Paralyzed by Fear: Rigid and Discrete Pricing under Demand Uncertainty *

Cosmin Ilut — Duke University & NBER
Rosen Valchev — Boston College
Nicolas Vincent — HEC Montreal

February 5, 2019

Abstract

We propose a new theory of price rigidity based on firms' Knightian uncertainty about their competitive environment. This uncertainty has two key implications. First, firms learn about the shape of their demand function from past observations of quantities sold. This learning gives rise to kinks in the expected profit function at previously observed prices, making those prices both sticky and more likely to reoccur. Second, uncertainty about the relationship between aggregate and industry-level inflation ensures nominal rigidity. We prove the main insights analytically and also quantify the effects of our mechanism. Our estimated quantitative model is consistent with a wide range of micro-level pricing facts that are typically challenging to match jointly, and implies significant monetary non-neutrality.

*Email addresses: Ilut cosmin.ilut@duke.edu, Valchev: valchev@bc.edu, Vincent: nicolas.vincent@hec.ca.
We would like to thank Etienne Gagnon, Han Hong, Nir Jaimovich, John Leahy, Virgiliu Midrigan, Martin Schneider and Luminita Stevens as well as seminar and conference participants at the Workshop on "Ambiguity and Robustness in Macroeconomics and Finance", on "Ambiguity and its Implications in Finance and Macroeconomics", "Behavioral Aspects of Macroeconomics and Finance" Conference, 9th Annual CIGS conference on Macroeconomic Theory and Policy, CIREQ Montreal Macro conference, Boston Macro Juniors Workshop, Duke, ECB, ESWC, NBER Summer Institute, NBER EFG meeting, Northwestern, SED, Stanford, UC San Diego, U Penn for helpful discussions and comments.

1 Introduction

Macroeconomists have long recognized that incomplete price adjustment plays a crucial role in the amplification and propagation of macroeconomic shocks. On the one hand, there is ample evidence that inflation responds only slowly to monetary shocks (e.g. Christiano et al. (2005)). On the other, numerous studies have shown that at the micro-level prices are not just sticky, but also display a number of other characteristics that could play a crucial macro role as well, and are thus important to match (e.g. Bils and Klenow (2004)).

In this paper, we propose a new theory of price rigidity based on firms' Knightian uncertainty about the demand for their product. This uncertainty endogenously generates an *as-if* kink in expected profits, and hence a first-order cost of moving away from a previously posted price. The mechanism not only leads to price stickiness, but also explains a number of additional micro-level pricing facts, and implies significant monetary non-neutrality.

Our economy is composed of a continuum of industries, each populated with monopolistic firms who face uncertainty about their competitive environment. In order to evaluate how demand changes as a function of the nominal price they post, firms need to jointly assess (1) the unknown *demand curve*, as a function of the relevant relative price; and (2) the *relative price* itself, which equals the firm's nominal price minus the unobserved industry price index. Uncertainty about both jointly leads to nominal rigidity.

Standard models abstract from such uncertainty, typically assuming that firms know the structure of the economy and observe competitors' prices. In contrast, we assume firms face specification doubts about the model of demand. We capture such doubts by drawing on the large experimental and theoretical work motivated by Ellsberg (1961) that distinguishes between risk (uncertainty with known odds) and ambiguity, or Knightian uncertainty (unknown odds). In particular, we model the aversion to ambiguity using the recursive multiple priors preferences axiomatized by Epstein and Schneider (2003), and characterize the firm's lack of confidence through a *set* of possible prior distributions over both the unknown demand shape *and* the unknown relative price.

We aim to put on equal footing the decision maker and an econometrician who is analyzing data from an unknown data generating process.[1] To this end, we assume the firm estimates its unknown demand function from past observations of quantities sold and prices. In doing so, the firm knows demand is a smooth, downward-sloping function, but is not confident (i) that it belongs to a particular parametric family of functions, and (ii) in a unique probability measure over the space of demand functions.

Moreover, the firm has two sources of information on the unknown industry price. One

[1]This equal-footing approach addresses a general desideratum proposed in Hansen (2014).

is unambiguous: it comes from periodically conducting marketing reviews that reveal the current value of the industry price level. Second, the firm has full access to the aggregate price, but perceives the link between industry and aggregate prices as ambiguous – while the firm understands that the two indices are cointegrated in the long run, it is not confident about their short-run relationship. In particular, over short horizons, observing a change in the aggregate price level does not convince the firm that the industry price has evolved in the same way. We model this lack of confidence as a set of potential relationships, resulting in a *set* of conditional beliefs about the current industry price.[2]

The firm knows that its demand is the sum of a price-sensitive component and a temporary shock, but faces a signal extraction problem because it does not observe each separately. The firm uses the history of observations of total quantity sold at past prices, together with its set of priors, to form a set of conditional beliefs about its demand function.

In the face of ambiguity about both the demand function and the relative price, the firm optimally selects a nominal price *as if* nature draws the joint prior distribution that implies the lowest (i.e. worst-case) *conditional* expected demand. A key result is that this joint worst-case belief changes *endogenously* around the level of previously posted prices adjusted for the revision in the industry price review signals – i.e. the change is around the *unambiguous* estimate of the firm's relative price. The reason is intuitive: a price increase sets in motion a concern for a "double whammy" – that nature draws (1) the most locally-elastic demand function allowed by the prior set and (2) the largest decrease in the unobserved industry price given the relevant set of conditional beliefs. Hence, the firm fears the increase in its relative price is larger than expected *and* that demand is especially sensitive to it. The opposite concern occurs in the case of a decrease in price – the firm fears that demand is inelastic and the industry price index rose.

This endogenous switch in the worst-case scenario is at the heart of our mechanism, and results in *kinks* in expected demand and thus price rigidity.[3] An unambiguous change in the relative price would move the firm away from the safety of previously accumulated information, and therefore expose it to increased uncertainty about the shape of demand. When interacted with ambiguity about the industry price, and therefore uncertainty about the relative price achieved by a specific choice of nominal price, the rigidity becomes nominal.

The key is that the optimal choice robust to the joint uncertainty is to price as-if short-run

[2]Using the BLS' most disaggregated 130 CPI indices as well as aggregate CPI, we present evidence that an econometrician would generally have very little confidence that short-run aggregate inflation is related to industry-level inflation, even though she can be confident that the two are cointegrated in the long-run.

[3]Such endogeneity is the defining feature of the Ellsberg experiment: when the agent evaluates a bet on either a black or a white ball from the ambiguous urn, he does so *as if* the probability of drawing that ball is less than 0.5 in either case. This behavior is inconsistent with any single probability measure on the associated state space, but can be explained by the multiple-priors model.

industry inflation is *not forecastable*, and thus keep nominal prices rigid to take advantage of the perceived kinks in demand. Intuitively, a directly observed change in the industry price index would lead to an immediate adjustment in the nominal price, as such an observation has an unambiguous effect on the relative price. In contrast, the effect of aggregate inflation on the underlying industry price level is ambiguous: if the firm assumes a positive link and responds by increasing its nominal price, this would be precisely the wrong action in case the industry price actually fell, and vice versa if it were to act under the belief that the two are negatively correlated. These fears make aggregate (or other) indexation sub-optimal.

An unambiguous change in the relative price away from a previously observed value incurs an endogenous, time-varying cost in terms of expected profits, whose properties we derive analytically. First, this cost is locally first-order, so that a firm has an incentive to keep its estimated relative price constant even when hit with marginal-cost shocks. Second, conditional on changing, the firm is inclined to repeat a price it has already posted in the past – such previously estimated relative prices are associated with kinks in expected profits, and become 'reference' price points. Third, the cost is perceived to be larger for prices that have been observed more often in the past, as higher signal-to-noise ratios deepen the kinks. Fourth, given the resulting time-variation in the first-order cost, the firm may find it optimal to implement small or large price changes. Fifth, the perceived cost of changing a posted price increases with the value of the demand shock at that price. Sixth, even though firms are forward-looking, the optimal experimentation strategy may in fact reinforce stickiness.

Since the worst-case belief is that aggregate inflation is uninformative about industry prices, it follows that between review periods, the firm faces a first-order cost of *nominal* adjustment with similar properties. This results in what looks like "price plans", where the price series generally tends to bounce around a few reference prices. When a new review signal arrives, the whole price plan shifts accordingly.

In addition to the analytical results, we evaluate the model quantitatively. We solve numerically for its stochastic steady state and estimate the parameters by targeting standard micro-level pricing moments from the IRI Marketing Dataset. We then show that our learning mechanism is quantitatively consistent with a rich set of additional moments that are typically considered challenging to match *jointly*: (i) memory in prices; (ii) co-existence of small and large price changes; (iii) pricing behavior over the product's life-cycle; (iv) downward-sloping hazard function of price changes; as well as a novel implication that (v) a price with a positive demand innovation is less likely to change.[4]

[4]Given the importance of controlling for unobserved heterogeneity in recovering the hazard function facts, and the novelty of the role of demand signals for pricing decisions, our detailed documentations of these two particular conditional moments is of independent empirical interest for the pricing literature.

We conclude the paper by showing that our quantitative model predicts large and persistent real effects from an aggregate nominal spending shock. These effects occur even though the model is consistent with the observed high frequency and large median absolute size of price changes, typically taken to imply low monetary non-neutrality in standard state-dependent models due to the selection effect analyzed in Golosov and Lucas (2007). This happens because the mechanism has additional important and empirically relevant implications for price dynamics. In particular, as in the data, our model interprets a large part of the observed price flexibility as a combination of (i) both small and large price changes; (ii) outcomes purely related to learning, such as life-cycle experimentation or realized demand signals, that are essentially decoupled from the standard targeting of real markups; and (iii) movements to and from between previously established kinks. These forces weaken the selection effect and allow our model to generate a high degree of money non-neutrality.

In Section 2 we review the literature. Section 3 derives analytical results in a real model, while Section 4 expands them to a nominal model. Section 5 quantifies the mechanism.

2 Relation to literature

By connecting learning under ambiguity to the problem of a firm setting prices, this paper relates to multiple literature strands. First is the extensive body of work on *theories of real and nominal price rigidity*. With respect to the former, it relates to work on kinked demand curves, including Stigler (1947), Stiglitz (1979), Ball and Romer (1990), Kimball (1995) and Dupraz (2016). The key novelty is that while in these models the kinks are a feature of the true demand curve, in our setup they arise only as a result of uncertainty about the shape of the demand, and an econometrician would not be expected to find evidence of actual kinks.

On nominal rigidity, we connect to a large and growing literature that emphasizes the role of imperfect information in generating prices that adjust slowly to aggregate nominal shocks, including early work such as Mankiw and Reis (2002), Sims (2003), Woodford (2003), Reis (2006) and Mackowiak and Wiederholt (2009). However, in order to generate prices that are constant over multiple periods, as observed in the data, these models typically require additional frictions, for example in the form of a menu cost or costly memory of calendar time.[5] In contrast, we show that imperfect information alone can lead to inaction.

We follow the spirit of a broad literature that documents stylized facts aimed at disci-

[5]Bonomo and Carvalho (2004) and Knotek and Edward (2010) are early examples of merging information frictions with a physical cost or an exogenous probability of price adjustment. Recent models of rational inattention, such as Woodford (2009) or Stevens (2014), assume that memory, including the passage of time, is subject to costly processing. Therefore, in periods when the firm is inattentive, it does not index to aggregate inflation as this requires paying costly attention to calendar time.

plining theoretical models of rigidity, such as Bils and Klenow (2004), Klenow and Kryvtsov (2008), Nakamura and Steinsson (2008), Klenow and Malin (2010) and Campbell and Eden (2014), among many. By testing our mechanism against a set of overidentifying restrictions, we connect to several subsets of the literature on *theoretical and quantitative* pricing models.

First, in our model, prices tend to return to previously observed values, giving rise to discreteness and memory in prices. As such, we relate to the work of Eichenbaum et al. (2011), Kehoe and Midrigan (2015) and Stevens (2014) who find ubiquitous evidence of 'reference prices' in micro price data and highlight the importance of this empirical regularity for macroeconomic fluctuations. In particular, they show that the degree of aggregate nominal rigidity is not only a function of the unconditional price flexibility, but also of how likely prices tend to return to previous 'reference' values, conditional on a change.

The prediction that prices have memory constitutes a challenge to standard models because typical state-dependent pricing theories rely on a fixed cost of change that does not depend on the specific location of the price. To address this challenge, Eichenbaum et al. (2011) assume a price-specific cost of changing it through an exogenously defined price plan, while Kehoe and Midrigan (2015) build a model with heterogeneous menu costs where it is more expensive to change the regular price than the sale price. A recent complementary approach based on rational inattention emphasizes the discrete nature of optimal signal structures, as in Matějka (2015) and Stevens (2014).[6] In these 'reference price' models, the resulting heterogeneous costs of changing associated with different prices lead not only to memory, but also to some other features shared with our framework. In particular, Matějka (2015) analyzes the theoretical prediction of a decreasing hazard function, while in the quantitative model of Stevens (2014) both small and large price changes may arise.

While there are many parallels to be drawn between our model and those of the 'reference price' family, our mechanism is distinct in two fundamental ways. The first difference lies in its novel micro-foundation, that is consistent with the four stylized facts discussed above (i.e. rigidity, memory, decreasing hazard and small price changes) even in the absence of discrete signal structures or technological differences in menu costs across prices. Second, our mechanism suggests further testable implications that speak to the role of information accumulation about demand. One prediction is that the frequency and size of price changes fall with the product's age, consistent with the evidence produced by Argente and Yeh (2017). Another is that the firm is more reluctant to change a price that experienced a positive demand shock. We find empirical support for this implication in the micro data,

[6] Given restrictions on the objective function and the prior uncertainty, that work studies how the firm may choose a discrete price distribution to economize on the costs of acquiring information about the unobserved states. We show that even when the firm conditions on signals drawn from standard continuous distributions, a lack of confidence in the firm's model of the world, and hence priors, leads to inaction and discreteness.

providing in the process a novel moment restriction to the literature on price setting.

The second set of related models is on pricing under demand uncertainty. The standard approach has been to analyze learning about a parametric demand curve under expected utility.[7] Unlike our environment, learning about parametric functions, such as linear demand curves, does not result in kinks in conditional beliefs, unless the assumed function is itself non-differentiable. Indeed, the objective in introducing learning in existing quantitative models has not been to generate nominal price stickiness or memory, but instead to match some additional facts, such as the hazard function shape (Bachmann and Moscarini (2011), Baley and Blanco (2018)) or the pricing behavior over the life-cycle (Argente and Yeh (2017)).

A third subset of related quantitative work extends the standard menu cost model to match some of the stylized facts that our model speaks to. Often, an extension is offered to address a few specific restrictions from the data. For example, Midrigan (2011) documents the co-existence of small and large price changes and explains it with economies of scope in a multi-product menu cost model. Alvarez et al. (2011) studies how a menu cost model with observation costs can generate both that empirical fact and a potentially decreasing hazard. Our work differs in the theoretical mechanism behind price stickiness in general, but is also simultaneously consistent with other important empirical observations, such as memory.

At its core, our theoretical framework fits within a large literature motivated by the classic work of Ellsberg (1961), and as such we build on previous contributions that include Gilboa and Schmeidler (1989), Dow and Werlang (1992), and Epstein and Schneider (2003). In this context, we are related to work in the industrial organization literature on ambiguity over demand. For example, Bergemann and Schlag (2011) studies a static pricing problem where the firm has multiple priors over the distribution of buyers' valuations, while Handel and Misra (2015) extends that analysis to a two-period model with maxmin regret that allows for various forms of consumer heterogeneity. In related work, Handel et al. (2013) uses a static model to inject ambiguity into a standard panel data discrete choice framework in order to incorporate partially identified preferences. Compared to this literature, we simplify the consumer's side of the market and instead develop a tractable learning environment to study how the accumulation of information about a non-parametric set of demand curves leads to pricing behavior that is empirically supported and of interest for macroeconomic models.

3 Analytical Model

In this section, we lay out and analyze the key mechanism in a smaller, analytically-tractable real model. We present the full nominal model in Section 4.

[7] An early contribution is Rothschild (1974), who frames the learning process as a two-arm bandit problem.

We study a monopolistic firm that each period sells a single good, facing a log demand:

$$q(p_t) = x(p_t) + z_t, \tag{1}$$

where p_t is the log of the real posted price. Demand consists of two components, the price-sensitive and price-insensitive parts, denoted by $x(p_t)$ and z_t, respectively. The firm's time-t realized profit is:

$$\upsilon_t = \left(e^{p_t} - e^{c_t}\right) e^{q(p_t)}, \tag{2}$$

where we have assumed a linear cost function, with c_t denoting the time-t log marginal cost.

The decomposition of demand in (1) serves two purposes. First, it generates a motive for signal extraction. In this respect we assume that the firm only observes total quantity sold, $q(p_t)$, but not the underlying $x(p_t)$ and z_t separately. Furthermore, we model z_t as iid, and thus past demand realizations $q(p_t)$ are noisy signals about the unknown function $x(p)$.

The second purpose is to differentiate between risk and ambiguity. We model z_t as purely risky, and give the firm full confidence that it is iid and drawn from the known Gaussian distribution $z_t \sim N(0, \sigma_z^2)$. On the other hand, the $x(p_t)$ component is ambiguous, meaning that the firm is not fully confident in the distribution from which it has been drawn and does not have a unique prior over it. Instead, the firm entertains a whole *set* of possible priors, Υ_0, which is not restricted to a given parametric family.

Each individual prior in the set Υ_0 is a Gaussian Process distribution, $GP(m(p), K(p, p_t))$, with mean function $m(p)$ and covariance function $K(p, p_t)$. A Gaussian Process distribution is the generalization of the Gaussian distribution to infinite-sized collections of real-valued random variables, and is a convenient choice of a prior for doing Bayesian inference on function spaces. It has the defining feature that any finite sub-collection of random variables has a multivariate Gaussian distribution.[8] Thus, for any finite vector of prices $\mathbf{p} = [p_1, ..., p_N]'$, the corresponding vector of demands $x(\mathbf{p})$ is distributed as

$$x(\mathbf{p}) \sim N\left(\begin{bmatrix} m(p_1) \\ \vdots \\ m(p_N) \end{bmatrix}, \begin{bmatrix} K(p_1, p_1) & \dots & K(p_1, p_N) \\ \vdots & \ddots & \vdots \\ K(p_N, p_1) & \dots & K(p_N, p_N) \end{bmatrix}\right),$$

where the mean function $m(\cdot)$ controls the average slope of the underlying functions $x(p)$, and the covariance function $K(\cdot, \cdot)$ controls their smoothness. In other words, this distribution is a cloud of functions dispersed around $m(p)$, according to the covariance function $K(\cdot, \cdot)$.

[8] Intuitively, we can think of a function as an infinite collection of variables, and the GP distribution defines a measure over such infinite-length random vectors by defining the distribution of any finite sub-collection.

We model ambiguity by assuming that all priors have the same covariance function, but different mean functions. We assume that the covariance function is of the widely-used squared exponential class (see Rasmussen and Williams (2006)):

$$K(p, p_t) = \text{Cov}(x(p), x(p_t)) = \sigma_x^2 e^{-\psi(p-p_t)^2}. \quad (3)$$

The function has two parameters: σ_x^2 measures the prior variance about demand at any given price, and $\psi > 0$ controls the extent to which information about demand at some price p is informative about its value at a different price p_t. The larger is ψ, the faster the correlation between quantity demanded at different prices declines with the distance between those prices.[9] This covariance function parsimoniously, yet flexibly, captures the natural prior view that there is an imperfect and declining correlation between demand at difference prices. Additionally, this prior puts zero probability on demand functions that are not infinitely differentiable. Thus, we have selected a very conservative class of priors for our purposes, as in this case any non-differentiability in the firm's perceptions about demand would be fully attributable to our ambiguity-aversion mechanism.

The multiple priors differ in their mean function $m(p)$. We assume that the set of entertained $m(p)$ is centered around the true DGP of a standard log-linear demand function,

$$x^{DGP}(p) = -bp, \quad (4)$$

so that the entertained $m(p)$ lie within the interval $x^{DGP}(p) \pm \gamma$,

$$m(p) \in [-\gamma - bp, \gamma - bp]. \quad (5)$$

The parameter $\gamma > 0$ gives the size of perceived ambiguity and captures the firm's lack of confidence in assigning probability assessments over the mean demand at a given price p.

In addition, to preclude any ex-ante built-in non-differentiability, we also bound the local variability of admissible $m(p)$. The firm only entertains differentiable $m(p)$ functions with a derivative that lies within an interval centered around the derivative of the true DGP,

$$m'(p) \in [-b - \delta, -b + \delta], \quad (6)$$

with $\delta > 0$ controlling the size of that interval. Throughout we assume that $\delta \leq b$, hence the firm is at least confident that demand is weakly downward-sloping.

[9] A Gaussian Process with a higher ψ has a higher rate of change (i.e. larger derivative) and its value is more likely to experience a bigger change for the same change in p. For example, it can be shown that the mean number of zero-crossings over a unit interval is given by $\frac{\psi}{\sqrt{2\pi}}$.

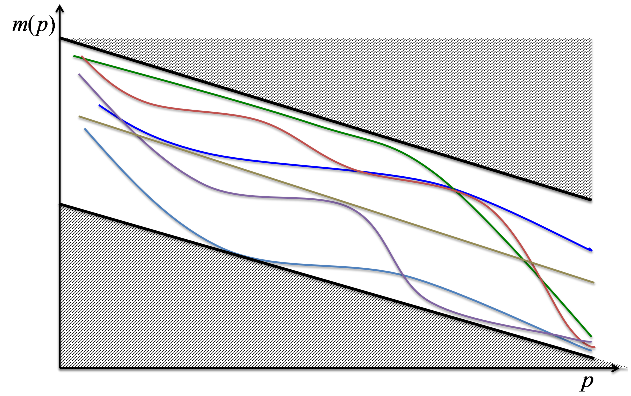

Figure 1: Illustrative set of priors for mean demand

Figure 1 provides an illustration of the set of admissible $m(p)$. The overall interpretation is that the firm has some a-priori information on the true demand, but is not confident in a single probabilistic weighting of the potential demand schedules (i.e. a single prior), nor is it able to restrict attention to a particular parametric family of demand functions. Still, while it faces uncertainty, its beliefs are centered around the truth.

The parametrization of ambiguity characterizing the sets (5) and (6) serves two purposes. First, it avoids overparameterizing Υ_0, so that we represent the ambiguity over a non-parametric family of functions using only two parameters, γ and δ. Second, it contains the minimal ingredients necessary for our main results. In particular, note that when $\gamma = 0$, the set Υ_0 collapses to a singleton, hence the firm has a unique prior and there is no ambiguity at all. On the other hand, with $\delta = 0$ the firm faces no ambiguity about the *shape* of the demand function, which is the key ingredient of our theory. The bounds on the derivative of $m(p)$ are not necessary, but they will help us show that the mechanism does not rely on discontinuities, but only on non-differentiability in the *perceived* expected demand, which we will see is the natural result of an endogenous switch in the worst-case demand elasticity.

3.1 Information and Preferences

The timing of choices and revelation of information is as follows: We assume that c_t is known at the end of $t-1$ and that it is a continuous random variable following a Markov process with a conditional density function $g(c_t|c_{t-1})$. The firm enters period t with information on the history of all previously-sold quantities $q^{t-1} = [q(p_1), ..., q(p_{t-1})]'$ and the corresponding prices at which those were observed $p^{t-1} = [p_1, ..., p_{t-1}]'$, where a superscript denotes history up to that time. It updates its beliefs about demand conditional on $\varepsilon^{t-1} = \{q^{t-1}, p^{t-1}\}$, observes c_t and posts a price p_t that maximizes its objective, specified further below. At the end of period t, the idiosyncratic demand shock z_t is realized and the firm updates its information set with the observed quantity sold $q(p_t)$ and cost c_{t+1}.

The firm uses the available data ε^{t-1} to update the set of initial priors Υ_0. Learning occurs through standard Bayesian updating, but measure-by-measure to account for the initial ambiguity. Thus, for each prior in the initial set Υ_0, the firm uses the new information and Bayes' Rule to obtain a posterior distribution. Given that there is a set of priors, the Bayesian update results in a set of posteriors. In particular, we denote by $x_{t-1}(p; m(p))$ the posterior Gaussian distribution of $x(p)$, conditional on ε^{t-1} and a particular prior $m(p)$. We denote the conditional mean and variance of demand as:

$$\widehat{x}_{t-1}(p_t; m(p)) := E\left[x(p)|\varepsilon^{t-1}; m(p)\right] \tag{7}$$

$$\widehat{\sigma}^2_{t-1}(p) := Var\left[x(p)|\varepsilon^{t-1}\right]. \tag{8}$$

While the conditional expectation depends on the prior $m(p)$, the variance is the same for all priors, as they differ only in their means. The evolution of beliefs is analytically tractable and follows the standard Bayesian-updating formulas derived in Online Appendix A.1.

The monopolistic firm is owned by an agent that is ambiguity-averse and has recursive multiple priors utility[10], so that she values the firm's profits as:

$$V\left(\varepsilon^{t-1}, c_t\right) = \max_{p_t} \min_{m(p)\in\Upsilon_0} E\left[v(\varepsilon_t, c_t) + \beta V\left(\varepsilon^t, c_{t+1}\right) \middle| \varepsilon^{t-1}, c_t \right], \tag{9}$$

where $v(\varepsilon_t, c_t)$ is the per-period profit defined in (2), a function of the beginning-of-period t price and end-of-period realized demand $q(p_t)$. The firm forms its conditional beliefs as well as evaluates the expected profits and continuation utility using the available information ε^{t-1} and the prior $m^*(p; p_t)$ that achieves the worst-case belief, given a pricing choice p_t.

Importantly, the minimization is conditional on an entertained choice of p_t. We conjecture and verify that the minimizing prior $m^*(p; p_t)$ is such that, for a given price p_t and history ε^{t-1}, it implies the lowest admissible expected demand $\widehat{x}_{t-1}(p_t; m^*(p; p_t))$ at that price p_t. Thus, for any price p_t, the firm worries that the underlying demand is low, given the data it has seen, and hence maximizes over p_t under the worst-case belief $\widehat{x}_{t-1}(p_t; m^*(p; p_t))$.

3.2 *As-if* kinks in demand from learning

To gain intuition on how updating and the basic mechanism work, we start by considering the simplest case, where the information set ε^{t-1} contains only observations of demand at a single price point p_0 that has been seen N_0 times, and has an associated average demand

[10]Epstein and Schneider (2003) develop axiomatic foundations for the recursive multiple priors utility.

realization $q_0 = x(p_0) + \frac{1}{N_0}\sum_{i=1}^{N_0} z_i$. For a given prior $m(p)$, the joint distribution of the signal and the unknown demand function $x(.)$ at any price p is:

$$\begin{bmatrix} x(p) \\ q_0 \end{bmatrix} \sim N\left(\begin{bmatrix} m(p) \\ m(p_0) \end{bmatrix}, \begin{bmatrix} \sigma_x^2 & \sigma_x^2 e^{-\psi(p-p_0)^2} \\ \sigma_x^2 e^{-\psi(p-p_0)^2} & \sigma_x^2 + \sigma_z^2/N_0 \end{bmatrix} \right).$$

The distribution of $x(p)$ conditional on q_0 is also Gaussian, and its expectation and variance are given by the familiar prior plus signal-updating formulas:

$$E(x(p)|q_0, m(p)) = m(p) + \alpha_{t-1}(p)\,[q_0 - m(p_0)] \tag{10}$$

$$Var(x(p)|q_0) = \sigma_x^2(1 - \alpha_{t-1}(p)), \tag{11}$$

where the signal-to-noise ratio used to update beliefs of demand at a given price p is

$$\alpha_{t-1}(p) = \frac{\sigma_x^2}{\sigma_x^2 + \sigma_z^2/N_0} e^{-\psi(p-p_0)^2}. \tag{12}$$

Thus, the Bayesian update of the conditional expectation in equation (10) combines the prior for demand at that price, $m(p)$, with the information revealed by the difference between the observed signal q_0, and the prior expected demand at that price, $m(p_0)$. Also note that with $\psi > 0$, the signal-to-noise ratio $\alpha_{t-1}(p)$ and the resulting reduction in uncertainty is largest right at the observed price p_0: as the correlation of quantity demanded at different prices decreases with the distance between them, the information obtained from the signal at p_0 is most useful in updating the firm's beliefs about demand around that price.

Worst-case prior

The firm minimizes the conditional expectation of demand over the priors $m(p) \in \Upsilon_0$. The resulting worst-case prior $m^*(p; p_t)$ depends on the price p_t at which the firm computes its expected demand. From equation (10) we see that the conditional expectation of demand at $p_t = p_0$ is decreasing in $m(p_0)$, since $\alpha_{t-1}(p) \in (0, 1)$. Hence, the worst-case belief corresponds to the prior with the lowest value of $m(p_0)$, so $m^*(p_0; p_t) = -\gamma - bp_0$ by equation (5).

When updating demand at a price $p_t \neq p_0$, the firm minimizes over $m(p_t)$ *and* $m(p_0)$, as both appear in the updating equation. It is useful to re-write equation (10) as

$$E(x(p_t)|q_0, m(p)) = \underbrace{(1 - \alpha_{t-1}(p_t))m(p_t)}_{\text{Prior demand at } p_t} + \underbrace{\alpha_{t-1}(p_t)(q_0 + m(p_t) - m(p_0))}_{\text{Signal at } p_0 + \Delta \text{ in Demand between } p_t \text{ and } p_0},$$

since it makes clear that uncertainty over the prior $m(p)$ affects both the overall level of expected demand (through the first term), and how the firm interprets its signal q_0 (second term). The uncertainty about the shape of the demand function implies a lack of confidence in how information about demand at p_0 translates into information about the quantity demanded at p_t. Consequently, the firm minimizes over both the prior level of demand at p_t *and* its likely change between p_t and p_0, the position of the observed signal.

First, minimizing over the prior at the entertained price, $m(p_t)$, is straightforward – the worst-case is that it lies at the lower bound of the set Υ_0, so that

$$m^*(p_t; p_t) = -\gamma - bp_t. \tag{13}$$

Second, and crucially, the firm is worried that demand falls as it changes its price from p_0 to p_t, so that the signal q_0 is bad news for demand at p_t. Since the worst-case $m(p_t)$ is at the lower bound of Υ_0, the worst-case for $m(p_0)$ is to be as high as possible given the constraints on the level and derivatives of the admissible $m(p)$. Given those constraints, the implied fall in demand does not have to be drastic; what is required is a switch in the worst-cast demand shape, depending on whether the firm considers a price increase or a decrease.

Hence, conditional on a price increase, $p_t > p_0$, the worst-case is that demand is elastic, since this generates a larger drop in demand. The drop from $m(p_0)$ to $m(p_t)$ is disciplined by the constraints on Υ_0, which restrict both the derivative (thus the highest local elasticity) of $m(p)$ at any price p, and the maximal level of $m(p_0)$. Therefore, the worst-case prior is

$$m^*(p_0; p_t) = \min\left[\gamma - bp_0, -\gamma - bp_t + (b + \delta)(p_t - p_0)\right]. \tag{14}$$

On the other hand, when the firm considers a price cut, $p_t < p_0$, it worries that demand is inelastic and that the price decrease generates as small of an increase in demand as possible. The worst-case is again restricted by the constraints on Υ_0, but in this case, the relevant derivative restriction is the lower bound in (6), as the firm worries demand is flat to the left of p_0. Hence, the worst-case $m^*(p_0; p_t)$ is now

$$m^*(p_0; p_t) = \min\left[\gamma - bp_0, -\gamma - bp_t + (b - \delta)(p_t - p_0)\right]. \tag{15}$$

Worst-case conditional expectation and kinks

Having characterized the worst-case prior, we can now plug it in equation (10) to obtain the worst-case conditional expectation at any entertained price p_t. Since the worst-case prior changes depending on whether p_t is above or below p_0, the conditional expectation

$\widehat{x}^*_{t-1}(p_t) \equiv E(x(p_t)|q_0, m^*(p; p_t))$ equals the following piecewise function

$$\widehat{x}^*_{t-1}(p_t) = \begin{cases} -\gamma - bp_t + \alpha_{t-1}(p_t)[q_0 - (-\gamma - bp_0)] - \alpha_{t-1}(p_t)\delta|p_t - p_0| & \text{if } p_t \in [\underline{p}, \bar{p}] \\ -\gamma - bp_t + \alpha_{t-1}(p_t)\left[q_0 - (\gamma - bp_0)\right] & \text{if } p_t \notin [\underline{p}, \bar{p}] \end{cases} \tag{16}$$

where $\underline{p} = p_0 - \frac{2\gamma}{\delta}$ and $\bar{p} = p_0 + \frac{2\gamma}{\delta}$. For prices $p_t \in [\underline{p}, \bar{p}]$, the worst-case prior demand at p_0 is obtained by moving away from $m^*(p_t; p_t) = -\gamma - bp_t$ along the steepest (flattest) possible demand curve, when p_t is higher (lower) than p_0. At the threshold prices $\underline{p}, \bar{p}$, moving along these worst-case elasticities intersects the upper bound of the set Υ_0, so the solution to the worst-case prior in equations (14) and (15) for prices p_t outside $[\underline{p}, \bar{p}]$ is given by $\gamma - bp_0$.

Thus, the multiple priors endogenously generate a kink in expected demand at p_0, as captured by the absolute value term $|p_t - p_0|$ in (16). In essence, the overall worst-case expectation is the result of splicing two different priors together – an elastic one to the right of p_0, and an inelastic one to its left – which creates a kink, even though all individual priors are differentiable. Panel (a) of Figure 2 illustrates the resulting, kinked worst-case expected demand, conditional on seeing a signal equal to the true DGP, at a single price point p_0.

Updating beliefs when ε^{t-1} contains observations at more than one price point is an extension of the discussion so far. Online Appendix A.1 describes the general formulas and an analytical approach to finding the worst-case prior. The intuition is the same as for the single observed price case: the worst-case is to set the prior at the entertained p_t equal to the lowest bound of Υ_0, and the level of prior demand at the other prices in ε^{t-1} as high as admissible, given the restrictions on Υ_0. The main difference is that because the endogenous switch in the worst-case priors now applies more generally at all previously-observed prices, the firm perceives kinks at all of them. For example, Panel (b) of Figure 2 shows the worst-case expectation when the firm has observed demand signals at two distinct prices.

The firm chooses its price to maximize expected profits under the worst-case beliefs. The problem is dynamic, as posting a price today affects not only current profits, but also next period's information set. Solving the full infinite horizon optimization problem is difficult numerically, because the size of the state space is unbounded, and explodes as the number of posted prices increases over time. For this reason, we split our analysis in three parts: In Section 3.3, we analyze a myopic problem that ignores the continuation value of information, but provides a tight analytical characterization of the first-order forces at play. Then in Section 3.4 we provide analytical results for a tractable approximation to the forward-looking problem, before numerically analyzing it extensively in Section 5.

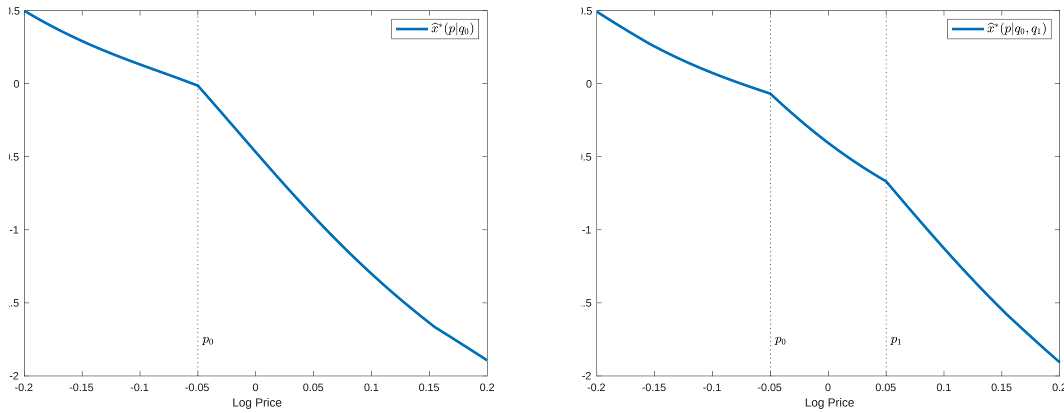

(a) One previously-observed price (b) Two previously-observed prices

Figure 2: Worst-case Expected Demand

3.3 An *as-if* cost of changing the price

A myopic firm chooses p_t to maximize time-t's worst-case expected profit

$$\max_{p_t} \min_{m(p)\in\Upsilon_0} E\left[\upsilon(\varepsilon_t, c_t)\middle|\varepsilon^{t-1}, c_t\right] = \max_{p_t} \underbrace{\left(e^{p_t} - e^{c_t}\right) e^{\widehat{x}_{t-1}(p_t; m^*(p;p_t)) + .5\widehat{\sigma}^2_{t-1}(p_t) + .5\sigma_z^2}}_{=\nu^*(\varepsilon^{t-1}, c_t, p_t)}. \qquad (17)$$

The optimal behavior crucially hinges on the history of observations ε^{t-1}, which is an endogenous object, as it depends on the past actions of the firm. In order to describe analytically the key mechanics of the model, in this section we take ε^{t-1} *as given*. We leave to Section 5 the analysis of the model when ε^{t-1} is endogenous.

We start with the simplest case as of the previous section: ε^{t-1} contains only a single price p_0, observed for N_0 number of times with an average quantity sold of q_0. As showed before, the key implication is a kink in the *as-if* expected demand. We now study the firm's incentives to keep its price rigid when faced with variations in costs. To do so, we perform a log-linear approximation of expected profits around p_0, evaluated at any given c_t, and describe a local first-order loss when moving away from p_0, as formalized in Proposition 1.

Proposition 1. *Define $\delta^* = \delta\,\text{sgn}\,(p_t - p_0)$. For a given realization of c_t, the difference in worst-case expected profits at p_t and p_0, up to a first-order approximation around p_0, is*

$$\ln \upsilon^*(\varepsilon^{t-1}, c_t, p_t) - \ln \upsilon_0^*(\varepsilon^{t-1}, c_t, p_0) \approx \left[\frac{e^{p_0}}{e^{p_0} - e^{c_t}} - (b + \alpha_{t-1}(p_0)\delta^*)\right](p_t - p_0)\,.$$

Proof. The sign switch in δ^* follows from the worst-case expected demand in (16). Also, the marginal effect $\frac{\partial \alpha_{t-1}(p_t)}{\partial p_t} = 0$ at $p_t = p_0$. For details, see Online Appendix A.2. □

Proposition 1 shows the locally-evaluated tradeoff of moving the price away from p_0. The first term in the squared brackets is the direct effect of a change in price, holding demand constant. The second term is the demand effect of a price change, by moving along the perceived demand elasticity. The fact that the elasticity switches by $\alpha_{t-1}(p_0)\delta^*$ around p_0, as indicated by the signum function, is the key mechanism in our model. We now describe the main results that stem from this property.

Result #1: *There exists an inaction region around previously-posted prices*

Given the first-order loss arising from the switch in elasticity, an implication of Proposition 1, is that, as derived explicitly in Corollary 1, there is a positive interval of c_t realizations, around $\overline{c}_0 = p_0 - \ln\left(\frac{b}{b-1}\right)$, for which the firm prefers to keep its current price fixed at $p_t = p_0$.

Corollary 1. *Under the first-order approximation in Proposition 1, p_0 is a local maximizer for any c_t in the interval $(\underline{c}_{t-1,0}, \overline{c}_{t-1,0})$, where $\underline{c}_{t-1,0} = \overline{c}_0 + \ln\left[\frac{b}{b-1}\frac{b-\alpha_{t-1}(p_0)\delta-1}{b-\alpha_{t-1}(p_0)\delta}\right]$ and $\overline{c}_{t-1,0} = \overline{c}_0 + \ln\left[\frac{b}{b-1}\frac{b+\alpha_{t-1}(p_0)\delta-1}{b+\alpha_{t-1}(p_0)\delta}\right]$.*

Proof. For any $c_t \in (\underline{c}_{t-1,0}, \overline{c}_{t-1,0})$ we have $\frac{e^{p_0}}{e^{p_0}-e^{c_t}} \in (b - \alpha_{t-1}(p_0)\delta, b + \alpha_{t-1}(p_0)\delta)$. Thus, the derivative in Proposition 1 is negative for $p_t > p_0$, when $\delta^* = \delta$, and positive for $p_t < p_0$, when $\delta^* = -\delta$. This gives the necessary and sufficient conditions for p_0 to be a local maximizer. □

To gain intuition, consider an increase in cost to some $c_t > \overline{c}_0$. This lowers the markup when keeping the price constant at p_0, which gives the firm a reason to consider an increase in the price. However, when the firm entertains a higher price $p_t > p_0$, it perceives a discrete increase in demand elasticity to $b + \alpha_{t-1}(p_0)\delta$, which lowers the target markup. As long as costs do not increase too much, so that $c_t \leq \bar{c}_{t-1,0}$, the implied markup at p_0 is in fact still higher than the new target markup. Hence, the firm finds it optimal to keep its price fixed and let the markup decline. If the cost eventually moves higher than that threshold, the fall in markup is too big, and this induces the firm to change its price.

The logic is similar for a decrease in cost from $\overline{c}_0$. As the firm entertains lowering its price from p_0, it perceives the discretely-flatter elasticity $b - \alpha_{t-1}(p_0)\delta$. Facing this decrease in elasticity, the firm finds it optimal to keep its price fixed and let the markup increase until c_t falls to the lower bound $\underline{c}_{t-1,0}$. Only for a cost realization below this threshold is the implied increase in markup big enough to incentivize the firm to lower its price and move along the flatter demand curve it perceives below p_0.

Proposition 1 implies that rigidity arises if and only if there is ambiguity about the demand *shape*. If that is not the case, i.e. $\delta = 0$, the interval of costs for which p_0 is the local optimizer is the singleton set $\{\bar{c}_0\}$. Since c_t is a continuous random variable, in this case the probability that p_0 is a local maximizer becomes zero. Instead, with ambiguity, this probability is strictly positive, and can be computed using the density function $g(c_t|c_{t-1})$ as:

$$\Pr\left[p_t^* = p_0 | p_0, N_0, q_0\right] = \int_{\underline{c}_{t-1,0}}^{\bar{c}_{t-1,0}} g(c_t|c_{t-1}) dc_t > 0, \tag{18}$$

where p_t^* is the optimal choice for p_t under the approximation in Proposition 1.

Unlike a fixed cost of changing the price, the *as-if* first-order perceived cost that emerges in our model is history-dependent. There are two fundamental dimensions along which past information matters for this perception, which we now turn our attention to.

Result #2: ***The inaction region widens as a price gets observed more often***

The first dimension is that the perceived demand loss of changing the price increases with the signal-to-noise ratio $\alpha_{t-1}(p_0)$ (see Proposition 1). Intuitively, increasing the precision of information available at p_0 makes the firm more confident in its estimate of $x(p_0)$, effectively amplifying the perceived increase in uncertainty when moving away from p_0. This translates in a larger difference between the worst-case demand elasticities on either side of p_0, which in turn raises the first-order loss of changing prices. Since $\alpha_{t-1}(p_0)$ increases with N_0, by equation (12), it follows that, holding everything else constant, having seen the price p_0 more often in the past leads to a larger inaction region, as summarized in Corollary 2.

Corollary 2. *The interval, defined in Corollary 1, of cost shock realizations c_t for which p_0 is a local maximizer widens with N_0 :*

$$\frac{\partial \underline{c}_{t-1,0}}{\partial N_0} < 0; \ \frac{\partial \bar{c}_{t-1,0}}{\partial N_0} > 0$$

Proof. Follows from Corollary 1 and from $\frac{\partial \alpha_{t-1}(p_t)}{\partial N_0} > 0$ in equation (12). □

A larger inaction region makes the probability of p_0 being a local maximizer, conditional on a history in which p_0 has been posted more often in the past $N_0' > N_0$, strictly larger:

$$\Pr\left[p_t^* = p_0 | \{p_0, N_0', q_0\}\right] > \Pr\left[p_t^* = p_0 | \{p_0, N_0, q_0\}\right].$$

Result #3: ***Prices display memory***

Another crucial property of history dependence is that when past information ε^{t-1} contains more than one unique price point, the general updating formulas discussed in

Section 3.2 imply that there exist kinks in the *as-if* expected demand around each previously observed price point $p_i \in \varepsilon^{t-1}$. These kinks lead to qualitatively similar first-order losses in the expected profit around all such prices. This result is formalized in Proposition 2.

Proposition 2. *Define* $\delta_i^* = \delta \operatorname{sgn}(p_t - p_i)$ *for all* $p_i \in \varepsilon^{t-1}$*. For a given realization of* c_t*, up to a first-order approximation around each such* $p_i \in \varepsilon^{t-1}$*:*

$$\ln \upsilon^*(\varepsilon^{t-1}, c_t, p_t) - \ln \upsilon_0^*(\varepsilon^{t-1}, c_t, p_i) \approx \left[\frac{e^{p_i}}{e^{p_i} - e^{c_t}} - (b + \alpha_{t-1}(p_i)\delta^*)\right](p_t - p_i).$$

Proof. The switch in δ^* follows from the construction of the worst-case expected demand detailed in Online Appendix A.1. □

Letting $\overline{c}_i = p_i - \ln\left(\frac{b}{b-1}\right)$ for all $p_i \in \varepsilon^{t-1}$, a direct counterpart to Corollary 1 follows.

Corollary 3. *Under the first-order approximation in Proposition 2, for each* $p_i \in \varepsilon^{t-1}$ *there exists the interval* $(\underline{c}_{t-1,i}, \overline{c}_{t-1,i})$*, where* $\underline{c}_{t-1,i} = \overline{c}_i + \ln\left[\frac{b}{b-1}\frac{b-\alpha_{t-1}(p_i)\delta-1}{b-\alpha_{t-1}(p_i)\delta}\right]$ *and* $\overline{c}_{t-1,i} = \overline{c}_i + \ln\left[\frac{b}{b-1}\frac{b+\alpha_{t-1}(p_i)\delta-1}{b+\alpha_{t-1}(p_i)\delta}\right]$*, such that for all* $c_t \in (\underline{c}_{t-1,i}, \overline{c}_{t-1,i})$ p_i *is a local maximizer.*

Proof. For any c_t in this interval, the first order derivative of the change in profits in Proposition 2 is negative for $p_t > p_i$ and positive for $p_t < p_i$, for all $p_i \in \varepsilon^{t-1}$. □

Propositions 1 and 2 imply that the firm is not only reluctant to change its current price, but is in general inclined to repeat a price it has already seen in the recent past, a form of 'memory' in prices. These previously observed past prices, at which there are kinks in the profit function, become 'reference' prices.

So far the analysis has been local, focusing on the first-order effect of price deviations around any of the $p_i \in \varepsilon^{t-1}$ and showing that the corner solution of keeping the price fixed is a local optimum. However, to find the global optimum of equation (17), we need to compare the worst-case expected profits at all such local optima against the interior optimum price.

This comparison involves a novel force arising from the intrinsic non-linearity of the expected demand in equation (16). In our setting, the signal-to-noise ratio $\alpha_{t-1}(p)$ declines with the distance between p and p_i. Intuitively, because the levels of demand at different prices are imperfectly correlated, the information about demand at a price p is most useful for updating beliefs at prices in its neighborhood. This naturally arises from the fact that demand does not come from a particular parametric family – when learning non-parametrically, information is inherently local, as it does not update beliefs about parameters that control the underlying function globally. The non-linearity of α is of second-order locally, but matters when thinking about the global maximum.

Let us consider the second derivative of the worst-case expected demand, derived in Proposition 3, where we define by $\widehat{z}_0 = q_0 - (-\gamma - bp_0)$ the perceived innovation at p_0.

Proposition 3. *The second partial derivative of $\widehat{x}_{t-1}(p_t; m^*(p; p_t))$ with respect to p_t is*

$$2\psi\alpha_{t-1}(p_t)\left[-\widehat{z}_0 + 2\delta(p_t - p_0)\right]\operatorname{sgn}(p_t - p_0).$$

Proof. Follows from the price derivatives in equations (12) and (16). □

There are two distinct economic forces at play, corresponding to the two terms in the squared brackets above. The first force, the δ term, simply reflects that as more distant prices are entertained, the concern that demand has changed for the worse from p_0 to p_t becomes less important, as the signal at p_0 loses informational content at far away prices.

Result #4: *Good demand signals make a price change less likely*

For the second effect, we formalize it in its own separate Corollary 4, as it forms the basis for a specific model implication that we test in the data.

Corollary 4. *The derivative of the worst-case expected demand to the right (left) of p_0 becomes more negative (positive) as the perceived innovation $\widehat{z}_0$ increases.*

Proposition 3 shows that the effects of the endogenous switch in the worst-case demand elasticity are amplified by the local effect of the perceived innovation $\widehat{z}_0$. Intuitively, positive demand news $\widehat{z}_0 > 0$ shift up the conditional belief about demand at all prices, but the update has a weight, $\alpha_{t-1}(p_t)$, that decreases with $|p_t - p_0|$. Thus, following positive news, beliefs about demand in the neighborhood of p_0 shift up the most and expected demand becomes effectively steeper (flatter) for prices to the right (left) of p_0.

3.4 Incorporating forward-looking behavior

Next, we consider how forward-looking behavior affects optimal pricing, and stickiness in particular. The current price choice and demand realization become state variables in next period's problem, as they get incorporated in the future information set ε^t. This gives rise to a new incentive: posting a price for the sake of obtaining new information.[11]

To characterize this exploration motive, we need to analyze the continuation value in (9). This presents a technical problem – the relevant state ε^{t-1} is the whole history of prices and demand realizations, which is infinitely long, thus making the general form of the dynamic

[11] Conceptually, our environment is related to the multi-arm bandit literature. Here the payoffs of the arms (i.e. price choices) are correlated since $\psi > 0$, and evaluated under multiple priors. See Bergemann and Valimaki (2008) for a survey of related applications of bandit problems studied under expected utility.

problem intractable. To get around this, we assume the firm understands that its action today (time t) will change its information set in the future, but thinks that none of its future pricing decisions ($t+k$) will affect its information set again – that is, $\varepsilon^{t+k} = \varepsilon^t$,$\forall k \geq 1$. We denote the resulting continuation value of the recursive problem from $t+1$ onward, when the firm does not face any more changes in the endogenous state ε^t, but still faces the fluctuations in exogenous cost process c_{t+k}, as $\tilde{V}(.)$.[12] Plugging it into (9), the firm solves

$$V(\varepsilon^{t-1}, c_t) = \max_{p_t} \min_{m(p)\in\Upsilon_0} E\left[\nu(\varepsilon_t, c_t) + \beta \int \tilde{V}(\varepsilon^t, c_{t+1}) g(c_{t+1}|c_t) dc_{t+1} \middle| \varepsilon^{t-1}\right] \tag{19}$$

This approximation makes the dynamic problem tractable, while featuring two important conceptual advantages. First, the firm is forward-looking into the discounted infinite future in terms of the cost process c_{t+k}, hence does not only consider the likely cost next period as it would in a simple two-period model. Second, the approximation leaves the history ε^{t-1} completely unrestricted. Thus, it avoids any ad-hoc assumptions limiting the firms' memory, which could lead to built-in conclusions on how firms learn and the resulting pricing decisions. Instead, leaving it unrestricted allows us to evaluate in Section 5 the long-run properties of the model at its stochastic steady state, where that history is fully endogenous.

In this section, however, we will focus on analyzing the qualitative features of the economic forces shaping the exploration motive and treat the history ε^{t-1} as given.

Option value

The experimentation motive is underpinned by an important option-value effect that is central to our analysis. To see its origin, notice that the worst-case expected demand at time $t+k$ is conditional on ε^t, hence depends on the choice of p_t and realization of q_t:

$$\widehat{x}_t(p_{t+k}; m^*(p; p_{t+k})) = \widehat{x}_{t-1}(p_{t+k}; m^*(p; p_{t+k})) + \alpha_t(p_{t+k}; p_t)(\underbrace{q_t - \widehat{x}_{t-1}(p_t; m^*(p; p_{t+k}))}_{=\widehat{z}_t})$$

where $\widehat{z}_t$ is the perceived innovation in the signal q_t. Conditional on time $t-1$ information, it is a mean-zero Gaussian variable with variance $\widehat{\sigma}^2_{t-1}(p_t) + \sigma^2_z$.

The stochasticity of $\widehat{z}_t$ underpins the option value of exploration. It makes the future *expected* demand (which incorporates the new signal q_t) uncertain at time t, with variance

$$\mathrm{Var}(\widehat{x}_t(p_{t+k}; m^*(p; p_{t+k}))|\varepsilon^{t-1}, p_t) = \alpha_t(p_{t+k}; p_t)^2(\widehat{\sigma}^2_{t-1}(p_t) + \sigma^2_z).$$

[12] $\tilde{V}(.)$ is the solution to the following recursive problem, with details presented in Online Appendix A.3

$$\tilde{V}(\varepsilon^t, c_{t+1}) = \max_{p_{t+1}} \min_{m(p)\in\Upsilon_0} E\left[\nu(\varepsilon_{t+1}, c_{t+1}) + \beta \int \tilde{V}(\varepsilon^t, c_{t+2}) g(c_{t+2}|c_{t+1}) dc_{t+2} \middle| \varepsilon^t\right]$$

The firm likes this variance to be high because if the new information about demand at p_t is bad (i.e. $\widehat{z}_t < 0$), it has the option of selecting a future price p_{t+k} away from p_t, lowering $\alpha_t(p_{t+k}; p_t)$ and minimizing the effect of the bad news on future profits – this option lowers the downside of new information. On the one hand, this gives the firm an incentive to choose a p_t in an unexplored part of the demand curve where uncertainty is the highest (so that the variance of $\widehat{z}_t$ is high). On the other hand, the weight $\alpha_t(p_{t+k}; p_t)$ on the signal q_t decreases in the distance between p_t and the future choice p_{t+k}. Intuitively, the firm values *relevant* information, i.e. signals that would affect beliefs about demand near prices that are likely to be posted in the future. The balance of these two forces, together with the location of prior information determines whether the exploration incentives lead to selection of a brand new price p_t or revisiting one of the previously observed prices.

Analytical results

To gain insight, we assume i) $\psi = \infty$, so beliefs about demand at different prices are uncorrelated, i.e. $\text{Cov}(x(p), x(p')) = 0$ if $p \neq p'$; and ii) perfect foresight about future costs, i.e. $c_{t+k} = c$ for all $k \geq 1$, where we keep the cost value $c > 0$ arbitrary. Under these assumptions, we characterize the expected continuation value $E\left[\tilde{V}(\{\varepsilon^{t-1}, p_t, q_t\}, c) \middle| \varepsilon^{t-1}, p_t\right]$ as a function of p_t (the expectation is over the realizations of the new signal q_t, whose uncertainty underpins the key option value effect), and show two analytical results that illustrate how the exploration incentive could be maximized either away from or exactly at one of the previously observed prices. Thus forward-looking behavior could both counter-act or reinforce the stickiness emerging from the static maximization discussed earlier. Which effect dominates and when is a quantitative question we take up in Section 5.

The key to whether the optimal exploration strategy is to stay put or try something new is the composition of the history of observations ε^{t-1}. To illustrate, we consider two cases that would help us understand the numerical results in Section 5 where ε^{t-1} is endogenous. First, let $\varepsilon^{t-1} = \varepsilon^0$ contain demand realizations at only one distinct price level p_0. To make the point stark, we assume that the observed signal q_0 is good enough (i.e. $q_0 > -\gamma - bp_0 + \frac{\sigma_x^2}{2}$), so that when $c = \overline{c}_0 = p_0 - \ln(\frac{b}{b-1})$, p_0 is not just locally optimal (recall Corollary 1), but that it is the global static maximizer conditional on ε^0. In Proposition 4 we characterize the current price p_t that maximizes the expected continuation value when $c = \overline{c}_0$.

Proposition 4. *The expected continuation value* $E\left[\tilde{V}(\{\varepsilon^0, p_t, q_t\}, \overline{c}_0) \middle| \varepsilon^0, p_t\right]$ *achieves its maximum at*

$$p_t^* = \arg\min_p (p - p_0)^2 \;\; s.t. \;\; p \neq p_0.$$

Proof. We provide intuition in the text below, see Online Appendix A.3 for details. □

Intuitively, choosing p_t^* today ensures that the new signal q_t will be informative about a price as close as possible to the ex-ante expected optimal p_0 – this makes the new information highly relevant. As a result, if the realization $\widehat{z}_t$ at the new signal is above a threshold $\bar{z}_t(p_t^*)$, characterized in the proof, then the firm will stick with this price in the future, set $p_{t+k} = p_t^*$, and take advantage of the unexpectedly high demand at that price while remaining near its ex-ante optimal markup level. On the other hand, if the signal realization happens to be bad, the firm can safely switch back to the ex-ante optimal p_0, where the belief about demand is not affected by $\widehat{z}_t$, and still offers lower uncertainty and the preferred markup.

The reason for not picking $p_t = p_0$ is that a bad realization of the new signal erodes the ex-ante best pricing option, p_0, while the firm does not have a good fall-back alternative, as it has no observations of demand at other prices. Because of this, it is best to experiment with a brand new price, though the desire for *relevant* information keeps the firm near p_0.

Proposition 4 describes a case where the value of new information is maximized away from p_0. However, next we show that this is not a general result, but depends on whether the firm has seen one or more distinct prices in the past. In particular, let $\varepsilon^{t-1} = \varepsilon^1$ contain demand realizations at two distinct prior prices, p_0 and p_1. Also, to simplify the exposition we assume that the information received at these prices is of the same quality – demand at each price has been observed the same number of times ($N_1 = N_0$), and the observed signals, q_0 and q_1, imply equally-good news, i.e. the same perceived innovation: $\widehat{z}_0 = \widehat{z}_1 = \widehat{z}$.

Proposition 5 shows that when the previously observed demand at p_0 and p_1 has been good enough, the continuation value is maximized at p_0 for a range of cost shocks around $\bar{c}_0$. Thus, forward-looking behavior *reinforces* the static stickiness result (Corollary 1).

Proposition 5. *There is a non-singleton interval of costs* $(\underline{c}, \bar{c})$ *around* $\bar{c}_0$*, and a threshold* $\chi > 0$*, such that if* $\widehat{z} > \chi$*, then for any* $c \in (\underline{c}, \bar{c})$*:*

$$p_0 = \arg\max_{p_t} E\left[\tilde{V}(\{\varepsilon^1, p_t, q_t\}, c)\middle|\varepsilon^1\right].$$

Moreover, the threshold χ *is decreasing in* $|p_1 - p_0|$*.*

Proof. We provide intuition in the text below, see Online Appendix A.3 for details. □

The reason for this result is two-fold. First, information about demand at p_0 is the most *relevant* since that is the price expected to be optimal in the future (with perfect foresight on future costs, the firm essentially faces a static maximization problem, hence Corollary 1 applies). Second, even if the firm receives a bad new signal q_t at p_0, it has a good fall-back option as it has also accumulated information (and thus reduced uncertainty) at p_1. Thus, the firm can set $p_t = p_0$ and further reduce uncertainty about demand at the most likely future price, safe with the knowledge that it has a good alternative in case the new information is bad. The value of the fall-back option is important – in particular, the perceived innovation in the average past demand realization at p_1 must exceed a threshold χ (which we characterize in the proof). This threshold is lower when p_0 and p_1 are closer to each other, because then their implied markups are more similar, making the two price choices closer substitutes, and thus p_1 a more attractive fall-back option.

Our analytical results show that forward-looking behavior can both counteract or reinforce the previous stickiness result derived from static maximization. The resulting overall effect depends crucially on the structure of the prior history ε^{t-1}, which highlights the importance of taking into account the endogeneity of that history. To that end, Section 5 numerically analyzes the stochastic steady state of a general version of our forward-looking model, with $\psi < \infty$ and stochastic cost shocks. We find that experimentation is not only consistent with significant price stickiness, but also helps generate an empirically relevant (i) life-cycle profile of pricing behavior and (ii) size distribution of price changes.

4 Quantitative Model and Nominal Rigidity

In this section we embed our mechanism in a typical macroeconomic model with monopolistic competition in which a typical firm needs to make a joint assessment about (i) its demand curve as a function of the relevant relative price and (ii) the relative price itself. We first show analytically that this two-dimensional uncertainty gives rise to *as-if* kinks in demand in terms of *nominal* prices. Second, we quantify the ability of our mechanism to match micro-level moments and generate monetary non-neutrality.

4.1 Economic Framework

There are two layers of demand. First, a competitive final-good producer buys from a continuum of industries indexed by j and sells to a representative household. Second, each industry itself is composed of a competitive final-good producer that aggregates over a continuum of intermediate monopolistic firms indexed by i. The motivation for having a layer of demand between the intermediate good firms and aggregate demand, is to capture the broad observation that the relevant competitors' price index for a specific firm is typically not the economy-wide aggregate price index.

The representative household consumes and works according to

$$\max_{c_{t+k},L_{i,t+k}} \sum_{k=0}^{\infty} E_t \left(\beta^{t+k} \left[c_{t+k} - \int L_{i,t+k} di \right] \right)$$

where c_t denotes log consumption of the aggregate good, subject to the budget constraint

$$\int e^{p_{j,t}+c_{j,t}}dj + E_t q_{t+1} d_{t+1} = d_t + e^{p_t+w_t}\int L_{i,t}di + \int \upsilon_{i,t}di,$$

where q_{t+1} is the stochastic discount factor, d_t are state contingent claims on the aggregate shocks, $\upsilon_{i,t}$ is the profit from the monopolistic intermediaries and w_t is the log real wage. The consumption basket and the associated aggregate price index are:

$$c_t = \frac{b}{b-1}\log\left(\int e^{c_{j,t}\frac{b-1}{b}}dj\right) \quad , \qquad p_t = \frac{1}{1-b}\log\left(\int e^{p_{j,t}(1-b)}dj\right), \tag{20}$$

where $p_{j,t}$ represent the log price indices of the distinct industries.

Each industry j has a representative final-good firm that produces by aggregating over the continuum of intermediate goods i with the technology

$$e^{c_{j,t}} = f_j^{-1}\left(\int f_j(e^{c_{i,j,t}})g_j(e^{z_{i,t}})di\right), \tag{21}$$

where $z_{i,t}$ is an idiosyncratic demand shock for good i, distributed as $N(0,\sigma_z^2)$. Each industry j has potentially different production functions f_j and g_j, and price index $p_{j,t}$ such that $e^{p_{j,t}+c_{j,t}} = \int e^{p_{i,t}+c_{i,j,t}}di$, where $c_{i,j,t}$ is the log amount of good i in the final product of industry j. Solving the cost-minimization problem of the final good firm in industry j yields

$$c_{i,j,t} = \log\left[f_j'^{-1}\left(e^{p_{i,t}-p_{j,t}}\frac{f_j'\left(e^{c_{j,t}}\right)}{g_j(e^{z_{i,t}})}\right)\right] \equiv h_j\left(e^{p_{i,t}-p_{j,t}}, c_{j,t}, z_{i,t}\right). \tag{22}$$

This is a generalization of the typical CES production structure, with the familiar result that the demand of industry j for a given intermediate good i is a function of the relative log price $p_{i,t}-p_{j,t}$, overall industry output $c_{j,t}$, and demand shocks $z_{i,t}$. We denote the effective demand function h_j and note that it is a transformation of the functions f_j and g_j. Lastly, we assume that each intermediate good producer i sells to only one industry j.[13]

An intermediate-good firm produces variety i using the production function $y_{i,t} = \omega_{i,t} + a_t + \log L_{i,t}$, where $\omega_{i,t}$ and a_t are an idiosyncratic and aggregate productivity shock, respectively, and $L_{i,t}$ is hours hired by firm i. The processes for these shocks are known:

$$\omega_{i,t} = \rho_\omega\omega_{i,t-1} + \varepsilon_{i,t}^\omega;\ a_t = \rho_a a_{t-1} + \varepsilon_t^a$$

where $\varepsilon_{i,t}^\omega$ is iid $N(0,\sigma_\omega^2)$ and ε_t^a is iid $N(0,\sigma_a^2)$. Therefore, using the household's labor supply

[13]As a result, firms are indexed by both i and j. However, for ease of notation we drop the j subscript with the understanding that each firm i is unique to a given industry.

decision $w_t = c_t$, and the market clearing condition $c_t = y_t$, the real flow profit of firm i is

$$\upsilon_{i,t} = \left(e^{p_{i,t}-p_t} - e^{y_t-\omega_{i,t}}\right) e^{y_{i,t}}. \tag{23}$$

Finally, log nominal aggregate spending $s_t = p_t + c_t$ follows a random walk with drift, $s_t = \mu + s_{t-1} + \epsilon_t^s$, where ϵ_t^s is iid $N(0, \sigma_s^2)$.

4.2 Information structure

We assume that the information available to intermediate-good firms is imperfect in three ways. First, firms do not know the functional forms of the industry-level production technologies f_j and g_j, and hence the effective demand function h_j that they face. Second, firms have imperfect information on the prices and quantities of their direct competitors – we assume that they observe the relevant industry-level prices and quantities $p_{j,t}$ and $c_{j,t}$ only infrequently, and do not directly observe the individual prices of all intermediate goods. They do, however, observe the full history of their own prices and quantities, $p_{i,t}$ and $y_{i,t}$, as well as the aggregate output and price level, y_t and p_t. Lastly, firms do not observe the idiosyncratic demand shocks $z_{i,t}$, but perfectly observe the productivity shocks $\omega_{i,t}$ and a_t.

The firm does not know the specific functional form of its demand function and estimates it using its observables. Here, for tractability, we assume the firm understands that the aggregate industry demand $c_{j,t}$ and the demand shocks $z_{i,t}$ enter multiplicatively in the unknown function h_j in equation (22).[14] Moreover, firms know the structure of the aggregate consumption basket, and can use that to substitute out industry output via $c_{j,t} = y_t - b(p_{j,t} - p_t)$, and thus obtain their own, individual good demand schedules

$$y_{i,t} = h_j(p_{i,t} - p_{j,t}) + y_t - b(p_{j,t} - p_t) + z_{i,t}. \tag{24}$$

4.2.1 Ambiguity about competition

Ambiguity about the demand function h_j is modeled as in equations (5) and (6). There is a set of multiple priors, each of which is a GP distribution with mean function $m(r_i)$ so that

$$m(r_i) \in [-\gamma - br_i, \gamma - br_i];\ \ m'(r_i) \in [-b - \delta, -b + \delta], \tag{25}$$

where we define the relative price $r_{i,t} \equiv p_{i,t} - p_{j,t}$. This price is not only the relevant argument of the unknown demand function h_j, but is also not perfectly observed itself – the firm

[14] Our learning framework extends to the case of learning about demand as a function of multiple variables without conceptual differences. We make this assumption to transparently focus on the main mechanism.

observes and sets its nominal price $p_{i,t}$, but is uncertain about the industry price level $p_{j,t}$.

The firm has two sources of information on the unknown industry price $p_{j,t}$. First, it conducts infrequent marketing reviews that reveal the current period industry price $p_{j,t}$. We model the frequency of reviews by assuming that they occur with some exogenous probability λ_T. We are implicitly assuming that there are some technological constraints on the ability to perform frequent reviews (e.g. the necessary data may not be observed every period), or simply that reviews are costly and so the firm does not want to perform them frequently.[15]

Second, between reviews, the firm estimates $p_{j,t}$ based on the observed aggregate information. Since its direct competitors form a small portion of the overall economy, the firm knows that $p_{j,t} \neq p_t$, where p_t is the aggregate, fully-observable price. Still, the firm understands that industry-level and aggregate prices are cointegrated, and hence there is information in p_t about $p_{j,t}$. However, since the firm does not know the exact industrial structure (i.e. h_j), it does not know the functional form of that cointegration relationship – different industry production functions imply different structural relationships between p_t about $p_{j,t}$. Lastly, due to the ambiguity about h_j, the firm is similarly not confident in any single cointegration relationship and entertains a non-singleton set of potential relationships $\phi(.)$ such that

$$p_{j,t} = \widetilde{p}_{j,t} + \phi(p_t - \widetilde{p}_{j,t}), \tag{26}$$

where $\widetilde{p}_{j,t}$ is the most recent review signal as of time t. Thus, the uncertainty about ϕ translates into uncertainty about how to best use p_t to forecast $p_{j,t}$.

Ambiguity about the cointegrating function is modeled with the same tools as the uncertainty about the demand function h_j. We assume that the priors on ϕ are GP distributions, with mean functions that lie in a set Ω_ϕ around the true DGP $\phi(p_t - \widetilde{p}_{j,t}) = p_t - \widetilde{p}_{j,t}$. For tractability, we focus on the limiting case in which the variance function of the GP distributions over ϕ goes to zero almost everywhere. Given the resulting Dirac priors, we can simplify notation and specify the set of priors directly as a set of ϕ's.

To model the idea that firms are uncertain about the short-run relationship between industry and aggregate prices, even though the two are cointegrated, we specify that for small $|p_t - \widetilde{p}_{j,t}|$, i.e. small inflationary pressure, the function ϕ lies in the interval

$$\phi(p_t - \widetilde{p}_{j,t}) \in [-\gamma_p, \gamma_p], \text{ for } |p_t - \widetilde{p}_{j,t}| \leq \Gamma. \tag{27}$$

And since firms correctly realize that aggregate and industry inflation are cointegrated

[15] As long as reviews do not happen every period, using deterministic or state-dependent review lags would not change our analysis significantly. The modeling advantage over deterministic timing is computational: we find that stochastic review times achieve faster convergence towards the stationary distribution. The advantage over state-dependent times is tractability, as it avoids modeling a cost-benefit analysis of reviews.

in the long-run, we make the set of potential ϕ grow with $p_t - p_{j,t}$ at higher inflation levels

$$\phi(p_t-\widetilde{p}_{j,t}) \in [-\gamma_p+p_t-\widetilde{p}_{j,t}-\Gamma\,\mathrm{sgn}(p_t-\widetilde{p}_{j,t}), \gamma_p+p_t-\widetilde{p}_{j,t}-\Gamma\,\mathrm{sgn}(p_t-\widetilde{p}_{j,t})], \text{ for } |p_t-\widetilde{p}_{j,t}| \geq \Gamma.$$

Note that all admissible priors imply that the price ratio $p_{j,t} - p_t$ is stationary with probability one, but allow for potentially complex, non-linear relationships locally. In Online Appendix B.1, we document that such uncertainty about the local relationship between aggregate and industry inflation is well supported by the data. Thus, our model essentially assumes that the firms have no special advantage over real-world econometricians and cannot ex-ante eliminate the uncertainty inherent to the short-run inflation relationship.

To make a transparent comparison between our model and the usual rational expectations benchmark we assume a simple true DGP where each industry j has the same CES functions f_j and g_j in (21): $f_j(e^{c_{j,t}}) = (e^{c_{j,t}})^{\frac{b-1}{b}}$ and $g_j(e^{z_{i,t}}) = e^{z_{i,t}/b}$. These aggregators lead to a standard demand schedule $c_{j,i,t} = z_{i,t} + c_{j,t} - b\,(p_{i,t} - p_{j,t})$. Substituting out the aggregate demand for industry output $c_{j,t}$, it follows that a RE firm has full knowledge that its demand is $y_{i,t} = y_t + b(p_t - p_{i,t}) + z_{i,t}$. The resulting optimal RE nominal price takes the familiar form $p_{i,t}^{RE} = \log\frac{b}{b-1} + p_t - \omega_{i,t}$, where the aggregate price (up to a constant) is $p_t^{RE} = s_t - a_t$.

4.3 Optimization problem

The firm enters period t with knowledge of the history of all its previous quantities sold, $y_i^{t-1} = [y_{i,1}, ..., y_{i,t-1}]'$, the corresponding nominal prices at which those quantities were observed $p_i^{t-1} = [p_{i,1}, ...p_{i,t-1}]'$, and its history of industry price review signals, $\widetilde{p}_j^{t-1} = [\widetilde{p}_{j,1}, ...\widetilde{p}_{j,t-1}]'$. In addition, the firm sees the history of aggregate prices $p^{t-1} = [p_1, ...p_t]'$ and output $y^{t-1} = [y_1, ..., y_t]'$. We denote the collection of all this information by ε^{t-1}. At the start of t, the idiosyncratic productivity $\omega_{i,t}$ and aggregates (p_t, y_t) are observed. Finally, the firm's current review signal $\widetilde{p}_{j,t}$ equals $p_{j,t}$ if a new review is conducted, and $\widetilde{p}_{j,t-1}$ otherwise.

Therefore, given a history of observables ε^{t-1} and current states $s_t = \{\omega_{i,t}, p_t, y_t, \widetilde{p}_{j,t}\}$, the firm's problem is to optimize over its action, $p_{i,t}$, taking into account the ambiguity about both the demand curve and the effective relative price:

$$V\left(\varepsilon^{t-1}, s_t\right) = \max_{p_{i,t}} \min_{m(r),\phi(p_t-\widetilde{p}_{j,t})} E\left[\upsilon(\varepsilon_t, s_t) + \beta V\left(\varepsilon^t, s_{t+1}\right)\right], \tag{28}$$

where $\upsilon(\varepsilon_t, s_t)$ is real profit $\left(e^{p_{i,t}-p_t} - e^{y_t-\omega_{i,t}}\right)e^{y_{i,t}}$ and log-demand is given by

$$y_{i,t} = h_j\left[p_{i,t} - \widetilde{p}_{j,t} - \phi(p_t - \widetilde{p}_{j,t})\right] - b\phi(p_t - \widetilde{p}_{j,t}) + y_t - b(\widetilde{p}_{j,t} - p_t) + z_{i,t}, \tag{29}$$

where we have substituted out $p_{j,t}$ in the demand equation (24) by using its law of motion in (26). At the end of period t, the idiosyncratic demand shock $z_{i,t}$ is realized and the firm updates its information set with the observed realized quantity sold $y_{i,t}$.

The aggregate price level p_t affects real profits through three channels. The first is the standard effect of deflating nominal profits by p_t. The second is as a demand shifter bp_t in (29): everything else constant, a larger p_t increases demand for the industry j's final good, which in turn translates into a higher demand for firm i. The third channel is that p_t affects beliefs about $p_{j,t}$ through its impact on $\phi(p_t - \widetilde{p}_{j,t})$, but this relationship is uncertain. Finally, the firm needs to conjecture a law of motion of p_t to forecast future profits. Here we assume that there is a measure zero of ambiguity-averse firms, while the rest of the economy is populated by rational-expectations firms, hence the equilibrium p_t is given by p_t^{RE}.[16]

4.3.1 Joint worst-case beliefs

To illustrate the key insights, for the rest of this section we focus on a myopic firm born at time $t = 0$ that is in its second period of life (i.e. $t = 1$). Hence, its information set contains $y_{i,0}$, the quantity sold at just one price point, $p_{i,0}$.[17] In addition, the firm observes the history of aggregates, $\{y_0, y_1, p_0, p_1\}$, and signals on the industry price level, $\{\widetilde{p}_{j,0}, \widetilde{p}_{j,1}\}$.

The firm forms beliefs about profits at any price $p_{i,1}$ under the worst-case scenario of low demand. Abstracting from the known component $y_1 - b(\widetilde{p}_{j,1} - p_1)$, expected demand is

$$m(r_{i,1}) - b\phi(p_1 - \widetilde{p}_{j,1}) + \alpha \{y_{i,0} - [m(r_{i,0}) - b\phi(p_0 - \widetilde{p}_{j,0})]\}, \tag{30}$$

where $\alpha = \frac{\sigma_x^2}{\sigma_x^2+\sigma_z^2}$.[18] Thus, expected demand depends on the joint worst-case priors for both the demand curve, $m(.)$, and the cointegration relationship $\phi(.)$.

It is convenient to rewrite equation (30) so as to separate the effects of uncertainty on the *level* of demand at the relative price $r_{i,1}$ and the *change* in demand between $r_{i,1}$ and $r_{i,0}$:

$$(1-\alpha)\left(\underbrace{m(r_{i,1}) - b\phi(p_1 - \widetilde{p}_{j,1})}_{\text{Prior of demand at } r_{i,1}}\right) + \alpha\left\{y_{i,0} - \left[\underbrace{m(r_{i,0}) - m(r_{i,1})}_{\text{Prior on change in demand}} - b\underbrace{(\phi(p_1 - \widetilde{p}_{j,1}) - \phi(p_0 - \widetilde{p}_{j,0}))}_{\text{Perceived change in industry price}}\right]\right\}$$

The analysis of the worst-case beliefs follows the same basic logic as in the real model.

[16]This benchmark provides an upper bound for the degree of price neutrality compared to the case of a measure one of ambiguity-averse firms, as it ignores strategic complementaries in price setting.

[17]In Online Appendix B.2 we show how the conceptual analysis is extended to multiple prices, and Section 5 shows numerical results for the general case.

[18]To simplify notation and the analysis, for the rest of this section we suppress the local information effects by working with $\psi = 0$. We relax this assumption in Section 5.

However, while in the real model the relative price was uniquely determined by the firm's action, here this is not the case. The relative price $r_{i,t} = p_{i,t} - \widetilde{p}_{j,t} - \phi(p_t - \widetilde{p}_{j,t})$ has two known components, $p_{i,t}$ and $\widetilde{p}_{j,t}$, but also an ambiguous one, in the form of $\phi\left(p_t - \widetilde{p}_{j,t}\right)$. The firm takes an action robust to both sources of ambiguity, hence it minimizes (30) jointly over the prior about the demand shape, $m(.)$, and the prior of the cointegration relationship $\phi(.)$ – which affects perceptions about both $p_{j,0}$ (the industry price at the time the signal $y_{i,0}$ was realized) and $p_{j,1}$ (the industry price that determines the current effective relative price).

The worst-case scenario for $m(r_{i,1})$ is straightforward – it is equal to the lower bound of (25). Notice that substituting in this worst-case form of $m(r_{i,1})$ cancels out $b\phi(p_1 - \widetilde{p}_{j,1})$ from the first term. Thus, what is left to do is to pick the joint worst-case for the change of demand between $r_{i,0}$ and $r_{i,1}$ and the perceived change in the industry price, as implied by the observation of the aggregate price and the cointegration relationship ϕ. Once the relevant constraints on the admissible functions m and ϕ are taken into account, determining the joint worst-case reduces to solving the following minimization problem:

$$\min_{\delta' \in [-\delta,\delta]} \min_{\phi(p_t - \widetilde{p}_{j,t}) \in [-\gamma_p,\gamma_p]} -\delta' \left\{(p_{i,1} - \widetilde{p}_{j,1}) - (p_{i,0} - \widetilde{p}_{j,0}) - [\phi(p_1 - \widetilde{p}_{j,1}) - \phi(p_0 - \widetilde{p}_{j,0})]\right\}. \quad (31)$$

The joint worst-case prior beliefs depend on whether the firm considers raising $p_{i,1} - \widetilde{p}_{j,1}$ relative to $p_{i,0} - \widetilde{p}_{j,0}$. Therefore, for notational purposes, we define $\widetilde{r}_{i,t} \equiv p_{i,t} - \widetilde{p}_{j,t}$ and note that this object is the *unambiguous estimate* of the relative price at time t. This estimate is based on two observables: the firm's own nominal price and the unambiguous signal of the unknown industry price index $p_{j,t}$. When the firm entertains increasing its estimated relative price $\widetilde{r}_{i,t}$, it sets in motion a concern that its effective demand is sensitive to this action. This concern manifests itself in a joint worst-case belief that both (i) the unknown demand curve $m(.)$ is steep, i.e. $\delta^* = \delta$, and that (ii) there was a decline in the unknown price index of its competition. Hence, if $\widetilde{r}_{i,1} \geq \widetilde{r}_{i,0}$, then the minimizing priors are

$$\delta^* = \delta;\ \phi^*(p_1 - \widetilde{p}_{j,1}) = -\gamma_p;\ \phi^*(p_0 - \widetilde{p}_{j,0}) = \gamma_p. \quad (32)$$

In contrast, when the firm entertains decreasing its estimated relative price $\widetilde{r}_{i,1}$, it worries about the opposite situation, where (i) its unknown demand curve is flat (i.e. $\delta^* = -\delta$), and (ii) it is facing an increase in the unknown price index of the competition, so that if $\widetilde{r}_{i,1} \leq \widetilde{r}_{i,0}$[19] the minimizing priors instead switch to

$$\delta^* = -\delta;\ \phi(p_1 - \widetilde{p}_{j,1}) = \gamma_p;\ \phi(p_0 - \widetilde{p}_{j,0}) = -\gamma_p. \quad (33)$$

[19]When $p_{i,1} - \widetilde{p}_{j,1} = p_{i,0} - \widetilde{p}_{j,0}$ the objective in equation (31) equals $-2\delta\gamma_p$, arising from (32) or (33).

Note that the worst-case belief about the change in the industry price index $p_{j,t}$ is not always that the firm's competition has lowered prices. The reason is that the industry price affects the firm's demand in two ways: i) it determines the relevant relative price (the argument of the ambiguous component of demand $h_j(.)$), and ii) acts as a demand shifter since lower overall prices in industry j boosts demand for all firms inside the industry (through the term $-bp_{j,t}$). These two effects go in opposite directions, and which one dominates depends on the perceived elasticity of the intra-industry demand function $h_j(.)$. As a result, the worst-case belief about $\phi(.)$ depends on the firm's action $\widetilde{r}_{i,1}$, and endogenously *changes* depending on whether the firm is contemplating a price increase or a price decrease.

The key implication from equations (32) and (33) is that the joint worst-case beliefs generate a kink in expected profits at the previous unambiguous relative price estimate $\widetilde{r}_{i,0}$, an important result that we state formally in Proposition 6 below.

Proposition 6. *Let* $\delta^* = \delta \operatorname{sgn}(\widetilde{r}_{i,1} - \widetilde{r}_{i,0})$. *The joint worst-case beliefs over the demand curve and the cointegration relationship induce a worst-case conditional demand schedule in the space of the unambiguous estimate of relative prices,* $\widetilde{r}_{i,t}$, *given by*

$$(-\gamma - b\widetilde{r}_{i,1}) + \alpha \left[y_{i,0} - (-\gamma - b\widetilde{r}_{i,0}) - \delta^* (\widetilde{r}_{i,1} - \widetilde{r}_{i,0}) - 2\delta\gamma_p \right] \tag{34}$$

Proof. Follows from substituting the worst-case beliefs, given by $m^*(r_{i,1}) = -\gamma - b\widetilde{r}_{i,1}$ and the joint solution in equations (32) and (33), in the expected demand of equation (30). □

There are two important results derived by this Proposition. First, it shows that the relevant argument of the worst-case expected demand is the unambiguous relative price $\widetilde{r}_{i,t}$. Intuitively, the firm is facing an identification problem, as it is uncertain about both the argument and the shape of the demand function. Proposition 6 proves that the robust solution is to estimate the demand curve in terms of the best available, unambiguous estimate of the relative price $\widetilde{r}_{i,t}$. Second, due to the uncertainty about the local shape of the demand function, there is a kink at the previously observed $\widetilde{r}_{i,0}$.

4.4 Learning and nominal rigidity

The kink in the worst-case expected demand (eq. (34)) leads to a first-order loss of having an estimated relative price $\widetilde{r}_{i,1}$ different from $\widetilde{r}_{i,0}$, and hence a first-order cost of posting a nominal price $p_{i,1}$ away from $\widetilde{p}_{j,1} + \widetilde{r}_{i,0}$. Between reviews of the industry price level, when $\widetilde{p}_{j,1} = \widetilde{p}_{j,0}$, this leads to a rigid optimal nominal price $p^*_{i,1}$. In particular, it may stay fixed at its previous value $p_{i,0}$ even as the aggregate price changes. We center our discussion around this result, but note that similar rigidity is obtained for changes in the other state variables.

4.4.1 Nominal prices respond infrequently to the aggregate price level

Due to the kink in (34), there is an interval of the aggregate price level p_1 for which it is optimal for the firm to keep the estimated relative price fixed at $\widetilde{r}_{i,0}$. In order to implement $\widetilde{r}_{i,1} = \widetilde{r}_{i,0}$, however, the firm needs to keep its nominal price $p_{i,1}$ equal to $\widetilde{p}_{j,1} + \widetilde{r}_{i,0}$. In contrast, under full information there is only one such value of the aggregate price level: $\overline{p}_1 = \widetilde{p}_{j,1} + \widetilde{r}_{i,0} + \omega_{i,1} - \ln\left(\frac{b}{b-1}\right)$.

Proposition 7. *The nominal price* $p_{i,1} = \widetilde{p}_{j,1} + \widetilde{r}_{i,0}$ *is a local maximizer of the worst-case expected profits for any aggregate price* $p_1 \in \left(\overline{p}_1 + \ln\left(\frac{b}{b-1}\frac{b-\alpha\delta-1}{b-\alpha\delta}\right), \overline{p}_1 + \ln\left(\frac{b}{b-1}\frac{b+\alpha\delta-1}{b+\alpha\delta}\right)\right)$.

Proof. Let $\upsilon^*(\varepsilon^0, s_1, p_{i,1})$ denote the expected profit, conditional on history ε^0, state $s_1 = \{\omega_{i,1}, p_1, y_1, \widetilde{p}_{j,1}\}$, and some $p_{i,1}$, with its associated worst-case demand in eq. (34). Then for any p_1 in the interval above, the derivative of $\ln \frac{\upsilon^*(\varepsilon^0, s_1, p_{i,1})}{\ln \upsilon^*(\varepsilon^0, s_1, \widetilde{p}_{j,1}+\widetilde{r}_{i,0})}$ is negative to the right of $\widetilde{p}_{j,1} + \widetilde{r}_{i,0}$ and positive to its left, due to the kink in (34) and change in sign of $\widetilde{r}_{i,1} - \widetilde{r}_{i,0}$. □

The intuition is akin to the one we have seen before. A value of the aggregate price $p_1 > \overline{p}_1$ lowers the real markup, which leads the firm to ponder raising its nominal price. However, such a raise would increase the estimated relative price $\widetilde{r}_{i,1}$ above the previously observed $\widetilde{r}_{i,0}$, which creates a perceived increase of the demand elasticity (to $b + \alpha\delta$) and thus a lower target markup. As long as the aggregate price does not increase too much, so that $p_1 \leq \overline{p}_1$, the implied real markup at $p_{i,1} = \widetilde{p}_{j,1} + \widetilde{r}_{i,0}$ (the nominal price which keeps $\widetilde{r}_{i,1} = \widetilde{r}_{i,0}$) is still higher than the markup the firm believes it could achieve if it were to increase its estimated relative price. Hence, it decides to keep $p_{i,1} = \widetilde{p}_{j,1} + \widetilde{r}_{i,0}$ and let the markup decline. If the aggregate price moves above $\overline{p}_1$ however, then the fall in markup is too large to bear and the firm adjusts its price accordingly. A similar logic of inaction applies for a decrease in the aggregate price level.

Consider now what happens in periods when a new review of the industry price does not occur, so that $\widetilde{p}_{j,1} = \widetilde{p}_{j,0}$. In that case, Proposition 7 immediately implies that if the firm finds it optimal to take advantage of the kink in estimated relative prices, it will do so by keeping its nominal price fixed, since $\widetilde{p}_{j,1} + \widetilde{r}_{i,0} = p_{i,0}$. Thus, when the current and previous review signals are the same, the desire to maintain an estimated relative price equal to $\widetilde{r}_{i,0}$ is achieved by posting the previous nominal price: $p_{i,1} = p_{i,0}$. This makes nominal prices rigid.

This result does not hold if a new industry price review occurs and the firm observes the industry price, i.e. $\widetilde{p}_{j,1} = p_{j,1}$. Then, the perceived kink at $\widetilde{r}_{i,0}$ implies that expected demand has a kink at the nominal price $p_{i,1} = p_{i,0} - \widetilde{p}_{j,0} + p_{j,1}$. To take advantage of the kink, the firm changes its nominal price in response to the information contained in $p_{j,1} - \widetilde{p}_{j,0}$, and unless $p_{j,1} = \widetilde{p}_{j,0}$, rigidity in the estimated relative price leads to a nominal adjustment.

4.4.2 The suboptimality of nominal price indexation

According to Proposition 7, a policy that sees the firm index its nominal price to the observed aggregate price level has to be suboptimal. Let us formalize this specific result.

Consider the firm setting its current nominal price $p_{i,1}^{index}$ equal to the targeted relative price $\widetilde{r}_{i,0}$ plus the aggregate price p_1. It follows that if $\widetilde{p}_{j,1} \neq p_1$, indexation leads to a different pricing action than the optimal nominal price of $\widetilde{r}_{i,0} + \widetilde{p}_{j,1}$. Hence, there must be a first-order loss in expected profits under the indexation rule, whenever $p_1 \neq \widetilde{p}_{j,1}$.

Proposition 8. *Let* $\delta^{index} = \delta \operatorname{sgn}(p_1 - \widetilde{p}_{j,1})$*. Up to a first-order approximation around* $p_1 = \widetilde{p}_{j,1}$*, the difference* $\ln \upsilon^*(\varepsilon^0, s_1, \widetilde{r}_{i,0} + p_1) - \ln \upsilon^*(\varepsilon^0, s_1, \widetilde{r}_{i,0} + \widetilde{p}_{j,1})$ *equals*

$$\left[\frac{e^{\widetilde{r}_{i,0}}}{e^{\widetilde{r}_{i,0}} - e^{y_1 - \omega_{i,1}}} - b - \alpha\delta^{index}\right](p_1 - \widetilde{p}_{j,1}) < 0.$$

Proof. See Online Appendix B.3. □

The reason is simple: when $p_1 \neq \widetilde{p}_{j,1}$, indexation leads to a change in the estimated relative price $\widetilde{r}_{i,1}$ away from $\widetilde{r}_{i,0}$. This moves the firm away from the perceived kink in demand, creating a loss in expected profits. Since this loss is first-order, it dominates the standard markup and aggregate-demand effects from a change in the aggregate price, leading to a drop in expected profits. This makes indexation strictly suboptimal.

Finally, let us consider a counter-factual economy where the firm is endowed with full confidence that the true DGP is the unique relationship $\phi(p_t - \widetilde{p}_{j,t}) = p_t - \widetilde{p}_{j,t}$, yet remains uncertain about the shape of its demand curve h. While the perceived kinks in expected profits in the space of the estimated relative prices remains, the firm is now confident that $p_{j,t} = p_t$. In this case, the perceived kink at $\widetilde{r}_{i,0}$ implies a kink at the nominal price $p_{i,1} = p_{i,0} + p_1 - p_0$. As a result, indexation is now optimal. Hence, while ambiguity about the shape of demand generates real rigidity, it is its interaction with uncertainty about the link between aggregate and industry prices that turns it into a nominal rigidity.[20]

4.4.3 Stickiness and memory in nominal prices

The previous insights carry through once we move beyond the example of a firm in its second period of life. With an unrestricted history of observations, past nominal prices and industry-price reviews $\{p_i^{t-1}, \widetilde{p}_j^{t-1}\}$ continue not to matter separately; the sufficient statistic is the history of estimated relative prices $\widetilde{r}_i^{t-1}$. Together with the quantities realized at those relative prices, it forms the information set used to update beliefs about demand.

[20]Propositions B.1 and B.2 in Online Appendix B.3 provide the analysis of this counterfactual economy, as a counterpart to Propositions 7 and 8 for the benchmark economy.

Proposition 6 extends in a straightforward fashion: given past history $\widetilde{r}_i^{t-1}$, worst-case beliefs feature kinks at all previously observed $\widetilde{r}_i \in \widetilde{r}_i^{t-1}$. When a review does not occur this period so that $\widetilde{p}_{j,t} = \widetilde{p}_{j,t-1}$, the kinks in expected demand occur at the same set of nominal prices as last period: $p_{i,t} = \widetilde{r}_m + \widetilde{p}_{j,t-1}$, for any $\widetilde{r}_m$ in the set $\{\widetilde{r}_1, ..., \widetilde{r}_n\}$. Thus, between industry price reviews, nominal prices display both rigidity and memory. In this general environment, stickiness in nominal prices manifests itself as "price plans", where the price series tends to bounce around a few points that look like "reference prices". When a new review signal arrives, the firm shifts the whole price plan accordingly.

5 Quantitative evaluation

Next, we evaluate quantitatively the empirical relevance of the model described in the previous section by testing its implications against a rich set of conditional and unconditional moments. This requires solving numerically the general decision problem of the ambiguity-averse firms given in equation (28). As discussed earlier, the dimensionality of the space grows with the length of the history ε^{t-1}, and to handle this problem we use the same $\tilde{V}$ approximation as outlined in Section 3.4. The advantage of this approach is that we can leave ε^{t-1} completely unrestricted, hence do not need to impose any ad-hoc assumptions limiting the memory of the firms. This way, we can evaluate the performance of our mechanism in the long-run, at the stochastic steady state of the model, where the history of observations ε^{t-1} is both endogenous, reflecting past optimal choices, and long.

5.1 Calibration

The model period is a week. We calibrate $\beta = 0.97^{(1/52)}$ to match an annual interest rate of 3%. The mean growth rate of nominal spending $\mu = 0.00046$ is set to match an annual inflation of 2.4%, and we pick the standard deviation $\sigma_s = 0.0015$ to generate an annual standard deviation of nominal GDP growth of 1.1%. Following the calibration in Vavra (2014) we set the persistence and standard deviation of aggregate productivity $\rho_a = 0.91^{(1/13)} = 0.9928$ and $\sigma_a = 0.0017$ to match the quarterly persistence and standard deviation of average labor productivity, as measured by non-farm business output per hour. We choose an elasticity of substitution of $b = 6$, implying a (flexible price) markup of 20%.

We choose the remaining parameters by targeting micro-level pricing moments from the IRI Marketing Dataset. The dataset consists of scanner data for the 2001 to 2011 period collected from over 2,000 grocery stores and drugstores in 50 U.S. markets. The products cover a range of almost thirty categories, mainly food and personal care products. For our purposes, we focus on nine markets and six product categories.[21] Because our model does not feature a rationale for sales, all reported moments are based on "regular price" series in which temporary sales are filtered out.[22]

Learning parameters

Our mechanism emphasizes non-parametric learning under ambiguity, which creates a rich learning environment characterized by six parameters $\{\delta, \gamma_p, \gamma, \sigma_x^2, \psi, \lambda_T\}$. With a focus on limiting the associated degrees of freedom, we set two of the learning parameters to values corresponding to natural limiting cases, and freely estimate the remaining four parameters.

First, regarding ambiguity over the demand function, we assume that the firm is confident that the mean demand function cannot be locally upward sloping, hence $\delta \leq b$. To minimize degrees of freedom, we fix $\delta = b$. Second, in terms of ambiguity over the unobserved industry price index, the parameter γ_p controls the size of the entertained set of cointegration relationships in equation (27). As detailed in Section 4, a positive γ_p is the reason why the joint worst-case beliefs about the demand function and the relative price lead to nominal rigidity in the short run. However, once the worst-case is determined and the firm engages in learning through the relative price $\widetilde{r}_i$, the specific value of γ_p only enters as a price-independent demand shifter in equation (34). Its quantitative role is therefore limited and thus we study the limit of $\gamma_p \to 0$.[23] This leaves four learning parameters $\{\gamma, \sigma_x^2, \psi, \lambda_T\}$ that we estimate by targeting micro-level moments, as detailed below.

Firm exit

The only modeling difference relative to the environment described in Section 4 is our assumption that with probability λ_ϕ, firm i exits and a newly-born firm takes its place in industry j. New firms have no information on the demand function beyond the time-zero prior, thus exit resets the information capital of firms.[24] This assumption serves two purposes. First, with an infinitely growing history of signals, conditional beliefs are non-stationary, making it difficult to evaluate behavior at the stochastic steady state. Second, it

[21] The markets are Atlanta, Boston, Chicago, Dallas, Houston, Los Angeles, New York City, Philadelphia and San Francisco. The categories are beer, cold cereal, frozen dinner entrees, frozen pizza, salted snacks and yogurt. A more complete description of the dataset is available in Bronnenberg et al. (2008).

[22] We use the methodology of Nakamura and Steinsson (2008) which aims to eliminate V-shaped sales. Also, as is usual with scanner datasets, we obtain the unit price by dividing weekly revenue by quantity sold. In order to minimize the probability that we identify spurious price changes due to middle-of-the-week repricing, the use of coupons, loyalty cards, etc., we take the conservative approach of eliminating any observations that feature a price with fractional cents.

[23] Relaxing the two assumptions on γ_p and δ would only allow the model to fit the data better, at the cost of leaving less room for additional testable implications.

[24] As such, we interpret reseting the informational capital as a broad concept, which includes any shock that makes the firm unsure that past observations are still informative, including major changes in the competitive landscape, the introduction of rival substitutes or technological change.

allows us to study pricing behavior over the firm's life-cycle, which serves as an additional moment restriction on our learning mechanism. Here, we set the exit probability $\lambda_\phi = 0.0075$, following Argente and Yeh (2017), who provide a detailed analysis of the duration of a UPC-store pair in the same IRI dataset that we use.

Demand and cost shocks

The firm's quantity sold is subject to demand shocks, with a standard deviation of σ_z. We calibrate this parameter by using empirical evidence on the accuracy of predicting one-period-ahead quantity. This involves estimating the demand regression:

$$q_{ijt} = \beta_0 + \beta_1 q_{i,j,t-1} + \beta_2 p_{ijt} + \beta_3 p_{ijt}^2 + \beta_4 cpi_t + week_t'\theta_1 + store_j'\theta_2 + item_i'\theta_3 + z_{ijt} \quad (35)$$

where q_{ijt} and p_{ijt} are quantities and prices in logs for item i in store j at time t; cpi_t is the (log) consumer price index for food and beverages; while $week_t$, $store_j$ and $item_i$ are vectors of week, store and item dummies respectively.[25,26] We then compute the empirical standard deviation of the residuals z_{ijt} leading us to set $\sigma_z = 0.613$.[27]

The firm also faces cost shocks. Since we do not have cost data, we estimate the persistence and volatility of idiosyncratic productivity (respectively ρ_w and σ_w).

Simulated method of moments

We estimate the six remaining parameters, $\{\rho_w, \sigma_w, \sigma_x, \psi, \lambda_T, \gamma\}$, via simulated method of moments, by targeting the six pricing moments described in Table 2. For the most part, these are basic pricing moments widely used in the literature to discipline price-setting models. Throughout, we define the 'reference price' as the modal price within a 13-week window period, as in Gagnon et al. (2012). The last moment, the mean duration of a pricing regime, appeals to the fact that in our model, the kinks in expected demand turn basic stickiness into price plans. In both actual and simulated data, we identify these price plans using the method in Stevens (2014).[28] Table 1 presents all parameters values, while Table 2 shows the outcomes for moments targeted in th estimation.[29] The model matches the targeted

[25]To remain consistent with all the other moments analyzed, the regression is only run on observations for which the posted price is the regular price: $p_{ijt} = p_{ijt}^{reg}$ and $p_{ijt-1} = p_{ijt-1}^{reg}$.

[26]Given the high (weekly) frequency of our data and the fact that we do not find evidence of middle-of-the-week price changes, endogeneity is unlikely to be a significant issue here.

[27]The regression is run and the volatility measure is computed first for each of the 54 category/market pairs, before being aggregated using revenue weights.

[28]The methodology modifies the Kolmogorov-Smirnov test to identify shifts in the distribution of price changes over time. In order to have enough observations from which to identify regimes when applying to the data, we ignore quote-lines that have missing price data or less than 104 weekly observations (2 years). We use Stevens (2014)' standard critical value of 0.61 throughout our regime identification exercises, for both actual and simulated data. Also, in both cases, we eliminate regular price changes of less than 1%.

[29]The estimation is based on a simulated panel of 5000 time periods with 1000 active firms in each period.

moments very well and, naturally, it does so through a positive ambiguity parameter γ, which is the necessary source for any price stickiness in the model.[30]

Table 1: Parameter Values

Calibrated Parameters							Estimated Parameters					
β	μ_s	σ_s	ρ_a	σ_a	σ_z	λ_ϕ	ρ_w	σ_w	σ_x	ψ	λ_T	γ
0.9994	0.00046	0.0015	0.993	0.0017	0.613	0.0075	0.998	0.008	0.691	4.609	0.018	0.614

Table 2: Targeted moments - Data vs model

	Data	Model
Frequency of regular price changes	0.108	0.105
Median size of absolute regular price changes	0.149	0.154
75th pctile of the distribution of non-zero absolute price changes	0.274	0.277
Fraction of non-zero price changes that are increases	0.537	0.533
Frequency of modal price changes (13-week window)	0.027	0.026
Mean duration of pricing regimes	29.90	30.54

5.2 Testable implications

Next, we analyze the ability of the model to match various features of the data that were not directly targeted in the estimation, yet speak to the mechanisms at the heart of our model. We start with moments that are more specific to models of the "reference price" family, such as the behavior of the reference price, as well as memory and discreteness characteristics. Then we turn our attention towards empirical features that represent a challenge for many price-setting mechanisms, such as the coexistence of small and large price changes, as well as the declining hazard function of price changes. Finally, we conclude with moments that have been less studied or are novel in the literature, such as the behavior of prices over the life-cycle of a product, or the relationship between demand realizations and price adjustment.

5.2.1 Reference price moments

Panel A of Table 3 shows that the model matches well the empirical behavior of reference prices. First, it correctly predicts that the typical modal price is generally also the highest price in a given 13-week window – the probability of that occurring in the data is 82% vs. 74% in the model. Second, within each 13-week window, we also compute the average

[30]Online Appendix C.1 shows that the estimated γ implies an empirically plausible amount of ambiguity, as it generates dispersion in prior demand forecasts that matches the evidence in Gaur et al. (2007).

Table 3: Untargeted moments - Data vs model

		Data	Model
Panel A	Prob. modal P is max P	0.819	0.740
	Fraction of weeks at modal P (13-week window)	0.828	0.880
	Prob. price moves to modal P	0.592	0.669
Panel B	Prob. visiting old price (26-week window)	0.48	0.414
	uni (26-week window)	0.792	0.822
Panel C	Avg hazard slope (LPM)	-0.011	-0.015
Panel D	Old vs. young prices - Average slope		
	Spell threshold: $\Gamma = 4$	-0.090	-0.212
	$\Gamma = 5$	-0.104	-0.189
	$\Gamma = 6$	-0.104	-0.173

fraction of weeks that the regular price spends at the reference (modal) price. The simulated moment compares favorably to its empirical counterpart: while in the data the regular price spends 83% of the time at the modal price, in the model this fraction equals 88%.

A worry could be that the above result is simply a mechanical artefact of the high degree of stickiness of modal prices. To show that this is not the case, we compute the probability that the change of a non-modal regular price ends at the modal price, and not at some other regular price. This probability is equal to 59% in the data, compared to 67% in the simulations. It confirms that the 'attractiveness' of the modal or reference price is not simply a by-product of pervasive price stickiness.

5.2.2 Discreteness and memory

In our model, the first-order perceived cost of moving away from any of the previously-observed prices implies that prices display memory (see Corollary 3). This prediction is clearly present in our dataset, even with temporary sales filtered out: Panel B of Table 3 reports the probability that, conditional on a price change, the firm posts a regular price that it has already visited within the last six months (26 weeks) is 48%, once weighted across markets and categories. In the model, that same probability is 41%, in line with the data. Note that a standard menu cost or Calvo model would feature no such price memory and the probability would be 0%, since the firm has no inherent reason to repost a previous price.

A related empirical observation is that firms tend to cycle through a relatively limited, discrete set of prices as opposed to posting a lot of new unique prices. For example, in the IRI dataset, the average number of unique prices observed in a window of 26 weeks is only 2.34. This simple statistic, however, is directly impacted by the degree of price stickiness. To ease interpretation, we produce another statistic: for each product i (a given UPC) sold in a specific store j, we compute the number of unique prices and price changes observed within the 26-week window centered around week t, and denote them by u_{ijt} and c_{ijt} respectively. We then define the ratio $uni_{ijt} \equiv u_{ijt}/(c_{ijt}+1)$.[31]

This statistic standardizes the number of unique prices observed by the number we would expect to see if the firm never revisited old prices, given the probability of price change. Thus, if all price changes end up at a price that had not been visited before within a specific window, the ratio uni_{ijt} would be equal to 1. This is the value we would expect from a standard price-setting mechanism. Yet, in the data we see that the average value of this moment is 0.792 across category/market pairs, as reported in Table 3. Our model does an excellent job of matching this feature of the data as well, implying a value of 0.822.

5.2.3 Size distribution of price changes

In our model, the perceived cost of changing prices is history-dependent and a function of the absolute size of the price change (see Proposition 3). For instance, if the firm is sitting at a price where it has accumulated a series of good demand signals, the adjustment away from that price is typically gradual, generating small price changes.[32] As a result, our model allows for the co-existence of large and small price changes.

This property is evident from Figure 3, which plots the distribution of the size of price changes in both the model and the data. Focusing on the left panel, we can see that there is a substantial amount of both small and large price changes in our model, *despite the fact* that there is no parameter heterogeneity and all firms are ex-ante identical. Arguably, introducing such heterogeneity would allow us to better match the empirical distribution (right panel), by smoothing out the distribution across the whole support. Still, our framework is intrinsically compatible with the presence of price changes of various sizes.

5.2.4 Hazard function of price changes

A related characteristic of our setup is that, all else equal, the firm is less willing to move away from a price that it has stayed at for longer and thus acquired more information about – this is because the perceived cost in terms of expected profits of changing the last posted price is increasing in the number of times it has been observed (see Corollary 2). This naturally gives rise to a declining hazard function of price changes: the probability of a price change conditional on the price having survived τ periods is decreasing in τ.

[31] In both the data and the model, we drop from the computation any window that features no price change. We thank an anonymous referee for suggesting this moment to us.

[32] More details on this behavior will be provided in Section 5.3.

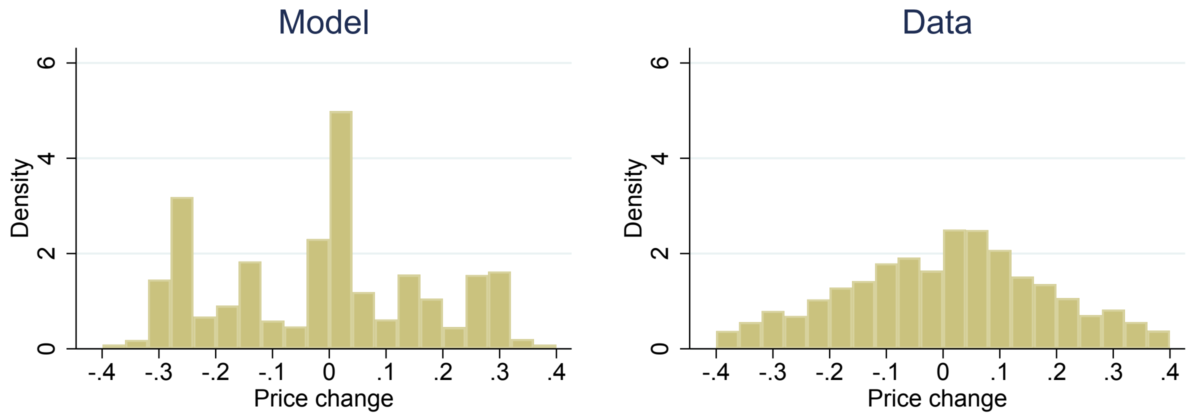

Figure 3: Distribution of the absolute size of price changes. Data vs. simulations.

The shape of the hazard function has been heavily discussed in the price-setting literature, and evidence of a declining hazard has been documented in many micro price datasets. Nakamura and Steinsson (2008), for example, estimate a downward-sloping hazard using U.S. CPI data, a characteristic that they consider represents a challenge to many popular price-setting mechanisms. Some, however, have argued that this empirical finding could be a by-product of not taking proper care of heterogeneity: as noted by Klenow and Kryvtsov (2008), "[t]he declining pooled hazards could simply reflect a mix of heterogeneous flat hazards, that is, survivor bias." In light of this word of caution, we use two approaches to confirm that it is indeed a robust feature of our dataset.[33]

Linear probability model

Our first exercise is to run a linear probability model (LPM) with a rich set of fixed effects to control for heterogeneity in unconditional price change frequencies that may mechanically generate downward-sloping hazard functions. The LPM is run separately for each of the 54 category/market pairs, allowing for different slopes of the hazard function. The use of a linear regression circumvents the incidental parameters problem that arises with the use of fixed effects in non-linear models, such as a proportional hazard framework or a probit. While this approach is similar in spirit to the one used by Klenow and Kryvtsov (2008) on CPI micro price data, we are much more aggressive in controlling for heterogeneity: in their case, only ten fixed effects are employed, one per decile of the price change frequency distribution. The much larger number of observations in our dataset allows us to control for product, store *and* time fixed effects. Note that in line with the literature, we drop all left-censored spells from the sample. Importantly, in Online Appendix C.2 we apply our econometric approach to panels of simulated data and show that it allows us to recover the

[33]Since we use regular instead of posted prices, it implies that temporary sales are not driving the results.

true value of the slope of the hazard function, *even* in the presence of pervasive heterogeneity.

For each category/market, we run a separate regression of the type:

$$\mathbb{1}(p_{i,j,t} \neq p_{i,j,t-1}) = \alpha + \beta\tau_{i,j,t} + \gamma_i + \gamma_j + \gamma_t + u_{i,j,t} \tag{36}$$

where the symbol $\mathbb{1}(.)$ denotes the indicator function.[34] Since $\tau_{i,j,t}$ is the length of the price spell (i.e. the number of weeks since the price has last been changed), the coefficient β therefore represents the estimate of the slope of the hazard function. Finally, γ_i, γ_j and γ_t are product, store and week fixed effects respectively. These shifters control for any systematic heterogeneity in the degree of stickiness across items, outlets and time which, as we discussed earlier, would bias downwards the slope of the hazard. We run the regression on spells shorter than or equal to 26 weeks, where the vast majority of observations lie.

A summary of our results is in Panel C of Table 3. For the data, we obtain a slope estimate $\widehat{\beta}$ of -0.011, once averaged across the 54 category/market pairs. The estimated slope coefficients are negative and statistically significant at the 1% level in *all* category/market pairs, whether we use unweighted or weighted observations.[35] This value implies that each additional week that a spell survives lowers the probability of observing a price change by about 1.1 percentage point.[36] This evidence of a downward-sloping hazard function is consistent with that from Nakamura and Steinsson (2008), albeit we apply a more aggressive treatment of heterogeneity by allowing product- and even store-level shifters.

To evaluate the model's ability to match the empirical hazard, we estimate the same LPM regression on the data simulated by the model. At -0.015, the slope of the simulated hazard is steeper, yet compares well with its empirical counterpart.

Product/store cells

How confident should we be that the fixed-effects method allows us to circumvent the survivor bias issue? The Monte Carlo analysis in Online Appendix C.2 confirms that the fixed-effects approach is able to recover the true slope of the hazard function, even in the presence of pervasive item- and store-level heterogeneity. Yet, one may be worried that heterogeneity is also present at the *item/store*-level, or that our estimates could be biased through some complex interaction, such as between the heterogeneity in price change frequencies and nonlinear hazard slopes.

To account for these potential biases, we complete our analysis by applying the approach

[34] We cluster standard errors at the store level, the cluster which yields the highest standard errors.

[35] In Figure C.1 of the Online Appendix C.3, we plot the distribution of coefficient estimates $\widehat{\beta}$ across the 54 category/market pairs.

[36] If we run regressions with sales-weighted observations, the result is very similar at -0.010 for all price spells (-0.011 unweighted).

suggested by Campbell and Eden (2014) to our dataset. Following their nomenclature, we define a "cell" as a specific product (UPC) sold in a given store that has at least 200 observations.[37] Then, for each cell, we compute the probability of observing a price change for young and old prices, where an old price is one that has survived for $\tau \geq \Gamma$ periods.

We find overwhelming evidence of declining hazards: we vary Γ between 4 and 6 and show that the vast majority of product/store pairs have price change frequencies that are lower for older than for younger prices. As indicated in Panel D of Table 3, the weighted average of the individual slopes ranges from -0.104 to -0.090. The left column of Figure C.2 in Online Appendix C.3 displays the empirical distributions of the hazard slope.

The differences between the price change frequencies of older and younger prices are statistically significant, too: the fraction of product/stores for which the slope is negative *and* statistically significant ranges from 77.1% to 83.3% across the values of Γ, while the cells with positive and significant slopes account for only 5.9% to 15.2% of the total.[38]

In line with our regression-based approach, we find that the model generates hazard slopes that tend to be steeper but still comparable to their empirical counterparts. From Panel E of Table 3, we can see that for the middle threshold of $\Gamma = 5$, the slope is -0.189 in the simulated sample versus -0.104 in the data. In addition, Figure C.2 in Online Appendix C.3 shows that the model is able to generate a similar distribution of slopes to that in the data, despite the fact that the model does not feature any ex-ante heterogeneity.

5.2.5 Pricing behavior over the product life-cycle

Recent work by Argente and Yeh (2017) documents interesting novel facts about the evolution of micro-level price dynamics over the life-cycle of the typical product. Using the same dataset as ours, they find that the frequency and size of price changes of the typical product both decline significantly as the product ages.

In our model, the price behavior over the life-cycle of the product/firm is shaped by the history dependence of the optimal pricing decision through the interaction of two forces. First, at the beginning of its life, the firm does not have much information about the demand curve of the product it sells and has therefore not yet established any deep perceived kink in expected demand. Second, the fact that the firm has very little information about demand increases the relative value of experimentation, as discussed in Section 3.4. Both of these forces imply that price flexibility decreases with age: newly-born firms tend to change prices more frequently than firms that have been in existence for a while and have accumulated

[37] Our conclusions are not affected by the choice of the minimal number of observations. In fact, raising the minimum number of observations tends to generate stronger quantitative results.

[38] We use the statistical test detailed under footnote 15 of Campbell and Eden (2014).

significant information capital at past prices. Similarly, the experimentation motive implies that the average *size* of price changes for young firms is larger than that of older firms.

We quantify the life-cycle properties of the frequency and size of price changes in our model by running the following two regressions on the simulated data:

$$\mathbb{1}(p_{i,t} \neq p_{i,t-1}) = \beta_0^{freq} + \beta_1^{freq} \mathbb{1}(age_{i,t} \leq 26) + \varepsilon_{i,t}$$

$$|\Delta p_{it}| = \beta_0^{size} + \beta_1^{size} \mathbb{1}(age_{i,t} \leq 26) + \varepsilon_{i,t},$$

where the coefficients of interest are β_1^{freq} and β_1^{size}, which capture respectively the frequency and size of price changes in the first 6 months of a firm's life relative to the next half a year.[39]

In both cases, we find positive and statistically-significant coefficients: $\widehat{\beta_1}^{freq} = 0.23$ and $\widehat{\beta_1}^{size} = 0.09$. In other words, our model predicts that both the frequency and size of price changes fall as a new product ages, in line with the evidence from Argente and Yeh (2017).

5.2.6 Past demand realizations and price-setting decisions

The focus so far has been on price-related moments, as is common in the literature. Yet, our model also has stark and unique implications about the relationship between quantities and prices. In particular, the perceived cost of changing the last posted price increases with the realized value of the demand shock at that price (see Corollary 4): a firm that observes a particularly good demand realization is more likely to stay put, while bad demand realizations raise the likelihood of a price reset. We test this prediction of the model by producing a novel set of results for both simulated and actual data.

As a first step, we extract demand innovations from the data by using regression (35), as described earlier. The object of interest is the residual z_{ijt}, the unexplained or "surprise" demand component for item i in store j at time t. We then construct two indices that capture, in light of our model, how attractive a given price may be from the perspective of the firm. We define the z-score of price p_{ijt} as:

$$zscore_{ijt} = \frac{\sum_{\tau=0}^{26} [valid_{ij,t-\tau} \times z_{ij,t-\tau}]}{\sum_{\tau=0}^{26} valid_{ij,t-\tau}}.$$

The indicator $valid_{ij,t-\tau} = 1$ if $p_{ijt} = p_{ij,t-\tau}$, that is, if the price at time $t - \tau$ is the same as the one we compute the z-score for. Conceptually, the z-score of price p_{ijt} corresponds to

[39] Focusing on first 12 months of life helps isolate the life-cycle effects. The estimates are even more pronounced if we do not censor on the right. As is typical with any moments on the size of price changes, the second regression only considers time periods with a non-zero price change.

the average of the demand innovations at that price.[40]

The $zscore_{ijt}$ includes information up to time t. It is useful to also define a version that only incorporates demand innovations up to $t-1$ (which is what informs price choice at t):

$$zscore^{lag}_{ijt} = \frac{\sum_{\tau=1}^{26} \left[valid_{ij,t-\tau} \times z_{ij,t-\tau} \right]}{\sum_{\tau=1}^{26} valid_{ij,t-\tau}}.$$

Finally, we define $wscore_{ijt}$, which captures how often a price has been posted in the past:

$$wscore_{ijt} = \sum_{\tau=0}^{26} valid_{ij,t-\tau}.$$

In order to test whether the firm is less (more) likely to move away from a price that experienced an unexpectedly good (bad) demand realization, we run the following regression:

$$\mathbb{1}(p_{i,j,t} \neq p_{i,j,t-1}) = \beta_0 + \beta_1(zscore_{ij,t-1} - zscore^{lag}_{ij,t-1}) + \beta_2 wscore_{ij,t-1} + f_{ij} + \varepsilon_{ijt}. \quad (37)$$

The LHS equals 1 when the price at t is different than at $t-1$, and 0 otherwise. The regressor of interest, $zscore_{ij,t-1} - zscore^{lag}_{ij,t-1}$, corresponds to the change in the z-score of the price posted at $t-1$: a positive value indicates that all else equal, the firm was hit by a relatively good demand realization at time $t-1$.[41] We also control for the w-score, which is the proper way of controlling for the declining hazard under the null hypothesis of our mechanism. When run on the actual data, the panel regression includes either category/market or product/store fixed effects f_{ij} in order to control for the heterogeneity in price change frequency.

For both the model and the data, we run two main regressions. The first one imposes no additional restrictions. The second uses only observations for which $wscore_{ij,t-1} \leq 12$, so that the price the firm is considering leaving has been posted for at most half of the periods within the backward-looking 26-week window. This distinction is driven by our model prediction that new demand realizations are less likely to influence the decision to change a price that has been observed more often in the past (high w-score).

Table 4 presents the results of running the regression in equation (37) on both the actual and simulated data. To ease the interpretation, the coefficients are reported as marginal effects: the impact of a one-standard-deviation deviation in the z- or w-score on

[40] We truncate the window to 26 weeks for two reasons. First, from a data standpoint, we want to avoid losing too many observations through left-censoring. Second, this truncation allows us to capture the idea that demand realizations very far back are likely to be of little value to the firm. We also tried to geometrically discount past observations; this has little impact on the results.

[41] To minimize the risk that changes in the z-score are driven by some complex non-linearity in the demand function, we focus on observations for which there was no price change at $t-1$, i.e. $p_{i,j,t-1} = p_{i,j,t-2}$. This implies that changes in the z-score are driven by recent demand realizations.

the likelihood of a price change. All coefficients are statistically significant at the 1% level. Three observations on the z-score effect are worth highlighting.

First, the effect is negative in all regressions: a good (bad) demand realization at the posted price that lifts (lowers) the z-score decreases (increases) the chance of moving away from that price. This is in contrast to most state-dependent mechanisms, such as a standard menu-cost model: in these environments, both positive and negative shocks make the firm more likely to reprice as they raise the gap between the current and optimal prices.

Second, the effect is indeed larger for more "recent" prices (low w-score): while a one-standard-deviation change in the z-score decreases the probability of a price change by between 80 to 90 basis points when we condition on $wscore_{ij,t-1} \leq 12$, the effect is only around 55 basis points with $wscore_{ij,t-1} \leq 25$. The effects are also economically meaningful, as a 80bp increase in the probability of a price change is about 10% of the unconditional probability of a price change in the data.

Third, the z-score effects in the data and the model are similar: for younger prices, the absolute impact on the price change frequency is 83bp, almost perfectly in line with the 86-87bp effect in the data. They also compare favorably when conditioning on $wscore_{ij,t-1} \leq 25$ (65bp vs. 57-58bp in the data).

Table 4: Results from the z-score regressions

	Data				Model	
$wscore_{ij,t-1} \leq x$	$x = 12$		$x = 25$		$x = 12$	$x = 25$
$zscore_{ij,t-1} - zscore^{lag}_{ij,t-1}$	-0.0087	-0.0086	-0.0058	-0.0057	-0.0083	-0.0065
$wscore_{ij,t-1}$	-0.0373	-0.0290	-0.0466	-0.0264	-0.0253	-0.0195
Category/market FE	X		X			
Product/store FE		X		X		

Note: The dependent variable equals 1 when $p_{i,j,t} \neq p_{i,j,t-1}$, 0 otherwise. The empirical regressions include either both category and market fixed effects, or item/store fixed effects. We report marginal effects: the impact of a one-standard-deviation in the independent variable on the likelihood of a price change. Standard errors are clustered at the category-market level. All coefficients are statistically significant at 1% level.

5.3 The typical pricing policy at the stochastic steady state

In this section, we analyze the optimal pricing policy at the stochastic steady state. This serves two main purposes. The first is to show that even after hundreds of periods of observations, firms still face significant uncertainty over demand – learning proceeds slowly in the model. The second is to visualize the typical pricing policy function, which helps explain how the model generates the moments highlighted in the previous section.

We start by noting that at any point in time, the equilibrium of our model is described by a whole distribution of beliefs over the unknown demand function, varying across firms. The reason is that firms have faced different histories of idiosyncratic shocks, and thus have made different pricing decisions, resulting in heterogeneous histories of signals. To understand the average behavior, here we analyze the action of a firm at the typical history of observations.

Since firms learn in terms of the estimated relative prices $\widetilde{r}_{it}$ (as per Section 4), the information sets of different firms are characterized by the unique $\widetilde{r}_{it}$ values seen in the past, together with the resulting demand signals at those prices. A striking characteristic is that even though the average life span of firms in our model is 133 periods, the histories contain only 6 unique estimated relative prices on average. Moreover, the most often posted $\widetilde{r}_{it}$ accounts, on average, for 74% of all past observations. Hence, the typical history features one dominant "reference" estimated relative price point that the firm tends to revert to.

To visualize this typical behavior, we average over the histories of observations of the different firms in order to come up with a "typical" history of observations - the precise details of the procedure are presented in Online Appendix C.4. We then compute the optimal pricing policy conditional on having observed this typical price history, as a function of the level of idiosyncratic productivity, keeping aggregate variables constant at their mean values. This is true in particular for the gap between the aggregate price level and the unambiguous signal of the industry price, $p_t - \widetilde{p}_{j,t}$, which is kept fixed at its average level. Under this assumption, the statements below about the estimated relative price $\widetilde{r}_{it}$ are also statements about the behavior of the nominal posted price p_{it} between industry price reviews.

The resulting pricing policy, plotted in Figure 4, exhibits several key characteristics. First, it features a large flat spot that covers the middle part of the support for idiosyncratic productivity (recall $E(w_{it}) = 0$) – this corresponds to the "dominant" estimated relative price point (the one that is on average posted 74% of the time) and it occurs at $\widetilde{r}_i = 0.11$. It is intuitive that the firm has established a large flat spot at a price that is optimal for productivity values w_{it} close to the mean, as they are the ones it is most likely to face.

Second, the policy also features five smaller flat spots corresponding to the other previously observed five price points. Those estimated relative prices are sticky and attractive, but because each is optimal for fewer and less likely w_{it} realizations, the firm tends to post these prices less often. Combined with infrequent observations of the industry price p_{jt}, these features of the policy function generate both stickiness and memory in nominal prices between reviews. The price is not only likely to be "stuck" at one of the flat spots but, even conditional on moving, the price is likely to go to one of the other flat spots (since a large

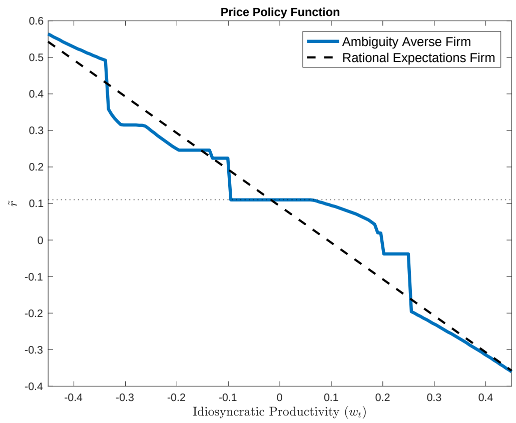

Figure 4: Optimal pricing policy function at the stochastic SS

part of the productivity support maps to one of them), thus revisiting past values.

Third, there are several jumps in the pricing policy, typically occurring as a switch from one flat spot to another. The largest jumps, however, correspond to a move from a flat spot to a brand new estimated relative price further out in the tails, and can be explained by the experimentation motive: the typical firm has not collected much information about demand at very high or low values of $\widetilde{r}_{it}$. Given this high uncertainty, the firm would generally not like to price in those regions, but large enough shocks will eventually force it to. However, the high remaining uncertainty about demand in those parts of the price space makes experimentation attractive, and rather than extending its pricing decision continuously, the firm finds it optimal to adjust a lot, thus learning more about the distant regions of the price space.

Fourth, in addition to the jumps, the pricing policy also features several continuous downward-sloping portions which are behind the small price changes seen in the simulations. The most pronounced of those continuous portions occurs immediately to the right of the main flat spot in the middle. Intuitively, when the firm experiences a moderate productivity shock, it remains in the neighborhood of its "safe" reference price that it knows best instead of exploring remote price points. This is due to the local nature of learning – the firm has reduced uncertainty not only right at the reference price, but also in its neighborhood, and would rather not move far away unless productivity changes by a substantial amount.

Lastly and importantly, the policy function also shows that the average firm has far from perfect information about its demand curve. This is evident from the significant difference between the typical policy function and the full information RE policy (dashed black line). The reason behind this substantial residual demand uncertainty is that the history of observations is *endogenously sparse*. The optimal policy leads the firm to often repeat estimated relative prices, resulting in a history of observations that provides a lot of information about the average level of demand at those select prices, but leaves the firm uncertain about the *shape* of its demand. Hence our mechanism, which operates specifically through the uncertainty about the local shape of demand, has a strong bite even at the steady state of the model, when firms have seen long histories of demand observations.[42]

5.4 Comparative statics

We now turn to comparative statics. A common theme throughout is the nuanced link between price flexibility and memory, which as we will see in Section 5.5, is an important determinant of how micro-data stickiness maps into the effects of monetary policy.

Table 5: Moments - Comparative statics

	Benchmark	$\beta = 0$	$\psi = 0$	Low δ	High σ_ω	High b
Freq. regular prices changes	0.105	0.075	0.064	0.207	0.160	0.199
Median size of abs. changes	0.154	0.007	0.015	0.108	0.123	0.015
Freq. modal price changes	0.026	0.028	0.029	0.041	0.037	0.056
Prob. visiting old price	0.414	0.237	0.469	0.444	0.502	0.488

Note: Moments are computed across versions of the model in which only the parameter in the column header is changed, while all others are kept at their benchmark value. 'Low' or 'High' means that we halve or double, respectively, the corresponding parameter compared to its benchmark value.

First, we model a myopic firm by setting $\beta = 0$, which eliminates all experimentation incentives. The key pricing moments under this parameterization are reported in Table 5, where we see a drop in both the frequency and median size of price changes. Without a reason to explore new parts of the demand curve, firms now have less incentives to change prices often or by large amounts. Furthermore, this leads the firm to concentrate its information accumulation in the middle range of productivity shocks, leading to a policy function with two large flat spots in the middle, but no other kinks. As a result, the frequency of modal price changes rises slightly, but memory falls significantly because there are no other attractive prices outside of those two. Moreover, in unreported results we find that this myopic version generates very few large price changes and does not match the product pricing life-cycle facts, as young firms no longer have an experimentation motive to change prices more often.

Second, we set $\psi = 0$ to eliminate the local nature of learning. In that case, each signal carries the same quantity of information for any other price point, irrespective of its distance from the current price. This setting also kills the experimentation motive (Proposition A.1 in

[42] In fact, because of the local nature of learning and the endogenous location of demand signals, learning proceeds so slowly that the mechanism survives even if firms live for thousands of periods. We explore this implication further in Online Appendix C.5 by setting $\lambda_\phi = 0$. In the same appendix we also show that the accumulation of new information could in fact change the optimal position of some of the reference prices.

the Online Appendix A.3), because the new information contained in a signal is not specific to the position of the price at which the signal was observed. Therefore, the resulting moments are mostly similar to the ones with $\beta = 0$, as can be seen in Table 5. The main difference is memory, which increases to 47%. This is due to the emergent ergodic policy function, which now features numerous, smaller kinks as opposed to just two large ones, increasing the probability of switching between kinks. The intuition can be seen from Proposition 3, which shows that when $\psi = 0$ the perceived demand loss of moving away from a kink to a new price is relatively steeper for larger price changes as compared to smaller adjustments. As a result, smaller price changes are perceived as relatively safer, leading the firm to establish several kinks in the same neighborhood, as opposed to just a single one.

Third, we decrease the degree of ambiguity by halving δ. By Proposition 1, this lowers the as-if cost of moving away from the previously posted price. As a result, price changes occur more often (both regular and modal), and the size of the resulting price changes is smaller. Interestingly, this increased flexibility implies more kinks and hence more (but smaller) flat spots in the pricing policy. The result is higher memory, as there is a higher number of attractive prices that were set previously.

Fourth, we double the standard deviation of the idiosyncratic productivity shocks, σ_ω. This raises the frequency of modal and posted price changes, an intuitive result that is shared with a number of other standard frameworks.[43] In our model, however, the increased price flexibility is also accompanied by higher memory. The reason is that with more frequent price changes, information accumulation is spread out over a larger set of individual prices, resulting in a policy function with more steps and thus increased memory. Hence, even though prices change more frequently, they are also more likely to revert to past price levels – which we will see can result in a less responsive aggregate price level.

Finally, we increase the average price elasticity of demand by doubling the value of b. The resulting higher sensitivity to deviations from the optimal markup, now at 9%, leads to a significantly higher frequency and a smaller absolute size of price changes, as documented in the last column of Table 5.[44] We find that, as in the δ and σ_w comparative statics, the increased flexibility comes with higher memory, from having more steps in the policy function. This positive correlation of frequency and memory is not mechanical, as shown by the $\psi = 0$ case where the two moments move in the opposite direction.

[43]See Klenow and Willis (2016) as an example of a discussion on the role played by the distribution of shocks in standard price-setting models.

[44]These results are consistent with Mongey (2018) who reports that products facing more competition are characterized by a larger frequency of posted prices and smaller absolute price changes.

5.5 Monetary non-neutrality

We quantify the degree of monetary non-neutrality by computing the impulse response of output produced by the measure-zero set of ambiguity-averse firms to an innovation in aggregate nominal spending, s_t. Note that because all other firms have rational expectations, our exercise arguably represents a lower bound on the size and persistence of monetary non-neutrality since it ignores potentially important strategic complementaries in price setting.

We estimate the impulse response via Jorda projections, an approach well suited to the high degree of non-linearity in our model. In the expression below, the LHS represents the $t + k$ output of ambiguity-averse firms, while the RHS corresponds to the nominal shock ϵ_t^s:

$$\ln\left(\int Y_{i,t+k} di\right) = \alpha_k + \beta_k \epsilon_t^s + u_{j,t+k}.$$

The series of coefficients β_k form the impulse response of output to the nominal shock, and are plotted with the blue line in Figure 5. The x-axis indicates the number of weeks since the shock, and the y-axis is scaled to represent the output response as a fraction of the shock. We see that for a 1% shock, real output increases by 0.36% on impact. The effect is persistent and declines gradually, with a full cumulative output effect of 6.7% after 52 weeks.

We contrast the response with simple menu cost and Calvo models, both calibrated to match the average frequency of price changes, while keeping the same shock processes.[45] For the menu cost version (dashed black line) we find a similar response on impact, but unlike our model, the effect dies down quickly, disappearing after 7 weeks. The Calvo model (dash-dotted black line) has a stronger impact effect, but also declines quicker than in our model. Overall, our cumulative output effect is similar to that in Calvo (at 6.5% after 52 weeks) and is six times larger than in the simple menu cost model (at 1.06% after 52 weeks).

Role of micro-level moments in shaping monetary non-neutrality

Our model features significant and persistent real effects of nominal spending shocks, despite the fact that it is consistent with two moments that are often taken to imply strong monetary neutrality: (i) a high frequency of price adjustment (10.5% at the weekly level) and ii) a large median size of price changes (15%). As we now show, this is due to the fact that our model is also consistent with two empirical facts that matter fundamentally for monetary policy effects: i) prices have memory, so that more than 41% of price resets end up at a previously visited price level; and ii) there are both large and small price changes.

[45] We leave to future work the analysis of a full general equilibrium version, which would take into account strategic complementaries that have been shown to magnify the effect of other pricing frictions. Therefore, in order to facilitate comparisons with standard models, we use versions of the menu cost and Calvo models that similarly ignore such strategic complementarities.

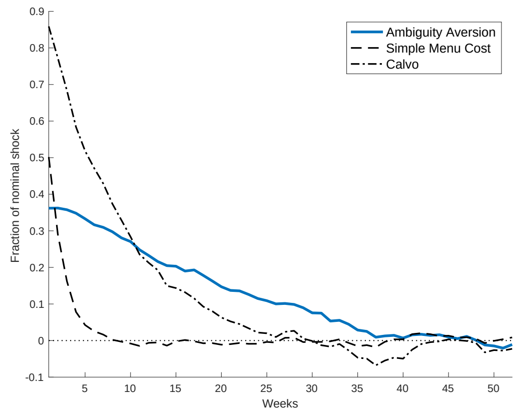

Figure 5: Impulse Response to a Monetary Shock

Consider first the impact effect. As we see in Figure 5, the output effect in the menu cost model is much smaller than for Calvo. This is due to the well-known Golosov and Lucas (2007) selection effect – the adjusting firms in the Calvo framework are chosen randomly, but are self-selected in the menu cost model and hence they adjust by a lot on average. Our model is also state-dependent, as the firms that sit at the threshold of adjustment respond to the fall in real markups brought upon by the positive nominal shock, as detailed in Proposition 7. In fact, the significantly larger price changes in our model relative to the menu cost counterfactual (15% versus 2.5% in absolute terms) might at first look suggest that the selection effect has even more bite in our case, as the typical price change is larger.

However, two forces weaken this selection, to the point that the effect on output on impact is similar to the menu cost model. First, a fraction of firms in our model are on the continuous portion of their policy function (recall Figure 4), and thus change prices only by a little. This force highlights the importance of having a model consistent with the presence of both large and small price changes, a point made for example by Midrigan (2011). Second, many large price changes in our model arise from experimentation motives, as discussed in Section 5.4, and thus do not depend directly on the change in the real markup.[46]

Let us now turn to the output effects in the periods following the impact. While in the Calvo model the effect declines at the exogenous frequency of price changes, in the menu cost model this persistence is significantly lower because the selection effect leads to large price adjustments in all periods after the shock. Our model differs from these two counterfactuals by predicting significantly stronger persistence. The fundamental reason is that optimal nominal prices have memory, a point emphasized by Eichenbaum et al. (2011) and Kehoe and Midrigan (2015). In the case of our model, memory generates long-lived real effects

[46]We share this weakening of the selection effect that arises from experimentation with the parametric learning model of Argente and Yeh (2017).

because price movements tend to happen between kinks that have formed prior to the shock.

By the selection effect logic, the firms that raise their prices due to the positive nominal shock were, to begin with, close to indifferent between the price they just left and their new one. Importantly, 41% of price changes land at a previously visited price (memory), where the firm perceives another kink that translates into a flat spot in the optimal policy function (see Figure 4). Since monetary policy shocks are small, a typical firm that lands on another kink would again be close to indifferent between the new and previous prices. Thus, as they face new idiosyncratic shocks in the following periods, these firms are likely to revert back to their old, lower prices. Essentially, the *as-if* cost of price change in our model is asymmetric for a firm that is almost indifferent between two flat spots in its policy function – it perceives a small cost of changing to the other kink, but a large cost to moving in the opposite direction. This asymmetry explains why a significant proportion of firms tend to undo their initial price change. As a result, the real effect of the nominal shock is long-lived, even as firms exhibit apparent flexibility in their nominal prices.

This discussion indicates that our theory is relevant not only because its micro-level predictions find strong empirical support, but also because they matter for shaping monetary policy effects. This makes the model well-suited for counterfactual analysis, a property we exploit by considering two additional exercises. First, we find that doubling the standard deviation of the idiosyncratic productivity shock, σ_ω, lowers the cumulative effect of the nominal shock to 3.6% (from 6.7%). Second, we make the set of demand functions steeper by doubling the value of b; this lowers the cumulative effect to 5.25%. These predictions are consistent with Boivin et al. (2009) and Kaufmann and Lein (2013) who empirically find that monetary non-neutrality decreases with both idiosyncratic volatility and competition.

6 Conclusion

In this paper we show how firms' specification doubts about their perceived model of demand leads to a novel theory of price stickiness. We find strong empirical support for the theory by subjecting the mechanism to a rich set of micro-level implications. The parsimony and quantitative relevance of the mechanism make it a promising step towards building macroeconomic models that can be used for counterfactual analysis.

References

Alvarez, F. E., F. Lippi, and L. Paciello (2011): "Optimal Price Setting With Observation and Menu Costs," *The Quarterly Journal of Economics*, 126, 1909–1960.

Argente, D. and C. Yeh (2017): "Product's Life Cycle, Learning, and Nominal Shocks," Manuscript, Minneapolis Fed.

Bachmann, R. and G. Moscarini (2011): "Business cycles and endogenous uncertainty," Manuscript, Yale University.

Baley, I. and J. A. Blanco (2018): "Firm uncertainty cycles and the propagation of nominal shocks," *AEJ: Macroeconomics*, forthcoming.

Ball, L. and D. Romer (1990): "Real rigidities and the non-neutrality of money," *The Review of Economic Studies*, 57, 183–203.

Bergemann, D. and K. Schlag (2011): "Robust monopoly pricing," *Journal of Economic Theory*, 146, 2527–2543.

Bergemann, D. and J. Valimaki (2008): "Bandit problems," The New Palgrave Dictionary of Economics, 2nd ed. Macmillan Press.

Bils, M. and P. Klenow (2004): "Some evidence on the importance of sticky prices," *Journal of Political Economy*, 112, 947–985.

Boivin, J., M. P. Giannoni, and I. Mihov (2009): "Sticky prices and monetary policy: Evidence from disaggregated US data," *American Economic Review*, 99, 350–84.

Bonomo, M. and C. Carvalho (2004): "Endogenous time-dependent rules and inflation inertia," *Journal of Money, Credit and Banking*, 1015–1041.

Bronnenberg, B., M. Kruger, and C. Mela (2008): "Database paper: The IRI Marketing Data Set," *Marketing Science*, 27, 745–748.

Campbell, J. R. and B. Eden (2014): "Rigid prices: Evidence from US scanner data," *International Economic Review*, 55, 423–442.

Christiano, L., M. Eichenbaum, and C. Evans (2005): "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 113.

Dow, J. and S. Werlang (1992): "Uncertainty Aversion, Risk Aversion, and the Optimal Choice of Portfolio," *Econometrica*, 60, 197–204.

Dupraz, S. (2016): "A Kinked-Demand Theory of Price Rigidity," Mimeo, Banque de France.

Eichenbaum, M., N. Jaimovich, and S. Rebelo (2011): "Reference Prices, Costs, and Nominal Rigidities," *American Economic Review*, 101, 234–62.

Ellsberg, D. (1961): "Risk, Ambiguity, and the Savage Axioms," *The Quarterly Journal of Economics*, 643–669.

Epstein, L. G. and M. Schneider (2003): "Recursive Multiple-Priors," *Journal of Economic Theory*, 113, 1–31.

Gagnon, E., D. López-Salido, and N. Vincent (2012): "Individual Price Adjustment along the Extensive Margin," *NBER Macroeconomics Annual*, 27, 235–281.

Gaur, V., S. Kesavan, A. Raman, and M. L. Fisher (2007): "Estimating demand uncertainty using judgmental forecasts," *Manufacturing & Service Operations Management*, 9, 480–491.

Gilboa, I. and D. Schmeidler (1989): "Maxmin Expected Utility with Non-unique Prior," *Journal of Mathematical Economics*, 18, 141–153.

Golosov, M. and R. E. Lucas (2007): "Menu costs and phillips curves," *Journal of Political Economy*, 115, 171–199.

Handel, B. R. and K. Misra (2015): "Robust new product pricing," *Marketing Science*, 34, 864–881.

Handel, B. R., K. Misra, and J. W. Roberts (2013): "Robust firm pricing with panel data," *Journal of Econometrics*, 174, 165–185.

Hansen, L. P. (2014): "Nobel lecture: uncertainty outside and inside economic models," *Journal of Political Economy*, 122, 945–967.

Kaufmann, D. and S. M. Lein (2013): "Sticky prices or rational inattention–What can we learn from sectoral price data?" *European Economic Review*, 64, 384–394.

Kehoe, P. and V. Midrigan (2015): "Prices are sticky after all," *Journal of Monetary Economics*, 75, 35–53.

Kimball, M. S. (1995): "The Quantitative Analytics of the Basic Neomonetarist Model," *Journal of Money, Credit, and Banking*, 27.

Klenow, P. J. and O. Kryvtsov (2008): "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent US Inflation?" *The Quarterly Journal of Economics*, 863–904.

Klenow, P. J. and B. A. Malin (2010): "Microeconomic Evidence on Price-Setting," *Handbook of Monetary Economics*, 3, 231–284.

Klenow, P. J. and J. L. Willis (2016): "Real rigidities and nominal price changes," *Economica*, 83, 443–472.

Knotek, I. and S. Edward (2010): "A Tale of Two Rigidities: Sticky Prices in a Sticky-Information Environment," *Journal of Money, Credit and Banking*, 42, 1543–1564.

Mackowiak, B. and M. Wiederholt (2009): "Optimal Sticky Prices under Rational Inattention," *American Economic Review*, 99, 769–803.

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.

Matějka, F. (2015): "Rationally inattentive seller: Sales and discrete pricing," *The Review of Economic Studies*, 83, 1125–1155.

Midrigan, V. (2011): "Menu costs, multiproduct firms, and aggregate fluctuations," *Econometrica*, 79, 1139–1180.

Mongey, S. (2018): "Market structure and monetary non-neutrality," Working Paper.

Nakamura, E. and J. Steinsson (2008): "Five facts about prices: A reevaluation of menu cost models," *The Quarterly Journal of Economics*, 123, 1415–1464.

Raman, A., M. Fisher, and A. McClelland (2001): *Supply chain management at World Co., Ltd*, Harvard Business School Boston, MA.

Rasmussen, C. E. and C. K. Williams (2006): *Gaussian processes for machine learning*, vol. 1, MIT press Cambridge.

Reis, R. (2006): "Inattentive producers," *The Review of Economic Studies*, 73, 793–821.

Rothschild, M. (1974): "A two-armed bandit theory of market pricing," *Journal of Economic Theory*, 9, 185–202.

Sims, C. A. (2003): "Implications of rational inattention," *Journal of monetary Economics*, 50, 665–690.

Stevens, L. (2014): "Coarse Pricing Policies," Manuscript, Univ. of Maryland.

Stigler, G. J. (1947): "The kinky oligopoly demand curve and rigid prices," *The Journal of Political Economy*, 432–449.

Stiglitz, J. E. (1979): "Equilibrium in product markets with imperfect information," *The American Economic Review*, 339–345.

Vavra, J. (2014): "Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation," *The Quarterly Journal of Economics*, 129, 215–258.

Woodford, M. (2003): "Imperfect Common Knowledge and the Effects of Monetary Policy," *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, 25.

——— (2009): "Information-constrained state-dependent pricing," *Journal of Monetary Economics*, 56, S100–S124.

Online Appendix

Table of Contents

A Appendix for Section 3 **55**

A.1 Updating with more observed prices 55

A.2 Proof of Proposition 1 56

A.3 Forward looking behavior 57

B Appendix for Section 4 **72**

B.1 Empirical link between aggregate and industry prices 72

B.2 Joint uncertainty over demand shape and relative price 73

B.3 Proofs on learning and nominal rigidity 76

C Appendix for Section 5 **80**

C.1 Dispersion of forecasts 80

C.2 Simulated hazards 81

C.3 Additional evidence on hazard functions 83

C.4 Constructing the typical history of observations 84

C.5 Speed of learning 85

A Appendix for Section 3

A.1 Updating with more observed prices

We can readily expand the updating formulas that we have developed in Section 3.2 for one observed price. Assume that firm has seen a whole vector of T previous signals, $\mathbf{y_0}$, with the corresponding vectors of prices $\mathbf{p_0}$ and number of times $\mathbf{N_0}$. The joint distribution with demand at any price p is again jointly Normal

$$\begin{bmatrix} x(p) \\ \mathbf{y_0} \end{bmatrix} \sim N\left(\begin{bmatrix} m(p) \\ m(\mathbf{p_0}) \end{bmatrix}, \Sigma(p, \mathbf{p_0})\right)$$

with

$$\Sigma(p, \mathbf{p_0}) = \begin{bmatrix} \sigma_x^2 & (\sigma_x^2, \ldots, \sigma_x^2) \\ (\sigma_x^2, \ldots, \sigma_x^2)' & \Sigma_x + diag(\frac{\sigma_z^2}{\mathbf{N_0}}) \end{bmatrix}$$

where $(\sigma_x^2, \ldots, \sigma_x^2)$ is a $1xT$ vector, and Σ_x is a TxT matrix with all entries equal to σ_x^2.

The resulting conditional expectation follows from applying the standard formula for conditional Normal expectations:

$$E(x(p)|\mathbf{y_0}) = m(p) + [\sigma_x^2, \ldots, \sigma_x^2](\Sigma_x + diag(\frac{\sigma_z^2}{\mathbf{N_0}}))^{-1}(\mathbf{y_0} - m(\mathbf{p0}))$$

The conditional expectation is again linear in the prior and a weighted sum of the demeaned signals. Expanding the above formula, we obtain

$$E(x(p)|\mathbf{y_0}) = m(p) + \alpha_0(y_{0,1} - m(p_{0,1})) + \cdots + \alpha_T(y_{0,T} - m(p_{0,T}))$$

where $y_{0,i}$ is the i-th element of the vector $\mathbf{y_0}$, and $\alpha_i \in (0, 1)$ is the i-th element of the $1xT$ vector $[\sigma_x^2, \ldots, \sigma_x^2](\Sigma_x + diag(\frac{\sigma_z^2}{\mathbf{N_0}}))^{-1}$.

Without loss of generality, assume the prices in $\mathbf{p_0}$ are sorted in ascending order, with the last element being the largest price. In building the worst case expectation, one can work from right to left and start with $p_t > p_{0,T}$, where $p_{0,i}$ denotes the i-th element of $\mathbf{p}_0$. The firm wants $m^*(p_t)$ to be the lowest possible so it sets it equal to the lower bound of the prior set, but sets the priors on all observed signals to their largest admissible value, so that for $p < p_t$

$$m^*(p) = \min\left[\gamma - bp, -\gamma - bp_t + (b + \delta)(p_t - p)\right] \tag{38}$$

Next consider, $p_t \in (p_{0,T-1}, p_{0,T}]$. The worst case $m^*(p_t)$ is again at the lower bound of the admissible set. Similarly, the worst-case is that all observe signals imply negative news about expected demand, hence it sets the prior at lower priors, $m^*(p|p < p_t)$, accordingly to the highest possible derivative $(b + \delta)$, and the prior at higher prices, $m^*(p|p > p_t)$ according to the lowest possible derivative. As a result

$$m(p;p_t) = \begin{cases} \min\left[\gamma - bp, -\gamma - bp_t + (b+\delta)(p_t - p)\right] & \text{for } p < p_t \\ \min\left[\gamma - bp, -\gamma - bp_t + (b-\delta)(p_t - p)\right] & \text{for } p \geq p_t \end{cases}$$

We can now confirm that there is a kink around any $p \in \mathbf{p}_0$, with the same properties as in the case of one previously observed price, which was analyzed in the main text.

A.2 Proof of Proposition 1

Proposition 1. *Define $\delta^* = \delta \operatorname{sgn}(p_t - p_0)$. For a given realization of c_t, the difference in worst-case expected profits at p_t and p_0, up to a first-order approximation around p_0, is*

$$\ln \upsilon^*(\varepsilon^{t-1}, c_t, p_t) - \ln \upsilon_0^*(\varepsilon^{t-1}, c_t, p_0) \approx \left[\frac{e^{p_0}}{e^{p_0} - e^{c_t}} - (b + \alpha_{t-1}(p_0)\delta^*)\right](p_t - p_0).$$

Proof. Consider $\ln \upsilon^*(\varepsilon^{t-1}, c_t, p_t)$ at some $p_t \in [p_0 - \frac{2\gamma}{\delta}, p_0 + \frac{2\gamma}{\delta}]$. When $p_t > p_0$, we have

$$\ln(e^{p_t} - e^{c_t}) + \left\{-\gamma - bp_t + \alpha_{t-1}(p_t)\widehat{z}_0 - \alpha_{t-1}(p_t)\delta\,(p_t - p_0) + .5\widehat{\sigma}_{t-1}^2(p_t) + .5\sigma_z^2\right\}, \tag{39}$$

while at $p_t < p_0$, this equals

$$\ln(e^{p_t} - e^{c_t}) + \left\{-\gamma - bp_t + \alpha_{t-1}(p_t)\widehat{z}_0 + \alpha_{t-1}(p_t)\delta\,(p_t - p_0) + .5\widehat{\sigma}_{t-1}^2(p_t) + .5\sigma_z^2\right\}. \tag{40}$$

In turn, $\ln \upsilon^*(\varepsilon^{t-1}, c_t, p_0)$ equals

$$\ln(e^{p_t} - e^{c_t}) + \left\{-\gamma - bp_0 + \alpha_{t-1}(p_t)\widehat{z}_0 + .5\widehat{\sigma}_{t-1}^2(p_0) + .5\sigma_z^2\right\}.$$

Fix some c_t and take a first-order approximation of $\ln \upsilon^*(\varepsilon^{t-1}, c_t, p_t)$ with respect to p_t, evaluated at p_0. Since this function is not differentiable at p_0, we analyze its right and left derivative. The former derivative equals

$$\frac{e^{p_0}}{e^{p_0} - e^{c_t}} - b - \alpha_{t-1}(p_0)\delta + \frac{\partial \alpha_{t-1}(p_t)}{\partial p_t}\left[\widehat{z}_0 - \delta\,(p_t - p_0)\right] + .5\frac{\partial \widehat{\sigma}_{t-1}^2(p_t)}{\partial p_t} \tag{41}$$

where the partial derivatives $\frac{\partial \alpha_{t-1}(p_t)}{\partial p_t}$ and $\frac{\partial \widehat{\sigma}_{t-1}^2(p_t)}{\partial p_t}$ are evaluated locally at p_0. In particular, given that

$$\alpha_{t-1}(p_t) = \frac{\sigma_x^2}{\sigma_x^2 + \sigma_z^2/N_0} e^{-\psi(p_t - p_0)^2};\ \widehat{\sigma}_{t-1}^2(p_t) = \sigma_x^2(1 - \alpha_{t-1}(p_t)),$$

then these two functions are differentiable p_0, with marginal effects equal to zero at p_0. Therefore, the local approximation to the right of p_0 simplifies to

$$\frac{e^{p_0}}{e^{p_0}-e^{c_t}} - \left[b+\alpha_{t-1}(p_0)\delta\right]. \tag{42}$$

The first term in the brackets reflects the effect of changing the price on profits, while the second captures the movement of demand along a curve with elasticity $-b$. The third term arises from the effect of demand of moving along a steeper demand curve, which is a characteristic of the worst-case belief about the demand elasticity.

Therefore, we obtain the local approximation to the right of p_0

$$\ln \upsilon^*(\varepsilon^{t-1}, c_t, p_t) - \ln \upsilon_0^*(\varepsilon^{t-1}, c_t, p_0) \approx \left[\frac{e^{p_0}}{e^{p_0}-e^{c_t}} - (b+\alpha_{t-1}(p_0)\delta)\right](p_t - p_0) \tag{43}$$

A similar derivation follows for the derivative to the left of p_0, where we obtain

$$\frac{e^{p_0}}{e^{p_0}-e^{c_t}} - \left[b-\alpha_{t-1}(p_0)\delta\right]$$

and therefore the local approximation to the left of p_0 is simply

$$\ln \upsilon^*(\varepsilon^{t-1}, c_t, p_t) - \ln \upsilon_0^*(\varepsilon^{t-1}, c_t, p_0) \approx \left[\frac{e^{p_0}}{e^{p_0}-e^{c_t}} - (b-\alpha_{t-1}(p_0)\delta)\right](p_t - p_0) \tag{44}$$

We obtain the result in Proposition 1 by putting together equations (43) and (44) and using the signum function to define $\delta^* = \delta \operatorname{sgn}(p_t - p_0)$. □

A.3 Forward looking behavior

We solve the recursive optimization problem in two steps. First, we compute the value function at time $t+1$. The key insight is that from this point onward the firm solves a series of static maximization problems because the endogenous state variable, the information set ε^t, remains the same from period to period. Still, the firm faces a dynamic, recursive problem because of the law of motion of the exogenous state variable, the cost shock c_t, which evolves according to its law of motion $g(c_{t+1}|c_t)$. Hence, the value function at $t+1$, which we label with $\tilde{V}(.)$ to differentiate from the time-t value function $V(.)$, is given by

$$\tilde{V}(\varepsilon^t, c_{t+1}) = \max_{p_{t+1}} \min_{m(p)\in\Upsilon_0} E\left[\nu(\varepsilon_{t+1}, c_{t+1}) + \beta \int \tilde{V}(\varepsilon^t, c_{t+2}) g(c_{t+2}|c_{t+1}) dc_{t+2} \middle| \varepsilon^t\right]$$

Since the information set is not growing over time, the state space for this problem is finite and tractable. As a result, we can solve for $\tilde{V}(\varepsilon^t, c_{t+1})$ through standard techniques and use it as the continuation value perceived by the firm at time t:

$$V(\varepsilon^{t-1},c_t)=\max_{p_t}\min_{m(p)\in\Upsilon_0}E\left[\nu(\varepsilon_t,c_t)+\beta\int\tilde{V}(\varepsilon^t,c_{t+1})g(c_{t+1}|c_t)dc_{t+1}\bigg|\varepsilon^{t-1}\right] \tag{45}$$

s.t.

$$\varepsilon^t=\{\varepsilon^{t-1},p_t,q_t\}.$$

Thus, at time t the firm fully takes into account that p_t, and the resulting new demand signal q_t, will serve as informative signals for future profit-maximization decisions. Importantly, this information is useful not only in the very next period, but propagates through the infinite future according to the law of motion of c_t.

For the following analytical results we work with the case where $\psi=\infty$ and the firm has perfect foresight on future costs, s.t. $c_{t+k}=c$ for all $k\geq 1$, for some constant c. In this case, the time $t+1$ value function is just the present discounted value of worst-case expected profits when the cost shock equals c:

$$\tilde{V}(\varepsilon^t,c)=\frac{\max_p\min_{m(p)\in\Upsilon_0}E\left[\nu(\varepsilon_{t+1},c)\Big|\varepsilon^t\right]}{1-\beta}$$

Hence, the only remaining uncertainty in $\tilde{V}(.)$ from the perspective of time t is the uncertainty about the realization of the time t signal q_t. Next, we turn to characterizing the expectation of $\tilde{V}$, given the time t information set ε^{t-1}.

For all analytical results below, we assume that (i) $\psi\rightarrow\infty$ and (ii) there is perfect foresight on future costs so that $c_{t+k}=c$ for some c.

Exploration makes prices more flexible when ε^{t-1} ***contains demand observations at only one previous price*** p_0

We start with the case where the time t information set, ε^{t-1}, contains only one price point, p_0, observed N_0 times with an average signal q_0. To be specific, call that information set ε^0. We will assume that the realization of the signal q_0 is good enough, so that when $c=\bar{c}_0=p_0-\ln(\frac{b}{b-1})$, p_0 is not just locally optimal (recall Corollary 1), but that it is the global maximizer conditional on ε^{t-1}. The relevant condition is

$$\widehat{z}_0=q_0-(-\gamma-bp_0)>\frac{\sigma_x^2}{2},$$

in which case

$$p_0=\arg\max_p\min_{m(p)\in\Upsilon_0}E\left[\nu(\varepsilon_{t+1},\bar{c}_0)\Big|\varepsilon^0,m(p)\right]$$

Hence in the absence of any new information, in future periods the firm will optimally set p_0, since it essentially faces a static problem with marginal cost equal to $\bar{c}_0$. The signal pair $\{p_t,q_t\}$

provides such new information and could lead to a different optimal action p_{t+k}.

Our first result is a characterization of the current price p_t that maximizes the expected continuation value when $c = \bar{c}_0$. It turns out that when the firm has collected prior information about demand only at p_0, then even even at that value of the cost the optimal exploration strategy is to deviate from p_0.

Proposition 4. *The expected continuation value* $E\left[\tilde{V}(\{\varepsilon^0, p_t, q_t\}, \bar{c}_0) \middle| \varepsilon^0, p_t\right]$ *achieves its maximum at*

$$p_t^* = \arg\min_p (p - p_0)^2 \ s.t. \ p \neq p_0.$$

Proof. In order to simplify notation, throughout the proofs we will use the standard expectation notation $E(.)$ to define the worst-case expectation of the firm.

The limiting case $\psi \to \infty$ simplifies the construction of the worst-case expected demand because $corr(x(p), x(p')) = 0$ for all $p \neq p'$. Thus, when updating beliefs about demand at any price p, only past signals observed at that particular price p matter. For future reference, it will be convenient to define the following notation for signal-to-noise ratios that will show up repeatedly

$$\alpha_0 \equiv \alpha_{t-1}(p_0; p_0) = \frac{\sigma_x^2}{\sigma_x^2 + \sigma_z^2/N_0}$$

$$\alpha_{t|0} \equiv \alpha_t(p_0; p_0 | p_t = p_0) = \frac{\sigma_x^2}{\sigma_x^2(N_0 + 1) + \sigma_z^2}$$

$$\alpha_t \equiv \alpha_{t-1}(p_t; p_t | p_t \neq p_0) = \frac{\sigma_x^2}{\sigma_x^2 + \sigma_z^2}$$

where the first is the signal-to-noise ratio of the signal q_0 conditional on ε^0 information, $\alpha_{t|0}$ and α_t are the (recursive) signal-to-noise ratios applicable to the new signal q_t given the signal q_0, in the two cases where $p_t = p_0$ and $p_t \neq p_0$ respectively. Since $p_0 = \ln(\frac{b}{b-1}) + \bar{c}_0$, it is the optimal myopic price for $c_{t+k} = \bar{c}_0$, which is the relevant case in the future. Thus, if its information set does not change, the firm will price $p_{t+k} = p_0$ in the future. The information set changes, of course, as a function of the current period pricing choice p_t and the resulting new signal q_t. For convenience, define the perceived innovations in the existing signal q_0 and the new signal q_t as

$$\widehat{z}_0 \equiv q_0 - (-\gamma - bp_0)$$

$$\widehat{z}_t \equiv q_t - (-\gamma - bp_t)$$

and the variance adjusted innovation of q_0 as

$$\widetilde{z}_0 \equiv \widehat{z}_0 - \frac{1}{2}\sigma_x^2.$$

Observe that since $c_{t+k} = \bar{c}_0$ with probability one, the only uncertainty over future profits is in the innovation of the new signal $\widehat{z}_t$. Hence, the expected continuation value is simply the expected discounted value of a stream of worst-case static profits at $c_{t+k} = \bar{c}_0$, after taking the expectation over the unknown $\widehat{z}_t$: $E\left[\tilde{V}(\{\varepsilon^0, p_t, q_t\}, \bar{c}_0)|\varepsilon^0, p_t\right] = \frac{\beta}{1-\beta}E\left[E(\nu(p^*_{t+k}, \bar{c}_0)|\{\varepsilon^0, p_t, q_t\})\Big|\varepsilon^0, p_t\right] = \frac{\beta}{1-\beta}E\left[\nu^*_{t+k}(p^*_{t+k}, \bar{c}_0)\Big|\varepsilon^0, p_t\right]$, where p^*_{t+k} is the resulting static optimal price, given the updated information set $\{\varepsilon^0, p_t, q_t\}$.

If $p_t = p_0$, this optimal price is still p_0 unless the information in the new signal q_t is particularly bad and sufficiently erodes the firm's beliefs about profits at p_0, in which case the firm switches to the interior optimal price p^{int}_{t+k} – the ex-ante second best option. To find this interior optimum, note that for all prices $p_{t+k} \neq p_0$ the worst-case demand is simply

$$\widehat{x}^*_t(p_{t+k}; m^*(p; p_{t+k})) = -\gamma - bp$$

hence the interior optimal price is

$$p^{int}_{t+k} = \min\{p|(p - p_0)^2 > 0\},$$

which gets you as close as possible the to optimal markup $\frac{b}{b-1}$ while still staying on the smooth portion of the firm's demand curve (recall: there is a kink in the worst-case belief at p_0, but is smooth everywhere else). Thus, if $p_t = p_0$, optimal p^*_{t+k} is equal to p_0 unless $\widehat{z}_t < \underline{z}_0$, where $\underline{z}_0$ is such that:

$$\frac{E_{t-1}(\nu^*_{t+k}(p_0, \bar{c}_0)|\varepsilon^0, p_t = p_0, \widehat{z}_t = \underline{z}_0)}{\lim_{p \to p_0} E(\nu^*_{t+k}(p, \bar{c}_0)|\varepsilon^0, p_t = p_0, \widehat{z}_t = \underline{z}_0)} = 1$$

Substituting in the relevant expressions and simplifying, we can derive

$$\underline{z}_0 = \frac{\sigma^2_x}{2}(1 - \alpha_0) - \frac{\alpha(p_0)}{\alpha_{t|0}}\widetilde{z}_0.$$

Hence if $p_t = p_0$, the optimal p^*_{t+k} is equal to p_0 as long as the innovation in the new signal is good enough – namely $\widehat{z}_t \geq \underline{z}_0$.

If $p_t \neq p_0$, p_0 remains the optimal price at $t + k$ unless the new signal q_t is *good enough* to convince the firm to deviate from its ex-ante optimum p_0 and move to the newly observed p_t itself. In the limiting case $\psi \to \infty$ we know that the only potential alternative is p_t, because q_t does not update beliefs anywhere else, and hence p_0 dominates all other prices. In particular, for every possible p_t there is an upper threshold for the innovation in q_t, such that $p^*_{t+k} = p_t$ if and only if $\widehat{z}_t > \bar{z}(p_t)$. This threshold $\bar{z}(p_t)$ satisfies

$$\frac{E(\nu^*_{t+k}(p_t, \bar{c}_0)|\varepsilon^0, p_t \neq p_0, \widehat{z}_t = \bar{z}(p_t))}{E(\nu^*_{t+k}(p_0, \bar{c}_0)|\varepsilon^0, p_t \neq p_0, \widehat{z}_t = \bar{z}(p_t))} = 1$$

Substituting in the respective expressions, and simplifying we can derive:

$$\bar{z}(p_t) = \frac{\alpha_0}{\alpha_t}\widetilde{z}_0 + \frac{\sigma_x^2}{2} - \frac{1}{\alpha_t}\left[\ln\left(\frac{\exp(p_t) - \exp(\bar{c}_0)}{\exp(p_0) - \exp(\bar{c}_0)}\right) + b(p_0 - p_t)\right]$$

With the two thresholds thusly characterized, we can conclude that the optimal pricing policy at time $t+k$ is given by:

$$p_{t+k}^* = \begin{cases} p_0 & \text{if } p_t = p_0 \text{ and } \widehat{z}_t \geq \underline{z}_0 \text{ or } p_t \neq p_0 \text{ and } \widehat{z}(p_t) \leq \bar{z}(p_t) \\ p_t & \text{if } p_t \neq p_0 \text{ and } \widehat{z}_t > \bar{z}(p_t) \\ p_{t+k}^{int} & \text{if } p_t = p_0 \text{ and } \widehat{z}_t < \underline{z}_0 \end{cases}$$

We can then evaluate the expected continuation value $E\left[\tilde{V}(\{\varepsilon^0, p_t, q_t\}, \bar{c}_0)\middle|\varepsilon^0, p_t\right]$ – we do so separately for the cases $p_t = p_0$ and $p_t \neq p_0$, since the expected continuation value (which we will denote by the short-hand $E_{t-1}(\tilde{V})$ to save space) is potentially discontinuous at $p_t = p_0$, so that $E_{t-1}(\tilde{V}|p_t = p_0) =$

$$= \Phi(\frac{\underline{z}_0}{\sqrt{\sigma_x^2(1-\alpha_0) + \sigma_z^2}})(\exp(p_0) - \exp(\bar{c}_0))\exp(-\gamma - bp_0 + \frac{1}{2}(\sigma_x^2 + \sigma_z^2))$$

$$+ (1 - \Phi(\frac{\underline{z}_0}{\sqrt{\sigma_x^2(1-\alpha_0) + \sigma_z^2}}))(\exp(p_0) - \exp(\bar{c}_0))\exp(-\gamma - bp_0 + \alpha_0\widehat{z}_0 + \frac{1}{2}(\sigma_x^2(1-\alpha_0) + \sigma_z^2))\frac{\Phi(\frac{\alpha_{t|0}(\sigma_x^2(1-\alpha_0)+\sigma_z^2) - \underline{z}_0}{\sqrt{\sigma_x^2(1-\alpha_0)+\sigma_z^2}})}{1 - \Phi(\frac{\underline{z}_0}{\sqrt{\sigma_x^2(1-\alpha_0)+\sigma_z^2}})}$$

$$= (\exp(p_0) - \exp(\bar{c}_0))\exp(-\gamma - bp_0\frac{1}{2}(\sigma_x^2 + \sigma_z^2))\left(\Phi(\frac{\alpha_{t|0}(\sigma_x^2(1-\alpha_0) + \sigma_z^2) - \underline{z}_0}{\sqrt{\sigma_x^2(1-\alpha_0) + \sigma_z^2}})\exp(\alpha_0\widetilde{z}_0) + \Phi(\frac{\underline{z}_0}{\sqrt{\sigma_x^2(1-\alpha_0) + \sigma_z^2}})\right)$$

while $E_{t-1}(\tilde{V}|p_t \neq p_0) =$

$$= P(\widehat{z}_t < \bar{z}(p_t))(\exp(p_0) - \exp(\bar{c}_0))\exp(-\gamma - bp_0 + \alpha_0\widehat{z}_0 + \frac{1}{2}(\sigma_x^2(1-\alpha_0) + \sigma_z^2))$$
$$+ P(\widehat{z}_t \geq \bar{z}(p_t))(\exp(p_t) - \exp(\bar{c}_0))\exp(-\gamma - bp_t + \frac{1}{2}(\sigma_x^2(1-\alpha_t) + \sigma_z^2))E(\exp(\alpha_t\widehat{z}_t)|\widehat{z}_t > \bar{z}(p_t))$$
$$= \Phi(\frac{\bar{z}(p_t)}{\sqrt{(\sigma_x^2 + \sigma_z^2)}})(\exp(p_0) - \exp(\bar{c}_0))\exp(-\gamma - bp_0 + \alpha_0\widehat{z}_0 + \frac{1}{2}(\sigma_x^2(1-\alpha_0) + \sigma_z^2))$$
$$+ \Phi(\frac{\alpha_t(\sigma_x^2 + \sigma_z^2) - \bar{z}(p_t)}{\sqrt{(\sigma_x^2 + \sigma_z^2)}})(\exp(p_t) - \exp(\bar{c}_0))\exp(-\gamma - bp_t + \frac{1}{2}(\sigma_x^2 + \sigma_z^2)))$$

where we use the fact that the firm perceives $\widehat{z}_t \sim N(0, \hat{\sigma}_{t-1}^2(p_t) + \sigma_z^2)$, and $\Phi(.)$ denotes the CDF of the standard normal distribution.

The first question of interest is if and when the expected continuation value is discontinuous at $p_t = p_0$. To answer this question, we evaluate the ratio $\frac{E_{t-1}(\tilde{V}|p_1 = p_0)}{\lim_{p_1 \to p_0} E_{t-1}(\tilde{V}|p_1 \neq p_0)}$. It is useful to first evaluate the denominator and collect terms, concluding that $\lim_{p_t \to p_0} E_{t-1}(\tilde{V}|p_t \neq p_0) =$

$$= (\exp(p_0) - \exp(\bar{c}_0))\exp(-\gamma - bp_0 + \frac{1}{2}(\sigma_x^2 + \sigma_z^2))\left(\Phi(\frac{\bar{z}(p_t)}{\sqrt{(\sigma_x^2 + \sigma_z^2)}})\exp(\alpha_0\widetilde{z}_0) + \Phi(\frac{\alpha_t(\sigma_x^2 + \sigma_z^2) - \bar{z}(p_t)}{\sqrt{(\sigma_x^2 + \sigma_z^2)}})\right)$$

It then follows that the ratio $\frac{E_{t-1}(\tilde{V}|p_t=p_0)}{\lim_{p_t\to p_0}E_{t-1}(\tilde{V}|p_t\neq p_0)} =$

$$= \frac{\Phi(\frac{\frac{\sigma_x^2}{2}(1-\alpha_0)+\frac{\alpha_0}{\alpha_{t|0}}\widetilde{z}_0}{\sqrt{\sigma_x^2(1-\alpha_0)+\sigma_z^2}})\exp(\alpha_0\widetilde{z}_0)+\Phi(\frac{(1-\alpha_0)\frac{\sigma_x^2}{2}-\frac{\alpha_0}{\alpha_{t|0}}\widetilde{z}_0}{\sqrt{\sigma_x^2(1-\alpha_0)+\sigma_z^2}})}{\Phi(\frac{\frac{\alpha_0}{\alpha_t}\widetilde{z}_0+\frac{\sigma_x^2}{2}}{\sqrt{(\sigma_x^2+\sigma_z^2)}})\exp(\alpha_0\widetilde{z}_0)+\Phi(\frac{\frac{\sigma_x^2}{2}-\frac{\alpha_0}{\alpha_t}\widetilde{z}_0}{\sqrt{(\sigma_x^2+\sigma_z^2)}})}$$

where we have substituted in the respective values of the thresholds $\underline{z}_0$ and $\bar{z}(p_t)$. The ratio limits to 1 as $\widetilde{z}_0\to\infty$, and it is below 1 at $\widetilde{z}_0=0$, as in this case

$$\frac{E_{t-1}(\tilde{V}|p_t=p_0)}{\lim_{p_t\to p_0}E_{t-1}(\tilde{V}|p_t\neq p_0)} = \frac{\Phi(\frac{\frac{\sigma_x^2}{2}(1-\alpha_0)}{\sqrt{\sigma_x^2(1-\alpha_0)+\sigma_z^2}})}{\Phi(\frac{\sigma_x^2}{2\sqrt{\sigma_x^2+\sigma_z^2}})} < 1$$

Next, we show that the derivative of the ratio in respect to $\widetilde{z}_0$ is positive for the relevant values $\widetilde{z}_0\geq 0$, which is enough to conclude that $\frac{E_{t-1}(\tilde{V}|p_t=p_0)}{\lim_{p_t\to p_0}E_{t-1}(\tilde{V}|p_t\neq p_0)}$ converges to 1 from below and hence is less than one for all finite $\widetilde{z}_0\geq 0$. The needed derivative,

$$\frac{\partial\frac{E_{t-1}(\tilde{V}|p_t=p_0)}{\lim_{p_t\to p_0}E(\tilde{V}|p_t\neq p_0)}}{\partial\widetilde{z}_0},$$

it is proportional to

$$\begin{aligned}
&\left(\underbrace{(\phi(\frac{\frac{\sigma_x^2}{2}(1-\alpha_0)+\frac{\alpha_0}{\alpha_{t|0}}\widetilde{z}_0}{\sqrt{\sigma_x^2(1-\alpha_0)+\sigma_z^2}})\exp(\alpha_0\widetilde{z}_0)-\phi(\frac{(1-\alpha_0)\frac{\sigma_x^2}{2}-\frac{\alpha_0}{\alpha_{t|0}}\widetilde{z}_0}{\sqrt{\sigma_x^2(1-\alpha_0)+\sigma_z^2}}))}_{=0}\frac{\alpha_0}{\alpha_{t|0}\sqrt{\sigma_x^2(1-\alpha_0)+\sigma_z^2}}+\Phi(\frac{\frac{\sigma_x^2}{2}(1-\alpha_0)+\frac{\alpha_0}{\alpha_{t|0}}\widetilde{z}_0}{\sqrt{\sigma_x^2(1-\alpha_0)+\sigma_z^2}})\exp(\alpha_0\widetilde{z}_0)\alpha_0\right)* \\
&\left(\Phi(\frac{\frac{\alpha_0}{\alpha_t}\widetilde{z}_0+\frac{\sigma_x^2}{2}}{\sqrt{(\sigma_x^2+\sigma_z^2)}})\exp(\alpha_0\widetilde{z}_0)+\Phi(\frac{\frac{\sigma_x^2}{2}-\frac{\alpha_0}{\alpha_t}\widetilde{z}_0}{\sqrt{(\sigma_x^2+\sigma_z^2)}})\right)-\left(\Phi(\frac{\frac{\sigma_x^2}{2}(1-\alpha_0)+\frac{\alpha_0}{\alpha_{t|0}}\widetilde{z}_0}{\sqrt{\sigma_x^2(1-\alpha_0)+\sigma_z^2}})\exp(\alpha_0\widetilde{z}_0)+\Phi(\frac{(1-\alpha_0)\frac{\sigma_x^2}{2}-\frac{\alpha_0}{\alpha_{t|0}}\widetilde{z}_0}{\sqrt{\sigma_x^2(1-\alpha_0)+\sigma_z^2}})\right)* \\
&\left(\underbrace{(\phi(\frac{\frac{\sigma_x^2}{2}+\frac{\alpha_0}{\alpha_t}\widetilde{z}_0}{\sqrt{\sigma_x^2+\sigma_z^2}})\exp(\alpha_0\widetilde{z}_0)-\phi(\frac{\frac{\sigma_x^2}{2}-\frac{\alpha_0}{\alpha_t}\widetilde{z}_0}{\sqrt{\sigma_x^2+\sigma_z^2}}))\frac{\alpha_0}{\alpha_1\sqrt{\sigma_x^2+\sigma_z^2}}}_{=0}+\Phi(\frac{\frac{\sigma_x^2}{2}+\frac{\alpha_0}{\alpha_1}\widetilde{z}_0}{\sqrt{\sigma_x^2+\sigma_z^2}})\exp(\alpha_0\widetilde{z}_0)\alpha_0\right) \\
&=\alpha_0\exp(\alpha_0\widetilde{z}_0)\left[\Phi(\frac{\frac{\sigma_x^2}{2}(1-\alpha_0)+\frac{\alpha_0}{\alpha_{t|0}}\widetilde{z}_0}{\sqrt{\sigma_x^2(1-\alpha_0)+\sigma_z^2}})\Phi(\frac{\frac{\sigma_x^2}{2}-\frac{\alpha_0}{\alpha_t}\widetilde{z}_0}{\sqrt{\sigma_x^2+\sigma_z^2}})-\Phi(\frac{\frac{\sigma_x^2}{2}(1-\alpha_0)-\frac{\alpha_0}{\alpha_{t|0}}\widetilde{z}_0}{\sqrt{\sigma_x^2(1-\alpha_0)+\sigma_z^2}})\Phi(\frac{\frac{\sigma_x^2}{2}+\frac{\alpha_0}{\alpha_t}\widetilde{z}_0}{\sqrt{\sigma_x^2+\sigma_z^2}})\right]
\end{aligned}$$

Thus, the derivative is positive if and only if

$$\frac{\Phi(\frac{\frac{\sigma_x^2}{2}-\frac{\alpha_0}{\alpha_t}\widetilde{z}_0}{\sqrt{\sigma_x^2+\sigma_z^2}})}{\Phi(\frac{\frac{\sigma_x^2}{2}+\frac{\alpha_0}{\alpha_t}\widetilde{z}_0}{\sqrt{\sigma_x^2+\sigma_z^2}})} > \frac{\Phi(\frac{\frac{\sigma_x^2}{2}(1-\alpha_0)-\frac{\alpha_0}{\alpha_{t|0}}\widetilde{z}_0}{\sqrt{\sigma_x^2(1-\alpha_0)+\sigma_z^2}})}{\Phi(\frac{\frac{\sigma_x^2}{2}(1-\alpha_0)+\frac{\alpha_0}{\alpha_{t|0}}\widetilde{z}_0}{\sqrt{\sigma_x^2(1-\alpha_0)+\sigma_z^2}})} \tag{46}$$

This inequality holds since

$$\frac{\Phi(\frac{\frac{\sigma_x^2}{2}-\frac{\alpha_0}{\alpha_t}\widetilde{z}_0}{\sqrt{\sigma_x^2+\sigma_z^2}})}{\Phi(\frac{\frac{\sigma_x^2}{2}+\frac{\alpha_0}{\alpha_t}\widetilde{z}_0}{\sqrt{\sigma_x^2+\sigma_z^2}})} > \frac{\Phi(\frac{\frac{\sigma_x^2}{2}-\frac{\alpha_0}{\alpha_{t|0}}\widetilde{z}_0}{\sqrt{\sigma_x^2+\sigma_z^2}})}{\Phi(\frac{\frac{\sigma_x^2}{2}+\frac{\alpha_0}{\alpha_{t|0}}\widetilde{z}_0}{\sqrt{\sigma_x^2+\sigma_z^2}})} > \frac{\Phi(\frac{\frac{\sigma_x^2}{2}(1-\alpha_0)-\frac{\alpha_0}{\alpha_{t|0}}\widetilde{z}_0}{\sqrt{\sigma_x^2(1-\alpha_0)+\sigma_z^2}})}{\Phi(\frac{\frac{\sigma_x^2}{2}(1-\alpha_0)+\frac{\alpha_0}{\alpha_{t|0}}\widetilde{z}_0}{\sqrt{\sigma_x^2(1-\alpha_0)+\sigma_z^2}})}$$

where the first inequality follows from $\alpha_{t|0} < \alpha_t$, and the second from the fact that

$$\frac{\partial \frac{\frac{\sigma_x^2}{2}(1-\tilde{\alpha}_0)-\frac{\alpha_0}{\alpha_{t|0}}\widetilde{z}_0}{\sqrt{\sigma_x^2(1-\tilde{\alpha}_0)+\sigma_z^2}}}{\partial\tilde{\alpha}_0} < \frac{\partial \frac{\frac{\sigma_x^2}{2}(1-\tilde{\alpha}_0)+\frac{\alpha_0}{\alpha_{t|0}}\widetilde{z}_0}{\sqrt{\sigma_x^2(1-\tilde{\alpha}_0)+\sigma_z^2}}}{\partial\tilde{\alpha}_0}$$

and the fact that the term

$$\frac{\partial\left(\frac{\Phi(\frac{\frac{\sigma_x^2}{2}(1-\tilde{\alpha}_0)-\frac{\alpha_0}{\alpha_{t|0}}\widetilde{z}_0}{\sqrt{\sigma_x^2(1-\tilde{\alpha}_0)+\sigma_z^2}})}{\Phi(\frac{\frac{\sigma_x^2}{2}(1-\tilde{\alpha}_0)+\frac{\alpha_0}{\alpha_{t|0}}\widetilde{z}_0}{\sqrt{\sigma_x^2(1-\tilde{\alpha}_0)+\sigma_z^2}})}\right)}{\partial\tilde{\alpha}_0}$$

equals

$$\phi(\frac{\frac{\sigma_x^2}{2}(1-\alpha_0)-\frac{\alpha_0}{\alpha_{t|0}}\widetilde{z}_0}{\sqrt{\sigma_x^2(1-\alpha_0)+\sigma_z^2}})\Phi(\frac{\frac{\sigma_x^2}{2}(1-\alpha_0)+\frac{\alpha_0}{\alpha_{t|0}}\widetilde{z}_0}{\sqrt{\sigma_x^2(1-\alpha_0)+\sigma_z^2}})\frac{\partial\frac{\frac{\sigma_x^2}{2}(1-\tilde{\alpha}_0)-\frac{\alpha_0}{\alpha_{t|0}}\widetilde{z}_0}{\sqrt{\sigma_x^2(1-\tilde{\alpha}_0)+\sigma_z^2}}}{\partial\tilde{\alpha}_0} - \Phi(\frac{\frac{\sigma_x^2}{2}(1-\alpha_0)-\frac{\alpha_0}{\alpha_{t|0}}\widetilde{z}_0}{\sqrt{\sigma_x^2(1-\alpha_0)+\sigma_z^2}})\phi(\frac{\frac{\sigma_x^2}{2}(1-\alpha_0)+\frac{\alpha_0}{\alpha_{t|0}}\widetilde{z}_0}{\sqrt{\sigma_x^2(1-\alpha_0)+\sigma_z^2}})\frac{\partial\frac{\frac{\sigma_x^2}{2}(1-\tilde{\alpha}_0)+\frac{\alpha_0}{\alpha_{t|0}}\widetilde{z}_0}{\sqrt{\sigma_x^2(1-\tilde{\alpha}_0)+\sigma_z^2}}}{\partial\tilde{\alpha}_0}$$
$$< 0$$

Thus, we can conclude that

$$\frac{E_{t-1}(\tilde{V}|p_t = p_0)}{\lim_{p_t\to p_0} E_{t-1}(\tilde{V}|p_t \neq p_0)} < 1$$

for all $\widetilde{z}_0 \geq 0$ meaning that there is discontinuous jump down in the continuation value at $p_t = p_0$.

Lastly, consider what value of p_t optimizes the expected continuation value. Since the discontinuity at p_0 (the only potential corner solution) is a jump down, the maximizing p_t must be the interior maximum, which satisfies the FOC condition that $\frac{\partial E_{t-1}(\tilde{V}|p_t\neq p_0)}{\partial p_t} = 0$. Taking the derivative, $\frac{\partial E_{t-1}(\tilde{V}|p_t\neq p_0)}{\partial p_t} =$

$$= \phi(\frac{\bar{z}(p_t)}{\sqrt{\sigma_x^2+\sigma_z^2}})(e^{p_0} - e^{\bar{c}_0})\exp(-\gamma - bp_0 + \alpha_0\widehat{z}_0 + \frac{1}{2}(\sigma_x^2(1-\alpha_0)+\sigma_z^2))\frac{\frac{\partial\bar{z}(p_t)}{\partial p_t}}{\sqrt{\sigma_x^2+\sigma_z^2}}$$
$$- \phi(\frac{\alpha_t(\sigma_x^2+\sigma_z^2)-\bar{z}(p_t)}{\sqrt{(\sigma_x^2+\sigma_z^2)}})(e^{p_t} - e^{\bar{c}_0})\exp(-\gamma - bp_t + \frac{1}{2}(\sigma_x^2(1-\alpha_t)+\sigma_z^2+\alpha_t^2(\sigma_x^2+\sigma_z^2)))\frac{\frac{\partial\bar{z}(p_t)}{\partial p_t}}{\sqrt{\sigma_x^2+\sigma_z^2}}$$
$$+ \Phi(\frac{\alpha_t(\sigma_x^2+\sigma_z^2)-\bar{z}(p_t)}{\sqrt{(\sigma_x^2+\sigma_z^2)}})\exp(-\gamma - bp_t + \frac{1}{2}(\sigma_x^2(1-\alpha_t)+\sigma_z^2+\alpha_t^2(\sigma_x^2+\sigma_z^2)))(e^{p_t} - b(e^{p_t} - e^{\bar{c}_0}))$$

The above expression limits to zero as $p_t \to p_0$. To see that, note that $\lim_{p_t \to p_0} \frac{\partial \bar{z}(p_t)}{\partial p_t} = 0$, thus the first 2 terms of the FOC expression above fall out. For the last term, using $p_0 = \ln(\frac{b}{b-1}) + c_0$ it follows that

$$(e^{p_0} - b(e^{p_0} - e^{\bar{c}_0})) = \frac{b}{b-1} e^{\bar{c}_0} - \frac{b}{b-1} e^{\bar{c}_0} = 0$$

Therefore, we can conclude that $\lim_{p_t \to p_0} \frac{\partial E_{t-1}(\tilde{V}|p_t \neq p_0)}{\partial p_t} = 0$, and thus the interior maximum of the expected continuation value is $p_t \to p_0$. □

Intuitively, $p_t^* = \arg\min_t (p - p_0)$ s.t. $p \neq p_0$, ensures that the new signal q_t will be informative about a price as close as possible to the ex-ante expected optimal p_0, and thus achieves almost the same markup – this makes the new information highly relevant. As a result, if the realization of $\widehat{z}_t$ happens to be good enough, i.e. $\widehat{z}_t$ is above a threshold $\bar{z}_t(p_t^*)$ that is characterized in the proof above, then the firm will stick with this price in the future, set $p_{t+k} = p_t^*$, and take advantage of the unexpectedly high demand at that price. On the other hand, if the signal realization happens to be bad, the firm can safely switch back to the ex-ante optimal p_0, where the belief about demand is not affected by $\widehat{z}_t$, and still offers lower uncertainty and a good perceived markup.

The reason for not picking $p_t = p_0$ is that a bad signal realization at p_0 erodes the ex-ante best available pricing option, p_0, and at the same time the firm does not have a good fall-back alternative, as it has no observations of demand at other prices. If in that case the realization of $\widehat{z}_t$ falls below the threshold $\underline{z}_0$, the news about $x(p_0)$ is bad enough to incentivize the firm to set p_{t+k} to a previously unvisited price. Due to this downside risk at p_0, there is a first-order gain of obtaining information at a new price, which manifests in the discontinuous jump down in the expected continuation value at p_0.

As shown in Proposition 4, the best forward-looking strategy is therefore to experiment by posting a new price. This exploration incentive could potentially overturn the rigidity result implied by the static maximization pricing choice analyzed earlier, but as we show next it turns out that this results is specific to the firm having seen only one price in the past. In more general situations, when the firm has seen more than one distinct price point in the past, forward-looking behavior can in fact *reinforce* the static rigidity incentives.

Exploration makes prices stickier, when ε^t contains observations at multiple prices

Proposition 5. *There is a non-singleton interval of costs $(\underline{c}, \bar{c})$ around $\bar{c}_0$, and a threshold $\chi > 0$, such that if $\widehat{z} > \chi$, then for any $c \in (\underline{c}, \bar{c})$:*

$$p_0 = \arg\max_{p_t} E\left[\tilde{V}(\{\varepsilon^1, p_t, q_t\}, c) \middle| \varepsilon^1, p_t\right].$$

Moreover, the threshold χ is decreasing in $|p_1 - p_0|$.

Proof. The proof follows a similar logic as the previous one. First, we characterize the optimal p_{t+k} for $c = \bar{c}_0$, but now conditional on ε^1, and then use it to compute the expected continuation value and show that it is maximized at $p_t = p_0$. Lastly, we appeal to continuity to conclude that $p_t = p_0$ is optimal for an interval of cost values around $\bar{c}_0$. In addition to the signal-to-noise ratio notation $\alpha_0, \alpha_{t|0}, \alpha_t$ defined in the previous proof, we define

$$\alpha_1 \equiv \alpha_{t-1}(p_1; p_1) = \frac{\sigma_x^2}{\sigma_x^2 + \sigma_z^2/N_1}$$

$$\alpha_{t|1} \equiv \alpha_t(p_1; p_1|p_t = p_1) = \frac{\sigma_x^2}{\sigma_x^2(N_1+1) + \sigma_z^2}$$

Similarly, we define the (variance corrected) innovation in the signal at p_1 as

$$\widetilde{z}_1 \equiv \widehat{z}_1 - \frac{1}{2}\sigma_x^2 = q_1 - (-\gamma - bp_1) - \frac{1}{2}\sigma_x^2$$

The optimal policy at $t+k$ follows a similar structure to the one described in the previous proof. Conditional on just ε^1 the optimal p_{t+k} is equal to p_0, and the way the new information contained in q_t affects the optimal p_{t+k} depends on the position of p_t. If $p_t = p_0$, then the firm stays at p_0 unless the new signal is too bad ($\widehat{z}_t < \underline{z}_0$). If $p_t = p_1$, then the firm *moves* to p_1 if the signal is good enough ($\widehat{z}_t > \bar{z}_1$) otherwise stays at p_0. And if $p_t \notin \{p_0, p_1\}$, then the firm again stays at p_0 unless the signal is too good, but compared to a different threshold: $\widehat{z}_t > \bar{z}(p_t)$. The key difference from the previous proof is what happens if $p_t = p_0$ and the signal is sufficiently bad to prompt a move ($\widehat{z}_t < \underline{z}_0$). There exists a $\chi_1 > 0$ such that if $\widehat{z}_1 > \chi_1$, then the firm does not move to the interior optimum p^{int}, but rather to p_1, which as another relatively good price at which the firm has built some information capital is a better option than the brand new p^{int} where the firm has not accumulated any information. To see this, note that

$$\frac{E(\nu^*_{t+k}(p_1, \bar{c}_0)|\varepsilon^1, p_t = p_0)}{\lim_{p \to p_0} E(\nu^*_{t+k}(p, \bar{c}_0)|\varepsilon^1, p_t = p_0)} = (b\exp(p_1 - p_0) - b + 1)\exp(-b(p_1 - p_0) + \alpha_1 \widetilde{z}_1) > 1$$

Note that the RHS is increasing in $\widetilde{z}_1$, and thus in $\widehat{z}_1$ and limits to infinity as $\widehat{z}_1 \to \infty$, hence there exists a constant $\chi_1 > 0$ such that the above ratio is strictly greater than one when $\widehat{z} > \chi_1$. For the rest of the proof we assume that $\widehat{z}_1 > \chi_1$ so that the above inequality holds. The relevant thresholds $\underline{z}_0, \bar{z}_1, \bar{z}(p_t)$ can be computed as before, by finding the value of the signal at which the firm is indifferent between p_0 and the respective alternative option:

$$\underline{z}_0 = \frac{\sigma_x^2}{2}(1 - \alpha_0) - \frac{1}{\alpha_{t|0}}(b(p_1 - p_0) - \ln(be^{(p_1 - p_0)} - b + 1))$$

$$\bar{z}_1 = \frac{\sigma_x^2}{2}(1 - \alpha_1) + \frac{1}{\alpha_{t|1}}(b(p_1 - p_0) - \ln(be^{(p_1 - p_0)} - b + 1))$$

$$\bar{z}(p_t) = \frac{\alpha_0}{\alpha_t}\widetilde{z}_0 + \frac{\sigma_x^2}{2} - \frac{1}{\alpha_t}\left[\ln\left(\frac{\exp(p_t) - \exp(\bar{c}_0)}{\exp(p_0) - \exp(\bar{c}_0)}\right) + b(p_0 - p_t)\right]$$

So the $t+k$ optimal pricing policy is:

$$p^*_{t+1} = \begin{cases} p_0 & \text{if } p_t = p_0 \text{ and } \widehat{z}_t \geq \underline{z}_0, \text{ or } p_t = p_1 \text{ and } \widehat{z}_t \leq \bar{z}_1 \text{ or } p_t \notin \{p_0, p_1\} \text{ and } \widehat{z}_t \leq \bar{z}(p_t) \\ p_1 & \text{if } p_t = p_1 \text{ and } \widehat{z}_t > \bar{z}_1 \text{ or } p_t = p_0 \text{ and } \widehat{z}_t < \underline{z}_0 \\ p_t & \text{if } p_t \notin \{p_0, p_1\} \text{ and } \widehat{z}_t > \bar{z}(p_t) \end{cases}$$

We can now use this result to characterize the expected continuation value and find its maximizer. Note that the value of p_t that maximizes $E(\left[\tilde{V}(\{\varepsilon^1, p_t, q_t\}, \bar{c}_0)\middle|\varepsilon^1, p_t\right])$ is either one of the two corner solutions p_0 and p_1, or the interior maximum. Moreover, we can appeal to the proof of Proposition 4 for the result that the expected continuation value achieves its interior maximum at the limit of $p_t \to p_0$. This follows because under $\psi \to \infty$ the additional signal q_1 only matters when updating beliefs at p_1 itself, hence at $p \neq p_1$ the expected continuation value is equivalent to the one conditional on ε^0, that we analyzed above. We proceed in two steps. First we show that the two corner solutions are in fact equivalent to each other, and then we conclude by showing that p_0 also dominates the interior solution p^{int}. The expected value $E(\left[\tilde{V}(\{\varepsilon^1, p_t, q_t\}, \bar{c}_0|\varepsilon^1, p_t = p_0)\right])$ is slightly different than before, because the fall back option (in case of a bad new signal q_t) is now p_1. Now, $E_{t-1}(\tilde{V}|p_t = p_0) =$

$$= \Phi(\frac{\underline{z}_0}{\sqrt{\sigma_x^2(1-\alpha_0)+\sigma_z^2}})(\exp(p_1) - \exp(\bar{c}_0))\exp(-\gamma - bp_1 + \alpha_1\widehat{z}_1 + \frac{1}{2}(\sigma_x^2(1-\alpha_1)+\sigma_z^2))$$
$$+ (1 - \Phi(\frac{\underline{z}_0}{\sqrt{\sigma_x^2(1-\alpha_0)+\sigma_z^2}}))(\exp(p_0) - \exp(\bar{c}_0))\exp(-\gamma - bp_0 + \alpha_0\widehat{z}_0 + \frac{1}{2}(\sigma_x^2(1-\alpha_0)+\sigma_z^2))\frac{\Phi(\frac{\alpha_{t|0}(\sigma_x^2(1-\alpha_0)+\sigma_z^2)-\underline{z}_0}{\sqrt{\sigma_x^2(1-\alpha_0)+\sigma_z^2}})}{1-\Phi(\frac{\underline{z}_0}{\sqrt{\sigma_x^2(1-\alpha_0)+\sigma_z^2}})}$$
$$= \frac{1}{b-1}\exp(\bar{c}_0 - \gamma - bp_0 + \alpha_0\widetilde{z}_0 + \frac{1}{2}(\sigma_x^2+\sigma_z^2))\left(\Phi(\frac{\alpha_{t|0}(\sigma_x^2(1-\alpha_0)+\sigma_z^2)-\underline{z}_0}{\sqrt{\sigma_x^2(1-\alpha_0)+\sigma_z^2}}) + \Phi(\frac{\underline{z}_0}{\sqrt{\sigma_x^2(1-\alpha_0)+\sigma_z^2}})(be^{p_1-p_0} - b + 1)e^{-b(p_1-p_0)}\right)$$

Similarly, $E(\left[\tilde{V}(\{\varepsilon^1, p_t, q_t\}, \bar{c}_0|\varepsilon^1, p_t = p_1)\right])$ can be computed as $E_{t-1}(\tilde{V}|p_t = p_1) =$

$$= P(\widehat{z}_t \leq \bar{z}_1)(\exp(p_0) - \exp(\bar{c}_0))\exp(-\gamma - bp_0 + \alpha_0\widehat{z}_0 + \frac{1}{2}(\sigma_x^2(1-\alpha_0)+\sigma_z^2))$$
$$+ P(\widehat{z}_t > \bar{z}_1)(\exp(p_1) - \exp(\bar{c}_0))\exp(-\gamma - bp_1 + \alpha_1\widehat{z}_1 + \frac{1}{2}(\sigma_x^2(1-\alpha_1)(1-\alpha_{t|1})+\sigma_z^2))E(\exp(\alpha_{t|1}\widehat{z}_t)|\widehat{z}_t > \bar{z}_1)$$
$$= \frac{1}{b-1}\exp(\bar{c}_0 - \gamma - bp_0 + \alpha_0\widetilde{z}_0 + \frac{1}{2}(\sigma_x^2+\sigma_z^2))\left[\Phi(\frac{\bar{z}_1}{\sqrt{(\sigma_x^2(1-\alpha_1)+\sigma_z^2)}}) + \Phi(\frac{\alpha_{t|1}(\sigma_x^2(1-\alpha_1)+\sigma_z^2)-\bar{z}_1}{\sqrt{(\sigma_x^2(1-\alpha_1)+\sigma_z^2)}})(be^{p_1-p_0} - b + 1)e^{-b(p_1-p_0)}\right]$$

Substituting in the expressions for $\underline{z}_0$ and $\bar{z}_1$ we obtain

$$E_{t-1}(\tilde{V}|p_t = p_0) = E_{t-1}(\tilde{V}|p_t = p_1)$$

Lastly, note that for $p_t \notin \{p_0, p_1\}$, $E(\left[\tilde{V}(c_0, \{\varepsilon^{t-1}, p_t, q_t\}|\varepsilon^1, p_t)\right])$ is the same as computed in the proof of Proposition 4 above. As a result, the interior maximum is achieved at $\lim p_t \to p_0$, hence to conclude our argument we need to compare $E_{t-1}(\tilde{V}|p_t = p_0)$ against $\lim_{p_t \to p_0} E_{t-1}(\tilde{V}|p_t \notin \{p_0, p_1\})$, which in turn equals

$$(\exp(p_0) - \exp(\bar{c}_0)) \exp(-\gamma - bp_0 + \frac{1}{2}(\sigma_x^2 + \sigma_z^2)) \left(\Phi(\frac{\bar{z}(p_t)}{\sqrt{(\sigma_x^2 + \sigma_z^2)}}) \exp(\alpha_0 \widetilde{z}_0) + \Phi(\frac{\alpha_t(\sigma_x^2 + \sigma_z^2) - \bar{z}(p_t)}{\sqrt{(\sigma_x^2 + \sigma_z^2)}}) \right)$$

Let $\hat{\theta} = (b(p_1 - p_0) - \ln(be^{(p_1-p_0)} - b + 1)) > 0$, then after substituting the expressions for $\underline{z}_0$ and $\bar{z}(p_t)$ and simplifying, the ratio of the two expected continuation values simplifies to:

$$\frac{E_{t-1}(\tilde{V}|p_t = p_0)}{\lim_{p_t \to p_0} E_{t-1}(\tilde{V}|p_t \notin \{p_0, p_1\})} = \frac{\Phi(\frac{\frac{\sigma_x^2}{2}(1-\alpha_0)+\frac{\hat{\theta}}{\alpha_{t|0}}}{\sqrt{\sigma_x^2(1-\alpha_0)+\sigma_z^2}}) + \Phi(\frac{\frac{\sigma_x^2}{2}(1-\alpha_0)-\frac{\hat{\theta}}{\alpha_{t|0}}}{\sqrt{\sigma_x^2(1-\alpha_0)+\sigma_z^2}}) \exp(-\hat{\theta})}{\Phi(\frac{\frac{\alpha_0}{\alpha_t}\widetilde{z}_0+\frac{\sigma_x^2}{2}}{\sqrt{(\sigma_x^2+\sigma_z^2)}}) + \Phi(\frac{\frac{\sigma_x^2}{2}-\frac{\alpha_0}{\alpha_t}\widetilde{z}_0}{\sqrt{(\sigma_x^2+\sigma_z^2)}}) \exp(-\alpha_0 \widetilde{z}_0)} \quad (47)$$

The denominator is decreasing in $\widetilde{z}_0$ and thus also in $\widehat{z}_0$, hence for every $\hat{\theta}$ there is a $\widehat{z}_0$ big enough such that the above ratio is strictly greater than 0. As a result, there exists a finite constant $\chi_0 > 0$ such that when $\widehat{z}_0 > \chi_0$ it follows that $p_t = p_0$ maximizes the expected continuation value. Finally, let $\chi = \max\{\chi_0, \chi_1\}$, then if $\widehat{z}_1 = \widehat{z}_0 > \chi$,

$$p_0 = \arg\max_{p_t} E\left[\tilde{V}(\{\varepsilon^1, p_t, q_t\}, \bar{c}_0)|\varepsilon^1, p_t\right]$$

Since $\tilde{V}$ is continuous in the cost shock c, it follows that there exists a non-singleton interval $(\underline{c}, \bar{c})$ around $\bar{c}_0$, such that if $c \in (\underline{c}, \bar{c})$, then

$$p_0 = \arg\max_{p_t} E\left[\tilde{V}(\{\varepsilon^1, p_t, q_t\}, c)|\varepsilon^1, p_t\right]$$

Lastly, we want to show that $\frac{\partial \chi}{\partial |p_0 - p_1|} < 0$. This follows directly form the facts that (i) the numerator of (47) is decreasing in $\hat{\theta}$, and that (ii) $\hat{\theta}$ is increasing in $(p_1 - p_0)$. Hence, as we decrease the distance between p_0 and p_1, we increase the RHS of (47), and thus we require a smaller $\widehat{z} = \widehat{z}_0 = \widehat{z}$ to make the ratio bigger than 1. □

Exploration incentives disappear as $\psi \to 0$

In this section we move away from the limiting case $\psi \to \infty$. Relaxing the assumption $\psi = \infty$ generally reduces the experimentation incentives of the firm, in the sense that it flattens the continuation value $\tilde{V}$. The reason is that when $\psi < \infty$, observing a signal q_t at a price p_t is informative no only about $x(p_t)$ itself, but also about other prices p round p_t, with the informativeness dropping to zero as the distance $|p - p_t|$ goes to infinity. Moreover, a higher ψ implies that the correlation between $x(p)$ and $x(p')$ at distinct p and p' decreases faster with the distance between p and p'. Hence, higher ψ increases the specificity of new information, making it more localized.

Lower ψ on the other hand, makes the information at a given p_t more useful at *any* p. As a result, this erodes the firm's incentive to experiment with new prices – it could learn most of the same information by repeating one of its established, safe prices anyways. Formally, this means that the continuation value function $\tilde{V}$ becomes flatter, and in fact, as Proposition A1 shows, in the limit $\psi \to 0$ the continuation value is a perfectly flat line.

Proposition A1. *The expected continuation value* $E\left[\tilde{V}(\{\varepsilon^0, p_t, q_t\}, \bar{c}_0)|\varepsilon^0, p_t\right]$ *becomes flat in respect to the time* t *price* p_t *as* $\psi \to 0$*:*

$$\lim_{\psi \to 0} \frac{\partial E\left[\tilde{V}(\{\varepsilon^0, p_t, q_t\}, \bar{c}_0)|\varepsilon^0, p_t\right]}{\partial p_t} = 0$$

Proof. First we will prove that with $\psi < \infty$, the expected continuation value $E\left[\tilde{V}(\{\varepsilon^0, p_t, q_t\}, \bar{c}_0)|\varepsilon^0, p_t\right]$ is differentiable. The key intuition is that if the firm selects a time t price away from p_0, thus obtaining a signal at a new price $p_t \neq p_0$, in expectation this would not create a second kind in the expected future worst-case demand. The only kink in the time t expectation of the future worst-case demand appears at the already observed p_0, since it evolves recursively as:

$$\widehat{x}_t^*(p) = \widehat{x}_{t-1}^*(p) + \alpha_t(p)(q_t - \widehat{x}_{t-1}^*(p))$$

where

$$\alpha_t(p) = \frac{(\sigma_x^2 + \sigma_z^2/N_0)\sigma_x^2 \exp(-\psi^2(p-p_t)^2) - \sigma_x^4 \exp(-\psi^2((p-p_0)^2 + (p_t - p_0)^2))}{\sigma_x^4(1 - \exp(-2\psi^2(p_t - p_0)^2)) + \sigma_x^2\sigma_z^2\frac{N_0+1}{N_0} + \sigma_z^4/N_0}$$

is the signal-to-noise ratio applicable to the new signal at p_t, when updating beliefs about $x(p)$ at some price p.

There is obviously a kink at $p_0$in $\widehat{x}_t^*(p)$, since $\widehat{x}_{t-1}^*(p)$ has a kink there. However, there is no other kink, because the firm correctly perceives that

$$q_t \sim N(\widehat{x}_{t-1}^*(p_t), \widehat{\sigma}_{t-1}^2(p_t)).$$

In other words, there is no possibility for a kink arising from the signal innovation term, since the signal is evaluated against the proper worst-case belief at time t, leaving only one kink in the expectation of the future worst-case demand. Of course, that is what happens only in expectation – once the signal is realized, and the firm perceives some surprise, the time $t+k$ worst-case will indeed feature two kinks. Still, in expectation, the kink is smoothed over, hence does not affect the time t pricing incentives of the firm.

We are going to use the notations for signal innovation level, $\widehat{z}_t$, and the signal-to-noise ratios defined above. Also recall that $E\left[\tilde{V}(\{\varepsilon^0, p_t, q_t\}, \bar{c}_0)|\varepsilon^0, p_t\right] = \frac{\beta}{1-\beta}E\left[\nu_{t+k}^*(p_{t+k}^*, \bar{c}_0)\middle|\varepsilon^0, p_t\right]$, where p_{t+k}^* is the resulting static optimal price, given the updated information set $\{\varepsilon^0, p_t, q_t\}$. And to simplify notation, we will again use the shorthand $E_{t-1}(\tilde{V})$ to denote the expected continuation value $E\left[\tilde{V}(\{\varepsilon^0, p_t, q_t\}, \bar{c}_0)|\varepsilon^0, p_t\right]$.

To show that the expected continuation value is differentiable, we will show two things. First, we show that the derivatives of $E_{t-1}(\tilde{V}|p_t > p_0)$ and $E_{t-1}(\tilde{V}|p_t < p_0)$ in respect to p_t exist everywhere. Second, we show that

$$\lim_{p_t \uparrow p_0} \frac{\partial E_{t-1}(\tilde{V}|p_t < p_0)}{\partial p_t} = \lim_{p_t \downarrow p_0} \frac{\partial E_{t-1}(\tilde{V}|p_t > p_0)}{\partial p_t}.$$

Let's start with showing that $\frac{\partial E_{t-1}(\tilde{V}|p_t > p_0)}{\partial p_t}$ exists everywhere. The firm has perfect foresight on $c_{t+k} = \bar{c}_0$, and since $p_0 = \ln(b/(b-1)) + \bar{c}_0$ absent any information in the new signal q_t the optimal price at $t+k$ would be p_0. Thus, the worst-case expected profit given a choice of $p_t > p_0$ can be written as:

$$\begin{aligned} E_{t-1}(\tilde{V}|p_t > p_0) = {} & \Phi(\underline{z}(p_t))E_{t-1}(\nu^*_{t+k}(p^*(p_t), \bar{c}_0|p_t > p_0, \widehat{z}_t < \underline{z}(p_t)) + (\Phi(\bar{z}(p_t)) - \Phi(\underline{z}(p_t)))E_{t-1}(\nu^*_{t+k}(p_0, \bar{c}_0|p_t > p_0) \\ & + (1 - \Phi(\bar{z}(p_t)))E_{t-1}(\nu^*_{t+k}(p^*(p_t)|p_t > p_0, \widehat{z}_t > \bar{z}(p_t)) \end{aligned}$$

where $\underline{z}(p_t)$ and $\bar{z}(p_t)$ are the threshold values for the innovation of the signal at p_t such that: (1) if $\widehat{z}_t > \bar{z}(p_t)$, the demand realization at p_t is so good that it pulls the optimal price away from p_0, and to an interior optimal price $p^*(p_t)$ closer to the new, good signal at p_t; (2) if $\widehat{z}_t < \underline{z}(p_t)$, the new demand realization is so bad that it pushes the optimal price away from both p_0 and p_1, to a new interior optimal $p(p_t)^* < p_0 < p_t$. For $\widehat{z}_t$ realizations in between these two threshold, the optimal price at time $t+k$ is at the kink p_0. We will prove that all of the components in the above expression are differentiable.

It is straightforward to show that the expected profit function (at any price p), $E_{t-1}(\nu^*_{t+k}(p, \bar{c}_0)|p_t > p_0)$, is differentiable in respect to p_t:

$$E_{t-1}(\nu^*_{t+k}(p, \bar{c}_0)|p_t > p_0) = (e^p - e^{\bar{c}_0}) \exp(\widehat{x}^*_{t-1}(p) + \alpha_t(p)\widehat{z}_t + \frac{1}{2}(\widehat{\sigma}^2_t(p) + \sigma^2_z))$$

The only components that are a function of p_t are the signal to noise ratio, $\alpha_t(p)$ and the posterior variance $\hat{\sigma}^2_t(p)$, and both of those are differentiable in respect to p_t everywhere. The signal-to-noise ratio $\alpha_t(p)$ was already defined above, and it is obviously differentiable, and the posterior variance can be obtained by the familiar recursive formula:

$$\widehat{\sigma}^2_t(p) = \sigma^2_x(1 - \alpha_0(p))(1 - \alpha_t(p))$$

where

$$\alpha_0(p) = \frac{\sigma^2_x}{\sigma^2_x + \sigma^2_z/N_0} e^{-\psi(p-p_0)^2}$$

is the signal-to-noise ratio applicable to the q_0 signal. This only depends on p_t through $\alpha_t(p)$, hence it is differentiable as well.

Next, consider the optimal interior price p^* – it satisfies the first order condition

$$p^* - (\bar{c}_0 + \ln(\frac{\widehat{x}^{*\prime}_{t-1}(p^*) + \alpha'_t(p^*)\widehat{z}_t + \frac{1}{2}\widehat{\sigma}^{2\prime}_t(p^*)}{1 + \widehat{x}^{*\prime}_{t-1}(p^*) + \alpha'_t(p^*)\widehat{z}_t + \frac{1}{2}\widehat{\sigma}^{2\prime}_t(p^*)})) = 0 \tag{48}$$

We can show that the derivative $\frac{\partial p^*}{\partial p_t}$ exists by using i) the implicit function theorem and ii) the fact that $\widehat{x}^*_{t-1}(p)$ has no kinks for $p > p_0$. To save on notation let

$$\theta^*(p^*, p_t) = \frac{\widehat{x}^{*\prime}_{t-1}(p^*) + \alpha'_t(p^*)\widehat{z}_t + \frac{1}{2}\widehat{\sigma}^{2\prime}_t(p^*)}{1 + \widehat{x}^{*\prime}_{t-1}(p^*) + \alpha'_t(p^*)\widehat{z}_t + \frac{1}{2}\widehat{\sigma}^{2\prime}_t(p^*)}$$

be the effective markup at the optimal price. By the implicit function theorem

$$\frac{\partial p^*}{\partial p_t} = -\frac{\frac{\partial \theta^*}{\partial p_t}}{1 - \frac{1}{\theta^*}\frac{\partial \theta^*}{\partial p^*}}$$

The derivative of $\frac{\partial \theta^*}{\partial p_t}$ is only a function of the derivatives $\alpha'_t(p)$ and $\widehat{\sigma}^{2\prime}_t(p)$ which exist everywhere since their expressions (as defined above) are infinitely differentiable. The derivative $\frac{\partial \theta^*}{\partial p^*}$ depends on the second derivatives of $\alpha_t(p)$ and $\widehat{\sigma}^2_t(p)$, and the time-t information worst-case demand, $\widehat{x}^*_{t-1}(p)$ – which is infinitely differentiable everywhere outside of $p_1 = p_0$. Hence, for $p_t > p_0$ the interior optimal price p^* is differentiable in respect to p_t.

Next, we work with the upper threshold $\bar{z}(p_t)$, which is implicitly defined by the equality

$$E_{t-1}(\nu^*_{t+k}(p_0|p_t > p_0, \widehat{z}_t = \bar{z}(p_t)) = E_{t-1}(\nu^*_{t+k}(p^*|p_t > p_0, \widehat{z}_t = \bar{z}(p_t))$$

$$\iff$$

$$(e^{p_0} - e^{\bar{c}_0})\exp(\widehat{x}^*_{t-1}(p_0) + \alpha_t(p_0)\bar{z}(p_t) + \frac{1}{2}(\widehat{\sigma}^2_t(p_0) + \sigma^2_z)) = (e^{p^*} - e^{\bar{c}_0})\exp(\widehat{x}_{t-1}(p^*) + \alpha_t(p^*)\bar{z}(p_t) + \frac{1}{2}(\widehat{\sigma}^2_t(p^*) + \sigma^2_z))$$

which can similarly be shown to be differentiable in respect to p_t by the implicit function theorem. Similar argument can be shown for the lower threshold $\underline{z}(p_t)$ as well.

Thus, we conclude that $\frac{\partial E_{t-1}(\tilde{V}|p_t > p_0)}{\partial p_t}$ exists everywhere. Similar arguments can be used to show that the mirror image derivative, $\frac{\partial E_{t-1}(\tilde{V}|p_t < p_0)}{\partial p_t}$ exists everywhere as well. Hence the only thing that remains to be shown, is that

$$\lim_{p_t \uparrow p_0} \frac{\partial E_{t-1}(\tilde{V}|p_t < p_0)}{\partial p_t} = \lim_{p_t \downarrow p_0} \frac{\partial E_{t-1}(\tilde{V}|p_t > p_0)}{\partial p_t}.$$

Note that outside of the limit $p_t \to p_0$ the thresholds $\underline{z}(p_t)$ and $\bar{z}(p_t)$ are different for the two cases i) $p_t > p_0$ and ii) $p_t < p_0$. Intuitively, the optimal interior price p^* could be different depending on whether the firm received a very good signal ($\widehat{z}_t > \bar{z}(p_t)$) for a price higher or lower than p_0. Importantly, the distance $|p^* - p_0|$ could also be different, because (at least locally) the slope of worst-case demand to the left of p_0 is different from that to the right of p_0. So resulting interior prices, and also the thresholds for $\widehat{z}_t$ at which they become optimal are different – i.e. the problem is not symmetric around p_0.

However, in the limit $p_t \to p_0$ the candidate interior prices and thresholds converge to the same values. The candidate interior price is given by the first-order condition (48), the minimum threshold $\lim_{p_t \to p_0} \underline{z}(p_t) = \underline{z}$ is defined as

$$E_{t-1}(\nu^*_{t+k}(p_0|p_t = p_0, \widehat{z}_t = \underline{z}) = E_{t-1}(\nu^*_{t+k}(p^*|p_t = p_0, \widehat{z}_t = \underline{z})$$

$$\iff$$

$$\begin{aligned}&(e^{p_0} - e^{\bar{c}_0}) \exp(\widehat{x}^*_{t-1}(p_0) + \alpha_t(p_0|p_0 = p_t)\underline{z} + \frac{1}{2}(\widehat{\sigma}^2_t(p_0|p_0 = p_t) + \sigma^2_z)) \\ &= (e^{p^*} - e^{\bar{c}_0}) \exp(\widehat{x}_{t-1}(p^*) + \alpha_t(p^*|p_0 = p_t)\underline{z} + \frac{1}{2}(\widehat{\sigma}^2_t(p^*|p_0 = p_t) + \sigma^2_z))\end{aligned}$$

and the upper threshold, $\bar{z}(p_t)$, converges to infinity – intuitively a new positive signal at p_0 only strengthens the desire to pick price p_0. New information will only destroy the kink at p_0 if it is sufficiently bad, while good new information will strengthen it.

With that in mind we can show

$$\begin{aligned}\lim_{p_t \uparrow p_0} \frac{\partial E_{t-1}(\tilde{V}|p_t < p_0)}{\partial p_t} &= \phi(\underline{z}) E_{t-1}(\nu^*_{t+k}(p^*(p_t)|p_t = p_0, \widehat{z}_t < \underline{z}))\frac{\partial \underline{z}}{\partial p_t} + \Phi(\underline{z}) \lim_{p_t \to p_0} \frac{\partial E_{t-1}(\nu^*_{t+k}(p^*(p_t)|p_t = p_0, \widehat{z}_t < \underline{z})}{\partial p_t} \\ &+ \phi(\underline{z}) E_{t-1}(\nu^*_{t+k}(p^*(p_t)|p_t = p_0, \widehat{z}_t \geq \underline{z}))\frac{\partial \underline{z}}{\partial p_t} + (1 - \Phi(\underline{z})) \lim_{p_t \to p_0} \frac{\partial E_{t-1}(\nu^*_{t+k}(p^*(p_t)|p_t = p_0, \widehat{z}_t \geq \underline{z})}{\partial p_t} \\ &= \lim_{p_t \uparrow p_0} \frac{\partial E_{t-1}(\tilde{V}|p_t > p_0)}{\partial p_t}\end{aligned}$$

which follows from (i) all limits exist and (ii) $\lim_{p_t \uparrow p_0} \underline{z}(p_t) = \lim_{p_t \downarrow p_0} \underline{z}(p_t) = \underline{z}$ as argued above. Lastly, we need to take the limit $\psi \to 0$. In this case, the signal-to-noise ratio function becomes flat, i.e. $\alpha_t(p) = \alpha_t$ for all p, and the same holds for the posterior variance $\widehat{\sigma}^2_t(p) = \widehat{\sigma}^2_t$, since now information at a price p' is equally useful at all prices p. As a result, it follows directly that $\lim_{\psi \to 0} \underline{z} = -\infty$ – i.e. since the signal realization erodes the expected profit equally at all prices, it does not make any price p^* better than p_0. By extension, $\lim_{\psi \to 0} \frac{\partial \underline{z}}{\partial p_t} = 0$. Lastly, since $\lim_{\psi \to 0} \frac{\partial \alpha_t(p)}{\partial p_t} = 0$, it also follows directly that $\lim_{\psi \to 0} \frac{\partial E_{t-1}(\nu^*_{t+k}(p^*(p_t)|p_t = p_0, \widehat{z}_t \geq \underline{z})}{\partial p_t} = 0$.

Essentially, the position of the new signal p_t no longer matters, as a result

$$\lim_{\psi \to 0} \frac{\partial E\left[\tilde{V}(\{\varepsilon^0, p_t, q_t\}, \bar{c}_0)|\varepsilon^0, p_t\right]}{\partial p_t} = 0$$

$\square$

B Appendix for Section 4

B.1 Empirical link between aggregate and industry prices

In this section, we use US CPI data to show that the relationship between aggregate and industry prices is time-varying and unstable over short-horizons. In particular, an econometrician would generally have very little confidence that short-run aggregate inflation is related to industry-level inflation, even though he can be confident that the two are cointegrated in the long-run. Thus, our assumption on the uncertainty over $\phi(.)$ puts the firm on an equal footing with an econometrician outside of the model.

Our analysis uses the Bureau of Labor Statistics' most disaggregated 130 CPI indices as well as aggregate CPI inflation. The empirical exercise consists of the following regression method. For a specific industry j, we define its inflation rate between $t-k$ and t as $\pi_{j,t,k}$ and similarly $\pi^a_{t,k}$ for aggregate CPI inflation. For each industry j, we run the rolling regressions:

$$\pi_{j,t,k} = \beta_{j,k,t}\pi^a_{t,k} + u_t$$

over three-year windows starting in 1995 and ending in 2010.[47] We repeat this exercise for k equal to 1, 3, 6, 12 and 24 months. Finally, for each of these horizons we compute the fraction of regression coefficients $\beta_{j,k,t}$ (across industries and 3-year regression windows) that are statistically different from zero at the 95% level.

We find that for 1-month inflation rates, only 11.4% of the relationships between sectoral and aggregate inflation are statistically significant. For longer horizons k, these fractions generally remain weak but do rise over time: 26.4%, 40.6%, 58.5% and 69.1% for the 3-, 6-, 12- and 24-month horizons respectively. This supports our assumption that while disaggregate and aggregate price indices might be cointegrated in the long run, their short-run relationship is weak.

In fact, not only is the relationship statistically weak in general, but it is highly unstable. This can be seen in Figure B.1, which shows the evolution of the coefficient $\beta_{j,k,t}$ for $k = 3$ for 3-year-window regressions starting in each month between 1995 and 2010, for four industries. Not only are there large fluctuations in the value of this coefficient over our sample, but sign reversals are common. In general, at any given date, there is little confidence that the near-future short-horizon industry-level inflation would be highly correlated with aggregate inflation, even though the data is quite clear that the two are tightly linked over the long-run.

[47] Results are very similar if we use windows of 2 or 5 years instead.

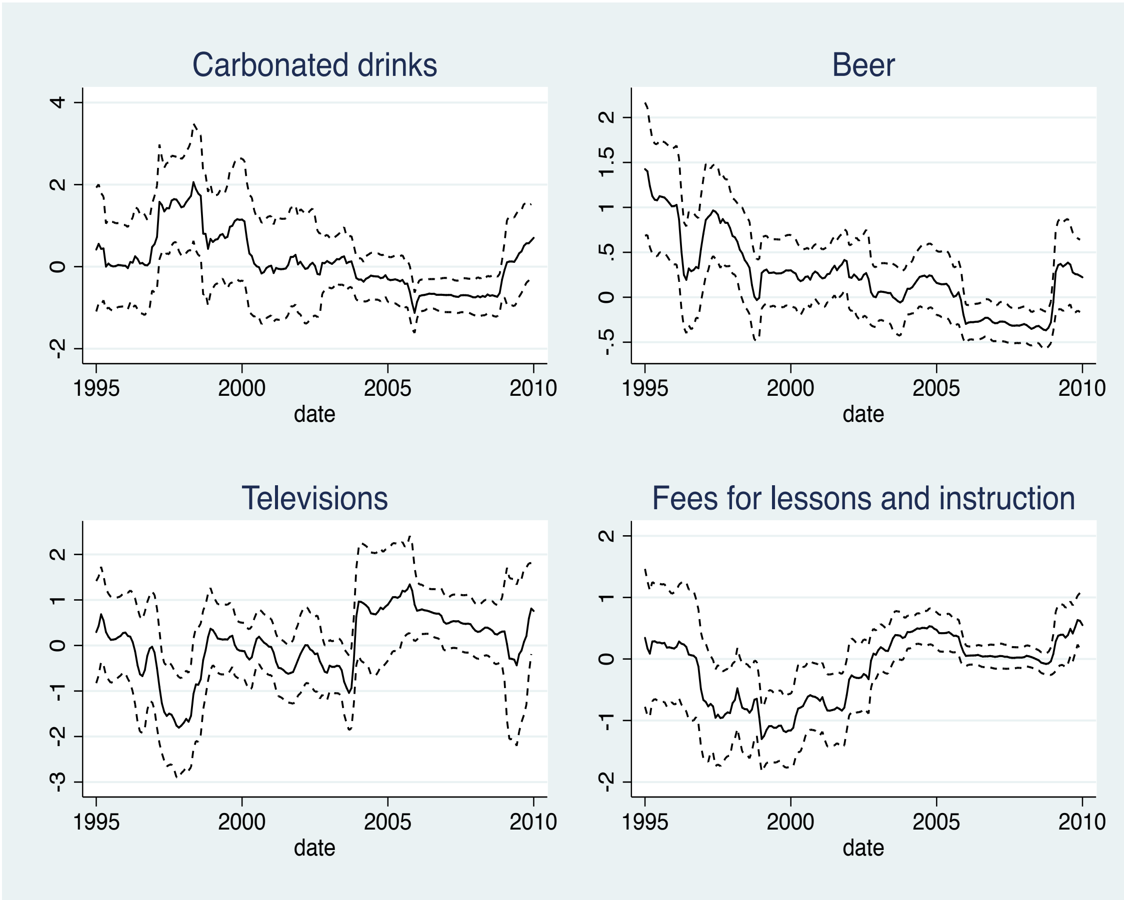

Figure B.1: 3-year rolling regressions of 3-month industry inflation on 3-month aggregate inflation for four categories. The solid line plots the point estimate of regression coefficient on aggregate inflation. The dotted lines plot the 95% confidence intervals.

B.2 Joint uncertainty over demand shape and relative price

In section 4.3.1 we have developed the solution to the worst-case beliefs when the firm observes one previous unambiguous estimated relative price, which here for brevity we call an estimated relative price. In this appendix we show how the analysis extends to multiple prices. The analysis follows the similar logic as in the real model, detailed in appendix A.1, with the added analysis of the worst-case belief of the unknown industry price. We do so by presenting details on the case of updating beliefs in the third period of life, when the firm has seen demand realizations at two previous prices $p_{i,0}$ and $p_{i,1}$, with corresponding quantities sold there $y_{i,0}$ and $y_{i,1}$. In addition, the firm observes the history of aggregates, $\{y_0, y_1, y_2, p_0, p_1, p_2\}$, and signals on the industry price level, $\{\widetilde{p}_{j,0}, \widetilde{p}_{j,1}, \widetilde{p}_{j,2}\}$. We will use the helpful $\widetilde{r}_{i,t} = p_{i,t} - \widetilde{p}_{j,t}$ notation for the unambiguously estimated relative price. In particular, without loss of generality, suppose that the prior observations imply unambiguously estimated relative price such that $\widetilde{r}_{i,0} < \widetilde{r}_{i,1}$.[48]

The firm is interested in updating beliefs at a current price $p_{i,2}$. Consider first a case where

[48]The opposite case is analogous.

$p_{i,2}$ implies an estimated relative price $\widetilde{r}_{i,2} > \widetilde{r}_{i,1}$. The expectation of demand is a function of the worst-case prior $m(r)$ at the true (unobserved) relative prices $r_{i,2}, r_{i,1}$, and $r_{i,0}$.

The worst-case prior at $r_{i,2}$ is again simply $m^*(r_{i,2}) = -\gamma - br_{i,2}$, (implying lowest prior level of demand at the current price). The resulting demand estimate ignoring all known aggregate effects, is given by

$$-\gamma - br_{i,2} + \alpha_0 y_{i,0} + \alpha_1 y_{i,1} - \alpha_0 \left[m(r_{i,0}) - b\phi(p_0 - \widetilde{p}_{j,0})\right] - \alpha_1 \left[m(r_{i,1}) - b\phi(p_1 - \widetilde{p}_{j,1})\right],$$

where α_0 and α_1 are weights on the perceived innovations in the signals $y_{i,0}$ and $y_{i,1}$, respectively. The prior belief about demand at $r_{i,0}$ and $r_{i,1}$ can be written as

$$m(r_{i,0}) = -\gamma - br_{i,0} + \delta_0'(r_{i,1} - r_{i,0});\; m(r_{i,1}) = -\gamma - br_{i,1} + \delta_1'(r_{i,2} - r_{i,1})$$

where δ_0', δ_1' are the local derivatives of the mean prior around $r_{i,0}$ and $r_{i,1}$ respectively (they do not have to be the same).

We can use the definition of $r_{i,t} \equiv p_{i,t} - p_{j,t}$ and substitute $p_{j,t}$ from equation (26) to simplify the portion of the demand estimate over which nature chooses the joint worst-case demand shapes δ_0' and δ_1', together with the short-run co-integrating relationship $\phi(p_t - \widetilde{p}_{j,t})$, as follows:

$$\min_{\delta_0', \delta_1'} \min_{\phi(p_t - \widetilde{p}_{j,t})} -\alpha_0\delta_0'(\widetilde{r}_{i,1} - \widetilde{r}_{i,0}) - \alpha_0\delta_0'\phi(p_0 - \widetilde{p}_{j,0}) - \alpha_1\delta_1'(\widetilde{r}_{i,2} - \widetilde{r}_{i,1}) + \alpha_1\delta_1'\phi(p_2 - \widetilde{p}_{j,2}) + \left(\alpha_0\delta_0' - \alpha_1\delta_1'\right)\phi(p_1 - \widetilde{p}_{j,1})$$

We obtain the solution for the joint worst-case

$$\delta_1^* = \delta_0^* = \delta;\; \phi^*(p_2 - \widetilde{p}_{j,2}) = -\gamma_p; \phi^*(p_0 - \widetilde{p}_{j,0}) = \gamma_p$$
$$\phi^*(p_1 - \widetilde{p}_{j,1}) = \gamma_p I(\alpha_0 < \alpha_1) - \gamma_p I(\alpha_0 > \alpha_1)$$

where $I(\alpha_0 < \alpha_1)$ denotes the indicator function of whether $\alpha_0 < \alpha_1$.

Intuitively, if the current entertained estimated relative price $\widetilde{r}_{i,2}$ is higher than the highest previously estimated relative price, then the joint worst-case beliefs over the demand shape and the unknown industry price index have the following characteristics. First, the prior demand shape between the three prices is steep. Second, the current industry price index is low and the price index at the lowest previously estimated relative price is high. In this way, the relative price between today and the lowest different estimated relative price is high, which, together with the steep demand curve, leads to the largest possible losses. Third, the worst-case belief about the industry price index at the previously estimated relative price that sits in the middle of the two extreme prices is a function of the updating weights. If these weights are the same then this belief is not determinate, as it does not matter for the posterior estimate.

Consider now the case where the entertained $p_{i,2}$ implies an unambiguously estimated relative price $\widetilde{r}_{i,2} < \widetilde{r}_{i,0}$. We follow the same steps as above to write the demand estimate and obtain the minimization objective

$$\min_{\delta_0',\delta_1'} \min_{\phi(p_t-\widetilde{p}_{j,t})} -\alpha_0\delta_0'(\widetilde{r}_{i,2}-\widetilde{r}_{i,0})+\alpha_0\delta_0'\phi(p_2-\widetilde{p}_{j,2})-\alpha_1\delta_1'\phi(p_1-\widetilde{p}_{j,1})-\alpha_1\delta_1'(\widetilde{r}_{i,1}-\widetilde{r}_{i,0})-(\alpha_0\delta_0'-\alpha_1\delta_1')\,\phi(p_0-\widetilde{p}_{j,0})$$

The joint worst-case beliefs are given by

$$\delta_1^* = \delta_0^* = -\delta;\ \phi^*(p_2-\widetilde{p}_{j,2}) = \gamma_p; \phi^*(p_1-\widetilde{p}_{j,1}) = -\gamma_p$$
$$\phi^*(p_0-\widetilde{p}_{j,0}) = \gamma_p I(\alpha_0<\alpha_1) - \gamma_p I(\alpha_0>\alpha_1)$$

Intuitively, if the current estimated relative price is lower than the lowest previously estimated relative price, then the worst-case prior demand is one with a flat shape between these three prices. In addition, the current unknown industry price index is high and the index at the highest previously estimated relative price is low. In this way, the relative price between today and highest different price is low, which together with the flat curve means the gain in demand is as low as possible. Finally, the belief about the industry price index at the intermediate price between the two extremes is a function of the updating weights. When these weights are the same then this belief is not determinate.

The final case is when the current entertained price $\widetilde{r}_{i,2}$ is between $\widetilde{r}_{i,0}$ and $\widetilde{r}_{i,1}$. The same steps as above deliver:

$$\min_{\delta_0',\delta_1'} \min_{\phi(p_t-\widetilde{p}_{j,t})} -\alpha_0\delta_0'(\widetilde{r}_{i,2}-\widetilde{r}_{i,0})-\alpha_0\delta_0'\phi(p_0-\widetilde{p}_{j,0})-\alpha_1\delta_1'\phi(p_1-\widetilde{p}_{j,1})-\alpha_1\delta_1'(\widetilde{r}_{i,1}-\widetilde{r}_{i,0})+(\alpha_0\delta_0'+\alpha_1\delta_1')\,\phi(p_2-\widetilde{p}_{j,2})$$

and the worst-case beliefs:

$$\delta_0^* = \delta; \delta_1^* = -\delta; \phi^*(p_0-\widetilde{p}_{j,0}) = \gamma_p; \phi^*(p_1-\widetilde{p}_{j,1}) = -\gamma_p$$
$$\phi^*(p_2-\widetilde{p}_{j,2}) = \gamma_p I(\alpha_0<\alpha_1) - \gamma_p I(\alpha_0>\alpha_1)$$

Intuitively, if the current price is in between the two previously estimated relative prices, then the worst-case prior demand is steep to the left and flat to the right. This concern for losing demand then also activates a concern that the industry price index is high at the left and low to the right. The belief about the current industry price index is a function of the updating weights. If these weights are the same then this belief does not matter. If the updating weight is larger on the previously low estimated relative price, then the worst-case is that the current index is low. This way the firm is worried about losing a lot of demand since it already acts as if it faces a steep part of the curve. If the weight is larger on the previously high estimated relative price, then the worst-case is that the current index is high. This way, the firm is concerned that it does not gain much demand since it already acts as if it faces a flat part of the demand curve.

By induction, we can build the worst-case belief of the firm in this fashion for any length of the previous history of observations, with the key result that the worst-case expected demand will have kinks around the unambiguous estimates of the previously observed prices $\widetilde{r}_{i,t}$.

B.3 Proofs on learning and nominal rigidity

Proposition 7. *The nominal price* $p_{i,1} = \widetilde{p}_{j,1} + \widetilde{r}_{i,0}$ *is a local maximizer of the worst-case expected profits for any aggregate price* $p_1 \in \left(\overline{p}_1 + \ln\left(\frac{b}{b-1}\frac{b-\alpha\delta-1}{b-\alpha\delta}\right), \overline{p}_1 + \ln\left(\frac{b}{b-1}\frac{b+\alpha\delta-1}{b+\alpha\delta}\right)\right)$.

Proof. Let $\upsilon^*(\varepsilon^0, s_1, p_{i,1})$ denote the worst-case expected profit, conditional on the history ε^0 and the current state $s_1 = \{\omega_{i,1}, p_1, y_1, \widetilde{p}_{j,1}\}$, evaluated at some nominal price $p_{i,1}$. Conditional on $p_{i,1} - \widetilde{p}_{j,1}$, the worst-case expected demand is given by equation (34). Take a first-order approximation of the change in profits, $\upsilon^*(\varepsilon^0, s_1, p_{i,1}) - \upsilon^*(\varepsilon^0, s_1, \widetilde{p}_{j,1} + \widetilde{r}_{i,0})$, evaluated around $p_{i,1} = \widetilde{p}_{j,1} + \widetilde{r}_{i,0}$. This equals

$$\left[\frac{e^{\widetilde{p}_{j,1}+\widetilde{r}_{i,0}-p_1}}{e^{\widetilde{p}_{j,1}+\widetilde{r}_{i,0}-p_1} - e^{y_1-\omega_{i,1}}} - (b + \alpha\delta^*)\right](p_{i,1} - \widetilde{p}_{j,1} - \widetilde{r}_{i,0}),$$

where $\delta^* = \delta \operatorname{sgn}(p_{i,1} - \widetilde{p}_{j,1} - \widetilde{r}_{i,0})$, as in Proposition 6.

It then follows that for any $p_1 \in (\underline{p}, \overline{p})$, where we define

$$\underline{p} = \overline{p}_1 + \ln\left(\frac{b}{b-1}\frac{b-\alpha\delta-1}{b-\alpha\delta}\right);\ \overline{p} = \overline{p}_1 + \ln\left(\frac{b}{b-1}\frac{b+\alpha\delta-1}{b+\alpha\delta}\right),$$

we have

$$\frac{e^{\widetilde{p}_{j,1}+\widetilde{r}_{i,0}-p_1}}{e^{\widetilde{p}_{j,1}+\widetilde{r}_{i,0}-p_1} - e^{y_1-\omega_{i,1}}} \in (b - \alpha\delta, b + \alpha\delta),$$

which makes the first-order derivative of the change in profits negative to the right of $\widetilde{p}_{j,1} + \widetilde{r}_{i,0}$ and positive to its left. This gives the necessary and sufficient conditions for $\widetilde{p}_{j,1} + \widetilde{r}_{i,0}$ to be a local maximizer. □

Proposition 8. *Let* $\delta^{index} = \delta \operatorname{sgn}(p_1 - \widetilde{p}_{j,1})$. *Up to a first-order approximation around* $p_1 = \widetilde{p}_{j,1}$, *the difference* $\ln \upsilon^*(\varepsilon^0, s_1, \widetilde{r}_{i,0} + p_1) - \ln \upsilon^*(\varepsilon^0, s_1, \widetilde{r}_{i,0} + \widetilde{p}_{j,1})$ *equals*

$$\left[\frac{e^{\widetilde{r}_{i,0}}}{e^{\widetilde{r}_{i,0}} - e^{y_1-\omega_{i,1}}} - b - \alpha\delta^{index}\right](p_1 - \widetilde{p}_{j,1}) < 0.$$

Proof. First, analyze the worst-case expected profit under a policy rule that implements indexation, i.e. $p_{i,1}^{index} = \widetilde{r}_{i,0} + p_1$, given by

$$\upsilon^*(\varepsilon^0, s_1, p_{i,1}^{index}) = \left(e^{\widetilde{r}_{i,0}} - e^{y_1-\omega_{i,1}}\right) e^{\widehat{x}_0^*(p_{i,1}^{index}, y_1, p_1, \widetilde{p}_{j,1})}$$

where $\widehat{x}_0^*(p_{i,1}^{index}, y_1, p_1, \widetilde{p}_{j,1})$ equals $.5\,(\widehat{\sigma}_0^2 + \sigma_z^2) + c_t - b\widetilde{r}_{i,0} - \gamma + \alpha\,[y_0 - (-\gamma - b\widetilde{r}_{i,0})]$ plus

$$\min_{\delta' \in [-\delta,\delta]} \min_{\phi(p_t - \widetilde{p}_{j,t}) \in [-\gamma_p, \gamma_p]} -\alpha\delta'(p_1 - \widetilde{p}_{j,1}) + \alpha\delta'\left[\phi(p_1 - \widetilde{p}_{j,1}) - \phi(p_0 - \widetilde{p}_{j,0})\right]$$

We know from Proposition 6 that the joint worst-case demand shape and co-integrating relationship are given by

$$\delta^{index} = \delta \operatorname{sgn}(p_1 - \widetilde{p}_{j,1})\,; \;\; \phi^{index}(p_1 - \widetilde{p}_{j,1}) - \phi^{index}(p_0 - \widetilde{p}_{j,0}) = -2\gamma_p \operatorname{sgn}(p_1 - \widetilde{p}_{j,1})\,.$$

Given the presence of the kink we compute a log-linear approximation of $\upsilon^*(\varepsilon^0, s_1, p_{i,1}^{index})$ around $p_1 = \widetilde{p}_{j,1}$. At its right we have

$$\frac{d \ln \upsilon^*(\varepsilon^0, s_1, p_{i,1}^{index})}{dp_1} = -\alpha\delta$$

while at its left, the derivative is

$$\frac{d \ln \upsilon^*(\varepsilon^0, s_1, p_{i,1}^{index})}{dp_1} = \alpha\delta$$

The constant term in the approximation is given by evaluating $\ln \upsilon^*(\varepsilon^0, s_1, p_{i,1}^{index})$ at $p_1 = \widetilde{p}_{j,1}$:

$$\ln\left(e^{\widetilde{r}_{i,1}^*} - e^{y_1 - \omega_{i,1}}\right) + c_t - b\widetilde{r}_{i,0} - \gamma + \alpha\left[y_0 - (-\gamma - b\widetilde{r}_{i,0})\right] - 2\alpha\delta\gamma_p. \tag{49}$$

Second, let us analyze the worst-case expected profit under the original policy, $p_{i,1}^* = \widetilde{r}_{i,0} + \widetilde{p}_{j,1}$, which targets the same $\widetilde{r}_{i,0}$ but by adjusting the nominal price to the review signal $\widetilde{p}_{j,1}$. We have

$$\upsilon^*(\varepsilon^0, s_1, p_{i,1}^*) = \left(e^{\widetilde{r}_{i,0} + \widetilde{p}_{j,1} - p_1} - e^{y_1 - \omega_{i,1}}\right) e^{\widehat{x}_0^*(p_{i,1}^*, y_1, p_1, \widetilde{p}_{j,1})}$$

where $\widehat{x}_0^*(p_{i,1}^*, y_1, p_1, \widetilde{p}_{j,1})$ equals $.5\left(\widehat{\sigma}_0^2 + \sigma_z^2\right) + c_t - b\left(\widetilde{r}_{i,0} + \widetilde{p}_{j,1} - p_1\right) - \gamma + \alpha\left[y_0 - (-\gamma - b\widetilde{r}_{i,0})\right]$ plus

$$\min_{\delta' \in [-\delta, \delta]} \; \min_{\phi(p_t - \widetilde{p}_{j,t}) \in [-\gamma_p, \gamma_p]} \alpha\delta' \left[\phi(p_1 - \widetilde{p}_{j,1}) - \phi(p_0 - \widetilde{p}_{j,0})\right] = -2\alpha\delta\gamma_p$$

Note that $\upsilon^*(\varepsilon^0, s_1, p_{i,1}^*)$ does not have a kink in the p_1 space. Approximate around $p_1 = \widetilde{p}_{j,1}$ to obtain a derivative is:

$$\frac{d \ln \upsilon^*(\varepsilon^0, s_1, p_{i,1}^*)}{dp_1} = -\frac{e^{\widetilde{r}_{i,0}}}{e^{\widetilde{r}_{i,0}} - e^{my_1}} + b$$

The constant term is given by evaluating $\ln \upsilon^*(\varepsilon^0, s_1, p_{i,1}^*)$ at $p_1 = \widetilde{p}_{j,1}$, as:

$$\ln\left(e^{\widetilde{r}_{i,0}} - e^{y_1 - \omega_{i,1}}\right) + c_t - b\widetilde{r}_{i,0} - \gamma + \alpha\left[y_0 - (-\gamma - b\widetilde{r}_{i,0})\right] - 2\alpha\delta\gamma_p. \tag{50}$$

We now compute the difference $\ln \upsilon^*(\varepsilon^0, s_1, p_{i,1}^{index}) - \ln \upsilon^*(\varepsilon^0, s_1, p_{i,1}^*)$, up to their first-order approximation:

$$\left(\frac{e^{\widetilde{r}_{i,0}}}{e^{\widetilde{r}_{i,0}} - e^{y_1 - \omega_{i,1}}} - b - \alpha\delta^{index}\right)(p_1 - \widetilde{p}_{j,1}) < 0$$

using the worst-case demand shape $\delta^{index} = \delta \operatorname{sgn}(p_1 - \widetilde{p}_{j,1})$ and Proposition 6. The latter shows that the condition for having the optimal price $\widetilde{r}_{i,1}$ be at the kink $\widetilde{r}_{i,0}$ is that the derivatives at the right, based on demand elasticity $-b-\delta$, and at the left, using the elasticity $-b+\delta$, are negative and, respectively, positive. $\square$

Proposition B1. *Consider a counterfactual economy, where the firm knows that the unique co-integrating relationship is* $\phi(p_t-\widetilde{p}_{j,t})=p_t-\widetilde{p}_{j,t}$, $\forall t$. *For a given realization of the current state* $s_1=\{\omega_{i,1},p_1,y_1,\widetilde{p}_{j,1}\}$, *the difference in worst-case expected profits* $\ln\upsilon^*(\varepsilon^0,s_1,p_{i,1})-\ln\upsilon^*(\varepsilon^0,s_1,p_1+r_{i,0})$, *up to a first-order approximation around* $p_1+r_{i,0}$, *is*

$$\left[\frac{e^{r_{i,0}}}{e^{r_{i,0}}-e^{y_1-\omega_{i,1}}}-(b+\alpha\delta^*)\right](p_{i,1}-p_1-r_{i,0}),$$

where $\delta^*=\delta\,\mathrm{sgn}\,(p_{i,1}-p_1-r_{i,0})$.

Proof. In this counterfactual economy the firm has the same ambiguity about demand shape as in the benchmark model but is endowed with the knowledge that

$$\phi(p_t-\widetilde{p}_{j,t})=p_t-\widetilde{p}_{j,t},\ \forall t. \tag{51}$$

Therefore this firm now knows that the unobserved industry price equals the observed aggregate price, since

$$p_{j,t}=\widetilde{p}_{j,t}+\phi(p_t-\widetilde{p}_{j,t})=p_t.$$

As a result, the estimated relative price simply equals

$$r_{i,t}=p_{i,t}-p_t. \tag{52}$$

Let us analyze the property of this economy in the simple two period model. The resulting worst-case expected profit is given by

$$\left(e^{p_{i,1}-p_1}-e^{mc_{i,1}}\right)e^{\widehat{x}_0^*(p_{i,1},y_1,p_1,\widetilde{p}_{j,1})}, \tag{53}$$

where the conditional payoff $\widehat{x}_0^*(p_{i,1},y_1,p_1,\widetilde{p}_{j,1})$ equals $.5\,(\widehat{\sigma}_0^2+\sigma_z^2)$ plus

$$\min_{\delta'\in[-\delta,\delta]}\exp\left\{y_1-b\,(p_{i,1}-p_1)-\gamma+\alpha\left[y_0-(-\gamma-br_{i,0})\right]-\alpha\delta'\,(p_{i,1}-p_1-r_{i,0})\right\} \tag{54}$$

The worst-case demand shape is therefore given by

$$\delta^*=\delta\,\mathrm{sgn}\,(p_{i,1}-p_1-r_{i,0}).$$

Having described the worst-case expected profit, the proof follows from taking the derivatives of expected profit in (53) and payoff in (54) with respect to the action $p_{i,1}$. $\square$

Different from the benchmark economy, we note that in this counterfactual the worst-case expected profit does not depend directly on the aggregate price. Indeed, the optimal choice of the relative price in equation (54) is independent of p_1. In this economy indexation is built in, as instructed per equation (52) where, holding constant the relative price, the nominal price moves one to one with p_1. Therefore, not surprisingly, a nominal price policy that deviates from indexation is suboptimal. To show this, consider a firm that lives in this counterfactual economy but does not index to the aggregate price. Instead, it targets the same $r_{i,0}$ but by setting $p_{i,1}^{noindex} = r_{i,0} + \widetilde{p}_{j,1}$. Put differently, this firm uses only the review signal as the source of relevant information for $p_{j,1}$ but targets the same relative price. Proposition B2 below details how the non-indexing policy is strictly suboptimal.

Proposition B2. *In the counterfactual economy, the difference* $\ln \upsilon^*(\varepsilon^0, s_1, p_1 + r_{i,0}) - \ln \upsilon^*(\varepsilon^0, s_1, \widetilde{p}_{j,1} + r_{i,0})$*, up to a first-order approximation around* $\widetilde{p}_{j,1}$*, equals*

$$\left(\frac{e^{r_{i,0}}}{e^{r_{i,0}} - e^{y_1 - \omega_{i,1}}} - b - \alpha \delta^{noindex} \right) (p_1 - \widetilde{p}_{j,1}) > 0$$

where $\delta^{noindex} = -\delta \, \mathrm{sgn}(p_1 - \widetilde{p}_{j,1})$.

Proof. The firm that sets $p_{i,1}^{noindex}$ is subject to the same informational assumption as the firm that indexes, and, therefore, it still knows that the co-integrating relationship is given by (51). Compared to the indexing policy, this firm simply follows a different nominal pricing policy. The resulting worst-case expected profit is

$$\upsilon^*(\varepsilon^0, s_1, p_{i,1}^{noindex}) = \left(e^{r_{i,0} + \widetilde{p}_{j,1} - p_1} - e^{y_1 - \omega_{i,1}} \right) e^{\widehat{x}_0^*(p_{i,1}^{noindex}, y_1, p_1, \widetilde{p}_{j,1})}$$

where $\widehat{x}_0^*(p_{i,1}^{noindex}, y_1, p_1, \widetilde{p}_{j,1})$ equals $.5\left(\widehat{\sigma}_0^2 + \sigma_z^2\right) + y_1 - b\left[r_{i,0} + \widetilde{p}_{j,1} - p_1\right] - \gamma + \alpha\left[y_0 - (-\gamma - br_{i,0})\right]$ plus

$$\min_{\delta' \in [-\delta, \delta]} -\alpha \delta' \left(r_{i,0} - (p_1 - \widetilde{p}_{j,1}) - r_{i,0} \right).$$

The worst-case demand shape is therefore simply

$$\delta^{noindex} = -\delta \, \mathrm{sgn}\left(p_1 - \widetilde{p}_{j,1}\right). \tag{55}$$

Compute now the log-linear approximation with respect to p_1, for this worst-case expected profit $\upsilon^*(\varepsilon^0, s_1, p_{i,1}^{noindex})$, evaluated to the right and left of $\widetilde{p}_{j,1}$. Those derivatives are

$$\frac{d \ln \upsilon^*(\varepsilon^0, s_1, p_{i,1}^*)}{dp_1} = -\frac{e^{r_{i,0}}}{e^{r_{i,0}} - e^{y_1 - \omega_{i,1}}} + b + \alpha \delta^{noindex}$$

The resulting $\ln \upsilon^*(\varepsilon^0, s_1, p_1 + r_{i,0}) - \ln \upsilon^*(\varepsilon^0, s_1, \widetilde{p}_{j,1} + r_{i,0})$, up to a first order approximation, is

$$\left(\frac{e^{r_{i,0}}}{e^{r_{i,0}} - e^{y_1 - \omega_{i,1}}} - b - \alpha \delta^{noindex} \right) (p_1 - \widetilde{p}_{j,1}) > 0$$

since when p_1 is larger (smaller) than $\widetilde{p}_{j,1}$, by the worst-case in (55) we have $\delta^{noindex} = -\delta$ or δ, respectively. Here we have used that the optimal $r_{i,1}$ sitting at the kink $r_{i,0}$ implies that

$$\frac{e^{r_{i,0}}}{e^{r_{i,0}} - e^{y_1 - \omega_{i,1}}} - b + \delta > 0 > \frac{e^{r_{i,0}}}{e^{r_{i,0}} - e^{y_1 - \omega_{i,1}}} - b - \delta.$$

□

C Appendix for Section 5

C.1 Dispersion of forecasts

Here we detail how we use empirical evidence from Gaur et al. (2007) on survey data to evaluate the size of our calibrated ambiguity parameter γ. Gaur et al. (2007) use item-level forecasts of demand data from a skiwear manufacturer, called the Sport Obermeyer dataset. The dataset contains style-color level forecasts for 248 short lifecycle items for a selling season of about three months. The forecasts are done by members of a committee specifically constituted to forecast demand, consisting of: the president, a vice president, two designers, and the managers of marketing, production, and customer service. Raman et al. (2001) provides details on the forecasting procedures and on the dataset.

Our model connects to the data in Gaur et al. (2007) as follows. They observe forecasts made prior to the product being introduced. Their statistic for the dispersion of these forecasts is reported as a coefficient of variation. Our model relates to this measure through the set of multiple priors. Indeed, in our model, prior to observing any realized demand signals, the firm entertains a set of forecasts about quantity sold. We connect this set to the dispersion of forecasts made by the committee described above. In particular, in our model the firm entertains the following time-zero set of forecasts on the level of demand

$$[\exp(-\gamma - bp + 0.5\sigma_z^2), \exp(\gamma - bp + 0.5\sigma_z^2)]$$

While in the data the set consists of only seven forecasters, we have a continuum. But we can compute the coefficient of variation (CV) of these forecasts and compare it against the reported statistic. In particular, using a uniform distribution over the forecasts in the set above, the CV, normalized by the average forecast, equals

$$CV = \frac{1}{\sqrt{3}} \frac{e^{\gamma} - e^{-\gamma}}{(e^{\gamma} + e^{-\gamma})} \tag{56}$$

Gaur et al. (2007) report in Table 4 that the average level of coefficient of variation, scaled by the average forecast, across the products in the dataset equals 37.6%. When we plug in the calibrated value of our ambiguity parameter $\gamma = 0.614$, we obtain a CV equal to 31.58%.

C.2 Simulated hazards

In this section, we use simulations to confirm that our econometric approach is appropriate and allows us to recover the true slope of the hazard function, even in the presence of pervasive heterogeneity.

We simulate panels of 500 price changes for 100,000 products. Each product i is characterized by a randomly-chosen unconditional price change probability, ξ_i, as well as a coefficient that determines the slope of its hazard function, ϕ_i. To make the comparison between the true and estimated slopes easier, we assume for this exercise that the hazard functions are linear at the product level. The slope of product i's hazard, s_i, is defined as:

$$s_i = (1 - \phi_i)\xi_i / 13$$

As a result, the probability of a price change after a spell of length τ smaller or equal than 13 is given by $\xi_i^\tau = \xi_i - \tau s_i$. In other words, the slope is not a function of τ. For $\tau > 13$, the probability is assumed to be constant at $\xi_i^\tau = \xi_i - 13 s_i$ (we will only estimate the hazard slopes for spells less than or equal to 13 periods).

Panels differ in the distributions of the baseline probabilities ξ_i and slope factors ϕ_i. We run the exact same code we use for actual data on the simulated panels, including regressions with and without product fixed effects:

$$Pr(p_{i,t} \neq p_{i,t-1}) = \alpha + \beta \tau_{i,t} + \gamma_i + u_{i,t} \tag{57}$$

The results are summarized in Table C.1. Each column of the table represents a different simulated panel. The top portion of the table describes the distribution of the baseline price change probabilities (ξ) and slope parameters (ϕ) across simulated products, as well as the average, known slope of the hazard function across products.[49] The middle and bottom parts report the slope estimates $\hat{\beta}$, the standard error of the coefficient estimate and the p-value against the null of a flat slope, for regressions without and with product fixed effects respectively.

The first column, A, shows estimates of the slope of the hazard function when there is no heterogeneity in either price change probabilities or slope parameters. Not surprisingly, the coefficient $\hat{\beta}$ correctly recovers the true value of the slope and leads us to correctly conclude that the hazards are flat, whether product fixed effects are included or not.

Next, we introduce heterogeneity in the unconditional price change frequencies ξ_i. We do, however, keep a homogenous, flat slope of the hazard function. Our simulations confirm the presence of the survivor-bias issue discussed in the literature: without fixed effects, the estimation

[49] Unless otherwise noted, all distributions are uniform.

Table C.1: Estimated slopes of the hazard function for various simulated panels

		A	B	C	D	E	F
ξ distribution		[0.15,0.15]	[0.01,0.3]	Empirical	[0.01,0.3]	[0.01,0.3]	[0.01,0.3]
ϕ distribution		[1,1]	[1,1]	[1,1]	[0.7,0.7]	[0.5,1.5]	[0.2,1.2]
Actual slope (avg)		0	0	0	-0.0036	0	-0.0036
w/o fixed effects	$\hat{\beta}$	0.00032	-0.00658	-0.00407	-0.00910	-0.00664	-0.00909
		(0.00016)	(0.00015)	(0.00015)	(0.00016)	(0.00016)	(0.00016)
	p-value	0.042	0.000	0.000	0.000	0.000	0.000
w/ fixed effects	$\hat{\beta}$	0.00032	0.00031	0.00026	-0.00360	0.00022	-0.00350
		(0.00016)	(0.00016)	(0.00015)	(0.00016)	(0.00016)	(0.00017)
	p-value	0.042	0.052	0.096	0.000	0.185	0.000

finds a hazard that is declining on average (column B), even if our simulation features no relationship between spell length and price change frequency. This is also true if we use a distribution of the price change probabilities ξ_i that mimics the empirical distribution from our dataset (column C).[50] The inclusion of product fixed effects, on the other hand, correctly leads us to conclude that the hazards are flat on average: controlling for product-specific hazard shifters circumvents the downward bias that arises from heterogenous price rigidity.

If we instead assume a homogenous *declining* slope, the regression manages to recover perfectly its value of -0.0036 once we include product fixed effects (column D). Without fixed effects, however, the hazard is estimated to be three times steeper than it actually is, at -0.0091.

Finally, we also allow for heterogeneity in the slope factors ϕ_i. The last part of Table C.1 shows results for regressions run on simulated panels with two different distributions of ϕ_i. Once again, the fixed-effects regression correctly finds a flat average hazard when the distribution of ϕ_i is centered at 1 (column E). Second, it is able to recover a declining hazard function when it should (column F), with an estimate of -0.0035 vs. the actual value of -0.0036. As we saw earlier, omitting product fixed effects would lead us to find a slope that is almost three times larger (in absolute value) than it actually is, at -0.0091.

To conclude, our simulation exercises confirm that our econometric approach allows us to drastically alleviate the well-known survivor bias that arises in the computation of hazards of price changes.

[50] We found that a χ^2 distribution with 5 degrees of freedom, scaled to match the mean frequency found in our dataset, provides a good fit.

C.3 Additional evidence on hazard functions

To complement the evidence on the slopes of the hazard functions presented in the main text, we produce a number of figures.

First, we plot in Figure C.1 the distributions of the estimated hazard slopes across the 54 category/market combinations. These estimates are obtained from our linear probability regression model with fixed effects of equation (36). The left panel shows the slope estimates from unweighted regressions, while results from weighted regressions are shown in the right panel.

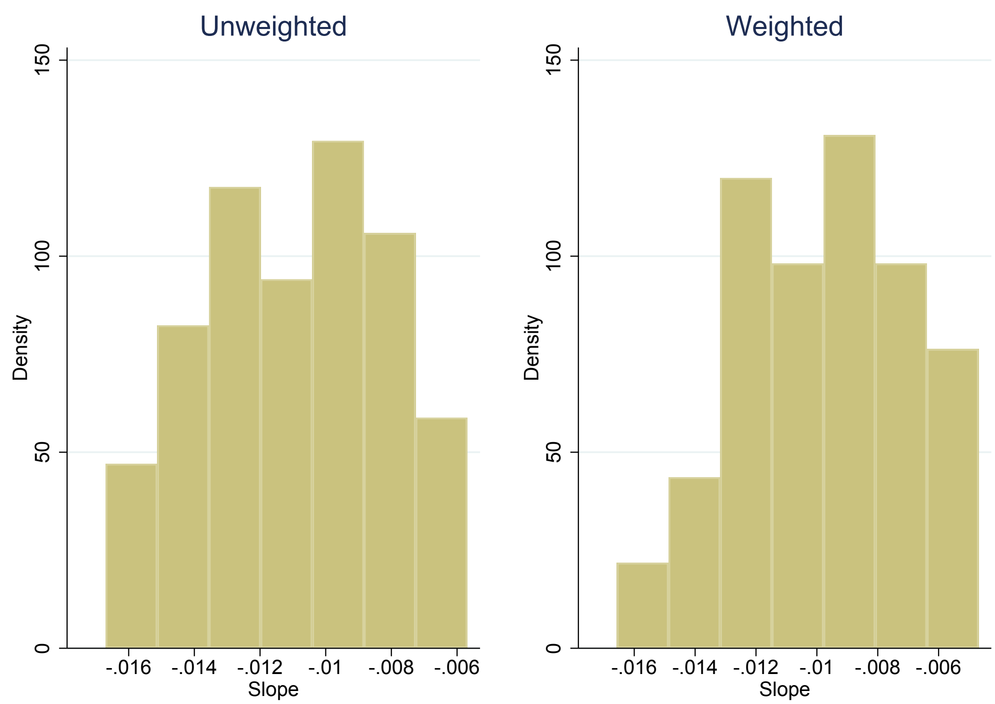

Figure C.1: Distribution of the slopes of hazard functions across 54 category/market pairs. Unweighted and weighted regressions.

In Figure C.2, we plot the distributions of cell-based slopes obtained using the approach of Campbell and Eden (2014). A cell is a specific product sold in a given store, while the slope is computed as the difference between the price change frequencies of older and younger prices. An "old" price is one that has survived at least Γ weeks. In order to obtain a more complete comparison between the data and the model simulations than just the average slope, we plot both the empirical (left column) and simulated (right) distributions of the cell-based hazard slopes, for $\Gamma = 4, 5, 6$.

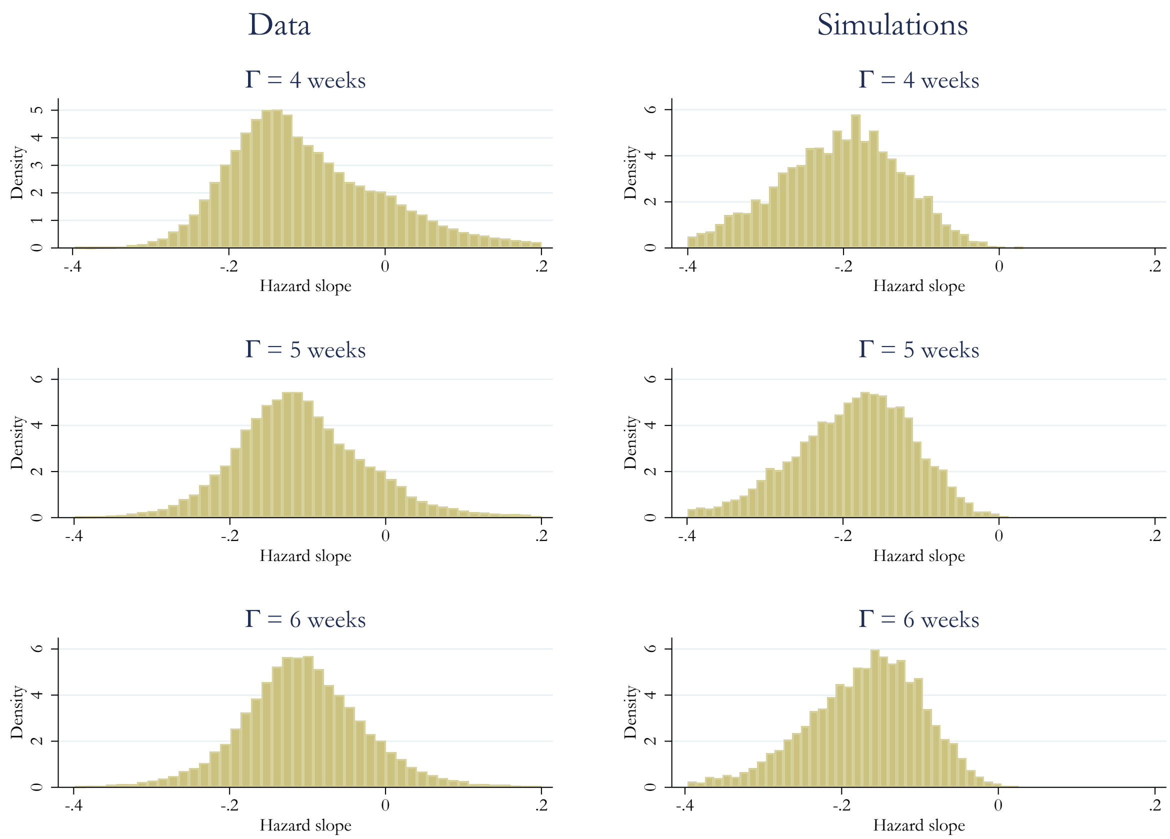

Figure C.2: Distributions of the cell-based hazard slopes. A slope is defined as the difference between the price change frequencies of old ($\tau \geq \Gamma$) and young ($\tau < \Gamma$) prices. Empirical (left) and simulated (right) distributions.

C.4 Constructing the typical history of observations

In the model, the price histories and demand realizations differ across firms. One reason is the idiosyncratic noise in demand realizations, but more importantly, the position of the demand signals is endogenous, because it depends on the past pricing decisions of the firm. With idiosyncratic productivity shocks, firms take different pricing decisions, and thus their information sets evolve differently. Let

$$\mathcal{I}_{it} = \left[\tilde{\mathbf{r}}_{it}^{uniq}, \mathbf{N}_{it}, \hat{\mathbf{y}}_{it}\right]$$

be the 3-column matrix that characterizes the information set of firm i at time t, where $\tilde{\mathbf{r}}_{it}^{uniq}$ is the vector of *unique* unambiguously estimated relative price points in the history of past price decisions, $\tilde{r}_i^t$, of firm i; $\mathbf{N}_{it}$ is the associated vector of the number of times each of those unique price points has been chosen in the past; and $\hat{\mathbf{y}}_{it}$ is the average, demeaned demand realization that the firm has seen at those unique price points. So each row of $\tilde{\mathbf{r}}_{it}^{uniq}$ is one of the unique price levels the firm has posted in the past, the corresponding row of $\mathbf{N}_{it}$ is the number of times this price has been seen in the past, and the corresponding row of $\hat{\mathbf{y}}_{it}$ is the average demeaned demand realizations the firm has experienced when choosing that price. The matrix $\mathcal{I}_{it}$ fully described the information set of the firm, and is the sufficient statistic needed to compute the worst-case expected demand $\hat{x}_{it}(\tilde{r})$.

As discussed in the main text, a striking characteristic of $\mathcal{I}_{it}$ is that the average cardinality of $\tilde{\mathbf{r}}_{it}^{uniq}$ is just six, hence the average firm tends to have chosen and thus seen only around six unique price levels in the past. Another interesting characteristic, is that the average firm has not seen each of those six price points equally often, but in fact the most often posted price accounts for 74% of all observations, on average. Moreover, the second most often chosen price accounts for another 19% of all observations. As a result, we observe that there is a clear hierarchy in the amount of information collected at the different price points observed in the past.

We want to preserve this hierarchical structure when averaging the price histories of different firms, hence we sort the rows of $\mathcal{I}_{it}$ based on the number of times each of the past price points has been visited (given in $\mathbf{N}_{it}$), and we call the sorted matrix $\mathcal{I}_{it}^{sorted}$. Next, we compute the cross-sectional average of $\mathcal{I}_{it}^{sorted}$ (element-wise) at each time period t, to come up with the information set of the average firm at time t:

$$\bar{\mathcal{I}}_t = \int \mathcal{I}_{it}^{sorted} di$$

Finally, we compute the time-average of $\bar{\mathcal{I}}_t$ to come up with the "typical" information set in the stochastic steady state of our model. Just as with all other moments we compute, we discard the first 1000 periods of our simulation, and focus on the remaining 4000 to give a chance to the model to converge to its stochastic steady state.

C.5 Speed of learning

The evolution of the pricing policy function over time

To further illustrate how learning and the resulting pricing policy evolve over time, panel a) of Figure C.3 shows how the policy function of one the longer-lived firms in the simulation changes from period one-hundred and fifty, to the three hundredth period of this firm's life. The blue line corresponds to the optimal policy at $t = 150$, and shows that by that period the firm had sampled a number of different prices, and established a fair number of kinks. While we might think that establishing such "special prices" happens once and for all, in fact the position of the kinks can move and they could even completely disappear as new information arrives. We can see that from the red line, which plots the policy at $t = 300$, and shows that by that period the two lowest flat spots in the policy became absorbed in a new, single flat spot at an intermediate price point.

Thus, the accumulation of new information could change the optimal position of some of the reference prices. Over time, it tends to be the case that any given neighborhood of the price space becomes associated with one special price, and the firm does not visit other prices nearby – this is another reason for the slow speed of learning.

A counterfactual economy with no firm exit

To showcase the slow nature of learning in our model, focus on the limiting case of no firm exit $\lambda_\phi = 0$, hence firms never stop accumulating new signals. As we show here, however, that by itself is not enough to ensure that firms eliminate demand uncertainty, because profit maximization incentives lead them to often repeat estimated relative prices $\tilde{r}_{it}$ that have already been visited in the past. Thus, the history of observations that the firm sees is endogenously sparse, concentrated in a handful of individual price points, as opposed to being distributed all over the support of the demand curve. As a result, the firm has good information about demand at several different price points, but remains uncertain about the shape in between those prices. Hence, our mechanism is preserved even in the very long run.

To illustrate, we note that the number of unique estimated relative prices that a firm has seen after 5000 periods is just 40 on average. Moreover, most of the signals have been observed at just 3 separate $\tilde{r}_{it}$ values, one of which accounts for 48% of all observations, and the other two for 33% and 12% respectively. As a result, even though the firm has accumulated a lot of signals, it remains uncertain about the overall shape of its demand. The accumulated signals are very informative about the average level of demand in the neighborhood of the few prices that the firm keeps repeating and collecting more information on, but this provides little guidance about the shape of the demand function between the observed prices. Thus, the mechanism we develop, which emphasizes uncertainty in the *local* shape of demand, remains present even after thousands of periods of observations. The key intuition behind this result is the endogeneity of the history of observations: the firms are not collecting an exogenous stream of observations randomly spread out over the whole demand curve, but are balancing the learning incentives with profit maximization.

As a result, even when firms are infinitely-lived and accumulate thousands observations about demand, the behavior of prices remains qualitatively similar to that in the benchmark model, with prices displaying both stickiness and memory. To understand this pricing behavior, we use our procedure to compute the typical optimal policy function (in terms of the estimated relative price $\tilde{r}_{it}$) from this simulation, with results plotted in panel b) of Figure C.3. As can be seen from the Figure, the policy function is qualitatively similar to that in the benchmark case, and is essentially a step function across the whole support of the price space. Again, this is because even though the firms have seen much longer histories of observations, they have concentrated their pricing, and thus information accumulation, in the set of previously observed estimated relative prices. This results in a pricing policy that is a step-function, generating both stickiness and memory in prices.

In the model with no exit ($\lambda_\phi = 0$), the frequency of changing posted nominal prices is 6.5%, and the frequency of changing modal prices is 2.8%. Meanwhile, the median size of price changes is 10.8%, and the probability of revisiting prices posted in the past (conditional on a price change) is 50% (most non-revisits in this case come from new industry price review signals). Hence, even without firm exit, the model shares many of the same characteristics as the benchmark model. We have chosen to include firm exit in the benchmark model purely out of numerical convenience, as exit introduces faster convergence to the stochastic steady state, with moments that are more stable at smaller simulation sizes. This helps make the estimation feasible.

Figure C.3: Optimal Pricing Policy Function

(a) Benchmark economy, at two intermediate points in time (b) Stochastic steady-state pricing policy function, $\lambda_\phi = 0$