Equilibrium Price Dispersion and Rigidity: A New Monetarist Approach*

Allen Head
Queen's University

Lucy Qian Liu
I.M.F.

Guido Menzio
University of Pennsylvania
NBER

Randall Wright
University of Wisconsin-Madison
Federal Reserve Bank of Minneapolis

September 2010

Abstract

Why do some sellers set prices in nominal terms that do not respond to changes in the aggregate price level? In many models, prices are sticky by assumption. Here it is a result. We use search theory, with two consequences: prices are set in dollars since money is the medium of exchange; and equilibrium implies a nondegenerate price distribution. When money increases, some sellers keep prices constant, earning less per unit but making it up on volume, so profit is unaffected. The model is consistent with the micro data. But, in contrast with other sticky-price models, money is neutral.

*A previous version of the paper was titled "Really, Really Rational Inattention: Or How I Learned to Stop Worrying and Love Sticky Prices." For comments and conversations we are grateful to Marios Angeletos, Ken Burdett, Jeff Campbell, Ben Eden, Christian Hellwig, Boyan Jovanovic, Peter Klenow, John Leahy, Dale Mortensen and Julio Rotemberg. We also thank participants at the MIT macro seminar, the Search and Matching conference (University of Pennsylvania, September 2009), the Economics and Music conference (Essex University, May 2010), the NBER Summer Institute (Cambridge, July 2010), and at the Minnesota Macroeconomic Workshop (Minneapolis, July 2010). Wright thanks the National Science Foundation and the Ray Zemon Chair in Liquid Assets at theWisconsin School of Buisiness. Head thanks the Social Sciences and Humanities Research Council of Canada. The usual disclaimer applies.

1 Introduction

Arguably one of the most difficult questions in macroeconomics is this: Why do some sellers do not adjust their nominal prices in response to changes in the aggregate price level, ostensibly flying in the face of fundamental microeconomic principles? Many popular macro models, including those used by most policy makers, argue that price stickiness is caused by the existence of technological restrictions to nominal price adjustment (Calvo fairies or Mankiw costs). Here, we argue that price stickiness is caused by the existence of search frictions in the product market. Moreover, in contrast with theories of price stickiness that rely on technological constraints, our theory predicts that money is neutral: the central bank cannot engineer a boom or end a slump simply by issuing currency. Hence, while our theory provides a microfoundation for the key ingredient at the core of Keynesian economics (i.e. price stickiness), it has very different policy implications.

We also show that our theory of price stickiness is more successful than existing theories at matching the salient empirical features of the dynamics of individual prices.[1] For example, our theory can account for the average duration of prices that we observe in the data. Moreover, our theory can account for the fact that on average price changes are large and yet many price changes are small. Also, our theory can account for the fact that prices change more frequently, and not just by larger amounts, when inflation is higher. In contrast, menu cost theories of price stickiness cannot easily account for the second empirical fact mentioned above. And Calvo theories of price stickiness cannot easily account for the third empirical fact.

These findings seem relevant for several reasons. First, despite the successes of the New Classical and Real Business Cycle paradigms, it is hard to deny that at least some nominal prices seem sticky in the sense defined above – they do not respond to changes in the aggregate price level. Moreover, this observation is one of main reasons why many Keynesians are Keynesian. Consider Ball and Mankiw (1994), who we consider representative, when

[1] A sample of empirical work on price stickiness includes Cecchetti (1985), Carlson (1986), Bils and Klenow (2005), Campbell and Eden (2007), Klenow and Kryvtsov (2008), Nakamura and Steinsson (2009), Eichenbaum, Jamovich and Rebelo (2009) and Gagnon (2009). See Klenow and Malin (2010) for a recent survey.

they say: “We believe that sticky prices provide the most natural explanation of monetary nonneutrality since so many prices are, in fact, sticky.” They go on to claim that “based on microeconomic evidence, we believe that sluggish price adjustment is the best explanation for monetary nonneutrality.” Furthermore, “As a matter of logic, nominal stickiness requires a cost of nominal adjustment.” Some people that one might not think of as Keynesian present similar positions. For instance, Golosov and Lucas (2003) assert that “menu costs are really there: The fact that many individual goods prices remain fixed for weeks or months in the face of continuously changing demand and supply conditions testifies conclusively to the existence of a fixed cost of repricing.”

We interpret the above claims as containing three points related, respectively, to empirics, theory, and policy. The first claim is that price stickiness is in the data. The quotations assert this, and it is substantiated in the numerous empirical studies mentioned above. We concede this point. The second claim is that price stickiness implies “as a matter of logic” the existence of some technological constraint to price adjustment. We prove this wrong: we describe equilibria that match not only the broad observation of price stickiness, but also some of the more detailed empirical findings, with recourse to no technological constraints. The third claim, to which at least Ball and Mankiw seem to subscribe, is that price stickiness implies that money is not neutral, and this rationalizes Keynesian policy advice. We also prove this wrong: our theory is consistent with the relevant observations, but rather than yielding Keynesian policy implications, money is neutral. Sticky prices simply do not constitute evidence that money is not neutral or that particular policy recommendations are warranted.[2]

To explain our approach, we begin by pointing out that the issues at hand concern monetary phenomena: Why are prices quoted in dollars in the first place? Why do they not all adjust to changes in the money supply? What does this imply about monetary policy? To answer these questions, it seems natural to use a monetary model. We choose to work with

[2] To be clear, our position is not that money is neutral is the real world, and we know of compelling examples (like the one in Lucas, 1972) where it is not neutral in theory. The point is to construct a coherent economic environment with two properties: (i) it is consistent with the sticky-price facts; and (ii) it nevertheless delivers neutrality. Money is not superneutral in the model – inflation matters, even if the money stock does not, as in true in many models – but this hardly rationalizes Keynesian policy prescriptions.

a version of the New Monetarist framework laid out in Lagos and Wright (2005), Nosal and Rocheteau (2010) or Williamson and Wright (2010a,b). In this framework, specialization and search frictions limit barter, while commitment and information frictions limit credit. Hence, money is essential for at least some exchange, and naturally prices are posted in dollars, because dollars are the objects being traded. To please those with a different taste for microfoundatrions, we could have worked with a cash-in-advance or a money-in-the-utility-function models.

Now consider the following observation. In many New Keynesian models, such as Clarida, Gali and Gertler (1999) or Woodford (2003), nominal price rigidities generate price dispersion only if there is inflation. To understand this point, suppose that two sellers set the same p_t at date t. At date $s_1 > t$, the first seller is visited by the Calvo fairy and sets its price to p_{s_1}. At date $s_2 > s_1$, the second seller is visited by the Calvo fairy and sets its price to p_{s_2}. Between dates s_1 and s_2, the prices of the two sellers are different if and only if inflation is strictly positive. But the data suggest that there is dispersion in prices even during episodes of low or zero inflation (see, e.g. Campbell and Eden 2007). For this reason, it seems important to work with a model that delivers price dispersion even with low or no inflation. Although there are many market structures that are known to generate price dispersion without inflation (e.g. Varian 1980 and Stahl 1989), we adopt the structure of Burdett and Judd (1983).

In order to understand Burdett and Judd's theory of equilibrium price dispersion, it helps to review the history of search theory. The earliest search models of McCall (1970) and Mortensen (1970) were partial equilibrium models, in the sense that they characterized the optimal search strategy of the buyers taking as given the distribution of prices posted by different firms. Diamond (1971) was the first general equilibrium search model, in the sense that the price distribution was derived from the profit maximization problem of the firms. Diamond proved that, as long as buyers can only sample one seller at a time, the equilibrium price distribution is degenerate. In particular, every seller finds it optimal to ask the pure monopoly price. Since buyers have no reason to search without price dispersion, Diamond's result was seen by many as the end of search theory. However, Burdett and Judd (1983) came to the rescue. They proved that the equilibrium price distribution in a product market with search frictions is non-degenerate as long as the probability that a buyer meets two (or more) sellers and the probability that he meets only one seller are both strictly positive. When the probability that a buyer meets two or more sellers approaches one, the distribution converges to the competitive price. When the probability that a buyer meets two or more sellers approaches zero, the distribution converges to the pure monopoly price. Intuitively, if buyers may meet one or multiple sellers with positive probability, the pricing problem of the firm is analogue to a first-price procurement auction in which the number of bidders is uncertain. In this kind of auction, the optimal bidding strategy involves mixing.

In this paper, we embed Burdett and Judd (1983) into a dynamic general equilibrium monetary model a la Lagos and Wright (2005). In particular, we assume that agents alternate between trading in a decentralized product market that operates as in Burdett and Judd, and a centralized market that operates as in Arrow-Debreu. In the decentralized market, buyers and sellers need to use money as a medium of exchange because search frictions preclude the use of barter and information frictions preclude the use of credit. In the centralized market, buyers and sellers can use either money or credit. Technically, our model is a version of Lagos and Wright (2005) in which prices in the decentralized market are not determined as the outcome of a bargaining game between buyers and sellers, but they are posted by the sellers.[3]

In equilibrium, identical sellers choose to post different prices in the decentralized market. In particular, at any date t, the distribution of prices in the decentralized market is given by a continuous cumulative distribution function F_t with support $[\underline{p}_t, \overline{p}_t]$, with $\underline{p}_t < \overline{p}_t$. While the equilibrium pins down the price distribution, it does not pin down the price of an individual seller. Indeed, every individual seller is indifferent between choosing any price between the lower and the upper bound of the support of the distribution. If the seller posts a low price,

[3] The framework is flexible in this regard, and several mechanisms other than Nash bargaining, have been studied. For instance, Rocheteau and Wright (2005) and Berentsen, Menzio and Wright (2010) introduce price posting with directed search and (Walrasian) price taking; Aruoba, Berentsen and Waller (2007) consider alternative bargaining solutions; Galenianos and Kircher (2008) use auctions. No one previously tried posting with random matching in Lagos-Wright, although it is used in the related model of Shi (1997) by Head and Kumar (2005) and Head, Kumar and Lapham (2010).

he earns a small profit per unit of output, but he sells many units of output. If the seller posts a high price, he earns a large profit per unit of output, but he sells fewer units. When the money supply increases from M_t to M_{t+1}, the equilibrium price distribution shifts to F_{t+1} with support $[\underline{p}_{t+1}, \overline{p}_{t+1}]$. This requires that some sellers change their prices, but not all of them. In particular, if a seller's price at date t, p_t, does not belong to the support of the price distribution at date $t+1$, $[\underline{p}_{t+1}, \overline{p}_{t+1}]$, then the seller must adjust its price. But, if p_t belongs to $[\underline{p}_{t+1}, \overline{p}_{t+1}]$, the seller may leave his nominal price unchanged. Hence, in our model, sellers can change price infrequently in the face of continuous movements in the aggregate price level, even though they are allowed to change them whenever they like and at no cost.

Our theory says that sellers can be rationally inattentive to the aggregate price level and monetary policy, within some range, since as long as $p_t \in [\underline{p}_t, \overline{p}_t]$, their place in this distribution does not matter. We show that a calibrated version of our theory can match quite well the empirical behavior of prices in the US retail market. First, our theory predicts an average price duration (11.6 months) that is close to the one observed in the data (8.6 months). Second, our theory generates a price change distribution that has the same shape and the same features of the empirical price change distribution (i.e. the average magnitude of price changes is large, there are many small price changes, and there are many negative price changes). Third, as it is observed in the data, in our model the probability and magnitude of price adjustments are approximately independent of the age of a price. Fourth, the theory correctly predicts that inflation increases both the frequency and the magnitude of price changes. Overall, our theory of price stickiness appears to be empirically reasonable. But, again, according to our theory, money is neutral.[4] If the central bank issues extra money at date t, the price distribution responds fully and immediately, and no real expansion is created. Hence, nominal price stickiness does not logically justify interventionalist policies

[4] There are several other interesting models where, despite price stickiness, money may be (sometimes approximately) neutral. These include Caplin and Spulber (1987), Eden (1994), and Golosov and Lucas (2007). Our approach differs in a number of respects. First, we start with a general equilibrium model where money is essential as a mediun of exchange. Second, money by design is exactly neutral, although not superneutral, in our environment. Third, stickiness arises entirely endogenously and robustly – it does not depend on particular functional forms, timing, the money supply process, etc. Fourth, the distribtion of prices is endogenous and derived from standard microfoundations (Burdett-Judd), instead of simply assuming, say, prices are distributed uniformly on some interval.

by central banks.

2 The Model

2.1 Preferences, technology and markets

Time is discrete and continues forever. In every period, two markets open sequentially. The first is a decentralized product market in which buyers and sellers come together through a frictional matching process. In this market, barter is not feasible because buyers do not carry goods that are valued by the sellers, and credit is not feasible because buyers are anonymous. Instead, exchange takes place with fiat money, which is supplied by the government. We will refer to this decentralized and anonymous market as the *Burdett-Judd market* (BJ market). The second market is a centralized product and labor market in which buyers and sellers are recognizable. In this market, exchange may take place using either money or credit. We will refer to this market as the *Arrow-Debreu market* (AD market).

The economy is populated by a continuum of households with measure 1. Each household has preferences described by the utility function

$$\sum\nolimits_{t=0}^{\infty} \beta^t [u(q_t) + v(x_t) - h_t], \tag{2.1}$$

where $\beta \in (0, 1)$ is the discount factor, u is a strictly increasing and strictly concave function defined over the consumption of the good traded in the BJ market (the BJ good), v is a strictly increasing and strictly concave function defined over the consumption of the good traded in the AD market (the AD good), and $-h$ is the disutility of working h hours. For the sake of concreteness, we assume $u(q) = q^{1-\gamma}/(1-\gamma)$ with $0 < \gamma < 1$.

The economy is also populated by a continuum of firms with measure $s > 0$. Each firm operates a technology that turns h hours of labor into $f(h)$ units of the BJ good, as well as a technology that turns h hours of labor into $g(h)$ units of the AD good. For concreteness, we assume that $f(h) = h/c$ and that $g(h) = h$, where $c > 0$ is the cost of producing a unit of the BJ good relative to the cost of producing the AD good. Firms are owned by the households through a balanced mutual fund.

In the BJ market, each firm posts a nominal price p, taking as given the amount of money, m_t, brought into the market by each household and the cumulative distribution function, $F_t(p)$, of nominal prices posted by all the other firms. Following Burdett and Judd (1983), we assume that each household observes the entire price distribution, but can only purchase the good from a random sample of firms. In particular, the household cannot purchase the good from any firm with probability $\alpha_0 \in [0, 1)$. The household can purchase the good from exactly one firm with probability $\alpha_1 \in (0, 1 - \alpha_0)$. And with probability $\alpha_2 = 1 - \alpha_0 - \alpha_1$, the household can purchase the good from two firms. As mentioned above, all transactions in the BJ market are carried out with money.

In the AD market, the government prints money and injects it into the economy through a lump-sum transfer to the households, T_t. Hence, $T_t = (\mu - 1)M_t$, where M_t is the quantity of money at the beginning of the period and $\mu > 1/\beta$ is the money growth rate. For the sake of simplicity, we assume that μ is constant over time. Each household receives the transfer from the government, T_t, and a nominal dividend payment from the mutual fund that owns the firms, D_t. Then, each household chooses how much to work, h_t, how much of the AD good to consume, x_t, and how much money to carry into the next BJ market, m_{t+1}, taking as given the wage and the price of the AD good. Similarly, each firm chooses how much labor to hire and how much of the AD good to produce, taking as given the wage and the price of the AD good. As mentioned above, exchange in the AD market may take place with either credit or money.

2.2 The problem of the household

First, consider a household who enters the AD market with m_t units of money. The lifetime utility of this household is given by

$$\begin{aligned} W_t(m_t) = & \max_{h_t, x_t, m_{t+1}} v(x_t) - h_t + \beta U_{t+1}(m_{t+1}), \\ \text{s.t.} \quad & w_t x_t + m_{t+1} \leq w_t h_t + m_t + D_t + T_t, \\ & x_t \geq 0,\ m_{t+1} \geq 0. \end{aligned} \tag{2.2}$$

The expression above is intuitive. The household chooses how much to work, h_t, how much to consume, x_t, and how much money to hold, m_{t+1}, so as to maximize the sum of its utility in the current AD market, $v(x_t) - h_t$, and its lifetime utility at the beginning of the next BJ market, $\beta U_{t+1}(m_{t+1})$. The household's choice of (h_t, x_t, m_{t+1}) must be affordable given its non-labor income, $m_t + T_t + D_t$, and the wage, w_t.

The household's optimal choices $(h_t^*, x_t^*, m_{t+1}^*)$ satisfy the conditions

$$\begin{aligned} h_t^* = x_t^* + w_t^{-1}\left(m_{t+1}^* - m_t - D_t - T_t\right), \\ v'(x_t^*) = 1, \quad \beta U_{t+1}(m_{t+1}^*) = w_t^{-1}. \end{aligned} \tag{2.3}$$

These optimality conditions are easy to interpret. The optimal choice for m_{t+1} is such that the disutility from earning an additional unit of money in the AD market, $1/w_t$, equates the utility from carrying an additional unit of money into the next BJ market, $\beta U'_{t+1}(m_{t+1})$. The optimal choice for x_t is such that the disutility from working an additional hour, 1, equates the utility from consuming an additional unit of the AD good, $v'(x_t)$. The optimal choice for h_t is such that the household's budget constraint holds with equality.

As in Lagos and Wright (2005), the optimal choices for m_{t+1} and x_t are independent of m_t, and the optimal choice for h_t is a linearly decreasing function of m_t with slope $-1/w_t$. That is, the amount of money with which the household enters the AD market does not affect the household's decision of how much to consume and how much money to carry to the next BJ market. The amount of money with which the household enters the AD market only affects the household's decision of how much to work and it does so linearly. Given the properties of the optimal choices $(h_t^*, x_t^*, m_{t+1}^*)$, it is immediate to verify that the household's lifetime utility W_t is a linear function of m_t with slope $1/w_t$. Notice that the independence of m_{t+1} from m_t greatly simplifies the analysis of the equilibrium as it guarantees that the distribution of money holdings across households in the BJ market is degenerate. Also, notice that the linearity of W_t with respect to m_t simplifies the analysis of the equilibrium because it guarantees that the distribution of money holdings across households in the AD market does not affect their valuation of the dividends paid out by the firms.

Next, consider a household who enters the BJ market with m_t units of money and purchases the good at the price p. The lifetime utility of this household is given by

$$\begin{aligned} V_t(m_t, p) = \quad & \max_{q_t} \frac{q_t^{1-\gamma}}{1-\gamma} + W_t\left(m_t - pq_t\right), \\ \text{s.t.} \quad & 0 \leq pq_t \leq m_t. \end{aligned} \tag{2.4}$$

The expression above is intuitive. The household chooses how much of the good to purchase, q_t, so as to maximize the sum of its utility in the current BJ market, $q_t^{1-\gamma}/(1-\gamma)$, and its lifetime utility from entering the next AD market with $m_t - pq_t$ units of money, $W_t(m_t - pq_t) = W_t(0) + (m_t - pq_t)/w_t$. The choice of q_t is subject to the cash constraint $pq_t \leq m_t$ because money is the only medium of exchange in the BJ market.

The household's optimal choice q_t^* is given by

$$q_t^*(p) = \begin{cases} m_t/p, & \text{if } p \leq \hat{p}_t, \\ (w_t/p)^{\frac{1}{\gamma}}, & \text{if } p > \hat{p}_t, \end{cases} \tag{2.5}$$

where $\hat{p}_t = w_t(w_t/m_t)^{\frac{\gamma}{1-\gamma}}$. If the price p is smaller than $\hat{p}_t$, the cash constraint is binding and the household purchases m_t/p units of the good. Otherwise, the cash constraint does not bind and the household purchases $(w_t/p)^{1/\gamma}$ units of the good. Figure I illustrates the household's demand for the BJ good. The black curve is the household's demand given that the cash constraint is binding. The grey curve is the household's demand given that the cash constraint is lax. The household's demand, q_t^*, is the lower envelope of the black and the gray curves.

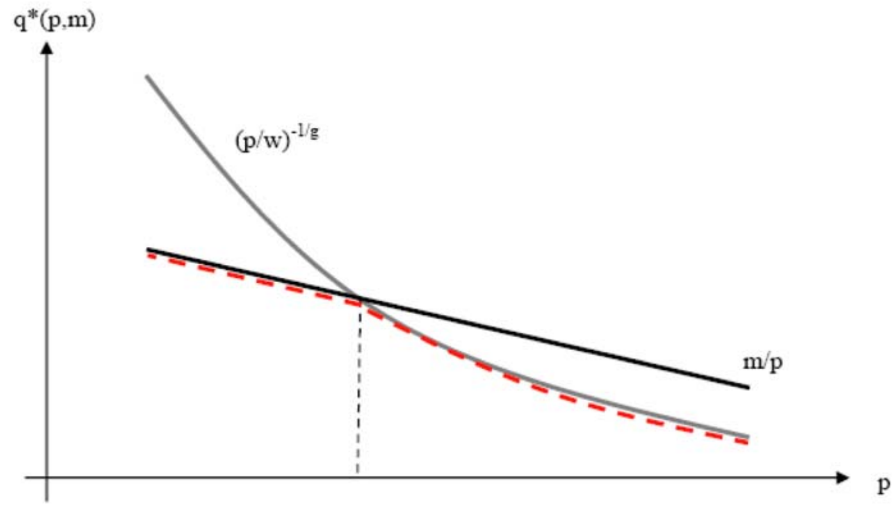

Figure I: Household's demand for the BJ good

Finally, the lifetime utility of a household who enters the BJ market with m_t units of money is given by

$$\begin{aligned} U_t(m_t) = \quad & \alpha_0 W_t(m_t) + \alpha_1 \int V_t(p, m_t) dF_t(p) \\ & + \alpha_2 \int V_t(p, m_t) d\left[1 - (1 - F_t(p))^2\right]. \end{aligned} \tag{2.6}$$

The expression above is easy to understand. With probability α_0, the household cannot purchase the BJ good and enters the next AD market with m_t units of money. With probability α_1, the household can purchase the good from exactly one firm. The price p charged by the firm is a random variable with distribution $F_t(p)$. With probability α_2, the household can purchase the good from two firms. The lowest price p charged by the two firms is a random variable with distribution $1 - [1 - F_t(p)]^2$.

After substituting (2.5) into (2.6) and differentiating U_t with respect to m_t, we can rewrite the optimality condition for m_{t+1} as

$$\begin{aligned} i = \quad & \alpha_1 \int_{\mathbf{1}[p \leq \hat{p}_t]} \left[\frac{w_t}{p} \left(\frac{m_t}{p} \right)^{-\gamma} - 1 \right] dF_t(p) \\ & + 2\alpha_2 \int_{\mathbf{1}[p \leq \hat{p}_t]} \left[\frac{w_t}{p} \left(\frac{m_t}{p} \right)^{-\gamma} - 1 \right] \left[1 - F_t(p)\right] dF_t(p), \end{aligned} \tag{2.7}$$

where i is the net nominal interest rate $\beta^{-1}(w_{t+1}/w_t) - 1$. The expression above is intuitive. The left-hand side of (2.7) is the marginal cost of carrying money, i.e. the net nominal interest rate. The right-hand side of (2.7) is the marginal benefit of carrying money, i.e. the value of relaxing the cash constraint in the BJ market. The optimal choice for money holdings, m_{t+1}^*, equates the marginal cost and the marginal benefit of carrying money.

2.3 The problem of the firm

If a firm posts the price p in the BJ market, its profits are given by

$$\Pi_t(p) = \frac{1}{s}\left[\alpha_1 + 2\alpha_2(1 - F_t(p)) + \alpha_2 n_t(p)\right] R_t(p), \tag{2.8}$$

where $n_t(p)$ and $R_t(p)$ are defined as

$$\begin{aligned} n_t(p) &= \lim\nolimits_{\epsilon \to 0^+} F_t(p) - F_t(p - \epsilon), \\ R_t(p) &= q_t^*(p)(p - cw_t). \end{aligned}$$

The first term on the right-hand side of (2.8) is the number of customers served by the firm. First, there are α_1/s households who purchase the good from the firm because they are not in contact with any other seller. Second, there are $2\alpha_2(1 - F_t(p))/s$ households who purchase the good from the firm because the other seller with whom they are in contact charges a price greater than p. Finally, there are $2\alpha_2 n_t(p)/s$ households who are in contact with the firm and with another seller that also charges the price p. Each of these households purchases the good from the firm with probability 1/2.

The second term on the right-hand side of (2.8) is the firm's profit per customer. The firm sells $q^*(p)$ units to each customer and each unit is sold at the price p and produced at the cost cw_t. In Figures II and III, we plot the firm's profits per customer as a function of the price p. The black curve is the plot of $(m_t/p)(p - cw_t)$, which us the profit per customer given that the customer is cash constrained. The gray curve is the plot of $(w_t/p)^{1/\gamma}(p - cw_t)$, which is the profit per customer given that the customer is not cash constrained. The firm's actual profit per customer is given by the lower envelope of the black and grey curves. Figure II illustrates a case in which the customer's money holdings are relatively high. In this case, the price that maximizes the firm's profit per customer is $cw_t/(1 - \gamma)$. Figure III illustrates the case in which the customer's money holdings are relatively low. In this case, the price that maximizes the firm's profit per customer is $\hat{p}_t$. Overall, the price that maximizes the firm's profit per customer is $p_t^m = \max\{cw_t/(1-\gamma), \hat{p}_t\}$. We shall refer to p_t^m as the *monopoly price*.

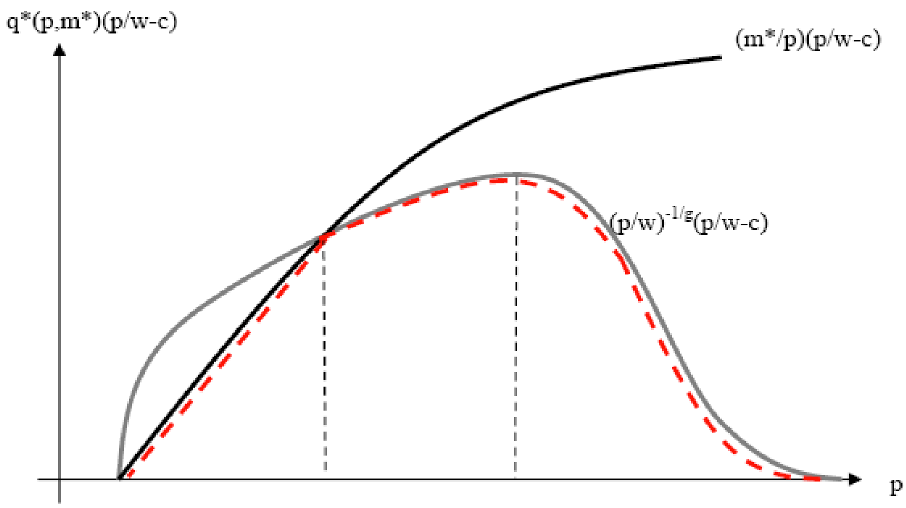

Figure II: Firm's profit per customer in the BJ market

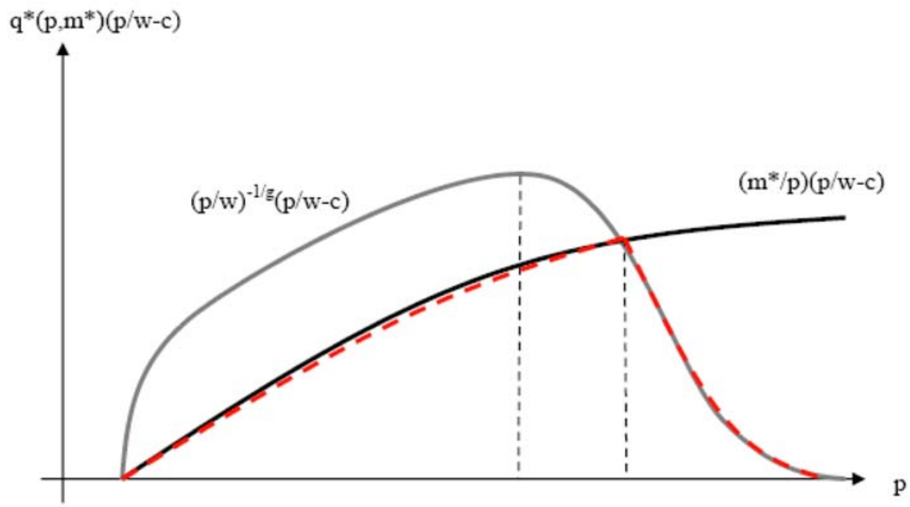

Figure III: Firm's profit per customer in the BJ market

Each firm chooses the price p so as to maximize the profit function $\Pi_t(p)$. Therefore, the price distribution F_t is consistent with the firms' pricing strategy only if the profit function $\Pi_t(p)$ is maximized by every price on the support of F_t, i.e.

$$\begin{array}{c} \Pi_t(p_0) = \Pi_t^* \text{ for all } p \in \text{supp}F, \\ \Pi_t^* \equiv \max_p \Pi_t(p). \end{array} \tag{2.9}$$

The following lemma makes use of condition (2.9) in order to characterize the price distribution F_t. The proof of the lemma adapts the arguments developed by Burdett and Judd (1983) to characterize the price distribution in a market for an indivisible good where buyers have deep-pockets to a market for a divisible good where buyers are cash-constrained.

Lemma 1 (Burdett and Judd, 1983): *The unique price distribution that is consistent with the firms' pricing strategy is*

$$F_t(p) = 1 - \frac{\alpha_1}{2\alpha_2}\left[\frac{R_t(p_t^m)}{R_t(p)} - 1\right]. \tag{2.10}$$

The price distribution F_t is continuous and its support is the connected interval $[\underline{p}_t, \overline{p}_t]$, where

$$\begin{array}{c} R_t(\underline{p}_t) = \dfrac{\alpha_1}{\alpha_1 + 2\alpha_2} R_t(p_t^m), \\ \overline{p}_t = p_t^m. \end{array} . \tag{2.11}$$

Proof: In Appendix A.

First, lemma 1 states that the price distribution is continuous. Intuitively, if the price distribution had a mass point at p_0, a firm that posts p_0 could increase its profits by charging $p_0 - \epsilon$ instead, as the deviation would leave the firm's profit per customer approximately constant but it would increase the firm's customer base by a discrete amount. Second, lemma 1 states that the support of the price distribution is connected. Intuitively, if the support of the price distribution had a gap between p_0 and p_1, a firm that posts p_0 could increase its profits by charging p_1 instead, as the deviation would not affect the firm's customer base and it would increase the firm's profits per customer. Third, lemma 1 states that the highest price on the support of the price distribution, $\overline{p}_t$, is equal to the monopoly price, p_t^m. Intuitively, if $\overline{p}_t$ were smaller than p_t^m, a firm that posts $\overline{p}_t$ could increase its profits by charging p_t^m instead, as the deviation would not affect the firm's customer base and it would increase the firm's profits per customer. If, on the other hand, $\overline{p}_t$ were greater than p_t^m, a firm that posts $\overline{p}_t$ could increase its profits by charging p_t^m, as the deviation would increase both the firm's customer base and the profit per customer.

Given the above properties of the price distribution and condition (2.9), we can derive the expression (2.10) for F_t and the expression (2.11) for $\underline{p}_t$. Since the price distribution has no mass points, the profit of a firm that charges the price p is given by

$$\Pi_t(p) = \frac{1}{s}\left[\alpha_1 + 2\alpha_2(1 - F_t(p))\right] R_t(p). \tag{2.12}$$

Since the monopoly price is on the support of the price distribution, the maximized profit of the firm is given by

$$\Pi_t^* = \frac{\alpha_1}{s} R_t(p_t^m) \tag{2.13}$$

Finally, since every price on the support of the price distribution must maximize the profit of the firm, it follows that for every $p \in [\underline{p}_t, \overline{p}_t]$

$$\frac{1}{s}\left[\alpha_1 + 2\alpha_2(1 - F_t(p))\right] R_t(p) = \alpha_1 R_t(p_t^m). \tag{2.14}$$

Solving (2.14) with respect to F_t leads to equation (2.10). In turn, setting $F_t(\underline{p}_t) = 0$ and solving for $\underline{p}_t$ leads to equation (2.11).

2.4 Equilibrium

We are now in the position to define an equilibrium.

Definition 2 (Stationary Monetary Equilibrium): *A stationary monetary equilibrium, Σ^*, is a sequence of quantities $\{m_t^*, x_t^*, h_t^*, q_t^*\}_{t=0}^{\infty}$ and prices $\{w_t^*, F_t^*\}$ that satisfy the following conditions:*

1. m^*_{t+1}, x^*_t *and* h^*_t *solve the household's problem in the AD market:*

$$i = \int_{\mathbf{1}[p \leq \hat{p}_t]} \left[\frac{w^*_t}{p} \left(\frac{m^*_t}{p} \right)^{-\gamma} - 1 \right] \left[\alpha_1 + 2\alpha_2 \left(1 - F^*_t(p)\right) \right] dF^*_t(p),$$
$$v'(x^*_t) = 1, \qquad h^*_t(m) = x^*_t + w^{*-1}_t \left(m^*_{t+1} - m - D_t - T_t \right);$$

2. q^*_t *solves the household's problem in the BJ market:*

$$q^*_t(p) = \begin{cases} m^*_t/p, & \text{if } p \leq \hat{p}_t, \\ (w^*_t/p)^{\frac{1}{\gamma}}, & \text{if } p > \hat{p}_t, \end{cases}$$

where $\hat{p}_t = w^*_t (w^*_t/m^*_t)^{\frac{\gamma}{1-\gamma}}$*;*

3. F^*_t *is consistent with the solution to the firm's problem in the BJ market:*

$$F^*_t(p) = 1 - \frac{\alpha_1}{2\alpha_2} \left[\frac{R_t(p^m_t)}{R_t(p)} - 1 \right] \quad \text{all } p \in [\underline{p}_t, \overline{p}_t],$$

where R_t *is given by (2.8) and* $\underline{p}_t$ *and* $\overline{p}_t$ *are given by (2.11);*

4. w^*_t *is such that the household's money demand equates the government's money supply:*

$$m^*_t = M_t;$$

5. *nominal variables grow at the rate* μ *and real variables are constant:*

$$m^*_{t+1} = \mu m^*_t, \quad w_{t+1} = \mu w_t, \quad F^*_{t+1}(\mu p) = F^*_t(p),$$
$$x^*_{t+1} = x^*_t, \quad h^*_{t+1} = h^*_t, \quad q^*_{t+1}(\mu p) = q^*_t(p).$$

To establish the existence of a stationary monetary equilibrium, we proceed in two steps. In the first step, we prove that the prices posted by the firms in the BJ market are decreasing (in the first order stochastic dominance sense) with respect to the amount of money that the households are expected to carry into the BJ market. The intuition for this result is simple. If the households carry more money, the cash constraint is relaxed, and the quantity of the good demanded by a customer at a low-price firm increases relative to the quantity demanded at a high-price firm. For this reason, the profits that a low-price firm makes on each customer increase relative to the profits made by a high-price firm. In order to keep the firms indifferent between posting low and high prices, the distribution of prices must fall so that the number of customers served by a low-price firm declines relative to the number of customers served by a high-price firm.

In the second step, we prove that the amount of money carried by the households in the BJ market is decreasing with respect to the prices posted by the firms in the BJ market. Intuitively, if the prices posted by the firms are higher (in the sense of first order stochastic dominance), the household has a lower probability of meeting a low-price seller and, consequently, a lower probability of being cash constrained. Therefore, if the prices posted by the firms are higher, the value to the household from carrying an additional unit of money in the BJ market falls.

From the above observations, it follows that the amount of money that the households choose to carry into the BJ market is a monotonic function of the amount of money that the firms expect the households to hold. Moreover, we can prove that the amount of money that the households choose to carry into the BJ market is bounded. Hence, from Tarski's fixed point theorem (Tarski, 1955), there exists an m_t^* such that: (i) m_t^* is the amount of money that solves the households' problem given that the price distribution is F_t^*, and (ii) F_t^* is the price distribution given that the firms expect the households to have an amount of money m_t^*. That is, from Tarski's fixed point theorem, it follows that there exist m_t^* and F_t^* that satisfy the equilibrium conditions (1) and (3). Given m_t^* and F_t^*, we can then find the h_t^*, x_t^*, q_t^* and w_t^* that satisfy the remaining equilibrium conditions.

Theorem 3 (Existence): *For any $\mu > \beta^{-1}$, a stationary monetary equilibrium exists.*
Proof: In Appendix B.

3 Price stickiness and monetary neutrality

3.1 Price stickiness

The equilibrium uniquely pins down the aggregate price distribution in the BJ market, but it does not pin down the price of an individual firm. In fact, an individual firm is indifferent between posting any price from the support of the aggregate price distribution, as any one of these prices is profit maximizing. Figure IV illustrates the implications of this property of equilibrium for the dynamics of the aggregate price distribution and for the dynamics of the price of individual firms in the case of an inflationary economy, i.e. $\mu > 1$. The black line is the aggregate price distribution in period t, F_t^*, and the red line is the aggregate price distribution in period $t+1$, F_{t+1}^*. All the firms in the vertically shaded area must change their price between the two periods as their period-t price does not maximize profits in period $t+1$. Each of the firms in the horizontally shaded area, however, is indifferent between keeping its price constant, posting a new price from the interval $[\underline{p}_{t+1}, \overline{p}_{t+1}]$, or randomizing between keeping its price and changing it. The only equilibrium restriction on the price dynamics of individual firms between period t and $t+1$ is that the aggregate price distribution in period $t+1$ has to be F_{t+1}^*. We formalize this restriction in the following definition.

Definition 4 (Supportable policies): *The stationary monetary equilibrium Σ^* supports the price policy $p_{t+1}^*(p)$ if, given that the price distribution in period t is $F_t^*(p)$ and all firms follow the policy $p_{t+1}^*(p)$, the price distribution in period $t+1$ is $F_{t+1}^*(p)$.*

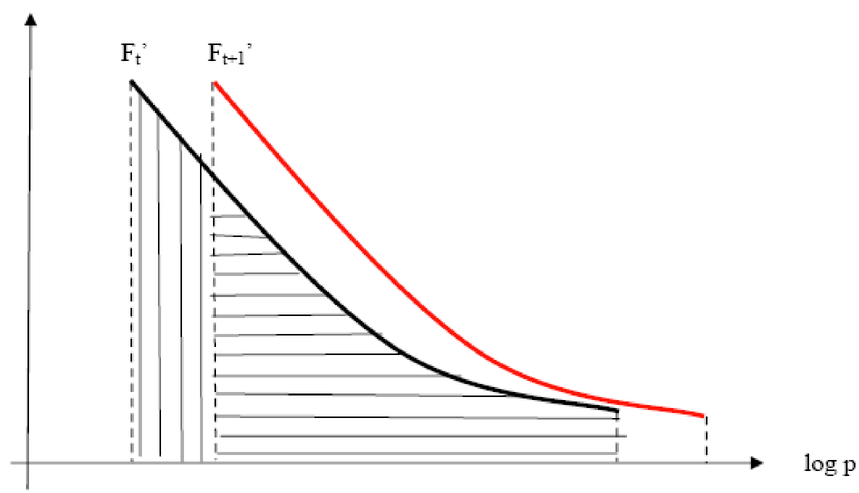

Figure IV: Equilibrium price distribution

In the remainder of the paper, we will restrict attention to the case of an inflationary economy and to pricing policies of the form

$$
\begin{aligned}
&p^*_{t+1}(p_t) = \left\{ \begin{array}{ll} p_t & \text{w.p. } \rho \\ p' & \text{w.p. } 1-\rho \end{array} \right. , && \text{if } p_t \geq \mu \underline{p}_t = \underline{p}_{t+1}, \\
&p^*_{t+1}(p_t) = p', && \text{if } p_t < \mu \underline{p}_t = \underline{p}_{t+1},
\end{aligned}
\tag{3.1}
$$

where ρ is a parameter between 0 and 1 and p' is a price randomly drawn from the cumulative distribution function

$$
G^*_{t+1}(p) = \left\{ \begin{array}{ll} \dfrac{F^*_t(p/\mu)}{F^*_t(\mu \underline{p}_t) + (1-\rho)\left[1 - F^*_t(\mu \underline{p}_t)\right]}, & \text{if } p < \mu \underline{p}_t, \\ \dfrac{F^*_t(p/\mu) - \rho F^*_t(p,\mu)}{F^*_t(\mu \underline{p}_t) + (1-\rho)\left[1 - F^*_t(\mu \underline{p}_t)\right]}, & \text{if } p \geq \mu \underline{p}_t. \end{array} \right.
\tag{3.2}
$$

According to (3.1), the firm posts a new price in period $t+1$ if the price it posted in the period t is outside of the support of the equilibrium price distribution for period $t+1$. Otherwise, the firm posts the same price in period $t+1$ as in period t with probability ρ and posts a different price with probability $1-\rho$. Whenever the firm changes price, it draws the new price from the distribution function $G^*_{t+1}(p)$ defined in (3.2). To see that $G^*_{t+1}(p)$ is a legitimate cumulative distribution function, notice that $G^*_{t+1}(p) = 0$ for all $p \leq \mu \underline{p}_t$, $G^*_{t+1}(p) = 1$ for all $p \geq \mu \overline{p}_t$, and $G'^*_{t+1}(p) > 0$ for all p in the interval between $\mu \underline{p}_t$ and $\mu \overline{p}_t$.

Given that the price distribution in period t is F^*_t and that all firms follow the pricing policy p^*_{t+1}, the price distribution in period $t+1$ is given by

$$
F_{t+1}(p) = \left\{ \begin{array}{ll} \left\{F^*_t(\mu \underline{p}_t) + (1-\rho)\left[1 - F^*_t(\mu \underline{p}_t)\right]\right\} G^*_{t+1}(p), & \text{if } p < \mu \underline{p}_t, \\ \rho F^*_t(p) + \left\{F^*_t(\mu \underline{p}_t) + (1-\rho)\left[1 - F^*_t(\mu \underline{p}_t)\right]\right\} G^*_{t+1}(p), & \text{if } p \geq \mu \underline{p}_t. \end{array} \right.
\tag{3.3}
$$

Using (3.2) to substitute out $G^*_{t+1}(p)$ in (3.3), we obtain $F_{t+1}(p) = F^*_t(p/\mu)$ which, in turn, is equal to $F^*_{t+1}(p)$. Hence, we have established the following result.

Proposition 5 (Supported policies): *The pricing policy (3.1) is supported by the stationary monetary equilibrium Σ^* for all $\rho \in [0,1]$.*

The class of pricing policies described by (3.1) is not exhaustive, but it is sufficient to capture a wide range of behaviors in a parsimonious way. For $\rho = 1$, the pricing policy (3.1) describes the extreme case in which a firm only changes its price when it is no longer profit maximizing. Clearly, for $\rho = 1$, the pricing policy (3.1) attains the smallest fraction of price changes and the highest average price duration that are consistent with equilibrium. For $\rho = 0$, the pricing policy (3.1) describes the opposite extreme case in which a firm changes its price in every period. Clearly, for $\rho = 0$, the pricing policy (3.1) attains the largest fraction of price changes and the lowest average price duration that are consistent with equilibrium. For ρ between 0 and 1, the pricing policy (3.1) describes the intermediate cases in which the firm changes a profit maximizing price with probability ρ. As the parameter ρ increases from 0 to 1, the frequency of price changes and the average duration of a price move from one extreme to the other.

For a given ρ, let us compute the average duration of a price. The cumulative distribution of new prices in period t is $G_t^*(p)$. Let N denote the largest integer such that $\mu^N \underline{p}_t \leq \overline{p}_t$. For $n = 1, 2, ...N$, a fraction $G_t^*(\mu^n \underline{p}_t) - G_t^*(\mu^{n-1} \underline{p}_t)$ of new prices lies in the interval $[\mu^{n-1} \underline{p}_t, \mu^n \underline{p}_t]$. Also, a fraction $1 - G_t^*(\mu^N \underline{p}_t)$ of new prices lies in the interval $[\mu^N \underline{p}_t, \overline{p}_t]$. Each price in the interval $[\mu^{n-1} \underline{p}_t, \mu^n \overline{p}_t]$ will be changed in period $t + i$ (and not before) with probability $\rho^{i-1}(1 - \rho)$, $i = 1, 2, ...n - 1$, and will be changed in period $t + n$ with probability ρ^{n-1}. Therefore, the average duration of prices in the interval $[\mu^{n-1} \underline{p}_t, \mu^n \overline{p}_t]$ is equal to $(1 - \rho) + 2\rho(1 - \rho) + ... + n\rho^{n-1} = (1 - \rho^n)/(1 - \rho)$ periods. Each price in the interval $[\mu^N \underline{p}_t, \overline{p}_t]$ will be changed in period $t + i$ with probability $\rho^{i-1}(1 - \rho)$, $i = 1, 2, ...N$, and will be changed in period $t + N + 1$ with probability ρ^N. Therefore, the average duration of prices in the interval $[\mu^N \underline{p}_t, \overline{p}_t]$ is equal to $(1 - \rho^{N+1})/(1 - \rho)$ periods. Overall, the average duration of a new price is

$$A(\rho) = \left\{ \sum_{n=1}^{N} \left[G_t^*(\mu^n \underline{p}_t) - G_t^*(\mu^{n-1} \underline{p}_t) \right] \frac{1 - \rho^n}{1 - \rho} \right\} + \left[1 - G_t^*(\mu^N \underline{p}_t) \right] \frac{1 - \rho^{N+1}}{1 - \rho} \qquad (3.4)$$

Notice that, since the ratio $(1 - \rho^n)/(1 - \rho)$ is increasing in ρ and n and the distribution function G_t^* is increasing in ρ (in the first order stochastic dominance sense), $A(\rho)$ is an increasing function of ρ.

Second, we compute the fraction of prices that change between period t and $t + 1$. The cumulative distribution of prices in period t is $F_t^*(p)$. A fraction $F_t^*(\mu \underline{p}_t)$ of prices lies in the interval $[\underline{p}_t, \mu\underline{p}_t]$, and each of these prices will be changed between period t and $t+1$ with probability 1. A fraction $1 - F_t^*(\mu\underline{p}_t)$ of prices lies in the interval $[\mu\underline{p}_t, \overline{p}_t]$, and each of these prices will be changed between period t and $t+1$ with probability $1-\rho$. Overall, the fraction of prices that change between t and $t+1$ is given by

$$FR(\rho) = F_t^*(\mu\underline{p}_t) + (1-\rho)\left[1 - F_t^*(\mu\underline{p}_t)\right]. \tag{3.5}$$

Clearly, $FR(\rho)$ is a decreasing function of ρ.

Third, we compute the distribution of the magnitude of price changes. The density of firms that post the price p in period t and a different price in period $t+1$ is given by $F_t'^*(p)/FR(\rho)$ if p is smaller than $\mu\underline{p}_t$, and by $(1-\rho)F_t'^*(p)/FR(\rho)$ if p is greater than $\mu\underline{p}_t$. Among the firms that post the price p in period t and a different price in period $t+1$, a fraction $G_{t+1}^*(p(1+x))$ increases its price by x percent or less. Therefore, the cumulative distribution function for the magnitude of price changes is given by

$$H_t(x,\rho) = \frac{1}{FR(\rho)} \int G_{t+1}^*(\rho(1+x))\left(1 - \rho\mathbf{1}[p \geq \mu\underline{p}_t]\right) dF_t'(p). \tag{3.6}$$

From (3.2) and (3.6), it is immediate to verify that the fraction of negative price changes, $H_t(0,\rho)$, is strictly positive for all $\rho < 1$.

The following theorem contains the main theoretical results of the paper.

Theorem 6 (Sticky prices): *The stationary monetary equilibrium Σ^*, together with the pricing policy p_{t+1}^*, is such that the average price duration is $A(\rho)$ and the frequency of a price change is $FR(\rho)$, where $A(\rho)$ is increasing and $FR(\rho)$ is decreasing in ρ.*

(i) There exists a $\mu^* > 1$ such that, for all $\mu \in (1, \mu^*)$ and all $\rho \in (0,1]$, the average price duration $A(\rho)$ is strictly greater than 1, and the frequency of a price change $FR(\rho)$ is strictly smaller than 1.

(ii) For all $\mu \in (1,\mu^*)$ and all $\rho \in [0,1)$, the fraction of negative price changes, $H(0,\rho)$, is strictly greater than 0.

Proof: In Appendix C.

Part (i) of the theorem shows that, unless the growth rate of money is too high, our model is consistent with the observation that some firms leave their price unchanged for weeks or months in the face of a continuously changing aggregate price level. Our model delivers this result not because there are technological restriction to price adjustment, but because, due to search frictions, there is an entire interval of prices for which the profits of the firm are maximized. Part (ii) of the theorem shows that our model is consistent with the observation that some firms change their price down in the face of a continuously increasing aggregate price level. Our model delivers this result because of search frictions, not because of idiosyncratic shocks to the firm's cost of production. More broadly, the theorem shows that one should be cautious in inferring the existence of menu costs or Calvo fairies from the observed stickiness of individual sellers prices. Similarly, the theorem shows that negative price changes in an inflationary economy are not necessarily caused by idiosyncratic productivity shocks.

3.2 Monetary neutrality

In our model, some firms may post the same nominal price for many periods in the face of a continuously increasing aggregate price level. The government, however, cannot increase short-run production or consumption through an unexpected increase in the monetary base. For example, consider what would happen in the government were to unexpectedly double the stock of money at the opening of the AD market. In response to this policy, the amount of money that the households carry into the BJ market would double, but so would the distribution of prices that the firms post in the BJ market. Hence, the quantity of the BJ good traded and produced would remain unchanged. Similarly, in response to this policy, the amount of money that the households spend in the AD market would double, but so would the prices charged by the firms in the AD market. Hence, the quantity of the AD good traded and produced would remain unchanged. Overall, expansionary monetary policy is completely neutral. Intuitively, money is neutral because, while the price posted by some sellers is rigid, the distribution of prices is perfectly flexible.

4 Quantitative evaluation

In the previous section, we developed a theory of price rigidity that does not rely on the existence of technological frictions to price adjustment (e.g. menu costs or Calvo fairies), but on the existence of search frictions in the product market. Unlike theories of price rigidity based on technological constraints to price adjustment, our theory implies that money is neutral because the equilibrium price distribution responds fully and instantaneously to increases in the stock of money. In this section, we want to find out whether our theory can account for the empirical facts about prices that have been documented by Klenow and Kryvtsov (2008).

4.1 Data and calibration

The household's preferences are described by the utility function for the BJ good, $u(q) = q^{1-\gamma}/(1-\gamma)$, the utility function for the AD good, $v(x)$, and the discount factor, β. The firm's technology is described by the production function for the BJ good, $f(h) = h/c$, and the production function for the AD good, $g(h) = h$. We restrict attention to the case in which the BJ and AD goods are produced using the same technology, i.e. $c = 1$. The firm's pricing behavior is described by the policy function $p^*_{t+1}(p, \rho)$. The search frictions in the BJ market are described by the probability distribution $\{\alpha_i\}_{i=0}^2$, where α_i is the probability that a household meets i firms. We restrict attention to the family of probability distributions $\{\alpha_i\}_{i=0}^2$ that would obtain if each household could search the BJ market twice and each search would lead to meeting a firm with probability λ. That is, we restrict attention to the family of probability distribution such that $\alpha_0 = (1-\lambda)^2$, $\alpha_1 = 2(1-\lambda)\lambda$ and $\alpha_2 = \lambda^2$. Finally, the monetary policy is described by the growth rate of money, μ. Overall, the model is fully characterized by the parameters $\{\gamma, \beta, \rho, \lambda, \mu\}$ and the utility function $v(x)$.

We calibrate the model to the US economy over the period 1988-2004, and we interpret the BJ market as the retail sector and the AD market as the intermediate goods sector. We choose the model period to be a month. We choose the parameter β so that the annual real interest rate in the model, β^{-12}, equals the average real interest rate in the data, 1.035. We choose the parameter μ so that the annual inflation rate in the model, μ^{12}, equals the average inflation rate in the data, 1.03. We choose the parameters γ and ρ so as to minimize the distance between the model-generated distribution of price changes in the BJ market, $H_t(x, \rho)$, and the empirical distribution of price changes observed in the retail sector (as measured by Klenow and Kryvtsov, 2008). Finally, we choose the parameter λ so that the average mark-up in the BJ market is 30 percent, which is a common estimate of the average mark-up in the retail sector (see Faig and Jerez, 2005). After having calibrated the parameters $\{\gamma, \beta, \rho, \lambda, \mu\}$, the predictions of the model regarding the behavior of prices in the BJ market are uniquely pinned down. Since we are interested in comparing these predictions of the model with the empirical behavior of prices in the retail sector, we do not need to calibrate the utility function for the AD good, $v(x)$.

There is a simple intuition behind our calibration strategy for γ and ρ. The parameter γ determines the elasticity of the firm's profit per customer, $R_t(p)$, with respect to the firm's price, p. Hence, the parameter γ affects the equilibrium price distribution, $F_t^*(p)$, and the equilibrium price change distribution, $H_t(x, \rho)$. Similarly, the parameter ρ determines the probability that a firm does not adjust its price when the firm is indifferent between adjusting and not adjusting. Hence, the parameter ρ affects the distribution of prices among firms that do not change their price and, consequently, the distribution of prices among firms that do change their price, $G_t^*(p)$, and the price change distribution, $H_t(x, \rho)$. Our calibration strategy for λ is also easy to understand. In fact, the equilibrium price distribution, $F_t^*(p)$, is decreasing (in the first-order stochastic dominance sense) with respect to the meeting probability λ.

4.2 The price of individual sellers

Using the calibrated model, we derive the predictions of the model regarding the behavior of prices and compare these predictions with the empirical findings reported in Klenow and Kryvtsov (2008). The bottom line is that our search theory of price rigidity can account quite well for the empirical behavior of prices.

According to the data analyzed by Klenow and Kryvtsov (2008), the average duration of a price in the retail sector is between 6.8 and 10.4 months, depending on whether temporary sales and product substitutions are interpreted as price changes or not. In particular, if temporary sales and product substitutions are both interpreted as price changes, the average duration of a price is 6.8 months. If product substitutions are considered as price changes but temporary sales are not, the average duration of a price is 8.6 months. If neither product substitutions nor temporary sales are considered price changes, the average duration of a price increases to 10.4 months. Following Klenow and Kryvtsov (2008), we will take the second case as our benchmark.

The average duration of a price predicted by the model is close to its empirical counterpart. In particular, given an average inflation rate of 3 percent and a calibrated value of ρ of 0.93, the model predicts that the average duration of a price is 11.6 months. Notice that, for higher values of ρ, the average price duration predicted by the model would be higher, up to a maximum of 34 months. Conversely, for lower values of ρ, the average price duration implied by the model would be lower, down to a minimum of 1 month. The model would generate the same average price duration as in the data for $\rho = 0.91$, a value very close to the one that is obtained with our calibration. Also, notice that the average price duration is decreasing with respect to inflation, as higher inflation increases the fraction of prices that exit the support of the equilibrium price distribution in each period. These properties of the average price duration are illustrated in Figure V.

Figure V: Average price duration

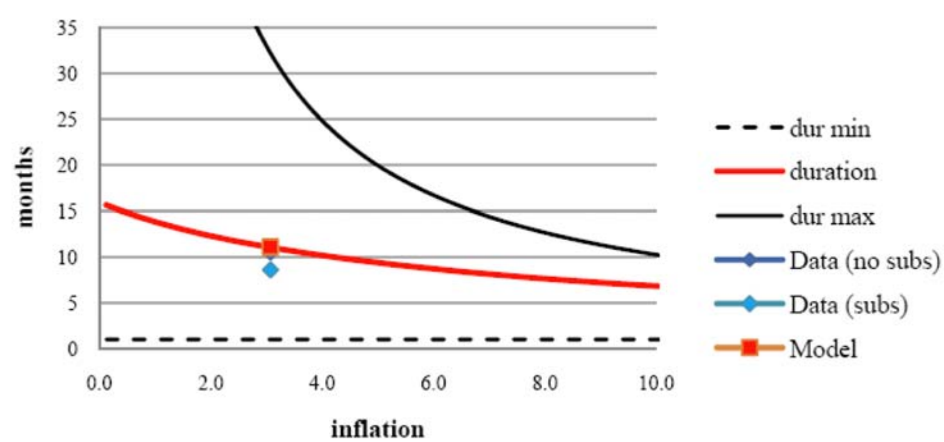

The blue histogram in figure VI is the empirical price change distribution estimated by Klenow and Ksrystof (2008). Three features of the price change distribution are worth stressing. First, on average, price changes are large. More specifically, the average of the absolute value of price changes is 11 percent. Second, many price changes are small. More specifically, 44 percent of all price changes are smaller than 5 percent in absolute value. Third, many price changes are negative. More specifically, 35 percent of all price changes are negative. Klenow and Krystof (2008), Golosov and Lucas (2007), and Midrigan (2006) interpret the existence of many small and negative price changes as evidence of large and frequent shocks to individual seller's idiosyncratic productivity.

The red histogram in figure VI is the price change distribution predicted by the model. One can immediately see that the model-generated price change distribution is very close to its empirical counterpart and shares the same features. More specifically, for the model-generated price change distribution, the average of the absolute value of price changes is 9 percent, the fraction of price changes between -5 and +5 percent is 43 percent, and the fraction of negative price changes is 35 percent. Interestingly, our model generates a good fit of the empirical price change distribution without seller-specific productivity shocks. According to our model, price changes are large because search frictions create a lot of dispersion in the equilibrium price distribution. For example, the price posted by a firm at the 90th percentile of the equilibrium price distribution is approximately 2 times larger than the price posted by a firm at the 10th percentile. Hence, the firms whose price exits the support of the distribution make, on average, a large price adjustment. According to our model, many price changes are small because there are many firms that change their price before it exits the support of the distribution. For the same reason, many price changes are negative.

Klenow and Kryvtsov (2008) estimate the relationship between the probability that a firm adjusts its price for a given item (i.e. the price-change hazard) and the time since the previous price adjustment (i.e. the age of the price). Moreover, they estimate the relationship between the absolute value of the size of a price adjustment (i.e. the price-change size) and the time since the previous price adjustment. After controlling for unobserved heterogeneity across different items, they find that the price-change hazard remains approximately constant during the first 11 months and increases significantly during the 12th month. Similarly, after controlling for unobserved heterogeneity, they find that the price-change size is approximately independent of the age of the price.

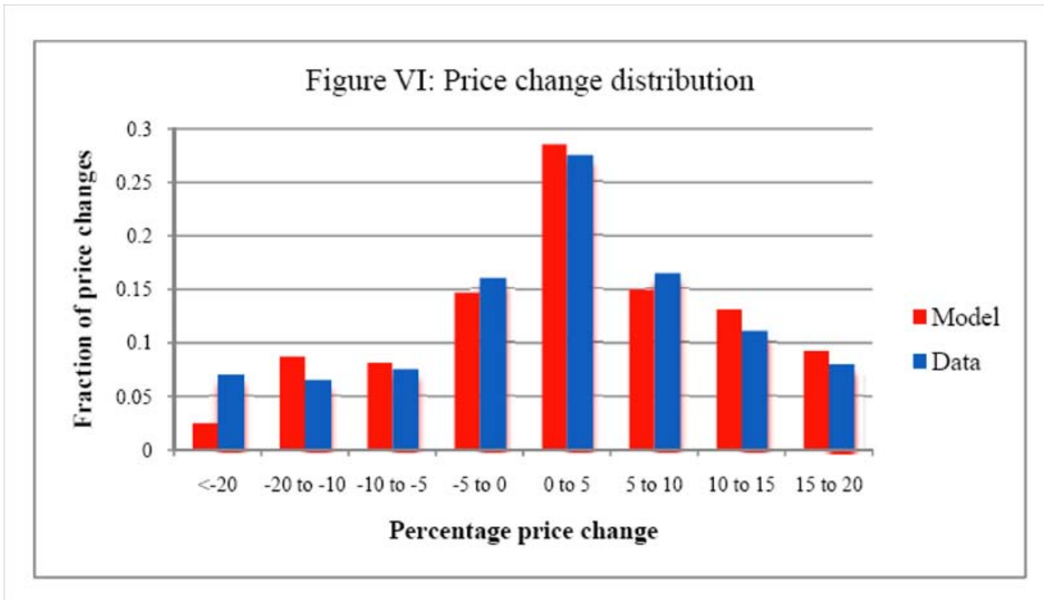

The red histogram in figure VII is the price-change hazard predicted by our model. Notice that, just like in the data, the price-change hazard is approximately constant for the first 11 months of the life of a price. However, unlike in the data, the price-change hazard does not increase on the 12th month of the life of a price. These findings are easy to explain. In our model, the equilibrium price distribution has a wide support. Therefore, during the 12 months after a price change, only few firms need to readjust their price because it is no longer profit maximizing. Instead, during the 12 months after a price change, the majority of firms change their price with probability ρ, a probability that is independent of the age of a price. The model does not predict a spike in the price-change hazard after 12 months because, unlike in the real world, firms in our model have no seasonal incentives to adjust their price. The blue histogram in figure 2 is the price-change size predicted by the model. For the same reasons that we mentioned above, the price-change size predicted by the model is approximately independent of the age of a price.

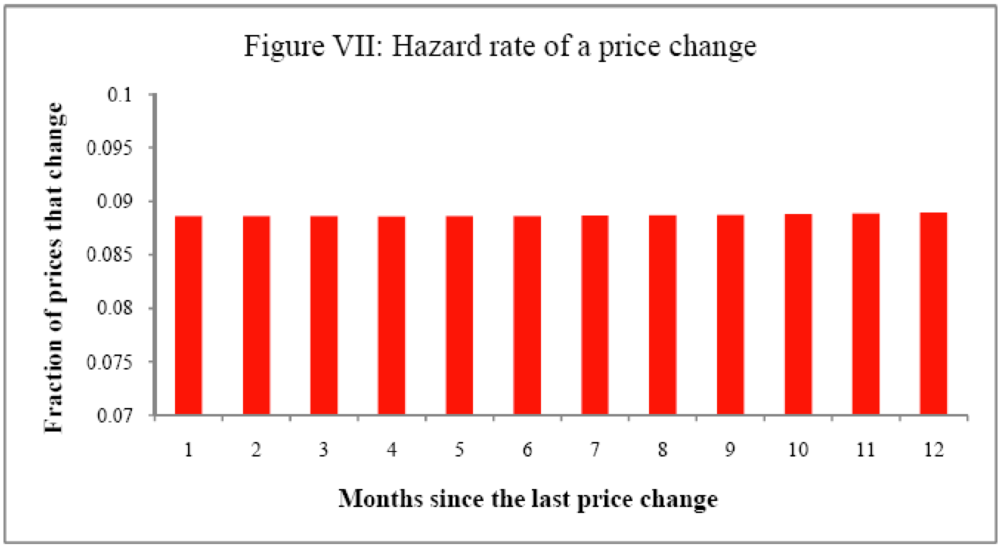

4.3 The effect of inflation

Using time-variation in the US inflation rate over the period 1988-2005, Klenow and Kryvtsov (2008) measure the effect of inflation on the frequency of price adjustments (i.e. the extensive marine of price adjustments) and on the magnitude of price adjustments (i.e. the intensive margin of price adjustments). They accomplish this task by estimating the coefficient on inflation in the regression of the frequency of price adjustments and in the regression of the magnitude of price adjustments. Their main finding is that inflation has a positive effect on both the frequency and the magnitude of price adjustments. More specifically, they find that a 1 percentage point increase in inflation increases the frequency of price adjustments by 2.38 percentage points and the magnitude of price adjustments by 3.55 percentage points.

Figure VIII illustrates the predictions of our model regarding the effect of inflation on the frequency and magnitude of price adjustments. Notice that, according to our model, an increase in inflation increases both the frequency and the magnitude of price adjustments. These findings are easy to explain. First, an increase in inflation leads to a decline in the real balances carried by the households in the BJ market and, in turn, the decline in real balances leads to a compression of the support of the equilibrium price distribution. Second, given the length of the support of the equilibrium price distribution, an increase in inflation reduces the time it takes for a price to exit the support. For both reasons, an increase in inflation increases the fraction of prices that are adjusted in every month. For similar reasons, an increase in inflation leads to an increase in the average price adjustment.

In our model, the effect of inflation on the extensive and intensive margins of price adjustment has the same sign as in the data but a different magnitude. For example, an increase in annual inflation from 3 to 4 percent increases the frequency of price adjustment by approximately 9 percentage points, and it increases the magnitude of price adjustments by approximately 5 percentage points. This discrepancy between the predictions of the model and the results of the regression analysis should not be too surprising nor too much of a concern. In the real world, fluctuations in the inflation rate are likely to be correlated with other shocks (e.g. productivity shocks) that are not controlled for in the regression. In our model, fluctuations in the inflation rate are the only shock.

Figure VIII: Fraction and size of price changes

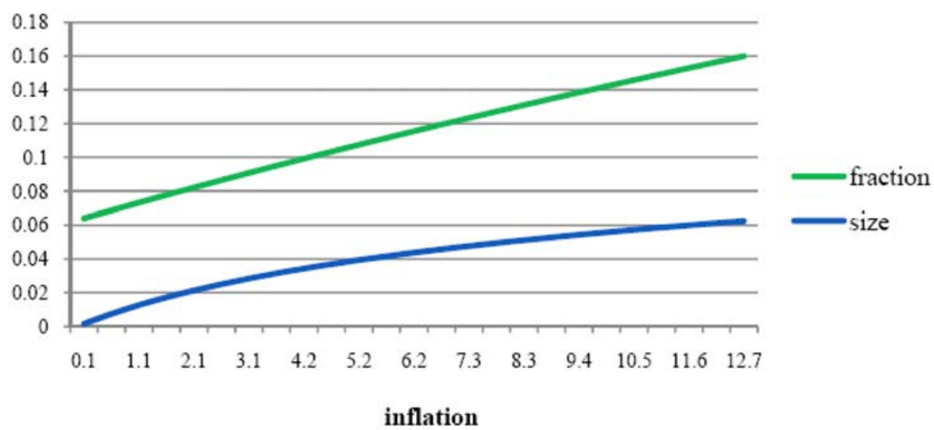

Finally, Klenow and Kryvstov (2008) measure the effect of inflation on the fraction of prices that increase (i.e. positive price adjustments) and on the fraction of prices that decrease (i.e. negative price adjustments). Again, they accomplish this task by estimating the coefficient on inflation in the regressions of the fraction of prices that increase and on the fraction of prices that decrease. Their main finding is that inflation has a positive effect on the fraction of price increases and a negative effect on the fraction of prices that decrease. More specifically, they find that a 1 percentage point increase in inflation increases the fraction of positive price changes by 5.48 percentage points and it decreases the fraction of negative price changes by 3.10 percentage points.

Figure IX illustrates the predictions of our model regarding the effect of inflation on positive and negative price changes. Like in the data, our model predicts that an increase in inflation increases the fraction of positive price adjustments, and it reduces the fraction of negative price adjustments. However, the magnitude of the effect of inflation on these variables is different than the one obtained from the regression analysis. In particular, according to our model, an increase in annual inflation from 3 to 4 percent increases the fraction of positive price changes by approximately 10 percentage points, and it decreases the fraction of negative price changes by approximately 2.5 percentage points. For the reasons discussed above, we do not find this discrepancy between the predictions of the model and the outcomes of the regression analysis too surprising.

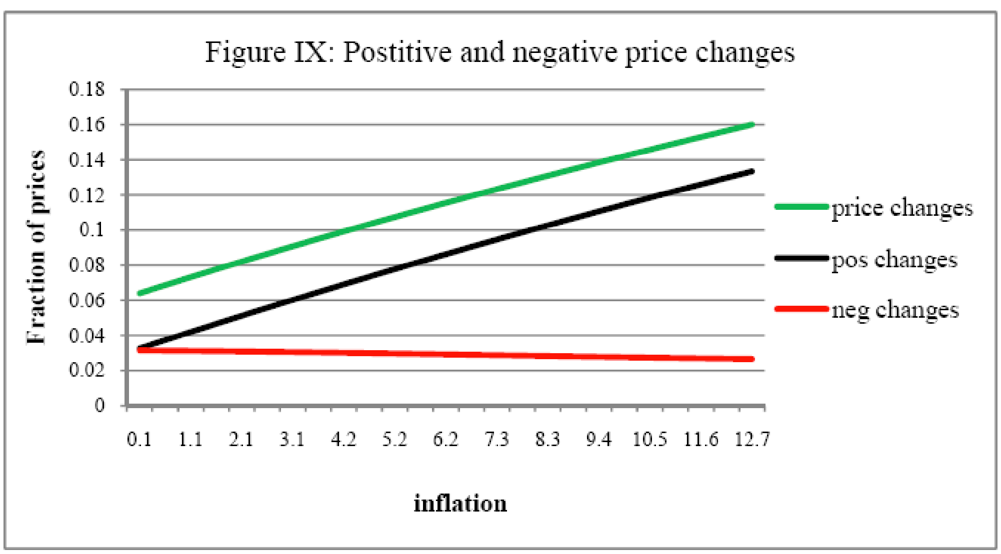

4.4 Summary of results

From the previous paragraphs, it is clear that our theory of price rigidity can account quite well for the empirical behavior of prices in the US retail market. First, the model predicts an average price duration (11.6 months) that is close to the one observed in the data (8.6 months). Second, the model generates a price change distribution that has the same shape and the same features of the empirical price change distribution (i.e. the average magnitude of price changes is large, there are many small price changes, and there are many negative price changes). Third, as it is observed in the data, in the model the probability and magnitude of price adjustments are approximately independent of the age of a price. Fourth, the model correctly predicts that inflation increases both the frequency and the magnitude of price changes. Finally, the model correctly predicts that inflation increases the fraction of positive price changes and reduces the fraction of negative price changes.

In contrast to our model, existing theories of price rigidity cannot account for all these features of the empirical behavior of prices. On the one hand, menu costs theories of price rigidity (e.g. Golosov and Lucas, 2007) cannot simultaneously account for the average duration of prices (which suggests that menu costs are large) and for the large fraction of price changes that are small (which suggests that menu costs are small). On the other hand, time-dependent theories of price rigidity (e.g. Calvo 1983 and Taylor 1980) cannot account for the effect of inflation on the frequency of price adjustment, because this frequency is a technological parameter. To the best of our knowledge, the only theory that matches the empirical behavior of prices as well as ours is the one by Midrigan (2006), which combines elements of state-dependent and time-dependent models. However, Midrigan's theory has very different implications than ours. In Midrigan's model money is not neutral. Hence, the central bank can engineer an expansion by injecting money into the economy. In our model, money is neutral and monetary policy cannot be used to generate short-run outbursts of production and consumption.

5 Conclusions

We have constructed a theory of nominal price stickiness that is consistent with the empirical evidence. Our theory does not impose ad hoc restrictions on firms' pricing decisions – they are free to change their prices when they like at no cost. Instead, it relies on the existence of standard search frictions in goods markets, which give rise to equilibrium price dispersion. Sticky prices are an obvious corollary of price dispersion. This would be true of relative prices in an economy with perfect credit, in which some individual sellers may keep the same price, measured in terms of a numeraire commodity, in the face of shocks to supply and demand conditions (i.e., technology and preferences). Since many of the claims and arguments in macroeconomics concern monetary matters, we develop the idea in the context of a monetary economy, and show it is also true that individual sellers may keep the same price, measured in terms of dollars, in the face of shocks not only to supply and demand but also to changes in the money supply.

Contrary to what one sees in other models that claim to be consistent with the facts, based on Calvo or Mankiw pricing, in our framework the monetary authority cannot engineer a boom or get us out of a slump simply by issuing currency. We prove that money is neutral in the model, although not superneutral (inflation can have real effects on prices and allocations). Of course, a model is only a model. We did not prove that a central bank could not engineer a real expansion merely by printing money in the real world. And, we did not in this paper get into the question of, even if they could, does this mean they should? All we did was to show that the observation of sticky nominal prices does not logically imply that we need menu costs, or related devices, in our models to match the facts, and that this observation also does not rationalize particular policy prescriptions as some people seem to believe.

References

[1] Aruoba, B, G. Rocheteau, and C. Waller. 2007. "Bargaining and the Value of Money", *Journal of Monetary Economics* 54: 2636–55.

[2] Baharad, E., and B. Eden. 2004. "Price Rigidity and Price Dispersion: Evidence From Micro Data." *Review of Economic Dynamics* 7: 613–41.

[3] Ball, L., and G. Mankiw. 1994. "A Sticky-Price Manifesto." NBER Working Paper 4677.

[4] Berentsen, A., G. Menzio, and R. Wright. 2010. "Inflation and Unemployment in the Long Run." *American Economic Review*.

[5] Bils, M., and P. Klenow. 2005. "Some Evidence on the Importance of Sticky Prices." *Journal of Political Economy* 112(5): 947-85.

[6] Burdett, K., and K. Judd. 1983. "Equilibrium Price Dispersion." *Econometrica* 51: 955-70.

[7] Campbell, J., and B. Eden. 2007. "Rigid Prices: Evidence from U.S. Scanner Data." Working Paper 2005-08, Federal Reserve Bank of Chicago, Chicago.

[8] Calvo, G. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics*, 12: 383–398.

[9] Caplin, A., and D. Spulber. 1987. "Menu Costs and the Neutrality of Money." *Quarterly Journal of Economics*, 102: 703–725.

[10] Carlton, D. 1986, "The Rigidity of Prices," *American Economic Review* 76, 637-658.

[11] Cecchetti, S. 1985. "Staggered Contracts and the Frequency of Price Adjustment," *Quarterly Journal of Economics*, 100: 935–959.

[12] Clarida, R., Gali, J., and Gertler, M. 1999. "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature* 37, 1661-1707.

[13] Dotsey, M., R. King, and A. Wolman. 1999. "State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output," *Quarterly Journal of Economics*, 114: 655–690.

[14] Eden, B. 1994. "The Adjustment of Prices to Monetary Shocks when Trade is Uncertain and Sequential." *Journal of Political Economy* 102(3): 493–509.

[15] Eichenbaum, M., N. Jaimovich, and S. Rebelo. 2009. "References Prices, Costs, and Nominal Rigidities," manuscript Northwestern University.

[16] Faig, M., and B. Jerez. 2005. "A Theory of Commerce." *Journal of Economic Theory*, 122: 60-99.

[17] Gagnon, E. 2007. "Price Setting During Low and High Inflation: Evidence from Mexico." International Finance Discussion Paper 2007-896, The Federal Reserve Board.

[18] Galenianos, M., and P. Kircher. 2007. "A Model of Money with Multilateral Matching." *Journal of Monetary Economics.*

[19] Golosov, M., and R.E. Lucas, Jr. 2003. Menu Costs and Phillips Curves", NBER Working Paper 10187.

[20] Golosov, M., and R.E. Lucas, Jr. 2007. "Menu Costs and Phillips Curves", *Journal of Political Economy*, 115(2): 171-99.

[21] Head, A., and A. Kumar. 2005. "Price Dispersion, Inflation, and Welfare." *International Economic Review* 46(2): 533-72.

[22] Head, A., A. Kumar, and B. Lapham. 2008. "Market Power, Price Adjustment, and Inflation." *International Economic Review.*

[23] Jean, K., S. Rabinovich, and R. Wright. 2010. "On the Multiplicity of Monetary Equilibria: Green-Zhou Meets Lagos-Wright." *Journal of Economic Theory.*

[24] Klenow, P., and O. Kryvtsov. 2008. "State-Dependent or Time-Dependent Pricing: Does it Matter for Recent U.S. Inflation?" *Quarterly Journal of Economics* 123(3): 863–904.

[25] Klenow, P., and B. Malin. 2010. "Microeconomic Evidence on Price-Setting," *Handbook of Monetary Economics*, forthcoming.

[26] Lagos, R., and R. Wright. 2005. "A Unified Framework for Monetary Theory and Policy Analysis." *Journal of Political Economy*, 113(2): 463–88.

[27] Lucas, R.E. 1972. "Expectations and the Neutrality of Money", *Journal of Economic Theory* 4: 103-124.

[28] Mankiw, G. 1985. "Small Menu Costs and Large Business Cycles: A Macroeconomic Model," *Quarterly Journal of Economics* 100, 529-38.

[29] McCall, J. 1970. "Economics of Information and Job Search," *Quarterly Journal of Economics* 84, 113-26.

[30] Mortensen, D. 1970. "A Theory of Wage and Employment Dynamics," in *Microeconomic Foundations of Employment and Inflation Theory*, E.S. Phelps, ed. New York: W.W. Norton.

[31] Midrigan, V. 2006. "Menu Costs, Multi-Product Firms, and Aggregate Fluctuations." Manuscript, New York University.

[32] Nakamura, E., and J. Steinsson. "Five Facts About Prices: A Reevaluation of Menu Cost Models." *Quarterly Journal of Economics* 123: 1415-1464.

[33] Nosal, E., and Rocheteau, G. 2010. *Money, Payments, and Liquidity*, MIT Press, forthcoming.

[34] Rocheteau, G., and R. Wright. 2005. "Money in Search Equilibrium, in Competitive Equilibrium, and in Competitive Search Equilibrium", *Econometrica* 73: 175–202.

[35] Shi, S. 1997. "A Divisible Model of Fiat Money," *Econometrica* 65: 75-102.

[36] Tarski, A. 1955. "A Lattice Theoretical Fixpoint Theorem and Its Applications," *Pacific Journal of Mathemathics* 5: 285–309.

[37] Taylor, J. 1980. "Aggregate Dynamics and Staggered Contracts," *Journal of Political Economy*, 88: 1–23.

[38] Williamson, S., and R. Wright. 2010a. "New Monetarist Economics: Methods," Manuscript, University of Winsconsin-Madison.

[39] Williamson, S., and R. Wright. 2010b. "New Monetarist Economics: Models," *Handbook of Monetary Economics*, forthcoming.

[40] Woodford, M. 2003. *Interest and Prices*, Princeton University Press, Princeton NJ.

Appendix

A Proof of Lemma 1

We take five steps to prove the lemma.

Claim 1: The firm's maximized profits, Π_t^*, are strictly positive.

Proof: For any $\lambda > 1$, the profits of a firm that posts the price $\lambda c w_t$ are given by

$$\begin{aligned}\Pi_t(\lambda c w_t) &= \frac{1}{s}\left[\alpha_1 + 2\alpha_2\left(1 - F_t(\lambda c w_t)\right) + \alpha_2 n_t(\lambda c w_t)\right] q_t^*(\lambda c w_t)(\lambda - 1) c w_t \\ &> \frac{\alpha_1}{s} q_t^*(\lambda c w_t)(\lambda - 1) c w_t.\end{aligned}$$

Since $q_t^*(\lambda c w_t) > 0$ and $\lambda > 1$, the previous inequality implies $\Pi_t(\lambda c w_t) > 0$. In turn, since $\Pi_t^* \geq \Pi_t(\lambda c w_t)$, we have $\Pi_t^* > 0$.

Claim 2: The price distribution F_t is continuous.

Proof: On the way to a contradiction, suppose that there exists a price p_0 such that $n_t(p_0) > 0$. The profits of a firm that posts p_0 are given by

$$\Pi_t(p_0) = \frac{1}{s}\left[\alpha_1 + 2\alpha_2\left(1 - F_t(p_0)\right) + \alpha_2 n_t(p_0)\right] R_t(p_0).$$

Now, consider any $p_1 < p_0$ such that (i) $R_t(p_1)$ is strictly positive and (ii) $\Delta \equiv R_t(p_0) - R_t(p_1)$ is strictly smaller than $\alpha_2 n_t(p_0) R_t(p_0)/(\alpha_1 + 2\alpha_2)$. A price p_1 that satisfies conditions (i) and (ii) exists because the function $R_t(p)$ is continuous in p and, since $\Pi_t(p_0) = \Pi_t^* > 0$, $R_t(p_0)$ is strictly positive. The profits of a firm that posts p_1 are given by

$$\begin{aligned}\Pi_t(p_1) &= \frac{1}{s}\left[\alpha_1 + 2\alpha_2\left(1 - F_t(p_1)\right) + \alpha_2 n_t(p_1)\right] R_t(p_1) \\ &\geq \frac{1}{s}\left[\alpha_1 + 2\alpha_2\left(1 - F_t(p_0)\right) + 2\alpha_2 n_t(p_0)\right]\left[R_t(p_0) - \Delta\right] \\ &\geq \Pi_t(p_0) + \alpha_2 n_t(p_0)\left[R_t(p_0) - \Delta\right] - (\alpha_1 + 2\alpha_2)\Delta\end{aligned} \tag{A.1}$$

Since $R_t(p_0) - \Delta > 0$ and $\Delta < \alpha_2 n_t(p_0) R_t(p_0)/(\alpha_1 + 2\alpha_2)$, it follows from (A.1) that $\Pi_t(p_1) > \Pi_t(p_0)$. However, from the fact that p_0 belongs to the support of the price distribution F_t, it follows that $\Pi_t(p_0) = \Pi_t^* \geq \Pi_t(p_1)$. Hence, we have reached a contradiction.

Claim 3: The monopoly price, p_t^m, is the highest price on the support of the distribution

F_t.

Proof: On the way to a contradiction, suppose that the highest price on the support of the distribution is $\overline{p}_t \neq p_t^m$. The profits of a firm that posts the price $\overline{p}_t$ are given by

$$\Pi_t\left(\overline{p}_t\right) = \frac{\alpha_1}{s} R_t(\overline{p}_t). \tag{A.2}$$

The profits of a firm that posts the price p_t^m are given by

$$\begin{aligned}\Pi_t\left(p_t^m\right) &= \frac{1}{s}\left[\alpha_1 + 2\alpha_2\left(1 - F_t\left(p_t^m\right)\right)\right] R_t(p_t^m) \\ &\geq \frac{\alpha_1}{s} R_t(p_t^m) \\ &> \frac{\alpha_1}{s} R_t(\overline{p}_t),\end{aligned} \tag{A.3}$$

where the third line makes use of the fact that p_t^m is the unique maximizer of R_t. From (A.2) and (A.3), it follows that $\Pi_t\left(p_t^m\right) > \Pi_t\left(\overline{p}_t\right)$. However, from the fact that $\overline{p}_t$ belongs to the support of the price distribution F_t, it follows that $\Pi_t\left(\overline{p}_t\right) = \Pi_t^* \geq \Pi_t\left(p_t^m\right)$. Hence, we have reached a contradiction.

Claim 4: The support of the price distribution F_t is connected.

Proof: On the way to a contradiction, suppose that p_0 and p_1 are two prices on the support of the distribution such that $p_0 < p_1$ and $F_t(p_0) = F_t(p_1)$. Since p_0 and p_1 belong to the support of the price distribution, they are greater than cw_t and smaller than p_t^m. Hence, $cw_t < p_0 < p_1 \leq p_1$. The profits of a firm that posts the price p_0 are given by

$$\Pi_t\left(p_0\right) = \frac{1}{s}\left[\alpha_1 + 2\alpha_2\left(1 - F_t\left(p_0\right)\right)\right] R_t(p_0).$$

The profits of a firm that posts the price p_1 are given by

$$\Pi_t\left(p_1\right) = \frac{1}{s}\left[\alpha_1 + 2\alpha_2\left(1 - F_t\left(p_1\right)\right)\right] R_t(p_1).$$

First, note that $\alpha_1 + 2\alpha_2\left(1 - F_t\left(p_1\right)\right)$ is equal to $\alpha_1 + 2\alpha_2\left(1 - F_t\left(p_0\right)\right)$ because $F_t\left(p_1\right) = F_t\left(p_0\right)$. Second, note that $R_t(p_1) > R_t(p_0)$ because $cw_t < p_0 < p_1 \leq p_1$ and the function $R_t(p)$ is strictly increasing for all $p \in [cw_t, p_t^m]$. From these observations, it follows that $\Pi_t\left(p_1\right) > \Pi_t\left(p_0\right)$. However, from the fact that p_0 and p_1 belong to the support of the price distribution F_t, it follows that $\Pi_t\left(p_0\right) = \Pi_t\left(p_1\right) = \Pi_t^*$. Hence, we have reached a contradiction.

Claim 5: The price distribution F_t is given by

$$F_t(p) = 1 - \frac{\alpha_1}{2\alpha_2}\left[\frac{R_t(p_t^m)}{R_t(p)} - 1\right]. \tag{A.4}$$

Proof: Since the price distribution has no mass points, the profit of a firm that charges the price p is given by

$$\Pi_t(p) = \frac{1}{s}\left[\alpha_1 + 2\alpha_2(1 - F_t(p))\right] R_t(p).$$

Since the monopoly price is on the support of the price distribution, the maximized profit of the firm is given by

$$\Pi_t^* = \frac{\alpha_1}{s} R_t(p_t^m).$$

Finally, since every price on the support of the price distribution must maximize the profit of the firm, it follows that for every $p \in [\underline{p}_t, \overline{p}_t]$

$$\frac{1}{s}\left[\alpha_1 + 2\alpha_2(1 - F_t(p))\right] R_t(p) = \alpha_1 R_t(p_t^m). \tag{A.5}$$

Solving (A.5) with respect to F_t leads to equation (A.4).

B Proof of Theorem 3

We find it convenient to express all the nominal variables in real terms (i.e. units of labor). Specifically, we let n_t denote the real balances m_t/w_t, we let z_t denote the real price p/w_t, and let H_t denote the distribution of real prices z. Then, a stationary monetary equilibrium can be defined as a tuple $\{n^*, x^*, h^*, q^*, H^*\}$ such that:

(i) n^* and x^* and h^* solve the household's problem in the AD market:

$$i = \int_{\underline{z}(n^*)}^{\hat{z}(n^*)} \left[\frac{1}{z}\left(\frac{n^*}{z}\right)^{-\gamma} - 1\right]\left[\alpha_1 + 2\alpha_2\left(1 - H^*(z|n^*)\right)\right] dH^*(z|n^*),$$

$$v'(x^*) = 1, \qquad h^*(n) = x^* + \left(\mu^{-1}n^* - n - D - T\right);$$

(ii) q^* solves the household's problem in the BJ market:

$$q^*(z,n) = \begin{cases} n/z, & \text{if } z \leq \hat{z}(n), \\ (1/z)^{\frac{1}{\gamma}}, & \text{if } z > \hat{z}(n); \end{cases}$$

(iii) H^* is consistent with the solution to the firm's problem in the BJ market:

$$H^*(z|n^*) = 1 - \frac{\alpha_1}{2\alpha_2}\left[\frac{R(z^m, n^*)}{R(z, n^*)} - 1\right] \quad \text{all } z \in [\underline{z}(n^*), \overline{z}(n^*)],$$

where $R(z, n^*) = q^*(z, n^*)(z - c)$.

To establish the existence of a stationary monetary equilibrium, it is sufficient to prove that there exists a $\hat{n}^*$ such that the households' optimal choice of real balances is $\hat{n}^*$ when the distribution of prices posted by the firms is $H^*(z|\hat{n}^*)$. That is, it is sufficient to prove that there exists a $\hat{n}^*$ such that

$$i = \int_{\underline{z}(n^*)}^{\hat{z}(n)} \left[\frac{1}{z}\left(\frac{n}{z}\right)^{-\gamma} - 1\right]\left[\alpha_1 + 2\alpha_2\left(1 - H^*(z|n^*)\right)\right] dH^*(z|n^*),$$

for $n = \hat{n}^*$. We carry out this task in three steps.

Claim 1: Let $\underline{n}^*$ and $\overline{n}^*$ be defined as

$$\hat{z}(\underline{n}^*) = \frac{c}{1-\gamma}, \ \ \hat{z}(\overline{n}^*) = c\left[1 - \frac{\alpha_1}{\alpha_1 + 2\alpha_2}\frac{1}{n^*}\left(\frac{c}{1-\gamma}\right)^{-\frac{1}{\gamma}}\left(\frac{\gamma c}{1-\gamma}\right)\right]^{-1}.$$

For all n_0^* and n_1^* such that $0 < n_0^* < n_1^* \leq \overline{n}^*$, the price distribution $H(z|n_0^*)$ first order stochastically dominates the price distribution $H(z|n_1^*)$. For all n_0^* and n_1^* such that $\overline{n}^* \leq n_0^* < n_1^*$, the price distribution $H(z|n_0^*)$ is equal to the price distribution $H(z|n_1^*)$.

Proof: For all $n^* \in (0, \underline{n}^*]$, the price distribution is given by

$$H(z|n^*) = 1 - \frac{\alpha_1}{2\alpha_2}\left[\frac{\overline{z}(n^*)^{-1} n^*\left(\overline{z}(n^*) - c\right)}{z^{-1} n^*\left(z - c\right)} - 1\right]$$

and its support is given by $[\underline{z}(n^*), \overline{z}(n^*)]$, where

$$\overline{z}(n^*) = n^{*-\frac{\gamma}{1-\gamma}},$$

$$\underline{z}(n^*) = c\left[1 - \frac{\alpha_1}{\alpha_1 + 2\alpha_2}\left(\frac{\overline{z}(n^*) - c}{\overline{z}(n^*)}\right)\right]^{-1}.$$

Consider any n_0^* and n_1^* such that $0 < n_0^* < n_1^* \leq \overline{n}^*$. Clearly, $\overline{z}(n_0^*) > \overline{z}(n_1^*)$ and $\underline{z}(n_0^*) > \underline{z}(n_1^*)$. For $z \geq \overline{z}(n_0^*)$, $H(z|n_0^*) = H(z|n_1^*) = 1$. For $z \in (\underline{z}(n_1^*), \overline{z}(n_0^*))$, $H(z|n_0^*) < H(z|n_1^*)$ because $\overline{z}(n_0^*) > \overline{z}(n_1^*)$. For $z \leq \underline{z}(n_1^*)$, $H(z|n_0^*) = H(z|n_1^*) = 0$. Hence, $H(z|n_0^*)$ first order stochastically dominates $H(z|n_1^*)$.

For all $n^* \in [\underline{n}^*, \overline{n}^*]$, the price distribution is given by

$$H(z|n^*) = \begin{cases} 1 - \dfrac{\alpha_1}{2\alpha_2}\left[\dfrac{\overline{z}(n^*)^{-\frac{1}{\gamma}}\left(\overline{z}(n^*) - c\right)}{z^{-\frac{1}{\gamma}}\left(z - c\right)} - 1\right], & \text{if } z \in [\hat{z}(n^*), \overline{z}(n^*)], \\ 1 - \dfrac{\alpha_1}{2\alpha_2}\left[\dfrac{\overline{z}(n^*)^{-\frac{1}{\gamma}}\left(\overline{z}(n^*) - c\right)}{z^{-1}n^*\left(z - c\right)} - 1\right], & \text{if } z \in [\underline{z}(n^*), \hat{z}(n^*)] \end{cases}$$

and its support is given by $[\underline{z}(n^*), \overline{z}(n^*)]$, where

$$\overline{z}(n^*) = \frac{c}{1-\gamma},$$

$$\underline{z}(n^*) = c\left[1 - \frac{\alpha_1}{\alpha_1 + 2\alpha_2}\frac{\overline{z}(n^*)^{-\frac{1}{\gamma}}\left(\overline{z}(n^*) - c\right)}{n^*}\right]^{-1}.$$

Consider any n_0^* and n_1^* such that $\underline{n}^* \leq n_0^* < n_1^* \leq \overline{n}^*$. Clearly, $\overline{z}(n_0^*) = \overline{z}(n_1^*)$, $\hat{z}(n_0^*) > \hat{z}(n_1^*)$ and $\underline{z}(n_0^*) > \underline{z}(n_1^*)$. For $z \geq \hat{z}(n_0^*)$, $H(z|n_0^*) = H(z|n_1^*)$. For $z \in [\hat{z}(n_1^*), \hat{z}(n_0^*))$, $H(z|n_0^*) < H(z|n_1^*)$ because $z^{-1}n_0^* < z^{-\frac{1}{\gamma}}$. For $z \in [\underline{z}(n_0^*), \hat{z}(n_1^*))$, $H(z|n_0^*) < H(z|n_1^*)$ because $n_0^* < n_1^*$. For $z \leq \underline{z}(n_0^*)$, $H(z|n_0^*) \leq H(z|n_1^*)$. Hence, $H(z|n_0^*)$ first order stochastically dominates $H(z|n_1^*)$.

For all $n^* \geq \overline{n}^*$, the price distribution is given by

$$H(z|n^*) = 1 - \frac{\alpha_1}{2\alpha_2}\left[\frac{\overline{z}(n^*)^{-\frac{1}{\gamma}}\left(\overline{z}(n^*) - c\right)}{z^{-\frac{1}{\gamma}}\left(z - c\right)} - 1\right]$$

and its support is given by $[\underline{z}(n^*), \overline{z}(n^*)]$, where

$$\overline{z}(n^*) = \frac{c}{1-\gamma},$$

$$\underline{z}(n^*) = c + \frac{\alpha_1}{\alpha_1 + 2\alpha_2}\left(\frac{\underline{z}(n^*)}{\overline{z}(n^*)}\right)^{-\frac{1}{\gamma}}\left(\overline{z}(n^*) - c\right).$$

In this case, the price distribution does not depend on n^*. Hence, $H(z|n_0^*) = H(z|n_0^*)$ for all n_0^* and n_1^* such that $\overline{n}^* \leq n_0^* < n_1^*$.

Claim 2: Given the price distribution $H(z|n^*)$, the unique solution for n to the equilibrium condition (i) is $\psi(n^*)$. The solution $\psi(n^*)$ has the following properties:

(i) For all n_0^*, n_1^* such that $0 < n_0^* < n_1^* \leq \overline{n}^*$, $\psi(n_0^*) \leq \psi(n_1^*)$;

(ii) For all n_0^*, n_1^* such that $\overline{n}^* \leq n_0^* < n_1^*$, $\psi(n_0^*) = \psi(n_1^*)$;

(iii) For all $n^* > 0$, $\psi(n^*) \in [\underline{\psi}, \overline{\psi})$, where $\underline{\psi} > 0$ and $\overline{\psi} = \overline{n}^*$.

Proof: Given the price distribution $H(z|n^*)$, the equilibrium condition (i) is

$$\begin{aligned} i = \quad & \alpha_1 \int_{\underline{z}(n^*)}^{\hat{z}(n)} \left[\left(\frac{z}{n}\right)^{\gamma} \frac{1}{z} - 1 \right] dH(z|n^*) \\ & + \alpha_2 \int_{\underline{z}(n^*)}^{\hat{z}(n)} \left[\left(\frac{z}{n}\right)^{\gamma} \frac{1}{z} - 1 \right] d\left[1 - (1 - H(z|n^*))^2\right]. \end{aligned}$$

Let $\phi(n, n^*)$ denote the right-hand side of the equation above. First, notice that $\lim_{n \to 0} \phi(n, n^*) = \infty$. Second, notice that $\phi(n, n^*)$ is strictly decreasing in n for all $n \in (0, \hat{z}^{-1}(\underline{z}(n^*)))$. Third, notice that $\phi(n, n^*) = 0$ for all $n \geq \hat{z}^{-1}(\underline{z}(n^*)) = \underline{z}(n^*)^{-\frac{1-\gamma}{\gamma}}$. From these observations, it follows that there exists a unique solution, $\psi(n^*)$, to the equation $\phi(n, n^*) = i$. Moreover, $\psi(n^*) > 0$ and $\psi(n^*) < \hat{z}^{-1}(\underline{z}(n^*))$.

Consider any n_0^* and n_1^* such that $0 < n_0^* < n_1^* \leq \overline{n}^*$. In this case, claim 1 implies that $H(z|n_0^*)$ first order stochastically dominates $H(z|n_1^*)$ and, consequently, that $1 - (1 - H(z|n_0^*))^2$ first order stochastically dominates $1 - (1 - H(z|n_1^*))^2$. From these observations and from the fact that $(z/n)^{\gamma}/z - 1$ is decreasing in z, it follows that $\phi(n, n_0^*) \leq \phi(n, n_1^*)$. Therefore, $\psi(n_0^*) \leq \psi(n_1^*)$. Moreover, it is straightforward to verify that $\psi(n_0^*) \geq \underline{\psi} > 0$.

Now, consider any n_0^* and n_1^* such that $\overline{n}^* \leq n_0^* < n_1^*$. In this case, claim 1 implies that $H(z|n_0^*)$ is equal to $H(z|n_0^*)$ and that $1 - (1 - H(z|n_0^*))^2$ is equal to $1 - (1 - H(z|n_1^*))^2$. From these observations, it follows that $\phi(n, n_0^*) = \phi(n, n_1^*)$ and, hence, $\psi(n_0^*) = \psi(n_1^*)$. Moreover, it is straightforward to verify that $\psi(n_1^*) < \overline{n}^*$.

Claim 3: There exists a $\hat{n}^* \in [\underline{\psi}, \overline{\psi})$ such that $\psi(\hat{n}^*) = \hat{n}^*$.

Proof: Claim 2 implies that: (a) $\psi(n^*)$ is an increasing function of n^*, (b) for all n^* in the interval $[\underline{\psi}, \overline{\psi}]$, $\psi(n^*)$ belongs to the interval $[\underline{\psi}, \overline{\psi})$. Hence, from Tarski's fixed point theorem, it follows that there exists a $\hat{n}^* \in [\underline{\psi}, \overline{\psi})$ such that $\psi(\hat{n}^*) = \hat{n}^*$. Moreover, claim 2 implies $\psi(n^*) \geq \underline{\psi}$ for all $n^* < \underline{\psi}$, and $\psi(n^*) < \overline{\psi}$ for all $n^* \geq \overline{\psi}$. Hence, there is no $\hat{n}^* \notin [\underline{\psi}, \overline{\psi})$ such that $\psi(\hat{n}^*) = \hat{n}^*$.

C Proof of Theorem 6

We take two steps to prove Theorem 6.

Claim 1: In any stationary monetary equilibrium, Σ^*,

$$\frac{\overline{p}_t - \underline{p}_t}{\overline{p}_t} \geq \Delta > 0, \tag{A.6}$$

where

$$\Delta = \frac{2\alpha_2\gamma}{2\alpha_2 + \alpha_1(1-\gamma)}.$$

Proof: Let Σ^* be an arbitrary stationary monetary equilibrium. From condition (1) in the definition of equilibrium, it follows that households are either cash constrained in all transactions in the BJ market, or they are cash constrained in a positive fraction of transactions and not constrained in the others.

First, suppose Σ^* is such that households are cash constrained in all transactions. In this case, we have

$$\overline{p}_t = w_t^* \left(\frac{w_t^*}{m_t^*}\right)^{\frac{\gamma}{1-\gamma}} \geq \frac{cw_t^*}{1-\gamma}, \tag{A.7}$$

$$\underline{p}_t = \frac{cw_t^*}{\left[1 - \dfrac{\alpha_1}{\alpha_1 + 2\alpha_2}\dfrac{\overline{p}_t - cw_t^*}{\overline{p}_t}\right]}. \tag{A.8}$$

From (A.7) and (A.8), it follows that

$$\begin{aligned} \frac{\overline{p}_t - \underline{p}_t}{\overline{p}_t} &= \frac{2\alpha_2(\overline{p}_t - cw_t^*)}{2\alpha_2\overline{p}_t + \alpha_1 cw_t^*} \\ &\geq \frac{2\alpha_2\gamma}{2\alpha_2 + \alpha_1(1-\gamma)} = \Delta, \end{aligned} \tag{A.9}$$

where the second line makes use of the fact that $\overline{p}_t \geq (1-\gamma)^{-1}cw_t^*$.

Next, suppose Σ^* is such that households are cash constrained in some but not all transactions. In this case, we have

$$\overline{p}_t = \frac{cw_t^*}{1-\gamma} \geq w_t^* \left(\frac{w_t^*}{m_t^*}\right)^{\frac{\gamma}{1-\gamma}}, \tag{A.10}$$

$$\underline{p}_t = \frac{cw_t^*}{\left[1 - \dfrac{\alpha_1}{\alpha_1 + 2\alpha_2}\left(\dfrac{\overline{p}_t}{w_t^*}\right)^{-\frac{1}{\gamma}}\left(\dfrac{\overline{p}_t - cw_t^*}{m_t^*}\right)\right]}. \tag{A.11}$$

From (A.10) and (A.11), it follows that

$$\begin{aligned}\frac{\overline{p}_t - \underline{p}_t}{\overline{p}_t} &= \frac{(\alpha_1 + 2\alpha_2)\gamma m_t^* - \alpha_1 \gamma \left(\frac{c}{1-\gamma}\right)^{\frac{\gamma-1}{\gamma}} w_t^*}{(\alpha_1 + 2\alpha_2) m_t^* - \alpha_1 \gamma \left(\frac{c}{1-\gamma}\right)^{\frac{\gamma-1}{\gamma}} w_t^*} \\ &\geq \frac{2\alpha_2 \gamma}{2\alpha_2 + \alpha_1(1-\gamma)} = \Delta,\end{aligned} \tag{A.12}$$

where the second line makes use of the fact that $m_t^* \geq [(1-\gamma)/c]^{\frac{1-\gamma}{\gamma}} w_t^*$. Combining (A.9) and (A.12), we obtain (A.6).

Claim 2: Let $\mu \in (1, \mu^*)$, where $\mu^* = (1-\Delta)^{-1}$. Then, for any stationary monetary equilibrium Σ^* together with the pricing policy $p_{t+1}^*(p, \rho)$, we have

$$\begin{array}{ll} FR(\rho) < 1 \text{ and } A(\rho) > 0, & \text{all } \rho \in (0,1], \\ H_t(0,\rho) > 0, & \text{all } \rho \in [0,1). \end{array}$$

Proof: From claim 1 and $\mu \in (1, \mu^*)$, it follows that

$$\mu \underline{p}_t < \frac{1}{1-\Delta} \underline{p}_t \leq \overline{p}_t. \tag{A.13}$$

For any $\rho \in (0,1]$, the fraction of prices that adjust is

$$\begin{aligned} FR(\rho) &= F_t^*(\mu \underline{p}_t) + (1-\rho)\left[1 - F_t^*(\mu \underline{p}_t)\right] \\ &< F_t^*((1-\Delta)^{-1}\underline{p}_t) + (1-\rho)\left[1 - F_t^*((1-\Delta)^{-1}\underline{p}_t)\right] \\ &< 1, \end{aligned}$$

where the second line makes use of (A.13), and the third line makes use of the fact that $F_t^*((1-\Delta)^{-1}\underline{p}_t) < F_t^*(\overline{p}_t) = 1$. Since the fraction of prices that adjust is less than 1, the average duration of a price must be greater than 1 period. That is, $A(\rho) > 1$. Finally, it is straightforward to verify that for all $\rho \in [0,1)$, the fraction of negative price changes, $H_t(0,\rho)$, is strictly positive.