

Logit price dynamics

James Costain*

*Banco de España

Anton Nakov†

†European Central Bank and CEPR

September 2013

Abstract

We model retail price stickiness as the result of errors due to costly decision-making. Under our assumed cost function for the precision of choice, the timing of price adjustments and the prices firms set are both logit random variables. Errors in the prices firms set help explain micro “puzzles” relating to the sizes of price changes, the behavior of adjustment hazards, and the variability of prices and costs. Errors in adjustment timing increase the real effects of monetary shocks, by reducing the “selection effect”. Allowing for both types of errors also helps explain how trend inflation affects price adjustment.

Keywords: Nominal rigidity, logit equilibrium, state-dependent pricing, near rationality, information-constrained pricing

JEL Codes: E31, D81, C73

1 Introduction¹

Economists seeking to explain price stickiness have often appealed to small fixed costs of nominal price changes, commonly called “menu costs” (Barro 1972). In theory, even small menu costs might make price adjustments infrequent and make aggregate dynamics differ substantially from the flexible-price optimum (Mankiw 1985). But quantitatively, Golosov and Lucas (2007) showed that fixed menu costs do little to generate aggregate price stickiness in a macroeconomic model with realistically large firm-specific shocks. The dynamics of their model are quite close to monetary neutrality, so fixed menu costs seem unpromising to explain the nontrivial real effects of monetary shocks observed in macroeconomic data (*e.g.* Christiano, Eichenbaum, and Evans, 1999). Moreover, detailed microeconomic evidence suggests that menu costs, as usually interpreted, are only a small fraction of the overall costs of price setting

¹We thank Hervé Le Bihan, Jeff Campbell, Anton Cheremukhin, Robert Gary-Bobo, Peter Karadi, Andrew Levin, Bartosz Mackowiak, Filip Matejka, Isabelle Méjean, Galo Nuño, Ricardo Reis, Luminita Stephens, Antonella Tutino, Carl Walsh, Mirko Wiederholt, Michael Woodford, and seminar participants at UC Santa Cruz, the Bank of Spain, CREST, FRB Richmond, the ECB, BI Business School, T2M 2012, ESSIM 2012, CEF 2012, EEA-ESEM 2012, and the 2013 Chicago NBER Monetary Economics meeting for helpful comments. We are grateful to Virgiliu Midrigan for making his data available to us, and to the James M. Kilts Center at the Univ. of Chicago GSB, which is the original source of those data. Views expressed here are those of the authors and do not necessarily coincide with those of the Bank of Spain, the Eurosystem, the ECB, or the CEPR.

(Zbaracki *et al.* 2004). A much larger part of the costs of price adjustment consists of managerial costs associated with information collection and decision making. This raises the question: can costs related to decision making explain microeconomic and macroeconomic evidence of price stickiness better than fixed menu costs do? And furthermore, how exactly should these costs be modeled?

This paper proposes a simple model of price stickiness based on costly decision-making, estimates its two free parameters, and shows by simulation that it is consistent with a wide variety of microeconomic and macroeconomic evidence. Two key considerations motivate our setup. First, if choice is costly, then decisions will typically be imperfect, that is, prone to errors. Thus it is natural to think of the decision outcomes as random variables, instead of treating actions as deterministic. Second, it is natural to assume that more precise decisions are more costly than imprecise ones. Motivated by these points, we adopt the “control cost” approach from game theory (see, for example, van Damme 1991). Formally, instead of modeling the choice of an optimal action directly, this approach describes the decision problem as the choice of a probability distribution over possible actions. The problem is solved subject to a cost function such that more precise decisions (more concentrated distributions) are more expensive. Making any given decision in a perfectly precise way is feasible, but this is usually not worth the cost. Therefore the action actually taken will be a random variable correlated with fundamentals, instead of being a deterministic function of fundamentals.

In the context of dynamic price setting, a firm faces two key margins of decision: *when* to change the price of a product it sells, and *what new price* to set. We allow for errors on both of these margins. The exact shape of the error distribution depends on the assumed cost function for precision. It happens to be particularly convenient to measure precision in terms of entropy, and to define costs as a linear function of the relative entropy of the distribution of actions, as compared with a uniform distribution. Under these functional forms, the distribution of actions is a multinomial logit. This has the desirable implication that the probability of taking any given action increases smoothly with the value of that action, compared with the values of other feasible actions. General equilibrium then takes the form of a logit equilibrium:² each decision maker plays a logit in which the values of actions are evaluated assuming that other decision makers play logits too. The decision costs backed out from our benchmark calibration do not seem excessive: firms spend roughly 0.9% of revenue on decision-making, and in addition incur a loss of roughly 0.5% of revenue due to suboptimal choices.

The fact that an entropy-related cost function can “microfound” a logit distribution of actions has been shown by many previous authors in game theory and economics (Stahl 1990; Marsili 1999; Mattson and Weibull 2002; Bono and Wolpert 2009; Matejka and McKay 2011).³ However, economics applications have typically focused on decisions taken at known, exogenously given points in time; it is not immediately obvious how to apply the logit framework to a context of intermittent adjustment where a key question is *when* changes should occur. We study how the derivation of logit choice behavior can

²Logit equilibrium is a commonly-applied parametric special case of quantal response equilibrium (see McKelvey and Palfrey, 1995, 1998).

³This reflects much older results in physics, where a related optimization problem gives rise to the Boltzmann distribution of particles in a gas.

be extended so that it is applicable to fully dynamic decisions of timing. We show that if the decision cost associated with a time-varying adjustment hazard is a linear function of its relative entropy, compared with a constant adjustment hazard, then the decision to adjust or not in a given time period is governed by a *weighted* binary logit. While a standard static logit model has a single free parameter representing the accuracy of decisions, the weighted logit in our dynamic setup has two free parameters, related to the *speed* and the *accuracy* of decision making. The inclusion of the speed parameter ensures that our model has a well-defined continuous-time limit, and thus clarifies how parameters must be adjusted if the frequency of the data or the model simulation is changed.

While it is reasonable to assume that the size and the timing of firms' adjustments are both subject to error, we run simulations that shut down one type of mistakes or the other in order to see what each one contributes empirically. We find that errors in the *size* of price changes help explain a number of observations from retail price microdata that represent "puzzles" for many standard models. In particular, unlike a fixed menu cost model, our setup implies that many large and small price changes coexist (Klenow and Kryvstov 2008; Midrigan 2011; Klenow and Malin 2010, "Fact 7"). It implies that the adjustment hazard is nearly flat, but slightly decreasing in the first few months, as found by empirical studies that control for heterogeneity in adjustment frequency (Nakamura and Steinsson 2008, "Fact 5"; Klenow and Malin 2010, "Fact 10"). Furthermore, we find that the standard deviation of price adjustments is mostly constant, independent of the time since last adjustment (Klenow and Malin 2010, "Fact 10"). Many alternative models, including the Calvo model, instead imply that price adjustments are increasing in size. Also, we find that exceptionally high or low prices are more likely to have been set recently than prices near the center of the distribution (Campbell and Eden 2010). Finally, prices are more volatile than costs, as documented by Eichenbaum, Jaimovich, and Rebelo (2011), whereas the opposite is true in both the Calvo and fixed menu cost models.

While errors in the size of price adjustments help reproduce patterns in microdata, by themselves they do not imply strong real effects of monetary policy. Indeed, since the model with errors only in the size of price adjustments has only one free parameter, it is hard to get it to match multiple features of the data simultaneously, and the degree of nonneutrality it implies is sensitive to the details of our calibration procedure. But whenever we include mistakes in the timing of price adjustments, our model implies substantial monetary nonneutrality (roughly halfway between the effects observed in the fixed menu cost model, and those observed in the Calvo model). The cause of the nonneutrality is the same as in the Calvo model: by decreasing the relation between the value of adjustment and the probability of adjustment, the "selection effect" highlighted by Caplin and Spulber (1987) and Golosov and Lucas (2007) is reduced. But in contrast with the Calvo setup, our model also does a good job in reproducing the effects of trend inflation on price adjustment. In particular, it is consistent with the effect of trend inflation on the typical size of price changes, and on the fraction of adjustments that are increases, which are both margins where the fixed menu cost model performs poorly.

1.1 Related literature

This paper connects with several areas of economic literature. A wave of recent research has documented the dynamics of price adjustment in new databases from the retail sector (key papers include Klenow and Kryvtsov, 2008; Nakamura and Steinsson, 2008; and Klenow and Malin, 2010; and Eichenbaum *et al.*, 2011). In response, many macroeconomists have simulated numerical models of pricing under fixed or stochastic menu costs in the presence of aggregate and firm-specific shocks, fitting them to microdata and then studying their macroeconomic implications. Some influential papers in this tradition include Golosov and Lucas (2007), Midrigan (2011), Dotsey, King, and Wolman (2013), Álvarez, Beraja, González, and Neumeyer (2011), Kehoe and Midrigan (2010), and Matejka (2011).⁴ While some of these recent generalized menu cost models can match many empirical price facts, they typically have far more free parameters than our model does. A particularly promising recent branch of the literature instead considers both a fixed cost of price adjustment and a fixed cost of acquiring information (Álvarez, Lippi, and Paciello, 2011; Demery, 2012). Like our own framework, these “menu cost and observation cost” models are highly empirically successful in spite of relying on only two free parameters to model the adjustment process.

While much recent work on state-dependent pricing assumes prices are set optimally subject to menu costs, we assume instead that price adjustment involves errors, and we do not assume any menu costs, at least not as they are usually interpreted. The fact that we allow for errors may seem like a break with standard practice in macroeconomics, but it is consistent with microeconometrics, where error terms are indispensable (though they are not always interpreted as mistakes). In representative-agent macroeconomic modeling, ignoring errors is not necessarily inconsistent with microeconometrics, since to a first approximation, errors might cancel out. But when calibrating a heterogeneous-agent macroeconomic model to the full distribution of adjustments in microdata, such an argument does not apply: if there are any errors at all, these are likely to increase the variance of observed adjustments, so that a calibration without errors would (for example) mistakenly overestimate the variance of the underlying exogenous shocks. In this sense, the microdata-based calibration strategies in most recent literature on state-dependent pricing may represent a more radical departure from previous microeconomic and macroeconomic methodology than our model does.

The logit equilibrium framework for modeling error-prone behavior has been widely applied in experimental game theory, where it has helped explain play in a number of games in which Nash equilibrium performs poorly, such as the centipede game and Bertrand competition games (McKelvey and Palfrey 1998; Anderson, Goeree, and Holt 2002). It has been much less frequently applied in other areas of economics; we are unaware of any application of logit equilibrium inside a dynamic general equilibrium macroeconomic model, other than our own work.⁵ Macroeconomists’ reluctance to allow for errors may

⁴This paper also builds on two related papers of our own: in Costain and Nakov (2011C) we study the microeconomic and macroeconomic implications of logit errors in price decisions, while one specification considered in Costain and Nakov (2011A) imposes logit errors on the timing of price adjustment.

⁵The logit choice function is a standard econometric framework for discrete choice, and has been applied in many microe-

derive partly from discomfort with the many degrees of freedom opened up by abandoning the benchmark of full rationality. However, since logit equilibrium is just a one- or two-parameter generalization of fully rational choice, it actually imposes much of the discipline of rationality on the model.⁶

While McKelvey and Palfrey defined logit equilibrium both for extensive form (1998) and normal form (1995) games, we found it necessary to extend their framework in order to deal with the *timing* of price adjustment.⁷ Our setup applies the same logic to decisions on the timing margin that it applies on the pricing margin. In a static context, logit choice is derived by penalizing the entropy of the random choice, relative to a uniform distribution. Likewise, we derive a weighted binary logit governing the timing of adjustment by penalizing the entropy of the random time of adjustment, relative to a uniform adjustment hazard. In other words, precision in the *size* of the adjustment is measured by comparing the price distribution to a uniform distribution; likewise, precision in the *timing* of adjustment is measured by comparing the state-dependent hazard rate to a *Calvo model*. The hazard we derive from this specification has the same functional form derived by Woodford (2008), though his microfoundations differ: he assumes firms face a constraint on information flow, plus a fixed cost of purchasing full information.

Woodford's (2008, 2009) papers form part of the "rational inattention" literature that follows Sims (2003) by assuming economic agents face costs associated with information flow in the sense of Shannon (1948). Measuring the precision of choices in terms of entropy is a feature our model shares with the rational inattention approach.⁸ The only difference is that in our setup, the probability distribution over a firm's actions is conditioned on the firm's true state, whereas under rational inattention the distribution of actions is conditional on the firm's *prior* about its true state. In other words, Sims assumes the true state of the world is never known with certainty, whereas under the "control cost" approach, making optimal choices is costly in spite of the fact that the true state of the world is known. Ultimately, the main reason we focus on costs of decision making *per se* instead of costs of information transmission is a practical one: it makes our model "infinitely" easier to solve than those of Sims (2003) and Matejka (2011), by dramatically reducing the dimensionality of the problem. A firm acting under rational inattention must condition on a prior over its possible productivity levels (a very high dimensional object), whereas in our setup, the firm just conditions on its true productivity level. Moreover, once one knows that entropy reduction costs imply logit, one can simply impose a logit function directly (and then subtract off the implied costs) rather than explicitly solving for the form of the error distribution. These facts make our approach entirely tractable in a DSGE context, as this paper will show.

conometric contexts. But logit *equilibrium*, in which each player makes logit decisions, based on payoff values consistent with other players' logit decisions, has to the best of our knowledge rarely been applied outside of experimental game theory.

⁶Haile, Hortaçsu, and Kosenok (2008) have argued that quantal response equilibrium, which has an infinite number of free parameters, is impossible to reject empirically. However, this criticism does not apply to logit equilibrium (the special case of quantal response equilibrium which has been most widely applied in practice) since it is very tightly parameterized.

⁷Initially we thought that an extensive form game with a choice between adjustment and nonadjustment at each point in time would suffice to model the timing decision. But such a framework turns out to be sensitive to the assumed time period: for a given logit rationality parameter, decreasing the model period eventually drives errors in the timing of adjustment to zero. Essentially, this approach fails because it does not allow for a free parameter measuring the speed of decision-making relative to the time scale of the model.

⁸Another recent application of entropy in economics is the "robust control" methodology of Hansen and Sargent (2008).

2 Model

This discrete-time model embeds near-rational price adjustment in an otherwise standard New Keynesian general equilibrium framework based on Golosov and Lucas (2007). Prices are set by a continuum of monopolistically competitive firms. In addition, there is a representative household, and a monetary authority that sets an exogenous growth process for the nominal money supply.

2.1 Household

The household's period utility function is $\frac{1}{1-\gamma}C_t^{1-\gamma} - \chi N_t + \nu \log(M_t/P_t)$, where C_t is consumption, N_t is labor supply, and M_t/P_t is real money balances. Utility is discounted by factor β per period. Consumption is a CES aggregate of differentiated products C_{it} , with elasticity of substitution ϵ :

$$C_t = \left\{ \int_0^1 C_{it}^{\frac{\epsilon-1}{\epsilon}} di \right\}^{\frac{\epsilon}{\epsilon-1}}. \quad (1)$$

The household's nominal period budget constraint is

$$\int_0^1 P_{it} C_{it} di + M_t + R_t^{-1} B_t = W_t N_t + M_{t-1} + B_{t-1} + T_t^M + T_t^D, \quad (2)$$

where $\int_0^1 P_{it} C_{it} di$ is total nominal consumption. B_t represents nominal bond holdings, with interest rate $R_t - 1$; T_t^M is a lump sum transfer from the central bank, and T_t^D is a dividend payment from the firms.

Households choose $\{C_{it}, N_t, B_t, M_t\}_{t=0}^{\infty}$ to maximize expected discounted utility, subject to the budget constraint (2). Optimal consumption across the differentiated goods implies

$$C_{it} = (P_{it}/P_t)^{-\epsilon} C_t, \quad (3)$$

so nominal spending can be written as $P_t C_t = \int_0^1 P_{it} C_{it} di$ under the price index

$$P_t \equiv \left\{ \int_0^1 P_{it}^{1-\epsilon} di \right\}^{\frac{1}{1-\epsilon}}. \quad (4)$$

The household's first-order conditions for labor supply, consumption, and money use can be written as follows:

$$\chi = C_t^{-\gamma} W_t / P_t, \quad (5)$$

$$R_t^{-1} = \beta E_t \left(\frac{P_t C_{t+1}^{-\gamma}}{P_{t+1} C_t^{-\gamma}} \right), \quad (6)$$

$$1 - \frac{\nu P_t}{M_t C_t^{-\gamma}} = \beta E_t \left(\frac{P_t C_{t+1}^{-\gamma}}{P_{t+1} C_t^{-\gamma}} \right). \quad (7)$$

2.2 Monopolistic firms

Each firm i produces output Y_{it} under a constant returns technology $Y_{it} = A_{it}N_{it}$. Here labor N_{it} is the only input, and $A_{it} \equiv \exp(a_{it})$ is an idiosyncratic productivity process. Log productivity a_{it} follows a time-invariant Markov process on a bounded set, $a_{it} \in \Gamma^a \subseteq [\underline{a}, \bar{a}]$, and productivity innovations are *iid* across firms. Thus, a_{it} is correlated with $a_{i,t-1}$, but it is uncorrelated with other firms' shocks. Firm i is a monopolistic competitor that sets a price P_{it} , facing the demand curve $Y_{it} = C_t P_t^\epsilon P_{it}^{-\epsilon}$. Note that since firms are infinitesimal, each firm i assumes that its own price P_{it} has no effect on the aggregate price level P_t . It hires in a competitive labor market at wage rate W_t , generating profits

$$U_{it} = P_{it}Y_{it} - W_tN_{it} = \left(P_{it} - \frac{W_t}{A_{it}}\right) C_t P_t^\epsilon P_{it}^{-\epsilon} \quad (8)$$

per period. Firms are owned by the household, so they discount nominal income between times t and $t+1$ at the rate $\beta \frac{P_t C_{t+1}^{-\gamma}}{P_{t+1} C_t^{-\gamma}}$, consistent with the household's marginal rate of substitution.

We assume each firm must fulfill all demand at its chosen price. Therefore, its only decisions are when to adjust its price, and what price to set upon adjustment. The firm may make errors in either of these choices. We discuss these two decisions in turn, beginning with the latter.

2.2.1 Choosing a new price

Our model develops the idea that nominal rigidities may be derived primarily from decision costs. One possible approach would be to assume that upon paying a fixed cost, a firm can make an optimal choice. But we find it more appealing and realistic to assume that firms can choose to spend more or less to make a better or worse decision. In equilibrium in our framework firms will typically prefer an interior degree of precision when they make choices. Therefore their chosen action will not always be the optimal one; instead, firms will sometimes (indeed, usually) make errors.

Consistent with this general description, we adopt the “control cost” approach from game theory (see van Damme, 1991, Chapter 4). A key implication of this approach is that we will model the price decision indirectly: we write the decision problem “as if” firms choose a probability distribution over prices, instead of choosing a price directly and deterministically.⁹ The decision problem incorporates a cost function that increases with precision: concentrating greater probability on a small range of prices increases costs. There are many possible measures of precision; we choose one based on relative entropy, also known as Kullback-Leibler divergence.

Kullback-Leibler divergence is a measure of distance between one probability distribution and another. For two distributions $\pi_1(x)$ and $\pi_2(x)$, for some random variable x with support $\mathcal{X}$, the Kullback-

⁹Luce (1959) and Machina (1985) are early advocates of analyzing decisions in terms of a probability distribution over alternatives; this approach is also adopted by Sims (2003). See Chapter 2 of Anderson *et al.* (1992) for discussion.

Leibler divergence $\mathcal{D}(\pi_1||\pi_2)$ of π_1 relative to π_2 is defined by

$$\mathcal{D}(\pi_1||\pi_2) = \int_{x \in \mathcal{X}} \pi_1(x) \ln \left(\frac{\pi_1(x)}{\pi_2(x)} \right) dx. \quad (9)$$

Following Stahl (1990) and Mattsson and Weibull (2002), we will assume that the decision cost is proportional to the Kullback-Leibler divergence of the chosen distribution, relative to a uniform distribution. This normalizes the cost of a perfectly random decision (a uniform distribution) to zero, and implies that any more precise decision has positive cost.

Consistent with typical US retail data, we assume that prices are set in nominal terms, remaining constant in nominal terms until a new adjustment occurs. However, to simplify notation we will define the support of the random price in real terms. Thus, let us define the log real price of firm i as

$$p_{it} \equiv \ln(P_{it}) - \ln(P_t). \quad (10)$$

We model the firm's choice of a new *nominal* price as the allocation of probabilities $\pi(p)$ to log *real* prices lying in a bounded set $p \in \Gamma^p \subseteq [\underline{p}, \bar{p}]$.¹⁰ We assume the set is sufficiently wide so that the real prices preferred at the minimal and maximal values of productivity, $\underline{a}$ and $\bar{a}$, lie strictly inside $[\underline{p}, \bar{p}]$. It is not necessary at this point to specify whether Γ^p is a discrete or continuous set; in the former case, $\pi(p)$ should be interpreted as a set of discrete probabilities summing to one over the elements of Γ^p ; and in the latter, it should be interpreted as a probability density function on Γ^p .¹¹

Given these preliminaries, we can now define our decision cost function. We assume that costs are denominated in units of time, since we regard managers' time as the main input to decision-making.

Assumption 1. The time cost of choosing a distribution $\pi(p)$, for $p \in \Gamma^p$, is $\kappa_\pi \mathcal{D}(\pi||u)$, where $\kappa_\pi > 0$ is a constant, and u represents a uniform distribution: $u(p) = \bar{u}$ for all $p \in \Gamma^p$.

Here κ_π represents the marginal cost of entropy reduction, in units of labor time. The cost function in Assumption 1 can also be written in terms of the constant density $\bar{u}$, as follows:

$$\kappa_\pi \mathcal{D}(\pi||u) = \kappa_\pi \left(\int_{p \in \Gamma^p} \pi(p) \ln \pi(p) dp - \ln(\bar{u}) \right). \quad (11)$$

This cost function is nonnegative and convex.¹² The upper bound on (11) is associated with any distribution that places all probability on a single price $p \in \Gamma^p$. The lower bound on (11) is zero, associated

¹⁰Alternatively, we could define the set from which new prices are chosen in nominal terms. But then, in the presence of a nonstationary aggregate price level, we would need to allow for a time-varying support, Γ_t^P . Later, to detrend the model and define its steady state, we would have to define the corresponding real support Γ^p . Exposition is simplified, but the model is unchanged, by starting directly from the real support Γ^p .

¹¹If Γ^p is assumed to be a discrete set, then the integral in (9) should be interpreted as a sum.

¹²Cover and Thomas (2006), Theorem 2.7.2.

with a uniform distribution. Thus, Assumption 1 implies that decision costs are maximized by perfect precision and minimized by perfect randomness.

Now consider the pricing decision under this cost function. Suppose the firm, at time t , has already decided to update its price, but has not yet chosen which new nominal price to set. We will write the value of its price decision problem as $\tilde{V}_t(a)$, where a is the firm's current log productivity. The value of its decision depends on the value of producing at each possible log real price $p \in \Gamma^p$, which we write as $V_t(p, a)$. The values $\tilde{V}_t(a)$ and $V_t(p, a)$ are both defined in nominal terms, and they are written with time subscripts to indicate that they depend on the aggregate state of the economy at time t . They are related by the following Bellman equation:

$$\tilde{V}_t(A) = \max_{\pi(p)} \int_{p \in \Gamma^p} \pi(p) V_t(p, a) dp - \kappa_\pi W_t \int_{p \in \Gamma^p} \pi(p) \ln \pi(p) dp + \kappa_\pi W_t \ln(\bar{u}) \quad \text{s.t.} \quad \int_{p \in \Gamma^p} \pi(p) dp = 1 \quad (12)$$

Thus, the firm chooses a price distribution that maximizes its value, net of computational costs (which we convert to nominal terms by multiplying by the wage). The first-order condition for $\pi(p)$ is

$$V_t(p, a) - \kappa_\pi W_t (1 + \ln \pi(p)) - \mu = 0,$$

where μ is the multiplier on the constraint. Some rearrangement yields:

$$\pi(p) = \exp \left(\frac{V_t(p, a)}{\kappa_\pi W_t} - 1 - \frac{\mu}{\kappa_\pi W_t} \right). \quad (13)$$

Since the probabilities sum to one, we have $\exp \left(1 + \frac{\mu}{\kappa_\pi W_t} \right) = \int \exp \left(\frac{V_t(p, a)}{\kappa_\pi W_t} \right) dp$. Therefore the optimal probabilities (13) reduce to the following logit formula:

$$\pi_t(p|a) \equiv \frac{\exp \left(\frac{V_t(p, a)}{\kappa_\pi W_t} \right)}{\int_{p' \in \Gamma^p} \exp \left(\frac{V_t(p', a)}{\kappa_\pi W_t} \right) dp'} \quad (14)$$

The parameter κ_π in the logit function can be interpreted as the degree of noise in the decision process; in the limit as $\kappa_\pi \rightarrow 0$, (14) converges to the policy function under full rationality, so that the optimal price is chosen with probability one.

By calculating the logarithm of $\pi_t(p|a)$ from (13), and plugging it into the objective, we obtain an analytical formula for the value function:

$$\tilde{V}_t(a) = \kappa_\pi W_t \ln \left(\bar{u} \int_{p \in \Gamma^p} \exp \left(\frac{V_t(p, a)}{\kappa_\pi W_t} \right) dp \right). \quad (15)$$

This solution gives the value of adjusting the current price, net of decision costs. It is optimal to adjust the current price if $\tilde{V}_t(a)$ exceeds the value of maintaining the price in place at the beginning of period t . For clarity, we will write the beginning-of-period log real price of firm i as $\tilde{p}_{it}$, to distinguish it from p_{it} ,

the price at which it sells its product at the end of period t . Thus, given any beginning-of- t log real price $\tilde{p}$, optimal adjustment depends on the difference $D_t(\tilde{p}, a)$ between the value of adjusting, and the value of keeping the current price fixed:

$$D_t(\tilde{p}, a) \equiv \tilde{V}_t(a) - V_t(\tilde{p}, a). \quad (16)$$

We next analyze the decision whether or not to adjust at time t .

2.2.2 Choosing the timing of adjustment

By defining costs in terms of the Kullback-Leibler divergence of the price distribution, relative to a uniform distribution, we are penalizing any variation in the probability of one price relative to another. Next, we set up an analogous cost function that *penalizes variation in the probability of adjusting at any given time, relative to another*. More precisely, we penalize variation in the *hazard rate* of the adjustment time. The relevant benchmark for comparison is therefore a Poisson process.

Now, suppose the time period is sufficiently short so that we can approximately ignore multiple adjustments within a single period. If the firm adjusts its price at time t , it obtains the value gain $D_t(\tilde{p}, a)$ defined in (16). Suppose it adjusts its price with probability λ_t . We measure the cost of this adjustment probability in terms of Kullback-Leibler divergence, relative to some arbitrary Poisson process with arrival rate $\bar{\lambda}$:

Assumption 2. The time cost incurred in period t by choosing to adjust in period t with probability $\lambda_t \in [0, 1]$ is $\kappa_\lambda \mathcal{D}((\lambda_t, 1 - \lambda_t) || (\bar{\lambda}, 1 - \bar{\lambda}))$, where $\kappa_\lambda > 0$ and $\bar{\lambda} \in [0, 1]$ are constants.

Here κ_λ is the marginal cost of entropy reduction in the timing decision, which might or might not equal the corresponding parameter κ_π from the pricing decision. Since the decision to adjust or not in any given period t is binary, Assumption 2 states that the decision cost in t depends on the relative entropy of a binary choice with probabilities $(\lambda_t, 1 - \lambda_t)$, relative to another with probabilities $(\bar{\lambda}, 1 - \bar{\lambda})$.¹³

In other words, our cost function benchmarks the state-dependent price adjustment process λ_t against the state-independent Calvo framework. This is a natural way to penalize variability in the distribution of a random time, just as comparing to a uniform distribution penalizes variability in the distribution of possible prices. Since a Calvo model can be defined at any arbitrary adjustment rate $\bar{\lambda}$, this setup

¹³At first blush, it might seem more natural to treat the timing choice as a binary application of the decision model from Sec. 2.2.1, benchmarking the probabilities $(\bar{\lambda}, 1 - \bar{\lambda})$ against the uniform distribution $(0.5, 0.5)$. But this formulation is not well behaved, because it lacks a free parameter corresponding to the modeler's choice of period length. It would imply an adjustment probability of 0.5 per period when indifferent *regardless of period length*, and if taken to its continuous-time limit would imply perfectly rational timing decisions *regardless of* κ_λ .

Another specification that might seem appealing would be to supplement the setup from Sec. 2.2.1 with one more option representing “do nothing now”. But this turns out to be exactly equivalent to the alternative that we have just discussed and rejected (because the grouping axiom of information theory implies that combining a two-step decision problem into a single step leaves overall entropy unchanged).

implies the existence of one free parameter that measures the speed of decision-making, in addition to the parameter κ_{λ}^{-1} that measures the accuracy of decision-making.

Given this cost function, which we rewrite using the definition (9) of Kullback-Leibler divergence, the optimal adjustment probability for a given log real price p and log productivity a satisfies

$$G_t(\tilde{p}, a) = \max_{\lambda_t} D_t(\tilde{p}, a) \lambda_t - \kappa_{\lambda} W_t \left[\lambda_t \ln \left(\frac{\lambda_t}{\bar{\lambda}} \right) + (1 - \lambda_t) \ln \left(\frac{1 - \lambda_t}{1 - \bar{\lambda}} \right) \right]. \quad (17)$$

Here, the value function G represents the expected gains from adjustment, net of the costs of the decision whether or not to adjust. The first order condition from (17) is

$$D_t(\tilde{p}, a) = \kappa_{\lambda} W_t [\ln \lambda_t + 1 - \ln \bar{\lambda} - \ln(1 - \lambda_t) - 1 + \ln(1 - \bar{\lambda})]. \quad (18)$$

Rearranging, we can solve (18) to obtain¹⁴

$$\lambda_t = \lambda \left(\frac{D_t(\tilde{p}, a)}{\kappa_{\lambda} W_t} \right) \equiv \frac{\bar{\lambda}}{\bar{\lambda} + (1 - \bar{\lambda}) \exp \left(\frac{-D_t(\tilde{p}, a)}{\kappa_{\lambda} W_t} \right)} \quad (19)$$

$$= \frac{\bar{\lambda} \exp \left(\frac{\tilde{V}_t(a)}{\kappa_{\lambda} W_t} \right)}{\bar{\lambda} \exp \left(\frac{\tilde{V}_t(a)}{\kappa_{\lambda} W_t} \right) + (1 - \bar{\lambda}) \exp \left(\frac{V_t(\tilde{p}, a)}{\kappa_{\lambda} W_t} \right)} \in [0, 1]. \quad (20)$$

This weighted binary logit hazard was also derived by Woodford (2008) from a model with a Shannon constraint.¹⁵ The free parameter $\bar{\lambda}$ measures the rate of decision making; concretely, the probability of adjustment in one discrete time period is $\bar{\lambda}$ when the firm is indifferent between adjusting and not adjusting, that is, at $(\tilde{p}, a)$ such that $D_t(\tilde{p}, a) = 0$.

Here again, we can explicitly solve for the value function. Rearranging the first-order conditions above, we have

$$\frac{1 - \lambda_t}{1 - \bar{\lambda}} = \frac{\lambda_t}{\bar{\lambda}} \exp \left(\frac{-D_t(\tilde{p}, a)}{\kappa_{\lambda} W_t} \right) = \left(1 - \bar{\lambda} + \bar{\lambda} \exp \left(\frac{D_t(\tilde{p}, a)}{\kappa_{\lambda} W_t} \right) \right)^{-1}. \quad (21)$$

Plugging these formulas into the objective function, the value G of problem (17) is

$$G_t(\tilde{p}, a) = \kappa_{\lambda} W_t \ln \left(1 - \bar{\lambda} + \bar{\lambda} \exp \left(\frac{D_t(\tilde{p}, a)}{\kappa_{\lambda} W_t} \right) \right). \quad (22)$$

¹⁴Note also that (19) has a well-defined continuous-time limit. If $\bar{\lambda}$ is a continuous-time constant hazard against which we benchmark the costs of a time-varying hazard λ_t , then the continuous-time analogue of (19) is $\lambda_t = \bar{\lambda} \exp \left(\frac{D_t(\tilde{p}, a)}{\kappa_{\lambda} W_t} \right)$.

¹⁵Woodford's (2009) paper only states a first-order condition like (18); his (2008) manuscript points out that the first-order condition implies a logit hazard of the form (19).

2.2.3 Value of production

We must still state the Bellman equation that defines the value $V_t(p, a)$ of producing at any log real price p , and any log productivity a . To do so, we write the firm's current profits as

$$U_t(p, a) = \left(P_t \exp(p) - \frac{W_t}{\exp(a)} \right) C_t \exp(-\epsilon p) \quad (23)$$

The Bellman equation for V can be stated most simply in terms of the expected gains G . Using (17), for any beginning-of-period log real price $\tilde{p}$, and any log productivity a , we have

$$V_t(\tilde{p}, a) + G_t(\tilde{p}, a) = \max_{\lambda} \left\{ (1 - \lambda) V_t(\tilde{p}, a) + \lambda \tilde{V}_t(a) - \kappa_{\lambda} W_t \mathcal{D}((\lambda, 1 - \lambda) | (\bar{\lambda}, 1 - \bar{\lambda})) \right\}. \quad (24)$$

Now, note that if a firm leaves its nominal price unchanged from one period to the next, its log *real* price p_{it} declines to $\tilde{p}_{i,t+1} \equiv p_{it} - \ln(P_{t+1}/P_t)$ at the beginning of time $t + 1$. Thus, discounting by the household's stochastic discount factor and using (24), the Bellman equation for V is

$$V_t(p, a) = U_t(p, a) + \beta E_t \left\{ \frac{P_t C_{t+1}^{-\gamma}}{P_{t+1} C_t^{-\gamma}} \left[V_{t+1} \left(p - \ln \frac{P_{t+1}}{P_t}, a' \right) + G_{t+1} \left(p - \ln \frac{P_{t+1}}{P_t}, a' \right) \right] \middle| a \right\}. \quad (25)$$

Here, the expectation is taken over the firm's log productivity next period, a' , conditional on its current log productivity, a , and also over the aggregate state next period, conditional on the aggregate state today.

The terms inside the expectation in the Bellman equation represent the value V of continuing without adjustment, plus the flow of expected gains G due to adjustment. Note that the function G is known analytically in terms of the function $D = \tilde{V} - V$, according to (22). Likewise, $\tilde{V}$ is known analytically in terms of the function V , as seen in (15). Thus, numerical backwards induction is especially simple in this context, because all the maximization steps can be performed analytically.

2.2.4 Extreme special cases

The decision framework defined by (12), (17), and (25) nests two special cases which we will compare with the general case in the simulations that follow. On one hand, we could allow for mistakes in the size of price adjustments, but assume that the timing of price adjustment is perfectly optimal, setting $\kappa_{\pi} > 0$ but $\kappa_{\lambda} = 0$. That is, we could assume that price resetting behavior is governed by (14), while for any $L \equiv \frac{D}{\kappa_{\lambda} W}$ the timing of resets is given by

$$\lambda(L) = \mathbf{1}(L \geq 0), \quad (26)$$

so that adjustment occurs if and only if it increases value. Since the potential for errors in (14) makes price adjustment risky, it means firms will avoid adjusting whenever they are sufficiently close to the optimum, which is why we have called this specification “precautionary price stickiness” in an earlier paper (Costain and Nakov 2011C).

At the opposite extreme, we could set $\kappa_\pi = 0$ but $\kappa_\lambda > 0$. In this case, any adjusting firm always sets the optimal price ($\pi_t(p^*|a) = 1$ if $p^* = \arg\max_p V_t(p, a)$, with probability zero for all other prices), but there may be “mistakes” in the timing of price adjustment, which is governed by the weighted logit (19). Such a framework exhibits near-rational price stickiness in the sense of Akerlof and Yellen (1985) and Costain and Nakov (2011A, B): the probability of price adjustment increases smoothly with the value of adjustment, so a firm is likely to leave its price unchanged when the value of adjustment is small. We will call the functional form (19) for the adjustment probability “Woodford’s logit”, because Woodford (2008) derived it as a consequence of a Shannon constraint on information flow together with a fixed cost of purchasing full information.¹⁶

2.3 Distributional dynamics

As firms respond to productivity shocks, managing their prices according to (14) and (19), the distribution of prices and productivities evolves over time. We now state the equations governing the dynamics of the distribution.

We will use the notation $\tilde{P}_{it}$ to refer to firm i ’s nominal price at the beginning of period t , prior to adjustment; this may of course differ from the price P_{it} at which it produces, because the price may be adjusted before production. Therefore we will distinguish the beginning-of-period distribution of prices and productivity, $\tilde{\Phi}_t(\tilde{P}_{it}, a_{it})$, from distribution of prices and productivity at the time of production, $\Phi_t(P_{it}, a_{it})$. Besides keeping track of nominal prices P_{it} , it will also be helpful to track log real prices p_{it} , defined by (10). In analogy to the nominal distributions, we define $\tilde{\Psi}_t(\tilde{p}_{it}, a_{it})$ as the real beginning-of-period distribution, and $\Psi_t(p_{it}, a_{it})$ as the real distribution at the time of production. Finally, we also use lower-case letters to represent the joint densities associated with these distributions, which we write as $\tilde{\phi}_t(\tilde{P}_{it}, a_{it})$, $\phi_t(P_{it}, a_{it})$, $\tilde{\psi}_t(\tilde{p}_{it}, a_{it})$, and $\psi_t(p_{it}, a_{it})$, respectively.¹⁷

Two stochastic processes drive the dynamics of the distribution. First, there is the Markov process for firm-specific productivity, which we can write in terms of the following *c.d.f.*:

$$S(a'|a) = \text{prob}(a_{i,t+1} \leq a' | a_{i,t} = a), \quad (27)$$

or in terms of the corresponding density function:

$$s(a'|a) = \frac{\partial}{\partial a'} S(a'|a). \quad (28)$$

Thus, suppose that the density of nominal prices and log productivities at the end of period $t - 1$ is

¹⁶Although we share Woodford’s functional form for the adjustment hazard, this special case of our model is not exactly the same as Woodford (2009). Since he considered a rational inattention framework, the gains from adjustment in his model are evaluated in terms of a prior over possible values of the current state, whereas in our model the gains from adjustment are evaluated in terms of the firm’s true state.

¹⁷Our notation in this section assumes that all densities are well-defined on a continuous support, but we do not actually impose this assumption on the model. With slightly more sophisticated notation we could allow explicitly for distributions with mass points, or with discrete support.

$\phi_{t-1}(P, a)$. This density is then affected by productivity shocks; the density at the beginning of t will therefore be

$$\tilde{\phi}_t(P, a') = \int s(a'|a)\phi_{t-1}(P, a)da. \quad (29)$$

But this equation conditions on a given nominal price P . Holding fixed a firm's nominal price, its real log price is changed by inflation, from $p_{i,t-1}$ to $\tilde{p}_{i,t} \equiv p_{i,t-1} - \log(P_t/P_{t-1})$. Therefore the density of real log prices and log productivities at the beginning of t is given by

$$\tilde{\psi}_t\left(p - \log \frac{P_t}{P_{t-1}}, a'\right) = \int s(a'|a)\psi_{t-1}(p, a)da, \quad (30)$$

and hence the cumulative distribution at the beginning of t , in real terms, is

$$\tilde{\Psi}_t(\tilde{p}, a') = \int^{\tilde{p}} \int^{a'} \left(\int s(b|a)\psi_{t-1}\left(q + \log \frac{P_t}{P_{t-1}}, a\right) da \right) db dq. \quad (31)$$

The second stochastic process that determines the dynamics is the process of real price updates, which we have defined in terms of a conditional density of logit form in (14). A firm with real log price $\tilde{p}$ and log productivity a at the beginning of period t adjusts its price with probability $\lambda \left(\frac{D_t(\tilde{p}, a)}{\kappa_\lambda W_t} \right)$, and its new real log price is distributed according to $\pi_t(p|a)$. Therefore, if the density of firms at the beginning of t is $\tilde{\psi}_t(\tilde{p}, a)$, the density at the end of t is given by

$$\psi_t(p, a) = \left(1 - \lambda \left(\frac{D_t(p, a)}{\kappa_\lambda W_t} \right) \right) \tilde{\psi}_t(p, a) + \int \lambda \left(\frac{D_t(\tilde{p}, a)}{\kappa_\lambda W_t} \right) \pi_t(p|a) \tilde{\psi}_t(\tilde{p}, a) d\tilde{p}. \quad (32)$$

The cumulative distribution at the end of t is simply given by integrating up this density:

$$\Psi_t(p, a) = \int^p \int^a \psi_t(q, b) db dq. \quad (33)$$

2.4 Monetary policy and aggregate consistency

The nominal money supply is affected by an AR(1) shock process z ,¹⁸

$$z_t = \phi_z z_{t-1} + \epsilon_t^z, \quad (34)$$

where $0 \leq \phi_z < 1$ and $\epsilon_t^z \sim i.i.d.N(0, \sigma_z^2)$. Here z_t represents the time t rate of money growth:

$$M_t/M_{t-1} \equiv \mu_t = \mu^* \exp(z_t). \quad (35)$$

¹⁸In related work (Costain and Nakov 2011B) we have studied state-dependent pricing when the monetary authority follows a Taylor rule. Our conclusions about the degree of state-dependence, microeconomic stylized facts, and the real effects of monetary policy were not greatly affected by the type of monetary policy rule considered. Therefore we focus here on the simple, transparent case of a money growth rule.

Seigniorage revenues are paid to the household as a lump sum transfer T_t^M , and the government budget is balanced each period, so that $M_t = M_{t-1} + T_t^M$.

Bond market clearing is simply $B_t = 0$. When supply equals demand for each good i , total labor supply and demand satisfy

$$N_t - K_t^\lambda - K_t^\pi = \int_0^1 \frac{C_{it}}{A_{it}} di = C_t \int \int \psi_t(p, a) \exp(-\epsilon p - a) da dp \equiv \Delta_t C_t, \quad (36)$$

where K_t^λ is total time devoted to deciding whether to adjust prices, and K_t^π is total time devoted to choosing which price to set by firms that adjust.¹⁹ Equation (36) also defines a measure of price dispersion, $\Delta_t \equiv P_t^\epsilon \int_0^1 P_{it}^{-\epsilon} A_{it}^{-1} di$, weighted to allow for heterogeneous productivity. As in Yun (2005), an increase in Δ_t decreases the goods produced per unit of labor, effectively acting like a negative aggregate productivity shock.

2.5 General equilibrium

At this point, all equilibrium conditions have been spelled out. We could now define general equilibrium in nominal terms, involving shock processes z_t and M_t governed by (34)-(35); a profit function, value functions, and policy functions U_t , V_t , $\tilde{V}_t$, D_t , G_t , π_t , and λ_t satisfying (23), (25), (15), (16), (22), (14), and (19); prices and policies P_t , C_t , W_t , and R_t satisfying (4)-(7); and densities and distributions $\tilde{\phi}_t$, ϕ_t , $\tilde{\Phi}_t$, and Φ_t consistent with the dynamics discussed in Sec. 2.3. Time subscripts on these equilibrium objects indicate that they are functions of the nominal state of the economy; for example, the value function should take the form $V_t(p, a) = V(p, a, \Omega_t)$, where Ω_t is the nominal state. Since the beginning-of-period distribution $\tilde{\Phi}_t$ is predetermined at t , we could conjecture the existence of a nominal equilibrium that depends on the nominal state $\Omega_t \equiv (z_t, M_t, \tilde{\Phi}_t)$.

However, an equilibrium defined in nominal terms will be nonstationary if the money supply is nonstationary. Therefore, it is more convenient to define equilibrium in real terms, deflating all nominal variables by the nominal price level P_t . To do so, we must identify an appropriate *real* state variable Ξ_t , which would allow us to define a real value function of the form

$$v(p, a, \Xi_t) \equiv V_t(p, a)/P_t. \quad (37)$$

We will detrend the other nominal value functions analogously: $\tilde{v}(a, \Xi_t) \equiv \tilde{V}_t(a)/P_t$, $d(p, a, \Xi_t) \equiv D_t(p, a)/P_t$, and $g(p, a, \Xi_t) \equiv G_t(p, a)/P_t$, and we will also define the real wage and real money balances as $w_t \equiv W_t/P_t$ and $m_t \equiv M_t/P_t$.

Note that a firm's real beginning-of-period price $\tilde{p}_{it} \equiv P_{it}/P_t$ is *not* predetermined at t , since it depends on the aggregate price level P_t , which is endogenous at t . Therefore the beginning-of-period real distribution $\tilde{\Psi}_t$ is likewise not predetermined at t , and hence cannot enter into the definition of the

¹⁹We will not actually need K_t^λ and K_t^π to define general equilibrium, so we omit the formulas here.

time- t real state Ξ_t . Instead, we will now show that the real state can be defined as $\Xi_t \equiv (z_t, \Psi_{t-1})$, since the time $t - 1$ distribution Ψ_{t-1} is predetermined at t .

While the price level will drop out of the real equation system, the inflation rate will still be present. Inflation at t will depend both on the time $t - 1$ state, and on the time t money supply shock, so we expect to find an inflation function of the form

$$i(\Xi_{t-1}, \Xi_t) \equiv \ln \left(\frac{P_t}{P_{t-1}} \right). \quad (38)$$

Substituting inflation in place of the nominal price level P_t will allow us to eliminate all reference to nominal variables in the model.

We can now restate the full equation system in detrended form. The real value of production is given by the following Bellman equation:

$$\begin{aligned} v(p, a, \Xi_t) = & \left(\exp(p) - \frac{w(\Xi_t)}{\exp(a)} \right) \frac{C(\Xi_t)}{\exp(\epsilon p)} \\ & + \beta E \left\{ \frac{C(\Xi_{t+1})^{-\gamma}}{C(\Xi_t)^{-\gamma}} \left[v(p - i(\Xi_t, \Xi_{t+1}), a', \Xi_{t+1}) + g(p - i(\Xi_t, \Xi_{t+1}), a', \Xi_{t+1}) \right] \middle| a, \Xi_t \right\}, \end{aligned} \quad (39)$$

where E represents an expectation over a' and Ξ_{t+1} , conditional on a and Ξ_t . This equation includes the real expected gains function g , which must satisfy the following relations for any beginning-of-period log real price $\tilde{p}$:

$$g(\tilde{p}, a, \Xi) \equiv \kappa_\lambda w(\Xi) \ln \left(1 - \bar{\lambda} + \bar{\lambda} \exp \left(\frac{d(\tilde{p}, a, \Xi)}{\kappa_\lambda w(\Xi)} \right) \right), \quad (40)$$

$$d(\tilde{p}, a, \Xi) \equiv \tilde{v}(a, \Xi) - v(\tilde{p}, a, \Xi), \quad (41)$$

$$\tilde{v}(a, \Xi) \equiv \kappa_\pi w(\Xi) \ln \left(\bar{u} \int_{\underline{p}}^{\bar{p}} \exp \left(\frac{v(p', a, \Xi)}{\kappa_\pi w(\Xi)} \right) dp' \right). \quad (42)$$

The adjustment probability associated with these value functions is

$$\lambda \left(\frac{d(\tilde{p}, a, \Xi)}{\kappa_\lambda w(\Xi)} \right) \equiv \frac{\bar{\lambda}}{\bar{\lambda} + (1 - \bar{\lambda}) \exp \left(\frac{-d(\tilde{p}, a, \Xi)}{\kappa_\lambda w(\Xi)} \right)}, \quad (43)$$

and the density of new real prices is

$$\pi(p|a, \Xi) \equiv \frac{\exp \left(\frac{v(p, a, \Xi)}{\kappa_\pi w(\Xi)} \right)}{\int_{\underline{p}}^{\bar{p}} \exp \left(\frac{v(p', a, \Xi)}{\kappa_\pi w(\Xi)} \right) dp'} = \bar{u} \frac{\exp \left(\frac{v(p, a, \Xi)}{\kappa_\pi w(\Xi)} \right)}{\exp \left(\frac{\tilde{v}(a, \Xi)}{\kappa_\pi w(\Xi)} \right)}. \quad (44)$$

The probability densities then evolve as follows:

$$\tilde{\psi}_t(\tilde{p}, a') = \int s(a'|a) \psi_{t-1}(\tilde{p} + i(\Xi_t, \Xi_{t-1}), a) da \quad (45)$$

$$\psi_t(p, a) = \left(1 - \lambda \left(\frac{d(\tilde{p}, a, \Xi_t)}{\kappa_\lambda w(\Xi_t)}\right)\right) \tilde{\psi}_t(p, a) + \int \lambda \left(\frac{d(\tilde{p}, a, \Xi_t)}{\kappa_\lambda w(\Xi_t)}\right) \pi(p|a, \Xi_t) \tilde{\psi}_t(\tilde{p}, a) d\tilde{p}. \quad (46)$$

These densities can be integrated up to give the following cumulative distributions:

$$\tilde{\Psi}_t(\tilde{p}, a) = \int^{\tilde{p}} \int^a \tilde{\psi}_t(q, b) db dq, \quad (47)$$

$$\Psi_t(p, a) = \int^p \int^a \psi_t(q, b) db dq. \quad (48)$$

Also, aggregate variables must satisfy the representative household's first-order conditions. The condition for labor supply is²⁰

$$w(\Xi_t) C(\Xi_t)^{-\gamma} = \chi. \quad (49)$$

The Euler equation for intertemporal consumption and the money demand equation can be combined (by eliminating the nominal interest rate) to give:

$$1 - \frac{v'(m(\Xi_t))}{u'(C(\Xi_t))} = \beta E \left(i(\Xi_t, \Xi_{t+1}) \frac{u'(C(\Xi_{t+1}))}{u'(C(\Xi_t))} \middle| \Xi_t \right). \quad (50)$$

where $m(\Xi_t) \equiv M_t/P(\Omega_t)$ is the real money supply. Its growth rate must be consistent with the growth rate of the nominal money supply, and inflation, which implies:

$$\frac{\mu \exp(z_t)}{\exp(i(\Xi_{t-1}, \Xi_t))} = \frac{m(\Xi_t)}{m(\Xi_{t-1})}. \quad (51)$$

Finally, the distribution of real prices must be consistent at all times with the definition of the aggregate price level, (4), which implies the following identity:

$$\int \int \exp((1 - \epsilon)p) \psi_t(p, a) da dp = 1. \quad (52)$$

We now have enough equations to define an equilibrium in terms of the real state variable Ξ . Firms' behavior is characterized by the value functions and policy functions v , g , d , $\tilde{v}$, λ , and π , which must satisfy the relationships (39)-(44). The densities and cumulative distributions $\tilde{\psi}$, ψ , $\tilde{\Psi}$, and Ψ , are governed by the dynamics (45)-(48). Finally, the functions w , c , m , and i must satisfy the four equations (49)-(52). A simultaneous solution of these relationships constitutes a real general equilibrium of this economy.

Computing equilibrium requires a high-dimensional calculation, because we must track the evolution of the distribution Ψ of prices and productivities across firms. We compute the model following the algorithm of Reiter (2009), as described in the Online Appendix.

²⁰Our assumption of linear labor disutility χN is helpful, because it allows us to calculate equilibrium without actually solving for N . But the general case of nonlinear labor disutility is also tractable. The equilibrium definition would require four more scalar equations to determine labor $N(\Xi)$, decision costs $K^\pi(\Xi)$ and $K^\lambda(\Xi)$, and price dispersion $\Delta(\Xi)$.

3 Results

We next describe the calibration of the model and report simulation results. We describe the model's steady state implications for microdata on price adjustments, both at a low inflation rate, and as the rate of trend inflation is substantially increased. We also analyze the macroeconomic implications for the effects of monetary policy shocks. The simulations are performed at monthly frequency, and all data and model statistics are monthly unless stated otherwise.

Our focus throughout is on understanding the implications of error-prone price setting. Therefore, to see how each margin of error affects the results, and to see how a pure logit equilibrium compares with a logit equilibrium derived from control costs, we report results for six specifications that turn on or shut down different aspects of the model one by one.²¹ Two specifications allow for errors in the size of price adjustments, but not in their timing, and are labelled “PPS”, for “precautionary price stickiness”. Two specifications allow for errors in the timing of price adjustments, but not in their size; these are labelled “Woodford”, since the adjustment hazard takes the functional form derived in Woodford (2008). The specifications with both types of errors are labelled “nested”. For all these cases, we report the model based on control costs, as well as a model that imposes errors of logit functional form exogenously, without deriving them from control costs. Whenever we refer to the “main model” or the “benchmark model”, we mean the nested control cost specification, in which both types of errors are present, and decision costs are subtracted out of profits, as described in (39)-(44).

3.1 Parameters

The key parameters related to the decision process are the rate and noise parameters $\bar{\lambda}$ and κ . We estimate these two parameters to match two steady-state properties of the price process: the average rate of adjustment, and the histogram of nonzero log price adjustments, in the Dominick's supermarket dataset described in Midrigan (2011).²² More precisely, let h be a vector of length $\#h$ representing the frequencies of nonzero log price adjustments in a histogram with $\#h$ fixed bins.²³ We choose $\bar{\lambda}$ and κ to minimize the following distance criterion:

$$\text{distance} = \sqrt{\#h} \|\lambda_{\text{model}} - \lambda_{\text{data}}\| + \|h_{\text{model}} - h_{\text{data}}\| \quad (53)$$

where $\|\bullet\|$ represents the Euclidean norm, λ_{model} and λ_{data} represent the average frequency of price adjustment in the simulated model and in the Dominick's dataset, and h_{model} and h_{data} are the vectors of bin frequencies for nonzero price adjustments in the model and the data.²⁴ Clearly these features of

²¹Alternatively, we could compare our main model to more familiar price adjustment models. But in Costain and Nakov (2011C) we already compared our “PPS” specification to the Calvo and menu cost models. We refer readers to that paper for comparable tables and graphs documenting those specifications.

²²The weekly adjustment rate in the Dominick's data is aggregated to a monthly rate for comparability with the model.

²³See Figure 3, which compares these histograms in the data and in all specifications of our model.

²⁴Since the Euclidean norm of a vector scales with the square root of the number of elements, we scale the first term by $\sqrt{\#h}$ to place comparable weights on the two components of the distance measure.

Table 1: Adjustment parameters.

	Woodford logit	Woodford control	PPS logit	PPS control	Nested logit	Nested control
$\bar{\lambda}$	0.044	0.045	–	–	0.083	0.22
κ_π	–	–	0.049	0.0044	0.013	0.018
κ_λ	0.0051	0.0080	–	–	0.013	0.018

the data are informative about the two parameters, since $\bar{\lambda}$ will shift the frequency of adjustment and κ will spread the distribution of price adjustments.

The rest of the parameterization is less crucial for our purposes. Hence, for comparability, we take our utility parameterization directly from Golosov and Lucas (2007). Thus, we set the discount factor to $\beta = 1.04^{-1/12}$. Consumption utility is CRRA, $u(C) = \frac{1}{1-\gamma} C^{1-\gamma}$, with $\gamma = 2$. Labor disutility is linear, $x(N) = \chi N$, with $\chi = 6$. The elasticity of substitution in the consumption aggregator is $\epsilon = 7$. Finally, the utility of real money holdings is logarithmic, $v(m) = \nu \log(m)$, with $\nu = 1$. We assume productivity is AR(1) in logs:


$$\log A_{it} = \rho \log A_{it-1} + \varepsilon_t^a, \quad (54)$$

where ε_t^a is a mean-zero, normal, *iid* shock. We take the autocorrelation parameter from Blundell and Bond (2000), who estimate it from a panel of 509 US manufacturing companies over 8 years, 1982-1989. Their preferred estimate is 0.565 on an annual basis, which implies ρ around 0.95 at monthly frequency. The variance of log productivity is $\sigma_a^2 = (1 - \rho^2)^{-1} \sigma_\varepsilon^2$, where σ_ε^2 is the variance of the innovation ε_t^a . We set the standard deviation of log productivity to $\sigma_a = 0.06$, which is the standard deviation of “reference costs” estimated by Eichenbaum *et al.* (2011). The rate of money growth is set to match the roughly 2% annual inflation rate observed in the Dominick’s dataset.

Parameter estimates for the six specifications we compare are reported in Table 1. Note that the PPS specification has only one free parameter: the level of noise κ_π in the pricing decision. The Woodford model has two free parameters: the rate parameter $\bar{\lambda}$, and the level of noise κ_λ in the timing decision. The nested model features the same two free parameters, except that the noise parameter now applies both to the timing and pricing decisions ($\kappa_\pi = \kappa_\lambda \equiv \kappa$).²⁵ The estimated parameters are similar across the logit and control cost specifications, except for the “PPS” case, where the estimated noise is much smaller under control costs than it is under an exogenous logit. We will see (Table 2) that this level of noise in the decision process implies only modest revenue losses. The rate parameter $\bar{\lambda}$ is estimated to be lower than the observed adjustment frequency in the Woodford specification, but is twice as high as the observed adjustment frequency in the main model, marked “nested control”. The combination of a high

²⁵It would also be interesting to allow the two noise parameters of the nested specification to differ, but we leave this for future work, since the simple cross-sectional statistics we are using may not suffice to identify these parameters separately.

Figure 1: Price change distributions and adjustment function: nested control cost model.



Notes:

First panel: price distribution conditional on cost.

Second panel: adjustment probability conditional on price and cost.

Third panel: contour plot of density of firms at time of production.

Fourth panel: contour plot of density of adjusting firms.

underlying adjustment rate, together with a low noise parameter, indicates a high degree of rationality in this estimate of the benchmark model.

3.2 Results: distribution of price adjustments

The steady state behavior of the main model is illustrated in Fig. 1. The first panel of the figure illustrates the distribution of prices chosen conditional on productivity, $\pi(p|a)$; the axes show prices and costs (inverse productivity), expressed in log deviations from their unconditional means. As expected, the mean price chosen increases roughly one-for-one with cost, but the smooth bell-shape of the distribution conditional on a reflects the presence of errors. Similarly, the second panel shows the probability $\lambda(d(p, a)/(\kappa w))$ of price adjustment conditional on beginning-of-period price and productivity. Near

Table 2: Model-Simulated Statistics and Evidence (2% annual inflation)

	Woodford logit	Woodford control	PPS logit	PPS control	Nested logit	Nested control	Data
<i>Adjustment frequency</i>							
Freq. of price changes	10.2	10.2	10.2	10.2	10.2	10.2	10.2
<i>Price change statistics</i>							
Mean absolute price change	4.88	4.68	14.0	6.72	8.11	7.51	9.90
Std of price changes	5.51	5.27	17.0	7.32	10.1	9.30	13.2
Kurtosis of price changes	2.24	2.22	2.58	2.37	3.48	3.40	4.81
Percent of price increases	62.7	63.3	55.2	62.3	58.3	58.8	65.1
Percent of changes $\leq 5\%$	47.9	49.7	16.5	27.9	31.5	33.6	35.4
<i>Variability of prices and costs</i>							
$100 \times \text{Std}(p)/\text{Std}(a)$	95.2	91.0	113	97.7	109	104	115**
<i>Costs of decisions and errors</i>							
Pricing costs*	0	0	0	0.174	0	0.509	
Timing costs*	0	0.167	0	0	0	0.361	
Loss relative to full rationality*	0.258	0.416	0.665	0.365	0.582	1.41	

Note: All statistics refer to regular consumer price changes excluding sales, and are stated in percent.


Quantities with an asterisk are stated as a percentage of monthly average revenues.

Dataset: Dominick's, except for double asterisk, which indicates Eichenbaum *et al.* (2011).

the 45°-line, the adjustment probability reaches a (strictly positive) minimum; moving away from the 45°-line, it increases smoothly towards one. The third panel is a contour plot of the end-of-period distribution of prices and productivities, $\Psi(p, a)$. Dispersion in the horizontal direction represents variation in idiosyncratic productivity over time; dispersion in the vertical direction represents deviation from the conditionally-optimal price, caused either by failures to adjust in response to productivity shocks, or by errors when adjustment occurs. This distribution spreads out horizontally at the beginning of the period when new productivity shocks hit. The resulting distribution of adjusting firms is illustrated by the contour plot in the last panel of the figure. The most frequently observed adjustments occur at firms whose prices deviate from their conditionally-optimal values by 5%-10%; firms with smaller deviations have little incentive to adjust, while firms with larger deviations are rare because adjustment usually occurs before a larger deviation is reached. The asymmetry observed in the density of adjustments reflects the fact that downward price errors (implying high sales at an unprofitably low price) are more costly than upward price errors.

Table 2 also compares other specifications of the model. It reports statistics from the steady state of each specification, and the corresponding statistics from the Dominick's data. All specifications successfully match the 10.2% monthly adjustment frequency observed in the data. But the typical size of the adjustments is too small in the Woodford model and in PPS-control, whereas it is too large in PPS-logit.

Figure 2: Distribution of price adjustments: comparing models.



Notes:

Comparing histogram of price adjustments across models.

Shaded area: histogram of price adjustments in Dominick's data.


Solid and dashed lines: histograms of price adjustments in various versions of the model.

In contrast, the two free parameters of the nested specification help it match both the frequency and the size of price changes simultaneously. Thus the main model is more consistent with the mean absolute change, the standard deviation of the adjustments, the fraction of small adjustments, and even the kurtosis of the data than the other specifications are. The only reported statistic where the nested model performs less well is the fraction of positive adjustments, which is matched very well by the Woodford specification and by PPS-control.

These differences in adjustment behavior can be further understood by graphing the histogram of nonzero log price adjustments. First of all, Figure 2 compares the distribution of price changes in the data (shaded blue bars) with the distributions implied by models with fixed menu costs (FMC) or Calvo (1983) adjustment behavior, and with that implied by our benchmark model. In the data, the distribution of nonzero adjustments exhibits a small peak of negative adjustments, a high peak of positive adjustments, and very fat tails. Under FMC, there are two sharp spikes in the histogram, representing small price increases or decreases occurring near the (S, s) bands. Under the Calvo specification, the distribution of adjustments is narrow and unimodal, with a sharp central peak around zero. In contrast, our benchmark “nested control” model implies a distribution with a widely-spread central peak, and also relatively fat tails. While the distribution in our model is unimodal, it follows the overall shape of the data fairly well both in the center and in the tails.

Next, in Fig. 3, we compare the price adjustment histograms associated with all six versions of our

Figure 3: Distribution of price adjustments: comparing models.



Notes:
Comparing histogram of price adjustments across models.
Shaded area: histogram of price adjustments in Dominick's data.
Solid lines: histograms of price adjustments in logit versions of the model.
Dashed lines: histograms of price adjustments in control cost versions of the model.

error-prone model. The vector of bin frequencies for the 81 bars in these histograms is the object that enters the second term of the distance criterion (53). For the PPS model (first panel of the figure), implications differ strongly between the exogenous logit and control cost specifications. As Table 1 showed, the estimated noise is much lower when control costs are included. *Ceteris paribus*, adjustment is less likely if it requires a decision cost; hence to match the same empirical frequency of adjustment in the logit and control cost specifications, price adjustment must be *less risky* (must have a lower κ) under control costs. Thus, our estimate of the control cost version of the PPS model has extremely low noise, resulting in behavior that is very close to full rationality. The implied distribution of price adjustments resembles the FMC case from Fig. 2, with two sharp spikes representing increases or decreases occurring near a pair of (S,s) bands. On average, this implies much smaller price adjustments than those in the data, with little mass in the tails of the distribution.

Compared with PPS-control, the exogenous logit version of the PPS model requires much more noise to produce the same average adjustment frequency, implying a smoother, wider, more bell-shaped distribution than that observed in the data. In summary, the single free parameter of the PPS framework provides insufficient flexibility to match both the average frequency and the average size of price adjustments. In Costain and Nakov (2011C), for a different dataset with a zero average inflation rate, we reported an estimate of the PPS model that matched both the frequency and size of price adjustments

well. But this finding was essentially coincidental; in the current dataset matching the mean adjustment frequency either implies price changes that are too small (assuming control costs) or too large (assuming an exogenous logit).

Since the Woodford specification has two free parameters, it might seem to have more potential to fit both the frequency and size of adjustments. However, with no errors in the chosen price, this specification implies a much tighter distribution of adjustments than those observed in the Dominick's data. While the data show some adjustments as large as $\pm 50\%$, our estimate of the Woodford specification implies no price changes larger than $\pm 20\%$. While a sufficiently high volatility of underlying costs would spread out the distribution of adjustments observed in this specification, by itself this would be unlikely to reproduce the fat tails of the empirical distribution of price changes. Indeed, while the standard deviation of adjustments in the Woodford specification is slightly larger than in the Calvo model from Fig. 2, the tails of the distribution drop off even more sharply in the Woodford case than they do in the Calvo case.


The nested specifications (in both the control cost and pure logit versions) are better able to match the price change distribution. Like the Woodford model, these models have only two free parameters (since we are constraining the noise in the timing decision to be the same as the noise in the pricing decision). But the pricing errors present in the nested model make it easier to generate a wide, fat-tailed distribution than it is in the Woodford model. At the same time, the parameter $\bar{\lambda}$ helps ensure that the nested model gets the adjustment frequency right. Stated differently, the restriction $\kappa_\pi = 0$ imposed by the Woodford specification strongly constrains its ability to match the data, whereas the restriction $\kappa_\pi = \kappa_\lambda$ that we have maintained when estimating the nested specification does not seem to be strongly rejected by the data. While the main peak is smoother than that observed in the data, the nested model is quite successful both in reproducing the average size of price adjustments and in generating relatively fat tails.

Another way to look at the adjustment process is to consider the losses generated by nonadjustment. Fig. 4 shows the distribution of losses $d(p, a)$ from nonadjustment, expressed as a percentage of average monthly revenue, at the beginning and end of the period, under the benchmark specification. The distribution of losses is strongly skewed out to the right: losses of up to 7% of revenue are visible in the histogram, but most of the mass is concentrated at the left, with a mode at *negative* 7%. The firms at the left end of this distribution are strictly better off not adjusting, because adjustment would require a decision cost, and would also imply a risk of setting the wrong price (this latter phenomenon is what we call “precautionary price stickiness”). Adjustment eliminates some, but not all, of the largest losses, so the beginning-of-period distribution (shaded blue bars) shifts slightly leftward (black line) before production and transactions occur. Adjustment fails to completely eliminate the right tail of the distribution for two reasons: some firms that would be expected to benefit from adjustment fail to adjust, and some firms that do adjust make costly errors.

Losses are also reported in the last few lines of Table 2. The last line of the table shows the average monthly gain from eliminating all decision costs and frictions, as a fraction of average monthly revenues.²⁶ The previous two lines decompose the losses, showing the costs K^π of choosing prices and

²⁶The table shows the gain from adjustment that would accrue to one infinitesimal firm if it could make perfect decisions

Figure 4: Losses from failure to adjust: nested control cost model.



Notes:
Loss from not adjusting, expressed as a percentage of monthly average revenues. Potential losses before adjustments occur (shaded blue) and realized losses after adjustments (black line).


the costs K^λ of deciding the timing of adjustment. The part of the loss reported in the last line that is not attributable to decision costs results from errors. The largest total loss occurs in the nested control costs model, where choosing prices cost firms half of one percent of revenues, choosing the timing of adjustment costs one-third of one percent of revenues, and errors eat up another half of a percent of revenues. In a case study of an industrial firm, Zbaracki *et al.* (2004) find that decision and negotiation costs associated with price adjustment eat up roughly 1.2% of revenues; this is larger than the decision costs, 0.87%, that we find for the nested control model.²⁷ They do not attempt to calculate the revenue loss caused by the suboptimality of the price process at the firm they study.

3.3 Results: some puzzles from microdata

Our model of price adjustment also performs well in reproducing several puzzling observations from microdata. First, note that our main (“nested”) model matches well the observation of Eichenbaum, Jaimovich, and Rebelo (2011) that prices are more volatile than costs (see Table 2). In their data, the ratio of the standard deviation of log reference prices to log reference costs is above unity (1.15), while both the menu cost and the Calvo model predict that this ratio should be less than one.²⁸ This is because

costlessly, holding fixed the behavior of all other firms.
²⁷Since consumers are price takers in our model, all management costs in price adjustment are related to decision-making rather than negotiation.
²⁸The “reference” prices and costs reported by Eichenbaum *et al.* eliminate “sales” and similar phenomena. For their alternative measure of “weekly” prices and costs; the ratio of standard deviations is 1.08.

Figure 5: Price adjustment hazard: comparing models.



Notes:

Adjustment probability as a function of time since last price change.

Shaded area: price adjustment hazard in Nakamura-Steinsson (2008) data.


Solid lines: price adjustment hazards in logit versions of the model.

Dashed lines: price adjustment hazards in control cost versions of the model.

in the menu cost and the Calvo model optimal prices anticipate mean reversion of productivity shocks; prices are set conservatively, taking into account future conditions. Likewise, prices are less volatile than costs in the Woodford version of our model, since it does not allow for pricing errors. However, in the nested and PPS-logit versions, price dispersion is augmented by the possibility of price errors, which results in a higher volatility of prices than of costs, as in the data.

Figures 5-7 show how the six specifications compare with some statistics from microdata that condition on the time since last adjustment. First, one might intuitively expect price adjustment hazards to increase with the time since last adjustment. But empirically, price adjustment hazards are *decreasing* with the time since adjustment, even after controlling for heterogeneity, as in Figure 5, where the blue shaded bars are the adjustment hazards reported by Nakamura and Steinsson (2008). Comparing the various versions of our model we see that under Woodford's logit the adjustment hazard increases over time, since newly set prices are conditionally optimal, and subsequent inflation and productivity shocks gradually drive prices out of line with costs. In contrast, under the PPS-logit specification the adjustment hazard decreases very strongly with the time since last adjustment. This is a consequence of the relatively noisy decisions implied by the estimated parameters for this specification— prices adjust again quickly after a large error occurs. A similar effect exists in PPS-control and the nested models— the possibility of errors in price setting makes the adjustment hazard downward sloping. But the downward slope is much milder than it was for PPS-logit, both because there is less noise in the pricing decision, and because

Figure 6: Mean adjustment and price duration: comparing models.




Notes:
Mean absolute size of price adjustment as function of time since last price change.
Shaded area: Klenow-Kryvtsov dataset.
Solid lines: logit versions of the model.
Dashed lines: control cost versions of the model.

errors in the *timing* of adjustment imply that firms do not always respond immediately when they err in the *size* of their adjustments. Thus, PPS-control and the nested models are the specifications that best fit the mildly negative slope of the empirical adjustment hazard.

The shaded blue bars in Figure 6 illustrate Klenow and Kryvtsov’s (2008) data on the average absolute price change as a function of the time since last adjustment. The size of the adjustment is largely invariant with the age of the current price, with a very slightly positive slope. Under Woodford’s hazard function, the size of the adjustment is instead strongly increasing with the time since last adjustment, since an older price is likely to be farther out of line with current costs. Under the PPS and nested specifications, the size of the adjustment varies less with the age of the price, although it is initially decreasing (due to the correction of recent large errors). It is unclear which of our specifications performs best relative to this phenomenon in the microdata.

Finally, Figure 7 illustrates the observation of Campbell and Eden (2010) that extreme prices tend to be young. The shaded blue bars represent their data, after controlling for sales; the figure shows the fraction of prices that are less than two months old, as a function of the deviation of the price from the mean price in the product group to which that price belongs. In the Campbell and Eden data, the fraction of young prices is around 50% for prices that deviate by more than 20% from the mean, whereas the fraction of young prices is only around 35% for a price equal to the mean. Extreme prices also tend to be young in the PPS and nested models; in these models extreme prices are likely to result from an extreme

Figure 7: Extreme prices tend to be young: comparing models.



Notes:

Fraction of prices set within the last two months, as a function of deviation from average price in product class.

Shaded area: Campbell-Eden dataset.

Solid lines: logit versions of the model.

Dashed lines: control cost versions of the model.


productivity draw compounded by an error, and are therefore unlikely to last long. However, the relation is much too strong under the PPS specification (with prices that deviate by more than 20% from the mean being around 90% young, and only 10% young prices at the mean). The nested specification shows a U-shaped relationship that is more quantitatively consistent with the data. In the Woodford specification the relationship is much flatter, though in that specification too a mild U-shape is observed.

3.4 Results: trend inflation

Next, in Figure 8 and Table 3, we study how each of our specifications performs under large changes in trend inflation. The first row of Fig. 8 shows how the frequency of price adjustment varies as trend inflation rises from 4% to 63% annually, which is the range of inflation rates documented by Gagnon (2009) for a Mexican dataset (price adjustment statistics for 63% annual inflation are also reported in Table 3). Given this increase in the inflation rate, the frequency of price adjustment in the Mexican data increased by a factor of 1.6.²⁹ In the Woodford and PPS specifications, the increase in the adjustment frequency is much too high, ranging from a factor of 2.6 for Woodford-logit to 3.1 for PPS-control. The best performance comes from the nested specifications, although the change is still excessive: the frequency rises by a factor of 2.2 in Nested-logit and by 2.3 for Nested-control.

²⁹In the figure, the adjustment frequency at the low 4% inflation rate is scaled to 100 in all cases, to better compare the changes in frequency across specifications.

Figure 8: Effects of trend inflation (4.3%, 29%, and 63% annually): comparing models.



Notes:
First row: Adjustment frequency as a function of trend inflation rate (normalized to frequency=100 at 4.3% annual inflation).
Second row: Standard deviation of price adjustments as a function of trend inflation rate (normalized to standard deviation = 100 at 4.3% annual inflation).
Third row: Price increases as a percentage of price adjustments, as a function of trend inflation rate.
Line with red stars: Gagnon (2009) Mexican dataset.
Solid lines: logit versions of the model.
Dashed lines: control cost versions of the model.

The second row of Figure 8 shows that the standard deviation of price adjustments changes very little with trend inflation. On this issue, the Woodford specification performs remarkably well, reproducing the data almost perfectly. Most of the specifications with errors in the size of price adjustments instead counterfactually show a small increase in this standard deviation as inflation increases. The exception is PPS-control, where the standard deviation of price adjustments falls by a factor of 0.74 as inflation increases.³⁰ Overall, though, none of these model specifications seem stongly inconsistent with the mildly nonmonotonic change in the standard deviation of price adjustments seen in Gagnon’s data.

³⁰This is because the PPS-control specification acts very much like a fixed menu cost model. As we show in Costain and Nakov (2011A), fixed menu costs imply a strongly bimodal distribution of price adjustments at a low inflation rate, which collapses to a single-peaked distribution as inflation rises, implying a large decrease in the standard deviation of adjustments.

Table 3: Model-Simulated Statistics and Evidence (63% annual inflation)

	Woodford logit	Woodford control	PPS logit	PPS control	Nested logit	Nested control
Freq. of price changes	29.0	29.6	32.1	35.5	24.2	26.0
Mean absolute price change	14.0	13.7	19.7	11.5	17.9	16.6
Std of price changes	4.87	4.86	20.1	5.26	11.7	10.9
Kurtosis of price changes	3.14	3.12	3.32	6.38	4.64	4.43
Percent of price increases	99.9	99.9	78.5	98.9	93.3	93.1
Percent of changes $\leq 5\%$	4.4	4.5	11.2	7.71	7.83	8.71
Pricing costs*	0	0	0	0.557	0	1.06
Timing costs*	0	1.00	0	0	0	0.66
Loss relative to full rationality*	0.752	1.65	1.92	1.00	1.72	3.25

Note: All statistics refer to regular consumer price changes excluding sales, and are stated in percent.

Quantities with an asterisk are stated as a percentage of monthly average revenues.

Dataset: Gagnon (2008) Mexican data.

In the last row, Fig. 8 shows how the fraction of price increases varies with trend inflation. All versions of the model except PPS-logit depart from similar fractions of price increases at 4% inflation. But as inflation rises, the nested models track the proportions of price increases and decreases much more accurately than the other model versions do. The nested models imply that even when annual inflation hits 63%, around 7% of price adjustments are still negative. In contrast, the Woodford and PPS-control specifications tend quickly to a corner solution: the fraction of price decreases is negligible (0.1% or 1.1%, respectively) when inflation reaches 63%.³¹ Finally, since PPS-logit implies much noisier choice than our other specifications, it still displays more than 20% price decreases at a 63% inflation rate.


3.5 Results: money supply shocks

Finally, we turn to the issue of monetary shocks. To begin with, Figure 9 contrasts our benchmark nested model with the Calvo and fixed menu cost models. As Golosov and Lucas (2007) and other recent papers have made clear, different models of price stickiness have remarkably different implications for monetary nonneutrality. After an increase in money supply, the Calvo model implies a small but very persistent rise in inflation; in the FMC model, there is instead a large inflation spike that is even less persistent than the money growth process itself. The almost immediate equilibration of prices in the FMC model means that there is only a small rise in output after the money supply shock, in contrast with the large and persistent output increase implied by the Calvo model.

Figure 10 instead compares the different versions of our model. Like Fig. 9, it shows the impulse responses of inflation and consumption to a 1% money growth rate shock with monthly autocorrelation

³¹Similarly, the fixed menu cost model (not shown here; see Costain and Nakov 2011C) implies that price decreases are almost completely eliminated at annual inflation rates of 29% or 63%.

Figure 9: Impulse responses to money growth shock: comparing models.




Notes:
Impulse responses of inflation and consumption to money growth shock with autocorrelation 0.8 (monthly).
Top row: logit specifications. Bottom row: control cost specifications.
Green lines with squares: Calvo model. Blue lines with circles: FMC model. Red lines: Nested control costs.

of 0.8. Somewhat surprisingly, the responses are quite similar across five of our six specifications, the exception being PPS-control. In the nested and Woodford specifications, the money supply shock leads to a fairly strong real expansion. Consumption rises by 1.8% on impact in response to a one percent money growth shock, and converges back to steady state with a half-life of roughly four months. This is a less persistent response than we reported for the “smoothly state-dependent pricing” specification of Costain and Nakov (2011B), but is still a much stronger effect on consumption than Fig. 9 showed for the fixed menu cost model. If we take the area under the consumption impulse response function as a measure of total nonneutrality, then the figure shows that our nested model has roughly twice the nonneutrality of the FMC case, while in turn the Calvo framework doubles the nonneutrality again.

Returning to Figure 10, we see that the Woodford specifications and nested specifications imply almost identical impulse response functions, both for consumption and inflation. This suggests that the timing errors in Woodford’s logit are the main factor responsible for the nonneutrality of the nested model too. Timing errors obviously help cause monetary nonneutrality since they imply that not all prices adjust immediately in response to a monetary shock. What is more surprising is that PPS-logit also exhibits a very similar nonneutrality. In this case, the real effects can be understood in terms of the large pricing errors implied by our estimate of the model. Given these noisy decisions, firms’ adjustments may be far from optimal responses to the money supply shock. They may therefore need to readjust; note that Fig. 5 shows an adjustment hazard of almost 50% immediately after a price change for this specification. Thus firms may require several attempts before setting a satisfactory price, which slows down adjustment of

Figure 10: Impulse responses to money growth shock: comparing models.



Notes:

Impulse responses of inflation and consumption to money growth shock with autocorrelation 0.8 (monthly).

Top row: logit specifications. Bottom row: control cost specifications.


Blue lines with circles: PPS versions. Green lines with squares: Woodford versions. Red lines: Nested versions.

the aggregate price level and leads to substantial monetary nonneutrality.

With much lower noise, the PPS-control framework behaves very differently. Errors in price setting are small, and timing is perfectly rational, so the small decision cost and risk associated with price adjustment in this specification basically act like a small menu cost. Thus, as we already saw in Figs. 2-3, the PPS-control model behaves very much like the fixed menu cost model. This is true of its impulse responses too: a money supply shock causes a strong initial inflation spike, due to the immediate price changes made by the firms that cross the lower (S, s) band when the money supply increases. Thus, prices adjust quickly and the response of consumption is correspondingly reduced, almost to that of the FMC case.

To see that the initial inflation spike is indeed a “selection effect” in the sense of Golosov and Lucas (2007), we decompose the inflation response in Fig. 11. To construct the decomposition, define the conditionally optimal price level $p_t^{*k} \equiv \arg\max_p v_t(p, a^k)$, and also $x_t^{*jk} \equiv \log(p_t^{*k}/p^j)$, the desired log price adjustment of a firm at time t with productivity a^k and real price p^j . The actual log price adjustment

Figure 11: Decomposition of inflation impulse responses: comparing models.



Notes:

Decomposition of inflation impulse response to money growth shock with autocorrelation 0.8 (monthly).

Top row: logit specifications. Bottom row: control cost specifications.

Blue lines with circles: PPS versions. Green lines with squares: Woodford versions. Red lines: Nested versions.

of such a firm (call it i) can thus be decomposed as $x_{it} = x_t^{*jk} + \epsilon_{it}$, where ϵ_{it} is an error, in logs. We can then write the average desired adjustment as $\bar{x}_t^* = \sum_{j,k} x_t^{*jk} \tilde{\Psi}_t^{jk}$, and write the fraction of firms adjusting as $\bar{\lambda}_t = \sum_{j,k} \lambda_t^{jk} \tilde{\Psi}_t^{jk}$, and write the average log error as $\bar{\epsilon}_t = \sum_{j,k,l} \pi_t^{lk} \log(p^l/p_t^{*k}) \lambda_t^{jk} \tilde{\Psi}_t^{jk}$. Then inflation can be written as

$$\Pi_t = \sum_{j,k} x_t^{*jk} \lambda_t^{jk} \tilde{\Psi}_t^{jk} + \bar{\epsilon}_t. \quad (55)$$

To a first-order approximation, we can decompose the deviation in inflation at time t as

$$\Delta \Pi_t = \bar{\lambda} \Delta \bar{x}_t^* + \bar{x}^* \Delta \bar{\lambda}_t + \Delta \sum_{j,k} x_t^{jk} (\lambda_t^{jk} - \bar{\lambda}_t) \tilde{\Psi}_t^{jk} + \Delta \bar{\epsilon}_t, \quad (56)$$

where terms without time subscripts represent steady states, and Δ represents a change relative to steady state.³²

The “intensive margin”, $\mathcal{I}_t \equiv \bar{\lambda} \Delta \bar{x}_t^*$, is the part of inflation due to changes in the average desired adjustment, holding fixed the fraction of firms adjusting. The “extensive margin”, $\mathcal{E}_t \equiv \bar{x}^* \Delta \bar{\lambda}_t$, is the part due to changes in the fraction adjusting, assuming the average desired change among those who adjust equals the steady-state average in the whole population. The “selection effect”, $\mathcal{S}_t \equiv \Delta \sum_{j,k} x_t^{jk} (\lambda_t^{jk} - \bar{\lambda}_t) \tilde{\Psi}_t^{jk}$, is the inflation caused by redistributing adjustment opportunities from firms desiring small (or

³²See Costain and Nakov (2011B) for further discussion of this decomposition.

Table 4: Variance decomposition and Phillips curves

<i>Correlated money growth shock</i> ($\phi_z = 0.8$)	Woodford logit	Woodford control	PPS logit	PPS control	Nested logit	Nested control	Data*
Freq. of price changes (%)	10.2	10.2	10.2	10.2	10.2	10.2	10.2
Std of money shock (%)	0.16	0.15	0.16	0.12	0.17	0.17	
Std of qtrly inflation (%)	0.25	0.25	0.25	0.25	0.25	0.25	0.25
% explained by μ shock alone	100	100	100	100	100	100	
Std of qtrly output growth (%)	0.41	0.37	0.34	0.20	0.45	0.43	0.51
% explained by μ shock alone	80	73	67	38	89	84	
Slope coeff. of Phillips curve*	0.32	0.29	0.31	0.15	0.38	0.35	
R ² of regression	0.96	0.94	0.999	0.85	0.99	0.98	

*The “slope coefficients” are 2SLS estimates of the effect of inflation on consumption

First stage: $\pi_t^q = \alpha_1 + \alpha_2 \mu_t^q + \epsilon_t$; second stage: $c_t^q = \beta_1 + \beta_2 \hat{\pi}_t^q + \varepsilon_t$, where the instrument μ_t^q is the exogenous growth rate of the money supply and the superscript q indicates quarterly averages.

Dataset: Dominick’s.

negative) price adjustments to firms desiring large (positive) adjustments, while fixing the total number adjusting. The last term, $\Delta \bar{\epsilon}_t$, is the change in the average log error. Figure 11 reports the inflation decomposition for our six specifications. We see that, indeed, the spike of inflation on impact in PPS-control is a selection effect. Interestingly, the majority of the inflation response is also attributed to the selection component in the nested specifications, but this selection effect is more spread out over time. The intensive margin is smaller, and the extensive margin and error margins are negligible, in all the specifications considered.³³

In Table 4, we provide an additional assessment of the degree of nonneutrality in our model by running two calculations from Golosov and Lucas (2007). Assuming for concreteness that money shocks are the only cause of macroeconomic fluctuations, we calibrate the standard deviation of the money shock for each specification to perfectly match the standard deviation of quarterly inflation (one quarter of one percent) in US data. We then check what fraction of the time variation in US output growth can be explained by those shocks. In the Woodford and nested specifications, these money shocks would explain around 80% or 90% of the observed variation in US output growth. In PPS-logit, they would explain 67% of output growth variation, while in PPS-control they would explain only 38%, consistent with the strong inflation spike and small output response observed in Fig. 10 for this specification. In the last line of the table, we also report “Phillips curve” coefficients, that is, estimates from an instrumental variables regression of the effect of inflation on output, instrumenting inflation by the exogenous money

³³Because of the asymmetry of the adjustment process (last panel of Fig. 1), the steady state average log error $\bar{\epsilon}$ is nontrivial. But time variation in the average pricing error is negligible.

supply process. The coefficient is more than twice as large for the nested, Woodford, and PPS-logit cases as it is for PPS-control.

4 Conclusions

This paper has modeled nominal price rigidity as a near-rational phenomenon. Price adjustment is costly, but the interpretation of the costs is not the usual one: they represent the cost of decision-making by management.

We operationalize this idea by adopting a common assumption from game theory: a “control cost” function that depends on the precision of the decision. Following Mattsson and Weibull (2002), we assume that precision is measured by relative entropy, and then show that decisions are random variables with logit form. This well-known game-theoretic result is directly applicable to the question of *which* price to set, once the firm has decided to make an adjustment. We show how to extend the logit result to the decision of *when* to adjust the price. Just as the cost of the price choice is assumed proportional to relative entropy compared with a uniform price distribution, the cost of the timing choice is assumed proportional to the relative entropy of the adjustment hazard, compared with a uniform adjustment hazard. The resulting model of near rational choice has just two parameters: a noise parameter measuring the accuracy of decisions, and a rate parameter measuring the speed of decisions.

We shut down the errors on each choice margin— the timing margin and the pricing margin— to see the role played by each type of error. The model with pricing errors, but perfect adjustment timing, implies that prices are sticky when they are near the optimum, because of the risk of choosing a worse price; therefore we call this specification “precautionary price stickiness” (PPS). This special case has only one free parameter— the degree of noise in the pricing decision. Our simulations show that noise in the pricing decision helps match a variety of features of the price adjustment microdata, but with only one free parameter the model cannot in general match both the typical size of adjustment and its typical frequency. We refer to the model with errors in timing but perfect pricing decisions as “Woodford’s logit”, because the functional form for the adjustment hazard is the same weighted logit derived by Woodford (2008) for a rational inattention model. Both the Woodford specification, and our general nested specification, have two free parameters: the decision accuracy and the decision rate. But with a few exceptions the nested specification fits the data far better than the Woodford specification does.

With just two parameters, the nested specification fits well both the timing and size of price adjustments. As microdata show, both large and small price adjustments coexist in the distribution. Both the adjustment hazard, and the average size of the adjustment, are largely independent of the time since last adjustment. Extreme prices are more likely to have been recently set. Prices are more volatile than costs. The nested model is well-behaved as inflation rises from 4% to 63% annually, and it performs better than the PPS or Woodford specifications in describing how the distribution of price adjustments changes with inflation (in light of the Mexican data of Gagnon, 2009). Both the nested model and the Woodford model imply a realistic degree of monetary nonneutrality in response to money growth shocks (though substan-

tially less than a Calvo model with the same average rate). While this paper has focused on showing what errors on each adjustment margin contribute to the overall performance of the main nested specification, we have also briefly compared the nested framework to the Calvo and fixed menu cost models to show that it outperforms those specifications too.

Standard practice in microeconometrics includes error terms in all behavioral equations. Most recent work on state-dependent pricing has instead modeled the full distribution of price adjustments as if firms' behavior were entirely error-free. Here, instead, we allow for mistakes, and interpret them structurally as the result of costly managerial decision-making. The payoff to this approach is that once we allow for errors, we can eliminate all other forms of frictions (including "menu costs" and exogenous probabilistic barriers to adjustment) but nonetheless match micro and macrodata at least as well as competing frameworks (most of which are less sparsely parameterized). While the present paper has focused on price adjustment, our framework also seems appropriate for other contexts in which a decision maker intermittently flips a switch or updates a number or a vector. Interesting potential applications include wage bargaining, hiring and firing decisions, inventory control, portfolio adjustment problems, lumpy investment problems, and adjustment of macroeconomic policy instruments.

References

- [1] Akerlof, George, and Janet Yellen (1985), "A near-rational model of the business cycle with wage and price inertia." *Quarterly Journal of Economics* 100 (Supplement): 823-38.
- [2] Álvarez, Fernando; Francesco Lippi, and Luigi Paciello (2011), "Optimal price setting with observation and menu costs." *Quarterly Journal of Economics* 126 (4): 1909-60.
- [3] Álvarez, Fernando; Martín González-Rozada; Andy Neumeyer; and Martín Beraja (2011), "From hyperinflation to stable prices: Argentina's evidence on menu cost models." Manuscript, Univ. of Chicago.
- [4] Anderson, Simon; André de Palma; and Jacques-François Thisse (1992), *Discrete Choice Theory of Product Differentiation*. MIT Press.
- [5] Anderson, Simon; Jacob Goeree, and Charles Holt (2002), "The logit equilibrium: a perspective on intuitive behavioral anomalies." *Southern Economic Journal* 69 (1): 21-47.
- [6] Barro, Robert (1972), "A theory of monopolistic price adjustment." *Review of Economic Studies* 39 (1): 17-26.
- [7] Blundell, Richard, and Stephen Bond (2000), "GMM Estimation with persistent panel data: an application to production functions." *Econometric Reviews* 19 (3): 321-40.

- [8] Bono, James, and David Wolpert (2009), “Statistical prediction of the outcome of a noncooperative game.” American University Economics Working Paper 2009-20.
- [9] Calvo, Guillermo (1983), “Staggered prices in a utility-maximizing framework.” *Journal of Monetary Economics* 12: 383-98.
- [10] Campbell, Jeffrey, and Benjamin Eden (2010), “Rigid prices: evidence from U.S. scanner data.” FRB Chicago WP2005-08 (revised July 2010).
- [11] Caplin, Andrew, and Daniel Spulber (1987), “Menu costs and the neutrality of money.” *Quarterly Journal of Economics* 102 (4): 703-25.
- [12] Christiano, Lawrence; Martin Eichenbaum, and Charles Evans (1999), “Monetary policy shocks: what have we learned and to what end?” In J. Taylor and M. Woodford, eds. (1999), *Handbook of Macroeconomics*, Volume 1A. Elsevier.
- [13] Costain, James, and Anton Nakov (2011A), “Price adjustments in a general model of state-dependent pricing.” *Journal of Money, Credit and Banking* 43 (2-3): 385-406.
- [14] Costain, James, and Anton Nakov (2011B), “Distributional dynamics under smoothly state-dependent pricing.” *Journal of Monetary Economics* 58 (6-8): 646-65.
- [15] Costain, James, and Anton Nakov (2011C), “Precautionary price stickiness.” European Central Bank Working Paper #1375.
- [16] Cover, Thomas, and Joy Thomas (2006), *Elements of Information Theory*, 2nd ed. Wiley Interscience.
- [17] Demery, David (2012), “State-dependent pricing and the non-neutrality of money.” *Journal of Macroeconomics* 34: 933-44.
- [18] Dotsey, Michael; Robert King; and Alexander Wolman (2013), “Inflation and real activity with firm-level productivity shocks.” FRB Philadelphia Working Paper 13-35.
- [19] Eichenbaum, Martin; Nir Jaimovich, and Sergio Rebelo (2011), “Reference prices, costs and nominal rigidities.” *American Economic Review* 101 (1): 234-62.
- [20] Gagnon, Etienne (2009), “Price setting under low and high inflation: evidence from Mexico.” *Quarterly Journal of Economics* 124: 1221-63.
- [21] Golosov, Mikhail, and Robert E. Lucas, Jr. (2007), “Menu costs and Phillips curves.” *Journal of Political Economy* 115 (2): 171-99.
- [22] Haile, Philip, Ali Hortaçsu, and Grigory Kosenok (2008), “On the empirical content of quantal response equilibrium.” *American Economic Review* 98 (1): 180-200.

- [23] Hansen, Lars, and Thomas Sargent (2008), *Robustness*. Princeton University Press.
- [24] Kehoe, Patrick, and Virgiliu Midrigan (2010), “Prices are sticky after all.” NBER Working Paper #16364.
- [25] Klenow, Peter, and Oleksiy Kryvtsov (2008), “State-dependent or time-dependent pricing: does it matter for recent US inflation?”, *Quarterly Journal of Economics* 123: 863-904.
- [26] Klenow, Peter, and Benjamin Malin (2010), “Microeconomic evidence on price setting.” In *Handbook of Monetary Economics*, Benjamin Friedman and Michael Woodford, eds., Vol. 3, no. 3.
- [27] Klein, Paul (2000), “Using the generalized Schur form to solve a multivariate linear rational expectations model.” *Journal of Economic Dynamics and Control* 24 (10): 1405-23.
- [28] Machina, Mark (1985), “Stochastic choice functions generated from deterministic preferences over lotteries.” *Economic Journal* 95: 575-94.
- [29] Mankiw, N. Gregory (1985), “Small menu costs and large business cycles: a macroeconomic model of monopoly.” *Quarterly Journal of Economics* 100 (2): 529-37.
- [30] Marsili, Matteo (1999), “On the multinomial logit model.” *Physica A* 269: 9-15.
- [31] Matejka, Filip, and Alisdair McKay (2011), “Rational inattention to discrete choices: a new foundation for the multinomial logit model.” CERGE-EI Working Paper Series 442.
- [32] Matejka, Filip (2011), “Rationally inattentive seller: sales and discrete pricing.” Manuscript, CERGE-EI.
- [33] Mattsson, Lars-Göran, and Jörgen Weibull (2002), “Probabilistic choice and procedurally bounded rationality.” *Games and Economic Behavior* 41: 61-78.
- [34] McKelvey, Richard, and Thomas Palfrey (1995), “Quantal response equilibrium for normal form games.” *Games and Economic Behavior* 10: 6-38.
- [35] McKelvey, Richard, and Thomas Palfrey (1998), “Quantal response equilibrium for extensive form games.” *Experimental Economics* 1: 9-41.
- [36] Midrigan, Virgiliu (2011), “Menu Costs, Multi-Product Firms and Aggregate Fluctuations.” *Econometrica* 79 (4): 1139-80.
- [37] Nakamura, Emi, and Jón Steinsson (2008), “Five facts about prices: a reevaluation of menu cost models.” *Quarterly Journal of Economics* 123 (4): 1415-64.
- [38] Reiter, Michael (2009), “Solving heterogeneous-agent models by projection and perturbation.” *Journal of Economic Dynamics and Control* 33 (3): 649-65.

- [39] Shannon, Claude (1948), "A mathematical theory of communication." *The Bell System Technical Journal*, 27: 379-423 and 623-56.
- [40] Sims, Christopher (2003), "Implications of rational inattention." *Journal of Monetary Economics* 50: 665-90.
- [41] Stahl, Dale (1990), "Entropy control costs and entropic equilibrium. " *International Journal of Game Theory* 19: 129-38.
- [42] Van Damme, Eric (1991), *Stability and Perfection of Nash Equilibrium*, 2nd edition. Springer Verlag.
- [43] Woodford, Michael (2008), "Inattention as a source of randomized discrete adjustment." Manuscript, Columbia University.
- [44] Woodford, Michael (2009), "Information-constrained state-dependent pricing." *Journal of Monetary Economics* 56: S100-S124.
- [45] Yun, Tack (2005), "Optimal Monetary Policy with Relative Price Distortions." *American Economic Review* 95 (1): 89-109.
- [46] Zbaracki, Mark; Mark Ritson; Daniel Levy; Shantanu Dutta; and Mark Bergen (2004), "Managerial and customer costs of price adjustment: direct evidence from industrial markets." *Review of Economics and Statistics* 86 (2): 514-33.

Appendix to “Logit price dynamics”: For online publication

J. Costain and A. Nakov, Sept. 2013

Outline of algorithm

Computing this model is challenging due to heterogeneity: at any time t , firms will face different productivity shocks A_{it} and will be stuck at different prices P_{it} . The Calvo model is popular because, up to a first-order approximation, only the average price matters for equilibrium. But this property does not hold in most models; here we must treat all equilibrium quantities as functions of the time-varying distribution of prices and productivity across firms.

We address this problem by implementing Reiter’s (2009) solution method for dynamic general equilibrium models with heterogeneous agents and aggregate shocks. As a first step, the algorithm calculates the steady-state general equilibrium in the absence of aggregate shocks. Idiosyncratic shocks are still active, but are assumed to have converged to their ergodic distribution, so the real aggregate state of the economy is a constant, Ξ . The algorithm solves for a discretized approximation to this steady state; here we restrict real log prices p_{it} and log productivities a_{it} to a fixed grid $\Gamma \equiv \Gamma^p \times \Gamma^a$, where $\Gamma^p \equiv \{p^1, p^2, \dots, p^{\#^p}\}$ and $\Gamma^a \equiv \{a^1, a^2, \dots, a^{\#^a}\}$ are both logarithmically spaced. We can then view the steady state value function as a matrix $\mathbf{V}$ of size $\#^p \times \#^a$, comprising the values $v^{jk} \equiv v(p^j, a^k, \Xi)$ associated with prices and productivities $(p^j, a^k) \in \Gamma$.³⁴ Likewise, the price distribution can be viewed as a $\#^p \times \#^a$ matrix Ψ in which the row j , column k element Ψ^{jk} represents the fraction of firms in state (p^j, a^k) at the end of any given period. Under this discrete representation, we can calculate steady state general equilibrium by guessing the wage w , then solving the firm’s problem by backwards induction on the grid Γ , then updating the conjectured wage, and iterating to convergence.

In a second step, Reiter’s method constructs a linear approximation to the dynamics of the discretized model, by perturbing it around the steady state general equilibrium on a point-by-point basis. That is, the value function is represented by a $\#^p \times \#^a$ matrix $\mathbf{V}_t$ with row j , column k element $v_t^{j,k} \equiv v(p^j, a^k, \Xi_t)$, thus summarizing the time t values at all grid points $(p^j, a^k) \in \Gamma$. Then, instead of viewing the Bellman equation as a functional equation that defines $v(p, a, \Xi)$ for all possible idiosyncratic and aggregate states p , a , and Ξ , we think of it as an expectational relation between the matrices $\mathbf{V}_t$ and $\mathbf{V}_{t+1}$. This amounts to a (large!) system of $\#^p \#^a$ first-order expectational difference equations that determine the dynamics of the $\#^p \#^a$ variables $v_t^{j,k}$. We linearize these equations numerically (together with the $\#^p \#^a$ equations that describe the evolution of the mass of firms at each grid point, and a few other scalar equations). We then solve for the saddle-path stable solution of the linearized model using the QZ decomposition, following Klein (2000).

This method combines linearity and nonlinearity in a way appropriate for models of price setting, where idiosyncratic shocks tend to be more relevant for firms’ decisions than aggregate shocks are (*e.g.* Klenow and Kryvtsov, 2008; Golosov and Lucas, 2007; Midrigan, 2011). When we linearize the model’s

³⁴In this appendix, bold face indicates matrices, and superscripts represent indices of matrices or grids.

aggregate dynamics, we recognize that changes in the aggregate shock z_t or in the distribution Ψ_t are unlikely to have a strongly nonlinear effect on the value function. Note that this smoothness does not require any “approximate aggregation” property, in contrast with the Krusell and Smith (1998) method; nor do we need to impose any particular functional form on the distribution Ψ . However, to allow for the importance of firm-specific shocks, Reiter’s method treats variation along idiosyncratic dimensions in a fully nonlinear way: the value at each grid point is determined by a distinct equation, which could in principle be entirely different from the equations governing the value at neighboring points.

The discretized model

In the discretized model, the value function $\mathbf{V}_t$ is a matrix of size $\#^p \times \#^a$ with elements $v_t^{jk} \equiv v(p^j, a^k, \Xi_t)$ where $(p^j, a^k) \in \Gamma$. The expected value of setting a new price is a column vector $\tilde{\mathbf{v}}_t$ of length $\#^a$, with k th element

$$\tilde{v}_t^k \equiv \kappa w(\Xi_t) \ln \left(\frac{1}{\#^p} \sum_{j=1}^{\#^p} \exp \left(\frac{v_t^{jk}}{\kappa w(\Xi_t)} \right) \right). \quad (57)$$

Other relevant $\#^p \times \#^a$ matrices include the adjustment values $\mathbf{D}_t$, the probabilities $\mathbf{\Lambda}_t$, and the expected gains $\mathbf{G}_t$, with (j, k) elements given by³⁵

$$d_t^{jk} \equiv \tilde{v}_t^k - v_t^{jk}, \quad (58)$$

$$\lambda_t^{jk} \equiv \lambda \left(d_t^{jk} / (\kappa w_t) \right) \quad (59)$$

$$g_t^{jk} \equiv \kappa w_t \left(1 - \bar{\lambda} + \bar{\lambda} \exp \left(\frac{d_t^{jk}}{\kappa w_t} \right) \right) \quad (60)$$

Finally, we also define a matrix of logit probabilities $\mathbf{\Pi}_t$, which has its (j, k) element given by

$$\pi_t^{jk} = \pi_t(p^j | a^k) \equiv \frac{\exp \left(v_t^{jk} / (\kappa w_t) \right)}{\sum_{n=1}^{\#^p} \exp \left(v_t^{nk} / (\kappa w_t) \right)}$$

which is the probability of choosing real log price p^j conditional on log productivity a^k if the firm decides to adjust its price at time t .

In this discrete representation, the productivity process (54) can be written in terms of a $\#^a \times \#^a$

³⁵ Actually, (59) is a simplified description of λ_t^{jk} . While (59) implies that λ_t^{jk} represents the function $\lambda(L)$ evaluated at the log price grid point p^j and log productivity grid point a^k , in our computations λ_t^{jk} actually represents the *average* of $\lambda(L)$ over all log prices in the interval $\left(\frac{p^{j-1} + p^j}{2}, \frac{p^j + p^{j+1}}{2} \right)$, given log productivity a^k . Calculating this average requires interpolating the function $d_t(p, a^k)$ between price grid points. Defining λ_t^{jk} this way ensures differentiability with respect to changes in the aggregate state Ω_t .

matrix $\mathbf{S}$, where the (m, k) element represents the following transition probability:

$$S^{mk} = \text{prob}(a_{it} = a^m | a_{i,t-1} = a^k).$$

It is helpful to introduce analogous Markovian notation for the price process. Let $\mathbf{R}_t$ be a $\#^p \times \#^p$ Markov matrix in which the row m , column l element represents the probability that firm i 's beginning-of-period log real price $\tilde{p}_{i,t}$ equals $p^m \in \Gamma^p$ if its log real price at the end of the previous period was $p^l \in \Gamma^p$:

$$R_t^{ml} \equiv \text{prob}(\tilde{p}_{it} = p^m | p_{i,t-1} = p^l).$$

Generically, the deflated log price $p_{i,t-1} - i(\Xi_{t-1}, \Xi_t)$ will fall between two grid points; then the matrix $\mathbf{R}_t$ must round up or down stochastically. Also, if $p_{i,t-1} - i(\Xi_{t-1}, \Xi_t)$ lies below the smallest or above the largest element of the grid, then $\mathbf{R}_t$ must round up or down to keep prices on the grid.³⁶ Therefore we construct $\mathbf{R}_t$ according to

$$R_t^{ml} = \text{prob}(\tilde{p}_{it} = p^m | p_{i,t-1} = p^l, i_t) = \begin{cases} 1 & \text{if } p^l - i_t \leq p^1 = p^m \\ \frac{p^l - i_t - p^{m-1}}{p^m - p^{m-1}} & \text{if } p^1 < p^m = \min\{p \in \Gamma^p : p \geq p^l - i_t\} \\ \frac{p^{m+1} - p^l + i_t}{p^{m+1} - p^m} & \text{if } p^1 \leq p^m = \max\{p \in \Gamma^p : p < p^l - i_t\} \\ 1 & \text{if } p^l - i_t > p^{\#p} = p^m \\ 0 & \text{otherwise} \end{cases} \quad (61)$$

Given this notation, we can now write the distributional dynamics in a more compact form. Equations (45) and (47) become

$$\tilde{\Psi}_t = \mathbf{R}_t * \Psi_{t-1} * \mathbf{S}', \quad (62)$$

where $*$ represents ordinary matrix multiplication. Note that exogenous shocks are represented from left to right in the matrix Ψ_t , so that their transitions can be treated by right multiplication, while policies are represented vertically, so that transitions related to policies can be treated by left multiplication. Next, to calculate the effects of price adjustment on the distribution, let $\mathbf{E}_{pp}$ and $\mathbf{E}_{pa}$ be matrices of ones of size $\#^p \times \#^p$ and $\#^p \times \#^a$, respectively. Equations (46) and (48) can then be rewritten as

$$\Psi_t = (\mathbf{E}_{pa} - \mathbf{\Lambda}) .* \tilde{\Psi}_t + \mathbf{\Pi}_t .* (\mathbf{E}_{pp} * (\mathbf{\Lambda} .* \tilde{\Psi}_t)), \quad (63)$$

where (as in MATLAB) the operator $.*$ represents element-by-element multiplication.

The same transition matrices $\mathbf{R}$ and $\mathbf{S}$ show up when we write the Bellman equation in matrix form.

³⁶In other words, we assume that any nominal price that would have a real log value less than p^1 after inflation is automatically adjusted upwards to the real log value p^1 (and when computing examples with deflation we must adjust down any real log price exceeding $p^{\#p}$). This assumption is made for numerical purposes only, and has a negligible impact on the equilibrium as long as we choose a sufficiently wide grid Γ^p .

Let $\mathbf{U}_t$ be the $\#^p \times \#^a$ matrix of current payoffs, with elements

$$u_t^{jk} \equiv \left(\exp(p^j) - \frac{w(\Xi_t)}{\exp(a^k)} \right) \frac{C(\Xi_t)}{\exp(\epsilon p^j)} \quad (64)$$

for $(p^j, a^k) \in \Gamma$. Then the Bellman equation can be written as

$$\mathbf{V}_t = \mathbf{U}_t + \beta E_t \left\{ \frac{C_{t+1}^{-\gamma}}{C_t^{-\gamma}} [\mathbf{R}'_{t+1} * (\mathbf{V}_{t+1} + \mathbf{G}_{t+1}) * \mathbf{S}] \right\}. \quad (65)$$

The expectation E_t in (65) refers only to the effects of the time $t + 1$ aggregate shock z_{t+1} , because the dynamics of the idiosyncratic state $(p^j, a^k) \in \Gamma$ are completely described by the matrices $\mathbf{R}'_{t+1}$ and $\mathbf{S}$. Note that since (65) iterates backwards in time, its transitions are represented by $\mathbf{R}'$ and $\mathbf{S}$, whereas (63) iterates forward in time and therefore involves $\mathbf{R}$ and $\mathbf{S}'$.

Next, we discuss how we apply the two steps of Reiter's (2009) method to this discrete model.

Step 1: steady state

In the aggregate steady state, the shocks are zero, and the distribution takes some unchanging value Ψ , so the state of the economy is constant: $\Xi_t \equiv (z_t, \Psi_{t-1}) = (0, \Psi) \equiv \Xi$. We indicate the steady state of all equilibrium objects by dropping the time subscripts and the function argument Ξ , so the steady state value function $\mathbf{V}$ has elements $v^{jk} \equiv v(p^j, a^k, \Xi)$.

Long run monetary neutrality in steady state implies that the rate of nominal money growth equals the rate of inflation:

$$\mu = \exp(i).$$

Thus, the steady-state transition matrix $\mathbf{R}$ is known, since it depends only on steady state inflation i . Moreover, the Euler equation reduces to

$$\exp(i) = \beta R.$$

We can then calculate general equilibrium as a one-dimensional root-finding problem: guessing the wage w , we have enough information to solve the Bellman equation and the distributional dynamics. Knowing the steady state aggregate distribution, we can construct the real price level, which must be one. Thus finding a value of w at which the real price level is one amounts to finding a steady state general equilibrium.

More precisely, for any w , we calculate $C = (w/\chi)^{1/\gamma}$. We can then construct the matrix $\mathbf{U}$, with elements

$$u^{jk} \equiv \left(\exp(p^j) - \frac{w}{\exp(a^k)} \right) \frac{C}{\epsilon p^j}. \quad (66)$$

We then find the fixed point of the value $\mathbf{V}$ (simultaneously with $\tilde{\mathbf{v}}$, $\mathbf{D}$, $\boldsymbol{\lambda}$, and $\mathbf{G}$):

$$\mathbf{V} = \mathbf{U} + \beta \mathbf{R}' * (\mathbf{V} + \mathbf{G}) * \mathbf{S}. \quad (67)$$

This allows us to calculate the matrix of logit probabilities Π , with elements

$$\pi^{jk} \equiv \frac{\exp(v^{jk}/(\kappa w))}{\sum_{n=1}^{\#p} \exp(v^{nk}/(\kappa w))}. \quad (68)$$

We can then find the steady state distribution as the fixed point of these two equations:

$$\Psi = (\mathbf{E}_{pa} - \Lambda) . * \tilde{\Psi} + \Pi . * (\mathbf{E}_{pp} * (\Lambda . * \tilde{\Psi})) \quad (69)$$

$$\tilde{\Psi} = \mathbf{R} * \Psi * \mathbf{S}' \quad (70)$$

Finally, we check whether

$$1 = \sum_{j=1}^{\#p} \sum_{k=1}^{\#a} \Psi^{jk} (p^j)^{1-\epsilon} \equiv p(w) \quad (71)$$

If $p(w) = 1$, then an equilibrium value of w has been found.

Step 2: linearized dynamics

Given the steady state, the general equilibrium dynamics can be calculated by linearization. To do so, we eliminate as many variables from the equation system as we can. We can then summarize the general equilibrium equation system in terms of the exogenous shock process z_t , the lagged distribution of idiosyncratic states Ψ_{t-1} , which is the endogenous component of the time t aggregate state; and finally the endogenous 'jump' variables including $\mathbf{V}_t$, Π_t , C_t , R_t , and i_t . The equation system reduces to

$$z_t = \phi_z z_{t-1} + \epsilon_t^z \quad (72)$$

$$\frac{\mu \exp(z_t)}{\exp i_t} = \frac{m_t}{m_{t-1}} \quad (73)$$

$$\Psi_t = (\mathbf{E}_{pa} - \Lambda_t) . * \tilde{\Psi}_t + \Pi_t . * (\mathbf{E}_{pp} * (\Lambda_t . * \tilde{\Psi}_t)) \quad (74)$$

$$\mathbf{V}_t = \mathbf{U}_t + \beta E_t \left\{ \frac{C_{t+1}^{-\gamma}}{C_t^{-\gamma}} [\mathbf{R}'_{t+1} * (\mathbf{V}_{t+1} + \mathbf{G}_{t+1}) * \mathbf{S}] \right\} \quad (75)$$

$$1 = \sum_{j=1}^{\#p} \sum_{k=1}^{\#a} \Psi_t^{jk} \exp((1-\epsilon)p^j) \quad (76)$$

If we now collapse all the endogenous variables into a single vector

$$\vec{X}_t \equiv (\text{vec}(\Psi_{t-1})', \text{vec}(\mathbf{V}_t)', C_t, m_{t-1}, i_t)'$$

then the whole set of expectational difference equations (72)-(76) governing the dynamic equilibrium becomes a first-order system of the following form:

$$E_t \mathcal{F} \left(\vec{X}_{t+1}, \vec{X}_t, z_{t+1}, z_t \right) = 0 \quad (77)$$

where E_t is an expectation conditional on z_t and all previous shocks.

To see that the variables in vector $\vec{X}_t$ are in fact the only variables we need, note that given i_t and i_{t+1} we can construct $\mathbf{R}_t$ and $\mathbf{R}_{t+1}$. Given $\mathbf{R}_t$, we can construct $\tilde{\Psi}_t = \mathbf{R}_t * \Psi_{t-1} * \mathbf{S}'$ from Ψ_{t-1} . Given $w_t = \chi C_t^\gamma$, we can construct $\mathbf{U}_t$, with (j, k) element equal to $u_t^{jk} \equiv \left(\exp(p^j) - \frac{w_t}{\exp(a^k)} \right) \frac{C_t}{\exp(\epsilon p^j)}$. Finally, given $\mathbf{V}_t$ and $\mathbf{V}_{t+1}$ we can construct $\mathbf{\Pi}_t$, $\mathbf{D}_t$, and $\mathbf{D}_{t+1}$, and thus $\mathbf{\Lambda}_t$ and $\mathbf{G}_{t+1}$. Therefore the variables in $\vec{X}_t$ and z_t are indeed sufficient to evaluate the system (72)-(76).

Finally, if we linearize system $\mathcal{F}$ numerically with respect to all its arguments to construct the Jacobian matrices $\mathcal{A} \equiv D_{\vec{X}_{t+1}} \mathcal{F}$, $\mathcal{B} \equiv D_{\vec{X}_t} \mathcal{F}$, $\mathcal{C} \equiv D_{z_{t+1}} \mathcal{F}$, and $\mathcal{D} \equiv D_{z_t} \mathcal{F}$, then we obtain the following first-order expectational difference equation system:

$$E_t \mathcal{A} \Delta \vec{X}_{t+1} + \mathcal{B} \Delta \vec{X}_t + E_t \mathcal{C} z_{t+1} + \mathcal{D} z_t = 0 \quad (78)$$

where Δ represents a deviation from steady state. This system has the form considered by Klein (2000), so we solve our model using his QZ decomposition method.³⁷

³⁷ Alternatively, the equation system can be rewritten in the form of Sims (2001). We chose to implement the Klein method because it is especially simple and transparent to program.