B Online Appendix B (Not for publication)

B.1 Exploration incentives

In Section 2.4 we assumed $\psi = \infty$ for analytical tractability. Relaxing that assumption generally reduces the experimentation incentives of the firm, in the sense that it flattens the continuation value $\tilde{V}$. The reason is that when $\psi < \infty$, observing a signal y_t at a price p_t is informative not only about $x(p_t)$ itself, but also about other prices p around p_t, with the informativeness dropping to zero as the distance $|p - p_t|$ goes to infinity. Moreover, a higher ψ implies that the correlation between $x(p)$ and $x(p')$ at distinct p and p' decreases faster with the distance between p and p'. Hence, higher ψ increases the specificity of new information, making it more localized.

Lower ψ on the other hand, makes the information at a given p_t more useful at *any* p. As a result, this erodes the firm's incentive to experiment with new prices – it could learn most of the same information by repeating one of its established, safe prices anyways. Formally, this means that the continuation value function $\tilde{V}$ becomes flatter. In fact, as we show in Proposition B1 below, in the limit $\psi \to 0$ the continuation value is a perfectly flat line.

Proposition B1. *The expected continuation value* $E\left[\tilde{V}(\{\varepsilon^0, p_t, y_t\}, c_0^*)|\varepsilon^0, p_t\right]$ *becomes flat in respect to the time t price* p_t *as* $\psi \to 0$*:*

$$\lim_{\psi \to 0} \frac{\partial E\left[\tilde{V}(\{\varepsilon^0, p_t, y_t\}, c_0^*)|\varepsilon^0, p_t\right]}{\partial p_t} = 0$$

Proof. First we will prove that with $\psi < \infty$, the expected continuation value $E\left[\tilde{V}(\{\varepsilon^0, p_t, y_t\}, c_0^*)|\varepsilon^0, p_t\right]$ is differentiable. The key intuition is that if the firm selects a time t price away from p_0, thus obtaining a signal at a new price $p_t \neq p_0$, in expectation this would not create a second kind in the expected future worst-case demand. The only kink in the time t expectation of the future worst-case demand appears at the already observed p_0, since it evolves recursively as:

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

where

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

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

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

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

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

We are going to use the notations for signal innovation level, $\widehat{z}_t$, and the signal-to-noise ratios defined above. Also recall that $E\left[\tilde{V}(\{\varepsilon^0, p_t, y_t\}, c_0^*)|\varepsilon^0, p_t\right] = \frac{\beta}{1-\beta} E\left[\nu^*_{t+k}(p^*_{t+k}, c_0^*)\middle|\varepsilon^0, p_t\right]$, where p^*_{t+k} is the resulting static optimal price, given the updated information set $\{\varepsilon^0, p_t, y_t\}$. And to simplify notation, we will again use the shorthand $E_{t-1}(\tilde{V})$ to denote the expected continuation value $E\left[\tilde{V}(\{\varepsilon^0, p_t, y_t\}, c_0^*)|\varepsilon^0, p_t\right]$.

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

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

Let's start with showing that $\frac{\partial E_{t-1}(\tilde{V}|p_t > p_0)}{\partial p_t}$ exists everywhere. The firm has perfect foresight on $c_{t+k} = c_0^*$, and since $p_0 = \ln(b/(b-1)) + c_0^*$ absent any information in the new signal y_t the optimal price at $t+k$ would be p_0. Thus, the worst-case expected profit given a choice of $p_t > p_0$ can be written as:

$$\begin{aligned} E_{t-1}(\tilde{V}|p_t > p_0) = {} & \Phi(\underline{z}(p_t))E_{t-1}(\nu^*_{t+k}(p^*(p_t), c_0^*|p_t > p_0, \widehat{z}_t < \underline{z}(p_t)) + (\Phi(\bar{z}(p_t)) - \Phi(\underline{z}(p_t)))E_{t-1}(\nu^*_{t+k}(p_0, c_0^*|p_t > p_0) \\ & + (1 - \Phi(\bar{z}(p_t)))E_{t-1}(\nu^*_{t+k}(p^*(p_t)|p_t > p_0, \widehat{z}_t > \bar{z}(p_t)) \end{aligned}$$

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

It is straightforward to show that the expected profit function (at any price p), $E_{t-1}(\nu^*_{t+k}(p, c_0^*)|p_t > p_0)$, is differentiable in respect to p_t:

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

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

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

where

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

is the signal-to-noise ratio applicable to the y_0 signal. This only depends on p_t through $\alpha_t(p)$, hence it is differentiable as well.

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

$$p^* - (c^*_0 + \ln(\frac{\widehat{x}^{*\prime}_{t-1}(p^*) + \alpha'_t(p^*)\widehat{z}_t + \frac{1}{2}\widehat{\sigma}^{2\prime}_t(p^*)}{1 + \widehat{x}^{*\prime}_{t-1}(p^*) + \alpha'_t(p^*)\widehat{z}_t + \frac{1}{2}\widehat{\sigma}^{2\prime}_t(p^*)})) = 0 \tag{38}$$

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

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

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

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

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

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

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

$$\Longleftrightarrow$$

$$(e^{p_0}-e^{c_0^*})\exp(\widehat{x}^*_{t-1}(p_0)+\alpha_t(p_0)\bar{z}(p_t)+\frac{1}{2}(\widehat{\sigma}_t^2(p_0)+\sigma_z^2)) = (e^{p^*}-e^{c_0^*})\exp(\widehat{x}_{t-1}(p^*)+\alpha_t(p^*)\bar{z}(p_t)+\frac{1}{2}(\widehat{\sigma}_t^2(p^*)+\sigma_z^2))$$

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

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

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

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

However, in the limit $p_t \to p_0$ the candidate interior prices and thresholds converge to the same values. The candidate interior price is given by the first-order condition (38), the minimum threshold $\lim_{p_t\to p_0}\underline{z}(p_t) = \underline{z}$ is defined as

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

$$\Longleftrightarrow$$

$$\begin{aligned}&(e^{p_0} - e^{c_0^*})\exp(\widehat{x}^*_{t-1}(p_0) + \alpha_t(p_0|p_0 = p_t)\underline{z} + \frac{1}{2}(\widehat{\sigma}_t^2(p_0|p_0 = p_t) + \sigma_z^2))\\ &= (e^{p^*} - e^{c_0^*})\exp(\widehat{x}_{t-1}(p^*) + \alpha_t(p^*|p_0 = p_t)\underline{z} + \frac{1}{2}(\widehat{\sigma}_t^2(p^*|p_0 = p_t) + \sigma_z^2))\end{aligned}$$

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

With that in mind we can show

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

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

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

$$\lim_{\psi \to 0} \frac{\partial E\left[\tilde{V}(\{\varepsilon^0, p_t, y_t\}, c_0^*)|\varepsilon^0, p_t\right]}{\partial p_t} = 0$$

□

B.2 Joint uncertainty over demand shape and relative price

In section 3.4 we have developed the solution to the worst-case beliefs when the firm observes one previous unambiguous estimated relative price, which here for brevity we call an estimated relative price. In this appendix we show how the analysis extends to multiple prices. The analysis follows the similar logic as in the real model, detailed in appendix A.1, with the added analysis of the worst-case belief of the unknown industry price. We do so by presenting details on the case of updating beliefs in the third period of life, when the firm has seen demand realizations at two previous prices $p_{i,0}$ and $p_{i,1}$, with corresponding quantities sold there $y_{i,0}$ and $y_{i,1}$. In addition, the firm observes the history of aggregates, $\{y_0, y_1, y_2, p_0, p_1, p_2\}$, and signals on the industry price level, $\{\widetilde{p}_{j,0}, \widetilde{p}_{j,1}, \widetilde{p}_{j,2}\}$. We will use the helpful $\widetilde{r}_{i,t} = p_{i,t} - \widetilde{p}_{j,t}$ notation for the unambiguously estimated relative price. In particular, without loss of generality, suppose that the prior observations imply unambiguously estimated relative price such that $\widetilde{r}_{i,0} < \widetilde{r}_{i,1}$, where the analysis for the opposite case is analogous.

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

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

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

where α_0 and α_1 are weights on the perceived innovations in the signals $y_{i,0}$ and $y_{i,1}$, respectively.

The prior belief about demand at $r_{i,0}$ and $r_{i,1}$ can be written as

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

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

We can use the definition of $r_{i,t} \equiv p_{i,t} - p_{j,t}$ and substitute $p_{j,t}$ from equation (16) to simplify the portion of the demand estimate over which nature chooses the joint worst-case demand shapes δ_0' and δ_1', together with the short-run co-integrating relationship $\phi(p_t - \widetilde{p}_{j,t})$, as follows:

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

We obtain the solution for the joint worst-case

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

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

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

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

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

The joint worst-case beliefs are given by

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

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

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

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

and the worst-case beliefs:

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

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

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

B.3 Counter-factual economies where indexation is optimal

Naturally, if either of our two key primitives on the structure of the economy or the structure of uncertainty is modified, then we recover full nominal flexibility. For example, if there is no industrial structure and firms understand they compete directly against all other firms in the economy, then the observed aggregate price p_t is the price index of the firm's direct competitors.

On the other hand, even if there is industrial structure with unknown industry-level demand functions, but the firms are somehow fully confident that movements in p_t translate one-to-one in movements in the underlying $p_{j,t}$, then p_t provides an unambiguous signal of the relevant $p_{j,t}$.

Below we analyze both of these alternative economies in detail. In particular, first we consider an economy where firm i directly competes against the aggregate price index p_t. Alternatively, on the information side, we assume that the firm still competes against the unobserved $p_{j,t}$, but is now endowed with full knowledge of the true DGP $\phi(p_t - \widetilde{p}_{j,t}) = p_t - \widetilde{p}_{j,t}$. Using the law of motion of $p_{j,t}$ in equation (16), the firm is now confident that $p_{j,t} = p_t$.

As in our benchmark model, in both of these alternative economies the uncertainty about the demand curve x_j retains the perceived kinks in expected profits at the unambiguous estimated relative prices. However, unlike in our benchmark model, in both cases the perceived kinks now lead the firm to change $p_{i,t}$ one-to-one in response to the observed p_t. Indeed, by equation (27) the perceived kink at the previous $\widetilde{r}_{i,0}$ implies a kink at the nominal price $p_{i,1} = p_{i,0} + p_1 - p_0$. As a result, indexation is now optimal. Hence, while ambiguity about the shape of demand generates real rigidity, it is its interaction with uncertainty about the link between aggregate and industry prices that turns it into a nominal rigidity.

Proposition B2. *Consider a counterfactual economy, where the firm knows that the unique co-integrating relationship is* $\phi(p_t - \widetilde{p}_{j,t}) = p_t - \widetilde{p}_{j,t}, \forall t$. *For a given realization of the current state* $s_1 = \{\omega_{i,1}, p_1, y_1, \widetilde{p}_{j,1}\}$, *the difference in worst-case expected profits* $\ln \upsilon^*(\varepsilon^0, s_1, p_{i,1}) - \ln \upsilon^*(\varepsilon^0, s_1, p_1 + r_{i,0})$, *up to a first-order approximation around* $p_1 + r_{i,0}$, *is*

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

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

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

$$\phi(p_t - \widetilde{p}_{j,t}) = p_t - \widetilde{p}_{j,t}, \ \forall t. \tag{39}$$

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

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

As a result, the estimated relative price simply equals

$$r_{i,t} = p_{i,t} - p_t. \tag{40}$$

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

$$\left(e^{p_{i,1}-p_1} - e^{mc_{i,1}}\right) e^{\widehat{x}_0^*(p_{i,1},y_1,p_1,\widetilde{p}_{j,1})}, \tag{41}$$

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

$$\min_{\delta' \in [-\delta,\delta]} \exp\left\{y_1 - b\left(p_{i,1} - p_1\right) - \gamma + \alpha\left[y_0 - \left(-\gamma - br_{i,0}\right)\right] - \alpha\delta'\left(p_{i,1} - p_1 - r_{i,0}\right)\right\} \tag{42}$$

The worst-case demand shape is therefore given by

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

Having described the worst-case expected profit, the proof follows from taking the derivatives of expected profit in (41) and payoff in (42) with respect to the action $p_{i,1}$. □

Different from the benchmark economy, we note that in this counterfactual the worst-case expected profit does not depend directly on the aggregate price. Indeed, the optimal choice of the relative price in equation (42) is independent of p_1. In this economy indexation is built in, as instructed per equation (40) where, holding constant the relative price, the nominal price moves one to one with p_1. Therefore, not surprisingly, a nominal price policy that deviates from indexation is suboptimal. To show this, consider a firm that lives in this counterfactual economy but does not index to the aggregate price. Instead, it targets the same $r_{i,0}$ but by setting $p_{i,1}^{noindex} = r_{i,0} + \widetilde{p}_{j,1}$. Put differently, this firm uses only the review signal as the source of relevant information for $p_{j,1}$ but targets the same relative price. Proposition A5 below details how the non-indexing policy is strictly suboptimal.

Proposition B3. *In the counterfactual economy, the difference* $\ln \upsilon^*(\varepsilon^0, s_1, p_1+r_{i,0}) - \ln \upsilon^*(\varepsilon^0, s_1, \widetilde{p}_{j,1}+r_{i,0})$, *up to a first-order approximation around* $\widetilde{p}_{j,1}$, *equals*

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

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

Proof. The firm that sets $p_{i,1}^{noindex}$ is subject to the same informational assumption as the firm that indexes, and, therefore, it still knows that the co-integrating relationship is given by (39). Compared to the indexing policy, this firm simply follows a different nominal pricing policy. The resulting worst-case expected profit is

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

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

plus

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

The worst-case demand shape is therefore simply

$$\delta^{noindex} = -\delta \,\mathrm{sgn}\left(p_1 - \widetilde{p}_{j,1} \right). \tag{43}$$

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

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

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

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

since when p_1 is larger (smaller) than $\widetilde{p}_{j,1}$, by the worst-case in (43) we have $\delta^{noindex} = -\delta$ or δ, respectively. Here we have used that the optimal $r_{i,1}$ sitting at the kink $r_{i,0}$ implies that

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

$\square$

B.4 Simulated hazards

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

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

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

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

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

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

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

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

Next, we introduce heterogeneity in the unconditional price change frequencies ξ_i. We do, however, keep a homogenous, flat slope of the hazard function. Our simulations confirm the presence of the survivor-bias issue discussed in the literature: without fixed effects, the estimation finds a hazard that is declining on average (column B), even if our simulation features no relationship between spell length and price change frequency. This is also true if we use a distribution of the price change probabilities ξ_i that mimics the empirical distribution from our dataset (column C). Here we found that a χ^2 distribution with 5 degrees of freedom, scaled to match the mean frequency found in our dataset, provides a good fit. The inclusion of product fixed effects, on the other hand, correctly leads us to conclude that the hazards are flat on average: controlling for product-specific hazard shifters circumvents the downward bias that arises from heterogeneous price rigidity.

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

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

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

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

absolute value) than it actually is, at -0.0091.

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

B.5 Additional evidence on hazard functions

In Figure B.1 the distributions of the estimated hazard slopes across the 54 category/market combinations. These estimates are obtained from our linear probability regression model with fixed effects of equation (32). The left panel shows the slope estimates from unweighted regressions, while results from weighted regressions are shown in the right panel.

B.6 Speed of learning

The evolution of the pricing policy function over time

To further illustrate how learning and the resulting pricing policy evolve over time, panel a) of Figure B.2 shows how the policy function of one the longer-lived firms in the simulation changes from period one-hundred and fifty, to the three hundredth period of this firm's life. The blue line corresponds to the optimal policy at $t = 150$, and shows that by that period the firm had sampled a number of different prices, and established a fair number of kinks. While we might think that establishing such "special prices" happens once and for all, in fact the position of the kinks can move and they could even completely disappear as new information arrives. We can see that from the red line, which plots the policy at $t = 300$, and shows that by that period the two lowest flat spots in the policy became absorbed in a new, single flat spot at an intermediate price point.

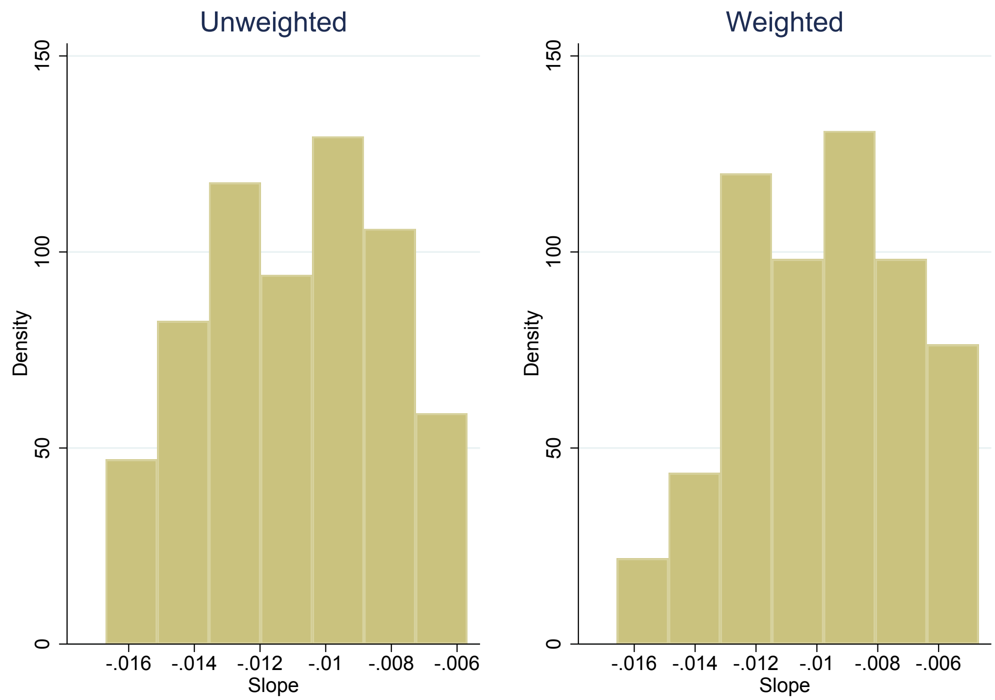

Figure B.1. Distribution of the slopes of hazard functions across 54 category/market pairs. Unweighted and weighted regressions.

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

A counterfactual economy with no firm exit

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

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

As a result, even when firms are infinitely-lived and accumulate thousands observations about demand, the behavior of prices remains qualitatively similar to that in the benchmark model, with prices displaying both stickiness and memory. To understand this pricing behavior, we use our procedure to compute the typical optimal policy function (in terms of the estimated relative price $\tilde{r}_{it}$) from this simulation, with results plotted in panel b) of Figure B.2. As can be seen from the Figure, the policy function is qualitatively similar to that in the benchmark case, and is essentially a step function across the whole support of the price space. Again, this is because even though the firms have seen much longer histories of observations, they have concentrated their pricing, and thus information accumulation, in the set of previously observed estimated relative prices. This results in a pricing policy that is a step-function, generating both stickiness and memory in prices.

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

B.7 Comparative statics

We now turn to comparative statics. A common theme throughout is the nuanced link between price flexibility and memory, which as we have shown in Section 4.3, is an important determinant of how micro-data stickiness maps into the effects of monetary policy.

As a first comparison, we consider a myopic firm by setting $\beta = 0$, which eliminates all experimentation incentives. The key pricing moments under this parameterization are reported in Table B.2, where we see a drop in both the frequency and median size of price changes. This is because without a reason to explore new parts of the demand curve, firms now have less incentives to change prices often or by large amounts. Further investigation shows that this leads firms to

Figure B.2. Optimal Pricing Policy Function

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

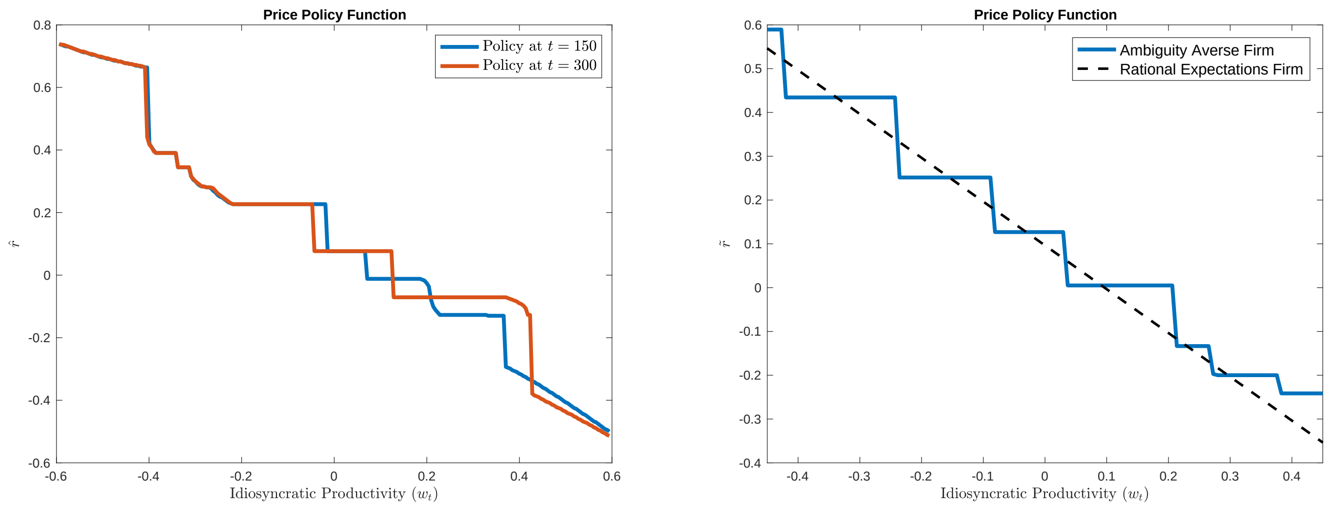

Table B.2. Moments - Comparative statics

	Benchmark	$\beta = 0$	$\psi = 0$	Low δ	High σ_ω	High b
Freq. regular prices changes	0.105	0.075	0.064	0.207	0.160	0.199
Median size of abs. changes	0.154	0.007	0.015	0.108	0.123	0.015
Freq. modal price changes	0.026	0.028	0.029	0.041	0.037	0.056
Prob. visiting old price	0.414	0.237	0.469	0.444	0.502	0.488
Real effect of s_t shock, (cumul. 52w)	7.22%	16%	21.9%	3.49%	5.52%	5.21%

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

concentrate their information accumulation in the middle range of productivity shocks, leading to an ergodic policy function with two large flat spots in the middle, but no other kinks. As a result, the frequency of modal price changes rises slightly, but memory falls significantly because there are no other attractive prices outside of those two. Moreover, in unreported results we find that this myopic version generates very few large price changes and does not match the product pricing life-cycle facts, as young firms no longer have an experimentation motive to change prices more often. Lastly, this version of the model implies a significantly stronger monetary non-neutrality, with a cumulative real output effect in the 52 weeks following a nominal shock rising to 16%. This is a combination of the fact that prices are less flexible overall, and that the lack of experimentation incentives also means that price change motives are more closely aligned with aggregate nominal shocks.

Next, we consider setting $\psi = 0$ to eliminate the local nature of learning. In that case, each signal carries the same quantity of information for any other price point, irrespective of its distance from the current price. This setting also kills the experimentation motive (Proposition A.1 in the Online Appendix A.3), because the new information contained in a signal is not specific to the position of the price at which the signal was observed. It is thus not surprising that the resulting moments are mostly similar to the ones with $\beta = 0$, as can be seen in Table B.2. The main difference is memory, which increases to 47%. This is due to the emergent ergodic policy function, which we find that now features numerous, smaller kinks as opposed to just two large ones, increasing the probability of switching between kinks. The intuition can be seen from section 2.3, which shows that when $\psi = 0$ the perceived demand loss of moving away from a kink to a new price is relatively steeper for larger price changes as compared to smaller adjustments. As a result, smaller price changes are perceived as relatively safer, leading the firm to establish several kinks in the same neighborhood, as opposed to just a single one. Lastly, the real output effect of this model is even bigger than in the case of $\beta = 0$, due to the lower frequency of price changes and higher memory as compared to the myopic version.

As a third comparison, we decrease the degree of ambiguity by halving δ. By Proposition 1, this lowers the as-if cost of moving away from the previously posted price. As a result, price changes occur more often (both regular and modal), and the size of the resulting price changes is smaller. Interestingly, this increased flexibility implies more kinks and hence more (but smaller) flat spots in the pricing policy. The result is higher memory, as there is a higher number of attractive prices that were set previously. Overall, the greatly increased price flexibility leads to a significantly smaller cumulative real output effect of 3.49%.

Fourth, we consider a version of the model with high costs volatility, and double the standard deviation of the idiosyncratic productivity shocks, σ_ω. This raises the frequency of modal and posted price changes, an intuitive result that is shared with a number of other standard frameworks (see Klenow and Willis (2016) for a discussion on the role of shocks' distribution in standard price-setting models). In our model, however, the increased price flexibility is also accompanied by higher memory. The reason is that with more frequent price changes, information accumulation is spread out over a larger set of individual prices, resulting in a policy function with more steps and thus increased memory. Hence, even though prices change more frequently, they are also more likely to revert to past price levels. The combination of increased price flexibility and memory nets out to a lower overall real output effect as compared to the benchmark model, but the fall in the real effect is smaller as compared to the case of lower δ, because of the counterbalancing increase in memory. The lower real output effect is consistent with the empirical evidence in Boivin et al. (2009) who find that monetary non-neutrality indeed decreases with idiosyncratic volatility.

Finally, we increase the average price elasticity of demand by doubling the value of b. The resulting higher sensitivity to deviations from the optimal markup, which is now just 9%, leads to a significantly higher frequency and a smaller absolute size of price changes, as documented in the last column of Table B.2. These results are consistent with Mongey (2018) who reports that products facing more competition are characterized by a larger frequency of posted prices and smaller absolute price changes. We find that, as in the δ and σ_w comparative statics, the increased flexibility comes with higher memory, from having more steps in the policy function. This positive correlation of frequency and memory is not mechanical, as shown by the $\psi = 0$ case where the two moments move in the opposite direction. Overall, the monetary non-neutrality in this version of the model is also weaker, with a cumulative real output effect of 5.21%, again owing to higher price flexibility balanced out with higher memory. This predictions of a smaller real output effect is consistent with Kaufmann and Lein (2013) who empirically find that monetary non-neutrality decreases with competition.