

Federal Reserve Bank of Minneapolis
Research Department

Sticky Prices: A New Monetarist Approach*

Allen Head, Lucy Qian Liu,
Guido Menzio, and Randall Wright

Working Paper 690

October 2011

*Head: Queen's University; Liu: International Monetary Fund; Menzio: University of Pennsylvania and NBER; Wright: University of Wisconsin-Madison, Federal Reserve Bank of Chicago, Federal Reserve Bank of Minneapolis, and NBER. This paper was delivered by Wright as the Marshall Lecture at the 2011 joint meetings of the Econometric Society and European Economic Association in Oslo, Norway. It was invited by the President of the Society, Chris Pissarides. We are honored and grateful for his support of our ongoing work, sufficiently so that we are submitting the flagship paper in the project, rather than a spinoff or summary (for the record, this paper has never been submitted for publication elsewhere). A previous version was titled "Really, Really Rational Inattention: Or How I Learned to Stop Worrying and Love Sticky Prices." We are grateful for input from Marios Angeletos, Mark Bilal, Ken Burdett, Jeff Campbell, V.V. Chari, Peter Diamond, Ben Eden, Chris Edmond, Mike Golosov, Christian Hellwig, Boyan Jovanovic, Peter Klenow, Oleksiy Kryvtsov, John Leahy, Dale Mortensen, Julio Rotemberg, Tom Sargent and Nancy Stokey. We also thank participants in presentations at MIT, Wisconsin, Yale, Ohio State, York, UCSD, BU, Rochester, Vienna, Paris II, NYU, the Federal Reserve Banks of Chicago, Minneapolis, Richmond, San Francisco and St. Louis, the Bank of Canada, the SED, the Canadian Macro Study Group, the Philadelphia Search and Matching Workshop, the Philly Fed - Penn Econ Conference on Money and Macro Economics, the Essex Conference on Economics and Music, the Minnesota Macro Workshop and the NBER. Wright thanks the National Science Foundation and the Ray Zemon Chair in Liquid Assets at the Wisconsin School of Business. Head thanks the Social Sciences and Humanities Research Council of Canada. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

ABSTRACT

Why do some sellers set nominal prices that apparently 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 the money supply increases, some sellers may keep prices constant, earning less per unit but making it up on volume, so profit stays constant. The calibrated model matches price-change data well. But, in contrast with other sticky-price models, money is neutral.

JEL numbers: E52, E31, E42

Keywords: Money; Sticky Prices; Monetary Policy; Neutrality

1 Introduction

Arguably the most difficult question in macroeconomics is this: Why do some sellers set prices in nominal terms that apparently do not adjust in response to changes in the aggregate price level. This seems to fly in the face of elementary microeconomic principles. Shouldn't every seller have a target relative price, depending on real factors, and therefore when the aggregate nominal price level increases by some amount, say due to an increase in the money supply, shouldn't every seller necessarily adjust his nominal price by the same amount? In many popular macro models, including those used by most policy makers, prices are sticky *by assumption*, in the sense that there are either restrictions on how often they can change, following Taylor (1980) or Calvo (1983), or there are real resource costs to changing them, following Mankiw (1985). We deliver stickiness *as a result*, in the sense that sellers set prices in nominal terms, and some may choose not to adjust in response to changes in the aggregate price level, even though we let them change whenever they like and at no cost. Moreover, in contrast with other theories with sticky prices, we construct our model so that money is neutral: the central bank cannot engineer a boom or end a slump simply by issuing currency. Hence, while our theory provides microfoundations for the core ingredient in Keynesian economics – sticky prices or nominal rigidities – it has rather different policy implications.

We emphasize at the outset that our objective here is *not* to establish that monetary policy is neutral or nonneutral in the real world. That is beside the point. Our objective is to show formally two other results: (1) one does not need to introduce technological restrictions or costs, as in Calvo- or Mankiw-style models, to generate price stickiness; and (2) the appearance of nominal rigidities in the real world does not logically imply that policy can exploit these rigidities, as some economists think. To explain our motivation by analogy to a rather famous paper, Lucas (1972) describes a microfounded monetary model consistent with the observation that, in the data, there is a positive correlation between the aggregate price level (or money supply) and output (or employment), but policy cannot systematically exploit the relationship. That is, increasing inflation by printing money at a faster rate will not increase average output or employment. We think this was a good lesson. We similarly

want to show that one can write down a microfounded monetary model consistent with some other observations, those concerning nominal price adjustments, but it is not possible for policy to exploit this. Monetary policy is neutral in the model *by design* – this is how we make point (2) above, that price stickiness does not logically imply nonneutrality.¹

Not only does our model provide counterexamples to some popular beliefs about monetary theory and policy, we also argue that it is empirically reasonable in the following sense. We show that our approach to price stickiness is successful, relative to alternative theories, at matching the salient features of the micro data on individual price dynamics.² We can account for the average duration of prices in the data, for the fact that price changes are large on average, even though many changes are small, and that prices change more frequently (and not just by larger amounts) when inflation is higher. In contrast, simple menu-cost theories cannot easily account for the second fact – that on average price changes are large even though many changes are small – while Calvo theories cannot easily account for the third – that the frequency of price changes increases with inflation. It is not our claim that somehow complicating or integrating existing theories cannot work, and there are some reasonably successful attempts in the literature, including e.g. Midrigan (2006). Our claim is that even a very basic version of our theory does a good job matching the facts.

We think these findings are relevant for several reasons. First, despite the successes of, say, the New Classical and Real Business Cycle paradigms, they seem to miss one basic feature of the data: at least some nominal prices seem sticky in the sense defined above (they do not respond to changes in the aggregate price level). One should want to know if this somehow invalidates these theories or their policy implications, and means the only valid theories and recommendation emanate from a Keynesian approach. It seems clear to us that the observation of price stickiness is one of main reasons why many Keynesians are Keynesian. Consider Ball and Mankiw (1994), who we think representative, when they say: “We believe that sticky prices provide the most natural explanation of monetary nonneutrality since so

¹To be clear, it is not the case that monetary policy in the model has no impact: changing the inflation or nominal interest rate has real consequences, as in any good monetary model, but this has nothing to do with nominal rigidities.

²Empirical work on price stickiness includes, e.g., 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 survey.

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.” And, “As a matter of logic, nominal stickiness requires a cost of nominal adjustment.” Some others that one might not think of as Keynesian present similar positions, including Golosov and Lucas (2003), who argue 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.”³

We interpret the above claims as containing three points related, respectively, to empirics, theory, and policy. The first claim is that price stickiness is a fact. The quotations assert this, and it is substantiated by numerous empirical studies. We concede the 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 do so by displaying 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 that this rationalizes Keynesian policy advice. We also prove this wrong. Our theory is consistent with the relevant observations, but money is neutral, which means that sticky prices simply do not constitute definitive evidence that money is nonneutral or that particular policy recommendations are warranted. To reiterate, that the point here is not about whether money is neutral in the real world, it is rather about constructing a coherent, and we think compelling, economic environment with two properties: (1) it matches the sticky-price facts; and (2) it nevertheless delivers neutrality.

³The point here is not to pick on any particular individuals, but to pick out some that apparently come from different macro camps, in order to convey a general feeling in the profession about the implications of price stickiness. Another example, from a leading macro theorist who recently became a monetary policy maker, is Kocherlakota (2009): “If the Federal Reserve injects a lot of money into the economy, then there is more money chasing fewer goods. This extra money puts upward pressure on prices. If all firms changed prices continuously, then this upward pressure would manifest itself in an immediate jump in the price level. But this immediate jump would have little effect on the economy. Essentially, such a change would be like a simple change of units (akin to recalculating distances in inches instead of feet). In the real world, though, firms change prices only infrequently. It is impossible for the increase in money to generate an immediate jump in the price level. Instead, since most prices remain fixed, the extra money generates more demand on the part of households and in that way generates more production. Eventually, prices adjust, and these effects on demand and production vanish. But infrequent price adjustment means that monetary policy can have short-run effects on real output.”

It is clear that the issues at hand concern monetary phenomena: Why are prices quoted in dollars? Why do they not all adjust to changes in the money supply? What does this imply about central bank policy? To study these questions, it seems natural to use a monetary model. We work with a version of the New Monetarist framework, recently surveyed by Williamson and Wright (2010*a,b*) and Nosal and Rocheteau (2011). This framework tries to be explicit about details of the trading process, so that one can ask, who trades with whom, and how? Thus, specialization and search frictions can limit barter, while commitment and information frictions can limit credit, making money essential for at least some exchange. Because the points we make are really quite general, we could also make them with other monetary models, including cash-in-advance, money-in-the-utility-function or overlapping-generations models, but we think a search-based approach is useful, for several reasons. First, it is the approach used by most people doing monetary theory (if not monetary policy) these days. Also, the framework has proved very tractable and easily generalizable in a variety of other applications. And, more significantly, a search-based approach not only can generate a role for money, it can generate endogenous price dispersion, which is an important element of our theory.

To explain this idea, first note that many New Keynesian models, such as those described in Clarida, Gali and Gertler (1999) or Woodford (2003), generate price dispersion if and only if there is inflation. Suppose in a stationary real environment a number of sellers set the same p_t at date t . Then, at $t_1 > t$, some seller is the first one allowed to change price and changes it to p_{t_1} , at date $t_2 > t_1$ a second seller is allowed to change, etc. This induces price dispersion if and only if inflation is not zero. But the data suggest that there is price dispersion even during periods of low or zero inflation (something we first noticed in Campbell and Eden 2007). This suggests that it is important to work with models that can deliver price dispersion even without inflation. There are several candidate models, including Varian (1980), Albrecht and Axel (1984) or Stahl (1989), but we use Burdett and Judd (1983). In Burdett-Judd models, search frictions deliver price dispersion, and these same frictions help generate a role for money. Hence, it is parsimonious in terms of assumptions to use a search-based framework for the issues at hand. Burdett-Judd has also proved useful in other applications, including the large literature on labor markets following Burdett and

Mortensen (1998) (see Mortensen and Pissarides 1999 for a survey).

To understand Burdett-Judd, it helps to give a very brief history of search theory. The earliest models of McCall (1970) and Mortensen (1970) were partial equilibrium models, in the sense that they characterized the optimal search strategy of a searcher taking as given the distribution of prices (or wages in labor applications) posted by firms, and were soundly criticized on this point (e.g., Rothschild 1973). Diamond (1971) set out to build a general equilibrium search model in which the price distribution was derived endogenously: first firms post prices, taken as given the prices of others; then individuals search over these firms as in the standard theory. What he found is that there is a unique equilibrium and it entails a degenerate price distribution. The proof is easy. Given any $F(p)$, individuals use a reservation price R , buying when they sample the first $p \leq R$. But then there is no reason for any firm to set anything other than $p = R$. This proves equilibrium must have a single price. Moreover, the single price turns out to be the pure monopoly price. Since there cannot be price dispersion, the result seemed bad for search theory, but also set off a wave of research on trying to generate dispersion.

The approach in Burdett and Judd (1983) is to make one, ostensibly minimal, change in the standard sequential search model: rather than sampling prices one at a time, suppose there is a positive probability of sampling two (or more) at once. Then it is not hard to see that equilibrium must entail a nondegenerate price distribution. We are more precise when we present the formal model, but the idea is this. Suppose all sellers in some set (with positive measure) set the same p . A buyer who samples two such sellers has to use some tie-breaking rule to pick one. This gives an individual seller a big incentive to lower price, to get the sale for sure. In fact, in equilibrium, all sellers charge different prices, and one can actually derive the closed-form solution for the distribution $F(p)$. The model captures standard results as special cases: when the probability that a buyer meets two or more sellers approaches 1, we converge to a single price and it is the perfectly competitive price; and when this probability approaches 0, we converge to Diamond's monopoly price.

We embed Burdett-Judd pricing into a dynamic New Monetarist model, where agents alternate between trading in centralized and decentralized markets, and in the former market

buyers use money as a medium of exchange because frictions preclude the use of credit.⁴ In equilibrium, sellers post prices in dollars, naturally, since this is how buyers are paying. As in the baseline Burdett-Judd model, at any date t , there is a continuous distribution of prices $F_t(p)$ with nondegenerate support $[\underline{p}_t, \bar{p}_t]$. While the equilibrium pins down the distribution, it does not pin down the price of an individual seller: every seller gets the same profit from any $p_t \in [\underline{p}_t, \bar{p}_t]$, because one that posts a low price earns less per unit but makes it up on the volume. When the money supply increases from M_t to M_{t+1} , the equilibrium distribution shifts to $F_{t+1}(p)$ with support $[\underline{p}_{t+1}, \bar{p}_{t+1}]$. For this to happen, some sellers must change their prices, but not all of them: if an individual seller's price is $p_t \notin [\underline{p}_{t+1}, \bar{p}_{t+1}]$ it must adjust; but if $p_t \in [\underline{p}_{t+1}, \bar{p}_{t+1}]$ it may not.

As regards the question with which we started – Shouldn't every seller have a target real price, and therefore when M_t increases shouldn't every seller adjust his nominal price by the same amount? – the answer is No! Sellers do not have a unique target price. Equilibrium requires a distribution of prices all of which yield the same profit. If you do not change your p_t when M_t increases, you indeed earn less profit per unit, but again you make it up on the volume. Hence, sellers can change prices infrequently in the face of continuous movements in the aggregate price level, even though they are allowed to change whenever they like at no cost. One might say that sellers can be rationally inattentive to the aggregate price level and monetary policy, within some range, since as long as $p_t \in [\underline{p}_t, \bar{p}_t]$, their place in this distribution does not matter. But policy cannot exploit this. The distribution of relative prices is pinned down uniquely, and if, say, M_t were to unexpectedly double, $F_t(p)$ must adjust to keep the real distribution the same, even if many individual prices do not adjust. Hence, the level of the money supply M_t is irrelevant, even if inflation, nominal interest or money growth rates matter. This is classical neutrality.⁵

⁴This alternating-market structure is taken from Lagos and Wright (2005), mainly because it is extremely tractable, but as we said any other monetary model could be used. Previous analyses in this framework have used several different pricing mechanisms, including various bargaining solutions, price posting with directed search, Walrasian price taking, and auctions (see the above-mentioned surveys). No one has previously tried Burdett-Judd pricing in the model, although it was used in the related model of Shi (197) by Head and Kumar (2005) and Head, Kumar and Lapham (2010).

⁵Although we focus in this application on changes in M_t , the same argument applies to what Golosov and Lucas (2003) call “continuously changing demand and supply conditions.” Any change in utility or cost functions can change the Burdett-Judd price distribution, but this simply does not imply that all sellers must adjust their individual prices.

We then show that a calibrated version of the model can match quite well the empirical behavior of prices in the US retail sector. First, the calibrated model predicts an average price duration that is reasonably close to what one sees in the data. Second, our theory generates a price change distribution that has the same shape and the features of the empirical price change distribution – e.g., the average price change is large, yet there are many small changes, and even many negative price changes. Third, in the model the probability and magnitude of price adjustments are approximately independent of the time since the last adjustment, as in the data. Fourth, the theory correctly predicts that inflation increases both the frequency and the magnitude of price changes. Overall, our model of price stickiness appears empirically reasonable. But, again, money is neutral. This demonstrates formally that nominal stickiness neither requires technological restrictions on price adjustment nor justifies particular interventions by central banks.⁶

2 The Model

Time is discrete and continues forever. In every period, two markets open sequentially. We call the first the Burdett-Judd market, or BJ for short, a decentralized market for a consumption good q_t in which buyers and sellers meet through a frictional matching process. Here barter is not feasible since buyers have nothing to offer by way of *quid pro quo*, and credit is not feasible because they are anonymous. Hence, exchange takes place using fiat money, supplied by the government according to the rule $M_{t+1} = \mu_t M_t$, where $\mu_t > 1/\beta$ is the money growth rate at t . After the BJ market closes, there convenes a centralized market where agents trade a different good x_t , as well as labor h_t and money m_t , called the Arrow-Debreu market, or AD for short. In AD households receive a lump sum transfer (or tax) T_t to accommodate increases (or decreases) in M_t . Also, in this market, as in standard

⁶There 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. Second, by design, money is exactly neutral. Third, stickiness arises entirely endogenously and robustly – it does not depend on particular functional forms, timing, the money supply process, etc. Fourth, the distribution of prices is endogenous, derived from standard microeconomics (Burdett-Judd), instead of simply assuming, say, prices are distributed uniformly (as in Caplin-Spulber). Also, we take our model to the data, as do some (e.g., Golosov and Lucas 2007) but not all (e.g., Caplin and Spulber 1987) of the above-mentioned studies

general equilibrium theory, we cannot say who trades with whom or how – the approach does not allow one to ask if they use barter, money or credit, only requiring that household satisfy their budget equations and markets clear.⁷

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

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

where $\beta \in (0, 1)$ is the discount factor, while u and v are strictly increasing and concave functions over the BJ good and AD good, respectively. There is a continuum of firms with measure s . Firms operate technologies for producing goods described as follows: producing a unit of x requires $h = x$ hours of labor, and producing a unit of q requires $h = cq$ hours of labor, so that c is the constant marginal cost of producing the BJ goods in terms of AD goods. As in standard general equilibrium theory, households own the firms, and receive profits as dividends D_t , in dollars, in the AD market.⁸

In the BJ market at t , each firm posts a nominal price p_t , taking as given the distribution of prices posted by all the other firms, described by the CDF $F_t(p)$, as well as the distribution of money across buyers in the market, in general, although in this model that is degenerate – i.e., along the equilibrium path, $m_t = M_t$ for each household in the BJ market. Households know the distribution $F_t(p)$, but only contact and hence can only purchase from a random sample of BJ sellers. A household generally contacts k firms with probability α_k . For simplicity we assume $\alpha_0 \in [0, 1)$, $\alpha_1 \in (0, 1 - \alpha_0)$ and $\alpha_2 = 1 - \alpha_0 - \alpha_1$, so that a household contacts at most two firms. One can easily generalize this in a variety of without changing the substantive results (e.g., as in Mortensen 2005); one can also allow households to choose

⁷The role of money in the BJ market is basically the same as Kiyotaki and Wright (1989); see Kocherlakota (1998) or Wallace (2011) for rigorous discussions. Also, one should not worry about the assumption that changes in M_t are accomplished via lump sum transfers or taxes. It is equivalent for all we do here to have the government to use increases or decreases in M_t to buy more or fewer AD goods (all that would change for the households is h_t).

⁸The baseline assumption is that q is produced in the AD market, and carried into the next BJ market by firms, who know exactly how much they will sell as a function of their price by the law of large numbers (and we do not dwell here on technicalities regarding the conditions needed for this law to apply). This allows us to interpret firms as simply technologies, owned by households, as in standard general equilibrium theory, but is usually equivalent if not more convenient to alternatively interpret firms as individuals who produce and consume (see Section 3).

endogenously how many sellers they sample at some cost (e.g., as in the original Burdett-Judd 1983 paper) and show that in equilibrium we get $\alpha_1, \alpha_2 \in (0, 1)$ and $\alpha_n = 0 \forall n > 2$. Also, although for ease of notation we assume all trade in the BJ market is monetary, it is easy to allow some credit trades, since for money to be essential we only need to have some BJ trade where credit is unavailable (see Head et al. 2011).

Before proceeding, we mention that we are aware that there are options for the types of mechanisms firms can post. In principle, they could post menus, where buyers can have any q , perhaps in some set Q , for a payment $P(q)$, but here we impose linearity, $P(q) = pq$.⁹ We experimented with alternatives, but thought the linear case was sufficiently interesting to focus on it for now. We also studied a version of the model where BJ goods are indivisible (see Liu 2010 or Head et al. 2011), which avoids the issue, since the only option is to post a p giving the price for an indivisible unit. That version is easier on some dimensions, although it also has some problems – e.g., as is well known, monetary models with indivisible goods and price posting admit a multiplicity of equilibria (see Jean, Rabinovich and Wright 2010 and references therein). At some level, this multiplicity does not matter, since all the equilibria are qualitatively similar, but it is slightly inconvenient, and so we use divisible goods here.

2.1 Households' Problem

Let $W_t(m_t)$ and $V_t(m_t)$ be the value functions for a household with m_t dollars in the AD and BJ market, respectively. Let ϕ_t be the value of money (the inverse of the nominal price level) in AD, where the price of x_t and the real wage are both 1 given our technology. Then the AD problem for a household is

$$W_t(m_t) = \max_{h_t, x_t, \hat{m}_t} \{v(x_t) - h_t + \beta V_{t+1}(\hat{m}_t)\} \text{ st } x_t = h_t + \phi_t (m_t - \hat{m}_t + D_t + T_t),$$

with nonnegativity constraints implicit. Eliminating h_t using the budget equation, we can reduce this to

$$W_t(m_t) = \phi_t (m_t + D_t + T_t) + \max_{x_t, \hat{m}_t} \{v(x_t) - x_t + \phi_t - \hat{m}_t + \beta V_{t+1}(\hat{m}_t)\}. \quad (2.2)$$

⁹Ennis (2001) and Dong and Jiang (2011), e.g., study related monetary models where nonlinear pricing is used by sellers to elicit private information about buyers' preferences.

The solution satisfies the FOC

$$v'(x_t) = 1 \text{ and } \beta V'_{t+1}(\hat{m}_t) = \phi_t, \quad (2.3)$$

plus the budget constraint, $h_t = x_t + \phi_t (\hat{m}_t - m_t - D_t - T_t)$. This implies: (1) \hat{m}_t and x_t are independent of m_t , so that in particular the equilibrium distribution money is degenerate across households entering the BJ market; and (2) W_t is linear with slope ϕ_t .

For a household in the BJ market with m_t dollars, conditional on sampling at least one price and the lowest price sampled being p , we define

$$U_t(p, m_t) = \max_{q_t} \{u(q_t) + W_t(m_t - pq_t)\} \text{ st } pq_t \leq m_t.$$

Thus, $q_t = q_t(p, m_t)$ solves an elementary demand problem with liquidity constraint $pq_t \leq m_t$. It is easy to show the difference between the slopes in (q, p) space of the unconstrained demand curve and the constraint at equality has the same sign as $1 - \gamma(q)$ when the curves cross, where $\gamma(q) = qu''(q)/u'(q)$ is the coefficient of relative risk aversion. It is convenient to have a single crossing, so that the constraint binds either for high p or for low p , and a sufficient condition for this is either $\gamma(q) > 1 \forall q$ or $\gamma(q) < 1 \forall q$. We assume constant relative risk aversion, $u(q) = q^{1-\gamma}/(1-\gamma)$, and assume $\gamma \in (0, 1)$, so that demand is constrained at low p (see Liu 2010 for the case $\gamma > 1$, and for results with a general function u).

With this specification, the conditional BJ problem is

$$U_t(p, m_t) = \max_{q_t} \left\{ \frac{q_t^{1-\gamma}}{1-\gamma} + W_t(m_t - pq_t) \right\} \text{ st } pq_t \leq m_t. \quad (2.4)$$

This is easily solved to get

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

where $\hat{p}_t = \phi_t^{\frac{1}{\gamma-1}} m_t^{\frac{\gamma}{\gamma-1}}$. If $p < \hat{p}_t$ households cash out; otherwise $pq_t(p, m_t) < m_t$, so they have money to spare and demand is unconstrained. This is shown in Figure 1, where constrained demand is given by the lower envelope of the two curves representing unconstrained demand and the constraint at equality.

The unconditional value function entering the BJ market, before potentially contacting sellers and observing prices, is

$$V_t(m_t) = \alpha_0 W_t(m_t) + \alpha_1 \int U_t(p, m_t) dF_t(p) + \alpha_2 \int U_t(p, m_t) d\{1 - [1 - F_t(p)]^2\}.$$

This expression is easy to understand: with probability α_0 the household contacts no seller and enters the next AD market with m_t unchanged; with probability α_1 the household contacts one seller posting a random draw from $F_t(p)$; and with probability α_2 the household contacts two firms and the lower of the two prices is a random draw from $1 - [1 - F_t(p)]^2$. Simple algebra reduces this to

$$V_t(m_t) = \alpha_0 W_t(m_t) + \int [\alpha_1 + 2\alpha_2 - 2\alpha_2 F_t(p)] U_t(p, m_t) dF_t(p). \quad (2.6)$$

Differentiating $V_t(m_t)$, the FOC $\phi_t = \beta V'_{t+1}(\hat{m}_t)$ becomes

$$\phi_t = \beta \phi_{t+1} \left\{ 1 + \int_0^{\hat{p}_{t+1}} [\alpha_1 + 2\alpha_2 - 2\alpha_2 F_{t+1}(p)] \left[\frac{1}{p\phi_{t+1}} \left(\frac{\hat{m}_t}{p} \right)^{-\gamma} - 1 \right] dF_{t+1}(p) \right\}.$$

Although we focus on stationary equilibria below, for now we do not impose this restriction. In general, the inflation rate is $\pi_t = \phi_t/\phi_{t+1}$, and the Fisher equation gives the nominal interest rate by $1 + i_t = (1 + r_t)(1 + \pi_t)$, where $1 + r_t = 1/\beta$ is the real interest rate, which is time invariant here due to quasi-linear utility. To be clear, as is standard, we can obviously price any asset in equilibrium, including real or nominal claims between the AD market at t and the AD market at $t+1$, even if these do not circulate in the BJ market (say, because they are not tangible assets, simply claims on numeraire goods or money in AD). This defines the above interest rates, and allows us to rewrite previous condition as

$$i_t = \int_0^{\hat{p}_{t+1}} [\alpha_1 + 2\alpha_2 - 2\alpha_2 F_{t+1}(p)] \left[\frac{1}{p\phi_{t+1}} \left(\frac{\hat{m}_t}{p} \right)^{-\gamma} - 1 \right] dF_{t+1}(p). \quad (2.7)$$

Heuristically, the LHS of (2.7) is the marginal cost of carrying cash between t and $t+1$, the nominal interest rate; and the RHS is the marginal benefit, the expected value of relaxing the liquidity constraint in the next BJ market which binds when $p < \hat{p}_{t+1}$.

2.2 Firms' Problem

If a firm posts p in the BJ market, profit is

$$\Pi_t(p) = \frac{1}{s} [\alpha_1 + 2\alpha_2 - 2\alpha_2 F_t(p) + \alpha_2 \xi_t(p)] R_t(p), \quad (2.8)$$

where $\xi_t(p) = \lim_{\varepsilon \rightarrow 0^+} F_t(p) - F_t(p - \varepsilon)$, and $R_t(p)$ is profit per buyer served, given that in equilibrium all buyers have M_t ,

$$R_t(p) = q_t(p, M_t)(p\phi_t - c).$$

The term in brackets in (2.8) is the number of customers served: α_1/s households purchase from the firm because this is their only contact; $2\alpha_2 [1 - F_t(p)]/s$ households purchase from the firm because they contact another seller with have a price above p ; and there are $2\alpha_2 \xi_t(p)/s$ households that contact the firm plus another with the same p , and in this case we can assume they randomize, although as we shall see this term vanishes in equilibrium because the probability two sellers set the same price is 0.

Figures 2 and 3 show two curves. One is $(M_t/p)(p\phi_t - c)$, which is profit in units of numeraire from selling to a buyer that is constrained. The other is $(p\phi_t)^{-1/\gamma}(p\phi_t - c)$, which is profit from selling to a buyer that is not liquidity constrained. Actual profit per customer is the lower envelope of these curves. Figure 2 illustrates the case in which the constraint $p q_t \leq m_t$ is not very tight, and the price that maximizes profit per customer is $p = c/\phi_t(1 - \gamma)$. Figure 3 illustrates the case in which the constraint is tighter, and the price that maximizes profit per customer is $p = \hat{p}_t$. The profit-maximizing price in general is $p_t^m = \max\{c/\phi_t(1 - \gamma), \hat{p}_t\}$, which we call the monopoly price. Each firm chooses p to maximize $\Pi_t(p)$. Therefore, a price distribution $F_t(p)$ is consistent with profit maximization by all firms when $\Pi_t(p)$ is maximized by every p on the support of F_t , denoted \mathcal{F}_t . In other words, profit maximization means

$$\Pi_t(p) = \Pi_t^* \equiv \max_p \Pi_t(p) \quad \forall p \in \mathcal{F}_t, \quad (2.9)$$

The following result characterizes F_t by adapting the arguments in Burdett and Judd (1983), which apply even though the environment different because the BJ good is divisible

and because buyers can be liquidity constrained. The proof is in Appendix A.

Proposition 1: *The unique price distribution consistent with profit maximization by all firms at t is*

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

with support $\mathcal{F}_t = [\underline{p}_t, \bar{p}_t]$, where

$$R_t(\underline{p}_t) = \frac{\alpha_1}{\alpha_1 + 2\alpha_2} R_t(p_t^m) \text{ and } \bar{p}_t = p_t^m. \quad (2.11)$$

The price distribution is continuous, intuitively, because if it had a mass point at some p_0 , say, a firm posting p_0 could increase profit by changing to $p_0 - \varepsilon$, as this leaves profit per customer approximately constant and increases sales by a discrete amount. The support \mathcal{F}_t is connected, intuitively, because if it had a gap between p_0 and p_1 , say, a firm posting p_0 could increase profits by changing to p_1 , as this does not reduce the number of sales and increases profit per sale. Since $F_t(p)$ has no mass points, $\xi_t(p) = 0$, and (2.8) reduces to

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

The closed form in (2.10) is derived as follows: $\Pi_t^* = (\alpha_1/s) R_t(p_t^m)$ since $p_t^m \in \mathcal{F}_t$; equating this to (2.12), we solve for $F_t(p)$.

2.3 Equilibrium

We are now in the position to define an equilibrium.

Definition 1: *Given a process $\{M_t\}$, an equilibrium Σ^* is a (bounded and nonnegative) sequence of AD quantities $\{h_t^*, x_t^*, \hat{m}_t^*\}$, BJ decision rules $\{q_t^*(p, \hat{m})\}$ and prices $\{\phi_t^*, F_t^*(p)\}$ satisfying the following conditions for all t :*

1. $(h_t^*, x_t^*, \hat{m}_t^*)$ solves the household's AD problem, and in particular \hat{m}_t^* satisfies (2.7);
2. $q_t^*(p, \hat{m})$ solves the household's BJ problem as described in (2.5);
3. $F_t^*(p)$ solves the firm's BJ problem as described in Proposition 1 with support $\mathcal{F}_t^* = (\underline{p}_t^*, \bar{p}_t^*)$;
4. ϕ_t implies market clearing, $\hat{m}_t^* = M_t$.

As mentioned above, we are mostly interested here in stationary outcomes, which makes sense when policy is stationary, $M_{t+1} = \mu M_t \forall t$ for some constant μ . Assuming this is the case, we have the following:

Definition 2: *A stationary monetary equilibrium is an equilibrium where all nominal variables grow at rate μ , all real variables at rate 0, and $\phi_t^* > 0$.*

Stationarity implies $F_{t+1}^*(\mu p) = F_t^*(p)$ and $q_{t+1}^*(\mu p) = q_t^*(p)$, which means that the real distribution of BJ prices and the BJ decision rule are time invariant. It also implies a constant inflation rate $\pi = \mu$ and nominal interest rate $1 + i = \mu/\beta$.

To define some terminology, classical neutrality means the following: suppose we have an equilibrium Σ , and we change M_t to $M'_t = \Theta M_t \forall t$ for some $\Theta > 0$. Then there exists an equilibrium Σ' where all nominal variables increase by a factor Θ – e.g., $p'_t = \Theta p_t$, $\phi'_t = \phi_t/\Theta$ etc. – while all real variables are the same – e.g., $q'_t = q_t$ etc. Clearly, in this model, equilibria (stationary or otherwise) display neutrality in this sense. This merely says that units do not matter. Later we consider another notion of neutrality, given an unexpected change in M_t . In any case, we emphasize that neutrality does not imply superneutrality: changing the growth rate μ in the rule $M_{t+1} = \mu M_t$ does have real effects. Also note that in a stationary monetary equilibrium it is equivalent to choose the money growth rate μ , the inflation rate π or the nominal interest rate i as a policy instrument.

We establish the existence of a stationary monetary equilibrium formally in Appendix B, but here we give the basic idea behind the argument. First, we show that prices posted in the BJ market are decreasing, in the sense of first-order stochastic dominance, with respect to the amount of money firms expect households to carry. Intuitively, if households have more money the liquidity constraint is relaxed, which increases profit at low-price firms relative to high-price firms, because the former are where the constraint binds; so, to keep firms indifferent between low and high prices, the distribution must shift to reduce the number of customers served by low- relative to high-price firms. Then we prove the amount of money carried by households is decreasing with respect to prices in the BJ market. Intuitively, if prices are higher, in the sense of first order stochastic dominance, a household has a lower probability of meeting a low-price seller and hence a lower probability of being liquidity constrained, so the value of money in the BJ market falls. It follows that the amount

of money households carry is a monotone function of the amount firms expect them to carry. Moreover, the amount of money households carry is bounded. Hence, from a fixed point theorem of Tarski (1955), there exists an m_t^* such that: (1) m_t^* solves the households' problem given F_t^* ; and (2) F_t^* is the BJ price distribution given m_t^* . Given m_t^* and F_t^* we easily get all the other endogenous variables.

Proposition 2: *A stationary monetary equilibrium exists.*

3 Sticky Prices

Equilibrium uniquely pins down the aggregate BJ price distributions for all t – both real and nominal – but not the price of any individual firm, since by definition equilibrium implies the same profit from any $p \in \mathcal{F}_t$. Figure 4 illustrates the implications for the dynamics of the distribution and individual prices when $\mu > 1$, by showing F_t^* and F_{t+1}^* . All firms with prices in the vertically shaded area must change between t and $t + 1$, since for them p_t does not maximize $\Pi_{t+1}(p)$, even though it did maximize $\Pi_t(p)$. The firms in the horizontally shaded area, however, are indifferent between keeping price constant and posting a new price in $[\underline{p}_{t+1}, \bar{p}_{t+1}]$. The only equilibrium restriction on the individual price dynamics between t and $t + 1$ is that the aggregate distribution at $t + 1$ has to be F_{t+1}^* .

Definition 3: *In a stationary monetary equilibrium, a repricing policy $p_{t+1}^*(p)$ is admissible if, when the distribution at t is $F_t^*(p)$ and all firms follow policy $p_{t+1}^*(p)$, the distribution at $t + 1$ is $F_{t+1}^*(p)$.*

In the remainder of the paper, we restrict attention to stationary outcomes, and the case $\mu > 1$. We also focus on repricing policies of the following form:

$$\begin{aligned} & \text{if } p_t \notin \mathcal{F}_{t+1} \text{ then } p_{t+1}^*(p_t) = p' \\ & \text{if } p_t \in \mathcal{F}_{t+1} \text{ then } p_{t+1}^*(p_t) = \begin{cases} p_t & \text{with prob } \rho \\ p' & \text{with prob } 1 - \rho \end{cases} \end{aligned} \quad (3.1)$$

where $p' \in \mathcal{F}_{t+1}$ is a profit-maximizing price at $t + 1$, determined as discussed below. The parameter ρ is a probability used as a tie-breaking rule: if you are indifferent between changing and not changing your price, you randomize. Although this may bear a superficial resemblance to Calvo pricing, we cannot emphasize strongly enough that it could not be more different. With Calvo pricing, firms may be *desperate* to change p , but are only allowed to

do so with some probability each period. Here, any firm that wants to change p can and will; only those who are genuinely indifferent may randomize.

The only additional structure we impose on repricing is symmetry. This means that, first, that all sellers use the same ρ , and second, that those who reprice between t and $t + 1$ all draw a new p' from the same distribution, say $G_{t+1}(p')$. We now show that once ρ is specified $G_{t+1}(p')$ is pinned down uniquely. To begin, note that in stationary equilibrium $F_{t+1}(\mu p) = F_t(p)$, which says that when inflation is μ the probability of finding a price below p today is the same as the probability of finding price below μp tomorrow. What kind of repricing distribution makes this happen?

Given $F_t(\cdot)$, and any $G_{t+1}(\cdot)$, we compute $F_{t+1}(\cdot)$ as follows: First, for $p \in (\underline{\mu p}_t, \bar{p}_t)$,

$$F_{t+1}(p) = F_t(\underline{\mu p}_t) G_{t+1}(p) + \left[1 - F_t(\underline{\mu p}_t)\right] (1 - \rho) G_{t+1}(p) + \left[F_t(p) - F_t(\underline{\mu p}_t)\right] \rho.$$

The first equality says that the measure of sellers below p evolves as follows: a measure $F_t(\underline{\mu p}_t)$ fall off the support between t and $t + 1$ and they all reprice using $G_{t+1}(p)$; a measure $1 - F_t(\underline{\mu p}_t)$ do not fall off the support and do not have to reprice, but do so anyway with probability $1 - \rho$; and a measure $F_t(p) - F_t(\underline{\mu p}_t)$ with price below p do not have to reprice and choose not to with probability ρ . Algebra yields

$$F_{t+1}(p) = \left[1 - \rho + \rho F_t(\underline{\mu p}_t)\right] G_{t+1}(p) + \left[F_t(p) - F_t(\underline{\mu p}_t)\right] \rho.$$

Similarly, for $p > \bar{p}_t$ we similarly have

$$F_{t+1}(p) = \left[1 - \rho + \rho F_t(\underline{\mu p}_t)\right] G_{t+1}(p) + \left[1 - F_t(\underline{\mu p}_t)\right] \rho.$$

We now impose stationarity, $F_{t+1}(\mu p) = F_t(p)$, and solve for the repricing distribution:

$$G_{t+1}^*(p) = \begin{cases} \frac{F_t^*(p/\mu) - \left[F_t(p) - F_t(\underline{\mu p}_t)\right] \rho}{1 - \rho + \rho F_t^*(\underline{\mu p}_t)} & \text{if } p \in (\underline{p}_t, \bar{p}_t/\mu) \\ \frac{F_t^*(p/\mu) - \left[1 - F_t(\underline{\mu p}_t)\right] \rho}{1 - \rho + \rho F_t^*(\underline{\mu p}_t)} & \text{if } p \in (\bar{p}_t/\mu, \bar{p}_t). \end{cases} \quad (3.2)$$

Given inflation μ , which is a policy variable, and any tie-breaking rule ρ , the unique repricing

distribution that keeps the real price distribution constant is (3.2). The equilibrium law of motion for the nominal price distribution is:

$$F_{t+1}(p) = \begin{cases} \left\{ F_t^*(\mu \underline{p}_t) + (1 - \rho) \left[1 - F_t^*(\mu \underline{p}_t) \right] \right\} G_{t+1}^*(p) & \text{if } p < \mu \underline{p}_t \\ \left[F_t^*(p) - F_t^*(\mu \underline{p}_t) \right] \rho + \left\{ F_t^*(\mu \underline{p}_t) + (1 - \rho) \left[1 - F_t^*(\mu \underline{p}_t) \right] \right\} G_{t+1}^*(p) & \text{if } p \geq \mu \underline{p}_t \end{cases} \quad (3.3)$$

We have established the following result.

Proposition 3: *The pricing policy (3.1), with all new prices drawn from $G_{t+1}^*(p')$ as given in (3.2) is consistent with stationary monetary equilibrium $\forall \rho \in [0, 1]$.*

The class of repricing policies (3.1) is not exhaustive, but it captures a wide range of behavior in a parsimonious way. For $\rho = 1$, (3.1) describes an extreme case in which firms only change p when it is no longer profit maximizing, giving the smallest fraction of price changes and highest average price duration consistent with equilibrium. For $\rho = 0$, we have the opposite extreme in which firms change p in every period, giving the largest fraction of changes and the lowest average duration consistent with equilibrium. As ρ increases from 0 to 1, the frequency of changes and the average price duration move from one extreme to the other. For any ρ , we now compute this frequency and average price duration.

The distribution of new prices in period t is $G_t^*(p)$. Let N denote the largest integer such that $\mu^N \underline{p}_t \leq \bar{p}_t$. For $n = 1, 2, \dots, N$, a fraction $G_t^*(\mu^n \underline{p}_t) - G_t^*(\mu^{n-1} \underline{p}_t)$ of new prices are in $[\mu^{n-1} \underline{p}_t, \mu^n \underline{p}_t]$, and a fraction $1 - G_t^*(\mu^N \underline{p}_t)$ are in $[\mu^N \underline{p}_t, \bar{p}_t]$. Each $p \in [\mu^{n-1} \underline{p}_t, \mu^n \underline{p}_t]$ changes at $t + i$ and not before with probability $\rho^{i-1}(1 - \rho)$, $i = 1, 2, \dots, n - 1$, and will change in period $t + n$ with probability ρ^{n-1} . Therefore, the average duration of prices in $[\mu^{n-1} \underline{p}_t, \mu^n \underline{p}_t]$ is $(1 - \rho) + 2\rho(1 - \rho) + \dots + n\rho^{n-1} = (1 - \rho^n)/(1 - \rho)$. Each $p \in [\mu^N \underline{p}_t, \bar{p}_t]$ will change at $t + i$ with probability $\rho^{i-1}(1 - \rho)$, $i = 1, 2, \dots, N$, and at $t + N + 1$ with probability ρ^N . Therefore, the average duration in the interval $[\mu^N \underline{p}_t, \bar{p}_t]$ is $(1 - \rho^{N+1})/(1 - \rho)$. The overall average duration of a new price is

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

Since $(1 - \rho^n)/(1 - \rho)$ is increasing in ρ and n , and G_t^* is increasing in ρ in the first-order stochastic dominance sense, $A(\rho)$ is increasing in ρ .

We now compute the fraction of prices that change between t and $t + 1$, starting from $F_t^*(p)$. A fraction $F_t^*(\mu \underline{p}_t)$ of prices are in $[\underline{p}_t, \mu \underline{p}_t]$, and each of these change between t and $t + 1$ with probability 1 . A fraction $1 - F_t^*(\mu \underline{p}_t)$ are in $[\mu \underline{p}_t, \bar{p}_t]$, and each of these change between t and $t + 1$ with probability $1 - \rho$. The overall fraction of prices that change between t and $t + 1$ is therefore

$$\Phi(\rho) = F_t^*(\mu \underline{p}_t) + (1 - \rho) \left[1 - F_t^*(\mu \underline{p}_t) \right], \quad (3.5)$$

with $\Phi(\rho)$ decreasing in ρ .

We next compute the distribution of the magnitude of price changes. The density of firms that post p at t and a different price at $t + 1$ is $F_t^*(p)/\Phi(\rho)$ if $p < \mu \underline{p}_t$ and $(1 - \rho)F_t^*(p)/\Phi(\rho)$ if $p > \mu \underline{p}_t$. Among the firms that post p a new p at $t + 1$, a fraction $G_{t+1}^*[p(1 + \delta)]$ increase p by δ percent or less. Therefore, the distribution for the magnitude of price changes is

$$H_t(\delta, \rho) = \frac{1}{\Phi(\rho)} \int G_{t+1}^*[\rho(1 + \delta)] \left(1 - \rho \mathbf{1} \left\{ p \geq \mu \underline{p}_t \right\} \right) dF_t^*(p). \quad (3.6)$$

From (3.2) and (3.6), it is immediate that $H_t(0, \rho) > 0$ for all $\rho < 1$.

Proposition 4: *A stationary monetary equilibrium Σ^* together with a repricing policy p_{t+1}^* yields an average price duration $A(\rho)$ and a frequency of price changes $\Phi(\rho)$, with $A(\rho)$ increasing and $\Phi(\rho)$ is decreasing in ρ . There is a $\mu^* > 1$ such that $\mu \in (1, \mu^*)$ and $\rho \in (0, 1]$, $A(\rho) > 1$ and $\Phi(\rho) < 1$. For all $\mu \in (1, \mu^*)$ and $\rho \in [0, 1)$, the fraction of negative price changes, is $H(0, \rho) > 0$*

Proof: See Appendix C.

The result tells us that , unless the growth rate of money is too high, the model is consistent with the observation that some firms stick to their prices for some time despite a constantly changing aggregate price level.¹⁰ Our model delivers this result not because there are technological restrictions on price adjustment, but because standard search frictions imply an interval of prices all of which maximize profit. It is also consistent with the observation that some firms lower their price despite a constantly increasing aggregate price level. It also delivers this result because of search frictions, and not because of idiosyncratic

¹⁰Obviously, if inflation is too high, all firm must reprice every period. If, e.g., we start at t with prices in $\mathcal{F}_t = [1, 2]$, and double the money supply between t and $t + 1$, the support moves to $\mathcal{F}_{t+1} = [2, 4]$, and the set of agents with $p_t \in \mathcal{F}_t \cup \mathcal{F}_{t+1}$ has measure 0.

shocks. More broadly, the results show that one should be cautious about making inferences concerning the existence or degree of menu costs and related restrictions on the timing of price changes from the observed stickiness of individual prices. Similarly, one should be cautious about making inferences concerning idiosyncratic productivity shocks from observed price changes.

Perhaps most importantly, one should be very cautious about making policy recommendations based on these observations. Some firms may well stick to the same nominal p for many periods, but this cannot be exploited by policy in our economy. Government cannot, e.g., increase short-run production or consumption through an unexpected increase in M . If we were to unexpectedly double the stock of money at the opening of the AD market, the m that each household carries into BJ would double, and so would the distribution of nominal prices in that market. Theory – i.e., utility maximization, profit maximization and equilibrium taken together – pins down uniquely the distribution of real prices here, and doubling M does not affect this. Similarly, the amount of money agents bring back to the AD market doubles, but the value of this money ϕ is cut in half. This is classical neutrality.

Expanding M is neutral, intuitively, because while the price posted by some sellers can be rigid in the short run, the aggregate distribution F_t is perfectly flexible. This contrasts sharply with what would happen if there were positive menu costs or if sellers were only allowed to change with probability less than 1. In these cases, if we unexpectedly double M , it is not possible in general to keep the distribution of real prices constant – e.g., suppose the support goes from $\mathcal{F}_t = [1, 2]$ to $[2, 4]$ after M doubles. This requires firms to change their prices with probability 1, and in our model they do. But if a fraction of sellers are not allowed to change p after a shock to M , as in Calvo-style models, or if some sellers have a high enough cost to changing p , as in Mankiw-style models, they are stuck with prices that are too low and do not maximize profit. This obviously does affect the real outcome and welfare. Without working through the details, it is clear that many households are going to find BJ goods going at bargain-basement prices and, in general will demand more, which might force the firms to supply more, depending on how one specifies the details.¹¹

¹¹A detail we mention here is that, in the above description of the environment, we said sellers buy inventories in AD and bring them to BJ, with expectations about how much they will sell that are correct with probability 1. This cannot happen if M doubles and the nominal distribution does not – but whether

Although the exact outcome may depend on details, the general conclusions are very robust. In Head et al. (2011), e.g., we present an indivisible goods version of the model, where there is no scope for changes in money to affect production or consumption on the intensive margin, but introduce a participation condition: households must pay a fixed cost to enter the BJ market, analogous to the free-entry condition for firms to enter the labor market in Pissarides (2000). With Calvo- or Mankiw-style pricing, a increase in M that catches sellers by surprise means many real prices too low from a profit maximizing perspective, and generally we expect this to increase entry of households into the BJ market. That is, when sellers cannot change their prices, even though they would like to, monetary policy can instigate a shopping spree by households in search (literally) of bargains, and this sets off a production boom if sellers are obliged to meet demand, as in most sticky-price models. Symmetrically, a fall in M can lead to a slump in Calvo- or Mankiw-style models. Neither a boom nor a slump occurs in our setup under these policy scenarios, where prices are reset quickly, even though in normal times many prices may be reset only gradually. Our economy has the property that nominal magnitudes do not matter, real magnitudes do.¹²

4 Quantitative Evaluation

We have a theory of nominal rigidities that relies on search frictions in product markets, not on the existence of technological frictions to repricing. In this section, we ask if the theory can account for the empirical evidence. While our model delivers equilibrium price distributions, we choose to look at equilibrium price-change distributions instead, since many macroeconomists have been focusing on the latter of late. Still it is worth mentioning that future work could analyze price distributions. The labor-market version of Burdett-Judd, the

this results in firms stocking out, or somehow producing additional output, is something we do not go into here. The point is simply that something other than the expected equilibrium has to happen. If we assume sellers can produce q (as opposed to selling out of inventory), and that they are obliged to do so for everyone that pays the posted price, as in most Keynesian models, then a surprise increase (decrease) in M can raise (lower) consumption and output with Calvo- or Mankiw-style models. By contrast, in our model, the real allocation is not affected by the policy under consideration.

¹²Two points bear repeating. First, the goal is not to prove that money is neutral in the real world; it is to prove that observed price stickiness does not logically imply nonneutralities, nor does it logically imply technological restrictions on repricing are needed for theories to match the data. Second, while the level of M does not matter, the growth rate μ does. In the interest of space we do not go into details, but see Wang (2011) for an extended analysis of the effects of inflation on allocations and welfare in a version of our model.

Burdett-Mortensen (1998) model, e.g., has been applied to study wage (not wage-change) distributions. While the simplest Burdett-Morrtsensen models do not fix the facts well, much has been learned from the effort, and the models have been adapted and extended to do much better. Something similar could help us learn about product markets. But for this project we instead look at the evidence on price changes, as described in a representative study by Klenow and Kryvtsov (2008) (see Klenow and Malin 2010 for a survey of related empirical work).

For our purposes, in terms of preferences and technology, we need to specify the discount factor β , the utility function for the BJ good $u(q) = q^{1-\gamma}/(1-\gamma)$, and its the marginal cost, which we normalize to $c = 1$. We do not need to specify utility for the AD good $v(x)$, although it may be needed to see how well the model fits observations other than those on which we focus. It can, e.g., affect the model-generated money demand curve – the relationship between i and real balances – which one can compare to the data. This is studied in an extension of the framework by Wang (2011), where the model does reasonably well on this dimensions, so we concentrate on other issues. In particular, we concentrate on repricing behavior, as described by a function $p_{t+1}^*(p, \rho)$ with parameter ρ . We also need to parameterize search frictions, as described by α_k , where α_k is the probability that a household contacts k firms, $k = 0, 1, 2$. We restrict attention to the case where each household attempts to solicit two price quotes from BJ market, each of which succeeds independently with probability λ . Thus, $\alpha_0 = (1-\lambda)^2$, $\alpha_1 = 2(1-\lambda)\lambda$ and $\alpha_2 = \lambda^2$. Finally, the monetary policy is described by the growth rate of money μ , but as we said above this is equivalent to targeting inflation or nominal interest rates.

We calibrate the model to the US economy over the period 1988-2004. We interpret the BJ market as the retail sector and the AD market as an intermediate goods sector. We choose the model period to be a month, and set β so that the annual real interest rate matches the average in the data, 1.035. We set μ so that the annual inflation rate in the model matches that in the data, 1.03. We then choose γ and ρ to minimize the distance between the model-generated distribution of price changes in the BJ market $H_t(\delta, \rho)$ and its empirical counterpart for the retail sector, as described by Klenow-Kryvtsov. Finally, we choose λ so that the average markup in the BJ market is 30 percent, which is an average

across retailers in the survey data discussed in Faig and Jerez (2005). After calibrating the parameters, the predictions of the model regarding price-changes in the BJ market are uniquely pinned down. There is a simple intuition behind our calibration strategy for γ and ρ . The parameter γ determines the elasticity of profit per customer $R_t(p)$. Hence, γ affects the distribution $F_t^*(p)$, and therefore the price-change distribution $H_t(\delta, \rho)$. Similarly, ρ determines the probability that a firm does not adjust its price when indifferent, and so affects the distribution of prices among firms that do not change, and hence the distribution among those that do $G_t^*(p)$, and hence the distribution of price changes $H_t(\delta, \rho)$.

4.1 Results

The bottom line is that our theory of price rigidity can account quite well for the empirical behavior of prices. According to the data analyzed by Klenow-Kryvtsov, 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: if both are both interpreted as price changes, the average duration of a price is 6.8; if product substitutions are interpreted as price changes but temporary sales are not, the average duration is 8.6; and if neither are interpreted as price changes, the average duration is 10.4. The average duration of a price predicted by the model, given an inflation rate of 3 percent and a calibrated value of $\rho = 0.93$, is 11.6 months. We did not calibrate to this number. Heuristically, the two parameters γ and ρ are set to try to match the price-change distribution, and the predicted duration happens to come out 11.6, which is on the high end of the range given by Klenow-Kryvtsov, but still very reasonable. Obviously, for higher values of ρ average duration increases, up to a maximum of 34 months, and for lower values of ρ average price duration falls, down to a minimum of 1. With $\rho = 0.91$ we generate an average duration at the lower end of the range, 6.8 months. Also, note that average duration is decreasing with respect to inflation, as inflation increases the fraction of prices that exit the support \mathcal{F}_t each period. See Figure 5.¹³

¹³At zero inflation, the model has a lot of indeterminacy depending on ρ (although, again, once ρ is specified there is a unique symmetric equilibrium repricing distribution G_t). Thus, with no inflation sellers can reprice each period or never. But as inflation increases the indeterminacy diminishes quickly, as can be seen in Figure 5 from the minimum price duration curve becoming fairly low even for moderate inflation rates.

One can consider the ability of the model to match the average duration of prices an independent check on the calibration, which was targeted to the histogram of price changes. One can also ask how well calibration matches this target. The blue histogram in Figure 6 is the empirical price-change distribution from Klenow-Kryvtsov, while the red histogram is the distribution predicted by the model. One can immediately see they are very close. Three features of the empirical distribution are worth emphasizing. First, the average price change is large, around 11 percent. Second, despite the large average, many price changes are small, with 44 percent smaller than 5 percent in absolute value. Third, many price changes are negative, around 35. This is problematic for a simple menu-cost story, since the large average change suggests large menu costs, but that is inconsistent with so many small and negative changes. Klenow and Krystov (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.¹⁴

For our model-generated price-change distribution, the average absolute value is 9 percent, the fraction of changes between -5 and $+5$ percent is 43 percent, and the fraction of negative price changes is 35 percent. We capture the empirical distribution quite well with no seller-specific productivity shocks. According to our theory, average price changes are large because search frictions create a lot of dispersion in the equilibrium distribution. The price posted by a firm at the 90th percentile of the distribution, e.g., is approximately double that posted at the 10th percentile. Hence, when p exits the support \mathcal{F}_t , on average firms make a large adjustment. Many price changes are small, however, because there are many firms that change p before it exits \mathcal{F}_t , and for the same reason many changes are negative. One can describe several other features of the data that the model matches well, including the fact that when two firms reprice at the same t they typically do not both adjust to the same p' , as predicted by at least simple menu-cost models. Of course one may be able to get a less-simple menu-cost model to do better; we only mention that we do not need any bells or whistles here, as the most basic version of the theory does fairly well.

¹⁴We are sympathetic to the idea that one may be able to account for some of these observations by sellers sometimes moving prices around to uncover information about demand, costs etc. (learning by experimenting). This has little to do with monetary neutrality, however, since presumably what they care about is real demand, costs etc.

Klenow-Kryvtsov 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 adjustment – i.e., the age of the price. Moreover, they estimate the relationship between the absolute value of price adjustments – i.e., the price-change size – and age. After controlling for unobserved heterogeneity across items, they find the price-change hazard remains approximately constant during the first 11 months and increases significantly during month 12, and the price-change size is approximately independent of age. We do not think these observations are especially puzzling since, e.g., perhaps at least some price changes are discussed before implementation at annual meetings. Nonetheless, in Figure 7 the red histogram shows the price-change hazard predicted by the model. As in the data it is approximately constant for the first 11 months in the life of a price, although it does not increase in month 12. This is because F_t in the model has a wide support. Therefore, during the 12 months after a change, few firms need to readjust, so the majority change only with probability ρ , independent of p 's age. We do not predict a spike after 12 months because our firms have no seasonal reason to adjust, like an annual meeting.

We now turn to the effects of inflation. Using time-variation over the period 1988-2005, Klenow-Kryvtsov measure the effect of inflation on the frequency of price adjustments (the extensive margin) and on the magnitude (the intensive margin). They accomplish this by estimating the coefficient on inflation in a regression of the frequency of price adjustments and in a regression of the magnitude. 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 percent and the magnitude of price adjustments by 3.55 percent. Figure 8 illustrates the predictions of our model. According to the model, an increase in inflation increases both the frequency and the magnitude of price adjustments. This is 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, to a compression of the support \mathcal{F}_t . Second, given the support, an increase in inflation reduces the time it takes for a price to exit \mathcal{F}_t . For both reasons, an increase in inflation increases the fraction of prices that adjust every month. For similar reasons, an increase in inflation leads to a greater average price adjustment.

It is obvious that at least the standard Calvo-style model cannot match these observations: the magnitude of price changes may be endogenous but the frequency is exogenous and as such cannot depend on inflation. Our model matches the stylized facts about the extensive and intensive margins qualitatively, but does not nail them quantitatively. An increase in inflation from 3 to 4 percent, e.g., increases the frequency of price adjustment by approximately 9 percentage points and the magnitude 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 of much of a concern. In reality, fluctuations in the inflation rate may be correlated with other shocks that are not in the regression. There is still some work to do on both measuring the impact of inflation on the two margins, including the study of other episodes and countries, as well as modeling in more detail price-setting behavior, obviously, but our framework gives one potentially interesting alternative in which to explore these issues.

Finally, Klenow-Kryvstov measure the effect of inflation on the fraction of prices that increase and the fraction of prices that decrease. Again, they accomplish this by estimating the coefficient on inflation in regressions of the fraction of prices that increase and on the fraction 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. This was not a foregone conclusion, since it could be, e.g., that inflation induces more positive changes but has little impact on negative changes, or vice-versa. They find that a 1 percent increase in inflation raises the fraction of positive price changes by 5.48 percent and it decreases the fraction of negative changes by 3.10 percent. Figure 9 illustrates the predictions of our model. As in the data, the model predicts increases in inflation raise the fraction of positive and lower the fraction of negative adjustments. However, again, the magnitude of the effect is different than in the regression analysis. According to the model, an increase in inflation from 3 to 4 percent increases the fraction of positive changes by approximately 10 percent and decreases the fraction of negative changes by approximately 2.5 percent. Although we do not match this exactly, we are encouraged by the ability of the model to get the facts qualitatively correct, and think it provides an avenue for further research.

4.2 Summary of Quantitative Findings

It is clear that our theory of price rigidity can account quite well for the empirical behavior of price changes. First, it predicts an average price duration of 11.6 months, which is at the high end of what one sees in the data but still, we think, extremely reasonable. Second, we generate a price-change distribution that has the same overall shape as, and matches the salient features of, the empirical distribution: the average magnitude of is large, yet there are many small price changes, etc. 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. We do not say these are the key features of the empirical price-change distribution because the model does well on these dimensions – these are what are reported to be the key features in the papers mentioned above. Our model also makes predictions about the data not emphasized in the existing literature, including the functional form of the price (as opposed to the price-change) distribution. We have not studied these predictions in detail, but in principle one can try to fit actual price distributions for different products, since the BJ distribution depends on micro parameters like utility, cost and search frictions in particular markets, as well as macro variables like inflation.¹⁵

Existing theories cannot account for all these features of the 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 and size of changes, which suggest that menu costs are large, and the large fraction of price changes that are small, which suggests that menu costs are not large. On the other hand, time-dependent theories of

¹⁵It has been suggested that our model makes the prediction that while p may not increase with inflation for an individual firm, if it does not the q must go up. However, this is also true of many New Keynesian theories, and at some level is nothing more than the law of demand. A much stronger test would be to see if q goes up by enough to keep profit constant, which is really the essence of the BJ model. Although this requires more data, and may be a tough test, we think it is not obvious that it would fail. After all, whenever a seller chooses whether to raise, lower or maintain a price, he should be aware that this affects expected sales. The decision to not raise prices when the average price increases is tantamount to lowering one's relative price, which sellers have to realize increases quantity (again this is the law of demand), and we do not think it is out of the question they believe these effects just offset.

price rigidity (e.g. Calvo 1983 or Taylor 1980) cannot account for the effect of inflation on the frequency of price adjustment, because this is a technological parameter. One theory that matches the empirical behavior of prices reasonably well is the one by Midrigan (2006), which combines elements of state-dependent and time-dependent theories. However, in his model money is not neutral. Based on Midrigan's results, one might conclude that one needs a model where money is not neutral to account for the data, and that certain policy prescriptions are therefore warranted. We show this is not correct, by providing a model consistent with both the facts in which money is exactly neutral.

5 Conclusion

We have provided in this paper a theory of sticky prices that does not impose ad hoc restrictions on repricing decisions. Firms are free to change prices when they like and at no cost. Still our framework is consistent with the sticky-price facts. The model relies on standard search frictions in goods markets, which give rise to equilibrium price dispersion, and (combined with some other assumptions about information) also deliver a genuine role for money. Hence, it is natural that our firms set prices in dollars, and it is permissible for some of them to sometimes not change these prices even when aggregate money and prices change, or when real factors related to preferences and technologies change. Stickiness is a simple corollary of price dispersion. This is true of relative prices in nonmonetary economies, and of nominal prices in monetary economies. Contrary to claims one sees in the context of other models, in our theory price rigidity does not require technological restrictions on repricing, and does not imply central banks can exploit particular policy options, since we designed it to deliver monetary neutrality. This does not imply superneutrality, obviously, and real effects do obtain from changes in money growth, inflation or nominal interest rates. This is relevant because it makes it harder to rule out classical neutrality in the data. How can one be sure, e.g., that any real effects we see result from changes in M , and not changes in (expected) μ , π or i ?

Many extensions of the framework could be considered. One issue is that here we only consider homogeneous sellers. It is well known that in Burdett-Judd models, if every seller has a different marginal cost, we still get price dispersion but firms are not indifferent: in

equilibrium, each seller does have a single target (profit-maximizing) price, depending on real factors. In this case, a change in M leads to a proportional change in p for every firm that preserves their ranking in the relative price distribution, eliminating our nominal rigidities. But if there are only a finite number of seller types, there will be a finite number of BJ distributions, and in the support of each distribution all firms are indifferent. This resurrects our nominal rigidity. We think it is reasonable to assume a finite number of firm types. In the model, as in the real world, the marginal cost of a retailer is the wholesale price – each bottle of shampoo you sell can simply be replaced by your supplier. Even though our retailers sell goods in frictional markets, they buy their inventories in centralized markets, where the wholesale price is the same for all by the law of one price for frictionless markets. There may be deviations from this law in reality – e.g., Walmart may get shampoo at a lower cost than many other sellers, perhaps due to quantity discounts – but we do not believe that this implies every single seller has a different marginal cost. Of course, different retailers may pay different rents, have different opportunity costs etc., but this need not affect their *marginal* cost of sales.

It may be interesting to flesh out the details of the above story, but we think the important part of the message will survive. It might also be worthwhile to analyze models with a finite number of BJ sellers, perhaps all with different marginal costs, under the assumption that they do not see each others' prices (although presumably they can make inferences from their own sales). Would this also deliver real or nominal rigidities? Finally, we took the extreme position that menu costs are literally 0 in our model. This gives rise to some indeterminacy, although we can still take the model to the data. One could study models with $\varepsilon > 0$ menu costs, and imagine a refinement that selects a particular outcome of our model as the result of taking $\varepsilon \rightarrow 0$. It would be interesting to incorporate $\varepsilon > 0$ menu costs in the model future exploration. Of course menu costs that are literally 0 are unrealistic. As Peter Diamond pointed out after the Marshall Lecture, our model misses one obvious fact in the data: it actually does cost something to change your price. Sure, just like it costs something to change your shoes, change your mind or change your password. It is not so obvious, however, that economists ought to base policy recommendations exclusively on theories that rest critically on one such cost to the exclusion of others, or on any combination – say, a cost of changing

your price and your facial expression, not your quantity or your shoes. At the very least, we hope our results can be understood as a cautionary tale about leaping to conclusions about theory or policy from sticky observations.

Appendix

A Proof of Proposition 1

We prove the result in the following five steps.

Claim 1: $\Pi_t^* > 0$.

Proof: For any $\tau > 1$, the profit from posting $p = \tau c / \phi_t$ is

$$\begin{aligned}\Pi_t(\tau c / \phi_t) &= \frac{1}{s} \{ \alpha_1 + 2\alpha_2 [1 - F_t(\tau c / \phi_t)] + \alpha_2 \xi_t(\tau c / \phi_t) \} q_t^*(\tau c / \phi_t) (\tau - 1) c / \phi_t \\ &> \frac{\alpha_1}{s} q_t^*(\tau c / \phi_t) (\tau - 1) c / \phi_t,\end{aligned}$$

where $\xi_t(\tau c / \phi_t)$ was defined immediately after (2.8). Since $q_t^*(\tau c / \phi_t) > 0$ and $\tau > 1$, $\Pi_t(\tau c / \phi_t) > 0$. Hence, $\Pi_t^* \geq \Pi_t(\tau c / \phi_t) > 0$.

Claim 2: F_t is continuous.

Proof: Suppose $\exists p_0 \in \mathcal{F}_t$ such that $\xi_t(p_0) > 0$, and

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

Given $R_t(p)$ is continuous in p , there is a $p_1 < p_0$ such that $R_t(p_1) > 0$ and $\Lambda \equiv R_t(p_0) - R_t(p_1) < \alpha_2 \xi_t(p_0) R_t(p_0) / (\alpha_1 + 2\alpha_2)$. Then

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

where the second line follows from $F(p_0) - F(p_1) \geq \xi_t(p_0)$. With $R_t(p_0) > \Lambda$ and $\Lambda < \alpha_2 \xi_t(p_0) R_t(p_0) / (\alpha_1 + 2\alpha_2)$, (A.1) implies $\Pi_t(p_1) > \Pi_t(p_0)$. This contradicts $p_0 \in \mathcal{F}_t$.

Claim 3: The p_t^m (monopoly price) is the highest price in \mathcal{F}_t .

Proof: Suppose $\bar{p}_t \neq p_t^m$ is the highest price in \mathcal{F}_t . Then

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

With $F_t(p_t^m) \geq 0$, profits at p_t^m satisfy

$$\Pi_t(p_t^m) = \frac{1}{s} \{ \alpha_1 + 2\alpha_2 [1 - F_t(p_t^m)] \} R_t(p_t^m) \geq \frac{\alpha_1}{s} R_t(p_t^m) > \frac{\alpha_1}{s} R_t(\bar{p}_t).\tag{A.3}$$

Now (A.2)-(A.3) imply $\Pi_t(p_t^m) > \Pi_t(\bar{p}_t)$. However, by the equal profit condition, $\Pi_t(\bar{p}_t) = \Pi_t^* \geq \Pi_t(p_t^m)$. This establishes the claim.

Claim 4: \mathcal{F}_t is connected.

Proof: Suppose $p_0, p_1 \in \mathcal{F}_t$ with $p_0 < p_1$ and $F_t(p_0) = F_t(p_1)$. Then

$$\begin{aligned}\Pi_t(p_0) &= \frac{1}{s} \{\alpha_1 + 2\alpha_2 [1 - F_t(p_0)]\} R_t(p_0) \\ \Pi_t(p_1) &= \frac{1}{s} \{\alpha_1 + 2\alpha_2 [1 - F_t(p_1)]\} R_t(p_1)\end{aligned}$$

Since $F_t(p_1) = F_t(p_0)$, we have $\alpha_1 + 2\alpha_2 [1 - F_t(p_1)] = \alpha_1 + 2\alpha_2 [1 - F_t(p_0)]$. Since $p_0, p_1 \in \mathcal{F}_t$, we have $c/\phi_t < p_0 < p_1 \leq p_t^m$. Given $R_t(p)$ strictly increasing $\forall p \in [c/\phi_t, p_t^m]$, $R_t(p_1) > R_t(p_0)$. Combining these results, $\Pi_t(p_1) > \Pi_t(p_0)$, which contradicts $\Pi_t(p_0) = \Pi_t(p_1) = \Pi_t^*$.

Claim 5: F_t is given by

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

Proof: Since F_t has no mass points,

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

At p_t^m , profit is maximized at $\Pi_t^* = \frac{\alpha_1}{s} R_t(p_t^m)$. By equal profit, we get

$$\frac{1}{s} \{\alpha_1 + 2\alpha_2 [1 - F_t(p)]\} R_t(p) = \alpha_1 R_t(p_t^m) \quad \forall p \in [\underline{p}_t, \bar{p}_t]. \quad (\text{A.5})$$

Solving (A.5) for F_t leads to (A.4). ■

B Proof of Proposition 2

It is convenient to rewrite the model in real terms: $n_t = \phi_t \hat{n}_t$ is real money taken out of AD and into BJ, $z_t = \phi_t p$ is the real price associated with nominal price p in BJ and $J_t(z, n)$ is the distribution of real prices in BJ, now written explicitly as depending on real balances in equilibrium. Equivalent to the concepts presented in the text, we now present:

Definition 4: A stationary monetary equilibrium in real variables is a list of AD quantities (h^*, x^*, n^*) , a BJ decision rule $q^*(z, n^*)$ and a BJ distribution $J^*(z, n^*)$ satisfying the following conditions:

1. (h^*, x^*, n^*) solves the household's AD problem, and in particular n^* satisfies the real analog of (2.7),

$$i = \int_0^{\hat{z}(n^*)} [\alpha_1 + 2\alpha_2 - 2\alpha_2 J^*(z, n^*)] \left[\frac{1}{z} \left(\frac{n^*}{z} \right)^{-\gamma} - 1 \right] dJ^*(z, n^*)$$

where $\hat{z}(n^*)$ is the real price below which households cash out in BJ;

2. $q^*(z, n^*)$ solves the household's BJ problem as described by the real analog of (2.5)

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

3. $J^*(z, n^*)$ solves the firm's BJ problem as described by the analog to Proposition 1

$$J^*(z, n^*) = 1 - \frac{\alpha_1}{2\alpha_2} \left[\frac{R(z^m, n^*)}{R(z, n^*)} - 1 \right],$$

where $R(z, n^*) = q^*(z, n^*)(z - c)$, and the support of J^* is $\mathcal{J}^* = [\underline{z}(n^*), \bar{z}(n^*)]$.

There is nothing analogous to ϕ in Definition 1, since the relative price of real balances in terms of x is 1 by construction.

To show equilibrium exists, we now show there exists n^* such that households choose $n = n^*$ given the distribution $J^*(z, n^*)$. This is equivalent to showing there exists n^* such that

$$i = \int_0^{\hat{z}(n)} \left[\frac{1}{z} \left(\frac{n}{z} \right)^{-\gamma} - 1 \right] [\alpha_1 + 2\alpha_2 - 2\alpha_2 J^*(z, n^*)] dJ^*(z, n^*)$$

is solved by $n = n^*$. We proceed in three steps.

Claim 1: Let \underline{n}^* and \bar{n}^* be defined by

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

For n_0^* and n_1^* such that $0 < n_0^* < n_1^* \leq \bar{n}^*$, $J(z, n_0^*)$ first-order stochastically dominates $J(z, n_1^*)$. For n_0^* and n_1^* such that $\bar{n}^* \leq n_0^* < n_1^*$, $J(z, n_0^*) = J(z, n_1^*)$.

Proof: For $n^* \in (0, \underline{n}^*]$,

$$J(z, n^*) = 1 - \frac{\alpha_1}{2\alpha_2} \left\{ \frac{\bar{z}(n^*)^{-1} n^* [\bar{z}(n^*) - c]}{z^{-1} n^* (z - c)} - 1 \right\}$$

with support $[\underline{z}(n^*), \bar{z}(n^*)]$, where

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

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

For $n^* \in [\underline{n}^*, \bar{n}^*]$,

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

with support $[\underline{z}(n^*), \bar{z}(n^*)]$, where

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

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

For $n^* \geq \bar{n}^*$,

$$J(z, n^*) = 1 - \frac{\alpha_1}{2\alpha_2} \left\{ \frac{\bar{z}(n^*)^{-\frac{1}{\gamma}} [\bar{z}(n^*) - c]}{z^{-\frac{1}{\gamma}} (z - c)} - 1 \right\}$$

with support $[\underline{z}(n^*), \bar{z}(n^*)]$, where

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

In this case, $J(z, n_0^*) = J(z, n_1^*)$ for all n_0^* and n_1^* such that $\bar{n}^* \leq n_0^* < n_1^*$.

Claim 2: Given $J(z, n^*)$, let the unique solution for n in the household's AD problem be $n = \psi(n^*)$. Then $\psi(n^*)$ has the following properties:

- $\forall n_0^*, n_1^*$ such that $0 < n_0^* < n_1^* \leq \bar{n}^*$, $\psi(n_0^*) \leq \psi(n_1^*)$;
- $\forall n_0^*, n_1^*$ such that $\bar{n}^* \leq n_0^* < n_1^*$, $\psi(n_0^*) = \psi(n_1^*)$;
- $\forall n^* > 0$, $\psi(n^*) \in [\underline{\psi}, \bar{\psi})$, where $\underline{\psi} > 0$ and $\bar{\psi} = \bar{n}^*$.

Proof: Given the distribution $J(z, n^*)$, the equilibrium condition for n is

$$i = \alpha_1 \int_0^{\hat{z}(n)} \left[\left(\frac{z}{n} \right)^\gamma \frac{1}{z} - 1 \right] dJ(z, n^*) + \alpha_2 \int_0^{\hat{z}(n)} \left[\left(\frac{z}{n} \right)^\gamma \frac{1}{z} - 1 \right] d\{1 - [1 - J(z, n^*)]^2\}.$$

Let $\chi(n, n^*)$ denote the RHS of this equation. Notice $\lim_{n \rightarrow 0} \chi(n, n^*) = \infty$ and $\chi(n, n^*)$ is strictly decreasing in $n \forall n \in (0, \hat{z}^{-1}[\underline{z}(n^*)])$. Also, $\chi(n, n^*) = 0 \forall n \geq \hat{z}^{-1}[\underline{z}(n^*)] = \underline{z}(n^*)^{-\frac{1-\gamma}{\gamma}}$. From these observations, there is a unique solution $n = \psi(n^*)$ to $i = \chi(n, n^*)$. Moreover, $0 < \psi(n^*) < \hat{z}^{-1}[\underline{z}(n^*)]$.

Consider any n_0^*, n_1^* such that $0 < n_0^* < n_1^* \leq \bar{n}^*$. Claim 1 implies that $J(z, n_0^*)$ first-order stochastically dominates $J(z, n_1^*)$ and, consequently, $1 - [1 - J(z, n_0^*)]^2$ first-order stochastically dominates $1 - [1 - J(z, n_1^*)]^2$. From this and the fact that $(z/n)^\gamma / z - 1$ is decreasing in z , it follows that $\chi(n, n_0^*) \leq \chi(n, n_1^*) \forall n$. Therefore, $\psi(n_0^*) \leq \psi(n_1^*)$. Moreover, it is straightforward to verify that $\psi(n_0^*) \geq \underline{\psi}$ for some $\underline{\psi} > 0$.

Now consider any n_0^*, n_1^* such that $\bar{n}^* \leq n_0^* < n_1^*$. Claim 1 implies that $J(z, n_0^*) = J(z, n_1^*)$ and $1 - [1 - J(z, n_0^*)]^2 = 1 - [1 - J(z, n_1^*)]^2$. It follows that $\chi(n, n_0^*) = \chi(n, n_1^*)$, and hence $\psi(n_0^*) = \psi(n_1^*)$. It is straightforward to verify $\psi(n_1^*) < \bar{n}^*$.

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

Proof: Claim 2 implies: $\psi(n^*)$ is increasing and $\psi(n^*) \in [\underline{\psi}, \bar{\psi}) \forall n^* \in [\underline{\psi}, \bar{\psi}]$. By Tarski's theorem, $\exists n^* \in [\underline{\psi}, \bar{\psi})$ such that $\psi(n^*) = n^*$. This establishes the existence of equilibrium. Moreover, Claim 2 implies $\psi(n^*) \geq \underline{\psi} \forall n^* < \underline{\psi}$ and $\psi(n^*) < \bar{\psi} \forall n^* \geq \bar{\psi}$. Hence, $\nexists n^* \notin [\underline{\psi}, \bar{\psi})$ such that $\psi(n^*) = n^*$. ■

C Proof of Proposition 4

We take two steps to prove the results.

Claim 1: In any stationary monetary equilibrium

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

where $\Lambda = 2\alpha_2\gamma / [2\alpha_2 + \alpha_1(1 - \gamma)]$.

Proof: Households are either constrained in all BJ transactions, or constrained in some but not others. First suppose they are constrained in all transactions. Then

$$\phi_t \bar{p}_t = (\phi_t \hat{m}_t^*)^{\frac{\gamma}{\gamma-1}} \geq \frac{c}{1-\gamma}, \quad (\text{A.7})$$

$$\phi_t \underline{p}_t = c \left(1 - \frac{\alpha_1}{\alpha_1 + 2\alpha_2} \frac{\phi_t \bar{p}_t - c}{\phi_t \bar{p}_t} \right)^{-1}. \quad (\text{A.8})$$

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

$$\frac{\bar{p}_t - \underline{p}_t}{\bar{p}_t} = \frac{2\alpha_2(\phi_t \bar{p}_t - cw)}{2\alpha_2\phi_t \bar{p}_t + \alpha_1 c} \geq \frac{2\alpha_2\gamma}{2\alpha_2 + \alpha_1(1 - \gamma)} = \Lambda. \quad (\text{A.9})$$

Now suppose households are constrained in only some transactions. Then

$$\phi_t \bar{p}_t = \frac{c}{1-\gamma} \geq (\phi_t \hat{m}_t^*)^{\frac{\gamma}{\gamma-1}}, \quad (\text{A.10})$$

$$\phi_t \underline{p}_t = c \left[1 - \frac{\alpha_1}{\alpha_1 + 2\alpha_2} (\phi_t \bar{p}_t)^{-\frac{1}{\gamma}} \left(\frac{\phi_t \bar{p}_t - c}{\phi_t \hat{m}_t^*} \right) \right]^{-1}. \quad (\text{A.11})$$

From (A.10)-(A.11),

$$\frac{\bar{p}_t - \underline{p}_t}{\bar{p}_t} = \frac{(\alpha_1 + 2\alpha_2)\gamma\phi_t \hat{m}_t^* - \alpha_1\gamma \left(\frac{c}{1-\gamma} \right)^{\frac{\gamma-1}{\gamma}}}{(\alpha_1 + 2\alpha_2)\phi_t \hat{m}_t^* - \alpha_1\gamma \left(\frac{c}{1-\gamma} \right)^{\frac{\gamma-1}{\gamma}}} \geq \frac{2\alpha_2\gamma}{2\alpha_2 + \alpha_1(1 - \gamma)} = \Lambda \quad (\text{A.12})$$

where the inequality uses $\phi_t \hat{m}_t^* \geq [(1 - \gamma)/c]^{\frac{1-\gamma}{\gamma}}$. Combining (A.9) and (A.12), we obtain (A.6).

Claim 2: Let $\mu \in (1, \mu^*)$, where $\mu^* = (1 - \Lambda)^{-1}$. Then, given repricing policy $p_{t+1}^*(p, \rho)$, we have: $\Phi(\rho) < 1$ and $A(\rho) > 0 \forall \rho \in (0, 1]$; and $H_t(0, \rho) > 0 \forall \rho \in [0, 1)$.

Proof: From Claim 1 and $\mu \in (1, \mu^*)$, we have

$$\mu \underline{p}_t < \frac{\underline{p}_t}{1 - \Lambda} \leq \bar{p}_t. \quad (\text{A.13})$$

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

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

using (A.13) and $F_t^* \left[(1 - \Lambda)^{-1} \underline{p}_t \right] < F_t^*(\bar{p}_t) = 1$. Since the fraction of prices that adjust is less than 1, $A(\rho) > 1$. Finally, it is easy to verify $J_t(0, \rho) > 0 \forall \rho \in [0, 1)$. ■

References

- [1] Albrecht, J. and B. Axel. 1984. "An Equilibrium Model of Search Unemployment." *Journal of Political Economy* 92, 824-40.
- [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] Bils, M., and P. Klenow. 2005. "Some Evidence on the Importance of Sticky Prices." *Journal of Political Economy* 112, 947-85.
- [5] Burdett, K., and K. Judd. 1983. "Equilibrium Price Dispersion." *Econometrica* 51, 955-70.
- [6] Burdett, K., and D.T. Mortensen. 1998. "Wage Differentials, Employer Size, and Unemployment." *International Economic Review* 39, 257-73.
- [7] Campbell, J., and B. Eden. 2007. "Rigid Prices: Evidence from U.S. Scanner Data." Working Paper 2005-08, Federal Reserve Bank of 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., J. Gali and M. Gertler. 1999. "The Science of Monetary Policy: A New Keynesian Perspective." *Journal of Economic Literature* 37, 1661-1707.
- [13] Jiang, J., and M. Dong. 2011. "Money and Price Posting under Private Information." Mimeo, Bank of Canada.

- [14] 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.
- [15] Eden, B. 1994. "The Adjustment of Prices to Monetary Shocks when Trade is Uncertain and Sequential." *Journal of Political Economy* 102, 493-509.
- [16] Eichenbaum, M., N. Jaimovich and S. Rebelo. 2009. "References Prices, Costs and Nominal Rigidities." Mimeo, Northwestern University
- [17] Ennis, H. 2008. "Search, Money, and Inflation under Private Information." *Journal of Economic Theory* 138, 101-131.
- [18] Faig, M., and B. Jerez. 2005. "A Theory of Commerce." *Journal of Economic Theory* 122, 60-99.
- [19] Gagnon, E. 2007. "Price Setting During Low and High Inflation: Evidence from Mexico." International Finance Discussion Paper 2007-896, The Federal Reserve Board.
- [20] Golosov, M., and R.E. Lucas, Jr. 2003. Menu Costs and Phillips Curves." NBER Working Paper 10187.
- [21] Golosov, M., and R.E. Lucas, Jr. 2007. "Menu Costs and Phillips Curves." *Journal of Political Economy* 115, 171-99.
- [22] Head, A., and A. Kumar. 2005. "Price Dispersion, Inflation, and Welfare." *International Economic Review* 46, 533-72.
- [23] Head, A., A. Kumar and B. Lapham. 2010. "Market Power, Price Adjustment, and Inflation." *International Economic Review*.
- [24] Head, A., L.Q. Liu, G. Menzies and R. Wright. 2011. "Commitment, Costly Credit, and Sticky Prices." Mimeo.
- [25] Jean, K., S. Rabinovich and R. Wright. 2010. "On the Multiplicity of Monetary Equilibria: Green-Zhou meets Lagos-Wright." *Journal of Economic Theory* 145, 392-401.

- [26] Kiyotaki, N., and R. Wright. 1989. "On Money as a Medium of Exchange." *Journal of Political Economy* 97, 927-954.
- [27] 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, 863-904.
- [28] Klenow, P., and B. Malin. 2010. "Microeconomic Evidence on Price-Setting." *Handbook of Monetary Economics*, forthcoming.
- [29] Kocherlakota, N. 1998. "Money is Memory." *Journal of Economic Theory* 81, 232-251.
- [30] Lagos, R., and R. Wright. 2005. "A Unified Framework for Monetary Theory and Policy Analysis." *Journal of Political Economy* 113, 463-88.
- [31] Liu, L.Q. 2010. *Essays on Inflation and Output: A Search Theoretic Approach*. Ph.D. Dissertation, Queen's University.
- [32] Lucas, R.E., Jr. 1972. "Expectations and the Neutrality of Money." *Journal of Economic Theory* 4, 103-124.
- [33] Mankiw, G. 1985. "Small Menu Costs and Large Business Cycles: A Macroeconomic Model." *Quarterly Journal of Economics* 100, 529-38.
- [34] McCall, J. 1970. "Economics of Information and Job Search." *Quarterly Journal of Economics* 84, 113-26.
- [35] Midrigan, V. 2006. "Menu Costs, Multi-Product Firms, and Aggregate Fluctuations." Manuscript, New York University.
- [36] Mortensen, D.T. 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.
- [37] Mortensen, D.T. 2005. "A Comment on 'Price Dispersion, Inflation, and Welfare'." *International Economic Review* 46, 573-578.

- [38] Mortensen, D.T., and C.A. Pissarides. 1999. “New Developments in Models of Search in the Labor Market.” In *Handbook in Labor Economics*. O. Ashenfelter and D. Card, eds. Amsterdam: North Holland.
- [39] Nakamura, E., and J. Steinsson. “Five Facts About Prices: A Reevaluation of Menu Cost Models.” *Quarterly Journal of Economics* 123, 1415-1464.
- [40] Nosal, E., and G. Rocheteau. 2011. *Money, Payments and Liquidity*. Cambridge: MIT Press.
- [41] Pissarides, C.A. 2000. *Equilibrium Unemployment Theory*. Oxford: Basil Blackwell.
- [42] Rothschild, M. 1973. “Models of Market Organization with Imperfect Information: A Survey.” *Journal of Political Economy* 81, 1283-308.
- [43] Shi, S. 1997. “A Divisible Model of Fiat Money.” *Econometrica* 65, 75-102.
- [44] Tarski, A. 1955. “A Lattice Theoretical Fixpoint Theorem and Its Applications.” *Pacific Journal of Mathematics* 5, 285-309.
- [45] Taylor, J. 1980. “Aggregate Dynamics and Staggered Contracts.” *Journal of Political Economy* 88, 1-23.
- [46] Wallace, N. 2010. “The Mechanism Design Approach to Monetary Economics.” In *Handbook of Monetary Economics*, B. Friedman and M. Woodford, eds. Elsevier.
- [47] Wang, L. 2011. *Essays on Monetary Economics with a Microfoundation, Price Dispersion and Inflation*. Ph.D. Dissertation, University of Pennsylvania.
- [48] Williamson, S., and R. Wright. 2010a. “New Monetarist Economics: Methods.” *FRB St. Louis Review* 92.
- [49] Williamson, S., and R. Wright. 2010b. “New Monetarist Economics: Models.” In *Handbook of Monetary Economics*, B. Friedman and M. Woodford, eds. Elsevier.
- [50] Woodford, M. 2003. *Interest and Prices*. Princeton: Princeton University Press.

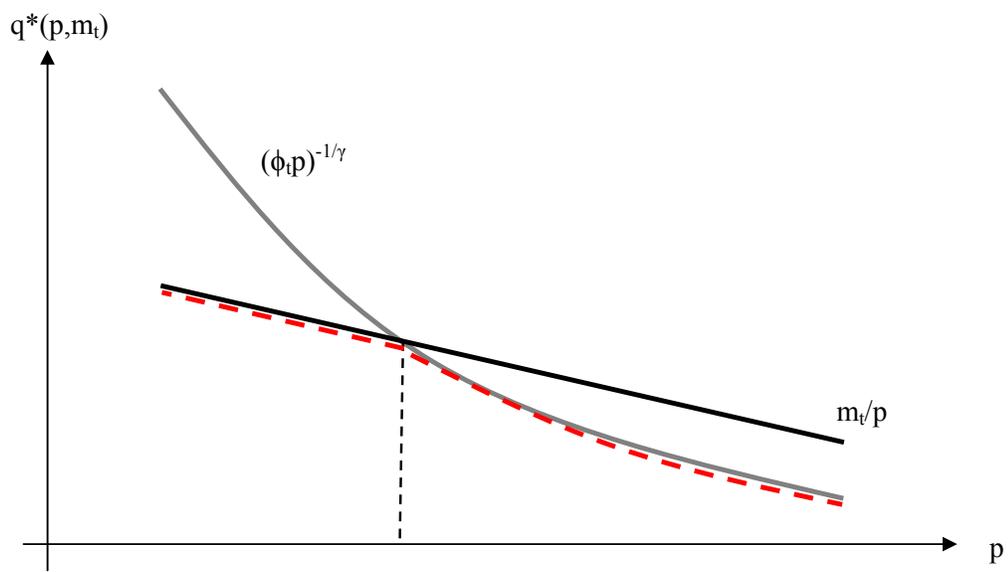


Figure 1: Household's demand for the BJ good

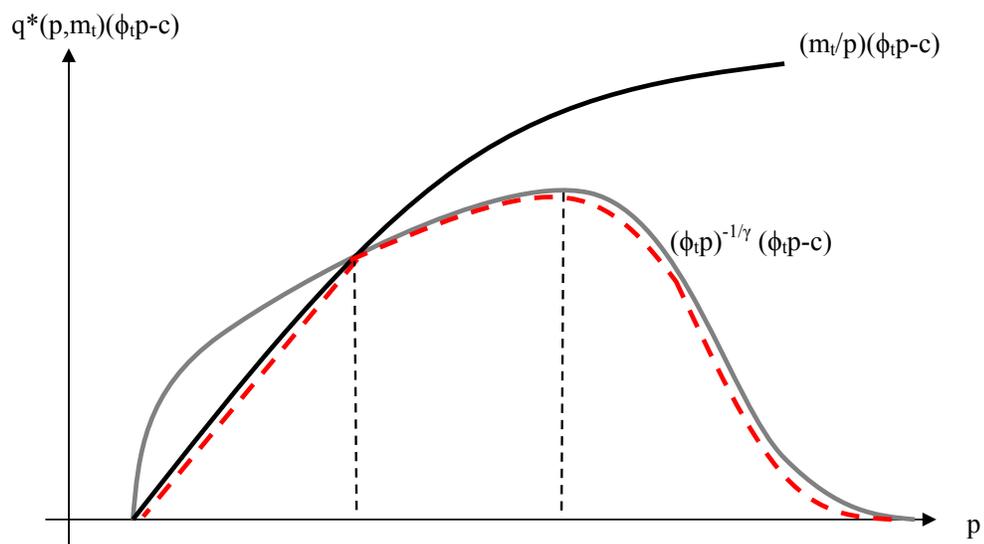


Figure 2: Firm's profit per customer in the BJ market

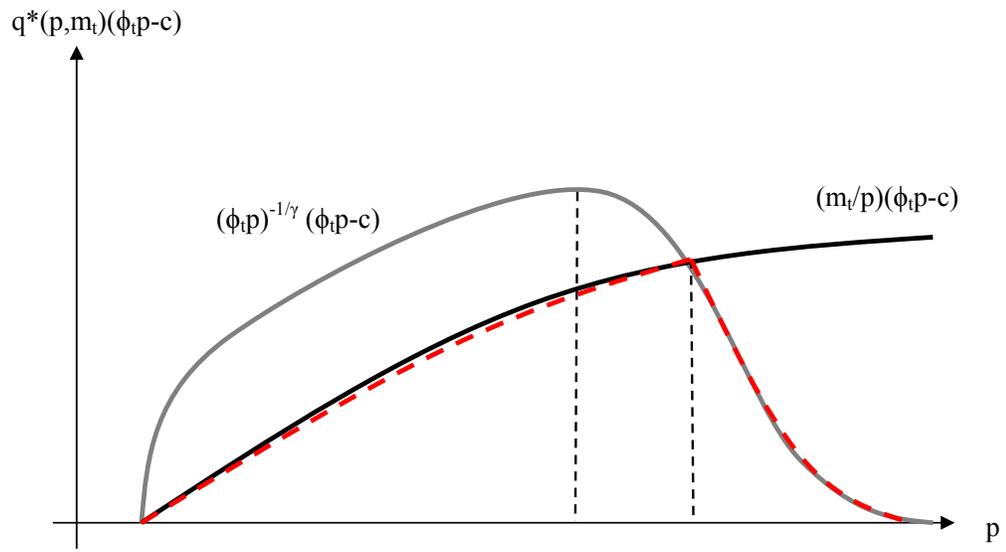


Figure 3: Firm's profit per customer in the BJ market

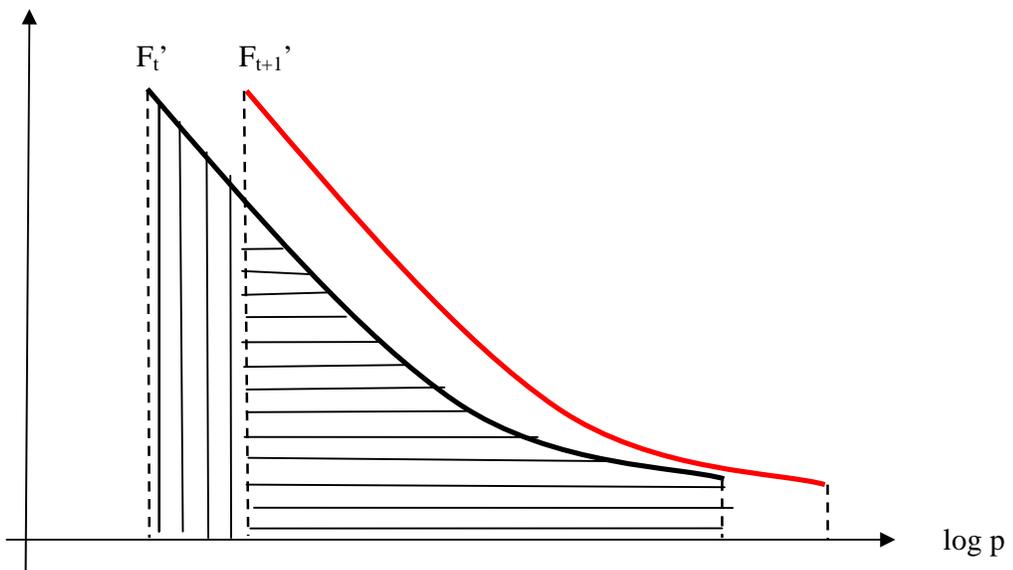


Figure 4: Equilibrium price distribution

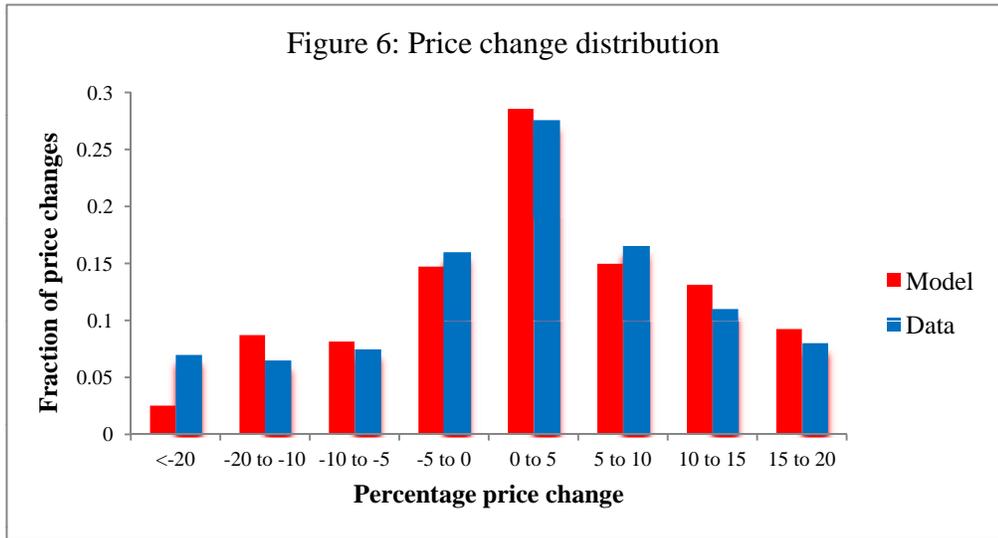
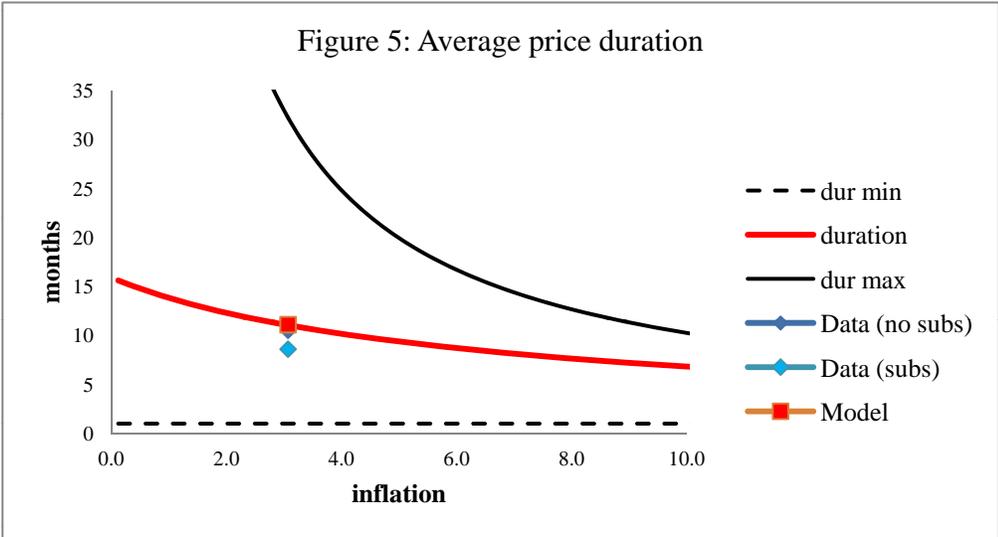


Figure 7: Hazard rate of a price change

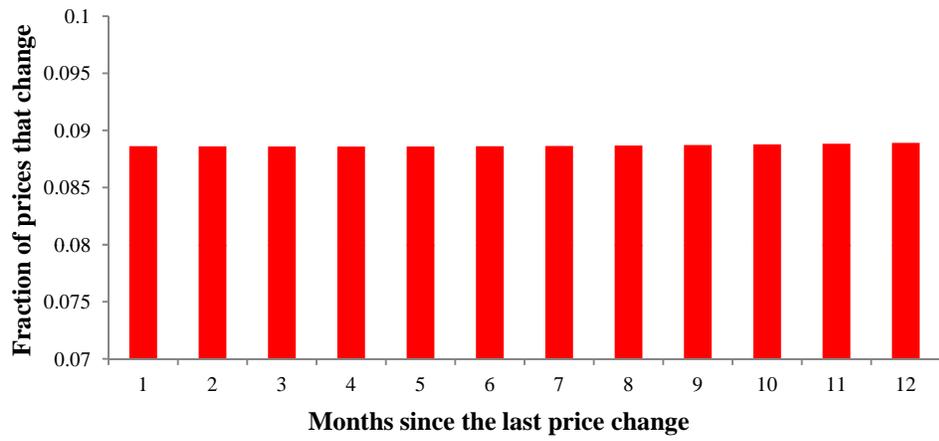


Figure 7: Hazard rate of a price change

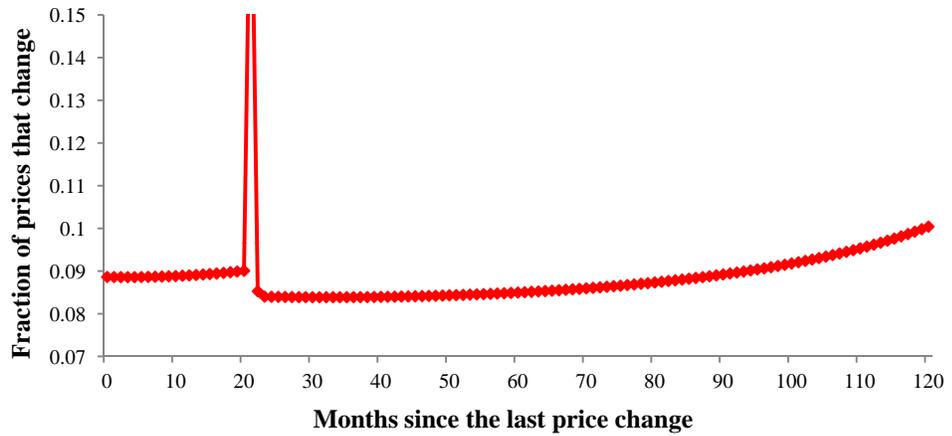


Figure 8: Fraction and size of price changes

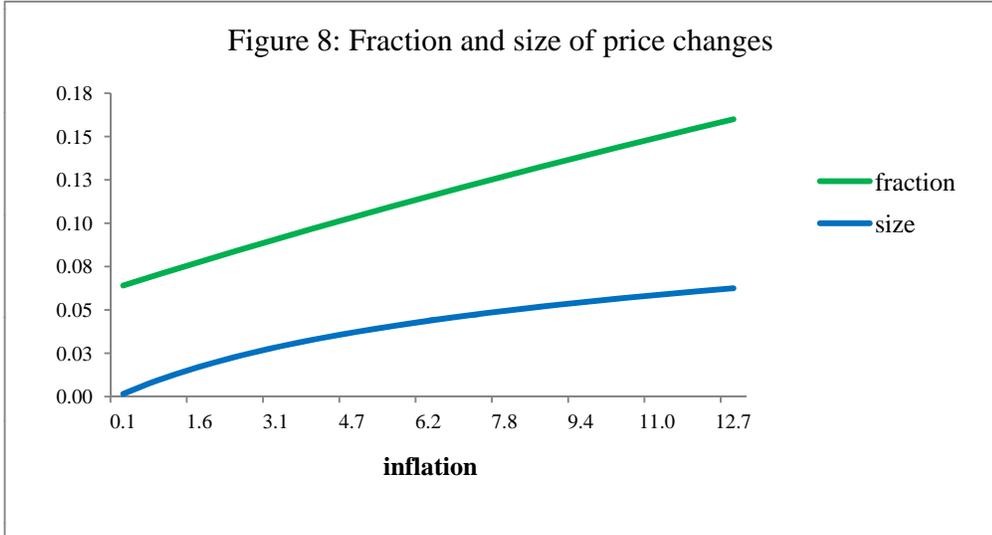


Figure 9: Positive and negative price changes

