

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

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

June 19, 2020

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 generates nominal rigidity. We prove the main insights analytically and 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. It also implies significantly more persistent monetary non-neutrality than in standard models, allowing it to generate large real effects from nominal shocks.

*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, Krisztina Molnar, Martin Schneider and Luminita Stevens, as well as seminar and conference participants at the 9th Annual CIGS conference on Macroeconomic Theory and Policy, 22nd Annual Research Conference of the Dutch Central Bank, Ambiguity and Robustness in Macroeconomics and Finance Workshop, Ambiguity and its Implications in Finance and Macroeconomics Conference, Behavioral Aspects of Macroeconomics and Finance Conference, Boston Macro Juniors Workshop, Chicago Fed, CIREQ Montréal Macro conference, Duke, ECB, ESWC, NBER Summer Institute, NBER EFG meeting, Northwestern, SED, Stanford, UC San Diego, U Penn, UT Austin, UVA for helpful discussions and comments. Vincent acknowledges financial support from the Fondation HEC Montréal.

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 aggregate 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 as sticky as the aggregates imply. They do, however, display other puzzling characteristics that could play a crucial macro role (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 price memory and a number of additional micro-level pricing facts. These features allow it to generate significant monetary non-neutrality despite prices changing relatively frequently, as in the data.

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 (i) the unknown *demand curve*, as a function of the relevant relative price; and (ii) 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 the price index of the competition. 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).¹ In particular, we model the aversion to ambiguity using the multiple priors preferences axiomatized by Gilboa and Schmeidler (1989), 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.

To this end, we assume that the firm, similar to an econometrician, estimates its unknown demand function from past observations of prices and quantities sold. 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

¹See Machina and Siniscalchi (2014) for a review of related theory and experiments. The latter confirm the basic conjecture in Ellsberg (1961) of prevalent aversion to ambiguity, and includes surveys and experiments specifically involving business managers, such as in Einhorn and Hogarth (1986), March and Shapira (1987), Kunreuther et al. (1993) and Maffioletti and Santoni (2005).

measure over the space of potential demand functions. In particular, while the firm knows that its demand is the sum of a price-sensitive component and a temporary shock, it faces a signal extraction problem because it does not observe each separately. The firm uses its history of quantities sold at past prices, together with its set of priors, to form a set of conditional beliefs about its demand function.

The firm has two sources of information on the unknown industry-wide price level. The first are periodic marketing reviews that fully reveal its current value. The second is the aggregate price level, which the firm observes freely, but is an imperfect signal of the firm's specific industry price, because the link between industry and aggregate prices is uncertain and ambiguous – while the firm understands that the two indices are cointegrated in the long run, it is not confident about their short-run relationship. Specifically, 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 given an observed value for the aggregate price level.²

In the face of ambiguity about both its demand function and its effective 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 relative to the firm's best, unambiguous estimate of industry inflation. The reason is intuitive: an unambiguous price increase sets in motion a concern for a “double whammy” – that nature draws (i) the most locally-elastic demand function allowed by the prior set and (ii) 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: it generates *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

²Using 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.

³Such 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.

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 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, since it 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 suboptimal.

In sum, a 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 marketing 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 tends to bounce around just a few repeated price points. One important difference with standard “price plan” models is that in our framework, the endogenous price plan evolves *gradually* over time, incorporating new prices one-by-one as the firm experiments and learns about demand at new price points. We document that this novel implication of gradual adjustment in price plans is prevalent in the data, and also show that it has important implications for the aggregate transmission of monetary shocks in the model.

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 Academic 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

Lastly, we show that our quantitative model predicts large and persistent real effects from a 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 Golosov and Lucas (2007) selection effect. The reason lies in the endogenous memory of prices, and in particular the slow adjustment of the effective “price plan”, an empirically supported feature unique to our model. Because it significantly slows down the transmission of nominal shocks, this type of memory delivers more persistent real effects than in standard price-plan models. As a result, our mechanism has important novel features that can arguably help fit the evidence of persistent monetary policy effects (Christiano et al. (2005)).

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

Relation to the literature

By connecting learning under ambiguity to the problem of a firm setting prices, our paper relates to multiple strands of the literature. 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). 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 demand, and an econometrician would not be expected to find evidence of actual kinks.

On nominal rigidity, we connect to the literature that emphasizes the role of imperfect information in generating slow adjustment to aggregate nominal shocks, including Mankiw and Reis (2002), Sims (2003), Woodford (2003), Reis (2006) and Mackowiak and Wiederholt (2009). However, while in order to obtain fully rigid prices these models typically require additional frictions, (e.g. a menu cost), we show that uncertainty alone can lead to inaction.⁵

In testing our mechanism against a rich set of overidentifying restrictions, we connect to

⁴Given 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.

⁵Bonomo 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 (e.g. Woodford (2009) or Stevens (2014)) assume that memory, including assessing the passage of time, is costly. Therefore, in periods when the firm is inattentive, it does not index to aggregate inflation.

several literatures on pricing models that have grappled with one or more of these facts.⁶

First, in our model, prices tend to return to previous values, giving rise to discreteness and memory. This empirical regularity has been well documented following the seminal work of Eichenbaum et al. (2011), as it presents a challenge to standard state-dependent pricing theories that rely on a single fixed cost of a price change. To address this, the literature has used exogenously defined price plans (Eichenbaum et al. (2011)), heterogeneous menu costs (Kehoe and Midrigan (2015)), and rational inattention, emphasizing the discrete nature of the optimal signal structure under certain conditions (Matějka (2015) and Stevens (2014)). As discussed in the introduction, our mechanism differs from these frameworks both in terms of its micro-foundations, testable predictions and aggregate implications.

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, this does not result in kinks in conditional beliefs at old prices, and thus price stickiness and memory. In fact, the objective in introducing learning in existing models has not been to generate stickiness, but instead to match other facts, such as the shape of hazard function (Bachmann and Moscarini (2011), Baley and Blanco (2018)) or the pricing behavior over the product life cycle (Argente and Yeh (2017)). Our model also matches these facts, in addition to others such as stickiness and memory.

At its core, our framework fits within the literature motivated by the classic work of Ellsberg (1961), such as Gilboa and Schmeidler (1989), Dow and Werlang (1992), and Epstein and Schneider (2003). In the field of industrial organization, Bergemann and Schlag (2011) studies a static pricing problem with 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 consumer heterogeneity. In contrast, 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 set of demand curves leads to pricing behavior that is empirically supported and of interest for macroeconomic models.

2 Analytical Model

In this section, we develop the key insights of our mechanism in the context of a simple, analytically-tractable model that does not distinguish between real and nominal prices. We present the full nominal model in Section 3.

⁶In this, we follow the spirit of a broad literature that documents micro-level facts aimed at disciplining 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.

⁷An 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 the log demand

$$y(p_t) = x(p_t) + z_t, \quad (1)$$

where p_t is the log price. Demand consists of two components – the price-sensitive $x(p_t)$ and a price-insensitive component captured by z_t . The firm’s time- t realized profit is:

$$v_t = (e^{p_t} - e^{c_t}) e^{y(p_t)}, \quad (2)$$

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, $y(p_t)$, but not the underlying $x(p_t)$ and z_t separately. Furthermore, we model z_t as iid, and thus past demand realizations serve as 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 thus a convenient choice of a prior for doing Bayesian inference on function spaces. It has the defining feature that for any finite sub-collection of function inputs, e.g. a vector of prices $\mathbf{p} = [p_1, \dots, p_N]'$ for some $N > 1$, the corresponding vector of quantities demanded $x(\mathbf{p})$ 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 $m(p)$ controls the average slope of the underlying functions $x(p)$, and the covariance function $K(p, p')$ controls their smoothness. In other words, this distribution is a cloud of functions dispersed around $m(p)$, according to the covariance function $K(p, p')$.

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)):

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

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' . The larger is ψ , the faster the correlation between quantity demanded at different prices declines with the distance between those prices.⁸ This covariance function parsimoniously, yet flexibly, captures the natural prior view that there is an imperfect and declining correlation between demand at different prices. Additionally, this prior puts zero probability on demand functions that are not infinitely differentiable – thus any non-differentiability in the firm’s eventual worst-case perceptions about demand are fully attributable to the 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, $x^{DGP}(p) = -bp$, so that the potential $m(p)$ lie within an interval of width 2γ around $x^{DGP}(p)$,

$$m(p) \in [-\gamma - bp, \gamma - bp]. \quad (3)$$

The parameter $\gamma > 0$ controls 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,

$$m'(p) \in [-b - \delta, -b + \delta], \quad (4)$$

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. As an illustration, the left panel of Figure 1 provides some examples of admissible demand schedules $m(p)$, out of the infinite set of functions that satisfy (3) and (4). Later, we explain how an ambiguity-averse firm extracts out of this set the kinked worst-case prior shown in the right panel.

The overall interpretation of our setup 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

⁸A 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 $\frac{\psi}{\sqrt{2\pi}}$.

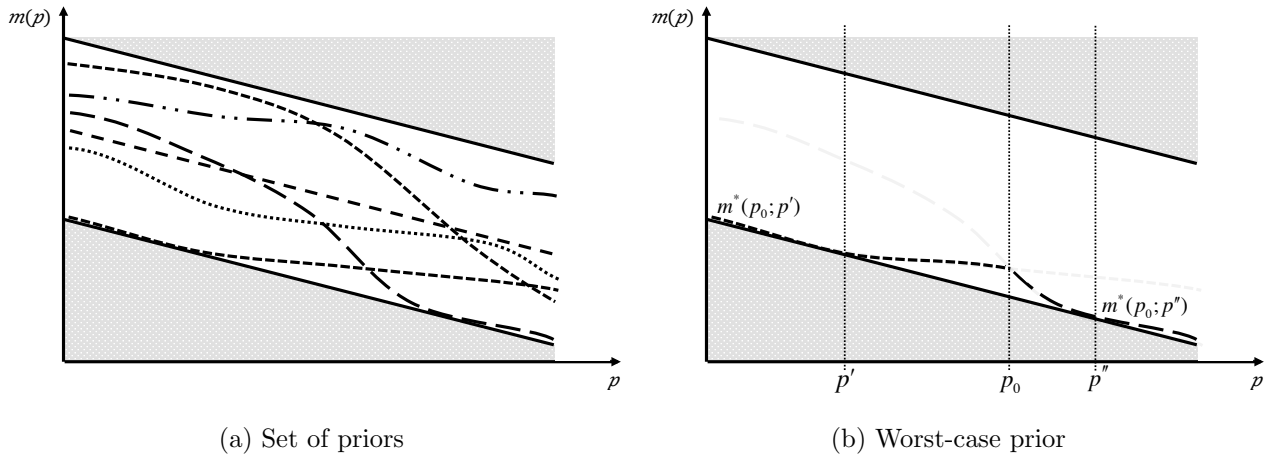


Figure 1. Illustrative set of priors and worst-case prior

demand schedules (i.e. a single prior), nor is it able to restrict attention to a particular parametric family of demand functions. The set of admissible beliefs may itself reflect the disagreement between heterogeneous, but otherwise unique, prior beliefs expressed by various agents inside the firm. The agent that takes the pricing decisions is not confident how to probabilistically weigh them as these beliefs are all entertained as reasonable priors.⁹

The parametrization of ambiguity characterizing the sets (3) and (4) 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, when $\gamma = 0$, the set Υ_0 collapses to a singleton, hence the firm has a unique prior and there is no ambiguity. 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. Uncertainty about the local elasticity of demand (i.e. $\delta > 0$) is at the heart of our mechanism.

2.1 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 $y^{t-1} = [y(p_1), \dots, y(p_{t-1})]'$ and the corresponding prices at which those were observed $p^{t-1} = [p_1, \dots, p_{t-1}]'$, where a superscript denotes history

⁹The connection between the set of beliefs about $m(p)$ to the dispersion of prior forecasts made by experts inside a firm allows us to empirically discipline the magnitude of ambiguity in the model's quantitative evaluation. The view that uncertainty is not primarily a probabilistic concept is consistent with the survey and experimental evidence involving business managers (see for example March and Shapira (1987)).

up to that time. The firm updates its beliefs about demand conditional on $\varepsilon^{t-1} = \{y^{t-1}, p^{t-1}\}$, observes c_t and posts a price p_t that maximizes expected profits, as detailed 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 resulting quantity sold $y(p_t)$, and the new 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, prior-by-prior – 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:

$$\begin{aligned}\hat{x}_{t-1}(p_t; m(p)) &:= E [x(p)|\varepsilon^{t-1}; m(p)] \\ \hat{\sigma}_{t-1}^2(p) &:= Var [x(p)|\varepsilon^{t-1}].\end{aligned}$$

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 firm is owned by an agent who is ambiguity averse and has recursive multiple priors utility (Epstein and Schneider (2003)), so that she values the firm’s profits as:

$$V(\varepsilon^{t-1}, c_t) = \max_{p_t} \min_{m(p) \in \Upsilon_0} E \left[v(\varepsilon_t, c_t) + \beta V(\varepsilon^t, c_{t+1}) \middle| \varepsilon^{t-1}, c_t \right], \quad (5)$$

where $v(\varepsilon_t, c_t)$ is the per-period profit defined in (2). The firm forms its conditional beliefs and 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 over the priors is conditional on the 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 $\hat{x}_{t-1}(p_t; m^*(p; p_t))$ at that price. Thus, for any price p_t the firm worries that, given the data it has seen, the underlying demand is low and hence maximizes over p_t under the worst-case belief $\hat{x}_{t-1}(p_t; m^*(p; p_t))$.

2.2 *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 realization $\bar{y}_0 = x(p_0) + \frac{1}{N_0} \sum_{i=1}^{N_0} z_i$. For a given prior $m(p)$, the joint distribution of the

signal and the unknown demand function x at any price p is:

$$\begin{bmatrix} x(p) \\ \bar{y}_0 \end{bmatrix} \sim N \left(\begin{bmatrix} m(p) \\ m(p_0) \end{bmatrix}, \begin{bmatrix} \sigma_x^2 & \sigma_x^2 e^{-\psi(p-p_0)^2} \\ \sigma_x^2 e^{-\psi(p-p_0)^2} & \sigma_x^2 + \sigma_z^2/N_0 \end{bmatrix} \right).$$

The distribution of $x(p)$ conditional on \bar{y}_0 is also Gaussian, and its expectation and variance are given by the familiar prior plus signal-updating formulas:

$$\begin{aligned} E(x(p)|\bar{y}_0, m(p)) &= m(p) + \alpha_{t-1}(p) [\bar{y}_0 - m(p_0)] \\ \text{Var}(x(p)|\bar{y}_0) &= \sigma_x^2(1 - \alpha_{t-1}(p)), \end{aligned} \quad (6)$$

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

$$\alpha_{t-1}(p) = \frac{\sigma_x^2}{\sigma_x^2 + \sigma_z^2/N_0} e^{-\psi(p-p_0)^2}. \quad (7)$$

Thus, the Bayesian update of the conditional expectation in equation (6) combines the prior for demand at that price, $m(p)$, with the information revealed by the difference between the observed signal realization \bar{y}_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 (6) 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 (3).

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 (6) as

$$E(x(p_t)|\bar{y}_0, m(p)) = \underbrace{(1 - \alpha_{t-1}(p_t))m(p_t)}_{\text{Prior demand at } p_t} + \underbrace{\alpha_{t-1}(p_t)(\bar{y}_0 + m(p_t) - m(p_0))}_{\text{Signal at } p_0 + \Delta \text{ in Demand between } p_t \text{ and } p_0},$$

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 \bar{y}_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

$$m^*(p_t; p_t) = -\gamma - bp_t.$$

Second, the firm is worried that demand changes for the worse as the price moves from p_0 to p_t , implying a low value of $m(p_t) - m(p_0)$. Thus, 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 $m(p)$ and the worst-case level for $m(p_t)$. Crucially, this implies a switch in the worst-case demand *shape* between p_t and p_0 , depending on whether the firm considers a price increase or a decrease.

Conditional on a price increase, i.e. $p_t > p_0$, 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 of $m(p)$ at any price p , and the maximal level of $m(p_0)$. Therefore, the worst-case prior for $m(p_0)$ when $p_t > p_0$ is

$$m^*(p_0; p_t) = \min[\gamma - bp_0, -\gamma - bp_t + (b + \delta)(p_t - p_0)]. \quad (8)$$

On the other hand, when the firm considers a price cut, i.e. $p_t < p_0$, 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 , and in particular the lower bound on the admissible derivative of demand in (4). Effectively, the firm worries demand is flat to the left of p_0 , hence, the worst-case $m^*(p_0; p_t)$ in this case is

$$m^*(p_0; p_t) = \min[\gamma - bp_0, -\gamma - bp_t + (b - \delta)(p_t - p_0)]. \quad (9)$$

Worst-case conditional expectation and kinks

Having characterized the worst-case prior, we can now plug it in equation (6) 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 $\hat{x}_{t-1}^*(p_t) \equiv E(x(p_t)|q_0, m^*(p; p_t))$ equals the following piecewise function

$$\hat{x}_{t-1}^*(p_t) = \begin{cases} -\gamma - bp_t + \alpha_{t-1}(p_t)[\bar{y}_0 - (-\gamma - bp_0)] - \alpha_{t-1}(p_t)\delta|p_t - p_0| & \text{if } p_t \in [\underline{p}, \bar{p}] \\ -\gamma - bp_t + \alpha_{t-1}(p_t)[\bar{y}_0 - (\gamma - bp_0)] & \text{if } p_t \notin [\underline{p}, \bar{p}] \end{cases} \quad (10)$$

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 (8) and (9) 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 (10). In essence, the overall worst-case expectation is the result of splicing two different priors together – an elastic one when evaluating prices to the right of p_0 , and an inelastic one to its left – which creates a kink, even though all individual priors are differentiable. Going back to Figure 1, panel (b) illustrates this splicing when entertaining setting some $p' < p_0$ or $p'' > p_0$, conditional on seeing a signal equal to the true DGP at a single price point p_0 and facing the set of priors in panel (a).

Putting everything together, the left panel of Figure 2 shows the resulting worst-case expected demand at any price p_t (i.e. it plots equation (10)). Extending these derivations to the case where ε^{t-1} contains observations at more than one price point is straightforward – 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 case of a single previously observed price: 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 the prior at the 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. To illustrate, panel (b) of Figure 2 plots the worst-case expectation when the firm has observed demand signals at two distinct price points p_0 and p_1 , and naturally the worst-case expectation is kinked around both of these prices.

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

When choosing its price to maximize expected profits under the worst-case beliefs, the problem of the firm is dynamic: 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 this section, 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 2.4 we provide analytical results for a tractable approximation to the forward-looking problem, before numerically analyzing it extensively in Section 4.

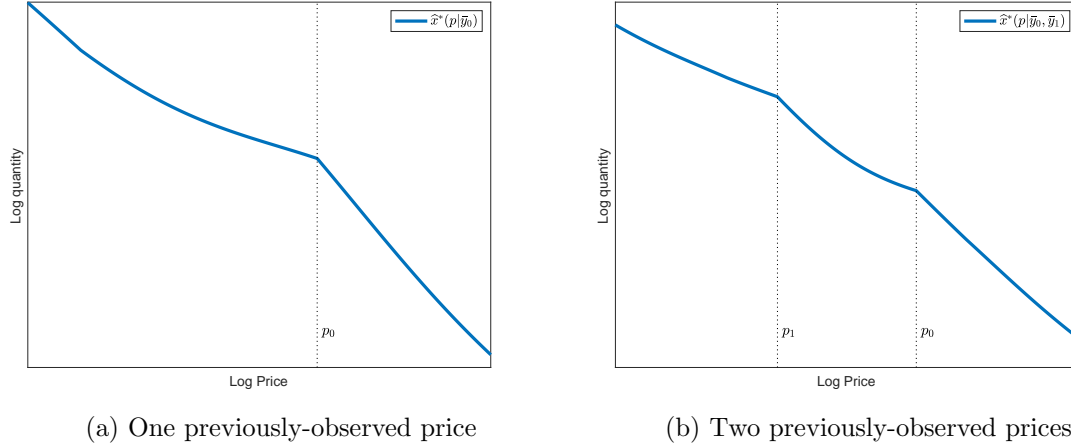


Figure 2. Worst-case Expected Demand

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

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

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 expand the analysis to the case where ε^{t-1} is endogenous in Section 4.

We start with the simplest case for the firm's information set and assume ε^{t-1} contains a single price p_0 , observed for N_0 number of times with an average quantity sold of \bar{y}_0 . As shown before, this results in a kink in the *as-if* expected demand, and in turn this provides the firm with an incentive to keep its price rigid even when faced with variations in costs. To show this insight analytically, in Proposition 1 we consider a log-linear approximation of expected profits around p_0 , which reveals a first-order loss of moving away from p_0 .

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

$$\ln v^*(\varepsilon^{t-1}, c_t, p_t) - \ln v_0^*(\varepsilon^{t-1}, c_t, p_0) \approx \left[\frac{e^{p_0}}{e^{p_0} - e^{c_t}} - (b + \alpha_{t-1}(p_0)\delta^*) \right] (p_t - p_0).$$

Proof. The switch in sign of δ^* follows from the worst-case expected demand in (10). 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. \square

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 around p_0 , a direct implication (as derived explicitly in Corollary 1) is that there is a positive interval of c_t realizations, around $c_0^* \equiv p_0 - \ln\left(\frac{b}{b-1}\right)$, for which the firm keeps its current price fixed at $p_t = p_0$.

Corollary 1. *Under the approximation in Proposition 1, p_0 is a local maximizer for any $c_t \in (\underline{c}_{t-1,0}, \bar{c}_{t-1,0})$, where $\underline{c}_{t-1,0} = c_0^* + \ln\left[\frac{b}{b-1} \frac{b-\alpha_{t-1}(p_0)\delta-1}{b-\alpha_{t-1}(p_0)\delta}\right]$ and $\bar{c}_{t-1,0} = c_0^* + \ln\left[\frac{b}{b-1} \frac{b+\alpha_{t-1}(p_0)\delta-1}{b+\alpha_{t-1}(p_0)\delta}\right]$.*

Proof. For any $c_t \in (\underline{c}_{t-1,0}, \bar{c}_{t-1,0})$ 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. \square

To gain intuition, consider an increase in cost to some $c_t > c_0^*$. This lowers the markup if the price remains 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 optimal markup the firm targets. 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 the threshold $\bar{c}_{t-1,0}$, the fall in markup would be too big, inducing the firm to change its price.

The logic is similar for a decrease in cost below 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 $\{c_0^*\}$, and thus the probability that p_0 is a local maximizer becomes zero. With ambiguity, this probability becomes strictly positive.

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 the 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 (7)), 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 :*

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

Proof. Follows from Corollary 1 and from $\frac{\partial \alpha_{t-1}(p_t)}{\partial N_0} > 0$ in equation (7). □

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 2.2 imply that there exist kinks in the *as-if* expected demand around each previously observed price level $p_i \in \varepsilon^{t-1}$. 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. *Let $\delta_i^* \equiv \delta \operatorname{sgn}(p_t - p_i)$ for all $p_i \in \varepsilon^{t-1}$. For a given realization of c_t , up to a first-order approximation around each such $p_i \in \varepsilon^{t-1}$:*

$$\ln v^*(\varepsilon^{t-1}, c_t, p_t) - \ln v_0^*(\varepsilon^{t-1}, c_t, p_i) \approx \left[\frac{e^{p_i}}{e^{p_i} - e^{c_t}} - (b + \alpha_{t-1,i}(p_i)\delta^* + A_i) \right] (p_t - p_i).$$

Proof. The switch in δ^* follows directly from the worst-case expected demand detailed in Online Appendix A.1. There we also define the term $\alpha_{t-1,i}(p)$ which denotes the weight put on the past demand realization at p_i in the conditional expectation, i.e. its effective signal-

to-noise ratio when updating with multiple signals. Lastly, A_i collects additional derivative terms that do not depend on the sign of $(p_t - p_i)$, for details see Online Appendix A.2. \square

Letting $c_i^* = p_i - \ln\left(\frac{b}{b-1}\right)$ 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 $(\underline{c}_{t-1,i}, \bar{c}_{t-1,i})$, where $\underline{c}_{t-1,i} = c_i^* + \ln\left[\frac{b}{b-1} \frac{b - \alpha_{t-1,i}(p_i)\delta + \hat{\alpha} - 1}{b - \alpha_{t-1,i}(p_i)\delta + \hat{\alpha}}\right]$ and $\bar{c}_{t-1,i} = c_i^* + \ln\left[\frac{b}{b-1} \frac{b + \alpha_{t-1,i}(p_i)\delta + \hat{\alpha} - 1}{b + \alpha_{t-1,i}(p_i)\delta + \hat{\alpha}}\right]$, such that for all $c_t \in (\underline{c}_{t-1,i}, \bar{c}_{t-1,i})$ 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 $p_i \in \varepsilon^{t-1}$. \square

Thus, Proposition 2 and Corollary 3 imply that the firm is not only reluctant to change its current price, but is in general inclined to repeat prices posted in the past, since there are kinks in the profit function there as well. This generates ‘memory’ in the price series.

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

Lastly, the analysis so far has focused on the first-order effect of price deviations around any of the $p_i \in \varepsilon^{t-1}$, showing that keeping the price fixed at previous levels is a local optimum. Next, we show that there is an additional interaction between the *level* of the past quantity sold \bar{y}_i and the change in the perceived demand slope around the price p_i at which that signal was observed. This interaction is of second-order, thus washes away in the local analysis above, but can matter for finding the global optimum, as we do in Section 4.

This interaction arises from the fact that the signal-to-noise ratio $\alpha_{t-1,i}(p)$ 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 some price p_i 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 $\alpha_{t-1,i}(p)$ is of second-order locally, but matters when thinking about the global maximum.

The second cross-derivative of the worst-case expected demand, with respect to price and the perceived innovation at p_i , denoted by $\hat{z}_i \equiv \bar{y}_i - (-\gamma - bp_i)$, is given by

$$\frac{\partial^2 \hat{x}_{t-1}^*(p_t; m^*(p; p_t))}{\partial p_t \partial \hat{z}_i} = -2\psi \alpha_{t-1,i}(p_t)(p_t - p_i).$$

Note that the derivative of the worst-case expected demand to the right (left) of p_i becomes more negative (positive) as the perceived innovation \hat{z}_i increases. Hence, a higher

signal innovation \hat{z}_i amplifies the effects of the endogenous switch in the worst-case demand elasticity. Intuitively, positive demand news shift up the conditional belief about demand at all prices, but the weight put on the signal decreases with $|p_t - p_i|$, so that beliefs about demand shift up the most locally. In Section 4 we investigate this interaction empirically and quantify how much a firm that observes a particularly good (bad) demand realization is more likely to keep (change) its posted price. This asymmetry in the effect of demand news on the probability of changing a price stands in contrast to most state-dependent mechanisms, such as a standard menu-cost model, where both positive and negative shocks make the firm more likely to reprice as they raise the gap between the current and optimal prices.

2.4 Incorporating forward-looking behavior

Next, we consider how forward-looking behavior affects optimal pricing, and stickiness in particular. The current price choice p_t and demand realization y_t become state variables in next period's problem, as they are 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.¹⁰

To characterize this exploration motive, we need to analyze the continuation value in (5). 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 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} .¹¹ Plugging it into (5), the firm solves

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

This approximation makes the dynamic problem tractable, while featuring two important conceptual advantages. First, the firm is forward-looking into the discounted infinite future

¹⁰Conceptually, 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.

¹¹ $\tilde{V}(\cdot)$ is the solution to the following recursive problem, with details presented in Online Appendix A.3

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

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 4 the long-run properties of the model at its stochastic steady state, where that history is fully endogenous and long. In this section, however, we will focus on analyzing the qualitative features of the economic forces shaping the exploration motive, and to this end we treat the history ε^{t-1} as given.

The experimentation motive is driven by a desire for information that is both new and relevant. On the one hand, the firm would like to obtain information on new parts of the demand curve where its ex-ante uncertainty is high. On the other hand, 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 determines whether the exploration incentives lead to the selection of a brand new price p_t or revisiting one of the previously observed prices. Which one dominates depends crucially on the information the firm enters the period with $-\varepsilon^{t-1}$.

To gain analytical insight into this trade-off, we consider the special case where i) $\psi = \infty$, so that beliefs about demand at different prices are uncorrelated; and ii) there is perfect foresight that future costs are constant at some arbitrary level $c > 0$, i.e. $c_{t+k} = c$ for all $k \geq 1$. Under these assumptions, we can characterize analytically the expected continuation value $E \left[\tilde{V}(\{\varepsilon^{t-1}, p_t, y_t\}, c) \middle| \varepsilon^{t-1}, p_t \right]$ as a function of p_t (the expectation is over the realizations of the new signal y_t), and prove two results that illustrate how the exploration incentive could be maximized either away from or exactly at one of the previously observed prices.

The composition of the history of observations ε^{t-1} is key to determining whether the optimal exploration strategy is to stay put or try something new. To illustrate, we consider two cases that would also help understand the numerical results in Section 4 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 starker, we assume that the realization of the observed signal \bar{y}_0 is good enough (i.e. $\bar{y}_0 > -\gamma - bp_0 + \frac{\sigma_x^2}{2}$), so that when the cost equals $c_0^* = p_0 - \ln(\frac{b}{b-1})$, p_0 is not just locally optimal, but is in fact the global static profit-maximizer conditional on ε^0 .

In Proposition 3, we characterize the current price p_t that maximizes the expected continuation value when $c = \bar{c}_0$.

Proposition 3. *The expected continuation value $E \left[\tilde{V}(\{\varepsilon^0, p_t, y_t\}, c_0^*) \middle| \varepsilon^0, p_t \right]$ achieves its maximum at*

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

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

Intuitively, choosing p_t^* today ensures that the new signal y_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 \hat{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 \hat{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 3 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, \bar{y}_0 and \bar{y}_1 , imply equally-good news, i.e. the same perceived innovation: $\hat{z}_0 = \hat{z}_1 = \hat{z}$.

Proposition 4 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 c_0^* . Thus, forward-looking behavior *reinforces* the static stickiness result (Corollary 1).

Proposition 4. *There is a non-singleton interval of costs (\underline{c}, \bar{c}) around c_0^* , and a threshold $\chi > 0$, such that if $\hat{z} > \chi$, then for any $c \in (\underline{c}, \bar{c})$:*

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

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. \square

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. Second, even if the firm receives a ‘disappointing’ new signal y_t at p_0 , it has a good fall-back option as it has also accumulated information (and thus reduced uncertainty) at the price level 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 *and* 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 4 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.

3 Quantitative Model and Nominal Rigidity

In this section, we embed our mechanism in a macroeconomic model with monopolistic competition. The key elements are that firms are uncertain about both (i) their demand curve and (ii) the competitors' price index. We first show analytically that this two-dimensional uncertainty gives rise to *as-if* kinks in demand in terms of *nominal* prices. Then, in the next section, we quantify the ability of our mechanism to match micro-level moments and generate monetary non-neutrality. In what follows, all lower case variables are in logs.

3.1 Structure of competition

The first primitive of the economic framework is the firm's set of *direct* competitors. We assume that firm i sells to a single industry j and in doing so, competes against a continuum of other monopolistically competitive firms who do the same. Each industry j has a representative final-good firm that aggregates the varieties i . Its cost-minimization problem implies a demand schedule $x_j(\cdot)$ for the good of firm i in industry j

$$y_{i,t} = x_j(p_{i,t} - p_{j,t}, y_{j,t}, z_{i,t}), \quad (11)$$

where $p_{i,t}$ is the log price set by firm i , and the log industry price index $p_{j,t}$ is such that $e^{p_{j,t}+y_{j,t}} = \int e^{p_{i,t}+y_{i,t}} di$.¹² The $z_{i,t}$ term is an idiosyncratic demand shock for good i which is unobserved by firm i but known to be distributed as $N(0, \sigma_z^2)$. The demand curve in equation (11) is a generalization of the typical CES structure, with the familiar result that the demand for a given intermediate good i is a function of the firm's price relative to the industry average, $p_{i,t} - p_{j,t}$; overall industry output $y_{j,t}$; and demand shocks $z_{i,t}$.

At the aggregate level, a representative household consumes a final good produced by a competitive firm that buys from the continuum of industries j . The household's consumption basket and the associated aggregate price index are given by the standard CES structures $y_t = \frac{b}{b-1} \ln \left(\int e^{y_{j,t} \frac{b-1}{b}} dj \right)$ and $p_t = \frac{1}{1-b} \ln \left(\int e^{p_{j,t}(1-b)} dj \right)$. Cost minimization by the final good producer implies a standard demand curve for the industry j composite good

$$y_{j,t} = y_t + b(p_t - p_{j,t}). \quad (12)$$

We denote the relative prices that enter as arguments in the demand curves for firm i in equation (11) and for industry j in equation (12), respectively, as

$$r_{i,t} \equiv p_{i,t} - p_{j,t}; \quad r_{j,t} \equiv p_t - p_{j,t}. \quad (13)$$

3.2 Information about competition

We model a firm that has Knightian uncertainty over the joint assessment of (i) its demand curve x_j as a function of its own relative price $p_{i,t} - p_{j,t}$, and (ii) the price index of its direct competitors $p_{j,t}$. Each firm i observes the full history of its own prices and quantities, $p_{i,t}$ and $y_{i,t}$, as well as the aggregate output and price levels, y_t and p_t . Intuitively, our framework is meant to capture the idea that since firms do not know the exact structure of the demand they face, they also do not know how to precisely aggregate over the prices of their direct competitors to build the relevant price index they compete against. Thus, uncertainty about the competitive environment manifests itself in uncertainty over the shape of the demand curve, but also the relevant price index that determines a specific firm's relative price.

¹²In the background, the technology is modeled as $e^{y_{j,t}} = f_j^{-1} \left(\int f_j(e^{y_{i,t}}) g_j(e^{z_{i,t}}) di \right)$, where each industry j has potentially different production functions f_j and g_j . Solving the cost-minimization problem of the final good firm in industry j yields $y_{i,t} = \ln \left[f_j'^{-1} \left(e^{p_{i,t} - p_{j,t}} \frac{f_j'(e^{y_{j,t}})}{g_j'(e^{z_{i,t}})} \right) \right]$. In equation (11) we summarize the effective demand curve as x_j and note that it is a transformation of the functions f_j and g_j .

Ambiguity about the demand curve

For tractability, we assume the firm understands that the industry demand $y_{j,t}$ and the demand shocks $z_{i,t}$ enter multiplicatively in the unknown function x_j in equation (11). Since firm i also knows the structure of the aggregate consumption basket, it can substitute out industry output $y_{j,t}$ from equation (12) to obtain the demand schedule

$$y_{i,t} = x_j(r_{i,t}) + br_{j,t} + y_t + z_{i,t}, \quad (14)$$

where the relative prices $r_{i,t}$ and $r_{j,t}$ are defined in (13).

Ambiguity about the demand curve x_j is modeled as in equations (3) and (4): there is a set of multiple priors, each of which is a GP distribution with mean function $m(r_i)$ so that

$$m(r_i) \in [-\gamma - br_i, \gamma - br_i]; \quad m'(r_i) \in [-b - \delta, -b + \delta]. \quad (15)$$

Ambiguity about the relative price

In our model, the firm does not directly observe its direct competitors' price index $p_{j,t}$. It does, however, have two relevant sources of information. These sources differ in the perceived ambiguity about their informational content. In particular, the firm is confident, i.e. perceives no ambiguity, about the first source, which consists of marketing reviews that perfectly reveal the value of $p_{j,t}$. We model reviews as occurring with some exogenous probability λ_T . Here we implicitly assume 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, leading the firm to perform them infrequently.¹³

In addition, the firm observes the aggregate price level p_t at all time. However, unlike in a rational expectations (RE) framework, we assume the firm is *not confident* about how p_t relates to the unknown $p_{j,t}$, and perceives their relationship as ambiguous. In particular, we assume that while the firm is certain that aggregate and industry prices are cointegrated and thus must keep pace with each other in the long-run, the firm is uncertain in the short-run structural relationship between the two. Putting together the firm's two sources of information, it perceives the evolution of p_{jt} as

$$p_{j,t} = \tilde{p}_{j,t} + \phi(p_t - \tilde{p}_{j,t}), \quad (16)$$

¹³As 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 a deterministic timing is computational: we find that stochastic review times achieve faster convergence towards the stationary distribution. The advantage over a state-dependent setup is tractability, as it avoids modeling a cost-benefit analysis of reviews.

where $\tilde{p}_{j,t}$ is the most recent review signal as of time t , and the function ϕ summarizes the unknown and ambiguous structural relationship between p_t and $p_{j,t}$. Indeed, our assumption that firms do not know the exact industrial structure (i.e. the function x_j) implies that they also do not know the exact equilibrium relationship between p_t and $p_{j,t}$ – different industry production functions imply different such structural relationships.

Ambiguity about ϕ is modeled with the same tools as the uncertainty about the demand function x_j . Specifically, we assume that the priors on ϕ are GP distributions, with mean functions that lie in a set Ω_ϕ around the true DGP $\phi(p_t - \tilde{p}_{j,t}) = p_t - \tilde{p}_{j,t}$. 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 possible ϕ 's the firm entertains.

In particular, we specify that for small inflationary pressure, i.e. when $|p_t - \tilde{p}_{j,t}|$ is less than some threshold Γ , the relationship is uncertain and the function ϕ lies in the interval

$$\phi(p_t - \tilde{p}_{j,t}) \in [-\gamma_p, \gamma_p], \text{ for } |p_t - \tilde{p}_{j,t}| \leq \Gamma. \quad (17)$$

This captures the idea that observing a small change in the aggregate price p_t does not convince the firm that the unobserved industry price has also changed. Online Appendix A.7, shows that such uncertainty about the local relationship between aggregate and industry inflation is well supported by the data. Lastly, to ensure that under all admissible priors $p_t - p_{j,t}$ is stationary, we make the set of potential ϕ grow with $p_t - p_{j,t}$ as inflation rises:

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

Unambiguous estimates of relevant relative prices and demand

The review signal $\tilde{p}_{j,t}$ is the only unambiguous estimate of $p_{j,t}$. The firm can use this signal to construct unambiguous estimates of the relative prices of interest, $r_{i,t}$ and $r_{j,t}$, as

$$\tilde{r}_{i,t} \equiv p_{i,t} - \tilde{p}_{j,t}; \quad \tilde{r}_{j,t} \equiv p_t - \tilde{p}_{j,t}. \quad (18)$$

Here, $\tilde{r}_{i,t}$ represents the firm's estimate of the relevant relative price driving its own demand curve, constructed using the firm's observed nominal price $p_{i,t}$ and the review signal $\tilde{p}_{j,t}$. In turn, $\tilde{r}_{j,t}$ is the estimate of the relative price that enters the industry j demand curve.

Using these expressions, we can decompose the relative prices $r_{i,t}$ and $r_{j,t}$ into a component over which the firm is confident and one that is perceived as ambiguous. Specifically, given the law of motion of $p_{j,t}$ in equation (16), the unknown relative prices $r_{i,t}$ and $r_{j,t}$

defined in (13), and their unambiguous estimates in (18), the decomposition is given by

$$r_{i,t} = \tilde{r}_{i,t} - \phi(\tilde{r}_{j,t}); \quad r_{j,t} = \tilde{r}_{j,t} - \phi(\tilde{r}_{j,t}). \quad (19)$$

Substituting the decompositions of the unobserved $r_{i,t}$ and $r_{j,t}$ in (19) into the demand equation (14) leads to

$$y_{i,t} = \underbrace{x_j(\tilde{r}_{i,t} - \phi(\tilde{r}_{j,t})) - b\phi(\tilde{r}_{j,t})}_{\text{ambiguous components of demand}} + b\tilde{r}_{j,t} + y_t + z_{i,t}, \quad (20)$$

which isolates the unambiguous and ambiguous components of demand. The former are given by the observed aggregate output y_t , the risky demand shock $z_{i,t}$ and the unambiguous estimates of the relative prices $(\tilde{r}_{i,t}, \tilde{r}_{j,t})$. The ambiguous components are due to ambiguity over (i) the demand curve x_j itself, with the set of priors given in equation (15); and over (ii) the relative prices $(r_{i,t}, r_{j,t})$, arising from the set of priors for ϕ in (17).

3.3 Optimization problem

Having described the competitive and information structures, we now turn our attention to the profit maximization problem. We assume that the production function of the intermediate-good firm i is given by $y_{i,t} = \omega_{i,t} + a_t + l_{i,t}$, where $\omega_{i,t}$ and a_t are an idiosyncratic and aggregate productivity shock respectively; and $l_{i,t}$ represents the (log) hours supplied by the household to firm i . The processes for these shocks are known:

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

where $\varepsilon_{i,t}^\omega$ is iid $N(0, \sigma_\omega^2)$ and ε_t^a is iid $N(0, \sigma_a^2)$. Because we are ultimately interested in how nominal shocks affect the firm's pricing decisions, nominal spending is the other aggregate shock in this economy: the log nominal aggregate spending $s_t = p_t + y_t$ is assumed to follow a random walk with drift, $s_t = \mu + s_{t-1} + \varepsilon_t^s$, where ε_t^s is iid $N(0, \sigma_s^2)$.

The aggregate side of the model is standard. We assume linear labor disutility for the representative household, which then equalizes the wage to consumption, which in turn is simply equal to aggregate output by market clearing (as detailed in the Online Appendix A.4). Given the log-demand in equation (20), the real flow profit of firm i becomes

$$v_{i,t} = (e^{p_{i,t} - p_t} - e^{y_t - \omega_{i,t}}) e^{y_{i,t}}. \quad (21)$$

The firm's problem can be summarized as follows: The firm enters period t with

knowledge of the history of all its previous quantities sold, the corresponding nominal prices at which those quantities were observed and its history of industry price review signals. In addition, the firm sees the history of aggregate prices and output. At the start of time t , the idiosyncratic productivity $\omega_{i,t}$ and aggregates (p_t, y_t) are observed. If a new review is conducted, $\tilde{p}_{j,t}$ equals $p_{j,t}$ and equals $\tilde{p}_{j,t-1}$ otherwise. Given the history of observables and current states, the firm optimizes over its action, $p_{i,t}$, taking into account the ambiguity over its demand in equation (20). At the end of period t , the demand shock $z_{i,t}$ is realized and the firm updates its information set with the observed quantity sold $y_{i,t}$ at price $p_{i,t}$.

Finally, the firm needs to conjecture a law of motion of p_t to forecast future profits. We close the model by assuming that the ambiguity-averse firms are measure zero, while the rest of the economy is populated by rational-expectations firms. This makes aggregate determination simple, as the equilibrium p_t is given by p_t^{RE} . The implication of this assumption is that by ignoring strategic complementarities in price setting, the quantitative benchmark of Section 4 can be seen as a lower bound on the degree of monetary non-neutrality.

For a transparent comparison between our model and the RE benchmark, we assume a simple DGP where each variety i faces the same demand function coming from industry j in (14), given by $x_j(r_{i,t}) = -br_{i,t}$. Using this knowledge of the demand function in equation (20), it follows that a RE firm has full knowledge that its demand is $y_{i,t}^{RE} = -b(p_{i,t} - p_t) + y_t + z_{i,t}$. 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 $p_t^{RE} = s_t - a_t$.

Aggregate price level and profits

Note that the aggregate price p_t affects real profits in equation (21) through three possible channels. The first is the standard effect of deflating nominal profits by p_t . The other two effects show up in the demand equation (20). On the one hand, when observing a higher p_t , holding $\tilde{p}_{j,t}$ constant, the firm estimates that the industry j 's composite good is relatively cheaper. As a result, the firm expects higher demand for industry j 's composite good, which in turn, holding everything else constant, translates into a higher demand for firm i . This demand shifter is given by $b\tilde{r}_{j,t}$ in equation (20). On the other hand, the same observation of a higher p_t may change the firm's perception of the unobserved $p_{j,t}$, through their structural relationship $\phi(\tilde{r}_{j,t})$. In particular, holding constant $\tilde{p}_{j,t}$, a larger p_t indicates to the firm that the aggregate price index is higher than the firm's unambiguous estimate of $p_{j,t}$. While in a RE model this observation would fully convince the firm that $p_{j,t}$ must have also risen, our ambiguity-averse firm lacks such confidence, as summarized in the set of beliefs in (17).

3.4 Joint worst-case beliefs

To illustrate the key intuition analytically, for the rest of this section we zero-in on the special case of 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 just one previous price point, $p_{i,0}$, and the quantity sold at that price, $y_{i,0}$.¹⁴ In addition, the firm observes the history of aggregates, $\{y_0, y_1, p_0, p_1\}$, and signals on the industry price level, $\{\tilde{p}_{j,0}, \tilde{p}_{j,1}\}$. These observables are used by the firm to form the unambiguous estimates of the relative prices, $\{\tilde{r}_{i,0}, \tilde{r}_{j,0}, \tilde{r}_{j,1}\}$, as defined in (18).

Abstracting from the unambiguous terms, expected demand is given by

$$m(\tilde{r}_{i,1} - \phi(\tilde{r}_{j,1})) - b\phi(\tilde{r}_{j,1}) + \alpha\{y_{i,0} - [m(\tilde{r}_{i,0} - \phi(\tilde{r}_{j,0})) - b\phi(\tilde{r}_{j,0})]\}, \quad (22)$$

where $\alpha = \frac{\sigma_x^2}{\sigma_x^2 + \sigma_z^2}$.¹⁵ The expectation is simply a combination of a prior belief about demand at the current relative price $r_{i,1}$ (i.e. $m(\tilde{r}_{i,1} - \phi(\tilde{r}_{j,1})) - b\phi(\tilde{r}_{j,1})$) and an update given the realized quantity $y_{i,0}$ at the previous price $r_{i,0}$. Since the industry prices $p_{j,0}$ and $p_{j,1}$ are uncertain, expected demand depends on the priors for both the demand curve, $m(\cdot)$, and the structural relationship between industry and aggregate prices, $\phi(\cdot)$.

The analysis of the worst-case priors is similar to the one in the real model in that it can be decomposed in jointly considering the worst-case priors for (i) the *level* of demand at the current relative price $r_{i,1}$ and (ii) the *change* in demand between $r_{i,1}$, and the previous $r_{i,0}$ (where the quantity sold $y_{i,0}$ was observed). However, the firm now needs to account for the joint ambiguity over both the demand shape, i.e. m , and the industry-aggregate price relationship ϕ , since the latter affects perceptions about the unknown $p_{j,0}$ (the industry price when $y_{i,0}$ was realized) and $p_{j,1}$ (the industry price that determines the current relative price).

The worst-case prior for the level of demand at the current price $r_{i,1}$ is straightforward – it is equal to the lower bound of (15), given by $-\gamma - b(\tilde{r}_{i,1} - \phi(\tilde{r}_{j,1}))$. Under this prior, the first two components of expected demand in (22) simplify to $-\gamma - b\tilde{r}_{i,1}$, as the $b\phi(\tilde{r}_{j,1})$ terms cancel out. This reflects our assumption that while the firm faces ambiguity about the local elasticity of its demand schedule, its beliefs are centered around the true $x_j(r_{i,t}) = -br_{i,t}$.

Still, the ambiguity about the shape of demand impacts the firm's evaluation of the change in demand from $r_{i,1}$ to $r_{i,0}$, and thus its update based on the realization of $y_{i,0}$. In

¹⁴In Online Appendix B.2 (on the authors' website) we show how the conceptual analysis is extended to having multiple past prices, and Section 4 shows numerical results for the general case of forward-looking firms with unrestricted history.

¹⁵To 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 4.

fact, there are two sources of ambiguity affecting this update, since the surprise in $y_{i,0}$ is

$$y_{i,0} - \underbrace{[m(\tilde{r}_{i,0} - \phi(\tilde{r}_{j,0})) - m(\tilde{r}_{i,1} - \phi(\tilde{r}_{j,1}))]}_{\text{Ambiguous change along firm's demand curve}} + \underbrace{b[\phi(\tilde{r}_{j,0}) - \phi(\tilde{r}_{j,1})]}_{\text{Ambiguous change in industry's relative price}} \quad (23)$$

The first source of ambiguity in the update is the firm's uncertainty over the shape of its demand curve between $r_{i,0}$ and $r_{i,1}$, governed by the set of priors m . The second source of ambiguity is due to the unobserved change in the industry price level from $p_{j,0}$ to $p_{j,1}$, and is governed by the set of priors on the industry-aggregate price structural relationship ϕ . Moreover, the unknown industry inflation impacts the firm's update in two places: (i) it affects the actual change in the relative price $r_{i,t}$ and enters the argument of the prior on the demand function m , and (ii) introduces a shift in the firm's demand schedule, as movements in $p_{j,t}$ against the aggregate price p_t change the economy's demand for the composite good of industry j , shifting demand for all firms in that industry (the second term in (23)).

Once the relevant constraints on the admissible functions m and ϕ are taken into account, determining the joint worst case for the ambiguous update in (22) reduces to

$$\min_{\delta' \in [-\delta, \delta]} \min_{\phi(\tilde{r}_{j,t}) \in [-\gamma_p, \gamma_p]} -\delta' \{ \tilde{r}_{i,1} - \tilde{r}_{i,0} - [\phi(\tilde{r}_{j,1}) - \phi(\tilde{r}_{j,0})] \}. \quad (24)$$

The joint worst-case prior beliefs depend on whether the firm considers a price that raises or lowers $\tilde{r}_{i,1}$ relative to $\tilde{r}_{i,0}$. When it entertains an action that increases $\tilde{r}_{i,t}$, 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 direct competition within industry j . Hence, if $\tilde{r}_{i,1} \geq \tilde{r}_{i,0}$, the minimizing priors are

$$\delta^* = \delta; \quad \phi^*(\tilde{r}_{j,1}) = -\gamma_p; \quad \phi^*(\tilde{r}_{j,0}) = \gamma_p. \quad (25)$$

In contrast, when the firm entertains decreasing its estimated relative price $\tilde{r}_{i,1}$, it worries about the opposite situation: that (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. Hence, if $\tilde{r}_{i,1} \leq \tilde{r}_{i,0}$,

$$\delta^* = -\delta; \quad \phi^*(\tilde{r}_{j,1}) = \gamma_p; \quad \phi^*(\tilde{r}_{j,0}) = -\gamma_p. \quad (26)$$

It is worth pointing out that the worst-case belief about the change in $p_{j,t}$ 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 x_j), and (ii) acts as a demand shifter since lower overall prices in industry j boost demand

for all firms inside the industry. These two effects go in opposite directions, and which one dominates depends on the perceived elasticity of the demand function x_j . When x_j takes on the average elasticity b , these two effects cancel out. However, when the firm is contemplating a price increase or decrease, the worst-case local elasticity of x_j changes away from b . As a result, the worst-case belief about ϕ endogenously depends on the firm's action $\tilde{r}_{i,1}$.

The key implication of equations (25) and (26) is that the joint worst-case beliefs over the demand curve and the industry-aggregate price relationship endogenously induce a kink in the worst-case conditional demand schedule around $\tilde{r}_{i,0}$, – notice they give rise to an absolute value term $-\alpha\delta|\tilde{r}_{i,1} - \tilde{r}_{i,0}|$ in (24). The reason behind the emergence of the kink is similar to the real model, and in particular equation (10). The difference is that while in the real model the relevant relative price was uniquely determined by the firm's action, here this is not the case. Intuitively, the firm is facing an identification problem, as it is uncertain about both the argument and the shape of the demand function. Faced with this joint ambiguity, it turns out that the robust solution is to estimate the demand curve in terms of the unambiguous estimate of the relative price $\tilde{r}_{i,t}$. Due to the uncertainty about the local shape of the demand function, a kink at the previously observed $\tilde{r}_{i,0}$ emerges.

3.5 Learning and nominal rigidity

The kink in the worst-case expected demand leads to a first-order expected loss of having an estimated relative price $\tilde{r}_{i,1}$ different from $\tilde{r}_{i,0}$. The firm can avoid this loss by posting a nominal price $p_{i,1}$ such that $\tilde{r}_{i,1} = \tilde{r}_{i,0}$, or

$$p_{i,1} - \tilde{p}_{j,1} = p_{i,0} - \tilde{p}_{j,0}. \quad (27)$$

Naturally, rigidity in $\tilde{r}_{i,t}$ has implications for the optimal nominal price.

First, consider the case in which a new review of the industry price does not occur at time $t = 1$, so that $\tilde{p}_{j,1} = \tilde{p}_{j,0}$. In that case, if the firm finds it optimal to take advantage of the kink in estimated relative prices, equation (27) shows that the firm will do so by keeping its nominal price fixed and set $p_{i,1} = p_{i,0}$. This makes the nominal price rigid.

Of particular importance is the result that the optimal nominal price $p_{i,1}^*$ may stay fixed at its previous value $p_{i,0}$ even as the aggregate price changes. The current aggregate price p_1 is one of the state variables that affects profits, as shown in equation (21) and discussed in Section 3.3, and thus affects pricing decisions. However, due to the kink in the worst-case expected demand, captured by the $-\alpha\delta|\tilde{r}_{i,1} - \tilde{r}_{i,0}|$ term discussed above, we show (see details in Proposition A1 in Online Appendix A.5) that there is a *range* of values of aggregate inflation for which it is optimal for the firm to keep $p_{i,1} = p_{i,0}$. A similar argument leads to

nominal rigidity conditional on changes in the other state variables as well.

Alternatively, consider periods when the firm observes the industry price through a review signal, i.e. $\tilde{p}_{j,1} = p_{j,1}$. To take advantage of the kink at $\tilde{r}_{i,0}$, the firm changes its nominal price away from $p_{i,0}$ in response to the information that is revealed by the difference $p_{j,1} - \tilde{p}_{j,0}$. Unless by chance $p_{j,1} = \tilde{p}_{j,0}$, rigidity in the estimated relative price leads to a nominal adjustment.

Indexation is suboptimal

Our model makes indexation to the aggregate price level suboptimal, even in the absence of external costs of changing prices. Intuitively, this is because the unambiguous estimate of the relevant relative price, $p_{i,t} - \tilde{p}_{j,t}$, is conceptually different from the real price, $p_{i,t} - p_t$. The firm optimally seeks to minimize exposure to ambiguity about its demand. Since demand is a function of the relative price $p_{i,t} - p_{j,t}$, the firm optimally keeps constant the unambiguous estimate of this relative price, not the real price as measured against the aggregate price. The two concepts of relative prices are materially different in our model because the observed aggregate price p_t is neither (i) the direct competitors' price index, nor (ii) an unambiguous estimate of it.

To see this, note that the observed aggregate price can be used to deflate all nominal prices to obtain real prices. We denote the resulting real versions of the price of the firm, the current industry j price, and the industry price observed at the last marketing review as

$$p_{i,t}^{\text{real}} \equiv p_{i,t} - p_t; \quad p_{j,t}^{\text{real}} \equiv p_{j,t} - p_t; \quad \tilde{p}_{j,t}^{\text{real}} \equiv \tilde{p}_{j,t} - p_t. \quad (28)$$

The key point is that as long as the competitive and information structure defined in Sections 3.1 and 3.2 are maintained, such re-normalization of prices based on variables in the firm's information set leads to the same optimal solution as derived above. Therefore, the nominal rigidity is not an artifact of a lack of properly deflating nominal prices or nominal illusion.

We can rewrite all of the basic relationships of our model using the deflated prices in (28). For example, the demand curve faced by the individual firm can be written as

$$y_{i,t} = x_j [p_{i,t}^{\text{real}} - \tilde{p}_{j,t}^{\text{real}} - \phi(-\tilde{p}_{j,t}^{\text{real}})] - b [\tilde{p}_{j,t}^{\text{real}} + \phi(-\tilde{p}_{j,t}^{\text{real}})] + y_t + z_{i,t}, \quad (29)$$

where the set of beliefs over the functions $x_j(\cdot)$ and $\phi(\cdot)$ is the same as in (15) and (17). Indeed, in the short run the firm still entertains that $\phi(-\tilde{p}_{j,t}^{\text{real}}) \in [-\gamma_p, \gamma_p]$. In our two-period case, equation (29) leads to a kink in the worst-case expected demand around the estimated

relative real price $p_{i,0}^{\text{real}} - \tilde{p}_{j,0}^{\text{real}}$. The local profit maximizer around that kink is to set

$$p_{i,1}^{\text{real}} - p_{i,0}^{\text{real}} = \tilde{p}_{j,1}^{\text{real}} - \tilde{p}_{j,0}^{\text{real}}. \quad (30)$$

But since aggregate inflation $p_1 - p_0$ cancels out on both sides of this equation, we obtain back the optimal nominal price solution given by equation (27).

Intuitively, when there is no review, the RHS of equation (30) implies that the observed aggregate inflation rate reduces the unambiguous estimate of the industry j real price one-to-one. To keep demand at the same estimated real price relative to the industry average, the firm finds it optimal to keep its nominal price unchanged, i.e. $p_{i,1} = p_{i,0}$, and thus let its real price also decrease one-to-one, as implied by the LHS of equation (30).

Suppose on the contrary that the firm were to index its nominal price using aggregate inflation and follow a pricing policy that sets $p_{i,1}^{\text{real}} = p_{i,0}^{\text{real}}$, and keeps its real price constant. Positive aggregate inflation implies a lower estimate of the real industry price $\tilde{p}_{j,1}^{\text{real}}$, compared to the last period. Keeping its own real price constant while that estimate has decreased is equivalent to an increase in the firm's real relative price vis-a-vis its direct within industry competitors, i.e. $p_{i,1}^{\text{real}} - \tilde{p}_{j,1}^{\text{real}}$ goes up. This is precisely what the firm wishes to avoid, as movements in this relative price expose it to first-order losses arising from the ambiguity about the demand shape, as discussed above.¹⁶ Thus, such indexation is suboptimal.

Stickiness and memory in nominal prices

The result of equation (27) extends in a straightforward fashion once we move beyond the example of a firm in its second period of life. The unrestricted history of estimated relative prices \tilde{r}_i^{t-1} and realized quantities forms the information set used to update beliefs about demand. Given past history \tilde{r}_i^{t-1} , the worst-case beliefs feature kinks at all previously observed $\tilde{r}_i \in \tilde{r}_i^{t-1}$. When a review does not occur this period so that $\tilde{p}_{j,t} = \tilde{p}_{j,t-1}$, the kinks in expected demand occur at the same set of nominal prices as last period.

These kinks make stickiness in nominal prices akin to “price plans”, where posted prices tend to bounce around a few “reference prices”. However, unlike other frameworks such as Eichenbaum et al. (2011), price plans adjust gradually over time: as shocks push the firm to visit a price it has not posted previously, it is added to its “price plan”. Importantly, in the future the firm is still likely to revisit the older prices, as the signals it has accumulated there

¹⁶Formally, these first-order losses dominate the standard markup and aggregate-demand effects of a change in p_t , as shown in Propositions A1 and A2, Online Appendix A.5. Naturally, we can modify either one of our two key primitives that (i) the economy is sub-divided in industries so that the set of firm's direct competitors is different from the set of all firms in the economy, and thus $p_{j,t} \neq p_t$, or that (ii) the firms are not confident that movements in p_t translate one-to-one in movements in $p_{j,t}$, in ways that would recover full nominal flexibility. Online Appendix B.3 (on the authors' website) has details on such alternative economies.

remain in its information set. As we discuss in Section 4.3, this feature has implications for the persistence of the real effects of monetary shocks.

4 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 infinite horizon decision problem of the ambiguity-averse firms. 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 2.4. 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.

4.1 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 Academic 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.¹⁷ Because our model does not feature a rationale for sales, all reported moments are based on “regular price” series in

¹⁷The 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).

which temporary sales are filtered out.¹⁸

Learning parameters and stochastic shocks

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 thus simply 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 (17). As detailed in Section 3, 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 perceived relative price \tilde{r}_i , the value of γ_p only enters as a price-independent demand shifter in the worst-case expectation. Its quantitative role is therefore limited and thus we study the limit of $\gamma_p \rightarrow 0$. This leaves four learning parameters $\{\gamma, \sigma_x^2, \psi, \lambda_T\}$ that we estimate by targeting micro-level moments, as detailed below.

The only modeling difference relative to the environment described in Section 3 is the 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.¹⁹ 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 allows us to study pricing behavior over the firm's life-cycle, which serves as an additional set of untargeted moment restrictions 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.

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-

¹⁸We 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.

¹⁹As such, we interpret resetting 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.

period-ahead quantity. This involves estimating the demand regression:

$$q_{ijt} = \beta_0 + \beta_1 q_{i,j,t-1} + \beta_2 p_{ijt} + \beta_3 p_{ijt}^2 + \beta_4 cpi_t + week'_i \theta_1 + store'_j \theta_2 + item'_i \theta_3 + z_{ijt} \quad (31)$$

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.²⁰ We then compute the empirical standard deviation of the residuals z_{ijt} leading us to set $\sigma_z = 0.613$.²¹

The firms also face idiosyncratic productivity shocks, whose persistence and volatility (respectively ρ_w and σ_w) we estimate via moment matching.

Simulated method of moments

We estimate the six free parameters, $\{\rho_w, \sigma_w, \sigma_x, \psi, \lambda_T, \gamma\}$, via simulated method of moments, by targeting the six pricing moments listed in Table II. 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).²² Table I presents all parameters values, while Table II shows the outcomes for moments targeted in the estimation.²³ The model matches the targeted 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.²⁴

4.2 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.

²⁰Given 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.

²¹The regression is run and the volatility measure is computed first for each of the 54 category/market pairs, before being aggregated using revenue weights.

²²The 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%.

²³The estimation is based on a simulated panel of 5000 time periods with 1000 active firms in each period.

²⁴Online Appendix A.6 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 I. 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	

All model moments are computed at the stochastic steady state. An important characteristic of the steady state is that the typical information set of a firm turns out to feature demand observations at a relatively small number of distinct price points. Because of this endogenous sparsity of the information set, the typical firm has several price points associated with low uncertainty that are both sticky and likely to be revisited, but also still faces substantial residual demand uncertainty away from those price points. Online Appendix A.8 provides a detailed discussion of the typical information set at the stochastic steady state.

We start by covering moments that have previously been analyzed in the literature, before turning to more novel features. The results are presented in Table III. All empirical moments are obtained by weighing results across markets and categories.

Table II. Targeted moments - Data vs model

	Data	Model
Frequency of regular price changes	0.107	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	0.535	0.533
Frequency of modal price changes (13-week window)	0.027	0.026
Mean duration of pricing regimes	29.90	30.54

Table III. 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)	0.829	0.880
Prob. price moves to modal P	0.592	0.646
Panel B Prob. revisiting old price (26-week window)	0.478	0.493
<i>uni</i> (26-week window)	0.774	0.822
Panel C Avg hazard slope	-0.011	-0.015
Panel D Kurtosis(ΔP distribution)	3.00	2.16
Prob. modal P revisits old modal P (52-week window)	0.085	0.112
Prob. revisiting price from before last modal P change	0.269	0.253
Prob. revisiting price from before brand new price	0.327	0.387

4.2.1 Reference prices, memory and size distribution

We first show that our model matches well a number of moments analyzed in the literature.

Reference prices. Panel A of Table III shows that the model 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 fraction of weeks that the regular price spends at the reference (modal) price: while in the data the regular price spends 83% of the time at the modal price, in the model this fraction equals 88%. We confirm that this is not simply a by-product of pervasive price stickiness: the probability that a non-modal regular price change ends at the modal price, and not at some other regular price, is 65% in the simulations compared to 59% in the data.

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). Panel B of Table III 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): this probability is 49% in our model, and 48% in the data. Note that a standard menu cost or Calvo model would feature no such price memory and the probability would be 0%.

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. To test this property, we produce a novel statistic that takes into account price stickiness: 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 $uni_{ijt} \equiv u_{ijt}/(c_{ijt} + 1)$.²⁵ Note that if all price changes end up at a price that had not been visited before within a specific window, the ratio uni_{ijt} is equal to 1. Yet, in the data we see that the average value of this moment is 0.774, as reported in Table III. This is very close to the simulated moment of 0.822.

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. 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 the model (left

²⁵In 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.

panel) and the data (right panel). While parameter heterogeneity across firms would allow for a smoother distribution, our model clearly can generate price changes of various sizes.




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

4.2.2 Hazard function of price changes

An important force in our setup is that, all else equal, a firm is less willing to move away from a price that it has stayed at for longer and thus acquired more information about. 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 .

The shape of the hazard function has been heavily discussed in the price-setting literature. 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 finding could be a by-product 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, our approach is to employ 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 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.²⁶

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

$$\mathbb{1}(p_{i,j,t} \neq p_{i,j,t-1}) = \alpha + \beta\tau_{i,j,t} + \gamma_i + \gamma_j + \gamma_t + u_{i,j,t} \quad (32)$$

²⁶In Online Appendix B.4 (on the authors’ website), we apply our econometric approach to panels of simulated data and show that it allows us to recover the true value of the slope of the hazard function, *even* in the presence of pervasive heterogeneity.

where the symbol $\mathbf{1}(\cdot)$ denotes the indicator function. 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 that would bias downwards the slope of the hazard. We run the regression on spells of 26 weeks or less, where the vast majority of observations lie.

Panel C of Table III shows that we obtain a slope estimate $\hat{\beta}$ of -0.011, once averaged across the 54 category/market pairs. 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. 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.²⁷

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.

In Online Appendix A.9 we apply the alternative approach of Campbell and Eden (2014) and reach the same conclusion: the hazard of regular price changes is downward sloping and of similar magnitude in both model and data.

4.2.3 Pricing behavior over the product life-cycle

Using the same dataset as ours, Argente and Yeh (2017) find that both the frequency and size of price changes decline significantly with the age of the typical the product. 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 2.4. 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 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

²⁷We cluster standard errors at the store level, the cluster which yields the highest standard errors. Also, in line with the literature, we drop all left-censored spells from the sample. Lastly, in Figure B.1 of the Online Appendix B.5 (on the authors’ website), we plot the distribution of coefficient estimates $\hat{\beta}$ across the 54 category/market pairs.

model by running the following two regressions on the simulated data:

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

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

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.²⁸

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).

4.2.4 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 Result 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.

To test this prediction, we first extract demand innovations in the data by using regression (31), which was 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 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 the average of the demand innovations at that price.²⁹ It is also useful to define a version of the $zscore$ that only incorporates demand innovations up to $t - 1$, and which therefore

²⁸Focusing 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.

²⁹We truncate the window to 26 weeks 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.

informs the price choice at t :

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

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:

$$\mathbb{1}(p_{i,j,t} \neq p_{i,j,t-1}) = \beta_0 + \beta_1(zscore_{ij,t-1} - zscore_{ij,t-1}^{lag}) + \beta_2wscore_{ij,t-1} + f_{ij} + \varepsilon_{ijt}. \quad (33)$$

The LHS equals 1 when the price at t is different than at $t-1$, and 0 otherwise. The regressor of interest, $zscore_{ij,t-1} - zscore_{ij,t-1}^{lag}$, 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 $t-1$.³⁰ 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 IV presents the results of running the regression in equation (33) 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 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

³⁰To 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. $p_{i,j,t-1} = p_{i,j,t-2}$.

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 IV. Results from the z -score regressions

	Data				Model	
	$x = 12$	$x = 25$	$x = 12$	$x = 25$	$x = 12$	$x = 25$
$wscore_{ij,t-1} \leq x$						
$zscore_{ij,t-1} - zscore_{ij,t-1}^{lag}$	-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.

4.3 Monetary non-neutrality

In this section, we argue that our theory is relevant not only due to its successful micro-level predictions, but also because the model alters the standard relationships highlighted in the literature between micro-level moments and the propagation of nominal shocks.

We quantify the degree of monetary non-neutrality by computing the impulse response of total output produced by the measure-zero set of ambiguity-averse firms to an innovation in aggregate nominal spending. 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 any strategic complementarities in price setting.




Figure 4. Nominal spending effects on real output.

We estimate the impulse responses via Jordá (2005) projections, an approach well suited to the high degree of non-linearity in our model. In the expression below, we regress the $t + k$ output of ambiguity-averse firms on 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 coefficients β_k represent the impulse response of output to the nominal shock, plotted by the solid line in Figure 4, panel (a). The response is shown as a fraction of the shock. We find that, for a 1% shock, real output increases by 0.38% on impact, and this increase is persistent, with full a cumulative output effect of 7.2% after 52 weeks.

Alternative models

To provide context for our results, we contrast our mechanism with three simple and widely used alternative mechanisms of nominal rigidity.³¹ The first is a Calvo model, where the firm can change its price with an exogenous probability. The second is a menu cost version, where the firm can change its price at any time by paying a fixed cost. The third model uses Kimball (1995) preferences which create a smoothed version of a kink in demand - a form of real rigidity. To generate nominal rigidity in that model, we follow the literature and assume

³¹Online Appendix B.7 (on the authors' website) details instead comparative statics within our model, focusing on the resulting nuanced link between price flexibility, memory, and non-neutrality.

that the firm must pay a small menu cost.³² Besides these differences, all specifications share the same economic framework: we study a similar measure-zero sector and assume the same cost processes as in our benchmark setup.

The Kimball model represents a particularly useful comparison because even if it shares a similar kinked-demand flavor, the underlying mechanism is very different from ours. For one, our framework generates stickiness without additional fixed costs since it features kinks that are not smoothed-out. Second, our model generates *perceived* kinks in demand; in a Kimball world, an econometrician would be expected to find evidence of actual kinks in demand schedules. This property is not innocuous as it has proven difficult to find evidence for the large super-elasticity values that are needed to jointly generate a significant persistence of monetary shocks and plausible micro price facts (see the evidence in Dossche et al. (2010) and the analysis in Klenow and Willis (2016)).

Our model of rigidity is also consistent with a range of facts, such as experimentation and effects of the level of past demand realizations on price setting, that are driven by our learning forces which are absent in the three considered alternative models. A crucial property here is that in the alternative models, conditional on a price change the gap towards the new optimal frictionless price is perfectly closed and the probability that the firm visits an old price is zero. Instead, in our framework, the firm does not typically eliminate this gap, preferring to return to a previously visited “safe” price due to ambiguity aversion. As we discuss below, this type of price memory has important implications for monetary non-neutrality.

Our first set of comparisons are based on calibrating the respective free parameter of each of the alternative models (Calvo probability or menu cost) to the same 10.5% frequency of regular price changes as in our ambiguity model. The impulse responses corresponding to these models are marked by the non-solid lines in Figure 4, panel (a). Overall, the cumulative real effect in our model is similar to that of the Calvo model (7.58%), and significantly larger than in the menu cost and Kimball models (1.12% and 1.81% respectively).

Impulse responses

The most striking difference is in the degree of persistence. In our model, the real effect only dies out after 48 weeks and has a half-life of 18 weeks. In contrast, in the menu cost and Kimball models the real effects have half-lives of less than 3 weeks, and disappears within 10 weeks, while in Calvo it has a half-life of 8 weeks and dies out 32 weeks in. The different persistence across these standard models is due to the Golosov and Lucas (2007) selection effect - the adjusting firms in the Calvo framework are chosen randomly, but are

³²See for example Klenow and Willis (2016), whom we follow to use a value of 10 for the demand super-elasticity (elasticity of the demand’s price elasticity) in the Kimball aggregator.

self-selected in the two other versions and hence adjust by a lot on average, leading to a quick transmission of the shock.

The fundamental reason for the significantly larger persistence in our model is that, consistent with the data, it features memory in nominal prices: a substantial proportion of price movements occurs between perceived demand kinks that have formed prior to the shock. As the firm revisits old price points, aggregate adjustment is slowed down, even though prices look very flexible at the micro level. Indeed, our environment features something akin to a “price plan” – a collection of low-uncertainty prices that the firm has visited in the past and switches between relatively flexibly.

This flexibility when switching between low-uncertainty, previously-visited prices accounts for the muted real effect on impact of the nominal spending shock. Intuitively, as noted by Alvarez and Lippi (2019), in models that have price plan-like behavior (like ours) firms actively use their free but imperfect margin of adjustment, leading to significant flexibility on impact. Yet, in our model, the blunted real effect on impact is more than made up for by the high persistence, leading to a large cumulative real effect, as well as to impulse response dynamics that are not well approximated by any of the three alternative models.

Frequency and kurtosis of regular price changes

Alvarez et al. (2016) show that the cumulative degree of nominal non-neutrality in a wide range of models, including the three alternative frameworks we study, is proportional to the kurtosis of the regular price change distribution. Intuitively, this proportionality arises because in this class of models the derivative of the density of price gaps at the adjustment thresholds, and implicitly the Golosov and Lucas (2007) selection effect, becomes small only for sufficiently leptokurtic distributions of price changes.

In contrast, as noted by Alvarez and Lippi (2019), frameworks that feature price plans, like ours, do not fit the class of models analyzed in Alvarez et al. (2016). Indeed, we find that our mechanism does not abide by this sufficient statistic relationship, in a manner that is quantitatively important. As we report in Panel D of Table III, the distribution of the regular price changes has a kurtosis of 3 in the data, while our model generates a kurtosis of 2.16. The standard analysis would predict that such a model-implied kurtosis, if anything an underprediction of its empirical value, would be associated with weak real effects of nominal shocks. However, even though the kurtosis in our model is significantly lower than in the Calvo version, (2.16 versus 5.7), the cumulative real effects are very similar (7.2% versus 7.6%). In fact, our kurtosis is much closer to that of the menu cost and Kimball models, at 1.23 and 1.79 respectively, yet our model implies a cumulative real effect that is six times larger than in the menu cost model.

The broad observation that price-plan-like behavior is important for mapping moments of regular price changes to the propagation of nominal shocks is not unique to our model – see previous discussion in Eichenbaum et al. (2011) (EJR) and Kehoe and Midrigan (2015). Yet, we argue next that the *nature* of price memory in our model is fundamentally different than in existing frameworks, in a way that has important implications for aggregate dynamics.

Gradual adjustments in price plans

If our mechanism was simply an endogenous version of the EJR “price plan”, then the persistence of the real output effect to a money shock would be governed by the frequency of changes in *reference* prices. In particular, Alvarez and Lippi (2019) show that a simple menu cost model calibrated to the frequency of reference price changes provides a useful upper bound on the effect of such a “price plan” model. But that is not true for our model – panel (b) of Figure 4 shows that our model’s real effects are significantly more persistent and as a result cumulatively larger (7.2% versus 5.5% and 6%, respectively) than those from a menu cost or Kimball model matched to the same frequency of reference price changes.

The reason is that the price plans in our model adjust *gradually*. In a standard price-plan model like EJR, once a firm decides to update its plan, it resets all prices within the plan. In our model, instead, the price plans evolve slowly, with new prices being added to the effective plan as the firm experiments with and learns about demand at new price points. The reason is that when shocks drive a firm to post a price outside of the set it has been visiting in the past, this does not destroy the information capital it has built up at its old price points. Hence, when idiosyncratic shocks mean-revert, the firm is likely to revisit those old price points. This is important: the fact that the firm returns to known prices, even after sampling a new portion of the price space, is what slows down the evolution of the aggregate price level beyond what is captured by the frequency of reference price changes.

We quantify the importance of these gradual price-plan adjustments in both the model and the data in three ways, which we report in Panel D of Table III. First, we consider memory in reference, as opposed to regular, prices: we compute the probability that a change in the reference price (i.e. the modal price in a 13-week window) revisits one of the *old* reference prices set in the past 52 weeks. The previous literature has not studied potential memory in reference prices; under the implicit assumption that reference price changes capture price plan resets, such as in EJR, such memory should be zero. Instead, we find that these probabilities are 8.5% and 11.2% in the data and model respectively, pointing at significant memory in reference prices.

Next, we try to identify gradual price plan changes. Consider that a reference price change occurs at time t . We compute the probability that the first price reset following the

reference price change leads to a price that had already been posted *before* period t . In standard memory models such as EJR, this probability would be 0%, as firms reset their whole price plan at once. Yet, we find this probability to be 27% in the data, compared to 26% time in the model. We also consider an alternative approach to capturing partial price plan changes. We first identify occurrences of “brand new prices”, which are defined as prices that have not been posted in the previous 52 weeks. We then compute the probability that the first price change that follows a time- t brand new price revisits a price that was posted before t (ostensibly from before the shift in the “price plan”). This probability is equal to 33% in the data, and 39% in the model.

Summary

Overall, our theory offers a micro-foundation for price memory that delivers novel implications at both the micro and macro levels. First, we show that our unified theory of rigidity and memory *jointly* matches a host of over-identifying restrictions on micro-level pricing behavior. These additional moments not only provide external validity of the underlying theory, but also imply a more nuanced picture of how the observed micro-level price flexibility maps into the aggregate propagation of nominal shocks. Second, our framework endogenously generates slow and gradually adjusting price plans. This feature implies much more persistent real effects from monetary shocks than in a benchmark menu cost model calibrated instead to the frequency of reference price changes. As a result, we believe that our mechanism may be better suited at capturing the high persistence of monetary policy effects estimated in the data (Christiano et al. (2005)).

5 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 our 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. Importantly, the theory has novel predictions about the way in which nominal shocks affect the aggregate economy.

References

ALVAREZ, F., H. LE BIHAN, AND F. LIPPI (2016): “The real effects of monetary shocks in sticky price models: a sufficient statistic approach,” *American Economic Review*, 106,

2817–51.

- ALVAREZ, F. AND F. LIPPI (2019): “Temporary Price Changes, Inflation Regimes and the Propagation of Monetary Shocks,” *American Economic Journal: Macroeconomics*, forthcoming.
- 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.
- 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.
- DOSSCHE, M., F. HEYLEN, AND D. VAN DEN POEL (2010): “The kinked demand curve and price rigidity: Evidence from scanner data,” *Scandinavian Journal of Economics*, 112, 723–752.
- 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.
- EINHORN, H. J. AND R. M. HOGARTH (1986): “Decision making under ambiguity,” *Journal of Business*, S225–S250.
- 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.
- JORDÁ, O. (2005): “Estimation and Inference of Impulse Responses by Local Projections,” *American Economic Review*, 95, 161–182.
- 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.

- KUNREUTHER, H., R. HOGARTH, AND J. MESZAROS (1993): “Insurer ambiguity and market failure,” *Journal of Risk and Uncertainty*, 7, 71–87.
- MACHINA, M. J. AND M. SINISCALCHI (2014): “Ambiguity and ambiguity aversion,” in *Handbook of the Economics of Risk and Uncertainty*, Elsevier, vol. 1, 729–807.
- MACKOWIAK, B. AND M. WIEDERHOLT (2009): “Optimal Sticky Prices under Rational Inattention,” *American Economic Review*, 99, 769–803.
- MAFFIOLETTI, A. AND M. SANTONI (2005): “Do trade union leaders violate subjective expected utility? Some insights from experimental data,” *Theory and Decision*, 59, 207–253.
- 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.
- MARCH, J. G. AND Z. SHAPIRA (1987): “Managerial perspectives on risk and risk taking,” *Management science*, 33, 1404–1418.
- MATĚJKA, F. (2015): “Rationally inattentive seller: Sales and discrete pricing,” *The Review of Economic Studies*, 83, 1125–1155.
- 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.