Pricing in a Frictional Product Market*

Leena Rudanko
Federal Reserve Bank of Philadelphia

February 15, 2017

Preliminary and Incomplete!

Abstract

Recent research argues that a key factor limiting firm growth is the gradual accumulation of product market demand. Motivated by this evidence, this paper studies firm price setting and growth in a frictional product market where firms accumulate customers over time. In this competitive search model of firm growth, firms face a trade-off in setting prices each period, between making profits on existing customers with high prices, and attracting new customers with low prices. Firms with more existing customers choose higher prices, attract fewer new customers, and grow more slowly, than those with less. I show that if a new firm has full commitment to future prices, it can attain efficient growth through its lifetime. This plan is not time consistent, however: if a firm with existing customers reoptimizes, it will choose a higher price today than planned. I study firm pricing and growth when the firm does not have commitment to future prices, focusing on Markov perfect equilibria. The baseline model considers an industry with a single monopolistically competitive firm, but I also consider the impact of a competitive fringe on the monopolist's pricing, a version with a continuum of firms within the industry that delivers an equilibrium theory of price dispersion, as well as the implications for firm responses to shocks to demand and costs.

Keywords: Product Market Search, Price Dynamics, Competitive Search, Limited Commitment.

*The views expressed here are not necessarily those of the Federal Reserve Bank of Philadelphia or the Federal Reserve System.

1 Introduction

New businesses start out small, and grow slowly over time. Recent research argues that a key factor limiting firm growth is the gradual accumulation of product market demand. Foster, Haltiwanger and Syverson (2008, 2016) argue that new businesses set lower prices than existing ones, and grow slowly despite of that, suggesting their growth is limited by frictions in the product market and the gradual accumulation of a customer base over time. What are the implications of such frictions for firm price-setting and growth? This paper studies this question a model where firms face search frictions in the product market and accumulate customers over time.

In the spirit of a framework of monopolistic competition (Dixit and Stiglitz 1977), I begin by considering a single industry with a monopolistically competitive producer. The producer can produce as many units of its good as it likes, but its sales are limited by the size of its customer base. The frictions in the product market are captured with a matching function, such that the measure of new customer relationships created depends on the queue length of customers the firm attracts. In the spirit of a competitive search equilibrium (Moen 1997), the firm's pricing influences the queue length of potential buyers the firm attracts. On the other side of the market, there is a continuum of potential buyers who face an opportunity cost of search. The measure of buyers searching is endogenous, and the buyers optimize in choosing which of the continuum of industries they choose to enter each period to search. In equilibrium, firms setting higher prices attract a shorter queue of potential buyers, and grow more slowly as a result. Matches last until a separation shock destroys the relationship.

The firms must set a single price each period for all their customers. This constraint gives a rise to a trade-off for the sellers in setting their prices each period, between making profits on their existing customers with high prices and building their customer base with low prices to attract new customers. This tension leads to an inefficiency, where firms tend to set their prices too high to profit on their existing customers, at the cost of new customer acquisition.[1]

In a dynamic infinite horizon model, I show that if a new firm has full commitment to future prices, it can attain efficient growth throughout its lifetime. This involves a constant price over time, together with a constant queue length and hence constant firm growth over time. This plan is not time consistent, however. If a firm with existing customers reoptimizes, it will choose a higher price today than planned. Having existing customers in the planning

[1] The competitive search equilibrium often leads to efficient outcomes, but here this is not the case (see, e.g., Rogerson, Shimer, and Wright (2005)).

period implies an incentive to raise prices to profit on those existing customers, distorting new customer acquisition.

What happens if the firm reoptimizes its prices each period? If it each period faces the incentive to raise prices too high to make profits on existing customers? To answer this, I study differentiable Markov perfect equilibria. I derive a generalized Euler equation for the firm, and compare it to the corresponding equations for the planner and the seller with commitment, to show the wedges that arise to distort allocations. I characterize the effects in the context of a parameterized model, showing that a larger firm chooses higher prices, attracts a shorter queue of potential buyers, and grows more slowly, than a smaller one. As a result, firm growth is characterized by rising prices, and a decreasing rate of growth over time. Both features are consistent with the evidence on firm dynamics: small firms grow faster and price lower than larger ones (see, e.g., Caves (1998) and Foster, Haltiwanger, and Syverson (2008)).[2]

In the model firm size remains permanently below efficient, as growth is stunted by the firm taking advantage of the monopoly power it has over its stock of existing customers, an effect that only becomes stronger as the firm grows. Note that in the standard monopolistic competition framework firms also set prices above efficient, limiting firm size, but for a different reason. They do so because they are monopoly producers facing a downward sloping demand curve for industry output. To isolate the model mechanism from this standard effect, in the baseline model I consider an infinitely elastic industry demand. However, I also show that the model can be extended to allow a downward sloping demand, to combine the two effects.

The lack of commitment leads to inefficiently high prices while full commitment potentially achieves efficient outcomes. This suggest the commitment horizon is important for outcomes. To consider the impact of the commitment horizon on outcomes, I show that the model can be extended to probabilistic repricing, where the firm each period faces a probability of being able to reprice. How much commitment do firms have in reality? On the empirical side, there exists a literature investigating the duration of prices observed at the micro level, with a relatively wide range of implied durations, from as little as three months to as high as ten, depending on methodologies. I demonstrate the impact of a longer duration of prices to reduce prices and increase output.

The monopoly producer sets prices too high, but what if the industry also has a compet-

[2]As discussed in Foster, Haltiwanger, and Syverson (2008), the rising path of prices could lead to the mistaken inference that the firm is becoming more productive as it grows, if productivity is measured with revenue productivity in the absence of information on firm level prices.

itive fringe of small producers? Will their presence discipline the dominant firm's pricing? To think about this, I introduce into the industry a continuum of atomistic sellers, such that a fraction of the potential buyers entering the industry are randomly be allocated to a fringe seller, rather than the dominant firm. The fringe sellers have no pricing power, but instead bargain prices with buyers individually. The dominant firm continues to control the average price in the industry, acknowledging the presence of the fringe sellers in its market. I show that introducing a competitive fringe tends to raise the industry price level further, as to the dominant firm raises its price as the costs of the higher industry price are now born partly by the fringe firms.[3]

An alternative interpretation of the model environment is also possible. Instead of following the monopolistic competition interpretation where each industry has a monopolist firm, which faces competition from other industries rather than other firms in the industry, one could also consider an interpretation where a single industry is populated by a continuum of firms which accumulate customers and compete with each other for buyers. Assuming firms probabilistically die and are replaced by new ones without customers gives rise to a new equilibrium model of price dispersion.[4]

Models with customer base dynamics were originally conceived with the dynamics of price setting in mind (Phelps and Winter 1970). Perhaps the key question is how prices respond to shocks when firms trade off making profits on existing customers versus attracting new customers? It is thus natural to consider in the case of firm level shocks here as well. Pozzi and Schivardi (2016) have recently made progress on the measurement of demand versus productivity shocks at the firm level, arguing that both appear important for understanding firm growth. They argue that some form of frictions appear to limit firm responses to both types of shocks, but find that the effect appears stronger for productivity than demand shocks. I show that the model with probabilistic pricing is consistent with this type of asymmetry, because in this case firm growth responds to productivity shocks only insofar as prices adjust to attract more customers.

The inefficiency in the model naturally raises the question of how best to correct it. An effective solution in the model would be for firms to price discriminate between new and existing customers. In this case firms would charge existing customers their willingness to pay, while charging new customers a low enough (possibly negative) price to encourage buyer

[3]See Carlton and Perloff (1994) for a textbook treatment of a competitive fringe, and Gowrisankaran and Holmes (2004) for an example of a modern application.

[4]Burdett and Coles (1997) offers an equilibrium theory of price dispersion with a related idea. Their theory applies the noisy search framework of Burdett and Judd (1983) in a continuous time environment, but does not discuss time inconsistency.

entry. If this could be implemented in an unconstrained way, the firm would attain efficient growth. Of course, it may not be possible for a firm to distinguish between the two types of customers in practice.

Literature This paper is motivated by a recent literature in macroeconomics and international trade considering the implications of product market frictions and customer base concerns for a variety of phenomena, developing ideas originally proposed by Phelps and Winter (1970). On the empirical side, Foster, Haltiwanger and Syverson (2008, 2016) provide empirical evidence from US manufacturing on gradual accumulation of demand limiting firm growth, while Fitzgerald, Haller, and Yedid-Levy (2016) provide related evidence from international trade.[5] Dinlersoz and Yorukoglu (2012), Gourio and Rudanko (2014) and Perla (2016) use them to explain firm valuation, investment dynamics and industry dynamics. Kleshchelski and Vincent (2009), Nakamura and Steinsson (2008) and Paciello, Pozzi, and Trachter (2013) use them to explain passthrough in pricing. Arkolakis (2010) and Drozd and Nosal (2012) show that these ideas help explain market penetration and pricing in international trade. The framework of this paper is most closely related to the competitive search model of Gourio and Rudanko (2014), which avoids the time inconsistency emphasized here by allowing different prices for new and existing customers. Within the search literature, Shi (2016) also studies long-term customer relationships in a framework with competitive search, focusing on the emergence of such relationships when sellers can choose to give priority to existing customers.[6]

The paper is related to a literature on switching costs surveyed by Klemperer (1995) and Farrell and Klemperer (2007). This strand of theoretical industrial organization discusses a tradeoff firms face between attracting new customers and making profit on existing ones, as well as a related time-inconsistency issue. This literature has mainly focused on models of strategic interaction in oligopolistic markets, typically formalized in the context of two-period models. Perhaps the closest paper in that literature is Beggs and Klemperer (1992), which develops a dynamic infinite horizon model and considers Markov perfect equilibria. In this paper I study the implications of these ideas for firm pricing and growth in an equilibrium model where the switching cost is endogenous.

Time inconsistency is also a feature of the Deep Habits framework developed for thinking

[5]See also Bronnenberg and Dube (2017) for a review of marketing literature focusing on persistence in demand.

[6]The literature on search models of money, such as Lagos and Wright (2005), also considers frictional product markets, but does not think about the implications of long-term relationships between buyers and sellers.

about customer base concerns by Ravn, Schmitt-Grohe and Uribe (2006, 2010).[7] As discussed by Nakamura and Steinsson (2008), that framework is typically solved assuming habits are external, because the case of internal habits involves a time-inconsistency issue complicating the analysis. Nakamura and Steinsson consider the implications of the version with internal habits for prices, focusing on repeated games with sustainable plans.

The paper is organized as follows. Section 2 considers a static model, to demonstrate key ideas in a simplified setting. Section 3 considers a dynamic infinite horizon model. Section 5 concludes.

2 Static Model

This section introduces the model in the context of a one-period framework, to highlight key ideas. The analysis is extended to a dynamic infinite horizon setting in Section 3.

Consider an industry with a monopoly producer that has access to a linear production technology for producing a unit of its good at cost c. The firm has measure n existing customers in the beginning of the period. The industry is one among a continuum of other similar industries, each with the same production cost but possibly differing n.

To make sales, each producer needs to both produce the goods, and acquire the customers to sell them to. On the other side of the market there is a large measure of atomistic potential buyers, who each get utility u from a unit of each producer's good. The product market is frictional, requiring buyers and firm representatives to meet to determine whether the firm's good is acceptable for that specific buyer. The frictions involved in this process are formalized with a matching function, such that the measure of new customer relationships for the firm can be written as $\eta(\theta)$, where θ is the measure of buyers arriving to inspect the firm's good and $\eta' > 0, \eta'' < 0$. Correspondingly, the probability of matching for a buyer is $\mu(\theta) = \eta(\theta)/\theta$. Potential buyers face a search cost w.

The queue of potential buyers arriving at a firm is endogenous, and satisfies

$$w = \mu(\theta)(u - p). \tag{1}$$

This equation equates the cost of search w with the return to search, determined by the probability of matching $\mu(\theta)$ and the gain from matching $u - p$, where p is the price charged by the firm. The equation implies that setting a higher price is associated with a shorter queue of potential buyers, other things equal.

[7] That framework has been used for example by Gilchrist, Schoenle, Sim, and Zakrajzek (2016b, 2016a) for thinking about recent macroeconomic events.

The firm chooses its price to maximizes profit:

$$\max_{p,\theta}(n+\eta(\theta))(p-c), \tag{2}$$

taking as given the buyer entry condition (1).

Solving for the price from (1) and substituting into the firm objective yields

$$\max_{\theta}(n+\eta(\theta))(u-c)-w\theta-nw/\mu(\theta), \tag{3}$$

which implies the first order condition:

$$\eta'(\theta)(u-c)=w[1-\frac{n\mu'(\theta)}{\mu(\theta)^2}]. \tag{4}$$

To understand the equation, note that as the firm raises its price, it makes more profit per customer, but attracts fewer potential buyers, and consequently fewer new customers. The firm raises its price to a point where the gains from additional matches, $\eta'(\theta)(u-c)$, equal the associated cost of search for the buyers w and profits lost on existing customers n. These losses are greater the more existing customers the firm has, and implying that the firm's price is increasing in the measure of existing customers.

To see what happens in the context of a concrete example, consider the matching function $\eta(\theta)=\theta/(1+\theta)$ and $\mu(\theta)=1/(1+\theta)$. In this case the first order condition reads $\frac{1}{(1+\theta)^2}\frac{u-c}{w}=1+n$, which implies that $\theta=\sqrt{\frac{u-c}{w}}\frac{1}{\sqrt{1+n}}-1$. The optimal queue length is a decreasing function of existing customers n. The corresponding price, $p=u-w(1+\theta)$ (from (1)), thus rises in existing customers n.

Efficient Benchmark What does the socially optimal outcome look like? The planner would maximize output, taking into account the costs of search:

$$\max_{\theta}(n+\eta(\theta))(u-c)-w\theta. \tag{5}$$

Here the first term reflects the gains from a total of $n+\eta(\theta)$ matches, while the second deducts the costs of search involved with creating $\eta(\theta)$ new matches.

The planner problem implies the first order condition

$$\eta'(\theta)(u-c)=w, \tag{6}$$

which balances the gains from additional matches with the costs of creating them.

In the context of our example, the planner optimality condition reads $\frac{1}{(1+\theta)^2}\frac{u-c}{w}=1$, and implies $\theta=\sqrt{\frac{u-c}{w}}-1$.

Comparing the firm's choice with planner's choice, note that the efficient queue length and price are independent of the existing customer base, while the firm's price is higher the more existing customers it has, meaning that the larger the firm is, the fewer new customers it attracts. The firm also generally chooses a higher price and shorter queue than the planner, unless it has no existing customers. The firm thus generally attracts fewer new customers than efficient, unless it has no existing customers.

Price Discrimination The competitive search equilibrium would seem to imply inefficient outcomes here, as the solutions to the two problems differ. Why does this happen? The answer lies in the constraint to charge all the customers the same price. If the firm can offer the new and existing customers different prices, $p - \varepsilon$ and p, respectively, we can write the firm problem as

$$\max_{p,\varepsilon,b} n(p-c) + \eta(\theta)(p-\varepsilon-c), \tag{7}$$

taking as given $w = \mu(\theta)(u - p + \varepsilon)$ and $p \leq u$. The last constraint is needed because no buyer would agree to pay more (and in this case it binds).

It is easy to see that the firm would set $p = u$, and $\varepsilon = w/\mu(\theta)$, where θ satisfies $\eta'(\theta)(u-c) = w$. The latter condition implies that the firm achieves the efficient allocation now. It sets the price for existing customers equal to their willingness to pay, but gives a discount to new customers such that customer acquisition is efficient.

In terms of redistribution: While price discrimination maximizes total surplus in the industry, note that both types of buyers are now left with a zero net gain, leaving them worse off than in the original firm problem, while the firm now gets the entire surplus and is thus strictly better off.

Intensive Margin The analysis assumes that buyers have unit demands for the good, with no explicit quantity choice. Including such a choice is straightforward, however, and does not change the basic observations. To see this, suppose the buyer's utility from the good is a concave function $u(q)$ of the quantity q, and the production cost is cq. The planner problem in this case reads:

$$\max_{\theta,q}(n+\eta(\theta))[u(q)-cq] - w\theta, \tag{8}$$

which implies the first order conditions

$$\begin{cases} u'(q) = c, \\ \eta'(\theta)[u(q)-cq] = w. \end{cases} \tag{9}$$

Note that the optimal quantity and queue of buyers are independent of the measure of customers.

As for market outcomes, the firm would solve

$$\max_{p,\theta,q}(n+\eta(\theta))(p-cq) \tag{10}$$

$$\text{s.t. } w=\mu(\theta)[u(q)-p]. \tag{11}$$

Substituting out prices, this reads

$$\max_{\theta,q}(n+\eta(\theta))(u(q)-cq)-w\theta-wn/\mu(\theta), \tag{12}$$

which implies the first order conditions

$$\begin{cases} u'(q)=c, \\ \eta'(\theta)[u(q)-cq]=w[1-\frac{n\mu'(\theta)}{\mu(\theta)^2}]. \end{cases} \tag{13}$$

The first condition specifies the same quantity the planner would choose, which is independent of the measure of customers. Given this choice, the second implies the same conclusions as above. If the firm has no existing customers, it achieves efficient customer acquisition, while with existing customers customer acquisition is distorted. Including an intensive margin of demand thus leaves the key observation unchanged. For these reasons, I will omit an explicit quantity choice in what follows.[8]

Substitution across Industries The standard monopolistic competition framework involves a downward sloping demand curve for industry output, which reflects limited substitutability of the goods across industries in combining the goods in the production of a final good. Here I have instead assumed a constant value u for the goods, which would be in line with the goods being perfectly substitutable. The downward sloping demand curve implies its own distortions when prices are set by a monopolist, separate from the mechanism discussed above. To allow this standard effect in the model, one can incorporate an increasing and concave valuation $\phi(.)$ into the analysis.

In this case the planner problem becomes

$$\max_{\theta}\phi\big(n+\eta(\theta)\big)-c(n+\eta(\theta))-w\theta, \tag{14}$$

which leads to the optimality condition

$$\eta'(\theta)[\phi'\big(n+\eta(\theta)\big)-c]=w. \tag{15}$$

[8]Note that q could also be interpreted as quality.

Note that the socially optimal queue of buyers now does depend on the measure of existing customers, due to a decreasing marginal value of the goods. The more existing customers, the shorter the queue, and the higher the price. To see the latter, note that if θ falls, $p = \phi'(\eta(\theta)) - w/\mu(\theta))$ rises due to both increased marginal valuation and lower costs of search.

As for market outcomes, the firm would solve

$$\max_{p,\theta}(n + \eta(\theta))(p - c), \tag{16}$$
$$\text{s.t. } w = \mu(\theta)[\phi'\big(n + \eta(\theta)\big) - p], \tag{17}$$

which implies the first order condition

$$\eta'(\theta)[\phi'\big(n + \eta(\theta)\big) + (n + \eta(\theta))\phi''\big(n + \eta(\theta)\big) - c] = w[1 - \frac{n\mu'(\theta)}{\mu(\theta)^2}]. \tag{18}$$

In this case we see that the firm would choose a lower queue than the planner even if it has no existing customers. This is not surprising because when a monopolist sets prices facing a downward sloping demand curve, the outcome should differ from the efficient one: output is lower and the price higher. To see this in the equation, note that the firm's valuation for the goods $\eta(\theta)\phi'(\eta(\theta))$ decreases faster than the planner's $\phi(\eta(\theta))$. Hence the firm will choose a smaller size and higher price than that associated with efficient outcomes. The price is higher, as above, due to both increased marginal utility and lower costs of search.

In the presence of existing customers, this standard effect combines with the mechanism discussed above: not only is customer acquisition depressed due to the decreasing demand curve for industry output, but the firm also develops internal market power over its existing customers in the sense discussed above. Both lead to lower customer acquisition, associated with a higher price than efficient.

In our example case, suppose $\phi(y) = y^{(1-\sigma)/\sigma}$. The efficient queue length would then satisfy $\frac{1}{(1+\theta)^2 w}[(n + \frac{\theta}{1+\theta})^{-1/\sigma} - c] = 1$, while the monopolist chooses instead $\frac{1}{(1+\theta)^2 w}[(1 - 1/\sigma)(n + \frac{\theta}{1+\theta})^{-1/\sigma} - c] = 1 + n$.

Decreasing Returns in Production Assuming a linear production cost is convenient for analysis, but we would also like to be able to consider decreasing returns in technology. It is straightforward to incorporate a convex cost function $g(.)$ with $g' > 0, g'' > 0$ as well.

In this case the planner's first order condition reads

$$\eta'(\theta)[u - g'(n + \eta(\theta))] = w, \tag{19}$$

implying the efficient queue length now also decreases in existing customers, due to the convex costs. In this case the price implementing the efficient allocation would thus rise in existing customers.

As for market outcomes, the firm's first order condition becomes

$$\eta'(\theta)[u - g'(n + \eta(\theta))] = w[1 - \frac{n\mu'(\theta)}{\mu(\theta)^2}], \tag{20}$$

implying that with no existing customers, the firm attains the efficient allocation (unlike in the case with the demand curve considered above). With existing customers, internal market power again leads to a departure of market outcomes from efficient ones, as discussed above: the price is too high and queue length too low, relative to what is efficient.

Competitive Fringe The monopolists high price limits customer acquisition. Would introducing a competitive fringe into the market help put downward pressure on prices?

Suppose the industry has, in addition to the dominant firm, a continuum of small producers each serving only a single customer. Suppose the dominant firm's market share is λ, and that the remaining fraction $1 - \lambda$ of the buyers entering the industry would randomly be allocated to the fringe sellers instead. The atomistic fringe sellers have no pricing power vis a vis the industry price level, but rather bargain prices with their customers individually. Meanwhile the dominant firm continues to control the industry price level, but is aware of the fringe firms' presence in the market when doing so.

The fringe firms' price p^f maximizes the product $(u - p^f)^\gamma (p^f - c)^{1-\gamma}$, where γ is the buyers' bargaining power and the fringe firms' production cost is assumed to be same as the dominant firm's. This implies $p^f = \gamma c + (1 - \gamma)u$, such that buyers get fraction γ of the match surplus $u - c$, and sellers the rest.

The average price in the industry, $\lambda p + (1 - \lambda)p^f$, remains in the control of the dominant firm through its price p. With this, the dominant firm's problem becomes

$$\max_{p,\theta}(n + \eta(\theta))(p - c), \tag{21}$$

$$\text{s.t. } w = \mu(\theta)[\lambda(u - p) + (1 - \lambda)\gamma(u - c)]. \tag{22}$$

Substituting out the price gives

$$\max_{\theta}(n + \eta(\theta))[\lambda(u - c) + (1 - \lambda)\gamma(u - c)] - w\theta - wn/\mu(\theta), \tag{23}$$

which yields the first order condition

$$\eta'(\theta)[\lambda(u - c) + (1 - \lambda)\gamma(u - c)] = w[1 - \frac{n\mu'(\theta)}{\mu(\theta)^2}]. \tag{24}$$

Note that this differs from the monopolist's problem only in that the match surplus $u-c$ is replaced by the average $\lambda(u-c)+(1-\lambda)\gamma(u-c)$. This means that: i) if $\gamma = 1$, the monopolists behavior is unchanged with the introduction of the fringe into the market, and ii) if $\gamma < 1$ the introduction of the fringe reduces the returns to creating new matches, meaning that the dominant firm will choose and even higher price and lower queue length in the presence of the fringe than the monopolist.

What explains this outcome, which seems surprising at the outset? One would expect that introducing competition into the market would reduce prices rather than increasing them further. To understand what happens, note that in this environment, when the dominant firm raises its price to profit from its existing customers, part of the costs of reduced customer acquisition are borne by the fringe firms instead. This means the dominant firm will tend to set a higher price than the monopolist. The fringe firms' bargaining power is important for outcomes, however. If their bargaining power is low, that simultaneously works to encourage buyer entry. In the limiting case where the fringe sellers' bargaining power is zero, the two effects exactly offset each other, and the dominant firm and monopolist act the same.

3 Dynamic Model

To think about how firm pricing evolves as the firm grows from a new firm with no customers to one with an established customer base, this section moves to a dynamic infinite horizon setting. The dynamic firm problem brings with it a time-inconsistency issue, which did not appear in the one period problem considered in the previous section.

Efficient Benchmark What would socially optimal industry growth look like in this environment? The planner problem, extended to the dynamic setting, reads:

$$\max_{\{\theta_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t[(n_t+\eta(\theta_t))(u-c)-w\theta_t] \tag{25}$$
$$\text{s.t. } n_{t+1}=(1-\delta)(n_t+\eta(\theta_t)),\ \forall t\geq 0$$

with n_0 given. The planner maximizes the present discounted value of the measure of customers times the per-period gains from matching $u-c$, net of the costs involved in creating those customer relationships.

The first order condition reads:

$$w=\eta'(\theta_t)\frac{u-c}{1-\beta(1-\delta)}, \tag{26}$$

for all $t \geq 0$. The planner chooses a queue length such that at the margin, the increase in new matches from an additional buyer times the present discounted gain from those new matches equals the search cost of the buyer. This implies that the efficient queue length of buyers is constant over time, $\theta_t = \theta$, causing the industry to grow at a constant rate over time, approaching the steady-state level of $n + \eta(\theta) = \eta(\theta)/\delta$ customers.

Monopolists price plan Consider the monopolist, then, making a price plan for the infinite future at time zero. The monopolist maximizes the present discounted value of profits, taking into account a law of motion for the customer base, as well as the constraint reflecting buyer entry into search each period:

$$\max_{\{p_t,\theta_t\}} \sum_{t=0}^{\infty} \beta^t (n_t + \eta(\theta_t))(p_t - c) \tag{27}$$

$$\text{s.t. } n_{t+1} = (1-\delta)(n_t + \eta(\theta_t)), \ \forall t \geq 0, \tag{28}$$

$$w = \mu(\theta_t) \sum_{k=t}^{\infty} \beta^{k-t}(1-\delta)^{k-t}(u - p_k), \ \forall t \geq 0, \tag{29}$$

with n_0 given. The buyer entry condition (29) is now forward-looking, with buyers entering up to a point such that the returns to search, given by the probability of matching times the present discounted value of gains from the match $u - p_t$, equal the cost of search w. The monopolist chooses a sequence of prices $\{p_t\}_{t=0}^{\infty}$, which pins down a sequence of queue lengths $\{\theta_t\}_{t=0}^{\infty}$ via (29), and a sequence of customer bases via the law of motion (28).

As in the one-period problem, one can use equation (29) to substitute out prices from this dynamic firm problem as well (see Appendix A), rewriting it as

$$\begin{aligned} &\max -\frac{wn_0}{\mu(\theta_0)} + \sum_{t=0}^{\infty} \beta^t[(n_t + \eta(\theta_t))(u-c) - w\theta_t] \\ &\text{s.t. } n_{t+1} = (1-\delta)(n_t + \eta(\theta_t)), \ \forall t \geq 0, \end{aligned} \tag{30}$$

with n_0 given. Comparing this problem with the planner problem reveals the two to be very similar, with the exception of an additional term in period zero in the monopolist's problem.

An interior solution to the monopolist's problem satisfies

$$\begin{cases} w[1 - \frac{n_0\mu'(\theta_0)}{\mu(\theta_0)^2}] = \eta'(\theta_0)\frac{u-c}{1-\beta(1-\delta)}, \ \text{for } t = 0, \\ w = \eta'(\theta_t)\frac{u-c}{1-\beta(1-\delta)}, \ \text{for } t = 1, 2, \ldots, \end{cases} \tag{31}$$

which imply for the initial period that $\theta_0 < \theta_1$ if $n_0 > 0$, as

$$1 - \frac{n_0\mu'(\theta_0)}{\mu(\theta_0)^2} = \frac{\eta'(\theta_0)}{\eta'(\theta_1)}. \tag{32}$$

The optimality conditions reveal that if the price plan is made when a new firm starts with no existing customers, then the monopolist achieves efficient industry growth throughout its lifetime. The queue length is constant and the same the planner would choose, and the monopolist implements it with a constant price over time. However, if the price plan is made when the monopolist already has some existing customers, then the monopolists behavior differs from the planner in the initial period. Customer acquisition is lower. But after the initial period the monopolist continues with the efficient queue length.

Note that the optimal price plan involves a constant price and queue length after the planning period, at the efficient levels. But if the monopolist were to reoptimize at some later date, the new plan it would make would generally involve a higher price than the efficient one in the new planning period. This means that the monopolist's price plan is not time consistent.

Repricing period-by-period What happens if the monopolist cannot commit to future prices, but expects to reoptimize each period? To think about this case, I consider Markov perfect equilibria, where the queue length in the reoptimization period is assumed to be a function of existing customers: $\Theta(n)$. This function is an equilibrium object, satisfying a recursive version of the firm problem (30), where the queue length is reoptimized each period:

$$\begin{aligned} \Theta(n) = \max_{\theta} & -\frac{wn}{\mu(\theta)} + (n+\eta(\theta))(u-c) - w\theta + \beta V(n') \\ \text{s.t. } & n' = (1-\delta)(n+\eta(\theta)), \\ & V(n) = (n+\eta(\theta))(u-c) - w\theta + \beta V(n'). \end{aligned} \tag{33}$$

Taking the first order condition and derivative of $V(n)$ with respect to n yields

$$\frac{wn\mu'(\theta)}{\mu(\theta)^2} - w + \eta'(\theta)(u - c + \beta(1-\delta)V'(n')) = 0 \tag{34}$$

and

$$V'(n) = -\frac{wn\mu'(\theta)}{\mu(\theta)^2}\Theta'(n) + u - c + \beta(1-\delta)V'(n'). \tag{35}$$

These two equations can then be combined into a Generalized Euler equation:

$$\begin{aligned} [1 - \frac{n_t\mu'(\theta_t)}{\mu(\theta_t)^2}]\frac{w}{\eta'(\theta_t)} = u - c + \beta(1-\delta)\big[(1 - \frac{n_{t+1}\mu'(\theta_{t+1})}{\mu(\theta_{t+1})^2})\frac{w}{\eta'(\theta_{t+1})} \\ - \frac{wn_{t+1}\mu'(\theta_{t+1})}{\mu(\theta_{t+1})^2}\Theta'(n_{t+1})\big]. \end{aligned} \tag{36}$$

Note that this equation is more complicated to solve than a standard Euler equation, because the derivative of the function $\Theta(n)$ appears explicitly in the equation. This means that even solving for a steady state involves more than simply evaluating variables at their steady-state levels and solving the equation together with the law of motion for customers.

[Add interpretation of equation, relate to planner and commitment versions.]

Repricing probabilistically The lack of commitment to future prices appears to lead to distortions reducing customer acquisition in the model. But what it the firm does not reprice each period, but only does it occasionally. To allow considering such longer repricing horizons, suppose the firm resets its prices probabilistically, with probability α each period. In repricing periods the seller chooses a new present value of prices P and queue length θ, such that $w = \mu(\theta)(U - P)$. In a Markov perfect equilibrium, these will be functions of existing customers n. In periods when the firm does not reprice, both remain fixed.

In periods when the firm does not reprice, θ is constant and firm value satisfies

$$\begin{aligned} V^f(n,\theta) =& (n+\eta(\theta))(u-c) - w\theta + \beta\alpha V^r(n') + \beta(1-\alpha)V^f(n',\theta) \\ \text{s.t. } n' =& (1-\delta)(n+\eta(\theta)), \end{aligned} \tag{37}$$

where $V^r(n)$ is the value of the firm when it does price, satisfying

$$\begin{aligned} V^r(n) =& (n+\eta(\Theta(n)))(u-c) - w\Theta(n) + \beta\alpha V^r(n') + \beta(1-\alpha)V^f(n',\Theta(n)) \\ \text{s.t. } n' =& (1-\delta)(n+\eta(\Theta(n))). \end{aligned} \tag{38}$$

The function $\Theta(n)$ is the equilibrium queue length the firm chooses when it reoptimizes, which solves the problem

$$\begin{aligned} &\max_{\theta} -\frac{wn}{\mu(\theta)} + (n+\eta(\theta))(u-c) - w\theta + \beta\alpha V^r(n') + \beta(1-\alpha)V^f(n',\theta) \\ &\text{s.t. } n' = (1-\delta)(n+\eta(\theta)), \end{aligned} \tag{39}$$

given n.

This extension allows considering the impact of the commitment horizon on outcomes, in Section 4.

Competitive fringe To think about the impact of a competitive fringe on the dominant firm's pricing, begin by considering the fringe firms. A fringe firm's value from being matched with a buyer satisfies $F_t = p_t^f - c + \beta(1-\delta)F_{t+1}$, while a buyer's value from being matched with a fringe firm satisfies $B_t = u - p_t^f + \beta(1-\delta)B_{t+1}$. The sum, $S_t = F_t + B_t$, satisfies $S_t = u - c + \beta(1-\delta)S_{t+1}$, which implies a constant surplus $S_t = (u-c)/(1-\beta(1-\delta))$ for all t. The fringe sellers bargain prices with buyers individually. Denoting the buyer's bargaining power by γ, this means the surplus is divided such that the buyer's get $B_t = \gamma S_t$ and the fringe firms' $F_t = (1-\gamma)S_t$, for all t. As a result, $p_t^f = \gamma u + (1-\gamma)c$ for all t.

With the fringe firms in the market, the buyer entry condition reads

$$w = \mu(\theta_t)[\lambda \sum_{k=t}^{\infty} \beta^{k-t}(1-\delta)^{k-t}(u-p_k) + (1-\lambda)\gamma S_t]. \tag{40}$$

With this, the dominant firm's problem becomes

$$\begin{aligned} &\max \sum_{t=0}^{\infty} \beta^t (n_t + \eta(\theta_t))(p_t - c) \\ &\text{s.t. } n_{t+1} = (1-\delta)(n_t + \eta(\theta_t)), \ \forall t \geq 0, \\ &\qquad w = \mu(\theta_t)[\lambda \sum_{k=t}^{\infty} \beta^{k-t}(1-\delta)^{k-t}(u-p_k) + (1-\lambda)\gamma S_t], \ \forall t \geq 0, \end{aligned} \tag{41}$$

with n_0 given.

One can again use equation (40) to substitute out prices from this dynamic firm problem (see Appendix A), rewriting it as:

$$\begin{aligned} &\max -\frac{wn_0}{\mu(\theta_0)} + \sum_{t=0}^{\infty} \beta^t[(n_t + \eta(\theta_t))(\lambda(u-c) + (1-\lambda)\gamma(u-c)) - w\theta_t] \\ &\text{s.t. } n_{t+1} = (1-\delta)(n_t + \eta(\theta_t)), \ \forall t \geq 0, \end{aligned} \tag{42}$$

with n_0 given. [Incorporate $P_t \leq u/(1-\beta(1-\gamma))$.]

Taking first order conditions for a dominant firm making a price plan for the infinite future, we have

$$\begin{cases} w[1 - \frac{n_0\mu'(\theta_0)}{\mu(\theta_0)^2}] = \eta'(\theta_0)\frac{\lambda(u-c)+(1-\lambda)\gamma(u-c)}{1-\beta(1-\delta)}, \text{ for } t = 0, \\ w = \eta'(\theta_t)\frac{\lambda(u-c)+(1-\lambda)\gamma(u-c)}{1-\beta(1-\delta)}, \text{ for } t = 1, 2, \ldots \end{cases} \tag{43}$$

which again imply that $\theta_0 < \theta_1$ if $n_0 > 0$.

Note that the dominant firm's problem is the same as that of the monopolist, but with the gains from matching $u - c$ replaced by the average $\lambda(u-c) + (1-\lambda)\gamma(u-c)$. This means that if $\gamma = 1$, the dominant firm behaves like the monopolist, and the industry attains efficient growth if the dominant firm can commit to its price plan for the infinite future when the industry is born and has size zero. In this case the dominant firm and the measure of fringe firms will both grow at the same efficient rate over time, with the dominant firm's market share remaining constant over time. On the other hand, if $\gamma < 1$, the dominant firm will not attain efficient growth even under these circumstances, but will tend to price too high due to the fringe firms in the market.

If the dominant firm instead reprices each period, its choice of queue length solves the problem

$$\Theta(n) = \max_{\theta} -\frac{wn}{\mu(\theta)} + (n+\eta(\theta))(\lambda(u-c) + (1-\lambda)\gamma(u-c)) - w\theta + \beta V(n') \tag{44}$$
$$\text{s.t. } n' = (1-\delta)(n+\eta(\theta)),$$
$$V(n) = (n+\eta(\theta))(\lambda(u-c) + (1-\lambda)\gamma(u-c)) - w\theta + \beta V(n').$$

Taking the first order condition and derivative of $V(n)$ with respect to n yields

$$\frac{wn\mu'(\theta)}{\mu(\theta)^2} - w + \eta'(\theta)(\lambda(u-c) + (1-\lambda)\gamma(u-c) + \beta(1-\delta)V'(n')) = 0 \tag{45}$$

and

$$V'(n) = -\frac{wn\mu'(\theta)}{\mu(\theta)^2}\Theta'(n) + \lambda(u-c) + (1-\lambda)\gamma(u-c) + \beta(1-\delta)V'(n'), \tag{46}$$

which can be combined to yield the Generalized Euler equation:

$$\begin{aligned}[1-\frac{n_t\mu'(\theta_t)}{\mu(\theta_t)^2}]\frac{w}{\eta'(\theta_t)} = \lambda(u-c) + (1-\lambda)\gamma(u-c) + \beta(1-\delta)\big[(1-\frac{n_{t+1}\mu'(\theta_{t+1})}{\mu(\theta_{t+1})^2})\frac{w}{\eta'(\theta_{t+1})} \\ - \frac{wn_{t+1}\mu'(\theta_{t+1})}{\mu(\theta_{t+1})^2}\Theta'(n_{t+1})\big].\end{aligned} \tag{47}$$

4 Quantitative Illustration

This section demonstrates what the dynamic model can deliver. I first consider the impact on firm growth, and how that depends on the duration of prices. I then consider the impact of the competitive fringe on market outcomes. Finally, I consider the implications for shock propagation. The focus is on the case of a linear production technology and industry demand.

4.1 Calibration and Solution Method

To calibrate the model, I first set the time period to one month, and $\beta = 1.05^{-12}$. I adopt a value for customer turnover rate and a target for the steady-state queue length from Gourio and Rudanko (2014), with $\delta = 0.0184$ (an annual turnover rate of 20%) and a target queue length of 0.25. I use the matching function in the text, $\eta(\theta) = \xi\theta/(1+\theta)$.

The steady-state firm size in the model is given by $\eta(\theta)/\delta$. Imposing the turnover rate and target queue length, this expression implies an increasing relationship between firm size and ξ – a measure of the degree of friction in the market. If we normalize firm size to one, ξ is pinned down as $\xi = \delta(1+\theta)/\theta = 0.0921$.

Finally, we can combine equations $w = \mu(\theta)\frac{u-p}{1-\beta(1-\delta)}$ and $w = \eta'(\theta)\frac{u-c}{1-\beta(1-\delta)}$ to arrive at: $\frac{1}{\eta'(\theta)} - \frac{1}{\mu(\theta)} = \frac{p-c}{c}\frac{c}{w}\frac{1}{1-\beta(1-\delta)}$. Normalizing the production cost c to one and adopting a target for the markup of 15% pins down the opportunity cost of search w, as $w = c\frac{p-c}{c}\frac{1}{1-\beta(1-\delta)}\frac{\xi}{(1+\theta)\theta} = 0.0975$. This means that $p = 1.15$ and $u = 1.75$.

The solution method extends the approach described in Krusell, Kuruscu, and Smith (2002) for solving the generalized Euler equation for this model. The solution method is described in more detail in Appendix B. The next sections describe the solution.

4.2 Prices, Output, and Productivity Mismeasurement

As the model section explains, the monopoly producer trades off making profits on existing customers versus attracting new ones in setting its price each period. This tradeoff leads it to set a higher price the more existing customers it has. The higher price also means the firm attracts fewer new customers the more existing customers it has. In the model this manifests itself in a queue length of buyers that declines in existing customers – a function $\Theta(n)$ which declines in n. Figure 2 demonstrates this relationship as it emerges from solving the dynamic model.

As the figure illustrates, the monopolist generally chooses a shorter queue length than the planner would. The monopolists queue also decreases in existing customers, whereas the planner's would not depend on customers at all (if technology and demand are linear).

The differences in queue translate into differences in firm size in the model. The firm grows by $\eta(\theta_t)$ customers in period t, and approaches a steady-state firm size of $\eta(\theta)/\delta$ over time. The monopolist firm thus grows more slowly than what the planner would choose, due to higher prices, and at a rate that decreases in the firm's size. This means that the dynamic distortion in pricing will lead to the monopolist growing more slowly and remaining smaller throughout its lifetime than what would be socially optimal. Figure 2 demonstrates the growth of the monopolist in the parameterized model. Given that the socially optimal output here is normalized to one, the figure shows that the model mechanism has a large effect

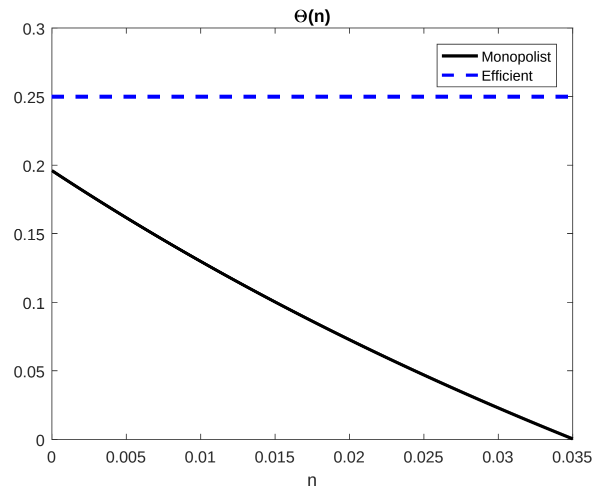

Figure 1: Function $\Theta(n)$

Notes: TBW

to depress firm size. Of course, relaxing the linearity of the model by allowing decreasing returns in production, for example, is likely to reduce the difference between the two.

The figure also illustrates productivity mismeasurement issue discussed in detail by Foster, Haltiwanger, and Syverson (2008). If productivity is measured as revenue productivity – in the absence of information on an individual firm's prices (using revenue as a proxy for output) – this firm would appear to be growing more productive over time. There is no underlying change in productivity in the model, however, just markups that grow with firm size.

As discussed in the text, the monopolist faces a commitment problem in setting its prices, while with full commitment to future prices, it could potentially even attain the socially optimal rate of growth. Thus, increasing the duration of prices should alleviate the distortions in the monopolized market. Figure 3 illustrates this relationship, by plotting solutions to the dynamic model with probabilistic pricing, where the monopolist resets its pricing with probability α each period. The implied duration of prices is thus $1/\alpha$. Based on a survey of the literature in Nakamura and Steinsson (2013), the median duration of prices is approximately 4 months for posted prices (including sales), and 7-8 months for regular prices.

As the expected, the figure indicates that a longer duration of prices is associated with lower prices and shorter queues. The quantitative effect is sizable in the figure, with output more than doubling as a result.

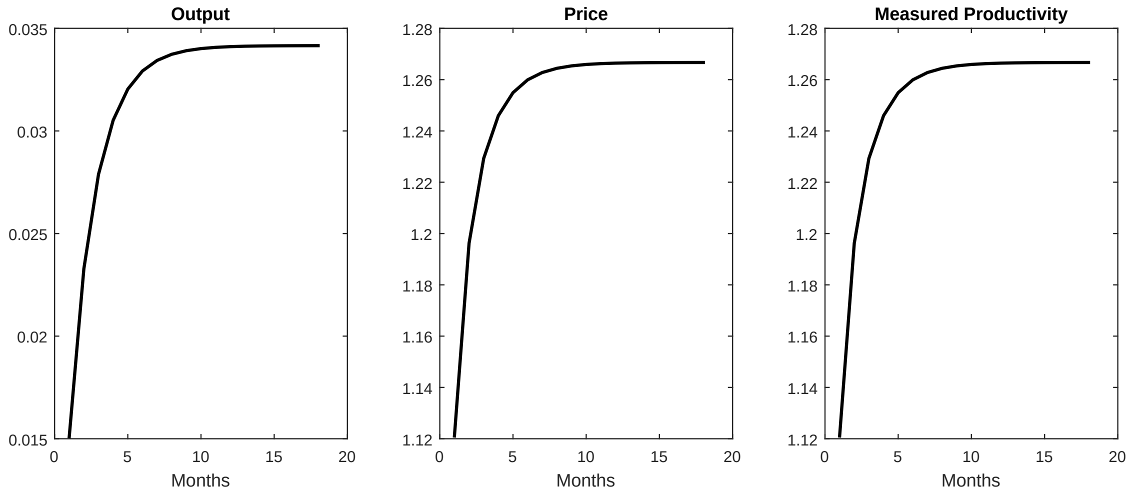

Figure 2: Firm Growth and Productivity Mismeasurement

Notes: TBW

4.3 Competitive Fringe

What happens to the monopolists pricing if we introduce a competitive fringe of atomistic producers into this frictional product market? It turns out that the answer depends in an important way on the bargaining power of the fringe sellers in their bilateral bargains with the buyers.

If the fringe sellers have low bargaining power, we arrive at an outcome where introducing the fringe is – at least weakly – welfare improving. The top panels of Figure 4 demonstrates how market outcomes vary as we vary the monopolists market share λ in this case. Note that as *lambda* falls from one, there is no change in output for a range of values of λ. This is the range where the monopolists chooses its price in an unconstrained way. In this range, as the monopolists market share falls, it raises its price, but at a rate that exactly offsets the drop in its market share. The average price – the average between the higher monopoly price and the lower fringe price – remains unchanged, implying that output remains unchanged. In this range introducing the fringe turns out to have no impact on the market – no change in average price, output, or welfare.

Eventually, as the dominant firm becomes small enough, its price becomes constrained by the buyers' willingness to pay. Reducing the market share further then only serves to raise output as it increases the share of the market where prices reflect the low bargaining power of the fringe firms. Eventually welfare does thus improve, but ultimately output actually becomes inefficiently high if the dominant firm vanishes completely.

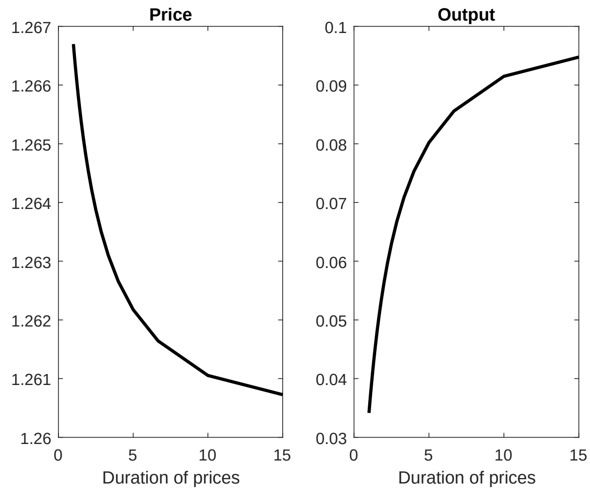

Figure 3: Duration

Notes: TBW

If the fringe sellers have higher bargaining power, introducing a small competitive fringe into the market is generally bad for welfare. The bottom panels of Figure 4 demonstrate this in the extreme case where introducing any competitive fringe is bad for welfare. As λ falls from one, the dominant firm's prices rise, as the firm recognizes that the costs of its high prices in terms of reduced customer acquisition will partly be borne by the fringe firms. The average price in the market rises, and hence output falls. Introducing a competitive fringe of any size thus decreases output and welfare in this market. More generally, there is at least a range of λ where increasing the size of the fringe reduces output, before it begins to increase, and approaches the level consistent with the sellers' bargaining power.

Seller bargaining power zero

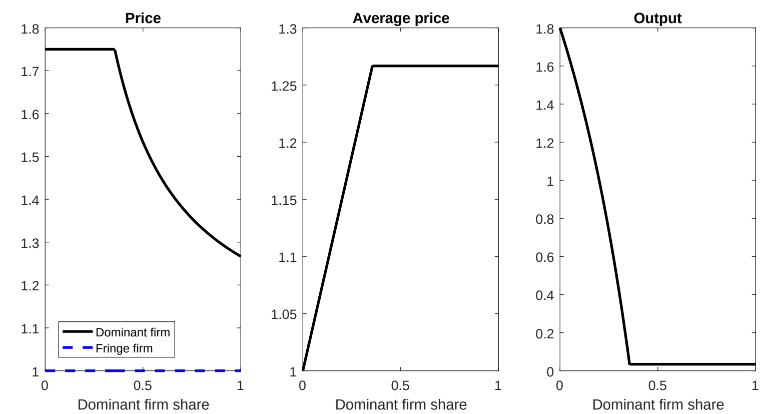

Seller bargaining power higher

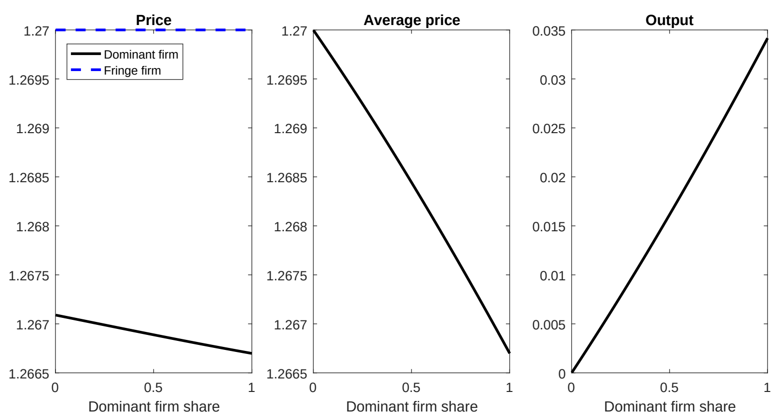

Figure 4: Competitive Fringe

Notes: TBW

4.4 Price Dispersion

So far we have interpreted this as a model of price setting for a monopolist or dominant firm in an industry. But we could also think of an industry with a cross-section of competitive firms, which face competition both from the buyer's outside option but also the other firms in the industry. In the style of the competitive search equilibrium, the basic firm problem would be unchanged. It would be natural to add to the model an exit shock for existing firms, which would then be replaced by new entrants of size zero in the subsequent period. This interpretation would yield a model of firm dynamics where small firms set lower prices and grow faster than larger firms. Both observations are consistent with the evidence on firm dynamics. Small firms grow faster than larger ones (see, e.g., Caves (1998)), and set lower prices (Foster, Haltiwanger, and Syverson 2008). But this extension allows the model to speak to further puzzle economists have sought to explain: price dispersion.

There exists a sizable theoretical literature developing theories of equilibrium price dispersion, seeking to capture the observation that the prices of individual goods appear to vary across sellers. There also exists a growing body of work documenting evidence of this dispersion in various markets. Model analysed in this paper offers one alternative theory for thinking about this dispersion. The theory is relatively tractable due to the competitive search equilibrium, in the sense that all one needs to do to solve the model is to solve the firm problem taking as given the buyers' outside option w. With the resulting function Θ in hand, one can then simulate the firms' behavior to characterize the degree of dispersion in prices, for example. Figure TBA illustrates the dispersion in prices emerging from simulating the model as described. In this simulation, the firms do not face shocks to demand or costs over time, which would naturally generate additional dispersion in outcomes as well. Rather, the dispersion in the figure is generated purely as a consequence of the model mechanism where large and small firms choose different prices and firm sizes evolve over time.

4.5 Shock Propagation: Productivity vs Demand

Given that models of customer markets were originally developed to think about the dynamics of price setting (Phelps and Winter 1970), it is natural to consider that here as well. In the context of firm-specific shocks to productivity and demand, the model produces what one might expect to see: In response to an increase in demand the firm raises prices, but customer acquisition accelerates nevertheless as the increase in prices does not fully make up for the increase in willingness to pay. Both prices and output thus rise. In response to a reduction in costs, customer acquisition accelerates because the firm reduces prices. Thus

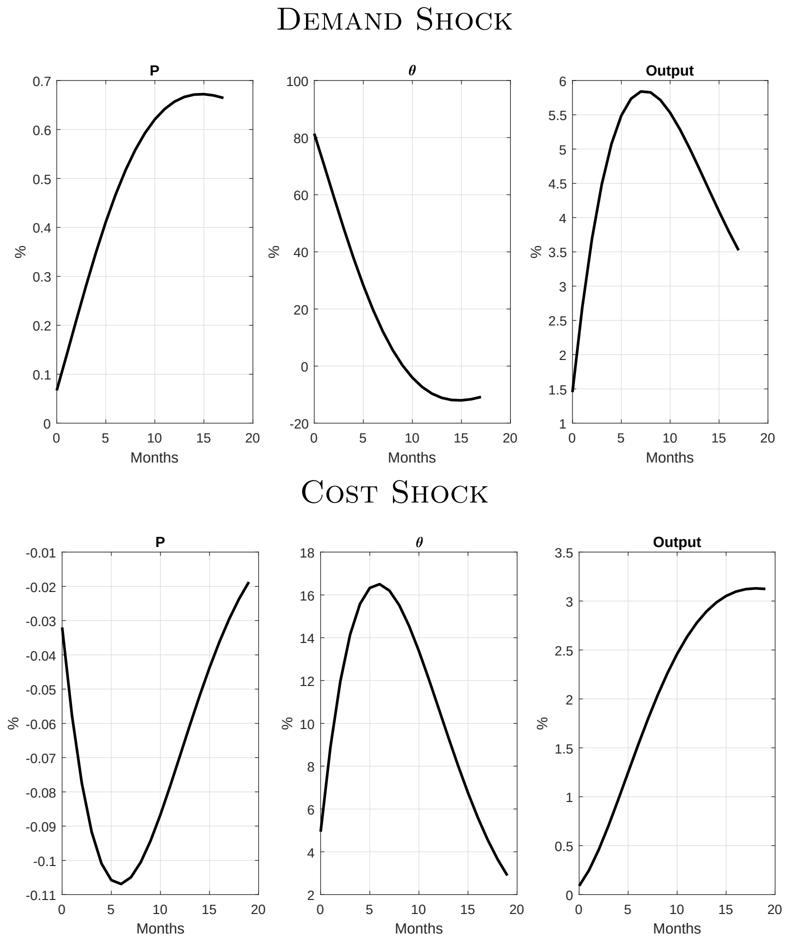

Figure 5: Cost and Demand

Notes: TBW

prices fall and output rises. In both cases the markup rises.[9]

In a recent empirical study, Pozzi and Schivardi (2016) make progress on the measurement of firm-level productivity and demand shocks, and determine that both appear to be important for firm growth. They find that some type of frictions appear to limit the response of output to both types of shocks, but that this effect is more pronounced in the case of productivity shocks. The model with probabilistic pricing generates this feature in a natural way, as the entry condition of buyers depends explicitly on demand but not productivity. If prices are sticky to respond to a shock, an increase in demand increases buyer entry, and hence the customer base and output of the firm. But an increase in productivity only increases buyer entry to the extent that the price falls, and hence the customer base

[9] Note that these implications concern responses to firm-level shocks. An aggregate shock should affect also the opportunity cost of search w, and an increase in this value would likely work to decrease markups.

and output only rise as the price adjusts. The figures below demonstrate these differing responses to shocks to demand and productivity. Both consider an expansionary shock – an increase in demand and decrease in costs, respectively. The price is sticky to adjust, but ultimately increases in response to the demand shock and decreases in response to the productivity shock.

5 Conclusions

To be written.

References

Arkolakis, C. (2010): "Market Penetration Costs and the New Consumers Margin in International Trade," *Journal of Political Economy*, 118, 1151–99.

Beggs, A., and P. Klemperer (1992): "Multi-Period Competition with Switching Costs," *Econometrica*, 60(3), 651–666.

Bronnenberg, B., and J.-P. Dube (2017): "The Formation of Consumer Brand Preferences," *Annual Review of Economics*, 9, X–X.

Burdett, K., and M. Coles (1997): "Steady State Price Distributions in a Noisy Search Equilibrium," *Journal of Economic Theory*, 72, 1–32.

Burdett, K., and K. L. Judd (1983): "Equilibrium Price Dispersion," *Econometrica*, 51(4), 955–970.

Carlton, D., and J. Perloff (1994): *Modern Industrial Organization, 2nd Ed.* Harper-Collins.

Caves, R. (1998): "Industrial Organization and New Findings on the Turnover and Mobility of Firms," *Journal of Economic Literature*, 36(4), 1947–1982.

Dinlersoz, E. M., and M. Yorukoglu (2012): "Information and Industry Dynamics," *American Economic Review*, 102, 884–913.

Dixit, A. K., and J. E. Stiglitz (1977): "Monopolistic Competition and Optimum Product Diversity," *American Economic Review*, 67(3), 297–308.

Drozd, L., and J. Nosal (2012): "Understanding International Prices: Customers as Capital," *American Economic Review*, 102, 364–95.

Farrell, J., and P. Klemperer (2007): "Coordination and Lock-In: Competition with Switching Costs and Network Effects," in *Handbook of Industrial Organization*, pp. 1967–2072.

Fitzgerald, D., S. Haller, and Y. Yedid-Levy (2016): "How Exporters Grow," Unpublished manuscript, FRB Minneapolis.

Foster, L., J. Haltiwanger, and C. Syverson (2008): "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?," *American Economic Review*, 98(1), 394–425.

——— (2016): "The Slow Growth of New Plants: Learning about Demand?," *Economica*, 83(329), 91–129.

Gilchrist, S., R. Schoenle, J. Sim, and E. Zakrajzek (2016a): "Financial Heterogeneity and Monetary Union," Unpublished manuscript, Boston University.

——— (2016b): "Inflation Dynamics During the Financial Crisis," *American Economic Review, forthcoming*.

Gourio, F., and L. Rudanko (2014): "Customer Capital," *Review of Economic Studies*, 81(3), 1102–36.

Gowrisankaran, G., and T. J. Holmes (2004): "Mergers and the Evolution of Industry Concentration: Results From the Dominant Firm Model," *RAND Journal of Economics*, 35, 561–82.

Klemperer, P. (1995): "Competition when Consumers have Switching Costs," *Review of Economic Studies*, 62, 515–539.

Kleshchelski, I., and N. Vincent (2009): "Market Share and Price Rigidity," *Journal of Monetary Economics*, 56(3), 344–352.

Krusell, P., B. Kuruscu, and A. Smith (2002): "Equilibrium Welfare and Government Policy with Quasi-Geometric Discounting," *Journal of Economic Theory*, 105, 42–72.

Lagos, R., and R. Wright (2005): "A Unified Framework for Monetary Theory and Policy Analysis," *Journal of Political Economy*, 113, 463–484.

Moen, E. R. (1997): "Competitive Search Equilibrium," *Journal of Political Economy*, 105(2), 385– 411.

Nakamura, E., and J. Steinsson (2008): "Price Setting in Forward-Looking Customer Markets," *Journal of Monetary Economics*.

——— (2013): "Price Rigidity: Microeconomic Evidence and Macroeconomic Implications," *Annual Review of Economics*, 5, 133–63.

Paciello, L., A. Pozzi, and N. Trachter (2013): "Price Setting with Customer Retention Concerns," Unpublished manuscript, Einaudi Institute for Economics and Finance.

Perla, J. (2016): "Product Awareness, Industry Life Cycles, and Aggregate Profits," Unpublished manuscript, University of British Columbia.

Phelps, E., and S. Winter (1970): "Optimal Price Policy under Atomistic Competition," in *Microeconomic Foundations of Employment and Inflation Theory*, pp. 309–337.

Pozzi, A., and F. Schivardi (2016): "Demand or Productivity: What Determines Firm Growth?," *RAND Journal of Economics*, 47(3), 608–630.

Ravn, M., S. Schmitt-Grohe, and M. Uribe (2006): "Deep Habits," *Review of Economic Studies*, 73(1), 195–218.

——— (2010): "Incomplete Cost Pass-Through Under Deep Habits," *Review of Economic Dynamics*, 13, 317–332.

Rogerson, R., R. Shimer, and R. Wright (2005): "Search-Theoretic Models of the Labor Market: A Survey," *Journal of Economic Literature*, 18, 959–988.

Shi, S. (2016): "Customer Relationship and Sales," *Journal of Economic Theory*, 166, 483–516.

Appendix

A Proofs

This section shows how to rewrite the firm problem (27) using the buyer entry condition (29), to arrive at the problem (30).

First, note that the firm's objective can be rewritten as

$$\sum_{t=0}^{\infty} \beta^t (n_t + \eta(\theta_t))(p_t - c) = n_0 \sum_{t=0}^{\infty} \beta^t (1-\delta)^t (p_t - c) + \sum_{t=0}^{\infty} \beta^t \sum_{k=0}^{t} (1-\delta)^{t-k} \eta(\theta_k)(p_t - c), \tag{48}$$

as the law of motion for the customer base implies that the total customer base in period t equals

$$n_t + \eta(\theta_t) = (1-\delta)^t n_0 + \sum_{k=0}^{t} (1-\delta)^{t-k} \eta(\theta_k). \tag{49}$$

The first term in (48) can then be rewritten as

$$n_0 \sum_{t=0}^{\infty} \beta^t (1-\delta)^t (p_t - c) = n_0 [\sum_{t=0}^{\infty} \beta^t (1-\delta)^t (u - c) - \frac{w}{\mu(\theta_0)}], \tag{50}$$

as the buyer entry condition implies that the present value of prices above can be expressed as

$$\sum_{t=0}^{\infty} \beta^t (1-\delta)^t p_t = \sum_{t=0}^{\infty} \beta^t (1-\delta)^t u - \frac{w}{\mu(\theta_0)}. \tag{51}$$

The second term in (48) can be rewritten as

$$\begin{aligned}
\sum_{t=0}^{\infty} \beta^t \sum_{k=0}^{t} (1-\delta)^{t-k} \eta(\theta_k)(p_t - c) &= \sum_{k=0}^{\infty} \beta^k \eta(\theta_k) \sum_{t=k}^{\infty} \beta^{t-k} (1-\delta)^{t-k} (p_t - c) \\
&= \sum_{k=0}^{\infty} \beta^k [\eta(\theta_k) \sum_{t=k}^{\infty} \beta^{t-k} (1-\delta)^{t-k} (u - c) - w\theta_k] \\
&= \sum_{t=0}^{\infty} \beta^t \sum_{k=0}^{t} (1-\delta)^{t-k} \eta(\theta_k)(u - c) - \sum_{t=0}^{\infty} \beta^t w \theta_t,
\end{aligned} \tag{52}$$

where the first equality follows from rearranging the terms in the sum, and the second again uses the buyers' entry condition to substitute out the present value of prices.

Combining the two terms in (50) and (52) and rearranging, the firm objective becomes

$$n_0[\sum_{t=0}^{\infty}\beta^t(1-\delta)^t(u-c)-\frac{w}{\mu(\theta_0)}]+\sum_{t=0}^{\infty}\beta^t[\sum_{k=0}^{t}(1-\delta)^{t-k}\eta(\theta_k)(u-c)-w\theta_t]$$

$$=\sum_{t=0}^{\infty}\beta^t[(n_t+\eta(\theta_t))(u-c)-w\theta_t]-\frac{wn_0}{\mu(\theta_0)}, \tag{53}$$

as stated in the text.

In the case with the competitive fringe, the derivation is little changed. The first term in (48) can be rewritten as

$$n_0\sum_{t=0}^{\infty}\beta^t(1-\delta)^t(p_t-c)=n_0[\sum_{t=0}^{\infty}\beta^t(1-\delta)^t(u-c+\frac{1-\lambda}{\lambda}\gamma(u-c))-\frac{w}{\lambda\mu(\theta_0)}], \tag{54}$$

as the buyer entry condition implies that the present value of prices above can be expressed as

$$\sum_{t=0}^{\infty}\beta^t(1-\delta)^tp_t=\sum_{t=0}^{\infty}\beta^t(1-\delta)^t(u+\frac{1-\lambda}{\lambda}\gamma(u-c))-\frac{w}{\mu(\theta_0)}. \tag{55}$$

The second term in (48) can be rewritten as

$$\sum_{t=0}^{\infty}\beta^t\sum_{k=0}^{t}(1-\delta)^{t-k}\eta(\theta_k)(p_t-c)=\sum_{k=0}^{\infty}\beta^k\eta(\theta_k)\sum_{t=k}^{\infty}\beta^{t-k}(1-\delta)^{t-k}(p_t-c)$$

$$=\sum_{k=0}^{\infty}\beta^k[\eta(\theta_k)\sum_{t=k}^{\infty}\beta^{t-k}(1-\delta)^{t-k}(u-c+\frac{1-\lambda}{\lambda}\gamma(u-c))-w\theta_k]$$

$$=\sum_{t=0}^{\infty}\beta^t\sum_{k=0}^{t}(1-\delta)^{t-k}\eta(\theta_k)(u-c+\frac{1-\lambda}{\lambda}\gamma(u-c))-\sum_{t=0}^{\infty}\beta^tw\theta_t, \tag{56}$$

where the first equality follows from rearranging the terms in the sum, and the second again uses the buyers' entry condition to substitute out the present value of prices.

Combining the two terms in (54) and (56) and rearranging, the firm objective becomes

$$n_0[\sum_{t=0}^{\infty}\beta^t(1-\delta)^t(u-c+\frac{1-\lambda}{\lambda}\gamma(u-c))-\frac{w}{\mu(\theta_0)}]$$

$$+\sum_{t=0}^{\infty}\beta^t[\sum_{k=0}^{t}(1-\delta)^{t-k}\eta(\theta_k)(u-c+\frac{1-\lambda}{\lambda}\gamma(u-c))-w\theta_t]$$

$$=\sum_{t=0}^{\infty}\beta^t[(n_t+\eta(\theta_t))(u-c+\frac{1-\lambda}{\lambda}\gamma(u-c))-w\theta_t]-\frac{wn_0}{\mu(\theta_0)}, \tag{57}$$

as stated in the text.

In the case with curvature in demand and/or costs, the derivation is also little changed. The first term in (48) can be rewritten as

$$n_0 \sum_{t=0}^{\infty} \beta^t (1-\delta)^t p_t = n_0 [\sum_{t=0}^{\infty} \beta^t (1-\delta)^t u'(n_t + \eta(\theta_t)) - \frac{w}{\mu(\theta_0)}], \tag{58}$$

as the buyer entry condition implies that the present value of prices above can be expressed as

$$\sum_{t=0}^{\infty} \beta^t (1-\delta)^t p_t = \sum_{t=0}^{\infty} \beta^t (1-\delta)^t u'(n_t + \eta(\theta_t)) - \frac{w}{\mu(\theta_0)}. \tag{59}$$

The second term in (48) can be rewritten as

$$\begin{aligned} \sum_{t=0}^{\infty} \beta^t \sum_{k=0}^{t} (1-\delta)^{t-k} \eta(\theta_k) p_t &= \sum_{k=0}^{\infty} \beta^k \eta(\theta_k) \sum_{t=k}^{\infty} \beta^{t-k} (1-\delta)^{t-k} p_t \\ &= \sum_{k=0}^{\infty} \beta^k [\eta(\theta_k) \sum_{t=k}^{\infty} \beta^{t-k} (1-\delta)^{t-k} u'(n_t + \eta(\theta_t)) - w\theta_k], \end{aligned} \tag{60}$$

where the first equality follows from rearranging the terms in the sum, and the second again uses the buyers' entry condition to substitute out the present value of prices.

Combining the two terms in (58) and (60) and rearranging, the firm profit becomes

$$\begin{aligned} & n_0 [\sum_{t=0}^{\infty} \beta^t (1-\delta)^t u'(n_t + \eta(\theta_t)) - \frac{w}{\mu(\theta_0)}] + \sum_{k=0}^{\infty} \beta^k [\eta(\theta_k) \sum_{t=k}^{\infty} \beta^{t-k} (1-\delta)^{t-k} u'(n_t + \eta(\theta_t)) - w\theta_k] \\ &= \sum_{t=0}^{\infty} \beta^t [(n_t + \eta(\theta_t)) u'(n_t + \eta(\theta_t)) - w\theta_t] - \frac{w n_0}{\mu(\theta_0)}, \end{aligned} \tag{61}$$

as stated in the text.

B Solution approach

The solution method extends that outlined in Krusell, Kuruscu, and Smith (2002) for the generalized Euler equation. We begin by describing the approach in the context of the generalized equation (36) (or with the fringe firms included in (47)), and then the system (37)-(39).

The generalized Euler equation is a functional equation in $\Theta(n)$, defined over a range of values of n encompassing the steady state. The idea is to calculate a Taylor polynomial approximating $\Theta(n)$ around its steady state. Calculating a k^{th} order polynomial involves first analytically differentiating the Euler equation k times with respect to n, acknowledging that $\Theta(n)$ and $N(n, \Theta(n))$ are functions of n. The resulting system of $k+1$ equations is then evaluated in steady state, setting Θ and its derivatives to their steady-state values $\theta, \theta', ..., \theta^{k+1}$, as well as replacing n with its steady-state value $\mu(\theta)(1-\delta)/(\delta+\mu(\theta)(1-\delta))$. This yields a system of $k+1$ equations in the unknowns $\theta, \theta', ..., \theta^{k+1}$. Setting θ^{k+1} to zero, one then arrives at a system of $k+1$ equations in $k+1$ unknowns—the coefficients of the polynomial approximating $\Theta(n)$.

The implementation has two stages. The first stage involves using an analytical solver, such as Mathematica, to calculate the analytical derivatives and to perform the substitutions needed to arrive at the system of equations in $\theta, \theta', ..., \theta^k$. The second stage involves using a numerical solver, such as Matlab, to solve this nonlinear system of equations. Solving the system can, in practice, require a good initial guess, so we approach the problem iteratively. Starting with a 0^{th} order Taylor polynomial, one can proceed to successively higher-order polynomials using the results from the previous step as initial guesses, looking for convergence in the coefficients of the polynomial along the way.

With probabilistic pricing, the same approach is applied directly to the system (37)-(39). Instead of taking derivatives of a Euler equation, derivatives are taken of the first order condition of the union objective in equation (39). These derivatives will involve the values and derivatives of the functions $V^f(n, \theta), V^r(n)$ as well, and must therefore be calculated in a separate step, to incorporate them into the equations.

Again, the solution approach is implemented in two stages. In the analytical derivation stage, we tackle the system (37)-(39) starting from the value equations (37)-(38), and then the first order condition for equation (39).

Beginning with the system (37)-(38): To solve for the derivatives of values $V^f(n, \theta), V^r(n)$ at steady state, differentiate the system with respect to (n, θ) (acknowledging that the recontracting tightness $\Theta(n)$ and $N(n, \Theta(n))$ are functions of n), evaluate the resulting equations at steady state, and solve the resulting system for steady-state values of the derivatives $V_n^f(n, \theta), V_n^r(n)$ as functions of the steady-state values and derivatives of $\Theta(n)$ and values of $V^f(n, \theta), V^r(n)$. To solve for higher order derivatives of $V^f(\theta, n), V^r(n)$, proceed along the same lines. The equations allow analytical solutions in each step.

Finally, we turn to the first-order condition for equation (39), proceeding to differentiate the equation k times, and evaluate the resulting equations in steady state. This yields $k+1$ equations in the steady-state value and derivatives of $\Theta(n)$, which depend also on those of

$V^f(n,\theta), V^r(n)$. To eliminate the latter from the equations, one can use the expressions derived above. This yields $k+1$ equations representing the $0^{th} - k^{th}$ derivatives of the first order condition in the unknowns $\theta, \theta', ..., \theta^{k+1}$. Setting θ^{k+1} to zero, one arrives at a system of $k+1$ equations in $k+1$ unknowns—the coefficients of the polynomial approximating $\Theta(n)$.

The numerical stage involves solving this nonlinear system of equations. In practice, solving the system can require a good initial guess, so the problem is approached iteratively. One starts with a zero-order Taylor polynomial and proceeds to successively higher-order polynomials using the results from the previous step as initial guesses, looking for convergence in the coefficients of the polynomial along the way.