Monetary policy implications
of
state-dependent prices and wages

James Costain
(Banco de España)

Anton Nakov
(European Central Bank and CEPR)

Borja Petit
(CEMFI)

PRELIMINARY AND INCOMPLETE
September 2017

Abstract

This paper studies the dynamic general equilibrium effects of monetary policy shocks in a "control cost" model of state-dependent retail price adjustment and state-dependent wage adjustment. Both suppliers of retail goods and suppliers of labor are monopolistic competitors subject to idiosyncratic productivity shocks and nominal rigidities. Nominal rigidity arises because precise choice is costly: decision-makers tolerate errors both in the timing of price and wage adjustments, and in the new level at which the price or wage is set, because achieving perfect precision in these decisions would be excessively costly.

The model is calibrated to match the size and frequency of price and wage adjustments. We find that sticky wages by themselves account for much of the nonneutrality that occurs in the model where both sticky wages and sticky prices are present. Hence, a model in which both prices and wages are sticky implies substantially larger real effects of monetary shocks than does a model with sticky prices only. Interestingly, when wages are sticky, decreasing price stickiness may increase these real effects.

Keywords: Sticky prices, sticky wages, state-dependent adjustment, logit equilibrium, near rationality, control costs

JEL Codes: E31, D81, C73

1 Introduction[1]

The nominal rigidity of prices and/or wages is a prominent assumption in monetary macroeconomics today. For reasons of analytical tractability, many models are based on Calvo's (1983) framework, in which the probability of adjustment is constant. But several influential papers have claimed that if nominal stickiness is derived from rational decision-making, instead of being imposed in an *ad hoc* way, then the real macroeconomic effects monetary policy are negligible (see for example the menu cost models of Caplin and Spulber, 1987, and Golosov and Lucas, 2007). This finding motivates a large new literature that is investigating how the conclusions of Calvo-style models and menu cost models hold up in a variety of state-dependent pricing frameworks that are closely calibrated to retail price microdata (*e.g.* Klenow and Kryvtsov, 2008; Gagnon, 2009; Matejka, 2010; Midrigan, 2011; Álvarez, González-Rozada, Neumeyer, and Beraja, 2011; Eichenbaum, Jaimovich, and Rebelo, 2011; Kehoe and Midrigan, 2012; Dotsey, King, and Wolman, 2013; Álvarez, Lippi, and Paciello, 2014; Costain and Nakov, 2011, 2015).

Much of this new literature concludes, to quote Kehoe and Midrigan, that "prices are sticky after all". That is, while stripped-down menu cost models like Golosov and Lucas (2007) imply that monetary policy is nearly neutral, related models that fit retail microdata better imply that price stickiness does in fact matter at the aggregate level, leading to nontrivial real effects of monetary policy.[2] This encouraging result appears to improve the link between microdata and modern macroeconomics. However, it may still be premature to draw strong quantitative conclusions from this literature, which for computational reasons revolves around models where sticky retail prices are the only friction considered. This contrasts with the current generation of empirical DSGE models that rely not only on nominal rigidity (à la Calvo, 1983, or Rotemberg, 1982) in prices and wages, but also on many other frictions, such as consumption habits, investment adjustment costs, and labor matching frictions. This suggests that further progress in understanding the quantitative relevance of nominal rigidity requires us to study models in which multiple frictions interact.

This paper takes a modest step forward by analyzing a model with one additional layer of state-dependent adjustment, allowing for wage stickiness as well as price stickiness.[3] A natural point of departure for our exercise is Erceg, Henderson, and Levin's (2000) study of monopolistic retail price setters and monopolistic wage setters under the Calvo framework. Following Erceg *et al.*, we set up the wage setters' problem so that it closely parallels the price setting problem, but we allow for state dependence in both decisions. More precisely, we compare a framework in which both price and wage

[1]Thanks to Isaac Baley, Jordi Galí, Erwan Gautier, Alok Johri, Julián Messina, Michael Reiter, Ernesto Villanueva, and seminar participants at CEF (2016 and 2017), EEA-ESEM (2016 and 2017), DYNARE 2016, T2M 2017, the Catalan Economic Society 2017, the 2017 Inflation Targeting Seminar of the Banco Central do Brasil, and at the Dutch Central Bank (2017) for helpful comments. 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.

[2]The reason for nonneutrality is that the microdata seem to favor specifications in which the "selection effect" is weaker than the Golosov and Lucas (2007) framework implies.

[3]We know of only one previous study of state-dependent prices and wages in a DSGE model, Takahashi (2017). But that paper abstracts from idiosyncratic shocks, so it cannot be closely calibrated to microdata. Takahashi's paper also differs from ours in that it analyzes a stochastic menu cost model (following Dotsey *et al.*, 1999) rather than a control cost model.

setters are constrained by the Calvo friction to a framework in which price and wage setters are both constrained by a state-dependent friction, and in addition we compare these with scenarios in which price setting and/or wage setting is almost perfectly flexible. We emphasize that the purpose of our paper is to compare different specifications of price stickiness and wage stickiness while abstracting from any other frictions that might affect the labor market (or other markets). While the interaction of nominal rigidities with labor market matching is a major area of recent research, our goal in this paper is to quantify the effects of state-dependent prices and wages by themselves, leaving their interaction with matching frictions for future work.

Our model of state-dependent adjustment is an extension of the "control cost" model of price stickiness proposed by Costain and Nakov (2015), henceforth CN15. Control costs are a modeling device from game theory designed to capture the idea that the costs of precise decision-making sometimes lead players to make some mistakes.[4] Under the control cost framework, a decision is regarded as a random variable defined over a set of feasible alternatives, and the decision-maker is assumed to face a cost function that increases with the precision of that random variable. Placing probability one on the optimal alternative is a very precise decision, so the decision-maker may instead economize on the costs of choice by tolerating some randomness (some errors) in the alternative chosen. CN15 models nominal rigidity by applying this framework *both* to the prices firms choose, and to firms' control of the *timing* of their adjustments. In equilibrium, in their model, managers of retail firms economize on the time devoted to decision-making by tolerating some low-cost errors in the prices they set, and some low-cost errors in the timing of their price adjustments.

There are a number of reasons why it seems interesting to extend the CN15 framework to other frictions, beyond price stickiness. First, it describes adjustment costs in a sparsely parameterized way; the benchmark scenario in CN15 simultaneously fits many "puzzling" features of retail price setting by calibrating only two free parameters in the decision cost function. Second, these costs have an appealing interpretation: the costs of price adjustment are interpreted as time devoted by management to decision-making. These may be plausibly larger than the menu-type fixed costs associated with the physical act of changing the price, and may be compared, at least roughly, to case studies on time use by management. Third, the model is no harder to solve numerically than comparable menu cost models, but it is far more tractable than "rational inattention" models in the tradition of Sims (2003). Fourth, the mathematical structure of the model— resetting a control variable at intermittent points of time— seems applicable to many decisions other than price adjustment, potentially allowing us to describe many margins of a general equilibrium model in a mutually consistent and mutually comparable way. Finally, since the calibration strategy in the recent state-dependent pricing literature involves matching many moments of the distribution of individual price adjustments, it stretches credulity to abstract from errors. When matching (for example) the standard deviation of observed price adjustments, inferences about the standard deviation of the underlying shocks may differ greatly depending on whether or not we insist that every single price adjustment represents a precisely optimal action.

[4] See Stahl (1990), Mattsson and Weibull (2002), or van Damme (1991), Ch. 4.

2 Model

This discrete-time model embeds near-rational price adjustment and wage adjustment in an otherwise standard New Keynesian general equilibrium framework that mixes elements of Erceg, Henderson, and Levin (2000) with others from Golosov and Lucas (2007). There is a continuum of retail firms and a continuum of workers; retail goods markets and labor markets are both monopolistically competitive. Each firm is the unique seller of a differentiated retail good, and resets its nominal price intermittently. Each worker is the unique seller of a differentiated type of labor, and resets its nominal wage intermittently. Price and wage adjustments are driven by idiosyncratic as well as aggregate shocks. Workers belong to a representative household; the budget constraint is defined at the household level. In addition, there is also a monetary authority that sets an exogenous growth process for the nominal money supply.

2.1 Household

The worker's period utility function is $u(C_t) - X(H_t) + v\,(M_t/P_t)$, where C_t is consumption, H_t is total time devoted to working *or* decision-making, and M_t/P_t is real money balances. The functions u and v are assumed increasing and concave. We assume the increasing, convex disutility function $X(H) = \frac{\chi}{1+\zeta}H^{1+\zeta}$. We will focus initially on the linear case $\zeta = 0$, implying $X(H) = \chi H$, which is easier to solve, but we will see below that the nonlinear specification $\zeta > 0$ is necessary to match wage adjustment data. Utility is discounted by factor β per period. Consumption is a CES aggregate of differentiated products C_{jt}, with elasticity of substitution ϵ:

$$C_t = \left\{ \int_0^1 C_{jt}^{\frac{\epsilon-1}{\epsilon}} dj \right\}^{\frac{\epsilon}{\epsilon-1}} . \tag{1}$$

The representative household consists of a continuum of workers, and aggregates their resources. Its period budget constraint, in nominal terms, is

$$\int_0^1 P_{jt}C_{jt}dj + M_t + R_t^{-1}B_t = \int_0^1 W_{it}H_{it}di + M_{t-1} + B_{t-1} + T_t^M + T_t^D. \tag{2}$$

Here $\int_0^1 P_{jt}C_{jt}dj$ is total nominal consumption, and $\int_0^1 W_{it}H_{it}di$ is total labor compensation received from supplying the differentiated labor varieties H_{it}. 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 $\{C_{jt}, B_t, M_t\}_{t=0}^{\infty}$ to maximize expected discounted utility, subject to the budget constraint (2).[5] The workers in each household set nominal wages intermittently, as we will discuss in Sec. 2.4, and they supply labor to fulfill the demand that arises given the nominal wages they have set.

[5] We use an abbreviated notation here for the sake of brevity. The time subscript on the household's decision variables should not be interpreted as indicating deterministic dependence on time; instead, it indicates dependence on the stochastic aggregate state of the economy.

Optimal consumption across the differentiated goods implies

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

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

$$P_t \equiv \left\{ \int_0^1 {P_{jt}}^{1-\epsilon} dj \right\}^{\frac{1}{1-\epsilon}}. \tag{4}$$

The first-order conditions for total consumption and for money use are:

$$R_t^{-1} = \beta E_t \left(\frac{P_t u'(C_{t+1})}{P_{t+1} u'(C_t)} \right), \tag{5}$$

$$1 - \frac{v'(M_t/P_t)}{u'(C_t)} = \beta E_t \left(\frac{P_t u'(C_{t+1})}{P_{t+1} u'(C_t)} \right). \tag{6}$$

2.2 Monetary policy

For simplicity, we consider a monetary authority that generates an exogenous process for the money growth rate. We assume the nominal money supply is affected by an AR(1) shock process g,[6]

$$g_t = \phi_g g_{t-1} + \epsilon_t^g, \tag{7}$$

where $0 \leq \phi_g < 1$ and $\epsilon_t^g \sim i.i.d.N(0, \sigma_g^2)$. Here g_t represents the time t rate of money growth:

$$M_t/M_{t-1} \equiv \mu_t = \mu^* \exp(g_t). \tag{8}$$

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 $M_t = M_{t-1} + T_t^M$.

One advantage of this specification of monetary policy is that we can analytically solve for the nominal interest rate process. STATE RESULT HERE...

2.3 Monopolistic firms

Each firm j produces output Y_{jt} under a constant returns technology $Y_{jt} = A_{jt} N_{jt}$. Efficiency units of labor, denoted N_{jt}, are the only input. A_{jt} represents an idiosyncratic productivity process that follows a time-invariant Markov process on a bounded set, $A_{jt} \in \Gamma^A \subseteq [\underline{A}, \overline{A}]$. Productivity innovations are *iid* across firms. Thus, A_{jt} is correlated with $A_{j,t-1}$, but it is uncorrelated with other firms' shocks. Firm j is a monopolistic competitor that sets a price P_{jt}, facing the demand curve $Y_{jt} = C_t P_t^{\epsilon} {P_{jt}}^{-\epsilon}$. We assume

[6] In related work (Costain and Nakov 2011) 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.

Figure 1: Sequencing of firms' decisions within the period.

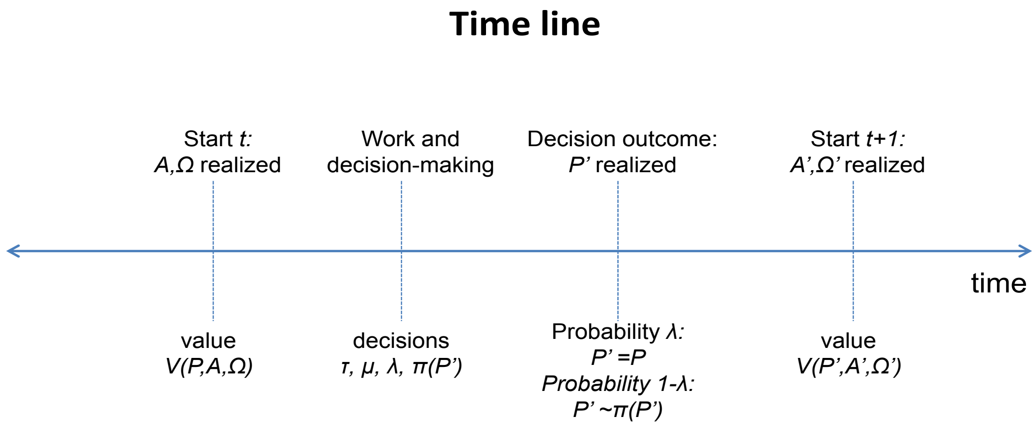

each firm must fulfill all demand at its chosen price. Since firms are infinitesimal, each firm j ignores the effect of its own price P_{jt} on the aggregate price level P_t. It hires labor at wage rate W_t, generating profits

$$U_{jt} = P_{jt}Y_{jt} - W_t N_{jt} = \left(P_{jt} - \frac{W_t}{A_{jt}}\right) C_t P_t^{\epsilon} P_{jt}^{-\epsilon} \tag{9}$$

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

It will help to distinguish value functions at several different points in time. First, let $V_t(P, A)$ be the value of a firm that begins period t with nominal price P and productivity A, prior to any time t decisions, and prior to time t output (see the timeline). We assume that choices take time, so if the firm decides in period t to adjust its price, the new price only becomes effective at time $t+1$.[7] Next, let $O_t(P, A)$ be the firm's continuation value, net of current profits, when it still has the option to adjust prices. That is,[8]

$$V_t(P, A) \;=\; \left(P - \frac{W_t}{A}\right) C_t P_t^{\epsilon} P^{-\epsilon} + O_t(P, A) \tag{10}$$

The continuation value $O_t(P, A)$ incorporates the value of the firm's two possible time-t decisions: whether to adjust its price, and if so, which new price P' to set for period $t+1$. The firm may make errors in either of these choices. We discuss these two decisions in turn, beginning with the latter.

[7] A one-period lag would be unrealistic if the time period were very long. When we calibrate the model, we will impose a very short time period, close enough to continuous time that a one-period lag is not restrictive.

[8] Again, we use succinct notation, where time subscripts on the value functions represent dependence on the aggregate state. Thus, if the aggregate state of the economy is Ω_t, we define $V_t(P, A) \equiv V(P, A, \Omega_t)$ and $O_t(P, A) \equiv O(P, A, \Omega_t)$, Time-subscripted variables in equation (10) represent aggregate quantities: $P_t \equiv P(\Omega_t)$ is the aggregate price level, $W_t \equiv W(\Omega_t)$ is the aggregate wage, and $C_t \equiv C(\Omega_t)$ is aggregate consumption demand.

2.3.1 Choosing a new price

Our model formalizes the idea that nominal rigidities may derive primarily from the costs of decision-making. While one might assume that by paying a fixed cost, the firm can make the optimal choice, this would amount to imposing a corner solution with perfect precision. We find it more appealing and more realistic to assume that firms can devote more or less time and resources to decision-making, in order to choose more or less precisely. In equilibrium in our framework firms will typically prefer to make choices with an interior degree of precision. Therefore their chosen action will not always be the one that would have been optimal in the absence of decision costs; instead, most choices will involve some degree of "error".

Consistent with this general description, we adopt the "control cost" approach from game theory (see van Damme, 1991, Chapter 4). A key feature of this approach is that we model the price decision indirectly: the firm's problem is written "as if" it chooses a probability distribution over prices, rather than choosing the price *per se*.[9] The problem incorporates a cost function that increases with precision: concentrating greater probability on a small range of prices increases costs. Many possible measures of precision could be used to define this cost function; we choose a definition based on relative entropy, also known as Kullback-Leibler divergence, which is a measure of the difference between one probability distribution and another. For two possible distributions $\pi_1(x)$ and $\pi_2(x)$ of some random variable x with support on set $\mathcal{X}$, the Kullback-Leibler divergence $\mathcal{D}(\pi_1||\pi_2)$ of π_1 relative to π_2 is defined by[10]

$$\mathcal{D}(\pi_1||\pi_2) = \int_{\mathcal{X}} \pi_1(x) \ln\left(\frac{\pi_1(x)}{\pi_2(x)}\right) dx. \tag{11}$$

Following Stahl (1990) and Mattsson and Weibull (2002), we assume that the decision cost is proportional to the Kullback-Leibler divergence of the chosen distribution, relative to an exogenous benchmark distribution. Thus, if no decision costs are paid, the action x is distributed according to the benchmark distribution. But by putting more effort into the decision process, the decision-maker can shrink the distribution of the action towards the most desirable alternatives.

We assume that decision costs are denominated in units of time, since we regard managers' time as the main input to decision-making. Consistent with typical US retail data, we will assume that when the firm sets a new nominal price $\tilde{P}$, this remains constant in nominal terms until a new adjustment occurs. We benchmark the cost of the decision process against an exogenous benchmark distribution $\eta_t(\widetilde{P})$ with support Γ_t^P. The time subscripts on η_t and Γ_t^P allow the benchmark price distribution to change over time, which allows the economy to have a nominal trend; later we detrend the model by restating it in real terms.

[9] Luce (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.

[10] While we write (11) with an integral, we can be agnostic at this point about whether $\mathcal{X}$ is a discrete or continuous set. If it is a continuous set, then π_1 and π_2 should be interpreted as density functions. If it is a discrete set, then π_1 and π_2 should be interpreted as vectors of probabilities, and the integral in (11) should be interpreted as a sum.

Assumption 1. The time cost of choosing a distribution $\pi(\widetilde{P})$ over nominal prices $\widetilde{P} \in \Gamma_t^P$ is $\kappa_\pi \mathcal{D}(\pi||\eta_t)$, where $\kappa_\pi > 0$ is a constant, and $\eta_t(\widetilde{P})$ is an exogenously-given benchmark distribution with support Γ_t^P.

Here κ_π represents the marginal cost of entropy reduction, in units of labor time. The cost function described in Assumption 1 is nonnegative and convex.[11] The upper bound on the cost function is associated with a distribution that places all probability on a single price $\widetilde{P}$ (concretely, costs are maximized when all probability is placed on one price that minimizes the benchmark probability $\eta_t(\widetilde{P})$). The lower bound on this cost function is zero, associated with choosing the distribution $\pi(\widetilde{P})$ equal to the benchmark distribution $\eta_t(\widetilde{P})$.

Now consider the pricing decision under this cost function. If the firm sets a new nominal price $\widetilde{P}$ at time t, this new price only becomes effective at $t+1$, so the value of setting $\widetilde{P}$ at t is

$$V_t^e(\widetilde{P}, A) \;\equiv\; E_t\left[\beta \frac{P_t u'(C_{t+1})}{P_{t+1} u'(C_t)} V_{t+1}(\widetilde{P}, A') \middle| A\right], \tag{12}$$

where $E_t[\bullet|A]$ represents an expectation over the time $t+1$ variables $\Omega' \equiv \Omega_{t+1}$ and $A' \equiv A_{j,t+1}$ conditional on the time t aggregate state Ω_t and firm j's time t productivity $A_{j,t} = A$. Following the control costs methodology, we assume the firm maximizes its value by allocating probability across possible nominal prices $\widetilde{P}$, taking account of decision costs, as follows:

$$\tilde{V}_t(A) = \max_{\pi(\widetilde{P})} \int \pi(\widetilde{P}) V_t^e(\widetilde{P}, A) d\widetilde{P} - \kappa_\pi W_t \int \pi(\widetilde{P}) \ln\left(\frac{\pi(\widetilde{P})}{\eta_t(\widetilde{P})}\right) d\widetilde{P} \quad \text{s.t.} \quad \int \pi(\widetilde{P}) d\widetilde{P} = 1 \tag{13}$$

Note that the decision costs in (13) are converted to nominal units by multiplying by the wage rate. We write the nominal value of the pricing decision as $\tilde{V}_t(A)$, where $A \equiv A_{jt}$ is the firm's current productivity.

The first-order condition for $\pi(\widetilde{P})$ in problem (13) is

$$V_t^e(\widetilde{P}, A) - \kappa_\pi W_t \left(1 + \ln\left(\frac{\pi(\widetilde{P})}{\eta_t(\widetilde{P})}\right)\right) - \mu = 0,$$

where μ is the multiplier on the constraint. Some rearrangement yields a weighted multinomial logit formula:

$$\pi_t(\widetilde{P}|A) \;\equiv\; \frac{\eta_t(\widetilde{P}) \exp\left(\frac{V_t^e(\widetilde{P}, A)}{\kappa_\pi W_t}\right)}{\int_{\Gamma^P} \eta_t(P') \exp\left(\frac{V_t^e(P', A)}{\kappa_\pi W_t}\right) dP'} \tag{14}$$

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

[11] Cover and Thomas (2006), Theorem 2.7.2.

an analytical formula for the value function:

$$\tilde{V}_t(A) = \kappa_\pi W_t \ln \left(\int \eta_t(\widetilde{P}) \exp \left(\frac{V_t^e(\widetilde{P}, A)}{\kappa_\pi W_t} \right) d\widetilde{P} \right). \tag{15}$$

This formula gives the nominal value of adjusting the current price, net of decision costs.

Some interpretive comments may be helpful at this point. First, although we write the firm's problem "as if" it chooses a probability distribution over prices, this should not be taken literally— actually choosing a distribution would be a complex, costly diversion from the true task of choosing the price itself. Rather, we define the decision as a choice of a mixed strategy because this is a way to incorporate errors into the model. And we describe it as an optimization problem because this disciplines the errors; it amounts to assuming that the firm devotes time and effort to avoiding especially costly mistakes. Aspects of the model that we do take seriously include (a) making decisions is costly in terms of time and other resources; (b) therefore decision-makers do not always take the action that would otherwise be optimal; (c) *ceteris paribus*, more valuable actions are more probable; (d) in a retail pricing context, these considerations apply both to the timing of price adjustment, and to the actual price chosen.

Second, the problem is written conditional on the true expected discounted values $V_t^e(\widetilde{P}, A)$ of the possible nominal prices $\widetilde{P}$, instead of conditioning on a prior, as a "rational inattention" model would. This reflects the fact that we are *not* assuming imperfect information. But this is different from saying that the firm "knows" the true values $V_t^e(\widetilde{P}, A)$. Instead, our interpretation is that the firm has sufficient information to calculate $V_t^e(\widetilde{P}, A)$. Even so, drawing correct conclusions from that information, and acting accordingly, may be costly.[12]

2.3.2 Choosing the timing of price adjustment

We next analyze, in an analogous manner, the decision whether or not to adjust at time t. As in the previous subsection, we define costs relative to a benchmark probability distribution over possible actions. But for this decision, at any t, there are only two options: adjust now, or not. Since the probability of these two actions must sum to one, effectively the relevant benchmark is just a single number, which we can interpret as an exogenous default hazard rate.

We suppose the time period is sufficiently short so that we can approximately ignore multiple adjustments within a single period. If the firm chooses not to adjust its current price P, then its nominal price next period will be unchanged: $\widetilde{P}' = P$; the expected value of this unchanged price, from the point of view of period t, is $V_t^e(P, A)$. If instead the firm adjusts its price at time t, then its expected value is $\tilde{V}_t(A)$, as given by (13) and (15). Now suppose it adjusts its price with probability λ. We measure the

[12]Since economists are accustomed to models of perfect rationality, they often equate observing a given information set with knowing all quantities that can be calculated from that information set. But when rationality is less than perfect, we cannot equate these two assumptions. Here, we assume firms can observe all relevant shocks and state variables, but we do not equate this with actually knowing $V_t^e(\widetilde{P}, A)$ or knowing the optimal action, and therefore we do not equate it with implementing the optimal action with probability one.

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 setting the price adjustment hazard $\lambda \in [0,1]$ in period t is $\kappa_\lambda \mathcal{D}((\lambda, 1-\lambda)||(\bar{\lambda}, 1-\bar{\lambda}))$, where $\kappa_\lambda > 0$ and $\bar{\lambda} \in [0,1]$ are exogenous parameters.

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.

Rewriting this cost function using definition (11), the optimal adjustment probability at time t solves the following Bellman equation:

$$O_t(P,A) \;=\; \max_{\lambda}\,(1-\lambda)V_t^e(P,A) + \lambda\tilde{V}_t(A) - \kappa_\lambda W_t\left[\lambda\ln\left(\frac{\lambda}{\bar{\lambda}}\right) + (1-\lambda)\ln\left(\frac{1-\lambda}{1-\bar{\lambda}}\right)\right]. \quad (16)$$

Recall that $O_t(P,A)$ represents the continuation value of the firm, net of decision costs, when it still has the option to adjust, or not to do so. The first order condition from (16) is

$$V_t^e(P,A) - \tilde{V}_t(A) = \kappa_\lambda W_t\left[\ln\lambda + 1 - \ln\bar{\lambda} - \ln(1-\lambda) - 1 + \ln(1-\bar{\lambda})\right]. \quad (17)$$

Rearranging, we can solve (17) to obtain[13]

$$\lambda_t(P,A) \;=\; \frac{\bar{\lambda}\exp\left(\frac{\tilde{V}_t(A)}{\kappa_\lambda W_t}\right)}{\bar{\lambda}\exp\left(\frac{\tilde{V}_t(A)}{\kappa_\lambda W_t}\right) + (1-\bar{\lambda})\exp\left(\frac{V_t^e(P,A)}{\kappa_\lambda W_t}\right)} \quad (18)$$

$$\;=\; \frac{\bar{\lambda}}{\bar{\lambda} + (1-\bar{\lambda})\exp\left(\frac{-D_t(P,A)}{\kappa_\lambda W_t}\right)}\,, \quad (19)$$

where $D_t(P,A)$ is the expected gain from adjustment:

$$D_t(P,A) \;\equiv \tilde{V}_t(A) - V_t^e(P,A). \quad (20)$$

The weighted binary logit hazard (18) was also derived by Woodford (2008) from a model with a Shannon constraint.[14] 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 (*i.e.* when $D_t(P,A) = 0$).[15]

[13]Note also that (18) 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 (18) is $\lambda_t(P,A) = \bar{\lambda}\exp\left(\frac{D_t(P,A)}{\kappa_\lambda W_t}\right)$.

[14]Woodford's (2009) paper only states a first-order condition like (17); his (2008) manuscript points out that the first-order condition implies a logit hazard of the form (18).

[15]This model nests Calvo price adjustment as a special case. If we set $\kappa_\pi = 0$ and $\kappa_\lambda = \infty$, then the firm always sets the optimal price, conditional on adjustment, and adjustment occurs with a constant probability $\bar{\lambda}$.

2.3.3 Deriving the Bellman equation

A two-step decision

Next, to calculate the value function $V_t(P,A)$, we concatenate the two decision steps described in Secs. 2.3.1-2.3.2. If the firm starts period t with nominal price P, then its value $V_t(P,A) \equiv V_t(P,A,\Omega_t)$ at the beginning of t satisfies:

$$\begin{aligned} V_t(P,A) \;=\; & \max_{\lambda,\pi(\widetilde{P})} \left(P - \frac{W_t}{A}\right) C_t P_t^{\epsilon} P^{-\epsilon} + (1-\lambda) V_t^e(P,A) + \lambda \int \pi(\widetilde{P}) V_t^e(\widetilde{P},A) d\widetilde{P} \qquad (21) \\ & - \lambda \kappa_\pi W_t \int \pi(\widetilde{P}) \ln\left(\frac{\pi(\widetilde{P})}{\eta_t(\widetilde{P})}\right) d\widetilde{P} \;-\; \kappa_\lambda W_t \left[\lambda \ln\left(\frac{\lambda}{\bar{\lambda}}\right) + (1-\lambda)\ln\left(\frac{1-\lambda}{1-\bar{\lambda}}\right)\right] \\ & \text{s.t.} \quad \int \pi(\widetilde{P}) d\widetilde{P} = 1. \end{aligned}$$

This Bellman equation subtracts of the two cost functions seen in the previous subsections. There is a time cost associated with monitoring whether or not a price adjustment is required, which we will call

$$\mu_t(P,A) \;\equiv\; \kappa_\lambda \left[\lambda \ln\left(\frac{\lambda}{\bar{\lambda}}\right) + (1-\lambda)\ln\left(\frac{1-\lambda}{1-\bar{\lambda}}\right)\right]. \qquad (22)$$

The time cost of choosing which new price to set is

$$\tau_t(P,A) \;\equiv\; \lambda \kappa_\pi \int \pi(\widetilde{P}) \ln\left(\frac{\pi(\widetilde{P})}{\eta_t(\widetilde{P})}\right) d\widetilde{P}. \qquad (23)$$

Finally, the time devoted to the actual production of goods will be written as

$$N_t(P,A) \equiv \frac{C_t}{A}\left(\frac{P_t}{P}\right)^{\epsilon}. \qquad (24)$$

Hence, the firm's total demand for labor hours is $N_t(P,A) + \mu_t(P,A) + \tau_t(P,A)$.

A one-step decision

While Sections 2.3.1-2.3.2 distinguished between the two steps of the choice process in order to analyze each one in detail, the firm's problem can alternatively (and equivalently) be written as a single decision. This decision is defined over an augmented set of alternatives $\Gamma_t^{P\dagger}$:

$$\Gamma_t^{P\dagger} \equiv \{\text{Don't adjust, Adjust to } \tilde{P}, \text{ where } \tilde{P} \in \Gamma_t^P\}. \qquad (25)$$

In other words, the feasible set includes the discrete alternative of not adjusting, together with the alternatives of adjusting to any $\tilde{P} \in \Gamma_t^P$, which could in principle be a continuum.

Bellman equation (21) can always be regarded as a single decision over the augmented choice set (25). But this representation is especially helpful in the special case $\kappa_\pi = \kappa_\lambda = \kappa$. Without loss of generality, suppose the firm chooses the first alternative with probability $1-\lambda$, and draws from the remaining alternatives with density $\pi^\dagger(\tilde{P}) = \lambda\pi(\tilde{P})$, where $1-\lambda+\int\pi^\dagger(\tilde{P})d\tilde{P} = 1$. Then the two cost terms in (21) can be combined as follows:

$$\begin{aligned}
&\lambda\kappa_\pi W_t \int \pi(\widetilde{P}) \ln\left(\frac{\pi(\widetilde{P})}{\eta_t(\widetilde{P})}\right) d\widetilde{P} \;+\; \kappa_\lambda W_t\left[\lambda\ln\left(\frac{\lambda}{\bar{\lambda}}\right) + (1-\lambda)\ln\left(\frac{1-\lambda}{1-\bar{\lambda}}\right)\right] \\
&= \kappa W_t\left\{(1-\lambda)\ln\left(\frac{1-\lambda}{1-\bar{\lambda}}\right) \;+\; \lambda\left[\int \pi(\widetilde{P})\ln\left(\frac{\pi(\widetilde{P})}{\eta_t(\widetilde{P})}\right)d\widetilde{P} \;+\; \ln\left(\frac{\lambda}{\bar{\lambda}}\right)\right]\right\} \\
&= \kappa W_t\left\{(1-\lambda)\ln\left(\frac{1-\lambda}{1-\bar{\lambda}}\right) \;+\; \lambda\left[\int \pi(\widetilde{P})\ln\left(\frac{\lambda\pi(\widetilde{P})}{\bar{\lambda}\eta_t(\widetilde{P})}\right)d\widetilde{P}\right]\right\} \\
&= \kappa W_t\left\{(1-\lambda)\ln\left(\frac{1-\lambda}{1-\bar{\lambda}}\right) \;+\; \left[\int \pi^\dagger(\widetilde{P})\ln\left(\frac{\pi^\dagger(\widetilde{P})}{\eta_t^\dagger(\widetilde{P})}\right)d\widetilde{P}\right]\right\}. \qquad (26)
\end{aligned}$$

The expression in (26) is a relative entropy measure over the set of alternatives $\Gamma_t^{P\dagger}$: it is the Kullback-Leibler divergence of the probabilities $(1-\lambda, \pi^\dagger(\tilde{P}))$ relative to the default distribution $(1-\bar{\lambda}, \eta^\dagger(\tilde{P}))$, where $\eta^\dagger(\tilde{P}) \equiv \bar{\lambda}\eta(\tilde{P})$.

Therefore, in the special case $\kappa_\pi = \kappa_\lambda = \kappa$, the Bellman equation (21) can be rewritten as:

$$\begin{aligned}
V_t(P,A) \;=\; &\max_{1-\lambda,\pi^\dagger(\widetilde{P})} \left(P - \frac{W_t}{A}\right) C_t P_t^\epsilon P^{-\epsilon} + (1-\lambda)V_t^e(P,A) + \int \pi^\dagger(\widetilde{P}) V_t^e(\widetilde{P},A)d\widetilde{P} \qquad (27) \\
&- \kappa W_t\left\{(1-\lambda)\ln\left(\frac{1-\lambda}{1-\bar{\lambda}}\right) \;+\; \left[\int \pi^\dagger(\widetilde{P})\ln\left(\frac{\pi^\dagger(\widetilde{P})}{\eta_t^\dagger(\widetilde{P})}\right)d\widetilde{P}\right]\right\} \\
&\text{s.t.}\;\; 1-\lambda+\int\pi^\dagger(\widetilde{P})d\widetilde{P} = 1.
\end{aligned}$$

An advantage of writing the problem this way is that the payoffs from the choice are a linear function of the probabilities $(1-\lambda, \pi^\dagger(\tilde{P}))$, and the costs are strictly convex, because relative entropy is a strictly convex function of probabilities. Therefore the Bellman equation maximizes a stictly concave function (expected payoffs minus costs) over a convex set (the simplex on which the probabilities integrate to one). Hence each backwards induction step has a unique solution, given by the unique solution to the first-order conditions.[16]

2.4 Labor market

We next construct a model of nominal wage rigidity analogous to our treatment of nominal price rigidity. We suppose each worker i is the monopolistic supplier of a specific type of labor H_{it}, sold at wage W_{it}

[16]When $\kappa_\pi \neq \kappa_\lambda$, a similar argument can instead be applied separately to the two decision steps discussed in Sections 2.3.1-2.3.2, to prove existence of a unique solution at each step. (((IS THIS SUBSECTION/FOOTNOTE REALLY NEEDED? MAYBE I SHOULD MOVE THIS FOOTNOTE TO LAB MKT SECTION??)))

per unit of time. The value of worker i's labor H_{it} is shifted by a shock process Z_{it}, which follows a time-invariant Markov process on a bounded set, $Z_{it} \in \Gamma^Z \subset [\underline{Z}, \overline{Z}]$. We will define $N_{it} = Z_{it} H_{it}$ as the "effective labor" of worker i. By this definition, we can say that the price of effective labor is $\frac{W_{it}}{Z_{it}}$. The idiosyncratic shock process Z_{it} represents worker-specific changes in the market value of labor, which may include various forms of human capital accumulation.

Firm j's labor input into goods production, N_{jt}, is defined as a CES aggregate across varieties of effective labor i, with elasticity of substitution ϵ_n. That is,

$$N_{jt} = \left\{ \int_0^1 N_{ijt}^{\frac{\epsilon_n - 1}{\epsilon_n}} di \right\}^{\frac{\epsilon_n}{\epsilon_n - 1}}. \tag{28}$$

It is straightforward to show that under this demand structure, the firm's optimal hiring satisfies

$$H_{ijt} \equiv \frac{N_{ijt}}{Z_{it}} = Z_{it}^{\epsilon_n - 1} \left(\frac{W_{it}}{W_t} \right)^{-\epsilon_n} N_{jt}, \tag{29}$$

when we define the wage index

$$W_t \equiv \left\{ \int_0^1 \left(\frac{W_{it}}{Z_{it}} \right)^{1-\epsilon_n} di \right\}^{\frac{1}{1-\epsilon_n}}. \tag{30}$$

Firm j's nominal wage bill for goods production is then

$$\int_0^1 W_{it} H_{ijt} di = W_t N_{jt}. \tag{31}$$

We assume that firms use the same CES mix of labor for decision making that they use for goods production. Then (29) implies that total demand for worker i's time is

$$H_{it} = Z_{it}^{\epsilon_n - 1} \left(\frac{W_{it}}{W_t} \right)^{-\epsilon_n} N_t, \tag{32}$$

where N_t represents aggregate labor demand by all firms. N_t includes demand for labor used in goods production, given by (24), and demand for labor in decision making, given by (22)-(23).

The worker adjusts her nominal wage W_{it} intermittently to maximize the value of labor income net of labor disutility. This decision is subject to control costs, both on the timing decision, and on the choice of which wage to set. We assume workers act in the interest of the households of which they form part, and that their consumption is fully insured by the household. Therefore they discount future income at the same rate $\beta \frac{P_t u'(C_{t+1})}{P_{t+1} u'(C_t)}$ that applies to the household and firm. Now let $L_t(W, Z)$ be the nominal value of a worker with wage W and productivity Z at the beginning of period t, before supplying labor, and before making any decisions. As in the case of price decisions, we assume that a wage adjustment in period t becomes effective in period $t+1$. Therefore the value of setting the nominal wage to some arbitrary new value $\widetilde{W}$ is

$$L_t^e(\widetilde{W}, Z) \;\equiv\; E_t\left[\beta\frac{P_t u'(C_{t+1})}{P_{t+1}u'(C_t)}L_{t+1}(\widetilde{W}, Z')\middle|\, Z\right].$$

We make two assumptions about the cost of workers' decision-making that are analogous to those we made in the case of the firm.

Assumption 3. The time cost of choosing a distribution $\pi^W(\widetilde{W})$ over nominal wages $\widetilde{W} \in \Gamma_t^W$ is $\kappa_w \mathcal{D}(\pi^W||\eta_t^W)$, where $\kappa_w > 0$ is a constant, and $\eta_t^W(\widetilde{W})$ is an exogenously-given benchmark distribution with support Γ_t^W.

Assumption 4. The time cost incurred in period t by setting the wage adjustment hazard $\rho \in [0,1]$ in period t is $\kappa_\rho \mathcal{D}((\rho_t, 1-\rho_t)||(\bar{\rho}, 1-\bar{\rho}))$, where $\kappa_\rho > 0$ and $\bar{\rho} \in [0,1]$ are exogenous parameters.

Now, let τ^w be the (expected) amount of time dedicated in period t to setting a new wage; let μ^w be the time dedicated to monitoring whether it is a good moment to reset the wage; and let $H_t(W,Z) = Z^{\epsilon_n - 1}N_t\left(\frac{W_t}{W}\right)^{\epsilon_n}$ be labor demand conditional on the current wage W and current productivity Z. We can then write the worker's wage setting problem in a form analogous to the pricing problem (21):

$$\begin{aligned}
L_t(W,Z) = & \max_{\tau^w,\mu^w,\rho,\pi^W(\widetilde{W})} WH_t(W,Z) - \frac{P_t}{u'(C_t)}X(H_t(W,Z)+\tau^w+\mu^w) + (1-\rho)L_t^e(Z,W) + \rho\int \pi^W(\widetilde{W})L_t^e(\widetilde{W},Z)d\widetilde{W} \\
& \text{s.t.} \quad \int \pi^W(\widetilde{W})d\widetilde{W} = 1, \\
& \rho\kappa_w \int \pi^W(\widetilde{W})\ln\left(\frac{\pi^W(\widetilde{W})}{\eta_t^W(\widetilde{W})}\right)d\widetilde{W} \;=\; \tau^w, \\
& \kappa_\rho\left[\rho\ln\left(\frac{\rho}{\bar{\rho}}\right) + (1-\rho)\ln\left(\frac{1-\rho}{1-\bar{\rho}}\right)\right] \;=\; \mu^w. \qquad (33)
\end{aligned}$$

Notice that (33) allows for a nonlinear labor disutility function X; this function is scaled by the factor $P_t/u'(C_t)$ to express the whole Bellman equation in nominal units.

Recall now that we stated the firm's decision in two separate steps, (16) and (13), representing the decision of whether or not to adjust prices, and the decision of what price to set conditional on adjustment, respectively. This decomposition of the pricing decision was possible because we assumed the firm could hire any quantity of labor at the (aggregate) wage rate W_t: that is, the firms' labor costs were a linear function of its labor demand. But imposing a linear cost function for a worker's time use would be much more restrictive. We will compute an example with a linear labor disutility function $X(h) = \chi h$ in Sec. 3.2, but we will find that a more general, nonlinear specification $X(h) = \chi\frac{h^{1+\zeta}}{1+\zeta}$ is needed to match wage adjustment data. But therefore we cannot simply condition on a constant marginal cost of labor: the quantity of labor supplied to the firm affects the marginal cost of time used for each type of decision-making, so the two decisions are analyzed simultaneously in the wage setting problem (33).

Nonetheless, the policy functions for wage setting and wage adjustment timing resemble the policy functions from the firm's problem. Following our previous calculations, we find that if the worker adjusts, she chooses the following density over nominal wages $\widetilde{W}$:

$$\pi_t^W(\widetilde{W}|W,Z) \;\equiv\; \frac{\eta_t^W(\widetilde{W})\exp\left(\frac{L_t^e(\widetilde{W},Z)}{\kappa_w x_t(W,Z)}\right)}{\int \eta_t^W(W')\exp\left(\frac{L_t^e(W',Z)}{\kappa_w x_t(W,Z)}\right)dW'}\,. \tag{34}$$

Here

$$x_t(W,Z) \;\equiv\; \frac{P_t}{u'(C_t)}X'(H_t(W,Z)+\tau_t^w(W,Z)+\mu_t^w(W,Z)) \tag{35}$$

denotes the marginal disutility of time in period t. Marginal disutility is computed from the three uses of the worker's time: labor hours $H_t(W,Z)$ demanded by the worker's employers, and the two components of the time use in the wage decision process, $\tau_t^w(W,Z)$ and $\mu_t^w(W,Z)$. Note also that (35) rescales disutility to nominal units for comparability with the value function L^e.

Likewise, if the worker's beginning-of-period wage and productivity are W and Z, her optimal adjustment probability must satisfy:

$$\rho_t(W,Z) \;=\; \frac{\bar{\rho}\exp\left(\frac{\tilde{L}_t(W,Z)}{\kappa_\rho x_t(W,Z)}\right)}{\bar{\rho}\exp\left(\frac{\tilde{L}_t(W,Z)}{\kappa_\rho x_t(W,Z)}\right)+(1-\bar{\rho})\exp\left(\frac{L_t^e(W,Z)}{\kappa_\rho x_t(W,Z)}\right)} \tag{36}$$

$$\;=\; \frac{\bar{\rho}}{\bar{\rho}+(1-\bar{\rho})\exp\left(\frac{-D_t^W(W,Z)}{\kappa_\rho x_t(W,Z)}\right)}\,, \tag{37}$$

where

$$D_t^W(W,Z) \;\equiv \tilde{L}_t(W,Z) - L_t^e(W,Z)\,. \tag{38}$$

Note that in these formulas, x_t is a function of the current wage W, except in the linear disutility case $X(h)=\chi h$, where it reduces to $x_t = P_t\chi/u'(C_t)$, which is independent of the idiosyncratic state. Therefore, although the worker's probability formulas (34) and (36) have the same functional form as the firm's policies (14) and (19), the worker's probabilities require a higher-dimensional calculation when labor disutility is nonlinear, since they then vary across W, Z, and $\widetilde{W}$. For this reason, Sec. 3.2 will first compute an example with linear disutility, before we attempt the higher-dimensional calculation of the nonlinear case in Sec. 3.3.

2.5 Detrending

Before we describe the dynamics of the distributions of firms and workers, it is helpful to remove the nominal trend from the description of firms' and workers' behavior. If we choose the default distributions for nominal prices and wages, $\eta_t^P(\widetilde{P})$ and $\eta_t^W(\widetilde{W})$, so that they can be interpreted as unchanging distributions $\eta^p(\widetilde{p})$ and $\eta^w(\widetilde{w})$ of *real* prices and wages, then the maximization problems of the firms and workers are homogeneous of degree one in nominal prices. Then the Bellman equations can be stated equivalently in real terms, rather than nominal terms as we did above.

Let Ω_t be a nominal aggregate state variable for this economy at time t. This implies that there exist functions P and W that define the nominal price level and the nominal wage level as a function of Ω_t:

$$P_t = P(\Omega_t), \tag{39}$$
$$W_t = W(\Omega_t). \tag{40}$$

We will define real variables by dividing by the aggregate price level, and we will treat all idiosyncratic real variables in logs. In particular, we define the following idiosyncratic quantities:

$$p_{jt} \equiv \ln P_{jt} - \ln P(\Omega_t), \tag{41}$$
$$\widetilde{p}_{jt} \equiv \ln \widetilde{P}_{jt} - \ln P(\Omega_t), \tag{42}$$
$$a_{jt} \equiv \ln A_{jt}, \tag{43}$$
$$w_{it} \equiv \ln W_{it} - \ln P(\Omega_t), \tag{44}$$
$$\widetilde{w}_{it} \equiv \ln \widetilde{W}_{it} - \ln P(\Omega_t), \tag{45}$$
$$z_{it} \equiv \ln Z_{it}, \tag{46}$$
$$\xi_{it} \equiv x(W_{it}, Z_{it}, \Omega_t)/P(\Omega_t). \tag{47}$$

Since we wish to define the default distributions of *real* prices and wages as time invariant, the default distributions of nominal variables are defined accordingly. Given $\widetilde{P} \equiv P(\Omega_t)e^{\widetilde{p}}$, we must have $\eta_t^P(\widetilde{P}) = \widetilde{P}^{-1}\eta^p(\widetilde{p})$. Likewise, given $W \equiv P(\Omega_t)e^{\widetilde{w}}$, we must have $\eta_t^W(\widetilde{W}) = \widetilde{W}^{-1}\eta^w(\widetilde{w})$.[17]

Now let Ξ_t be the real variable constructed by replacing all nominal state variables that are included in Ω_t by their log real counterparts, and by likewise replacing any distributions of nominal idiosyncratic state variables that are included in Ω_t by the corresponding distributions of log real state variables. It is reasonable to conjecture that Ξ_t is a valid real aggregate state variable for this economy at time t. If so, there must exist functions m, w, x, and i that determine the real money supply, the real aggregate wage, and the inflation rate in terms of Ξ:

$$m_t \equiv M_t/P(\Omega_t) = m(\Xi_t), \tag{48}$$
$$w_t \equiv W(\Omega_t)/P(\Omega_t) = w(\Xi_t), \tag{49}$$
$$i_t \equiv \ln P(\Omega_t) - \ln P(\Omega_{t-1}) = i(\Xi_t, \Xi_{t-1}). \tag{50}$$

Note that the variable i_t represents the consumer price inflation rate between periods $t-1$ and t. Likewise,

[17] To see this, when we say that there is an unchanging distribution of $\widetilde{p}$, we mean that $cdf_t^P(\widetilde{P}) = cdf^p(\widetilde{p})$, evaluated at the point $\widetilde{P} = P_t e^{\widetilde{p}}$. Using the chain rule, this implies $\frac{\partial cdf_t^P}{\partial P}(\widetilde{P})P_t e^{\widetilde{p}} = \frac{\partial cdf^p}{\partial p}(\widetilde{p})$. Then since $\eta_t^P(\widetilde{P}) \equiv \frac{\partial cdf_t^P}{\partial P}(\widetilde{P})$ and $\eta^p(\widetilde{p}) \equiv \frac{\partial cdf^p}{\partial p}(\widetilde{p})$ we obtain $\eta_t^P(\widetilde{P}) = \widetilde{P}^{-1}\eta^p(\widetilde{p})$.

it must be possible to write aggregate consumption and labor as functions of the real state, so that

$$c(\Xi_t) \;=\; C_t \;\equiv\; C(\Omega_t), \tag{51}$$

$$n(\Xi_t) \;=\; N_t \;\equiv\; N(\Omega_t), \tag{52}$$

and firm-specific labor demand can be written as

$$h(w, z, \Xi_t) \;\equiv\; H(P(\Omega_t)e^w, e^z, \Omega_t) \;=\; e^{z(\epsilon_n - 1)} n(\Xi_t) w(\Xi_t)^{\epsilon_n} e^{-\epsilon_n w}. \tag{53}$$

We will now argue that if such a real state variable exists, then the Bellman equations of the firms and workers can be rewritten in terms of real value functions v and v^e that satisfy the identities[18]

$$v(p, a, \Xi) \;\equiv\; \frac{V(P(\Omega)e^p, e^a, \Omega)}{P(\Omega)}, \tag{54}$$

$$v^e(p, a, \Xi) \;\equiv\; \frac{V^e(P(\Omega)e^p, e^a, \Omega)}{P(\Omega)} \;=\; \beta E\left\{ \frac{u'(c(\Xi_{t+1}))}{u'(c(\Xi_t))} v(p - i_{t+1}, a', \Xi_{t+1}) \middle| a, \Xi_t \right\}. \tag{55}$$

We see in (55) that in the absence of any nominal price adjustment, a log real price p at time t becomes $p - i_{t+1}$ at time $t + 1$.[19] Now, assuming that the real value functions v and v^e exist, (21) becomes:

$$\begin{aligned} v(p, a, \Xi_t) \;=\; \max_{\lambda, \pi^p(\tilde{p})} & \left(e^p - \frac{w(\Xi_t)}{e^a}\right) c(\Xi_t) e^{-\epsilon p} + (1 - \lambda) v^e(p, a, \Xi_t) + \lambda \int \pi^p(\tilde{p}) v^e(\tilde{p}, a, \Xi_t) d\tilde{p} \\ & - \; \lambda \kappa_\pi w(\Xi_t) \int \pi^p(\tilde{p}) \ln\left(\frac{\pi^p(\tilde{p})}{\eta^p(\tilde{p})}\right) d\tilde{p} \;-\; \kappa_\lambda w(\Xi_t) \left[\lambda \ln\left(\frac{\lambda}{\bar{\lambda}}\right) + (1 - \lambda) \ln\left(\frac{1 - \lambda}{1 - \bar{\lambda}}\right)\right] \\ & \text{s.t.} \quad \int \pi^p(\tilde{p}) d\tilde{p} = 1. \end{aligned} \tag{59}$$

Obviously, the worker's Bellman equation (33) can be detrended in analogy with that of the firm. To do so, we postulate real value functions l and l^e that satisfy the identities

$$l(w, z, \Xi) \;\equiv\; \frac{L(P(\Omega)e^w, e^z, \Omega)}{P(\Omega)}, \tag{60}$$

$$l^e(w, z, \Xi) \;\equiv\; \frac{L^e(P(\Omega)e^w, e^z, \Omega)}{P(\Omega)} \;=\; \beta E\left\{ \frac{u'(c(\Xi_{t+1}))}{u'(c(\Xi_t))} l(w - i_{t+1}, z', \Xi_{t+1}) | z, \Xi_t \right\}. \tag{61}$$

[18]Here we are not yet describing which variables are included in the real state Ξ. We will identify a candidate for the real state Ξ in the next subsections, as we describe the real distributional dynamics.

[19]To prove (55), note that

$$P(\Omega) v^e(p, a, \Xi_t) \equiv V^e(P(\Omega)e^p, e^a, \Omega) \;=\; E\left\{ \beta \frac{P(\Omega_t) u'(C(\Omega_{t+1}))}{P(\Omega_{t+1}) u'(c(\Omega_t))} V(P(\Omega_t)e^p, A', \Omega_{t+1}) \middle| A, \Omega_t \right\} \tag{56}$$

$$= E\left\{ \beta \frac{P(\Omega_t) u'(C(\Omega_{t+1}))}{P(\Omega_{t+1}) u'(C(\Omega_t))} V(P(\Omega_{t+1}) e^{p - i_{t+1}}, A', \Omega_{t+1}) \middle| A, \Omega_t \right\} \tag{57}$$

$$= P(\Omega_t) E\left\{ \beta \frac{u'(c(\Xi_{t+1}))}{u'(c(\Xi_t))} v(p - i_{t+1}, a', \Xi_{t+1}) \middle| a, \Xi_t \right\}. \tag{58}$$

The worker's Bellman equation can then be rewritten in real terms as follows:

$$\begin{aligned}
l(w,z,\Xi_t) &= \max_{\tau^w,\mu^w,\rho,\pi^w(\tilde{w})} e^w h(w,z,\Xi_t) - \frac{X(h(w,z,\Xi_t)+\tau^w+\mu^w)}{u'(c(\Xi_t))} + (1-\rho) l_t^e(w,z,\Xi_t) + \rho \int \pi^w(\tilde{w}) l^e(\tilde{w},z,\Xi_t) d\tilde{w} \\
&\quad \text{s.t.} \quad \int \pi^w(\tilde{w}) d\tilde{w} = 1, \\
&\quad \rho\kappa_w \int \pi^w(\tilde{w}) \ln\left(\frac{\pi^w(\tilde{w})}{\eta^w(\tilde{w})}\right) d\tilde{w} \;=\; \tau^w, \\
&\quad \kappa_\rho \left[\rho \ln\left(\frac{\rho}{\bar{\rho}}\right) + (1-\rho)\ln\left(\frac{1-\rho}{1-\bar{\rho}}\right)\right] \;=\; \mu^w.
\end{aligned} \tag{62}$$

Analyzing (62), it is straightforward to show that the chosen distribution of wages takes the form[20]

$$\pi_t^w(\widetilde{w}|w,z) \;\equiv\; \frac{\eta^w(\widetilde{w}) \exp\left(\frac{l_t^e(\widetilde{w},w)}{\kappa_w \xi_t(w,z)}\right)}{\int \eta^w(w') \exp\left(\frac{l_t^e(w',z)}{\kappa_w \xi_t(w,z)}\right) dw'}, \tag{63}$$

where

$$\xi_t(w,z) \;\equiv\; \frac{X'(h_t(w,z)+\tau_t^w(w,z)+\mu_t^w(w,z))}{u'(C_t)} \tag{64}$$

is the worker's marginal disutility of time spent working, expressed in units of consumption goods. Similarly, using the first-order condition for ρ, it can be shown that the adjustment hazard takes the following form:

$$\rho_t(w,z) \;=\; \frac{\bar{\rho}\exp\left(\frac{\tilde{l}_t(w,z)}{\kappa_\rho \xi_t(w,z)}\right)}{\bar{\rho}\exp\left(\frac{\tilde{l}_t(w,z)}{\kappa_\rho \xi_t(w,z)}\right) + (1-\bar{\rho})\exp\left(\frac{l_t^e(w,z)}{\kappa_\rho \xi_t(w,z)}\right)}. \tag{65}$$

Thus, the decision noise in both the timing choice and the wage-setting choice is proportional to the worker's marginal disutility of labor.

Indeed, to characterize the worker's decision in a given state (w,z,Ξ), it suffices to find the unique value of $\xi_t(w,z)$ that solves (64).[21] The amounts of time devoted to the timing decision and the wage-setting decision are

$$\mu_t^w(w,z) \;= \kappa_\rho \rho_t(w,z) \log\left(\frac{\rho_t(w,z)}{\bar{\rho}}\right) + (1-\rho_t(w,z))\log\left(\frac{1-\rho_t(w,z)}{1-\bar{\rho}}\right) \tag{66}$$

$$\tau_t^w(w,z) \;= \kappa_w \rho_t(w,z) \int \pi^w(\widetilde{w}|w,z) \ln\left(\frac{\pi^w(\widetilde{w}|w,z)}{\eta^w(\tilde{w})}\right) d\tilde{w}. \tag{67}$$

[20] As we showed earlier for the worker's problem, it is possible to rewrite this problem in terms of a single entropy cost term (a convex function) and a linear objective function. Since labor disutility is also convex, a unique well-defined solution exists for the maximization problem involved in a single backwards induction step. MOVE/SUPPRESS FOOTNOTE??

[21] This discussion refers to the calculations required to solve a backwards induction step of the Bellman equation (62). Therefore we regard the future value $\tilde{l}_t^e(w,z)$ as a known function.

These can be calculated using formulas (63) and (65), and their sum is strictly decreasing as a function of ξ.[22] Since marginal disutility increases strictly with total time use (and since $h_t(w,z)$ is predetermined, so it does not depend on ξ), the right-hand side of (64) can be regarded as a strictly decreasing function of ξ. Therefore (64) can be solved by bisection to give a unique solution $\xi_t(w,z) \geq 0$ in any given state (w, z, Ξ_t).

2.6 Distributional dynamics

The distribution of firms' prices and productivities, and likewise that of workers' wages and productivities, evolves over time as firms and workers respond to idiosyncratic and aggregate shocks. We first state the equations governing the dynamics of the distribution across firms.

We continue to use the notation P_{jt} to refer to the nominal price at which firm j produces in period t, prior to adjustment. This may of course differ from its price $\widetilde{P}_{jt}$ at the end of t, when price adjustments are realized. Therefore we will distinguish the beginning-of-period distribution of prices and log productivities, $\Phi_t(P_{jt}, a_{jt})$, from the distribution of prices and log productivities at the end of t, $\widetilde{\Phi}_t(\widetilde{P}_{jt}, a_{jt})$. But instead of tracking nominal prices P_{jt}, it is simpler to focus on log real prices p_{jt}. Therefore, in analogy to the nominal distributions, we define $\Psi_t(p_{jt}, a_{jt})$ as the real distribution at the beginning of t, when production takes place, and $\widetilde{\Psi}_t(\widetilde{p}_{jt}, a_{jt})$ as the real distribution at the end of t. Finally, we also use lower-case letters to represent the joint densities associated with these distributions, which we write as $\phi_t(P_{jt}, a_{jt})$, $\widetilde{\phi}_t(\widetilde{P}_{jt}, a_{jt})$, $\psi_t(p_{jt}, a_{jt})$, and $\widetilde{\psi}_t(\widetilde{p}_{jt}, a_{jt})$, respectively.[23]

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

$$S(a'|a) \ = \ prob(a_{j,t} \leq a'|a_{j,t-1} = a), \tag{68}$$

or in terms of the corresponding density function:

$$s(a'|a) \ = \ \frac{\partial}{\partial a'} S(a'|a). \tag{69}$$

Thus, suppose that the density of nominal prices and log productivities at the end of period $t-1$ is $\tilde{\phi}_{t-1}(\tilde{P}, a)$. This density is then affected by productivity shocks; the density at the beginning of t will therefore be

$$\phi_t(\tilde{P}, a') \ = \ \int s(a'|a) \widetilde{\phi}_{t-1}(\tilde{P}, a) da. \tag{70}$$

[22] Again, the easiest way to prove this is to recognize that the timing choice and the wage-setting choice can be written as a single decision problem. Then $\mu^w + \tau^w$ simply represents the relative entropy of the decision, multiplied by the scalar κ. Relative entropy is a strictly decreasing function of the decision noise ξ (for a proof, see Lemma 1e, Costain 2017). Total decision time is also a decreasing function of ξ in the general case when $\kappa_\rho \neq \kappa_w$; see Prop. 4, Costain 2017. MOVE/SUPPRESS FOOTNOTE??

[23] Our 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.

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

$$\psi_t\left(\widetilde{p} - \log\frac{P_t}{P_{t-1}}, a'\right) \;=\; \int s(a'|a)\widetilde{\psi}_{t-1}(\widetilde{p}, a)da, \tag{71}$$

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

$$\Psi_t(p, a') \;=\; \int^p \int^{a'} \left(\int s(b|a)\widetilde{\psi}_{t-1}\left(q + i_t, a\right)da\right) db\, dq. \tag{72}$$

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 p and log productivity a at the beginning of period t adjusts its price with probability $\lambda\left(\frac{d_t(p,a)}{\kappa_\lambda w_t}\right)$, where

$$d_t(p, a) \;\equiv\; \widetilde{v}_t(a) - v^e_t(p, a). \tag{73}$$

Upon adjustment, its new real log price is distributed according to $\pi_t(\tilde{p}|a)$. Therefore, if the density of firms at the beginning of t is $\psi_t(p, a)$, the density at the end of t is given by

$$\widetilde{\psi}_t(\widetilde{p}, a) \;=\; \left(1 - \lambda\left(\frac{d_t(\widetilde{p}, a)}{\kappa_\lambda w_t}\right)\right)\psi_t(\widetilde{p}, a) + \int \lambda\left(\frac{d_t(p, a)}{\kappa_\lambda w_t}\right)\pi_t(\widetilde{p}|a)\psi_t(p, a)dp. \tag{74}$$

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

$$\widetilde{\Psi}_t(p, a) \;=\; \int^{\widetilde{p}} \int^{a} \widetilde{\psi}_t(q, b)db\, dq. \tag{75}$$

The dynamics of wages and worker productivities is analogous; it suffices to go directly to the real log dynamics, without developing notation for the nominal dynamics. Let $\Psi^w_t(w_{it}, z_{it})$ be the distribution of real log prices and log worker productivities at the beginning of the period, when production takes place, and let $\widetilde{\Psi}^w_t(\widetilde{w}_{it}, z_{it})$ as the corresponding distribution at the end of the period. We write the densities associated with these distributions as $\psi^w_t(w_{it}, z_{it})$ and $\widetilde{\psi}^w_t(\widetilde{w}_{it}, z_{it})$, respectively.

We assume worker productivity is driven by the Markov process S^z:

$$S^w(z'|z) \;=\; prob(z_{i,t+1} \leq z'|z_{i,t} = z), \tag{76}$$

with the following density function:

$$s^z(z'|z) \;=\; \frac{\partial}{\partial z'}S(z'|z). \tag{77}$$

Meanwhile, holding fixed a worker's nominal wage, her real log wage is changed by inflation, from

$\widetilde{w}_{i,t-1}$ at the end of $t-1$, to $w_{i,t} \equiv \widetilde{w}_{i,t-1} - i_t$. Therefore the density of real log wages and log worker productivities at the beginning of t is given by

$$\psi_t^w\left(\widetilde{w} - i_t, z'\right) \;=\; \int s^z(z'|z)\tilde{\psi}_{t-1}^w(\widetilde{w}, z)dz. \tag{78}$$

Hence the corresponding cumulative distribution at the beginning of t is

$$\Psi_t^w(w, z') \;=\; \int^w \int^{z'} \left(\int s^z(b|z)\tilde{\psi}_{t-1}^w\left(q + i_t, z\right)dz\right) db\, dq. \tag{79}$$

Next, a worker with real log wage w and log productivity z at the beginning of period t adjusts her wage with probability $\rho\left(\frac{d_t^w(w,z)}{\kappa_\rho \xi_t(w,z)}\right)$, where

$$d_t^w(w, z) \;\equiv\; \widetilde{l}_t(w, z) - l_t^e(w, z). \tag{80}$$

Upon adjustment, her new real log wage is distributed according to $\pi_t^w(\tilde{w}|w, z)$. Therefore, if the density of workers at the beginning of t is $\psi_t^w(w, z)$, the density at the end of t is given by

$$\widetilde{\psi}_t^w(\widetilde{w}, z) \;=\; \left(1 - \rho\left(\frac{d_t^w(\widetilde{w}, z)}{\kappa_\rho \xi_t(\widetilde{w}, z)}\right)\right)\psi_t^w(\widetilde{w}, z) + \int \rho\left(\frac{d_t^w(w, z)}{\kappa_\rho \xi_t(w, z)}\right)\pi_t^w(\widetilde{w}|w, z)\psi_t^w(w, z)dw. \tag{81}$$

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

$$\widetilde{\Psi}_t^w(\widetilde{w}, z) \;=\; \int^{\tilde{w}} \int^{z} \psi_t(q, b)db\, dq. \tag{82}$$

2.7 Aggregate consistency

When supply equals demand for each good j, total supply and demand of effective labor satisfy

$$N_t - \mu_t - \tau_t = \int_0^1 \frac{C_{jt}}{A_{jt}}dj = C_t \int\int \psi_t(p, a)\exp(-\epsilon p - a)da\, dp \equiv \Delta_t C_t. \tag{83}$$

Here μ_t is total time devoted to deciding whether to adjust prices, and τ_t is total time devoted to choosing which price to set by firms that adjust:

$$\mu_t \;=\; \int\int \psi_t(p, a)\mu_t(p, a)dadp \tag{84}$$

$$\tau_t \;=\; \int\int \psi_t(p, a)\tau_t(p, a)dadp \tag{85}$$

where firm-specific decision times are given by (22)-(23). Equation (83) also defines a measure of price dispersion, $\Delta_t \equiv P_t^\epsilon \int_0^1 P_{jt}^{-\epsilon} A_{jt}^{-1} dj$, 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.

In nominal terms, the price level and wage level are given as follows

$$\int\int P^{1-\epsilon}\phi_t(P,A)dA\,dP \;=\; P(\Omega_t)^{1-\epsilon}. \tag{86}$$

$$\int\int \left(\frac{W}{Z}\right)^{1-\epsilon_N}\phi_t^W(W,Z)dZ\,dW \;=\; W(\Omega_t)^{1-\epsilon_N}. \tag{87}$$

Given (86), the real price level is one by definition:

$$\int\int \exp((1-\epsilon)p)\psi_t(p,a)da\,dp \;=\; 1. \tag{88}$$

The real wage level satisfies

$$\int\int \exp((1-\epsilon_N)(w-z))\psi_t^W(w,z)dz\,dw \;=\; w(\Xi_t)^{1-\epsilon_N}. \tag{89}$$

Since nominal prices are predetermined under the timing we have assumed here, it is natural to conjecture that the nominal state of the economy can be summarized by the following state variable:

$$\Omega_t \;\equiv\; (M_t, g_t, \Phi_t, \Phi_t^w). \tag{90}$$

Since the model is homogeneous of degree one in nominal variables, the corresponding real state variable would be:

$$\Xi_t \;\equiv\; (g_t, \Phi_t, \Phi_t^w). \tag{91}$$

We will show that this is a valid state variable for the economy by constructing an equilibrium in terms of this state.

3 Results

3.1 Parameters

We simulate the model at monthly frequency. Utility from consumption and money holdings, and disutility from labor are parameterized as $u(C) = \frac{1}{1-\gamma}(C^{1-\gamma}-1)$, $v(m) = \nu\log(m)$, and $X(h) = \frac{\chi}{1+\zeta}h^{1+\zeta}$, respectively. Following Golosov and Lucas (2007), we set $\gamma = 2$, $\nu = 1$, and $\chi = 6$. The elasticity of substitution $\epsilon = 7$ across differentiated goods is likewise taken from Golosov and Lucas (2007); we assume the same elasticity across varieties of labor, $\epsilon_N = 7$.

As we discussed in Sec. 2.4, the worker's problem is much simpler to compute when labor disutility is linear ($\zeta = 0$); therefore we will begin by studying the $\zeta = 0$ case before we move on to the nonlinear case in Sec. 3.3.

The productivity processes for firms and workers are assumed to follow discretized approximations of the following AR(1) processes:

$$
\begin{aligned}
a_{jt} &= \rho_a a_{jt-1} + \epsilon_t^a, && (92)\\
z_{it} &= \rho_z z_{it-1} + \epsilon_t^z, && (93)
\end{aligned}
$$

where ϵ_t^a and ϵ_t^z are *i.i.d.* normal shocks with mean zero. Thus the variances of a_{jt} and z_{it} are $\sigma_a^2 = \frac{\sigma_{\epsilon a}^2}{1-\rho_a^2}$ and $\sigma_z^2 = \frac{\sigma_{\epsilon z}^2}{1-\rho_z^2}$, where $\sigma_{\epsilon a}^2$ and $\sigma_{\epsilon z}^2$ are the variances of the innovations ϵ_t^a and ϵ_t^z, respectively.

3.2 Numerical results under linear labor disutility

We now simulate and compare several versions of the model with varying degrees of noise in the pricing and and wage-setting processes, under the assumption of linear labor disutility, $\zeta = 0$. At the micro level, we study how decision costs on each margin affect the frequency and the distribution of price and wage adjustment; at the macro level, our main aim is to identify which noise margin contributes most to the non-neutrality of monetary shocks. To that end, we compare six parameterizations of our model (listed in Table 1), in which we vary the noise parameters κ_π, κ_λ, κ_w, and κ_ρ, holding the remaining parameters fixed.

Table 1: Adjustment parameters.

	V1	V2	V3	V4	V5	V6
$\kappa_\pi = \kappa_\lambda$	0.017	0.0017	0. 00017	0.017	0.017	0.00017
$\kappa_w = \kappa_\rho$	0.017	0.017	0.017	0.0017	0.00017	0.00017

Note: Baseline noise $\kappa_0 \equiv 0.017$ is estimated in CN15 by fitting retail price change data.

The parameterizations described in Table 1 are variations on the benchmark noisy specification V1, in which all four noise parameters are set to $\kappa_\pi = \kappa_\lambda = \kappa_w = \kappa_\rho = \kappa_0 \equiv 0.017$, implying substantial nominal rigidity both for prices and for wages. The benchmark noise level $\kappa_0 = 0.017$ is the estimate of CN15, who found that this value, together with $\bar{\lambda} = 0.2$, gave the best fit to the frequency and distribution of retail price adjustments under the constraint $\kappa_\pi = \kappa_\lambda$. Following CN15, we set both $\bar{\lambda} = 0.2$ and $\bar{\rho} = 0.2$ in all versions V1-V6.

Specifications V2-V6 vary the noise parameters while fixing all remaining parameters. Specifications V2 and V3 reduce price stickiness relative to the benchmark V1, lowering both κ_π and κ_λ first to $\kappa_0/10 = 0.0017$ and then to $\kappa_0/100 = 0.00017$, which effectively means that prices are almost perfectly flexible. Specifications V4 and V5 instead reduce wage stickiness relative to the benchmark V1, lowering both κ_w and κ_ρ first to $\kappa_0/10$ and then to $\kappa_0/100$, which makes wages almost perfectly flexible. Version V6 assumes both margins are very flexible, setting all noise parameters to $\kappa_0/100$.

Table 2 and subsequent results will compare the various parameterizations of our model to microdata on price and wage adjustments. The price data we consider come from the Dominick's supermarket dataset documented by Midrigan (2011).[24] These data represent weekly regular price changes, excluding temporary sales. We aggregate weekly adjustment rates to monthly rates for comparability with some of the other data sources we consider. Our motivation for ignoring sales is that recent literature has found that monetary nonneutrality depends primarily on the frequency of "regular" or "non-sale" price changes (see for example Eichenbaum *et al.*, 2011; Guimaraes and Sheedy, 2011; or Kehoe and Midrigan, 2014).

Our wage change data are from the International Wage Flexibility Project (IWFP). The histograms shown as blue shaded areas in our Fig. 2 are taken from Fig. 2a of Dickens *et al.* (2007), which documents the results of the IWFP. The figure aggregates histograms of price adjustment across multiple countries. While most of the underlying national data are drawn from surveys of firms, they refer to annual nominal wage changes of individual workers. The IWFP focused on annual changes because it observed a widespread tendency for wages to change once a year for many workers in many countries, which in turn means that much of the available survey data addresses annual changes. Clearly this makes our data on wage changes less than perfect for comparison with our price change data, which are at weekly frequency. Nonetheless, to try to get a quantitative benchmark for our theoretical model, we will take the IWFP data at face value. Therefore in Table 2 we report that the monthly frequency of nominal wage adjustment is 1/12=0.083, and we calculate statistics about nominal wage changes directly from the IWFP histogram.

Table 2 compares steady-state statistics on price and wage adjustments as the noise parameters vary. Note first that as we move from the benchmark V1 to the low-noise specification V6, aggregate consumption increases while total labor hours decrease. This is to be expected, as the overall efficiency of the economy increases when there are less frictions. The price adjustment frequency more than quintuples, rising from 10.1% to 54.4% per month. The wage adjustment frequency also rises, but much less sharply, from 6.02% to only 6.95% monthly.

It is also interesting to observe the cross effects of price rigidity on wage adjustment, and *vice versa*. While decreasing wage rigidity tends to increase the price adjustment frequency (it rises from 10.1% in V4 to 10.4% in V5), a decrease in price rigidity may instead *decrease* the frequency of wage changes (which falls from 7.28% in V5 to 6.95% in version V6).

As price rigidity decreases (comparing V1, V2, and V3), the absolute size of price changes falls from 8.57% to 4.76%. Likewise, their standard deviation falls from 10.6% to 5.30%, and their kurtosis decreases from 3.20 to 1.94. Price resetting and timing costs fall as a fraction of revenues, and the losses relative to the fully flexible case are greatly decreased, to less than half a percent of revenues. The effects of decreased wage rigidity are analogous. Comparing V1 with V4 and V5, the absolute size of wage changes falls from 6.14% to 1.98%; their standard deviation falls dramatically, from 8.53% to 1.26%, and their kurtosis likewise falls from 10.0 to 3.62. The costs associated with wage resetting and wage

[24] We are grateful to Virgiliu Midrigan for making his price 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.

reset timing almost vanish in specifications V5 and V6.

Again, there are contrasting cross-effects between prices and wages. Decreased wage rigidity has no observable effect on the absolute size of price adjustments. Decreased price rigidity is instead observed to *increase* the size of wage changes (compare V5 and V6).

The statistics in Table 2 are clarified by observing the steady-state distributions of nonzero price and wage adjustments, plotted as histograms in Figure 2. The predicted distributions from the different versions of the model are shown as black lines; the distributions from microdata are given by the blue shaded areas. The left column of the figure shows how the distribution of nonzero price changes varies as we decrease the noise in the price decision, comparing versions V1, V2, and V3. The variance of the price adjustment distribution in the benchmark version V1 is similar to that of the data, though the model-generated distribution is smoother and less bimodal than the data. As κ_π and κ_λ decrease, the histogram of price changes becomes much more bimodal than the data, displaying the two sharp spikes generated by the menu cost model of Golosov and Lucas (2007). In contrast, the right column of the figure shows that the distribution of nonzero wage changes remains unimodal, becoming ever more concentrated around a single sharp peak as we decrease the noise in the wage decision from version V1 to V4 and V5. The peak of the wage adjustment histogram lies above zero, reflecting the upward aggregate nominal trend in our simulations; likewise, the mean price adjustment is positive.

Thus, as decision noise decreases, price adjustments increasingly resemble the familiar (S, s) behavior associated with a menu cost model. Errors in pricing and timing smooth out the distribution of changes in version V1, but as noise is reduced, the preponderance of price changes occur around two upper and lower thresholds. Very small changes are rare, because it is not worth paying the cost of changing the price when it is already near its target value. Rather than a menu cost, in our model the cost of price adjustment is the decision cost $\kappa_\pi \mathcal{D}(\pi||\eta)$ of choosing which price to set. This cost decreases as κ_π declines; this is why the distance between the two peaks of the price change histogram (the analogues in or model of the (S, s) bands) decreases as we move down the left panels of Fig. 2 from version V1 to V2 and V3. Overall, it appears that our model's fit could be improved by choosing a noise level in between those of V1 and V2.

The reason the wage adjustment pattern differs here is that labor disutility is linear. Note that if decision-making were perfectly costless, labor supply would be infinitely elastic at the wage $w_t = \chi/u'(C_t)$: worker i would respond to a rise in idiosyncratic productivity $z_{i,t}$ by supplying all the additional labor demanded at the initial wage, instead of setting a higher wage. Error-prone choice spreads out wages slightly around this frictionless optimum, as we see in version V1 (top, right panel of Fig. 2). But as the noise in wage adjustment decreases (moving down the right panels of the figure from V1 to V4 and V5), price changes are ever more tightly concentrated around a single peak slightly above zero.

The single peak of the wage adjustment distribution corresponds to small intermittent upward adjustments in response to the nominal trend of the model. Although the worker faces idiosyncratic shocks, it is not optimal to respond to them by adjusting the wage (given linear disutility). Again, this resembles the behavior of a fixed menu cost model, when idiosyncratic shocks are absent or when responding to idiosyncratic shocks is suboptimal. While there are implicitly two "(S,s) bands", the only observed adjustments are the upward bumps that occur when the nominal trend drives the worker's real wage down past the lower threshold. But clearly this would change if utility were sufficiently nonlinear: workers would then prefer to respond to increased idiosyncratic productivity by setting a higher wage. This suggests a possible way forward for matching the wage data. On one hand, nonlinear labor disutility would give the worker an incentive to vary the wage, which would help spread out the distribution. But the data show far more upward adjustments than downward adjustments. While our model does not imply downward rigidity *per se*, it might reproduce this pattern if we parameterize it so that most adjustments represent upward shifts from the "lower threshold". This suggests that we should favor a relatively low degree of nonlinearity in labor disutility. It might also help to impose some demographic turnover in the model, such that workers expect a positive trend of idiosyncratic productivity during their lifetimes, and are eventually replaced by younger, low productivity workers.

The factors discussed above can also be seen at work in Figure 3, which illustrates the logit policy functions from the benchmark case V1. The left panels display the logit probabilities of each price (wage), conditional on cost, while the right panels show the adjustment probabilities conditional on the current price (wage) and cost. Each function is graphed in two ways, for greater clarity: as a surface plot (first and third rows) and as multiple overlaid cross-sections (second and fourth rows). The upper left panel of the graph shows a surface plot of the logit probabilities $\pi(p|a)$ as a function of the firm's cost shock $-a$ and its possible prices p. Just below this, in the second row, we see the smooth, bell-shaped probability distributions $\pi(p|a)$ corresponding to each possible productivity level a. If the firm's cost shock is high (*i.e.* a is low, shown in red in the graph) then its chosen probability distribution shifts towards higher prices. Looking to the right column, we see that the adjustment probability $\lambda(p, a)$ approaches zero for any p that is near the modal value of $\pi(p|a)$.

The bottom panels of Fig. 3 are analogous, but instead show the worker's policy functions $\pi^w(w|z)$ and $\rho(w, z)$. Notice that the worker's logit probabilities $\pi^w(w|z)$ are concentrated around *the same* w, regardless of z (see the bottom left panel). Regardless of her productivity shock, the worker prefers the same real wage, which explains the tight unimodal distribution of wage changes seen earlier in Fig. 2.

We now turn to the macroeconomic implications of the model, comparing impulse responses to money shocks across versions V1, V3, V5, and V6 in Figure 4. The figure shows the impulse responses to a 1% money growth shock, with monthly autocorrelation 0.8, on consumption, labor, price and wage inflation, and the real wage. In the benchmark specification V1 (black with circles), consumption and labor rise by more than 2% on impact, then revert smoothly and gradually with a half-life of roughly six months. Price inflation and wage inflation both rise persistently to a rate of roughly 0.5%. Wage inflation is slightly higher than price inflation, causing the real wage to peak at roughly 0.4% above steady state after four months.

In contrast, in the flexible specification V6 (green), both price and wage inflation spike on impact, with an 4% jump in prices and wages. Consumption and labor increase by half a percent in the period of impact only, then return to their steady state levels. In summary, current and expected money growth feeds rapidly into prices, and its real impact is small and transitory.

It is particularly interesting to compare the results of specifications V3 (red, with sticky wages but very flexible prices) and V5 (blue, with sticky prices but very flexible wages). The key takeaway is seen in the response of consumption – version V3, with wage stickiness only, comes very close to the baseline model V1 with both price and wage stickiness. The reason is that wage stickiness keeps firms' marginal costs from adjusting rapidly, so even though prices are much more flexible in version V3 than V1, the impulse response of inflation is very similar in both cases. Both wages and prices adjust gradually in version V3, leading to a real impact on consumption and output that is almost as large and persistent as we saw in version V1.

This conclusion is reinforced by considering version V5. With sticky prices and flexible wages, specification V5 implies an immediate burst of wage inflation when the money supply shock hits– prices rise 3% on impact. In spite of price stickiness, this rise in nominal marginal costs also causes prices to increase by 1.3% on impact, more than they do in cases V1 and V3. Even so, the overall effect is a large increase in real wages, which discourages labor demand and thus greatly reduces the persistence of the real effects of the money shock.

In other words, wage stickiness appears to be a more important source of money non-neutrality than price stickiness alone. Nonetheless, the behavior of real wages suggests that there is a role for both forms of stickiness in explaining the data. While version V5 shows a large increase in real wages after a money shock, V3 instead displays a small *decrease* in real wages, as price inflation jumps slightly more than wage inflation. In this sense the benchmark specification V1 appears most consistent the preponderance of empirical evidence, which suggests that there is a small increase in the real wage after a money shock (see for example Christiano *et al.*, 2005; McCallum and Smets, 2006; or Olivei and Tenreyro, 2007).

Finally, to isolate the effects of state-dependence in prices and wages, we compare the impulse responses of Fig. 4 to those derived from an otherwise identical economy with Calvo pricing. Thus, we hold fixed the theoretical structure of the economy (the model, the parameters, and the finite grid approximation) except that we assume firms' price adjustments and workers' wage adjustments are governed by the Calvo (1983) mechanism. Thus, firms (workers) reset their prices (wages) with a constant, exogenously-fixed probability per month, and the new price (wage) is optimally chosen (it is optimal taking into account the fact that future adjustments will take place at random times in the future). For comparability with the simulations reported previously, we define the Calvo adjustment hazards for versions V1C-V6C equal to the steady-state adjustment hazards derived from the previous state-dependent specifications V1-V6.

The results of the Calvo simulations are displayed in Figure 5. Three findings stand out. First, the Calvo model generates much more persistence than the state-dependent adjustment models seen in Fig. 4; the half-life of the impulse responses of consumption and labor rises to roughly 15 months (note that the time horizon shown on the horizontal axis of Fig. 5 is twice as long as that shown in Fig. 4). Thus, the "intensive", "extensive", and "selection" effects present in our state-dependent model jointly imply a more flexible aggregate price level than does the Calvo model, which incorporates the "intensive" margin only. Second, the impulse responses of the Calvo specifications V1C, V3C, V5C, and V6C do not differ very much from one another (except in their implications for real wages). This may reflect the fact that the frequency of wage adjustment was very similar in our state-contingent simulations V1-V6, and therefore the assumed Calvo wage adjustment hazards does not differ much across our Calvo simulations V1C-V6C. Finally, the qualitative behavior of (all) the Calvo specifications is similar to that of our benchmark state-dependent model V1: consumption, labor, price inflation, and wage inflation all jump on impact following a money supply shock, and then smoothly revert to their means. The real wage rises, but by much less than it does in state-dependent version V5. Thus, the dynamic predictions of the Calvo specification appear quite consistent with those of a state-dependent model, as long as we bear in mind that the Calvo setup exaggerates the persistence of a nominal shock.

Table 2: **Evaluating the linear disutility model with different values of κ_π, κ_λ, κ_w and κ_ρ.**

	Data		Sticky		Decreasing price stickiness				Decreasing wage stickiness				Flexible	
			V1		V2		V3		V4		V5		V6	
Consumption			0.3496		0.3508		0.3514		0.3501		0.3489		0.3522	
Labor			0.3530		0.3493		0.3481		0.3535		0.3523		0.3488	
Wage			0.8576		0.8638		0.8666		0.8576		0.8576		0.8666	
	Prices	Wages	Prices	Wages	Prices	Wages	Prices	Wages	Prices	Wages	Prices	Wages	Prices	Wages
Frequency of adj.	10.2	8.3	10.1	6.02	22.5	6.03	54.4	6.04	10.1	6.41	10.41	7.28	54.4	6.95
Mean change	1.60	5.10	1.68	2.83	0.76	2.82	0.31	2.82	1.68	2.66	1.68	2.34	0.31	2.45
Std of prices/wages			5.26	3.07	5.63	3.08	6.06	3.08	5.26	1.21	5.26	0.75	6.06	0.96
Mean abs(Change)	9.90	6.47	8.57	6.14	6.80	6.16	4.76	6.16	8.57	2.70	8.57	1.98	4.76	2.29
Std of changes	13.2	6.52	10.6	8.53	7.50	8.53	5.30	8.52	10.6	2.10	10.6	1.26	5.26	1.79
Kurtosis of changes	4.81	4.39	3.20	10.0	1.82	9.77	1.94	9.66	3.20	5.21	3.20	3.62	1.94	3.33
Pct increases	65.1	86.5	58.5	72.1	55.5	72.0	54.0	71.9	58.5	92.2	58.5	99.1	54.0	92.0
Changes $\leq$5%	35.5	43.0	28.8	48.1	27.0	47.9	60.7	47.9	28.8	91.7	28.8	100	60.8	99.9
Changes $\leq$2.5%	12.0	11.8	14.2	24.7	8.90	24.6	19.9	24.6	14.2	57.4	14.2	81.1	20.0	61.6
Resetting cost as % Rev.			0.51	0.13	0.19	0.14	0.07	0.14	0.51	0.03	0.51	0.004	0.07	0.004
Timing cost as % Rev.			0.37	0.14	0.11	0.15	0.03	0.15	0.37	0.02	0.37	0.004	0.03	0.003
Loss. relative to Flex			1.78	1.62	0.56	1.73	0.13	1.71	1.78	1.14	1.78	1.18	0.13	1.17

Figure 2: Distribution of nonzero price and wage changes: varying stickiness ($\zeta = 0$).

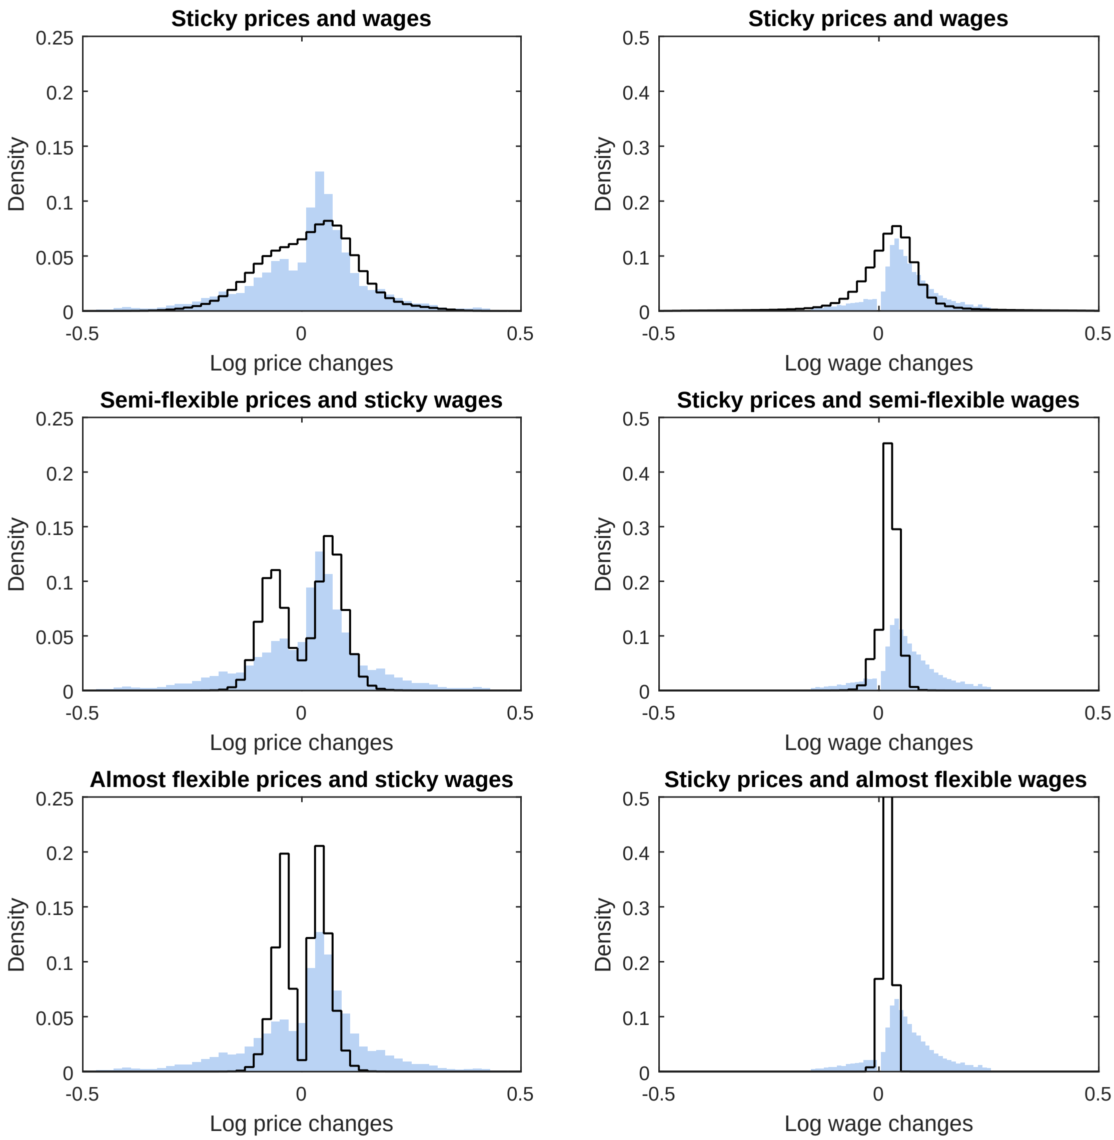

Notes: Black lines: model versions with linear labor disutility ($\zeta = 0$). Blue shaded areas: Data.
Left column: Effect of decreasing price stickiness (versions V1, V2, V3) on distribution of nonzero price adjustments.
Right column: Effect of decreasing wage stickiness (versions V1, V4, V5) on distribution of nonzero wage adjustments.

Figure 3: Adjustment behavior. Benchmark model (V1) with sticky prices and sticky wages ($\zeta = 0$).

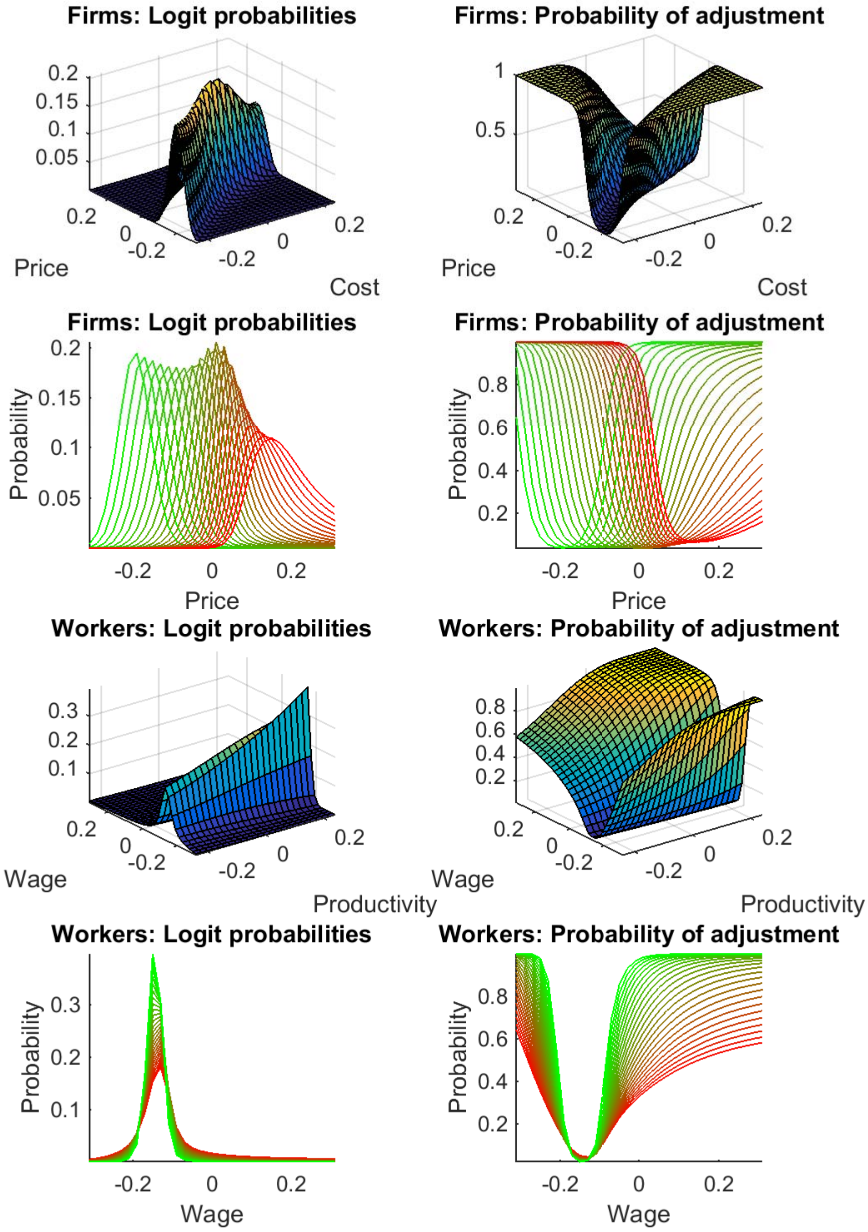

Notes: Distribution of adjustments and adjustment probability for prices (top four panels) and wages (bottom four panels) under linear labor disutility ($\zeta = 0$).

Left panels: 3d plots show price (wage) choice probabilities, conditional on cost (productivity).

Left panels: 2d plots show price (wage) choice probabilities, conditional on each possible cost (productivity).

Right panels: 3d plots show adjustment probabilities, conditional on current price (wage) and cost (productivity).

Right panels: 2d plots show adjustment probabilities, conditional on each possible cost (productivity).

Colors in 2d plots: For firms, green represents low cost (high a). For workers, green represents high productivity (high z).

Figure 4: Money growth shock: effects of nominal rigidity. Error-prone adjustment, $\zeta = 0$.

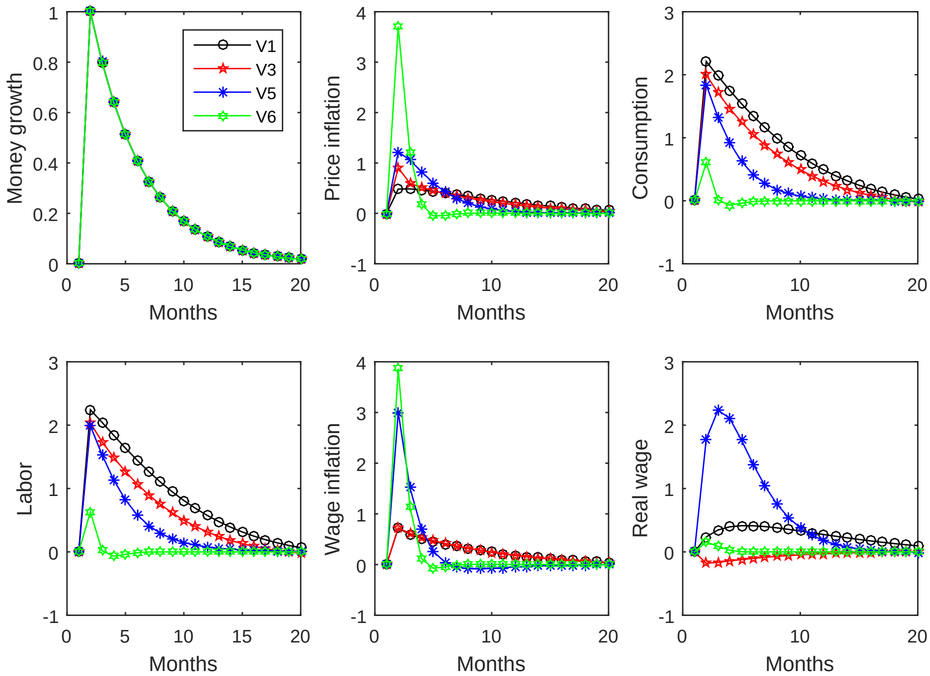

Notes:
Impulse responses of inflation and consumption to money growth shock with autocorrelation 0.8 (monthly), under linear labor disutility ($\zeta = 0$).
Black: Benchmark (V1), both prices and wages sticky. Red: V3, flexible prices and sticky wages.
Blue: V5, sticky prices and flexible wages. Green: V6: both prices and wages flexible.

3.3 Results: Nonlinear disutility

As we have seen above, generating a nontrivial wage distribution requires nonlinear disutility of labor. Therefore, we now compute a nonlinear specification, setting $X(h) = \frac{\chi}{1+\zeta}h^{1+\zeta}$, with $\zeta = 0.5$. Table 2 reports steady-state statistics for this parameterization, considering all the combinations of noise parameters defined in Table 1; we denote these versions of the model as V1N-V6N. As in the linear case, decreased noise in price setting or wage setting makes adjustment more frequent. As before, in the case of prices, lower noise implies smaller absolute price changes, a lower standard deviation and kurtosis of price changes, and more of the smallest changes (less than 5% or less than 2.5%). However, lower noise in the wage setting process leads to *larger* absolute wage changes and a higher standard deviation of wage changes, and makes the smallest adjustments less frequent.

The reason for this contrasting behavior of prices and wages is clear when we inspect Figure 6.

Figure 5: Money growth shock: effects of nominal rigidity. Calvo adjustment, $\zeta = 0$.

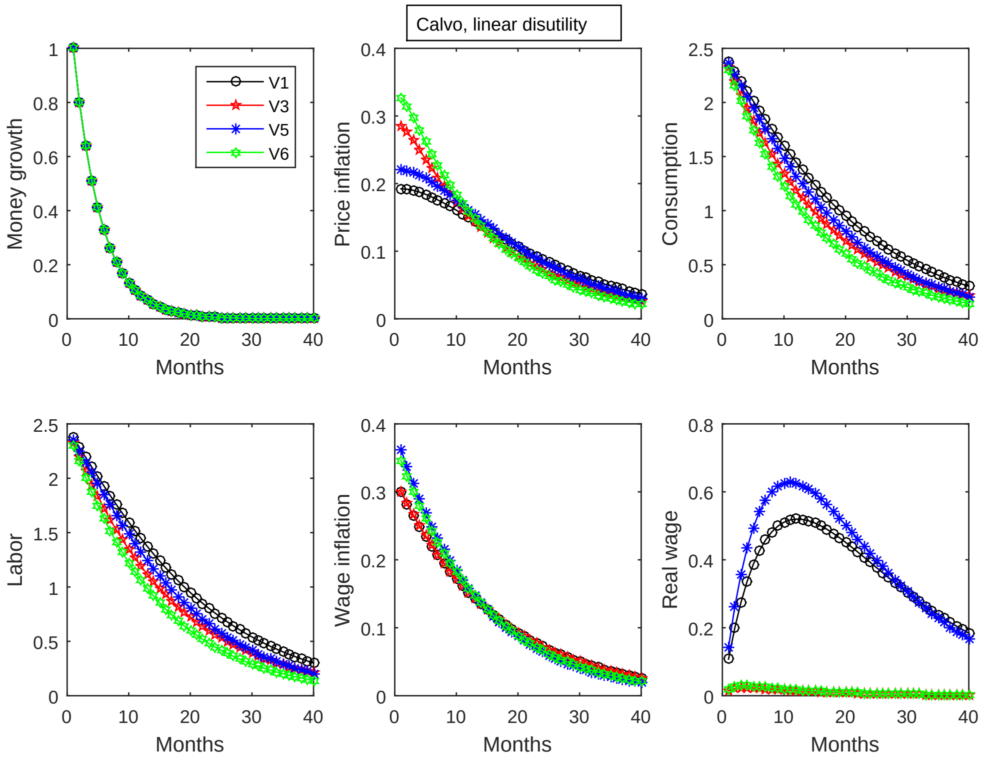

Notes:

Impulse responses of inflation and consumption to money growth shock with autocorrelation 0.8 (monthly), under Calvo adjustment with linear labor disutility ($\zeta = 0$).

Black: Benchmark (V1C), both prices and wages sticky. Red: V3C, flexible prices and sticky wages.

Blue: V5C, sticky prices and flexible wages. Green: V6C: both prices and wages flexible.

Firms in version V1N often make large price adjustments as they attempt to respond to their idiosyncratic productivity shocks. As noise decreases (versions V2N and V3N), this tracking becomes more accurate, increasingly resembling an (S, s) model in which the firm makes relatively frequent small upward or downward adjustments to track productivity. But workers in version V1N make relatively little attempt to track productivity; the histogram of wage adjustments is dominated by a single sharp spike representing small upward adjustments that track the trend of the nominal money supply. Decision-making is simply too noisy under parameterization V1N to justify changes in response to (transitory) idiosyncratic shocks. But as the noise in the wage-setting process decreases (versions V4N and V5N) it becomes worthwhile to respond to individual productivity. Therefore we begin to see the two (S, s) bands, as workers begin to pay the decision costs of tracking the upward and downward movements in their productivity processes. Wage adjustments are therefore not only more frequent, but also larger and more variable, in V4N and V5N compared with V1N.

These observations are also consistent with the behavior of aggregate consumption and labor across specifications. Decreasing pricing noise from version V1N to V2N and V3N, consumption rises and aggregate labor falls, consistent with a wealth effect from reduced frictions. Decreasing wage setting noise from version V1N to versions V4N and V5N, consumption rises sharply and aggregate labor rises too, consistent with a greater average incentive to work when wages are more efficiently conditioned on productivity.

Figure 7 depicts several steady-state objects for the nonlinear benchmark V1N: the logit probabilities governing price resets and wage resets (left panels) and firms' and workers' adjustment hazards (right panels). In each case the probabilities are shown as functions of the price-cost (resp. wage-productivity) pairs. As in Fig. 3, firms prefer higher prices when costs are higher, and the probability of adjustment rises smoothly as firms deviate from the prices they prefer (conditional on costs). But now, in contrast with Fig. 3, we see also that workers tend to set higher wages when their productivity rises. Wage variation remains somewhat limited, but the preferred wage now varies by roughly $\pm 8\%$ as worker productivity varies between its maximum and minimum values in the Tauchen (1986) grid approximation, which differ by $\pm 24\%$.[25]

Finally, we return to the issue of monetary non-neutrality. Figure 8 shows the effects of an autocorrelated money growth shock with monthly persistence equal to 0.8. The figure depicts the responses of price and wage inflation, consumption, hours and the real wage for models V1N, V3N, V5N and V6N. As before, the sticky-price, sticky-wage specification implies substantial real effects, with consumption and labor rising 2.5% on impact, with a half-life of seven or eight months. The version with reduced wage stickiness (V5N) has similar real effects on impact, but much lower persistence, because it implies a large and persistent increase in real wages that offsets firms' incentive to demand more labor. As expected the model with the smallest real impact is V6, which has very low persistence, as in the Golosov-Lucas (2007) menu cost model.

But surprisingly, the version with the biggest real effects is now V3N, with sticky wages but *flexible* prices.[26] The consumption and hours responses are similar on impact to the other versions, but persistence is greatly magnified, reaching a half-life of almost fifteen months.

This finding is particularly striking if we compare what happens in a Calvo model calibrated to the same adjustment frequencies as models V1N-V6N; we denote these Calvo specifications as V1CN-V6CN. Figure 9 shows that the parameterization that delivers largest real effects from the nominal shock is then V1CN: the one where both wages and prices are sticky. This Calvo specification implies a half-life of roughly 20 months for consumption and labor. This suggests that there is something special

[25] While the firms' logit probabilities in the second row, first column of Figure 7 are vary smoothly with p and a, workers' logit probabilities (fourth row, first column) are rather spiky. This suggests that our simulations are not sensitive to the finite grid approximation of the pricing decision, but obtaining a robust approximation of the wage setting decision may require a finer wage grid.

[26] A variety of checks demonstrate that this is a robust result. In particular, our impulse responses for version V2N (not shown) are intermediate between V1N and V3N. Version V2N, with sticky wages and a degree of price flexibility intermediate between V1N and V3N, implies *larger* and more persistent real effects than V1N (but not so large and persistent as V3N). Also, when we recompute scenarios V1N-V6N on a wider grid, our findings are unchanged.

about state-dependence (e.g. the "selection effect") that reverses the effects of price stickiness when wages are sticky. Note that both wage inflation and price inflation are greatly reduced in the impulse responses for the state-dependent specifications V3N compared with V1N. This is not the case when we compare the Calvo specifications V3CN and V1CN. One possible explanation is that the selection effect may strengthen the strategic complementarity between wage dynamics and price dynamics in a state-contingent model, compared with a Calvo model.

The key takeaway after these findings remains unchanged: under state-dependent pricing, stickiness in wage setting appears to be more important for the non-neutrality of money shocks than stickiness of price setting. This reinforces our finding in the previous section for the case of linear labor disutility. But in contrast to our results in Sec. 3.2, and in contrast to the implications of the Calvo framework, it now appears that price flexibility may reinforce, rather than weaken, wage stickiness.

Table 3: Evaluating the nonlinear disutility model with different values of κ_π, κ_λ, κ_w and κ_ρ.

	Data		Sticky		Decreasing price stickiness				Decreasing wage stickiness				Flexible	
			V1N		V2N		V3N		V4N		V5N		V6N	
Consumption			0.4113		0.4117		0.4122		0.4265		0.4282		0.4308	
Labor			0.4155		0.4112		0.4102		0.4307		0.4324		0.4286	
Wage			0.8559		0.8603		0.8621		0.8560		0.8560		0.8621	
	Prices	Wages	Prices	Wages	Prices	Wages	Prices	Wages	Prices	Wages	Prices	Wages	Prices	Wages
Frequency of adj.	10.2	8.3	7.74	7.44	18.6	7.83	49.6	8.34	7.78	23.5	7.78	22.1	50.4	22.0
Mean change	1.60	5.10	2.20	2.29	0.92	2.18	0.34	2.04	2.19	0.72	2.19	0.77	0.34	0.77
Std			4.13	1.79	4.41	1.73	4.84	1.62	4.10	2.84	4.10	3.11	4.84	3.11
Mean abs change	9.90	6.47	7.30	3.26	5.77	3.07	4.02	2.85	7.25	3.45	7.24	3.85	3.99	3.85
Std of changes	13.2	6.52	9.10	3.20	6.31	3.09	4.42	2.86	9.01	3.67	8.99	4.10	4.39	4.10
Kurtosis of changes	4.81	4.39	4.25	6.77	1.82	7.66	1.89	8.71	4.22	1.76	4.22	1.74	1.90	1.74
Pct increases	65.1	86.5	63.4	83.8	57.5	87.4	53.9	87.9	63.4	60.0	63.4	60.0	53.9	60.0
Changes $\leq$5%	35.5	43.0	35.2	83.4	37.3	83.8	74.3	86.8	35.3	88.2	35.3	79.3	74.9	79.2
Changes $\leq$2.5%	12.0	11.8	17.2	45.8	11.8	57.3	30.4	65.4	17.2	36.8	17.2	26.6	31.3	26.6
Resetting cost			0.49	0.96	0.18	1.09	0.07	1.18	0.47	0.29	0.47	0.03	0.06	0.03
Timing cost			0.41	0.67	0.10	0.72	0.03	0.77	0.40	0.14	0.40	0.01	0.03	0.01
Loss. relative to Flex			1.87	4.07	0.85	4.36	0.51	4.51	1.84	1.45	1.83	1.13	0.51	1.13

Figure 6: Distribution of nonzero price and wage changes: varying stickiness ($\zeta = 0.5$).

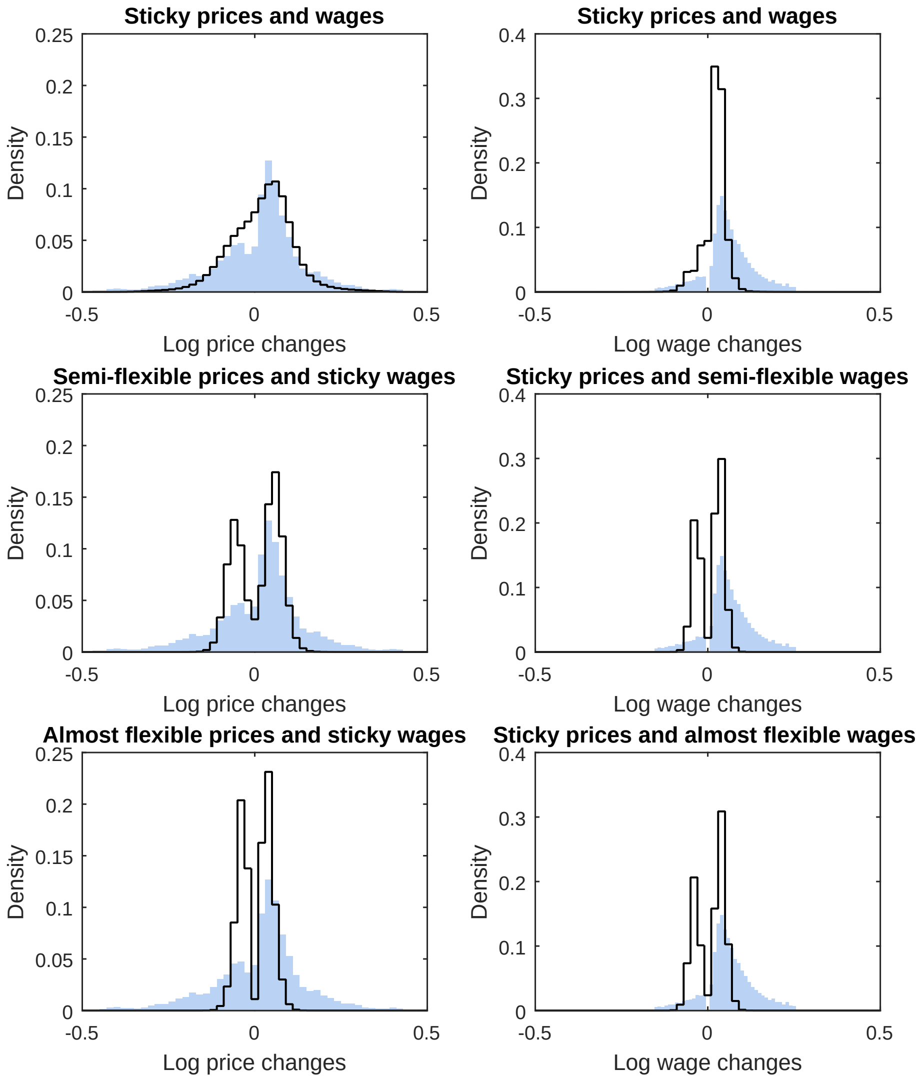

Notes: Black lines: model versions with nonlinear labor disutility ($\zeta = 0.5$). Blue shaded areas: Data.
Left column: Effect of decreasing price stickiness (versions V1N, V2N, V3N) on distribution of nonzero price adjustments.
Right column: Effect of decreasing wage stickiness (versions V1N, V4N, V5N) on distribution of nonzero wage adjustments.

Figure 7: Adjustment behavior. Benchmark model (V1N) with sticky prices and sticky wages ($\zeta = 0.5$).

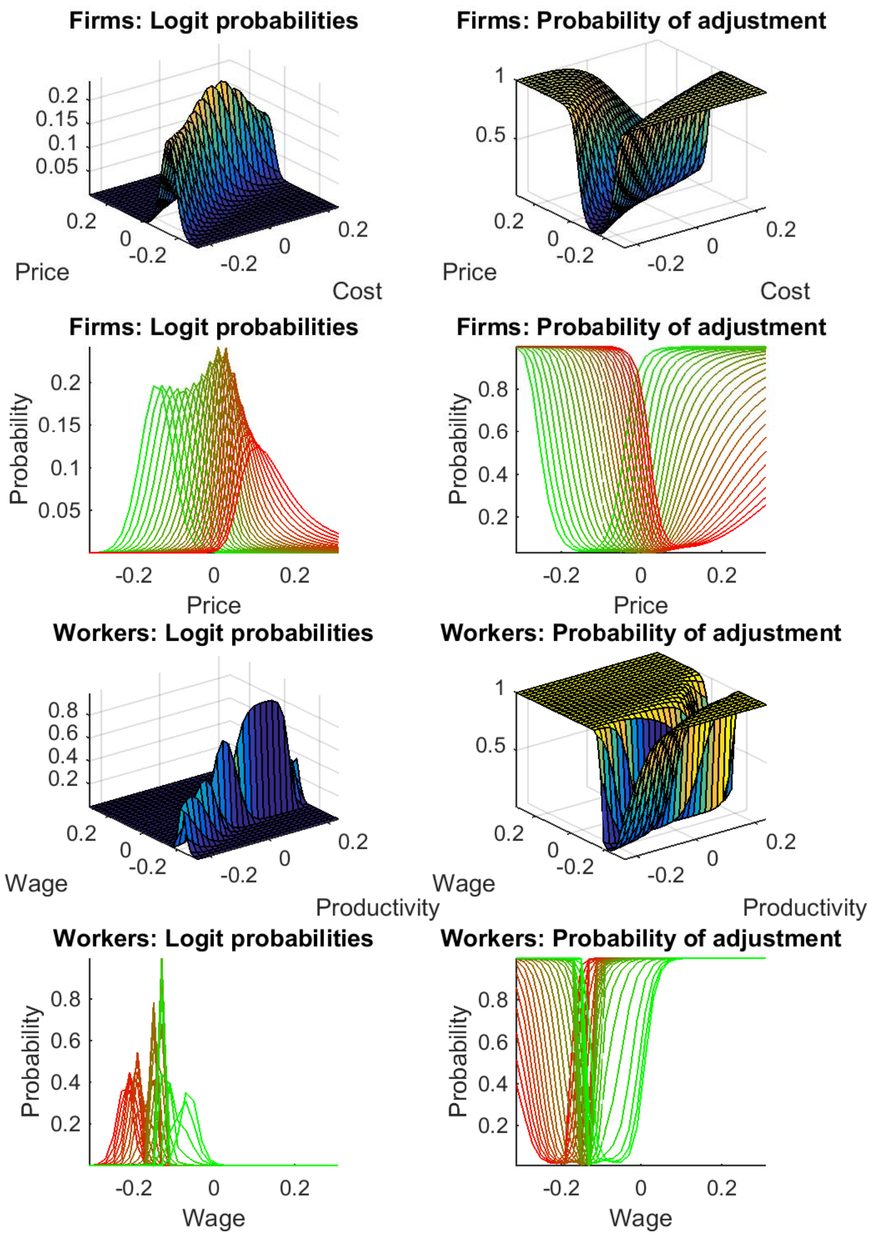

Notes: Distribution of adjustments and adjustment probability for prices (top four panels) and wages (bottom four panels) under nonlinear labor disutility ($\zeta = 0.5$).

Left panels: 3d plots show price (wage) choice probabilities, conditional on cost (productivity).

Left panels: 2d plots show price (wage) choice probabilities, conditional on each possible cost (productivity).

Right panels: 3d plots show adjustment probabilities, conditional on current price (wage) and cost (productivity).

Right panels: 2d plots show adjustment probabilities, conditional on each possible cost (productivity).

Colors in 2d plots: For firms, green represents low cost (high a). For workers, green represents high productivity (high z).

Figure 8: Money growth shock: effects of nominal rigidity. Error-prone pricing, $\zeta = 0.5$.

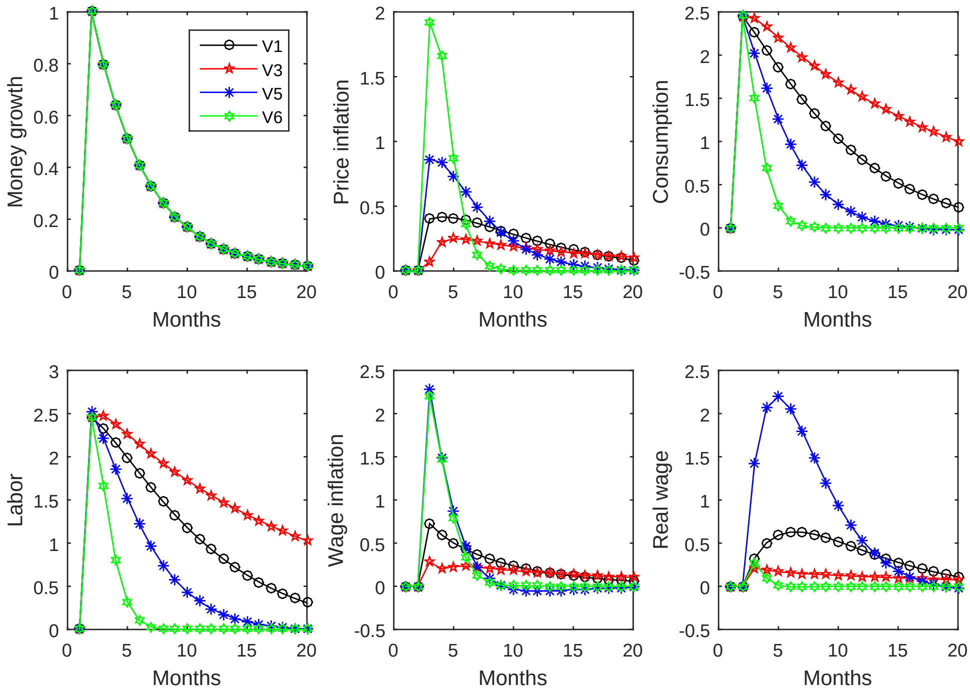

Notes:

Impulse responses of inflation and consumption to money growth shock with autocorrelation 0.8 (monthly), under nonlinear labor disutility ($\zeta = 0.5$).

Black: Benchmark (V1N), both prices and wages sticky. Red: V3N, flexible prices and sticky wages.

Blue: V5N, sticky prices and flexible wages. Green: V6N: both prices and wages flexible.

Figure 9: Money growth shock: effects of nominal rigidity. Calvo pricing, $\zeta = 0.5$.

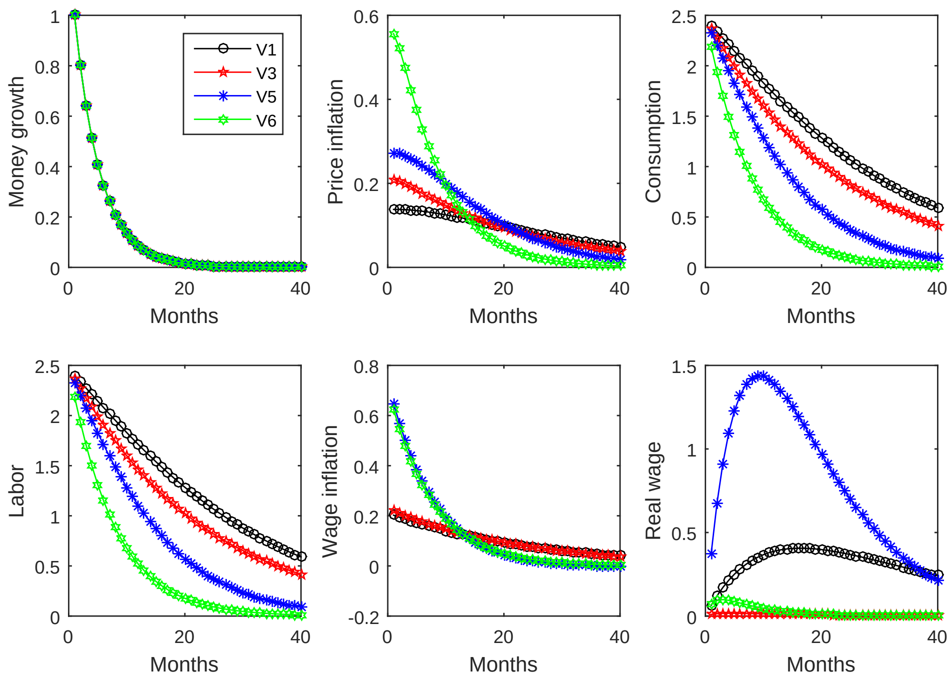

Notes:

Impulse responses of inflation and consumption to money growth shock with autocorrelation 0.8 (monthly), under Calvo adjustment with nonlinear labor disutility ($\zeta = 0.5$).

Black: Benchmark (V1CN), both prices and wages sticky. Red: V3CN, flexible prices and sticky wages.

Blue: V5CN, sticky prices and flexible wages. Green: V6CN: both prices and wages flexible.

4 Conclusions

Outline

- We study a DSGE model with state-dependent prices and state-dependent wages
- Combines monopolistic competition in goods and labor inputs, following Erceg, Henderson, and Levin (1999), with nominal rigidity derived from costly decision-making, following Costain and Nakov (2015)
- First paper to study state dependence in prices and wages in a model with idiosyncratic shocks, for comparison to microdata (Takahashi 2017 studied a model with representative worker and representative firm instead)
- Labor can be costlessly reallocated across firms at any time; our focus is the interaction of nominal price stickiness and nominal wage stickiness, abstracting from matching frictions or other forms of labor specificity

Micro findings

- We compare different calibrations to see the role of price stickiness and wage stickiness for micro and macro results
- Assuming linear labor disutility makes the model much easier to solve, but implies that the wage never varies in response to individual productivity shocks; allowing for nonlinear disutility improves the fit to wage data

Macro findings

- We find that wage stickiness is more likely to cause persistent effects of monetary shocks than price stickiness is
- Huang and Liu (2002) reported the same finding for a time-dependent model; we are the first to study this issue in a state-dependent model
- With linear labor disutility, or with the Calvo model, we find that version with wage stickiness only has almost as much non-neutrality as version with wage and price stickiness together
- With nonlinear labor disutility, we find that *decreasing* price stickiness, in the presence of sufficient wage stickiness, *increases* persistence of real effects of money shocks

Next steps?

- Computational challenge prevents us from estimating the model; exploring robustness is slow
- Now considering possible computational improvements in hopes of exploring model further

References

[1] Álvarez, Fernando; Francesco Lippi, and Luigi Paciello (2014), "Monetary shocks in models with inattentive producers." NBER Working Paper 20817.

[2] 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.

[3] Avouyi-Dovi, Sanvi; Denis Fougère; and Erwan Gautier (2013), "Wage rigidity, collective bargaining, and the minimum wage: Evidence from French agreement data." *Review of Economics and Statistics* 95 (4): 1337-1351.

[4] Barattieri, Alessandro; Susanto Basu; and Peter Gottschalk (2014), "Some evidence on the importance of sticky wages." *American Economic Journal: Macroeconomics* 6 (1): 70-101.

[5] Barro, Robert (1972), "A theory of monopolistic price adjustment." *Review of Economic Studies* 39 (1): 17-26.

[6] Blundell, Richard, and Stephen Bond (2000), "GMM Estimation with persistent panel data: an application to production functions." *Econometric Reviews* 19 (3): 321-40.

[7] Calvo, Guillermo (1983), "Staggered prices in a utility-maximizing framework." *Journal of Monetary Economics* 12: 383-98.

[8] Caplin, Andrew, and Daniel Spulber (1987), "Menu costs and the neutrality of money." *Quarterly Journal of Economics* 102 (4): 703-25.

[9] 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.

[10] Christiano, Lawrence; Martin Eichenbaum, and Charles Evans (2005), "Nominal rigidities and the dynamic effects of a shock to monetary policy." *Journal of Political Economy*, 113 (1), 1-45.

[11] Costain, James (2017), "Costly decisions and sequential bargaining. " Banco de España Working Paper 1729.

[12] Costain, James, and Anton Nakov (2011), "Distributional dynamics under smoothly state-dependent pricing." *Journal of Monetary Economics* 58 (6-8): 646-65.

[13] Costain, James, and Anton Nakov (2015), "Logit price dynamics." CEPR Discussion Paper #10731.

[14] Cover, Thomas, and Joy Thomas (2006), *Elements of Information Theory*, 2nd ed. Wiley Interscience.

[15] Dickens, William; Lorenz Goette; Erica Groshen; Steinar Holden; Julian Messina; Mark Schweitzer; Jarkko Turunen; and Melanie Ward (2007), "How wages change: Micro evidence from the International Wage Flexibility Project." *Journal of Economic Perspectives* 21 (2): 195-214.

[16] Dotsey, Michael; Robert King; and Alexander Wolman (2013), "Inflation and real activity with firm-level productivity shocks." FRB Philadelphia Working Paper 13-35.

[17] Eichenbaum, Martin; Nir Jaimovich, and Sergio Rebelo (2011), "Reference prices, costs and nominal rigidities." *American Economic Review* 101 (1): 234-62.

[18] Erceg, Christopher, Dale Henderson, and Andrew Levin (2000), "Optimal monetary policy with staggered wage and price contracts." *Journal of Monetary Economics* 46: 281-313.

[19] Gagnon, Etienne (2009), "Price setting under low and high inflation: evidence from Mexico." *Quarterly Journal of Economics* 124: 1221-63.

[20] Golosov, Mikhail, and Robert E. Lucas, Jr. (2007), "Menu costs and Phillips curves." *Journal of Political Economy* 115 (2): 171-99.

[21] Huang, Kevin, and Zheng Liu (2002), "Staggered price-setting, staggered wage-setting, and business cycle persistence." *Journal of Monetary Economics* 49, pp. 405-433.

[22] Huang, Kevin; Zheng Liu; and Louis Phaneuf (2004), "Why does the cyclical behavior of real wages change over time?" *American Economic Review* 94 (4), pp. 836-856.

[23] Kehoe, Patrick, and Virgiliu Midrigan (2012), "Prices are sticky after all." FRB Minneapolis Staff Report 413.

[24] 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.

[25] 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.

[26] 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.

[27] Le Bihan, Hervé; Jérémi Montornès; and Thomas Heckel (2012), "Sticky wages: Evidence from quarterly microeconomic data." *American Economic Journal: Macroeconomics* 4 (3): 1-32.

[28] Matejka, Filip, and Alisdair McKay (2015), "Rational inattention to discrete choices: a new foundation for the multinomial logit model." *American Economic Review* 105 (1): 272-298.

[29] Matejka, Filip (2010), "Rationally inattentive seller: sales and discrete pricing." CERGE-EI Working Paper 408.

[30] Mattsson, Lars-Goram, and Jorgen Weibull (2002), "Probabilistic choice and procedurally bounded rationality." *Games and Economic Behavior* 41: 61-78.

[31] McCallum, Andrew, and Frank Smets (2007), "Real wages and monetary policy transmission in the euro area." Kiel Working Paper 1360.

[32] Midrigan, Virgiliu (2011), "Menu Costs, Multi-Product Firms and Aggregate Fluctuations." *Econometrica* 79 (4): 1139-80.

[33] Nakamura, Emi, and Jón Steinsson (2008), "Five facts about prices: a reevaluation of menu cost models." *Quarterly Journal of Economics* 123 (4): 1415-64.

[34] Olivei, Giovanni, and Silvana Tenreyro (2007), "The timing of monetary policy shocks." *American Economic Review* 97 (3): 636-663.

[35] Reiter, Michael (2009), "Solving heterogeneous-agent models by projection and perturbation." *Journal of Economic Dynamics and Control* 33 (3): 649-65.

[36] Rotemberg, Julio (1982), "Monopolistic price adjustment and aggregate output." *Review of Economic Studies* 49, pp. 517-31.

[37] Sigurdsson, Jósef, and Rannveig Sigardurdottir (2016), "Time-dependent or state-dependent wage-setting? Evidence from periods of macroeconomic instability." *Journal of Monetary Economics* 78: 50-66.

[38] Sims, Christopher (2003), "Implications of rational inattention." *Journal of Monetary Economics* 50: 665-90.

[39] Stahl, Dale (1990), "Entropy control costs and entropic equilibrium. " *International Journal of Game Theory* 19: 129-38.

[40] Takahashi, Shuhei (2017), "State dependency in price and wage setting." *International Journal of Central Banking* 13 (1), pp. 151-189.

[41] Tauchen, George (1986), "Finite state Markov chain approximations to univariate and vector autoregressions." *Economics Letters* 20: 177-181.

[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.

Appendix A: Supplementary figures

Figure 10: Additional equilibrium objects. Benchmark model (V1) with sticky prices and sticky wages ($\zeta = 0.5$).

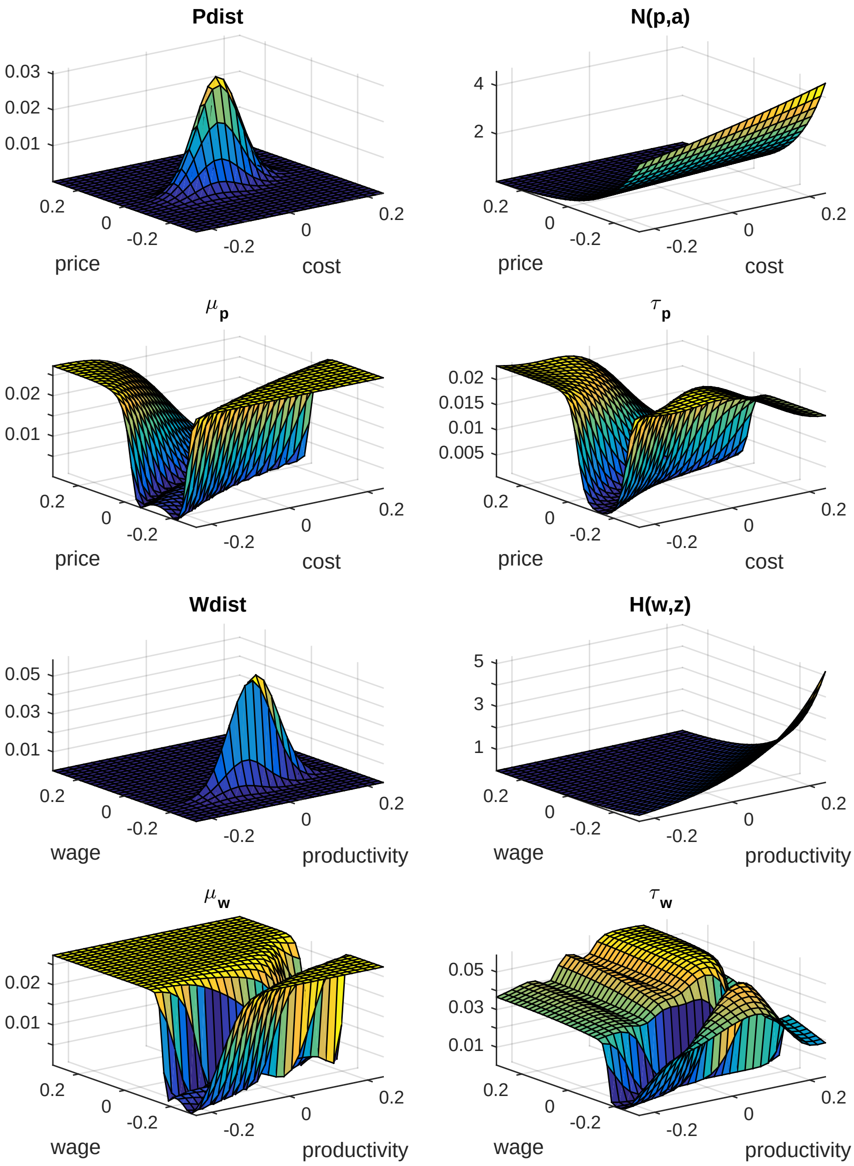

Notes: Additional equilibrium objects, assuming nonlinear labor disutility ($\zeta = 0.5$).
Top row: Steady-state distribution of prices and productivity $\Psi(p, a)$ and firms' demand for effective labor $N(p, a)$.
Second row: Firms' time devoted to decisions: monitoring $\mu(p, a)$ and price setting $\tau(p, a)$.
Third row: Steady-state distribution of wages and labor productivity $\Psi^w(w, z)$ and demand for worker's time $H(w, z)$.
Bottom row: Workers' time devoted to decisions: monitoring $\mu^w(w, z)$ and wage setting $\tau^w(w, z)$.

Figure 11: Money supply shock: additional IRFs. Error-prone pricing, $\zeta = 0.5$.

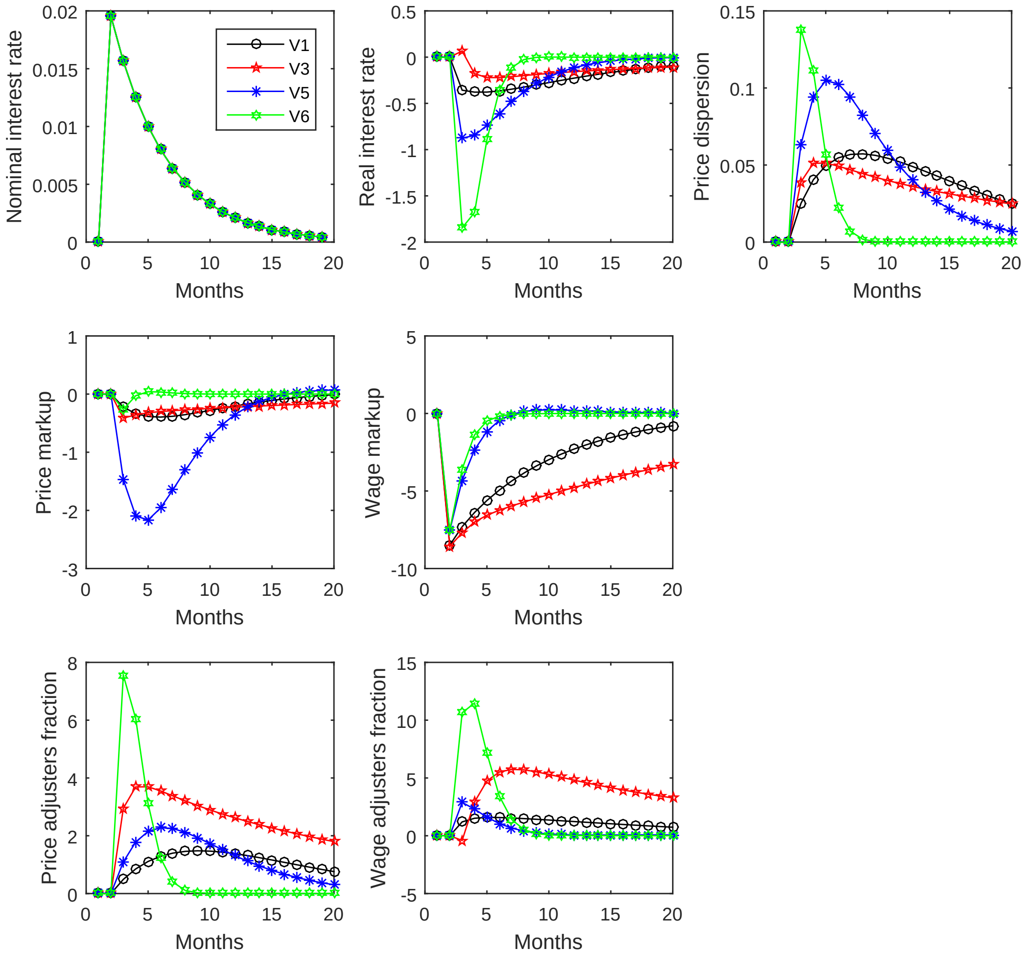

Notes:

Impulse responses to money growth shock with autocorrelation 0.8 (monthly).

Black: Benchmark (V1), both prices and wages sticky. Red: V3, flexible prices and sticky wages.

Blue: V5, sticky prices and flexible wages. Green: V6: both prices and wages flexible.

Appendix B: Frictionless limits

As a basis for comparison, we also calculate versions of our model where price adjustment and/or wage adjustment are perfectly frictionless. These versions can also be used to generate initial guesses for solving our main model, in which both price adjustment and wage adjustment are subject to control costs.

In the following calculations, we allow for aggregate dynamics, due to aggregate shocks or transition dynamics. However, we assume that the idiosyncratic shock processes a_{jt} and z_{it} have already converged to their ergodic distributions.

Implications of flexible prices

When prices are frictionless, firms choose P_{jt} to optimize static profits $P_{jt}^{1-\epsilon} C_t P_t^{\epsilon} - \frac{W_t P_{jt}^{-\epsilon}}{A_{jt}} C_t P_t^{\epsilon}$ each period, implying the first-order condition $P_{jt} = \left(\frac{\epsilon}{\epsilon-1}\right) \frac{W_t}{A_{jt}}$. In real terms, this becomes

$$p_{jt} = \log(P_{jt}/P_t) = \log\left(\frac{\epsilon(W_t/P_t)}{\epsilon-1}\right) - a_{jt} = \log\left(\frac{\epsilon w_t}{\epsilon-1}\right) - a_{jt}, \tag{94}$$

so the distribution of prices is an affine translation of the distribution of productivity at all times.

The aggregate price identity can be written as $E\exp((1-\epsilon)p) = 1$. Plugging in the formula for the idiosyncratic price, the aggregate price identity implies

$$\left(\frac{\epsilon}{\epsilon-1}\right)^{1-\epsilon} w_t^{1-\epsilon} E\exp((\epsilon-1)a) \;=\; 1,$$

which can be rearranged to give

$$w_t \;=\; w_{pf} \;\equiv\; \frac{\epsilon-1}{\epsilon}\left\{E\exp((\epsilon-1)a)\right\}^{\frac{1}{\epsilon-1}} \;=\; \frac{\epsilon-1}{\epsilon}\left\{\int A_{jt}^{\epsilon-1} dj\right\}^{\frac{1}{\epsilon-1}}. \tag{95}$$

Hence the aggregate real wage is constant ($w_t = w_{pf}$) if prices are flexible. This formula can be simplified further if a_{jt} is normally distributed,[27] but we will not impose this assumption here since in our model a_{jt} is governed by a discrete Tauchen (1986) approximation, which is not exactly normal.[28]

In the limit where price adjustment is costless, labor market clearing is simply $N_t = \Delta_t C_t$, where $\Delta_t = E\exp(-\epsilon p_{jt} - a_{jt})$. We then have

$$-\epsilon p_{jt} - a_{jt} \;=\; -\epsilon \log\left(\frac{\epsilon w_t}{\epsilon-1}\right) + (\epsilon-1)a_{jt}.$$

[27]If $a \sim N(0, \sigma_a^2)$, then $E\exp((\epsilon-1)a) = \exp\left((\epsilon-1)^2\sigma_a^2/2\right)$. So in this case price flexibility implies

$$w_{pf} \;=\; \frac{\epsilon-1}{\epsilon}\exp\left(\frac{(\epsilon-1)\sigma_a^2}{2}\right).$$

[28]The steady state distribution of a is given by the largest eigenvector of the Markov matrix that governs the dynamics of a.

Hence we can calculate the price dispersion measure Δ_t as

$$\Delta_t \;=\; E\exp(-\epsilon p_{jt} - a_{jt}) \;=\; \left(\frac{\epsilon w_t}{\epsilon - 1}\right)^{-\epsilon} E\exp((\epsilon-1)a).$$

But we already know that $\frac{\epsilon w_t}{\epsilon-1} = \{E\exp((\epsilon-1)a)\}^{\frac{1}{\epsilon-1}}$, so the formula for Δ_t reduces to

$$\Delta_t \;=\; \Delta_{pf} \;\equiv\; \{E\exp((\epsilon-1)a)\}^{\frac{-\epsilon}{\epsilon-1}+1} \;=\; \{E\exp((\epsilon-1)a)\}^{\frac{1}{1-\epsilon}} \;=\; \frac{\epsilon-1}{\epsilon w_{pf}}, \tag{96}$$

which is unchanging over time.[29] This explicit formula for the price dispersion measure implies that if prices are flexible, labor demand is a constant multiple of consumption: $N_t = \Delta_{pf} C_t$.

General equilibrium with flexible prices

With flexible prices, the aggregate real wage will equal w_{pf}. If we guess steady-state aggregate labor demand N, we can then calculate labor demand for each worker, $h(w_{it}, z_{it}) \equiv N w_{pf}^{\epsilon_n} \exp((\epsilon_n - 1)z_{it} - \epsilon_n w_{it})$. Using $h(w_{it}, z_{it})$, and imposing the consumption level $C = N/\Delta_{pf}$ and the steady state inflation rate $i = \log(\mu^*)$, we can solve the worker's problem by backwards iteration on (61)-(62).

By solving the worker's problem, we obtain the steady-state value and policy functions $l(w,z)$, $l^e(w,z)$, $\xi(w,z)$ $\pi^w(\tilde{w}|w,z)$, $\rho(w,z)$, $\mu^w(w,z)$, and $\tau^w(w,z)$. Using the policy functions π^w and ρ, and the steady-state inflation rate i, we can solve the dynamics of the wage distribution forward (equations (76)-(82)), to obtain the distribution Ψ^w.

Finally, we check the real wage identity (89). If it is satisfied, the steady-state general equilibrium level of labor demand N has been found. So steady state can be calculated as a one-dimensional root-finding problem in N.

To calculate the dynamics under flexible prices, we solve the same equation system that applies to our benchmark model, but we suppress the firm's problem (and the associated distributional dynamics), and we suppress the aggregate consistency condition for prices. The labor market clearing condition is replaced by $C_t = \Delta_{pf} N_t$, and the aggregate consistency condition for wages must hold, evaluated at the constant real wage level $w_t = w_{pf}$.

[29] When a is normal, this reduces to $\Delta_{pf} = \exp\left((1-\epsilon)\sigma_a^2/2\right)$.

Implications of flexible wages

Likewise, when wages are frictionless, workers optimize their static payoffs, which can be written in real terms as $w_{it}h_t(w_{it}, z_{it}) - \frac{X(h_t(w_{it}, z_{it}))}{u'(C_t)}$, where $h_t(w_{it}, z_{it}) \equiv N_t w_t^{\epsilon_n} \exp((\epsilon_n - 1)z_{it} - \epsilon_n w_{it})$. (Since the wage decision is costless, only time spent working enters as an argument of the disutility function X.) The first-order condition is

$$\exp(w_{it}) \;=\; \frac{\epsilon_n}{\epsilon_n - 1} C_t^{\gamma} \chi h_t(w_{it}, z_{it})^{\eta} \;=\; \frac{\epsilon_n \chi}{\epsilon_n - 1} C_t^{\gamma} \left[N_t w_t^{\epsilon_n} \exp((\epsilon_n - 1)z_{it} - \epsilon_n w_{it}) \right]^{\eta}.$$

Taking logs on both sides, and rearranging, the first-order condition becomes

$$w_{it} \;=\; \frac{1}{1 + \eta\epsilon_n} \log\left(\frac{\epsilon_n \chi}{\epsilon_n - 1} C_t^{\gamma} N_t^{\eta} w_t^{\eta\epsilon_n} \right) + \frac{\eta(\epsilon_n - 1)}{1 + \eta\epsilon_n} z_{it}. \tag{97}$$

By definition, the aggregate real wage satisfies

$$w_t^{1-\epsilon_n} \;=\; E \exp\left((1 - \epsilon_n)(w_{it} - z_{it}) \right)$$

at all times, where

$$(1 - \epsilon_n)(w_{it} - z_{it}) \;=\; \frac{1 - \epsilon_n}{1 + \epsilon_n \eta} \log\left(\frac{\epsilon_n \chi}{\epsilon_n - 1} C_t^{\gamma} N_t^{\eta} w_t^{\eta\epsilon_n} \right) + \frac{(\epsilon_n - 1)(1 + \eta)}{1 + \epsilon_n \eta} z_{it}.$$

Hence the wage index must satisfy

$$w_t^{1-\epsilon_n} \;=\; \left(\frac{\epsilon_n \chi}{\epsilon_n - 1} C_t^{\gamma} N_t^{\eta} w_t^{\eta\epsilon_n} \right)^{\frac{1-\epsilon_n}{1+\epsilon_n\eta}} E \exp\left(\frac{(1 - \epsilon_n)(1 + \eta)}{1 + \epsilon_n \eta} z \right).$$

Raising each side to the power $\frac{1+\epsilon_n\eta}{1-\epsilon_n}$ and rearranging, we obtain[30]

$$w_t \;=\; \left(\frac{\epsilon_n \chi}{\epsilon_n - 1} C_t^{\gamma} N_t^{\eta} \right) \left\{ E \exp\left(\frac{(1 - \epsilon_n)(1 + \eta)}{1 + \epsilon_n \eta} z \right) \right\}^{\frac{1+\epsilon_n\eta}{1-\epsilon_n}}. \tag{98}$$

Thus, if wages are flexible, the aggregate real wage varies over time, but it is a constant markup over $\frac{X'(N_t)}{u'(C_t)} = \chi N_t^{\eta} C_t^{\gamma}$. Each worker's hours $h_t(w_{it}, z_{it})$ are proportional to labor demand N_t (although the factor of proportionality varies with z_{it}), so each worker's marginal disutility is proportional to $\chi N_t^{\eta} C_t^{\gamma}$.

[30] If z is normally distributed, (98) reduces to

$$w_t \;=\; \left(\frac{\epsilon_n \chi}{\epsilon_n - 1} C_t^{\gamma} N_t^{\eta} \right) \exp\left(\frac{(1 - \epsilon_n)(1 + \eta)^2 \sigma_z^2}{2(1 + \epsilon_n \eta)} \right).$$

General equilibrium with flexible wages

If labor disutility is linear, the flexible aggregate wage relation (98) reduces to

$$w_t = \left(\frac{\epsilon_n \chi}{\epsilon_n - 1} C_t^{\gamma}\right) \{E \exp\left((1-\epsilon_n)z\right)\}^{1-\epsilon_n}. \qquad (99)$$

Therefore, with linear labor disutility, guessing the steady-state real wage w allows us to infer steady-state consumption C. From w and C we can calculate current real profits $u(p, a) = (e^p - we^{-a})\, Ce^{-\epsilon p}$ for any firm-specific state (p, a).

If labor disutility is nonlinear, we must instead guess both w and C to construct profits $u(p, a)$. In either case, we can then use steady-state inflation $i = \log(\mu^*)$ to solve the firm's problem by backwards iteration on (55) and (59). Thus we obtain the steady-state value and policy functions $v(p, a)$, $v^e(p, a)$, $\pi(p|a)$, $\lambda(p, a)$, $\mu(p, a)$, and $\tau(p, a)$. Using the policy functions π and λ, and the steady-state inflation rate i, we can solve the dynamics of the price distribution forward (equations (**??**)-(75)), to obtain the distribution steady-state Ψ.

Once we know the distribution Ψ, we can calculate price dispersion as $\Delta_t = \int\int \exp(-\epsilon p - a)\psi(p, a) da dp$. If labor disutility is linear, we can calculate labor demand from the labor market clearing condition (83). In the nonlinear case, we can instead calculate labor demand from the flexible aggregate wage condition (98).

Finally, we check the real price identity (88) in the linear disutility case; if it is satisified, we have found the steady-state real wage w. In the nonlinear case we must instead check both the price identity (88) and the labor market clearing condition (83); if both are satisfied, we have found the steady-state values of w and C. So steady-state general equilibrium can be calculated either as a one- or two-dimensional root-finding problem.

To calculate the dynamics under flexible wages, we solve the same equation system that applies to our benchmark model, but we suppress the worker's problem (and the associated distributional dynamics), imposing instead the wage equation (98).

When both prices and wages are flexible

When prices are flexible, we have $w_t = \frac{\epsilon \Delta_t}{\epsilon - 1}$ and labor market clearing implies $N_t = \Delta_t C_t$. When wages are flexible, the aggregate real wage satisfies (98). Combining these equations, we obtain a single equation for C_t:

$$\frac{\epsilon \Delta_t}{\epsilon - 1} = \left(\frac{\epsilon_n \chi}{\epsilon_n - 1} C_t^{\gamma} [\Delta_t C_t]^{\eta}\right) \left\{E \exp\left(\frac{(1-\epsilon_n)(1+\eta)}{1+\epsilon_n \eta} z\right)\right\}^{\frac{1+\epsilon_n \eta}{1-\epsilon_n}}.$$

We can solve this equation to show that frictionless consumption is a constant ($C_t = C_{flex}$), given by

$$\begin{aligned} C_{flex}^{\gamma+\eta} &= \left(\frac{\epsilon_n - 1}{\epsilon_n \chi}\right) \Delta_t^{-\eta} w_{pf} \left\{ E \exp\left(\frac{(1-\epsilon_n)(1+\eta)}{1+\epsilon_n \eta} z\right)\right\}^{\frac{1+\epsilon_n \eta}{\epsilon_n - 1}} \\ &= \left(\frac{\epsilon(\epsilon_n - 1)}{(\epsilon - 1)\epsilon_n \chi}\right) \Delta_t^{1-\eta} \left\{ E \exp\left(\frac{(1-\epsilon_n)(1+\eta)}{1+\epsilon_n \eta} z\right)\right\}^{\frac{1+\epsilon_n \eta}{\epsilon_n - 1}}, \end{aligned} \tag{100}$$

where Δ_t is given by (96). In other words, money shocks are neutral when both prices and wages are flexible. (If we instead included aggregate technology shocks in the model, the labor market clearing condition would have to be restated to take account of the shocks, so C_t and N_t would vary over time.)

Knowing steady-state consumption C_{flex}, all other steady-state quantities can be calculated, and since money is neutral real variables are always at their steady state.

Appendix C. Computation

Outline of algorithm

Computing this model is challenging due to heterogeneity. At any time t, firms face different productivity shocks A_{jt} and are stuck at different prices P_{jt}; likewise productivity and wages vary across workers. 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, restricting all idiosyncratic state variables to discrete grids. That is, real log prices p_{jt} lie at all times on a fixed grid $\gamma^p \equiv \{p^1, p^2, ...p^{\#^p}\}$; real log wages w_{it} lie in $\gamma^w \equiv \{w^1, w^2, ...w^{\#^w}\}$; and likewise for log productivities of firms and workers: $a_{jt} \in \gamma^a \equiv \{a^1, a^2, ...a^{\#^a}\}$ and $z_{it} \in \gamma^z \equiv \{z^1, z^2, ...z^{\#^z}\}$. The four grids γ^p, γ^w, γ^a, and γ^z are all assumed to have constant step sizes (in logs) between grid points. Moreover, we assume (only for numerical convenience) that the step size in γ^w equals that in γ^p, and also that the number of grid points is the same in these two grids: $\#^w = \#^p$.

We can then view firms' 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 $\left(p^j, a^k\right) \in \gamma^p \times \gamma^a$.[31] Similarly, the distribution of firms at the beginning (or end) of any given period can be viewed as a $\#^p \times \#^a$ matrix $\mathbf{\Psi}$ (or $\widetilde{\mathbf{\Psi}}$) in which the row j, column k element Ψ^{jk} (or $\widetilde{\Psi}^{jk}$) represents the fraction of firms in state (p^j, a^k) at the beginning (or end) of any given period. Likewise, the workers' steady-state value function and the beginning- and end-of-period distributions of workers can be represented by matrices $\mathbf{L}$, $\mathbf{\Psi^w}$ and $\widetilde{\mathbf{\Psi}}^\mathbf{w}$ of size $\#^w \times \#^z$. While these matrices are large objects, we can nonetheless solve for a steady-state general equilibrium as a low-dimensional root-finding problem. By guessing the steady-state values of C and N, we can set up the Bellman equations of the workers and firms, and solve for their fixed points $\mathbf{L}$ and $\mathbf{V}$; given optimal policies, we can describe the dynamics of the distributions, and thus solve for the steady-state distributions $\mathbf{\Psi^w}$, $\widetilde{\mathbf{\Psi}}^\mathbf{w}$, $\mathbf{\Psi}$, and $\widetilde{\mathbf{\Psi}}$; knowing the distributions, we will show that we can construct two scalar equations that suffice to check the values of C and N.

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 firms' value function is represented by a $\#^p \times \#^a$ matrix $\mathbf{V}_t$ with row j, column k element $v_t^{jk} \equiv v(p^j, a^k, \Xi_t)$, thus summarizing the time t values at all grid points $(p^j, a^k) \in \gamma^p \times \gamma^a$. Then, instead of viewing the Bellman equation as a functional equation that defines $v(p, a, \Xi)$ for all possible

[31] In this appendix, bold face indicates matrices, and (most) superscripts represent indices of matrices or grids.

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^{jk}. In addition, there will be a relation between the workers' values $\mathbf{L}_t$ and $\mathbf{L}_{t+1}$ at times t and $t+1$, which can also be seen as a system of $\#^w\#^z$ scalar equations in $\#^w\#^z$ unknowns. Finally, the distribution of firms at time $t+1$, $\mathbf{\Psi}_{t+1}$ is derived from the distribution at time t, $\mathbf{\Psi}_t$, which amounts to $\#^p\#^a$ scalar equations; and the distributional dynamics of workers links the distributions $\mathbf{\Psi^w}_t$ and $\mathbf{\Psi^w}_{t+1}$ with a matrix equation that is equivalent to a system of $\#^w\#^z$ scalar equations.[32]

We linearize these equations numerically (together with a handful of scalar equations, including first-order conditions for some aggregate variables). We then solve for the saddle-path stable solution of the linearized model using the QZ decomposition, following Klein (2000). It is crucial to note here that our problem is tractable because we have separated the two sticky decisions in our model between two different classes of decision-makers. In a model where a single decision-maker adjusted p and w in response to the shocks a and z, the value function and distributional dynamics would both have to be evaluated over $\#^p\#^w\#^a\#^z$ grid points. Solving for dynamic general equilibrium would require solving a system of slightly more than $2\#^p\#^w\#^a\#^z$ equations. Instead, since we have assumed prices and wages are set by different agents, we will have to solve slightly more than $2\#^p\#^a + 2\#^w\#^z$ equations, which is a vastly smaller problem.[33]

The discretized model

Firms' values are summarized by matrices $\mathbf{V}_t$ and $\mathbf{V}_t^e$, of size $\#^p \times \#^a$, and the vector $\tilde{\mathbf{v}}_t$, of length $\#^a$. Workers' values are described by the matrices, $\mathbf{L}_t$, $\mathbf{L}_t^e$, and $\widetilde{\mathbf{L}}_t$, of size $\#^w \times \#^z$. The elements of $\mathbf{V}_t$ are $v_t^{jk} \equiv v(p^j, a^k, \Xi_t)$, and the elements of $\mathbf{V}_t^e$ are $v_t^{e,jk} \equiv v^e(p^j, a^k, \Xi_t)$, for $\left(p^j, a^k\right) \in \gamma^p \times \gamma^a$. Likewise, $\mathbf{L}_t$ has elements $l_t^{jk} \equiv l(w^j, z^k, \Xi_t)$, and $\mathbf{L}_t^e$ has elements $l_t^{e,jk} \equiv l^e(w^j, z^k, \Xi_t)$, for $\left(w^j, z^k\right) \in \gamma^w \times \gamma^z$. The expected values of setting a new price or wage are given by vectors $\tilde{\mathbf{v}}_t$ and $\tilde{\mathbf{L}}_t$, with elements $\tilde{v}_t^k \equiv \tilde{v}(a^k, \Xi_t)$ and $\tilde{l}_t^{jk} \equiv \tilde{l}(w^j, a^k, \Xi_t)$.

Related matrices include the probability matrices of firms and workers, $\mathbf{\Lambda}_t$ and $\mathbf{R}_t$. The (j, k) ele-

[32] Here we are assuming that we can substitute out the steps that define the end-of-period distributions $\widetilde{\mathbf{\Psi}}_t$ and $\widetilde{\mathbf{\Psi}}_t^{\mathbf{w}}$. If not, our system will contain an additional $2\#^w\#^z$ equations.

[33] In other words, computational complexity under our approach scales exponentially with the number of sticky decisions if these decisions are all taken by the same agent, but scales linearly in the number of sticky decisions if different decisions are controlled by different agents. (Actually, the same principle is true in models of fully flexible decisions, but the issue is more relevant here because *stickiness creates heterogeneity*— while prices and wages are jump variables in flexible models, in the presence of nominal rigidity they become state variables.)

ments of these matrices are given by[34]

$$\lambda_t^{jk} \equiv \lambda\left(\frac{\tilde{v}_t^k - v_t^{jk}}{\kappa_\lambda w_t}\right), \quad \rho_t^{jk} \equiv \rho\left(\frac{\tilde{l}_t^k - l_t^{jk}}{\kappa_\rho \xi_t^{jk}}\right). \tag{101}$$

Finally, we also define the logit probabilities $\mathbf{\Pi}_t$ (a matrix) and $\mathbf{\Pi^w}_t$ (a 3d array). The elements of these matrices are

$$\pi_t^{jk} = \pi_t(p^j|a^k) \equiv \frac{\eta^j \exp\left(v_t^{jk}/(\kappa_\pi w_t)\right)}{\sum_{n=1}^{\#p} \eta^n \exp\left(v_t^{nk}/(\kappa_\pi w_t)\right)}, \tag{102}$$

$$\pi_t^{w,jkn} = \pi_t^w(w^n|w^j, z^k) \equiv \frac{\eta^{w,n} \exp\left(l_t^{nk}/(\kappa_w \xi_t^{jk})\right)}{\sum_{m=1}^{\#w} \eta^{w,m} \exp\left(l_t^{mk}/(\kappa_w \xi_t^{jk})\right)}. \tag{103}$$

Here π_t^{jk} is the probability that a firm which has decided to adjust its price at time t chooses real log price p^j, conditional on log productivity a^k; $\pi_t^{w,jkn}$ is a worker's corresponding probability of choosing the real log wage w^n, conditional on current log real wage w^j and log productivity z^k. The default probabilities for log real prices $p \in \gamma^p$ are $\boldsymbol{\eta} \equiv (\eta^1, \ldots, \eta^{\#^p}) \equiv (\eta(p^1), \ldots, \eta(p^{\#^p}))$, and $\boldsymbol{\eta}^{\mathbf{w}} \equiv (\eta^{w,1}, \ldots, \eta^{w,\#^w}) \equiv (\eta^w(w^1), \ldots, \eta^w(w,^{\#^w}))$ is the analogous vector for log real wages $w \in \gamma^w$.

In this discrete representation, the productivity processes (68) and (76) can be summarized by matrices $\mathbf{S}$ and $\mathbf{S^z}$ of size $\#^a \times \#^a$ and $\#^z \times \#^z$. The (m,k) elements of these matrices represent the following transition probabilities, respectively:

$$S^{mk} = prob(a_{jt} = a^m | a_{j,t-1} = a^k), \quad S^{z,mk} = prob(z_{it} = z^m | z_{i,t-1} = z^k). \tag{104}$$

It is helpful to introduce analogous Markovian notation to describe the deflation of real prices and wages as the aggregate price level rises. Let $\mathbf{T}_t$ be a $\#^p \times \#^p$ Markov matrix in which the row m, column l element represents the probability that firm j's beginning-of-period log real price $\tilde{p}_{jt}$ equals $p^m \in \gamma^p$ if its log real price at the end of the previous period was $p^l \in \gamma^p$:

$$T_t^{ml} \equiv prob(\widetilde{p}_{jt} = p^m | p_{j,t-1} = p^l). \tag{105}$$

Generically, the deflated log price $\tilde{p}_{jt} \equiv p_{j,t-1} - i_t \equiv p_{j,t-1} - i(\Xi_t, \Xi_{t-1})$ will fall between two grid points; then the matrix $\mathbf{T}_t$ must round up or down stochastically. Also, if $p_{j,t-1} - i_t$ lies below the smallest or above the largest element of the grid, then $\mathbf{T}_t$ must round up or down to keep prices on the

[34]Actually, (101) is a simplified description of λ_t^{jk}. While (101) implies that λ_t^{jk} represents the function $\lambda(\bullet)$ 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 $\lambda(\bullet)$ 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 Ξ_t.

grid.[35] Therefore we construct $\mathbf{T}_t$ according to

$$T_t^{ml} = prob(\widetilde{p}_{jt} = p^m | p_{j,t-1} = p^l, i_t) = \begin{cases} 1 & \text{if } p^l - i_t \leq p^1 = p^m \\ \frac{p^l - i_t - p^{m-1}}{p^m - p^{m-1}} & \text{if } p^1 < p^m = \min\{p \in \Gamma^p : p \geq p^l - i_t\} \\ \frac{p^{m+1} - p^l + i_t}{p^{m+1} - p^m} & \text{if } p^1 \leq p^m = \max\{p \in \Gamma^p : p < p^l - i_t\} \\ 1 & \text{if } p^l - i_t > p^{\#^p} = p^m \\ 0 & \text{otherwise} \end{cases} \tag{106}$$

Furthermore, recall that we have assumed that the price and wage grids γ^p and γ^w have the same step size, and the same number of grid points. Note that in this case, the transition probabilities mapping real log wages from one period to the beginning of the next are the same as those for real log prices. In other words, for all m and l,

$$prob(\widetilde{w}_{it} = w^m | w_{j,t-1} = w^l) \;=\; prob(\widetilde{p}_{jt} = p^m | p_{j,t-1} = p^l) \;=\; T_t^{ml}. \tag{107}$$

Thus we can describe the distributional dynamics of wages using exactly the same matrix $\mathbf{T}_t$ that we used from prices.

Given this notation, we can now write the distributional dynamics in a more compact form. The time t distributions of firms and workers are derived from the distributions at the end of $t-1$ as follows:

$$\mathbf{\Psi}_t = \mathbf{T}_t \widetilde{\mathbf{\Psi}}_{t-1} \mathbf{S}', \qquad \mathbf{\Psi}_t^w = \mathbf{T}_t \widetilde{\mathbf{\Psi}}_{t-1}^w (\mathbf{S}^{\mathbf{z}})'. \tag{108}$$

Note that exogenous shocks are represented from left to right in the matrices $\widetilde{\mathbf{\Psi}}_t$ and $\widetilde{\mathbf{\Psi}}_t^w$, so that their transitions can be treated by right multiplication, while sticky decision variables are represented vertically, so that transitions related to choice variables can be described by left multiplication. Next, to calculate the effects of price adjustment on the distribution, let $\mathbf{1}_{pp}$, $\mathbf{1}_{pa}$, $\mathbf{1}_{ww}$, and $\mathbf{1}_{wz}$ be matrices of ones of size $\#^p \times \#^p$, $\#^p \times \#^a$, $\#^w \times \#^w$, and $\#^w \times \#^z$, respectively. After production occurs at time t, as new real prices are set, the price distribution adjusts as follows:

$$\widetilde{\mathbf{\Psi}}_t \;=\; (\mathbf{1}_{pa} - \mathbf{\Lambda}_t) \odot \mathbf{\Psi}_t + \mathbf{\Pi}_t \odot (\mathbf{1}_{pp}(\mathbf{\Lambda}_t \odot \mathbf{\Psi}_t)). \tag{109}$$

where the operator $\odot$ represents element-by-element multiplication (the Hadamard product). The matrix notation does not carry over to the wage dynamics, because the distribution of new wages varies with the *current* wage, so instead we state the dynamics one row at a time:

$$\widetilde{\mathbf{\Psi}}_t^{w,j} \;=\; (\mathbf{1}_z - \mathbf{\Lambda}_t^j) \odot \mathbf{\Psi}_t^j + \; \mathbf{1}_w'(\mathbf{\Pi}_t^{w,j} \odot \mathbf{\Lambda}_t \odot \mathbf{\Psi}_t). \tag{110}$$

[35] In 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.

Here $\widetilde{\mathbf{\Psi}}_t^{w,j}$ and $\mathbf{\Lambda}_t^j$ are the jth rows of $\widetilde{\mathbf{\Psi}}_t^w$ and $\mathbf{\Lambda}_t$, respectively, while $\mathbf{\Pi}_t^{w,j}$ is the matrix representing the probability of choosing wage w^j conditional on each possible state: $prob(w^j|w^i, z^k, \Xi_t)$. $\mathbf{1}_z$ and $\mathbf{1}_w$ are conformable matrices of ones.

The same transition matrices $\mathbf{T}_t$, $\mathbf{S}$, and $\mathbf{S^z}$ show up when we write the Bellman equations in matrix form. The discounted values of choosing each possible real price $\tilde{p}$ are

$$\mathbf{V}_t^e = \beta E_t \left\{ \frac{C_{t+1}^{-\gamma}}{C_t^{-\gamma}} \mathbf{T}'_{t+1} \mathbf{V}_{t+1} \mathbf{S} \right\}, \qquad \mathbf{L}_t^e = \beta E_t \left\{ \frac{C_{t+1}^{-\gamma}}{C_t^{-\gamma}} \mathbf{T}'_{t+1} \mathbf{L}_{t+1} \mathbf{S^z} \right\}. \tag{111}$$

Here the expectation E_t refers only to the effects of the time $t+1$ aggregate shock g_{t+1}, because the dynamics of the idiosyncratic states (p_{jt}, a_{jt}) and (w_{it}, z_{it}) are completely described by the matrices $\mathbf{T}'_{t+1}$, $\mathbf{S}$, and $\mathbf{S^z}$.

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

$$u_t^{jk} \equiv \left(\exp(p^j) - \frac{w_t}{\exp(a^k)} \right) \frac{C_t}{\exp(\epsilon p^j)} \tag{112}$$

for $\left(p^j, a^k\right) \in \gamma^p \times \gamma^a$. The define the current payoffs of the workers, let $\mathbf{H}_t$ be the $\#^w \times \#^z$ matrix containing the elements $h_t^{jk} \equiv h_t(w^j, z^k)$, representing labor demand in state (w^j, z^k, Ξ_t). Also define $\mathbf{W}$ as a conformable matrix with all the elements of row j equal to $\exp w^j$, and $\mathbf{X}_t$ as a matrix containing the elements $\frac{X(h_t^{jk} + \tau_t^{jk} + \mu_t^{jk})}{u'(c(\Xi_t))}$ representing total disutility of time use in state (w^j, z^k, Ξ_t). Then we can calculate the value functions as

$$\mathbf{V}_t = \mathbf{U}_t + \mathbf{\Lambda}_t \odot (E^\pi \mathbf{V^e}_t - \mathbf{K}_t^\pi) + (\mathbf{1}_{pp} - \mathbf{\Lambda}_t) \odot \mathbf{V^e}_t - \mathbf{K}_t^\lambda \tag{113}$$

$$\mathbf{L}_t = \mathbf{W} \odot \mathbf{H}_t - \mathbf{X}_t + \mathbf{R}_t \odot E^\pi \mathbf{L^e}_t + (\mathbf{1}_{ww} - \mathbf{R}_t) \odot \mathbf{L^e}_t \tag{114}$$

In order to check labor market clearing it will be helpful to define several summary statistics related to labor time use. First, let K_t^λ and K_t^π and be total time use for choosing the timing of the price decision, and actually choosing prices:

$$\begin{aligned}
K_t^\lambda &= \sum_{j=1}^{\#p} \sum_{k=1}^{\#a} \psi_t^{jk} \left(\lambda_t^{jk} \ln \left(\frac{\lambda_t^{jk}}{\bar{\lambda}} \right) + (1 - \lambda_t^{jk}) \ln \left(\frac{1 - \lambda_t^{jk}}{1 - \bar{\lambda}} \right) \right), && (115) \\
K_t^\pi &= \sum_{j=1}^{\#p} \sum_{k=1}^{\#a} \psi_t^{jk} \lambda_t^{jk} \left(\sum_{i=1}^{\#p} \pi_t^{ik} \ln \left(\frac{\pi_t^{ik}}{\eta^k} \right) \right), && (116) \\
\Delta_t &= \sum_{j=1}^{\#p} \sum_{k=1}^{\#a} \psi_t^{jk} \exp(-\epsilon p^j - a^k). && (117)
\end{aligned}$$

Note that in the second equation, the time K_t^π devoted to choosing prices is weighted by the fraction adjusting, λ_t^{jk}. In the third equation, Δ_t represents a price dispersion measure that relates time devoted to production to total goods produced.

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, aggregate shocks are zero; the distribution of firms takes some unchanging value $\mathbf{\Psi}$, and the distribution of workers takes some unchanging value $\mathbf{\Psi^w}$. Thus the aggregate state of the economy is constant: $\Xi_t \equiv (g_t, \Psi_{t-1}, \Psi^w_{t-1}) = (0, \mathbf{\Psi}, \mathbf{\Psi^w}) \equiv \Xi$. 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 $\mathbf{T}$ is known, since it depends only on steady state inflation i. Morever, the Euler equation reduces to

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

which simply serves to determine the nominal interest rate R.

We can then calculate general equilibrium as a three-dimensional root-finding problem, by guessing consumption C, labor demand N, and the aggregate wage level w. On one hand, knowing $c(\Xi)$ and $w(\Xi)$ we can construct the firm's profit function $u(p, a, \Xi) = (e^p - w(\Xi)e^{-a})c(\Xi)e^{-\epsilon p}$. Knowing the profit function, we can solve the firm's problem by backwards induction, which yields the value functions v, v^e, and $\tilde{v}$, and the policy functions λ and π. Given the firm's policy functions, we can calculate the distributional dynamics to find the steady-state distribution of prices and productivities, $\Psi(p, a)$. From the firm's problem and the steady-state distribution we can also calculate the time firms devote to decision-making (K_t^λ and K_t^π), and the efficiency wedge Δ.

On the other hand, knowing $n(\Xi)$ and $w(\Xi)$ we can construct the labor demand function $h(w, z, \Xi) = e^{z(\epsilon_n - 1)}n(\Xi)w(\Xi)^{\epsilon_n}e^{-\epsilon_n w}$, and given $c(\Xi)$ we can also calculate worker's utility value of labor income, $u'(c(\Xi))e^w h(w, z, \Xi)$. We can then solve the worker's Bellman equation by backwards induction. This yields the vaue functions l, l^e, and $\tilde{l}$, and the policy functions ρ, and π^w, as well as the time use function τ and μ, and the worker's marginal value of time ξ. Given the worker's policy functions, we can calculate the distributional dynamics to find the steady-state distribution of wages and productivities, $\Psi^w(w, z)$.

With these distributions in hand, we can then check whether the guessed values of C, N, and w are consistent with an equilibrium. Then we check the following three scalar equations:

$$
\begin{aligned}
1 &= \sum_{j=1}^{\#^p}\sum_{k=1}^{\#^a} \psi^{jk} \exp\left((1-\epsilon)p^j\right), && (118)\\
w &= \left\{\sum_{j=1}^{\#^p}\sum_{k=1}^{\#^a} \psi^{w,jk} \exp\left((1-\epsilon_n)(w^j - z^k)\right)\right\}^{\frac{1}{1-\epsilon_n}}. && (119)\\
N &= \Delta C + \kappa_\pi K^\pi + \kappa_\lambda K^\lambda. && (120)
\end{aligned}
$$

The first two equations are the aggregate price and wage identities; the last is the labor market clearing condition. If these three equations are satisfied with sufficient accuracy, then a steady-state general equilibrium has been found.

Step 2: linearized dynamics

We now conjecture that nominal and real state variables take the form $\Omega_t \equiv (M_t, g_t, \Phi_t, \Phi_t^w)$ and $\Xi_t \equiv (g_t, \Phi_t, \Phi_t^w)$, respectively. We will show that this is a valid state variable for the economy by constructing an equilibrium in terms of this state.

Given the steady state, the general equilibrium dynamics can be calculated by linearization. To reduce the size of the Jacobian, we will eliminate many variables from the equation system. Thus, we calculate the end-of-period distributions as an intermediate step, without explicitly counting them in the equation system:

$$\widetilde{\mathbf{\Psi}}_t = (\mathbf{1}_{pa} - \mathbf{\Lambda}_t) \odot \mathbf{\Psi}_t + \mathbf{\Pi}_t \odot (\mathbf{1}_{pp}(\mathbf{\Lambda}_t \odot \mathbf{\Psi}_t)) \tag{121}$$

$$\widetilde{\mathbf{\Psi}}_t^{w,j} = (\mathbf{1}_z - \mathbf{\Lambda}_t^j) \odot \mathbf{\Psi}_t^j + \mathbf{1}'_w(\mathbf{\Pi}_t^{w,j} \odot \mathbf{\Lambda}_t \odot \mathbf{\Psi}_t) \tag{122}$$

Having thus calculated $\widetilde{\mathbf{\Psi}}_t$ and $\widetilde{\mathbf{\Psi}}_t^w$, the following two equations can be counted as determining the dynamics of the distributions $\mathbf{\Psi}$ and $\mathbf{\Psi^w}$ from periods t to $t+1$:

$$\mathbf{\Psi}_{t+1} = \mathbf{T}_{t+1}\widetilde{\mathbf{\Psi}}_t\mathbf{S}', \tag{123}$$

$$\mathbf{\Psi^w}_{t+1} = \mathbf{T}_{t+1}\widetilde{\mathbf{\Psi}}_t^{\mathbf{w}}\mathbf{S}'. \tag{124}$$

Similarly, we do not count the expected values V^e and L^e explicitly in our equation system, but we evaluate them in an intermediate step as follows:

$$\mathbf{V}_t^e = \beta E_t\left\{\frac{C_{t+1}^{-\gamma}}{C_t^{-\gamma}}\mathbf{T}'_{t+1}\mathbf{V}_{t+1}\mathbf{S}\right\} \tag{125}$$

$$\mathbf{L}_t^e = \beta E_t\left\{\frac{C_{t+1}^{-\gamma}}{C_t^{-\gamma}}\mathbf{T}'_{t+1}\mathbf{L}_{t+1}\mathbf{S^z}\right\} \tag{126}$$

Given the expected values V^e_t and L^e_t, which can be used to calculate the probabilities Λ_t, Π_t, and so forth, we then count the following two Bellman equations, which determine the dynamics of the value functions V_t and L_t:

$$\mathbf{V}_t = \mathbf{U}_t + \mathbf{\Lambda}_t \odot (E^\pi \mathbf{V^e}_t - \mathbf{K}^\pi_t) + (1 - \mathbf{\Lambda}_t) \odot \mathbf{V^e}_t - \mathbf{K}^\lambda_t \tag{127}$$

$$\mathbf{L}_t = \mathbf{W} \odot \mathbf{H}_t - \mathbf{X}_t + \mathbf{R}_t \odot E^\pi \mathbf{L^e}_t + (1 - \mathbf{R}_t) \odot \mathbf{L^e}_t \tag{128}$$

We also include the following six scalar equations in our system:

$$1 = \sum_{j=1}^{\#^p} \sum_{k=1}^{\#^a} \Psi^{jk}_t \exp((1-\epsilon)p^j) \tag{129}$$

$$w_t^{1-\epsilon_n} = \sum_{j=1}^{\#^p} \sum_{k=1}^{\#^a} \Psi^{w,jk}_t \exp((1-\epsilon_n)(w^j - z^k)) \tag{130}$$

$$N_t \; = \; \Delta_t C_t + \kappa_\pi K^\pi_t + \kappa_\lambda K^\lambda_t \tag{131}$$

$$\frac{\mu \exp(g_t)}{\exp i_t} = \frac{m_t}{m_{t-1}} \tag{132}$$

$$1 - \frac{\nu}{m_t C_t^{-\gamma}} = \beta E_t \left(\frac{C^{-\gamma}_{t+1}}{i_{t+1} C_t^{-\gamma}} \right) \tag{133}$$

$$g_{t+1} = \phi_g g_t + \epsilon^g_{t+1} \tag{134}$$

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

$$\overrightarrow{X}_t \equiv \big(vec\,(\mathbf{\Psi}_t)',\ vec\,(\mathbf{\Psi^w}_t)',\ \ m_{t-1},\ w_t,\ i_t,\ vec\,(\mathbf{V}_t)',\ vec\,(\mathbf{L}_t)',\ C_t,\ N_t \big)'$$

then the four matrix equations (123), (124), (127), and (128), together with the six scalar equations (129)-(134), amount to first-order system of the following form:

$$E_t \mathcal{F} \left(\overrightarrow{X}_{t+1}, \overrightarrow{X}_t, g_{t+1}, g_t \right) = 0 \tag{135}$$

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

Since the number of equations matches the number of variables included in the system $\mathcal{F}$, we can linearize the system numerically with respect to all its arguments to construct the Jacobian matrices $\mathcal{A} \equiv D_{\overrightarrow{X}_{t+1}} \mathcal{F}$, $\mathcal{B} \equiv D_{\overrightarrow{X}_t} \mathcal{F}$, $\mathcal{C} \equiv D_{g_{t+1}} \mathcal{F}$, and $\mathcal{D} \equiv D_{g_t} \mathcal{F}$. Thus we obtain the following first-order expectational difference equation system:

$$E_t \mathcal{A} \Delta \overrightarrow{X}_{t+1} + \mathcal{B} \Delta \overrightarrow{X}_t + E_t \mathcal{C} g_{t+1} + \mathcal{D} g_t = 0 \tag{136}$$

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. When applying this method, note that $\mathbf{\Psi}_t$, $\mathbf{\Psi^w}_t$, m_{t-1}, w_t, and i_t are all predetermined at t, while $\mathbf{V}_t$, $\mathbf{L}_t$, C_t, and N_t are jump variables.