Price Rigidities in U.S. Business Cycles *

Camilo Morales-Jiménez
Federal Reserve Board

Luminita Stevens
University of Maryland

March 19, 2025

We present a structurally estimated time series of US aggregate price rigidities from 1978 to 2023. Our estimation uses a novel generalized model of price setting with frictions in both timing of price changes and reset price choices. We microfound these frictions with information costs and menu costs, which are jointly estimated with other shocks and frictions using macro series and the evolution of the cross-sectional distribution of price changes over time. Estimated menu costs are small, while information costs are larger and more volatile. Price rigidities stem primarily from inaccurate rather than infrequent adjustment and show substantial medium-cycle volatility. The importance of frictions in reset prices challenges the conventional wisdom on nominal rigidities and monetary stabilization policy.

*The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of anyone else associated with the Federal Reserve System. We thank our discussants Anton Nakov and Gee He Hong, as well as Larry Christiano, Mark Gertler, Cosmin Ilut, Chris Wolf, Mike Woodford, and colleagues at the Bank of Canada, the Board of Governors, the European Central Bank, Chicago Booth, Duke, Maryland, MIT, NYU, Northwestern, Notre Dame, Indiana, the Barcelona Summer Forum, Midwest Macro Meetings, NBER Summer Institute, and other institutions and conferences for helpful feedback and suggestions. Contact: camilo.moralesjimenez@frb.gov, stevens7@umd.edu.

1 Introduction

How strong are nominal rigidities in the US economy? Have prices become more responsive to shocks over time? Does price rigidity vary over the business cycle? These questions are at the core of monetary economics: Rigidities in prices change how the economy adjusts in response to supply and demand shocks, and they determine the extent to which monetary policy can stabilize fluctuations in inflation and economic activity. Consider the debate regarding the likelihood of a soft or hard landing for the US economy following the inflation surge of 2021-2022. That likelihood depends in part on how flexibly prices adjust – first to the inflationary shocks themselves, and second to the interest rate increases undertaken by the Federal Reserve in its efforts to lower inflation. Despite the large literature measuring and modeling price rigidities, uncertainty and disagreement persist, reflecting both the difficulty of constructing model-free empirical estimates and the lack of clarity regarding what frictions are most relevant when modeling this rigidity.

This paper provides a structurally estimated time series for the degree of nominal price rigidities (NPR) in the United States, documenting that it has varied substantially in recent decades. We proceed in three steps: First, we propose a generalized model of nominal rigidities that nests existing paradigms. Next, we undertake a Bayesian estimation where variation in nominal frictions is identified using the dynamics of the empirical distribution of price changes underlying the US Consumer Price Index (CPI). Finally, using the frictions extracted from the Bayesian estimation, we compute two time-varying measures of NPR: the consumption response to a monetary policy shock and the sacrifice ratio of lowering inflation by one percentage point via a contraction in demand.

Our model allows for rigidity and imprecision in both the timing of price changes and the choice of what reset price to set when adjusting. In this way, we depart from prior quantitative models in which the nominal rigidity arises exclusively from the infrequent adjustment of prices. We microfound these frictions in a unified framework that combines rational inattention (Sims, 2003; Woodford, 2009) variable costs with fixed menu costs, allowing firms to optimize both the timing and precision of their choices. Firms are rational in that they fully understand their environment and maximize well-defined objectives; but they can only learn about the realized state in real time at cost. In our context, a firm's choice is what price to charge, which in turn determines its demand and production inputs. Since repricing is subject to a fixed cost, firms first decide if they want to change their existing price. They do so based on an imprecise awareness of the state of the economy, which generates not only infrequent adjustment, but a smooth adjustment hazard à la Caballero & Engel (1999, 2007), with the decision to change prices only probabilistically tied to the value of adjusting prices in a given state.[1]

When adjusting, firms decide how much to learn about the right price to set. Unlike in standard models with sticky prices, we do not assume that firms know the state perfectly when adjusting. Instead, they choose how much to learn about the optimal reset price, which amounts to choosing how strongly the reset price conditions on the state in real time. As a result, reset prices are only probabilistically tied to the full information optimal price. This generates state-dependent inefficient price dispersion even conditional on price adjustment, giving rise to additional nominal rigidity. Overall, the model spans pricing behavior from full state contingency to no state contingency in terms of both timing of price changes and reset prices, and we let the data pin down the degree of state dependence along each margin.[2]

Allowing imprecision in reset prices is motivated by a large body of evidence that economic choices are based on dispersed, imprecise beliefs and are only partially related to optima in a variety contexts. Many studies have documented dispersion and stochasticity in actions and forecasts *conditional on adjustment*, in both surveys and incentivized controlled laboratory experiments, from the seminal work of Carroll (2003), Mankiw, Reis & Wolfers (2003), Coibion & Gorodnichenko (2012), to the more recent work of Cavallo, Cruces & Perez-Truglia (2017), Magnani, Gorry & Oprea (2016), Khaw, Stevens & Woodford (2017), Coibion, Gorodnichenko & Kumar (2018), Coibion, Gorodnichenko & Ropele (2020), Dean & Neligh (2023), Khaw, Li & Woodford (2021), Fuster, Perez-Truglia, Wiederholt & Zafar

[1]See Caballero & Engel (2007); Costain & Nakov (2011) for early discussions of the desirability of a smooth adjustment hazard for prices, and Blanco, Boar, Jones & Midrigan (2024); Gagliardone, Gertler, Lenzu & Tielens (2025) for applications to the 2021 inflation surge.

[2]Another class of models with stochastic reset prices is the class of search-theoretic models of imperfect competition, which can generate both price dispersion and price rigidity (Burdett & Menzio, 2017; Head, Liu, Menzio & Wright, 2012). This alternative class of models has very different implications for non-neutrality and monetary policy, and yet a discussion of the identification of search frictions versus rational inattention remains an open question.

(2022), Angeletos, Huo & Sastry (2021).[3] Our model allows us to assess the relevance of such inaccuracy in U.S. price data, and its consequences for monetary non-neutrality.

We back out the implied pricing frictions by incorporating the dynamics of the distribution of price changes in the estimation, in addition to data on real economic activity, interest rates, and inflation, and a rich set of shocks. The distributional moments provide rich information that helps identify the time-varying pricing frictions in the model. This approach contrasts with prior work, which uses the averages of various pricing moments as targets, yielding an *average* degree of NPR. The distributional moments are built from the micro data underlying the U.S. CPI and cover the period from January 1978 to March 2023. These series were created by Nakamura, Steinsson, Sun & Villar (2018) and extended to 2023 by Montag & Villar (2023).[4] We incorporate this cross-sectional variation in our estimation, much like the work on heterogeneous households has begun using household income and wealth distributional data to inform models of the aggregate economy.[5] Our approach opens the door to further work on heterogeneous firms with other frictions and distributional moments (*e.g.*, investment costs or financial frictions).

Matching the time series of the distributional price moments requires variation in the pricing frictions over time. A wide range of shocks to firms' desired prices are simply not enough to generate the kind of volatility in the distribution of price changes that we see in the data. Most of the time, the distributional data suggest a menu cost model with time-varying errors in repricing. Moreover, volatility in the cost of repricing is strongly correlated with the measures of exogenous uncertainty of Jurado, Ludvigson & Ng (2015) and Ludvigson, Ma & Ng (2021): In times of rising aggregate uncertainty, firms start paying close attention to the prices they set, which reduces monetary non-neutrality. Conversely, when fundamental aggregate uncertainty is low, firms are less attentive to the environment, and monetary non-neutrality rises. We interpret this as variability in the attention and efficiency with which firms process information about the state of the economy (*e.g.*, Flynn & Sastry, 2024).

[3]See also the reviews of Assenza, Bao, Hommes & Massaro (2014) and Fuster & Zafar (2023). Our friction is also related to models of misallocation and misoptimization due to information frictions, *e.g.*, David, Hopenhayn & Venkateswaran (2016); David & Venkateswaran (2019); Flynn & Sastry (2024).

[4]We thank Daniel Villar for sharing the time series of these moments.

[5]For example, Auclert, Rognlie & Straub (2020); Bayer, Born & Luetticke (2020); Bilbiie, Primiceri & Tambalotti (2023)

Variability in the estimated frictions in turn generates variation in NPR. For each period in our sample, we compute what the cumulative response of consumption would have been in each period, in reaction to a monetary policy shock, given the pricing frictions we estimate for that period. To our knowledge, this is the first such structurally estimated series. According to our results, a shock to the monetary policy rule of of 25 basis points yields a cumulative change in consumption equal to 0.12 percent of annual steady-state consumption, on average. This represents 80% of the response a Calvo model would predict when calibrated to the same frequency and size of price changes. Underlying this average, we find considerable fluctuations, with nominal rigidities peaking in the mid-1990s and mid-2010s. We find no consistent pattern during recessions, casting some doubt on the hypothesis of increased price flexibility during recessions, and also no clear trend.

Another way to understand the severity of these frictions is via the sacrifice ratio, which asks what is the output loss associated with lowering inflation by one percentage point via contractionary monetary policy. A high sacrifice ratio indicates a high cost of disinflation. We find a time-varying sacrifice ratio that shows considerable volatility. Most of the time, the sacrifice ratio fluctuates below 0.5 in our sample. But two episodes stand out: First, we estimate an unusually high sacrifice ratio (above 2) during the 2011-2016 years of the first ELB period. During that time, price rigidities were very high, making output highly sensitive to monetary policy. Second, by the time of the 2021 inflation surge, the sacrifice ratio had fallen to a mere 0.03. Pricing frictions had come down dramatically, making monetary policy highly potent in controlling inflation.

The sharp increase in price flexibility that we estimate starting in 2016 is particularly noteworthy. This change coincides with both the Federal Reserve's liftoff of its policy rate from the effective lower bound and with an acceleration of the use of machine learning algorithms to determine real-time pricing (Adams, Fang, Liu & Wang, 2024).[6] Price rigidities had been falling well before the COVID-19 shocks and had already fallen below historical lows by January 2021, when CPI inflation was only 1.5%. In hindsight, this evidence suggests that we might have expected any inflationary shocks to be met with a sharper inflation

[6] While the former change suggests transitory heightened attention to price-setting, the latter could be a sign of a structural change worth monitoring with more disaggregated data on AI pricing practices.

response post-2016 versus pre-2016. It may explain why inflation surged so rapidly in 2021 and then declined so sharply in 2022. It also points to the high value of monitoring distributional pricing data in real time. According to our estimates, when the Federal Reserve started targeting higher interest rates in the first quarter of 2022, the cross-sectional distribution of price changes in the CPI implied that lowering inflation from its level at the time (approximately 8.5%) down to 2% would, all else equal, be associated with a contraction of less than 0.2 percentage points of consumption. But due to the varying severity of the pricing frictions, the same disinflation would have cost 2.7 percentage points of consumption in 1994. In contrast, according to the Calvo model, disciplined only by a time-varying frequency of price changes, the consumption loss would have been 0.5 percentage points in 2022 and 2 percentage points in 1994.

Underlying our estimated NPR series are several results that shed new light on the sources and dynamics of nominal rigidity. These results depart meaningfully from the conventional wisdom embedded in standard DSGE models with nominal frictions and reflect the interaction between firms' timing decisions and their repricing choices. First, we estimate a strongly state-dependent probability of adjustment (especially when prices are high relative to the expected reset price) and a modest menu cost that accounts for only a small fraction of both adjustment costs and the total degree of price rigidity. Conversely, we estimate substantial inaccuracy in repricing that drives most of the NPR on average and over time. In the parlance of literature, firms only partially close their price gaps when adjusting. It seems that firms are able to determine quite accurately when their prices have become obsolete, but they have more difficulty determining the right reset price.

Second, firm-level mistakes do not average out with aggregation. When firms adjust, they do not fully close the gap to the optimal price, so that the intensive margin of adjustment moves less than one-for-one with the shock. This incomplete closing of the pricing gap challenges standard models that assume perfect repricing conditional on adjustment. It breaks the connection between *adjustment* and *flexibility*: Even if prices are not very "sticky," in the sense that they are changing over time, they nevertheless only partially respond to economic conditions. Underscoring this dichotomy, we show that under some conditions, a higher frequency of price changes can be associated with a higher degree of NPR, generating a paradox of flexibility. Moreover, we show that this implies that the Calvo model no longer determines the upper bound on monetary non-neutrality.

Third, we find that imprecision in repricing is itself a source of infrequent adjustment: Most of the inaction in prices reflects uncertainty about the right price to set, rather than an unwillingness to pay the adjustment cost. In our model, firms understand that they risk picking the wrong price, so they often choose to forgo price changes altogether. This uncertainty provides a potent micro-foundation for inaction that is quantitatively significant: Introducing a very small degree of errors in pricing can nearly halve the frequency of adjustment.[7]Moreover, the degree to which firms tolerate errors in pricing interacts with their tolerance for errors in the timing of price changes. Under certain conditions, this interaction can rationalize Calvo-like price-setting as an optimal way to economize on repricing costs.

Lastly, our estimation method contributes methodologically to the literature that has sought to introduce heterogeneity in DSGE models. To our knowledge, this is the first Bayesian estimation of a model with rationally inattentive firms using the sequence-space Jacobian (SSJ) method of Auclert, Bardóczy, Rognlie & Straub (2021). Moreover, our sample includes two periods in which the effective lower bound (ELB) was binding on the federal funds rate. To handle this, we also show how to do maximum likelihood evaluation with SSJ and an occasionally binding constraint on the nominal interest rate, by adapting the methods proposed by Guerrieri & Iacoviello (2015) and Kulish, Morley & Robinson (2017).

Related Literature We build on work that has sought to use **moments from the micro pricing data** to develop micro-founded models of nominal rigidities. For example, while the frequency of price adjustment is a sufficient statistic for the canonical Calvo (1983) model, Alvarez, Le Bihan & Lippi (2016) prove that under certain conditions, frequency relative to kurtosis pins down non-neutrality in a wide class of menu cost models, Berger & Vavra (2018) argue for the additional empirical relevance of the standard deviation of price changes, and Luo & Villar (2021) suggest also taking into account the skewness of price changes. Empowered by detailed empirical analyses of micro pricing patterns starting with the seminal work of Bils & Klenow (2004), Nakamura & Steinsson (2008), and Campbell & Eden (2014), a

[7]Imprecision driving inaction is also discussed in Costain & Nakov (2015) and Ilut, Valchev & Vincent (2020), though arising via different mechanisms (namely, control costs and ambiguity aversion).

wave of menu cost models (*e.g.*, Alvarez & Lippi, 2014; Golosov & Lucas Jr, 2007; Midrigan, 2011; Nakamura & Steinsson, 2010; Vavra, 2013) have studied the contribution to NPR of different moments of the price change distribution. Our results underscore the importance of studying how time variation in these moments relates to time-varying price flexibility.

Second, our framework nests both models that generate non-neutrality via **infrequent** price adjustment and models that generate non-neutrality via the **incomplete** response of individual prices to shocks. In the first category, our model belongs to the class of generalized Ss models (Caballero & Engel, 2007), in which the probability of adjustment varies smoothly with the value of adjusting, as in Dotsey, King & Wolman (1999); Woodford (2009), and Karadi, Amann, Bachiller, Seiler & Wursten (2023). In the second category, our model belongs to the class of models with imprecise price-setting (Woodford, 2003), in particular work that operationalizes tools from information theory (Afrouzi, 2020; Afrouzi & Yang, 2021; Maćkowiak & Wiederholt, 2009; Matějka, 2015; Stevens, 2020; Turen, 2023). We nest these two classes and allow the estimation to speak to their relative importance in generating non-neutrality. Two important precursors to our work are the control cost pricing model of Costain & Nakov (2019) and the inattentive forecasting model of Khaw et al. (2017). Khaw et al. (2017) model rationally inattentive adjustment in both the timing of adjustment and the choice of a new forecast for individual decision-makers tracking the realizations of a slow-moving random variable. The model is then estimated on individual data from a controlled laboratory experiment. Costain & Nakov (2019) model price-setting firms that are subject to control costs in timing and repricing, and they use steady state moments of the distribution of price changes to estimate the severity of control costs on average. However, since the control costs that introduce errors in pricing centered around the optimal price, they largely average out with aggregation, playing a limited role in generating NPR, unlike here.

Finally, by allowing both information frictions and nominal adjustment frictions to play potentially distinct roles in generating nominal rigidity in response to shocks, our paper relates to work that bridges these two approaches to endogenizing pricing frictions: Angeletos & La'O (2009) and Nimark (2008) study the interaction between Calvo (1983) price-setting and dispersed information a la Woodford (2003), while Klenow & Willis (2007) models a sticky-information version of menu cost pricing. Melosi (2014) estimates that imperfect common knowledge a la (Woodford, 2003) fits U.S. inflation and output time series better than a model with Calvo frictions alone, and Alvarez, Lippi & Paciello (2011) present a theoretical analysis of price adjustment in the presence of menu costs and (fixed) information costs a la Reis (2006), and they also emphasize the interaction between the two sources of nominal rigidity. We also contribute to work that seeks to measure the costs of pricing frictions (Anderson, Jaimovich & Simester, 2015; Bandeira, Castillo-Martınez & Wang, 2024; Gorodnichenko & Weber, 2016) and the severity of information frictions (Coibion & Gorodnichenko, 2015). Our work is also complementary to Carvalho, Dam & Lee (2020), who study the degree of real rigidities and heterogeneity in price stickiness.

2 Theoretical Framework

To provide a credible structural estimation of the degree of nominal frictions over time, we need a model that can accommodate different types of uniquely identifiable pricing frictions. This section presents such a model, which allows for frictions in both the timing of price adjustment and the repricing decision. Information costs and menu costs are the sources of these frictions and generate volatile, infrequently updated prices. We place the information and adjustment frictions on monopolistically competitive retailers while retaining the assumption of full information, flexible adjustment for other agents in the economy. We first describe the retailers' problem, and then close the economy with a representative household, competitive intermediate goods producers, and fiscal and monetary authorities.

2.1 Pricing Frictions

A continuum of retailers j sell differentiated varieties and are monopolistically competitive price-setters in their product market and competitive price-takers in the market for their production input.

Operating Profits Each retailer faces demand $y_{jt} = p_{jt}^{-\varepsilon_t} Y_t$, with $\varepsilon_t > 1$ denoting the time-varying elasticity of substitution, Y_t final aggregate demand, and $p_{jt} = P_{jt}/P_t$ the good's relative price, with $\left(\int p_{jt}^{1-\varepsilon_t} dj\right)^{\frac{1}{1-\varepsilon_t}} = 1$. The production function is $y_{jt} = e^{a_{jt}} x_{jt}$, where a_{jt}

is an AR(1) process for idiosyncratic productivity and x_{jt} is a homogeneous intermediate input with real price p_t^x. Real operating profit per period is

$$\pi_{jt}^r = p_{jt} y_{jt} - p_t^x x_{jt}. \tag{1}$$

Price Setting Costs In each period, firms decide whether or not to update their price, and if so, what price to set. They then produce whatever quantity is needed to satisfy demand at the current price. These decisions are made subject to both information costs and a fixed price adjustment cost κ. For tractability, the decision of whether or not to adjust prices and the decision of what price to set conditional on adjustment are treated as two separate decisions, and the information used to make the first decision is not freely available to inform the second decision. The cost of information is assumed to be linear in the information acquired in order to make each decision,

$$\mathcal{C}_{jt}^a = \theta^a \mathcal{I}_{jt}^a \qquad \text{and} \qquad \mathcal{C}_{jt}^p = \theta^p \mathcal{I}_{jt}^p, \tag{2}$$

where θ^a is the unit cost of making a more informed decision about whether or not to adjust prices, θ^p is the unit cost of making a more informed reset price choice when adjusting prices, and $\mathcal{I}_{jt}^a$ and $\mathcal{I}_{jt}^p$ measure how much information is acquired for each decision. We allow for (but do not impose) potentially different unit costs, since they may reflect different managerial marginal costs of attention.

Value of the Firm Firms acquire information and make pricing decisions to maximize

$$\mathcal{V}_0 = E_{j0} \sum_{t=0}^{\infty} M_{0,t} \left[\pi_{jt}^r - \mathcal{C}_{jt}^a - \delta_{jt} \left(\kappa + \mathcal{C}_{jt}^p \right) \right], \tag{3}$$

where $M_{0,t}$ is the stochastic discount factor used to discount real profit streams from date t to date 0, δ_{jt} is an indicator equal to 1 if the firm picks a new price in period t and 0 otherwise, and κ is the fixed cost of repricing. If the firm does not revise its price in the period, it continues with its existing nominal price. Note that if θ^a and θ^p are zero, the firm's problem collapses to a canonical (full information) menu cost model. On the other hand, if θ^a and κ are zero, it collapses to a version of Woodford's (2003) imperfect common knowledge model, with variable signal precision.

Acquiring Information Firms fully understand the structure of their environment (payoffs, shock processes, markets), but they must expend resources to learn the realizations of stochastic variables in real time. Acquiring information about real-time market conditions is formalized as a choice that can be optimized using tools from information theory (Shannon, 1948, 1959), subject to a cost per unit of uncertainty reduction.[8] The choice of how much information to obtain amounts to choosing how much each decision conditions on the realized state in real time, relative to the decisions the firm could make based on beliefs it has for free, given its knowledge about the structure of its environment.

The amount of information that the firm expects to use in order to make its adjustment decision in some period t is

$$\mathcal{I}^a_{jt} = E_t\left\{\mathcal{D}\Big(\Lambda_{jt} \parallel \bar{\Lambda}\Big)\right\}, \tag{4}$$

$$\mathcal{D}\big(\Lambda \parallel \bar{\Lambda}\big) = \Lambda \ln\left(\frac{\Lambda}{\bar{\Lambda}}\right) + \big(1-\Lambda\big)\ln\left(\frac{1-\Lambda}{1-\bar{\Lambda}}\right), \tag{5}$$

where Λ_{jt} denotes the probability that the firm adjusts its price in period t, *after* obtaining information about the realized state, $\bar{\Lambda}$ is the *reference probability* of adjustment, based on the firm's beliefs *before* obtaining current information, $\mathcal{D}$ is the Kullback-Leibler (KL) divergence of the choice distribution from the reference distribution, and expectations integrate over the joint distribution of idiosyncratic and aggregate states that the firm could face in period t.[9] Hence, the contribution to the firm's cost of conditioning the adjustment decision on a period's realized state is proportional to the divergence of Λ from $\bar{\Lambda}$ in that state. The trade-off facing the firm reflects the fact that the more Λ conditions on the realized state, the more it deviates from $\bar{\Lambda}$, and hence the higher is its cost.

Analogously, for the pricing decision, the amount of information that the firm expects to

[8]While Sims's (2003) original formulation of rational inattention assumed optimization subject to a fixed cap on information, we use here the variable cost setup popularized by Woodford (2009).

[9]The KL divergence gives a measure of how "far off" one would be, on average, if they assumed the first distribution when the true distribution were in fact the second distribution.

acquire in some period t in order to decide what price to set, conditional on adjustment, is

$$\mathcal{I}_{jt}^{p} = E_t \left\{ \mathcal{D}\Big(f_{jt}(p) \parallel \bar{f}(p) \Big) \right\}, \tag{6}$$

$$\mathcal{D}\left(f(p) \parallel \bar{f}(p) \right) = \int f(p) \ln \left(\frac{f(p)}{\bar{f}(p)} \right) dp, \tag{7}$$

where $f_{jt}(p)$ is the probability that the firm sets its price equal to p conditional on the information it acquires about the realized state, and $\bar{f}(p)$ is the reference probability of setting price p, based on the firm's beliefs about the right price to set prior to obtaining current information. As is the case for the adjustment decision, $\mathcal{D}$ is the KL divergence of the choice distribution from the reference distribution, and expectations integrate over the joint distribution of prices, productivities, and aggregate states. Hence, the contribution to the total information flow of conditioning the pricing decision in a period on that period's state is proportional to the divergence of f_{jt} from $\bar{f}$ in that state. Note that $\mathcal{D}\left(\bar{f} \parallel \bar{f}\right) = 0$, so the cost of using the reference distribution is zero.

Reference Distributions We assume that the firms' reference probability of adjustment $\bar{\Lambda}$ in any period is the equilibrium frequency of adjustment in the steady state,

$$\bar{\Lambda} = \int \Lambda_{ss}(\tilde{p}, a)\, \tilde{\Omega}_{ss}(\tilde{p}, a)\, da\, d\tilde{p}, \tag{8}$$

which integrates the steady state adjustment probability Λ_{ss} over the (endogenous) steady-state joint distribution of firm prices and productivities, $\tilde{\Omega}_{ss}$, which we will derive later, and which is the distribution that arises before price review decisions are made, but after idiosyncratic shocks are realized in each period.[10]

Similarly, letting $f_{ss}(p|a)$ denote the steady state probability with which a firm with idiosyncratic productivity a sets price p when adjusting, the reference distribution for prices is set to the steady state distribution of prices, after price adjustments have been made,

$$\bar{f}(p) = \int f_{ss}(p \mid a)\, \Omega_{ss}(p, a)\, da, \tag{9}$$

[10] Going forward, we will index firms by p, a rather than j to make explicit the dependence on current price and idiosyncratic productivity.

where Ω_{ss} is the (endogenous) joint distribution of productivities and prices post-adjustment, which we will also derive later, as a function of the optimal adjustment and pricing policies.

The assumption that firms use the steady state cross-sectional distributions as their reference is motivated by the idea, plausible to us, that decision-makers with prior experience across a range of states may find it "easier" to have as references rules that they have observed work well on average, across many states. By constraining the reference distributions in this way, we are using a slightly inefficient information structure that should nevertheless be quite close to the optimal reference distributions that would be predicted by the pure RI solution, which would require the reference distributions to also be optimized, by being set to the discounted expected values of the choice distributions across the states expected to be encountered over the expected life of the policy. The Online Appendix discusses alternatives.

Recursive Formulation We now define the firm's problem recursively and solve for each element of the optimal policy. The value charging price p in a particular state is

$$V_t(p,a) = \pi_t^r(p,a) + E_t\Big\{ M_{t,t+1}\, V_{t+1}^*\left(p', a'\right) \Big\}, \tag{10}$$

where the subscript t indicates dependence on the aggregate state, expectations condition on the current state, V_{t+1}^* is the maximum attainable value the firm can expect in the next period (assuming optimal choices henceforth), $p' \equiv pP/P'$ is the real price with which the firm *begins* the next period (eroded by inflation), and a' is next period's idiosyncratic state. For a firm that begins period t with real price $\tilde{p}$ and productivity a, the firm's choices solve

$$V_t^*(\tilde{p},a) = \max_{\Lambda_t \in (0,1)} \Big\{ \Lambda_t \cdot \left[V_t^a(a) - \kappa\right] + (1-\Lambda_t)\cdot V_t(\tilde{p},a) - \theta^a\, \mathcal{D}\left(\Lambda_t \parallel \bar{\Lambda}\right) \Big\} \tag{11}$$

$$V_t^a(a) = \max_{f_t \in (0,1)} \Big\{ \int V_t(p,a)\, f_t(p\,|\,a)\, dp - \theta^p\, \mathcal{D}\left(f_t(p\,|\,a) \parallel \bar{f}(p)\right) \Big\} \tag{12}$$

s.t. $\int f_t(p\,|\,a)\, dp = 1\ \forall a$, where in the first equation we have suppressed the arguments of Λ_t to ease notation. The firm either adjusts to a new price, with probability $\Lambda_t\left(\tilde{p},a\right)$, or continues with its current price, which occurs with probability $1-\Lambda_t\left(\tilde{p},a\right)$. In either case, it pays the cost of conditioning this period's adjustment probability on this period's state.

If the firm continues with its existing price, it obtains $V_t(\tilde{p}, a)$, which consists of the flow operating profit at this price plus the expected discounted continuation value of entering the next period with this price. If, instead, the firm adjusts its price, it pays the menu cost κ and can expect to obtain $V_t^a(a)$, the expected value under the optimal repricing policy, net of the cost of using a policy that deviates from the default $\bar{f}$.

Optimal Choice Distributions The optimality condition for the choice of $\Lambda_t(\tilde{p}, a)$ equates the marginal value of a more accurate adjustment decision to its marginal cost, state by state. This yields an expression for the optimal log odds of adjustment given by

$$\ln\left(\frac{\Lambda_t(\tilde{p}, a)}{1-\Lambda_t(\tilde{p}, a)}\right) = \ln\left(\frac{\bar{\Lambda}}{1-\bar{\Lambda}}\right) + \frac{1}{\theta^a}\left[V_t^a(a) - V_t(\tilde{p}, a) - \kappa\right], \tag{13}$$

with $\bar{\Lambda}$ given by equation (8). The model predicts a linear relationship between the conditional log odds, the reference log odds, and the net gain from adjusting the firm's price, with the unit cost θ^a governing the sensitivity of the adjustment decision to the real-time value of adjusting in each state. In the limit, $\theta^a \to \infty$ implies a constant probability of adjustment, independent of the state, as in the Calvo model. As $\theta^a \to 0$ the adjustment decision converges to a deterministic Ss adjustment rule, with the firm adjusting its price with certainty outside the Ss bounds and not adjusting inside these bands. For intermediate values of θ^a, the adjustment probability is smoothly increasing in the value of adjusting, yielding a stochastically state-dependent adjustment decision, as in (Woodford, 2009) and similar to random menu costs models (Dotsey et al., 1999).

Now consider the optimal choice for the probability of charging a particular price in a particular state. This choice too can be optimized state by state and is given by

$$f_t(p \mid a) = \frac{\bar{f}(p)\,\exp\left\{\frac{V_t(p,a)}{\theta^p}\right\}}{\int \bar{f}(\hat{p})\,\exp\left\{\frac{V_t(\hat{p},a)}{\theta^p}\right\}\,d\hat{p}} \tag{14}$$

for each p charged with positive probability in the steady state, with $\bar{f}$ given by equation (9). A price is charged with a higher probability in a particular state if it yields a higher value in that state compared to the average value across all possible prices under the reference distribution. The value of deviating from the reference must be high enough, relative to the marginal cost of information θ^p, to compensate for the higher information expenditure. A lower attention cost θ^p enables finer differentiation across states. As the cost approaches zero, the firm's repricing policy approaches a degenerate distribution for each state, concentrating all the probability mass at the optimal full-information reset price. Conversely, as θ^p increases, the firm increasingly relies on the reference distribution. It does so first in states in which its continuation value is not too price sensitive, and eventually across all states.

Equations (13) and (14) give an optimization-based approach to generalizing the fixed adjustment cost model to a stochastic version. The firm acts *probabilistically* in both its decision about whether or not to change its price and its choice about what price to set. But the degree of stochasticity in each is the result of a cost-benefit analysis that is made state by state and period by period: In states of the world where mispricing is very costly, the firm devotes more resources to information acquisition. Conversely, in states where its payoffs are not very sensitive to pricing accuracy, the firm economizes on information costs, reallocating these savings to improve decision accuracy in more profit-sensitive states.

Price Distributions The law of motion for the joint distribution of prices and idiosyncratic states after all pricing decisions have been made is given by

$$\Omega_t(p,a) = [1 - \Lambda_t(p,a)] \cdot \tilde{\Omega}_t(p,a) + \left[\int \Lambda_t(\hat{p},a)\, \tilde{\Omega}_t(\hat{p},a)\, d\hat{p} \right] \cdot f_t(p \,|\, a), \tag{15}$$

where $\tilde{\Omega}_t(p,a)$ is the joint distribution at the beginning of the period, before any pricing decisions have been made, but after the realization of all shocks in the period. It is given by the post-adjustment distribution in the previous period $\Omega_{t-1}(p,a)$, with prices eroded by inflation and idiosyncratic states transitioned to period-t values,

$$\tilde{\Omega}_t(p,a) = \int \Omega_{t-1}(pP_t/P_{t-1}, \hat{a}) \cdot T_a(\hat{a},a)\, d\hat{a}, \tag{16}$$

where T_a gives the probability of transitioning from productivity $\hat{a}$ to a.

Price Dispersion and Resource Cost To complete the exposition of the pricing block of the model, we next define the total demand for the intermediate input,

$$x_t^d \equiv \int x_{jt}\,dj = \int p_{jt}^{-\varepsilon_t}\,e^{-a_{jt}}\,Y_t\,dj = Y_t\,\Delta_t, \tag{17}$$

where Δ_t is the equilibrium level of price dispersion in the economy and is pinned down by the distribution of prices:

$$\Delta_t = \int p^{-\varepsilon_t} e^{-a} d\Omega_t(p,a). \tag{18}$$

Given aggregate and idiosyncratic conditions, retailers make their pricing choices, generating a level of aggregate price dispersion Δ_t. In the economy with no pricing frictions, price dispersion would arise only due to differences in idiosyncratic productivity. With frictions, there is additional, inefficient price dispersion that causes misallocation across firms.

Finally, the resource cost of the pricing frictions is

$$F_t = \int \left\{ \theta^a \mathcal{D}\left(\Lambda_t(\tilde{p},a) \parallel \bar{\Lambda}\right) + \Lambda_t(\tilde{p},a)\left[\kappa + \theta^p \mathcal{D}\left(f_t(\cdot|a) \parallel \bar{f}\right)\right]\right\} d\tilde{\Omega}_t(\tilde{p},a), \tag{19}$$

where the first term integrates over all review costs and the second term adds the repricing costs of all firms that adjust in the period.

2.2 Closing the Model

Intermediate Goods Producers The supply of intermediate inputs is determined by a continuum of competitive producers who choose labor to maximize static real profits, $\pi_t^x = p_t^x x_t - w_t L_t$, subject to the production function $x_t = e^{a_t + z_t} L_t$, where L_t is labor input, w_t is the real wage, a_t is an AR(1) process for log aggregate productivity, and z_t is a random walk process that grows at the rate γ_{zt}, itself an AR(1) process. Optimization by the intermediate goods producer yields the real price for the intermediate input, $p_t^x = w_t\,e^{-(a_t+z_t)}$, and market clearing for the intermediate input yields total labor demand $L_t = Y_t\,\Delta_t\,e^{-(a_t+z_t)}$.

Households and Fiscal Authority The representative household chooses streams of consumption C_t, labor supply L_t, and the real value of the risk-free bonds B_t purchased in period

t, to maximize lifetime utility,

$$\mathcal{U} = E_0 \sum_{t=0}^{\infty} \beta^t \zeta_t \left[\ln\left(C_t - hC_{t-1}\right) - \xi_t \cdot \left(\frac{L_t^{1+\frac{1}{\nu}}}{1+\frac{1}{\nu}} \right) + \chi_t \cdot B_t \right] \tag{20}$$

subject to the sequence of flow budget constraints

$$C_t + B_t \; = \; w_t L_t + D_t - T_t + B_{t-1} \frac{i_{t-1}}{\pi_t}, \tag{21}$$

and a no-Ponzi condition, where $\beta \in (0, 1)$ is the discount factor, ζ_t is a discount factor shock, $h \in [0, 1)$ is the degree of habit in consumption, $\nu \geq 0$ is the Frisch elasticity of labor supply, ξ_t is a shock to the relative disutility of working with mean $\bar{\xi}$, χ_t is a shock to the household's preference for bonds with mean $\bar{\chi}$, D_t are firm dividends, T_t are lump-sum fiscal taxes net of transfers, i_{t-1} is the gross nominal interest rate earned on bonds between $t-1$ and t, and $\pi_t \equiv \frac{P_t}{P_{t-1}}$ is the gross inflation rate between $t-1$ and t.

The discount factor shock ζ_t affects the intertemporal Euler equation and has been shown by Justiniano, Primiceri & Tambalotti (2010) and others to be a useful (reduced-form) driver of consumption fluctuations. It is also often used to drive the economy to the effective lower bound on the nominal policy rate of the monetary authority (Eggertsson & Woodford, 2003). Shocks to labor disutility ξ_t are introduced to affect the firm's marginal cost function. Lastly, including the real value of bond holdings in the household's utility function allows for reduced form risk preference shocks χ_t, following Krishnamurthy & Vissing-Jorgensen (2012), Fisher (2015), and Campbell, Fisher, Justiniano & Melosi (2017).

The final good is used for consumption C_t, government spending G_t, and to pay for the costs associated with pricing frictions F_t,

$$Y_t = C_t + G_t + F_t, \tag{22}$$

where government spending is funded by lump-sum consumer taxes and is a fixed share g of steady state output net of pricing frictions, $G_t = g(Y_{ss} - F_{ss})$.

Wage Rigidity To avoid over-stating the role of pricing frictions and the degree of marginal cost flexibility, we include a simple reduced-form wage rigidity given by $w_t = \delta^w \bar{w}^* + (1 - \delta^w) w_t^*$, where $\bar{w}^*$ is the steady state real wage and w_t^* is the competitive real wage.

Monetary Authority The monetary authority follows a Taylor rule featuring deviations of output growth from long-run growth, as in Sims (2013). The rule also features interest rate smoothing and is subject to a zero lower bound. When not constrained by the lower bound, monetary policy implements

$$i_t^a = \rho_i \, i_{t-1}^a + (1 - \rho_i) \left[i_{ss}^a + \phi_\pi \left(\pi_t^a - \pi_{ss}^a \right) + \phi_y \left(dy_t^a - \gamma_{ss}^a \right) \right] + e_t^\pi + e_t^i \,, \tag{23}$$

where i_t^a is the annual nominal rate in month t, π_t^a is the inflation rate over the most recent 12 months, dy_t^a is the real output growth over the most recent 12 months, γ_{ss}^a is steady-state annual output growth, and i_{ss}^a is the steady state nominal rate associated with steady state inflation π_{ss}^a. Policy is subject to a Gaussian AR(1) shock e_t^π and a Gaussian i.i.d. shock e_t^i.

Balanced Growth The source of long-run growth is labor-augmenting technological progress z_t, which grows at rate γ_{zt}. We detrend aggregate variables by z_t and require that $\kappa_t, \theta_t^r, \theta_t^p$ grow at the same rate as z_t.

Shocks We include a range of fundamental shocks, to avoid overstating the role the pricing frictions play in generating aggregate volatility. The shocks are to aggregate TFP (a_t), impatience (ζ_t), disutility of labor supply (ξ_t), bond demand (χ_t), markups (ε_t), the Taylor rule (e_t^π, e_t^i), and permanent productivity growth (γ_{zt}).

3 Steady State Frictions

In this section we report what the pricing statistics over recent decades suggest about the nature and severity of pricing frictions in the United States, on average.

TABLE I: Data and Baseline Model Moments

	Full Sample 1978-2023.Q1	Post-1984 1984-2023.Q1	Moderation 1984-2007.Q2	Baseline Model
Frequency	0.1131	0.1092	0.1005	0.0972
Size	0.0735	0.0744	0.0740	0.0750
Std. deviation	0.129	0.133	0.129	0.122
Skew	-0.131	-0.166	-0.142	-0.141
Kurtosis	11	10	11	11
Frequency cuts	0.031	0.032	0.029	0.030
Size cuts	0.074	0.077	0.076	0.080
GDP growth rate	0.0151	0.0158	0.0204	0.0204
Inflation rate	0.0354	0.0281	0.0308	0.0308
Federal funds rate	0.0462	0.0354	0.0532	0.0532

The pricing statistics report averages for moments constructed for the monthly distributions of log-price changes. *Size* is the size of the absolute value of price changes. *Frequency cuts* and *Size cuts* are the frequency and size of price cuts. Per-capita real GDP growth, CPI inflation, and average effective federal funds rate are annual rates. The shaded column shows the moments targeted in the steady state estimation. Sources: Daniel Villar, FRED.

3.1 Moments and Parameters

We use statistics from the U.S. Consumer Price Index (CPI) to estimate the steady-state pricing frictions. These statistics are based on the individual price quotes underlying the CPI, as constructed by Nakamura et al. (2018) for the sample starting in January 1978 and ending in December 2014, and extended by Montag & Villar (2023) to March 2023. We thank Daniel Villar for sharing the time series of these moments with us. Table I reports the average values of seven pricing moments for the full sample period and for two sub-samples: post-1984, which may be of independent interest since it represents a period of modern monetary policy, and the Great Moderation, which the Federal Reserve dates between January 1984 and June 2007, and which we target for our steady-state. The table also reports GDP growth, inflation, and the effective federal funds rate, which we also target.

We parameterize the steady state by targeting averages over the Great Moderation period since that was a period of relative macroeconomic stability. The estimation of aggregate shocks away from the steady state will then make use of the entire data, including the volatile

TABLE II: Steady State Parameters

Parameter		Baseline	$\theta^p \doteq 0$	$\kappa \doteq 0$	$\theta^a \doteq 0$	$K \doteq 5$
Annual real growth rate	γ^a_{ss}	1.0204				
Annual inflation rate	π^a_{ss}	1.0308				
Gov't share G/C	g_c	0.25				
Frisch elasticity	ν	2				
Monthly discount factor	β	0.9583				
Mg. disutility of labor	ξ	1.7	1.1	1.8	2.1	1.3
Mg. utility of bonds	χ	0.058	0.059	0.059	0.059	0.057
Repricing accuracy cost	θ^p	1.07	**0**	0.57	1.09	1.59
Fixed cost	κ	0.026	0.18	**0**	0.038	0.055
Timing accuracy cost	θ^a	0.097	2.58	0.29	**0**	0.0008
EOS among varieties	ε	11	5.8	11	11	6
SD(idios. shocks)	σ	0.077	0.19	0.05	0.09	0.067
Persistence(id. shocks)	ρ	0.94	0.60	0.98	0.95	0.95

The estimation is at the monthly frequency. Parameters are set at conventional values or to match Great Moderation averages. The bottom six parameters are estimated jointly, targeting seven pricing moments. The 'Baseline' column shows the best fit. The $K \doteq 5$ column shows the estimation that targets a steady state kurtosis of price changes of 5. The remaining columns show parameter values when we re-estimate the model shutting down one pricing friction at a time.

periods at the beginning and end of the sample. Prices showed considerable volatility even during the Great Moderation, as has been pointed out in previous work (*e.g.*, Bils & Klenow, 2004; Golosov & Lucas Jr, 2007; Nakamura & Steinsson, 2008). In particular, the coexistence of large price cuts and price increases during a period of low inflation volatility points to the importance of idiosyncratic shocks. Hence, as is common in the pricing literature, we solve for a stochastic steady state with idiosyncratic shocks to firms' marginal costs.[11]

Table II presents the steady-state parameter values and the last column of Table I reports the model-implied values for the targeted moments. The empirical pricing moments are constructed using monthly price changes, so we estimate the model at the monthly frequency, to avoid having to make assumptions about price rigidities in the mapping from monthly to quarterly moments. We set monthly productivity growth in the model to target the average

[11]Results are similar in an alternative parameterization that includes idiosyncratic demand shocks and decreasing returns to scale, as, for example, in Burstein & Hellwig (2008).

annual real GDP growth per capita ($\gamma^a_{ss} = 1.0204$), the steady state gross inflation rate to match the realized annual average ($\pi^a_{ss} = 1.0308$), and government spending to 25% of steady state consumption, per the realized sample average. The Frisch elasticity of labor supply is $\nu = 2$ and the relative disutility of working is $\xi = 1.7$, set to normalize steady-state employment. Following Michaillat & Saez (2021), we parameterize wealth in the utility by setting the monthly discount factor to $\beta = 0.9583$ and internally calibrating the parameter governing the marginal utility of bonds ($\chi = 0.058$) such that the steady state nominal interest rate matches the average federal funds rate ($i^a_{ss} = 1.0532$). This specification helps the estimation when shocks push the economy to the effective lower bound on nominal interest rates (see also Cuba-Borda & Singh, 2019).

The seven pricing moments are used to jointly determine the six parameters that govern the price-setting frictions, ε, σ, ρ, κ, θ^a, θ^p, by minimizing the sum of residuals, in percentage terms. A notable feature of the estimated parameters is the high degree of mispricing conditional on adjustment, coupled with accurate timing of adjustment: We estimate $\theta^p = 1.07$ and $\theta^a = 0.097$, implying that firms determine quite accurately when their prices have become obsolete, but they have more difficulty determining the right price level to set. The coexistence of price cuts that are larger than price increases, excess kurtosis, negative skew, and a large standard deviation of price changes can only be reconciled with substantial imprecision in reset prices. On the other hand, we estimate only a modest menu cost ($\kappa = 0.026$, about half the value typically estimated in models with perfect repricing).

Matching the pricing moments also requires a high elasticity of substitution ($\varepsilon = 11$, larger than the value usually estimated in menu cost models), and volatile and persistent idiosyncratic shocks (which is typical in this literature, given the large volatility of price changes). Since we target more moments than we have parameters, the fit between model and data is not perfect, but it is quite close. The match is particularly notable given the Gaussian idiosyncratic shocks (rather than the leptokurtic shocks typically used to target the standard deviation and kurtosis of price changes, *e.g.*, Karadi & Reiff (2019)).[12]

[12] We have also computed results for alternative parameterizations that vary the moments targeted. For example, if we remove the frequency and size of price cuts as targets, and set the elasticity of substitution among varieties to a more frequently used value ($\varepsilon = 6$), the model can match the remaining five moments with slightly lower pricing frictions. One could proceed with this parameterization, which is quite common in the second-generation menu cost literature. However, the frequency and size of price cuts offer useful

3.2 Optimal Policies and Incentives for Information Acquisition

The accuracy with which firms make pricing decisions is shown in Figure 1, which plots the state-dependent probability of adjustment Λ and the state-dependent distribution of reset prices conditional on adjustment, f. When a firm's existing price is close to the price it would expect to set upon repricing, the probability of adjustment is essentially zero. Unlike in the canonical menu cost models, the probability of adjustment is never exactly zero or exactly one. Hence, instead of an area of inaction, there is an area of infrequent price adjustment, for which the probability of adjustment is still positive, but very close to zero. This area is shown in dark in the figure. For prices that are farther from the optimum, the probability of adjustment rises rapidly, generating strong state-dependence. The area of infrequent adjustment is *funnel-shaped*: narrow in the region of low marginal costs, and widening and flattening as marginal costs rise. Low marginal costs provide an opportunity to generate significantly higher profits than average, especially given a high elasticity of substitution across varieties. Hence, firms have strong profit incentives to accurately identify and capitalize on these states. At higher marginal costs, profits become less sensitive to mispricing, as long as prices are not too low. As a result, the area of infrequent adjustment fans out asymmetrically: the hazard function triggers adjustment with near certainty when the existing price is relatively low, protecting the firm from having to satisfy a lot of demand at a high cost, but only gradually increases when the firm's current price is above the expected reset price, helping the firm economize on repricing costs at little profit loss.

Conditional on adjustment, reset prices are drawn from an imprecise, weakly state-dependent distribution. The distribution of reset prices also features state-dependent pass-through of marginal cost to price, conditional on adjustment: At low marginal cost, the expected reset price tracks the optimal reset price closely and the dispersion of the pricing policy is low since, as noted above, the firm has strong profit incentives to acquire enough information to track the optimum closely, because the profit gains are so large. As marginal costs rise, the reset price distribution becomes more dispersed and less sensitive to marginal cost. For high enough marginal costs, the pricing policy does not meaningfully differentiate

clues not just about the distribution of idiosyncratic shocks, as previously discussed in the literature, but also about the severity and nature of mispricing, as we discuss in the identification section.

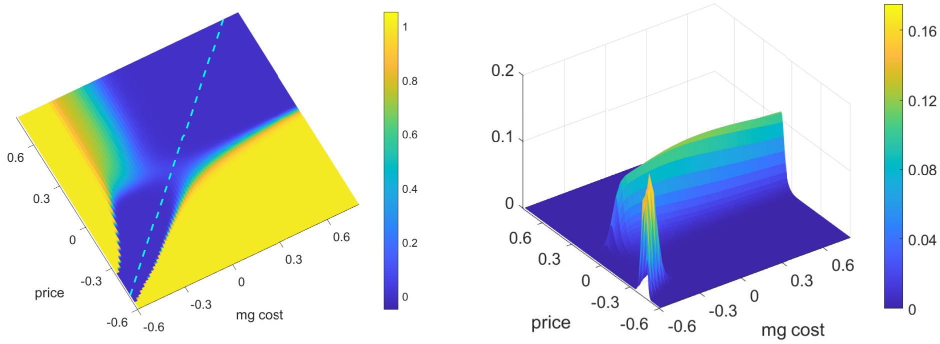

(a) Hazard function

(b) Pricing policy

Figure 1: **Optimal Policy in the Baseline Model**

Notes: (a) Estimated adjustment probability Λ and (b) optimal pricing policy f in steady state. The dashed line in panel (a) indicates the full information reset price.

across states, saving in repricing costs. Models that assume reset prices based on perfect information yield a single state-dependent optimal reset price, rather than a distribution, potentially misstating any possible state-dependence in the responsiveness of prices to shocks.

It is important to underscore that timing and repricing accuracy are chosen to be jointly optimal, and hence they interact to optimize the use of information. On the one hand, the possibility of mistakes in price-setting makes firms more careful when changing prices so that the timing of adjustments becomes more state-dependent. On the other hand, mistakes in timing make the firm pay more attention to the price it sets, reducing the dispersion in the distribution of reset prices. Finally, both types of mistakes are sensitive to the asymmetry of the profit function, which features larger losses from under-pricing than from over-pricing. This asymmetry shapes the state dependence in the attention of a firm trying to decide how to learn most efficiently about what price to set.[13]

3.3 Alternative Parameterizations

Tables II and III also report alternative parameterizations and the resulting pricing moments for versions of the model in which we re-estimate the steady state shutting down one pricing cost at a time. The results show that each pricing friction plays an important role in matching

[13]This asymmetry, which has been discussed in the menu cost literature before, is also why it is important to work with the actual profit function rather than quadratic approximations that are usually employed in the information frictions literature.

TABLE III: Pricing Moments: Baseline & Alternative Estimations

	Baseline	$\theta^p \doteq 0$	$\kappa \doteq 0$	$\theta^a \doteq 0$	$K \doteq 5$
Freq. of price changes	0.097	0.075	0.088	0.092	0.095
Absolute size	0.075	0.079	0.085	0.079	0.079
Standard deviation	0.122	0.098	0.116	0.145	0.107
Skew	-0.142	-0.139	-0.098	-0.127	-0.141
Kurtosis	11	9	7	12.6	5
Freq. of price cuts	0.030	0.027	0.030	0.025	0.065
Size of price cuts	0.080	0.063	0.082	0.098	0.078

Alternatives re-estimate the model when shutting down one friction at a time. The last column re-estimates the full model targeting a steady state kurtosis of 5.

the distribution of price changes. Importantly, the worst performing alternative is the one that imposes no frictions in reset prices ($\theta^p \doteq 0$). To approach the empirical targets, this version requires near-Calvo price adjustment ($\theta^a = 2.6$), a much higher menu cost ($\kappa = 0.18$), and implausibly large and transitory idiosyncratic shocks ($\sigma = 0.19$, $\rho = 0.6$). The finding that without imprecision in reset prices the data require weaker state dependence in the timing of price changes is consistent with the large literature on second-generation menu cost models that has found it necessary to augment menu cost models with Calvo-like features.[14] The fact that the re-estimated model cannot generate price cuts larger than price increases and yields a standard deviation, skew, and kurtosis of price changes are too small relative to the data is in line with prior work that has found it necessary to use Poisson arrival of idiosyncratic shocks (Midrigan, 2011) or a mixture of Gaussian shocks with different volatilities (Karadi & Reiff, 2019). Inaccuracy in reset prices instead delivers the needed dispersion and skew in the distribution of price changes without the need to change the underlying distribution of fundamental shocks, which has been found by the firm dynamics literature to be well approximated by AR(1) Gaussian processes.

The model with no menu cost ($\kappa \doteq 0$) continues to deliver infrequent price adjustment, arising solely from the information friction: Firms understand that they risk picking the wrong price, so they often choose to forgo price changes altogether. This uncertainty provides

[14]Two common ways of weakening state dependence are either assuming a menu cost that is zero with some probability, *e.g.*, Nakamura & Steinsson (2010), Gautier & Le Bihan (2022), or using multi-product firms with some economies of scope in price setting, *e.g.*, Midrigan (2011), Alvarez & Lippi (2014).

a potent micro-foundation for inaction: despite no menu cost, the frequency of price changes is only 8.8%. Matching the remaining pricing moments proves difficult however, without a fixed cost: The re-estimated model yields either too little or too much skew and kurtosis given the targets for the frequency and size of price changes. The best fitting alternative requires weaker state-dependence in the timing of price changes ($\theta^a = 0.29$ vs. 0.097 in the baseline) that is partially offset by more accurate reset prices ($\theta^p = 0.57$ vs. 1.07 in the baseline), but still misses the data along key dimensions: The resulting distribution of price changes has too little dispersion, skew, and kurtosis, price changes that are too large on average, and price cuts that are, counterfactually, smaller than price increases. Hence, even though the menu cost is not necessary for generating infrequent adjustment, it is needed for matching the shape of the distribution of price changes *conditional* on adjustment.

Imposing $\theta^a \doteq 0$, as in a menu cost model with errors in pricing, results in modest changes: We continue to estimate large repricing errors ($\theta^p = 1.09$) and a modest, albeit larger menu cost ($\kappa = 0.039$ vs. 0.026). Overall, the errors-in-pricing menu cost model delivers the broad patterns of the data, including larger price cuts than price increases. This is consistent with the fact that the estimated friction in the timing of reviews is small to begin with in the baseline model ($\theta^a = 0.097$). But matching the size and frequency of price changes without any errors in the timing of adjustment requires idiosyncratic shocks that are quite large, generating too much dispersion and excess kurtosis in the resulting distribution of price changes. We conclude that all three costs factor into firms' pricing decisions, each having a distinct role in shaping the distribution of price changes.

A Lower Kurtosis Among all the targeted moments, kurtosis deserves additional discussion, since the value for the US CPI data is much higher than the values from US scanner price data or datasets from Europe or Canada. This difference may be due to the broader range of goods in the CPI than in scanner data sets and it may reflect a combination of cross-sectional heterogeneity (in both pricing frictions and the distribution of desired price changes) and measurement error (since higher order moments are much more difficult to estimate accurately). How important is this higher than usual level of kurtosis for the results? The last columns of Tables II, IV, and III report the estimation results when targeting a

TABLE IV: Steady State Outcomes Across Estimations

	Baseline	$\theta^p \doteq 0$	$\kappa \doteq 0$	$\theta^a \doteq 0$	$K \doteq 5$
Spending on repricing (share of revenues)					
Fixed cost ($\kappa\,\bar{\Lambda}$)	0.003	0.014	0	0.004	0.006
Review cost ($\theta^a\, I^a_{ss}$)	0.008	0.002	0.009	0	0.0003
Repricing cost ($\theta^p\, I^p_{ss}$)	0.016	0	0.014	0.020	0.004
Total spending (F_t)	0.026	0.016	0.023	0.024	0.010
Outcomes (relative to flex-price outcomes)					
Consumption	0.93	0.85	0.96	0.94	0.97
Employment	1.05	1.02	1.04	1.04	1.07
Wages	0.95	0.86	0.97	0.96	0.99
Output	0.95	0.87	0.98	0.96	0.98
Price Dispersion	1.11	1.18	1.06	1.09	1.09

Spending on pricing decisions is reported as a share of steady state revenues. The reported numbers may not add up due to rounding. Aggregate outcomes are reported relative to the outcomes in a flexible price economy that is otherwise identically parameterized.

kurtosis of 5. This yields a higher value for the cost of repricing accuracy ($\theta^p = 1.59$ instead of 1.07), a lower elasticity of substitution ($\varepsilon = 6$ instead of 11), and double the menu cost ($\kappa = 0.055$). Since the estimated imprecision in reset prices is even stronger, we conclude that the high kurtosis in the US CPI data is not driving our key finding that repricing errors are a major source of inefficient price dispersion.

3.4 Aggregate Outcomes

Table IV reports the breakdown of price-setting costs and the steady state aggregate outcomes. We estimate that firms annually spend approximately 2.6% of revenues on total information and repricing costs: 0.8% on determining whether a price change is warranted, 1.6% on determining the right price to charge, and only 0.3% of revenues on the menu cost. Compared with the flexible-price steady state, the economy with pricing frictions delivers significantly lower welfare. Steady state consumption is 7% lower, employment is 5% higher, and wages are 5% lower. Consumers work harder for less. Equilibrium price dispersion is 10% higher than the (efficient) price dispersion that would be warranted given the productivity heterogeneity. This dispersion arises due to both inaction and inaccuracy in reset prices. The steady state features a distribution of reset prices that is most dispersed in the middle, reflecting the fact that firms fine-tune their pricing accuracy depending on the state. The firm's optimal distribution of reset prices is concentrated at low marginal costs because low costs are highly profitable opportunities that are worth capturing accurately. On the other hand, it is also narrow at high marginal costs, because they present opportunities to save on information and adjustment costs by not differentiating prices across states too much. Most of the time, however, firms are in the middle range with the widest price dispersion. As a result, these frictions aggregate to considerable price dispersion.

Across the versions of the model with different pricing frictions, the level and composition of total spending on price setting change, as do the resulting steady state outcomes. For example, the model that imposes no repricing errors ($\theta^p \doteq 0$) requires Calvo-like price adjustment to match the pricing moments and as a result, delivers far lower steady state consumption (85.3% of the flexible-price benchmark), due to a much higher level of inefficient price dispersion (7 percentage points higher than the baseline model).

3.5 Identification and Sensitivity Analysis

Perhaps it is not surprising that adding another friction to price setting helps better match the various pricing moments. What is surprising, however, is that this friction turns out to dominate the other two frictions included. To better understand why, we now turn to steady state simulations around our baseline estimates. The panels in Figure 2 plot how the pricing moments vary with each of the three key pricing parameters when fixing all other parameters at their estimated steady state levels.[15]

The first key finding of the simulations is that pricing moments are highly sensitive to the repricing cost θ^p (in red). Moreover, θ^p has a non-monotonic relationship with multiple moments. Introducing small frictions in reset prices generates substantial inaction, sharply lowering the frequency of price changes and increasing the size and dispersion. Hence, models may over-estimate the size of the adjustment costs needed to match the frequency of price

[15]In the interest of space, we omit the panel for the frequency of price cuts, which shows the same patterns as the frequency of overall price changes.

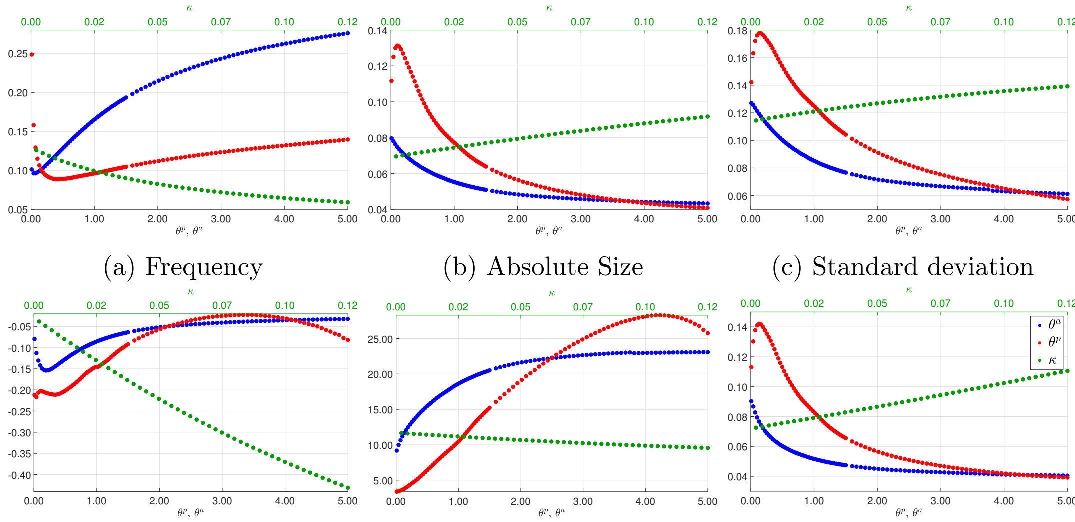

(a) Frequency (b) Absolute Size (c) Standard deviation

(d) Skew (e) Kurtosis (f) Abs. size of price cuts

Figure 2: **Pricing Moments in Simulations**

Note: This figure plots pricing moments for different values of the pricing parameters θ^p (red dots), θ^a (blue dots), and κ (green dots and top axis), while fixing the other parameters at their baseline values.

changes, if they abstract from the possibility of mistakes in repricing. Inaccuracy in reset prices incentivizes the firm to devote more attention to accurately choosing the timing of its price adjustments and to widen its area of low probability adjustment: The firm now avoids adjusting when the perceived value of adjusting is relatively small, so as not to risk setting an even worse reset price. Once θ^p exceeds 0.4 however, inaccuracy in repricing becomes high enough to generate more price changes. Beyond this threshold, θ^p monotonically raises the frequency and lowers the size and standard deviation of price changes.

The second key finding is that the dynamics of the pricing moments in response to higher information costs are very different not only from each other, but also from those in response to higher menu costs. Higher menu costs lead to less frequent, larger, and more dispersed price changes. They also slightly lower kurtosis, and sharply increase the negative skew of the price change distribution. On the other hand, higher information costs lower firms' incentives to track market conditions in real time. As a result, price changes become more frequent. In particular, a higher cost governing the timing of adjustments θ^a sharply *increases* the frequency of price changes: If firms are more uncertain about when to review their prices, they choose to review them more often, thus raising the overall frequency of adjustment. We conclude that menu costs should *not* be considered a stand-in for information frictions.

3.6 From Frictions to Non-Neutrality

The goal of estimating the frictions that govern pricing dynamics is to determine how flexibly prices respond to shocks. The standard New Keynesian assumption is that firms stand ready to produce whatever quantity is needed to satisfy demand at their posted prices. As a result, any rigidity in prices translates into inefficient adjustments of factor inputs and production in response to shocks, leading to an inefficiently low level of output and misallocation across firms. A simple way to summarize this effect is to compute the cumulative response of output or consumption to a monetary policy shock: In the absence of pricing frictions, the response should be zero. The larger the response of real quantities, the more severe are the nominal rigidities affecting the economy's adjustment to shocks.

Table V reports the cumulative impulse response (CIR) of consumption, as a percent of annual steady state consumption, to a 25 bp monetary policy shock, for our baseline estimation and for alternatives that turn off different pricing frictions. All models use the same Taylor rule for monetary policy, hence all responses will be dampened by the monetary authority responding to the gaps that open in response to the shock. Our baseline estimation implies substantial non-neutrality, despite the strong state-dependence in the timing of price changes: the CIR is 88% of the CIR of a Calvo model parameterized to match the same frequency and size of price changes. Frictions in reset prices amount to negative selection on the *intensive* margin that works to offset the positive selection along the extensive margin.[16]

The CIRs for alternative parameterizations yield additional takeaways: First, ceteris paribus, eliminating repricing frictions in the baseline model virtually eliminates non-neutrality (CIR = 0.008%), demonstrating that inaccuracy in reset prices is the dominant friction in the model. On the other hand, re-estimating the model by imposing no repricing frictions requires much less state-dependence in the timing of price changes in order to fit the data, potentially over-stating the degree of non-neutrality (CIR = 0.135% vs. 0.122% in the baseline). It also suggests a way to rationalize Calvo-like behavior: Firms that can learn what

[16] For discussions of the importance of the extensive margin of adjustment in time-dependent and state-dependent models, see Caballero & Engel (2007), Golosov & Lucas Jr (2007), Auclert, Rigato, Rognlie & Straub (2022), and Gagliardone, Gertler, Lenzu & Tielens (2024).

TABLE V: Consumption Responses to Monetary Shock

	CIR (%)	
Baseline	0.122	
	Ceteris Paribus	*Re-Estimated*
No menu cost ($\kappa \doteq 0$)	0.084	0.098
No timing errors ($\theta^a \doteq 0$)	0.131	0.081
No repricing frictions ($\theta^p \doteq 0$)	0.008	0.135
Calvo (same freq + size)		0.138

CIR of consumption, as a percent of annual steady state consumption, to a 25-bp impulse to the Taylor rule, for baseline model and versions that shut down one friction at a time, ceteris paribus (first column) and re-estimating the other parameters (second column). The Calvo model is parameterized to match the same frequency and size of price changes as the baseline.

the right price is very easily do not need to worry about timing their price changes. They can change prices with some constant probability, as is assumed in the Calvo model. In practice, our estimates suggest that firms may find it easier to learn they are setting the wrong price than to know how to fix it. Relatedly, it may seem surprising that the menu cost model with errors in pricing ($\theta^a \doteq 0$) generates higher non-neutrality that the baseline, which also allows for imprecision in the timing of adjustments (CIR = 0.131% vs. 0.122%). But this result also illustrates how frictions affect firms' incentives to acquire information: Since in the case of $\theta^a \doteq 0$, firms know exactly when the value of adjusting is big enough to warrant an adjustment, they save on information costs for repricing, paying little attention to the prices they set, since if they make a costly mistake, they will promptly learn and adjust. In the aggregate, this adds up to more overall rigidity.

How does the CIR vary with repricing frictions? With perfect repricing, varying the severity of information frictions regarding the *timing* of price adjustment spans the degree of state dependence in price setting, from the menu cost model (when θ^a approaches zero) to Calvo (1983) (as θ^a becomes very large), as shown by Woodford (2009). However allowing for errors in the reset price itself adds a new source of mispricing, and hence non-neutrality. We illustrate this in Figure 3: The CIR increases monotonically with θ^p. Notice that this breaks the conventional wisdom that ties the frequency of price changes to the degree of nominal rigidities: For small to moderate values of θ^p, nominal rigidities increase with θ^p while

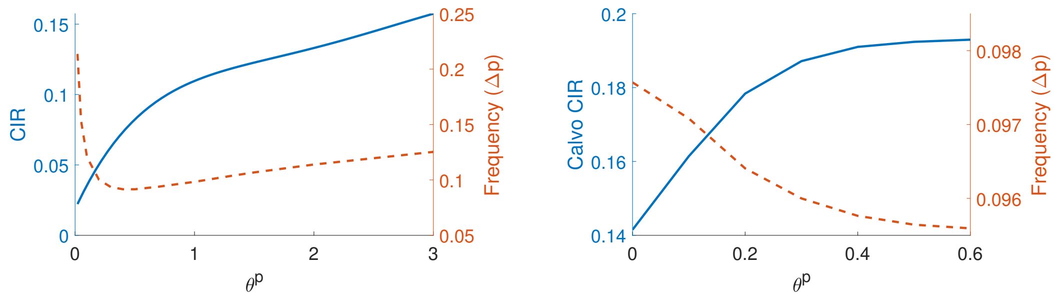

Figure 3: **CIRs and the Frequency of Price Changes**

Variation as a function of θ^p: (a) Baseline Model, (b) Calvo model.

frequency declines. But beyond this threshold, both frequency and rigidities *increase*. This *paradox of flexibility* implies that we can no longer take for granted that a higher frequency of price changes is necessarily associated with higher price flexibility, since it may well reflect lower or higher frictions in reset prices. Ultimately, what determines non-neutrality is the information content of prices, not the fact that they are changing *per se*.

The second panel of Figure 3 shows that the monotonic relationship between θ^p and non-neutrality also holds in the Calvo model augmented with inaccuracy in reset prices. As a result, the canonical Calvo model is no longer the upper bound on price rigidity: In the Calvo model, inaccurate reset prices substantially amplify nominal rigidities, while only modestly lowering the frequency of price changes.

4 Estimation

In this section, we report results from our Bayesian estimation of nominal pricing frictions over time. To our knowledge, this is the first use of Bayesian techniques to estimate a model featuring rationally-inattentive heterogeneous agents and an occasionally binding constraint on the nominal interest rate.

4.1 Estimation Approach

Our estimation exploits the time variation in the distribution of price changes to estimate fluctuations in pricing frictions and, as a consequence, in monetary non-neutrality over time. Figure 4 plots the time series for the pricing moments used in the estimation. There is

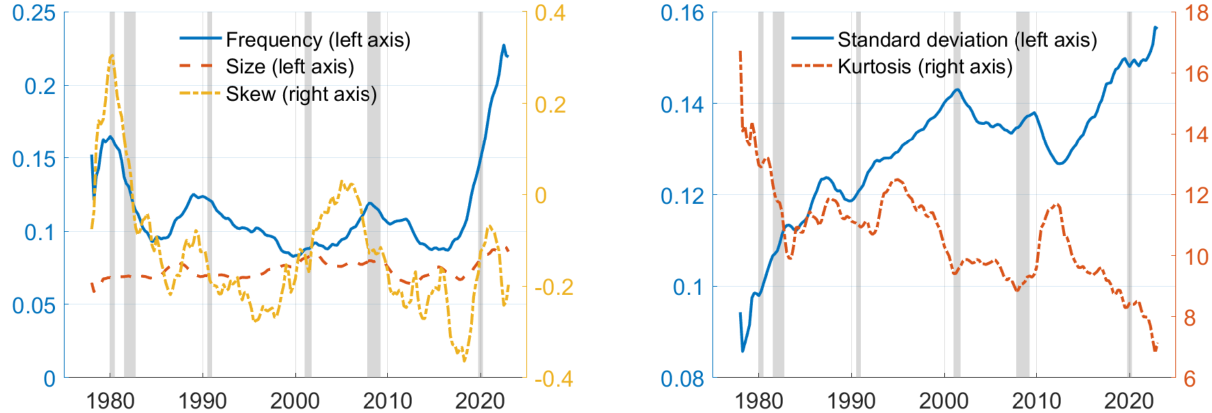

Figure 4: **Pricing Moments Over Time**

Note: This figure plots the smoothed pricing series used for the model estimation, from January 1978 to March 2023. Shaded areas mark NBER recession dates.

considerable volatility in the frequency and skew of price changes, with the monthly frequency ranging between 10% and over 20% and skew ranging between -0.4 and 0.3. There is also a large increase in the standard deviation of price changes, from around 9% at the beginning of the sample to more than 15% by 2023, and moreover, movements in standard deviation have been strongly negatively correlated with changes in kurtosis over time.

The jump in the frequency of price changes post-2020 is particularly striking. We caution that it may reflect, in part, pandemic-related changes in the BLS's data collection methods, in particular, greater reliance on online data due to lockdowns, data imputations due to missing price quotes, and changes in sample availability, especially for travel-related categories.[17] The extent to which these changes affected the aggregate pricing statistics during the 2020-2021 period remains unclear. But what is clear is that in fact the frequency, skew, and standard deviation of price changes all started increasing, and kurtosis started decreasing, several years *before* 2020. While neither of these moments is sufficient on its own to pin down non-neutrality, these synchronized movements already suggest that we might expect a change in the degree of price flexibility starting as early as 2016.

For the aggregate macro series, we use quarterly real GDP growth per capita, quarterly CPI inflation, the quarterly average of the federal funds rate, and the two-year and five-year Treasury yields. Following Kulish et al. (2017), we include these yields as observables

[17]See Montag & Villar (2023) for a discussion.

to aid identification at the effective lower bound (ELB) on the nominal interest rate. The estimation includes yield-specific shocks (η_t^2, η_t^5) and a common yield shock (η_t^y), in addition to the fundamental shocks listed in the model description.[18]

We compute the equilibrium dynamics using the sequence-space Jacobian (SSJ) method of Auclert et al. (2021). We extend this method to handle occasionally binding constraints, since our sample period includes two episodes during which the lower bound was binding on the federal funds rate. Appendix A describes how we adapt the methods proposed by Guerrieri & Iacoviello (2015) and Kulish et al. (2017) to handle the occasionally binding constraint when solving the model dynamics and evaluating the likelihood function. Even though we gain speed and accuracy with SSJ, the two ELB periods make the likelihood evaluation time-consuming, especially because we have a monthly model with a large number of observations (543 months). To reduce computational time, we only estimate the shock processes. Table B.1 in the Appendix reports the prior distributions for the estimated shock parameters, along with the posterior mode and standard deviation. These results are obtained by fixing habit formation and wage rigidity at conventional values, and running a separate Bayesian estimation for the parameters of the monetary authority's Taylor rule that best fit the behavior of the federal funds rate during the estimation period. Specifically, we set habit formation in consumption to $h = 0.67$, based on the meta-study of DSGE models of Havranek, Rusnak & Sokolova (2017) and we set the wage rigidity parameter to $\delta_w = 0.083$, which implies that wages change once a year on average. The Bayesian estimation for the parameters of the Taylor rule yields a value for persistence $\rho_i = 0.925$, an inflation coefficient $\phi_\pi = 1.7$, and a coefficient on GDP growth $\phi_y = 1.0$.

4.2 Pricing Frictions Over Time

Estimating the model by imposing constant pricing frictions yields a negligible amount of variation in the pricing moments over time: Shocks to firms' desired prices are simply not enough to generate the kind of volatility in the distribution of price changes that we see in the data. Hence, we allow for variation in the pricing frictions themselves: In addition to the shocks listed in Section 2, we allow for unanticipated shocks to the menu cost and to

[18]The Online Appendix provides a detailed description of data sources and transformations.

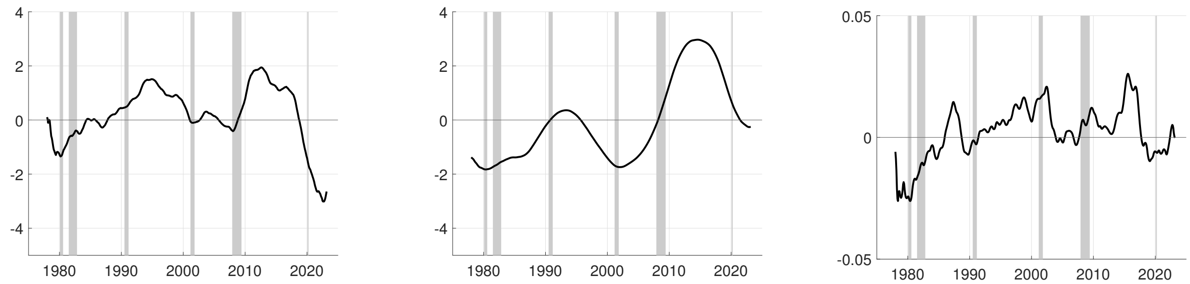

(a) $\widehat{\theta^p}_t$ (b) $\widehat{\theta^a}_t$ (c) $\widehat{\kappa}_t$

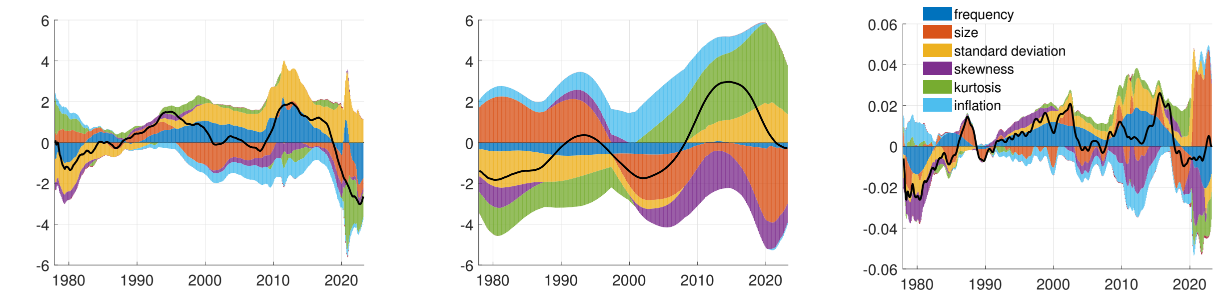

(d) $\widehat{\theta^p}_t$ decomposition (e) $\widehat{\theta^a}_t$ decomposition (f) $\widehat{\kappa}_t$ decomposition

Figure 5: **Estimated Pricing Frictions**

Note: The upper panels show the estimated series for θ^p, θ^a, and κ and the lower panels show the data decomposition when the model parameters are evaluated at the posterior mode. θ^p and θ^a are in log-deviations with respect to steady state. κ is in deviations with respect to steady state. Shaded areas represent NBER recession dates.

the marginal costs of acquiring information. Variation in these costs over time is interpreted as variation in the efficiency with which firms can process information and implement price changes. To avoid overstating the role of the pricing frictions, we also allow for unanticipated variation in σ_t, the standard deviation of idiosyncratic shocks.

Figure 5 plots the estimated pricing frictions series, $\widehat{\theta}^p_t$, $\widehat{\theta}^a_t$, in log-deviations from their steady state values, and $\widehat{\kappa}_t$, in deviations, when the model parameters are evaluated at their posterior mode. The lower panel of the figure shows the contribution of different pricing moments over time, and Figure 6 shows the evolution of the implied reference distribution $\bar{f}_t$. Overall, we find a large and variable repricing cost θ^p_t that has remained above 1 (0 in log-deviations from the steady state) for most of the sample period, indicating substantial inaccuracy in reset prices. Conversely, menu costs have been relatively small and the timing accuracy has been high.[19] The lower panels of the figure show the data decomposition

[19]Even though log deviations for θ^p appear large, note that they are deviations from a very small steady state value. With the exception of the 2011-2016 period, these fluctuations in $\widehat{\theta}^a_t$ do not significantly change

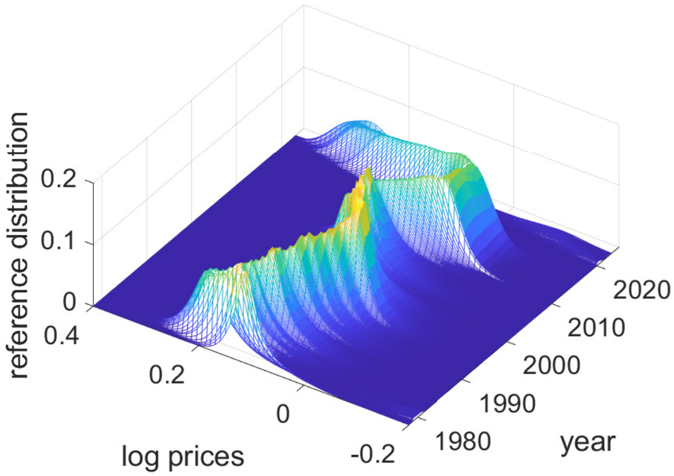

Figure 6: **Estimated Distribution of Prices Over Time**

for our pricing parameters, which allows us to gauge what features of our data inform the dynamics of our pricing parameters. All moments used in the estimation contain meaningful information about the dynamics of the pricing parameters, but the relative importance of different moments of the price change distribution varies over time. For example, while the frequency and standard deviation of log-price changes informed the spikes in θ^p during the mid-2010s, the decline in θ^p post-2020 was driven by the size and kurtosis of price changes, with frequency only contributing at the very end of the sample period.

We identify two episodes of heightened frictions: the mid-1990s and the mid 2010s, both periods when all costs–but especially the repricing cost–were highly elevated. In the 2012-2016 period, for much of the first ELB episode, we estimate that prices were characterized by near Calvo-like adjustment, with very little state-dependence in the timing of adjustments. Coupled with the heightened frictions for reset prices, this period saw very high nominal frictions. Notably, pricing frictions started falling precipitously in 2016 and reached historically low levels by 2021. How can we interpret these fluctuations in pricing frictions? Figure ?? plots $\widehat{\theta}^p_t$ and $\widehat{\theta}^a_t$ along with the time series of exogenous aggregate uncertainty of Jurado et al. (2015) and Ludvigson et al. (2021) (henceforth the JLN series). The series are strongly negatively correlated, suggesting a plausible economic interpretation to the variability we estimate: In periods of low uncertainty, when the benefits of accurate pricing are relatively

the degree of state dependence in the timing of price changes.

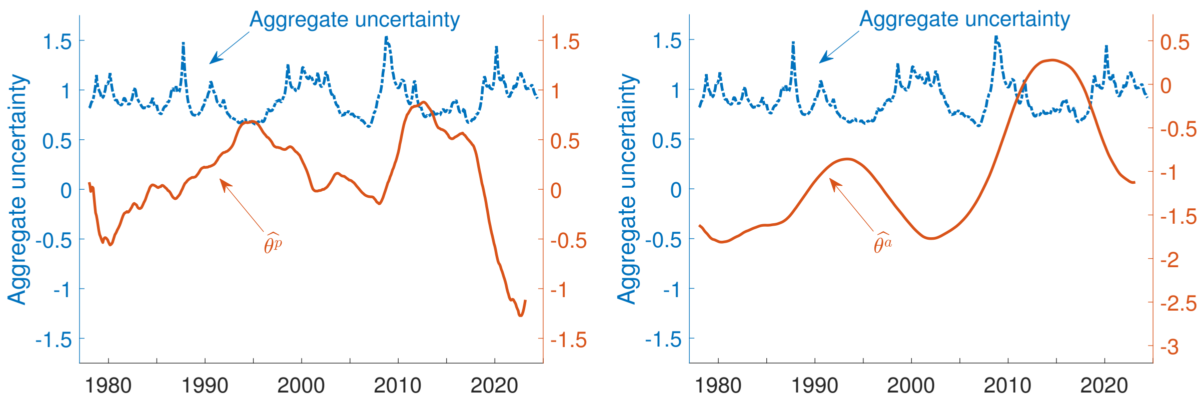

Figure 7: **Information Frictions and JLN Aggregate Uncertainty**

low, fewer resources are devoted to price-setting. This results in an empirical distribution of price changes that can only be rationalized with relatively high pricing frictions. Conversely, when uncertainty is high, mistakes in repricing can be quite costly, and so firms endogenously choose to pay more attention to market conditions, generating a distribution of price changes that our model rationalizes with low pricing frictions.

4.3 NPR Variability

What do these dynamics implied for how responsive the economy is to shocks? Figure 8 reports the estimated degree of NPR over time measured by the cumulative impulse response (CIR) of consumption to a 25 basis points shock to the monetary policy rate, given the estimated pricing parameters for each period, in percentage points of annual steady state consumption. A larger value of the CIR implies stronger monetary non-neutrality. For each point in time, we solve the model using the pricing parameters at that time and recalculating the reference probability of price adjustment $\bar{\Lambda}$ and the reference distribution of prices $\bar{f}$ while keeping all other parameters at their baseline values.

On average, we estimate a cumulative consumption response of 0.12 percent of annual steady-state consumption. This is a considerable degree of non-neutrality, representing 88% of the response a Calvo model would predict when calibrated to the same frequency and size of price changes. Surprisingly, we find no trend over time in the CIR, despite the fact that technology has arguably made repricing easier. We interpret this as suggestive of the complexity that goes hand in hand with technological progress and data abundance: more

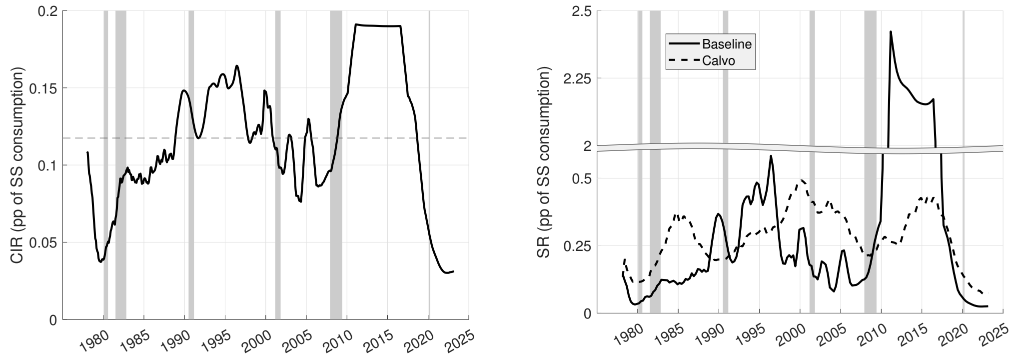

(a) CIR (b) Sacrifice Ratio

Figure 8: **Implied Nominal Rigidity Over Time**

Note: This figure plots the model-implied degree of nominal price rigidities over time: (a) CIR of consumption to a 25 bp monetary policy shock, in percentage points of steady state consumption. (b) Sacrifice ratio per percentage point of inflation reduction via contractionary monetary policy. Shaded areas represent NBER recession dates. The series are smoothed with a centered moving average of 13 months.

data do not always make decisions easier (Veldkamp & Chung, 2024). We also do not find consistent evidence of non-neutrality declining in recessions. A growing literature has argued that monetary policy may be less effective during downturns. This argument is based on evidence that price dispersion increases during downturns, which, viewed through the lens of models with frictionless reset prices, implies more price flexibility. However, in our model, the relationship between price dispersion and non-neutrality is not monotonic: While higher accuracy in reset prices always implies lower non-neutrality, it can imply either higher or lower price dispersion, depending on the initial level of pricing errors. Our results suggest a mixed record for US recessions. For example, we estimate that price rigidities rose during the Great Recession.

More importantly, we find substantial medium-cycle volatility in the degree of NPR over time. Our estimation shows two big cycles, one that peaked in the mid-1990s and another that peaked in the mid-2010s. Price rigidities reached a maximum value in the mid-2010s, a period of highly inattentive and inaccurate price setting. Our estimation also suggests that price rigidities started to decline well before the COVID-19 pandemic, pointing to the risk of higher inflation even before the pandemic-related shocks. In particular, the CIR started declining in mid-2016, concurrent with the Federal Reserve's liftoff of nominal rates from the first ELB period. By January 2021, it had already reached historical lows, even though inflation was only 1.5% at the time. This decline may explain why inflation surged then declined so rapidly in 2021 and 2022. It is possible that the 2016 liftoff induced pricing managers to refocus on inflation risks, thereby increasing the attention devoted to setting prices accurately. At the same time, the mid-2010s were also a period of rapid advances in AI pricing, with firms in retail, travel, and other industries increasingly using machine learning algorithms to devise optimal real time pricing strategies. Adams et al. (2024) document the rise of AI pricing at the aggregate level, and show a sharp acceleration starting in 2016 (see their Figure 1). The introduction of machine learning may turn out to be a structural shift in the rigidity of prices. Our series ends in March 2023, when we observe a slight uptick in rigidities, but further real-time data on the cross-sectional distribution of price changes would be needed to establish if 2016 was indeed a structural break.

The right panel of Figure 8 shows our estimated sacrifice ratio: how much would consumption have to decline in order for inflation to be lowered by one percentage point via contractionary policy. This measure offers an intuitive way to summarize the severity of pricing frictions. Most of the time, the sacrifice ratio fluctuations between 0.1 and 0.5 percent of consumption per percentage point of inflation decline. It was low during the early 1980s, suggesting that other factors beyond Volcker's disinflation contributed to the economic downturns of those years. It was high during most of the 1990s, a period of fast technological growth when managers may have been more focused on expansion of product markets than on accurate price setting. It was also extraordinarily high in the aftermath of the Great Recession, when nominal interest rates were at the effective lower bound. In 2016 it started declining as pricing frictions started falling and by the time of the 2021 inflation surge, the sacrifice ratio had fallen to a mere 0.03. Pricing frictions had come down dramatically, making monetary policy highly potent in controlling inflation. We estimate that when the Federal Reserve started raising interest rates in the first quarter of 2023, reducing inflation from 8.5%, its level at the time, to 2% would have been associated with a consumption cost of only 0.2%. Conversely, a similar disinflation in 1994 would have triggered a severe recession, costing approximately 2.7% of consumption.

Lastly, we also contrast our estimate sacrifice ratio series with the series implied by a

Calvo model parameterized to match the frequency of price changes over time. The Calvo model implies a 0.5 sacrifice ratio on average, but deviates significantly from the series predicted by our generalized model, serving as an imprecise guide to the cost of disinflation in different periods.

5 Conclusion

This article estimates a time-varying degree of nominal price rigidity in the U.S. economy over time. A key finding is that variation in the accuracy of reset prices is a major driver of both the average level and the volatility of nominal rigidities. We identify costly information as the main friction that prevents firms from adjusting prices more flexibly in response to shocks, with information about the right price to charge, conditional on adjustment, being the most significant friction. These results contribute to our understanding of how efficiently the U.S. economy has adjusted to shocks in recent decades, and how effectively policymakers have stabilized aggregate demand. On net, what do our results suggest for inflation and monetary control going forward? We emphasize the endogeneity and variability in the degree of state-dependence in price setting: First, our estimation results give great weight to firms' choices of how much attention to devote to choosing prices accurately. They suggest that while firms generally know with fairly high accuracy when their prices are outdated, they are much less certain about what the right price to charge is. Second, we find that firms' attention to market conditions is variable over time. This variability implies state-dependence in the cost of disinflation over time.

More work is needed to measure the attention firms devote to price setting versus other operational decisions. But our finding that mispricing is a major driver of monetary non-neutrality connects models of nominal rigidities to the much broader literature that has documented stochasticity in choice in a wide range of contexts. While stochastic choice may appear at odds with classic principles of optimization of well-specified objective functions, in this paper, we microfound it with rational acquisition of costly information. But it is worth separating the stochasticity result from the model through which we endogenize this stochasticity. We leave to future research other possible sources of randomness in decision-making (*e.g.*, deliberate, or exploratory randomization, model uncertainty, or other forms of bounded rationality). The important message is that whatever its source, the consequence of stochastic choice is often a systematic bias in the response of the aggregate price level to shocks. Stochasticity need not be divorced from but can rather be understood as a cause of bias (Woodford, 2020).

References

Adams, Jonathan J, Min Fang, Zheng Liu & Yajie Wang (2024), "The Rise of AI Pricing: Trends, Driving Forces, and Implications for Firm Performance," .

Afrouzi, Hassan (2020), "Strategic inattention, inflation dynamics, and the non-neutrality of money," CESifo working paper.

Afrouzi, Hassan & Choongryul Yang (2021), "Dynamic rational inattention and the Phillips curve," *SSRN 3770462* .

Alvarez, Fernando, Hervé Le Bihan & Francesco Lippi (2016), "The real effects of monetary shocks in sticky price models: a sufficient statistic approach," *American Economic Review* 106(10): 2817–51.

Alvarez, Fernando & Francesco Lippi (2014), "Price setting with menu cost for multiproduct firms," *Econometrica* 82(1): 89–135.

Alvarez, Fernando, Francesco Lippi & Luigi Paciello (2011), "Optimal price setting with observation and menu costs," *The Quarterly Journal of Economics* 126(4): 1909–1960.

Anderson, Eric, Nir Jaimovich & Duncan Simester (2015), "Price stickiness: Empirical evidence of the menu cost channel," *Review of Economics and Statistics* 97(4): 813–826.

Angeletos, George-Marios, Zhen Huo & Karthik A Sastry (2021), "Imperfect macroeconomic expectations: Evidence and theory," 35(1): 1–86.

Angeletos, George-Marios & Jennifer La'O (2009), "Incomplete information, higher-order beliefs and price inertia," *Journal of Monetary Economics* 56: S19–S37.

Assenza, Tiziana, Te Bao, Cars Hommes & Domenico Massaro (2014), "Experiments on expectations in macroeconomics and finance," in *Experiments in Macroeconomics*, vol. 17, pp. 11–70, Emerald Group Publishing Limited.

Auclert, Adrien, Bence Bardóczy, Matthew Rognlie & Ludwig Straub (2021), "Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models," *Econometrica* 89(5): 2375–2408.

Auclert, Adrien, Rodolfo D Rigato, Matthew Rognlie & Ludwig Straub (2022), "New Pricing Models, Same Old Phillips Curves?" Tech. rep., National Bureau of Economic Research.

Auclert, Adrien, Matthew Rognlie & Ludwig Straub (2020), "Micro jumps, macro humps: Monetary policy and business cycles in an estimated HANK model," Tech. rep., National Bureau of Economic Research.

Bandeira, Miguel, Laura Castillo-Martınez & Shiyuan Wang (2024), "Frictionless Inflation," Tech. rep., Working Paper.

Bayer, Christian, Benjamin Born & Ralph Luetticke (2020), "Shocks, frictions, and inequality in U.S. business cycles," Forthcoming, The American Economic Review.

Berger, David & Joseph Vavra (2018), "Dynamics of the US price distribution," *European Economic Review* 103: 60–82.

Bilbiie, Florin O, Giorgio Primiceri & Andrea Tambalotti (2023), "Inequality and business cycles," Tech. rep., National Bureau of Economic Research.

Bils, Mark & Peter J Klenow (2004), "Some evidence on the importance of sticky prices," *Journal of Political Economy* 112(5): 947–985.

Blanco, Andres, Corina Boar, Callum Jones & Virgiliu Midrigan (2024), "Non-linear inflation dynamics in menu cost economies," Tech. rep.

Burdett, Kenneth & Guido Menzio (2017), "The (Q, S, s) pricing rule: a quantitative analysis," *Research in Economics* 71(4): 784–797.

Burstein, Ariel & Christian Hellwig (2008), "Welfare costs of inflation in a menu cost model," *American Economic Review* 98(2): 438–443.

Caballero, Ricardo J & Eduardo MRA Engel (1999), "Explaining investment dynamics in US manufacturing: a generalized (S, s) approach," *Econometrica* 67(4): 783–826.

Caballero, Ricardo J & Eduardo MRA Engel (2007), "Price stickiness in Ss models: New interpretations of old results," *Journal of Monetary Economics* 54: 100–121.

Calvo, Guillermo A (1983), "Staggered prices in a utility-maximizing framework," *Journal of Monetary Economics* 12(3): 383–398.

Campbell, Jeffrey R & Benjamin Eden (2014), "Rigid prices: Evidence from us scanner data," *International Economic Review* 55(2): 423–442.

Campbell, Jeffrey R, Jonas DM Fisher, Alejandro Justiniano & Leonardo Melosi (2017), "Forward guidance and macroeconomic outcomes since the financial crisis," *NBER Macroeconomics Annual* 31(1): 283–357.

Carroll, Christopher D (2003), "Macroeconomic expectations of households and professional forecasters," *The Quarterly Journal of Economics* 118(1): 269–298.

Carvalho, Carlos, Niels Arne Dam & Jae Won Lee (2020), "The cross-sectional distribution of price stickiness implied by aggregate data," *Review of Economics and Statistics* 102(1): 162–179.

Cavallo, Alberto, Guillermo Cruces & Ricardo Perez-Truglia (2017), "Inflation expectations, learning, and supermarket prices: Evidence from survey experiments," 9(3): 1–35.

Coibion, Olivier & Yuriy Gorodnichenko (2012), "What can survey forecasts tell us about information rigidities?" *Journal of Political Economy* 120(1): 116–159.

Coibion, Olivier & Yuriy Gorodnichenko (2015), "Information rigidity and the expectations formation process: A simple framework and new facts," 105(8): 2644–78.

Coibion, Olivier, Yuriy Gorodnichenko & Saten Kumar (2018), "How do firms form their expectations? New survey evidence," *American Economic Review* 108(9): 2671–2713.

Coibion, Olivier, Yuriy Gorodnichenko & Tiziano Ropele (2020), "Inflation expectations and firm decisions: New causal evidence," *The Quarterly Journal of Economics* 135(1): 165–219.

Costain, James & Anton Nakov (2011), "Distributional dynamics under smoothly state-dependent pricing," *Journal of Monetary Economics* 58(6-8): 646–665.

Costain, James & Anton Nakov (2015), "Precautionary price stickiness," *Journal of Economic Dynamics and Control* 58: 218–234.

Costain, James & Anton Nakov (2019), "Logit price dynamics," *Journal of Money, Credit and Banking* 51(1): 43–78.

Cuba-Borda, Pablo & Sanjay R Singh (2019), "Understanding persistent stagnation," *Available at SSRN 3247350* .

David, Joel M, Hugo A Hopenhayn & Venky Venkateswaran (2016), "Information, misallocation, and aggregate productivity," *The Quarterly Journal of Economics* 131(2): 943–1005.

David, Joel M & Venky Venkateswaran (2019), "The sources of capital misallocation," *American Economic Review* 109(7): 2531–2567.

Dean, Mark & Nathaniel Neligh (2023), "Experimental tests of rational inattention," *Journal of Political Economy* 131(12): 3415–3461.

Del Negro, Marco & Frank Schorfheide (2008), "Forming priors for DSGE models (and how it affects the assessment of nominal rigidities)," *Journal of Monetary Economics* 55(7): 1191–1208, ISSN 0304-3932.

Dotsey, Michael, Robert G. King & Alexander L. Wolman (1999), "State-dependent pricing and the general equilibrium dynamics of money and output," *Quarterly Journal of Economics* 114: 655–690.

Eggertsson, Gauti B & Michael Woodford (2003), "The zero bound on interest rates and optimal monetary policy," *Brookings Papers on Economic Activity* 2003(1): 139–233.

Ferroni, Filippo, Stefano Grassi & Miguel A. León-Ledesma (2019), "Selecting structural innovations in DSGE models," *Journal of Applied Econometrics* 34(2): 205–220.

Fisher, Jonas DM (2015), "On the structural interpretation of the Smets–Wouters risk premium shock," *Journal of Money, Credit and Banking* 47(2-3): 511–516.

Flynn, Joel P & Karthik Sastry (2024), "Attention cycles," Tech. rep., National Bureau of Economic Research.

Fuster, Andreas, Ricardo Perez-Truglia, Mirko Wiederholt & Basit Zafar (2022), "Expectations with endogenous information acquisition: An experimental investigation," *Review of Economics and Statistics* 104(5): 1059–1078.

Fuster, Andreas & Basit Zafar (2023), "Survey experiments on economic expectations," in *Handbook of Economic Expectations*, pp. 107–130, Elsevier.

Gagliardone, Luca, Mark Gertler, Simone Lenzu & Joris Tielens (2024), "Anatomy of the Phillips Curve: Micro Evidence and Macro Implications," Tech. rep., National Bureau of Economic Research.

Gagliardone, Luca, Mark Gertler, Simone Lenzu & Joris Tielens (2025), "Micro and Macro Cost-price Dynamics during inflation surges versus normal times," Tech. rep.

Gautier, Erwan & Hervé Le Bihan (2022), "Shocks versus menu costs: Patterns of price rigidity in an estimated multisector menu-cost model," *Review of Economics and Statistics* 104(4): 668–685.

Golosov, Mikhail & Robert E Lucas Jr (2007), "Menu costs and Phillips curves," *Journal of Political Economy* 115(2): 171–199.

Gorodnichenko, Yuriy & Michael Weber (2016), "Are sticky prices costly? Evidence from the stock market," *American Economic Review* 106(01): 165–199.

Guerrieri, Luca & Matteo Iacoviello (2015), "OccBin: A toolkit for solving dynamic models with occasionally binding constraints easily," *Journal of Monetary Economics* 70: 22–38.

Havranek, Tomas, Marek Rusnak & Anna Sokolova (2017), "Habit formation in consumption: A meta-analysis," *European Economic Review* 95: 142–167.

Head, Allen, Lucy Qian Liu, Guido Menzio & Randall Wright (2012), "Sticky prices: A new monetarist approach," *Journal of the European Economic Association* 10(5): 939–973.

Ilut, Cosmin, Rosen Valchev & Nicolas Vincent (2020), "Paralyzed by fear: Rigid and discrete pricing under demand uncertainty," *Econometrica* 88(5): 1899–1938.

Jurado, Kyle, Sydney C Ludvigson & Serena Ng (2015), "Measuring uncertainty," *American Economic Review* 105(3): 1177–1216.

Justiniano, Alejandro, Giorgio E Primiceri & Andrea Tambalotti (2010), "Investment shocks and business cycles," *Journal of Monetary Economics* 57(2): 132–145.

Karadi, Peter, Juergen Amann, Javier Sánchez Bachiller, Pascal Seiler & Jesse Wursten (2023), "Price setting on the two sides of the Atlantic-Evidence from supermarket scanner data," *Journal of Monetary Economics* 140: S1–S17.

Karadi, Peter & Adam Reiff (2019), "Menu costs, aggregate fluctuations, and large shocks," 11(3): 111–46.

Khaw, Mel Win, Ziang Li & Michael Woodford (2021), "Cognitive imprecision and small-stakes risk aversion," *The review of economic studies* 88(4): 1979–2013.

Khaw, Mel Win, Luminita Stevens & Michael Woodford (2017), "Discrete adjustment to a changing environment: Experimental evidence," *Journal of Monetary Economics* 91: 88–103.

Klenow, Peter J & Jonathan L Willis (2007), "Sticky information and sticky prices," *Journal of Monetary Economics* 54: 79–99.

Krishnamurthy, Arvind & Annette Vissing-Jorgensen (2012), "The aggregate demand for treasury debt," *Journal of Political Economy* 120(2): 233–267.

Kulish, Mariano, James Morley & Tim Robinson (2017), "Estimating DSGE models with zero interest rate policy," *Journal of Monetary Economics* 88: 35–49.

Ludvigson, Sydney C, Sai Ma & Serena Ng (2021), "Uncertainty and business cycles: exogenous impulse or endogenous response?" *American Economic Journal: Macroeconomics* 13(4): 369–410.

Luo, Shaowen & Daniel Villar (2021), "The skewness of the price change distribution: A new touchstone for sticky price models," *Journal of Money, Credit and Banking* 53(1): 41–72.

Maćkowiak, Bartosz A. & Mirko Wiederholt (2009), "Optimal sticky prices under rational inattention," *The American Economic Review* 99(3): 769–803.

Magnani, Jacopo, Aspen Gorry & Ryan Oprea (2016), "Time and state dependence in an Ss decision experiment," *American Economic Journal: Macroeconomics* 8(1): 285–310.

Mankiw, N Gregory, Ricardo Reis & Justin Wolfers (2003), "Disagreement about inflation expectations," 18: 209–248.

Matějka, Filip (2015), "Rationally inattentive seller: Sales and discrete pricing," *The Review of Economic Studies* 83(3): 1125–1155.

Melosi, Leonardo (2014), "Estimating models with dispersed information," *American Economic Journal: Macroeconomics* 6(1): 1–31.

Michaillat, Pascal & Emmanuel Saez (2021), "Resolving New Keynesian anomalies with wealth in the utility function," *Review of Economics and Statistics* 103(2): 197–215.

Midrigan, Virgiliu (2011), "Menu costs, multiproduct firms, and aggregate fluctuations," *Econometrica* 79(4): 1139–1180.

Montag, Hugh & Daniel Villar (2023), *Price-Setting During the Covid Era*, FEDS Notes.

Nakamura, Emi & Jón Steinsson (2008), "Five facts about prices: A reevaluation of menu cost models," *The Quarterly Journal of Economics* 123(4): 1415–1464.

Nakamura, Emi & Jón Steinsson (2010), "Monetary non-neutrality in a multisector menu cost model," *The Quarterly Journal of Economics* 125(3): 961–1013.

Nakamura, Emi, Jón Steinsson, Patrick Sun & Daniel Villar (2018), "The elusive costs of inflation: Price dispersion during the US Great Inflation," *The Quarterly Journal of Economics* 133(4): 1933–1980.

Nimark, Kristoffer (2008), "Dynamic pricing and imperfect common knowledge," *Journal of Monetary Economics* 55(2): 365–382.

Reis, Ricardo (2006), "Inattentive producers," *The Review of Economic Studies* 73(3): 793–821.

Reiter, Michael (2009), "Solving heterogeneous-agent models by projection and perturbation," *Journal of Economic Dynamics and Control* 33(3): 649–665, ISSN 0165-1889.

Shannon, Claude E (1948), "A mathematical theory of communication," *The Bell System Technical Journal* 27(3): 379–423.

Shannon, Claude E. (1959), "Coding theorems for a discrete source with a fidelity criterion," *IRE Nat. Conv. Rec* 4(142-163): 1.

Sims, Christopher A. (2003), "Implications of rational inattention," *Journal of Monetary Economics* 50(3): 665–690.

Sims, Eric (2013), "Growth or the Gap? Which Measure of Economic Activity Should be Targeted in Interest Rate Rules?" *Mimeograph, University of Notre Dame* .

Stevens, Luminita (2020), "Coarse pricing policies," *The Review of Economic Studies* 87(1): 420–453.

Turen, Javier (2023), "State-dependent attention and pricing decisions," *American Economic Journal: Macroeconomics* 15(2): 161–189.

Vavra, Joseph (2013), "Inflation dynamics and time-varying volatility: New evidence and an Ss interpretation," *The Quarterly Journal of Economics* 129(1): 215–258.

Veldkamp, Laura & Cindy Chung (2024), "Data and the aggregate economy," *Journal of Economic Literature* 62(2): 458–484.

Woodford, Michael (2003), "Imperfect common knowledge and the effects of monetary policy," in *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, J. Stiglitz P. Aghion, R. Frydman & M. Woodford, eds., pp. 25–58, Princeton University Press, Princeton, NJ.

Woodford, Michael (2009), "Information-constrained state-dependent pricing," *Journal of Monetary Economics* 56(S): 100–124.

Woodford, Michael (2020), "Modeling imprecision in perception, valuation, and choice," *Annual Review of Economics* 12: 579–601.

A Estimation: SSJ with Occasionally Binding ELB

For linear models written in a recursive formulation, Kulish et al. (2017) and Guerrieri & Iacoviello (2015) show that, for an expected ELB duration, the law of motion of the economy can be written as a time-varying linear policy function. As a result, the likelihood evaluation can be computed based the Kalman filter with time-varying coefficients. For example, Kulish et al. (2017) show that DSGE models can be estimated for sample periods including the ELB period by replacing the federal funds rate with a time series of expected ELB durations as an observable. The expected ELB duration does not have to be model consistent. In other words, in absence of any other shocks, the federal funds rate could be expected to be above the ELB at a time period different than the implied by the expected ELB duration, adding another form of monetary policy shocks. In the context of this paper, where we use SSJ to solve and estimate the model, how can we estimate the model during the ELB period? We show that one possibility is to recover the time-varying and recursive formulation of the model based on SSJ IRFs for different ELB durations. Then, we can find the (time-varying) MA representation of the model to compute the log-likelihood of the model. However, this approach is time consuming as the matrix operations can be computationally demanding for large number of state variables and a large number of draws.

Instead, we propose a new and efficient way of using SSJ to compute the log-likelihood during the ELB period for a given expected ELB duration. Our approach is to model the expected ELB duration as a sequence of anticipated monetary policy shocks. For each month that the federal funds rate is at the ELB, given a sequence of shocks up to that month, and given the expected duration of the ELB in that month, agents in the economy expected a sequence of anticipated monetary policy shocks for as long as they expected the ELB to bind. First, we show how to compute impulse responses with SSJ when the ELB is expected to bind for m periods, and we show that those responses are identical to what we would get based on the recursive formulation of the model based on Reiter (2009) and Kulish et al. (2017). Second, we show how to get the time-varying recursive formulation of the policy functions based on the IRFs for different ELB durations, and we show that the responses from this method also match the previous methods. Next, we show that these responses are also equivalent to assuming a sequence of anticipated monetary policy shocks for m periods. We then combine these results to show how to evaluate the likelihood of the data when the ELB is binding during some periods. These steps are detailed in the Online Appendix.

B Prior and Posterior Distributions

The estimation includes the shocks listed in the model description plus shocks to yield curve $(\eta_t^2, \eta_t^5, \eta_t^y)$, to aid the estimation at the ELB. Table B.1 reports the prior distributions for these parameters along with the posterior mode and standard deviation. We use a normal distribution for shocks affecting the pricing frictions, long-term interest rates, and labor supply. Following Ferroni, Grassi & León-Ledesma (2019), we selected this prior, whose support includes zero, to avoid estimating "nonexistent" shocks, as it does not force the exogenous processes to be different from zero. For all other standard deviations, we select a standard inverse gamma distribution. Following Del Negro & Schorfheide (2008), we impose an implicit prior over the model-implied variance of the observable variables and the covariance between the federal funds rate, GDP growth, and inflation. Our prior states that those specific model-implied-covariances follow a normal distribution with parameters determined by the data moments.[20] Finally, to avoid exploring unreasonable areas of the parameter space, we impose that the the filtered value of the menu cost is not below zero and that the filtered log-deviations for θ_p and θ_a are not greater than three in absolute value.

[20]For example, we imposed that the model-implied variance for the federal funds rate is normal with mean (μ_i) and standard deviation (σ_i) equal to the mean and ω times standard deviation of the quadratic deviation of the federal funds rate with respect to the model's steady state, where ω is a scalar that we set to 0.25. Hence, $\mu_i = \sum_t x_t^i/T$ and $\sigma_i^2 = \omega^2 \sum_t \left(x_t^i - \mu_i\right)^2/T$, where $x_t^i = \left(i_t^{data} - i_{ss}\right)^2$. Appendix A presents formally this implicit prior.

TABLE B.1: Model Priors and Posteriors

Parameter Name		Prior			Posterior			
		Dist	Mode	SD	Mode	SD	90% HPD	
Standard deviation of shock innovations (x100)								
Price setting cost	θ^p	norm	1.000	0.500	3.538	0.077	3.445	3.649
Price review cost	θ^a	norm	1.000	0.500	1.795	0.100	1.710	1.987
Menu cost	κ	norm	0.026	0.013	0.091	0.001	0.090	0.093
Idiosyncratic risk	σ	norm	0.008	0.002	0.034	0.001	0.033	0.034
Markup	ε	invg	0.055	0.014	5.694	0.025	5.659	5.725
Permanent TFP	γ	invg	0.190	0.048	1.434	0.049	1.362	1.479
Impatience	ζ	invg	1.118	0.280	2.524	0.028	2.491	2.562
Transitory TFP	a	invg	0.165	0.041	19.147	0.120	19.022	19.330
Bond demand	χ_t	invg	2.333	0.583	99.960	0.315	99.248	99.966
Labor supply	ξ	norm	0.000	0.041	0.004	0.024	0.005	0.065
Trend inflation	e_t^π	invg	0.116	0.029	0.117	0.041	0.084	0.194
Monetary policy	e_t^i	invg	0.083	0.021	0.574	0.038	0.520	0.615
Term premia	η^y	norm	0.000	1.000	0.667	0.036	0.626	0.721
2yrs yield	η^2	norm	1.211	0.605	1.480	0.089	1.404	1.632
5yrs yield	η^5	norm	1.138	0.569	0.011	0.047	0.011	0.131
Autocorrelation								
Price setting cost	θ^p	beta	0.500	0.150	0.959	0.001	0.957	0.961
Price review cost	θ^a	beta	0.500	0.150	0.980	0.002	0.977	0.981
Menu cost	κ	beta	0.500	0.150	0.854	0.003	0.850	0.857
Idiosyncratic risk	σ	beta	0.500	0.150	0.958	0.001	0.956	0.959
Markup	ε	beta	0.500	0.150	0.002	0.001	0.001	0.005
Permanent TFP	γ	beta	0.025	0.150	0.003	0.024	0.003	0.054
Impatience	ζ	beta	0.500	0.150	0.974	0.001	0.973	0.975
Transitory TFP	a	beta	0.900	0.150	0.001	0.000	0.000	0.002
Bond demand	χ_t	beta	0.900	0.150	0.502	0.003	0.498	0.505
Labor supply	ξ	beta	0.025	0.150	0.032	0.120	0.028	0.347
Trend inflation	e_t^π	beta	0.500	0.150	0.401	0.148	0.362	0.798
Term premia	η^y	beta	0.500	0.150	0.976	0.003	0.971	0.980

Notes: norm and invg refer to the normal and inverse gamma distributions, respectively. HPD: High Probability Density.

ONLINE APPENDIX: Price Rigidities in U.S. Business Cycles*

Camilo Morales-Jiménez
Federal Reserve Board

Luminita Stevens
University of Maryland

March 19, 2025

A Data Description

We set our model to a monthly frequency and aggregate the model-simulated series to a quarterly frequency to align the model-simulated data with the U.S. data. In these notes, the time subscript t refers to a month, and $\tilde{x}$ is variable x detrended by trend-productivity.

Pricing Moments The pricing moments we use are the frequency, size, standard deviation, skewness, and kurtosis of log-price changes, and the frequency and size of log-price increases and decreases. We take the quarterly averages for each series and we smooth them using an 11-term moving average (MA), corresponding to a smoothing window equal to 5 on each side of a central observation. At the beginning and end of the sample, where fewer observations are available, a lower degree of smoothing is applied by reducing the smoothing window and ensuring all data points are included without truncation. The online appendix includes a figure with the raw and smoothed time series of the different pricing moments we use. The microdata underlying these series consist of approximately 80,000 monthly price quotes for products grouped into about 305 categories, or "entry-level items" (ELIs), which are then further aggregated into 13 major groups. Authors with access to the microdata can use the individual price quotes to construct empirical distributions of price changes for each month, from which various pricing statistics are then calculated. For example, to calculate the frequency of price changes in each period, one computes the fraction of nonzero price changes across products within each ELI, and then the expenditure-weighted median across ELIs. Similarly, conditional on a price change, the absolute size of price changes is computed by taking the average log price change across products within each ELI and then the

*The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of anyone else associated with the Federal Reserve System. Contact: camilo.moralesjimenez@frb.gov, stevens7@umd.edu.

expenditure-weighted median across ELIs. Skewness and kurtosis are computed in a similar way, but by pooling data within each *major group* rather than within each ELI, and then taking an expenditure-weighted average across the 13 major groups.[1]

Growth rate of real per-capita GDP Data for quarterly real GDP of chained 2012 dollars and seasonally adjusted was retrieved from FRED (series GDPC1). This series is divided by the quarterly average of the monthly civilian noninstitutional population, 16 yr+ available at FRED (series CNP16OV). Growth rates are computed by log differences (quarterly growth rate). The quarterly GDP growth is linked to the model variable as follows:

$$dy_t^q = \log\left(\tilde{y}_t + \frac{\tilde{y}_{t-1}}{\gamma_{zt}} + \frac{\tilde{y}_{t-2}}{\gamma_{zt-1}}\right) - \log\left(\tilde{y}_{t-3} + \frac{\tilde{y}_{t-4}}{\gamma_{zt-3}} + \frac{\tilde{y}_{t-5}}{\gamma_{zt-4}}\right) + \log\left(\gamma_{zt}\gamma_{zt-1}\gamma_{zt-2}\right) \tag{1}$$

where $\tilde{y}_t \equiv \tilde{c}_t - \tilde{G}_t$.

Quarterly CPI inflation rate Data for the Consumer Price Index for All Urban Consumer was retrieved from FRED (series CPIAUCSL). We take the quarterly average of this monthly series and compute quarterly inflation as the ratio of the quarterly CPI index between two months minus one. The quarterly inflation rate is linked to the model variable as follows:

$$\pi_t^q = \pi_t \cdot \pi_{t-1} \cdot \pi_{t-2} - 1. \tag{2}$$

Quarterly Federal Funds Rate We retrieve the Federal Funds rate from FRED (series DFF) and compute the quarterly average. The quarterly federal funds rate in the data (i_t^q) is linked to the model variables as follows:

$$i_t^q = \frac{1}{3}\left(i_t^{12} + i_{t-1}^{12} + i_{t-2}^{12}\right). \tag{3}$$

Quarterly 2 year and 5 year Treasury Yields We retrieve data for the Market Yield on U.S. Treasury Securities at 2-Year and 5-Year Constant Maturity from FRED (series DGS2). We take the quarterly average of this series. We map the long-term interest rates to our model based on the expectations hypothesis by relating yields on long-term bonds with

[1] Luo & Villar (2021) discuss the need to compute higher moments at the group level, due to sample size constraints at the ELI level: Since ELIs are narrowly defined categories, they sometimes have only a small number of observations. Higher moments are particularly sensitive to outliers, which is why a small number of observations is insufficient to compute them reliably within each ELI. On the other hand, pooling across ELIs introduces more heterogeneity across the products pooled.

agent's beliefs about the future path of the federal funds rate:

$$
\begin{aligned}
\text{2 year treasury yield} =& c^2 + \frac{1}{24} E_t \left[\sum_{j=0}^{23} i^{12}_{t+j} \right] + \eta^2_t + \eta^y_t && (4) \\
\text{5 year treasury yield} =& c^5 + \frac{1}{60} E_t \left[\sum_{j=0}^{59} i^{12}_{t+j} \right] + \eta^5_t + \eta^y_t && (5)
\end{aligned}
$$

where c^2 and c^5 are yield specific, time invariant components set to match the average difference between the federal funds rate and these yields, η^2, and η^5 are i.i.d. yield specific shocks, and η^y is an exogenous AR(1) process, common to all yields.

Recession Dates The NBER dating committee lists dates for peaks and troughs in economic activity. In the NBER's convention, the first month of a recession is the month following the peak, and the last month of a recession is the month of a trough. Therefore, we define the start month of a recession as peak plus one month and the end month of a recession as the trough. For example, in 2020, the peak economic activity was reached in February 2020 and the trough was reached in April 2020, yielding a two-month recession: March and April 2020.

The Effective Lower Bound The federal funds rate was at the effective zero lower bound (ELB) twice during our sample period: between January 2009 and December 2015, following the Great Recession, and between March 2020 and February 2022, following the Covid Recession. We use the Blue Chip data between 2008 and 2010, and the survey of primary dealers since January 2011 to construct a series with the expected ELB duration, in months, for each month in which the federal funds rate was at the ELB. Figure 1 plots the expected ELB duration series used in the Bayesian estimation.

This series is constructed as follows: Using the Blue Chip microdata, we compute the expected ELB duration (in quarters) for each forecaster and month, and then we take the median for each month across forecasters. Blue Chip is a monthly survey, but participants are asked for the expected federal funds rate in quarter intervals. We compute the expected duration in months assuming a lag of one week between data collection and publication and using the FOMC meetings calendar to determine the expected month of the liftoff. If a quarter has two FOMC meetings, we take the average of the expected duration associated with those two meetings. After 2010, we use the Survey of Primary Dealers. Between January 2011 and 2015, the survey asks for the *"Most Likely Quarter and Year of First Target Rate Increase"*. Subsequently, it asks for the specific FOMC meeting in which the target

Figure 1: Expected ELB duration

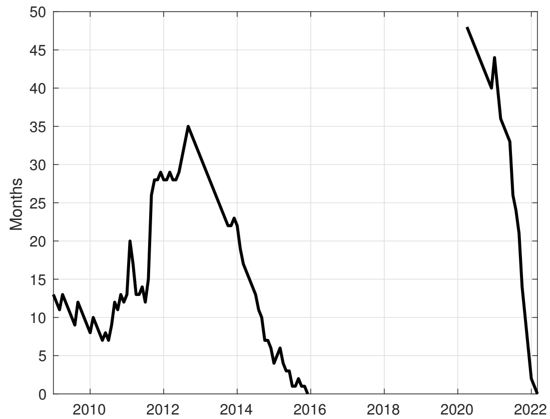

Note: This Figure plots the expected ELB duration for each month used in the Bayesian estimation of the model. We use the BlueChip survey between 2008 and 2010, and the survey of primary dealers since 2011.

rate will be increased. Based on the median answer for those questions, the date in which the survey was received, and the FOMC meetings calendar, we compute the expected ELB duration in months. As with the Blue Chip survey, we take the simple average among the expected durations associated with the meetings in each quarter, and we get a monthly series by interpolating and rounding.

B Household Optimality Conditions

Let λ_t denote the Lagrange multiplier on the flow budget constraint. Household optimization yields

$$\left(\frac{1}{C_t - hC_{t-1}}\right) - E_t\left[\left(\frac{\beta\zeta_{t+1}}{\zeta_t}\right)\left(\frac{h}{C_{t+1} - hC_t}\right)\right] = \lambda_t, \tag{6}$$

$$\xi_t L_t^{\frac{1}{\nu}} = \lambda_t w_t, \tag{7}$$

$$\zeta_t\,\chi_t + E_t\left[\beta\,\zeta_{t+1}\,\lambda_{t+1}\left(\frac{i_t}{\pi_{t+1}}\right)\right] = \zeta_t\,\lambda_t, \tag{8}$$

along with thr real discount factor between t and $t+1$,

$$M_{t,t+1} \equiv \frac{\beta\zeta_{t+1}\lambda_{t+1}}{\zeta_t\lambda_t}. \tag{9}$$

C Alternative Reference Distributions

One novelty of the model we consider is the fact that firms use the equilibrium distribution of prices as their reference distribution. This feature introduces an interesting layer of strategic complementarities (SC): Whereas the typical SC story is that each firm has incentives to keep its price close to that of its competitors, here each firm benefits from keeping the *distribution* from which it draws its prices close to the distribution of its competitors' prices. How does this fixed point relationship between individual policies and the aggregate distribution of prices differ from existing alternatives?

One alternative is the control cost model of stochastic choice. As the name suggests, these are models of costly control, rather than costly information. They allow for errors in the implementation of actions, when the optimal action is known in each state. They feature a reference distribution for each period that is uniform around the optimal action in that period. Exerting effort to deviate from this default amounts to choosing a more concentrated distribution to draw an action and entails a cost proportional to the divergence of the chosen distribution from that uniform. For example, Costain & Nakov (2019) apply control costs to price-setting in a general equilibrium monetary model. Since the reference distribution is always centered on the optimal price, featuring control costs on pricing helps match micro volatility moments well, but results in no meaningful additional rigidity: On average, reset prices adjust fully to aggregate conditions, and the frequency and state-dependence in the timing of price changes remain the main determinants of non-neutrality. Conversely, our model generates endogenous bias in repricing, which makes pricing mistakes highly relevant to aggregate outcomes.

On the other hand, rational inattention (RI) models of costly information imply not only an endogenous reference distribution, but one that is furthermore optimized to the decision-maker's individual problem at hand. Rational decision-makers have strong incentives to develop sophisticated reference or default probabilities. A well-chosen default distribution can lower both the relative value of conditioning actions on the state in real time, as well as the cost of doing so. Hence, a rational decision-maker would want to use knowledge about the structure of the economy, the laws of motion of the shocks, and the shape of the objective function to choose well-adapted reference distributions that can serve as no-cost defaults. What does a well-adapted no-cost default look like? In the RI model, it is one that gets as close as possible to conditioning on the state in real time, without actually doing so. Formally, the optimal reference distribution minimizes the choice distribution's average KL divergence from it, integrating over the distribution of possible states of the world that the decision-maker can expect to encounter. The pure RI model does not introduce any new parameters, but it places a strong restriction on what the reference distribution can be, by imposing that it minimizes the KL divergence of the choice distribution that the firm uses over the life of its policy. In our steady state estimations, we found that this formulation does not fit the data well: the reference distribution is too concentrated, yielding too little dispersion. We also found that trying to run the Bayesian estimation became intractable with the KL-divergence minimizing reference distribution.

Instead, we propose using less efficient, though still endogenous reference distributions that take into account both the structure of the economy and the actions of others. It is in this sense that the model is behavioral RI model.[2] In addition to fitting the steady state price dispersion better and making the dynamic estimation feasible, this assumption implies that the reference distributions do not fully discount states in which the current policy is likely to be revised.

D Estimating the Steady State

In the steady state estimation of the pricing parameters, we make use of the fact that pricing rigidities are determined by the value of the pricing parameters relative to steady state output. Hence, in our steady state estimation, without loss of generality, we normalize $Y = 1$ to estimate the pricing parameters $theta^p$, θ^a, and κ:

Proposition 1. *In steady state, for a given real input price p^x and aggregate demand Y, if $\Lambda(\tilde{p}, a)$ and $f(p \,|\, a)$ solve the firm's problem for pricing parameters θ^p, θ^a, and κ, then $\Lambda(\tilde{p}, a)$ and $f(p \,|\, a)$ also solve the firm's problem for real input price p^x, aggregate demand $\tilde{Y}$, and for any set of pricing parameters $\tilde{\theta}^p$, $\tilde{\theta}^a$, and $\tilde{\kappa}$ such that:*

$$\frac{\tilde{\theta}^p}{\tilde{Y}} = \frac{\theta^p}{Y}, \quad \frac{\tilde{\theta}^a}{\tilde{Y}} = \frac{\theta^a}{Y}, \quad \frac{\tilde{\kappa}}{\tilde{Y}} = \frac{\kappa}{Y}. \tag{10}$$

Proof. First, we show that the value of the firm is homogeneous of degree 1 in aggregate output as long as the pricing parameters (θ^p, θ^a, and κ) are constant relative to aggregate output. Using this result, we show that the optimal choice distributions are homogeneous of degree 0 in aggregate output as long as the pricing parameters are constant relative to aggregate output.

[2] See Woodford (2012) and Khaw, Stevens & Woodford (2019) for alternative deviations of the default distributions from the RI optima.

Value of a firm Denoting $\bar{x} = \frac{x}{Y}$ for $x = \theta^p, \theta^a, \kappa$, the value of a firm in steady state is

$$V(p_j, a_j; Y) = \pi(p_j, a_j; Y) + E\Big\{ M\, V^*\left(\tilde{p}_j, a_j; Y\right) \Big\} \tag{11}$$

where $\tilde{p} = \frac{p}{\pi}$, $\pi(p_j, a_j; Y) = Y p_j^{-\varepsilon}\left(p_j - p^x\right)$, and

$$V^*(p, a; Y) = \max_{\Lambda} \Big\{ \Lambda(a) \cdot \left[V^a(a; Y) - \bar{\kappa} Y\right] + (1 - \Lambda(a)) \cdot V(p, a; Y) - \bar{\theta}^a Y\, \mathcal{D}\left(\Lambda(a) \parallel \bar{\Lambda}\right) \Big\}, \tag{12}$$

$$V^a(a; Y) = \max_{f} \left\{ \int f(p \,|\, a) V(p, a; Y)\, dp - \bar{\theta}^p Y \mathcal{D}\left(f(p \,|\, a) \parallel \bar{f}(p)\right) \right\} \tag{13}$$

$$\text{s.t.} \quad \int f(p \,|\, a)\, dp = 1. \tag{14}$$

Substituting, and denoting the optimal distributions by Λ^* and $f^*(p)$,

$$\begin{aligned} V(p_j, a_j; Y) = {} & \pi(p_j, a_j; Y) \\ & + E\Big\{ M \Lambda^* \cdot \left[\int f(p \,|\, a) V(p, a; Y)\, dp - \bar{\theta}^p Y \mathcal{D}\left(f(p \,|\, a) \parallel \bar{f}(p)\right) - \bar{\kappa} Y \right] \Big\} \\ & + E\Big\{ M\left(1 - \Lambda^*\right) \cdot V(\tilde{p}, a; Y) - M \bar{\theta}^a Y\, \mathcal{D}\left(\Lambda^* \parallel \bar{\Lambda}\right) \Big\} \end{aligned} \tag{15}$$

Hence, note that the value of the firm is H(1) in aggregate output,

$$V(p_j, a_j; Y) = Y \cdot V(p_j, a_j; 1), \tag{16}$$

$$\begin{aligned} V(p_j, a_j; 1) = {} & \pi(p_j, a_j; 1) \\ & + E\Big\{ M \Lambda^* \cdot \left[\int f(p \,|\, a) V(p, a; 1)\, dp - \bar{\theta}^p \mathcal{D}\left(f(p \,|\, a) \parallel \bar{f}(p)\right) - \bar{\kappa} \right] \Big\} \\ & + E\Big\{ M\left(1 - \Lambda^*\right) \cdot V(\tilde{p}, a; 1) - M \bar{\theta}^a\, \mathcal{D}\left(\Lambda^* \parallel \bar{\Lambda}\right) \Big\} \end{aligned} \tag{17}$$

Substituting this result in $V^a(a; Y)$, we can also verify that $V^a(a; Y) = Y \cdot V^a(a; 1)$.

Optimal choice distribution The review distribution in steady state is

$$\ln\left(\frac{\Lambda(\tilde{p}, a)}{1 - \Lambda(\tilde{p}, a)}\right) = \ln\left(\frac{\bar{\Lambda}}{1 - \bar{\Lambda}}\right) + \frac{1}{\theta^a}\left[V^{a;Y}(a) - V(\tilde{p}, a; Y) - \kappa\right] \tag{18}$$

Given that the value of the firm is H(1) in aggregate output, and substituting $\bar{x} = \frac{x}{Y}$ for

$x = \theta^p, \theta^a, \kappa, \Lambda(\tilde{p}, a)$ is H(0) in aggregate output:

$$\ln\left(\frac{\Lambda(\tilde{p},a)}{1-\Lambda(\tilde{p},a)}\right) = \ln\left(\frac{\bar{\Lambda}}{1-\bar{\Lambda}}\right) + \frac{1}{\bar{\theta}^a Y}\left[V^a(a;1)Y - V(\tilde{p},a;1)Y - \bar{\kappa}Y\right]$$
$$= \ln\left(\frac{\bar{\Lambda}}{1-\bar{\Lambda}}\right) + \frac{1}{\bar{\theta}^a}\left[V^a(a;1) - V(\tilde{p},a;1) - \bar{\kappa}\right] \tag{19}$$

Analogously, the repricing distribution is also H(0) in aggregate output:

$$f(p\,|\,a) = \frac{\bar{f}(p)\,\exp\left\{\frac{V(p,a;Y)}{\theta^p}\right\}}{\int \bar{f}(\hat{p})\,\exp\left\{\frac{V(\hat{p},a;Y)}{\theta^p}\right\}\,d\hat{p}} = \frac{\bar{f}(p)\,\exp\left\{\frac{V(p,a;1)Y}{\bar{\theta}^p Y}\right\}}{\int \bar{f}(\hat{p})\,\exp\left\{\frac{V(\hat{p},a;1)Y}{\bar{\theta}^p Y}\right\}\,d\hat{p}} \tag{20}$$
$$= \frac{\bar{f}(p)\,\exp\left\{\frac{V(p,a;1)}{\bar{\theta}^p}\right\}}{\int \bar{f}(\hat{p})\,\exp\left\{\frac{V(\hat{p},a;1)}{\bar{\theta}^p}\right\}\,d\hat{p}}.$$

□

E Taylor Rule Estimation

TABLE I: Taylor Rule Paramters: Priors and Posterior

		Prior		Bounds	
Parameter	Mode	Mean	Std	Lower	Upper
ϕ_π	1.6977	2.0000	0.2500	1.00	3.00
ϕ_y	1.004	0.5000	0.2500	0.00	3.00
ρ_i	0.7914	0.5000	0.0500	0.00	0.95
ρ_π	0.2288	0.1200	0.0500	0.03	0.99
σ_i	0.3500	0.0625	0.0625	0.00	1.00
σ_π	0.0934	0.0625	0.0625	0.00	1.00

Notes: The table reports the posterior mode of the estimated Taylor rule parameter along with their prior mean, prior standard deviation, and lower and upper bounds. The prior distribution on all parameters is normal.

To estimate the parameters governing the Taylor rule, we run a Bayesian estimation on the federal funds rate during the Great Moderation period (1984Q1 to 2007Q2). The Taylor rule is given by

$$
\begin{aligned}
i_t &= \rho_i i_{t-1} + (1-\rho_i)\left[i_{ss} + \phi_\pi \left(\pi_t^a - \bar{\pi}^*\right) + \phi_y (dy_t^a - dy_{ss})\right] + \epsilon_t^i + e_t^\pi \\
e_t^\pi &= \rho_\pi e_{t-1}^\pi + \epsilon_t^\pi \\
\epsilon_t^i &\sim \mathcal{N}(0, \sigma_i), \quad \epsilon_t^\pi \sim \mathcal{N}(0, \sigma_\pi)
\end{aligned}
$$

where i_t is the federal funds rate at quarter t, π_t^a is the annual CPI inflation rate at quarter t, and dy_t^a is the annual quarterly GDP growth. ϵ^i is an i.i.d. monetary policy shock, and e_t^π is a persistent shock to the inflation target (π^*). We set i_{ss}, dy_{ss}, and $\bar{\pi}$ to the sample averages of the Federal Funds rate, annual quarterly GDP growth, and annual CPI inflation. For a set of parameters $\Theta = \{\rho_i, \phi_\pi, \phi_y, \rho^\pi, \sigma^i, \sigma^\pi\}$, we implement the Kalman filter on the Taylor rule and evaluate the likelihood. We impose a normal prior distribution on all these parameters, as shown in Table I. We restrict the Taylor rule parameter on inflation (ϕ_π) to be between 1 and 3, and our prior is that π_π distributes normal with mean equal to 2 and standard deviation equal to 0.25. Similarly, we impose the Taylor rule parameter on GDP growth (ϕ_y) to be between 0 and 3, and our prior is that ϕ_y distributes normal with mean equal to 0.5 and standard deviation equal to 0.25. We constraint the interest smoothing parameters (ρ_i) to be between 0 and 0.95, and we set the mean and stand deviations of its prior to 0.5 and 0.05. The predicted rate equals:

$$
i_t^p = \rho_i i_{t-1}^p + (1-\rho_i)\left[i_{ss} + \phi_\pi \left(\pi_t^a - \bar{\pi}^*\right) + \phi_y (dy_t^a - dy_{ss})\right], \ \ i_0^p = i_0.
$$

In general, the predicted rate tracks well the movements in the actual Federal Funds Rate, as shown in Figure 2, which plots the federal funds rate (solid line) and the predicted rate with the parameters evaluated at their mode (dashed line).

F Model Estimation: SSJ with Occasionally Binding ELB Constraint

We propose a new and efficient way of using SSJ to compute the log-likelihood during the ELB period for a given expected ELB duration. Our approach is to model the expected ELB duration as a sequence of anticipated monetary policy shocks. For each month that the federal funds rate is at the ELB, given a sequence of shocks up to that month, and given the expected duration of the ELB in that month, agents in the economy expected a sequence of anticipated monetary policy shocks for as long as they expected the ELB to bind. First, in

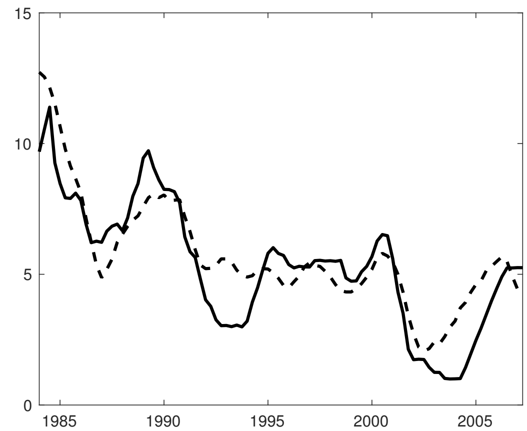

Figure 2: Federal Funds Rate and Fit

Note: This Figure plots the Federal Funds rate (solid line) along with the predict federal funds rate when the parameters are evaluated at their posterior mode (second column of Table I).

section F.1, we show how to compute impulse responses with SSJ when the ELB is expected to bind for m periods, and we show that those responses are identical to what we would get based on the recursive formulation of the model based on Reiter (2009) and Kulish, Morley & Robinson (2017). Second, in section F.2 we show how to get the time-varying recursive formulation of the policy functions based on the IRFs for different ELB durations, and we show that the responses from this method also match the previous methods. In section F.3, we show that these responses are also equivalent to assuming a sequence of anticipated monetary policy shocks for m periods. In section F.4, we combined these results to show how to evaluate the likelihood of the data when the ELB is binding during some periods.

F.1 IRFs when the ELB is binding using SSJ

Assuming that the ELB is not binding, the linearized system of equations describing the equilibrium in SSJ can be written as:

$$F_x dX + F_z dZ = 0 \tag{21}$$

where X and Z represent the paths of the endogenous and exogenous variables, respectively. The Jacobians F_x and F_z are stacked matrices that represent different parts of the model.

Hence, at the core of our model, the Jacobian F_x is given by:

$$F_x = \left[F_x^P; F_x^{px}; F_x^{ARC}; F_x^{Euler}; F_x^{Taylor}\right] \tag{22}$$
$$F_z = \left[F_z^P; F_z^{px}; F_z^{ARC}; F_z^{Euler}; F_z^{Taylor}\right] \tag{23}$$

where subscripts refer to the equations related to the aggregate price index (P), the aggregate marginal cost (px), the aggregate resource constraint (ARC), the euler equation (Euler), and the Taylor rule (Taylor). For example, F_x^{Taylor} and F_z^{Taylor} are the Jacobians describing the path of the linearized Taylor rule in response to a path of the endogenous variables (dX) and exogenous variables (dZ). F_x^{Taylor} is a matrix of size T by nT, and F_z^{Taylor} is a matrix of size T by eT, where T is a large horizon for which the Jacobian is computed, n is the number of endogenous variables, and e is the number of exogenous variables. In our case, the endogenous variables are the inflation rate (π), output (Y), consumption (C), interest rates (i), and the marginal cost (px). Hence, in SSJ, the responses of the endogenous variables to unanticipated shocks is given by:

$$dX = -F_x^{-1} F_x dZ \tag{24}$$

How to compute the responses when the interest rate is expected to bind for m periods? In this case, we would have to modify the Jacobian F_x^{Taylor} to reflect that the interest rate is at the ELB for the first m periods. Hence, in this case, the linearized system of equations describing the equilibrium is given by:

$$F_x^* dX + F_z^* dZ + C^* = 0 \tag{25}$$

where C^* is a column vector of size nT, and F_x^* and F_z^* are identical to F_x and F_z except for those rows describing the Taylor rule. Those Jacobians describing the path of the Taylor rule for the first m periods should reflect that in deviations with respect to the steady state:

$$\hat{i}_t = -i^{ss} \tag{26}$$

Hence, $F_x^{*Taylor}$ will be equal to the identity matrix for the first m periods, and then equal to F_x^{Taylor} for periods between $m+1$ and T. $F_z^{*Taylor}$ is a matrix of zeros, and C^{Taylor} is equal to $-i^{ss}$ for the first m periods and zero everywhere else.[3] In this case, the responses

[3] Here we assume that the ELB for the (net) interest rate is equal to 0. If the ELB is different than zero, $\hat{i}_t = -i^{ss} + \underline{i}$, where $\underline{i}$ is the lower bound in the interest rate.

to unanticipated shocks are given by:

$$dX = -F_x^{-1}\left(C^* + F_z^* dZ\right) \tag{27}$$

Figure 3 plots the interest rate and consumption responses to a 1% productivity shock when the interest rate is expected to be at 0 for 20 periods. The solid lines plot the responses based on the method described in this section (which we denote by "direct"), and the dash lines represent the same responses but computed based on the recursive formulation of the model (using the Reiter (2009) method) and then employing the method proposed by Kulish et al. (2017) for the ELB. These lines are almost identical and are on top of each other.[4][5]

Figure 3: IRFs at the Zero Lower Bound.

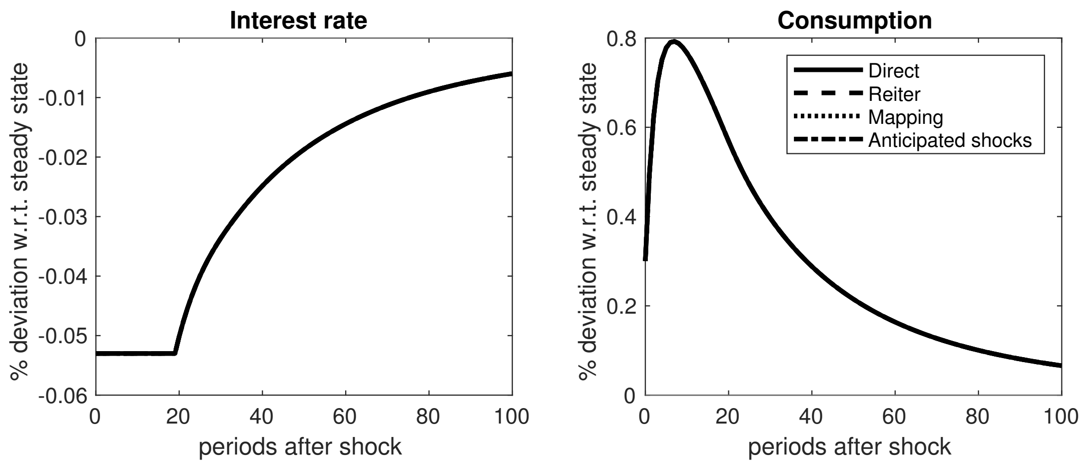

Note: This Figure plots the impulses responses to a 1% increase in productivity when the interest rate is expected to be at 0 for 20 periods. "Direct" refers to the approach described in section F.1. "Reiter" refers to the responses computed using the Reiter (2009) Method and Kulish et al. (2017). "Mapping" refer to the computation of the responses using the time-varying and recursive formulation of the model recovered from the SSJ solution as explained in section F.1.

F.2 Recovering time-varying recursive policy functions

Now, given that the goal is to evaluate the likelihood function, how can we use the Kalman filter and SSJ for this purpose? Auclert, Bardóczy, Rognlie & Straub (2021) show how to recover the unconditional recursive formulation of the policy functions given the IRFs

[4]The responses are also identical for all other endogenous variables in response to all exogenous shocks. We also checked that the responses are identical in both cases when the interest rate is expected to be at 0 between periods j and $j+m$, where $j > 1$. Additional graphs are available up to request.

[5]The results presented in this Appendix are based on a smaller grid for the marginal cost, as the computations with the Reiter method were taking hours with a grid size of 41 points for the marginal cost. In this Appendix, we used a grid with 5 points for the marginal cost.

computed based on SSJ. However, when the ELB is binding, the recursive formulation of the policy functions is time-varying. When the ELB is binding, we show that we can use the IRFs for different ELB durations computed in the previous section to recover the time-varying policy functions. In particular, denote the recursive formulation of the policy functions as:

$$X_t = C_t + P_t X_{t-1} + D_t E_t, \tag{28}$$

where matrices C_t, P_t, and D_t describe the policy rules at time t when the expected ELB duration is equal to $duration_t$. To recover these matrices from the SSJ solution when the ELB is expected to bind for m periods, one needs to (1) compute the IRFs for an expected duration equal to 1, (2) recover C_1, P_1, and D_1 from that solution following Auclert et al. (2021), and (3) repeat for each duration equal to 2, 3... m.

Dotted lines in Figure 3 represent the responses computed based on the time-varying and recursive formulation of the model recovered from the SSJ solution, which we denote by "mapping". These three lines are almost identical and are on top of each other.

F.3 ELB and anticipated monetary policy shocks

In this section, we show that responses presented above are equivalent to assuming anticipated monetary policy shocks that guarantee (in expectation) that the interest rate will be at 0 for m periods. While this method is less efficient to compute specific responses, this method is very efficient for the likelihood computation.

The basic idea of this method is the following: suppose that the monetary authority follows a Taylor-type policy rule. But, during ELB periods, the monetary authority activates anticipated shocks: today it announces future changes the the interest rate. Those anticipated shocks are such that the interest rate will be expected to be 0 for m periods.

Formally this procedure works as follows. The linearized system of equations describing the equilibrium in SSJ is given by:

$$F_x dX + F_z dZ + F_n dN = 0 \tag{29}$$

where, compared to (21), N is the path of the anticipated monetary policy shocks. dN is a column vector of size $T \cdot a$, where a is the number of anticipated monetary policy shocks.[6] F_n is the Jacobian of the endogenous equations with respect to the anticipated monetary

[6]For example, if the monetary authority announces shocks for the next 4 periods, $a = 4$.

policy shocks.[7] Then, the responses of the endogenous variables to shocks is given by:

$$dX = -F_x^{-1}\left(F_z dZ + F_n dN\right) \tag{30}$$

From 30, we can extract the responses of the interest rate, which we can denote by:

$$di = AdZ + BdN \tag{31}$$

Now, we want to find the sequence of anticipated monetary policy shocks (dN) that make the expected interest rate be 0 for m periods. Hence, based on (31), we can solve for dN such that:

$$di = C^i = AdZ + BdN \tag{32}$$

$$dN = -B^{-1}\left(C^i + AdZ\right) \tag{33}$$

where C^i is a column vector of size T, with the first m entries equal to $-i^{ss}$.[8] Now, substituting (33) into (30)

$$dX = -F_x^{-1}\left[F_z dZ - F_n B^{-1}\left(C^i + AdZ\right)\right] \tag{34}$$

$$dX = -F_x^{-1}\left[-F_n B^{-1} C^i + \left(F_z - F_n B^{-1} A\right) dZ\right] \tag{35}$$

Note the similarities between (35) and (27). Both expression would be identical if:

$$C^* = -F_n B^{-1} C^i \tag{36}$$

$$F_z^* = F_z - F_n B^{-1} A \tag{37}$$

In fact, Figure 3 also includes the responses to a 1% increase in productivity when the interest rate is expected to be at 0 for 20 periods based on the "anticipated shocks" method. This method deliver identical responses to the previous methods.

While the anticipated shocks method involves more operations than the "direct" method presented in the previous section, we will shock in the next section that the anticipated shocks method facilitates the computation of the likelihood. Note that recovering the time-varying and recursive formulation of the policy functions is time consuming to evaluate the likelihood function because it implies multiple large matrix operations for each set of

[7] In our case, this Jacobian is zero for all endogenous equations except for the rows associated with the Taylor rule.

[8] As before, we assume that the ELB for the interest rate is equal to 0. If the ELB is different than zero, the first m entries of C^i will be $-i^{ss} + \underline{i}$, where $\underline{i}$ is the lower bound in the interest rate.

parameters. However, given a set of parameters, recovering this approach is very efficient to compute stochastic simulations with ELB periods, because the matrix operations have to the executed only once. In contrast, for stochastic simulations, the anticipated shocks method can be time consuming.

F.4 Algorithm: Steps to Compute Likelihood when ELB is Binding

The challenging part of computing the likelihood consists of combining the fact that, during ELB period, there is sequence of past shocks and there is an expected duration of the ELB. In the previous two sections, where we presented the “direct” and “anticipated shocks” methods, we implicitly assumed that we were departing from the steady state of the economy. In other words, there was not a sequence of past shocks. In this section, we borrow the intuition from the “anticipated shocks” method to solve this problem. In particular, we can now make the anticipated sequence of shocks not only a function of the expected ELB duration and current shocks but also a function of the sequence of *past* aggregate shocks. To achieve this result, we: (1) get the MA representation of the economy based on (30). This give us the response of the endogenous variables to *past* shocks. (2) Based on the MA representation of the economy we create a vector describing the expected path of the interest rate. (3) We solve for the current set of anticipated shocks such that the interest rate is expected to be at zero for m periods. Note, again, that the difference with respect the IRFs presented in the previous two sections is that we are also conditioning on the *past* realizations of the shocks.

Computing the likelihood

1. Solving the model using SSJ, include L anticipated monetary policy shocks, where L should be as large as the maximum expected ELB duration. This results in an MA representation of the endogenous variables (dX) of the form:

$$dX = MA\epsilon^{t-1} + \tilde{\alpha}_0\tilde{\epsilon}_t + \beta_0\mu_t \tag{38}$$

where:

$$\epsilon_t = [\tilde{\epsilon}_t \quad \mu_t]' \tag{39}$$

$$\epsilon^t = [\epsilon_t \quad \epsilon_{t-1} \quad ... \quad \epsilon_{t-T}]' \tag{40}$$

$$\mu_t = [\varepsilon_{0,t} \quad \varepsilon_{1,t} \quad ... \quad \varepsilon_{L,t}]' \tag{41}$$

$$MA_j = \begin{bmatrix} \alpha_j & \alpha_{j+1} & ... & \alpha_T & \underbrace{0 \quad 0 \quad \cdots \quad 0}_{\text{j times}} \end{bmatrix} \quad \forall j = 0 \cdots T \tag{42}$$

$$\alpha_j = [\tilde{\alpha}_j \quad \beta_j]' \quad \forall j = 0 \cdots 1...T \tag{43}$$

μ_t is the vector of unanticipated and anticipated monetary policy shocks at time t, $\tilde{\epsilon}_t$ is the vector of all other shocks in the economy at time t, ϵ_t is the combination of the those two (i.e. all aggregate shocks at time t), and ϵ^t is the history of ϵ_t until time t. β_j is the response of the endogenous variables at time t to anticipated and unanticipated monetary policy shocks at time $t-j$. Similarly, $\tilde{\alpha}_j$ is the response of endogenous variables at time t to all other aggregate shocks at time $t-j$. MA_j is the endogenous variables response to all shocks up to period t.[9]

2. Extract the MA representation for the Federal Funds rate distinguishing between monetary policy shocks (anticipated and non-anticipated) and all other shocks in the economy. Hence, the federal funds rate at time t is given by:

$$\hat{i}_t = MA_1^i \epsilon^{t-1} + \tilde{\alpha}_0^i \tilde{\epsilon}_t + \beta_0^i \mu_t \tag{44}$$

where:

$$MA_j^i = \begin{bmatrix} \alpha_j^i & \alpha_{j+1}^i & ... & \alpha_T^i & \underbrace{0 \quad 0 \quad \cdots \quad 0}_{\text{j times}} \end{bmatrix} \quad \forall j = 0 \cdots T \tag{45}$$

$$\alpha_j^i = \begin{bmatrix} \tilde{\alpha}_j^i & \beta_j^i \end{bmatrix}' \quad \forall j = 0 \cdots 1...T \tag{46}$$

β_j^i is the response of the interest (i) at time t to anticipated and unanticipated monetary policy shocks at time $t-j$. Similarly, $\tilde{\alpha}_j^i$ is the response of the interest (i) at time t to all other aggregate shocks at time $t-j$. MA_j^i is the interest rate response to all shocks up to period t.

[9]Because the solution is truncated at horizon T in SSJ, the last j elements of MA_j are equal to zero.

3. Get the implied vector of monetary policy shocks for time t, as a function of past monetary policy shocks and as a function of the other aggregate shocks (current and past). At time t, the expected interest rate for $t+1$ is:

$$E_t\left[\hat{i}_{t+1}\right] = \begin{bmatrix}\alpha^i_2 & \alpha^i_3 & ... & \alpha^i_T & 0\end{bmatrix}\epsilon^{t-2} + \tilde{\alpha}^i_1\tilde{\epsilon}_t + \beta^i_1\mu_t \tag{47}$$

We can then group, the expected interest rate between t and $t+L$ as:

$$\begin{bmatrix}\hat{i}_t \\ E_t\left[\hat{i}_{t+1}\right] \\ \vdots \\ E_t\left[\hat{i}_{t+L}\right]\end{bmatrix} = \underbrace{\begin{bmatrix}\alpha^i_1 & \alpha^i_2 & \alpha^i_3 & \cdots & \alpha^i_{T-2} & \alpha^i_{T-1} & \alpha^i_T \\ \alpha^i_2 & \alpha^i_3 & \alpha^i_4 & \cdots & \alpha^i_{T-1} & \alpha^i_T & 0 \\ \alpha^i_3 & \alpha^i_4 & \alpha^i_5 & \cdots & \alpha^i_T & 0 & 0 \\ \vdots & \vdots & \vdots & & \vdots & \vdots & \vdots \\ \alpha^i_{L+1} & \alpha^i_{L+2} & \alpha^i_{L+3} & \cdots & 0 & 0 & 0\end{bmatrix}}_{\Omega_t}\epsilon^{t-1} + \underbrace{\begin{bmatrix}\tilde{\alpha}^i_0 \\ \tilde{\alpha}^i_1 \\ \tilde{\alpha}^i_2 \\ \vdots \\ \tilde{\alpha}^i_L\end{bmatrix}}_{\omega_t}\tilde{\epsilon}_t + \underbrace{\begin{bmatrix}\beta^i_0 \\ \beta^i_1 \\ \beta^i_2 \\ \vdots \\ \beta^i_L\end{bmatrix}}_{\lambda_t}\mu_t \tag{48}$$

Now, if the ELB is expected to bind for L periods:

$$\begin{bmatrix}\hat{i}_t \\ E_t\left[\hat{i}_{t+1}\right] \\ \vdots \\ E_t\left[\hat{i}_{t+L}\right]\end{bmatrix} = \underbrace{\begin{bmatrix}-i^{ss} \\ -i^{ss} \\ \vdots \\ -i^{ss}\end{bmatrix}}_{\zeta_t} \tag{49}$$

Hence, we can get the implied monetary policy shocks for time t as:

$$\mu_t = \lambda_t^{-1}\left[\zeta_t - \Omega_t\epsilon^{t-1} - \omega_t\tilde{\epsilon}_t\right] \tag{50}$$

Matrices, Z, ω, Ω, and λ are denoted with a subscript t as they are a funciton of the expected ELB duration at time t. We can re-group the elements in equation (50) such that:

$$\mu_t = \lambda_t^{-1}\zeta_t + \phi_t\tilde{\epsilon}^t + \psi_t\mu^{t-1} \tag{51}$$

where the first columns in ϕ_t correspond to the column in $-\lambda_t^{-1}\Omega_t$ associated with $\tilde{\epsilon}^{t-1}$, and the last columns correspond to $-\lambda_t^{-1}\omega_t$, and the columns of $-\lambda_t^{-1}\Omega_t$ associated with μ^{t-1} correspond to ψ_t.

4. Get the implied sequence of monetary policy shocks as a function of current and past aggregate shocks other than anticipated monetary policy shocks. Assuming that the ELB binds between periods t and $t+E$, stack

$$\begin{bmatrix} \mu_t \\ \mu_{t+1} \\ \vdots \\ \mu_{t+E} \end{bmatrix} = \underbrace{\begin{bmatrix} \lambda_t^{-1}\zeta_t \\ \lambda_{t+1}^{-1}\zeta_{t+1} \\ \vdots \\ \lambda_{t+E}^{-1}\zeta_{t+E} \end{bmatrix}}_{\mathcal{C}} + \underbrace{\begin{bmatrix} 0 & 0 & \cdots & 0 & \phi_t^0 & \phi_t^1 & \cdots & \phi_t^{T-1} & \phi_t^T \\ 0 & 0 & \cdots & \phi_{t+1}^0 & \phi_{t+1}^1 & \phi_{t+1}^2 & \cdots & \phi_{t+1}^T & 0 \\ \vdots \\ \phi_{t+E}^0 & \phi_{t+E}^1 & \cdots & 0 & 0 & 0 & \cdots & 0 & 0 \end{bmatrix}}_{\Phi} \begin{bmatrix} \tilde{\epsilon}_{t+E} \\ \tilde{\epsilon}_{t+E-1} \\ \vdots \\ \tilde{\epsilon}_{t+1} \\ \tilde{\epsilon}_t \\ \tilde{\epsilon}_{t-1} \\ \vdots \\ \tilde{\epsilon}_{t-T} \end{bmatrix}$$

$$+ \underbrace{\begin{bmatrix} 0 & 0 & \cdots & 0 & 0 & 0 \\ \psi_{t+1}^1 & 0 & \cdots & 0 & 0 & 0 \\ \psi_{t+2}^2 & \psi_{t+2}^1 & \cdots & 0 & 0 & 0 \\ \vdots \\ 0 & 0 & \cdots & \psi_{E+1}^2 & \psi_{E+1}^1 & 0 \end{bmatrix}}_{\Psi} \begin{bmatrix} \mu_t \\ \mu_{t+1} \\ \vdots \\ \mu_{t+E-1} \\ \mu_{t+E} \end{bmatrix} \tag{52}$$

where

$$\phi_t = \begin{bmatrix} \phi_t^0 & \phi_t^1 & \phi_t^2 & \cdots & \phi_t^T \end{bmatrix} \tag{53}$$
$$\psi_t = \begin{bmatrix} \psi_t^1 & \psi_t^2 & \cdots & \psi_t^T \end{bmatrix} \tag{54}$$

Hence, the implied sequence of monetary policy shocks is given by:

$$\mu^{t+E} = (I - \Psi)^{-1} \left[\mathcal{C} + \Phi \tilde{\epsilon}^{t+E}\right] \tag{55}$$

which implies that for each period t:

$$\mu_t = \mathcal{A}_t^{\mu} + \mathcal{B}_t^{\mu} \tilde{\epsilon}^t \tag{56}$$

Note that in our model $\mu_t = 0$ for all time periods before the first ELB episode. But this analysis is easy to extend to a model with active anticipated monetary policy shocks in non-ELB periods.

5. Plug the vector of monetary policy shocks (56) into the (38) to get the time t MA representation for the endogenous variables of the form:

$$dX_t = \mathcal{A}_t + \mathcal{B}_t \tilde{\epsilon}^t \tag{57}$$

6. Using (57) compute the likelihood of as in Auclert et al. (2021) (section 5.3). In particular, given the vector dX_t^{obs} of n_{obs} observables and a sample size of T_{obs}:

$$dX_t^{obs} = B \cdot dX_t + u_t \tag{58}$$

where u_t is a vector of measurement errors, and B is a selector matrix. We then stack the covariances of these observables in a large symmetric matrix V of size $n_{obs} \cdot T_{obs}$. Then, the likelihood function is given by:

$$\mathcal{L}\left(d\mathbf{X}^{obs} \mid \Theta\right) = (2\pi)^{-\frac{T_{obs}}{2}} \mid V \mid^{-\frac{1}{2}} exp\left\{\left[d\mathbf{X}^{obs} - \mathbf{A}\right]' V^{-1} \left[d\mathbf{X}^{obs} - \mathbf{A}\right]\right\} \tag{59}$$

where $d\mathbf{X}^{obs}$ is the stacked vector of observables, $\mathbf{A}$ is the staced vector of constants $(B \cdot \mathcal{A}_t)$, and Θ is the vector of model parameters.[10]

F.5 Overall Prior Distribution

In our Bayesian estimation, in addition to the standard informative priors over the model parameters, we specified an implicit prior over selected business cycle moments generated by our model, following Del Negro & Schorfheide (2008). Denote μ^{obs} as the vector of selected business cycle moments from the data, and $\mu(\Theta)$ as the model generated moments given the set of parameters Θ. $\mu(\Theta)$ relates to the data business cycle moments as follows:

$$\mu^{obs} = \mu(\theta) + \eta^{obs} \tag{60}$$

where η^{obs} is a vector of measurement errors that distributes normal with matrix of variance covariance equal to $\Sigma^{\eta^{obs}}$. We express (60) in terms of a conditional density (likelihood function) and use Bayes theorem in combination with a marginal density $\pi(\Theta)$ to generate

[10] Auclert et al. (2021) discuss how to quickly evaluate the determinant of V as well as the quadratic form.

a distribution that reflects our beliefs about the selected business cycle moments:

$$p\left(\Theta \mid \mu^{obs}\right)=\mathcal{L}\left(\mu(\Theta) \mid \mu^{obs}\right) \cdot \pi(\Theta) \tag{61}$$

where $\pi(\Theta)$ refers to the informative priors listed in Table **??**. Hence, our posterior distribution is given by:

$$\begin{aligned} p(\Theta \mid d\mathbf{X}^{obs}, \mu^{obs}) \propto &\quad \mathcal{L}(d\mathbf{X}^{obs} \mid \Theta) p\left(\Theta \mid \mu^{obs}\right) && (62)\\ \propto &\quad \mathcal{L}(d\mathbf{X}^{obs} \mid \Theta) \mathcal{L}\left(\mu(\Theta) \mid \mu^{obs}\right) \cdot \pi(\Theta) && (63)\end{aligned}$$

where $\mathcal{L}(d\mathbf{X}^{obs} \mid \Theta)$ is the likelihood function of the data defined in (59). The selected business cycle moments are: the variance of the observables, the covariance between the federal funds rate and GDP growth and CPI inflation, and the covariance between GDP growth and inflation. We assume that $\left(\Sigma^{\eta^{obs}}\right)^{-\frac{1}{2}}$ is a diagonal matrix with entries equal to ω times the standard deviation of the data moment, where ω is a scalar that we set to 0.25. Hence, for each moment m, we assume that the standard deviation of the measurement error associated with that moment equals σ^m:

$$\sigma^m = \omega \left[\sum_{t=1}^{T^{obs}} \frac{\left(x_t^m - \bar{x}^m\right)^2}{T^{obs}}\right]^{\frac{1}{2}} \tag{64}$$

where x_t^m refers to the data moment at time t, and $\bar{x}$ is the sample average of x. For example, $x_t^m = (dy_t^q - \bar{dy}^q)^2$ for the variance of GDP growth, and $x_t^m = (dy_t^q - \bar{dy}^q)(\pi_t^q - \bar{\pi}^q)$ for the covariance between GDP growth and quarterly inflation.

To compute the model generated moments $\mu(\Theta)$, we make use of the time-varying MA representation presented in the Online Appendix. Given the time-varying MA representation of the observables:

$$dX_t^{obs} = \mathcal{A}_t^{obs} + \mathcal{B}_t^{obs} \epsilon^t \tag{65}$$

we stack them to get:

$$d\mathbf{X}^{obs} = \mathbf{A} + \mathbf{B}\epsilon \tag{66}$$

Given that the steady state of our model matches the sample average of our observables, we are interested in computing:

$$E\left\{\left[d\mathbf{X}^{obs}\right]'\left[d\mathbf{X}^{obs}\right]\right\} = \Sigma(\Theta) \tag{67}$$

$$= \mathbf{A}'\mathbf{A} + \mathbf{B}'\Sigma^{\epsilon}\mathbf{B} \tag{68}$$

where Σ^{ϵ} equals to the matrix of variance covariance of the model shocks.[11] Note that $\Sigma(\Theta)$ is a big square and symmetric matrix of size $n^{obs} \cdot T^{obs}$.

Define matrix $\sigma(\Theta)_t$ as the square and symmetric matrix of size $nobs$ form by the rows and columns of $\Sigma(\Theta)$ starting in row and column $(t-1)nobs+1$ and ending in row and column $t \cdot nobs$. Matrix $\sigma(\Theta)_t$ is the expected sample covariance of the observables in period t or, in other words, the expected quadratic deviation of the observables from their sample average in period t. Then, the expected matrix of variance covariance generated by the model is given by:

$$\bar{\sigma}(\Theta) = \sum_{t=1}^{T^{obs}} \frac{\sigma(\Theta)_t}{T^{obs}} \tag{69}$$

Hence, our selected model generated moments used in our prior $(\mu(\Theta))$ correspond to the associated elements of matrix $\bar{\sigma}(\Theta)$.

Note that our model generated moments condition on the ELB episodes. In other words, our model generated moments are conditional on the ELB periods and their expected durations.

Other parameter restrictions: In addition to these priors, to avoid exploring unreasonable areas of the parameter space, we get the filtered series for each draw of parameters and discard those draws that imply a value of the menu cost below zero or log-deviations for θ_p and θ_a greater than three in absolute value.

F.6 Posterior Sampler

We use a standard Metropolis Hasting Algorithm with similar specification as in Kulish et al. (2017). The algorithm is the following:

Algorithm: At the beginning of each iteration j, given a set of parameters Θ_{j-1} :

1. Randomly select how many parameters to update from a uniform distribution between $\lceil 0.2 n_{\Theta} \rceil$ and n_{Θ}, where n_{Θ} is the number of parameters in Θ.

[11] Note that the last term in (68) equals to matrix V of Appendix F.

2. Randomly select which parameters to update.

3. Construct a proposed set of parameters Θ_j^p. To do this, we use a multivariate Student t distribution with $12+p$ degrees of freedom centered at Θ_{j-1} and with a matrix of variance covariance given by Σ^{Θ^p}, which will be specified below. p is te number of parameters being updated.

4. Compute the acceptance ratio (AR):

$$AR_j^{\Theta} = \frac{p(\Theta_j^p \mid d\mathbf{X}^{obs}, \mu^{obs})}{p(\Theta_{j-1} \mid d\mathbf{X}^{obs}, \mu^{obs})} \tag{70}$$

   set $AR_j = 0$ if the proposal includes inadmissible values.

5. Accept the proposal with probability $min\{AR_j, 1\}$ and set $\Theta_j = \Theta_j^p$. Otherwise, set $\Theta_j = \Theta_{j-1}$.

Matrix of variance covariance Σ^{Θ}: To compute the matrix of variance covariance Σ^{Θ}, we run a chain of 150,000 draws using a matrix of variance covariance equal to the diagonal matrix of the Hessian at the optimization mode scaled by κ^{mcmc}, which we set to 0.0575. Then, we drop the first 50,000 draws and compute the matrix of variance covariance resulting from this chain, which we denote by Σ^{Θ}. Then, for the proposal density, we can decomposte matrix Σ^{Θ} as follows:

$$\Sigma^{\Theta} = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{bmatrix} \tag{71}$$

where the first block refers to the fixed parameters at iteration j. Hence, the specific matrix of variance covariance used in iteration j equals to:

$$\Sigma^{\Theta^p} = \omega^{mcmc}\Sigma_{22|1} = \Sigma_{22} - \Sigma_{21}\Sigma_{11}^{-1}\Sigma_{12} \tag{72}$$

where ω^{mcmc} is a scaling parameter that we set to 0.2 to target an acceptance rate of 30%.

Sample chain: We used a single chain of 500,000 draws with an average acceptance rate of 30%. We drop the first 100,000 draws to compute statistic on the the posterior distribution.

F.7 Prior and Posterior distributions

Figure 4 plots the prior and posterior distributions and Figure 5 plots the trace plots. Solid black lines represent the posterior mode.

Figure 4: Prior and posterior distribution

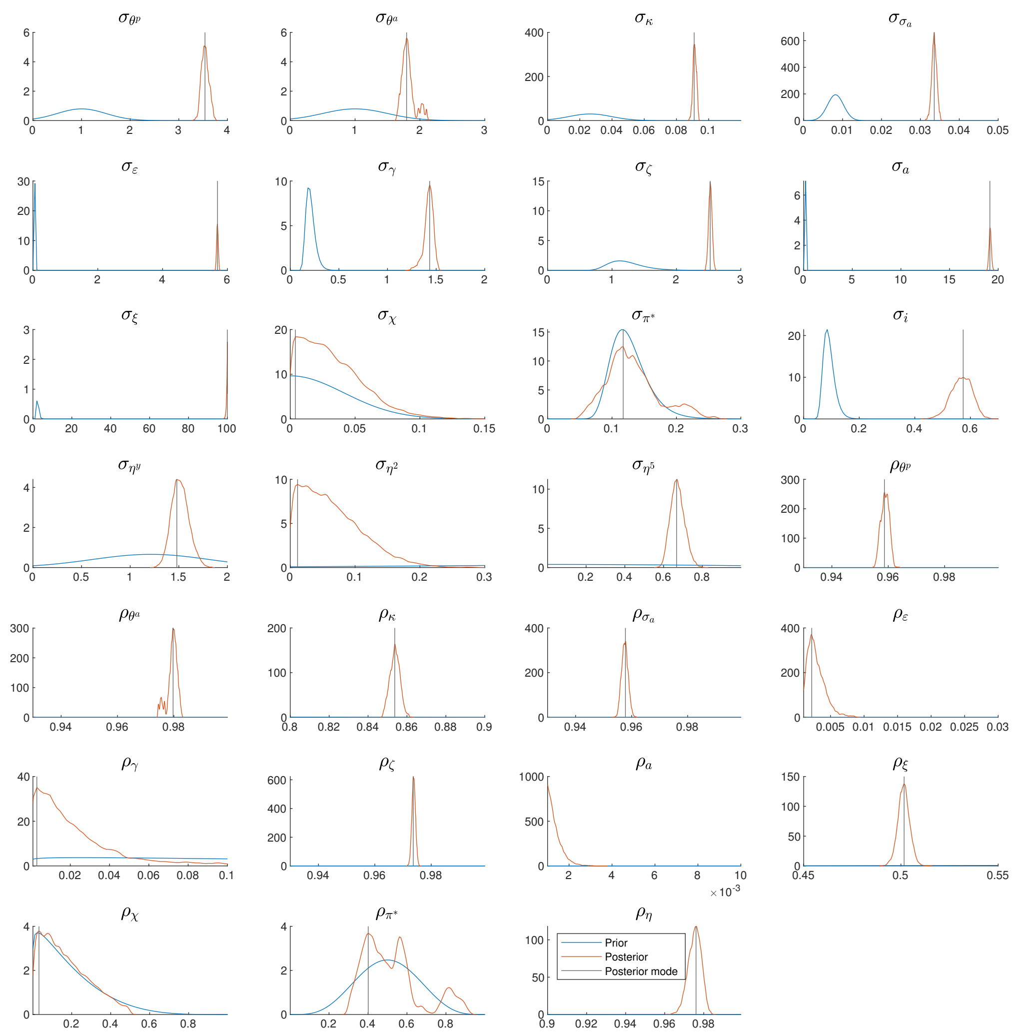

Figure 5: Trace plots

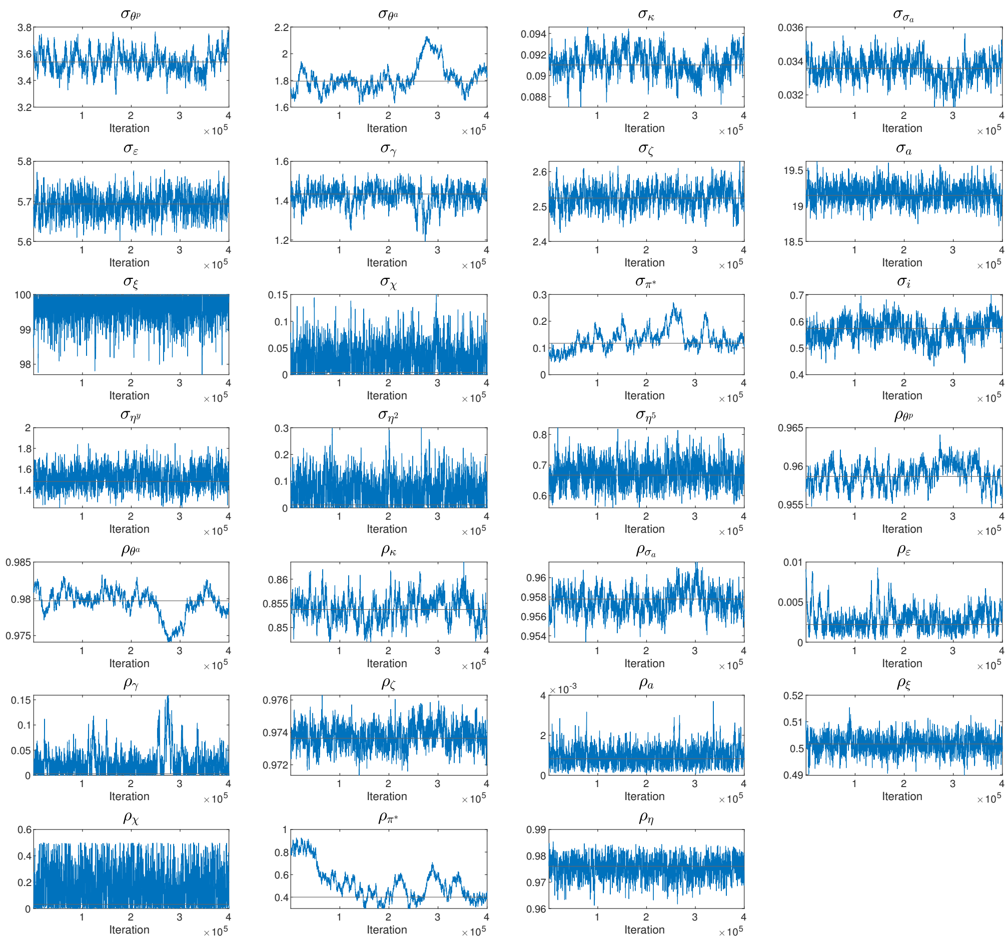

Note: Horizontal line represents posterior mode.

G Detrended Model Equations

Aggregate price index:

$$P_t = 1 = \left[\int p_{jt}^{1-\varepsilon} \Omega_{jt} dj \right]^{\frac{1}{1-\varepsilon}} \tag{73}$$

Aggregate resource constraint:

$$\Delta_t \tilde{Y}_t = \tilde{C}_t + \tilde{G}_t + \tilde{F}_t \tag{74}$$

where

$$\tilde{F}_t = \int \left[\tilde{\kappa} + \tilde{\mathcal{C}}^a_{jt} + \lambda_{jt} \tilde{\mathcal{C}}^p_{jt} \right] \tilde{\Omega}_{jt} dj \tag{75}$$

Euler equation:

$$\zeta_t \, \chi_t + E_t \left[\frac{\beta}{\gamma_{t+1}} \, \zeta_{t+1} \, \tilde{\lambda}_{t+1} \left(\frac{i_t}{\pi_{t+1}} \right) \right] = \zeta_t \, \tilde{\lambda}_t \tag{76}$$

where

$$\tilde{\lambda}_t = z_t \lambda_t = \left(\frac{1}{\tilde{C}_t - h \frac{\tilde{C}_{t-1}}{\gamma_t}} \right) - E_t \left[\left(\frac{\beta \zeta_{t+1}}{\zeta_t} \right) \left(\frac{h}{\gamma_{t+1} \tilde{C}_{t+1} - h \tilde{C}_t} \right) \right] \tag{77}$$

Marginal cost:

$$e^{a_t} \tilde{p}^x_t = \delta^w e^{a_t} \frac{\xi_t \left(\frac{\tilde{Y}_t}{e^{a_t}} \Delta_t \right)^{\frac{1}{\nu}}}{\tilde{\lambda}_t} + (1 - \delta^w) \tilde{p}^{x*} \tag{78}$$

where we made use of:

$$p^x_t = \frac{w_t}{e^{a_t}} \tag{79}$$

$$w^*_t = \frac{\xi_t L_t^{\frac{1}{\nu}}}{\lambda_t} \tag{80}$$

$$L_t = \frac{\tilde{Y}_t}{e^{a_t}} \Delta_t \tag{81}$$

and where

$$\Delta_t = \int_j e^{-a_{jt}} p_{jt}^{-\varepsilon} \Omega_{jt} dj \tag{82}$$

Taylor rule:

$$i_t = i_t^{target} + \epsilon_{rt}, \tag{83}$$

$$(i_t^{target})^p = \rho_i (i_{t-1})^p + (1-\rho_i)\left[(i_{ss})^p + \phi_\pi \left(\Pi_{j=0}^{p-1} \pi_{t-j} - (\pi_t^*)^p\right) + \phi_y (dy_t - dy_{ss})\right], \tag{84}$$

H Data Series

Figure 6 plots the raw and smoothed time series of the different pricing moments we use.

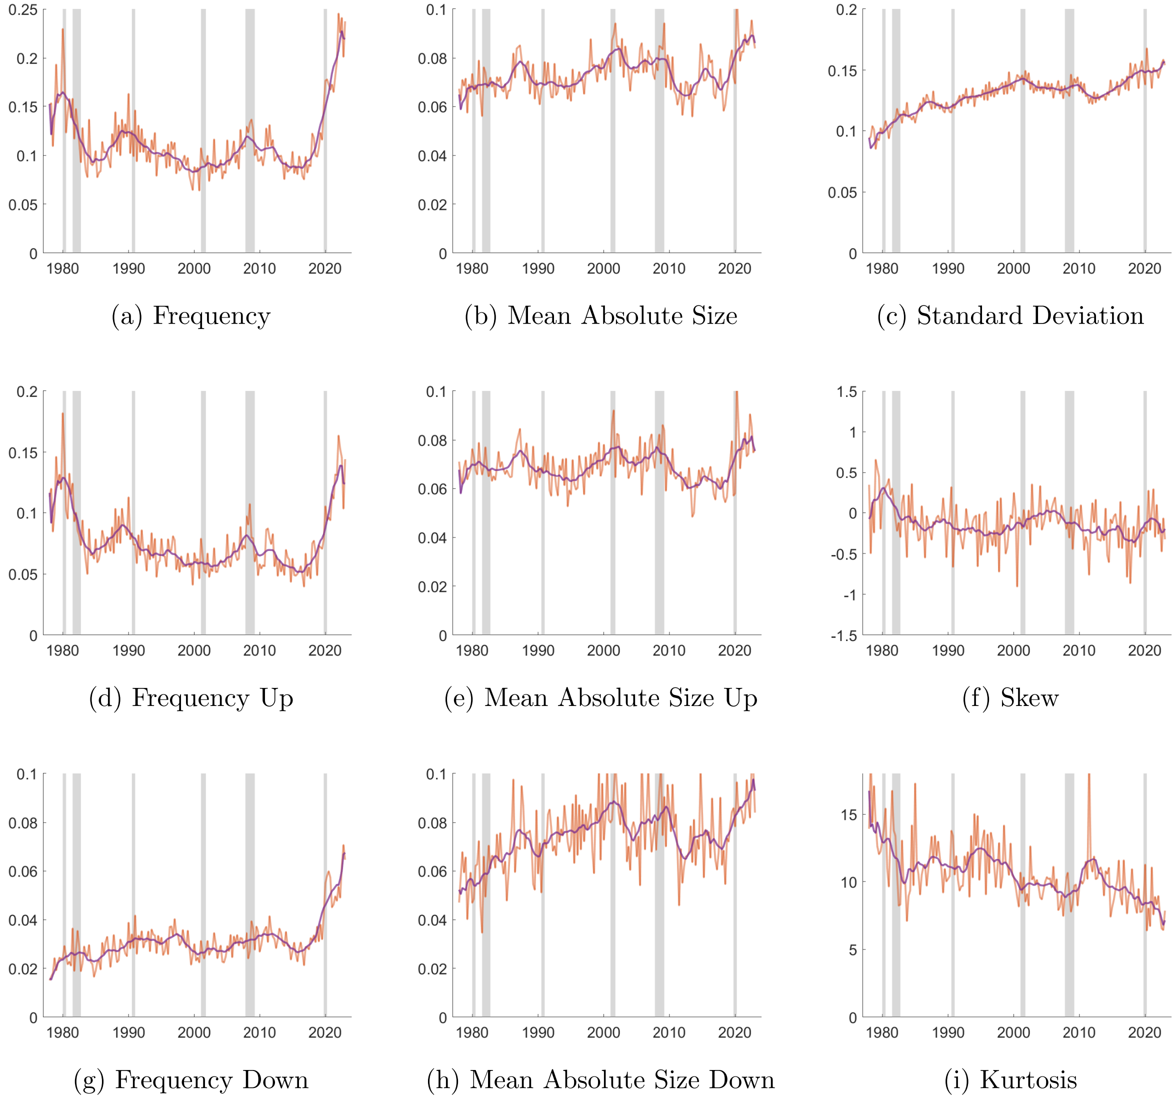

(a) Frequency (b) Mean Absolute Size (c) Standard Deviation

(d) Frequency Up (e) Mean Absolute Size Up (f) Skew

(g) Frequency Down (h) Mean Absolute Size Down (i) Kurtosis

Figure 6: Data Pricing Moments

Note: These panels plot the raw and MA-smoothed pricing series from 1978 to 2023.Q1. Shaded areas represent NBER recession dates.

References

Auclert, Adrien, Bence Bardóczy, Matthew Rognlie & Ludwig Straub (2021), "Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models," *Econometrica* 89(5): 2375–

2408.

Costain, James & Anton Nakov (2019), "Logit price dynamics," *Journal of Money, Credit and Banking* 51(1): 43–78.

Del Negro, Marco & Frank Schorfheide (2008), "Forming priors for DSGE models (and how it affects the assessment of nominal rigidities)," *Journal of Monetary Economics* 55(7): 1191–1208, ISSN 0304-3932.

Khaw, Mel Win, Luminita Stevens & Michael Woodford (2019), "Empirical models of inattentive choice," Working Paper.

Kulish, Mariano, James Morley & Tim Robinson (2017), "Estimating DSGE models with zero interest rate policy," *Journal of Monetary Economics* 88: 35–49.

Luo, Shaowen & Daniel Villar (2021), "The skewness of the price change distribution: A new touchstone for sticky price models," *Journal of Money, Credit and Banking* 53(1): 41–72.

Reiter, Michael (2009), "Solving heterogeneous-agent models by projection and perturbation," *Journal of Economic Dynamics and Control* 33(3): 649–665, ISSN 0165-1889.

Woodford, Michael (2012), "Inattentive valuation and reference-dependent choice," Unpublished Manuscript.