Online Appendices for "Logit price dynamics", July 2018

James Costain (Banco de España and European Central Bank)

Anton Nakov (European Central Bank and CEPR)

ONLINE APPENDIX A: COMPUTATION

Outline of algorithm

Heterogeneity is a challenge when computing our model: at any time t, productivities A_{it} and prices P_{it} will differ across firms. 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 sticky-price models, in which equilibrium quantities depend on the whole time-varying distribution of prices and productivity across firms.

To address this issue, we apply 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 a discretized approximation of the underlying model; here we restrict real log prices p_{it} and log productivities a_{it} to a fixed grid $\Gamma \equiv \Gamma^p \times \Gamma^a$, where $\Gamma^p \equiv \{p^1, p^2, ...p^{\#^p}\}$ and $\Gamma^a \equiv \{a^1, a^2, ...a^{\#^a}\}$ are both uniformly spaced (in logs). We can then view the steady state value function as a matrix $\mathbf{V}$ of size $\#^p \times \#^a$, comprising the values $v^{jk} \equiv v(p^j, a^k, \Xi)$ associated with prices and productivities $\left(p^j, a^k\right) \in \Gamma$.[1] Likewise, the price distribution can be viewed as a $\#^p \times \#^a$ matrix $\mathbf{\Psi}$ in which the row j, column k element Ψ^{jk} represents the fraction of firms in state (p^j, a^k) at the end of any given period. To calculate steady state general equilibrium, we can guess the wage w, then

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

solve the firm's problem by backwards induction on the grid Γ, then update the conjectured wage, and iterate to convergence.

The second step 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. The value function is represented by a $\#^p \times \#^a$ matrix $\mathbf{V}_t$ with row j, column k element $v_t^{jk} \equiv v(p^j, a^k, \Xi_t)$, summarizing the time t values at all grid points $(p^j, a^k) \in \Gamma$. Then, instead of treating the Bellman equation as a functional equation that defines $v(p, a, \Xi)$ for all possible idiosyncratic and aggregate states p, a, and Ξ, we view it as a difference equation linking the matrices $\mathbf{V}_t$ and $\mathbf{V}_{t+1}$. This amounts to a (large!) system of $\#^p\#^a$ first-order expectational difference equations governing the $\#^p\#^a$ variables v_t^{jk}. We linearize these equations numerically (together with the $\#^p\#^a$ equations that govern the distribution $\mathbf{\Psi}_t$, and a few other scalar equations). We solve the linearized model using the QZ decomposition, following Klein (2000).

This method combines linearity and nonlinearity in a way appropriate for models of price setting, where idiosyncratic shocks tend to be more relevant for firms' decisions than aggregate shocks are. By linearizing the aggregate dynamics, we recognize that changes in the aggregate shock z_t or in the distribution $\mathbf{\Psi}_t$ are unlikely to have a highly nonlinear impact on the value function. This smoothness does not require any "approximate aggregation" property, in contrast with the Krusell and Smith (1998) method; nor do we need to impose any particular functional form on the distribution $\mathbf{\Psi}$. However, to allow for the strong impact of firm-specific shocks, the method treats variation along idiosyncratic dimensions in a fully nonlinear way: the value at each grid point is determined by a distinct equation.

The discretized model

In the discretized model, the value $\mathbf{V}_t$ is a $\#^p \times \#^a$ matrix with elements $v_t^{jk} \equiv v(p^j, a^k, \Xi_t)$ for $\left(p^j, a^k\right) \in \Gamma$. A uniform default distribution θ allocates probability $1/\#^p$ to each price in Γ^p. Solving a single Bellman step analytically, the expected value of setting a new price is a row vector $\tilde{\mathbf{v}}_t$ of length $\#^a$, with kth element

$$\tilde{v}_t^k \;\equiv\; \kappa_\pi w_t \ln\left(\frac{1}{\#^p}\sum_{j=1}^{\#^p}\exp\left(\frac{v_t^{jk}}{\kappa_\pi w_t}\right)\right). \tag{40}$$

The value function $\mathbf{O}_t$ is also a $\#^p \times \#^a$ matrix, as is the hazard policy $\mathbf{\Lambda}_t$ and the logit price probabilities policy $\mathbf{\Pi}_t$; their (j,k) elements are given by[2]

$$o_t^{jk} \equiv \kappa_\lambda w_t \left(\bar{\lambda}\exp\left(\frac{\tilde{v}_t^k}{\kappa_\lambda w_t}\right) + (1-\bar{\lambda})\exp\left(\frac{v_t^{jk}}{\kappa_\lambda w_t}\right)\right), \tag{41}$$

$$\lambda_t^{jk} \equiv \bar{\lambda}\left(\bar{\lambda} + (1-\bar{\lambda})\exp((v_t^{jk} - \tilde{v}_t^k)/(\kappa w_t))\right)^{-1}, \tag{42}$$

$$\pi_t^{jk} \equiv \frac{\exp\left(v_t^{jk}/(\kappa w_t)\right)}{\sum_{n=1}^{\#p}\exp\left(v_t^{nk}/(\kappa w_t)\right)}. \tag{43}$$

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

In this discrete representation, the productivity process (34) can be written as a $\#^a \times \#^a$ matrix $\mathbf{S}$, where the (m,k) element represents the following transition probability:

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

Likewise, we can write the impact of inflation on real prices in Markovian notation. Let $\mathbf{R}_t$ be a $\#^p \times \#^p$ matrix in which element (m,l) represents the probability that firm i's beginning-of-t

[2]Equation (42) is a simplified description of λ_t^{jk}. While (42) implies that λ_t^{jk} represents the function $\lambda_t(p^j, a^k)$ evaluated at the log price grid point p^j and log productivity grid point a^k, in our computations λ_t^{jk} in fact represents the *average* of $\lambda_t(\tilde{p}, a^k)$ 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 $v_t(\tilde{p}, a^k)$ between price grid points. Defining λ_t^{jk} this way ensures differentiability with respect to changes in the aggregate state Ω_t.

log real price $\tilde{p}_{it}$ equals $p^m \in \Gamma^p$, if its log real price at the end of $t-1$ was $p^l \in \Gamma^p$:

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

Generically, the deflated log price $p_{i,t-1} - i_{t-1,t}$ will fall between two grid points; then the matrix $\mathbf{R}_t$ must round up or down stochastically.[3] Also, if $p_{i,t-1} - i_{t-1,t}$ lies below the smallest or above the largest element of the grid, then $\mathbf{R}_t$ must round up or down to keep prices on the grid.[4] Unbiased rounding results if $\mathbf{R}_t$ is constructed as:

$$R_t^{ml} = prob(\widetilde{p}_{it} = p^m | p_{i,t-1} = p^l, i_t) = \begin{cases} 1 & \text{if } p^l - i_t \leq p^1 = p^m \\ \frac{p^l - i_t - p^{m-1}}{p^m - p^{m-1}} & \text{if } p^1 < p^m = \min\{p \in \Gamma^p : p \geq p^l - i_t\} \\ \frac{p^{m+1} - p^l + i_t}{p^{m+1} - p^m} & \text{if } p^1 \leq p^m = \max\{p \in \Gamma^p : p < p^l - i_t\} \\ 1 & \text{if } p^l - i_t > p^{\#^p} = p^m \\ 0 & \text{otherwise.} \end{cases} \tag{44}$$

The distributional dynamics can now be written in compact matrix form; eq. (23) becomes:

$$\widetilde{\mathbf{\Psi}}_t = \mathbf{R}_t * \mathbf{\Psi}_{t-1} * \mathbf{S}', \tag{45}$$

where $*$ represents ordinary matrix multiplication. Productivity shocks are represented by right multiplication, while transitions in the real price level are represented by left multiplication. Next, to calculate the effects of price adjustment on the distribution, let $\mathbf{E}_{pp}$ and $\mathbf{E}_{pa}$ be matrices of ones of size $\#^p \times \#^p$ and $\#^p \times \#^a$, respectively. Eq. (24) is then:

$$\mathbf{\Psi}_t = (\mathbf{E}_{pa} - \mathbf{\Lambda}) .* \widetilde{\mathbf{\Psi}}_t + \mathbf{\Pi}_t .* (\mathbf{E}_{pp} * (\mathbf{\Lambda} .* \widetilde{\mathbf{\Psi}}_t)), \tag{46}$$

[3] If instead the firm's control variable were its real price, then $\mathbf{R}_t$ would simply be an identity matrix.

[4] In other words, any nominal price leading to a real log price below p^1 after inflation is automatically rounded up to the real log price p^1 (and to compute examples with deflation we must shift 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 Γ^p is sufficiently wide.

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

The same transition matrices $\mathbf{R}$ and $\mathbf{S}$ appear in the matrix form of the Bellman equation. Let $\mathbf{U}_t$ be the $\#^p \times \#^a$ matrix of current payoffs, with elements

$$u_t^{jk} \equiv \left(\exp(p^j) - \frac{w_t}{\exp(a^k)}\right) \frac{C_t}{\exp(\epsilon p^j)} \tag{47}$$

for $\left(p^j, a^k\right) \in \Gamma$. Then Bellman equation (14) becomes:

$$\mathbf{V}_t = \mathbf{U}_t + \beta E_t \left\{ \frac{C_{t+1}^{-\gamma}}{C_t^{-\gamma}} \left[\mathbf{R}'_{t+1} * \mathbf{O}_{t+1} * \mathbf{S}\right] \right\}. \tag{48}$$

The expectation E_t in (48) refers only to the effects of the time $t+1$ aggregate shock z_{t+1}, because the expectation over idiosyncratic states $(p^j, a^k) \in \Gamma$ is represented by multiplying by $\mathbf{R}'_{t+1}$ and $\mathbf{S}$. Note that since (48) iterates backwards in time, its transitions are governed by $\mathbf{R}'$ and $\mathbf{S}$, whereas (45) iterates forward in time, involving $\mathbf{R}$ and $\mathbf{S}'$.

We now discuss how we apply Reiter's (2009) two-step method to this discrete model.

Step 1: steady state

In the aggregate steady state, aggregate shocks are zero, and the distribution is in a steady state $\mathbf{\Psi}$, so the state of the economy is constant: $\Xi_t \equiv (z_t, \Psi_{t-1}) = (0, \mathbf{\Psi}) \equiv \Xi$. We indicate steady states of all equilibrium objects by dropping 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 implies that nominal money growth rate equals the inflation rate in steady state: $\mu = \exp(i)$. Thus, the steady-state transition matrix $\mathbf{R}$ is known, since it depends only on inflation i, and the Euler equation reduces to $\exp(i) = \beta R$.

We can then calculate general equilibrium as a one-dimensional root-finding problem in w.

Given w, we calculate $C = (w/\chi)^{1/\gamma}$, and then construct matrix $\mathbf{U}$, with elements

$$u^{jk} \equiv \left(\exp(p^j) - \frac{w}{\exp(a^k)}\right) \frac{C}{\exp(\epsilon p^j)} \,. \tag{49}$$

We can then find the fixed point of the value $\mathbf{V}$ (simultaneously with $\tilde{\mathbf{v}}$ and $\mathbf{O}$):

$$\mathbf{V} = \mathbf{U} + \beta \mathbf{R}' * \mathbf{O} * \mathbf{S}. \tag{50}$$

This allows us to calculate the logit matrix $\mathbf{\Pi}$, with elements

$$\pi^{jk} \equiv \frac{\exp\left(v^{jk}/(\kappa w)\right)}{\sum_{n=1}^{\#p} \exp\left(v^{nk}/(\kappa w)\right)}. \tag{51}$$

Likewise, we calculate the hazard matrix $\mathbf{\Lambda}$. We can then find the steady state distribution by iterating on the two-step distributional dynamics:

$$\mathbf{\Psi} = (\mathbf{E}_{pa} - \mathbf{\Lambda}) \,.*\, \widetilde{\mathbf{\Psi}} + \mathbf{\Pi} \,.*\, (\mathbf{E}_{pp} * (\mathbf{\Lambda} \,.*\, \widetilde{\mathbf{\Psi}})) \tag{52}$$

$$\widetilde{\mathbf{\Psi}} = \mathbf{R} * \mathbf{\Psi} * \mathbf{S}' \tag{53}$$

Finally, we check whether

$$1 \;=\; \sum_{j=1}^{\#^p} \sum_{k=1}^{\#^a} \Psi^{jk} \exp\left((1-\epsilon)p^j\right) \equiv p(w) \tag{54}$$

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

Step 2: linearized dynamics

Given the steady state, the general equilibrium dynamics can be calculated by linearization. First, we eliminate as many variables from the equation system as we can, summarizing the dynamics in terms of the exogenous shock process z_t, the lagged distribution of idiosyncratic states $\mathbf{\Psi}_{t-1}$, and the endogenous "jump" variables including $\mathbf{V}_t$, $\mathbf{\Pi}_t$, C_t, m_{t-1}, and i_t. The equation system reduces to

$$z_t = \phi_z z_{t-1} + \epsilon_t^z \tag{55}$$

$$\frac{\mu \exp(z_t)}{\exp i_t} = \frac{m_t}{m_{t-1}} \tag{56}$$

$$\mathbf{\Psi}_t = (\mathbf{E}_{pa} - \mathbf{\Lambda}_t) \mathbin{.*} \widetilde{\mathbf{\Psi}}_t + \mathbf{\Pi}_t \mathbin{.*} (\mathbf{E}_{pp} * (\mathbf{\Lambda}_t \mathbin{.*} \widetilde{\mathbf{\Psi}}_t)) \tag{57}$$

$$\mathbf{V}_t = \mathbf{U}_t + \beta E_t \left\{ \frac{C_{t+1}^{-\gamma}}{C_t^{-\gamma}} \left[\mathbf{R}'_{t+1} * \mathbf{O}_{t+1} * \mathbf{S} \right] \right\} \tag{58}$$

$$1 = \sum_{j=1}^{\#^p} \sum_{k=1}^{\#^a} \Psi_t^{jk} \exp((1-\epsilon)p^j) \tag{59}$$

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

$$\overrightarrow{X}_t \equiv \left(vec\left(\mathbf{\Psi}_{t-1}\right)', \quad vec\left(\mathbf{V}_t\right)', \quad C_t, \quad m_{t-1}, \quad i_t \right)'$$

then the whole set of expectational difference equations (55)-(59) governing the dynamic equilibrium becomes a first-order system of the following form:

$$E_t \mathcal{F}\left(\overrightarrow{X}_{t+1}, \overrightarrow{X}_t, z_{t+1}, z_t \right) = 0 \tag{60}$$

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

To see that the vector $\overrightarrow{X}_t$ in fact contains all the variables we need, note that given i_t and i_{t+1} we can construct $\mathbf{R}_t$ and $\mathbf{R}_{t+1}$. Given $\mathbf{R}_t$, we can construct $\widetilde{\mathbf{\Psi}}_t = \mathbf{R}_t * \mathbf{\Psi}_{t-1} * \mathbf{S}'$ from $\mathbf{\Psi}_{t-1}$. Given $w_t = \chi C_t^{\gamma}$, we can construct $\mathbf{U}_t$, with (j,k) element equal to $u_t^{jk} \equiv \left(\exp(p^j) - \frac{w_t}{\exp(a^k)} \right) \frac{C_t}{\exp(\epsilon p^j)}$. Finally, given $\mathbf{V}_t$, and $\mathbf{V}_{t+1}$ we can construct $\mathbf{\Pi}_t$ and $\tilde{\mathbf{v}}_t$, and thus $\mathbf{\Lambda}_t$ and $\mathbf{O}_{t+1}$. Therefore the variables in $\overrightarrow{X}_t$ and z_t are indeed sufficient to evaluate the system (55)-(59).

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

$$E_t\mathcal{A}\Delta\overrightarrow{X}_{t+1} + \mathcal{B}\Delta\overrightarrow{X}_t + E_t\mathcal{C}z_{t+1} + \mathcal{D}z_t = 0 \tag{61}$$

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.

ONLINE APPENDIX B: SEQUENTIAL STATEMENT OF THE OPTIMIZATION PROBLEM

In this appendix, we discuss a sequential representation of the firm's partial-equilibrium decision problem. In this representation, we think of the firm as choosing a plan contingent on any possible history up to a given time T. In particular, it must consider histories $(\tilde{p}^T, a^T, \Xi^T)$ incorporating its own prices and productivity, and aggregate states, where superscripts indicate time series: $\tilde{p}^T \equiv (\tilde{p}_0, \tilde{p}_1, \ldots, \tilde{p}_T)$, and likewise for a^T and Ξ^T.

By repeated substitution in the recursive problems (13)-(14), we can derive the following sequential optimization problem:

$$o_0(\tilde{p}_0, a_0) = \max_{\pi_t^\dagger \in \Delta(\Gamma^p(\tilde{p}_t))} E_0 \sum_{t=0}^{\infty} q_{0,t} \left[\int u_t(p_t, a_t) d\pi_t^\dagger(p_t) - \kappa w_t \mathcal{D}\left(\pi_t^\dagger || \theta^\dagger(\cdot|\tilde{p}_t)\right) \right] \tag{62}$$

$$\text{s.t. } \tilde{p}_t = p_{t-1} - i_{t-1,t}, \quad \text{and} \quad \int d\pi_t^\dagger(p) = 1 \text{ for all } t. \tag{63}$$

Here E_0 refers to an expectation calculated under the dynamics of the firm's productivity shock a_t and the dynamics of the aggregate state Ξ_t. The discount factor is $q_{0,t} \equiv \Pi_{s=0}^{t-1} q_{s,s+1}$, where $q_{t,t+1} \equiv \beta \frac{P_t C_{t+1}^{-\gamma}}{P_{t+1} C_t^{-\gamma}}$; we assume discount factors satisfy $\sum_{t=0}^{\infty} q_{0,t} < \infty$.

The choice problem here is to be understood as choosing a function $\pi_t^\dagger(\tilde{p}^t, a^t, \Xi^t)$ conditional on each history $(\tilde{p}^t, a^t, \Xi^t)$ of length t. Here $\tilde{p}_t = p_{t-1} - i_{t-1,t}$ represents the log real price at the beginning of t, prior to the choice of $\pi_t^\dagger$, and $\Delta(\Gamma^p(\tilde{p}_t))$ is the set of increasing functions f satisfying $f(\min \Gamma^p(\tilde{p}_t)) \geq 0$; constraint (63) ensures that f is a *c.d.f.* Notice that if two functions π_t^1 and π_t^2 both lie in the set $\Delta(\Gamma^p(\tilde{p}_t))$ and both integrate to one over $\Gamma^p(\tilde{p}_t)$, then so does any convex combination of those functions. The same argument can be made for history-contingent plans. Therefore we conclude that the choice set in problem (62)-(63) is convex.

Moreover, the objective function contains two terms (for each t): expected profits, which are a linear function of $\pi_t^\dagger$, minus decision costs, which are a convex function of $\pi_t^\dagger$. Therefore the objective function of (62)-(63) is concave.

We should also be careful to check that the constraint set of problem (62)-(63) is non-empty, and that the objective function is finite-valued. Note that the strategy $\pi_t^\dagger = \theta^\dagger(\cdot|\tilde{p}_t)$ is feasible, and has zero decision cost, attaining the value

$$\underline{o}_0(\tilde{p}_0, a_0) \;\equiv\; E_0 \sum_{t=0}^{\infty} \int q_{0,t} u_t(p_t, a_t) d\theta^\dagger(p_t|\tilde{p}_t) \;>\; -\infty.$$

This lower bound is finite because sets Γ^a and Γ^p are assumed bounded, and the profit function in (8) is continuous. On the other hand, the value of (62)-(63) is bounded above by

$$\overline{o}_0(a_0) \;\equiv\; E_0 \sum_{t=0}^{\infty} q_{0,t} \max_p u_t(p, a_t) \;<\; \infty.$$

Therefore the value of (62) is bounded: $\underline{o}_0(\tilde{p}_0, a_0) < o_0(\tilde{p}_0, a_0) < \overline{o}_0(a_0)$.

Thus (62)-(63) maximizes a concave function over a non-empty, convex set, attaining a finite value. Hence there can be at most one solution to the first-order conditions, and if such a solution is found, it represents a solution to the optimization problem (62)-(63). Indeed, the first-order conditions yield the same solution that we found for the recursive problem (13)-(14).[5]

[5] Deriving the first-order conditions of (62)-(63) is tedious, so we omit them here, but they closely resemble those of (13)-(14). Where the value $v_t(p_t, a_t)$ appears in (15), we instead find a term of the form $E_t \sum_{s=0}^{\infty} q_{t,t+s} prob(p_{t+s} = p_t - i_{t,t+s}) u_{t+s}(p_t - i_{t,t+s}, a_{t+s})$, where $i_{t,t+s} \equiv \ln(P_{t+s}/P_t)$. This term represents the discounted sum of profits at all future times $t+s$ conditional on the nominal price set at time t remaining unadjusted at time $t+s$.

ONLINE APPENDIX C: ROBUSTNESS OF THE RESULTS

C.1 Changing the decision cost function

In our model, the precision of price choices is measured by comparing the firm's chosen price distribution to an exogenously fixed "default" distribution. We have run extensive simulations to explore whether our results are robust to changes in the assumed default distribution. In summarizing our conclusions, it is useful to distinguish two properties of the default distribution – its functional form, and its standard deviation. We find that our results are qualitatively and quantitatively robust to changes in both of these properties. There is a simple reason for this: we estimate that decision costs are low. Since deviating from the default distribution is not very costly, the precise form of that default has little impact on our results.

If we generalize the uniform default probabilities assumed in our benchmark parameterization, then the weighted logit (21) no longer reduces to the unweighted logit (43). To see how the results differ, compare the blue and black lines in the price change histogram (left panel) and impulse responses (middle and right panels) shown in Figure C.1. The black lines represent the benchmark uniform specification; the blue lines instead assume a truncated normal default distribution, with the same standard deviation. The results are almost identical. We have computed several other examples which show that the form of the default distribution has very little effect. In other words, firms' optimization, represented by $\exp(v/(\kappa w))$ in the logit formula (21), is powerful enough that reweighting by a different distributional form θ hardly matters.

The fact that we define the default distribution on a discrete grid is likewise irrelevant for the results. Computation on a discrete grid is a matter of numerical necessity. However, making this grid much finer has entirely negligible effects, both on the steady state and on the dynamic implications of the model.

Another change that might seem especially relevant would be to allow the default distribution to vary over time by recentering it on the previous nominal price. We simulated a speci-

Figure C.1: Uniform versus normal default distributions.

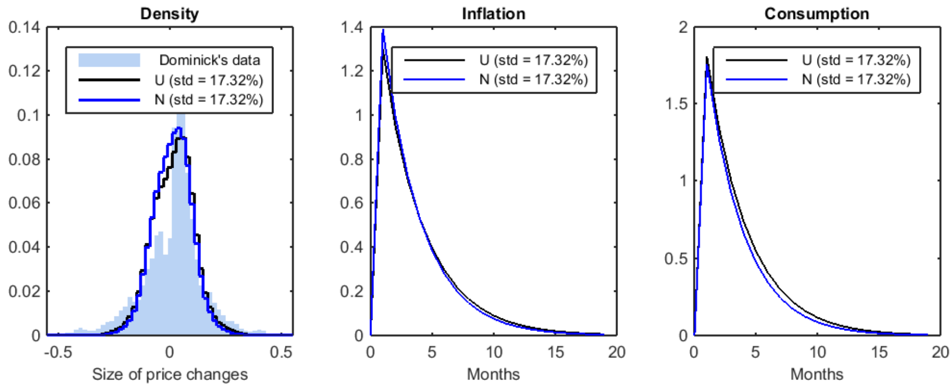

Notes:
Comparing benchmark uniform default distribution (as in paper) with truncated normal default distribution.
Left panel: Histogram of nonzero log price adjustments (Dominick's data shown as blue bars).
Middle and right panels: impulse responses to money growth shock with monthly autocorrelation 0.8.

fication of this type, assuming a truncated normal default $\theta(p|\tilde{p})$ centered around $\tilde{p}$, the price prior to adjustment. Figure C.2 compares this specification to the unchanging uniform distribution used in our benchmark calculations.[6] The recentered normal specification (blue) implies a somewhat more symmetric histogram, a small increase in the adjustment hazard, and a resulting small decrease in the real effects of a money shock. But overall the differences are minor.

The effects of increasing the standard deviation of the default distribution are shown in Figure C.3. The benchmark results are shown in black; the effects of making the grid 100% wider (while fixing the grid step size) are shown in green.[7] Doubling the grid width makes the histogram more strongly bimodal (improving fit in the center while making it worse in the tails); the frequency of adjustment decreases from 10.2% to 6.6% monthly (because choosing from a wider range of prices amounts to a more difficult decision problem) and therefore the real effects of the money shock increase slightly.

Thus, changing the standard deviation of the default distribution has a small but nontrivial

[6] Figure C.2, like Figure C.1, changes the form of the default distribution without altering its standard deviation.
[7] We have also studied the effects of increasing the standard deviation of the default distribution when the default is a truncated normal. The (small) effects are similar to those shown in Figure C.3 for the uniform case.

Figure C.2: Uniform versus recentered normal default distributions.

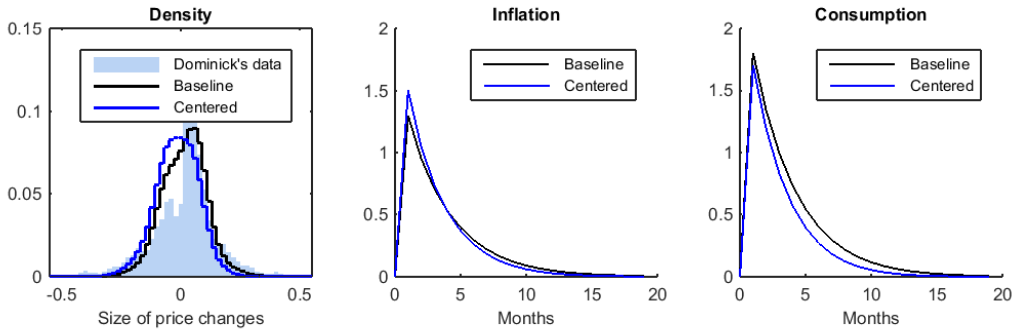

Notes:
Comparing benchmark uniform default distribution (as in paper) with truncated normal default distribution $\theta(p|\tilde{p})$ centred around pre-adjustment price $\tilde{p}$.
Left panel: Histogram of nonzero log price adjustments (Dominick's data shown as blue bars).
Middle and right panels: impulse responses to money growth shock with monthly autocorrelation 0.8.

impact on the results. But would it make any difference if we treated the standard deviation of the default as another free parameter to be estimated? This question is also addressed in Figure C.3, where the red line shows the results of jointly estimating κ, $\bar{\lambda}$, and the width of the price grid Γ^p to maximize our estimation criterion (33). The estimation favors a harder decision problem than we assumed in our benchmark calibration (the preferred grid is slightly more than twice as wide as the benchmark grid) but this is compensated by a slightly less error-prone and substantially quicker decision process ($\kappa = 0.17$ and $\bar{\lambda} = 0.35$, in contrast to the previous values $\kappa = 0.18$ and $\bar{\lambda} = 0.22$). The implied adjustment frequency is again 10.2% monthly, with the result that the impulse responses are almost indistinguishable from those in the benchmark specification (the black benchmark IRF is almost invisible under the red IRF resulting from estimating the width of the grid; likewise the black and red price adjustment histograms are almost identical). So while widening the grid, *ceteris paribus*, slightly increases monetary nonneutrality, estimating the grid width jointly with our other parameters gives results nearly identical to our benchmark parameterization.

Why are the impulse responses unchanged when we reestimate the model? All of the robust-

Figure C.3: Widening the support of the uniform default distribution.

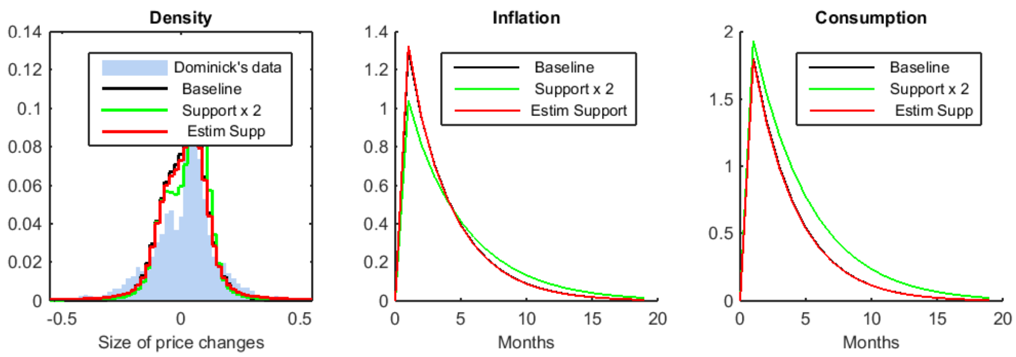

Notes:
Black: uniform default distribution (benchmark from paper). Green: uniform default with 100% wider support.
Red: reestimating model with width of support of default as a free parameter. (Grid step size fixed.)
Left panel: Histogram of nonzero log price adjustments (Dominick's data shown as blue bars).
Middle and right panels: impulse responses to money growth shock with monthly autocorrelation 0.8.

ness exercises that we have run suggest that as long as our model matches the 10.2% adjustment hazard of our estimation criterion, the degree of monetary nonneutrality is virtually unchanged. Obviously this does not mean that *all* models with a 10.2% adjustment hazard are equivalent; the Calvo model implies much larger real effects, as our paper shows. But our error-prone model is very robust to changes in the specification of the default distribution, as long we parameterize the model to fit the estimation criterion.

To summarize, changing the shape of the default distribution is not quantitatively relevant for our results. Neither is treating its standard deviation as a free parameter to be estimated. What matters is that price decisions are somewhat noisy (helping fit the microdata) and timing decisions are also somewhat noisy (diminishing the selection effect and generating higher nonneutrality). The precise form of the noise is not at all essential for these conclusions.

C.2 Extending the model

Throughout the paper we have studied a stripped-down general equilibrium structure in order to focus primarily on the role of price stickiness. However, our framework can readily be extended to incorporate a more complete macroeconomic environment. Building a full medium-scale DSGE model is beyond the scope of this paper, but in this section we consider two especially relevant extensions. Our main conclusions about state-dependent nominal rigidity are unaltered.

On one hand, there is no need to restrict monetary policy to a money growth rule. Here we instead consider a Taylor-style interest rate rule of the form

$$i_t = \phi_i i_{t-1} + (1-\phi_i)\phi_\pi \left[\pi_t - \ln(\mu)\right] + \epsilon_t^i,$$

where i_t is the net nominal interest rate, π_t is the net inflation rate; μ is the steady-state target inflation rate, ϕ_i is an interest rate smoothing parameter; ϕ_π controls the strength of monetary policy reaction to inflation; and ϵ_t^i is an interest rate shock. This rule replaces the money supply equation (26).

Second, while nominal rigidity is central to generating real responses to purely nominal disturbances, it is likely that multiple forms of real rigidity also play a role in propagating shocks. In particular, Blanco (2017) imposes a production function with intermediate inputs, as a realistic and tractable source of real rigidities that can reinforce the effects of nominal rigidity. Here we modify the production function as follows:

$$Y_{it} = A_{it} N_{it}^{1-\eta} M_{it}^{\eta},$$

where M_{it} denotes the intermediate inputs used in the production of the differentiated final

Figure C.4: Impulse responses to a Taylor rule shock, with intermediate inputs in production.

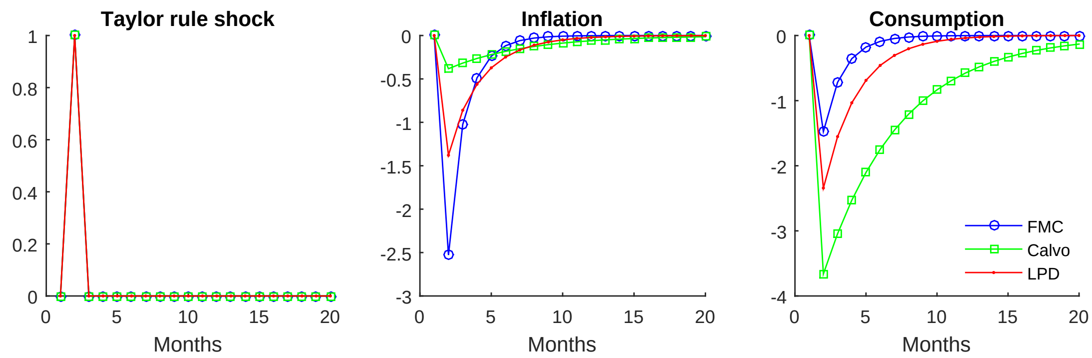

Notes:
Impulse responses of inflation and consumption to an *i.i.d.* interest rate shock (percentage points). Comparing models of nominal rigidity.

goods Y_{it}. The goods market clearing condition becomes

$$Y_t = C_t + \int M_{it} di\,;$$

C_t is now replaced by Y_t in equation (31).

We set $\phi_i = 0.9$, $\phi_\pi = 2$, and $\eta = 1/3$. The impulse responses to a monetary policy shock are shown in Figure C.4, which compares our nested benchmark specification (marked LPD), with FMC and Calvo models that likewise allow for a Taylor rule and real rigidities. The main purpose of this exercise is simply to show that the central results of our paper go through in this extended version. Namely, the nested specification produces real effects that fall between those of FMC and Calvo, and the relative degree of nonneutrality across these frameworks is quantitatively similar to what we found previously. When calibrating a model for applied purposes, real rigidities and a more realistic description of monetary policy are relevant elements to include in the analysis. But the difference in nonneutrality implied by our framework, relative to alternative models nominal rigidity, appears robust to these extensions.