

Price Rigidities in U.S. Business Cycles*

LATEST VERSION HERE

Camilo Morales-Jiménez
Federal Reserve Board

Luminita Stevens
University of Maryland

April 17, 2024

We provide a structurally estimated time series for the degree of nominal price rigidities in the United States between 1978 and 2023. To model the price rigidity, we allow for stochastic state dependence in both the timing of price changes and the choice of what price to charge. We give a cost-based micro-foundation to this stochasticity, modeling firms that face information costs and menu costs when making decisions regarding their prices. Estimating the model on time series of moments from the distribution of price changes over time—in addition to time series of real economic activity and inflation—we find considerable monetary non-neutrality with medium-cycle volatility. Underlying our estimated series are a number of results that shed new light on the sources and dynamics of price rigidity. In particular, the model attributes most of the rigidities in price setting not to infrequent, but rather to inaccurate price adjustment. The timing of adjustments has been accurate, especially for price decreases. Menu costs have been small but growing. Information frictions are large and volatile. This volatility has implications for the trade-offs involved in 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 colleagues at numerous institutions and conferences for helpful conversations. Contact: camilo.moralesjimenez@frb.gov, stevens7@umd.edu.

1 Introduction

How severe are price rigidities in the U.S. economy? Have prices become more responsive to shocks over time? Does price rigidity vary over the business cycle, complicating the stabilization efforts of the Federal Reserve? These questions are at the core of monetary economics. Rigidities in prices change how the economy adjusts in response to any shock, be it a supply shock or a demand shock. They also determine to what extent monetary policy can stabilize fluctuations in inflation and real economic activity, and they determine the welfare costs of inflation. Consider the question of a soft or a hard landing for the U.S. economy following the inflation burst that began in 2021. The answer to that question 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 rigidity, uncertainty and disagreement regarding the severity of price rigidities persist in the literature, reflecting both the difficulty of extracting model-free empirical estimates of price rigidity from the data, and the lack of clarity regarding what frictions are most relevant for this rigidity.¹

This paper provides a structurally estimated time series for the degree of nominal price rigidity (NPR) in the United States between January 1978 and March 2023. The measure of NPR in each period is given by what the model predicts the cumulative response of consumption would have been 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 time series.

The estimation of the pricing frictions themselves uses a model of the aggregate economy that is estimated on time series of the moments from the distribution of price changes, in addition to time series on real economic activity and inflation. The pricing moments are based on the micro data underlying the U.S. Consumer Price Index (CPI).² Incorporating the dynamics of the distribution of price changes in the estimation itself is a novel use of the pricing micro data, much like the

¹We discuss these challenges further in the literature review.

²We thank Daniel Villar for sharing the time series of key pricing moments with us. These series are constructed by [Nakamura, Steinsson, Sun & Villar \(2018\)](#) and extended to 2023 by [Montag & Villar \(2023\)](#).

literature on heterogeneous agent models has increasingly used household income and wealth distributional data to inform models of the aggregate economy (for example, [Bayer, Born & Luetticke, 2020](#); [Bilbiie, Primiceri & Tambalotti, 2023](#), in the growing HANK literature). The use of time variation in the distribution of price changes enables us to characterize if and how price rigidities change over time.

We model price frictions flexibly, allowing for the possibility of errors and rigidity in both the timing of price changes and the repricing itself (the choice of what price to set when adjusting). We give a cost-based micro-foundation to this inaccuracy, modeling firms that face both information costs (modeled as in [Sims, 2003](#); [Woodford, 2009](#)), and menu costs, and thus nesting the two main ways of endogenizing nominal price rigidities.

In the model, firms choose how much attention to pay to market conditions given the cost of obtaining more information and the cost of revising their current price. Since changing prices is costly, firms first decide if they want to change their price. They do so based on an imprecise awareness of the state of the economy, as in [Woodford \(2009\)](#). How accurate their timing decisions are is a choice that firms make state by state, weighing the marginal benefit vs/ the marginal cost in each state. Second, if they do decide to change their price, firms then need to decide what price to set. Unlike in prior pricing models, this repricing decision is also based on an imprecise awareness of market conditions. The degree to which firms tolerate errors in pricing interacts with their tolerance for errors in the timing of price changes, and we provide a discussion of this interaction and show under what conditions it can rationalize Calvo-like behavior.

Quantitatively, the model can span pricing behavior from fully flexible to fully random in terms of both timing and price levels, and we let the estimation on U.S. data pin down the degree of inaccuracy along each margin, which in turn pins down monetary non-neutrality.

We estimate a sizable degree of non-neutrality: on average, a 25 bp shock to the federal funds rate results in a cumulative change in consumption equal to 0.12% of annual steady-state consumption. This represents approximately 80% of the response the Calvo model would predict when calibrated to the same frequency and size of price changes. Between 1978 and 2023, the degree of non-neutrality

exhibits no clear trend. This is surprising since we might imagine that technology has made both information gathering and repricing less costly. We interpret this as suggestive of the complexity that goes hand in hand with technological progress and data abundance: more is not always easier.

Despite no clear trend, we find substantial volatility in the degree of NPR over time. We estimate above-average rigidity in the mid-1980s and mid-1990s. Rigidity increased slightly during the Great Recession and continued to rise after the recession, remaining elevated above average until 2016. In 2016, it began a steady decline that continued until the end of our sample in 2023. We do not find consistent patterns across other recessions, which casts some doubt on the hypothesis of increased price flexibility during recessions. The increased price flexibility that we estimate starting around 2016 is particularly interesting. In hindsight, it suggests that we might have expected any inflationary shocks, should they occur, to be met with a sharper inflation response, rather than a sharper output response post-2016 versus pre-2016. This may explain why inflation surged so rapidly in 2021, and it points to the value of using distributional pricing data in real-time.

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. At the heart of these implications is the interaction between the timing decision and the repricing choice.

First, we estimate substantial inaccuracy in the repricing decision. This is at odds with standard models that assume perfect repricing conditional on adjustment. Our finding of a relatively high degree of errors in pricing breaks the connection between *adjustment* and *flexibility*: Even if prices are not sticky, in the sense that they are changing over time, they nevertheless only partially respond to economic conditions. Underscoring this dichotomy, we estimate only a weak correlation between the frequency of price changes and the estimated NPR.

Inaccuracy in pricing is consistent with a large body of evidence that economic choices are based on dispersed beliefs and are imprecisely related to optima in many contexts. Many studies have documented dispersion in actions and forecasts *conditional on adjustment*, both in survey data and in incentivized con-

trolled laboratory experiments, including [Mankiw, Reis & Wolfers \(2003\)](#); [Carroll \(2003\)](#); [Coibion & Gorodnichenko \(2012\)](#); [Magnani, Gorry & Oprea \(2016\)](#); [Cav-allo, Cruces & Perez-Truglia \(2017\)](#); [Khaw, Stevens & Woodford \(2017\)](#); [Angeletos, Huo & Sastry \(2021\)](#), among many others.

Allowing for the possibility of mistakes in repricing also implies that Calvo is also no longer an upper bound on the degree of nominal rigidity. Varying the severity of information frictions regarding the timing of price adjustment spans the degree of state dependence in price setting, with the menu cost model at one end (when the information friction approaches zero) and [Calvo \(1983\)](#) at the other end (when the information friction is strong enough that the firm acquires no information to decide when to adjust its prices), as shown by [Woodford \(2009\)](#). But this result applies to models featuring perfect repricing. Adding errors in the repricing itself adds another layer of nominal rigidity. As a result, the model can feature larger non-neutrality than a Calvo model parameterized to have the same frequency of price adjustment.

Second, even though prices change infrequently (with an 11% monthly frequency on average), most of the inaction reflects uncertainty about the right price to set, rather than an unwillingness to pay the fixed menu cost to readjust. We estimate a small menu cost that accounts for only a small fraction of both adjustment costs and the total degree of price rigidity. 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 new micro-foundation for inaction that is quantitatively important. In contrast, without errors in repricing, the model would require larger menu costs or an exogenously low probability of adjustment to generate sufficiently infrequent price changes.

Third, the timing of price increases is fairly accurate, as indicated by the small estimated cost of accurately determining when price changes are warranted. This yields a strongly state-dependent probability of adjustment. It seems that firms are able to determine quite accurately when their prices have become obsolete, but they have more difficulty determining the right price level to set. This finding is consistent with recent work on selection in U.S. price data ([Carvalho & Kryvtsov, 2021](#); [Karadi, Schoenle & Wursten, 2022](#)). At the same time, it is at odds with the conventional wisdom, which posits that in order to get meaningful non-neutrality,

as we see in the data, one needs weak state dependence in the timing of price changes. For example, with perfect repricing, menu cost models need some auxiliary features to weaken the selection in terms of who adjusts in response to a shock. But with mistakes in repricing, weak state-dependence in the timing of price changes is no longer necessary for non-neutrality. In the terminology of [Caballero & Engel \(2007\)](#), repricing errors mute the selection on the intensive margin, reducing the need to rely on weak selection on the extensive margin.

Fourth, in the time series estimation, we find that the dynamics of the CIR are significantly impacted by large and variable repricing costs that were particularly elevated in the 1990s and after the Great Recession, but have declined significantly since 2016. By 2023, the cost of setting prices accurately reached its lowest level over the entire sample period. The cost of accurately timing price changes was low for most of the sample period, but that changed with the Great Recession: inaccuracy in timing rose starting in the Great Recession, reaching Calvo-like levels by 2012. It too has been declining since 2016 and in 2022 it returned to its steady state value. Finally, the menu cost has been small, but it too exhibits some variability: it rose in the 1980s and 1990s and has since fluctuated at a medium frequency, but at lower magnitudes of the two information costs.

In terms of magnitudes, we estimate that firms annually spend approximately 0.03% of revenues on the menu cost, which is an order of magnitude smaller than estimates from models with perfect repricing. Information costs are both larger and more volatile over time than menu costs. The estimation implies annual information expenditures of about 2.5% of sales, of which one third is spent on determining whether a price change is warranted and two thirds are spent determining the right price to charge.

Finally, we also want to highlight our estimation method, which contributes 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, and the first application of the sequence-space Jacobian (SSJ) method of [Auclert, Bardóczy, Rognlie & Straub \(2021\)](#) to a model with heterogeneous information. Moreover, since our estimation sample includes two periods in which the effective lower bound (ELB) was binding on the federal funds rate, we also show how to handle occasionally binding constraints with SSJ, by adapting the methods

proposed by [Guerrieri & Iacoviello \(2015\)](#) and [Kulish, Morley & Robinson \(2017\)](#).

2 Additional Related Literature

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. \(2017\)](#). [Costain & Nakov \(2019\)](#) model price-setting firms 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. [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.

More broadly, our results build on several strands of the literature. First, 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 frequency relative to kurtosis pins down non-neutrality in a wide class of menu cost models, [Berger & Vavra \(2018\)](#) argue for the additional 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 \(2004\)](#) and [Nakamura & Steinsson \(2008\)](#), a wave of menu cost models (e.g., [Golosov & Lucas Jr \(2007\)](#); [Nakamura & Steinsson \(2010\)](#); [Midrigan \(2011\)](#); [Alvarez & Lippi \(2014\)](#); [Vavra \(2013\)](#)) have studied the contribution to non-neutrality of different moments of the price change distribution. There is, nonetheless, an ongoing debate concerning the informativeness of various pricing moments for the degree of non-neutrality as well as the degree to which the degree of non-neutrality varies over time, and whether or not it is procyclical. Our results underscore that steady-state pricing moments may not be sufficient for pinning down non-neutrality, at least in models in which repricing itself is frictional. They also emphasize that inferences regarding the degree of aggregate price flexibility remain, at least for now, quite model-dependent.

Second, our framework nests models that generate non-neutrality via **infre-**

quent price adjustment with those 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 of infrequent adjustment such as [Dotsey, King & Wolman \(1999\)](#), [Caballero & Engel \(2007\)](#), and [Woodford \(2009\)](#), in which the probability of adjustment varies smoothly with the value of adjusting. In the second category, our model belongs to the class of models with imprecise price-setting ([Woodford, 2003](#)), in particular work that operationalizes the imprecision using tools from information theory ([Maćkowiak & Wiederholt, 2009](#); [Matějka, 2015](#); [Turen, 2023](#); [Afrouzi, 2020](#); [Afrouzi & Yang, 2021](#)). The first group of models, which make a firm’s nominal price sticky over time, assume perfect repricing: Once a firm has the opportunity to reprice, the newly chosen price is a deterministic, full-information optimal choice. The models in the second group remove the impediments to changing prices every period, but instead relax the assumption of perfect repricing. We nest these two cases and allow the estimation to speak to their relative importance in generating monetary non-neutrality.

Third, 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. [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, and also find support for state dependence in the acquisition of information for price-setters.

Fourth, we contribute to the strand of literature that has studied how NPR varies with inflation. Empirical work has shown that once inflation exceeds high single digits, price rigidity starts to decline with inflation, rapidly reaching near-flexibility ([Alvarez, Beraja, Gonzalez-Rozada & Neumeyer, 2019](#); [Gagnon, 2009](#)) and our model generates this state-dependence through the endogenous adjustment

of information acquisition.

Fifth, our paper also relates to work that seeks to estimate the severity of information frictions over time more generally, such as [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. Our results also add to the literature on parameter stability in DSGE models, e.g., [Fernández-Villaverde, Rubio-Ramírez, Cogley & Schorfheide \(2007\)](#).

3 Model

This section presents an economy with errors in both the timing of price adjustment and the repricing decision. Information costs are the source of these errors and, together with menu costs, generate noisy, 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. In addition to the retailers, the economy consists of competitive goods producers, a representative household, and fiscal and monetary authorities.

3.1 Final Goods

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

$$Y_t = C_t + G_t + F_t. \tag{1}$$

This good is an aggregator of differentiated varieties j ,

$$Y_t = \left[\int y_{jt}^{\frac{\varepsilon-1}{\varepsilon}} dj \right]^{\frac{\varepsilon}{\varepsilon-1}}. \tag{2}$$

where $\varepsilon > 1$ is the elasticity of substitution. Demand for each variety is

$$y_{jt} = p_{jt}^{-\varepsilon} Y_t, \tag{3}$$

with $p_{jt} = P_{jt}/P_t$ denoting the good's relative price, and

$$\left(\int p_{jt}^{1-\varepsilon} dj \right)^{\frac{1}{1-\varepsilon}} = 1. \quad (4)$$

3.2 Monopolistically Competitive Retailers

The differentiated varieties are sold by a continuum of retailers who are monopolistically competitive price-setters in their product market and price-takers in the market for their production input. They optimize subject to both information costs and menu costs.

Operating Profits The retailers' production function is

$$y_{jt} = e^{a_{jt}} x_{jt}, \quad (5)$$

where a_{jt} is an AR(1) process for idiosyncratic productivity and x_{jt} is the homogeneous input. Given its price, the retailer purchases whatever quantity of the intermediate good is needed to satisfy demand at that price.

A retailer's real operating profit per period is

$$\pi_{jt}^r = p_{jt} y_{jt} - p_t^x x_{jt} = \left[p_{jt}^{1-\varepsilon} - p_{jt}^{-\varepsilon} \left(\frac{p_t^x}{e^{a_{jt}}} \right) \right] Y_t, \quad (6)$$

where $p_t^x = P_t^x/P_t$ is the intermediate input's relative price.

Information Costs Firms are rationally inattentive (RI) to market conditions (Sims, 2003). They are rational, in that they optimize based on a complete understanding of the structure of their environment (payoff functions, shock processes, markets), but they must expend resources to learn the realizations of stochastic variables in real time.

Information acquisition is modeled as a choice that can be quantified and optimized using tools from information theory (Shannon, 1948, 1959). However, we assume that rather than being endowed with a fixed information capacity (like the communication channels that are modeled in the information theory literature),

firms can choose how much information to obtain, subject to a variable cost, as in [Woodford \(2009\)](#).

In our context, given the fixed menu cost of price adjustment, in each period firms must decide whether or not to update their price, and if so, what price to set. As is common in the RI literature, we assume that the costs of making the adjustment and pricing decisions contingent on the realized states are linear in the information acquired in order to make each decision,

$$\mathcal{C}_{jt}^a = \theta^a \mathcal{I}_{jt}^a, \quad (7)$$

$$\mathcal{C}_{jt}^p = \theta^p \mathcal{I}_{jt}^p, \quad (8)$$

where θ^a is the unit cost of making a more informed decision about whether or not to change prices, θ^p is the unit cost of making a more informed price choice when adjusting prices, and \mathcal{I}_{jt}^a and \mathcal{I}_{jt}^p measure how much information is acquired for each decision (which we discuss further below). We allow for potentially different unit costs, since they may reflect different managerial marginal costs of attention. If θ^a and θ^p are zero, the firm's problem collapses to a full information menu cost model.

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

$$\max \left\{ E_{j0} \sum_{t=0}^{\infty} M_{0,t} \left[\pi_{jt}^r - \mathcal{C}_{jt}^a - \delta_{jt} (\kappa + \mathcal{C}_{jt}^p) \right] \right\}, \quad (9)$$

subject to (6),(7),(8), 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 change its price in the period, it continues with its existing nominal price (there is no price indexation).

Acquiring Information A firm's choice of how much information to obtain amounts to choosing how much its decisions condition on each realized state, relative to the best decisions the firm could make based on beliefs it has for free.

For each decision, the amount of information acquired is measured by Shannon's

mutual information. For the adjustment decision, this is given by

$$\mathcal{I}_{jt}^a = E_t \left\{ \mathcal{D} \left(\Lambda_{jt} \parallel \bar{\Lambda}_{jt} \right) \right\}, \quad (10)$$

$$\mathcal{D}(\Lambda \parallel \bar{\Lambda}) = \Lambda \ln \left(\frac{\Lambda}{\bar{\Lambda}} \right) + (1 - \Lambda) \ln \left(\frac{1 - \Lambda}{1 - \bar{\Lambda}} \right), \quad (11)$$

where Λ_{jt} denotes the probability that the firm adjusts its price in period t , after obtaining information about the realized state, $\bar{\Lambda}_{jt}$ is the *reference probability* of adjustment, based on the firm's beliefs *before* obtaining current information, \mathcal{D} is the Kullback-Leibler 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 .³ Hence, the contribution to the firm's cost of conditioning the adjustment decision on a period's realized state is proportional to the probability of the firm finding itself in that state times the divergence of Λ from $\bar{\Lambda}$ in that state. The trade-off facing the firm captures 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. It also reflects the fact that all else equal, paying attention to more frequent states will cost more.

Analogously, for the pricing decision, the amount of information obtained in order to decide *what price* to set is

$$\mathcal{I}_{jt}^p = E_t \left\{ \mathcal{D} \left(f_{jt}(p) \parallel \bar{f}_{jt}(p) \right) \right\}, \quad (12)$$

$$\mathcal{D} \left(f(p) \parallel \bar{f}(p) \right) = \int f(p) \ln \left(\frac{f(p)}{\bar{f}(p)} \right) dp, \quad (13)$$

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 $\bar{f}_{jt}(p)$ is the reference probability of setting the price equal to 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 Kullback-Leibler divergence of the choice distribution from the

³The 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. Shannon's mutual information between two random variables x and y is the KL divergence of the joint distribution from the product of the marginal distributions.

reference distribution, and expectations integrate over the joint distribution of idiosyncratic 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 equal to the probability of the firm finding itself in that state times the divergence of f_{jt} in that state from \bar{f}_{jt} .

Reference Distributions We assume the firms’ reference distributions are the equilibrium distributions in steady state. Let $\Lambda_{ss}(p, a)$ denote the steady-state probability of adjustment of a firm with price p and idiosyncratic productivity a . The reference adjustment probability $\bar{\Lambda}$ is the equilibrium frequency of adjustment in the steady state,

$$\bar{\Lambda} = \int \Lambda_{ss}(p, a) \tilde{\Omega}_{ss}(p, a) da dp, \quad (14)$$

which integrates the adjustment probability over the steady-state distribution of firms before price review decisions have been made, but after the idiosyncratic shocks have been realized.

Similarly, the reference probability of charging each p in the set of possible prices is the steady state distribution of prices, but after adjustments have been made. 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

$$\bar{f}(p) = \int f_{ss}(p|a) \Omega_{ss}(p, a) da, \quad (15)$$

where $\tilde{\Omega}_{ss}$ is the steady state joint distribution of idiosyncratic productivities and prices prior to the adjustment decision, and Ω_{ss} is the joint distribution of productivities and post-adjustment prices, and are defined further below.

Discussion of Reference Distributions The assumption that firms use equilibrium 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.⁴

⁴See [Woodford \(2012\)](#) and [Khaw, Stevens & Woodford \(2019\)](#) for alternative deviations of the default distributions from the RI optima.

By constraining the reference distributions to be the cross-sectional equilibrium distributions, we are using a slightly inefficient information structure relative to the pure RI solution. How does it compare to benchmark alternatives?

One benchmark is the control cost (CC) model of stochastic choice. These models assume a reference distribution for each period that is uniform around the optimal action in that period. Choosing a more concentrated distribution entails a cost proportional to the divergence of the chosen distribution from the uniform. Such models allow for unbiased errors in the implementation of actions when the optimal action is known in each state, generating volatile, noisy actions. For example, [Costain & Nakov \(2019\)](#) apply control costs to price-setting in a general equilibrium monetary model. However, as the name suggests, they are models of costly control, rather than costly information.

On the other hand, RI models of costly information derive an endogenous reference distribution that is optimal, given the decision-problem at hand. Notice that 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.

Here we are 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.

Recursive Formulation and Optimal Choice Distributions We now solve for each element of the firm's optimal policy. For a given reference, the choice distribution maximizes total firm value net of the information cost. First, consider

the choice of the adjustment probability $\Lambda_t(\tilde{p}, a)$ for a firm that begins a period with real price \tilde{p} and idiosyncratic productivity a . This choice solves:

$$V_t^*(\tilde{p}, a) = \max_{\Lambda_t} \left\{ \Lambda_t \cdot [V_t^a(a) - \kappa] + (1 - \Lambda_t) \cdot V_t(\tilde{p}, a) - \theta^a \mathcal{D}(\Lambda_t \parallel \bar{\Lambda}) \right\}, \quad (16)$$

where the subscript t indicates dependence on the aggregate state and we have suppressed the arguments of the adjustment probability to ease notation. The retailer either adjusts to a new price, with probability $\Lambda_t(\tilde{p}, a)$, or continues with its current price, which occurs with probability $1 - \Lambda_t(\tilde{p}, a)$. 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,

$$V_t(p_{jt}, a_{jt}) = \pi_t(p_{jt}, a_{jt}) + E_t \left\{ M_{t,t+1} V_{t+1}^*(\tilde{p}_{j,t+1}, a_{j,t+1}) \right\}, \quad (17)$$

where expectations condition on the current state, $M_{t,t+1}$ is the real discount factor between the two periods, $\tilde{p}_{j,t+1} = p_{jt}P_t/P_{t+1}$ is the real price at the beginning of the next period given the current-period real price p_{jt} and the aggregate price levels, and V_{t+1}^* is the maximum attainable value the firm can expect in the next period (assuming optimal choices henceforth), which takes the form of equation (16).

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 pricing policy, net of the information cost associated with it,

$$V_t^a(a) = \max_{f_t} \left\{ \int f_t(p|a) V_t(p, a) dp - \theta^p \mathcal{D}(f_t(p|a) \parallel \bar{f}(p)) \right\} \quad (18)$$

$$\text{s.t. } \int f_t(p|a) dp = 1. \quad (19)$$

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

$$\ln \left(\frac{\Lambda_t(\tilde{p}, a)}{1 - \Lambda_t(\tilde{p}, a)} \right) = \ln \left(\frac{\bar{\Lambda}}{1 - \bar{\Lambda}} \right) + \frac{1}{\theta^a} \left[V_t^a(a) - V_t(\tilde{p}, a) - \kappa \right]. \quad (20)$$

The model predicts a linear relationship between the conditional log odds, the unconditional 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 state of the economy. Moreover, the firm can choose how much to deviate from the reference distribution state by state, which means it can have state-dependent accuracy if the value of having more accurate timing of adjustments differs across states. The contribution to the total information flow of conditioning the adjustment decision on a period's state is equal to the probability of the firm finding itself in that state times the KL divergence of $\Lambda_t(\tilde{p}, a)$ from $\bar{\Lambda}$,

In the limit, $\theta^a \rightarrow \infty$ implies a constant probability of adjustment, as in the Calvo model, while as $\theta^a \rightarrow 0$ the adjustment decision converges to the deterministic menu cost adjustment rule. For intermediate values of this cost, the adjustment probability is stochastically state-dependent, as in the information model of [Woodford \(2009\)](#) and as in the random menu costs model of [Dotsey et al. \(1999\)](#).

Now consider the optimal choice in (18)-(19) for the probability of charging a particular price in a particular state. This choice too can be made independently for each state and satisfies

$$f_t(p|a) = \frac{\bar{f}(p) \exp \left\{ \frac{V_t(p,a)}{\theta^p} \right\}}{\int \bar{f}(\hat{p}) \exp \left\{ \frac{V_t(\hat{p},a)}{\theta^p} \right\} d\hat{p}} \quad (21)$$

for each p charged with positive probability in the steady state.

In equation (21), the value of charging a certain price in a particular state consists of both its current payoff (in terms of the period operating profit) and the expected continuation value of entering the following period with this price. A price is charged with a higher probability in a particular state than unconditionally if it yields a higher value in that state compared with the average value across all possible prices, such that the gain in value is enough to compensate for the increase in information expenditure needed to increase the probability of charging the price

in that particular state. A lower attention cost θ^p enables more differentiation across states in terms of the price being charged. As the cost approaches zero, the firm's repricing approaches a degenerate distribution centered on the optimal full-information reset price.

Taking the reference distributions as given, the optimal choice distributions are given by (20) and (21). This gives us an optimization-based approach to generalizing the menu cost model to a stochastic version. Entropy reduction generates stochastic decisions: shrinking uncertainty to a degenerate distribution is often too costly, so the decision-maker is left with some residual uncertainty about the optimal course of action. In our context, this means that the firm acts probabilistically both in its decision about whether or not to change its price and in its decision about which price to charge. But the degree of randomness in choice is the result of a cost-benefit analysis and, as long as information is not infinitely costly, the firm will be more likely to adjust when the value of adjusting is higher and more likely to set a particular price when its continuation value at that price is higher compared with other possible prices. Moreover, the firm can specify the accuracy of its decisions in each state. For instance, in some states, it may not be worthwhile to expend resources on very precise information about market conditions, since perhaps in those states, the firm's payoffs are not very sensitive to having the correct price in place. On the other hand, other states of the world may make mispricing very costly, in which case the firm will want to condition its decisions more strongly on those states and be willing to pay the extra cost associated with that accuracy.

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

$$\Omega_t(p, a) = [1 - \Lambda_t(p, a)] \cdot \tilde{\Omega}_t(p, a) + \left[\int \Lambda_t(\hat{p}, a) \tilde{\Omega}_t(\hat{p}, a) d\hat{p} \right] \cdot f_t(p | a). \quad (22)$$

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, which is given by last period's post-adjustment distribution with all real prices eroded by inflation and idiosyncratic states transitioned to new values ac-

ording to the law of motion for the idiosyncratic state.

In (22), the first term captures the mass of firms that start the period in state $p \times a$ and do not adjust their price, while the second term captures the mass of firms with idiosyncratic state a that adjust from any price \hat{p} to end up with price p .

This completes the exposition of the pricing block of the model. Given aggregate and idiosyncratic conditions, retailers make their pricing choices, generating a level of aggregate price dispersion Δ_t , given by

$$\Delta_t \equiv \int e^{-a_{jt}} p_{jt}^{-\varepsilon} dj. \quad (23)$$

This measure of price dispersion in turn is the only input into the non-pricing block of the model, allowing us to separate the information problem from the rest of the economy.

Given the solution to the retailers' pricing problem, total demand for the intermediate input is

$$x_t^d \equiv \int x_{jt} dj = \int p_{jt}^{-\varepsilon} e^{-a_{jt}} Y_t dj = Y_t \Delta_t. \quad (24)$$

3.3 Intermediate Goods Producers

The supply of intermediate inputs is determined by a continuum of competitive producers who choose labor to maximize static real profits

$$\pi_t^x = p_t^x x_t - w_t L_t \quad (25)$$

subject to the production function

$$x_t = e^{a_t + z_t} L_t, \quad (26)$$

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_z .

Optimization by the intermediate goods producer yields

$$p_t^x = e^{-(a_t+z_t)} w_t \quad (27)$$

and market clearing for the intermediate input yields total labor demand

$$L_t = Y_t \Delta_t e^{-(a_t+z_t)}. \quad (28)$$

3.4 Households

The representative household chooses streams of consumption C_t , labor supply L_t , and the real value of the risk-free bonds purchased in period t , B_t , to maximize lifetime utility,

$$E_0 \sum_{t=0}^{\infty} \beta^t \zeta_t \left[\ln(C_t - hC_{t-1}) - \xi_t \cdot \left(\frac{L_t^{1+\frac{1}{\nu}}}{1 + \frac{1}{\nu}} \right) + \chi_t \cdot B_t \right], \quad (29)$$

subject to the sequence of flow budget constraints

$$C_t + B_t = w_t L_t + D_t - T_t + B_{t-1} \frac{1 + i_{t-1}}{1 + \pi_t}, \quad (30)$$

and a no-Ponzi condition, where $\beta \in (0, 1)$ is the household's discount factor, ζ_t is a discount factor shock, $\nu \geq 0$ is the Frisch elasticity of labor supply, ξ_t is a shock to the relative disutility of working, and $h \in [0, 1)$ is the degree of habit in consumption, D_t are firm dividends, and T_t are lump-sum taxes net of transfers, and $\pi_t \equiv \frac{P_t}{P_{t-1}} - 1$ denotes the net inflation rate between period $t - 1$ and period t . The household supplies labor to the intermediate goods producers, earns a real wage w_t , and can invest in one-period bonds that earn a nominal rate i_t between period t and period $t + 1$.

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

optimization yields

$$\text{FOC}-C_t : \left(\frac{1}{C_t - hC_{t-1}} \right) - E_t \left[\left(\frac{\beta\zeta_{t+1}}{\zeta_t} \right) \left(\frac{h}{C_{t+1} - hC_t} \right) \right] = \lambda_t \quad (31)$$

$$\text{FOC}-L_t : \xi_t L_t^{\frac{1}{\nu}} = \lambda_t w_t \quad (32)$$

$$\text{FOC}-B_t : \chi_t + E_t \left[\left(\frac{\beta\zeta_{t+1}}{\zeta_t} \right) \left(\frac{1 + i_t}{1 + \pi_{t+1}} \right) \lambda_{t+1} \right] = \lambda_t. \quad (33)$$

The exogenous discount factor shock affects the intertemporal Euler equation and has been shown by [Justiniano, Primiceri & Tambalotti \(2010\)](#) and others to be an important (reduced-form) driver of consumption fluctuations. It is also often used to drive the economy to the effective lower bound on the monetary authority’s nominal policy rate ([Eggertsson & Woodford, 2003](#)). Shocks to the disutility of labor 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 “flight to quality” shocks χ_t . This is particularly relevant in an extension of the model that features investment in capital, since these shocks generate variation in the demand for risk-free bonds that allows the interest rate controlled by the monetary authority to deviate from the return on other assets. Originally, [Krishnamurthy & Vissing-Jorgensen \(2012\)](#) proposed introducing “convenience” assets, namely highly liquid, very safe assets, such as U.S. Treasuries, in the utility function, analogously to a money-in-the-utility specification.⁵ [Fisher \(2015\)](#) shows how this specification endogenizes the risk preference shocks that [Smets & Wouters \(2007\)](#) introduce to the consumption Euler equation to help generate comovement between investment and consumption, and [Campbell, Fisher, Justiniano & Melosi \(2017\)](#) use it in a DSGE model that quantifies the effects of forward guidance on the U.S. economy since the Great Recession.

⁵[Krishnamurthy & Vissing-Jorgensen \(2012\)](#) point to theoretical results demonstrating how assets with superior liquidity and perceived safety can command a premium. They also demonstrate the empirical relevance of this premium. Using U.S. data on corporate bond-Treasury spreads, [Krishnamurthy & Vissing-Jorgensen \(2012\)](#) estimate an average convenience yield of 73 basis points from 1926 to 2008, nearly two thirds of which represents liquidity convenience, with the remainder representing safety convenience. This spread implies that the Treasury yield is lower than the actual risk-free rate and suggests care when parameterizing the risk-free rate in models that do not explicitly model the preference for convenience assets.

3.5 Balanced Growth

The source of long-run growth in the model is labor-augmenting technological progress, z_t , which grows at a constant rate γ . To guarantee balanced growth, we require that the variables $b_t, \kappa_t, \theta_t^r, \theta_t^p$ grow at the same rate as z_t . Hence, we define $x_t \equiv x \cdot z_t$ for each element in this set.

3.6 Wage Rigidity

We include reduced-form real wage rigidity given by

$$w_t = \delta^w \tilde{w}^* + (1 - \delta^w) w_t^*, \quad (34)$$

where \tilde{w}^* is the steady state real wage and w_t^* is the competitive real wage.

3.7 Monetary and Fiscal Authorities

The monetary authority follows a Taylor rule that features interest-rate smoothing, responds to deviations of inflation from its (potentially time-varying) target and to deviations of output growth from long run growth, subject to a zero lower bound. When not constrained by the zero lower bound, monetary policy implements

$$i_t = i_t^{target} + \epsilon_{rt}, \quad (35)$$

$$i_t^{target} = \rho_i i_{t-1} + (1 - \rho_i) [i_{ss} + \phi_\pi (\pi_t - \pi_t^*) + \phi_y (\gamma_t - \gamma_{ss})], \quad (36)$$

where π_t and γ_t are inflation and GDP growth over the latest 12 months, π_t^* is a potentially time-varying inflation target that allows for persistent deviations of the policy rule from targeting a constant inflation rate, and ϵ_{rt} is an i.i.d. monetary policy shock.⁶

Finally, government spending is characterized by an exogenous spending pol-

⁶Sims (2013). Here, we define GDP as output net of the pricing frictions cost: $GDP = Y_t - F_t = C_t + G_t$.

icy g on output net of pricing frictions,

$$G_t = g(Y_t - F_t), \tag{37}$$

and is funded by lump-sum consumer taxes.

3.8 Shocks

We include a range of fundamental shocks, to avoid overstating the role the pricing frictions play in generating aggregate volatility. The aggregate exogenous shocks are to aggregate TFP (a_t), impatience (ζ_t) labor supply (ξ_t), bond demand (χ_t), markups (ε_t), the Taylor rule, trend inflation, and to permanent productivity growth. In addition, we allow for unanticipated shocks to the menu cost, the marginal costs of acquiring information, and to the standard deviation of idiosyncratic shocks. Variation in the pricing costs over time is interpreted as variation in the efficiency with which firms can process information and implement price changes. Allowing for this variation enables us to uncover to what extent nominal frictions move with the economy's cycle, and if there have been trends in pricing costs over time.

4 Steady State Frictions

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

4.1 Steady State Data

We use statistics on price-setting patterns from the U.S. Consumer Price Index (CPI) to estimate the steady-state parameters. 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 pricing moments with us.

The microdata underlying these statistics consist of approximately 80,000 monthly price quotes for products grouped into roughly 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 log 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 entry-level item (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 aggregating it to the expenditure-weighted median across ELIs. Higher moments are computed in a similar way, but by pooling data within each *major group* rather than within each ELI, and by taking an expenditure-weighted average across the 13 major groups.⁷

Table I reports the averages of several key pricing moments, for the full sample period (from January 1978 until March 2023), for the full sample and two subsamples: the post-Volcker period, following the end of Chairman Volcker’s term in August 1987, which may be of independent interest since it represents a period of established modern monetary policy, and for the Great Moderation, from January 1984 to June 2007, which we target in our steady state calibration. In addition to pricing moments, the table also reports the average values for inflation, GDP growth, and federal funds rate, which we also target in our steady state estimation, and which we take from the Federal Reserve Bank of St. Louis FRED.

We calibrate the model to target the Great Moderation period for the steady state since that was a period of modest macroeconomic instability. During the Great Moderation, the frequency of price changes averaged 10% per month, the mean absolute size averaged 7.4%, and the standard deviation of price changes was almost 13%. The distribution of price changes was negatively skewed and had fat tails. Approximately two thirds of price changes were price increases, and these

⁷Higher moments are particularly sensitive to outliers, which is why a small number of observations would be insufficient to compute them reliably. Luo & Villar (2021) discuss sample size constraints: since ELIs are narrowly defined categories, there are often not enough observations to estimate higher moments of the price change distribution at the ELI level.

TABLE I: DATA MOMENTS

	Full Sample 1978-2023.Q1	Post-Volcker 1987.Q4-2023.Q1	Great Moderation 1984-2007.Q2
Frequency	0.1131	0.1103	0.1005
Size (absolute value)	0.0735	0.0746	0.0740
Standard deviation	0.1293	0.1350	0.1295
Skew	-0.131	-0.168	-0.142
Kurtosis	10.6	10.2	10.9
Frequency up	0.0758	0.0713	0.0683
Size up	0.0690	0.0690	0.0697
Federal funds rate	0.0462	0.0307	0.0532
Inflation rate	0.0353	0.0275	0.0307
GDP growth rate	0.0153	0.0145	0.0206

Notes: The pricing statistics report average values for moments constructed for the monthly distributions of log-price changes. *Frequency up* and *Size up* report the frequency and size of price increases. In the data, price increases are more frequent than price cuts, but also smaller. The average effective federal funds rate, CPI inflation, and per capita real GDP growth for each sub-sample are annual rates. The first column of numbers reports statistics for the full sample, ending in March 2023. The second column reports statistics for the post-Volcker period, starting in the fourth quarter of 1987. The last column reports statistics for the Great Moderation, which the Federal Reserve dates between January 1984 and June 2007. Sources: Daniel Villar and FRED.

increases were on average nearly half a percentage point smaller than price cuts.

As discussed in prior work (e.g., [Bils & Klenow, 2004](#); [Golosov & Lucas Jr, 2007](#)), the pricing moments indicate high pricing volatility, despite a relatively stable macroeconomic environment. In particular, given the low variability of inflation, the coexistence of large and frequent price cuts and price increases points to the importance of fairly large and frequent idiosyncratic shocks. Hence we will solve for a stochastic steady state with idiosyncratic shocks to firms' desired prices.

4.2 Steady State Parameters

We parameterize the model's stochastic steady state by targeting averages over the Great Moderation period. The estimation of shocks away from the steady

TABLE II: STEADY STATE PARAMETERS

Parameter	Value	Target
Discount factor	$\beta = 0.6^{1/12}$	} $i^{ss} = 5.32\%$
Marginal utility of bonds	$\chi = 0.0578$	
Frisch elasticity of labor supply	$\nu = 2$	
Marginal disutility of labor supply	$\xi = 1.68$	Normalization
Gov't spending (% private consumption)	$g_c = 0.25$	} GM averages
Steady state inflation rate	$\pi^{ss} = 1.0307^{1/12}$	
Steady state productivity growth	$\gamma^{ss} = 1.0206^{1/12}$	
Elasticity of substitution among varieties	$\varepsilon = 11$	} GM pricing moments
Standard deviation of idiosyncratic shocks	$\sigma_a = 0.23$	
Persistence of idiosyncratic shocks	$\rho_a = 0.94$	
Fixed menu cost	$\kappa = 0.026$	
Marginal cost of timing accuracy	$\theta^a = 0.097$	
Marginal cost of repricing accuracy	$\theta^p = 1.07$	

Notes: The model is estimated at the monthly frequency and parameters are set either at conventional values or to match targets for the Great Moderation (GM) period. The bottom six parameters are estimated jointly, targetting seven steady state pricing moments.

state will then make use of the entire data, including the volatile periods at the beginning and end of the sample.

Since the pricing moments aggregate monthly price changes, we set our model to a monthly frequency. In this way, we avoid having to map pricing moments based on monthly data into quarterly moments, which would require us to make assumptions about price rigidities in the transformation.⁸

Table II presents the calibrated and estimated steady state parameter values. We set the steady state inflation rate to the realized average ($\pi^{ss} = 1.0307^{1/12}$),

⁸For example, a Calvo model and a menu cost model parameterized to deliver the same frequency and size of price changes at the monthly frequency yield different quarterly moments, because the menu cost model features selection in price adjustment, whereas the Calvo model does not. Hence, mapping the monthly size of price changes to a quarterly frequency requires making assumptions about price adjustment selection, which we seek to avoid.

and we set productivity growth to target the average real GDP growth per capita ($\gamma^{ss} = 1.0206^{\frac{1}{12}}$). Following [Michaillat & Saez \(2021\)](#), we parameterize wealth in the utility by setting the annual discount factor to 60% ($\beta = 0.6^{\frac{1}{12}}$) and then internally calibrating the parameter governing the marginal utility of bonds ($\chi = 0.0578$) such that the steady state nominal interest rate is equal to the average federal funds rate during the Great Moderation ($i^{ss} = 1.0532$). This specification helps the later estimation, when shocks push the economy to the effective lower bound on nominal interest rates.⁹ The Frisch elasticity of labor supply is $\nu = 2$ and the relative disutility of working is $\xi = 1.68$, set to normalize employment in the steady state. Finally, we set government spending to 25% of consumption, per the realized sample average.

We estimate a vector of six parameters that jointly affect price-setting frictions, $\Theta = \{\varepsilon, \sigma_a, \rho_a, \kappa, \theta^a, \theta^p\}$, to target the averages of seven pricing moments: frequency, size, standard deviation, skew, and kurtosis of price changes, and frequency and size of price increases. We estimate Θ as follows:

$$\Theta = \arg \min_{\mathcal{X}} (\mu(\mathcal{X}) - \mu^{data}) W (\mu(\mathcal{X}) - \mu^{data})' \quad (38)$$

where $\mu(\mathcal{X})$ is the vector of model moments for a given pricing parameter \mathcal{X} , μ^{data} is the vector data moments, and our weighting matrix (W) equals a diagonal matrix with the inverse of the data moments.

We find a modest menu cost ($\kappa = 0.026$), a larger, though still moderate variable cost of information to determine if a price change is warranted ($\theta^a = 0.097$), and a large cost of accuracy in repricing ($\theta^p = 1.07$). Together, these costs imply a steady state level of expenditure on repricing that is 2.6% of steady state revenues.

[Table III](#) reports the breakdown of price-setting costs: Firms spend almost 0.3% of steady state sales on the fixed cost of price reviews. They spend approximately 0.76% of revenues on acquiring information to decide if the value of adjusting is high enough to warrant an adjustment, and roughly twice that amount (1.6% of revenues) to determine what price to charge, conditional on adjustment.

Compared with the flexible price steady state with the same distribution of

⁹See also [Cuba-Borda & Singh \(2019\)](#).

TABLE III: STEADY STATE OUTCOMES

<i>Spending (share of revenues)</i>	
Fixed cost ($\kappa \bar{\Lambda}$)	0.0028
Review cost ($\theta^a I_{ss}^a$)	0.0076
Repricing cost ($\theta^p I_{ss}^p$)	0.0159
Total expenditure	0.0263
<i>Outcomes (relative to flex-price outcomes)</i>	
Consumption	0.9272
Employment	1.0514
Wages	0.9507
Output	0.9523
Price Dispersion	1.1041

Note: Baseline steady state estimation. Spending on pricing decisions is reported as a share of steady state revenues. Aggregate outcomes are reported as a share of the aggregate outcomes in an economy without pricing costs, but otherwise identically parameterized.

idiosyncratic shocks and the same elasticity of substitution among varieties, the economy with pricing frictions delivers significantly lower welfare. Steady state consumption is 7.3% lower, employment is 5.1% higher, and wages are 4.9% lower. Acquiring information and adjusting prices does not bring prices to the flexible price levels, and as a result, equilibrium price dispersion is 10.4% higher than the price dispersion that would be warranted given heterogeneity in productivities alone. Hence, although relatively modest at the level of each individual firm, these frictions aggregate to considerable losses for consumers.

Table IV presents the match between model and data for the pricing moments. Since we target more moments than we have parameters, the fit is not perfect, but it is quite close. In addition to the pricing costs, matching these moments also requires a high elasticity of substitution ($\varepsilon = 11$, which is higher than the value usually estimated in menu cost models, but closer to the values used in the trade literature), and highly volatile and persistent idiosyncratic shocks (which is typical in this literature, given the large volatility of price changes). The match

TABLE IV: PRICING MOMENTS

	Data	Model
Freq. of price changes	0.1005	0.0976
Absolute size	0.0740	0.0751
Standard deviation	0.1295	0.1223
Skew	-0.1419	-0.1419
Kurtosis*	10.9	11.1
Freq. of price increases	0.0683	0.0676
Size of price increases	0.0697	0.0728

* Reflecting in part cross-sectional heterogeneity (which may be due to differences across ELIs in either technologies or costs of repricing), kurtosis is much higher than the values typically reported in the literature, which tend to range between 3.5 and 5.5. See text for additional discussion.

is particularly notable since we use normally distributed shocks (rather than the leptokurtic shocks that are usually used in order to match the standard deviation and kurtosis of price changes). The first two moments (frequency and size) are the standard targets in Calvo and first-generation menu cost models, where they pin down the menu cost and the variance of idiosyncratic shocks needed to generate large price changes, in excess of the rate of inflation. The remaining moments are used in various second-generation menu cost models, as they help pin down the degree of asymmetry in the profit function and the degree of state-dependence and selection in price adjustment in response to aggregate shocks.

Among all the targeted moments, which are quite standard, kurtosis deserves additional discussion, since the value for our data is much higher than the values typically reported in the literature, which tend to range between 3.5 and 5.5. This difference may be due to the broader range of goods in the CPI than in other data sets, and may reflect a combination of cross-sectional heterogeneity (for instance, differences across ELIs in repricing costs or in technologies) and measurement error (since higher order moments are much more difficult to estimate accurately). How much are results driven by this higher than usual level of kurtosis in the distribution of price changes? Estimating the model parameters targeting a kurtosis of 5 would




Figure 1: Pricing Policy

Note: This figure plots the optimal adjustment probability Λ and the optimal pricing policy f , in steady state, as functions of log price and log marginal cost.

yield a significantly higher value for the cost of repricing accuracy ($\theta^p = 1.6$ instead of 1.1), a much lower elasticity of substitution ($\varepsilon = 6$ instead of 11), higher menu cost ($\kappa = 0.055$ vs. 0.026), and lower values for θ_p , σ_a , and ρ_a . In short, a lower kurtosis requires somewhat smaller shocks and higher markups, and a different mix of fixed cost and variable reviewing costs, but the main result that errors in pricing are the primary driver of pricing frictions is just as strong if not stronger.

4.3 Discussion of Steady-State Policies

What do the estimated parameters imply for the accuracy of pricing decisions? Inaccuracy in pricing arises along two dimensions: the decision of whether or not to review the existing price, as well as the decision of what price to set when deciding to change the existing price. The accuracy with which firms make these two decisions is captured in their adjustment probability and in their pricing policy, both of which are plotted in Figure (1).

We estimate fairly high accuracy in the timing of price adjustments, as indicated by the strongly state-dependent probability of adjustment shown in the left panel of Figure (1). When a firm's existing price is close to the price it would expect to set upon repricing, the probability of adjustment is very close to zero. But for prices that are farther away, the probability of adjustment rises rapidly, and as a result, prices that are far from the optimum do not survive long. The adjustment occurs especially quickly when prices are relatively low, resulting in an asymmetric

survival probability. Prices that are above average do not yield very large losses and as a result, are not worth reoptimizing. The probability of adjustment is nearly zero, unless marginal costs are particularly low, in which case, it rises gradually. Conversely, the probability of adjustment is almost always near one for prices below average. This asymmetry has been documented as a feature of rationally inattentive price adjustment (Woodford, 2009), and we find it is also a strong feature in historical U.S. data. In effect, not only is the profit function asymmetric, with the cost of under-pricing significantly rising much faster than that of over-pricing, but historical inflation levels, although not very high, enable firms to further save on repricing costs by pricing high and letting inflation erode the price away.

Conditional on adjustment, we estimate substantial errors in pricing. As shown in the right panel of Figure (1), prices are drawn from an imprecise, weakly state-dependent distribution. Two features of this policy stand out: First, there is significant price dispersion. For a given marginal cost, there is a wide range of prices that the firm could charge with a significant probability, reflecting a high degree of uncertainty about the right price to charge. Models that assume reset prices based on perfect information would yield a degenerate distribution, thus potentially significantly overstating the responsiveness of prices to shocks, conditional on adjustment. Second, a pattern of occasional sales arises endogenously in the optimal pricing policy, purely as a way to capture highly profitable demand opportunities. The endogenous adjustment probability ensures that if the firm sets prices that are too low, it will change these prices quickly – note the steep increase in the right side of the hazard function. Hence, the review policy effectively insures against pricing too low for too long. Given this insurance, firms choose to occasionally set low prices, to capture significant demand from competitors (the elasticity of substitution we estimate is high), knowing that if they are wrong in their assessment of demand or marginal cost, they will promptly get an accurate signal to revise prices, such that the profit loss will be short-lived.

Overall, we estimate a strongly state-dependent price adjustment probability but inaccurate prices. It seems firms can determine quite accurately when their prices have become obsolete, but they have more difficulty determining the right price level to set. The data suggest a decomposition of the information frictions

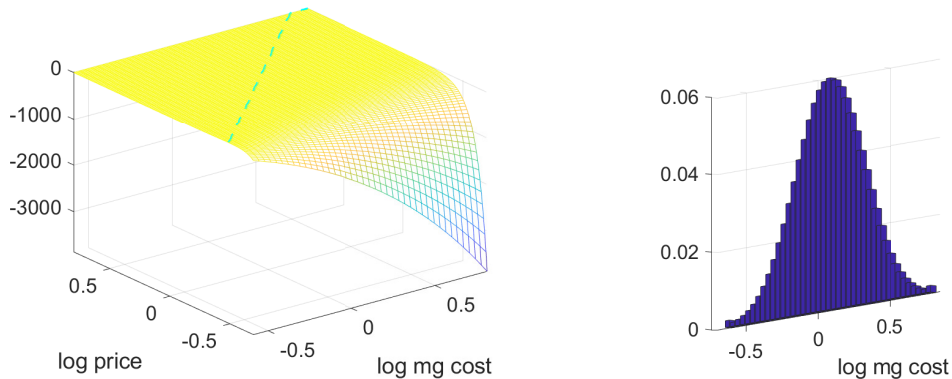


Figure 2: Firm Incentives

Note: This figure plots the profit function in the left panel and the distribution of marginal costs in the right panel. The near-diagonal line marks the full information flexible price optimum for each marginal cost value.

that is particularly severe along the pricing dimension, rather than along the dimension of whether or not to adjust one's price. This finding departs from the conventional models of price rigidity, which assume perfect but infrequent repricing.

4.4 Incentives for Information Acquisition

Figure (2) plots firms' operating profit function and the distribution of marginal costs for the baseline parameterization. These two objects shape firms' incentives to acquire information and set prices accurately in the steady state. The former governs the losses from mispricing, while the latter captures the incidence of these losses across states. In particular, the firm's operating profit per period – the flow profit before spending on information and adjustment costs – is asymmetric. The asymmetry shapes the attention of a firm trying to decide how to learn most efficiently about what price to set. Different considerations are at play, depending on where a firm finds itself in the distribution of marginal costs. Low current marginal costs provide an opportunity to generate significantly higher profits than average, especially given a high elasticity of substitution. Hence, firms have strong profit incentives to capitalize on these low-cost states. On the other hand, underpricing (relative to marginal cost) can generate large profit losses. Hence, firms have in-

centives to err on the side of over-pricing unless they are quite certain their costs are low.

Hence, just based on the shape of the firm's profit function, we would expect firms to have a tendency to (i) price high when they do not have much information, and (ii) pay more attention when profit opportunities are high. Prior work has emphasized the shape of the profit function, in the context of discussions regarding the strength of real rigidities as sources of price rigidity *distinct* from the nominal rigidity. Here, the sensitivity of a firm's profits to prices and how it varies across states in turn affects its willingness to set prices accurately, thus generating an *interaction* between real and nominal rigidities.

The incentive to set high prices is dampened by the frequency with which the firm expects to find itself in different states. With normally distributed shocks to marginal cost, the firm expects to spend more time closer to the center of the distribution, as shown in the second panel of Figure (2), which plots the distribution of marginal costs. This force induces firms to charge moderate prices more frequently, as these are likely to be close to optimal more often, thereby helping firms economize on information and adjustment costs.

4.5 Menu Cost Model Alternatives

To illustrate the incentives to price accurately in different states, we now compare our model to three menu cost alternatives: with perfect repricing, with repricing errors, and with timing errors. To show the differences across models more sharply, we begin with a higher menu cost than estimated, setting $\kappa = 0.07$.

Consider first the choices of firms that have perfect information for free, and are subject only to the menu cost when changing prices. Due to the fixed cost, there is a range of inaction in which firms keep their prices unchanged. Firms whose prices fall outside this range adjust to a dynamically optimal reset price. As has been extensively discussed in prior work, the resulting adjustment policy features asymmetric Ss bands of adjustment.

As shown in Figure (3), for low marginal costs, the Ss bands are very narrow and the optimal reset price, which maximizes the firm's continuation value, closely tracks the flexible price (though it is slightly above it): Despite the low incidence




Figure 3: Policies Across Alternative Models

Note: This figure plots the adjustment policy and repricing policy for four models: (a) the menu cost model with perfect repricing ($\kappa = 0.07$), (b) the menu cost model with errors in pricing ($\kappa = 0.07$, $\theta^p = 0.1$), in the top right panel, (c) the menu cost model with errors in timing ($\kappa = 0.07$, $\theta^a = 0.1$), in the bottom left panel, and (d) the model with all three costs. In all panels, the cyan dashed line marks the full information flexible price optimum for each marginal cost value, and the red solid line marks the optimal reset price. Panels (b) and (d) (for which $\theta^p > 0$) plot the weighted average of the conditional reset price distribution, f , and also super-impose this reset distribution.

of the low cost states, it is nonetheless worth it to the firm to price very accurately in these states, because the profit gains are so large. On the other hand, for higher marginal costs, the firm chooses instead a moderate price that applies to a wider range of states and enables it to save on the menu cost: The Ss bands widen (especially for prices above the optimal reset price), and the reset price becomes less and less sensitive to the marginal cost.

The top right panel of Figure (3) shows the optimal policy for a new type

of menu cost model that features repricing errors ($\theta^p > 0$). Instead of setting a deterministic reset price that maximizes its continuation value, the firm now sets a price probabilistically, drawn from an optimal distribution that maximizes its expected continuation value. The optimal distribution is given by an expression of the same form as in equation (21), with the appropriately adjusted value function.

Compared to the pure menu cost model, the errors-in-pricing model features additional endogenous price rigidity: The Ss bands widen and the frequency of price changes drops significantly. The possibility of mistakes in repricing arises as a new source of price rigidity, beyond the menu cost itself. A higher θ^p pushes toward a lower frequency of adjustment, biasing the firm toward inaction. But this is a rational response, given the possibility of mistakes in pricing. As a result, models may over-estimate the size of adjustment costs needed to match the frequency of price changes, if they abstract from the possibility of mistakes in repricing. Furthermore, compared to the optimal reset price in the pure menu cost model, the weighted average price charged in different states now becomes even less sensitive to marginal costs. The repricing distribution features strong dampening of price responsiveness to marginal costs that are above average, while maintaining the downward flexibility of prices for low marginal costs. The pricing policy also features stronger over-pricing at low marginal cost values, as a way for the firm to protect itself against underestimating its marginal cost.

The errors-in-pricing model features a distribution of reset prices that is most dispersed in the middle, reflecting the fact that the firm fine-tunes its pricing accuracy depending on the state. The firm's optimal conditional distribution of reset prices is tight at low marginal costs because low costs are highly profitable opportunities that are worth capturing accurately. On the other hand, it is also tight at high marginal costs, but for an entirely different reason: high marginal costs 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.

The bottom left panel of Figure (3) shows the optimal policy for a menu cost model with errors in timing, similar to the model of [Woodford \(2009\)](#). In this case, $\theta^p = 0$, and the firm sets the optimal reset price deterministically, but $\theta^a > 0$, such that the firm adjusts its price *probabilistically*, with a continuous adjustment

probability $\Lambda \in (0, 1)$, that takes a form similar to that in equation (20).

Compared to the pure menu cost model, the probability of adjustment is significantly above zero everywhere, even at the optimal reset price, since the firm is never certain of the state, and hence sometimes adjusts even when it should not. Away from the optimal reset price, the probability of adjustment increases gradually and asymmetrically: low costs are worth identifying accurately since, as before, they offer highly profitable opportunities, while the rest of the time, adjustment is Calvo-like even for modest information frictions. As θ^a increases, the probability of adjustment Λ becomes more and more Calvo-like, approaching a constant probability of adjustment $\bar{\Lambda}$.

Finally, the bottom right panel of Figure (3) reproduces the optimal policy for the menu cost models with errors in both timing and pricing. Timing and repricing accuracy are now chosen to be jointly optimal, and hence they interact to optimize the use of information. Mistakes in price-setting make firms more careful when changing prices so that the timing of adjustments becomes more state-dependent. Relative to the model of Woodford (2009), the introduction of mistakes in pricing makes the probability of adjustment Λ more well-shaped and with a wider region of near-inaction. On the other hand, mistakes in timing make the firm pay more attention to the price it sets, reducing over-pricing in low-cost states, slightly increasing sensitivity to marginal cost in high-cost states, and overall reducing the dispersion in the distribution of reset prices.

These interactions provide a new way to rationalize Calvo-like behavior. Firms that can learn what 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. But if figuring out the right price is very difficult, firms should pay close attention to when their prices are wrong, thus making their timing decision more state-dependent. 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, as indicated by the relative sizes of the two information costs.

TABLE V: NON-NEUTRALITY

Model	CIR
Our baseline model	0.112
The Calvo model	0.138
Woodford (2009)	0.097
Menu cost model	0.003

Notes: The table reports the cumulative response of consumption, as a percent of annual steady state consumption, in response to a 25-bp impulse to the federal funds rate.

4.6 Implied Non-neutrality

What do these pricing frictions imply for the degree of monetary non-neutrality? Table V reports the cumulative impulse response (CIR) of consumption to a 25 bp shock to the federal funds rate for our model and for alternatives parameterized to match the same steady state frequency and size of price adjustment.

The model deviates from the conventional wisdom, which posits that in order to get meaningful non-neutrality, one needs weak state dependence in the timing of price changes. For example, with perfect repricing, menu cost models need some auxiliary features to weaken the selection in terms of who adjusts in response to a shock. But with mistakes in repricing, firms can end up with a suboptimal price even if they correctly decide when to adjust. Hence mistakes in the timing of price changes are no longer strictly necessary for non-neutrality.

Mistakes in repricing mean that in the aggregate there is a lower effective frequency of adjustment. We can connect this discussion to the discussions of selection in time-dependent and state-dependent models ([Caballero & Engel, 2007](#); [Goloso & Lucas Jr, 2007](#); [Auclert, Rigato, Rognlie & Straub, 2022](#); [Gagliardone, Gertler, Lenzu & Tielens, 2023](#)). In effect, our model features negative selection on the *intensive* margin that either amplifies the lack of selection along the extensive margin in the Calvo model or works to offset the positive selection along the extensive margin in the menu cost model. As a result, the menu cost model




Figure 4: Exceeding the Calvo Bound on Non-neutrality

Note: This figure plots the CIR of consumption to a monetary policy shock in the Calvo model with errors in pricing, as we increase the cost of pricing accurately (θ^p) for (a) low elasticity of substitution and (b) high elasticity of substitution.

may overstate flexibility both because of over-estimating selection in terms of who adjusts and because of not considering the trade-offs between paying attention to *both* the pricing and the timing decisions.

4.7 Calvo and the Bounds on Non-neutrality

Firms with an accurate adjustment policy will not devote as much effort to implementing a very precise pricing decision: they will tolerate mistakes in price levels because the adjustment probability is steep so that these mistakes won't survive long. On the other hand, firms that have a more imprecise adjustment policy that conditions on the state more weakly, will choose to pay more attention to the prices they charge, so that even if they mistakenly decide to change prices, there is a relatively high probability that they will choose a suitable price.

This interaction illustrates how both the Calvo model and the menu cost model can overstate price flexibility. Information frictions suggest that if the adjustment is Calvo-like, the firm will try to make the pricing more accurate, while if the adjustment is accurate like in the menu cost model, the pricing will be imprecise.

With perfect repricing, varying the severity of information frictions regarding the *timing* of price adjustment spans the degree of state dependence in price setting, with the menu cost model at one end (when the information friction approaches

zero) to Calvo (1983) at the other end (when the information friction is strong enough that the firm acquires no information to decide when to adjust its prices), as shown by Woodford (2009). However allowing for errors in the repricing decision itself adds a new source of mispricing, and hence non-neutrality. As a result, Calvo need no longer be the upper bound on price rigidity: information could be so costly that both the adjustment and the pricing decision are imprecise.

Figure (4) illustrates how non-neutrality can exceed that implied by the Calvo model with perfect repricing, by plotting the cumulative impulse response of consumption to a monetary policy shock as we increase θ^p from 0 to 1.

As the cost of setting accurate prices increases, mistakes in repricing increase, and as a result, non-neutrality increases. However, the relationship between non-neutrality and θ^p is non-monotonic, and the turning point depends on the degree of strategic complementarities. This can be seen more clearly in the second panel of Figure (4), which plots the CIR for a parameterization with stronger strategic complementarities.

The non-monotonicity in the CIR reflects the non-monotonicity in price dispersion, which is a large contributor to the CIR. As θ^p increases, price dispersion initially increases: each firm's pricing policy becomes more imprecise. This feeds into the cross-sectional distribution of prices \bar{f} , which in turn feeds into the individual pricing policy, f . At a certain point, θ^p gets large enough that price dispersion starts to decrease, plateauing at the no-information limit.

This nonmonotonicity shows another way in which the errors-in-pricing model diverges from models with perfect repricing: in the Calvo or menu cost model, the degree of real rigidities is an independent amplifier of nominal rigidities. Here, however, they are no longer independent.

5 Statistics and Simulations

In the model, the pricing costs pin down the degree of monetary non-neutrality. What is the relationship between the model's pricing costs and moments of the price change distribution? Is there a sufficient statistic that captures the frictions in price setting in this information-constrained model of price setting, as is the case with a wide class of menu cost models (Alvarez et al., 2016)?

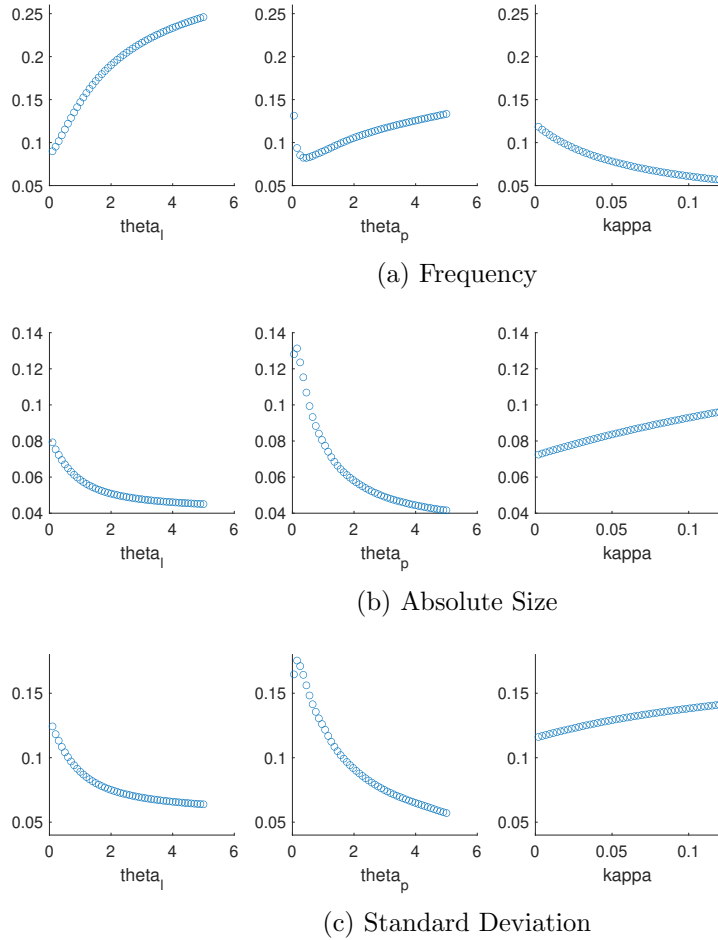


Figure 5: Identification of Policy Parameters in Simulations (A)

5.1 Identification

How do pricing moments identify parameters in our model? Figures (5) and (6) plot how the pricing moments vary with the three key pricing parameters (θ_p , θ_a , and κ).

As has already been discussed in the menu cost literature, the frequency of adjustment is the key parameter that pins down the size of the menu cost. To match the size, standard deviation, and kurtosis, menu cost models then rely on properties of the distribution of idiosyncratic shocks (for example, large shocks Golosov & Lucas Jr (2007) and fat tails as in Midrigan (2011)).

We find that frequency is also highly informative for the identification of the

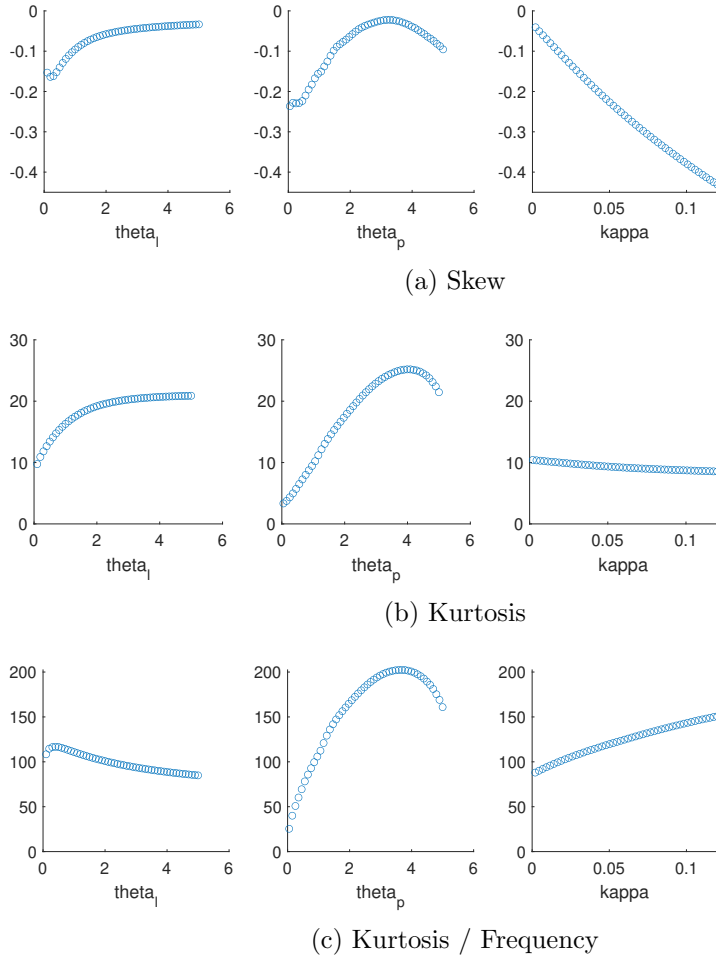


Figure 6: Identification of Policy Parameters in Simulations (B)

information costs, θ^p and θ^a . However, these parameters also have large effects on other pricing moments: the size and dispersion of price changes are particularly sensitive to θ^p , while θ^a is relevant for skew and kurtosis.

5.2 From Prices to Non-neutrality

How does the model's implied degree of nominal rigidities, as measured by the CIR of consumption to a monetary policy shock, vary with the pricing costs? We find that the CIR is particularly sensitive to mistakes in repricing driven by the repricing cost θ^p . These generate substantial price dispersion and also significant expenditure on information acquisition, both of which lower consumption. The

Figure 7: Simulation Results: Non-Neutrality and Pricing Costs

Note: The figure shows scatter plots of the cumulative response of consumption to a monetary policy shock for different model simulations parameterized to a range of pricing costs. The top panel varies the menu cost κ , the middle panel varies the review cost θ^a , and the lower panel varies the repricing cost θ^p . The panels also show how the CIR varies with the Alvarez et al. (2016) index (defined as kurtosis/frequency) for different parameterizations.

Figure 8: Simulation results varying all three costs, κ , θ^a , and θ^p .

Note: The figure shows scatter plots of the cumulative response of consumption to a monetary policy shock for different model simulations parameterized to a range of pricing costs. The figure shows how the CIR varies with the Alvarez et al. (2016) index (defined as kurtosis/frequency) for different parameterizations.

CIR is also sensitive — though to a lesser extent — to the value of the fixed cost κ and it is about half as sensitive to the accuracy-in-timing cost θ^a as it is to κ .

Figure (7) also plots simulated data to illustrate the strong positive relationship in our model between the Alvarez et al. (2016) index (kurtosis / frequency) and non-neutrality as we vary each pricing cost. Nevertheless, while strong, the relationship does not make ALL a sufficient statistic. This reflects the fact that even 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. It is nonetheless important to note that even with substantial errors in pricing, kurtosis and frequency remain highly informative moments about the CIR, along with standard deviation and size.

Figure (8) shows a scatterplot of the CIR against the ALL index when we allow all three pricing costs to vary. The nonlinear interaction between timing and pricing decisions as we vary all three parameters results in a weaker positive relationship between the CIR and the ALL index, as now the dispersion of price changes becomes more informative as well.

6 Estimation

In this section, we report results from our Bayesian estimation. To our knowledge, this is the first use of Bayesian techniques to estimate a model featuring rationally inattentive agents.

6.1 Estimation Approach

To gain both speed and accuracy, we compute the equilibrium dynamics using the sequence-space Jacobian (SSJ) method of [Auclert et al. \(2021\)](#). So far, the SSJ approach has been successfully used in a variety of models with heterogeneity across households; here we show how it can also be very effectively used in models with heterogeneous information.

We extend the SSJ 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. The appendix describes how we adapt the methods proposed by [Guerrieri & Iacoviello \(2015\)](#) and [Kulish et al. \(2017\)](#) to handle the constraint when solving the model dynamics using the SSJ method.

The estimation includes fundamental shocks to preferences, technologies, and policies (whose realizations are not observable for free to firms) as well as shocks to the pricing frictions themselves, interpreted as shocks to attention or efficiency of information processing and implementation of decisions.

The presence of the two ELB periods makes the model evaluation time-consuming. In addition, we evaluate the model many times since we have a monthly model with a large number of observations (541 months). Hence, to reduce computational time, we only estimate the shocks. We set habit formation in consumption to $h = 0.67$, based on the meta-study of DSGE models of [Havranek, Rusnak & Sokolova \(2017\)](#), the rigidity in real wages to $\delta_w = 0.083$, based on typical values, and we run a separate Bayesian estimation for the parameters of the Taylor rule for the monetary authority. [Table VI](#) reports the parameter values we use and the appendix describes the estimation of the Taylor rule coefficients.

TABLE VI: CALIBRATED PARAMETERS

Parameter	Value	
Habit in consumption	$h = 0.67$	Havranek et al. (2017)
Wage rigidity	$\delta_w = 0.08$	Monthly rigidity
Taylor Rule		Separately estimated
Persistence	$\rho_i = 0.925$	
Inflation coefficient	$\phi_\pi = 1.7$	
GDP growth coefficient	$\phi_y = 1.0$	

Notes: The table reports the values of the parameters that are calibrated rather than estimated in the dynamic estimation. The procedure used to estimate the Taylor rule parameters is detailed in the appendix.

Tracking the Distribution of Price Changes over Time We use the time series for five pricing moments to keep track of how the distribution of price changes evolves over time. These are the frequency of price changes, the mean absolute size, the standard deviation, the skew, and the kurtosis of the distribution of price changes in the Consumer Price Index.

Figure 9 plots the time series for these pricing moments for the estimation sample period. Of particular interest, note that the frequency, size and standard deviation of price changes all started increasing and kurtosis started decreasing well before the 2020 recession. While neither of these moments is sufficient on its own to pin down non-neutrality, these synchronized movements already suggest that we might expect an increase in the degree of price flexibility starting around 2016. As a result, we might expect any inflationary shocks, should they occur, to be met with a sharper inflation response, rather than a sharper output response post-2026 versus pre-2016.

Other notable patterns in these pricing series are (i) a relatively high and volatile frequency of price adjustment, ranging between 10% and over 20% monthly, (ii) a more stable size of price changes, ranging between 6% and 8% per month, (iii) a rising dispersion in price changes, measured by the standard deviation of price changes, which increases from less than 10% at the beginning of the sample to more than 15% by the end of the sample, (iv) volatile skew and kurtosis, both trending down over time.

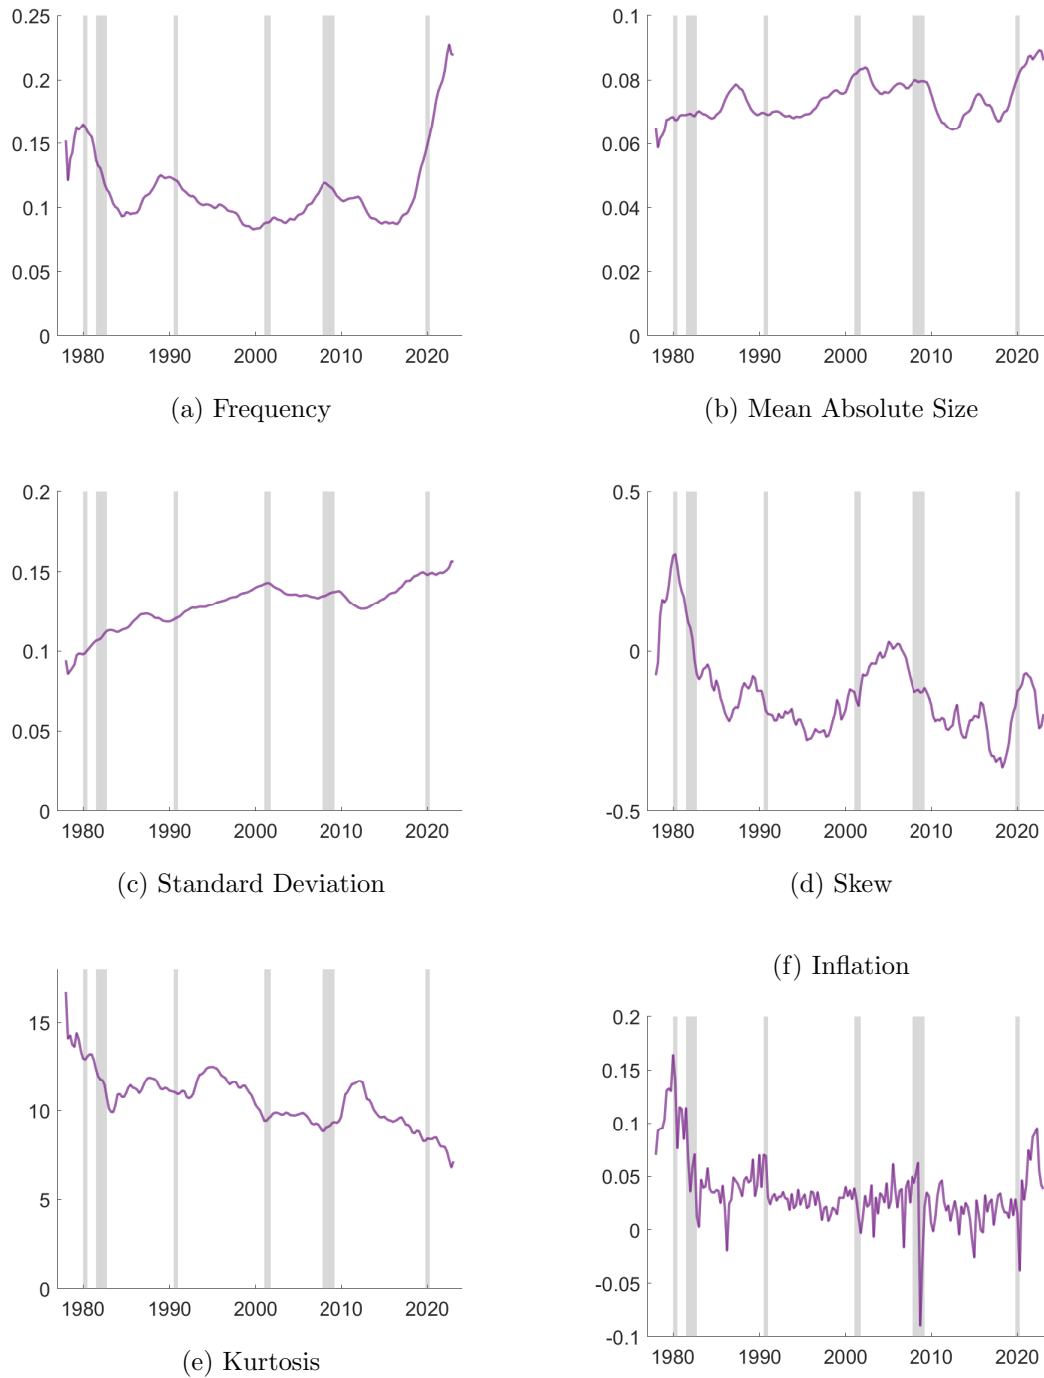


Figure 9: Pricing Moments Over Time

Note: This figure plots the smoothed pricing series used for the model estimation and the annualized CPI inflation rate. Shaded areas mark NBER recession dates.

In the last panel of the figure, we also plot the ratio of kurtosis over frequency, which [Alvarez et al. \(2016\)](#) demonstrate is a sufficient statistic for nominal flexibility in a wide range of state-dependent pricing models. This statistic also exhibits variability and was relatively high both during the early 1980s and during the Great Recession.

Macro Series We complement the pricing data with quarterly real GDP growth per capita, the monthly series for the federal funds rate, and the monthly CPI inflation rate. We also include long-term interest rates in our estimation as in [Kulish et al. \(2017\)](#), to aid with identification at the Effective Lower Bound on the nominal interest rate. Appendix A provides details about the data sources and transformations.

Measurement Equations For each of the five pricing moments x , the measurement equation we use is of the form

$$\text{CPI moment in quarter } q = \frac{1}{3}(x_{q,1} + x_{q,2} + x_{q,3}) + c_x, \quad (39)$$

where we average the monthly model moments for the months in each quarter, and where c_x is a constant we add to center the model's mean at the data mean over the estimation sample period.

The remaining measurement equations are

$$\text{Federal funds rate} = i_t^{12} \quad (40)$$

$$\text{Quarterly CPI inflation rate} = \pi_t \pi_{t-1} \pi_{t-2} \quad (41)$$

$$\begin{aligned} \text{Real GDP growth} = & \log(\tilde{Y}_t + \tilde{Y}_{t-1} + \tilde{Y}_{t-2}) - \log(\tilde{Y}_{t-3} + \tilde{Y}_{t-4} + \tilde{Y}_{t-5}) \\ & + \log(\gamma_t + \gamma_{t-1} + \gamma_{t-2}) \end{aligned} \quad (42)$$

$$\text{Two-year Treasury yield} = \frac{1}{24} \sum_{s=0}^{24} i_t^{12} + \eta_t + \eta_{2t} + c_2 \quad (43)$$

$$\text{Five-year Treasury yield} = \frac{1}{60} \sum_{s=0}^{60} i_t^{12} + \eta_t + \eta_{5t} + c_5 \quad (44)$$

where GDP in month t is $\tilde{Y}_t = C_t + G_t$, and c_2 and c_5 are centering constants such that the mean yields are equal to the empirical average yields over the estimation sample period.

Shocks We include a range of fundamental shocks, to avoid overstating the role the pricing frictions play in generating aggregate volatility.

The aggregate exogenous shocks are to aggregate TFP (a_t), impatience (ζ_t) labor supply (ξ_t), bond demand (χ_t), markups (ε_t), the Taylor rule, trend inflation, and permanent productivity growth. To aid with the estimation at the ELB, following [Kulish et al. \(2017\)](#), we include shocks to the term structure of interest rates: a common AR(1) yield shock, and i.i.d. shocks to the two-year and five-year Treasuries yields. In addition, we allow for unanticipated shocks to the menu cost, the marginal costs of acquiring information, and to the standard deviation of idiosyncratic shocks. Variation in the pricing moments over time allows us to disentangle the contribution of each of these shocks to price rigidity over time.

6.2 Estimation Results

Table [VII](#) shows the list of parameters we estimate along with the prior and the optimization mode.

We present the estimated series for the three costs in [Figure 10](#). This figure plots the series for θ^p , θ^a , and κ when the model parameters are evaluated at the optimization mode. Jointly, these parameter values imply a fairly volatile degree of nominal price rigidities over time. [Figure 11](#) plots our estimated series for the NPR between January 1978 and March 2023.

TABLE VII: ESTIMATION OF SHOCKS

Name	Prior	Mean	SD	Optimization Mode
σ_{θ^p}	\mathcal{IG}	1	0.5	2.40
σ_{θ^a}	\mathcal{IG}	1	0.5	1.44
σ_{κ}	\mathcal{IG}	0.05	0.025	0.01
σ_{risk}	\mathcal{IG}	0.05	0.0125	0.02
σ_{ε}	\mathcal{IG}	0.05	0.0125	5.62
σ_{γ}	\mathcal{IG}	0.19	0.0475	0.75
σ_{ζ}	\mathcal{IG}	0.05	0.0125	5.05
σ_a	\mathcal{IG}	0.05	0.0125	20.11
σ_{χ}	\mathcal{IG}	0.05	0.0125	5.55
σ_{π^*}	\mathcal{IG}	0.117	0.0292	0.11
σ_i	\mathcal{IG}	0.11	0.0275	0.69
σ_{yield2}	\mathcal{IG}	1.2	0.6	1.79
σ_{yield5}	\mathcal{IG}	1.1	0.55	0.03
σ_{yield}	\mathcal{IG}	0	1	0.75
σ_{ξ}	\mathcal{IG}	0.05	0.0125	2e-06
ρ_{θ^p}	\mathcal{B}	0.5	0.15	0.956
ρ_{θ^a}	\mathcal{B}	0.5	0.15	0.989
ρ_{κ}	\mathcal{B}	0.5	0.15	0.963
ρ_{risk}	\mathcal{B}	0.5	0.15	0.968
ρ_{ε}	\mathcal{B}	0.5	0.15	0.001
ρ_{γ}	\mathcal{B}	0.025	0.15	0.546
ρ_{ζ}	\mathcal{B}	0.5	0.15	0.896
ρ_a	\mathcal{B}	0.9	0.15	0.034
ρ_{χ}	\mathcal{B}	0.9	0.15	0.931
ρ_{ξ}	\mathcal{B}	0.025	0.15	0.025
ρ_{π^*}	\mathcal{B}	0.5	0.15	0.323
ρ_{yield}	\mathcal{B}	0.5	0.15	0.980

Notes: We use Inverse Gamma prior distributions for the for standard deviations of shocks and Beta prior distributions for the autoregressive coefficients.

The degree of NPR in each period is given by what the cumulative response of real output would have been in reaction to a monetary policy shock in that period. We estimate a fairly large degree of non-neutrality. Surprisingly, we find no trend over time, despite the fact that technology has arguably made repricing easier. We do however find substantial volatility in the degree of NPR over time.

A growing literature has argued that monetary policy may be less effective during downturns. This argument rests on evidence that price dispersion rises during downturns, which existing models predict implies more price flexibility. Our results suggest that this chain of reasoning did not apply during recent U.S. recessions. We estimate that price rigidities rose modestly during the Great Recession, and otherwise remain largely unchanged during downturns. Instead, price rigidities seem to vary at a medium-term frequency.

Our estimation also suggests that price rigidities started declining well before the Covid-19 pandemic, pointing to the risk of higher inflation even before the subsequent shocks.

One interpretation of our results is that they show that moments of the price distribution are not sufficient to pin down the effectiveness of monetary policy. In a way this is an expansion of the point originally made by the early literature on menu costs, which showed that—contrary to popular belief—the frequency of price changes is not a sufficient statistic for the degree of nominal rigidity. We push this point further by showing that different models that match a range of moments of the price distribution but have different assumptions regarding the information that is available to decision makers when making price setting decisions also have different implied degrees of monetary non-neutrality over the business cycle.

How to reconcile our findings with the growing literature that has documented diminished monetary policy effectiveness in downturns? We believe the answers lies in the distinction between the transmission of a central bank’s policy to the short term real interest rate and the effect that rate has on consumption, investment, employment, and other macroeconomic variables. In this article, we concentrate on the first part, while the cyclical and state-dependence of the second part has been the focus of papers such as those of [Berger, Milbradt, Tourre & Vavra \(2021\)](#); [Eichenbaum, Rebelo & Wong \(2022\)](#); [McKay & Wieland \(2021\)](#).

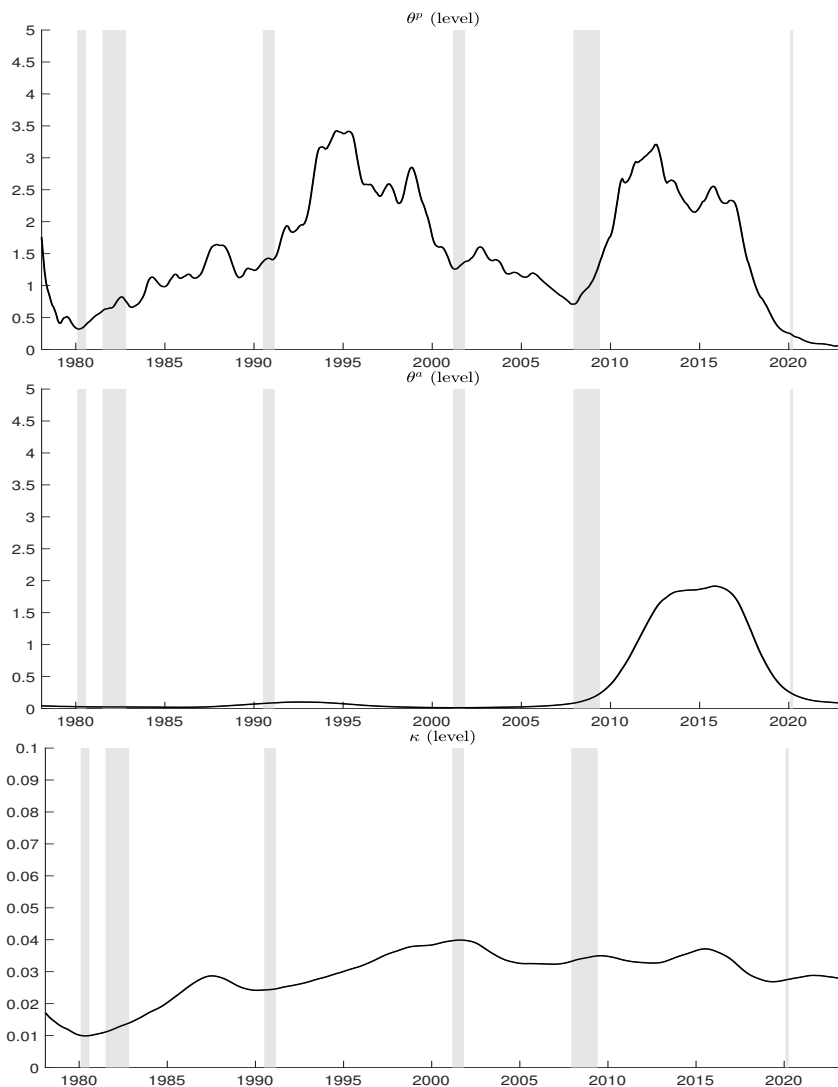


Figure 10: Estimated Pricing Frictions

Note: This figure plots the estimated series for θ^p , θ^a , and κ when the model parameters are evaluated at the optimization mode. Shaded areas represent NBER recession dates.

7 Conclusion

This article estimates the degree of nominal price rigidity in the U.S. economy: its trend and variability over time. We identify costly information as the main friction that prevents firms from adjusting prices more flexibly in response to shocks,




Figure 11: Implied Nominal Rigidity Over Time

Note: This figure plots the model implied degree of nominal rigidities over time. For each point in time, we solve the model and compute the cumulative impulse response to a monetary policy shock, using the pricing parameters at that time. For this exercise, we fix the reference probability of price adjustment $\bar{\lambda}$ and the reference distribution of prices \bar{f} at their baseline value and keep all other parameters constant at the optimization mode value. Shaded areas represent NBER recession dates.

with information about the right price to charge, conditional on adjustment, being the most significant driver of price rigidity. 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 responding to consumer constraints). 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

- 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, Martin Beraja, Martin Gonzalez-Rozada & Pablo Andrés Neumeyer (2019), “From hyperinflation to stable prices: Argentina’s evidence on menu cost models,” 134(1): 451–505.
- 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 multi-product 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.
- 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.
- 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.
- Bayer, Christian, Benjamin Born & Ralph Luetticke (2020), "Shocks, frictions, and inequality in U.S. business cycles," Forthcoming, *The American Economic Review*.
- Berger, David, Konstantin Milbradt, Fabrice Tourre & Joseph Vavra (2021), "Mortgage prepayment and path-dependent effects of monetary policy," *American Economic Review* 111(9): 2829–78.
- 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.
- 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, 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.
- Carvalho, Carlos & Oleksiy Kryvtsov (2021), "Price selection," *Journal of Monetary Economics* 122: 56–75.

- 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.
- 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* .
- 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.
- Eichenbaum, Martin, Sergio Rebelo & Arlene Wong (2022), “State-dependent effects of monetary policy: The refinancing channel,” *American Economic Review* 112(3): 721–61.
- Fernández-Villaverde, Jesús, Juan F Rubio-Ramírez, Timothy Cogley & Frank Schorfheide (2007), “How structural are structural parameters?[with comments and discussion],” *NBER Macroeconomics Annual* 22: 83–167.
- 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.
- Gagliardone, Luca, Mark Gertler, Simone Lenzu & Joris Tielens (2023), “Anatomy of the Phillips Curve: Micro Evidence and Macro Implications,” Tech. rep., National Bureau of Economic Research.
- Gagnon, Etienne (2009), “Price setting during low and high inflation: Evidence from Mexico,” *The Quarterly Journal of Economics* 124(3): 1221–1263.
- Golosov, Mikhail & Robert E Lucas Jr (2007), “Menu costs and Phillips curves,” *Journal of Political Economy* 115(2): 171–199.
- 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.
- Justiniano, Alejandro, Giorgio E Primiceri & Andrea Tambalotti (2010), “Investment shocks and business cycles,” *Journal of Monetary Economics* 57(2): 132–145.
- Karadi, Peter, Raphael Schoenle & Jesse Wursten (2022), “Price Selection in the Micro-data,” Working Paper.
- Khaw, Mel Win, Luminita Stevens & Michael Woodford (2017), “Discrete adjustment to a changing environment: Experimental evidence,” *Journal of Monetary Economics* 91: 88–103.
- Khaw, Mel Win, Luminita Stevens & Michael Woodford (2019), “Empirical models of inattentive choice,” Working Paper.
- 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.
- 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.
- McKay, Alisdair & Johannes F Wieland (2021), “Lumpy durable consumption demand and the limited ammunition of monetary policy,” *Econometrica* 89(6): 2717–2749.
- 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.
- 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* .
- Smets, Frank & Rafael Wouters (2007), “Shocks and frictions in US business cycles: A Bayesian DSGE approach,” *American economic review* 97(3): 586–606.
- 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.
- 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 (2012), “Inattentive valuation and reference-dependent choice,” Unpublished Manuscript.

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

A SSJ and Model Estimation During ELB Periods

For linear models written in a recursive formulation, [Kulish et al. \(2017\)](#) and [?](#) show that, for an expected ELB duration, the law of motion of the economy can be written as a time-varying linear function. As a result, the likelihood of the data can be computed based the Kalman filter with time-varying coefficients. Hence, 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. [Kulish et al. \(2017\)](#) even propose to estimate the ELB duration.

In the context of this paper, where we use SSJ to solve and estimate the model, how can estimate the model during the ELB period? One possibility is to (1) recover the recursive formulation of the model, as shown by [Auclert et al. \(2021\)](#), (2) compute the time varying formulation of the problem as shown by [Kulish et al. \(2017\)](#), (3) find the (time-varying) MA representation of the model, and (4) compute the log-likelihood of the model accordingly. This is time consuming as the matrix operations of step (2) can be computationally demanding for large number of state variables.

Instead, we propose a new and efficient way of computing the log-likelihood of the model during the ELB period for a given expected ELB duration using SSJ. This procedure is equivalent to the aforementioned possibility, but our approach is faster. We first present the procedure, we then explain each step in details, and we finally show its consistency/mapping with the method proposed by [Kulish et al. \(2017\)](#) and [?](#) for models written in recursive formulation.

A.1 Steps

1. Compute the MA representation of the model for the nx shocks, using SSJ.
2. Solving the model using SSJ, include L anticipated monetary policy shocks, where L should be as large as the expected ELB duration.
3. 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:

$$\hat{i}_t = MA_{i,1}\epsilon^{t-1} + \tilde{\alpha}_{i,0}\tilde{\epsilon}_t + \beta_{i,0}\mu_t \quad (\text{A.1})$$

where:

$$\epsilon_t = [\tilde{\epsilon}_t \quad \mu_t]' \quad (\text{A.2})$$

$$\epsilon^t = [\epsilon_t \quad \epsilon_{t-1} \quad \dots \quad \epsilon_{t-T}]' \quad (\text{A.3})$$

$$\mu_t = [\varepsilon_{0,t} \quad \varepsilon_{1,t} \quad \dots \quad \varepsilon_{L,t}]' \quad (\text{A.4})$$

$$MA_{i,j} = \begin{bmatrix} \alpha_{i,j}1 & \alpha_{i,j+1}2 & \dots & \alpha_{i,T} & \underbrace{0 \quad 0 \quad \dots \quad 0}_{j \text{ times}} \end{bmatrix} \quad \forall j = 0 \dots T \quad (\text{A.5})$$

$$\alpha_{i,j} = [\tilde{\alpha}_{i,j} \quad \beta_{i,j}]' \quad \forall j = 0 \dots 1 \dots T \quad (\text{A.6})$$

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 , and *epsilon* _{t} is the combination of the those two (i.e. all aggregate shocks at time t).

Therefore, at time t , the expected federal funds rate for $t + 1$ is:

$$E_t [\hat{i}_{t+1}] = [\alpha_2 \quad \alpha_3 \quad \dots \quad \alpha_T 0] \epsilon^{t-2} + \tilde{\alpha}_1 \tilde{\epsilon}_t + \beta_1 \mu_t \quad (\text{A.7})$$

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

$$\begin{bmatrix} i_t \\ E_t [\hat{i}_{t+1}] \\ \vdots \\ E_t [\hat{i}_{t+L}] \end{bmatrix} = \underbrace{\begin{bmatrix} \alpha_1 & \alpha_2 & \alpha_3 & \dots & \alpha_{T-2} & \alpha_{T-1} & \alpha_T \\ \alpha_2 & \alpha_3 & \alpha_4 & \dots & \alpha_{T-1} & \alpha_T & 0 \\ \alpha_3 & \alpha_4 & \alpha_5 & \dots & \alpha_T & 0 & 0 \\ \vdots & \vdots & \vdots & & \vdots & \vdots & \vdots \\ \alpha_{L+1} & \alpha_{L+2} & \alpha_{L+3} & \dots & 0 & 0 & 0 \end{bmatrix}}_{\Omega} \epsilon^{t-1} + \underbrace{\begin{bmatrix} \tilde{\alpha}_0 \\ \tilde{\alpha}_1 \\ \tilde{\alpha}_2 \\ \vdots \\ \tilde{\alpha}_L \end{bmatrix}}_{\omega} \epsilon_t + \underbrace{\begin{bmatrix} \beta_0 \\ \beta_1 \\ \beta_2 \\ \vdots \\ \beta_L \end{bmatrix}}_{\lambda} \mu_t \quad (\text{A.8})$$

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

$$\begin{bmatrix} i_t \\ E_t [\hat{i}_{t+1}] \\ \vdots \\ E_t [\hat{i}_{t+L}] \end{bmatrix} = \underbrace{\begin{bmatrix} i-1 \\ i-1 \\ \vdots \\ i-1 \end{bmatrix}}_Z \quad (\text{A.9})$$

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

$$\mu_t = \lambda^{-1} [Z - \Omega\epsilon^{t-1} - \omega\epsilon_t] \quad (\text{A.10})$$

This implies that the MA representation for the economy at period t is given by:

$$MA_t = MA \quad (\text{A.11})$$

B Data Description

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.

To calibrate the expected ELB duration at each point in time, we use the Blue Chip data and the survey of primary dealers. For each survey, we compute the expected ELB duration, in weeks, for each week as described below. Then, our weekly series correspond to the BlueChip series between 2008 and January 2011, and to the survey of professional forecasters since January 18, 2011.

Using the Blue Chip microdata, we compute the expected ELB duration (in quarters) for each forecaster and month.¹⁰ Then, we compute the median expected ELB duration for each month across forecasters. By construction, this expected duration is top-coded, as the Blue Chip survey only asks for the expected federal funds rate value for the next 5 quarters. However, before 2011, the median expected ELB duration is less than 5 quarters. We compute the expected duration in weeks by assuming a lag of one week between data collection and publication and by using the FOMC meetings calendar. If a quarter has two FOMC meetings, we take the simple average of the expected duration associated with those two meetings. Finally, we get the expected ELB duration in weeks for each weeks by interpolating.

¹⁰Blue Chip is a monthly survey, but they ask for the expected federal funds rate in quarter intervals. For example, the expected federal funds rate value in 2010Q1.

Since January 2011, the survey of primary dealers asks for the “*Most Likely Quarter and Year of First Target Rate Increase*”, later in our sample, they asked for the specific FOMC meeting. 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 weeks. 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 weekly series by interpolating.

C Additional Data Plots

We smooth the pricing series using a centered moving average filter to reduce volatility associated with measurement error. Figure (C.1) plots the time series of the different pricing moments we use, both raw and smoothed. The estimation uses the smoothed series.

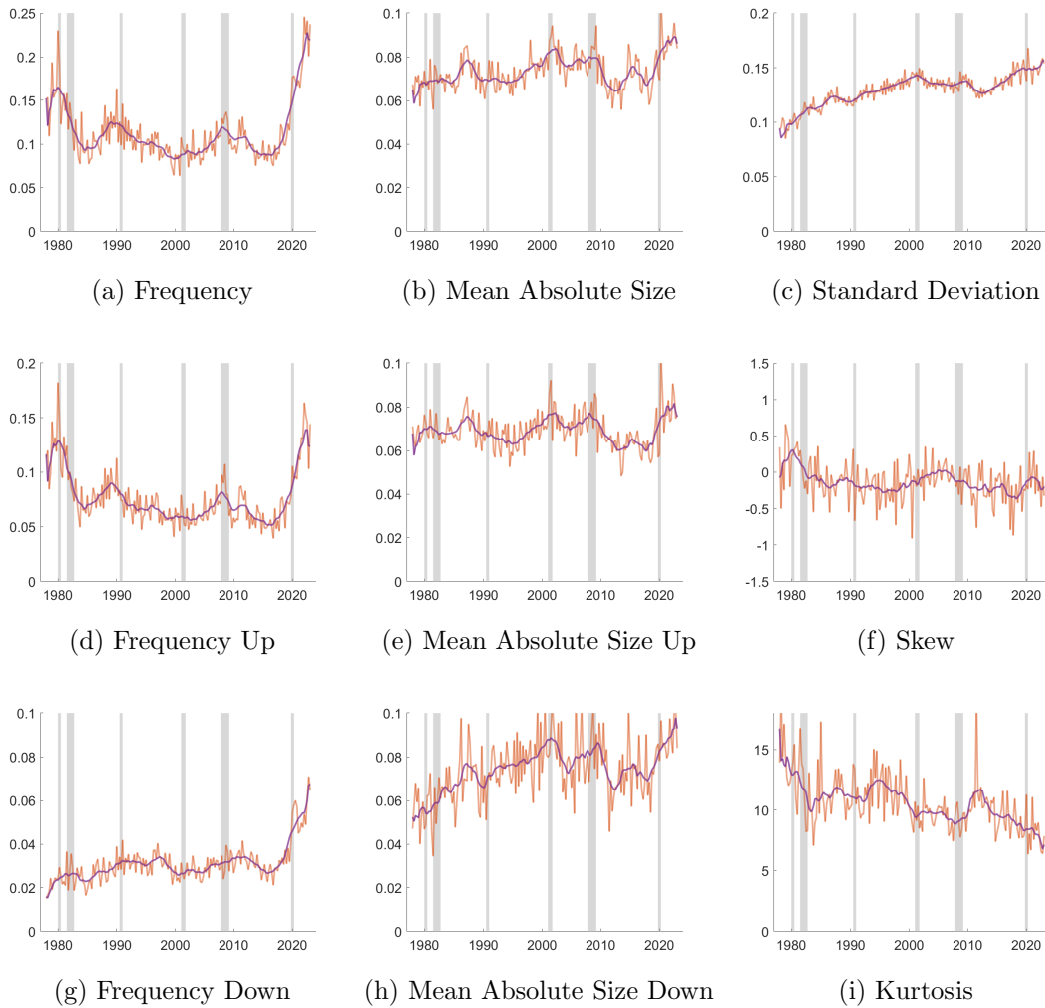


Figure C.1: Data Pricing Moments

Note: These panels plot the raw and MA-smoothed series for the pricing moments based on the U.S. CPI data from 1978 to 2023.Q1. Shaded areas represent NBER recession dates.