Price Stability and Markup Dispersion without Frictions: Inventory Management under Demand Uncertainty
PRELIMINARY AND INCOMPLETE

Olga A. Timoshenko *

Temple University

March 25, 2026

Abstract

This paper shows that inventory management under demand uncertainty generates empirical predictions that previously required strong theoretical assumptions. In each period, firms produce before observing the current demand shock, then sell after observing it, with the option to carry unsold goods forward. The joint analysis of quantities and prices delivers results that are absent from models without inventory or without learning. Price stability emerges as an optimal outcome without menu costs or information frictions. Markup dispersion arises without variable demand elasticity or productivity heterogeneity. Both results follow from a single mechanism: the piecewise option value of carry-over inventory. The paper offers parsimonious foundations for patterns attributed to separate frictions, opening new questions about inventory and firm-level pricing behavior.

Keywords: Bayesian learning, inventory, demand uncertainty, firm dynamics, markup dispersion, overproduction.

1 Introduction

Inventories are a central feature of firm behavior, yet their role over the firm life cycle is not well understood. For firms that produce to stock, a fundamental challenge is that production decisions must be committed to before the period's demand is observed, while the quantity sold and the price charged are determined only after demand is realized. Under pure demand learning (Jovanovic, 1982; Arkolakis et al., 2018), firms form expectations about demand,

*Temple University, Department of Economics, 1115 Polett Walk, Philadelphia, PA 19122, United States. E-mail: olga.timoshenko@temple.edu

produce accordingly, and sell all output within the period at the market-clearing price. This paper considers firms that have an additional option of carrying unsold output forward as inventory and deferring the sales and pricing decision to a future period with a different demand realization. This paper embeds both mechanisms into a single dynamic model of the firm life cycle, and shows that their joint solution delivers predictions about production, sales, pricing and markups that neither mechanism generates on its own and that previously required strong theoretical assumptions.

I embed an inventory management problem along the lines of Alessandria et al. (2010a,b) into a model of monopolistic competition and Bayesian demand learning following Arkolakis et al. (2018). At the start of each period, a firm forms expectations about its idiosyncratic unobserved demand and chooses how much to produce on the basis of those expectations. Demand shocks are then realised and the firm decides how much to sell and at what price, with the option to carry unsold goods into the next period at a constant rate of depreciation. The firm's optimal production decision targets an inventory level determined solely by expected demand. A firm that arrives under-stocked produces until that target is reached; a firm that arrives over-stocked does not produce and proceeds directly to the sales and pricing decision.

This timing gives rise to a marginal value of inventory that is non-constant and depends on the firm's stock position. When the firm arrives under-stocked, an extra unit of inventory saves one unit of variable production cost, so the marginal value is constant and equal to the replacement cost. When the firm arrives over-stocked, no production takes place and the marginal value of an additional unit is determined instead by the expected marginal value of future sales, which varies with the size of the excess stock.

In the canonical model of monopolistic competition (Melitz, 2003), a firm equates marginal revenue to a constant marginal cost of production, since all output must be sold within the period it is produced. A dynamic learning model without inventory (Arkolakis et al., 2018) preserves this structure, because the absence of storage commits the firm to selling its entire output each period.

The present model departs from both benchmarks by replacing the constant marginal cost with the option value of carry-over inventory, which is a non-monotone, piecewise function of the quantity sold, reflecting the qualitatively distinct continuation scenarios that arise depending on where next-period inventory falls relative to the production target. This non-constancy drives the paper's central results. The optimal sales and pricing decisions are characterised by three regimes, each with distinct implications for how quantities and prices respond to demand shocks and to the inherited stock of inventory.

The optimal policy partitions the demand space into three regimes determined by the

size of the demand shock relative to two endogenous thresholds that depend on beliefs, inventory, and the continuation value. In the *high-demand regime*, the inventory constraint binds and the firm sells its entire stock. The quantity is fixed by the stock level and all of the demand adjustment falls on the price, so that a higher demand shock raises the price while a larger stock lowers it. The implied markup exceeds the CES benchmark, because the firm is extracting a scarcity rent from its own depleted inventory.

In the *moderate-demand regime*, the firm sells the standard CES quantity at a price that is independent of both the demand shock and the inventory level. All variation in demand is absorbed on the quantity margin. The price is pinned at a constant markup over the replacement cost of inventory rather than over the production cost, reflecting the fact that the opportunity cost of selling a unit is the cost of replenishing it in the next period. This price stability does not follow mechanically from CES preferences. In a static CES model the markup is applied to a constant marginal cost of production, delivering a price that is invariant to demand by assumption. In this model the markup is applied to the option value of carry-over inventory, which is endogenously constant in the moderate regime because the next-period firm will produce to its target stock regardless of the quantity carried forward. It is this flatness of the option value, rather than the preference structure, that pins the price and makes quantity the sole margin of adjustment.

In the *low-demand regime*, carry-over inventory is so abundant that the next-period firm arrives over-stocked and would not produce regardless of how much is sold in the current period. The option value of carry-over inventory therefore lies below the replacement cost, because an extra unit of inventory does not save a production cost at the margin but instead lands in a region of diminishing returns to the continuation value. This is the only regime in which both price and quantity respond to demand shocks. A larger inherited stock pushes next-period inventory deeper into the over-stocked region, further eroding the marginal value of carry-over, which lowers the price and raises the quantity sold.

Taken together, these three regimes imply that price is a non-monotone function of the demand shock. The dominant adjustment margin shifts from prices to quantities as the firm moves from the tails of the demand distribution toward the interior, with quantity adjustment alone characterizing the moderate-demand regime and price adjustment dominating at the boundaries.

A further implication is that markup dispersion across firms with identical productivity arises endogenously from heterogeneity in inventory positions and the demand regime each firm faces. This is distinct from the markup dispersion generated by variable demand elasticity in models with non-CES preferences (Melitz and Ottaviano, 2008), where dispersion reflects productivity differences interacting with market toughness. Here markup disper-

sion arises even under CES preferences and among firms with identical productivity, driven entirely by differences in inventory and demand realizations.

The results connect the model to several distinct empirical literatures. A branch of empirical literature documents that prices change infrequently relative to quantities, with price changes concentrated in large, discrete adjustments rather than continuous drift (Nakamura and Steinsson, 2008). The theoretical literature attributes price rigidity to menu costs (Golosov and Lucas Jr, 2007), customer relationships (Rotemberg, 2005), or information frictions (Mankiw and Reis, 2002). In all of these frameworks, rigidity is a constrained outcome arising from a friction that prevents price adjustment. This present paper offers a complementary perspective. In the moderate-demand regime, the firm optimally leaves its price unchanged and absorbs all demand variation through quantity adjustment, not because adjusting the price is costly, but because the shadow cost of selling is constant and equal to the replacement cost of inventory. Price stability is therefore an equilibrium outcome of the inventory management problem rather than a deviation from the frictionless benchmark.

A related body of work documents substantial and rising dispersion in markups across firms (De Loecker et al., 2020). The theoretical literature offers several explanations for markup heterogeneity across firms. Under variable demand elasticity, more productive firms produce more output, face a more elastic residual demand, and charge higher markups, so that productivity dispersion translates directly into markup dispersion (Melitz and Ottaviano, 2008). In oligopolistic markets, firms with larger market shares exercise greater market power and charge higher markups, so that markup heterogeneity reflects differences in firm size and competitive position across segmented markets (Atkeson and Burstein, 2008). The present model offers an alternative mechanism. Markups vary across firms with identical productivity purely because firms face different inventory positions and occupy different demand regimes, a source that is orthogonal to both productivity and market structure. This suggests that part of the markup dispersion documented by De Loecker et al. (2020) may reflect heterogeneity in inventory positions across firms rather than differences in market power or productive efficiency.

The rest of the paper proceeds as follows. Section 2 presents the model, Section 3 characterizes the optimal solution, and Section 4 concludes. All proofs are relegated to Appendix A.

2 The Model

This section introduces the inventory management problem of a firm by Alessandria, Kaboski, and Midrigan (2010b) into the learning model of firm growth by Arkolakis, Papageor-

giou, and Timoshenko (2018).

2.1 Set-up

Time is discrete and denoted by t. The economy is populated by heterogeneous firms. Each firm supplies a differentiated variety of a good, and firms compete in a monopolistic competition environment. Firms can hold inventory of goods and carry the inventory over time.

2.2 Demand

At time t each firm faces a constant elasticity of substitution demand for its variety given by

$$q_t = e^{a_t} p_t^{-\sigma} Y_t P_t^{\sigma-1}, \tag{1}$$

where σ is the elasticity of substitution across varieties, a_t is the demand shock for a firm's variety, p_t is the price of that variety, Y_t is the aggregate income level, and P_t is the aggregate price level. A firm-level subscript is omitted for compactness of notation.

2.3 Firms

Firms are heterogeneous in terms of their time invariant idiosyncratic productivity, φ, and demand, θ. A firm's demand level θ is not observed by the firm, and is subject to transient independently and identically distributed preference shocks, ϵ_t. Firms form expectations about their demand levels by observing realizations of demand shocks, a_t. A demand shock is given by the sum of the demand level and the transient preference shock: $a_t = \theta + \epsilon_t$.

Timing At the beginning of period t, a firm observes its beginning-of-period stock of inventory, productivity, and expected demand level, and decides whether to exit, or to stay and produce. If the firm decides to stay, it makes a production decision to maximize the expected present discounted values of the stream of profits. Once production decisions are made, demand shocks are realized. After the firm observes the realization of its demand shock, the firm chooses how much of its current stock of inventory to sell to the market. Prices clear the goods markets. The firm updates its demand expectations and carries into the next period the stock of unsold inventory. Inventory depreciates at rate ρ.

2.4 The Problem of the Firm

Denote by s the beginning-of-period stock of inventory of a firm. Next, denote by $\bar{a}$ and n the mean and the number of the observed demand shocks respectively.

Stage 0: Entry and Exit

A firm enters a period with a state triplet given by $(\bar{a}, n, s)$ and decides whether to exit or stay, produce and sell. The value of exit is normalized to zero. If the firm decides to stay, it incurs fixed production cost f, denominated in the units of labor, and earns a continuation value of the production stage. The value function at the beginning of a period is therefore given by

$$V_0(\bar{a}, n, s) = \max\{0; -wf + V_1(\bar{a}, n, s)\}. \tag{2}$$

The problem of a continuing firm consists of two stages. Stage 1 is the Production Stage, and Stage 2 is the Sales Stage.

Stage 1: Production Stage

At Stage 1 the firm chooses how much of the good to produce. To produce i units of the final good, the firm incurs $i\frac{w}{\varphi}$ of the variable production cost. Once production has taken place, the firm carries over into Stage 2 the middle-of-period stock of inventory, s', given by the sum of the beginning-of-period stock of inventory, s, and the subsequent production, i, if any: $s' = s + i$.

Denote by $V_1(\bar{a}, n, s)$ the value of the firm at Stage 1.[1] Value function $V_1(\bar{a}, n, s)$ is hence given by

$$\begin{aligned} V_1(\bar{a}, n, s) \quad = \quad & \max_{i \geq 0} -i\frac{w}{\varphi} + \mathbb{E}_{a|\bar{a},n} V_2(\bar{a}, n, s', a) \\ & \text{s.t. } s' = s + i, \end{aligned} \tag{3}$$

where the expectation in equation (3) is taken with respect to the realized value of the demand shock, a, and is given by a Normal distribution with mean and variance determined by $(\bar{a}, n)$. Function $V_2(\bar{a}, n, s', a)$ is the value of the firm at Stage 2, which I describe next.

Stage 2: Sales Stage

At Stage 2 the firm chooses how much of the good to sell up to the amount given by the middle-of-period stock of inventory, s'. The good's market clearing price is implicitly determined by the downward sloping demand curve in equation (1). By selling q units of the good, the firm therefore earns revenue in the amount of $p(q)q = e^{a/\sigma} q^{\frac{\sigma-1}{\sigma}} Y^{1/\sigma} P^{\frac{\sigma-1}{\sigma}}$. After selling q units of the good, the firm carries over into the next period the depreciated value of the unsold inventory in the amount of $(1-\rho)(s'-q)$. The firm subsequently updates

[1] The value function of the firm also depends on the productivity of the firm. For compactness of notation I omit explicitly stating such dependence since productivity of a firm does not vary over time. This section therefore characterizes the problem of a firm conditional on a given productivity level.

its number of the observed demand shocks from n to $n+1$ and the mean of the observed demand shocks according to $\bar{a}' = (\bar{a}n + a)/(n+1)$.

Denote by $V_2(\bar{a}, n, s', a)$ the value function of the firm at Stage 2. Value function $V_2(\bar{a}, n, s', a)$ is hence given by

$$\begin{aligned} V_2(\bar{a}, n, s', a) &= \max_{0<q\leq s'} e^{a/\sigma} q^{\frac{\sigma-1}{\sigma}} Y^{1/\sigma} P^{\frac{\sigma-1}{\sigma}} + \beta(1-\delta) V_0(\bar{a}', n+1, s'') \\ &\quad \text{s.t. } \bar{a}' = \frac{\bar{a}n + a}{n+1} \text{ and } s'' = (1-\rho)(s'-q), \end{aligned} \tag{4}$$

where β denotes the discount rate, and δ denotes the exogenous death rate of firms.

To summarize, the firm's problem is described by the system of equations (2)-(4).

3 Properties of the Optimal Solution

This section describes properties of the firm's optimal production and sales decisions.

3.1 Entry, Exit, and the Value of Continuation

Proposition 1 *Properties of the entry and exit decision.*

Holding all else constant, for a given $(\bar{a}, n)$ there exits $\hat{s}(\bar{a}, n)$ such that
(i) if $s < \hat{s}(\bar{a}, n)$ a firm exits and

$$V_0(\bar{a}, n, s) = 0, \tag{5}$$

(ii) if $s \geq \hat{s}(\bar{a}, n)$ a firm stays and

$$V_0(\bar{a}, n, s) = -wf + V_1(\bar{a}, n, s). \tag{6}$$

Provided a solution exits, $\hat{s}(\bar{a}, n)$ is implicitly determined by

$$V_1(\bar{a}, n, \hat{s}(\bar{a}, n)) = wf. \tag{7}$$

If $V_1(\bar{a}, n, s) > wf \ \forall s$, then $\hat{s}(\bar{a}, n) = 0$.

Proofs to all propositions are included in Appendix A.

Proposition 1 states that, for a given set of demand expectations summarized by $(\bar{a}, n)$, there exists a minimum starting inventory $\hat{s}(\bar{a}, n)$ that makes it profitable for a firm to continue operations.

3.2 The Production Stage

Proposition 2 describes properties of the optimal production decision at Stage 1 of the firm's problem.

Proposition 2 *Properties of the production decision.*

Holding all else constant, for a given $(\bar{a}, n)$ *there exits* $s^*(\bar{a}, n)$ *such that*
(i) if $s < s^*(\bar{a}, n)$ *then*

$$\begin{aligned} s' &= s^*(\bar{a}, n), \\ i^*(\bar{a}, n, s) &= s^*(\bar{a}, n) - s, \text{ and} \\ V_1(\bar{a}, n, s) &= -(s^*(\bar{a}, n) - s)\frac{w}{\varphi} + \mathbb{E}_{a|\bar{a},n} V_2(\bar{a}, n, s^*(\bar{a}, n), a), \end{aligned}$$

(ii) if $s \geq s^*(\bar{a}, n)$ *then*

$$\begin{aligned} s' &= s, \\ i^*(\bar{a}, n, s) &= 0, \text{ and} \\ V_1(\bar{a}, n, s) &= \mathbb{E}_{a|\bar{a},n} V_2(\bar{a}, n, s, a). \end{aligned}$$

Proposition 2 states that for a given set of beliefs summarized by $(\bar{a}, n)$ there exits a corresponding production given by $s^*(\bar{a}, n)$, that is optimal for a firm to carry into the Sales Stage. Hence, at Stage 1, the firm will produce up to the optimal amount, or will not engage in production if the beginning-of-period stock of inventory is above the optimal production level.

Notably, the optimal firm size at Stage 1, $s^*(\bar{a}, n)$, is determined solely by the firm's productivity and expected demand level, and does not depend on the beginning-of-period stock of inventory. This can be seen from the first order condition to the firm's problem in equation (3) that yields

$$\frac{w}{\varphi} = \frac{\partial \mathbb{E}_{a|\bar{a},n} V_2(\bar{a}, n, s^*, a)}{\partial s^*}. \tag{8}$$

Equation (8) shows that the optimal firm size, $s^*(\bar{a}, n)$, is determined by equating the marginal variable cost of production to the expected marginal benefit, and is independent of the beginning-of-period stock of inventory, s.

Given that the optimal firm size at Stage 1 does not depend on the the beginning-of-period stock of inventory, the marginal value of a unit of inventory at the start of a period depends on whether the initial stock is above or below the optimal level of $s^*(\bar{a}, n)$. Proposition 2 implies that the marginal value of a unit of inventory given by

$$\frac{\partial V_1(\bar{a}, n, s)}{\partial s} = \begin{cases} \frac{w}{\varphi} & \text{if } s < s^*(\bar{a}, n), \\ \frac{\partial \mathbb{E}_{a|\bar{a},n} V_2(\bar{a}, n, s, a)}{\partial s} & \text{if } s \geq s^*(\bar{a}, n). \end{cases} \tag{9}$$

When production occurs at Stage 1, as per part (i) of Proposition 2, an extra unit of inventory at the start of a period reduces the production by one unit and therefore saves

the firm the variable production cost. When production does not occur at Stage 1, as per part (ii) of Proposition 2, an extra unit of inventory can subsequently be sold at Stage 2. Therefore, the unit's marginal value equals to the marginal expected value from the Sales Stage.

Notice from equation (9) that in the region where a firm engages in production, i.e. $s < s^*(\bar{a}, n)$ or equivalently $i^*(\bar{a}, n, s) > 0$, the value function is linear in s and can be written as

$$V_1(\bar{a}, n, s) = \frac{w}{\varphi} s + V_1(\bar{a}, n, 0), \tag{10}$$

where

$$V_1(\bar{a}, n, 0) = -\frac{w}{\varphi} s^*(\bar{a}, n) + \mathbb{E}_{a|\bar{a},n} V_2(\bar{a}, n, s^*(\bar{a}, n), a). \tag{11}$$

The first stage value function admits the following form

$$V_1(\bar{a}, n, 0) = -\frac{w}{\varphi} s^*(\bar{a}, n) + \mathbb{E}_{a|\bar{a},n} V_2(\bar{a}, n, s^*(\bar{a}, n), a) \tag{12}$$

$$V_1(\bar{a}, n, s) = \frac{w}{\varphi} s + V_1(\bar{a}, n, 0) \text{ if } 0 < s < s^*(\bar{a}, n) \tag{13}$$

$$V_1(\bar{a}, n, s) = \mathbb{E}_{a|\bar{a},n} V_2(\bar{a}, n, s, a) \text{ if } s \geq s^*(\bar{a}, n) \tag{14}$$

3.3 The Sales Stage

The firm enters the sales stage with stock s' and observes demand shock a. It chooses sales quantity $0 < q \leq s'$, leaving carry-over inventory stock in the amount of $s'' = (1-\rho)(s'-q)$. The interior first-order condition in problem (4) is given by

$$\underbrace{\frac{\sigma - 1}{\sigma} e^{a/\sigma} q^{-1/\sigma} Y^{1/\sigma} P^{\frac{\sigma-1}{\sigma}}}_{MR(q,a)\text{: marginal revenue today}} = \underbrace{\beta(1-\delta)(1-\rho)\frac{\partial V_0(\bar{a}', n+1, s'')}{\partial s''}}_{OV(q,a)\text{: option value of a unit tomorrow}}. \tag{15}$$

Notice, the left-hand side of equation (15) is the standard marginal revenue function under CES preferences. The optimal quantity is therefore determined by equating marginal revenue to the shadow cost of selling a unit.

To understand the intuition behind this first order condition, consider the following two benchmark models: the static Melitz (2003) model and a dynamic Arkolakis et al. (2018) model of learning without inventory. In a static model with CES preferences the right-hand side of equation (15) reduces to the marginal cost of production, w/φ, and equation (15) becomes

$$\frac{\sigma - 1}{\sigma} e^{a/\sigma} q^{-1/\sigma} Y^{1/\sigma} P^{\frac{\sigma-1}{\sigma}} = \frac{w}{\varphi}. \tag{16}$$

This yields the familiar CES optimal quantity.

In a dynamic learning model of Arkolakis et al. (2018), where every unit produced must be sold in the current period, the optimal quantity is similarly determined by equating *expected* marginal revenue to marginal cost and equation (15) becomes

$$\frac{\sigma-1}{\sigma}\mathbb{E}_{a|(\bar{a}',n)}\left(e^{a/\sigma}\right)q^{-1/\sigma}Y^{1/\sigma}P^{\frac{\sigma-1}{\sigma}}=\frac{w}{\varphi}. \tag{17}$$

When goods can be stored, the right-hand side of equation (15) captures the *option value* of a unit of a good: what the firm gives up tomorrow by selling one more unit today rather than storing it. Formally,

$$OV(q,a)\equiv\beta(1-\delta)(1-\rho)\frac{\partial V_0(\bar{a}',n+1,s'')}{\partial s''}. \tag{18}$$

Unlike in the benchmark cases, the option value is no longer constant. As q increases from 0 to s', carry-over inventory s'' falls from $(1-\rho)s'$ to 0, and $OV(q)$ is determined by where the next-period stock of inventory, s'', falls relative to the next-period exit, $\hat{s}(\bar{a}',n+1)$, and production, $s^*(\bar{a}',n+1)$, thresholds. Using equations (5)-(6) and (12)-(14), the option value is given by

$$OV(q,a)=\beta(1-\delta)(1-\rho)\times\begin{cases}\frac{\partial\mathbb{E}_{\tilde{a}|\bar{a}',n+1}V_2(\bar{a}',n+1,s'',\tilde{a})}{\partial s''} & \text{if } s''>s^*(\bar{a}',n+1) \quad,\\ w/\varphi & \text{if } \hat{s}(\bar{a}',n+1)\le s''\le s^*(\bar{a}',n+1) \quad,\\ 0 & \text{if } s''<\hat{s}(\bar{a}',n+1) \quad,\end{cases} \tag{19}$$

Figure 1 depicts $OV(q,a)$ as a function of q holding all else constant.

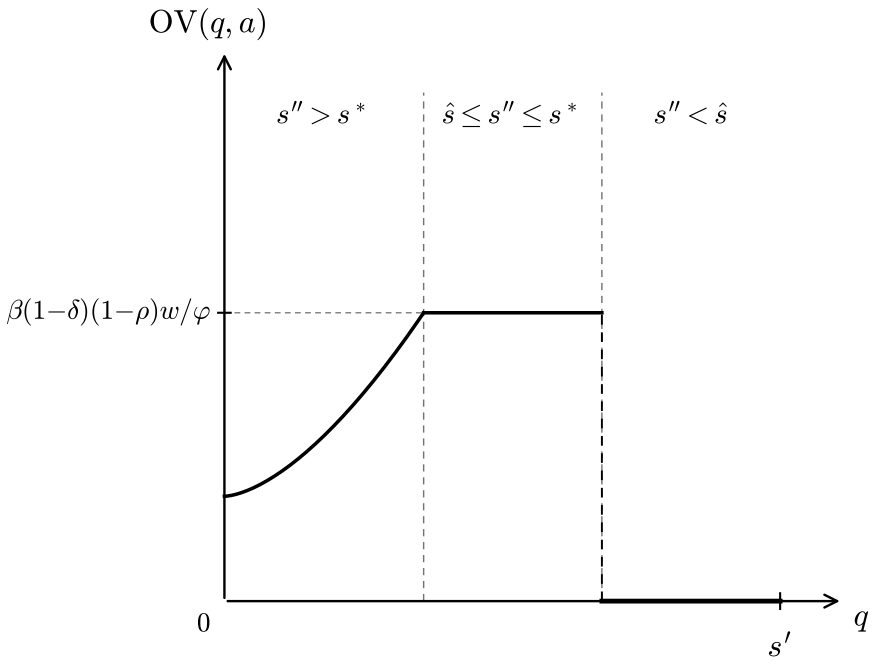
When it is optimal to sell little in a given period, optimal q is small and the carry-over inventory is sufficiently large such that $s''>s^*$, the firm arrives into the next period being over-stocked and will not produce. In this case, the value of an additional unit of inventory is given by the marginal value of that unit in the sales stage, namely $\partial\mathbb{E}(V_2)/\partial s''$. Since $\mathbb{E}(V_2)$ is strictly concave in s'', $\partial\mathbb{E}(V_2)/\partial s''$ is declining in s'' and, consequently, increasing in q.

When the carry-over inventory is moderate ($\hat{s}\le s''\le s^*$), the firm will produce up to the respective optimal level of inventory s^* in the next period. In this case, an extra unit of inventory saves exactly w/φ in production cost, so OV is flat in this region and equal to the replacement cost $\beta(1-\delta)(1-\rho)w/\varphi$.

When the carry-over inventory falls below the next-period exit threshold ($s''<\hat{s}$), an extra unit of inventory is worthless. Hence, $OV=0$ in this region.

The optimal firm-level quantity and price are determined by the intersection of the de-

Figure 1: Option Value of Carry-Over Stock of Inventory, $OV(q,a)$



Notes: The figure plots the option value $OV(q,a)$ of a unit of inventory in a given period as a function of current sales q, holding all else constant. s' is the current stock of inventory, $s'' = (1-\rho)(s'-q)$ is the stock of inventory carried over into the next period, s^* is the next-period production target, and $\hat{s}$ is the next-period exit threshold.

clining marginal revenue curve with the piecewise $OV(q,a)$ curve. As the demand shock a rises, MR curve shifts upward, the intersection moves through the three respective regimes. and generates three qualitatively distinct regimes for optimal sales and prices, which are characterized in Proposition 3 below.

Proposition 3 *Properties of the optimal sales and prices.*

Denote by s'' the beginning-of-period stock of inventory in the next period. Let $B \equiv (\beta(1-\delta)(1-\rho))^{-\sigma}(\varphi/w)^{\sigma}((\sigma-1)/\sigma)^{\sigma}YP^{\sigma-1}$. Holding all else constant, there exit $a_H^*(\bar{a},n,s')$ and $a_L^*(\bar{a},n,s')$ such that

(i) **High demand:** *if $a \geq a_H^*(\bar{a},n,s')$ then*

$$q^*(\bar{a},n,a,s') = s', \tag{20}$$
$$p^*(\bar{a},n,a,s') = e^{a/\sigma}(s')^{-1/\sigma}Y^{1/\sigma}P^{(\sigma-1)/\sigma}, \tag{21}$$
$$s'' = 0. \tag{22}$$

(ii) **Moderate demand:** *if* $a_L^*(\bar{a}, n, s') < a < a_H^*(\bar{a}, n, s')$ *then* $0 < q^*(\bar{a}, n, a, s') < s'$ *and*

$$q^*(\bar{a}, n, a, s') = Be^a, \tag{23}$$

$$p^*(\bar{a}, n, a, s') = \frac{\sigma}{\sigma - 1} \cdot \beta(1-\delta)(1-\rho)\frac{w}{\varphi}, \tag{24}$$

$$s'' = (1-\rho)(s' - q^*(\bar{a}, n, a, s')). \tag{25}$$

(iii) **Low demand:** *if* $a \leq a_L^*(\bar{a}, n, s')$ *then* $0 < q^*(\bar{a}, n, a, s') < s'$ *and*

$$q^*(\bar{a}, n, a, s') = Be^a \left(\frac{\partial \mathbb{E}_a V_2(\bar{a}', n+1, s''(s', q^*), a)}{\partial s''} \right)^{-\sigma}, \tag{26}$$

$$p^*(\bar{a}, n, a, s') = \frac{\sigma}{\sigma - 1} \cdot \beta(1-\delta)(1-\rho) \cdot \frac{\partial \mathbb{E}_a V_2(\bar{a}', n+1, s''(s', q^*), a)}{\partial s''}, \tag{27}$$

$$s'' = (1-\rho)(s' - q^*(\bar{a}, n, a, s')). \tag{28}$$

I discuss each regime in detail.

By part (i) in Proposition 3, in the **high-demand regime** ($a \geq a_H$), the inventory constraint binds. When the demand shock is sufficiently high, the firm exhausts its entire stock of inventory. The corner solution $q^* = s'$ arises because the marginal revenue exceeds the option value for all $q < s'$, hence it is optimal for the firm to sell everything.

In this regime the price is determined entirely by the demand curve evaluated at $q = s'$ as shown in equation (21). Both, the demand shock and the stock of inventory, affect the price directly through the demand curve. A higher demand shock a raises the price for a given stock, since the firm is rationed by inventory and cannot expand sales to meet demand: all of the demand response falls on the price. A larger stock s lowers the price, because the firm supplies more units from its existing inventory at the prevailing demand level: the price falls with $s'^{1/\sigma}$.

The implicit markup in this regime exceeds the CES benchmark of $\frac{\sigma}{\sigma-1}$. In this region $MR > OV$ for all $q < s'$ and the marginal revenue satisfies $MR = (\sigma - 1)/\sigma \cdot p$. Rearranging, the implied markup satisfies $(\sigma - 1)/\sigma \cdot p^* > OV$, such that $p^* > \frac{\sigma}{\sigma-1} \cdot OV$. The inventory shortage during the high demand period therefore generates market power beyond what CES preferences alone would imply: the firm is in effect extracting a scarcity rent from its own depleted stock.

By part (ii) in Proposition 3, in the **moderate-demand regime** ($a_L < a < a_H$) the MR curve intersects OV on its flat portion, where the shadow cost of carry-over inventory equals the constant replacement cost $\beta(1-\delta)(1-\rho)w/\varphi$. The firm sells to the point where

the revenue from the marginal unit today exactly compensates for the cost of replenishing it tomorrow, regardless of how many units it currently holds in stock.

The independence of optimal q^* from the stock of available inventory s' is the central result of this regime. The result means that within this demand range the firm's sales decision is completely decoupled from its inventory position. A firm with twice the stock sells exactly the same quantity as one with half the stock, provided both face the same demand shock a. The excess stock is simply carried forward. This decoupling arises because on the flat portion of OV the marginal value of carry-over inventory is constant and the next-period production decision absorbs the entire difference at the same shadow price $\beta(1-\delta)(1-\rho)w/\varphi$.

The price implication is striking. As can be seen from equation (24), the price is a constant markup over the option value and is independent of both the demand shock a and the inventory level s'. All variation in demand is absorbed entirely on the quantity margin: when a rises, the firm sells more at the same price; when a falls, it sells less and carries more inventory forward at that same price. The inventory buffer acts as a perfect shock absorber for the pricing decision.

The same logic governs the response of price to the inventory level. A firm that arrives with a large stock and one that arrives nearly depleted charge identical prices in the moderate-demand regime, because both face the same replacement cost at the margin. The larger stock simply means more units sold at that price — the firm with deeper inventory operates further to the right along its demand curve at the same price point, satisfying more demand without adjusting the price.

By part (iii) in Proposition 3, in the **low-demand regime** ($a \leq a_L$) the firm finds itself holding more inventory than it expects to need in the current and the next period. Importantly, in this region, regardless of the quantity sold, the firm will arrive over-stocked into the next period and no production will take place at the margin ($s'' > s^*$). The marginal unit of inventory does not save a production cost and will land in a region of diminishing returns to inventory, such that its value is $\beta(1-\delta)(1-\rho)\cdot\frac{\partial \mathbb{E}_a V_2(\bar{a}', n+1, s''(s', q^*), a)}{\partial s''} < \beta(1-\delta)(1-\rho)w/\varphi$. Hence, in this region the option value of a unit of inventory is below the replacement cost. Since $p^* = \frac{\sigma}{\sigma-1}\cdot OV$ and OV is below the replacement cost, prices in the low-demand regime are necessarily below those in both the moderate and high-demand regimes. The firm is cheap to sell precisely because its inventory is so abundant that the shadow value of carry-over is depressed.

The response of quantity and price to the demand shock a in this regime is qualitatively different from either of the other two cases, and the contrast is instructive. In the high-demand regime, an increase in a raises the price while leaving the quantity fixed at the stock constraint (equation (20)): all of the demand response falls on the price because the

quantity margin is exhausted. In the moderate-demand regime, an increase in a raises the quantity by equation (23) while leaving the price completely unchanged (equation (24)): all of the demand response falls on the quantity margin because the option value is flat and the shadow cost of selling does not change. In the low-demand regime, both margins are active. A higher a shifts the MR curve upward, raising q^*. Simultaneously, OV is not flat — it is increasing in q as higher sales reduce s'' and raise $\frac{\partial \mathbb{E}_a V_2}{\partial s''}$ by concavity of $\mathbb{E}_a V_2$ — selling more today raises the option value and therefore also raises the price. Therefore, both q^* and p^* increase with a.

The intuition is that in this regime the firm still has room to adjust carry-over inventory, so it does not face the hard quantity ceiling of the high-demand regime; but unlike the moderate regime, adjusting carry-over changes its marginal value, so the price cannot stay fixed. The low-demand regime is therefore the only one in which demand shocks generate a simultaneous response in both quantity and price.

Similarly, the effect of an increase in inventory is the direct opposite of what holds in the other two regimes. In the high-demand regime a larger stock s' lowers the price, because the firm supplies more units from its inventory at the prevailing demand level. In the moderate-demand regime inventory is irrelevant: both q^* and p^* are independent of s'. In the low-demand regime, for a given a, a larger s' raises carry-over inventory $s'' = (1-\rho)(s'-q)$ at any fixed q, pushing s'' further above s^* and reducing $\frac{\partial \mathbb{E}_a V_2}{\partial s''}$ by concavity. This lowers OV, so the right-hand side of the first order condition (15) falls, requiring a higher q^* to restore equality with the now-lower MR. The firm therefore sells more when it arrives with a larger stock, and since lower OV also means lower p^*, it does so at a lower price.

In the low-demand regime therefore, firms with deeper buffers sell more and charge lower prices. The mechanism is that an abundant stock pushes the next-period firm deeper into the over-stocked region, further eroding the continuation value of carry-over inventory, and therefore lowering the opportunity cost of selling aggressively today. This is entirely a consequence of the diminishing-returns structure of $\mathbb{E}_a V_2$: the more inventory the firm carries forward, the less each additional unit is worth, and the more freely it sells from current stock.

4 Conclusion

This paper develops a dynamic model of firm behavior in which Bayesian demand learning and inventory management interact over the firm life cycle. The central novelty is that the model generates empirical predictions that have previously required strong theoretical assumptions, but here arise endogenously from the optimal behavior of firms managing inventory under demand uncertainty. Price stability across a wide range of demand realizations emerges as an optimal outcome without menu costs, information frictions, or customer rela-

tionship assumptions. Markup dispersion across observationally similar firms arises without departing from CES preferences, without variable demand elasticity, and without differences in productivity or market structure. Both sets of predictions follow from a single mechanism: the piecewise structure of the option value of carry-over, which partitions the demand space into regimes with qualitatively different implications for quantities, prices, and markups.

The model therefore offers parsimonious theoretical foundations for empirical patterns that the existing literature attributes to a collection of separate and unrelated frictions. By showing that inventory management under demand uncertainty is sufficient to generate these patterns jointly, the paper opens a new perspective on the sources of price rigidity and markup heterogeneity, and raises a set of quantitative and empirical questions about the role of inventory in shaping firm-level pricing behavior over the life cycle.

References

Alessandria, G., J. Kaboski, and V. Midrigan (2010a): "The Great Trade Collapse of 2008–09: An Inventory Adjustment?" *IMF Economic Review*, 58, 254–294.

——— (2010b): "Inventories, Lumpy Trade, and Large Devaluations," *American Economic Review*, 100, 2304–2339.

Arkolakis, C., T. Papageorgiou, and O. A. Timoshenko (2018): "Firm learning and growth," *Review of Economic Dynamics*, 27, 146–168.

Atkeson, A. and A. Burstein (2008): "Pricing-to-Market, Trade Costs, and International Relative Prices," *American Economic Review*, 98, 1998–2031.

De Loecker, J., J. Eeckhout, and G. Unger (2020): "The Rise of Market Power and the Macroeconomic Implications," *The Quarterly Journal of Economics*, 135, 561–644.

Golosov, M. and R. E. Lucas Jr (2007): "Menu Costs and Phillips Curves," *Journal of Political Economy*, 115, 171–199.

Jovanovic, B. (1982): "Selection and the Evolution of Industry," *Econometrica*, 50, 649–670.

Mankiw, N. G. and R. Reis (2002): "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *The Quarterly Journal of Economics*, 117, 1295–1328.

Melitz, M. J. (2003): "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71, 1695–1725.

Melitz, M. J. and G. I. P. Ottaviano (2008): "Market Size, Trade, and Productivity," *Review of Economic Studies*, 75, 295–316.

Nakamura, E. and J. Steinsson (2008): "Five Facts about Prices: A Reevaluation of Menu Cost Models," *The Quarterly Journal of Economics*, 123, 1415–1464.

Rotemberg, J. J. (2005): "Customer Anger at Price Increases, Changes in the Frequency of Price Adjustment and Monetary Policy," *Journal of Monetary Economics*, 52, 829–852.

A Proofs of Propositions

Proposition 2 *Properties of the production decision.*

Holding all else constant, for a given $(\bar{a}, n)$ *there exits* $s^*(\bar{a}, n)$ *such that*
(i) if $s < s^*(\bar{a}, n)$ *then*

$$\begin{aligned} s' &= s^*(\bar{a}, n), \\ i^*(\bar{a}, n, s) &= s^*(\bar{a}, n) - s, \text{ and} \\ V_1(\bar{a}, n, s) &= -(s^*(\bar{a}, n) - s)\frac{w}{\varphi} - wf + \mathbb{E}_{a|\bar{a},n} V_2(\bar{a}, n, s^*(\bar{a}, n), a), \end{aligned}$$

(ii) if $s \geq s^*(\bar{a}, n)$ *then*

$$\begin{aligned} s' &= s, \\ i^*(\bar{a}, n, s) &= 0, \text{ and} \\ V_1(\bar{a}, n, s) &= -wf + \mathbb{E}_{a|\bar{a},n} V_2(\bar{a}, n, s, a). \end{aligned}$$

Proof: Substitute equation (4) into equation (3) to obtain

$$V_1(\bar{a}, n, s) = \max_{i \geq 0} -i\frac{w}{\varphi} - wf + \mathbb{E}_{a|\bar{a},n} V_2(\bar{a}, n, s + i, a).$$

Take the first order conditions with respect to i to obtain

$$\frac{w}{\varphi} = \frac{\partial \mathbb{E}_{a|\bar{a},n} V_2(\bar{a}, n, s + i, a)}{\partial i}. \tag{29}$$

Denote by $x = s + i$. Hence, equation (29) can be written as

$$\frac{w}{\varphi} = \frac{\partial \mathbb{E}_{a|\bar{a},n} V_2(\bar{a}, n, x, a)}{\partial x}. \tag{30}$$

Notice that, for a given wage w and productivity φ, the solution for x in equation (30) depends only on a firm's state $(\bar{a}, n)$. Denote the solution by $x^* = s^*(\bar{a}, n)$. Hence, $s + i = s^*(\bar{a}, n)$. If $s < s^*(\bar{a}, n)$ then $i^* = s^*(\bar{a}, n) - s > 0$ and $s' = s + i = s^*(\bar{a}, n)$. If $s > s^*(\bar{a}, n)$ then a solution to equation (30) such that $i^* > 0$ does not exit. Therefore, the constraint $i \geq 0$ binds, and in this case $i^* = 0$ and $s' = s$. ■

TO BE COMPLETED