Altruistic Dynamic Pricing with Customer Regret

Julio J. Rotemberg*

January 19, 2010

Abstract

A model is considered where firms internalize the regret costs that consumers experience when they see an unexpected price change. Regret costs are assumed to be increasing in the size of price changes and this can explain why the size of price increases is less sensitive to inflation than in models with fixed costs of changing prices. The latter predict unrealistically large responses of price changes to inflation for firms that do not frequently reduce their prices. Adjustment costs that depend on the size of price changes also raise the variability on the size of price increases. Lastly, it is argued that the common practice of announcing price increases in advance is much easier to rationalize with regret concerns by consumers than with more standard approaches to price rigidity. JEL: E31, L11, D11

*Harvard Business School, Soldiers Field, Boston, MA 02163, jrotemberg@hbs.edu. I wish to thank the Harvard Business School Division of Research for research support as well as Virgiliu Midrigan and seminar participants at the Kansas City Federal Reserve for comments.

Price changes can trigger consumer regret. Noticing a price increase for a storable good can lead consumers to regret not having purchased earlier while noticing a price decline can lead them to regret not having waited. Prices of goods whose consumption is intertemporally substitutable should give rise to similar regrets. As a final example, an individual who learns inside a store that a price has just increased might regret not having visited an alternate store instead.

While it does not directly absorb scarce physical resources, the regret induced by a price change is a cost, and the paper compares regret costs with the fixed costs of price adjustment postulated in the influential papers of Sheshinski and Weiss (1977) and Golosov and Lucas (2007). In particular, the paper emphasizes the existence of observations that are inconsistent with simple models of fixed costs of changing prices while being consistent with pricing by firms that incorporate consumer regret into their optimization because they act altruistically. The costs induced by regret can be expected to differ from fixed costs in two ways. First, regret is presumably larger when prices are changed by larger amounts, so regret costs depend on the size of price changes. Second, regret costs can be reduced if people are told about future price changes in advance.

The dependence of regret costs on the size of price adjustment makes firms who act as if they cared about these costs less willing to institute large price changes. This matters in two contexts. The first is the effect of inflation on the size of price changes. In the Sheshinki and Weiss (1977) model, an increase in inflation leads firms to post substantially larger price increases whenever they raise their prices. In practice, however, several papers have shown that the actual size of price increases rises only modestly when inflation rises. The lack of dependence of price changes on inflation is visible already in the early work on magazine prices by Checchetti (1986). It is brought to light even more clearly in Kashyap's (1995) study of catalog prices, in Goette, Minsch and Tyran (2005) analysis of Swiss restaurant data and in the Mexican and Norwegian price index studies carried out by Gagnon (2007) and Wulfsberg (2008) respectively.

Gagnon (2007) argues that the lack of substantial changes in the size of price increases when inflation rises does not constitute evidence against fixed costs of changing prices. He shows, in particular, that a model with fixed costs predicts a modest effect of inflation on the size of price increases when firm productivity is random, as in Golosov and Lucas (2007). However, he derives this result under the supposition that all firms adjust their prices with the same frequency and are equally likely to lower their prices whenever they do change them. Neither is true in practice: firms differ greatly in the rigidity of their prices and declines are much more common among firms that change their prices frequently.

This heterogeneity matters for two reasons. First, the size of price increases is more sensitive to inflation for firms that adjust their prices less frequently. Second, this cross-sectional relationship is convex, with a one percent change in annual inflation having only a negligible effect on the individual price changes of firms that adjust their prices continuously. As a result of this convexity, inflation raises the size of the average price increase by more than it raises the price increases of firms whose prices are adjusted at the average frequency.

Relative to fixed costs, regret costs reduce the extent to which firms with infrequent price changes let the size of their price increases respond to inflation. The reason is that the postponement of a price change requires a larger price increase (and hence more regret) when inflation is higher. As a result, increases in inflation make such delays less attractive to altruistic firms, and this dampens the size of their price increases. Interestingly, this effect can be so large that increases in inflation lead firms to reduce the size of their price increases. This result my be of empirical relevance because Wulfsberg (2008) shows that the size of many Norwegian price increases rose when inflation fell after the 1980's.

When firms see larger price changes as more costly, their prices have another desirable property. This is that the size of their price changes is more variable. This result may seem surprising since it might be felt that the desire to avoid regret leads all price changes to be consistently smaller. This is true for any given stochastic environment faced by firms. However, to be consistent both with the average magnitude of price increases and with the proportion of prices changes that are price reductions, one must make the environment faced by firms with regret costs more volatile. What happens, then, is that firms with regret costs institute only small price changes when inflation erodes their price for a given level of their real costs. On the other hand, they are forced to make larger price changes when they are hit by the large changes in real costs that are necessary to account for their behavior.

A somewhat different advantage of interpreting the costs of price adjustment as customer costs of regret is that one can then explain why firms warn customers in advance of price changes. By allowing people to plan for a price increase and even to buy before it occurs, much of this regret should be avoided. Since this paper does not derive regret costs from first principles, it cannot address the question of how much regret costs fall when firms announce their price changes in advance. What the paper does show is that altruistic firms do sometimes benefit from pre-announcements if these reduce regret costs. At the same time, the paper emphasizes that pre-announcements of price increases reduce firm profits in more traditional models of price rigidity. What happens in these models is that firms make more profits at the "new" prices than at the old, so they should not encourage their customers to stockpile goods at the old price.[1]

This paper studies implications of a model whose assumptions are meant to be intuitively appealing, and does not discuss the psychological evidence underlying the consumer attitudes that motivate the analysis. Rotemberg (2008) provides a general discussion of the empirical evidence of regret in purchase situations. Heidhues and Köszegi (2005) and Heidhues and Köszegi (2008) present related models where consumers suffer utility losses when they pay

[1]Encouraging customers to stockpile items in advance of price increases also seems inconsistent with models where firms keep their prices rigid because they lack sufficient information to change them (see, for example Mankiw and Reis (2002) and Woodford (2008)). Once the firm learns its price is no longer optimal, there seems to be little reason to help customers buy more at a non-profit maximizing price.

prices in excess of their "reference prices." Reference prices are purely forward-looking in these models, whereas I suppose consumers experience regret when they face prices that are higher than past prices at which they could have bought. This difference may well matter for the results, although verifying this will require further research because Heidhues and Köszegi (2005, 2008) do not investigate the role of inflation and let costs be drawn from time invariant distributions.[2]

While the current analysis can be interpreted as being a "reduced form," it is motivated by Rotemberg's (2010) model of the transmission of customer emotions such as regret to firm actions. In that model, firms follow consumers' wishes for fear of being found to be insufficiently altruistic. One empirical advantage of this interpretation is that consumers do sometimes lash out against firms, and firms do seem to take actions to avoid this.

The paper proceeds as follows. The next section considers deterministic models to show that the dependence of adjustment costs on the size of price adjustment can explain why inflation has modest effects on the size of price increases. Section 2 turns its attention to a model with stochastic technology inspired by Golosov and Lucas (2007). It contains two parts. The first discusses the extent to which such a model can explain the empirical link between inflation and the size of price increases. The second shows that the dependence of costs of adjustment on the size of price changes can, for a given degree of price rigidity, increase the variability of price increases. Section 3 studies price pre-announcements and Section 4 offers some concluding remarks.

[2]With purely forward-looking reference prices, inflation could lead to continuous readjustment of prices as firms incorporate into prices the rational expectations of inflation held by consumers.

1 The size of price increases with deterministic inflation and technology

1.1 Continuous time

Most of the results presented in this paper are derived from numerical exercises carried out with discrete time models. Nonetheless, it is worth starting with a deterministic continuous time model that is close to Sheshinski and Weiss (1977) because this model is analytically more tractable and can therefore provide intuition for the numerical results that follow.

Let p_t be the nominal price charged by a firm at t and let p_{t-} be the limit of the price charged at time x when x approaches t from below. Consumers are assumed to incur the cost $\ell(p_t, p_{t-})$ whenever p_t is not equal to p_{t-}. Consistent with the discussion in the introduction, these costs are allowed to depend on the size of price changes.

Purchases are assumed to be made continuously with no possibility of storage (this is relaxed below). Consumers obtain the following utility by purchasing a sequence of q_t units of a particular good and of z_t units of a numeraire good

$$\int_0^\infty e^{-rt}\left(\frac{\theta}{\theta-1}q_t^{\frac{\theta-1}{\theta}} + z_t\right)dt - \sum_i e^{-r\hat{t}_i}\ell(p_{\hat{t}_i}, p_{\hat{t}_i-}),$$

where r and θ are parameters and the $\hat{t}_i$ represent the dates where the firm changes its prices. The presence of the numeraire good makes it easier to isolate what occurs in a single market; a more complete model would treat all goods symmetrically. The price of the numeraire good, p_{zt} grows at the instantaneous rate μ while consumers have access to an asset with an instantaneous nominal rate of interest of i. Letting A denote the consumers' assets and $\dot{A}$ their time derivative,

$$\dot{A} = iA - p_t q_t - p_{zt} z_t + I_t$$

where I_t represents non-asset income. It follows that, unless $i = r + \mu$ individuals do not consume z smoothly over time. If this condition is satisfied, by contrast, individuals are indifferent as to when they consume z. They can, for example, reach their maximum utility by setting z_t equal to $\bar{z} - p_t q_t / p_{zt}$ where $\bar{z}$ equals $r[A_0 + \int_0^\infty e^{-i\tau} I_\tau d\tau]/p_{z0}$, the annuity value of lifetime wealth in terms of good z. Utility is thus

$$\int_0^\infty e^{-rt} \left(\frac{\theta}{\theta - 1} q_t^{\frac{\theta-1}{\theta}} - \frac{p_t q_t}{p_{zt}} + \bar{z} \right) dt - \sum_i e^{-r\hat{t}_i} \ell(p_{\hat{t}_i}, p_{\hat{t}_i -}). \tag{1}$$

Consumer optimization then implies that $q_t = (p_t/p_{zt})^{-\theta}$ so that θ is the elasticity of demand. With a mass N of consumers, total demand is given by $Q_t = N(p_t/p_{zt})^{-\theta}$. Given this level of purchases, (1) implies that each individual's instantaneous utility from consuming the non-numeraire good is given by

$$\frac{\theta}{\theta - 1} q_t^{\frac{\theta-1}{\theta}} - \frac{p_t q_t}{p_{zt}} = \frac{1}{\theta - 1} \left(\frac{p_t}{p_{zt}} \right)^{1-\theta}. \tag{2}$$

The function ℓ is assumed to have a positive limit as p_t goes to p_{t-} from above, even though $\ell(x, x) = 0$. The existence of such fixed psychological costs allows one to interpret the rigidity of prices as due exclusively to consumer psychology. If fixed psychological costs are regarded as implausible, the model can be interpreted as one that has both administrative and psychological costs of price changes, with the former being fixed and the latter being variable. Because administrative costs of changing prices are unlikely to depend on the size of the price change, the aspects of the model that hinge on this variability seem most easily interpreted as being due to the psychological forces that I stress.

The firm is assumed to act as if it had an altruism parameter of λ towards its consumers. With a constant marginal cost of production of cp_{zt}, its objective is

$$\int_0^\infty e^{-rt} W(p_t/p_{zt}) dt + N\lambda \sum_i e^{-r\hat{t}_i} \ell(p_{\hat{t}_i}, p_{\hat{t}_i -}) \tag{3}$$

$$\text{where} \qquad W(y) \equiv N \left\{ \frac{\theta - 1 + \lambda}{\theta - 1} y^{1-\theta} - c y^{-\theta} \right\}. \tag{4}$$

To ensure stationarity, let ℓ depend on $(p_t - p_{t-})/p_{t-}$. This implies that the time between price changes, τ, remains constant over time. Each time the firm sets a new price, it chooses the same real price $S = p_t/p_{zt}$ so that its price rises by $100(e^{\mu\tau} - 1)$ percent. The variables S and τ maximize

$$\frac{1}{1-e^{-r\tau}} \int_0^\tau e^{-rt} W(Se^{-\mu t}) dt + \lambda e^{-r\tau} \ell (e^{\mu\tau} - 1).$$

The first order condition of this maximization problem with respect to S is

$$\int_0^\tau W'(Se^{-\mu t}) e^{-(r+\mu)t} dt = 0, \tag{5}$$

so that the present value of the benefit of raising the price slightly has to be equal to zero over the period in which price is kept constant. The first order condition for τ is

$$\frac{e^{-r\tau}}{1-e^{-r\tau}} \left\{ W(Se^{-\mu\tau}) - \frac{r \int_0^\tau e^{-rt} W(Se^{-\mu t}) dt + r\ell\lambda}{1-e^{-r\tau}} - \lambda \ell' \mu s e^{\mu\tau} \right\} = 0. \tag{6}$$

Integrating $\int_0^\tau e^{-rt} W(Se^{-\mu t}) dt$ by parts and using (5), (6) implies that

$$r\lambda\ell - \lambda\mu S(1-e^{-r\tau})\ell' = W(S) - W(Se^{-\mu\tau}). \tag{7}$$

When $\ell' = 0$, this is identical to an equation in Sheshinski and Weiss (1977) which states that the difference between firm welfare at the reset price S and firm welfare at the terminal price $s = Se^{-\mu\tau}$ equals the interest rate times the cost of price adjustment (which is $\lambda\ell$ when $\ell' = 0$). The intuition for this is that the product of the interest rate and the cost of adjustment is the benefit of postponing the adjustment of prices by a small amount of time, and that, at an optimum, this ought to equal the cost of doing so.

Sheshinski and Weiss (1977) prove that, quite generally, increases in inflation raise the size of price adjustments. To see this graphically, Figure 1 depicts the objective function W as a function of price. For inflation rate μ_0, the firm lets inflation drive down the price to s_0 before it resets it to S_O, which is above the price p^* that maximizes W. The vertical distance between the initial and final W equals $\Delta = r\lambda\ell$ when $\ell' = 0$. If inflation were to rise to μ_1 and the firm maintained the initial and terminal prices, the size of its price change would be constant. However, starting at S_0, a higher inflation rate would reduce the time it takes for prices to reach p^*. After this point, W' is positive, and the fact that these points are reached faster implies that they receive a higher weight in (5), the present value of W' that must be zero at an optimal price. Since the limits of integration are unchanged regardless of how long prices stay fixed, the firm now wants to start with a price S_1 higher than S_0. Given the required relation between profits at the starting and ending price, the firm must end with a lower price so that the size of price increases is higher.

Now consider the case where ℓ' is positive. The initial and terminal prices $\{s_0, S_0\}$ remain valid for $\mu = \mu_0$ as long as Δ now equals the left hand side of (7) at this inflation rate. Now, however, an increase in inflation to μ_1 has the additional effect of lowering the left hand side of (7) for a given τ. Thus the difference between the W levels before and after adjustment is lower. This is represented by the dash-dotted line in Figure 1, which yields more frequent adjustments (of smaller size) than the line with short dashes. The intuition for this effect is simple. When inflation is higher, postponing a price increase by a given amount of time raises regret costs ℓ by more (because a postponement by a given amount of time dt requires a larger increase in price). All else equal, an altruistic firm is thus more discouraged from postponing price increases. As a result, increases in inflation have a smaller tendency to increase the size of price adjustments than they do when adjustment costs are fixed.

1.2 Discrete time

This effect of inflation can be quantitatively important. To demonstrate this, I turn to a version of the model where decisions are made once per time period and time periods have discrete length. The variables p_t, q_t, i and p_{zt} continue to represent, respectively, the price and individual consumption of the good under study, the one period interest rate and the price of the numeraire at t. Letting ρ be the rate at which consumers discount payoffs one period into the future, each consumer's lifetime utility at t is

$$\sum_{j=0}^{\infty}\rho^j\left(\frac{\theta}{\theta-1}q_{t+j}^{\frac{\theta-1}{\theta}}+z_{t+j}-\ell(p_{t+j}/p_{t+j-1})\right),$$

while each consumer's assets at t, A_t equal

$$A_t=(1+i)A_{t-1}-p_t(\tilde{q}_t+\hat{q}_{t+1})-p_{zt}z_t,$$

where $\tilde{q}_t$ are the purchases of the good at t for consumption at t and $\hat{q}_{t+1}$ are the purchases of the good at t for use at $t+1$. For the moment, I set $q_t=\tilde{q}_t$ and $\hat{q}_t=0$ so that purchases for inventory are ignored. This is relaxed when I consider pre-announcements below. To avoid a strict preference for zero consumption of z_t at certain points, it must be the case that

$$\rho(1+i)=(1+\mu), \tag{8}$$

where μ is the one period rate of growth of p_{zt}, and I assume this from now on. This condition ensures that consumers are indifferent as to when they consume good z. Consumer demand q_t is then equal to $(p/p_z)^{-\theta}$ once again and single-period utility from having access to this good at price p/p_{zt} equals $(p/p_z)^{1-\theta}/(\theta-1)$. With a constant real marginal cost of production c, a firm which behaves as if it cared λ times as much about consumer utility as about its own profits has the same one-period objective as before. I now write it as

$$N\frac{\theta+\lambda-1}{\theta-1}\left\{W\left(\frac{p_t}{p_{zt}}\right)-L\left(\frac{p_t-p_{t-1}}{p_{t-1}}\right)\right\}, \tag{9}$$

where

$$W\left(\frac{p_t}{p_{zt}}\right)=\left(\frac{p_t}{p_{zt}}\right)^{1-\theta}-\frac{c(\theta-1)}{\theta+\lambda-1}\left(\frac{p_t}{p_{zt}}\right)^{-\theta}\qquad L\equiv\frac{\theta-1}{\theta+\lambda-1}\lambda\ell\left(\frac{p_t-p_{t-1}}{p_{t-1}}\right). \tag{10}$$

I normalize N so that $N(\theta+\lambda-1)/(\theta-1)$ equals one. If the firm keeps its price constant for J periods starting in period 0, it incurs its next adjustment cost in period J. Supposing it raises its real price to S whenever it changes it, the present value of its welfare is

$$\begin{aligned} U &= \frac{\sum_{j=0}^{J-1} \rho^j W(S/(1+\mu)^j)}{1-\rho^J} - \frac{\rho^J L}{1-\rho^J} \\ &= \left\{ \sum_{j=0}^{J-1} \rho^j (1+\mu)^{(\theta-1)j} S^{1-\theta} - \frac{c(\theta-1)}{\theta+\lambda-1} \sum_{j=0}^{J-1} \rho^j (1+\mu)^{\theta j} S^{-\theta} - \rho^J L \right\} \Big/ \Big(1-\rho^J\Big) . \end{aligned} \quad (11)$$

The firm sets S to maximize the sum of the first two terms, which gives

$$S(J) = \frac{\theta c}{\theta+\lambda-1} \frac{\sum_{j=0}^{J-1} \rho^j (1+\mu)^{\theta j}}{\sum_{j=0}^{J-1} \rho^j (1+\mu)^{(\theta-1)j}}.$$

It is convenient at this point to normalize c by setting $\theta c/(\theta+\lambda-1)$ equal to one. This has the advantage that the optimal price equals one in the absence of inflation, and that departures from one are a measure of the effect of inflation on S. Using this normalization and substituting $S(J)$ back into (11) implies that the firm's objective function is

$$U = \frac{1}{\theta} \frac{\left[\sum_{j=0}^{J-1} \rho^j (1+\mu)^{\theta j}\right]^{1-\theta} \left[\sum_{j=0}^{J-1} \rho^j (1+\mu)^{(\theta-1)j}\right]^{\theta} - \rho^J L}{1-\rho^J}. \quad (12)$$

This equation allows one to find the numerical values of J that maximize this objective for given parameter values. I conduct several such experiments for different values of inflation and for different degrees of sensitivity of regret to the size of price increases. The normalizations ensure that, for given θ, λ affects the firm's price only through L so that its main role here is to determine the extent to which the firm cares about the regret cost of its customers. Notice also that in (9), L is in the same units as the one period revenues that the firm derives from one customer $(p_t/p_{zt})^{1-\theta}$. This facilitates the interpretation of this cost.

The two remaining parameters of the model are ρ and θ. In the simulations, these are set to the values used in Nakamura and Steinsson (2007) so that ρ equals .96 at annual rates and θ equals 4. In this section, a period is taken to be a day (so that firms can in principle change their prices daily). The results is that the ρ is $.96^{1/365}$, and similarly single period inflation μ satisfies $(1+\mu) = (1+\pi)^{1/365}$ where π is the annual inflation rate.

Consistent with the idea that both price increases and decreases cause some distress, the firm's perceived cost of price adjustment, L is given by

$$L_t = L_0 I_t + \left| \left(\frac{p_t - p_{t-1}}{p_{t-1}} \right) \right| L_1 = \left(L_0 + \left| (1+\mu)^J - 1 \right| L_1 \right) I_t, \tag{13}$$

where L_0 and L_1 are parameters and I_t is an indicator variable that takes the value of 1 if p_t differs from p_{t-1}. The assumption that this function is fully symmetric is made for simplicity.

Results from simulating this model are reported in Figures 2-3. Each figure contains four specifications for these costs, where these specifications differ both in L_0 and L_1. What the specifications in each figure have in common is the size of price changes at the baseline inflation rate of 2.4%. Thus, the L_0's in each figure can be thought of as having been chosen so that, for the L_1's being considered, each induces the same price changes when inflation is 2.4%. Since price changes are common at this baseline inflation rate, the figures allow one to understand the implications of different L's for the effect of inflation increases on price changes. The baseline price increases used as illustrations in the figures have, in turn, been chosen because they have been observed in empirical studies.

Figure 2 considers specifications where price increases equal 23.5% when inflation equals 2.4%. This specification is inspired by Cecchetti's (1986) study of magazine prices. His data show that, on average, price increases for his sample of magazines equalled 23.5% in the 1960's when inflation averaged 2.4%. Cecchetti (1986) reports that an average of 7 years elapsed between price adjustments and that the size of the adjustments he observed matched closely the aggregate inflation that took place since the last time these prices were adjusted. This suggests that, if there were any price declines at all, they must have been extremely rare. While this does not justify using a deterministic model to try to match his observations, it at least suggests a benefit of trying to explain them with models that, like those in this section, imply the absence of price declines.

As Figure 2 indicates, one can explain the price rigidity of magazines with a fixed cost of price adjustment equal to 35.6 times (daily) revenue.[3] Since the price in question is the newstand price of magazines, and magazines also receive revenue from subscriptions and advertisements, this represents a much smaller fraction of total magazine revenue. Still, it represents a substantial fraction of expenditure on newstand magazines. With a λ equal to one, consumers would have to suffer disappointment costs from price increases that are essentially the same as the monthly price of a magazine, and these psychological costs need to be larger still if λ is lower. One potential explanation for such large disappointment costs is that price increases may lead consumers to regret not having obtained a subscription or not having brought alternate reading material with them. While uncomfortably high, these costs may be more plausible than administrative costs of the same size. As Cecchetti (1986) argues, these are likely to be low for magazines because their price is literally printed anew in each issue.

If one views these costs as fixed administrative costs, an additional problem emerges. This is that the size of prices increases did not rise substantially in the 1970's. The average inflation rate in this period was 7.1%. According to the figure, price increases should thus have risen to equal 35.9% if this model were valid with fixed costs of changing prices. Instead, Cecchetti (1986) shows that price increases rose only to 25.3%. The figure also displays the predicted changes in the size of price increases when costs of price adjustment depend on the size of price increases. The bottom-most line in the figure displays the prediction of setting L_1 equal to the value of 1000, and this parameter actually leads predicted price increases to decline. The line with $L_1 = 800$ predicts a price increase of 25.4% when inflation is 7.1%, so a parameter in this range can explain the behavior of magazine price increases. Still the implied level of regret costs are very sensitive to the size of price adjustments. The particular functional form for regret costs in (13) implies that the elasticity of regret costs with respect

[3]This value does depend on the other parameters of the model. It is lower if the elasticity of demand θ is lower, for example.

to $(1+\mu)^J$ is $L_1((1+\mu)^J/(L_0+L_1((1+\mu)^J-1))$. This expression equals about 4.6 for the line with $L_1 = 800$ when these costs are evaluated at the 2.4% inflation rate.

Figure 3 focuses on somewhat more flexible prices, where price increases equal 9.9% at the baseline inflation rate. This is the size of price increases found by Wulfsberg (2008) in Norwegian CPI data from 1990 to 2004 (when inflation in Norway equaled 2.4% on average). It is also close to the size of price increases reported in Golosov and Lucas (2007). Because prices are adjusted more frequently, these simulations involve a lower fixed cost L_0. In addition, a somewhat lower value of L_1, namely 550, can now rationalize Wulfsberg's (2008) finding that the size of typical price increases was about 2.3% lower in the period 1975-1989 when average inflation in Norway rose to 8.4%. Unfortunately, this implies that the required elasticity of regret costs with respect to $(1+\mu)^J$ evaluated at an inflation of 2.4% equals about 10.9. While it is uncertain how this elasticity can be estimated, 10.9 seems unreasonably high.

2 Stochastic costs of production

As emphasized by Golosov and Lucas (2007), the existence of price declines is inconsistent with models where price changes are due exclusively to aggregate inflation. To explain this, it is reasonable to suppose that costs sometimes fall. Golosov and Lucas (2007) propose to model this as a mean reverting stochastic process for the technology of each individual firm, and this approach has been followed, among others, by Midrigan (2008), Gagnon (2007) and Nakamura and Steinsson (2008). In this section, I use a variant of this model to study two issues. The first is the connection between inflation and the mean size of price increases while the second concerns the variability of price changes.

2.1 The connection between inflation and the size of price increases with fixed costs

Gagnon (2007) shows that a model quite close to Golosov and Lucas (2007) implies that the size of price increases rises only very slightly when inflation rises. Gagnon (2007) verifies this prediction with Mexican data. Unlike the Norwegian data of Wulfsberg (2008), Gagnon's (2007) Mexican data does not show the size of price increases declining with inflation. Still, the effect of inflation is much more modest than is implied by the deterministic model considered in the previous section and Gagnon (2007) rightly points out that the capacity to explain this fact is an impressive accomplishment for the model with random technology.

What this evidence shows is that the changes in the size of price increases are consistent with a model where all firms have the same stochastic technology and fixed costs of changing prices. One difficulty with this approach, however, is that it has been widely recognized that firms differ a great deal in the extent to which their prices are rigid. For firms whose prices are flexible, one does not need a model of price rigidity. One is then left with the question whether the evidence is consistent with a model where *the firms whose price is relatively rigid* have fixed costs of changing prices (and stochastic technology). This section takes a modest step towards answering this question. It shows that, even with stochastic technology, firms that face substantial fixed costs of changing prices behave in a manner that is quite similar to the firms considered in the previous section: their prices always increase and never fall, and the size of their price increases is quite sensitive to inflation.

Following Golosov and Lucas (2007), marginal cost is now assumed to equal c/a_t where a_t is an index of technology that evolves according to

$$\log(a_t) = \delta \log(a_{t-1}) + \epsilon_t^a, \tag{14}$$

and ϵ_t^a is an i.i.d. normal random variable with standard deviation σ_a, while δ is a coefficient smaller than 1. This implies that W, the one period payoff to the firm leaving outside adjustment costs, is now equal to

$$W\left(\frac{p_t}{p_{zt}}, a_t\right) = \left(\frac{p_t}{p_{zt}}\right)^{1-\theta} - \frac{\theta - 1}{\theta a_t}\left(\frac{p_t}{p_{zt}}\right)^{-\theta}, \tag{15}$$

where this payoff has been written so that it incorporates the normalizations $N = (\theta - 1)/(\theta + \lambda - 1)$ and $c = (\theta + \lambda - 1)/\theta$. The firm arrives at t with a pre-existing real price p_{t-1}/p_{zt} and its value function can be written as

$$V\left(\frac{p_{t-1}}{p_{zt}}, a_t\right) = \max_{p_t}\left[W\left(\frac{p_t}{p_{zt}}, a_t\right) - L_t + E_t\rho V\left(\frac{p_t}{p_{zt+1}}, a_{t+1}\right)\right], \tag{16}$$

where the cost L_t is given by (13) so that it equals zero if p_t is set equal to p_{t-1}. This optimization is solved by value function iteration on a grid. To keep the optimization problem manageable, the length of the period is set equal to a month. Adjacent points on the (log) grid for real prices differ by .002, which is the baseline deterministic inflation rate of p_{zt} and corresponds to an annual inflation rate of 2.4%.[4]

Following Nakamura and Steinsson (2008), δ is set equal to .66. The parameters that still need to be calibrated are then σ_a, L_0 and L_1. To consider the case of fixed costs, I abstract from L_1 at first. I then set σ_a and L_0 so that the model reproduces two key statistics. These are the fraction of price changes that are increases, which is used in the Nakamura and Steinsson's (2008) calibration, and the average size of price increases, which is used in the calibration of Golosov and Lucas (2007). The average size of price increases is set at 9.9% once again[5] and the fraction of price changes that are increases is set to the value of .65 found by Nakamura and Steinsson (2008). The resulting values of σ_a and L_0 as well as some additional statistics from this baseline simulation are reported in the first column of Table 1. One additional dimension in which the simulation performs well is that prices are predicted to change in 8.2% of the observations, which is close to the value of 8.4% found

[4]The programs to carry out this optimization were adapted from those used by Nakamura and Steinsson (2008).

[5]Golosov and Lucas (2007) use 9.5%.

by Nakamura and Steinsson (2008) in US CPI data.[6]

The second column reports the effects of raising the annual inflation rate to 10% while all parameters stay at their baseline values. Consistent with the findings of Gagnon (2007), the average size of price increases rises only modestly. Here it rises by about 1 percentage point. Meanwhile, there are more substantial increases in both the overall frequency of price changes and the fraction of these changes that is made up of price increases. This may suggest that the implications of the deterministic model of section 1 are irrelevant.

Column 3 shows, however, that this result depends on the supposition that every price experiences the same number of increases and declines. As suggested earlier, the Cecchetti (1986) magazine price data appears to include few if any price declines. Similarly, price declines appear to be effectively absent in the restaurant data presented in Goette, Minsch and Tyran (2005). Like magazines, these prices are quite rigid with an elapsed time between price changes of around six quarters. In fact, the model with fixed costs of changing prices that I have been studying does predict that, for a given stochastic process for a_t, price declines should essentially disappear if the cost of changing prices is sufficiently high.

This is demonstrated in columns 3 and 4 of Table 1, where I simulate a firm that is subject to the same stochastic process for a as the firms in columns 1 and 2 but whose L_0 is set equal to .125. This is a value that is close to the minimum one that ensures that prices only rise whether yearly inflation equals 2.4 or 10%. Column 3 considers the case where this firm faces a 2.4% inflation rate. While it never lowers its prices, the size of its average price increases is essentially the same as that for firms in column 1, where the adjustment cost is much smaller. As adjustment costs increase from column 1 to column 3, the firm becomes more reluctant to lower prices and this further reduces its incentive to raise prices (because of the fear of being stuck with a price that is too high if a rises).

[6]This match is not entirely surprising since Nakamura and Steinsson (2008) use this overall frequency to calibrate their parameters and I use the δ from their study. The simulation statistics are obtained by constructing a stochastic sample path of 10 million observations.

Column 4 then simulates the actions of firms with $L_0 = .125$ in an environment where p_{zt} grows at 10% per year. This change in inflation raises the size of price increases from 9.9% to 15.9% very much in line with the deterministic results.[7] Thus, Goette, Minsch and Tyran's (2005) evidence that Swiss restaurants (who rarely if ever cut prices) kept the size of their price changes constant when inflation changed also constitutes evidence against this stochastic version of a fixed cost of price adjustment model.[8]

What keeps price increases from rising with inflation when $L_1 = 0$ as in the specifications in columns 1 and 2 appears to be the fact that the firm has the option of eliminating its price reductions when inflation rises. This is shown more generally in Figure 4. This Figure depicts the connection between a firm's adjustment frequency, its fraction of price declines, and the responsiveness of its price increases to inflation. The Figure is constructed by keeping the demand and technology parameters the same and considering firms that differ only in their L_0. Higher values of L_0 lead to a lower frequency of price adjustment, and this frequency (at 2.4% annual inflation) is used as the x-axis for the plots. The bottom plot shows that firms with higher adjustment costs are less likely to cut their prices (again at 2.4% annual inflation.) Since inflation is positive, they prefer letting their real prices erode by doing nothing. The top plot, meanwhile shows that the change in the size of the price increase induced by going for a 2.4 to a 10% annual inflation rate. This shows that firms that cut their prices frequently have price increases that are nearly unaffected by inflation. For firms that raise their prices between 50 and 65% of the time that they institute a price change, the induced rise in the size of their price increases hovers between 0 and 1%.

By contrast, as the proportion of price decreases falls below 35%, firms with fewer price reductions increasingly let the size of their price increases be affected by inflation. The effect

[7]In Figure 3 a change in inflation from 2.4 to 10% leads the size of price increases to go from 9.9% to 16.3%.

[8]The same is true for the evidence in Kashyap (1995). Declines constitute only 8 percent of his sample of price changes and yet he observes no difference between the size of price increases in the 1970's and the size of price increases in the 50's, 60's or 80's.

of inflation is particularly dramatic for firms with very infrequent price adjustments. Given the convexity of this plot, it seems fair to conclude that the average across firms of the increase in the size of price increases should exceed the increase of a firm whose frequency of price adjustment is the average one. To determine the exact prediction of the model one would have to know how many firms fall in each category.

Unfortunately, we do not even know the behavior of firms whose prices are adjusted at the average frequency. In the simulations of the calibrated model, the typical firm is assumed to have the average frequency of price changes as well as the average fraction of price increases. It is worth emphasizing, however, that the overall fraction of price changes that are decreases (35%) is significantly influenced by firms whose prices are flexible, since their prices are more likely to be observed changing. So, it is possible that a very large fraction of firms mostly raise their prices but that they change their prices sufficiently rarely that they contribute a disproportionately small amount to the overall volume of changes. In this case, the model would predict very substantial increases in the average size of price when inflation rises.

2.2 Adjustment costs that rise with the size of the adjustment

This section reintroduces $L_1 > 0$, this time into the model with stochastic a. I start by considering the case where $L_1 = .5$. Because this is a monthly model, the corresponding values for a daily model like that of Section 1 would be around 15, which is still smaller than most values considered in that section. The first column of Table 2 uses the parameters L_0 and σ_a that fulfilled the two calibration criteria in the case where L_1 was set to zero. I start with these parameters because, by allowing for a simple comparison with the case where $L_1 = 0$, they help provide intuition for the effect of L_1

The first column of Table 2 indicates that adding the cost $L_1 = .5$ while keeping all other parameters the same reduces the size of price increases while also reducing the overall frequency of price adjustment. This combination of effects may seem surprising, so Figure 5 provides some intuition.

This figure depicts slices of the policy functions that result from setting L_1 equal to either zero or .5. Both panels of this figure show the price that firms would charge as a function of the price they inherit when $\log(a)$ is equal to .09. Recall that overall mean of $\log(a)$ is zero so these slices involves relatively favorable technology. Two differences are apparent. The first is that the model with $L_1 = .5$ has two different reset prices instead of one. A firm with $L_1 > 0$ that is changing its price because it inherits a price that is too low does not set the same price as a firm that finds itself with a price that is too high. The reason, of course, is that both these firms reduce their costs by making smaller adjustments.

The second difference is that the band of inaction is somewhat larger in the case where $L_1 = .5$, with firms allowing their price to climb higher before they lower it. Particularly when a is temporarily high, so that the desired price is high, small price reductions appear not to be that valuable. These firms thus institute them only when their current price is further from their desired price.

The second of these features implies both that price decreases are less common (because firms wait longer to institute them) and that price increases are less common (because it is less likely that firms will use price increases to offset recent price reductions that were followed by declines in technology). Price increases are also made less common by the fact that price declines, when they occur, are smaller. On the other hand, the fact that price increases are smaller means that a price increase is likely to be followed sooner by the need to raise price again. This last effect, however, appears to be smaller than the other two since the actual frequency of price increases also declines somewhat when $L_1 = .5$.

The net effect of all this is that the simulations in the first column of Table 2 do not satisfy the target criteria: the size of price increases is too small and the fraction of price increases is too large. Keeping σ_a the same, one can fit the size of price increases by raising L_0. This increase in the cost of changing prices leads firms to be even more unwilling to cut prices, however, so that the fraction of price declines diminishes further. As already indicated, this is not really drawback from an empirical point of view since sectors like magazines have both infrequent price declines and large price increases. Still, if one wants all firms to have the same parameters while keeping the fraction of price increases equal to 65%, one must increase the variability of technology. The set of parameters σ_a and L_0 that matches the two target moments is displayed in column 2 of Table 2 and σ_a is now considerably larger.

The need to increase σ_a as L_1 is increased so as to keep the fraction of price increases at 65% limits the possibility of conducting numerical exercises with large values of L_1. The reason is that increases in L_1 now require larger grids of prices, and the size of the resulting grids quickly creates numerical problems. The result is that, for the values of L_1 that I was able to study, the effect of inflation on the size of price increases remains modest. This is shown in column 3 of Table 2, which demonstrates that the size of price increases does not change significantly when inflation is raised to 10% per annum.

While the requirement that σ_a rise to ensure that prices decline when $L_1 > 0$ limits the scope of the analysis, it has a potentially important consequence, namely that it makes the size of price changes more variable. This can be seen by comparing column 2 of Table 2 with column 1 of Table 1, both of which fit the target moments when inflation is 2.4%. The latter, however, exhibits a 15% larger standard deviation of price increases and nearly a doubling in the (admittedly small) proportion of prices increases that exceed 15%.

Among the 12 products considered in Kashyap (1988), a substantial fraction had price increases of less than 3% about 30% of the time while more than 10% of their increases exceeded 15%. As he notes, the dispersion of price increases observed in his data represented a challenge to models of fixed costs of changing prices.[9] If firms sometimes raise their prices by small amounts, and these small amounts represent the size of their bands of inaction, then they should keep their prices constant only when they are subject to minuscule changes in

[9] See also Carlton (1986).

cost. This seems difficult to reconcile with the observation of large price changes, unless costs are quite variable. But, if costs are so variable, why are prices typically constant for such long periods of time. The fixed cost model thus seems inconsistent with long periods of price rigidity that are interrupted by price changes of extremely variable size. Some solutions have been offered, including that costs of adjustment vary randomly over time (Dotsey, King and Wolman (1999)), that firms use a stochastic device to learn when price adjustments might be appropriate (Woodford 2008), that some costs of changing prices are "free" (Midrigan (2008)), or that costs of production are subject to leptokurtic disturbances (Midrigan (2008) and also Gertler and Leahy (2006)). These channels may well be necessary to explain the observations even after the role of L_1 is taken into account. Still, it is interesting that costs of adjustment that depend on the size of the price adjustment can also contribute to the variability of price changes.[10]

This effect is particularly stark in the case of the illustrative parameters considered in column 4 of Table 2. By simultaneously raising L_1, so it equals 1.0, and σ_a, so it equals .19, one obtains significantly more volatile changes in the size of price increases. In particular, the fraction of price increases smaller than 3% is now 13% while that over 15% equals 19%. Prices are now considerably less rigid, since they adjust on average every 3 months. While helpful in raising the variability of prices, this example does not succeed in reproducing the findings of Kashyap (1988). In particular, Kashyap (1988) finds an even larger proportion of small changes and longer durations of constant prices. Moreover, the example in column 4 features many more price declines than are found by Kashyap (1988).

While purely illustrative, the simulations in column 4 can usefully be compared to those

[10]While this paper is concerned with price changes as opposed to with the response of output to nominal disturbances, it is worth noting that these issues are closely linked, particularly in models like Golosov and Lucas (2007). Midrigan (2008), in particular, shows a mechanism through which a higher variability of price changes is connected with a higher response of output to monetary shocks. The idea is that the timing of relatively large price changes is unaffected by monetary policy because these changes are due to idyosincratic cost shocks. Thus, the observation of relatively large price changes suggests that, as in Calvo (1983), the timing of many price changes is insensitive to monetary shocks.

in column 5. Here, L_1 is lowered to zero while L_0 and σ_a are set so that fraction of price increases under 3% and the fraction of price increases above 15% are the same in both columns. The average size of price increases is also the same, while the overall standard deviation of price changes is quite comparable.

The two simulations do differ in one crucial respect, however, and this illustrates an advantage of considering a model with positive L_1: the simulation with $L_1 = 1$ has prices that change much less frequently. The reason for this is the (relative) reluctance of firms to lower prices when $L_1 > 0$. This cuts the frequency of price adjustment directly by reducing the number of price reductions. It also cuts indirectly the number of price increases because price reductions when $L_1 = 0$ require subsequent price increases when a suddenly falls.

3 Preannouncing price increases

Salespeople sometimes warn customers of impending price increases while, at other times, such increases are officially announced in advance. Searching for articles with the terms "price increases," "announced" and "effective" in *Business Wire* for the period 10/02 to 10/04 yielded 44 stories pertaining to companies who made such announcements. Of these, 14 (32%) announced price increases over one month in advance, 25 (57%) announced them less than one month in advance but over 10 days in advance and only 5 announced that these would affect shipments that would take place in the next ten days.

In several instances, firms specifically told their customers they could continue to place orders at the pre-increase price for some time. For example, the September 15, 2004 announcement by GrafTech that it was increasing electrode prices, explicitly stated this price increase would only apply to *orders* received after October 1. More generally, announcements made with a large degree of advance notice, such as Kimberly-Clark's announcement in March 2004 that it would increase its Kleenex prices by midsummer, give customers the capacity to respond by stocking up. The targets of these published announcements are businesses, rather than final consumers. The evidence for pre-announcements of consumer goods price increases is more anecdotal, with Starbucks, for example warning customers 10 days notice before raising its prices in September 2004.[11]

It is possible to imagine scenarios where such pre-announcements have only a minimal effect on quantities purchased while still being desirable to consumers, perhaps because these enjoy feeling that prices are predictable. For many goods, however, such a pre-announcement would also tend to drive up sales at the low price and drive down sales at the high price. This would be particularly true for goods whose purchases are intertemporally substitutable, but might also be true of other goods (if, for example, announced price increases lead to more search than surprise price increases).

In this section, I study the attractiveness of announcing prices in advance in the deterministic discrete time model of subsection 1.2 and in the stochastic model of Section 2. In either case, pre-announcements tend not to be attractive unless they reduce the cost of changing prices as they plausibly do in the case of regret cost. That fixed costs of changing prices can make pre-announcements unprofitable is already suggested by Benabou (1989). He shows that optimization by firms selling storable goods leads them to randomize the time of their price changes so as to dampen speculation, thereby establishing that they would lose from announcing this timing in advance.

To capture the effect of pre-announcements on purchases, I suppose that these lead a fraction α of customers to become aware of an impending price change, while the others remain unaware. As in Section 1.2, the consumption of an aware consumer at t, q_t, is the

[11] A story suggesting that a furniture store warned a customer in advance of a price increase, with the customer buying just before this took effect can be found under the headline "Homo Economicus," in the Washington Post of July 3, 2008. Retailers do, on occasion, respond to customer complaints after a price increase by charging them "the old price." The analysis of this should be similar to what is done here: it seems equally inconsistent with fixed costs of changing prices and consistent with the avoidance of regret costs.

sum $\hat{q}_t + \tilde{q}_t$ where $\tilde{q}_t$ is purchased at t while $\hat{q}_t$ is purchased at $t-1$ for t. Since consumers can borrow at rate i, they prefer buying all their goods for t at $t-1$ if $p_{t-1}(1+i) < p_t$. Otherwise they buy them at t. In the model of section 1, prices changes equal $(1+\mu)^J$ where J is the period of price rigidity. This condition is thus satisfied if $(1+\mu)^J > (1+\mu)/\rho$ where I have used (8) to substitute for $(1+i)$.

If consumers do purchase in advance, their purchases for t, $\hat{q}_t$ maximize

$$\frac{\theta}{\theta-1}\hat{q}_t^{(\theta-1)/\theta} + \bar{z}_t - \frac{\hat{q}_t p_{t-1}(1+i)}{p_{zt}}.$$

Using (8), this implies that aware consumers set $\hat{q}_t$ equal to $(p_{t-1}/\rho p_{zt-1})^{-\theta}$ and their welfare from buying the good is $(p_{t-1}/\rho p_{zt-1})^{1-\theta}/(\theta-1)$.

Ignoring costs of changing prices, a firm acting as if it were altruistic would behave as if its real payoff at $t-1$ from selling to these consumer were equal to

$$\frac{\theta-1+\lambda}{\theta-1}\left(\frac{p_{t-1}}{p_{zt-1}}\right)\left(\frac{p_{t-1}}{\rho p_{zt-1}}\right)^{-\theta} - c\left(\frac{p_{t-1}}{\rho p_{zt-1}}\right)^{\theta}.$$

Normalizing N and c, the firm's benefit from making all its sales for t at $t-1$ equals

$$\rho^{\theta}\left[\left(\frac{p_{t-1}}{p_{zt-1}}\right)^{1-\theta} - \frac{\theta-1}{\theta}\left(\frac{p_{t-1}}{p_{zt-1}}\right)^{-\theta}\right] = \rho^{\theta}W\left(\frac{p_{t-1}}{p_{zt-1}}\right), \tag{17}$$

where the equality is based on (10). If, instead, it sells all its goods for t at time t, the present value of its benefits as of $t-1$ is $\rho W(p_t/p_{zt})$. Since $\rho^{\theta} < \rho$, the expression in (17) is lower when $W(p_t/p_{zt})$ is equal to $W(p_{t-1}/p_{zt-1})$. The reason is that, even if these W's were the same, the firm would sell less in period $t-1$ because consumers have to pay the real interest rate to carry the goods forward in time.

Moreover, in all the simulations I conducted with the model of section 1, the value of W in the period before price adjustment was below that in the period with the new price. Thus, if firms adjust their prices at t, $W(p_{t-1}/p_{zt-1})$ is less than $W(p_t/p_{zt})$. This establishes that, in the model of section 1.2, pre-announcements are not attractive for firms with fixed adjustment costs. The next step is to study whether a firm would be willing to pre-announce if it could thereby save some of its consumer's regret costs.

It seems reasonable to suppose that customers who are able to buy at the earlier price should not experience any regret (and may instead experience additional utility from having obtained better terms than less aware consumers). I thus focus on the case where the pre-announcement eliminates a fraction α of the regret costs (which are incurred when the price changes at t, as before). In this case, pre-announcing is worthwhile, so that no equilibrium without pre-announcements exists, if

$$\rho(W(S) - L(\mu^J - 1)) < \rho^\theta W(S/(1+\mu)^{J-1}) \tag{18}$$

For the cases depicted in Figures 2-3, this conditions is always satisfied. These Figures are constructed by letting periods last one day. This analysis thus shows that firms would be willing to let people buy one day's worth of goods at the old price to eliminate adjustment costs. A one day pre-announcement might have only a modest effect and, perhaps for this reason, pre-announcements tend to involve longer periods.

To lengthen the period, it suffices to change the discount and inflation rates. In a monthly model with L_1 set equal to zero, (18) continues to be satisfied for products whose price changes are equal to either 9.9 or 23.5%. On the other hand, the condition is violated for products that change price every year so that their price change equals 2.4%. The reason is that L_0 is substantially smaller when prices are rigid for only one year. Once L_0 is small, the firm has less to gain by pre-announcing its price increases.

When condition (18) is satisfied, there is no equilibrium without pre-announcements. This means, however, that the equilibrium that does exist satisfies somewhat different equations. When setting its new price, the firm has to recognize both that only a fraction $1 - \alpha$ of its customers buy in the first period and that a fraction α of its customers buy in the last period for consumption one period hence. Thus, the firm's present discounted value of benefits from setting a price of S every J periods becomes

$$U = \frac{D_1 S^{1-\theta} - D_2 S^{-\theta} - \rho^J (1-\alpha) L}{1-\rho^J} \tag{19}$$

where

$$\begin{aligned} D_1 &= \sum_{j=0}^{J-1} \rho^j (1+\mu)^{(\theta-1)j} - \alpha \left(1 - \rho^{(J-1+\theta)} (1+\mu)^{(\theta-1)(J-1)}\right) \\ D_2 &= \frac{\theta-1}{\theta} \left[\sum_{j=0}^{J-1} \rho^j (1+\mu)^{\theta j} - \alpha \left(1 - \rho^{(J-1+\theta)} (1+\mu)^{\theta(J-1)}\right) \right]. \end{aligned}$$

The optimization of U yields a reset price S equal to $\theta D_2/(1-\theta)D_1$ and this can be plugged back into (19) to obtain the optimal J. For small enough values of α, the resulting optimum is very close to the one obtained without pre-announcements, so that (18) continues to hold and there are indeed pre-announcements each time the price is changed.

I now return to the case of stochastic costs and let a_t follows the process in (14). Pre-announcing price increases for $t+1$ at t now has the obvious disadvantage that the firm knows its marginal cost of production at t but does not know it for $t+1$. In spite of this disadvantage, consumer's ability to store the good at t for consumption at $t+1$ can make pre-announcing price increases more attractive when a is variable. This is true, in particular, when a firm expects its future costs to be higher than its current costs. By pre-announcing a price increase this firm induces more customers to buy while its costs are relatively low, and this can be profitable. Notice that this pre-announcement is only attractive if a firm finds itself simultaneously with low costs and a desire to raise prices, a combination which may not manifest itself very frequently. To obtain an estimate of this frequency, I study how often this occurs in simulated settings where, for simplicity, L_1 is set to zero.

If the firm announces a new price for t at $t-1$, let $\hat{p}_t$ represent this price. The value at t of a firm that has made such an announcement is $V^c(\hat{p}_t/p_{zt}, a_t)$. Given that the costs of announcing prices are independent of the price that is set, an optimizing firm sets $\hat{p}_t$ to maximize $E_{t-1}V^c(\hat{p}_t/p_{zt}, a_t)$. It thus sets $\hat{p}_t/p_{zt}$ as a function of only a_{t-1}, which contains all the information the firm has at $t-1$ about future a's. Let $\hat{S}(a_{t-1})$ denote this optimal real price, while $\hat{V}^c(a_{t-1})$ denotes $E_{t-1}V^c(\hat{S}(a_{t-1}), a_t)$. The value at t of a firm that has not pre-announced a price is $V^u(p_{t-1}/p_{zt-1}, a_t)$. For ease of exposition, let p_{t-} denote the price that the firm has inherited at t, where this equals either p_{t-1} or $\hat{p}_t$ depending on whether it announced a price in advance.

At an optimum, the value functions V^u and V^c must satisfy

$$V^u\left(\frac{p_{t-}}{p_{zt}}, a_t\right) = \max{}^*\left(\left\{(1+\alpha\rho^{-\theta})W\left(\frac{p_{t-}}{p_{zt}}, a_t\right) + \rho(\hat{V}^c(a_t) - \hat{L}_t)\right\}, \left\{\max_{p_t}\left(W\left(\frac{p_t}{p_{zt}}, a_t\right) + \rho E_t V^u\left(\frac{p_t}{p_{zt+1}}, a_{t+1}\right) - L_t\right)\right\}\right) \quad (20)$$

$$V^c\left(\frac{p_{t-}}{p_{zt}}, a_t\right) = \max{}^*\left(\left\{(1-\alpha+\alpha\rho^{-\theta})W\left(\frac{p_{t-}}{p_{zt}}, a_t\right) + \rho(\hat{V}^c(a_t) - \hat{L}_t)\right\}, \left\{(1-\alpha)W\left(\frac{p_{t-}}{p_{zt}}, a_t\right) + \rho E_t V^u\left(\frac{p_{t-}}{p_{zt+1}}, a_{t+1}\right)\right\}\right). \quad (21)$$

where $\hat{L}_t$ represents the adjustment costs that result from pre-announcements. In these equations, the max* operator gives the maximum of the two terms in braces except when p_{t-}/p_{zt} is greater $\hat{S}(a_t)/(1+i)$. When it is greater, aware consumers prefer waiting until $t+1$ to buy at the real price $S(a_t)$ rather than buying at p_{t-}. To ensure that the firm gains nothing from pre-announcing in this case, the operator is then set equal to the second element in braces. These equations can be solved by value function iteration. This involves using the existing value of V^c and V^u at each iteration, first to compute $\hat{V}^c$, and then using (20) and (21) to compute the next round of V^c and V^u.

Once this procedure converges, one is left with two indicator functions $I^u(p_{t-1}/p_{zt}, a_t)$ and $I^c(p_{\hat{t}}/p_{zt}, a_t)$. The first equals one when the maximum in (20) is the first expression and equals zero otherwise. Similarly, the second takes the value of one when the maximum in (21) is the first expression and equals zero otherwise. One is also left with the function $F(p_{t-1}/p_{zt}, a_t)$, which gives the real price p_t/p_{zt} that maximizes $W(p_t/p_{zt}, a_t) +$

$\rho E_t V^u(p_t/p_{zt+1}, a_{t+1}) - L_t$. The price p_t in a particular sample is then given by

$$\frac{p_t}{p_{zt}} = I_t^a \hat{S}_t(a_{t-1}) + (1 - I_t^a)\left[I^u\left(\frac{p_{t-1}}{p_{zt}}, a_t\right)\frac{p_{t-1}}{p_{zt}} + \left(1 - I^u\left(\frac{p_{t-1}}{p_{zt}}, a_t\right)\right) F\left(\frac{p_{t-1}}{p_{zt}}, a_t\right)\right],$$

where the indicator I_{t+1}^a, which takes a value of one if a price is pre-announced, is

$$I_{t+1}^a = I_t^a I^c(\hat{S}(a_{t-1}), a_t) + (1 - I_t^a) I^u\left(\frac{p_{t-1}}{p_{zt}}, a_t\right).$$

To simulate paths for p_t, I first compute value and policy functions using the same grid as in the case without pre-announcements. For each sample path, I compute the ratio of the number of times that I_t^a is equal to one (so that there is a pre-announcement) to the number of times that prices is increased. This ratio is displayed for various parameter combinations in Table 3. This table considers two cases. In the first, the present value of adjustment costs is the same whether prices are announced in advance or not so that $\hat{L}_t = \rho L_0$ when there is a pre-announcement. In the second, pre-announcements eliminate the regret of aware consumers so that $\hat{L}_t$ equals $(1 - \alpha)\rho L_0$ when a new price is pre-announced.

The first line shows that for the baseline parameters with $\alpha = .15$, there are pre-announcements with regret costs but not with fixed costs. The second line shows that increasing α to .2 leads firms with fixed costs of adjustment to occasionally make pre-announcements, though firms with regret costs do so more frequently. The third line continues to let $\alpha = .2$ and considers the case where $L_0 = .125$, the value which eliminates price reductions in the analysis of Section 2. This higher L_0 reduces pre-announcements. As lines 4 and 5 show, increases in fixed costs do not always have this effect. When $\sigma_a = .03$, an increase in L_0 from .0528 to .125 raises pre-announcements considerably (from half a percent to 2% of price increases). The variability of a thus affects the dependence of pre-announcements on the size of adjustment costs.

When L_0 is relatively low, increasing the variability of a makes it more likely that a firm with low costs will want to raise its price, and it is attractive to pre-announce such increases. Thus, the variability of a induces pre-announcements with low L_0. By contrast, the variability of a does not induce pre-announcements with high L_0 because firms become unwilling to cut prices and this also makes them reluctant to raise prices when costs are low.

With $L_0 = .125$ and regret costs, pre-announcements become quite frequent when α is increased. Line 6, for example, shows that such firms pre-announce about one third of their price increases if α is set equal to .3. By contrast, α needs to be greater than or equal to .9 to induce any pre-announcements by firms for whom costs are fixed, and the resulting fraction of price increases that are announced in advance is infinitesimal.

4 Conclusions

When firm managers are asked why they keep their prices rigid, their predominant response is that consumers react antagonistically to price changes (Blinder *et al.* (1998), Fabiani *et al.* (2006)). At the same time, most of the formal literature deriving price rigidity from more basic frictions has emphasized administrative menu costs that have no direct connection with the psychological states of consumers. This paper suggests that this may be a mistake.

Administrative menu costs have three implications that seem problematic. These are that the size of the price increases of firms that rarely drop their prices should rise substantially when inflation rises, that the volatility of price increases should be low, and that firms would rarely if ever voluntarily encourage their customers to stock up products by announcing the date of a price increase in advance. These counterfactual predictions are avoided if one interprets costs of adjustment as being due to the regret experienced by consumers as they face price increases. Moreover, the notion that consumers suffer losses from price adjustment fits well with the idea that consumers complain when they observe price increases. These complaints are to be expected if, as in Rotemberg (2010), consumers regard it as unfair when firms fail to act somewhat altruistically towards them.

The model of consumer psychology considered here, and of the transmission of consumers' psychological costs to firms, is still fairly crude. This reflects in part the lack of a consensus on how to model social preferences and how to model emotions that are not directly related to the amounts that people consume. Still, the suggestion that psychological considerations of this sort may help explain empirical pricing practices will hopefully encourage further research on these issues.

5 References

Benabou, Roland, "Optimal Price Dynamics and Speculation with a Storable Good," *Econometrica,* 57, January 1989, 41-80.

Blinder, Alan S., Elie R.D. Canetti, David E. Lebow and Jeremy B. Rudd *Asking about prices: a new approach to understanding price stickiness,* Russell Sage Foundation, New York, 1998.

Calvo, Guillermo A., "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics,* 12, 1983, 383–398.

Carlton, Dennis. W., "The Rigidity of Prices," *American Economic Review,* 76, 1986, 637–658.

Cecchetti, Stephen G., "The Frequency of price Adjustment, A Study of Newsstand Prices of Magazines," *Journal of Econometrics,* 31, 1986, 255-74.

Dotsey, Michael, Robert King and Alexander Wolman, "State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output," *Quarterly Journal of Economics,* 114, 1999, 655–690.

Fabiani, Silvia, Martine Druant, Ignacio Hernando, Claudia Kwapil, Bettina Landau, Claire Loupias, Fernando Martins, Thomas Mathä, Roberto Sabbatini, Harald Stahl and Ad Stockman, "What firms' surveys tell us about price-setting behavior in the Euro area," *International Journal of Central Banking,* 2, September 2006, 3–48.

Gagnon, Etienne, "Price Setting Under Low and High Inflation: Evidence from Mexico," Board of Governors of the Federal Reserve System, International Finance Discussion Papers Number 896, 2007.

Gertler, Mark and John Leahy, "A Phillips Curve with an Ss Foundation," *NBER Working Paper 11971,* 2006.

Goette, Lorenz, Rudolf Minsch and Jean-Robert Tyran, "Micro evidence on the adjustment of sticky-price goods: Its how often, not how much" Mimeo, 2005.

Golosov, Mikhail and Robert E. Lucas Jr., "Menu Costs and Phillips Curves," *Journal of Political Economy,* 115, 2007, 171-199.

Heidhues, Paul and Botond Köszegi, "The Impact of Consumer Loss Aversion on Pricing," mimeo, 2005.

—, "Competition and Price Variation When Consumers Are Loss Averse," *American Economic Review,* 98, September 2008, 1245–68.

Kashyap, Anil K.,"Sticky Prices: New Evidence from Retail Catalogues," *Quarterly Journal of Economics,* 1995, 245-74.

Mankiw, N. Gregory and Ricardo Reis, "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *Quarterly Journal of Economics,* 117, November 2002, 1295–1328.

Midrigan, Virgiliu, "Menu Costs, Multi-Product Firms, and Aggregate Fluctuations," mimeo, New York University, 2008.

Nakamura, Emi and Jon Steinsson, "Five Facts About Prices: A Reevaluation of Menu Cost Models," Mimeo, Columbia University, 2008

Rotemberg, Julio J., "Behavioral Aspects of Price Setting, and Their Policy Implications," *NBER Working Paper 13754*, January 2008.

—, "Fair Pricing," *Journal of the European Economics Association*, 2010 (forthcoming).

Sheshinski, Eytan and Yoram Weiss, "Inflation and Costs of Price Adjustment," *Review of Economic Studies,* 44, 1977, 287-303.

Woodford, Micahel, "Information-Constrained State-Dependent Pricing," *National Bureau of Economic Research Working Paper 14620,* December 2008.

Wulfsberg, Fredrik, "Price Adjustments and InflationEvidence from Consumer Price Data in Norway 19752004," mimeo 2008.

Table 1
Inflation and fixed costs of price adjustment with stochastic costs of production

	1	2	3	4
Inflation rate (% annual)	2.4	10.0	2.4	10.0
L_0 (Fixed adjustment cost)	.0285	.0285	.125	.125
σ_a (S.D. of shocks)	.0528	.0528	.0528	.0528
Overall adjustment frequency (%)	8.2	10.7	2.0	5.0
Fraction of increases	.65	.85	1.0	1.0
Mean price increase (%)	9.9	11.0	9.9	15.9
S.D. of price increases	.027	.024	.016	.027
Fraction increases $< 3\%$	0	0	0	0
Fraction increases $>15\%$	.04	.06	.002	.60

Table 2
Variable costs of price adjustment with stochastic costs of production: the effect of varying the parameters

	1	2	3	4	5
Inflation rate (% annual)	2.4	2.4	10.0	2.4	2.4
L_0 (Fixed adjustment cost)	.0285	.0390	.0390	.0017	$6.2e^{-6}$
L_1 (Slope of adjustment cost)	.5	.5	.5	1.0	0
σ_a (S.D. of shocks)	.0528	.097	.097	.19	.104
Overall adjustment frequency(%)	3.5	8.7	10.6	33.7	94.0
Fraction of increases	.988	.65	.85	.61	.49
Mean price increase (%)	5.8	9.9	10.6	9.5	9.5
S.D. of price increases	.013	.031	.032	.072	.070
Fraction of increases $< 3\%$	0	0	0	.13	.13
Fraction of increases $>15\%$	0	.07	.12	.19	.19

Table 3
Proportion of price increases that are pre-announced

	Parameters			Proportion	
	α	L_0	σ_a	with fixed cost	with regret cost
1	0.15	0.0285	0.0528	0	0.005
2	0.2	0.0285	0.0528	0.0037	0.02
3	0.2	0.125	0.0528	0	0.01
4	0.2	0.0285	0.03	0	0.005
5	0.2	0.125	0.03	0	0.83
6	0.3	0.125	0.0528	0	0.32
7	0.6	0.125	0.0528	0	0.91
8	0.9	0.125	0.0528	0.0001	*

Note: * denotes that convergence was not achieved in value function iteration

Figure 1: Continuous time model of fixed adjustment cost

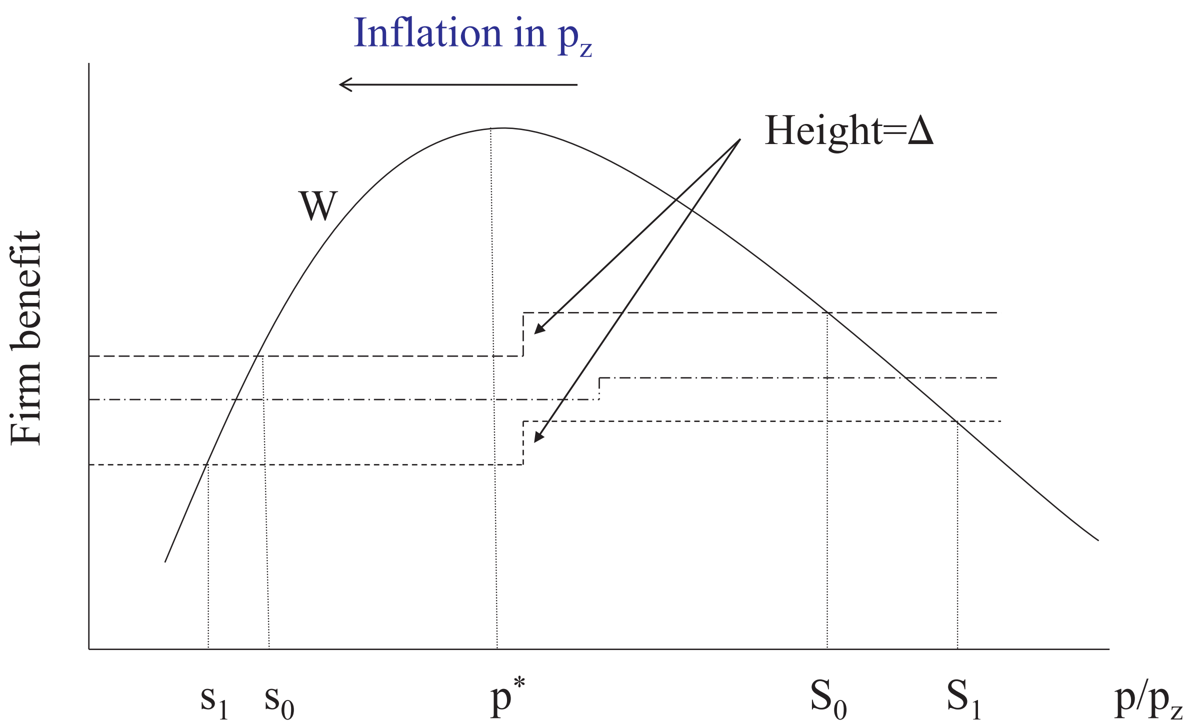

Figure 2: Size of price increases as one varies inflation and L_1. The case where prices are raised 23.5% under 2.4% inflation

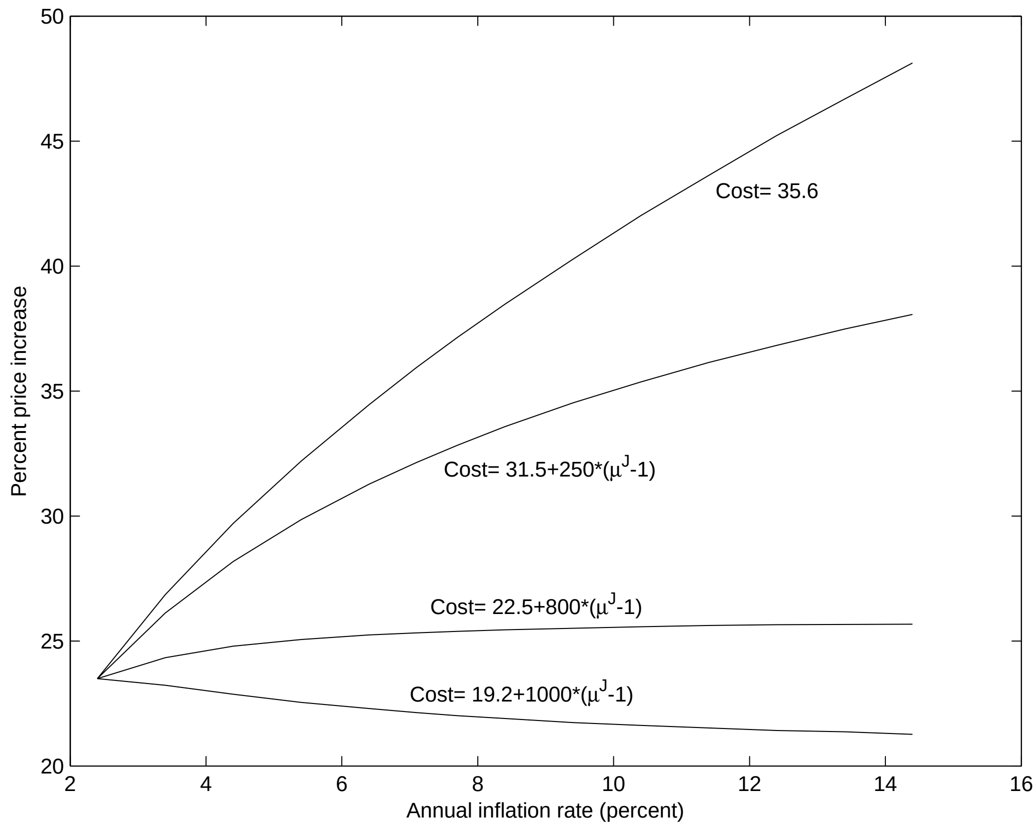

Figure 3: Size of price increases as one varies inflation and L_1. The case where prices are raised 9.9% under 2.4% inflation

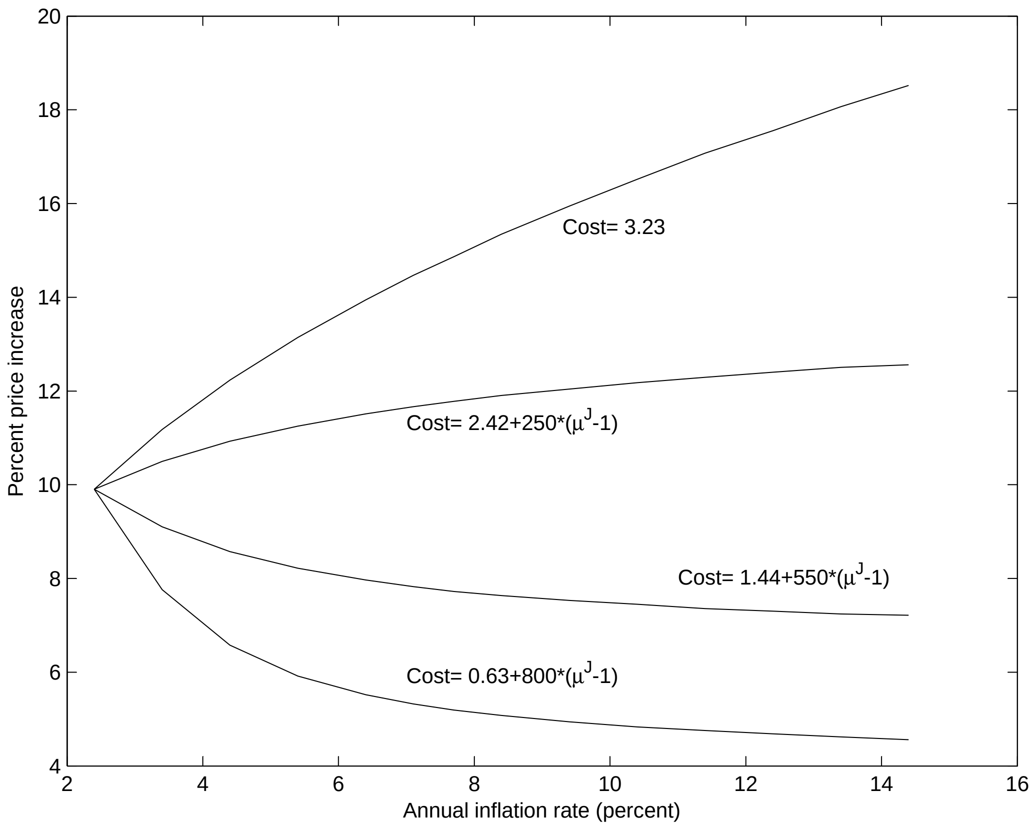

Figure 4: The effect of raising inflation from 2.4 to 10% on the size of price increases for firms that differ in their frequency of adjustment at 2.4% inflation

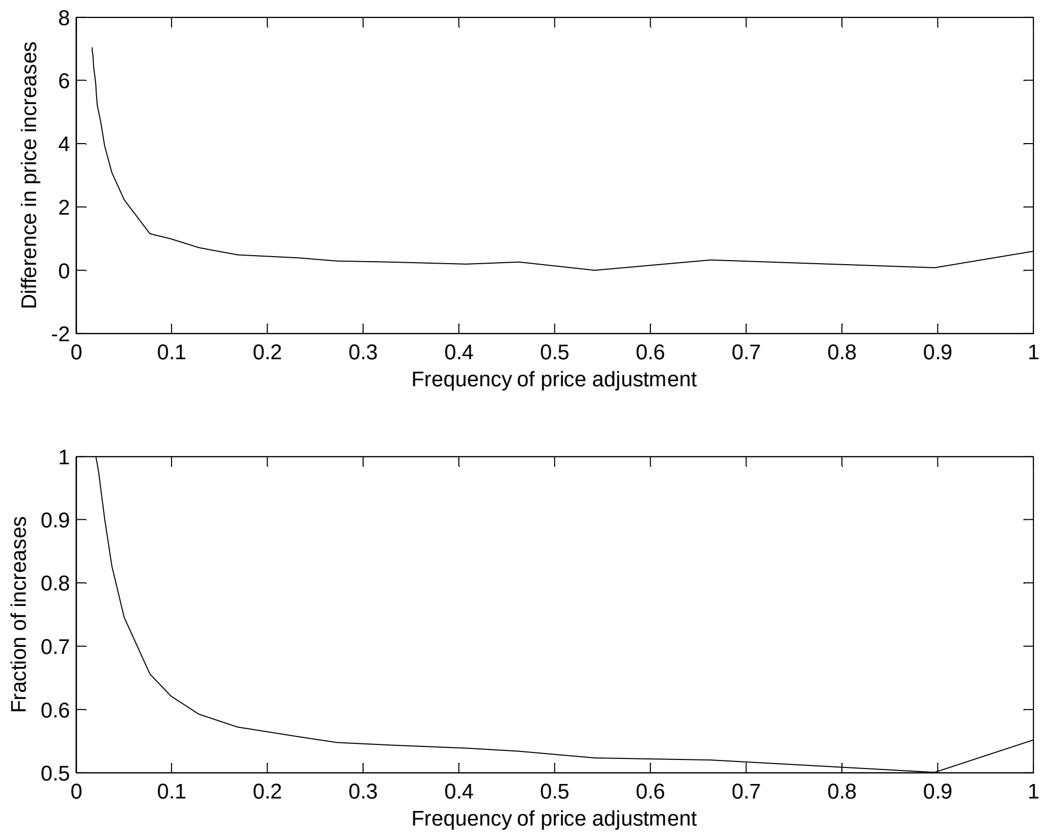

Figure 5: Policy functions with $L_1 = 0$ and $L_1 = .5$