

FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

**A Dynamic Model of Price Signaling, Consumer
Learning, and Price Adjustment**

Matthew Osborne
University of Toronto

Adam Hale Shapiro
Federal Reserve Bank of San Francisco

November 2014

Working Paper 2014-27

<http://www.frbsf.org/economic-research/publications/working-papers/wp2014-27.pdf>

The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of San Francisco or the Board of Governors of the Federal Reserve System.

A Dynamic Model of Price Signaling, Consumer Learning, and Price Adjustment*

Matthew Osborne[†] and Adam Hale Shapiro[‡]

November 22, 2014

Abstract

We examine a model of consumer learning and price signaling where price and quality are optimally chosen by a monopolist. Through numerical solution and simulation of the model we find that price signaling causes the firm to raise its prices, lower its quality, and dampen the degree to which it passes on cost shocks to price. We identify two mechanisms through which signaling affects pass-through. The first is static: holding quality fixed, price signaling increases the curvature of demand relative to the case where quality is known, which ultimately acts to dampen how prices respond to changes in cost. The second is dynamic: a firm that engages in signaling recognizes that changing prices today affects consumer beliefs about the relationship between prices and quality in the future. We also find that signaling can lead to asymmetric pass-through. If the cost of adjusting quality is sufficiently high, then cost increases pass through to a greater extent than cost decreases.

*We thank Heski Bar-Isaac, Michael Grubb and Tarcisio Da Graca for feedback on earlier drafts. We also thank audiences at the International Industrial Organization Conference 2012, Econometric Society 2014, EARIE 2014, and the Bureau of Economic Analysis for comments. The views expressed in this paper are those of the authors and not of the Federal Reserve Bank of San Francisco or the Federal Reserve System.

[†]University of Toronto, Institute for Management and Innovation and Rotman School of Management. email: matthew.osborne@rotman.utoronto.ca

[‡]Federal Reserve Bank of San Francisco. email: adam.shapiro@sf.frb.org

1 Introduction

We develop a model of consumer learning and price signaling where price and quality are optimally chosen by a monopolist. Price signaling may occur when consumers have imperfect information about product quality. To infer quality, consumers may rely on previous experience or may use some of the product's observable characteristics, such as the product's *price*. We examine the scenario whereby the firm can endogenously change consumers' beliefs about the product's quality by altering both the price and quality of its product. Our main findings are that, in this type of setting, price signaling causes the firm to raise its price, lower its quality, and dampen the degree to which it responds to cost shocks. If the cost of adjusting quality is sufficiently high, the dampening effect is pronounced in the downward direction, meaning that price signaling causes prices to respond less to cost decreases than cost increases.

In our model, consumers have beliefs about the joint distribution of price and quality for the firm's product. Given no prior information on quality and price, consumer beliefs about a firm's quality may or may not be correct. Before making a purchase, consumers do not observe the quality of the product but make an inference about it using their beliefs. After purchase, the product's quality is revealed and consumers update their beliefs through a quasi-Bayesian learning process. Thus, over time consumers will observe realized price and quality, and will update their beliefs about a firm's policy given their past experience. In each period, a forward-looking monopolist chooses the product's price and quality, accounting for the fact that its choice will affect consumer beliefs and future profits. We solve the model using numerical methods and then simulate it under a range of parameters.

To our knowledge, this study is the first to assess the effects of price signaling on price adjustment with endogenous quality and a forward-looking monopolist.¹ Previous studies

¹Concurrent work by L'Huillier (2013) shows in a macroeconomic model with exogenous monetary shocks and consumer learning that signaling can lead to price rigidity. In that model, the firm may have an incentive to hide whether its product is cheap or expensive relative to the numeraire good by keeping its price rigid. Two important differences between the two models are that in L'Huillier (2013) firms are not forward-looking and the firm does not choose its quality. Our work accounts for the impact of reputation on quality provision and price rigidity.

on price signaling, such as Judd and Riordan (1994) and Wolinsky (1983) have focused on examining static equilibrium price levels and have assumed either a one- or two-period setting. Our model extends this line of research by assessing the effects of cost changes on price adjustment, but it also introduces a dynamic framework with consumer learning. We show that consumer learning may be important if a firm's optimal strategy includes endogenously manipulating consumers' expectations about the relationship between price and quality. In particular, we find that this type of forward-looking behavior further moderates the effect of cost changes on optimal prices. Due to the complexity of our model, analytical solutions to the firm's policy functions are not available. Thus, we turn to simulation of a calibrated model to draw our conclusions.

Relative to a full information model, where a bit more than 50% of either a cost increase or a cost decrease is passed on to prices, we find a signaling model implies that only 43% of a cost increase and only 42% of a cost decrease is passed on. We solve our model at parameter values that are calibrated to match aggregate margins by U.S. firms. However, our finding that price signaling results in dampened pass-through is robust to perturbations of these parameters. We also find evidence of asymmetric pass-through. If the cost of adjusting quality is sufficiently high, a firm will pass on more of a cost increase than a decrease, which is consistent with empirical findings (Peltzman 2000).

The results of our model have important insights regarding the literature on price stickiness and cost pass-through, which covers a number of fields including marketing, macroeconomics, and international economics. There are a number of explanations for nominal price rigidity. Rotemberg (2005) proposes that price changes may be muted because they may anger consumers. Cabral and Fishman (2012) put forth a model showing that price changes may induce consumers to search for competitors' products. A few studies have shown that, more generally, prices will be more rigid the more curved the demand function (see for example, Kimball (1995) and Klenow and Willis (2006)). Indeed, in a recent study on cost pass-through in the coffee market, Nakamura and Zerom (2010) find that this type of demand curvature is responsible for a large portion of wholesale coffee price rigidity. Understanding what economic forces can drive price rigidity is important from a policy perspective because price rigidities provide scope for monetary policy. In macroeconomics, a large body of research is devoted to exploring

informational drivers of price stickiness (see Lucas (1972), Mankiw and Reis (2002), Reis (2006), Mackowiak and Wiederholt (2009), and L’Huillier (2013)).

Building on prior work by Weyl and Fabinger (2013), we mathematically derive two drivers of the impact of signaling on cost pass-through, which we denote as the static and dynamic effects. The static mechanism is in line with studies by Kimball (1995), Klenow and Willis (2006), and Nakamura and Zerom (2010). Specifically, we show that if the steady-state ratio of quality to price is not too high, signaling increases the curvature of the demand function, which ultimately acts to dampen how prices respond to changes in cost. The demand function becomes more curved because, under price signaling, raising the price may signal to the consumer that the product is of higher quality. Thus, demand does not necessarily monotonically fall with prices. This finding may help explain the relatively large degree of demand curvature needed to explain incomplete pass-through by Nakamura and Zerom (2010).

The dynamic effect arises because a forward-looking firm recognizes that if it changes the price of a product today, it affects consumer beliefs about the relationship between price and quality in the future. Adjusting prices has two effects on future demand. First, adjusting prices alters consumer beliefs about the relationship between price and quality, changing the effectiveness of the signal. Second, adjusting prices changes the precision of the signal, "locking in" their beliefs in the future. If the firm’s value function is sufficiently concave in price then signaling will lead to dampened pass-through. Our simulation results suggest that the dampened pass-through rates from signaling are due to both the static and dynamic effects. The static effect is much larger: signaling overall dampens pass-through by about 8 or 9%, and about 7 to 8% of that is due to the impact of signaling on the shape of the demand curve. To examine the generality of our finding that signaling moderates pass-through, we derive conditions on the ratio of price to quality under which signaling increases demand curvature. As long as this ratio is above the consumer’s willingness to pay for the product, and below an upper bound that increases in the share of income spent on the product, signaling will increase the demand curvature. Our findings suggest that the dampening effect of signaling on pass-through is more likely to be prevalent in more expensive products.

It is important to note that our modeling framework differs from prior theoretical

work on signaling such as Wolinsky (1983). This literature has assumed that consumers know the functional form of demand and costs, and can therefore directly compute the firm’s optimal pricing strategy. While this approach is tractable in the static settings typically examined in the signaling literature, it is intractable in a dynamic setting such as ours. Therefore, a strong assumption of our model is that consumers do not base their beliefs on the optimal solution of the firm’s problem. Similar issues have arisen in the estimation of dynamic games.² In our context, consumers are following a behavioral rule that is consistent with the notion of equilibrium behavior proposed by Fershtman and Pakes (2012). A consequence of our assumption is that, in our model, consumers will never be able to learn the optimal policy of the firm exactly as they would in more standard models. However, due to the Bayesian updating process they will, on average, be able to correctly predict the firm’s quality conditional on price.

In addition to contributing to the literature on price stickiness, our paper also contributes to the literature that models a firm’s reputation for quality. Klein and Leffler (1981) and Shapiro (1983) assess the conditions under which firms do not “cheat” consumers by offering a quality that is less than contracted for. Both studies consider a forward-looking firm and find that prices are set at a premium above cost in order to compensate for the firm’s initial investment in quality.³ Similar to our study, these studies consider forward-looking firms; however, they differ from ours in that they do not explicitly model the beliefs of consumers. These authors assume that once a firm cheats (by offering a low quality product at a high price) all consumers cease to purchase from the firm—a type of trigger strategy. More recent studies on reputation such as Board and Meyer-ter Vehn (2013) model quality as a function of past investments in quality, implicitly allowing the firm to control consumers’ beliefs through persistent investment decisions. In our model, a firm has an incentive to cheat consumers because they do not immediately observe a product’s quality and must learn about it over time. Consumers in our model learn according to a modified Bayesian learning process where consumer forgetting can occur: observations that occurred far in the past may be downweighted

²Recent work, for example Fershtman and Pakes (2012), has proposed that, rather than being fully rational, some agents are assumed to use past outcomes to forecast current outcomes.

³Specifically, a large discounted future flow of profits can then be a better strategy than pocketing the short-term gain of reducing quality and losing future customers.

relative to more recent data (see, for example, Orphanides and Williams (2004)).⁴ We find that a firm has more incentive to cheat consumers if they forget the past more easily; if the firm produces a poor quality product beliefs will update quickly. In contrast, if consumer beliefs update slowly then there is scope for a firm to offer a high quality product for a while to raise consumer perceptions, and then to cheat consumers for some time by producing a low quality product at a high price.

An outline of the paper is as follows. In Section 2, we introduce our model of price signaling and consumer learning. Section 3 describes the techniques used to numerically solve the model, and the steady state of our model at our preferred parameterization. Section 4 shows the simulated pass-through from our preferred parameterization, the mathematical derivations of pass-through, and the breakdown of pass-through into static and dynamic components. Section 6 shows the robustness of our results to alternative parameterizations, and Section 7 concludes.

2 Model

2.1 Demand

We assume there exists a continuous measure of infinitely lived consumers indexed by i , each endowed with wealth $w_i \sim U[0, \bar{w}]$. The total measure of consumers is normalized to 1. Under an additively linear direct utility function, it follows that consumer i 's indirect utility upon purchase of a product with quality x_t and price p_t in time period t is:

$$U_i = x_t - \alpha p_t + \epsilon_{it}. \tag{1}$$

The error term ϵ_{it} is an i.i.d. taste shock that reflects consumer heterogeneity in tastes for the firm's product. ϵ_{it} is observed to the consumer at time t but unobserved to the firm. We assume that the ϵ_{it} follows a type 1 Extreme value distribution. α measures the consumer's marginal utility of income.⁵ If a consumer does not purchase the product,

⁴In standard Bayesian updating, all observations receive equal weight.

⁵In this model we implicitly assume that the consumer has a utility function that is quasilinear in two goods, the monopolist's product and a numeraire good z . If the consumer purchases the monopolist's

she consumes the outside good. We normalize the price and quality of the outside good to zero.

The case we explore in this study consists of the scenario where the consumer does not know the true quality of the product, x_t , and thus must infer the quality based on her information set \mathfrak{S}_t (that is, her beliefs).⁶ Consumer i purchases the product if her expected utility for the product is positive, and if the price is lower than her wealth w_i . Under these assumptions, the aggregate demand for the good will be

$$D(p_t, \mathfrak{S}_t) = \begin{cases} \left(1 - \frac{p_t}{\bar{w}}\right) \left(\frac{\exp(E(x_t|\mathfrak{S}_t) - \alpha p_t)}{1 + \exp(E(x_t|\mathfrak{S}_t) - \alpha p_t)}\right) & \text{if } p_t \leq \bar{w} \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

We assume that the consumers believe the expected quality conditional on the observed price is linear in the price:

$$E(x_{it}|p_{it}) = b_0 + b_1 p. \quad (3)$$

Because in our model we interpret the outside good as a product with quality 0 and price 0, in order for consumer expectations to be consistent we normalize $b_0 = 0$. We additionally assume that consumers do not know b_1 , but learn about it over time. At the beginning of period t , consumers observe only the price and condition their beliefs about product quality on price and on their beliefs about b_1 . We denote the consumer's expected value of b_1 as β , and the variance of their belief around β as σ^2 . At the end of the period, the quality of the firm's product is revealed to the entire market. The information revelation could occur through magazine ratings such as Consumer Reports, ratings on Internet websites such as Amazon, or word of mouth.⁷ We assume that consumers form their beliefs in period t using a weighted regression:

$$\beta_t = \min_{\beta} \sum_{\tau=1}^t \lambda^{(t-\tau)} (x_{\tau} - \beta p_{\tau})^2. \quad (4)$$

good then her consumption of the numeraire good is $z = w_i - p$.

⁶For reasons we will clarify below, in our model the information set will be the same for all consumers.

⁷Allowing for consumers to differ in their beliefs would complicate the modeling considerably. The firm would have to track an infinite dimensional state variable, the distribution of beliefs, over time.

The term $\lambda \leq 1$ is an *exponential forgetting factor* that downweights earlier observations. If $\lambda = 1$, then consumers have perfect recall; $\lambda < 1$ assumes that consumers eventually forget early observations. Note that if $\lambda = 1$, then consumer uncertainty about the value of b_1 will approach zero. This can lead to the somewhat unrealistic scenario where, if a firm were to choose a high quality and a high price for a long period of time, consumer beliefs would not update much if the firm suddenly dropped its quality to zero. If $\lambda < 1$, consumers will always remain somewhat uncertain about the relationship between price and quality and a large change in x will have an impact on consumer beliefs, even in the limit ($t = \infty$).

When we solve the problem, it is convenient to write consumer belief updating in recursive form as a Kalman filter following Bersetkas and Tsitsiklas (1996). If the variance of beliefs at time t is σ_{t-1}^2 then the updating formulas for β_t and σ_t^2 can be written as:

$$\begin{aligned}\sigma_t^2 &= \frac{1}{\frac{\lambda}{\sigma_{t-1}^2} + p_t^2} \\ \beta_t &= \beta_{t-1} + \sigma_t^2 p_t (x_t - \beta_{t-1} p_t).\end{aligned}\tag{5}$$

Note that if $\lambda = 1$, it can be shown that $s^2 \rightarrow 0$ as $t \rightarrow \infty$.

The effect of changing prices on consumer beliefs leads to the dynamic effects of signaling which we discuss in the introduction. Consumer beliefs will update so that, on average, they are correct. For a given quality level of the firm, a price increase will tend to lower the parameter β_t , weakening consumer beliefs about the relationship between price and quality. Additionally, a higher price provides a stronger signal to consumers, and will lead to a lower value of σ_{t-1}^2 , as shown in equation (6). To the firm, the advantage of having a lower value of σ_{t-1}^2 is that consumers will be more locked in to their beliefs in the future, and their beliefs will update more slowly to future price and quality changes.

Before proceeding we make some comments on our assumptions about consumer beliefs and the process by which beliefs update. First, we note that we assume consumers believe the expected value of quality is linear in price. This assumption is attractive due to its tractability and has been used in learning papers in macroeconomics (Wieland 2000). One way to interpret this assumption is that, although the actual function $E(x|p)$ may be nonlinear in price, consumers are boundedly rational and have a limited ability

to process information. One could approximate $E(x|p)$ with any degree of accuracy using a Taylor expansion; the higher is the order of this Taylor expansion, the more parameters one has to track. Our assumption would be justified by the assumption that consumers use a linear approximation. To check the robustness of our findings to this assumption, in Appendix 11 we consider an alternative model where prices and quality take on discrete values, and consumer beliefs are not constrained to be linear. Our overall findings appear to be robust to the specification of beliefs. We prefer a specification of beliefs that allows for continuous prices, because observed prices are continuous and pass-through derivations are more tractable in a continuous setting.

A second assumption underlying consumer behavior is that consumers do not solve the firm’s dynamic problem and base their beliefs on the optimal solution of the firm problem. Prior theoretical work on signaling such as Wolinsky (1983), has assumed that consumers know the functional form of demand and costs, and can compute a firm’s optimal pricing strategy. Consumers make purchase decisions knowing this optimal strategy and taking it into account. This approach is tractable in the static settings that are typically examined in the signaling literature, but would not be tractable in a setting such as ours. Similar issues have arisen in the estimation of dynamic games, and recent work in that area has proposed that rather than modeling all agents as being fully rational, one can assume that some agents use outcomes experienced in past periods to form forecasts of outcomes in current periods (Fershtman and Pakes 2012). In our context, consumers are following a behavioral rule that is consistent with the notion of equilibrium behavior proposed by Fershtman and Pakes (2012).

2.2 Firm

A monopolist maximizes profits in a dynamic sense, in that it chooses price and quality, keeping in mind how this joint decision alters consumers’ expectation of quality. We assume that the firm’s marginal cost, \mathbf{c}_t , is a stochastic vector (c_{0t}, c_{1t}) . Each draw of \mathbf{c}_t is i.i.d. across time with distribution $F_c(\cdot)$. The stochastic nature of cost can include a multitude of exogenous factors such as process innovations, weather disruptions, or factory malfunctions. Firm profits will be a function of price, quality, cost draws, and consumer beliefs at time t , which we denote as $\mathbf{S}_t = (\beta_t, \sigma_t^2)$:

$$\pi(p_t, x_t, \mathbf{S}_t, x_{t-1}, \mathbf{c}_t) = D(p_t, x_t, \mathbf{S}_t)[p_t - c_{0t} - c_{1t}x_t^\nu] - \gamma(x_t - x_{t-1})^2.$$

The cost $c_{0t} \geq 0$ reflects the marginal cost of producing a quality 0 product, while $c_{1t} \geq 0$ allows for higher quality products to have higher production costs. $\nu \geq 0$ determines the rate at which the marginal cost of producing higher quality increases with quality. γ is an adjustment cost, which incorporates the idea that quality is not as easy to change as price: for example, creating a higher or lower quality product may involve retooling production plants. D represents aggregate demand, as specified in equation (26).

We assume that the monopolist is forward-looking, which means that it will choose prices and quality every period to maximize the expected present discounted sum of profits⁸,

$$\max_{\{p_{t+i}, x_{t+i}\}_{i=0}^{\infty}} E \left[\sum_{i=0}^{\infty} \delta^i \pi(p_{t+i}, x_{t+i}, \mathbf{S}_{t+i}, x_{t+i-1}, \mathbf{c}_{t+i}) | \mathbf{S}_t \right], \quad (6)$$

subject to the evolution of the state variable \mathbf{S}_t , which we write in vector form below as a mapping from $L : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ as follows:

$$\mathbf{S}_{t+1} = L(x_t, p_t, \mathbf{S}_t) = \left(\beta_{t-1} + \frac{p_t(x_t - \beta_{t-1}p_t)}{\frac{\lambda}{\sigma_{t-1}^2} + p_t^2}, \frac{1}{\frac{\lambda}{\sigma_{t-1}^2} + p_t^2} \right). \quad (7)$$

The firm's problem in (6) can be written in Bellman equation form as follows:

$$V(\mathbf{S}, x_-) = \int \max_{p, x} \{ \pi(p, x, \mathbf{S}, x_-, \mathbf{c}) + \delta V(L(p, x, \mathbf{S}), x) \} dF_{\mathbf{c}}(\mathbf{c}). \quad (8)$$

In each period t , the firm will optimally choose p and x , accounting for how those choices will affect consumer beliefs the subsequent period. The firm faces a trade-off in the sense that raising p and lowering x will raise current-period profits, but may also diminish future profits. A higher price will signal higher quality to consumers, and a lower x will lower current production costs, thus increasing current profits. However, when

⁸The expectation is taken over the cost draws.

consumers learn that quality is low they will lower their belief about the effectiveness of price signaling. Note that the impact of p and x on future beliefs will be greater the smaller are λ and σ_{t-1}^2 . If λ is close to 1 and σ_{t-1}^2 is close to zero, consumers will be relatively certain of the relationship between price and quality, and will update their beliefs slowly in response to new information. The firm will have an incentive to exploit the price signal, until consumer beliefs about β_t adjust to a sufficiently low value. A low value of λ keeps the firm in line in the sense that consumers will never be certain about their belief about β_t , so if the firm tries to exploit consumers today it will face a large change to its reputation tomorrow. Existence of a fixed point for the Bellman equation rests on the fact that the per-period profit function is continuous and bounded in \mathbf{S} (see Rust (1996)). The optimal choice of price and quality will come from the solution to the Bellman equation. Policy and value functions can be obtained from iterating the Bellman equation at an initial guess. We describe our algorithm for solving for the value and policy functions in more detail in Section 11.3.

3 Solving the Model and Numerical Methods

We solve for the firm's optimal policy and value functions using value function iteration combined with policy function iteration (Judd 1998). Because our state variables are continuous, we solve for the value function on a grid of points and interpolate the value function everywhere else using simplicial interpolation (Weiser and Zarantonello 1988). Simplicial interpolation is a generalization of linear interpolation to multiple dimensions. We describe the simplicial interpolation algorithm in Appendix 8.

The details of the algorithm are as follows. Our problem has two continuous state variables, which are β_t and σ_t^2 , the mean and precision of consumer beliefs about the relationship between price and quality. First, we split the continuous state space into a finite number of grid points, and interpolate the value function in between points. We choose a regular 51 by 51 grid for interpolation. We solve for the value and policy functions at each (β_t, σ_t^2) point on the grid; we index each such state point from $i = 1, \dots, N_s = 2601$, and denote each (β_t, σ_t^2) combination as \mathbf{s}_i . The bounds on the grid for β_t are 0 and 20, and the bounds on the grid for σ_t^2 are 0 and 16. Our simulated values

of β_t and σ_t^2 lie well within these bounds.

The value and policy function iteration then proceeds as follows. Indexing each iteration with n , at $n = 1$ we begin with a guess that the value function V_n is 0 at all points \mathbf{s}_i . We then solve for the optimal policy function $(x_n(\mathbf{s}_i, \mathbf{c}), p_n(\mathbf{s}_i, \mathbf{c}))$ at each state space point, \mathbf{s}_i , and possible cost draw \mathbf{c} :

$$(x_n(\mathbf{s}_i, \mathbf{c}), p_n(\mathbf{s}_i, \mathbf{c})) = \arg \max_{x,p} \{\pi(p, x, \mathbf{s}_i, \mathbf{c}) + \delta V_n(L(p, x, \mathbf{s}_i))\}. \quad (9)$$

We then solve for the value function that would be obtained if these policies were fixed. Operationally, we iterate on the value function contraction mapping with fixed policy functions until convergence. Denoting a policy iteration by n_p , we start by setting $V_1^p = V_n$ and update using the equation

$$V_{n_p+1}^p(\mathbf{s}_i) = E_{\mathbf{c}} \left[\pi(p_{n_p}(\mathbf{s}_i, \mathbf{c}), x_{n_p}(\mathbf{s}_i, \mathbf{c}), \mathbf{s}_i, \mathbf{c}) + \delta V_{n_p}^p(L(p_{n_p}(\mathbf{s}_i, \mathbf{c}), x_{n_p}(\mathbf{s}_i, \mathbf{c}), \mathbf{s}_i)) \right]. \quad (10)$$

We assume that the policy step has converged when $\max_{i=1, N_s^j} \|V_{n_p+1}^p(\mathbf{s}_i^j) - V_{n_p}^p(\mathbf{s}_i^j)\| < \epsilon_p$, where we choose $\epsilon_p = 1e - 4$. The policy iteration step in equation (31) converges very quickly. Once the policy step converges, we set $V_n = V^p$, and solve for the optimal policies at step $n + 1$ using equation (30). The algorithm converges when $\max_{i=1, N_s} \|V_{n+1}(\mathbf{s}_i) - V_n(\mathbf{s}_i)\| < \epsilon_v$, where we set $\epsilon_v = 1e - 3$. We have found that this algorithm converges much more quickly than standard value function iteration, where one would solve for the optimal policy every time one updated the value function.

3.1 Model Simulation

We simulate equilibrium prices and qualities for two versions of this model, one with zero adjustment cost and one with infinite adjustment cost. These two cases allow us to infer how price signaling affects the firm's pricing decisions in different types of situations. An infinite adjustment cost model might approximate the case in which a firm initially decides whether it wants to sell in the luxury or the economy market. A zero adjustment cost model, on the other hand, might represent a barber who can adjust quality by simply exerting different amounts of effort. The case with an infinite adjustment cost is simpler

in that the firm never adjusts the quality of its product, only its price. Therefore, once quality is initially chosen it can be considered exogenous thereafter. In the case, of zero adjustment costs, the firm is free to adjust both its price and quality without any cost. That is, quality is considered endogenous.

Our preferred parameter values for numerical solution of the model are shown in Table 1. Maximum income is normalized to 1, and the firm’s discount factor is chosen to be consistent with a yearly interest rate of 5%. The cost parameter values are chosen to produce prices such that marginal costs match aggregate labor share of income in the U.S. economy, which is 60% (this approach to calibrating markups is often used in papers in macroeconomics, for example Gali and Gertler (1999)). In this sense our parameter values are realistic in that they imply average markups that are consistent with what is observed for an average firm. The parameter on the convexity of marginal costs in quality is chosen to be 2. To get a sense of the robustness of our conclusions to changes in the parameters, we will also solve the model for cost distributions that are lower and higher than those in the table, and for values of λ that are lower and higher than 0.99. We discuss these perturbations to the model in Section 6.

The solution to the firm’s problem involves computing an optimal price p and an optimal quality x . To simplify the solution of this problem we discretize the x values and assume that the firm can choose x in increments of 0.1 between 0.5 and 10. The distribution of marginal costs is assumed to be uniform, and we integrate over costs using Gauss-Legendre quadrature with five quadrature points.

3.1.1 Fixed Quality: Infinite Adjustment Costs

The policy function under the infinite adjustment cost case is depicted in Figure 1. We graph the policy function conditional on the quality $x = 3.6$; as we show below, this is the level of quality that maximizes profits if consumers have full information and know the quality of the product with certainty. The figure shows that at very low values of β , price is increasing in β . At higher values of β , the optimal price becomes relatively flat with respect to β .

The non-monotonicity of the pricing policy function in β seems counterintuitive but is attributable to the fact that increasing β has two effects on demand. First, for a given

Table 1: Parameter Values for Base Simulation

Variable	Value
δ	0.95
\bar{w}	1
α	1
c_0	U[0.35, 0.45]
c_1	0.002
ν	2
λ	0.99

price, a larger β signals that the product is of higher quality, so demand is higher. All else equal, this “demand effect” should induce the firm to raise its price. Second, a larger β means that a change in price will have a larger effect on the change in demand, which acts to increase the price elasticity of demand. All else equal, this “elasticity effect” will induce the firm to lower its price. When β is low, an increase in β has a relatively larger demand effect than an elasticity effect, so the price increases in β . However, at higher levels of β the elasticity effect dominates slightly, which is why the pricing policy function eventually starts decreasing in β . In Appendix 9, Figure 3, we show how the optimal price for a myopic firm is affected by changes in β .

We assume the firm chooses a level of quality in period 0 to maximize the expected present discounted value of profits, and cannot change it afterwards. The optimal level of x depends on consumer beliefs. In the steady state of this model, beliefs and prices will always adjust to the same values conditional on x , so the effect of the starting points will die out. We simulated firm decisions at starting values of $\beta = 3.5$ and $\sigma^2 = 1$, which produced an optimal x of 3.6, which is the same as the optimal x in the full information (i.e. a model with no signaling) case.⁹

⁹In our simulations we have found that the optimal level of x is decreasing in the initial level of β , and is also decreasing in σ^2 . The higher is β and the lower is σ^2 , the lower will be the optimal x . Producing a higher quality product is costly for the firm in the short run; however, the benefits to being




Figure 1: Pricing policy function for $x_{t-1} = 3.6$, $\gamma = \infty$

Table 2 shows the steady-state values of price, quality, firm profit, and consumer surplus.¹⁰ To assess how price signaling impacts price and quality over and above the standard textbook case, the table also shows the case where the consumer knows the quality of the product with certainty (i.e., the case of full information). The general takeaway from the table is that information asymmetry (i.e., price signaling) acts to raise the equilibrium price relative to full information and also lowers consumer surplus.

3.1.2 Flexible Quality: Zero Adjustment Costs

The policy functions under the zero adjustment cost case are depicted in Figure 2. Here, there are two policy functions (one for x and one for p) since the choice of x is high quality are only realized in the long run as consumers adjust their beliefs about the relationship between price and quality. If β is initially higher, the firm has more incentive to produce a lower quality product since it will take longer for consumers to figure out the firm is bad; similar results arise if σ^2 is lower.

¹⁰The values of outcome variables are the averages over the simulation, starting with the 500th iteration and ending in the 1000th. The first 500 iterations are removed for burn-in. For a graphical illustration of the steady state values see Appendix 10.

endogenous. The shape of the pricing policy looks similar to the $\gamma = \infty$ case, with some subtle differences. This has to do with how the choice of x impacts the choice of p . First, note that the optimal choice of quality is generally decreasing in β . The intuition here is that as β becomes smaller, the consumer is less perceptive that a higher price signals higher quality. Thus, when β decreases, the firm attempts to raise the price signal by choosing to offer a high quality product at a high price. As β becomes larger this effect is weaker, hence the optimal price now falls with β when β is large. It is interesting to note how the optimal choice of x is impacted by σ^2 . The optimal choice of x is generally increasing in σ^2 . The intuition here is that, the smaller is σ^2 , the more apt the firm is to exploit the signal by cheating. For instance, if the consumer is very confident that a high price signals a high quality, areas where σ^2 is very low, the firm will cheat the consumer by offering a very low quality product at a relatively high price.

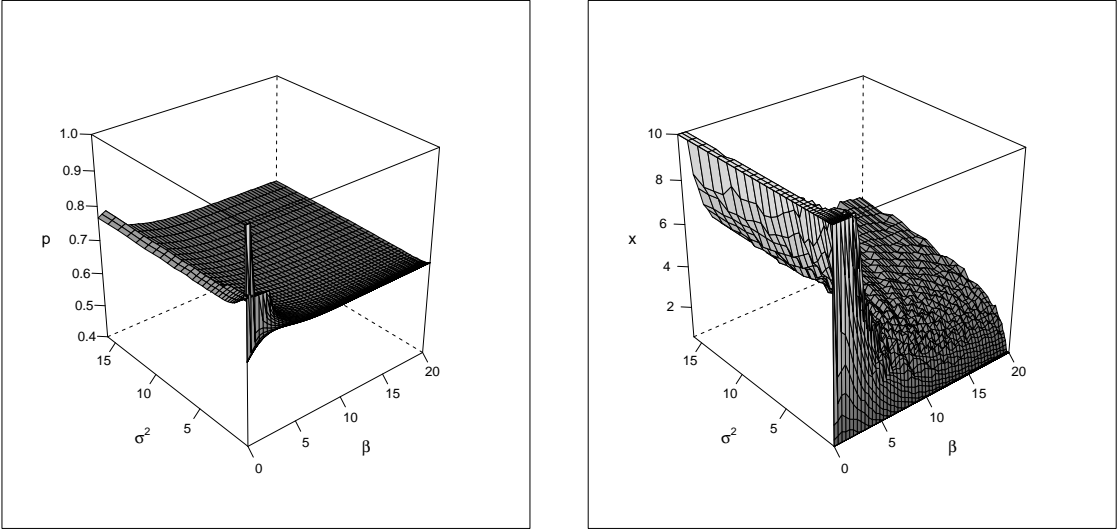


Figure 2: Policy functions for $\gamma = 0$

As shown in Table 2, the effect of signaling on price is exaggerated in the case of flexible quality. Specifically, the firm chooses to offer a lower quality product, $x = 2.3$, at the similarly higher price found in the fixed quality setting.¹¹ We should note that,

¹¹We simulate the firm's policy and start out at state variables of $\beta = 3.2$ and $s = 0.019$, which are

although average steady-state profits fall a bit moving from $\gamma = 0$ to $\gamma = \infty$, the present discounted value of future profits (not shown in the table) rises.¹² These results are robust to different parameters, as we discuss in Section 6.

Table 2: Comparison of Full Information and signaling Steady States

	$\gamma = \infty$		$\gamma = 0$	
	Full Info	signaling	Full Info	signaling
Price	0.71	0.72	0.71	0.72
Quality	3.6	3.6	3.64	2.3
Firm Profit	0.08	0.08	0.08	0.07
Consumer Surplus	0.85	0.82	0.86	0.5

In both full information cases, all the outcome variables (profits, prices, etc.) are averaged over costs. The values of outcome variables for the signaling model are the averages over the simulation, starting with the 500th iteration and ending with the 1000th. The first 500 iterations are removed for burn-in.

4 The Impact of Price Signaling on Cost Pass-through

We assess how price signaling affects price adjustment by simulating the model under a temporary cost shock. We begin at a base cost of 0.40, and for a decrease we drop the cost to 0.35, while for an increase we raise the cost to 0.45. We solve for optimal prices both before and after the cost change. Simulating both a cost increase and a cost decrease allows us to examine whether price signaling has any impact on the symmetry of pass-through.

the long-run values when x is fixed at 3.6. The steady state of the $\gamma = 0$ specification does not appear to be affected by the starting points of the simulation — we have always gotten to the same steady state for different initial beliefs.

¹²When the firm can choose quality, it initially chooses a slightly lower quality product to exploit beliefs. It subsequently raises quality over time to 2.3 as beliefs update (i.e., as β falls). In the $\gamma = \infty$ case, profits seem to be relatively constant over time.

The results of this exercise are depicted in Table 3. Specifically, this table shows the percentage impact of a cost shock on prices. We depict the case for $x = 2.3$ —the equilibrium quality under the case of zero adjustment costs—as well as the case for $x = 3.6$ —the equilibrium quality under the case of infinite adjustment costs. For completeness, we also show the infinite adjustment cost case under the assumption that $x = 2.3$. The table also shows cost pass-through when we shut down the dynamic component of the model (i.e., when we set $\delta = 0$), which we label the myopic model.

Relative to the full-information case, price signaling acts to considerably decrease cost pass-through. This effect is more pronounced for a cost decrease relative to a cost increase. For instance, in the case of $x = 3.6$, price signaling causes pass-through to fall from 51% to 46% for a cost increase, but to 45% for a cost decrease. As we will discuss later, most of the overall effect from signaling is coming from the myopic component of the model—the impact of signaling on the shape of the demand curve.

The asymmetry of cost pass-through seems to reverse when quality is chosen endogenously (i.e., $\gamma = 0$.) In this case, pass-through falls from 52% to 44% for a cost increase, but to 45 percent for a cost decrease. This reversal of the asymmetry likely arises due to the fact that the firm also adjusts quality in response to cost shocks. When cost decreases, the firm drops quality from 2.3 to 2.1. However, when cost increases, quality rises, but by a smaller amount to 2.4. Prices will tend to move with quality, and since quality drops more in response to cost decreases, prices do as well. Another interesting result is that pass-through seems to be lower for higher quality goods.

5 Discussion of Pass-Through Results

Our results indicate that price-signaling generates a considerable amount of price rigidity. In the following subsection we will discuss the role of demand-curvature and how this mechanism explains the link between price signaling and cost pass-through. We explain the intuition analytically under the simpler case of fixed quality, and then show that the intuition holds in the case where the firm can choose quality (i.e., flexible quality).

Table 3: The Impact of a Cost Shock on Prices

		$x = 2.3$			$x = 3.6$		
		Full Info	Signaling		Full Info	Signaling	
			Dynamic	Myopic		Dynamic	Myopic
$\gamma = \infty$	Cost Decrease	52	42	43	51	45	45
	Cost Increase	52	43	44	51	46	46
$\gamma = 0$	Cost Decrease	52	45		51		
	Cost Increase	52	44		51		

Notes: This table shows the percentage impact of a cost increase or decrease on prices. For all the simulations, we begin at a base cost of 0.40, and for a decrease we drop the cost to 0.35. For an increase we raise the cost to 0.45. In the signaling cases, we simulate the model for 1000 periods, holding cost fixed at its new value. That is, in period 1000 the cost is 0.40, and in the following period cost is 0.35 or 0.45.

5.1 Fixed Quality: Infinite Adjustment Costs

The simplest case considers a high adjustment cost (γ) so that the firm chooses quality in period 1 and never changes it after that. We will derive the impact of an idiosyncratic shock to c_0 on price. Recall that we assume that the stochastic process for the evolution of costs is

$$c_{jt} \sim U[\underline{c}_j, \bar{c}_j],$$

where c_{jt} is i.i.d. over time and j . To explain the intuition behind how cost shocks affect prices, we first assume that the firm is myopic $\delta = 0$, that γ is high so x is fixed over time, and that costs are nonstochastic: $c_{jt} = \bar{c}_j = \underline{c}_j, \quad \forall j, t$. We can derive the impact of an increase in c_0 on prices following analysis outlined in Weyl and Fabinger (2013) and Jaffe and Weyl (2013). We write the firm first-order conditions as

$$f(p) = -\frac{\partial D^{-1}}{\partial p} D(p) - (p - c_0 - c_1 x^\alpha) = 0. \quad (11)$$

If we add a tax t to the marginal cost, the first-order condition becomes

$$f(p) + t = 0.$$

By implicitly differentiating this first-order condition we get the equation for static pass-through:

$$\frac{\partial p}{\partial t} = \frac{-1}{f'(p)}. \quad (12)$$

Note that $f'(p) = D \frac{\partial^2 D}{\partial p^2} / \left(\frac{\partial D}{\partial p} \right)^2 - 2$ depends only on the first and second derivatives of demand, and is a measure of the curvature of the demand curve. This provides one way for signaling to affect pass-through. For the case of full information, where x is known to consumers, we could compute pass-through from equation (12). If signaling increases the curvature of demand relative to full information, then that increase in curvature will dampen pass-through. Intuitively, we would expect signaling to dampen pass-through in the short term, especially for cost decreases. If the monopolist observes a cost decrease, it has a disincentive to drop the price because a lower price signals lower quality.

The dynamics in our model provide an additional avenue for pass-through. Suppose that the firm is forward-looking, so $\delta > 0$. Then we can write the firm's first-order conditions in the same way as equation (11):

$$\underbrace{-\frac{\partial D^{-1}}{\partial p} D(p) - (p - c_0 - c_1 x^\alpha)}_{f(p)} - \delta \underbrace{\frac{\partial D^{-1}}{\partial p} \frac{\partial V'}{\partial \mathbf{S}} \frac{\partial L}{\partial p}}_{g(p)} + t = 0^{13}$$

An addition term $g(p)$ arises which is the derivative of the value function with respect to price.¹⁴ Note that this value function derivative is a function of t directly, since if we change the tax on costs, we change firm future profits. The dynamic analogue to our static pass-through equation (12) is

$$\frac{\partial p}{\partial t} = -\frac{1}{f'(p) + \frac{\partial g}{\partial p}}. \quad (13)$$

¹³The second term, g , is not a function of t because we are imagining a temporary cost shock. If the shock to costs were permanent this term would be a function of t as well.

¹⁴The partial of V is the Jacobian of V with respect to its arguments, and the L partial is a Jacobian as well. The prime indicates transpose, not a future value.

The term $\frac{\partial g}{\partial p}$ captures the dynamic effect of signaling on pass-through; if this term is negative, then signaling further dampens pass-through (in full information the firm is not forward-looking, so $\frac{\partial g}{\partial p} = 0$). The sign of this derivative will depend on the signs of the first and second derivatives of the value function with respect to price, and on the curvature of demand. To see how, note that if we express the value function as a function of price (subsuming the state transition) we can write

$$\frac{\partial g}{\partial p} = \frac{\frac{\partial^2 D}{\partial p^2}}{\left(\frac{\partial D}{\partial p}\right)^2} \frac{\partial V'}{\partial p} - \frac{1}{\frac{\partial D}{\partial p}} \frac{\partial^2 V'}{\partial p^2}. \quad (14)$$

Note that the first term in equation (5.1) contains the demand curvature from the static pass-through equation. We have previously argued that this term will be negative under signaling. Furthermore, for a sufficiently high price $\frac{\partial D}{\partial p}$ will be negative as well under signaling.¹⁵ For equation (5.1) to be negative under pass-through, the value function must be increasing and concave in price.

Table 4: Analytical Derivatives that Determine Pass-Through ($\gamma = \infty$)

	$x=2.3$		$x=3.6$	
	Full Info	Signaling	Full Info	Signaling
$f'(p)$	-1.93	-2.18	-1.98	-2.11
$\frac{\partial g}{\partial p}$		-0.028		-0.010

We found in Table 3 in the previous section that when x was fixed, prices were more rigid for a forward-looking firm than for a myopic one. We can relate this finding to the analytical equations (12) and (5.1) by computing the numerical equivalent in our simulation. These are shown in Table 4. Because $f'(p)$ falls from -1.93 to -2.18, signaling increases the curvature of demand acting to dampen pass-through. Additionally, because

¹⁵It is possible for $\frac{\partial D}{\partial p}$ to be positive under signaling for very low prices, because higher prices imply higher quality, increasing demand. The inclusion of the term $1 - \frac{p}{w}$ in our demand equation ensures that even if signaling is very strong, demand will eventually be decreasing in price.

$\frac{\partial g}{\partial p} < 0$ the firm's forward-looking behavior acts to further dampen pass-through. The magnitude of the static effect is much larger than the dynamic effect.

5.2 Flexible Quality: Zero Adjustment Cost

We now consider the case where x may also respond to changes in cost. In our simulations we discretized x , so if the cost shock is sufficiently small the firm will not change x , and the analysis should be similar to the previous section. However, for a large enough cost shock, x will change. When thinking about the change in x , the interesting case is where the firm is forward-looking—if the firm is myopic, it will set x as low as possible because increasing x does not affect current demand, but does increase marginal cost. The math is similar to the previous case, although there are some complications. For this exercise, we will assume that the set of x 's the firm can choose is continuous: $x \in [0, \infty]$.

When the firm can adjust x , there are two first-order conditions. The first one is the same as before, but we will relabel f as f_1 and g as g_1 . The second equation will have two similar terms,

$$\begin{aligned} f_2(p, x) &= -D(p)\alpha c_1 x^{\alpha-1} \\ g_2(p, x) &= \delta \frac{\partial V'}{\partial \mathbf{S}} \frac{\partial L}{\partial x}, \end{aligned}$$

and we can write the system of equations that defines the pass-through rate as

$$\begin{aligned} f_1(p) + g_1(p, x) + t &= 0 \\ f_2(p, x) + g_2(p, x) &= 0. \end{aligned}$$

It is straightforward to show that two pass-through rates arise from this system, one for p and one for x . They are

$$\frac{\partial p}{\partial t} = \frac{-1}{f'_1 + g_{1p} - g_{2p}g_{1x}/(f_{2x} + g_{2x})} \quad (15)$$

$$\frac{\partial x}{\partial t} = \frac{f_{2p} + g_{2p}}{(f'_1 + g_{1p})(f_{2x} + g_{2x}) - g_{2p}g_{1x}}. \quad (16)$$

The price pass-through equation is similar to the infinite adjustment cost case, except that the denominator contains an extra term that involves the derivatives of the first order condition with respect to x . Note that if the firm is myopic, the g terms will be zero in equation (15), and the equation will collapse to equation (5.1). Similarly, equation (16) collapses to $-\frac{\partial p}{\partial t} \frac{f_{2p}}{f_{2x}}$ which is negative.¹⁶

The logic of the previous paragraph suggests that one can break down the impact of pass-through into static and dynamic components as we did in Table 4. Under signaling, $f'_1(p) = -2.17$, and the static pass-through rate under signaling will be $1/2.17 = 0.46$. Since the overall pass-through rate under signaling is between 0.44 and 0.45 (Table 3), we can conclude that the optimal policy of a forward-looking firm is to dampen pass-through more than a myopic firm would.¹⁷

5.3 Generality of the Impact of Signaling on Pass-through

The results in Table 3, and our derivations of $f'(p)$ in the previous section, show that the most important impact of signaling on pass-through occurs through shape of the demand curve, measured by the term $f'(p)$. A follow-up question we now consider is: under what general conditions will signaling reduce the term $f'(p)$? Note that in Table 3, $f'_{FullInfo}(p)$ is somewhat larger than -2 and $f'_{Signaling}(p)$ is lower than -2. Hence, two sufficient conditions for signaling to result in dampened pass-through are

¹⁶Since $\frac{\partial p}{\partial t} > 0$ and $f_{2p} < 0$ and $f_{2x} < 0$, we should get $\frac{\partial x}{\partial t} < 0$, which is intuitive.

¹⁷Although, as with $\gamma = \infty$, the bulk of the impact of signaling on pass-through operates through the static effect.

$$\begin{aligned}
f'_{FullInfo}(p) + 2 &= D \frac{\partial^2 D}{\partial p^2} / \left(\frac{\partial D}{\partial p} \right)^2 > 0 \\
f'_{Signaling}(p) + 2 &= D \frac{\partial^2 D}{\partial p^2} / \left(\frac{\partial D}{\partial p} \right)^2 < 0.
\end{aligned}$$

It can be shown that in the full information model the static pass-through term $f'(p)$ reduces to

$$f'_{FullInfo}(p) + 2 = \bar{w}\alpha \frac{\left(1 - \frac{p}{\bar{w}}\right) \frac{1}{1+\exp(x-\alpha p)} \left(2 + \bar{w}\alpha \left(1 - \frac{p}{\bar{w}}\right) \frac{1-\exp(x-\alpha p)}{1+\exp(x-\alpha p)}\right)}{\left(1 + \bar{w}\alpha \left(1 - \frac{p}{\bar{w}}\right) \frac{1}{1+\exp(x-\alpha p)}\right)^2}, \quad (17)$$

and in the signaling model this is

$$f'_{Signaling}(p) + 2 = - \frac{(\beta - \alpha) \left(1 - \frac{p}{\bar{w}}\right) \frac{1}{1+\exp((\beta-\alpha)p)} \left(\frac{2}{\bar{w}} - (\beta - \alpha) \left(1 - \frac{p}{\bar{w}}\right) \frac{1-\exp((\beta-\alpha)p)}{1+\exp((\beta-\alpha)p)}\right)}{\left(\frac{1}{\bar{w}} - (\beta - \alpha) \left(1 - \frac{p}{\bar{w}}\right) \frac{1}{1+\exp((\beta-\alpha)p)}\right)^2}. \quad (18)$$

First consider equation (17). This equation will be positive as long as

$$2 + \bar{w}\alpha \left(1 - \frac{p}{\bar{w}}\right) \frac{1 - \exp(x - \alpha p)}{1 + \exp(x - \alpha p)} > 0.$$

At our chosen normalization of $\alpha = 1$ and $\bar{w} = 1$, this inequality will always hold because $\frac{1-\exp(x-\alpha p)}{1+\exp(x-\alpha p)} > -1$. In general, as long as the term α that measures the marginal utility of income is not too large relative to the maximum wealth, the inequality will hold.

Now let us turn to equation ((18)). A sufficient condition for this to be negative is if

$$(\beta - \alpha) \left(\frac{2}{\bar{w}} - (\beta - \alpha) \left(1 - \frac{p}{\bar{w}}\right) \frac{1 - \exp((\beta - \alpha)p)}{1 + \exp((\beta - \alpha)p)}\right) > 0.$$

This result suggests that signaling is more likely to dampen pass-through for relatively expensive goods—goods that take up a significant share of income. To see this, note that this inequality will hold as long as β is in the range $[\alpha, 2/(\bar{w} - p) + \alpha]$. In the steady state

of the signaling model, β will be close to x/p , the ratio of quality to price. What this inequality tells us is that the quality to price ratio should be larger than the consumer's willingness to pay for quality, α , but not too much larger as it depends on the inverse of the consumer's remaining budget— $\bar{w} - p$. That is, the upper bound on β will rise as the share of income spent on the monopolist's product increases.¹⁸

6 Pass-through at Alternative Parameters

In this section, we investigate the robustness of our results by altering the parameters of the model. Specifically, we assess the model under alternate distributions of costs and the exponential forgetting factor λ . In Table 5 we show the full information and signaling steady states for four different parameterizations of the model: a version with a lower cost distribution of $U[0.2, 0.3]$, a higher cost distribution of $U[0.5, 0.6]$, a low value of the exponential forgetting factor, $\lambda = 0.95$, and a high value of $\lambda = 0.999$.

Overall, Table 5 shows an emerging pattern similar to that found in our previous analysis. Relative to the case of full information, signaling causes consumer surplus to fall—the firm lowers the quality of its product and sells it at a higher or similar price. This is especially true in the case where the firm can freely adjust quality. An exception is the case where the cost distribution is high and quality is fixed (i.e., the case of infinite adjustment costs). In this instance, the signaling and full-information cases result in similar steady states.

Table 6 shows how pass-through responds to changes in the cost distribution and λ . There are three findings that are robust to changes in the cost distribution and λ : First, signaling results in dampened pass-through relative to full information. Second,

¹⁸At our chosen parameter values, note that the bounds are approximately $[1, 7.7]$, and the quality to price ratio is around 3 to 5 (Table 2). The inequality that defines how signaling affects pass-through looks like it depends on the units used to measure wealth due to the inclusion of the term $2/(\bar{w} - p)$. This is not the case however, for two reasons. First, if we scaled up wealth and prices by some factor κ , we would have to scale down the marginal utility of income, α , by the same amount. Additionally, since the term β measures how many units of utility a consumer expects to get from a product of price p , β would also scale by the same factor in the steady state. The inequality would then become $\beta/\kappa \in [\alpha/\kappa, 2/(\kappa\bar{w} - \kappa p) + \alpha/\kappa]$.

Table 5: Robustness: Comparison of Full Information and Signaling Steady States

	$\gamma = \infty$		$\gamma = 0$	
	Full Info	Signaling	Full Info	Signaling
<hr/> $c_{0t} \sim U[0.2, 0.3]$ <hr/>				
Price	0.64	0.65	0.64	0.65
Quality	3.8	3.6	3.76	2.35
Firm Profit	0.12	0.12	0.12	0.12
Consumer Surplus	1.16	1.04	1.15	0.65
<hr/>				
$c_{0t} \sim U[0.5, 0.6]$ <hr/>				
Price	0.79	0.79	0.79	0.79
Quality	3.5	3.6	3.44	2.03
Firm Profit	0.04	0.04	0.04	0.04
Consumer Surplus	0.6	0.6	0.58	0.32
<hr/>				
$\lambda = 0.95$ <hr/>				
Price	0.71	0.72	0.71	0.71
Quality	3.6	3.6	3.64	3.14
Firm Profit	0.08	0.08	0.08	0.08
Consumer Surplus	0.85	0.82	0.86	0.72
<hr/>				
$\lambda = 0.999$ <hr/>				
Price	0.71	0.72	0.71	0.71
Quality	3.6	3.6	3.64	0.95
Firm Profit	0.08	0.08	0.08	0.06
Consumer Surplus	0.85	0.82	0.86	0.23

In both full information cases, all the outcome variables (profits, prices, etc.) are averaged over costs. The values of outcome variables for the signaling model are the averages over the simulation, starting with the 500th iteration and ending with the 1000th. The first 500 iterations are removed for burn-in.

pass-through rates are slightly higher when the firm is myopic. Third, pass-through is asymmetric when the cost of changing quality is high. In contrast, with low adjustment costs, the asymmetry of pass-through varies with λ and with the cost distribution.

Recall that we found in the last section that when adjustment costs are low, the firm passes on more of cost decreases than increases. We can see in the second column of Table 6 that for low and high cost distributions the firm passes on more of cost increases than decreases. When λ is low, the firm passes on more of cost decreases than increases; this finding reverses when λ is 0.999. Overall, it appears that the asymmetry of pass-through depends on both the cost distribution and the speed at which consumer beliefs update. For extreme values of these parameters firms pass on more of cost increases than decreases. For values that are more in the middle we observe the opposite.

7 Conclusion

Our results indicate that price signaling affects the optimal price, quality, and the rate at which the optimal price adjusts after a cost shock. Our result that price signaling acts to raise the equilibrium price relative to the case of full information, corresponds well with classical works on price signaling, for example, Wolinsky (1983). Our model, however, extends this line of literature by allowing the firm to choose quality, and therefore dynamically affect the beliefs of the consumer about the price signal. We find that under this type of setting the firm chooses to offer a lower quality product at higher price relative to the full information case. We contribute to the literature examining the drivers of price stickiness by deriving conditions under which signaling leads to dampened pass-through.

There are many ways in which our work could be extended. One possibility would be to examine how results might change under competition. If more competition tends to result in decreased prices and increased quality, then the impact of signaling on pass-through could be mitigated. It could also be interesting to examine the robustness of our conclusions to the functional form used for the demand curve. We think that our proposed functional form is a reasonable starting point because it is based on the logit demand model, which is standard in empirical work. Finally, another step would be to

Table 6: Robustness: The Impact of a Cost Shock on Prices

		$x = 2.3$			$x = 3.6$		
		Full Info	signaling		Full Info	signaling	
			Dynamic	Myopic		Dynamic	Myopic
$c_{0t} \sim U[0.2, 0.3]$							
$\gamma = \infty$	Cost Decrease	52	39	40	51	42	43
	Cost Increase	52	40	41	51	43	44
$\gamma = 0$	Cost Decrease	52	36		51		
	Cost Increase	52	45		51		
$c_{0t} \sim U[0.5, 0.6]$							
$\gamma = \infty$	Cost Decrease	52	45	45	51	47	47
	Cost Increase	52	46	46	51	47	47
$\gamma = 0$	Cost Decrease	52	47		51		
	Cost Increase	52	49		51		
$\lambda = 0.95$							
$\gamma = \infty$	Cost Decrease	52	42	43	51	44	45
	Cost Increase	52	43	44	51	45	46
$\gamma = 0$	Cost Decrease	52	50		51		
	Cost Increase	52	40		51		
$\lambda = 0.999$							
$\gamma = \infty$	Cost Decrease	52	43	43	51	45	45
	Cost Increase	52	44	44	51	46	46
$\gamma = 0$	Cost Decrease	52	45		51		
	Cost Increase	52	46		51		

Notes: This table shows the percentage impact of a cost increase or decrease on prices. For all the simulations, we begin at a base cost at the midpoint of the c_{0t} distribution, and for a decrease we drop cost to the distribution's lower bound. For an increase we raise cost to its upper bound. In the signaling cases, we simulate the model for 1000 periods, holding cost fixed at its new value. That is, in period 1000 the cost is $0.5 * (\bar{c} + \underline{c})$ and in the following period cost is \bar{c} or \underline{c} .

take our model to the data and to determine the prevalence of signaling in actual markets. This question has received attention in the marketing literature (Erdem, Keane, and Sun 2000).

References

- Bersetkas, D. and J. Tsitsiklas (1996). *Neuro-Dynamic Programming*. Belmont, MA: Athena Scientific.
- Board, S. and M. Meyer-ter Vehn (2013). Reputation for quality. *Econometrica* 81(6), 2381–2462.
- Cabral, L. and A. Fishman (2012). Business as usual: A consumer search theory of sticky prices and asymmetric price adjustment. *International Journal of Industrial Organization* 30(4), 371–376.
- Erdem, T., M. Keane, and B. Sun (2000). A dynamic model of brand choice when price and advertising signal product quality. *Marketing Science* 27(6), 1111–1125.
- Fershtman, C. and A. Pakes (2012, November). Dynamic games with asymmetric information: A framework for empirical work. *Quarterly Journal of Economics* 127(4), 1611–1661.
- Gali, J. and M. Gertler (1999). Inflation dynamics: A structural econometric analysis. *Journal of Monetary Economics* 44, 195–222.
- Jaffe, S. and G. Weyl (2013). The first order approach to merger analysis. *American Economics Journal: Microeconomics* 5(4), 188–218.
- Judd, K. L. (1998). *Numerical Methods in Economics*. London, England: The MIT Press.
- Judd, K. L. and M. H. Riordan (1994). Price and quality in a new product monopoly. *The Review of Economic Studies* 61(4), 773–789.
- Kimball, M. S. (1995). The quantitative analytics of the basic neomonetarist model. *Journal of Money, Credit and Banking* 27(4), 1241–77.
- Klein, B. and K. B. Leffler (1981). The role of market forces in assuring contractual performance. *Journal of Political Economy* 89(4), 615–641.
- Klenow, P. J. and J. L. Willis (2006). Real rigidities and nominal price changes. Research working paper rwp 06-03, Federal Reserve Bank of Kansas City.

- L’Huillier, J.-P. (2013). Consumers’ imperfect information and price rigidities. Working paper.
- Lucas, R. E. (1972). Expectations and the neutrality of money. *Journal of Economic Theory* 4, 103–124.
- Mackowiak, B. and M. Wiederholt (2009). Optimal sticky prices under rational inattention. *American Economic Review* 99(3), 769–803.
- Mankiw, G. and R. Reis (2002). Sticky information versus sticky prices: a proposal to replace the new keynesian phillips curve. *Quarterly Journal of Economics* 117(4), 1295–1328.
- Nakamura, E. and D. Zerom (2010, July). Accounting for incomplete pass-through. *Review of Economic Studies* 77(3), 1192–1230.
- Null, B. (2008). The nested dirichlet distribution: Properties and applications. Working paper.
- Orphanides, A. and J. Williams (2004). *Imperfect Knowledge, Inflation Expectations, and Monetary Policy*. ”National Bureau of Economic Research, Inc”. in *The Inflation-Targeting Debate*.
- Peltzman, S. (2000, June). Prices rise faster than they fall. *The Journal of Political Economy* 108(3), 466–502.
- Reis, R. (2006). Inattentive producers. *Review of Economic Studies* 73, 793821.
- Rotemberg, J. J. (2005). Customer anger at price increases, changes in the frequency of price adjustment and monetary policy. *Journal of Monetary Economics* 52(4), 829–852.
- Rust, J. (1996). Numerical dynamic programming in economics. In H. Amman, D. Kendrick, and J. Rust (Eds.), *Handbook of Computational Economics*. Elsevier, North Holland.
- Shapiro, C. (1983). Premiums for high quality products as returns to reputations. *The Quarterly Journal of Economics* 98(4), 659–680.
- Weiser, A. and S. Zarantonello (1988). A note on piecewise linear and multilinear table interpolation in many dimensions. *Mathematics of Computation* 50, 189–196.

- Weyl, G. and M. Fabinger (2013). Pass-through as an economic tool: Principles of incidence under imperfect competition. *Journal of Political Economy* 121(3), 528–583.
- Wieland, V. (2000, August). Monetary policy, parameter uncertainty and optimal learning. *Journal of Monetary Economics* 46(1), 199–228.
- Wolinsky, A. (1983). Prices as signals of product quality prices as signals of product quality. *The Review of Economic Studies* 50(4), 647–658.

Appendix

8 Description of the Multilinear Interpolation Algorithm

We apply the algorithm in Weiser and Zarantonello (1988) to interpolate the value function over the two dimensional continuous part of the state space. The algorithm proceeds as follows. As we have described earlier, we split the two dimensional unit square into an N_g by N_g regular grid, where $N_g = 21$. First, given a vector $(\tilde{\alpha}_H, \tilde{\alpha}_L)$, we figure out which grid points contain the vector: we find the integer i such that $(i - 1)/(N_g - 1) \leq \tilde{\alpha}_H \leq i/(N_g - 1)$, and the integer j such that $(j - 1)/(N_g - 1) \leq \tilde{\alpha}_L \leq j/(N_g - 1)$. We then scale the square up to a $[0, 1]$ by $[0, 1]$ square, transforming $(\tilde{\alpha}_H, \tilde{\alpha}_L)$ to

$$(x_1, x_2) = \left((N_g - 1) \left(\tilde{\alpha}_H - \frac{(i - 1)}{N_g - 1} \right), (N_g - 1) \left(\tilde{\alpha}_L - \frac{(j - 1)}{N_g - 1} \right) \right).$$

We also relabel the save value functions at the vertices of the square as $V_n(l, m)$, where $(l, m) \in \{0, 1\} \times \{0, 1\}$. We then figure out which simplex of the unit square contains (x_1, x_2) . To do this, we find a permutation of $(1, 2)$, $(p(1), p(2))$ such that $x_{p(1)} \leq x_{p(2)}$. The interpolated value function can then be constructed as a linear combination of $V_n(l, m)$ at the vertices using the following algorithm:

1. Start with $s_0 = (1, 1)$
2. Let $\hat{V} = V_n(s_0)$
3. Let $i = 1$
4. Let $s_i = s_{i-1} - e_{p(i)}$, where $e_{p(i)}$ is a 2-vector with 1 in position $e_{p(i)}$ and 0 everywhere else.
5. Let $\hat{V} = \hat{V} + (1 - x_{p(i)})(V_n(s_i) - V_n(s_{i-1}))$
6. Increment i by 1

7. If $i \leq 2$, go to step 4. Otherwise, return the interpolated value function \hat{V} .

Note that this algorithm can easily be extended to continuous state spaces with $N > 2$ dimensions. It can also be extended to include nonregular grids. For a nonregular grid, the only change is that we find the grid points containing $(\tilde{\alpha}_H, \tilde{\alpha}_L)$, and scale up the grid points to the unit hypercube.

9 The Effect of Changes in β on the Optimal Price (Myopic Firm)

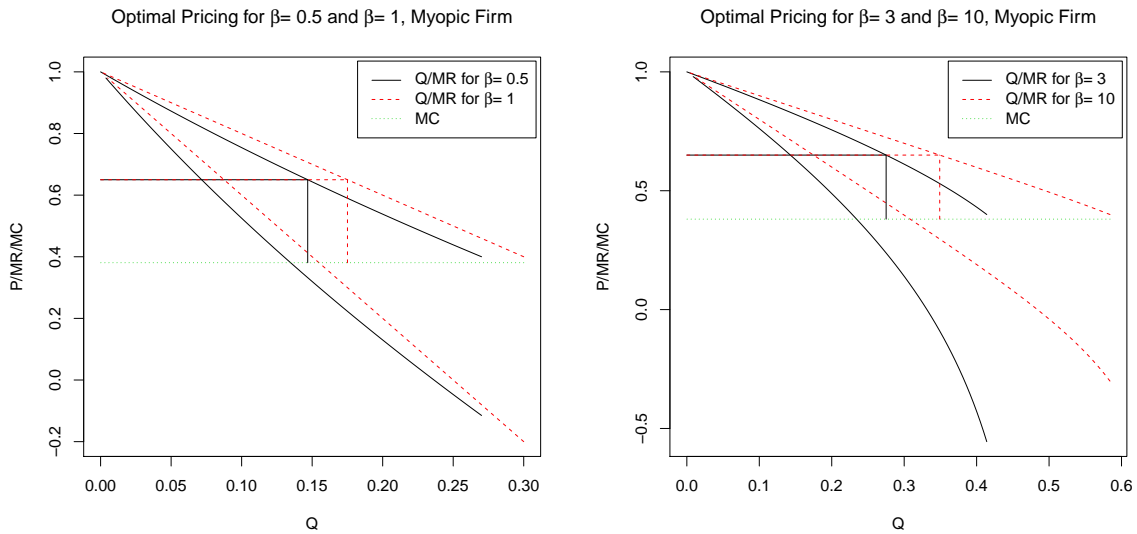


Figure 3: Optimal Myopic Price for Different Values of β

10 Steady-State Simulation

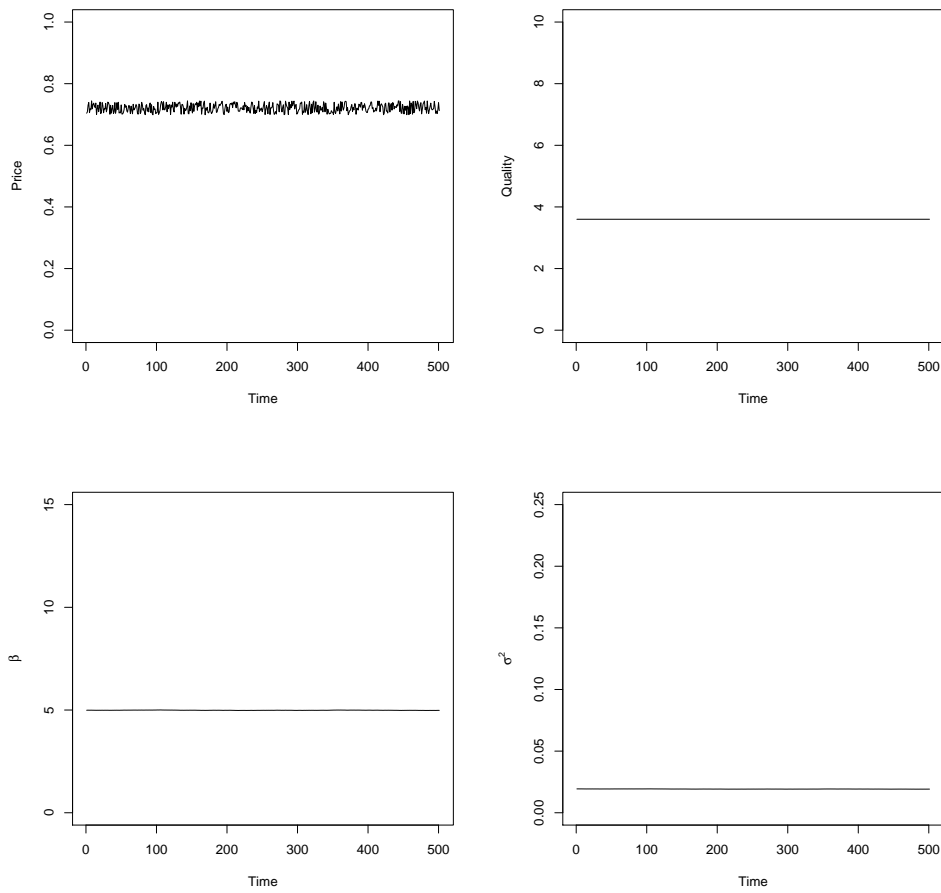


Figure 4: Simulation Results, $\gamma = \infty$

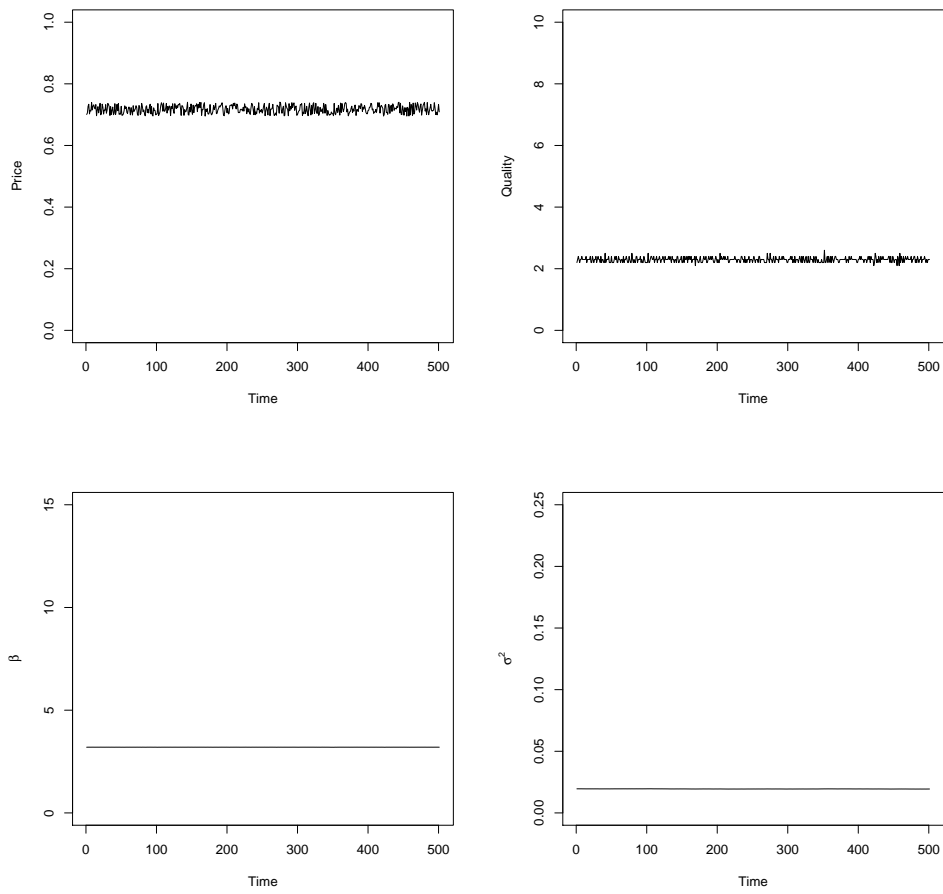


Figure 5: Simulation Results, $\gamma = 0$

11 Robustness: Discrete Price Model

In this section we present a robustness check against the linear functional form of our consumer learning process. For this section we assume that quality, x_t , can take on one of two values $\{x^L, x^H\}$, where $x^H > x^L$ and that the firm can only choose a discrete number of prices: $p_t \in \{p^1, \dots, p^K\}$. Consumers correctly believe that the joint distribution of p_t and x_t are discrete, and update those beliefs using a multinomial-Dirichlet Bayesian updating process.

11.1 Demand

We assume there exists a continuous measure of consumers indexed by i , each with taste for quality $v_i \sim U[0, \bar{v}]$ and endowed with wealth $w_i \sim U[0, \bar{w}]$. Under an additively linear direct utility function, it follows that consumer i 's indirect utility upon purchase of a product with vertical quality x_t and price p_t is:

$$U_i = v_i x_t - p_t, \quad (19)$$

The more general case consists of the scenario where the consumer does not know the true quality of the product, x_t , and thus must infer the quality based on her information set (that is, her beliefs). We assume that the consumer believes the relationship between price and quality takes the form:

$$Prob(x_t = x^j, p_t = p^k) = \mu_{jk}, \quad (20)$$

where $j, k \in \{H, L\}$. Therefore, upon observing the good's price, p_t , the consumer forms an expectation of its quality based on the conditional probability:

$$E[x_t | p_t = p^k] = \left(\frac{\mu_{Hk}}{\mu_{Hk} + \mu_{Lk}} \right) x^H + \left(\frac{\mu_{Lk}}{\mu_{Hk} + \mu_{Lk}} \right) x^L. \quad (21)$$

where $\mu_{Hk} + \mu_{Lk}$ is the probability of observing price p^k . Consumers' beliefs about the unknown parameters, μ_{jk} , are modeled as Dirichlet:

$$f(\mu_{jk}) = Dir(\alpha_{HH}, \alpha_{HL}, \alpha_{LH}, \alpha_{LL}), \quad (22)$$

where $\alpha_{jk} > 0$. Note that under this assumption,

$$E[\mu_{jk}] = \frac{\alpha_{jk}}{\sum_{j,k \in \{H,L\}} \alpha_{jk}}$$

and

$$E\left[\frac{\mu_{jk}}{\mu_{Hk} + \mu_{Lk}}\right] = \frac{\alpha_{jk}}{\sum_{j \in \{H,L\}} \alpha_{jk}}.^{19}$$

If consumers hypothetically live for an infinite amount of time, they will use all past information of prices and quality over the course of the good's life to infer $\boldsymbol{\alpha} = (\alpha_{HH}, \alpha_{HL}, \alpha_{LH}, \alpha_{LL})$. The information in period t can be summarized in the vector $\boldsymbol{\alpha}_t = (a_{HH,t}, a_{HL,t}, a_{LH,t}, a_{LL,t})$, where consumers believe that the distribution of $\boldsymbol{\alpha}$ at period t is $Dir(\boldsymbol{\alpha}_t)$. By the end of period t , consumers have observed both price and quality, and will update their beliefs using a multinomial-Dirichlet Bayesian updating process:

$$\alpha_{jk,t+1} = \begin{cases} \alpha_{jk,t} + 1 & \text{if } x_t = x^j \text{ and } p_t = p^k \\ \alpha_{jk,t} & \text{otherwise} \end{cases} \quad (23)$$

By updating their beliefs each period, consumers learn the relationship between price and quality. One problematic feature of assuming an infinitely lived consumer is that, in the limit, her beliefs about the relationship between price and quality become fixed. That is, $\boldsymbol{\alpha}$ will eventually be learned with zero standard error. In this case, any alteration of the stochastic process of price and quality will not cause beliefs to update.²⁰

To allow for beliefs to continue updating indefinitely, we propose that consumers recall the previous period's information with error. In our multinomial-Dirichlet Bayesian updating problem, demand for the product will be a function of consumer beliefs about

¹⁹Because the μ_{jk} 's are Dirichlet, the random variables $\frac{\mu_{jk}}{\mu_{Hk} + \mu_{Lk}}$ also follow a Dirichlet distribution (see Null (2008) and citations therein).

²⁰We have solved and simulated a version of our model with infinitely-lived consumers, and we find that when consumers become sure enough of their beliefs, the firm will charge a high price but lower product quality, and beliefs will update so slowly that the firm can maintain a low quality. In the limit, as consumers become completely sure of their beliefs, the firm can keep quality low and price high and consumers will continue to believe that *ex-ante* quality will be high.

quality conditional on price (see Equation (21)). We therefore assume that the consumers keep track of the conditional probabilities of observing x^H given p_t :

$$Prob(x_t = x^H | p_t = p^k, \boldsymbol{\alpha}_t) = \frac{\alpha_{Hk,t}}{\alpha_{Hk,t} + \alpha_{Lk,t}} = \tilde{\alpha}_{k,t}.$$

The memory error, which we denote ν_t , subsequently affects this conditional probability by introducing noise into the consumer's belief about quality upon observing price. We assume that there are two error terms, ν_{Ht} and ν_{Lt} , each corresponding to conditioning on $p_t = p^H$ or $p_t = p^L$, respectively. We assume that the ν_{kt} 's are i.i.d. across time and k , and that they follow a discrete distribution,

$$\nu_{kt} = \begin{cases} \bar{\nu} & \text{with probability } \pi_\nu \\ \underline{\nu} & \text{with probability } 1 - \pi_\nu \end{cases} \quad (24)$$

We further assume that consumers have limited memory in the sense that the consumer's information set only includes the last N_k observations, where $\alpha_{Hk,t} + \alpha_{Lk,t} = N_k$. In this sense, N_k determines the weight the consumer places on the observed current price in affecting her beliefs about quality. The modified Bayesian updating process in our overlapping generations type model is

$$\tilde{\alpha}_{k,t+1} = l(p_t, x_t, \tilde{\alpha}_{k,t}) = \begin{cases} \frac{N_k}{N_k+1} \tilde{\alpha}_{k,t} + \frac{1}{N_k+1} + \frac{N_k}{N_k+1} \nu_{kt} & \text{if } x_t = x^H \text{ and } p_t = p^k \\ \frac{N_k}{N_k+1} \tilde{\alpha}_{k,t} + \frac{N_k}{N_k+1} \nu_{kt} & \text{if } x_t = x^L \text{ and } p_t = p^k \\ \tilde{\alpha}_{k,t} + \nu_{kt} & \text{otherwise.} \end{cases} \quad (25)$$

The intuition behind this quasi-Bayesian updating formula is as follows. We can interpret the $\alpha_{jk,t}$'s as the number of times each quality and price combination has been observed. Consumers have a limited amount of memory, in the sense that they can only remember the past N_k observations for each value of p^k . Every period the α_{jk} 's are perturbed by an error that approximately preserves the average conditional probability $\tilde{\alpha}_{k,t}$; if the firm does not charge p^k in period t , then $E[\tilde{\alpha}_{k,t}] = E[\tilde{\alpha}_{k,t+1}]$. If the firm does charge p^k , then an observation must be added to either the high quality or low quality bin, depending on the value of x that is chosen. The multiplication by $\frac{N_k}{N_k+1}$ in this case keeps the total number of observations for p^k fixed at N_k . Note that if the firm chooses

to produce the high quality good, for example, consumer beliefs about the probability of a high quality good given p^k will increase on average by a factor of $\frac{1}{N_k+1}$, which is the same as how much the probability would go up under the standard Bayesian updating procedure. Note that the size of N_k will determine how quickly consumer beliefs adjust. If N_k is large, beliefs will adjust slowly in response to new information; if it is small, new information will receive greater weight.

Thus, in the limit, there always exists uncertainty about the underlying parameter vector α . In every period, it is only necessary for consumers to track the state variable $\tilde{\alpha}_t = (\tilde{\alpha}_{H,t}, \tilde{\alpha}_{L,t})$. It follows from (19) that consumer i will purchase the product if both $E(x_t|p_t, \tilde{\alpha}_t) - p_t > 0$ and $p_t < w_i$. Hence, the demand curve for each period that the firm faces is:

$$D(p_t, x_t, \tilde{\alpha}_t) = \begin{cases} \left(1 - \frac{p_t}{\bar{w}}\right) \left(1 - \frac{p_t}{\bar{v}E(x_t|p_t, \tilde{\alpha}_t)}\right) & \text{if } p_t \leq \bar{v}E(x_t|p_t, \tilde{\alpha}_t) \text{ and } p_t < \bar{w} \\ 0 & \text{otherwise} \end{cases} . \quad (26)$$

We define the mapping $L : \mathbb{R}^4 \rightarrow \mathbb{R}^2$ as follows:

$$\tilde{\alpha}_{t+1} = L(x_t, p_t, \tilde{\alpha}_t) = (l(x_t, p_t, \alpha_{H,t}), l(x_t, p_t, \alpha_{L,t})) \quad (27)$$

Equation (27) defines the evolution of beliefs in vector form, conditional on period t price and quality.

11.2 Firm

A monopolist maximizes profits in a dynamic sense, in that it chooses price and quality keeping in mind how this joint decision alters consumers' expectation of quality. We assume that the firm's marginal cost, \mathbf{c}_t , is a stochastic vector (c_{0t}, c_{1t}, c_{2t}) . Each draw of \mathbf{c}_t is i.i.d. across time with distribution $F_c(\cdot)$. The stochastic nature of cost can include a multitude of exogenous factors such as process innovations, weather disruptions, or factory malfunctions. Let total current period profits be represented as $\pi(p_t, x_t, \tilde{\alpha}_t, x_{t-1}, \mathbf{c}_t) = D(p_t, x_t, \tilde{\alpha}_t)[p_t - c_{0t} - c_{1t}x_t] - c_{2t}x_t - \gamma \mathbf{1}\{x_t \neq x_{t-1}\}$. The cost $c_{0t} \geq 0$ reflects the marginal cost of producing a quality 0 product, while $c_{1t} \geq 0$ allows for higher quality products to have higher production costs. c_{2t} can be interpreted as

a fixed cost of producing a higher quality product, and γ is an adjustment cost. The adjustment cost incorporates the idea that quality is not as easy to change as price: for example, creating a higher or lower quality product may involve retooling production plants.

It follows that the firm will choose price and quality to maximize its discounted stream of profits:

$$\max_{\{p_{t+i}, x_{t+i}\}_{i=0}^{\infty}} E \left[\sum_{i=0}^{\infty} \delta^i \pi(p_{t+i}, x_{t+i}, \tilde{\alpha}_{t+i}, x_{t+i-1}, \mathbf{c}_{t+i}) | \tilde{\alpha}_t \right], \quad (28)$$

subject to the evolution of the state variable $\tilde{\alpha}_t$ in equation (27). Specifically, the state variable $\tilde{\alpha}_t$ will evolve according to how consumers update their beliefs upon the firm's choices of price and quality. This can be intuited from writing the firm's problem in the form of a Bellman equation:

$$V(\tilde{\alpha}, x_-) = \int \max_{p, x} \{ \pi(p, x, \tilde{\alpha}, x_-, \mathbf{c}) + \delta \int V(L(p, x, \tilde{\alpha}), x) dF_\nu(\nu) \} dF_c(\mathbf{c}). \quad (29)$$

where f_ε , f_ν , and f_γ are the distributions of the respective shocks and x_- represents the previous period's quality choice. Existence of a fixed point rests on the fact that the per-period profit function is continuous and bounded in $\tilde{\alpha}$ (see Rust (1996)). The optimal choice of price and quality will come from the solution to the Bellman equation. Policy and value functions can be obtained from iterating the Bellman equation at an initial guess. We describe our algorithm for solving for the value and policy functions in more detail in Section 11.3.

11.3 Simulation Setup and Numerical Methods

Our procedure for solving for the value and policy functions is similar to that outlined in the body of the paper. Our problem has two continuous state variables, which are the probabilities people believe the product quality is high conditional on a high or low price, and last period's quality choice, which directly affects profits through the adjustment cost. First, we split the continuous part of the state space into a finite numbers of grid

points, and interpolate the value function in between points. We choose a regular 21 by 21 grid for interpolation. We solve for the value and policy functions at each $(\tilde{\alpha}_H, \tilde{\alpha}_L)$ point on the 21 by 21 grid, for each value of x_{t-1} ; we index each such state point from $i = 1, \dots, N_s = 441$, and denote each $(\tilde{\alpha}_H, \tilde{\alpha}_L, x_{t-1})$ combination as \mathbf{s}_i .

The value and policy function iteration then proceeds as follows. Indexing each iteration with n , at $n = 1$ we begin with a guess that the value function V_n is 0 at all points \mathbf{s}_i . We then solve for the optimal policy function $(x_n(\mathbf{s}_i, \mathbf{c}), p_n(\mathbf{s}_i, \mathbf{c}))$ at each state space point, \mathbf{s}_i , and possible cost draw \mathbf{c} :

$$(x_n(\mathbf{s}_i, \mathbf{c}), p_n(\mathbf{s}_i, \mathbf{c})) = \arg \max_{x,p} \{ \pi(p, x, \mathbf{s}_i, \mathbf{c}) + \delta E_{\nu} V_n(L(p, x, \mathbf{s}_i)) \} \quad (30)$$

We then solve for the value function that would be obtained if these policies were fixed. Operationally, we iterate on the value function contraction mapping with fixed policy functions until convergence. Denoting a policy iteration by n_p , we start by setting $V_1^p = V_n$ and update using the equation

$$V_{n_p+1}^p(\mathbf{s}_i) = E_{\mathbf{c}} \left[\pi(p_n(\mathbf{s}_i, \mathbf{c}), x_n(\mathbf{s}_i, \mathbf{c}), \mathbf{s}_i, \mathbf{c}) + \delta E_{\nu} V_{n_p}^p(L(p_n(\mathbf{s}_i, \mathbf{c}), x_n(\mathbf{s}_i, \mathbf{c}), \mathbf{s}_i)) \right]. \quad (31)$$

We assume that the policy step has converged when $\max_{i=1, N_s^j} \|V_{n_p+1}^p(\mathbf{s}_i^j) - V_{n_p}^p(\mathbf{s}_i^j)\| < \epsilon_p$, where we choose $\epsilon_p = 1e - 4$. The policy iteration step in equation (31) converges very quickly. Once the policy step converges, we set $V_n = V^p$, and solve for the optimal policies at step $n + 1$ using equation (30). The algorithm converges when $\max_{i=1, N_s} \|V_{n+1}(\mathbf{s}_i) - V_n(\mathbf{s}_i)\| < \epsilon_v$, where we set $\epsilon_v = 1e - 4$. We have found that this algorithm converges much more quickly than standard value function iteration, where one would solve for the optimal policy every time one updated the value function. The parameter values where we solve our problem are laid out in Table 7. Unlike the model in the body of the paper, in this version of the model we assume that costs are constant over time.

11.4 Simulation of a permanent shock to the marginal cost of production

In this section we investigate the extent to which price signaling may induce price stickiness, in the sense that a high quality firm may not drop its price in response to a

Table 7: Simulation Parameters

Variable	Value
δ	0.95
\bar{v}	0.6
\bar{w}	1
x^L	0.25
x^H	2
c_0	0.05
c_1	0
c_2	0.001
p^L	0.25
p^H	0.5
$\underline{\nu}$	-0.001
$\bar{\nu}$	0.001
π_ν	0.5
γ	0.003

surprise cost shock, even though in the full information case dropping the price would be optimal. To do this, we solve our model again, assuming that the marginal cost of production, c_0 , is 0 rather than 0.05. If there is no price signaling, it is optimal for the firm to lower its price in response to a cost shock like this. We lay out the optimal prices, costs and profits for the full information case in Table 8. Note that there are no dynamics in this case, as we assume that consumers know product quality upon entering the market.

Table 8: Optimal Price and Quality In Response to a Cost Shock, Full Information

	$c_0 = 0.05$	$c_0 = 0$
Optimal Price	0.5	0.25
Optimal Quality	2	2
Profits	0.129	0.146

To see what happens in the price signaling case, we first simulated the model for both $N_k = 50$ and $N_k = 1000$ for 10,000 periods. For a cost of 0.05, the steady-state price is 0.5 in each case, and the steady state-quality is 2. Additionally, in our simulations the steady-state $\tilde{\alpha}_{Ht}$ is close to 1, while $\tilde{\alpha}_{Lt}$ is low, settling at around 0.05. We see what the impact of a cost shock in period 10,000 would be by looking at the firm pricing policy functions, which are shown for the high cost in Figure 7 and at the low cost in Figure 6 for $N_k = 50$. If c_0 dropped unexpectedly at period 10,000, the firm's optimal policy would be to keep the price high, since the optimal policy for high $\tilde{\alpha}_{Ht}$ and low $\tilde{\alpha}_{Lt}$ is high from Figure 6. Intuitively, the firm does not want to lower the price because doing so in response to this cost shock would signal to consumers that the product's quality was low. We have also investigated the effect of a cost shock in the $N_k = 1000$ case. The results are similar to the $N_k = 50$ case. If $\tilde{\alpha}_{Ht}$ is sufficiently high, and $\tilde{\alpha}_{Lt}$ is sufficiently low, the firm will not drop the price in response to a cost shock.





Figure 6: Policy functions for $x_{t-1} = x^H$, $N_k = 50$





Figure 7: Policy functions for $x_{t-1} = x^H$, $N_k = 50$