v^*) dF(v) = s.$$

Since the cancellation policy are never taken up in the initial periods, gains from trade are unaffected. This is in sharp contrast with models of demand uncertainty, whereby the probability that a contract is cancelled is decreasing in the date at which the contract is signed. In addition, cancellation policies under demand uncertainty should be better understood as an add-on clause in the original contract that allows sellers to price discriminate among buyers (which are common in flight tickets), while cancellation policies in the context of search and horizontally differentiated products are better understood as a compensation for the losses that a seller incurs should he be excluded from the market (such as the one that Airbnb uses).

This insight can also be used to derive a necessary condition for cancellations to be taken up in any equilibrium, since it must be that a buyer is willing to search following the purchase of any good with a utility draw marginally below $\{v_t\}$. Since $v_1 = H(v_1)$, $v_2 > v_1$, this condition becomes approximately,

$$q \int_{2v_1}^{\infty} (u - 2v_1) dF(v) \geq s, \tag{14}$$

which in the case of the exponential distribution becomes $s \leq qe^{-2}\sigma$. Hence, for $s \in (qe^{-2}\sigma, qe^{-1}\sigma)$, the equilibrium involves search but no cancellations.[23] If cancellations do occur, they lead to a speeding up of trade and may have a significant impact on the distribution of the gains from trade.

Figure 6 depicts the equilibrium in the case of exponentially distributed valuations with mean 1, $q = 1$ and $s = 0.006$. It can be shown that cancellations are only

[23] For the uniform distribution, the markup equals $\sigma + \mu - w_t$ so that in every period $w_t + H(w_t) = 1$, and, therefore, cancellations never occur. For the Pareto distribution with $1 - v^{-2}$, the equilibrium involves cancellations only if $2s < q$, while there is an equilibrium without cancellations if $s < q$.

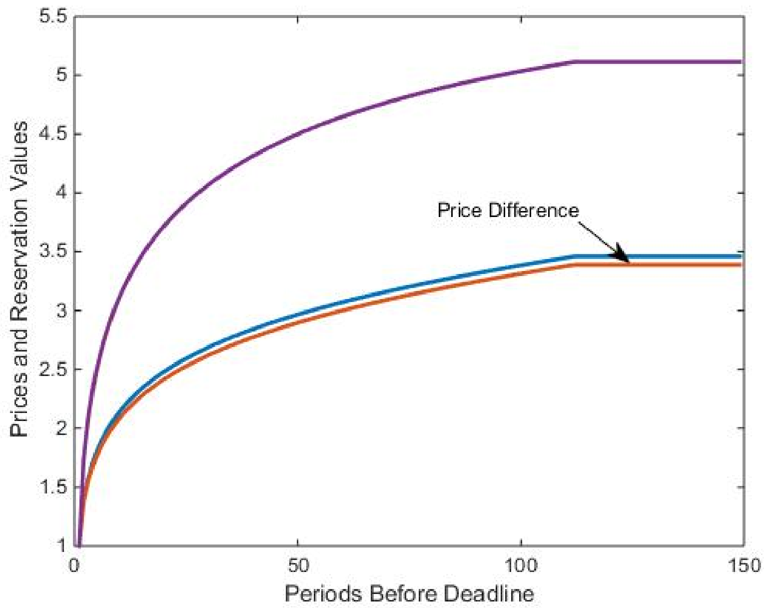

Figure 6: Prices and Reservation Values with and without Cancellation policies. The price with cancellation (blue) is always above the one without possibilities of cancellation (red).

relevant in period $t = 2$, where buyers accept matches with a utility draw $u > 1.66$ while they continue searching if $u < 1.73$, so that in period 1, approximately 1.5% of buyers hold an offer from a previous period. As it turns out, this leads to a sharp reduction in the gains from trade that buyers appropriate in the last period, which leads to faster trading and higher profits for sellers in earlier periods. As a result, the distribution of the gains from trade changes substantially, so that sellers' share of the pie increases by 3%.

7 Extensions

We now discuss several ways in which the basic model can be extended in order to incorporate some realistic features of many markets. The first extension corresponds to the case in which a single buyer may visit various sellers in different periods, introducing within period competition. We then highlight the role that strategic buyers play in our analysis by studying the cases in which buyers are the long side of the market, and, therefore, obtain limited rents and also the case in which buyers behave myopically.

7.1 Contemporaneous Seller Competition

A crucial feature of the model is that buyers are matched to at most one buyer within a given period. Thus, the current seller only competes against future sellers. This assumption allows us to focus on dynamic pricing while keeping the analysis simple but a fixed-sample search within each period can be easily incorporated. To this end, assume that, upon paying a search cost s, each buyer gets matched with probability $q = 1/2$ with two different sellers.[24] Given the sequence of prices, she buys from the seller who offers the highest surplus today provided that $u_{i,j} - p_j > w_{t-1} - p_{t-1}$. Thus, the only difference concerns the case in which at both sellers provided higher surplus than w_{t-1}. In this case, the gains from search are

$$q \int_{w}^{\infty} (v - w) dF^2(x) - s,$$

where $F^2(x)$ is the distribution of the highest of two draws. Let w^* be the associated cutoff value. Sellers choose prices to maximize

$$\Pi_{t-1} + (p_j - \Pi_{t-1})) \int_{w_t - p_t + p_j}^{\infty} F(v + p_t - p_j) f(v) dv.$$

which yields $p_t = \Pi_{t-1} + \mu(w_t)(w_{t-1})$, for some $\mu(w_t)$. It should be noted that $\mu(w_t) \leq H(w_t)$ if and only if the distribution has a non-increasing hazard rate. For our benchmark example with exponentially distributed valuations and $\sigma = 1$, this equation simplifies so that $p_t = \Pi_{t-1} + 1$. In this case, profits can be written as

$$\Pi_t = \frac{1}{2} \sum_{\tau=1}^{t-1} e^{-2w_\tau} (2e^{w_\tau} - 1).$$

On the other hand, continuation values are given by

$$V_t = q \frac{1}{2} \sum_{\tau=1}^{t-1} e^{-2w_\tau} (4e^{w_\tau} - 1) - ts.$$

So that the reservation values simplify to

$$w_t = 1 + \sum_{\tau=1}^{t-1} e^{-2w_\tau} (2e^{w_\tau} - (3/4)) - (t-1)s.$$

[24]The general case of k sellers is analogous but the expressions become considerably more complex.

Buyers search as long as $w_t < w^*$ with

$$w^* = \ln\left(\frac{q(1+\sqrt{1-s})}{2s}\right).$$

Although we can no longer solve for p_{T^*} using $w_{T^*} = w^*$ in closed form, one can show that all these expressions are rather close to the ones derived in the benchmark model.

The numerical results are presented in Figures 7 and 8. In Figure 7, the red and blue lines represent the reservation values and prices as a function of time in the model with contemporaneous seller competition, exponentially distributed valuations ($\lambda = 1$) and search cost $s = 0.0001$. The purple and yellow lines depict the evolution of reservation values and prices in the benchmark model, with $q = 1/2$ and a search cost such that both models have the same w^*. As we can see the market evolution over time is quite different but the distribution of the surplus in the initial period (which is what determines the gains from trade) is remarkably similar. Figure 8 depicts the effect of competition on the distribution of the surplus for different values of the search cost. As we saw earlier, higher search costs help sellers in both setups but the gains are smaller in the case of contemporaneous competition.

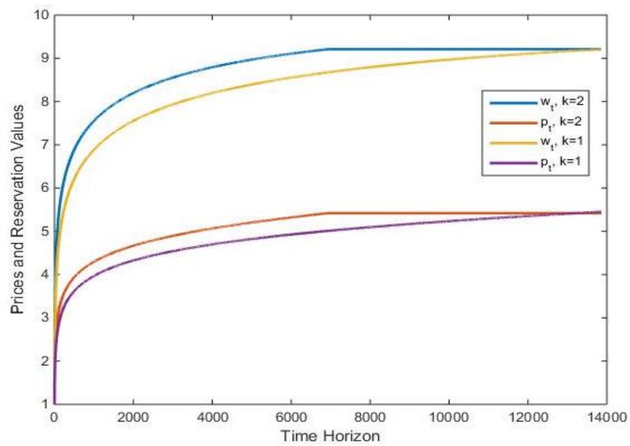

Figure 7: Prices and Reservation Values. The red and blue lines correspond to the case of within period competition ($k = 2$).

7.2 Sellers Market

Dynamic pricing is particularly relevant in markets with scarcity, like tickets for popular events (Sweeting, 2012). In this case, it is buyers rather than sellers who become pessimistic over time, which corresponds to $m > 1$.

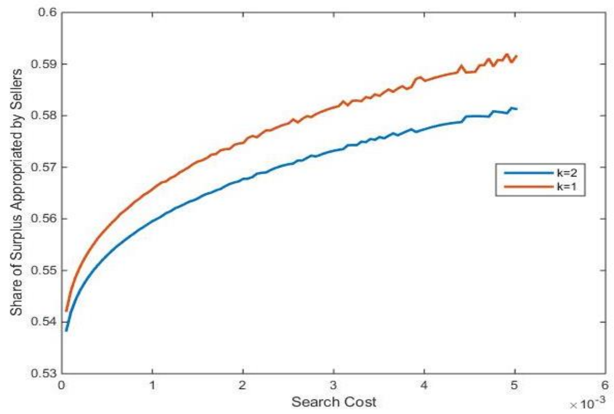

Figure 8: Share of the rents appropriated by sellers as a function of search costs with (blue line) and without (red line) within period competition.

The problem for buyers is now substantially more complex, because the gains from search crucially depend on how many other buyers are searching. In any given period, a consumer searches only if

$$q_t \int_{w_t}^{\infty} (v - w_t) dF(v) \geq s \tag{15}$$

where $q_t = q/(m_t \eta_t)$ and η_t is the proportion of consumers searching. Therefore, if

$$\frac{q}{m_t} \int_{w_t}^{\infty} (v - w_t) dF(v) > s$$

then $\eta_t = 1$. Otherwise, $\eta_t < 1$ adjusts so that (15) holds with equality. This leads to two novel effects. First, it must still be true that $q_{T^*} = 1$, for otherwise a buyer could deviate and visit a retailer in period $T^* + 1$ and observe a price arbitrarily close to p_{T^*} and match with probability one and then wait in the following period, obtaining higher expected utility.[25] Hence, trade always starts slowly in the sense that for a number of periods some buyers abstain from participating and, therefore, derive no rents from the market.

Second, it may also occur that in some later period t', $\eta_{t'} < 1$ because m_t dropped faster than the expected gains from search.[26] Furthermore, if m is large

[25]The fact that the price does not change significantly follows from the fact that (a) no firm expects to sell in period $T^* + 1$ and so $\Pi_{T^*+1} = \Pi_{T^*}$ and (b) buyers do not expect to make substantial rents in period T^* so that their reservation value in $T^* + 1$ is similar.

[26]For instance, for the case of exponentially distributed valuations and $q = 1$, if $1 - F(w_t) > m_t - \sqrt{m_t(m_t - 1)}$, the expected gains from search are decreasing in t if $\eta = 1$.

enough, in every single period buyers are indifferent between participating or not, and, therefore, derive zero rents in expectation. In this case, we have that $w_t = p_t$. A sufficient condition for this to hold is that

$$\frac{q}{m} \int_{p^m}^{\infty} (v - p^m) dF(v) \leq s.$$

Notice that $w_t \geq p^m$ and $m_t \geq m$ so that this establishes an upper bound on the gains from search in any particular period. For instance, in the case of exponential valuations with $\sigma = 1$ and $q = 1$, this holds if $s \in [\frac{1}{me}, \frac{1}{e}]$.

We now turn our attention to welfare. The average gains from trade per unit available in the market can be written as

$$W = w^* - H(w^*) + (m-1)(w^* - p_{T^*}).$$

Thus, and similarly to the case of scarce buyers, per-unit welfare is larger than $w^* - H(w^*)$. In this case, restricting entry of some buyers, leads to an increase in V_{T^*}, so that buyers delay entry, reducing total welfare.[27] Furthermore, we have that total surplus must still be bounded above by the second-best gains from trade, i.e. $W \leq w^*$. Therefore, $V_{T^*} < \frac{1}{m-1} H(w^*)$. Hence, we have the following result.

Proposition 6. *In a sellers' market ($m > 1$), as search costs vanish*

1. *If the distribution is log-concave, for every $m > 1$, buyers' surplus represents a negligible share of total gains from trade.*

2. *There exists some $1 < m^*(q, s) < \infty$ such that if $m > m^*$, $V_{T^*} = 0$, $p_{T^*} = w^*$ and buyer participation (η_t) increases as t approaches the deadline.*

The proof is similar to that of Proposition 4 and thus, omitted. From the perspective of an outsider, a market in which $m >> 1$ looks rather similar to one in which buyers are myopic. In that case, η_t can be thought of as the arrival rate. As we shall see in the next Subsection, the property that the arrival rate is increasing over time also carries over to that environment.

7.3 Myopic Buyers

In many settings, it may be convenient to assume that buyers are myopic. This assumption is common in the theoretical literature on dynamic pricing in monopoly (Talluri and Van Ryzin, 2006) and in many empirical applications (Sweeting, 2015).

[27] A simple way to see this here is that if $m-1$ buyers are restricted to enter, active buyers know they will match in every future period and will bid up their reservation values, leading to less trade.

A buyer participating in period t will buy if and only if their valuation exceeds the price so that their demand is given by $1 - F(p)$. The continuation value for the seller remains unchanged. Namely,

$$\Pi_t^M = \sum_{\tau=1}^{t-1} (1 - F(p_\tau^M))H(p_\tau^M),$$

so that $p_t^M - p_{t-1}^M = H(p_t^M) - F(p_{t-1}^M)H(p_{t-1}^M)$ describes the evolution of prices. Buyers are indifferent whether to participate or not in period t if $p_t = w^*$ and so let T^M denote the associated number of trading periods. In equilibrium, gains from trade are given (approximately) by

$$W^M = w^* - H(w^*) + S(w^*) \tag{16}$$

where $S(w^*)$ are the average gains of a consumer in this market. Notice that $S(w^*) > 0$ since in every period $t < T^M$ search yields positive expected value. Because consumers do not have the option to wait and the equilibrium with long-run buyers involved inefficient delay, the myopic equilibrium is more efficient.[28]

Proposition 7. *Suppose that the Hazard Rate is non-decreasing. Then, in an equilibrium with myopic buyers, trade starts at an earlier date $T^M \geq T^*$ and gains from trade are $W^M > W^*$.*

A caveat of this analysis is that it assumes that the number of trading periods will be determined by the indifference condition of buyers. Since buyers are myopic, this is no longer a natural restriction but it is the one that (i) maximizes welfare and (ii) yields a natural benchmark for the dynamic model.

A final remark regarding identification. Suppose that we have data on quantities and prices in each date but we do not have data on individual buyers' behavior or their arrival rates in the market, can we identify whether buyers are myopic in equilibrium? If we knew the distribution of valuations, the answer would naturally follow from observed prices, once we estimate expected profits in each period. Indeed, if buyers are myopic we have that $p_t = H(p_t) + \Pi_{t-1}$ which can be directly verified from the data. If, on the other hand, we do not now the distribution of valuations, these models cannot be distinguished from each other. In particular, suppose that buyers are fully strategic, let $F(v)$ be the true distribution of valuations and let $\{w_t\}$ be

[28]Notice that if sellers are myopic, the equilibrium corresponds to the case where $m = 0$. In this case, the equilibrium welfare is exactly the same as the case of looking-forward sellers. If sellers are forward-looking but cannot adjust prices over time, however, equilibrium welfare is given by $w^* - p^*(1 - Q(p^*))$, where Q is the total quantity sold over time.

the sequence of reservation values. Then, prices and quantities are fully consistent with a model in which buyers are myopic and the distribution of valuations is $F'(v)$ and arrival rate λ_t if $(1 - F'(p_t))\lambda_t = 1 - F(w_t)$ and $H(p_t) = H(w_t)$.

Figure 8 depicts the implied λ_t for the case of exponentially distributed valuations so that the markup is constant. Naturally, $\lambda_t \leq 1$ since forward-looking buyers are pickier. Second, as time approaches the deadline, $w_t \to p_t$ and so $\lambda_t \to 1$. Finally, λ_t is decreasing around T^* if $\xi \leq 0$

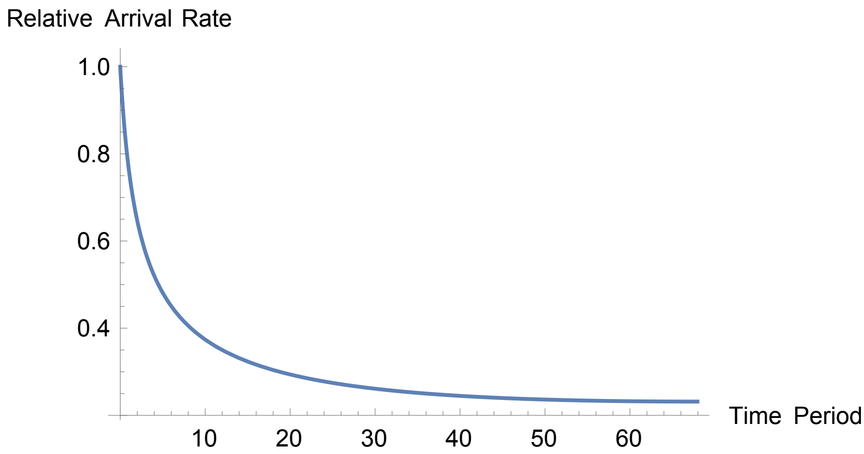

Figure 9: Implied arrival rates for myopic buyers for the case of $\sigma = 1$, $\xi = 0$, $q = 1$ and $s = 0.01$.

8 Conclusions

In this paper we have provided a tractable, yet rich, model of dynamic pricing in markets with search frictions. The equilibrium of the model is consistent with many empirical patterns of real-world markets. First, although trade occurs over a long-horizon, a large share of transactions occur in the last few periods. Second, prices decrease over time at increasing rates and approach the static optimum only in the very last period. Third, the division of rents crucially depends on the buyer-to-seller ratio so that if sellers are relatively abundant, and become even more abundant over time, they appropriate a negligible source of the surplus.

The model delivers a number of welfare implications. First, the degree of inefficiency in the market depends mostly on the markup in the initial period (which does not translate directly to profits). For some distributions, this markup may be substantial even as search frictions vanish. Second, if buyers and sellers are similar in number, reducing search frictions increases rents for buyers but also profits. This is because sellers benefit from the activation of an extensive margin of search that induces more trade and increases the value of holding a unit of the good. Importantly, a model that fails to account for strategic buyers and for an endogenous time span of the market may actually obtain the opposite answers to these questions. Finally, and from a market-design perspective, we find that taxes, subsidies and cancellation policies have a very limited impact on gains from trade but a very substantial impact on its distribution.

The model can be extended in many dimensions but two obvious features are missing. First, we have not included endogenous participation of buyers and sellers, although our results regarding the distribution of the gains from trade as a function of the relative sizes of both sides of the market suggest that this channel may induce equilibria with an even measure of each. Second, the current version of the model assumes that all sellers are ex-ante homogeneous but some heterogeneity in either capacity or quality may be relevant in some applications. These assumptions allow us to give analytical insights on a very rich model but are not crucial once the model is solved numerically for the purpose of its empirical implementation.

References

Armstrong, Mark and Jidong Zhou, "Search deterrence," *The Review of Economic Studies*, 2015, p. rdv036.

Ausubel, Lawrence M, Peter Cramton, Marek Pycia, Marzena Rostek, and Marek Weretka, "Demand reduction and inefficiency in multi-unit auctions," *The Review of Economic Studies*, 2014, p. rdu023.

Board, Simon and Andy Skrzypacz, "Revenue Management with Forward-Looking Buyers," *Journal of Political Economy*, 2016, *124* (4), 1046–1087.

Bulow, Jeremy I and Paul Pfleiderer, "A note on the effect of cost changes on prices," *Journal of political Economy*, 1983, *91* (1), 182–185.

Coey, Dominic, Bradley Larsen, and B Platt, "A Theory of Discounts and Deadlines in Retail Search," Technical Report, Working paper, Stanford University 2016.

de Albéniz, Victor Martínez and Kalyan Talluri, "Dynamic price competition with fixed capacities," *Management Science*, 2011, *57* (6), 1078–1093.

Deneckere, Raymond and James Peck, "Dynamic competition with random demand and costless search: A theory of price posting," *Econometrica*, 2012, *80* (3), 1185–1247.

Dudey, Marc, "Dynamic edgeworth-bertrand competition," *The Quarterly Journal of Economics*, 1992, *107* (4), 1461–1477.

Egorov, Aleksandr Ivanovich, *Riccati equations* number 5, Pensoft Publishers, 2007.

Einav, Liran, Chiara Farronato, Jonathan Levin, and Neel Sundaresan, "Auctions versus Posted Prices in Online Markets," Technical Report, Working paper, Stanford University 2016.

Garcia, Daniel and Sandro Shelegia, "Consumer Search with Observational Learning," Technical Report, University of Vienna, Department of Economics 2015.

_ , Gaston Llanes, and Joaquin Poblete, "Platforms and Consumer Search," 2015.

Hörner, Johannes and Larry Samuelson, "Managing strategic buyers," *Journal of Political Economy*, 2011, *119* (3), 379–425.

Janssen, Maarten CW, Alexei Parakhonyak, and Anastasia Parakhonyak, "Non-reservation price equilibria and consumer search," *Higher School of Economics Research Paper No*, 2014, *51.*

_ and _ , "Consumer search markets with costly revisits," *Economic Theory*, 2014, *55* (2), 481–514.

Lauermann, Stephan, "Dynamic matching and bargaining games: A general approach," *The American Economic Review*, 2013, *103* (2), 663–689.

Meisner, Vincent, "Competing for Strategic Buyers," 2016.

Moraga-González, José L and Makoto Watanabe, "Competitive sequential search equilibrium," 2016.

Moraga-González, José Luis, Zsolt Sándor, and Matthijs R Wildenbeest, "Prices and Heterogeneous Search Costs," 2015.

Podczeck, Konrad, "On existence of rich Fubini extensions," *Economic Theory*, 2010, *45* (1), 1–22.

Sweeting, Andrew, "Dynamic pricing behavior in perishable goods markets: Evidence from secondary markets for major league baseball tickets," *Journal of Political Economy*, 2012, *120* (6), 1133–1172.

— , "A Model of Non-Stationary Dynamic Price Competition with an Application to Platform Design," *Available at SSRN*, 2015.

Talluri, Kalyan T and Garrett J Van Ryzin, *The theory and practice of revenue management*, Vol. 68, Springer Science & Business Media, 2006.

Weyl, E Glen and Michal Fabinger, "Pass-through as an economic tool: Principles of incidence under imperfect competition," *Journal of Political Economy*, 2013, *121* (3), 528–583.

Xu, Ningxiong, "Multi-period dynamic supply contracts with cancellation," *Computers & Operations Research*, 2005, *32* (12), 3129–3142.

A Omitted Proofs

Proof of Lemma 1. We first establish that if $v - H(v)$ is increasing, then p_t follows (5). To see this notice that the Second Order Condition is

$$-f'(w_t - p_t + p_j)(p_j - \Pi_{t-1}) - 2f(w_t - p_t + p_j) \leq 0$$

since in a symmetric equilibrium $p_j = p_t$ and $p_j = \Pi_{t-1} + H(w_t)$ we have,

$$f'(w_t)(1 - F(w_t)) + 2f^2(w_t) \leq 0,$$

which holds iff $H'(w) \leq 1$.

We now show that there cannot be an equilibrium in which buyers are indifferent between entering or not in multiple periods and a positive mass of them participate. In such a case, $V_t = 0$ so that $w_t = w^*$. Since $m > 0$, $\Pi_t - \Pi_{t-1} > 0$ if there is a positive measure of buyers who participate, so that $p_t > p_{t-1}$. But this means that if $w_t = w^*$, then $w_{t-1} < w^*$ and $V_{t-1} > 0$. Thus, all equilibria are outcome equivalent to one in which trade starts at some T^*.

$$w_t = w_{t-1} + H(w_t) - H(w_{t-1}) + qm_{t-1}H(w_{t-1})(1 - F(w_{t-1})) + \int_{w_{t-1}}^{\infty} (v - w_{t-1})dF(v) - s$$

Since $H'(w) < 1$, $H(w_t) - H(w_{t-1}) < (w_t - w_{t-1})$, then $w_t > w_{t-1}$. $\square$

Proof of Proposition 2. The main argument is sketched in the main text. Here we simply establish that $w_t^l > w_t^{l+1}$ if $\tau(l) > l$ and $\tau(l+1) > l$. To see this notice that w_1^l is independent of l and

$$\begin{aligned} w_2^l = w_1^l + \psi(w_1^l) + \Delta(w_1^l, l) &= w_1^{l+1} + \psi(w_1^{l+1}) + \Delta(w_1^{l+1}, 1, l) \\ &> w_1^{l+1} + \psi(w_1^{l+1}) + \Delta(w_1^{l+1}, 1, l+1), \end{aligned}$$

where $\psi(w)$ represents the returns from search with a reservation value of w and $\Delta(w, t, l)$ represents the expected change in price assuming full participation from period l and a current reservation value of w. Thus, assume that $w_t^l > w_t^{l+1}$. It follows that,

$$\begin{aligned} w_{t+1}^l = w_t^l + \psi(w_t^l) + \Delta(w_t^l, l) &> w_t^{l+1} + \psi(w_t^{l+1}) + \Delta(w_t^{l+1}, t, l) \\ &> w_t^{l+1} + \psi(w_t^{l+1}) + \Delta(w_t^{l+1}, t, l+1). \end{aligned}$$

Therefore, $\tau(l) < \tau(l+1)$, so that if $\tau(l) \geq l$, $\tau(l+1) \geq l$, $\tau(l) \geq \tau(l+1)$. Since $\tau(\infty) = T_1 < \infty$, there must exist some lowest l^* such that $\tau(l^*) \leq l^*$, while

$\tau(l^*-1) > l^*-1$. If $\tau(l^*) = l^*$, this constitutes a fixed-point and the associated sequence corresponds to a pure-strategy equilibrium. Otherwise, we construct an equilibrium by assuming that buyers randomize in period l^* whether they participate or not. If no buyer participates, then $m_{l^*-1} = m_{l^*-1}^{l^*-1}$ and the continuation sequence of reservation values must be that associated with an initial period of l^*-1, so that $w_{l^*-1} < w^*$. If every buyer participates, we know that $w_{l^*-1} > w^*$ since $\tau(l^*) > l^*$. Thus, there exists some $\eta \in [0,1)$ such that if η proportion of buyers participate in period l^*, m_{l^*-1} is such that $w_{l^*} = w^*$. Because $w_{l^*} = w^*$, buyers are willing to randomize.

In order to complete the proof, notice that in an un-dominated equilibrium $w_{T^*} + \psi(w_{T^*}) + \Delta(w_{T^*}, T^*, T^*) > w^*$, for otherwise a deviating buyer would find it profitable to search at T^*+1. Hence $w_{T^*} \to w^*$ as $q \to 0$. Further, if $m = 1$ or $m = 0$, $\{w_t^l\}$ are independent of l so that an undominated equilibrium is necessarily unique. Similarly, in the case of exponential valuations, $H(v)$ is constant so that there is at most one fixed point of $\tau(l)$.

$\square$

Proofs of Proposition 3. Most of the argument follows from the main text for $(q,s) \to 0$. We now derive exact bounds for welfare as a function of s (and, crucially, independently of q), so that as we take $(s,q) \to 0$, the results hold exactly. Recall that,

$$\begin{aligned} W &= mV_{T^*} + \Pi_{T^*} \\ &= m\big(w_{T^*} - p_{T^*}\big) + \big(p_{T^*} - H(w_{T^*})\big) + qm\big(1 - F(w_{T^*})\big)H(w_{T^*}). \end{aligned}$$

Since in any un-dominated equilibrium $w_{T^*} + \psi(w_{T^*}) + \Delta(T^*) \geq w^* \geq w_{T^*}$, then

$$\begin{aligned} W &\geq m\big(w^* - \psi(w_{T^*}) - \Delta(T^*) - p_{T^*}\big) + \big(p_{T^*} - H(w_{T^*})\big) + \Delta(T^*) + H(w_{T^*}) - H(w_{T^*+1}) \\ &\geq m\big(w^* - \psi(w_{T^*}) - \Delta(T^*) - p_{T^*}\big) + \big(p_{T^*} + \Delta_{T^*} - H(w_{T^*+1})\big). \end{aligned}$$

If $m = 1$, we have,

$$W \geq w^* - \psi(w_{T^*}) - H(w_{T^*+1}) \geq w^* - s - H(w_{T^*+1}).$$

Therefore, for $H(w)$ decreasing, $w^* - H(w^*) \leq W \geq w^* - s - H(w^*)$. It then follows that as $s \to 0$, $W \to \lim_{s\to 0} w^* - H(w^*)$ as required. For the case in which $F(v)$ is

Pareto, we have,

$$W = w_{T^*} - H(w_{T^*}) + qm\big(1 - F(w_{T^*})\big)H(w_{T^*}).$$

Since $w - H(w)$ is increasing, $W \leq w_{T^*+1} - H(w_{T^*+1}) + qm(1 - F(w_{T^*}))H(w_{T^*})$. For the case of the Pareto distribution, we have that $qm\big(1 - F(w_{T^*})\big)H(w_{T^*}) < ms$ (see Section 4). Therefore, $w_{T^*} - H(w_{T^*}) + s \leq W^* \leq w_{T^*+1} - H(w_{T^*+1}) + s$. Since $H(w) = \alpha w$, we have that $\lim_{s\to 0} W \to \frac{\alpha}{1-\alpha}w^*$. □

Proofs of Proposition 4 and Corollary 1. The analysis is similar to that of Proposition 3. Suppose that $F(v)$ is log-concave. Then, $H(v) \geq w_1$ by definition of the monopoly price. We have,

$$\begin{aligned} W &\geq m\big(w^* - \psi(w_{T^*}) - \Delta(T^*) - p_{T^*}\big) + \big(p_{T^*} - H(w_{T^*})\big) + \Delta(T^*) + H(w_{T^*}) - H(w_{T^*+1}) \\ &= m(w^* - \psi(w_{T^*}) - m(p_{T^*} + \Delta(T^*)) + \big(p_{T^*} + \Delta(T^*) - H(w_{T^*+1})\big) \\ &\geq m(w^* - s - H(w_{T^*+1})) + (1-m)(p_{T^*} + \Delta_{T^*} - H(w_{T^*+1})). \end{aligned}$$

But, $W \leq mw^*$ by definition. Hence,

$$(1-m)\Bigg(\Pi_{T^*+1} - (H(w_{T^*}) - H(w_{T^*+1}))\Bigg) - m(s + H(w_{T^*+1})) \leq 0$$

. Thus,

$$\Pi_{T^*+1} \leq \frac{m}{1-m}(s + H(w_{T^*+1})) + H(w_{T^*} - H(w_{T^*+1}) \leq \frac{m}{1-m}H(w^*) + \gamma s.$$

for some $\gamma \leq \frac{w_1}{1-m}$. Using implicit differentiation in (3) and since $\frac{s}{1-F(w^*)} = MRL(w^*)$, where $MRL(w^*)$ is the mean-residual life of w^* which is decreasing for $F(w)$ log-concave, we have that

$$H'(w^*)\frac{dw}{ds} = 2MRL(w^*) + \frac{f'(w^*)}{f^2(w^*)}s < 2MRL(w^*) < 2\int vdF(v)$$

. Thus,

$$H(w^*) < \lim_{s\to 0} H(w^*) + 2s\int vdF(v).$$

Summarizing, we have that for every $\epsilon > 0$, $\Pi_{T^*+1} < \epsilon$ if

$$\frac{2m}{1-m}\int vdF(v)\lim_{s\to 0}(H(w^*) + s) + \gamma s < \epsilon.$$

So that if $\lim_{s\to 0} H(w^*) = 0$, for every ϵ, $\Pi_{T^*+1} < \epsilon$ if $s < s(m, \epsilon)$. For the case in which $H(w)$ converges to a constant, we have that $\Pi_{T^*+1}/w^* \to 0$ since $w^* \to \infty$. Finally, Corollary 3 follows from Proposition 3 since $\lim_{s\to\bar{s}} \Pi_{T^*} = qmp_1(1 - F(p_1))$ is independent of s. □

Proof of Proposition 5. The first statement follows from $w_t - w_{t-1} > 0$ and $p_t - p_{t-1} > 0$ being decreasing functions of w_{t-1}. The second statement follows from the text. To prove the third statement, consider the equilibria in two markets with the same distribution function but which differ in s, so that in a given market the search cost is s_1, while in the other the search cost is $s_2 > s_1$ and let $w^i(t)$ be the reservation value of at instant t in each of the markets (similarly for $p^i(t)$, etc.) Notice that in both markets $w^1(0) = \frac{\sigma-\xi\mu}{1-\xi}$, independent on s, but $\dot{w}^1(0) > \dot{w}^2(0)$. Hence, for $t < t^*$, $w^1(t) > w^2(t)$, for some t^* sufficiently small. Both $w^1(t)$ and $w^2(t)$ are continuous functions so that if at some time $t < T^*(s_2)$, $w^1(t) < w^2(t)$, it must be that at some earlier time t^{**}, $w^1(t^{**}) = w^2(t^{**})$. But then,

$$\begin{aligned} w^1(t^{**}) &= \int_0^{t^{**}} \dot{w}^1(t)dt = \int_0^{t^{**}} \frac{1}{1-\xi}\Big((\sigma - \xi\mu + \xi w^1(t))(1 - F(w^1(t))) - s_1\Big)dt \\ &> \int_0^{t^{**}} \frac{1}{1-\xi}\Big((\sigma - \xi\mu + \xi w^2(t))(1 - F(w^2(t))) - s_1\Big)dt \\ &> \int_0^{t^{**}} \frac{1}{1-\xi}\Big((\sigma - \xi\mu + \xi w^2(t))(1 - F(w^2(t))) - s_2\Big)dt = w^2(t^{**}). \end{aligned}$$

The first line follows by definition, the second line follows from $w^1(t) > w^2(t)$ for all t and the fact that of $(\sigma-\xi\mu+\xi w)(1-F(w))$ is decreasing in w and the third follows from $s_1 < s_2$. Hence, if $w^1(t^{**}) = w^2(t^{**})$ it must be the case that $t^{**} > T^*(s_2)$. Using the definition of $p(t)$, it follows that $p^1(t) > p^2(t)$. The last two comparative statics follow directly from the definitions of those variables with the appropriate assumption on the elasticity of total expenditure on search costs. □

Proof of Proposition 8. From (16) we need to show that $w^* - H(w^*) > S(w^*) > 0$. Notice that

$$S(w^*) \approx \sum_{\tau=1}^{T^*-1} \prod_{r=\tau+1}^{T^*} F(p_r)\Big(\int_{p_\tau}^{\infty} (v - p_\tau)dF(v) - s\Big),$$

since all the elements of the sum are positive, $S(w^*) > 0$. For the second part, we first prove that at every period t, $p_t^M \le w_t$, that is, prices (and reservation values) move slower with myopic buyers. To see this notice that $p_1^M = w_1$ and assume that for every $\tau < t$, $p_\tau^M \leq w_\tau$. It follows that

$$\begin{aligned}
p_t^M &= p_{t-1}^M + H(p_t^M) - F(p_{t-1}^M)H(p_{t-1}^M) \\
&\leq w_{t-1} + H(p_t^M) - F(p_{t-1}^M)H(p_{t-1}^M) \leq w_{t-1} + H(w_t) - F(w_{t-1})H(w_{t-1}) \\
&\leq w_{t-1} + +H(w_t) - F(w_{t-1})H(w_{t-1}) + \int_{w_{t-1}}^{\infty} (v - w_{t-1})dF(v) - s = w_t
\end{aligned}$$

where all steps follow from the definition except for the second inequality that follows from Non-Decreasing hazard rate. Thus, if $w_{T^*-1} < w^*$, $p_{T^*-1}^M < w^*$ so that $T^M \geq T^*$. $\square$

Details of the Examples in Section 4.

Exponential

We begin with part the exponential distribution. The associated differential equation is

$$\dot{w} = 2\sigma q e^{-\frac{w(t)}{\sigma}} - s$$

with initial condition $w(0) = 1$. This yields a solution

$$w(t) = \sigma \ln\left(2q\sigma - e^{-\frac{ts}{\sigma}}(2q\sigma - es)\right) - \sigma \ln(s).$$

and

$$\begin{aligned}
p(t) &= \sigma + q\sigma \int_0^t e^{-\frac{w(y)}{\sigma}} dy \\
&= \frac{\sigma}{2}\left(1 - \ln(s) + \ln\left(2\sigma q(e^{st}\sigma - 1) + es\right)\right.
\end{aligned}$$

This function is increasing in s for fixed $t < T^*$, while

$$p_{T^*} = \frac{\sigma}{2}\left(1 - \ln(s) + \ln(2q\sigma - es)\right)$$

is decreasing in s. Solving for $w(T^*) = w^*$ gives (12). On the other hand, the share of surplus that accrues to buyers is

$$\frac{V_{T^*}}{w^* - H(w^*)} \approx \frac{1}{2}\frac{(w^* - H(w^*)) - \sigma \ln\left(2 - \frac{es}{q\sigma}\right)}{w^* - H(w^*)} < \frac{1}{2}.$$

Since $w^* - H(w^*)$ decreases in s and the search expenditure increases in s, the share of gains from search going to buyers decrease in s.

Uniform

To solve the differential equation we simply use the integration method in

$$\frac{dw}{dt} = \frac{3}{4\sigma}q(\sigma + \mu - w)^2 - \frac{1}{2}s.$$

we get

$$\int_{w(0)}^{w^*} \frac{dw}{\frac{3}{4\sigma}q(\sigma + \mu - w)^2 - \frac{1}{2}s} = \int_0^{w^*} dt = T^*,$$

with $w(0) = \frac{\sigma+\mu}{2}$ and $w^* = \sigma + \mu - \sqrt{\frac{2s\sigma}{q}}$. Since T is homogeneous of degree one in (q, s) we normalize $q = 1$. Factoring the polynomial and integrating by parts we obtain,

$$T^* = \sqrt{\frac{8\sigma}{3s}}\frac{1}{2}\ln\left(\frac{\sqrt{2}\left(3+\sqrt{3}\right)(\mu+\sigma) - 4\left(1+\sqrt{3}\right)\sqrt{s}\sqrt{\sigma}}{\sqrt{2}\left(3-\sqrt{3}\right)(\mu+\sigma) + 4\left(\sqrt{3}-1\right)\sqrt{s}\sqrt{\sigma}}\right). \tag{17}$$

Since total expenditure is T^*s, it is of the order of $s^{1/2}$ for s small enough, although it is hump-shaped for the whole range of s such that $T^* > 1$. The approximation is however poor for $s > 10^{-3}$ since the total number of periods is very small. Since σ and μ scale up the gains from search, the derivative of total effort with respect to σ and μ is positive. Figure 10 below represents the total expenditure for $\sigma = 1$ and $\mu = 0$ as a function of search costs in the continuous-time limit and the numerical simulation of the discrete-time model. Finally, to see that p_{T^*} is increasing in s notice that,

$$\frac{\partial p_{T^*}}{\partial s} = \frac{4\sigma}{3\sigma - 8s} + \frac{1}{6\sqrt{s}}\sqrt{2}\sqrt{\sigma}\left(3 - \sqrt{3}\ln\left(\frac{2\sqrt{2}\left(3+\sqrt{3}\right)\sqrt{s} - 3\left(1+\sqrt{3}\right)\sqrt{\sigma}}{2\sqrt{2}\left(\sqrt{3}-3\right)\sqrt{s} - 3\left(\sqrt{3}-1\right)\sqrt{\sigma}}\right)\right)$$

which is increasing in s.[29] Therefore,

$$\begin{aligned}\frac{\partial p_{T^*}}{\partial s} &< \frac{-4}{3} - \frac{1}{6\sqrt{s}}\sqrt{2\sigma}(3 - \ln(2+\sqrt{3})) \\ &< \frac{4}{3} + \frac{1}{6\sqrt{s}}\sqrt{2\sigma}(3 - 2\ln(2)) < 0.\end{aligned}$$

[29]For simplicity, we took $\mu = 0$ but μ is only a location coefficient that does not affect the sign of the derivative.

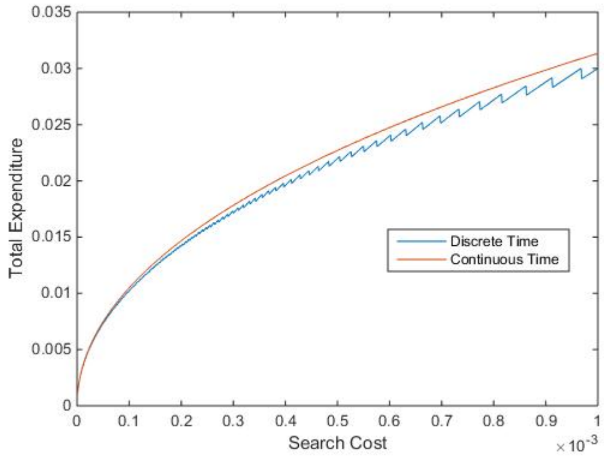

Figure 10: Expenditure in search effort as a function of search costs for $F(v) = v$.

Pareto

The differential equation,

$$\dot{w} = \frac{-3q}{w} + 2s = 0$$

admits as a solution,

$$w(t) = \frac{3\left(qW\left(-\frac{e^{-\frac{4s^2(t-c_1)}{3q}-1}}{3q}\right)+q\right)}{2s}$$

where $W(x)$ is the Lambert function and c_1 is a constant. Imposing $w(1) = 1$, yields,

$$c_1 = \frac{3q\ln\left(e^{\frac{2s}{3q}}(3q-2s)\right)}{4s^2}$$

and solving for $w(T^*) = w^* = \frac{q}{s}$, yields,

$$T^* = \frac{3q\ln\left(3-\frac{2s}{q}\right)-2(q-s)}{4s^2}$$

and so $T^*s \approx E(q,s)/s$ with $E(q,s) > 0$ decreasing in s, $E(q,0) = (1/4)q(3\ln(3) - 2) < q$ and $E(q,q) = 0$. To see that the gains from trade accruing to buyers are bounded above by 0.45 notice that

$$\frac{V_{T^*}}{V_{T^*}+\Pi_{T^*}} = \frac{2}{3q}\left(\frac{3}{2}(q-s)-\frac{3}{4}q(\ln(3q-2s)-\ln(q))\right)$$

which is decreasing in s for $s < q$. Hence, the upper-bound of the share of the gains from trade accruing to buyers can be obtained when $s \to 0$, where it approaches $1 - \frac{1}{2}\ln(3) \approx 0.45$.

B Partial Commitment

Suppose now that sellers can commit to a single price for the next k periods. For simplicity, we assume that $m = 1$ and $q = 1$. Then, their problem is simply to maximize,

$$\Pi_t = \max_{p_j} \sum_{l=0}^{k} \left(\prod_{\tau=0}^{l-1} F(w_{t+\tau} - p_t + p_j) \right) (1 - F(w_{t+l} - p_t + p_j)) p_j + \prod_{\tau=0}^{k-1} F(w_{t+\tau} - p_t + p_j) \Pi_{t+k}.$$

Intuitively, the firm faces a sequence of consumers with different demand functions (as indexed by w_t) and has a unit capacity to sell to each of them plus a terminal payoff Π_{t+k} if none of them buy. This expression simplifies to,

$$\Pi_t = \max_{p_j} \left(1 - \prod_{l=0}^{k-1} F(w_{t+l} - p_t + p_j) \right) p_j + \prod_{\tau=0}^{k-1} F(w_{t+\tau} - p_t + p_j) \Pi_{t+k}.$$

In a symmetric equilibrium, the condition for optimality is

$$\left(1 - \prod_{l=0}^{k-1} F(w_{t+l}) \right) - (p_j - \Pi_{t+k}) \left(\sum_{l=0}^{k-1} f(w_{t+l}) \prod_{\tau \neq l} F(w_{t+\tau}) \right),$$

or

$$p_j = \Pi_{t+k} + \frac{\left(1 - \prod_{l=0}^{k-1} F(w_{t+l}) \right)}{\left(\sum_{l=0}^{k-1} f(w_{t+l}) \prod_{\tau \neq l} F(w_{t+\tau}) \right)}. \tag{18}$$

The second term in Eq. (18) is the average inverse elasticity of demand in the next k periods.

The problem of buyers is unchanged except for the fact that in periods in which sellers do not change prices,

$$w_t - w_{t-1} = \int_{w_{t-1}}^{\infty} (u - w_{t-1}) dF(u) - s$$

while in periods when they change prices,

$$w_t - w_{t-1} = (\Pi_{t-k-1} - \Pi_{t-1}) + \Delta_H(t, t-k) + \int_{w_{t-1}}^{\infty} (u - w_{t-1}) dF(u) - s,$$

where $\Delta_H(t, t-k)$ measures the change in elasticity between periods. The evolution of reservation values is no longer guaranteed to be monotone under the assumption that $w - H(w)$ is increasing. Nevertheless, log-concavity of $1 - F(w)$ is still sufficient.