Finance and Economics Discussion Series

Federal Reserve Board, Washington, D.C.
ISSN 1936-2854 (Print)
ISSN 2767-3898 (Online)

How Firms Form Beliefs and the Implications for Inflation

Robert Minton and Hugo Monnery

2026-053

Please cite this paper as:
Minton, Robert, and Hugo Monnery (2026). "How Firms Form Beliefs and the Implications for Inflation," Finance and Economics Discussion Series 2026-053. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2026.053.

How Firms Form Beliefs and the Implications for Inflation

Robert Minton* and Hugo Monnery†

July 1, 2026

Abstract

Using survey data from U.S. firms, we study the primitive beliefs for price-setting: firms' forecasts of their own marginal costs. These forecasts are disconnected from CPI expectations, (over)react to current and past costs systematically, and underreact to aggregate shocks until costs move. We show that under empirically realistic cost beliefs the New Keynesian Phillips curve is steeper and less forward-looking. Supply shocks are more inflationary because they hit costs quickly. Demand shocks are less inflationary because firms fail to anticipate future wage pressure. Forward guidance weakens at long horizons but strengthens in the near term.

Firms' forecasts of their own marginal costs are central to all models of price setting in the presence of adjustment frictions. These forecasts are therefore a key link in the propagation of macroeconomic shocks into inflation. Despite an extensive literature on firms' CPI expectations, we do not know whether firms forecast their *own* costs optimally, and if not, what systematic mistakes they make. An empirically grounded model of these beliefs is key to a microfounded, behavioral model of inflation, which can be used to understand business cycles and inform monetary policy.

To develop such a model, we use microdata from the Atlanta Fed's Business Inflation Expectations survey. This survey has a relatively long time dimension, with some firms reporting their realized and expected cost growth each month since 2011. Additionally, the survey occasionally elicits CPI beliefs as well, allowing us to test whether CPI beliefs (which appear central in the standard New Keynesian Phillips curve) behave similarly to own cost beliefs, the primitive expectation in the New Keynesian price setting problem. We leverage this rare confluence of reports in survey data to establish five novel facts about how firms form own-cost expectations. Our empirical approach exploits both the cross-sectional and panel dimensions of the data. We use the

*Federal Reserve Board. Email: robert.j.minton@frb.gov. The views in this paper do not represent those of the Board of Governors of the Federal Reserve System or the Federal Reserve System.

†Harvard University. Email: hmonnery@g.harvard.edu.

This paper was Hugo's job market paper; we thank his advisors: Gabriel Chodorow-Reich, Xavier Gabaix, and Ludwig Straub. We also thank Andrew Figura, Tilman Graff, Helene Hall, Laura Nicolae, and Jieying Zhang. We thank Hassan Afrouzi for a great discussion of the paper. The paper has benefited greatly from workshop and seminar participants at Harvard, the Federal Reserve Board, UCLA Anderson, the Federal Reserve Bank of Atlanta, the 2025 Behavioral Economics Annual Meeting, the University of Oxford, the Federal Reserve Bank of Boston, the Spring 2026 Monetary Economics NBER meeting, the 2026 Midwest Macroeconomics Meeting, and the 2026 CMT Conference on Expectations and Behavior. We are grateful to the Federal Reserve Bank of Atlanta—and, in particular, Brent Meyer and Kevin Foster—for sharing the Business Inflation Expectations Survey microdata with us. They are the source and owner of this data. Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the Federal Reserve Bank of Atlanta.

cross-section to determine how information treatments induce different movements in firms' CPI and own-cost forecasts. We use the panel to compute own-cost forecast errors at the firm level and relate them to recent cost growth, its lags, and its idiosyncratic components, additionally conducting heterogeneity analysis using firm demographics. We estimate firms' heterogeneous cost exposures to aggregate shocks and study whether firms' beliefs react as a function of how quickly shocks affect their costs. The five facts we establish are summarized below.

Fact 1 Own-cost and CPI forecasts are disconnected. Providing firms with a professional CPI forecast in an information treatment, we find revisions in CPI beliefs do not cause meaningful revisions in own-cost beliefs.

Fact 2 Firms' own-cost beliefs react—and indeed overreact—to their own past costs, even after isolating idiosyncratic cost variation. Because firms' forecast errors are predictable from information provided by firms, this reflects a failure of rational expectations rather than a lack of information.

Fact 3 The overreaction of own-cost forecasts to recent cost growth is ubiquitous, holding across sectors, firm size, and time.

Fact 4 Firms' own-cost beliefs exhibit substantial inertia. Forecasts depend not only on cost growth from the current month, but also from a year ago.

Fact 5 Firms' own-cost beliefs react to oil prices—a salient public signal—primarily when their own costs are affected. The beliefs underreact to oil prices before costs move.

Motivated by these facts, we present and calibrate a simple model that disciplines the adaptive component of inflation (Friedman (1968), Phelps (1967)) by directly estimating how firms extrapolate from past costs. Our model nests full-information rational expectations (FIRE), while also allowing firms to learn gradually from their own past costs. We call this latter feature "Adaptive Learning" (AL). A key distinction of our model is that firms learn about the myriad macroeconomic shocks via their own costs, rather than by observing signals directly about the shocks themselves. We show that, in contrast to other models in the literature that focus on learning from shocks, our model can match the five facts established above.

We then formally calibrate our belief model, taking advantage of the panel structure of the BIE survey data to use estimated firm-specific loadings on macroeconomic shocks and signals—average cost growth in the BIE, CPI inflation, and monetary policy shocks. The calibration provides an equation characterizing firms' own-cost forecasts that can be directly entered into existing models. We find that firms place a positive but small weight on the FIRE forecast: firms do learn from information outside their costs, but the extent of such learning is limited. Firms place significant weight on Adaptive Learning, believing nominal marginal cost *growth* to be highly persistent (0.82 quarterly). And firms are backward looking, relying on past cost growth at the decaying rate 0.64.

We then attempt to explain why firms are too responsive to their own costs and insufficiently responsive to other news. We find that a lack of information about own costs or aggregate dynamics cannot explain firms' errors by themselves. However, failures of rational expectations, such as slow learning about cost persistence and confusion about exposures to macroeconomic shocks, can. First, we show that the overreaction to past costs may result from firms using an average cost persistence over a much longer horizon than is covered by our survey data, which starts in 2011. Aggregate data on the GDP deflator shows that costs were much more persistent in earlier decades, and we find that if firms learn very slowly, we can rationalize the overreaction observed in our sample.[1] Second, we show that the historical pass-through of CPI to average costs, at 0.68, is four times the 0.15 pass-through of CPI to own-cost beliefs found in **Fact 1**. This independently corroborates our calibration's estimate of a small weight on the FIRE forecast and also speaks to the mechanism driving underreaction: firms do not merely lack public information; even when informed about aggregates, they fail to understand the implications for their own costs.

We next explore the implications of Adaptive Learning for pricing decisions.[2] Our model predicts very little price response in anticipation of cost movements but high pass-through once costs move. This differs from many existing models in the literature. For example, sticky information models (such as Mankiw and Reis (2002)) predict limited anticipation, but they also predict prices will react less even when costs move. Diagnostic expectations (see Bordalo, Gennaioli and Shleifer (2018)) generate a stronger price pass-through than under FIRE, but only when news of the shock is revealed, rather than when costs start moving. Angeletos, Huo and Sastry (2021) allow for initial underreaction and delayed overreaction but still feature learning from shocks directly, so the reaction profile is not a function of costs. We show that in our model, sufficiently persistent cost shocks are underreacted to, while sufficiently transitory cost shocks are overreacted to.

We extend these results to the New Keynesian Phillips curve (NKPC), which we derive under our model of beliefs. In the special case of fully Adaptive Learning, our NKPC becomes the Friedman-Phelps Adaptive Expectations Phillips curve, where inflation depends on real marginal costs (or the output gap) today and the history of inflation. However, in our case this Phillips curve is microfounded from the firm-level pricing problem and our model of firms' beliefs. The microfoundation allows us to understand the parameters driving the shape of the Phillips curve and to use microdata to help calibrate them.

In particular, our calibration predicts a steeper and less forward-looking NKPC than under FIRE. We also show that firms' average beliefs about nominal marginal cost growth enter like a cost-push shock, driving up inflation for a given output gap. However, despite these average nominal marginal cost beliefs themselves being inertial (**Fact 4**), our calibrated NKPC features little inertia: there are limited effects on inflation in the future from one-off shocks to real marginal

[1] This explanation would suggest that over a much longer sample, or over other time periods, we may not find unconditional overreaction, as in **Fact 2**. However, many of the qualitative implications of Adaptive Learning beliefs remain even when firms extrapolate correctly on average from their costs; see Appendix D.5.

[2] We do so in the benchmark New Keynesian model with Calvo pricing and CES demand, without strategic complementarities. This paper will not discuss how firms gather information or form beliefs about competitors' prices. We leave these interesting questions for future research.

costs today. This is because real marginal cost or output gap shocks do not cause persistent *nominal* marginal cost *growth*.[3] We show that the degree of inertia in the Phillips curve is increasing in price *flexibility*, via Adaptive Learning behavior: as prices become more flexible, real marginal cost shocks induce larger long run increases in the price level (i.e. more cumulative nominal marginal cost growth), and the cost-push term in our NKPC becomes operative. Through this purely behavioral channel, shocks that generate substantial aggregate inflation, and thereby increase the frequency of price adjustment, would boost their own persistence.

We test our predictions for the NKPC by estimating an Adaptive Expectations Phillips curve using the state-level inflation data and methodology from Hazell, Herreño, Nakamura and Steinsson (2022). We find a steep slope and limited inertia; moreover, what little inertia we do find is driven by the pre-1990 period, when the frequency of price adjustment was higher.

To further explore shock propagation, we set our Phillips curve in a textbook New Keynesian model with sticky prices and wages. We show that supply shocks are more inflationary, while demand shocks are less inflationary.[4] The key feature driving this result is that (in the presence of sticky wages) demand shocks are in large part news about future cost movements, whereas supply shocks affect costs more quickly. This asymmetric reaction to supply and demand shocks also holds in a medium-scale New Keynesian model in which these shocks are estimated in the data. And the core intuition holds in a wide range of economic settings: upstream shocks are made less inflationary than downstream shocks; and the inflationary response to pure news shocks is muted. This asymmetric prediction—with some shocks amplified and others dampened—is driven *essentially* by the fact that firms rely on their own costs when forecasting.[5] We also show the qualitative prediction remains in a model closer to limited-information rational expectations (LIRE), where firms do not overreact unconditionally to costs.[6]

We then analyze optimal monetary policy. Forward guidance at long horizons is weaker in our model, as it is difficult to move firms' beliefs without first moving their costs. However, near-term forward guidance is more powerful. This is because firms react more strongly to current costs, which still respond to news about future rates via forward looking households.[7] Next, we turn to the optimal policy problem under commitment and discretion. We find a very limited role for commitment: it is not effective for the central bank to "tie their hands". Under FIRE, the ability to commit to future interest rates enables a smooth policy response to inflationary shocks; under Adaptive Learning, the central bank needs to react rapidly.

Finally, we ask whether the central bank has an incentive to act in anticipation of future infla-

[3] This is true even in the standard New Keynesian model under FIRE. A transitory real marginal cost shock first raises but then lowers nominal marginal costs. The rise is larger than the decline—and more so when prices are more flexible. However, the decline is more recent and therefore plays a more prominent role in our Adaptive Learning model.

[4] Kohler, L'Huillier, Phelan and Weiß (2025) also argue for higher inflationary pass-through from supply shocks. They consider fully informed, rational firms interacting strategically with uninformed consumers.

[5] For example, Angeletos et al. (2021) predict a different NKPC for each shock and therefore do not generate the asymmetry result.

[6] In a model of LIRE in which firms observe only their own costs, firms will initially underreact to the most persistent shocks, and overreact to the least persistent. However, in this model, beliefs must be correct on average, in contrast to **Fact 2** above. Importantly, this LIRE model has only aggregate supply and demand shocks. If the LIRE model additionally had idiosyncratic shocks that dominate the variance in firms' costs, firms might underreact to both supply and demand shocks, in contrast to our model.

[7] In fact, under Adaptive Learning, household beliefs are more important as they act as a mediator between news about the future and firms' current costs, which are passed through strongly into inflation.

tion. In the New Keynesian model under FIRE it does not, since a one-off monetary policy shock today does not have lasting effects.[8] Under Adaptive Learning, early action by the central bank can be optimal, but only if the Phillips curve is inertial. As such, while we find little support for the benefits of such early interventions on average in the United States in recent decades, we show that there are conditions under which monetary policy would have these lasting effects: for example, when the frequency of price adjustment rises, as occurred during Covid.

Related literature. Our paper contributes to a now substantial literature characterizing and explaining the empirical forecasts of economic agents. This literature had a strong initial focus on the Survey of Professional Forecasters (SPF). Coibion and Gorodnichenko (2015) showed underreaction of average beliefs, in accordance with noisy information models, while Bordalo, Gennaioli, Ma and Shleifer (2020) found individual-level overreaction, suggesting a failure of rational expectations. Broer and Kohlhas (2022) found the same result of overreaction in forecast revisions; but they additionally show that, while some public information is overreacted to, other public information is underreacted to. Angeletos et al. (2021) reconcile these findings in a model with deviations from both rational expectations and full-information. They find an additional prediction, borne out in the data, of delayed overshooting: average beliefs initially underreact and then overreact. Compared to this literature, we focus on firms' beliefs. But we too find evidence for both overreaction and underreaction, as well as delayed overshooting. The key difference is that we predict that these features depend on the time profile of costs. For example, if costs increase gradually, there will be initial underreaction; but if costs peak on impact and then sharply decline, there will be initial overreaction.

The literature studying firms' beliefs has focused most on CPI forecasts. Kumar, Afrouzi, Coibion and Gorodnichenko (2015) find that firms have large disagreement and volatile beliefs about future CPI and systematically misreport even past CPI inflation.[9] Information treatments have also been used to test the causal effect of inflation expectations on firm decisions, with mixed findings.[10] There has also been significant interest in the higher order moments of firm beliefs.[11] [12]

In contrast to this literature, we focus on firms' beliefs about their own nominal marginal costs, the primitive object for setting prices in the New Keynesian model. Meyer and Sheng (2025) present evidence from a special question in the BIE that firms report unit costs as substantially more important for pricing decisions than aggregate inflation, and Andrade, Dietrich, Leer, Schoenle, Tang and Zakrajsek (2026) find survey evidence that firms' expected future costs pass through to their current pricing decisions.[13] While one might have thought that aggregate in-

[8]The standard New Keynesian model with sticky prices but flexible wages, and without capital, is a purely forward-looking model.

[9]See also Weber, D'Acunto, Gorodnichenko and Coibion (2022) and Candia, Coibion and Gorodnichenko (2023).

[10]See Coibion, Gorodnichenko and Ropele (2020), Coibion, Gorodnichenko and Kumar (2018), and Savignac, Gautier, Gorodnichenko and Coibion (2024).

[11]See Coibion, Gorodnichenko, Kumar and Ryngaert (2021) and Kumar, Gorodnichenko and Coibion (2023).

[12]There are also some papers looking at other firm outcomes; for example, Bloom, Codreanu and Fletcher (2025), Ma, Ropele, Sraer and Thesmar (2020), and Bellifemine, Couturier and Tozzo (2025) find systematic mistakes when firms forecast sales growth.

[13]De Fiore, Lombardi and Mangiante (2025) argue that unexpected changes in input costs pass through into firms' actual and

flation beliefs capture the persistent component relevant for unit cost beliefs, Coibion et al. (2018) document that firms' unit cost beliefs and aggregate inflation beliefs appear only weakly correlated. Our information treatment shows that changes in CPI beliefs pass through very little to unit cost beliefs. Baumann, Ferrando, Georgarakos, Gorodnichenko and Reinelt (2024) find a similarly small pass-through, except to the average wages of current employees.[14] Further, we document that firms' past costs drive both their CPI and unit cost beliefs but in different ways.[15]

We also contribute to a growing literature focusing on how households and firms use local information when forming their beliefs. Taubinsky, Butera, Saccarola and Lian (2024) show that household beliefs are too responsive to household level shocks, violating rational expectations. D'Acunto, Malmendier, Ospina and Weber (2021) find consumers use the price of groceries they purchase when forecasting aggregate inflation. Turning to firms, Andrade, Coibion, Gautier and Gorodnichenko (2022) present evidence that firms' beliefs about CPI are correlated with sector prices. Born, Enders, Menkhoff, Müller and Niemann (2022) argue that expectations about the aggregate economy overreact to micro news and underreact to macro news. Baumann et al. (2024) show that euro area firms extrapolate from local inflation conditions – regional or national – to euro area inflation. Dovern, Müller and Wohlrabe (2023) show German firms extrapolate from local conditions to expected aggregate growth rates. In contrast to these papers, our paper establishes that firms' own costs are a key driver of beliefs. The closest paper to ours empirically is Albagli, Grigoli and Luttini (2022). The authors show that Chilean firms extrapolate from their own cost growth to expectations about aggregate CPI inflation. We differ from this paper by focusing on own-cost forecasts and exploring the macroeconomic implications of these beliefs.

There exist a wide array of models of belief formation in macroeconomics, and we briefly highlight those most relevant to our paper. Mankiw and Reis (2002) showed how informational frictions in price setting can result in a Phillips curve. Many more recent papers explore the role of information frictions in pricing and inflation.[16] Gabaix (2020) spells out the implications of one particular form of myopia – cognitive discounting – in the New Keynesian model.[17] An important implication is that the forward guidance puzzle is resolved. Our empirical model of firms' beliefs also dampens forward guidance significantly, but not at short horizons. Bordalo et al. (2018) show how diagnostic expectations generate extrapolation and overreaction seen in financial markets. As discussed above, Angeletos et al. (2021) present a model with failures of both full information and rational expectations to explain mixed findings of under- and overreaction. We compare the predictions of our behavioral theory to those from these alternative models and demonstrate that

expected future prices.

[14] Baumann et al. (2024) find a pass-through of 0.15 to "The average wage of your current employees", though a higher pass-through of 0.67 to "The average prices of production inputs (non-labor costs such as materials and energy)". It is not straightforward to compute the pass-through to marginal costs from these components.

[15] The BIE elicits multiple firm beliefs, allowing for rich characterizations. For example, Meyer, Prescott and Sheng (2022) use the BIE to argue that, in the early part of the COVID period, firms saw the pandemic mostly as a negative demand shock, lowered their CPI expectations, but had stable long run beliefs.

[16] E.g., see Afrouzi and Yang (2021), La'O and Tahbaz-Salehi (2022), and Gagliardone and Tielens (2026). Afrouzi (2023) argues that firms with fewer competitors should pay less attention to macroeconomic conditions and confirms this in the data.

[17] Angeletos and Huo (2021) show that information frictions have the effect of making agents less forward-looking (myopia) and more backward-looking (anchoring).

none of these alternatives are jointly consistent with all five facts of firms' own-cost beliefs.

In our model of Adaptive Learning, firms forecast the future using past observations and a (potentially) misspecified model. This is in the spirit of the substantial adaptive learning literature à la Evans and Honkapohja (2001). For a recent overview, see Eusepi and Preston (2018).[18] Our paper is unique in focusing on the primitive expectation of interest for inflation—firms' forecasts of their own costs—and disciplining our belief model using survey data on these forecasts.

1 How do firms forecast their costs?

In this section, we introduce our data on firms' own-cost forecasts and present evidence that (1) firms' own-cost beliefs are disconnected from their CPI beliefs, (2) firms' beliefs (over-)react to their past costs, (3) firms' beliefs depend on own costs in a stable way across sectors, firm size, and time, (4) firms learn from their costs slowly over time, and (5) firms underreact to oil prices until their costs move.

1.1 The Business Inflation Expectations Survey

The Federal Reserve Bank of Atlanta sends the Business Inflation Expectations Survey (hereafter, BIE) each month to executives and managers in the Sixth Federal Reserve District, which covers Florida, Alabama, Georgia, and parts of Tennessee, Louisiana, and Mississippi. Meyer and Sheng (2025) discuss how the sample selection criteria were chosen to make the sample as nationally representative as possible. The survey consists of core questions asked every month or quarter and special questions asked discretionarily. Our sample includes all participants' responses to core questions from October 2011 to September 2024 and to a subset of special questions pertaining to variables of interest in our study. Our sample also classifies each firm into 14 sectors and 9 employee count categories. In a typical month, around 200 firms respond to the survey, and the average duration in the survey for a firm is 30 periods.[19]

The key variables of interest in our study are firms' realized and expected nominal marginal cost growth. As part of the monthly core questions, firms are asked the percent change in unit costs compared with the same month last year.[20] The possible responses are captured by the set $\{-2\%, 0\%, 2\%, 4\%, 6\%\}$. Designating firm i's log nominal marginal cost as mc_{it}, we therefore measure

$$\Delta_{12}\mathrm{mc}_{it} \equiv \mathrm{mc}_{it} - \mathrm{mc}_{i,t-12}. \tag{1}$$

[18]For example, Eusepi and Preston (2011) embeds household learning in an RBC model, and Slobodyan and Wouters (2012) finds improved empirical performance in a medium scale DSGE model where beliefs are formed under learning.

[19]However, there is a long right tail, with over 100 firms responding for more than 84 months.

[20]As discussed in Meyer and Sheng (2025), and following Blinder (1994), firms largely do not use the term marginal cost. The BIE targets eliciting marginal costs and settled on the term "unit costs" after cognitive interviews with C-suite executives. We show in Appendix A.10 that reported unit costs exhibit statistically significant but small decreasing returns to scale; if fixed costs were a large component of unit costs, we would expect unit costs to exhibit increasing returns to scale, as fixed costs get averaged over more units when sales expand. Finally, Appendix A.10 also shows that, even with non-constant returns to scale, marginal cost dynamics can still mirror average cost dynamics.

Firms also report their beliefs for unit cost growth over the next twelve months by assigning probabilities to the same five scenarios: $\{-2\%, 0\%, 2\%, 4\%, 6\%\}$. The survey question, as firms see it, is shown in Appendix A.1. Firms' responses appear to be thoughtful and do not fall into simple patterns. For example, no single cost forecast distribution makes up even 4% of total responses, suggesting no strong default response across all firms. Nor does there seem to be a strong default response even within firms: in only 16% of cases do firms report the same cost forecast distribution in consecutive responses. Furthermore, firms do not simply forecast exactly their past costs; in only 7% of cases do firms put a 100% probability on the past cost growth they just reported.

Taking an expectation over the reported cost forecast distribution, we construct firm i's expected nominal marginal cost growth over the next year

$$\mathbb{E}_{it}\left[\Delta_{12}\mathrm{mc}_{i,t+12}\right] \equiv \mathbb{E}_{it}\left[\mathrm{mc}_{i,t+12}\right] - \mathrm{mc}_{it}. \tag{2}$$

As an external check on firms' responses, and as previously argued in Meyer and Sheng (2025), we show in Appendix A.3 that averages of BIE responses correlate strongly with published aggregate inflation statistics and expectations. We find a strong fit between BIE costs and the GDP deflator, except in 2022 when inflation rose above the 6% upper bound.[21] Average expectations comove closely with the one-year ahead SPF CPI forecast.

Beyond arguing that firms report their cost expectations seriously, we further show that firms perceive exactly these cost expectations as highly relevant to their pricing decisions. To do so, we use Special Questions in the BIE concerning firms' pricing, asked at much lower frequency. In Appendix A.4, we show that (controlling for past prices and past costs) firms that expect high cost growth expect to raise their prices.

Exploiting the panel structure of the data, we can also compute the forecast error, using a forecast in period t about period $t+12$ and a reported realized outcome in period $t+12$

$$\begin{aligned}\mathrm{FE}_{i,t,t+12} &\equiv \mathrm{mc}_{i,t+12} - \mathbb{E}_{it}[\mathrm{mc}_{i,t+12}] \\ &= \Delta_{12}\mathrm{mc}_{i,t+12} - \mathbb{E}_{it}[\Delta_{12}\mathrm{mc}_{i,t+12}].\end{aligned} \tag{3}$$

In extensions to our core results, we make use of three further questions in the BIE survey. Each quarter, firms are asked to forecast their annualized cost growth five to ten years in the future. Firms also report whether their unit sales are above or below normal. And in certain special questions, firms forecast CPI growth over the next year.

[21] In Appendix A.5 we show that the proportion of past cost reports at the 6% upper limit rises to 50% of firms at the peak. This still implies that half of firms are unconstrained, leaving some variation, and we show below that our results are robust to dropping reports at the boundaries. In other periods, this proportion lies below 10%. We also show that our results hold in the pre-Covid period where the boundary is rarely binding.

1.2 Five facts about firms' cost expectations

In the workhorse New Keynesian model, when firms set prices the central driver is their own expected future nominal marginal costs

$$p^*_{i,t} = \left(1 - \beta\theta_p\right) \sum_{\tau=0}^{\infty} \left(\beta\theta_p\right)^{\tau} \mathbb{E}_{i,t}\left[\mathrm{mc}_{i,t+\tau}\right] \tag{4}$$

Under full-information rational expectations (FIRE), we can derive from (4) the standard New Keynesian Phillips curve,

$$\pi_t = \kappa^{\mathrm{FIRE}} \mathrm{mc}^{\mathrm{real}}_t + \beta \mathbb{E}_t\left[\pi_{t+1}\right] \tag{5}$$

Equation (5) suggests a large role for CPI beliefs in determining inflation today. However, when firms' beliefs deviate from FIRE, (5) no longer generically holds. The apparent importance of CPI expectations disappears.[22] Instead, as (4) makes clear, the primitive expectations are over a firm's own costs.

Nevertheless, CPI forecasts might be an effective *proxy* for own-cost beliefs.[23] Below we show this is not the case: knowing how much a shock leads firms to revise their CPI forecasts is not sufficient to understand how much it leads to a revision of own-cost beliefs. Furthermore, we show that past costs are a key driver of firms' cost expectations.

Fact 1: Firms' CPI and own-cost beliefs are disconnected. To test the causal connection between firms' CPI and own-cost forecasts, we ran a Special Question in the May 2025 BIE survey. This exercise confirms the importance of studying firms' own-cost beliefs, as we find a substantial disconnect between CPI and own-cost beliefs.

We first asked firms to report their expectations for their own cost growth over the next year, $\mathbb{E}^{\mathrm{pre}}_{i,t}\left[\Delta \mathrm{mc}_{i,t+12}\right]$, and their expectation for CPI inflation over the next year, $\mathbb{E}^{\mathrm{pre}}_{i,t}\left[\pi^{\mathrm{CPI}}_{t+12}\right]$. We then told *all* firms the latest Blue Chip forecast (from April 2025) for inflation over the next year.[24] Finally, we asked firms, in light of this, to report their updated forecast for their own costs, $\mathbb{E}^{\mathrm{post}}_{i,t}\left[\Delta \mathrm{mc}_{i,t+12}\right]$, and their updated forecast for CPI, $\mathbb{E}^{\mathrm{post}}_{i,t}\left[\pi^{\mathrm{CPI}}_{t+12}\right]$. The exact wording of the questions is shown in Appendix A.2. The own-cost responses are given by assigning weight to bins (as is standard in the BIE survey, and as discussed in Section 1.1), whereas we elicit point estimates for the CPI forecasts. To prevent a mechanical bias, we ex-post truncate the CPI estimate to the range $[-2\%, +6\%]$, though our results are robust to this truncation assumption.

Denote the change in CPI beliefs, from before to after the information treatment, by

$$d\mathbb{E}_{i,t}\left[\pi^{\mathrm{CPI}}_{t+12}\right] \equiv \mathbb{E}^{\mathrm{post}}_{i,t}\left[\pi^{\mathrm{CPI}}_{t+12}\right] - \mathbb{E}^{\mathrm{pre}}_{i,t}\left[\pi^{\mathrm{CPI}}_{t+12}\right]$$

[22] In a similar spirit, Werning (2022) argues that it would be misleading to read off the causal contribution of inflation expectations from (5), which suggests a coefficient of β, as it generically holds only under FIRE.

[23] For example, this might hold if the only persistent components of costs are, or are perceived to be, aggregate.

[24] We told firms a CPI forecast, rather than a CPI realization, because Baumann et al. (2024) find the forecast information treatment is more powerful for moving firms' CPI forecasts.

Table 1: Effect of a change in CPI beliefs on own cost beliefs

	Change in own-cost forecast, $d\mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right]$		
	Main result	No truncation	Extensive margin
$d\mathbb{E}_{i,t}\Delta_{12}\pi^{\text{CPI}}_{t+12}$	0.15^{*}	0.09^{*}	
	(0.08)	(0.05)	
$\mathbb{1}\left\{\text{Changed CPI forecast}\right\}$			0.27^{***}
			(0.05)
Observations	349	349	349

Note: *p<0.10, **p<0.05, ***p<0.01. Heteroskedasticity-robust standard errors.

and the change in own-cost beliefs by

$$d\mathbb{E}_{i,t}\left[\Delta\text{mc}_{i,t+12}\right] \equiv \mathbb{E}^{\text{post}}_{i,t}\left[\Delta\text{mc}_{i,t+12}\right] - \mathbb{E}^{\text{pre}}_{i,t}\left[\Delta\text{mc}_{i,t+12}\right]$$

We estimate the regression

$$d\mathbb{E}_{i,t}\left[\Delta\text{mc}_{i,t+12}\right] = \text{constant} + \beta^{\text{Experiment}} \cdot d\mathbb{E}_{i,t}\left[\pi^{\text{CPI}}_{t+12}\right] + u_i$$

The results are shown in Table 1. A 1 percentage point increase in the forecast for CPI growth over the next year induces only a 0.15 percentage point increase in the forecast for own cost growth over the next year. Our finding suggests that, just because an information treatment affects a firm's CPI beliefs, it need not meaningfully affect their own-cost beliefs.[25]

We carry out two robustness exercises. First, we show the result where we do not truncate the CPI responses. Second, we show the results for the extensive margin only: we ask how much a firm updating its CPI forecast increases the probability that firm updates its own-cost forecast. We continue to find small effects.

This result is consistent with more reduced form evidence suggesting a disconnect. For example, the unconditional correlation between CPI and own-cost beliefs in the BIE is 0.36.[26] Coibion et al. (2018) find a correlation of just 0.01 in (a single cross-section of) different survey data.

This experiment tells us that to explore the effect of macroeconomic shocks on inflation via the cost expectations channel, it is necessary to study own-cost beliefs. Given the causal disconnect, the response of CPI expectations to a macroeconomic shock does not reveal the response of own-cost beliefs—the true primitive for price setting.[27] We will show in Section 2.4 that the find-

[25] This finding also appears striking from a theoretical point of view. For example, in Appendix B.12 we show that in the canonical Maćkowiak and Wiederholt (2009) model, we would expect a pass-through greater than one.

[26] Here we again truncate the CPI responses to lie in the range $[-2\%, +6\%]$. Without truncation, the correlation is 0.005. Additionally including individual fixed effects, in an attempt to identify the dynamic relationship between these beliefs, we find a correlation of 0.25.

[27] In a model with strategic complementarities in price setting, beliefs beyond a firm's own costs are relevant; but it is still important

Figure 1: Cost Growth is Associated with Expected and Realized Year-ahead Cost Growth

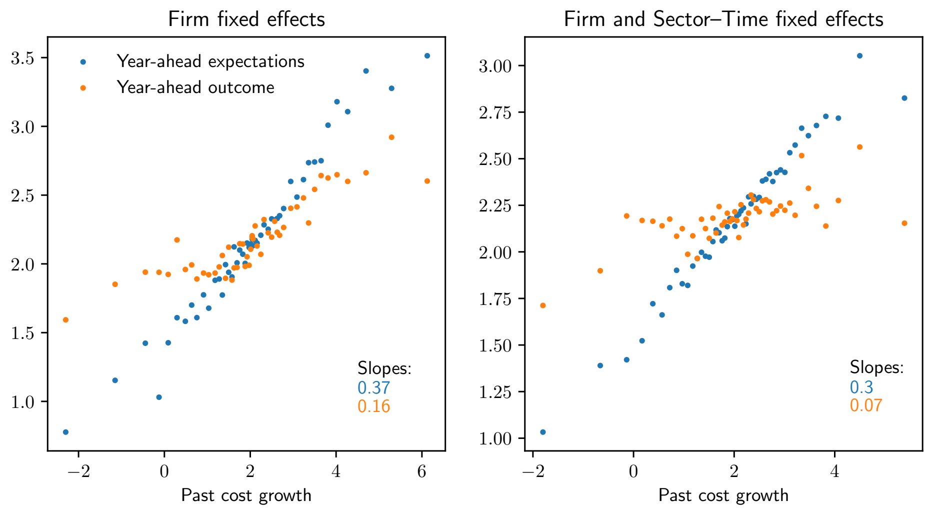

Note: 50 bins, each contain an equal number of observations. Data is residualized on the fixed effects then binned. We add back in the unconditional mean for each variable.

ings from this experiment suggest a degree of unresponsiveness to macroeconomic information consistent with the model of beliefs we develop below.

Fact 2: Firms (over)react to their own costs. Next, we relate firms' cost growth over the last year to their expectations for their cost growth over the next year. We visualize this relationship in a binned scatter plot, shown by the blue dots in Figure 1. There is a strong, positive, and linear relationship between costs and forecasts at the firm level (controlling for firm fixed effects).[28] This relationship persists when we also control for sector-time fixed effects. Firms believe that their nominal cost growth is persistent. In the orange dots, we evaluate this belief by relating cost growth over the last year to the actual cost growth over the next year, reported in a year's time. We find a positive relationship, though a weaker one than for beliefs: firms believe cost growth is more persistent than it truly is.

While binscatters speak to the functional form of the conditional expectation function, they do not reveal whether costs explain a small or large share of variation in beliefs. As one approach to test this, we take the R-squared from a regression of residualized year-ahead expectations onto residualized cost growth over the last year. We find an R-squared of 28% controlling only for firm fixed effects, and 20% controlling for both firm and time fixed effects. This suggests contemporaneous cost growth has significant explanatory power for beliefs, particularly since the measurement error due to binned survey responses implies the maximal R-squared from fundamentals is less

to understand how much own-cost beliefs are moved as distinct from beliefs about competitors' prices.

[28] In all regressions in this section we include firm fixed effects. As such our results can be thought of as capturing learning around firm-specific steady state cost inflation. In Appendix A.6 we show that firms appear to hold fairly accurate beliefs on average over time. Cattaneo, Crump, Farrell and Feng (2024) suggest jointly computing binned means and controlling for fixed effects, rather than residualizing first. We do not do so here because we want to test for nonlinearity in deviations from each firm's steady state. However, in Appendix A.7 we show that our conclusions hold when following their procedure too.

than 100%.[29]

We next turn to linear regressions to quantify the significance of these relationships and their robustness. Our specifications are

$$\mathbb{E}_{it}[\Delta_{12}\mathrm{mc}_{i,t+12}] = \alpha_i + \mathrm{Controls}_{it} + \beta_{\mathrm{expectations}}\Delta_{12}\mathrm{mc}_{it} + u_{i,t}^{\mathrm{expectations}} \tag{6}$$

$$\Delta_{12}\mathrm{mc}_{i,t+12} = \alpha_i + \mathrm{Controls}_{it} + \beta_{\mathrm{outcome}}\Delta_{12}\mathrm{mc}_{it} + u_{i,t}^{\mathrm{outcome}} \tag{7}$$

$$\mathrm{FE}_{i,t,t+12} = \alpha_i + \mathrm{Controls}_{it} + \beta_{\mathrm{FE}}\Delta_{12}\mathrm{mc}_{it} + u_{i,t}^{\mathrm{FE}} \tag{8}$$

where α_i is a firm fixed effect and the u_{it} are error terms. Table 2 shows the OLS estimates of $\beta_{\mathrm{expectations}}$, β_{outcome}, and β_{FE} under four specifications for controls: no controls, a time fixed effect, a sector-time fixed effect, and the first four principal components of cost growth (additional details in Appendix A.8). Standard errors are clustered at the firm level.

Table 2: Firms (over) react to their own costs

	Outcomes		
Controls	$\mathbb{E}_{i,t}\left[\Delta_{12}\mathrm{mc}_{i,t+12}\right]$	$\Delta_{12}\mathrm{mc}_{i,t+12}$	$\mathrm{FE}_{i,t,t+12}$
None	0.36^{***}	0.16^{***}	-0.20^{***}
	(0.02)	(0.02)	(0.02)
Time FE	0.30^{***}	0.09^{***}	-0.21^{***}
	(0.02)	(0.02)	(0.03)
Sector Time FE	0.29^{***}	0.06^{***}	-0.22^{***}
	(0.02)	(0.02)	(0.02)
PCA factors	0.23^{***}	-0.08^{***}	-0.31^{***}
	(0.02)	(0.02)	(0.02)

Note: *<0.10, **<0.05, ***<0.01. Standard errors, in brackets, are clustered at the firm level. The sample is consistent across regressions, and contains 19,785 observations.

The first column in Table 2 shows the regression coefficient for expectations of year-ahead cost growth on cost growth over the last year. Firms extrapolate from their costs, even accounting for the first four principal components of cost growth, which together explain 71% of cost growth: firms' cost forecasts depend on idiosyncratic cost movements. These coefficients say that a firm whose costs typically rise by 2% per year, but whose costs rose 3% over the last year, would increase their forecast from 2% to about 2.3%.

The second column shows the regression coefficient for realized year-ahead cost growth. As in the binned scatter plot, we find evidence of some persistence. However, as we isolate more id-

[29] By comparison, the R-squared of own-cost beliefs on year-on-year CPI growth (and not own costs), and controlling for firm fixed effects, is 13%.

iosyncratic variation, past cost growth predicts a decline in future cost growth, suggesting slight mean reversion. We see that the perceived persistence of costs is consistently above the true persistence. Under full-information rational expectations (FIRE), no information available when the forecast was made should be able to predict the forecast error. But more than this, since cost growth is reported to us by the firms, and therefore is in the firms' information sets, any finding that β_{FE} differs from zero is inconsistent with rational expectations. The third column formalizes that firms overreact to cost growth.

Overall, these two facts suggest that a model of firms' cost beliefs must feature a role for learning from own costs and that the model should not impose rational expectations. We now turn to a series of robustness and extension exercises.

In Appendix A.9, we show that, as in Albagli et al. (2022), firms use their past costs to forecast CPI growth over the next year. Nevertheless, in Appendix A.11, we show that firms still overreact to their own cost growth conditional on their CPI beliefs, reinforcing the finding in **Fact 1** that firms forecast own-cost beliefs differently from CPI beliefs.

In addition, we show that firms use their past costs to forecast their cost growth five to ten years ahead. As Table 2 shows low cost growth persistence even at a one-year horizon, this extrapolation to five to ten years ahead is likely further evidence of overreaction.[30]

In Appendix A.12, we show that our results still hold when we instrument cost growth with its lag, which accounts for the possibility of classical measurement error or measurement error that is correlated across outcome variables.[31] In Appendix A.12, we also show that our results are robust to the truncation in our data by removing all observations equal to the boundary values of -2% or 6%.[32] Firms could also use information on sales to make inferences about future marginal costs—particularly when returns to scale are not constant. In Appendix A.11, we include two measures of sales and a measure of profits from the BIE survey and continue to find our result. Appendix A.13 shows that our results do not meaningfully change if, rather than including firm fixed effects, we first residualize with respect to a variable's firm-level average prior to and including period t. Intuitively, this suggests a regression that each firm could carry out themselves at any date t, and from which we obtain predictable forecast errors, confirming our claim of a failure of rational expectations.

Fact 3: Firms' beliefs depend on their own costs in a stable way. Does the dependence of firms' forecasts on their own costs vary with observables? We can use these results to understand why firms extrapolate from past costs and to guide our modeling. For example, if firms only overreact to their costs in the post-Covid period, it would suggest the overreaction was state contingent. Instead, we find remarkable stability in how firms use their costs to inform their beliefs.

[30] In Appendix B.9, we show there is overreaction in both the short-run and the long-run component.

[31] Note that measurement error poses an even larger issue for regressions of forecast errors on forecast revisions—a common regression run in this literature—than it does for us here, since those regressions contain the current forecast on the left-hand side (negatively) and on the right-hand side (positively). This induces a mechanical negative correlation that suggests overreaction.

[32] Recall also that FIRE requires firms to hold correct beliefs about the entire distribution of future costs, so truncation and binning cannot explain away the deviations from FIRE we found in any case. Firms are asked to assign probabilities and report realizations of the same variable over the same intervals.

Figure 2: Heterogeneity analysis

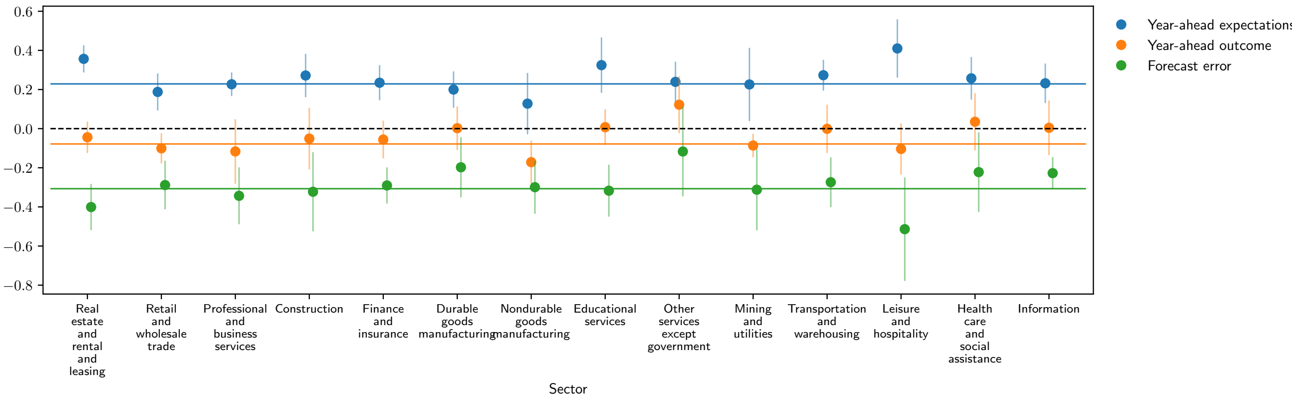

(a) Comparison across sectors

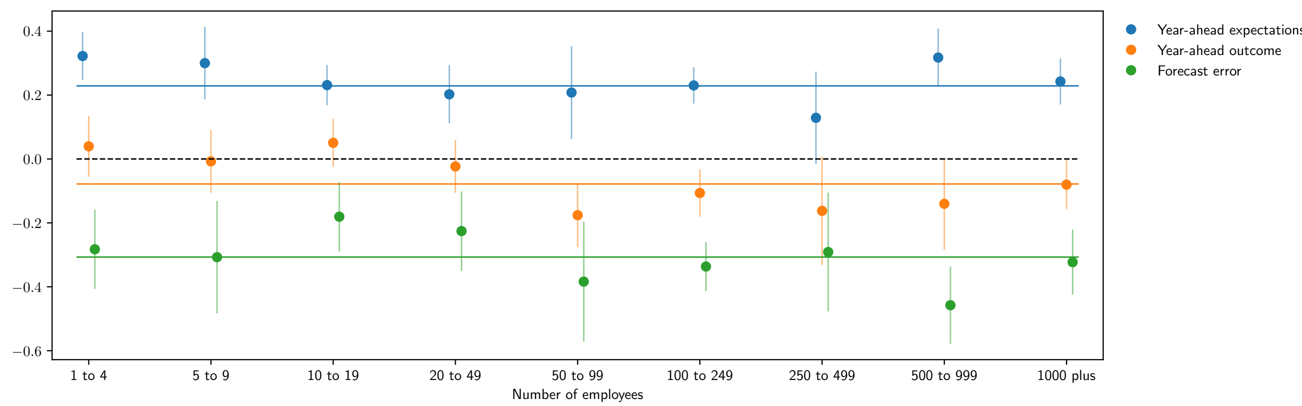

(b) Comparison across firm size distribution

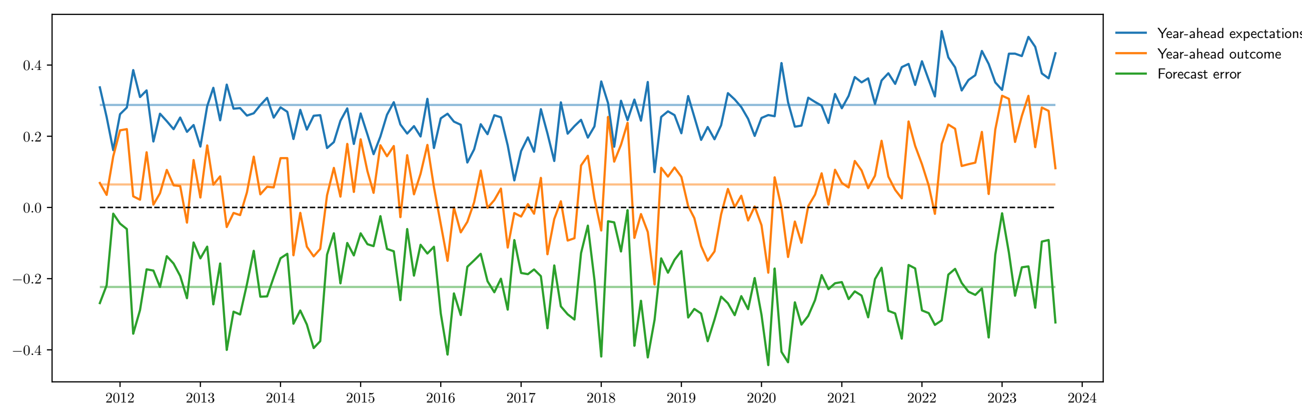

(c) Comparison across time

Note: The vertical bars in Figures (a) and (b) are 95% confidence intervals based on standard errors clustered at the firm level. The horizontal lines in all cases are the coefficients from the full sample regression. We include PCA controls in (a) and (b), and sector-time fixed effects in (c).

In Figure 2, we explore our results across three different cuts of the data: across the 14 sectors, across the nine size categories, and across the entire time series of the BIE data. We uniformly find that beliefs are positively related to costs, and that these costs are overreacted to. The point estimates are also generally close to the result found in the whole sample. This suggests that the way in which firms learn from their own costs is fairly stable, and that our (over)reaction finding is first order.

Another margin along which we might expect differential belief formation is frequency of price adjustment. Firms that adjust their prices daily have much less incentive to accurately forecast their costs a year ahead than firms that adjust prices annually. However, in Appendix A.14, we show that there is little variation in overreaction to own costs along this margin.

While **Fact 2** showed overreaction to costs on average, Appendix A.15 argues that learning from costs does depend in part on a firm's own cost process: firms with more persistent cost growth overreact less. If firms' costs are not persistent, they overreact even more than in our baseline result; but firms may happen to extrapolate correctly if their costs are persistent enough. Our finding mirrors the result of Afrouzi, Kwon, Landier, Ma and Thesmar (2023), who find in an experiment that individuals overreact more for less persistent processes.[33]

Fact 4: Firms learn from their costs slowly over time. So far we have only related cost forecasts to cost growth over the last *year*. However, this does not tell us whether firms use only the latest month's cost growth, or cost growth from over the entire year, when forming these beliefs.

The year-on-year structure of the survey necessitates some assumption to uncover the monthly dynamics of learning. We assume that beliefs can be written as some weighted sum of past monthly cost growth[34]

$$\mathbb{E}_{it}[\Delta_{12}\text{mc}_{i,t+12}] = \alpha_i + \text{Controls}_{it} + \sum_{\ell=0}^{L} \beta_\ell^{\text{expectations}} \Delta\text{mc}_{i,t-\ell} + u_{i,t} \tag{9}$$

We cannot run the regression equation (9) since we do not actually observe month-on-month cost changes. However, we can transform (9) into a regression on observables

$$\mathbb{E}_{it}[\Delta_{12}\text{mc}_{i,t+12}] - \mathbb{E}_{i,t-12}[\Delta_{12}\text{mc}_{i,t}] = \Delta_{12}\text{Controls}_{it} + \sum_{\ell=0}^{L} \beta_\ell^{\text{expectations}} \Delta\left(\Delta_{12}\text{mc}_{i,t-\ell}\right) + \Delta_{12}u_{i,t}$$

In Figure 3, we show $\beta_\ell^{\text{expectations}}$ for $\ell = 0, 1, \ldots, L$ with $L = 11$, for both no controls and sector-date fixed effects. We see that the coefficients remain well above zero throughout the year. This suggests firms may be slowly learning about an underlying, persistent component of cost growth. Furthermore, this result is highly suggestive that firms overreact to past costs; we show this for-

[33] Note that, even if firms are unconditionally correct about the autocorrelation of their cost growth, they can still make conditional forecast errors when different shocks generate different paths of cost growth, as we show in Section D.5.

[34] We can therefore motivate our previous regression under the assumption of homogeneous coefficients over the last twelve months: $\beta_\ell^{\text{expectations}} = \beta_{\text{expectations}}$ for $\ell = 0, \ldots, 11$, else $\beta_\ell^{\text{expectations}} = 0$.

Figure 3: Firm costs take time to drive beliefs

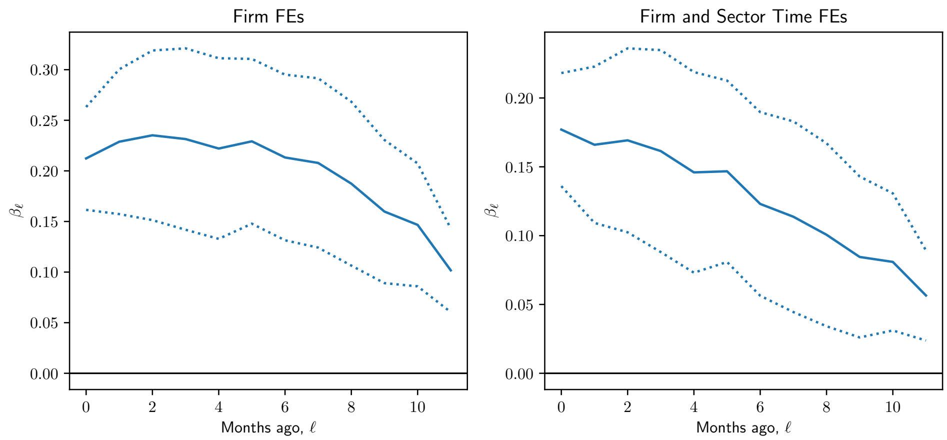

Note: Dotted lines are 95% confidence intervals, standard errors clustered at firm level.

mally in Appendix A.16.

Fact 5: Beliefs underreact until costs move. So far, we have focused on the pass-through of a firm's cost growth into beliefs about future costs. But how do firms use other information to form beliefs? This same question, albeit focused on the Survey of Professional Forecasters, has a complicated answer in the current literature; Broer and Kohlhas (2022) find beliefs overreact to some public signals and underreact to others. Furthermore, firms' beliefs may overreact and underreact to the same shock at different lags. Our hypothesis to explain these mixed findings is that firms consistently respond too little to public signals but respond too strongly to their private costs, as found above.[35] This suggests a particular time profile to learning about aggregate shocks. There is initial underreaction while news about the shock is contained only in public signals, followed by a move toward overreaction once firms' costs are affected. This time profile is akin to the "delayed overshooting" effect in Angeletos et al. (2021), except that for us the delayed overshooting is driven by the slow transmission of shocks to an endogenous economic object: firms' costs. As such, we also argue that while "delayed overreaction" is possible, so is overreaction with no delay, or beliefs that underreact at all dates.

We look at how firms react to a change in the Brent oil price.[36] As shown in Minton and Wheaton (2023), persistent oil price movements filter through the supply chain gradually and have substantial predictive power for future sectoral inflation. We first project each firm's cost growth onto a one-quarter change in the log Brent oil price, restricting to firms with a time series of at least 30 observations. We equally split firms into those relatively more affected after a year (slow exposure to oil) and those relatively more affected within a year (fast exposure to oil). The average cost impulse in each of these groups is shown in red in Figure 4. While these lines differ by

[35] When own costs are omitted, therefore, it is possible to find over- or underreaction to a variety of public signals, depending on how they correlated with firms' costs.

[36] International Monetary Fund, Global price of Brent Crude [POILBREUSDM], retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/POILBREUSDM, June 26, 2026.

Figure 4: Underreaction and overreaction to Oil Prices

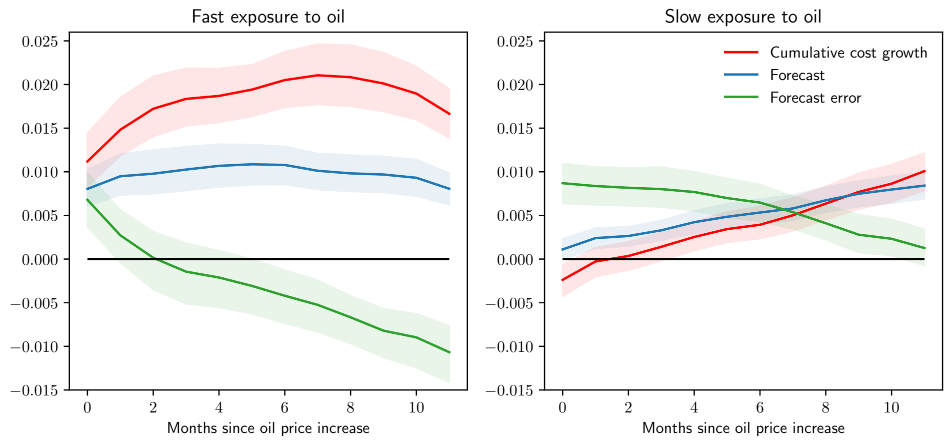

Note: 95% confidence intervals from a firm-level (cluster) bootstrap with 2000 draws. In each draw, we resample firms with replacement and re-compute the Fast/Slow split on the bootstrap sample.

construction, our theory suggests that expectations should move proportionately to costs; rather than jumping on impact to the optimal forecasts of costs twelve months ahead. This dynamic is borne out by the blue line: a local projection of beliefs onto the oil price change. Finally, to formally show that this behavior is not consistent with the path of costs beyond the one-year horizon, we also plot in green a local projection of the forecast error. Firms that are quickly exposed to oil see their costs rise, then start to fall, as the one-off cost shock falls out of cost growth. This leads to a pattern of under- and then overreaction. However, for firms only slowly exposed to oil, they underreact through the entire first year, as they are continually surprised by the increasing path of costs. This suggests that "delayed overshooting," as found in Angeletos et al. (2021), can occur but depends on the endogenous evolution of costs over time.

2 A calibrated model of firms' beliefs

In this section, we propose and calibrate a model of belief formation that is consistent with the evidence established in Section 1. To motivate our model, we first outline where existing models in the literature do not jointly match the facts we established above.[37]

In **Fact 2** we demonstrated overreaction to own-costs on average, violating FIRE. Another broad class of models consider limited-information while maintaining rational expectations (LIRE), for example Mankiw and Reis (2002) and Maćkowiak and Wiederholt (2009). In such models, firms do not necessarily know their current costs. This allows for overreaction to a firm's true costs.[38] However, LIRE models remain inconsistent with **Fact 2** because we find overreaction to *reported* own costs. There are also models that deviate from rational expectations but by damp-

[37] Table B.11 in Appendix B.11 summarizes these findings.

[38] For example, Afrouzi and Yang (2021) argue firms will pay more attention to more persistent components of costs, as they are more valuable to learn about dynamically. This may generate such overreaction to true costs. To further generate overreaction to reported costs, a further friction is needed.

ening beliefs. For example, cognitive discounting, as used in Gabaix (2020), attenuates forecasts back towards steady state, leading to *less* extrapolation from own costs.

By contrast, diagnostic expectations, as in Bordalo et al. (2018), can generate the overreaction we see. But diagnostic expectations predicts overreaction to surprises (relative to a FIRE forecast made in the previous period), and so does not explain that firms are insufficiently sensitive to their true cost persistence, as in **Fact 3**.[39] Furthermore, Bordalo et al. (2018) predict only overreaction to surprises relative to FIRE last period.[40] This contrasts with the inertia in **Fact 4**. In particular, we found forecast errors are predictable from costs in previous months.

Angeletos et al. (2021) combines aspects of a noisy signals model with diagnostic expectations, thereby matching overreaction while allowing for inertia in beliefs. In response to an aggregate shock, their model generates initial underreaction as firms only slowly process the news, followed by an overreaction, as they hold misspecified beliefs about the true signal processes. Since firms in their model learn by directly observing (with noise) the exogenous shocks, they predict a uniform "delayed overshooting" pattern for all firms, no matter the profile of their costs in response to the shock. However, in **Fact 5**, we argued that the pattern of a firm's underreaction and overreaction to a shock depends on how quickly that shock affects the firm's costs.

Next, we will propose a simple Adaptive Learning model that matches our five facts. In Section 3.1 below we show quantitatively how our model's predictions differ from those of other theories.

2.1 A model of Adaptive Learning

We propose a model in which firms use their own costs as a signal to learn about a persistent underlying process, and we also nest full-information rational expectations (FIRE).[41] In particular, when a firm makes a forecast, they place a weight α^{FIRE} on the FIRE forecast.[42] And they place a weight $1 - \alpha^{\text{FIRE}}$ on a behavioral forecast, denoted $\tilde{\mathbb{E}}_{i,t}$, that is entirely driven by a firm's own costs. Motivated by the stability of our results (across time, sector, and firm size) in **Fact 3**, our model features firms extrapolating from costs in a constant way; we call these "Adaptive Learning" (AL) beliefs.[43] These beliefs are formed as follows. Suppose that a firm believes that its nominal cost

[39] See in particular Section A.15 for further details.

[40] Chodorow-Reich, Guren and McQuade (2024) present a model where the reference period is delayed. Maxted (2024) presents a model in which diagnostic expectations (allowing for overreactive extrapolation) is combined with a slower moving belief component ("sentiment").

[41] The idea of mixing extrapolation from the current situation with a dampened rational reaction is captured in Keynes (1936): "It is reasonable, therefore, to be guided to a considerable degree by the facts about which we feel somewhat confident, even though they may be less decisively relevant to the issue than other facts about which our knowledge is vague and scanty. For this reason the facts of the existing situation enter, in a sense disproportionately, into the formation of our long-term expectations; our usual practice being to take the existing situation and to project it into the future, modified only to the extent that we have more or less definite reasons for expecting a change."

[42] In Appendix B.1, we show how this behavior results from firms finding it costly to compute the FIRE forecast. And in Theorem 1 below, we show that the aggregate implications are identical if we instead consider a two-agent model, in which some firms are fully FIRE, and others fully Adaptive Learners.

[43] Our focus is on matching average firms' beliefs to make macroeconomic predictions. Nevertheless, in Appendix B.2, we allow firms to choose a firm-specific $\tilde{\rho}_i$, allowing for heterogeneity as well as endogenizing the choice of $\tilde{\rho}_i$. Our findings there are approximately captured by the model presented in the main text.

growth has a persistent and a transitory component following the processes

$$\Delta \text{mc}_{i,t} = \Delta \tilde{\text{mc}}^{p}_{i,t} + \tilde{\eta}_{i,t}, \qquad \tilde{\eta}_{i,t} \overset{\text{iid}}{\sim} \mathcal{N}\left(0, \tilde{\sigma}^2_{\eta}\right) \tag{10}$$

$$\Delta \tilde{\text{mc}}^{p}_{i,t} = \tilde{\rho} \Delta \tilde{\text{mc}}^{p}_{i,t-1} + \tilde{\varepsilon}_{i,t}, \qquad \tilde{\varepsilon}_{i,t} \overset{\text{iid}}{\sim} \mathcal{N}\left(0, \tilde{\sigma}^2_{\varepsilon}\right) \tag{11}$$

where the tildes emphasize that these are the parameters only in the firm's *perceived* law of motion for costs. We focus here on deviations from a given permanent rate of nominal marginal cost growth.[44] Given the perceived law of motion implied by (10) and (11), the firm forecasts their costs as follows

$$\tilde{\mathbb{E}}_{i,t}\left[\Delta \text{mc}_{i,t+\tau}\right] = \tilde{\rho}^{\tau} \tilde{\mathbb{E}}_{i,t}\left[\Delta \tilde{\text{mc}}^{p}_{i,t}\right] \tag{12}$$

$$\tilde{\mathbb{E}}_{i,t}\left[\Delta \tilde{\text{mc}}^{p}_{i,t}\right] = \tilde{\mathcal{K}} \Delta \text{mc}_{i,t} + \left(1 - \tilde{\mathcal{K}}\right) \tilde{\rho} \tilde{\mathbb{E}}_{i,t-1}\left[\Delta \tilde{\text{mc}}^{p}_{i,t-1}\right] \tag{13}$$

where $\tilde{\mathcal{K}} \equiv 1 - \frac{2\tilde{\sigma}^2_{\eta}}{\tilde{\sigma}^2_{\varepsilon} + (1+\tilde{\rho}^2)\tilde{\sigma}^2_{\eta} + \sqrt{\left((1-\tilde{\rho}^2)\tilde{\sigma}^2_{\eta} - \tilde{\sigma}^2_{\varepsilon}\right)^2 + 4\tilde{\sigma}^2_{\eta}\tilde{\sigma}^2_{\varepsilon}}}$ is the Kalman gain.

Overall then the period t forecast of τ period ahead nominal cost growth is a weighted-average between this Adaptive Learning forecast and the FIRE forecast,

$$\mathbb{E}_{i,t}\left[\Delta \text{mc}_{i,t+\tau}\right] = \alpha^{\text{FIRE}} \mathbb{E}_{t}\left[\Delta \text{mc}_{i,t+\tau}\right] + \left(1 - \alpha^{\text{FIRE}}\right) \tilde{\mathbb{E}}_{i,t}\left[\Delta \text{mc}_{i,t+\tau}\right] \tag{14}$$

To calibrate firms' beliefs we therefore need to estimate three parameters: α^{FIRE}, $\tilde{\rho}$, and $\tilde{\mathcal{K}}$. α^{FIRE} captures firms learning from other signals and allows our model to nest FIRE. $\tilde{\rho}$ is the firm's perceived cost persistence. And $\tilde{\mathcal{K}}$ tells us whether firms use only their costs today, or also their past costs, when forming these beliefs. In theorems below, we will generally allow the parameter space to be $\tilde{\mathcal{K}} \in [0, 1]$, $\tilde{\rho} \in [0, 1)$, and $\alpha^{\text{FIRE}} \in [0, 1]$. As we will see, our calibration lies in this space.

This model is very similar to that in Angeletos, Huo and Sastry (2021) (AHS) with two key differences.[45] First, while we also model beliefs as an AR(1), we do not impose that the true process necessarily also follows an AR(1). Second, and more importantly, firms learn about many shocks from just their own costs. As we will see below, this leads to a single Phillips curve that determines the inflation response to all shocks. Under fully Adaptive Learning ($\alpha^{\text{FIRE}} = 0$), this Phillips curve is the adaptive expectations Phillips curve of Friedman (1968) and Phelps (1967). In AHS, the learning process differs in the true process of every shock, so there is one Phillips curve per shock, each of which has its own calibration. The existence of a separate Phillips curve per shock is exactly why AHS *can* find delayed overshooting in response to all shocks.

[44] With zero inflation in steady state and no permanent shocks, the model here is appropriate. In Appendix B.3, we furthermore show that if firms perfectly observe any permanent shocks to their rate of nominal cost growth, then the model here captures learning around this steady state. If instead we allow firms to learn about permanent shocks to cost growth via their own past costs, but view such shocks as unlikely, then firms would eventually learn about any permanent changes in their cost growth, with very little quantitative effect on our short-run dynamic results below.

[45] Our model may also be related, though less explicitly, to models with both diagnostic expectations and belief inertia (Chodorow-Reich et al. (2024), Maxted (2024)). As with AHS, our core distinction from such models is limited information: firms observe their costs but not the myriad shocks in the economy.

2.2 Calibrating the Model

The BIE survey asks firms each month to report cost growth over the last year and to forecast cost growth over the next year. The timing that aligns best with the data is therefore an annual frequency. Casting the model from Section 2.1 in annual time and combining equations (12), (13), and (14), beliefs depend on FIRE beliefs and cost growth according to

$$\mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right] = \alpha^{\text{FIRE}}\mathbb{E}_t\left[\Delta_{12}\text{mc}_{i,t+12}\right] + \left(1-\alpha^{\text{FIRE}}\right)\tilde{\rho}^{\text{ann}}\tilde{\mathcal{K}}^{\text{ann}}\sum_{\ell=0}^{\infty}\left[\left(1-\tilde{\mathcal{K}}^{\text{ann}}\right)\tilde{\rho}^{\text{ann}}\right]^{\ell}\Delta_{12}\text{mc}_{i,t-12\ell}. \tag{15}$$

Realized cost growth in the future is the rational expectation of future cost growth plus the realization of future shocks occurring after period t and up to period $t+12$:

$$\Delta_{12}\text{mc}_{i,t+12} = \mathbb{E}_t[\Delta_{12}\text{mc}_{i,t+12}] + \epsilon^{\text{FIRE}}_{i,t+1,t+12}. \tag{16}$$

Substituting the rational expectations relationship (16) into the belief model (15) and absorbing most lags of cost growth into the error term gives the regression equation

$$\mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right] = \alpha^{\text{FIRE}}\Delta_{12}\text{mc}_{i,t+12} + \left(1-\alpha^{\text{FIRE}}\right)\tilde{\rho}^{\text{ann}}\tilde{\mathcal{K}}^{\text{ann}}\Delta_{12}\text{mc}_{i,t} + \left(1-\alpha^{\text{FIRE}}\right)\tilde{\rho}^{\text{ann}}\tilde{\mathcal{K}}^{\text{ann}}\left(1-\tilde{\mathcal{K}}^{\text{ann}}\right)\tilde{\rho}^{\text{ann}}\Delta_{12}\text{mc}_{i,t-12} + \epsilon_{it}, \tag{17}$$

where $\epsilon_{it} = -\alpha^{\text{FIRE}}\epsilon^{\text{FIRE}}_{i,t+1,t+12} + \left(1-\alpha^{\text{FIRE}}\right)\tilde{\rho}^{\text{ann}}\tilde{\mathcal{K}}^{\text{ann}}\sum_{\ell=2}^{\infty}\left[\left(1-\tilde{\mathcal{K}}^{\text{ann}}\right)\tilde{\rho}^{\text{ann}}\right]^{\ell}\Delta_{12}\text{mc}_{i,t-12\ell}$.

Future cost growth $\Delta_{12}\text{mc}_{i,t+12}$ is correlated with $\epsilon^{\text{FIRE}}_{i,t+1,t+12}$, the shocks that have occurred between period $t+1$ and period $t+12$.[46] Note that, because $\Delta_{12}\text{mc}_{i,t}$ and $\Delta_{12}\text{mc}_{i,t-12}$ are known at time t, they are uncorrelated with the FIRE component of the error, $\epsilon^{\text{FIRE}}_{i,t+1,t+12}$. So we need an instrument for $\Delta_{12}\text{mc}_{i,t+12}$. Exogeneity and relevance are satisfied by any information from period $t-\ell$, with $\ell \geq 0$, that meaningfully affects $\Delta_{12}\text{mc}_{i,t+12}$.[47]

Our baseline calibration instruments future cost growth with average current cost growth in the BIE. In particular, we interact this time series with a firm-level indicator, thereby using a firm's loading on macroeconomic cost variation to forecast its costs. We include firms with at least 48 observations. We also test whether our results are robust to other instruments: firm-level indicators interacted with (1) twelve-month CPI inflation[48] and (2) Nakamura and Steinsson (2018) monetary policy shocks.[49] Using cross-sectional instruments helps us to deal with weak instruments issues

[46] Further, marginal cost growth may be autocorrelated, meaning regressors may also be correlated with the lagged cost components of ϵ_{it}. However, Table 2 showed autocorrelation of cost growth was very low at an annual frequency. Further, we can assess how much our estimates change when including additional lags of cost growth in our estimation. Appendix B.4 shows that our calibration does not meaningfully change when we include another lag of historical marginal cost growth, and the coefficients on this additional lag are small and statistically insignificant.

[47] Note that we do not need to causally identify the effect of a *shock* on $\Delta_{12}\text{mc}_{i,t+12}$; the instrument merely needs to be relevant for future costs and not include information after period t.

[48] U.S. Bureau of Labor Statistics, Consumer Price Index for All Urban Consumers: All Items in U.S. City Average [CPIAUCSL], retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/CPIAUCSL, June 26, 2026.

[49] We use the updated series from Acosta, Brennan and Jacobson (2024).

common to this setting.[50] Our first-stage gives large Kleibergen Paap F-statistics. We also check that our results are similar under limited information maximum likelihood (LIML) estimation and increasing the threshold for the minimum observations per firm.[51] We also show our results are robust to using sectoral, instead of firm-level, loadings.

The first stage is the panel local projection

$$\Delta_{12}\text{mc}_{i,t+12} = \mu_i + \lambda_t + \delta_{1i} z_t + \delta_2 X_{it} + v_{it}, \tag{18}$$

where z_t is our instrument, and X_{it} are the exogenous second-stage regressors $\Delta_{12}\text{mc}_{i,t}$ and $\Delta_{12}\text{mc}_{i,t-12}$.

As in our analysis above, we will also include firm fixed effects in our regressions. Firms have different steady state levels of cost growth, and this is not a difference that our model seeks to explain. Further, we include time fixed effects to isolate how our time series instruments with firm-level loadings induce differential responses of costs across firms. A particular concern that time fixed effects address in our setting is that aggregate sentiment or narrative-driven shocks may feed back into aggregate costs in general equilibrium, generating omitted variables bias. The final form of our calibration second stage is therefore

$$\mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right] = \mu_i + \lambda_t + \beta_1 \Delta_{12}\text{mc}_{i,t} + \beta_2 \Delta_{12}\text{mc}_{i,t-12} + \beta_3 \Delta_{12}\text{mc}_{i,t+12} + \epsilon_{it}, \tag{19}$$

where our model parameters are exactly identified from the regression coefficients

$$\beta_1 = \left(1-\alpha^{\text{FIRE}}\right)\tilde{\rho}^{\text{ann}}\tilde{\mathcal{K}}^{\text{ann}}, \quad \beta_2 = \left(1-\alpha^{\text{FIRE}}\right)\tilde{\rho}^{\text{ann}}\tilde{\mathcal{K}}^{\text{ann}}\left(1-\tilde{\mathcal{K}}^{\text{ann}}\right)\tilde{\rho}^{\text{ann}}, \quad \beta_3 = \alpha^{FIRE} \tag{20}$$

Standard errors are constructed via the delta method. We use Driscoll and Kraay (1998) standard errors with bandwidth 12.[52] The results are shown in Table 3.

We reject the hypothesis of full-information rational expectations, in which we would expect to find the coefficient on future cost growth to be one ($\alpha^{FIRE} = 1$) and the coefficients on recent and lagged cost growth to be zero. Instead, across all instruments and 2SLS and LIML estimation, we find that future cost growth and realized cost growth over the last year are about equally important for firms' beliefs about cost growth over the next year. Even when the rational expectation predicts cost growth to return to normal levels in the next year, recent cost growth above the norm will lead firms to expect elevated cost growth in the future. Further, though cost growth from more than a year ago is not very important for beliefs, firms do use this information in forming beliefs. In terms of model parameters, these findings translate into firms extrapolating from the perceived persistent component of recent cost growth, formalized by a statistically significant $\tilde{\rho}^{ann} > 0$, and using information from more than a year ago, formalized by a statistically significant $\tilde{\mathcal{K}}^{ann} < 1$.

[50] See Mavroeidis, Plagborg-Møller and Stock (2014) for discussion of weak instruments in the estimation of New Keynesian Phillips curves. Gagliardone, Gertler, Lenzu and Tielens (2023) is a recent example using panel variation to generate sufficiently powerful instruments to explore pricing.

[51] As in Nakamura and Steinsson (2014), we use LIML to address the use of many instruments. LIML is thought to be more robust to instrument weakness than 2SLS (see, e.g. Chao and Swanson (2005) and Stock and Yogo (2005)).

[52] Recall that we have monthly data on twelve-month changes, suggesting a mechanical MA(11) error process.

Table 3: Annual Calibration

Panel A: 2SLS estimates

		$\mathbb{E}_{i,t}\left[\Delta_{12}\mathrm{mc}_{i,t+12}\right]$		
Regression Estimates	$\Delta_{12}\mathrm{mc}_{i,t}$	0.236***	0.243***	0.236***
		(0.022)	(0.021)	(0.020)
	$\Delta_{12}\mathrm{mc}_{i,t-12}$	0.040***	0.034**	0.039***
		(0.015)	(0.015)	(0.015)
	$\Delta_{12}\mathrm{mc}_{i,t+12}$	0.251***	0.200***	0.246***
		(0.043)	(0.036)	(0.059)
Implied Parameters	α^{FIRE}	0.251	0.200	0.246
		(0.043)	(0.036)	(0.059)
	$\tilde{\mathcal{K}}^{\text{ann}}$	0.650	0.687	0.654
		(0.105)	(0.112)	(0.098)
	$\tilde{\rho}^{\text{ann}}$	0.484	0.442	0.479
		(0.059)	(0.058)	(0.065)
	Instrument	Av. Costs × Firm id	CPI × Firm id	NS × Firm id
	Min obs per firm	4 years	4 years	4 years
	Observations	8762	8762	8762
	Kleibergen Paap F stat	4356	4258	3143

Panel B: LIML estimates

		$\mathbb{E}_{i,t}\left[\Delta_{12}\mathrm{mc}_{i,t+12}\right]$		
Regression Estimates	$\Delta_{12}\mathrm{mc}_{i,t}$	0.296***	0.292***	0.278***
		(0.030)	(0.031)	(0.025)
	$\Delta_{12}\mathrm{mc}_{i,t-12}$	0.022	0.026	0.042*
		(0.017)	(0.018)	(0.022)
	$\Delta_{12}\mathrm{mc}_{i,t+12}$	0.093	0.127	0.245*
		(0.171)	(0.148)	(0.127)
Implied Parameters	α^{FIRE}	0.093	0.127	0.245
		(0.171)	(0.148)	(0.127)
	$\tilde{\mathcal{K}}^{\text{ann}}$	0.814	0.787	0.709
		(0.121)	(0.125)	(0.107)
	$\tilde{\rho}^{\text{ann}}$	0.401	0.425	0.518
		(0.086)	(0.082)	(0.113)
	Instrument	Av. Costs × Firm id	CPI × Firm id	NS × Firm id
	Min obs per firm	6 years	6 years	6 years
	Observations	6691	6691	6691
	Kleibergen Paap F stat	740	1233	318

Note: *p<0.10, **p<0.05, ***p<0.01. All regressions include Firm and Time fixed effects. Driscoll Kraay standard errors with bandwidth 12.

In Appendix B.5, we show how the annual calibration can be mapped to a quarterly calibration, which we use in our model. As is typical, this approach essentially entails taking fractional powers of the annual coefficients. The implied quarterly calibration, which we use below, is

$$\alpha^{\text{FIRE}} = 0.25, \qquad \tilde{\rho} = 0.82, \qquad \tilde{\mathcal{K}} = 0.22$$

A key contribution of our paper is to empirically characterize firms' beliefs. We are now in a position to state our calibrated beliefs model in full for the quarterly calibration.

Calibrated model of quarterly cost forecasts. *Firm i's forecast of its own future nominal marginal cost growth is given by the equation*

$$\mathbb{E}_{i,t}\left[\Delta \mathrm{mc}_{i,t+\tau}\right] = 0.25 \times \mathbb{E}_t\left[\Delta \mathrm{mc}_{i,t+\tau}\right] + 0.16 \times 0.82^{\tau} \times \sum_{\ell=0}^{\infty} 0.64^{\ell} \Delta \mathrm{mc}_{i,t-\ell} \tag{21}$$

2.2.1 Robustness and external validity

While we estimate α^{FIRE} here across multiple shocks, finding consistent results (and below we show our calibration is consistent with the oil response in **Fact 5**), we are not claiming that α^{FIRE} must be identical for all shocks.[53] Instead, we note substantial myopia on average and for typical macroeconomic shocks. Moreover, our emphasis in this paper is on the importance of the Adaptive Learning channel, with FIRE beliefs acting as a control and quantitative addition to the model.

A more specific concern might be that, as our calibration regression includes time fixed effects, we identify α^{FIRE} using cross-sectional variation in costs and cost beliefs.[54] However, while the cost belief outcome is therefore stripped of its common component, we still learn how firms forecast aggregate components of cost growth, since these beliefs are not constant across firms.[55] Additionally, since our instruments are not idiosyncratic shocks, but firm-specific loadings on aggregate shocks, we argue that our α^{FIRE} speaks to the belief reaction to aggregate shocks. For example, many rational inattention models find that firms should pay more attention to firm-specific conditions than aggregate conditions, making it unlikely that firms would under-appreciate how relevant a shock is for them and yet understand how relevant the shock is for the economy in the aggregate.

And other evidence supports the view that α^{FIRE} is not substantially higher for the common component of aggregate shocks in particular. First, the broader literature on firms' beliefs about aggregates reveals substantial mistakes, with firms not knowing current and past CPI and making predictable forecast errors about future CPI—overreacting to micro news and underreacting to macro news (Kumar et al. (2015), Andrade et al. (2022), Born et al. (2022)). Second, our experiment

[53] For example, a labor-intensive firm may be much more responsive than this if predicting the effect on their future costs of a pre-announced change in the minimum wage next month.

[54] Formally, variables are residualized with respect to their common components, so no variation would be left if all firms' variables moved in the same way.

[55] We showed this for CPI in Appendix A.9, and this echoes similar results in Andrade et al. (2022) and Albagli et al. (2022).

in **Fact 1** shows that the average perceived relevance of CPI to a firm's costs (in absolute, not relative, terms) is low; in Section 2.4 below, we derive an α^{FIRE} estimate from the experiment and show that it aligns with our calibration.

We also compare our calibration to a number of external moments. In Appendix B.6, we show that our calibrated beliefs model provides a good fit to unconditional average beliefs over time across all the sectors in our data. In Appendix B.7, we show that our calibrated model provides a much closer fit than FIRE to the empirical behavior of beliefs seen in **Fact 5**, generating them as an untargeted moment.

Finally, in Appendix Table B.6, we show that our results are robust to the use of sector-level loadings instead of firm-level loadings. With only 14 sectors, this regression uses many fewer instruments. And in this case, we can include all firms in the BIE sample, including firms present only for a short period of time. In Appendix Table B.7, we show that our 2SLS estimates are robust to increasing the threshold to drop firms that were not present in the BIE for at least 72 months – these results are estimated on the same sample as the LIML results in Panel B of Table 3.

2.3 Why do firms overreact to their own costs?

Fact 2 showed that firms believed cost growth was more persistent than it truly was *over our sample period*. We argue below that such behavior is consistent with an adaptive learning model (as in Evans and Honkapohja (2001)) in which firms learn about cost persistence sufficiently slowly, because cost growth was more persistent historically. Therefore, our model can be thought of as a mis-specified adaptive learning model, under the limiting approximation of no learning, which is highly accurate at business cycle frequencies.

Consider a model corresponding to that in Section 2.1 but with $\tilde{\sigma}_\eta = 0$ and cast in annual time. So the perceived law of motion for costs is

$$\Delta_{12}\text{mc}_{i,t} = a_i + \tilde{\rho}_{i,t}^{\text{ann}}\Delta_{12}\text{mc}_{i,t-12} + \tilde{\varepsilon}_{i,t} \tag{22}$$

We suppose that each firm updates the perceived cost persistence, $\tilde{\rho}_{i,t}^{\text{ann}}$, as new data comes in, under constant gain learning at rate γ.[56]

In order to evaluate learning over a longer horizon, we exploit the tight relationship between aggregate unit cost growth in the BIE and the GDP deflator over our sample period (as we show in Figure A.2). We leverage this relationship to simulate cost series back to 1949.[57] Given this simulated data and a gain parameter, γ, we can compute what the learning model predicts for $\tilde{\rho}_{i,t}^{\text{ann}}$ at each point in time in (22).

[56] Least squares learning is captured by the standard recursive least squares formulas

$$\beta_t = \beta_{t-1} + \frac{1}{t}R_t^{-1}x_t\left(y_t - x_t'\beta_{t-1}\right), \qquad R_t = R_{t-1} + \frac{1}{t}\left(x_t x_t' - R_{t-1}\right),$$

whereas constant gain learning replaces the $\frac{1}{t}$ terms with a constant gain, γ. In our setting x_t is a vector containing 1 and $\Delta_{12}\text{mc}_{i,t-12}$, and $y_t = \Delta_{12}\text{mc}_{i,t}$.

[57] For more details, see Appendix B.8.

Over our sample period in the BIE, as shown in Table 2, we estimated a constant $\tilde{\rho}^{\text{ann, static}}$ of 0.36. We can generate this belief in a constant-gain learning model when γ is 0.01 quarterly, i.e. if firms update just 1% of their beliefs each quarter. Even though true cost persistence post-2011 was much below 0.36, this belief is generated because of the historically higher persistence of cost growth, proxied by the GDP deflator. This very slow learning is not out of line with estimates in the learning literature.[58]

A different explanation for overreaction could be that firms assign too much probability to permanent cost shocks. Using firms' forecasts about longer horizons, Appendix B.9 shows that firms' beliefs about both short- and long-run cost dynamics overreact to past costs: permanent shocks alone cannot fully explain overreaction.[59]

2.4 Why do firms underreact to public signals?

In this section, we decompose α^{FIRE} into a full information and a rational expectation (RE) component, and use **Fact 1** to argue that a major reason firms underreact to public signals is a failure of RE. For example, in the context of **Fact 5**, this would imply that firms initially underreacted to oil prices not only because they didn't observe oil prices, but also because they didn't understand their dynamic exposure to them. Finally, we use our finding to inform an upper bound on α^{FIRE} that independently corroborates the estimate from our calibration above.

Recall the experiment from **Fact 1**. We now transform the estimated CPI to own-cost passthrough into a quantitative measure of the deviation from rational expectations. Suppose that year-on-year cost growth is given by

$$\Delta_{12}\text{mc}_{i,t} = \rho_{\text{idio}}\Delta_{12}\text{mc}_{i,t-12} + \psi_{\text{CPI}}\Delta_{12}p_t^{\text{CPI}} + \varepsilon_{i,t} \tag{23}$$

We allow firms to have limited information and to hold incorrect beliefs about the pass-through of each term into costs, so that beliefs are

$$\mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right] = \tilde{\rho}_{\text{idio}}\Delta_{12}\text{mc}_{i,t} + \tilde{\psi}_{\text{CPI}}\mathbb{E}_{i,t}\left[\Delta_{12}p_{t+12}^{\text{CPI}}\right]$$

Let $\tilde{\psi}_{\text{CPI}} = \alpha^{\text{RE}}\psi_{\text{CPI}}$ so that α^{RE} captures how much of the true effect of CPI on own costs is understood by the firm. Our experiment estimated $\tilde{\psi}_{\text{CPI}} \approx 0.15$. We now estimate the true ψ_{CPI} using (23). As this is essentially a time series regression, and the BIE sample only begins in October 2011, we instead take the integral over (23) and use the year-on-year change in the GDP deflator in place of $\int \Delta_{12}\text{mc}_{i,t+12}\text{d}i$. As we saw in Figure A.2 above, this measure provides a very close

[58] Eusepi and Preston (2011) for example estimate a gain of 0.002; though they note a more typical range in the literature is 0.007–0.05. Farmer, Nakamura and Steinsson (2024) show Bayesian learning about parameters may also be very slow.

[59] Recall that nominal cost forecasts beyond the date of the next price reset are irrelevant to a firm's optimal reset price today. Since almost all firms reset prices at least annually, mistakes about further out horizons are not pivotal for pricing decisions (though they may be for other decisions such as investment). In Appendix B.3, we motivated our main belief model under the assumption that firms are fully informed about permanent shocks. In Appendix Figure D.16, we show that even if all learning was about the long run (the $\tilde{\rho} = 1$ panel), inflation would continue to show limited anticipation and high pass-through – but there is more inertia than in our base case.

fit over the BIE sample period. The results from this exercise are shown in Appendix B.10. We estimate $\rho_{\text{idio}} = 0.17$ and $\psi_{\text{CPI}} = 0.68$, suggesting $\alpha^{\text{RE}} \approx 0.25$.

We can also use this framework to decompose our α^{FIRE} estimate from Section 2.2. Suppose there is some news about CPI inflation over the next year (with no effect on current costs), denoted s_t^{CPI}. We can decompose the effect of this on a firm's cost forecast into a failure of full-information and a failure of rational expectations, times the FIRE forecast,

$$\frac{d\mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right]}{ds_t^{\text{CPI}}} = \underbrace{\alpha^{\text{FI}} \cdot \alpha^{\text{RE}}}_{=\alpha^{\text{FIRE}}} \cdot \frac{d\mathbb{E}_t^{\text{FIRE}}\left[\Delta_{12}\text{mc}_{i,t+12}\right]}{ds_t^{\text{CPI}}}$$

where $\alpha^{\text{FI}} \equiv \frac{d\mathbb{E}_{i,t}\left[\Delta_{12}p_{t+12}^{\text{CPI}}\right]}{ds_t^{\text{CPI}}} \Big/ \frac{d\mathbb{E}_t^{\text{FIRE}}\left[\Delta_{12}p_{t+12}^{\text{CPI}}\right]}{ds_t^{\text{CPI}}}$ captures how much of the news is incorporated into CPI beliefs, compared to under FIRE.[60] We take two implications from this. First, if $\alpha^{\text{FI}} \leq 1$, our α^{RE} estimate represents an upper bound estimate for α^{FIRE}. Second, taking both estimates together suggests that a large portion of the deviation from FIRE is explained by a failure of rational expectations, rather than of full-information. Clearly, however, there are failures of full-information too: the very fact our information treatment moved CPI beliefs tells us so.

3 Implications for the Phillips Curve

In this section, we incorporate our calibrated beliefs into a Calvo pricing model. We characterize theoretically how nominal marginal cost shocks affect prices. We compare these effects to those generated by full-information rational expectations (FIRE) and alternative models of expectations formation in the literature. Moving to general equilibrium, we derive the New Keynesian Phillips curve in our model and show it is less forward-looking but features higher pass-through relative to FIRE. In the special case of completely Adaptive Learning ($\alpha^{\text{FIRE}} = 0$), our NKPC collapses to the traditional Friedman-Phelps adaptive expectations Phillips curve but with less inertia. Finally, we provide an empirical test of our findings, following the methodology in Hazell et al. (2022).

3.1 Pricing under Adaptive Learning

Assume a standard Calvo sticky pricing model under monopolistic competition and with constant returns to scale. Firms are able to reset prices each period with probability $1 - \theta_p$, and the elasticity of substitution between varieties is $\sigma > 1$.

Theorem 1. *To first order around a zero inflation steady state, the overall price level is*

$$p_t = \theta_p p_{t-1} + \left(1 - \theta_p\right) \alpha^{FIRE} p_t^{*,FIRE} + \left(1 - \theta_p\right) \left(1 - \alpha^{FIRE}\right) p_t^{*,AL} \tag{24}$$

[60] So α^{FI} could also reflect RE mistakes in mapping the news about CPI into CPI beliefs; our focus is to distinguish frictions in learning about aggregates (here, α^{FI}) and frictions in mapping beliefs about these aggregates into variables of direct interest to the firm (here, α^{RE}).

in log deviations from steady state. The FIRE optimal reset price is the standard expression

$$p_t^{*,FIRE} = \left(1-\beta\theta_p\right)\sum_{\tau=0}^{\infty}\left(\beta\theta_p\right)^{\tau}\mathbb{E}_t\left[mc_{t+\tau}\right], \tag{25}$$

and the Adaptive Learning optimal reset price is

$$p_t^{*,AL} = mc_t + \frac{\beta\theta_p\tilde{\rho}\tilde{\mathcal{K}}}{1-\beta\theta_p\tilde{\rho}}\sum_{\ell=0}^{\infty}\left[\left(1-\tilde{\mathcal{K}}\right)\tilde{\rho}\right]^{\ell}\Delta mc_{t-\ell}. \tag{26}$$

Proof. Under Calvo updating, the log-linearized aggregate price evolves as

$$p_t = \left(1-\theta_p\right)\int p_{i,t}^{*}\mathrm{d}i + \theta_p p_{t-1}. \tag{27}$$

Firm i's optimal reset price is given by

$$p_{i,t}^{*} = \left(1-\beta\theta_p\right)\sum_{\tau=0}^{\infty}\left(\beta\theta_p\right)^{\tau}\mathbb{E}_{i,t}\left[\mathrm{mc}_{i,t+\tau}\right]. \tag{28}$$

Substituting in the marginal cost beliefs in (12), (13), and (14), we obtain the result. □

Theorem 1 implies that, for aggregate implications, our model is equivalent to a two-agent model, in which some firms are fully FIRE, and some are fully adaptive learners. We also see that the optimal reset price under Adaptive Learning features no true anticipation; firms forecast the future but do so only with past costs. Prices will therefore overreact (relative to FIRE) to transitory shocks and underreact to persistent shocks, since both affect costs today, but persistent shocks affect the future more. This intuition is formalized in Theorem 2. Finally, we see how inertia in beliefs translates into pricing: prices will be high when past *nominal cost growth* was high.

Theorem 2. *Consider an AR(1) shock to nominal marginal cost growth, with true persistence ρ,*

$$\Delta mc_t = \rho\Delta mc_{t-1}, \qquad \Delta mc_0 = 1 \tag{29}$$

The on-impact price response is greater in our model than under FIRE if and only if

$$\tilde{\rho} > \frac{\rho}{\left(1-\beta\theta_p\rho\right)\tilde{\mathcal{K}} + \rho\beta\theta_p} \tag{30}$$

Under no backward lookingness, $\tilde{\mathcal{K}} = 1$, prices are higher than under FIRE exactly when the perceived persistence is too high, $\tilde{\rho} > \rho$,

$$p_t = p_t^{FIRE} + \left(1-\alpha^{FIRE}\right)\underbrace{\frac{\left(1-\theta_p\right)\beta\theta_p}{\left(1-\beta\theta_p\tilde{\rho}\right)\left(1-\beta\theta_p\rho\right)}\frac{\rho^{t+1}-\theta_p^{t+1}}{\rho-\theta_p}}_{>0}\cdot\left(\tilde{\rho}-\rho\right) \tag{31}$$

When the perceived persistence is correct, $\tilde{\rho} = \rho$, prices are always lower than under FIRE, if there is any degree of backward lookingness, $\tilde{\mathcal{K}} < 1$,

$$p_t = p_t^{FIRE} - \left(1 - \alpha^{FIRE}\right) \underbrace{\frac{(1-\theta_p)\,\beta\theta_p\rho}{(1-\beta\theta_p\rho)} \frac{\left((1-\tilde{\mathcal{K}})\,\rho\right)^{t+1} - \theta_p^{t+1}}{(1-\tilde{\mathcal{K}})\,\rho - \theta_p}}_{>0} (1-\tilde{\mathcal{K}}) \tag{32}$$

Proof. From (24), the difference between the price under our beliefs and under FIRE is

$$p_t - p_t^{\text{FIRE}} = \theta_p \left(p_{t-1} - p_{t-1}^{\text{FIRE}}\right) + (1-\theta_p)\left(1-\alpha^{\text{FIRE}}\right)\left(p_t^{*,\text{AL}} - p_t^{*,\text{FIRE}}\right) \tag{33}$$

From (25), the optimal reset price under FIRE, facing this AR(1) shock, is

$$p_t^{*,\text{FIRE}} = \frac{1-\rho^{t+1}}{1-\rho} + \frac{\beta\theta_p\rho}{1-\beta\theta_p\rho}\rho^t \tag{34}$$

And from (26), the optimal reset price under Adaptive Learning is

$$p_t^{*,\text{AL}} = \frac{1-\rho^{t+1}}{1-\rho} + \frac{\beta\theta_p\tilde{\rho}\tilde{\mathcal{K}}}{1-\beta\theta_p\tilde{\rho}} \frac{\rho^{t+1} - \left((1-\tilde{\mathcal{K}})\,\tilde{\rho}\right)^{t+1}}{\rho - (1-\tilde{\mathcal{K}})\,\tilde{\rho}} \tag{35}$$

The results in (30), (31), and (32) follow from combining these three equations – (33), (34), and (35) – under the stated parameter restrictions. □

Theorem 2 shows that, under our model of beliefs, macroeconomic shocks may be either under- or overreacted to depending on the persistence of the cost growth they generate. $\tilde{\mathcal{K}} < 1$ captures the key story at the heart of noisy information models: firms learn slowly about shocks, and so underreact. This dampening force is over-turned when the perceived persistence of cost growth sufficiently exceeds the true persistence. Theorem 2 formalizes the cutoff at which these forces balance on impact.[61]

Comparing to other models of beliefs. We now introduce four alternative theories of expectation formation and compare the models' predictions for pricing. We consider models with sticky information Mankiw and Reis (2002), sticky expectations Carroll, Crawley, Slacalek, Tokuoka and White (2020), cognitive discounting Gabaix (2020), and diagnostic expectations Bordalo et al. (2018).[62] The exact formulations used are detailed in Appendix B.11. We calibrate the models as follows. Under sticky information and sticky expectations, we keep 90% of expectations unchanged each period, $\theta^{\text{MR}} = \theta^{\text{SE}}$. Under cognitive discounting, we attenuate forecasts towards steady state at the rate $m^{\text{CD}} = 0.9$. Under diagnostic expectations, we place additional weight on

[61] Setting $\beta = 0.99$ and $\theta_p = 0.9$, and using the calibration for $\tilde{\rho}$ and $\tilde{\mathcal{K}}$ from Section 2.2, Theorem 2 implies that the on impact response is equal under FIRE and our calibrated model when $\rho = 0.42$.

[62] In all cases, we maintain the assumption of Calvo pricing and use these models to determine how expectations are formed.

Figure 5: Inflation response to 1pp permanent change in nominal marginal costs at date s

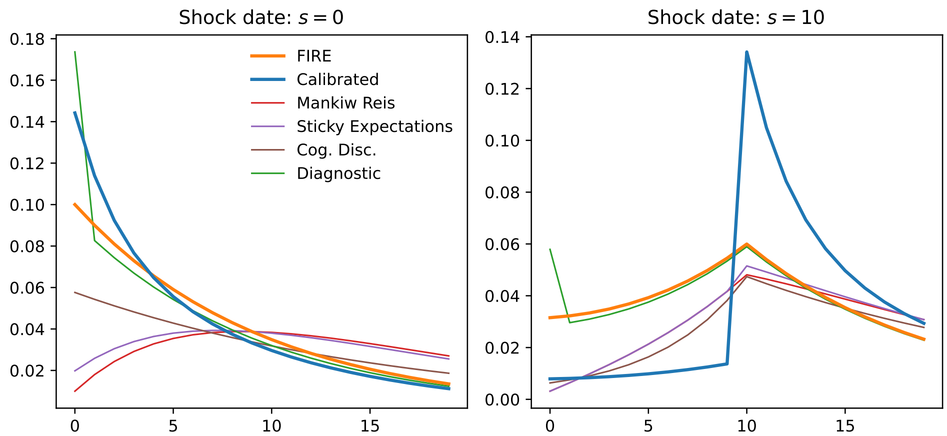

Note: Impulse response of inflation to a permanent rise in nominal marginal costs under six models. Calibrated is our model with α^{FIRE}, $\tilde{\rho}$ and $\tilde{\mathcal{K}}$ calibrated as in Section 2.2

news today of $\theta^{\text{DE}} = 1$. For our model we use the quarterly calibration of α^{FIRE}, $\tilde{\rho}$ and $\tilde{\mathcal{K}}$ from Section 2.2.

In Figure 5, we plot the response of inflation to a 1 percentage point, permanent increase in nominal marginal costs at date 0 (in the left-hand panel) and at date 10 (in the right-hand panel). Since the nominal marginal costs are given, the integral over inflation paths must sum to the same amount under all the models; the differences are entirely about timing.[63] Relative to FIRE, our model generates faster pass-through when the shock is immediate, as do diagnostic expectations, whereas the other three ("dampening") models suggest very slow pass-through to prices. By contrast, when the shock is in the future, our model implies the least anticipation, followed by strong pass-through once costs start moving. Overall, we agree with different elements of both these groups of theories. Once costs move, we argue that beliefs (and so prices) will adjust quickly and may overreact. However, we also agree with the "dampening" theories that prior to costs moving, there will be little reaction.

3.2 The New Keynesian Phillips curve

In Theorem 3, we derive the New Keynesian Phillips curve in our model: how real marginal costs drive inflation.

Theorem 3 (New Keynesian Phillips Curve.)**.** *The New Keynesian Phillips curve in our model takes the form*

$$\pi_t = \frac{\left(1 - \beta\delta\theta_p\right)\left(1 - \frac{\theta_p}{\delta}\right)}{\theta_p} \cdot \sum_{\tau=0}^{\infty} (\beta\delta)^{\tau}\, \mathbb{E}_t \left[mc_{t+\tau}^{real} \right] + \frac{1-\delta}{\delta} \cdot mc_t^{real} + \omega_b \cdot b_t \tag{36}$$

[63]In Appendix D.1, we show the long run of the price level under all the beliefs models is indeed the same.

where the average belief of the persistent component of own costs, $b_t \equiv \int \mathbb{E}^{AL}_{i,t}\left[\Delta mc^p_{i,t}\right] di$, depends on past nominal marginal cost growth

$$b_t = \tilde{\mathcal{K}}\Delta mc_t + \left(1-\tilde{\mathcal{K}}\right)\tilde{\rho} b_{t-1} \tag{37}$$

and the weight placed on this belief term is

$$\omega_b = \left(1-\alpha^{FIRE}\right)\frac{\left(1-\theta_p\right)\beta\tilde{\rho}}{1-\beta\theta_p\tilde{\rho}}\frac{1-\beta\tilde{\rho}\left(1-\tilde{\mathcal{K}}\right)\theta_p}{1-\beta\tilde{\rho}\left(1-\tilde{\mathcal{K}}\right)\delta} \tag{38}$$

The additional discounting of the future is captured by the term δ which solves the fixed point

$$\delta = \alpha^{FIRE} + \left(1-\alpha^{FIRE}\right)\left(\theta_p - \frac{\left(1-\theta_p\right)\beta\tilde{\rho}\tilde{\mathcal{K}}}{1-\beta\theta_p\tilde{\rho}}\frac{\theta_p-\delta}{1-\beta\tilde{\rho}\left(1-\tilde{\mathcal{K}}\right)\delta}\right) \in \left(\theta_p, 1\right) \tag{39}$$

Proof. From the law of motion for prices (27) and the optimal reset price (28), under Calvo,

$$\pi_t = \frac{1-\theta_p}{\theta_p}\left(\mathrm{mc}^{\mathrm{real}}_t + \sum_{\tau=1}^{\infty}\left(\beta\theta_p\right)^\tau \int \mathbb{E}_{i,t}\left[\Delta \mathrm{mc}_{i,t+\tau}\right] \mathrm{d}i\right) \tag{40}$$

Combining this with the belief model in (12), (13), and (14),

$$\pi_t = \frac{1-\theta_p}{\theta_p}\mathrm{mc}^{\mathrm{real}}_t + \left(1-\alpha^{\mathrm{FIRE}}\right)\frac{\beta\tilde{\rho}\left(1-\theta_p\right)}{1-\beta\theta_p\tilde{\rho}} b_t + \alpha^{\mathrm{FIRE}}\frac{1-\theta_p}{\theta_p}\sum_{\tau=1}^{\infty}\left(\beta\theta_p\right)^\tau \mathbb{E}_t\left[\Delta \mathrm{mc}_{t+\tau}\right] \tag{41}$$

where b_t is defined in (37). In Lemma 3, we show that this system is equivalent to the recursion

$$\pi_t - \omega_b b_t = \left(\frac{1-\theta_p}{\theta_p} + \beta\left(\theta_p-\delta\right)\right)\mathrm{mc}^{\mathrm{real}}_t - (1-\delta)\beta\mathbb{E}_t\left[\mathrm{mc}^{\mathrm{real}}_{t+1}\right] + \beta\delta\mathbb{E}_t\left[\pi_{t+1}-\omega_b b_{t+1}\right] \tag{42}$$

where ω_b satisfies (38) and δ is the stable root solving the fixed point in (39). Since $\beta\delta \in (0,1)$, iterating forward (42), for bounded solutions, we get (36). □

We can build the intuition for this NKPC by exploring each of these three terms:

(1) **Future marginal costs are discounted more quickly.** The first term corresponds to a standard FIRE NKPC but with additional discounting, δ, given by (39).[64] When $\delta < 1$, sufficiently distant shocks are less important relative to FIRE.

(2) **There is higher inflation pass-through on impact.** Since $\frac{1-\delta}{\delta} \geq 0$, the second term shows that current real marginal costs are given an outsized importance. This is because firms can perfectly observe their own current costs.

[64] See that with $\delta = 1$, this first term collapses to the FIRE NKPC: $\pi^{\mathrm{FIRE}}_t = \frac{(1-\beta\theta_p)(1-\theta_p)}{\theta_p}\sum_{\tau=0}^{\infty}\beta^\tau\mathbb{E}_t\left[\mathrm{mc}^{\mathrm{real}}_{t+\tau}\right]$.

(3) **High past cost growth acts like an inflationary cost-push shock.** The average belief in the persistent cost component, b_t, satisfies the recursion (37), and so depends positively on the history of average nominal marginal cost growth.[65] The weight on this backward looking term is ω_b, and this is always weakly positive.

The standard NKPC is recovered immediately by setting $\alpha^{\text{FIRE}} = 1$, and with the standard proportional relationship between real marginal costs and the output gap. In Appendix D.2, we plot the jacobian of the NKPC: the response of inflation at date t to a real marginal cost shock at date s, announced at date 0, and compare it to the NKPC under other models. We again find our model features particularly weak anticipation, but strong pass-through on impact.

Despite average nominal marginal cost beliefs themselves being inertial (**Fact 4**), our calibrated NKPC features little inertia: there are limited effects on inflation in the future from one-off shocks to real marginal costs today. This results from the fact that b_t in (36) does not remain persistently high following a real marginal cost shock under our calibration. To see how this is possible, note that a completely transitory real marginal cost shock first raises and then lowers nominal marginal costs. Under very sticky prices, the negative cost growth is as large as the positive but occurred more recently—thereby playing a larger role in beliefs and pushing against inertia.[66]

The Adaptive Expectations Phillips curve In the special case where firms are fully Adaptive Learners, we can characterize inflation purely as a function of current real marginal costs and past inflation. This is shown in Theorem 4.

Theorem 4. *Under $\alpha^{FIRE} = 0$, our NKPC becomes*

$$\pi_t = \kappa^{AL} mc_t^{real} + \omega_b^{AL} \sum_{\ell=1}^{\infty} \left(\lambda^{AL}\right)^{\ell} \pi_{t-\ell} \tag{43}$$

where $\kappa^{AL} \equiv \frac{1}{1-\mathcal{M}}\left(\frac{1-\theta_p}{\theta_p} + \mathcal{M}\right)$, $\omega_b^{AL} \equiv \frac{\frac{\mathcal{M}}{\theta_p}\left(\frac{(1-\tilde{\mathcal{K}})\tilde{\rho}}{1-\mathcal{M}} - \theta_p\right)}{\frac{1-\theta_p}{\theta_p}(1-\tilde{\mathcal{K}})\tilde{\rho} + \mathcal{M}}$, *and* $\lambda^{AL} \equiv \frac{\frac{1-\theta_p}{\theta_p}(1-\tilde{\mathcal{K}})\tilde{\rho} + \mathcal{M}}{\frac{1-\theta_p}{\theta_p} + \mathcal{M}}$, *and* $\mathcal{M} \equiv \frac{(1-\theta_p)\beta\tilde{\rho}\tilde{\mathcal{K}}}{(1-\beta\theta_p\tilde{\rho})}$ *is a multiplier term.*

Proof. This follows from (41), setting $\alpha^{\text{FIRE}} = 0$, and substituting in for b_t from (37). $\square$

The functional form of (43) matches the "Friedman-Phelps" adaptive-expectations Phillips curve, derived in Appendix C.2. We provide a modern microfoundation for this Phillips curve:

[65] We can also solve out the NKPC to depend only on current and future real marginal costs and b_{t-1}, to emphasize the role of beliefs as a state variable.

[66] This ambiguous prediction for inertia applies to shocks that affect inflation via marginal costs. Markup shocks, by contrast, generally have persistent effects on inflation, because they raise nominal marginal costs, without a subsequent reversal where nominal costs decline. We derive the Phillips curve with markup shocks under Adaptive Learning in Theorem A2, from which it can be shown that transitory markup shocks much more generally have persistent effects on inflation. (Specifically, so long as there is any degree of backward lookingness in belief formation $\tilde{\mathcal{K}} < 1$ and any degree of extrapolation $\tilde{\rho} > 0$.)

our firms optimally set prices subject to a nominal friction, and form beliefs according to an empirically motivated model of individual beliefs. An important implication of this is that the coefficients on current real marginal costs and past inflation depend on underlying structural parameters.

In particular, we show in Appendix C.3 that as the frequency of price adjustment rises (θ_p falls), the Phillips curve becomes steeper (κ^{AL} rises) and more inertial (the weight on past inflation rises). Under more flexible prices, real marginal cost shocks have larger lasting effects on nominal marginal costs.[67] This translates into persistently higher beliefs about cost growth.

In the post-2010s United States, we have estimates for the parameters. Using our calibration from Section 2.2, $\beta = 0.99$ and $\theta_p = 0.75$, we obtain

$$\kappa^{\text{AL}} \approx 0.5, \qquad \omega_b^{\text{AL}} \approx 0.0, \qquad \lambda^{\text{AL}} \approx 0.7$$

So we predict larger pass-through than under FIRE ($\kappa^{\text{FIRE}} \approx 0.1$). We also predict no inflation inertia.

This Phillips curve was derived with only a persistent and transitory cost component (though it nests the case of a permanent and transitory cost component). In Appendix C.4, we show that if firms are informed about permanent shocks, (43) characterizes the path of inflation around the long-run level of inflation.[68]

3.3 Testing the NKPC in pricing data

Our calibrated NKPC is steeper and less forward-looking than the FIRE NKPC. At the same time, it has much less inertia than the traditional Adaptive Expectations Phillips curve. We now recall the result in Hazell et al. (2022) (hereafter, HHNS) that the Phillips curve slope is much steeper (and in fact consistent with our calibration) when estimated under a null of no forward-lookingness. Further, we show that an Adaptive Expectations Phillips curve estimated using the HHNS methodology exhibits little inertia over our calibration sample period, also consistent with our calibrated NKPC.

HHNS develop an instrumental variables approach to estimate the NKPC in state-level data. We start by using their approach to estimate the Adaptive Expectations Phillips curve (AEPC) in (43). As we have merely changed the regression equation but used the same instruments as HHNS, we leave the details in Appendix E.1. We find, perhaps surprisingly, that the slope of the AEPC is quite high – consistent with the slope in our calibrated model – and that there is limited inertia. We discuss these results in turn.

Our calibration suggests essentially no inertia in the Phillips curve, at least over the BIE sample period starting in 2011. But the degree of inertia in our model varies with key model parameters:

[67] This holds true even in the FIRE NKPC, which features no inflation inertia.

[68] In Appendix D.3, we show that if firms learn about permanent cost shocks by observing their past costs, this results in the same long-run inflation-output trade-off as under FIRE. Such a model also generates a larger weight on the sum of the past inflation terms, though each individual weight may remain small.

the frequency of price adjustment, the perceived persistence of costs, the speed of belief updating. In Appendix E.1, we split the HHNS sample at the year 1990 and find inertia only in the pre-1990 sample, with no inflation inertia post-1990. The lack of inflation inertia post-1990 is consistent with our calibrated model, and our model further suggests that the decline in the frequency of price adjustment over this period (see Nakamura, Steinsson, Sun and Villar (2018)) could help explain the decline in inertia over time.

The surprisingly steep slope we estimated is already discussed in HHNS's original work. HHNS identify the slope of the NKPC in a regression of inflation on a discounted sum of future unemployment rates. The authors discuss how the estimated slope is highly sensitive to the choice of discount factor when computing the unemployment rate regressor; higher discount rates (i.e. less forward lookingness) yield higher slope estimates. HHNS do not estimate the degree of forward-lookingness, and under myopia as substantial as our α^{FIRE} estimate suggests, their results are consistent with those in this paper.[69] Notably, Gagliardone and Tielens (2026) estimate the degree of myopia in price setting directly and find a lack of forward lookingness just as substantial as suggested by our α^{FIRE} estimates.[70]

4 Aggregate implications of Adaptive Learning

4.1 Supply and Demand shocks in a Simple New Keynesian model

We now explore the implications of our calibrated beliefs in a simple macroeconomic model. This corresponds to the "basic New Keynesian model" in Galí (2015) in the special case of constant returns to scale, but with the addition of sticky wages and assuming a constant real rate monetary rule.

Consumption, and hence output, follows a standard Euler equation

$$y_t = -\text{eis} \cdot r_t + \mathbb{E}_t y_{t+1} \tag{44}$$

Deviations in the real interest rate are entirely driven by monetary policy shocks,[71]

$$r_t = i_t^{\text{shock}} \tag{45}$$

which have persistence ρ_{MP}: $i_t^{\text{shock}} = \rho_{\text{MP}} i_{t-1}^{\text{shock}} + \varepsilon_t^{\text{MP}}$. Nominal wages, w_t, are set by unions subject

[69] Ideally, we would test forward-lookingness directly in HHNS's framework, but this task would require an additional instrument meeting F-stat requirements for predicting state-level recessions at least one quarter in advance.

[70] Gagliardone and Tielens (2026) estimate a monthly discount factor of 0.82, or 0.09 annualized. This is even more myopic than the α^{FIRE} estimate of 0.25 we obtained from our calibration in section 2.2 (which was on an annual basis).

[71] The Taylor rule condition for inflation responsiveness to ensure determinacy is the same in our model as under FIRE: $\phi_\pi > 1$. The real rate rule employed here is useful for holding fixed the paths of output across models but involves an exactly unit weight on expected future inflation. To ensure determinacy, one can view this rule as the limiting case of a rule that responds slightly more than one-to-one.

Figure 6: Inflation response to supply and demand shocks

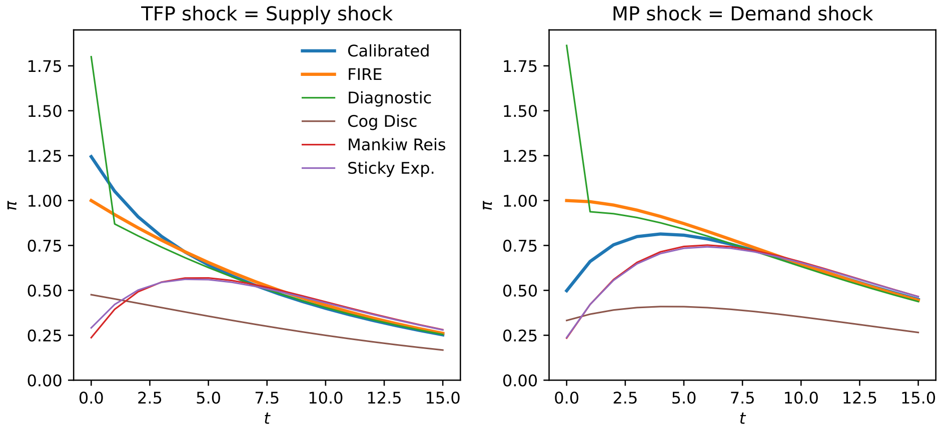

Note: The response of inflation to a TFP shock (left) and to a monetary policy shock (right) under different models of firm beliefs. Impulse responses normalized to 1 under FIRE. Our model, calibrated in Section 2.2 above, is shown in blue.

to a Calvo wage setting friction,

$$\pi_t^w = \kappa^w \left(\frac{1}{\text{frisch}} n_t + \frac{1}{\text{eis}} y_t - w_t^{\text{real}} \right) + \beta \mathbb{E}_t \left[\pi_{t+1}^w \right] \tag{46}$$

where $y_t = a_t + n_t$. Nominal marginal costs are

$$\text{mc}_t = w_t - a_t \tag{47}$$

where TFP has persistence ρ_a: $a_t = \rho_a a_{t-1} + \varepsilon_t^a$. Prices are set subject to a Calvo friction, so that inflation is

$$\pi_t = \frac{1-\theta_p}{\theta_p} \left(p_t^* - p_t \right) \tag{48}$$

where the optimal reset price is

$$p_t^* = \left(1 - \beta\theta_p\right) \sum_{\tau=0}^{\infty} \left(\beta\theta_p\right)^\tau \mathbb{E}_t^f \left[\text{mc}_{t+\tau}\right] \tag{49}$$

To close the model, we simply need to specify average firm beliefs, $\mathbb{E}_t^f$.

The model has two shocks: a TFP shock ε_t^a and a monetary policy shock $\varepsilon_t^{\text{MP}}$. The essential feature of this model is that, in the presence of nominal wage rigidity, monetary policy shocks are slow to propagate into costs, while TFP shocks materialize more quickly in costs. To calibrate the model we set $\beta = 0.99$, frisch $= 0.5$, eis $= 1$, $\theta_p = 0.9$, and we set wage stickiness so that the slopes of the wage and price Phillips curves are equal under FIRE: $\kappa^w = \kappa^{p,\text{FIRE}}$.

In Figure 6 we show the impulse response of inflation under our calibrated model of firm beliefs (in blue), and under FIRE (in orange).[72] In our model, firms "overreact" to the TFP shock, moving prices very quickly on impact. This is because the shock's implications for cost growth are in fact relatively transitory, and firms overestimate the persistence. On the other hand, firms underreact to the monetary policy shock. Monetary policy affects costs via wages, so in the presence of sticky wages, much of the initial response under FIRE is because firms appreciate that costs will rise in the future. However, our firms on the whole do not adjust their beliefs until their costs move, and so initially prices move less. While there is a clear difference in inflation timing shown in the impulse responses across models, the cumulative amount of inflation can also vary. Intuitively, faster price responses set off a wage price spiral, and so can lead to a larger overall movement in the price level. The long-run price level is higher under FIRE than our Calibrated model for demand shocks but lower for supply shocks.

In Figure 6 we also show the impulse responses under the other models of firm beliefs we introduced in Section 3.1 above. The "dampening" theories generate smaller reactions to both shocks, while the "amplifying" theory generates larger on impact reactions to both. Our model is unique in its asymmetric prediction: demand shocks propagate into inflation slower than under FIRE, but supply shocks propagate faster.

This result is much more general than the simple New Keynesian model we have presented here and provides an intuition for differential shock pass-through in many cases. In a production network, upstream shocks are mediated by the sticky prices of the network, leading them to resemble our demand shock example. In contrast, downstream shocks can still rapidly affect final-good producing firms' costs, leading to amplified effects on CPI. In Appendix D.7, we also show that our results hold for supply and demand shocks in a medium scale New Keynesian model and with estimated persistences.

Nonlinear shock propagation is also amplified by our mechanism. For example, suppose that small demand shocks lead to limited but also slow responses of costs. However, very large demand shocks lead to shortages along the production network so that, even without any adjustment from wages, costs rapidly start rising. In this scenario, we would argue that while small demand shocks have a dampened and slowed effect on inflation, large ones have an amplified and accelerated effect on inflation. Relatedly, decreasing returns to scale in production would provide an offsetting force to the demand shock underreaction in our model; expansionary monetary policy affects costs immediately by moving production to a higher point on an upward sloping supply curve, and overreaction to this force counteracts firms' underreaction to future wage increases.[73]

A central feature driving the differential reaction to supply and demand shocks in our model is that firms react largely as a function of their costs moving, and that they do not sufficiently vary this reaction across shocks. In Appendix D.5, we show that this prediction can also be consistent with a model of limited-information rational expectations (LIRE) in which firms observe only their own current and past costs but are otherwise rational. The model will be identical to the one

[72] Output is identical across all our models for these shocks due to the real rate rule. In Appendix D.4 we also show our results if the central bank follows a typical Taylor rule.

[73] However, in Appendix A.10 we showed that production appears close to constant returns to scale in the BIE.

above, but instead of calibrating the behavior of beliefs, we impose that firms choose the values of $\tilde{\rho}$ and $\tilde{\mathcal{K}}$ implied by (constrained) rational expectations.

4.2 Strong near-term forward guidance

The demand shock example in Figure 6 shows that monetary policy operates differently in our model than under FIRE. In Theorem 5, we focus on the effects of forward guidance. To simplify the analysis, we compare FIRE to a model with fully Adaptive Learning, and we make wages flexible. It follows that real marginal costs are proportional to output: $\text{mc}_t^{\text{real}} = \varphi y_t$, where $\varphi \equiv \frac{1}{\text{frisch}} + \frac{1}{\text{eis}}$. We find that forward guidance about the near term is *more* powerful, while forward guidance about the long term is less powerful and does not feature the inflation puzzle associated with FIRE, despite households remaining fully rational and informed.

Theorem 5. *Consider the standard New Keynesian model set out in Section 4.1, and with flexible wages. Suppose that at date* 0*, the central bank announces a one-off interest rate cut at date h. The effect on inflation at date* 0*, under FIRE and Adaptive Learning (AL), respectively, are*

$$-\frac{d\pi_0^{FIRE}}{dr_h} = eis \cdot \varphi\kappa^{FIRE}\frac{1-\beta^{h+1}}{1-\beta}, \qquad -\frac{d\pi_0^{AL}}{dr_h} = eis \cdot \varphi\kappa^{AL}.$$

Proof. Iterating the IS equation forward and imposing the boundary condition $\lim_{H\to\infty} y_H = 0$, we have $y_t = -\text{eis}\sum_{\tau=t}^{\infty} r_\tau$. Therefore, $\frac{dy_\tau}{dr_h} = -\text{eis}\cdot 1_{\tau\leq h}$. The FIRE NKPC in period 0 can be solved forward using the boundary condition $\lim_{H\to\infty}\beta^H\pi_H = 0$, yielding $\pi_0 = \varphi\kappa^{FIRE}\sum_{\tau=0}^{\infty}\beta^\tau y_\tau$. Taking the derivative, $\frac{d\pi_0}{dr_h} = \varphi\kappa^{FIRE}\sum_{\tau=0}^{\infty}\beta^\tau\frac{dy_\tau}{dr_h} = -\text{eis}\cdot\varphi\kappa^{FIRE}\frac{1-\beta^{h+1}}{1-\beta}$. The Adaptive Learning NKPC in period 0 is not forward looking and need not be solved forward. It straightforwardly yields $\frac{d\pi_0}{dr_h} = -\text{eis}\cdot\varphi\kappa^{AL}$. □

As households are still FIRE, we have not solved the output component of the forward guidance puzzle: there is no discounting in the IS equation, and announced interest rate cuts far in the future cause output to jump on impact of the news. Given that the magnitude of this jump is invariant to firms' beliefs, the steeper NKPC in our model is therefore a force for stronger forward guidance. This is a key difference from a model in which firms are completely myopic in their pricing behavior – while AL firms are not rationally forward-looking, their extrapolative beliefs generate a steeper NKPC. However, there is a countervailing force. Rationally forward-looking firms understand that, as the central bank increases the horizon of forward guidance, output will remain high for longer, and this force additionally boosts inflation today. Under Adaptive Learning, this forward-looking channel is shut down, and the inflation response to forward guidance is invariant to the horizon.[74]

Comparing our model to FIRE, short horizon forward guidance is more powerful under Adaptive Learning, and long horizon forward guidance is dampened. To derive the short horizon re-

[74] This invariance property is shared by a model in which firms are completely myopic, but inflation beyond period 0 would still differ under Adaptive Learning due to inflation inertia.

sult, use Theorem 5 evaluated at $h = 0$, which intuitively implies that the response of inflation under Adaptive Learning is larger than under FIRE if and only if $\kappa^{AL} > \kappa^{FIRE}$. For the long horizon, it also follows from Theorem 5 that, so long as $\beta \geq 1 - \kappa^{FIRE}/\kappa^{AL}$, there exists h large enough that $-\frac{d\pi_0^{\text{FIRE}}}{dr_h} \geq -\frac{d\pi_0^{\text{AL}}}{dr_h}$.[75] As $\beta \approx 0.99$ and $1 - \kappa^{FIRE}/\kappa^{AL} \approx 0.8$ on a quarterly basis, both the short-run and long-run conditions hold.

In Appendix Figure D.19, we show that Theorem 5 remains true when we make wages imperfectly flexible and use our fully calibrated beliefs model: policy is more effective under our beliefs than under FIRE up to a horizon of 4 quarters but is weaker thereafter. Our result is also robust to the addition of household myopia in the IS equation. The degree of household forward-lookingness plays a central role in forward guidance, and, in fact, the degree of household myopia is even more important in our model than under FIRE—suggesting a powerful complementarity between behavioral biases across agents. Nevertheless, in Appendix Theorem A1, we show that household myopia weakens but does not overturn our result that near-term forward guidance is stronger with Adaptive Learning firms.

4.3 Optimal monetary policy

We now turn to optimal monetary policy. To consider the canonical setting, we allow wages to be flexible. We break divine coincidence in our model by introducing cost-push shocks exactly as in Galí (2015) Chapter 5.2. See Appendix C.6 for further details. In Appendix Theorem A2, we derive the first-order conditions for optimal monetary policy under Adaptive Learning relative to FIRE.

We showed that our model of firm's beliefs had subtle predictions for the power of forward guidance. The predictions for commitment are starker. Since firms are not forward looking under Adaptive Learning, central bank promises about future output do not affect firms' inflation expectations today. Therefore the optimal policy does not vary whether solved under commitment or discretion.[76]

We now quantitatively compare the optimal policy reaction in our fully calibrated model and under FIRE. The frequency is quarterly and we keep the same calibration as above, except for making wages flexible. In Figure 7, we show the responses of inflation, output, and the real interest rate to a cost push shock, given the optimal monetary policy reaction. In order to focus on our model's implications for how the central bank should react for a *given* forecast of inflation and output, we normalize the markup shocks across models so that they induce the same inflation path if there is no active monetary response ($r_t = 0$). We choose $\{\mu_t\}$ to achieve an AR(1) path for inflation under $r_t = 0$, with persistence 0.9. And we set the central bank's weight on the output gap in the welfare function to be $\vartheta = 0.05$.

We see that under commitment, the output response under FIRE (orange, solid) is much more

[75] Note first that $-\frac{d\pi_0^{\text{FIRE}}}{dr_h}$ is monotonically increasing in h. Second, $\lim_{h\to\infty} -\frac{d\pi_0^{\text{FIRE}}}{dr_h} = \frac{eis\cdot\varphi\kappa^{FIRE}}{1-\beta}$. Comparing this limit to $-\frac{d\pi_0^{\text{AL}}}{dr_h} = \text{eis}\cdot\varphi\kappa^{AL}$ completes the proof.

[76] While such invariance contrasts starkly with FIRE, it is shared with a model in which firms are completely myopic.

Figure 7: Optimal inflation and output response to a persistent markup shock

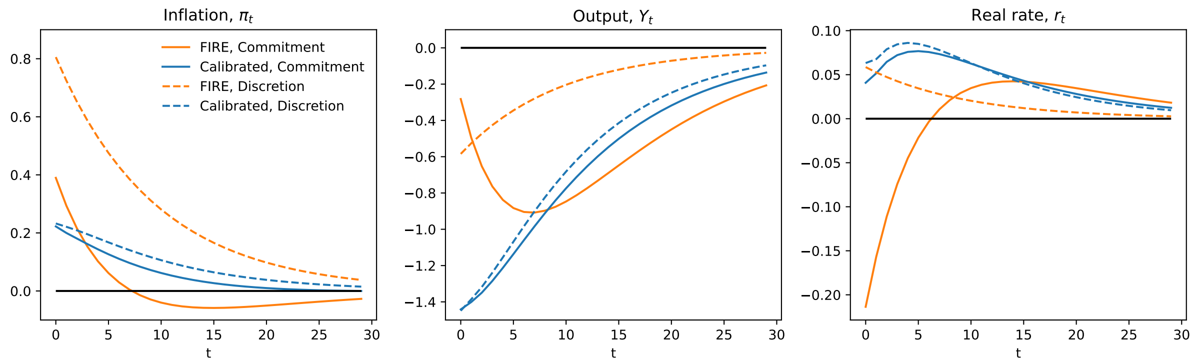

Note: Impulse responses for inflation, output, and the real interest rate in response to a markup shock. The responses are normalized so that inflation is identical and equal to an AR(1) with persistence 0.9 in all models if the central bank holds the real interest rate constant.

delayed than that under our Calibrated model (blue, solid).[77] Turning to discretion makes the response under FIRE much smaller and more front-loaded (orange, dashed), but has little effect on optimal policy in our Calibrated model (blue, dashed). Overall, our model suggests limited benefits of commitment over discretion and stresses the importance of a rapid monetary policy response.

Finally, we turn to the effects of interest rate shocks today on future inflation. As such shocks amount to a one-off shock to output in period 0 via the IS equation, the inflation effects of such policies can be read straight off the Phillips curve. The following result is therefore essentially a corollary of the FIRE NKPC and our derivation of the Adaptive Learning Phillips curve in Theorem 4.

Corollary 1. *Consider the standard New Keynesian model set out in Section 4.1 and with flexible wages. Suppose that the central bank raises the real rate at date 0 only. The effect on inflation at date $h > 0$, under FIRE and Adaptive Learning (AL), respectively, are*

$$\frac{d\pi_h^{FIRE}}{dr_0} = 0 \quad \text{and} \quad \frac{d\pi_h^{AL}}{dr_0} = -eis \cdot \varphi \kappa^{AL} \left(\kappa^{AL,lag}\right)^h,$$

with $\kappa^{AL,lag} \equiv \omega_b^{AL} \lambda^{AL} = \frac{\mathcal{M}}{\frac{1-\theta_p}{\theta_p} + \mathcal{M}} \left(\frac{(1-\tilde{\mathcal{K}})\tilde{\rho}}{\theta_p(1-\mathcal{M})} - 1\right)$. Inducing a recession today helps to fight inflation tomorrow, $\frac{d\pi_h^{AL}}{dr_0} < 0$, if and only if beliefs are sufficiently backward looking,

$$1 - \tilde{\mathcal{K}} > \frac{(1-\mathcal{M})\,\theta_p}{\tilde{\rho}}$$

Proof. Follow the same proof as in Theorem 5, with $y_0 = -eis \cdot 1_{h=0}$. The FIRE result follows immediately from the FIRE NKPC, and the AL result follows from iterating the AEPC (43) backwards. □

[77] Also note that the traditional result of price level targeting holds under FIRE beliefs but breaks in our behavioral model.

Under Adaptive Learning beliefs, the central bank may be able to lean against anticipated future inflation by hiking rates today. This result does not hold under FIRE. The idea is that by creating a recession and hence low inflation today, firms enter the next period with low inflation expectations, helping to fight against inflationary forces. This is intuitively the same story as emerges from an adaptive expectations model in which $\tilde{\mathbb{E}}_t \pi_{t+1} = \pi_{t-1}$. However, as we saw in Section 3, microfounding this mechanism uncovered a countervailing force. Firms extrapolate in growth rates of *nominal* marginal costs, and engineering a transitory recession with the real rate entails sequentially offsetting movements in nominal marginal costs – with a greater offset when prices are stickier. Over the sample period in our survey data, we found little evidence for inertia, suggesting the central bank should typically not lean against future inflation for this reason.

In Section 3.3 we did find evidence for some inflation inertia pre-1990, when prices were more flexible, exactly as our model would predict. Further, several papers argue that the frequency of price adjustment rose during Covid (see Montag and Villar Vallenas (2025), Cavallo, Lippi and Miyahara (2024), and Gagliardone, Gertler, Lenzu and Tielens (2025)). Our model suggests this may have led shocks during these periods to have longer-lasting effects but likewise allowed monetary policy to do the same.

5 Conclusion

We provide novel empirical results on a key object for macroeconomic propagation: firms' forecasts of their own future costs. We show that these forecasts differ from CPI forecasts and are overly reliant on a firm's own past costs, leading firms to overreact to shocks that affect costs quickly and to underreact to shocks that affect costs slowly. In terms of policy, the power of forward guidance is more potent in the near term than in the standard model but grows much less over time. Meanwhile, commitment is uniformly less powerful.

References

Acosta, Miguel, Connor M Brennan, and Margaret M Jacobson, "Constructing high-frequency monetary policy surprises from SOFR futures," *Economics Letters*, 2024, *242*, 111873.

Afrouzi, Hassan, "Strategic inattention, inflation dynamics, and the non-neutrality of money," 2023.

___ **and Choongryul Yang**, "Dynamic rational inattention and the Phillips curve," 2021.

___ **, Spencer Y Kwon, Augustin Landier, Yueran Ma, and David Thesmar**, "Overreaction in expectations: Evidence and theory," *The Quarterly Journal of Economics*, 2023, *138* (3), 1713–1764.

Albagli, Elias, Francesco Grigoli, and Emiliano Luttini, *Inflation Expectations and the Supply Chain*, International Monetary Fund, Research Department, 2022.

Andrade, Philippe, Alexander Dietrich, John Leer, Raphael Schoenle, Jenny Tang, and Egon Zakrajsek, "The Future in Today's Prices: Evidence from a Survey of U.S. Firms," 2026.

___ **, Olivier Coibion, Erwan Gautier, and Yuriy Gorodnichenko**, "No firm is an island? How industry conditions shape firms' expectations," *Journal of Monetary Economics*, 2022, *125*, 40–56.

Angeletos, George-Marios and Zhen Huo, "Myopia and anchoring," *American Economic Review*, 2021, *111* (4), 1166–1200.

___ **,** ___ **, and Karthik A Sastry**, "Imperfect macroeconomic expectations: Evidence and theory," *NBER Macroeconomics Annual*, 2021, *35* (1), 1–86.

Baumann, Ursel, Annalisa Ferrando, Dimitris Georgarakos, Yuriy Gorodnichenko, and Timo Reinelt, "SAFE to Update Inflation Expectations? New Survey Evidence on Euro Area Firms," Working Paper 32504, National Bureau of Economic Research May 2024.

Bellifemine, Marco, Adrien Couturier, and Jacopo Tozzo, "The Sentiment Channel of Fiscal Policy," *Available at SSRN 5954294*, 2025.

Blinder, Alan S, "On sticky prices: academic theories meet the real world," in "Monetary policy," The University of Chicago Press, 1994, pp. 117–154.

Bloom, Nicholas, Mihai A Codreanu, and Robert A Fletcher, "Rationalizing Firm Forecasts," Technical Report, National Bureau of Economic Research 2025.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, "Diagnostic expectations and credit cycles," *The Journal of Finance*, 2018, *73* (1), 199–227.

___ , ___ , Yueran Ma, and Andrei Shleifer, "Overreaction in macroeconomic expectations," *American Economic Review*, 2020, *110* (9), 2748–2782.

Born, Benjamin, Zeno Enders, Manuel Menkhoff, Gernot J Müller, and Knut Niemann, "Firm expectations and news: micro v macro," 2022.

Broer, Tobias and Alexandre N Kohlhas, "Forecaster (mis-) behavior," *Review of Economics and Statistics*, 2022, pp. 1–45.

Cagan, Phillip, "The Monetary Dynamics of Hyperinflation," *Studies in the Quantity Theory of Money*, 1956, pp. 25–117.

Candia, Bernardo, Olivier Coibion, and Yuriy Gorodnichenko, "The macroeconomic expectations of firms," in "Handbook of Economic Expectations," Elsevier, 2023, pp. 321–353.

Carroll, Christopher D, Edmund Crawley, Jiri Slacalek, Kiichi Tokuoka, and Matthew N White, "Sticky expectations and consumption dynamics," *American economic journal: macroeconomics*, 2020, *12* (3), 40–76.

Cattaneo, Matias D, Richard K Crump, Max H Farrell, and Yingjie Feng, "On binscatter," *American Economic Review*, 2024, *114* (5), 1488–1514.

Cavallo, Alberto, Francesco Lippi, and Ken Miyahara, "Large shocks travel fast," *American Economic Review: Insights*, 2024, *6* (4), 558–574.

Chao, John C and Norman R Swanson, "Consistent estimation with a large number of weak instruments," *Econometrica*, 2005, *73* (5), 1673–1692.

Chodorow-Reich, Gabriel, Adam M Guren, and Timothy J McQuade, "The 2000s housing cycle with 2020 hindsight: A neo-kindlebergerian view," *Review of Economic Studies*, 2024, *91* (2), 785–816.

Coibion, Olivier and Yuriy Gorodnichenko, "Information rigidity and the expectations formation process: A simple framework and new facts," *American Economic Review*, 2015, *105* (8), 2644–2678.

___ , ___ , and Saten Kumar, "How do firms form their expectations? new survey evidence," *American Economic Review*, 2018, *108* (9), 2671–2713.

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

___ , ___ , Saten Kumar, and Jane Ryngaert, "Do you know that I know that you know…? Higher-order beliefs in survey data," *The Quarterly Journal of Economics*, 2021, *136* (3), 1387–1446.

D'Acunto, Francesco, Ulrike Malmendier, Juan Ospina, and Michael Weber, "Exposure to grocery prices and inflation expectations," *Journal of Political Economy*, 2021, *129* (5), 1615–1639.

Dovern, Jonas, Lena Sophia Müller, and Klaus Wohlrabe, "Local information and firm expectations about aggregates," *Journal of Monetary Economics*, 2023, *138*, 1–13.

Driscoll, John C. and Aart C. Kraay, "Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data," *The Review of Economics and Statistics*, 1998, *80* (4), 549–560.

Eusepi, Stefano and Bruce Preston, "Expectations, learning, and business cycle fluctuations," *American Economic Review*, 2011, *101* (6), 2844–2872.

___ and ___ , "The science of monetary policy: An imperfect knowledge perspective," *Journal of Economic Literature*, 2018, *56* (1), 3–59.

Evans, George W and Seppo Honkapohja, *Learning and expectations in macroeconomics*, Princeton University Press, 2001.

Farmer, Leland E, Emi Nakamura, and Jón Steinsson, "Learning about the long run," *Journal of Political Economy*, 2024, *132* (10), 3334–3377.

Fiore, Fiorella De, Marco Jacopo Lombardi, and Giacomo Mangiante, "The Asymmetric and Heterogeneous Pass-Through of Input Prices to Firms' Expectations and Decisions," CEPR Discussion Paper 20835, Centre for Economic Policy Research 2025.

Friedman, Milton, "The role of monetary policy," *The American economic review*, 1968, *58* (1), 1–17.

Gabaix, Xavier, "A behavioral New Keynesian model," *American Economic Review*, 2020, *110* (8), 2271–2327.

Gagliardone, Luca and Joris Tielens, "Dynamic Pricing under Information Frictions: Evidence from Firm-level Subjective Expectations," *Available at SSRN 6277779*, 2026.

___ , Mark Gertler, Simone Lenzu, and Joris Tielens, "Anatomy of the Phillips Curve: Micro Evidence and Macro Implications," Technical Report, National Bureau of Economic Research 2023.

___ , ___ , ___ , and ___ , "Micro and macro cost-price dynamics in normal times and during inflation surges," Technical Report, National Bureau of Economic Research 2025.

Galí, Jordi, *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*, Princeton University Press, 2015.

Hazell, Jonathon, Juan Herreño, Emi Nakamura, and Jón Steinsson, "The slope of the Phillips Curve: evidence from US states," *The Quarterly Journal of Economics*, 2022, *137* (3), 1299–1344.

Keynes, John Maynard, *The General Theory of Employment, Interest, and Money*, London: Macmillan, 1936.

Kohler, Thomas, Jean-Paul L'Huillier, Gregory Phelan, and Maximilian Weiß, "Why Do Supply Disruptions Lead to Inflation? Survey Evidence from the COVID Pandemic," 2025.

Kumar, Saten, Hassan Afrouzi, Olivier Coibion, and Yuriy Gorodnichenko, "Inflation targeting does not anchor inflation expectations: Evidence from firms in New Zealand," *Brookings papers on economic activity*, 2015, *2015* (2), 151–225.

——, Yuriy Gorodnichenko, and Olivier Coibion, "The effect of macroeconomic uncertainty on firm decisions," *Econometrica*, 2023, *91* (4), 1297–1332.

La'O, Jennifer and Alireza Tahbaz-Salehi, "Optimal monetary policy in production networks," *Econometrica*, 2022, *90* (3), 1295–1336.

Ma, Yueran, Tiziano Ropele, David Sraer, and David Thesmar, "A quantitative analysis of distortions in managerial forecasts," Technical Report, National Bureau of Economic Research 2020.

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

Mankiw, N Gregory and Ricardo Reis, "Sticky information versus sticky prices: a proposal to replace the New Keynesian Phillips curve," *The Quarterly Journal of Economics*, 2002, *117* (4), 1295–1328.

Mavroeidis, Sophocles, Mikkel Plagborg-Møller, and James H Stock, "Empirical evidence on inflation expectations in the New Keynesian Phillips Curve," *American Economic Journal: Journal of Economic Literature*, 2014, *52* (1), 124–188.

Maxted, Peter, "A macro-finance model with sentiment," *Review of Economic Studies*, 2024, *91* (1), 438–475.

Meyer, Brent H and Xuguang Simon Sheng, "Unit cost expectations: Firms' perspectives on inflation," *European Economic Review*, 2025, p. 104955.

——, Brian Prescott, and Xuguang Simon Sheng, "The impact of the COVID-19 pandemic on business expectations," *International Journal of Forecasting*, 2022, *38* (2), 529–544.

Minton, Robert and Brian Wheaton, "Delayed Inflation in Supply Chains: Theory and Evidence," *Available at SSRN 4470302*, 2023.

Montag, Hugh and Daniel Villar Vallenas, "Post-pandemic price flexibility in the US: evidence and implications for price setting models," 2025.

Nakamura, Emi and Jón Steinsson, "Fiscal stimulus in a monetary union: Evidence from US regions," *American Economic Review*, 2014, *104* (3), 753–792.

—— and ——, "High-frequency identification of monetary non-neutrality: the information effect," *The Quarterly Journal of Economics*, 2018, *133* (3), 1283–1330.

——, ——, Patrick Sun, and Daniel Villar, "The elusive costs of inflation: Price dispersion during the US great inflation," *The Quarterly Journal of Economics*, 2018, *133* (4), 1933–1980.

Phelps, Edmund S, "Phillips curves, expectations of inflation and optimal unemployment over time," *Economica*, 1967, pp. 254–281.

Samuelson, Paul A and Robert M Solow, "Analytical aspects of anti-inflation policy," *The American economic review*, 1960, *50* (2), 177–194.

Savignac, Frédérique, Erwan Gautier, Yuriy Gorodnichenko, and Olivier Coibion, "Firms' inflation expectations: New evidence from France," *Journal of the European Economic Association*, 2024, p. jvae015.

Slobodyan, Sergey and Raf Wouters, "Learning in a medium-scale DSGE model with expectations based on small forecasting models," *American Economic Journal: Macroeconomics*, 2012, *4* (2), 65–101.

Smets, Frank and Rafael Wouters, "Shocks and frictions in US business cycles: A Bayesian DSGE approach," *American economic review*, 2007, *97* (3), 586–606.

Stock, James and Motohiro Yogo, "Asymptotic distributions of instrumental variables statistics with many instruments," *Identification and inference for econometric models: Essays in honor of Thomas Rothenberg*, 2005, *6*, 109–120.

Taubinsky, Dmitry, Luigi Butera, Matteo Saccarola, and Chen Lian, "Beliefs about the economy are excessively sensitive to household-level shocks: Evidence from linked survey and administrative data," Technical Report, National Bureau of Economic Research 2024.

Weber, Michael, Francesco D'Acunto, Yuriy Gorodnichenko, and Olivier Coibion, "The subjective inflation expectations of households and firms: Measurement, determinants, and implications," *Journal of Economic Perspectives*, 2022, *36* (3), 157–184.

Werning, Iván, "Expectations and the Rate of Inflation," Technical Report, National Bureau of Economic Research 2022.

Online Appendix

A Robustness and extensions in the BIE data

A.1 Example question from the BIE

In Figure A.1 we show the survey question about future cost forecasts.

Figure A.1: BIE question eliciting cost growth expectations over the next year

Projecting ahead, to the best of your ability, please assign a percent likelihood to the following changes to UNIT COSTS over the next twelve months. (Values should sum to 100%)

For example, if you think each of these is equally likely, you might answer 20% for each:

Unit costs down (less than -1%)	20
Unit costs about unchanged (-1% to 1%)	20
Unit costs up somewhat (1.1% to 3%)	20
Unit costs up significantly (3.1% to 5%)	20
Unit costs up very significantly (more than 5%)	20
Total	100

Unit costs down (less than -1%)	0	%
Unit costs about unchanged (-1% to 1%)	0	%
Unit costs up somewhat (1.1% to 3%)	0	%
Unit costs up significantly (3.1% to 5%)	0	%
Unit costs up very significantly (more than 5%)	0	%
Total	0	%

Source: FRBA Business Inflation Expectations (BIE) Survey.

A.2 Special question in the BIE

Firms are first asked, as part of the regular monthly BIE survey, to report their expectations for own cost growth over the next year:

Projecting ahead, to the best of your ability, please assign a percent likelihood to the following changes to UNIT COSTS over the next twelve months. (Values should sum to 100%)

After answering some other regular questions, they are asked the Special Questions. First we ask:

The **Consumer Price Index (CPI)** is a measure of the average change over time in the prices paid by urban consumers for a market basket of consumer goods and services.

Projecting ahead, to the best of your ability, what do you expect the aggregate rate of inflation, as measured by the CPI, will be over the next four quarters?

We then give firms the following information:

The **Blue Chip forecast** is the average forecast of more than 50 economists employed across some of the largest and most well-regarded U.S. manufacturers, banks, insurance companies, and brokerage firms.

- The Blue Chip four-quarters-ahead inflation forecast is computed from the April Blue Chip forecasts of quarterly inflation in Q2, Q3, and Q4 of 2025 and Q1 of 2026.

The April release of the Blue Chip forecast predicts CPI inflation will be 3.4% over the next four quarters.

Once firms have been given this information, we ask again about their CPI and own-cost forecasts, to see if there is a change. We first ask about CPI:

Does this information affect your beliefs for CPI inflation over the next four quarters?

Only if they respond 'Yes', we then ask

Your previous forecast for CPI inflation over the next four quarters was x%.

Please give your updated forecast in the box below.

where x is filled in from their previous response. We then ask about their own-cost forecast:

Earlier, you gave us estimates for your firm's unit cost expectations.

Does the professional forecast that CPI inflation will be 3.4% over the next four quarters affect your forecast about your firm's unit cost growth over the next 12 months?

Only if they respond 'Yes', we then ask

Below, you will see the same unit cost question as earlier in this survey, but with your previous answers in parentheses.

What are your new forecasts for changes in your firm's unit costs over the next 12 months?

Figure A.2: Average BIE responses align well with macroeconomic data

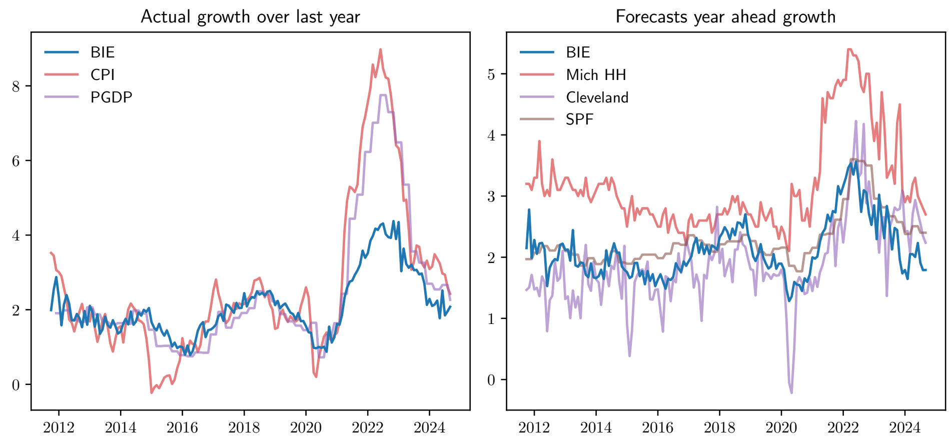

Note: BIE is employment-weighted cross-sectional average of $\Delta_{12}\text{mc}_{i,t}$ (LHS) and $\mathbb{E}_{i,t}[\Delta_{12}\text{mc}_{i,t+12}]$ (RHS). CPI and PGDP are the year on year percent change. Mich HH is the median expected price change over next year from Michigan Surveys of Consumers. Cleveland is 1 year expected inflation from the Federal Reserve Bank of Cleveland. SPF is 1 year ahead CPI forecast.

A.3 Comparing the BIE to aggregate time series

As previously argued in Meyer and Sheng (2025), we show that averages of BIE responses correlate strongly with published aggregate inflation statistics and expectations. We visualize these results for our sample in Figure A.2. In the left panel, we show that average reported year-on-year cost growth in the BIE correlates strongly with published year-on-year growth in the CPI[78] and GDP deflator.[79] In particular, the GDP deflator and the BIE co-move very strongly in the pre-Covid period, where the truncation is less binding. In the right panel, we show that average expected cost growth over the next year correlates strongly with analogous expectations measured in the University of Michigan's Survey of Consumers,[80] the Federal Reserve Bank of Philadelphia's Survey of Professional Forecasters,[81] and the Federal Reserve Bank of Cleveland's inflation expectation measure.[82]

A.4 Cost beliefs and pricing in the BIE

As part of the special questions in the BIE, firms were occasionally asked to forecast their price growth over the next twelve months, and to report their price growth over the past twelve months.[83] A typical question framing was

[78] U.S. Bureau of Labor Statistics, Consumer Price Index for All Urban Consumers: All Items in U.S. City Average [CPIAUCSL], retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/CPIAUCSL, June 26, 2026.

[79] U.S. Bureau of Economic Analysis, Gross Domestic Product: Implicit Price Deflator [GDPDEF], retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/GDPDEF, June 26, 2026.

[80] Surveys of Consumers, University of Michigan, University of Michigan: Inflation Expectation [MICH], retrieved from FRED, Federal Reserve Bank of St. Louis https://fred.stlouisfed.org/series/MICH/

[81] Federal Reserve Bank of Philadelphia, Survey of Professional Forecasters, median forecast of CPI inflation over the next year; https://www.philadelphiafed.org/surveys-and-data/real-time-data-research/survey-of-professional-forecasters

[82] Federal Reserve Bank of Cleveland, 1-Year Expected Inflation [EXPINF1YR], retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/EXPINF1YR, June 26, 2026.

[83] In February 2024, near the end of our sample, these became regular quarterly questions.

By roughly what percentage do you expect the price of the product, product line or service responsible for the largest share of your revenue to increase/decrease over the next 12 months?

and

By roughly what percentage did you increase/decrease the price of the product, product line or service responsible for the largest share of your revenue over the last 12 months?

Firms freely entered numerical responses. To mirror the cost data, we winsorize at -2% and +6%.[84] We regress expected price growth onto expected future cost growth, controlling for past price and cost growth[85]

$$\begin{aligned}\mathbb{E}_{i,t}\left[\Delta_{12}p_{i,t+12}\right] =& \alpha_i + \text{Controls}_{i,t} \\ &+ \beta^{\text{future}}\mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right] + \beta^{\text{past}}\Delta_{12}\text{mc}_{i,t} + \delta\Delta_{12}p_{i,t} + u_{i,t} \end{aligned} \tag{50}$$

It is conceptually important to control for past prices and costs, as these may lead to expected future price changes, given sticky prices, even absent any expected cost growth. (In practice, if we do not include these controls, we estimate a slightly larger β^{future}.)

The results are shown in Table A.1. We find a pass-through around one-half, and highly significant. This is true overall, but also when controlling for date and sector-date fixed effects. This suggests that firms view exactly the cost expectations that they report to us as highly significant for their own pricing decisions.

Table A.1: Firms' price expectations depend on cost expectations

	$\mathbb{E}_{i,t}\left[\Delta_{12}p_{i,t+12}\right]$			
$\mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right]$	0.51	0.47	0.46	0.44
	(0.07)	(0.07)	(0.07)	(0.07)
$\Delta_{12}\text{mc}_{i,t}$	0.02	−0.01	−0.02	−0.04
	(0.05)	(0.05)	(0.05)	(0.05)
$\Delta_{12}p_{i,t}$	0.18	0.17	0.18	0.17
	(0.03)	(0.03)	(0.04)	(0.03)
$\text{Controls}_{i,t}$	None	Date FEs	Sector-Date FEs	PCA controls
N	2163	2163	2154	2163

Note: Standard errors, in parentheses, are clustered at the firm level. Price responses winsorized to lie in the range $[-2\%, +6\%]$.

[84] The results are very similar if we bin the observations into the set $\{-2\%, 0\%, 2\%, 4\%, 6\%\}$, exactly like the cost reports. We also get significant results, though with about half the pass-through, if we drop any price responses outside the range $[-2\%, +6\%]$.

[85] This exercise is in the spirit of Meyer and Sheng (2025), who also argue the importance of cost expectations for pricing.

A.5 Proportion of cost responses in each bin over time

On the left side panel of Figure A.3, we show the proportion of firms reporting their costs in a particular bin at each date. We see that in general, the number of responses at -2% or $+6\%$, which might suggest truncation, is very small. However, in the post-Covid high inflation period, more firms were reporting 6% cost growth, suggesting that the truncation was more substantially binding for this period. On the right side panel, we show the average weight placed in each bin by firms forecasting their costs over the next year, with similar findings. The peak weight on 6% is slightly smaller, suggesting truncation may be less binding, and this is even more true when we consider that this includes firms placing only some weight at 6%.[86]

Figure A.3: Proportion of cost responses in each bin over time

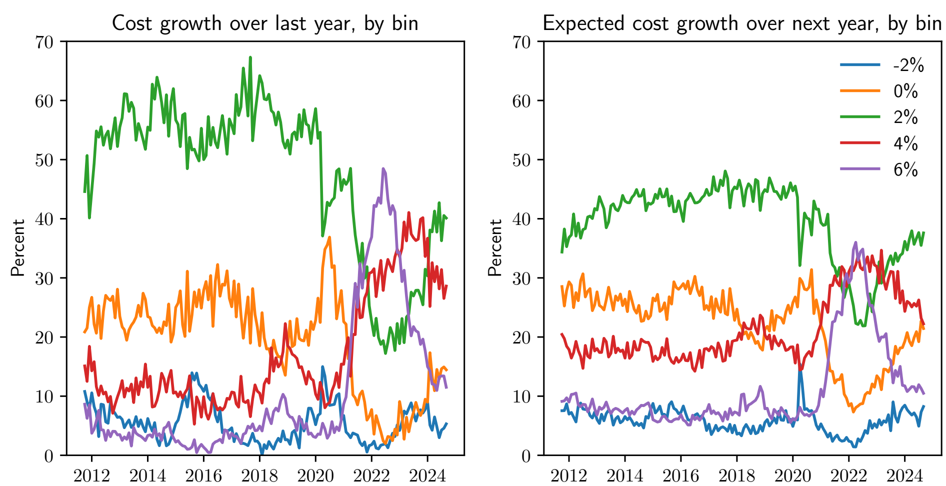

Note: Left hand side breaks down past cost growth response into five bins. Right hand side is average share of probability weight placed on the given bin when forecasting costs.

A.6 Firms' beliefs on average over time

In Figure A.4 we take all firms with at least 30 observations of costs and beliefs, and average these over the time series. We then show a scatter plot of these firm average costs and average forecasts. The beliefs lie close to, though not exactly over, the 45 degree line, suggesting firms do not typically make systematic mistakes about their average cost growth. For example, firms in industries with high cost inflation over many years do not substantially under-predict their cost growth on average.

[86] If we consider only firms placing a 100% probability for cost growth over the next year at 6% instead, then the peak proportion lies at 16% of firms.

Figure A.4: Average beliefs and average costs over time for each firm

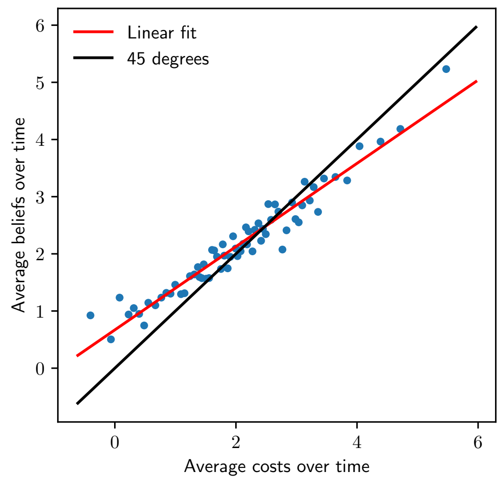

Note: Binscatter plot where each bin contains 5 firms. Only firms with at least 30 observations are included.

A.7 Binscatter

Following Cattaneo et al. (2024), we jointly estimate the mean at each of our bins and the fixed effect controls. For example, with beliefs as the dependent variable and with only firm fixed effects, we run

$$\mathbb{E}_{it}[\Delta_{12}\mathrm{mc}_{i,t+12}] = \alpha_i + \sum_{j \in \{-2,0,2,4,6\}} \beta_j 1_{\{\Delta_{12}\mathrm{mc}_{it}=j\}} + u_{i,t}$$

The point evaluated at j is then $\overline{\hat{\alpha}_i} + \hat{\beta}_j$. The results are shown in Figure A.5. We continue to find a strong, positive, and fairly linear relationship between costs and beliefs, and a weak, positive relationship between past and future costs.

In the main text, we do not take this approach, and instead attempt to recover the conditional expectation function between beliefs and costs around a firm-specific steady state.

A.8 Principal component analysis (PCA) of cost growth

To run the PCA we first need to create a balanced panel. Whenever an observation is missing, we fit its value from firm and time fixed effects estimated through the panel, so

$$(\Delta_{12}\mathrm{mc}_{i,t})^{\text{fitted}} = \begin{cases} \Delta_{12}\mathrm{mc}_{i,t} & \text{if not missing} \\ \widehat{\alpha}_i + \widehat{\alpha}_t & \text{if missing} \end{cases}$$

where $\widehat{\alpha}_i$ and $\widehat{\alpha}_t$ are estimated from the regression $\Delta_{12}\mathrm{mc}_{i,t} = \alpha_i + \alpha_t$. With no missing observations, we can now run a standard PCA to obtain the decomposition

$$\Delta_{12}\mathrm{mc}_{i,t} = \sum_k \Lambda_{k,i} F_{k,t}$$

Figure A.5: Firms (over) react to their own costs: Binscatter robustness

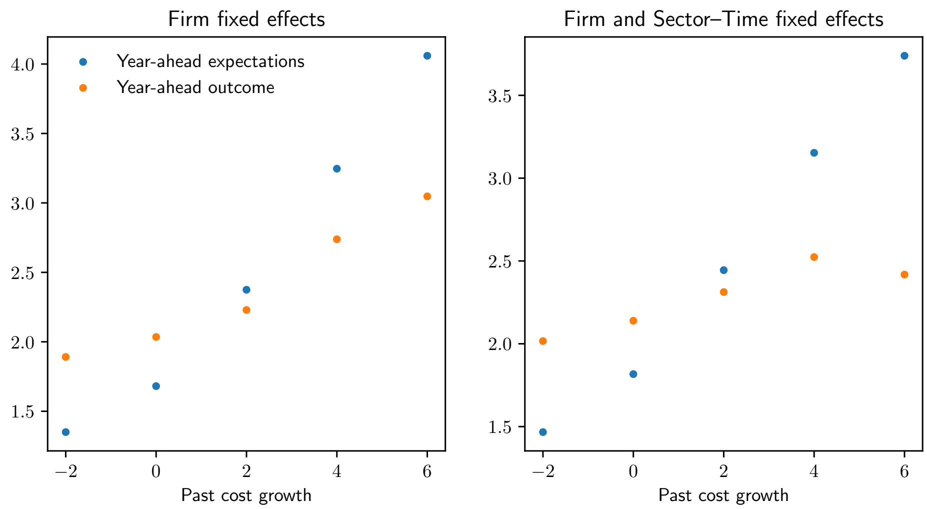

Note: Binscatter following the methodology in Cattaneo et al. (2024)

where we sort factors by the share of the variance they explain. Therefore our PCA controls are: $\Lambda_{1,i}F_{1,t}$, $\Lambda_{2,i}F_{2,t}$, $\Lambda_{3,i}F_{3,t}$, and $\Lambda_{4,i}F_{4,t}$.

A.9 Other belief outcomes

In Figure A.6, we show the coefficients from regressing other beliefs onto costs. In particular, firms appear to rely on movements in their own costs also when forecasting the very long run (their own cost growth five to ten years ahead) and forecasting aggregates (CPI inflation over the next year). We see the biggest decrease in the coefficient on CPI beliefs from introducing time fixed effects; this suggests that when learning about CPI itself, firms may be more inclined to use aggregate signals than when trying to forecast their own costs. This points towards a disconnect we further emphasize in Section 2.4.

A.10 Average and marginal costs

First, we show that when there are no fixed costs and production is homogeneous of any degree $\alpha > 0$, log changes in average and marginal cost are identical. Let $\mathcal{C}(\mathbf{P}, Y) = \min_{\mathbf{X}} \mathbf{P}'\mathbf{X}$ subject to $\mathcal{F}(\mathbf{X}) = Y$ be the variable cost function, where $\mathbf{P}$ is the vector of input prices and Y is output. Define average cost $AC = \mathcal{C}(\mathbf{P}, Y)/Y$ and marginal cost $MC = \partial\mathcal{C}(\mathbf{P}, Y)/\partial Y$, with logs ac and mc.

Lemma 1. *If production is homogeneous of degree $\alpha > 0$ and there are no fixed costs, then $dac = dmc$.*

Proof. It is a standard result in microeconomic theory that homogeneity of degree α in production implies that the variable cost function is homogeneous of degree $1/\alpha$ in output, so $\mathcal{C}(\mathbf{P}, Y) =$

Figure A.6: Costs drive firms' long-run and CPI beliefs

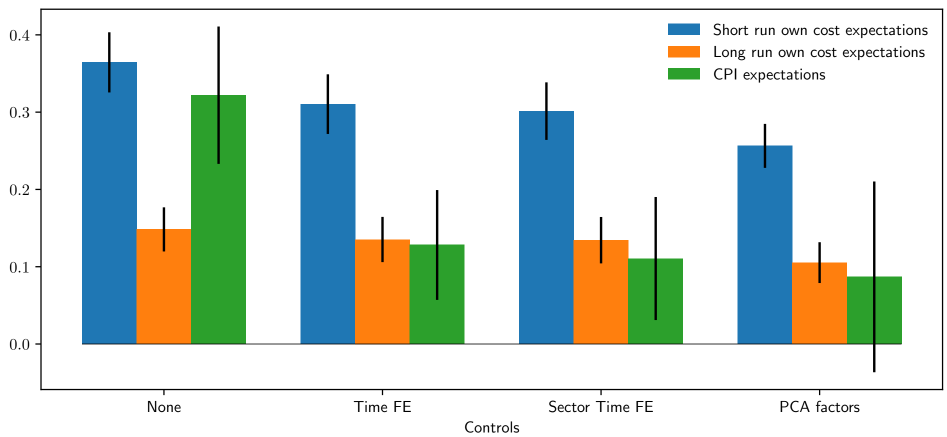

Note: Black lines are 95% confidence intervals, standard errors clustered at firm level.

$Y^{1/\alpha}\mathcal{C}(\mathbf{P},1)$. Hence

$$AC = Y^{\frac{1-\alpha}{\alpha}}\mathcal{C}(\mathbf{P},1), \qquad MC = \frac{1}{\alpha}Y^{\frac{1-\alpha}{\alpha}}\mathcal{C}(\mathbf{P},1) = \frac{1}{\alpha}AC.$$

Taking logs, $mc = -\ln\alpha + ac$. Since α is constant, $dac = dmc$. □

The result in Lemma 1 says that, without fixed costs, we need not worry about whether firms interpret "unit cost" in the BIE to mean average cost or marginal costs. With fixed costs, Lemma 1 no longer holds exactly, but we now show it holds approximately when the fixed-cost share of total cost is small. When the fixed cost share is large, there is a strong force for average cost to fall with output, which we then rule out by estimating a mild degree of decreasing returns in our data—i.e, unit costs rise slightly with output. To allow for fixed costs, define total cost as $TC = \mathcal{C}(\mathbf{P},Y) + F$ and redefine $AC = TC/Y$.

Lemma 2. *If production is homogeneous of degree $\alpha > 0$ and there is a fixed cost of production F, then, to a first order, $dac = s_C dmc - s_F dy$, where $s_C = \mathcal{C}(\mathbf{P},Y)/TC$ is the share of variable costs in total costs, and $s_F = 1 - s_C$ is the share of fixed costs in total costs. Further, $dac = s_C d\ln\mathcal{C}(\mathbf{P},1) + \frac{s_C-\alpha}{\alpha}dy$, so average cost rises with output if and only if $s_C > \alpha$.*

Proof. Marginal cost is unchanged from Lemma 1, since $\partial F/\partial Y = 0$. Using the definition of total cost, average cost is

$$AC = \frac{\mathcal{C}(\mathbf{P},Y)}{Y} + \frac{F}{Y}.$$

Therefore, to a first order,

$$dac = s_C d\ln\frac{\mathcal{C}(\mathbf{P},Y)}{Y} + s_F d\ln\frac{F}{Y}.$$

By Lemma 1, $d\ln\frac{\mathcal{C}(\mathbf{P},Y)}{Y} = dmc$, and, since F is constant, $dac = s_C dmc - s_F dy$. Substituting $dmc =$

$\frac{1-\alpha}{\alpha} dy + d \ln \mathcal{C}(\mathbf{P}, 1)$ from the proof of Lemma 1 and using $s_F = 1 - s_C$,

$$dac = s_C d \ln \mathcal{C}(\mathbf{P}, 1) + \left(s_C \frac{1-\alpha}{\alpha} - (1 - s_C) \right) dy = s_C d \ln \mathcal{C}(\mathbf{P}, 1) + \frac{s_C - \alpha}{\alpha} dy. \qquad \square$$

Now, we explore how unit cost growth depends on output growth in the BIE data. We have both qualitative and quantitative data on how sales are deviating from normal levels in the BIE (see the note in Table A.2 for more details). We take the twelve month change in these measures, which provide qualitative and quantitative (in percentage points) measures of how much sales increased over the last year. The percent change in sales is not quite the percent change in output, since many firms may interpret sales as revenues rather than quantities sold, but under sticky prices, our sales measure should provide a reasonable proxy to changes in log output. We estimate the regression model

$$\Delta_{12} mc_{i,t} = \beta_{\text{Intercept}} + \beta_{\text{Sales}} \text{Sales}_{i,t} + \text{Controls}_{i,t} + \epsilon_{i,t}$$

with a variety of fixed effects as controls and where $\Delta_{12} mc_{i,t}$ are the unit cost growth measures in our BIE data.

The results are shown in Table A.2. In panel (b), where we use the quantitative measure for sales, sales being 100pp higher than normal predicts, at most (column 2), 0.6pp higher unit cost growth. Using the estimates from the 5-point scale of the qualitative sales question (panel a), unit cost growth is at most 0.3pp higher than normal when sales are "much greater than normal" (2 scale steps higher than normal in column 2).

Turning to the result in Lemma 2, if we assume that unit costs actually measure average costs rather than marginal costs, our quantitative estimates in panel (b) suggest that $\frac{s_C - \alpha}{\alpha}$ is at least positive. Given that much of the literature estimates the central tendency of firm-level returns to scale to be near constant, i.e. α is not much smaller than 1, our estimate suggests that variable cost is the dominant cost for the firms in our data; for example, if we assume $\alpha > 0.9$, then $\frac{s_C - \alpha}{\alpha} \geq 0 \iff s_C \geq 0.9$. In this case, Lemma 2 says it is approximately true that changes in log marginal cost track changes in log average cost. Alternatively, if fixed costs do play a dominant role in costs for the firms in our data, our empirical results in Table A.2 reject the idea that firms' unit cost growth responses reflect their average cost growth instead of their marginal cost growth. In either case, the unit cost growth reported in our data should be very close to marginal cost growth.

A.11 Additional controls

First, we include three other firm responses as controls. Each month, firms qualitatively report if their sales are above or below normal (on a five-point scale), and likewise for profit margins. Each quarter, firms additionally report by what percent unit sales are above or below normal. In Figure A.7, we show the regression results when including all three of these responses as controls,

Table A.2: Cost growth and sales

(a) Qualitative sales measure

	$\Delta_{12}\text{mc}_{i,t}$			
$\Delta_{12}\text{sales}_{i,t}$	0.15***	0.09***	0.08***	0.03**
	(0.02)	(0.02)	(0.02)	(0.01)
$\text{Controls}_{i,t}$	None	Date FEs	Sector-Date FEs	PCA controls
N	19890	19890	19786	19890

(b) Quantitative sales measure

	$\Delta_{12}\text{mc}_{i,t}$			
$\Delta_{12}\text{sales}_{i,t}$	0.006***	0.002*	0.002	0.000
	(0.002)	(0.001)	(0.001)	(0.001)
$\text{Controls}_{i,t}$	None	Date FEs	Sector-Date FEs	PCA controls
N	5138	5138	5092	5138

Note: *p<0.10, **p<0.05, ***p<0.01. Standard errors, in parentheses, are clustered at the firm level. The qualitative sales question is asked monthly, and reported on a five point scale, from "Much less than normal" through to "Much greater than normal". The quantitative sales question is asked quarterly, and firms report the percent deviation of sales from normal.

though the results are similar with any subset of these three controls.

Next, we show that our effect is not undone by controlling for CPI beliefs. We show the results in Figure A.8. So although we do find a role for firms learning about CPI via their own costs in Section A.9 below – a result suggested by the existing literature – we note that understanding this channel alone is not sufficient to understand how firms use their own costs to forecast their own future costs, a key object for pricing decisions. Standard errors are wider because we only observe firms' CPI beliefs during months when those beliefs are elicited in special questions – a relatively small subset of our main sample.

A.12 Mismeasurement

Suppose that the reported survey variables reflect the true values plus measurement error,

$$\mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right] = \left(\mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right]\right)^{\text{true}} + \eta_{i,t}^{\text{expectation}}$$
$$\Delta_{12}\text{mc}_{i,t} = \left(\Delta_{12}\text{mc}_{i,t}\right)^{\text{true}} + \eta_{i,t}^{\text{outcome}}$$

Classical measurement error, i.e. $\text{Var}\left(\eta_{i,t}^{\text{outcome}}\right) > 0$, would attenuate our coefficients. So our effects might be even stronger than the results above suggest. On the other hand, if respondents to the survey have correlated measurement error across responses in the survey, i.e. $\text{Cov}\left(\eta_{i,t}^{\text{expectation}}, \eta_{i,t}^{\text{outcome}}\right) > 0$, then this could make our results appear too strong. We now show that our results survive instrumenting costs with lagged costs.

Figure A.7: Including sales and profits as controls

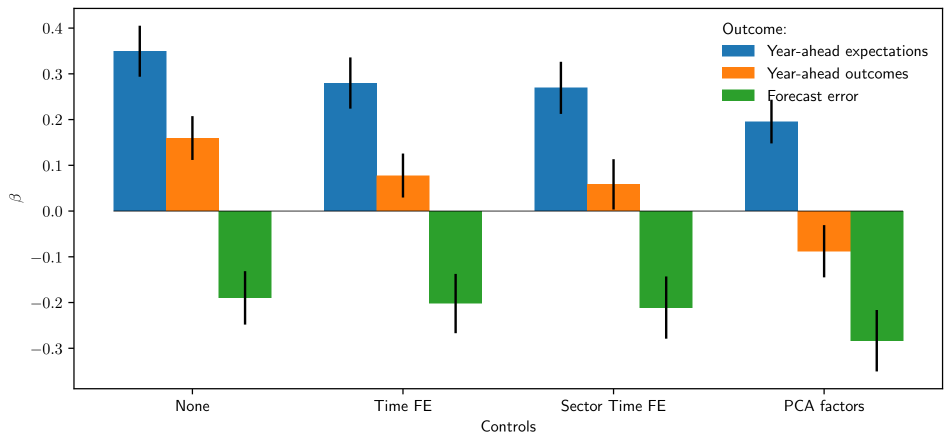

Note: Black lines are 95% confidence intervals, standard errors clustered at firm level.

Figure A.8: Including CPI beliefs as a control

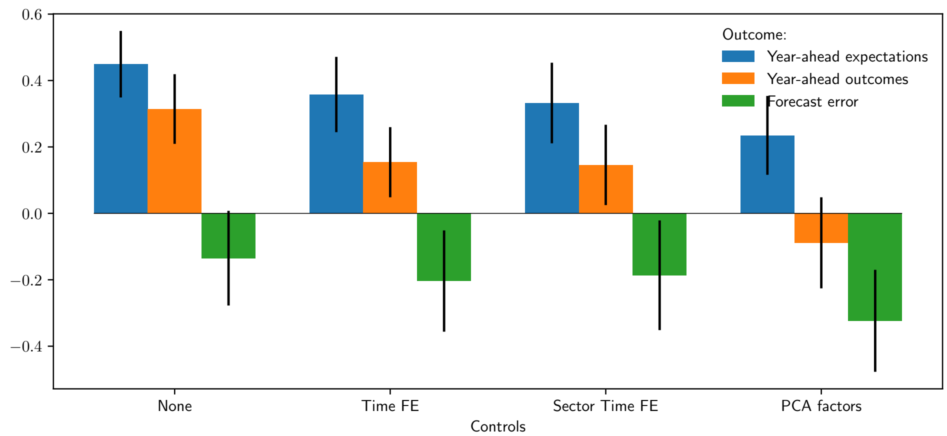

Note: Black lines are 95% confidence intervals, standard errors clustered at firm level.

Figure A.9: Robustness to instrumenting costs

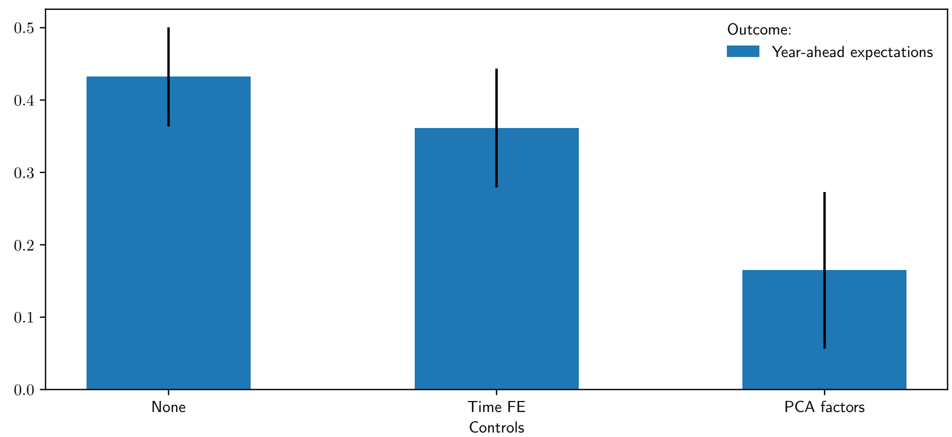

Note: Black lines are 95% confidence intervals, standard errors clustered at firm level. F-stats all larger than 250.

We repeat our analysis but now instrumenting costs with their lag. We run

$$\mathbb{E}_{it}[\Delta_{12}\text{mc}_{i,t+12}] = \alpha_i + \text{Controls}_{it} + \beta_{\text{expectations}}\widehat{\Delta_{12}\text{mc}_{it}} + u_{i,t}^{\text{expectations}} \tag{51}$$

$$\Delta_{12}\text{mc}_{i,t+12} = \alpha_i + \text{Controls}_{it} + \beta_{\text{outcome}}\widehat{\Delta_{12}\text{mc}_{it}} + u_{i,t}^{\text{outcome}} \tag{52}$$

$$\text{FE}_{i,t,t+12} = \alpha_i + \text{Controls}_{it} + \beta_{\text{FE}}\widehat{\Delta_{12}\text{mc}_{it}} + u_{i,t}^{\text{FE}} \tag{53}$$

where the instrumented costs are fitted from the first-stage regression

$$\Delta_{12}\text{mc}_{i,t} = \alpha_i + \text{Controls}_{it} + \beta_{\text{FS}}\Delta_{12}\text{mc}_{i,t-1} + u_{i,t}^{\text{FS}}$$

The results are shown in Figure A.9.

Next we show our results are not driven by truncation. In Figure A.10, we remove all boundary values (equal to -2% or +6%) and continue to find our results. Note also that **Fact 3** shows our results are similar pre- and post-Covid; from Appendix A.5, we know the truncation was far less binding in the earlier period.

A.13 De-meaning only with past data

We estimate our firm-level fixed effects over the entire sample. However, this may reflect future information in finite samples. To address this, we consider the regression

$$\widetilde{\text{FE}}_{i,t,t+12} = \beta_{\text{FE}}\widehat{\widetilde{\Delta_{12}\text{mc}}_{it}} + u_{i,t}^{\text{FE}}$$

where tildes represent de-meaning using only observations up to period t. We find that β_{FE} moves from -0.20 in our baseline regression to -0.23 when demeaning with only past observations.

Figure A.10: Robustness to truncation of the data

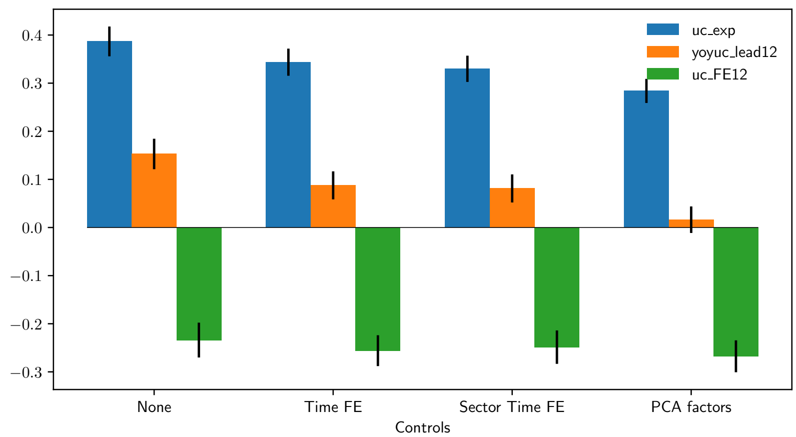

Note: Black lines are 95% confidence intervals, standard errors clustered at firm level.

A.14 Heterogeneity by frequency of price adjustment

Firms who adjust prices less often may have more incentive to forecast costs well. In Figure A.11 we show that our results are stable across different frequencies of price adjustment.

We construct the frequency of price adjustment measure by combining three *special questions* asked over the history of the BIE.[87] We create three groups for frequency of price adjustment: "At least monthly", "Quarterly or semiannually", and "Annually or less frequent". Since some firms respond multiple times over the sample, we ask how consistent responses are across the 2013, 2019, and 2023 waves. Responses are quite stable, with 68% remaining in the same group as the last response.[88] We observe the frequency of price adjustment at only three dates (at most), so in order to run our regressions, we fill in the values at all dates, always using the latest firm response.[89]

A.15 Do firms who should learn more from costs do so?

We first run a firm-level regression of current costs on past costs

$$\Delta_{12}\text{mc}_{i,t} = \alpha_i + \rho_i \Delta_{12}\text{mc}_{i,t-12} + u_{i,t} \tag{54}$$

[87] In 2013, firms were asked "What percentage of your products and/or services change price Daily/Weekly/Monthly/Quarterly/Semiannually/Annually/Less often than annually?". Firms assign percentages to each case; we use the mode of this distribution. Once in 2019 and once in 2023, firms were asked "How often do you typically change the price of the product or service responsible for the largest share of your sales revenue?", and had to choose from the options: Daily, Weekly, Monthly, Quarterly, Semiannually, Annually, Every 1-3 years, Less often than every 3 years, Irregularly/in response to a specific event, and Unsure. We do not include the last two responses in our analysis.

[88] Only 7% of consecutive responses jump from "At least monthly" to "Annually or less", or vice versa. i.e. 93% of responses move at most to the adjacent group.

[89] We fill backwards from the first response a firm gives. We fill forward up to the next time a firm responds.

Figure A.11: Heterogeneity by frequency of price adjustment

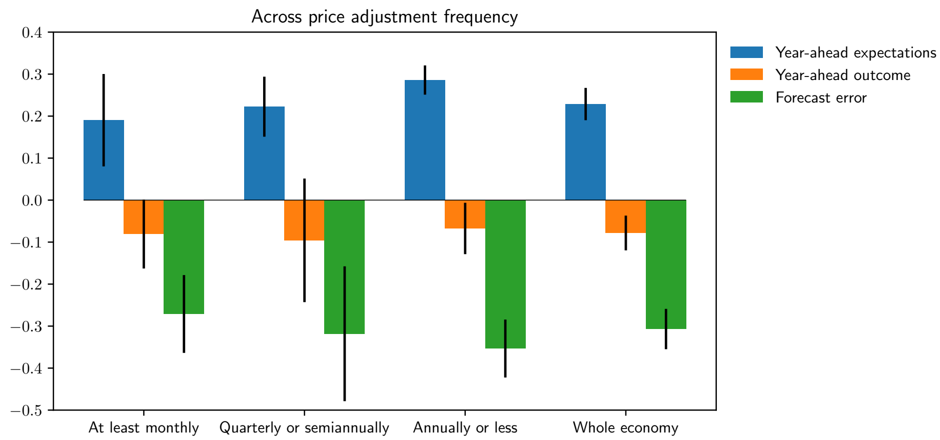

Note: Vertical lines are 95% confidence intervals, standard errors clustered at firm level. Includes PCA controls.

Figure A.12: Belief react too much to costs on average but too little to individual cost persistence

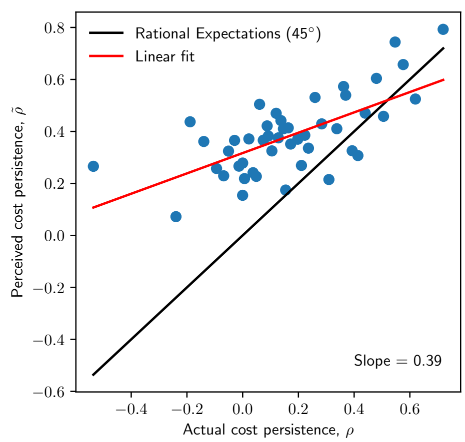

Note: Binned scatter plot of $\hat{\tilde{\rho}}_i$ and $\hat{\rho}_i$, recovered from regression equations (54) and (55) respectively. We run these time series regressions only for firms with at least 30 observations. Our results are robust to other cutoffs. Each bin contains five observations.

to estimate firm i's cost persistence. We next run the firm-level regression of beliefs about the future on current costs

$$\mathbb{E}_{it}[\Delta_{12}\text{mc}_{i,t+12}] = \alpha_i + \tilde{\rho}_i \Delta_{12}\text{mc}_{i,t} + u_{i,t} \tag{55}$$

to estimate firm i's perceived cost persistence.

In Figure A.12, we show a scatter plot of $\tilde{\rho}_i$ against ρ_i, including only firms with at least 30 observations.[90] A 45 degree *slope* would reflect each firm correctly judging how much more persistent their costs are than the average. If firms also correctly judged the *level* of cost persistence, we would recover the 45 degree line passing through the origin. Instead, we find a much shal-

[90]In Appendix B.2, we formally estimate the slope between $\tilde{\rho}_i$ and ρ_i, and test robustness of our results to using a cutoff of 96 observations per firm.

lower slope (with slope coefficient 0.39, rather than 1) suggesting that firms whose costs are in fact particularly transitory are not sufficiently aware of this.

A.16 Firms overreact to costs strictly in the past

Here, we run the regression

$$\mathrm{FE}_{i,t,t+12} = \alpha_i + \mathrm{Controls}_{it} + \beta_{\mathrm{FE}} \Delta_{12} \mathrm{mc}_{i,t-6} + u^{\mathrm{FE}}_{i,t}$$

This differs from (8) only in using the cost growth from six months ago. The results are shown in Table A.3. We continue to find significantly negative estimates for β_{FE}, telling us that firms do not only overreact in period t to surprises (relative to FIRE) in that same period.

Table A.3: Firms over-react to cost growth from six months ago

	Forecast error, $\Delta_{12}\mathrm{mc}_{i,t+12} - \mathbb{E}_{i,t}\left[\Delta_{12}\mathrm{mc}_{i,t+12}\right]$			
β_{FE}	-0.18***	-0.15***	-0.13***	-0.21***
	(0.03)	(0.03)	(0.02)	(0.02)
$\mathrm{Controls}_{i,t}$	None	Date FEs	Sector-Date FEs	PCA controls
Observations	15580	15580	15412	15580

Note: $^{*} < 0.10$, $^{**} < 0.05$, $^{***} < 0.01$. Standard errors clustered at the firm level.

B Belief Model and Calibration

B.1 Motivating part-FIRE, part-Adaptive Learning

Suppose that each firm anchors on the Adaptive Learning forecast, so that their prior is that the optimal forecast for the own costs is

$$\mathbb{E}_t\left[\Delta \mathrm{mc}_{i,t+\tau}\right] \sim \mathcal{N}\left(\tilde{\mathbb{E}}^{\mathrm{AL}}_{i,t}\left[\Delta \mathrm{mc}_{i,t+\tau}\right], \sigma^2_{\mathrm{AL}}\right) \tag{56}$$

Firms can attempt to improve this noisy prior by working out what FIRE implies. However, this is costly, as it requires learning about many shocks, and computing the firm's exposure to each shock. We model this as the firm receiving a noisy signal centered on FIRE,

$$s^{\mathrm{FIRE}}_{i,t} = \mathbb{E}_t\left[\Delta \mathrm{mc}_{i,t+\tau}\right] + \nu_{i,t}, \qquad \nu_{i,t} \overset{\text{i.i.d.}}{\sim} \mathcal{N}\left(0, \sigma^2_\nu\right) \tag{57}$$

Firms can pay some cognitive cost $C^{\text{FIRE}}\left(1/\sigma_{\nu}^{2}\right)$ to reduce the variance of the noise. In this case, the optimal forecast is

$$\mathbb{E}_{i,t}\left[\Delta \text{mc}_{i,t+\tau}\right]=\alpha^{\text{FIRE}}\mathbb{E}_{t}\left[\Delta \text{mc}_{i,t+\tau}\right]+\left(1-\alpha^{\text{FIRE}}\right)\tilde{\mathbb{E}}_{i,t}^{\text{AL}}\left[\Delta \text{mc}_{i,t+\tau}\right]+\alpha^{\text{FIRE}}\nu_{i,t} \tag{58}$$

where $\alpha^{\text{FIRE}} \equiv \frac{\sigma_{\text{AL}}^{2}}{\sigma_{\text{AL}}^{2}+\sigma_{\nu}^{2}}$. If the cost of increasing the FIRE signal's precision is steep enough, the firm will not fully recover FIRE, so $\alpha^{\text{FIRE}}<1$.

To first order, for given prior variance, σ_{AL}^{2}, and given ease of computing the FIRE forecast $C^{\text{FIRE}}(\cdot)$, the weight on FIRE, α^{FIRE}, does not vary over time or by shock. We focus our analysis on this case. However, α^{FIRE} may increase if firms (i) believe their Adaptive Learning prior to have become more noisy, or (ii) find it easier to learn about a shock or compute their exposure to it. In Section 2.2, we discuss the empirical stability of α^{FIRE}.

B.2 Firm specific perceived persistence

Consider a static Adaptive Learning model ($\tilde{\mathcal{K}}=1$) cast in annual time, but now allowing for a firm-specific perceived persistence, so that

$$\mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right]=\tilde{\rho}_{i}^{\text{ann}}\Delta_{12}\text{mc}_{i,t} \tag{59}$$

Suppose that the firm chooses $\tilde{\rho}_{i}^{\text{ann}}$ to minimize the mean squared forecast error, subject to a cost function over the choice of $\tilde{\rho}_{i}^{\text{ann}}$,

$$\min_{\tilde{\rho}_{i}^{\text{ann}}}\frac{1}{2}\mathbb{E}\left[\left(\Delta_{12}\text{mc}_{i,t+12}-\tilde{\rho}_{i}^{\text{ann}}\Delta_{12}\text{mc}_{i,t}\right)^{2}\right]+\mathcal{C}\left(\tilde{\rho}_{i}^{\text{ann}}\right)$$

where for simplicity we set $\mathcal{C}\left(\tilde{\rho}_{i}^{\text{ann}}\right)=\frac{\psi\text{Var}(\Delta_{12}\text{mc}_{i,t})}{2}\left(\tilde{\rho}_{i}^{\text{ann}}-\overline{\rho}^{\text{ann}}\right)^{2}$. This implies that the perceived persistence is a weighted average of some prior, $\overline{\rho}^{\text{ann}}$, and the best fit, ρ_{i}^{ann},

$$\tilde{\rho}_{i}^{\text{ann}}=\left(1-\alpha^{\text{RE}}\right)\overline{\rho}^{\text{ann}}+\alpha^{\text{RE}}\rho_{i}^{\text{ann}} \tag{60}$$

where $\rho_{i}^{\text{ann}} \equiv \frac{\text{Cov}(\Delta_{12}\text{mc}_{i,t+12},\Delta_{12}\text{mc}_{i,t})}{\text{Var}(\Delta_{12}\text{mc}_{i,t})}$ and $\alpha^{\text{RE}} \equiv \frac{1}{1+\psi}$. This implies the beliefs equation (59) can be written as

$$\mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right]=\left(1-\alpha^{\text{RE}}\right)\overline{\rho}^{\text{ann}}\cdot\Delta_{12}\text{mc}_{i,t}+\alpha^{\text{RE}}\cdot\rho_{i}^{\text{ann}}\Delta_{12}\text{mc}_{i,t}$$

Recall that Figure A.12 intuitively tests this equation, and found an intercept and a slope both around a third. We test this more formally, also accounting for estimation uncertainty, in Table B.4. We confirm that $\alpha^{\text{RE}}\approx\frac{1}{3}$, and that the default perceived persistence is too high, at around $\overline{\rho}^{\text{ann}}\approx 0.4$. We also control for aggregate components of the data, and also find that firms are too *in*sensitive to their idiosyncratic cost persistence. Finally in Table B.5, we show our findings are robust to increasing the minimum number of observations per firm from 48 to 96.

So the α^{RE} estimate here lies reasonably close to our α^{FIRE} estimate in Section 2.2 above. And recall that since the main text model mixes FIRE and Adaptive Learning, it predicts, proportional to α^{FIRE}, variation in firm-specific extrapolation from own-costs proportional to the true variation. Therefore, it seems that the main text model provides a good empirical approximation of a model with firm-specific endogenous perceived persistence.

Table B.4: Calibration with firm-specific perceived persistence

	Year-ahead cost beliefs, $\mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right]$			
α^{RE}	0.45^{***}	0.39^{**}	0.39^{***}	0.31^{**}
	(0.13)	(0.15)	(0.14)	(0.13)
$(1-\alpha^{RE})\overline{\rho}^{\text{ann}}$	0.28^{***}	0.26^{***}	0.25^{***}	0.24^{***}
	(0.05)	(0.04)	(0.04)	(0.02)
$\text{Controls}_{i,t}$	None	Date FEs	Sector-Date FEs	PCA controls
Min. obs. per firm	4 years	4 years	4 years	4 years
Observations	12976	12976	12769	12976

Note: $^{*} < 0.10$, $^{**} < 0.05$, $^{***} < 0.01$. Standard errors clustered at the firm level.

Table B.5: Calibration with firm-specific perceived persistence

	Year-ahead cost beliefs, $\mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right]$			
α^{RE}	0.55^{***}	0.39^{*}	0.30^{*}	0.27^{*}
	(0.17)	(0.21)	(0.18)	(0.14)
$(1-\alpha^{RE})\overline{\rho}^{\text{ann}}$	0.26^{***}	0.26^{***}	0.25^{***}	0.25^{***}
	(0.07)	(0.06)	(0.05)	(0.03)
$\text{Controls}_{i,t}$	None	Date FEs	Sector-Date FEs	PCA controls
Min. obs. per firm	8 years	8 years	8 years	8 years
Observations	6081	6081	5630	6081

Note: $^{*} < 0.10$, $^{**} < 0.05$, $^{***} < 0.01$. Standard errors clustered at the firm level.

B.3 Beliefs facing permanent cost growth shocks

Suppose now that firms believe costs to be driven by three components

$$\Delta\text{mc}_{i,t} = \Delta\tilde{\text{mc}}_{i,t}^{\text{Permanent}} + \Delta\tilde{\text{mc}}_{i,t}^{\text{Persistent}} + \Delta\tilde{\text{mc}}_{i,t}^{\text{Transitory}} \tag{61}$$

where the permanent component is given by

$$\Delta\tilde{\text{mc}}_{i,t}^{\text{Permanent}} = \Delta\tilde{\text{mc}}_{i,t-1}^{\text{Permanent}} + \tilde{\varepsilon}_{i,t}^{\text{Permanent}}, \qquad \tilde{\varepsilon}_{i,t}^{\text{Permanent}} \overset{\text{iid}}{\sim} \mathcal{N}\left(0, \tilde{\sigma}_{\text{Permanent}}^2\right) \tag{62}$$

the persistent component is given by

$$\Delta\tilde{\text{mc}}_{i,t}^{\text{Persistent}} = \tilde{\rho}\Delta\tilde{\text{mc}}_{i,t-1}^{\text{Persistent}} + \tilde{\varepsilon}_{i,t}^{\text{Persistent}}, \qquad \tilde{\varepsilon}_{i,t}^{\text{Persistent}} \overset{\text{iid}}{\sim} \mathcal{N}\left(0, \tilde{\sigma}_{\text{Persistent}}^2\right) \tag{63}$$

and the transitory component is driven by

$$\Delta\tilde{\text{mc}}_{i,t}^{\text{Transitory}} = \tilde{\varepsilon}_{i,t}^{\text{Transitory}}, \qquad \tilde{\varepsilon}_{i,t}^{\text{Transitory}} \overset{\text{iid}}{\sim} \mathcal{N}\left(0, \tilde{\sigma}_{\text{Transitory}}^2\right) \tag{64}$$

and where $\tilde{\varepsilon}_{i,t}^{\text{Permanent}}$, $\tilde{\varepsilon}_{i,t}^{\text{Persistent}}$, and $\tilde{\varepsilon}_{i,t}^{\text{Transitory}}$ are furthermore mutually independent.

Perfectly observed permanent shocks Now suppose that firms are perfectly informed about the permanent cost component. So the information set of firm i and time t is $\mathcal{I}_{i,t} = \left\{\Delta\text{mc}_{i,t-\ell}, \Delta\tilde{\text{mc}}_{i,t-\ell}^{\text{Permanent}}\right\}_{\ell=0}^{\infty}$.

Then the signal extract problem reduces to

$$\Delta\widehat{\text{mc}}_{i,t} = \Delta\tilde{\text{mc}}_{i,t}^{\text{Persistent}} + \Delta\tilde{\text{mc}}_{i,t}^{\text{Transitory}} \tag{65}$$

where the firm observes

$$\Delta\widehat{\text{mc}}_{i,t} \equiv \Delta\text{mc}_{i,t} - \Delta\tilde{\text{mc}}_{i,t}^{\text{Permanent}} \tag{66}$$

and given (63) and (64). The two key beliefs are then

$$\mathbb{E}_{i,t}\left[\Delta\tilde{\text{mc}}_{i,t}^{\text{Permanent}}\right] = \Delta\tilde{\text{mc}}_{i,t}^{\text{Permanent}} \tag{67}$$

and

$$\mathbb{E}_{i,t}\left[\Delta\tilde{\text{mc}}_{i,t}^{\text{Persistent}}\right] = \tilde{\mathcal{K}}\Delta\widehat{\text{mc}}_{i,t} + \left(1-\tilde{\mathcal{K}}\right)\tilde{\rho}\mathbb{E}_{i,t-1}\left[\Delta\tilde{\text{mc}}_{i,t-1}^{\text{Persistent}}\right] \tag{68}$$

B.4 Robustness to other instruments

The results using sector loadings are shown in Table B.6. The results using firm-level loadings and a cutoff of at least 72 observations per firm are shown in Table B.7. The results when accounting for the two-year lag of own costs are shown in Table B.8.

B.5 Quarterly Calibration from Annual

From our annual calibration parameters, we can derive a quarterly calibration under the typical assumptions made that a variable's movement over a given period be equally divided into movements over the component subperiods. Denote $A(x,\ell) = (x^{1-\ell} + x^{2-\ell} + x^{3-\ell} + x^{4-\ell})$. We can

Table B.6: Annual Calibration with sectoral loadings

Panel A: 2SLS estimates

		$\mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right]$		
Regression Estimates	$\Delta_{12}\text{mc}_{i,t}$	0.289***	0.291***	0.268***
		(0.023)	(0.025)	(0.018)
	$\Delta_{12}\text{mc}_{i,t-12}$	0.005	0.003	0.031***
		(0.011)	(0.013)	(0.012)
	$\Delta_{12}\text{mc}_{i,t+12}$	0.038	0.024	0.245**
		(0.096)	(0.095)	(0.104)
Implied Parameters	α^{FIRE}	0.038	0.024	0.245
		(0.096)	(0.095)	(0.104)
	$\tilde{\mathcal{K}}^{\text{ann}}$	0.945	0.963	0.753
		(0.112)	(0.147)	(0.067)
	$\tilde{\rho}^{\text{ann}}$	0.319	0.309	0.472
		(0.045)	(0.048)	(0.079)
	Instrument	Av. Costs × Sector id	CPI × Sector id	NS × Sector id
	Observations	13855	13855	13855
	Kleibergen Paap F stat	29.1	15.3	18.7

Panel B: LIML estimates

		$\mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right]$		
Regression Estimates	$\Delta_{12}\text{mc}_{i,t}$	0.291***	0.293***	0.265***
		(0.023)	(0.026)	(0.019)
	$\Delta_{12}\text{mc}_{i,t-12}$	0.003	0.000	0.035**
		(0.012)	(0.016)	(0.014)
	$\Delta_{12}\text{mc}_{i,t+12}$	0.022	0.001	0.274**
		(0.126)	(0.128)	(0.121)
Implied Parameters	α^{FIRE}	0.022	0.001	0.274
		(0.126)	(0.128)	(0.121)
	$\tilde{\mathcal{K}}^{\text{ann}}$	0.966	0.996	0.735
		(0.135)	(0.180)	(0.070)
	$\tilde{\rho}^{\text{ann}}$	0.308	0.295	0.497
		(0.062)	(0.065)	(0.097)
	Instrument	Av. Costs × Sector id	CPI × Sector id	NS × Sector id
	Observations	13855	13855	13855
	Kleibergen Paap F stat	29.1	15.3	18.7

Note: *p<0.10, **p<0.05, ***p<0.01. All regressions include Firm and Time fixed effects. Sample includes all firms matched to a sector. Driscoll Kraay standard errors with bandwidth 12.

Table B.7: Annual Calibration with loadings on firms with at least 72 observations

		$\mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right]$		
Regression Estimates	$\Delta_{12}\text{mc}_{i,t}$	0.297***	0.296***	0.289***
		(0.020)	(0.021)	(0.018)
	$\Delta_{12}\text{mc}_{i,t-12}$	0.021	0.023	0.029*
		(0.014)	(0.015)	(0.015)
	$\Delta_{12}\text{mc}_{i,t+12}$	0.087*	0.096***	0.149***
		(0.046)	(0.036)	(0.029)
Implied Parameters	α^{FIRE}	0.087	0.096	0.149
		(0.046)	(0.036)	(0.029)
	$\tilde{\mathcal{K}}^{\text{ann}}$	0.819	0.811	0.771
		(0.107)	(0.113)	(0.099)
	$\tilde{\rho}^{\text{ann}}$	0.397	0.403	0.441
		(0.038)	(0.043)	(0.047)
	Instrument	Av. Costs × Firm id	CPI × Firm id	NS × Firm id
	Min obs per firm	6 years	6 years	6 years
	Observations	6691	6691	6691
	Kleibergen Paap F stat	740	1233	318

Note: *p<0.10, **p<0.05, ***p<0.01. Estimate by 2SLS. All regressions include Firm and Time fixed effects. Driscoll Kraay standard errors with bandwidth 12.

Table B.8: Annual Calibration with additional lag

		$\mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right]$		
Regression Estimates	$\Delta_{12}\text{mc}_{i,t}$	0.235***	0.243***	0.230***
		(0.024)	(0.023)	(0.022)
	$\Delta_{12}\text{mc}_{i,t-12}$	0.045***	0.036**	0.052***
		(0.017)	(0.016)	(0.017)
	$\Delta_{12}\text{mc}_{i,t-24}$	-0.019	-0.019	-0.019
		(0.018)	(0.019)	(0.018)
	$\Delta_{12}\text{mc}_{i,t+12}$	0.258***	0.193***	0.304***
		(0.039)	(0.031)	(0.054)
Implied Parameters	α^{FIRE}	0.258	0.193	0.304
		(0.039)	(0.031)	(0.054)
	$\tilde{\mathcal{K}}^{\text{ann}}$	0.624	0.672	0.596
		(0.106)	(0.116)	(0.095)
	$\tilde{\rho}^{\text{ann}}$	0.508	0.448	0.555
		(0.068)	(0.065)	(0.079)
	Instrument	Av. Costs × Firm id	CPI × Firm id	NS × Firm id
	Min obs per firm	4 years	4 years	4 years
	Observations	7345	7345	7345
	Kleibergen Paap F stat	5491	5570	6617

Note: *p<0.10, **p<0.05, ***p<0.01. Estimate by 2SLS. All regressions include Firm and Time fixed effects. Driscoll Kraay standard errors with bandwidth 12. Implied parameters computed using same three regression parameters as main text.

write the year ahead cost forecast in the quarterly model as

$$\tilde{\mathbb{E}}_{it}\Delta_{12}\mathrm{mc}_{i,t+12} = A(\tilde{\rho},0)\tilde{\mathcal{K}}\sum_{\ell=0}^{\infty}(\tilde{\rho}(1-\tilde{\mathcal{K}}))^{\ell}\Delta_{3}\mathrm{mc}_{i,t-3\ell}.$$

Subtracting a weighted, one-year lag of the forecast, we have

$$\tilde{\mathbb{E}}_{it}\Delta_{12}\mathrm{mc}_{i,t+12} - (\tilde{\rho}(1-\tilde{\mathcal{K}}))^{4}\tilde{\mathbb{E}}_{i,t-12}\Delta_{12}\mathrm{mc}_{i,t} = A(\tilde{\rho},0)\tilde{\mathcal{K}}\sum_{\ell=0}^{3}(\tilde{\rho}(1-\tilde{\mathcal{K}}))^{\ell}\Delta_{3}\mathrm{mc}_{i,t-3\ell}.$$

If we assume that all quarterly changes with the year are equivalent and equal to $\Delta_{12}\mathrm{mc}_{it}/4$, then

$$\tilde{\mathbb{E}}_{it}\Delta_{12}\mathrm{mc}_{i,t+12} - (\tilde{\rho}(1-\tilde{\mathcal{K}}))^{4}\tilde{\mathbb{E}}_{i,t-12}\Delta_{12}\mathrm{mc}_{it} = A(\tilde{\rho},0)\tilde{\mathcal{K}}A(\tilde{\rho}(1-\tilde{\mathcal{K}}),1)\frac{\Delta_{12}\mathrm{mc}_{it}}{4}.$$

This is exactly analogous to the annual equation

$$\tilde{\mathbb{E}}_{it}\Delta_{12}\mathrm{mc}_{i,t+12} - (\tilde{\rho}^{\mathrm{ann}}(1-\tilde{\mathcal{K}}^{\mathrm{ann}}))\tilde{\mathbb{E}}_{i,t-12}\Delta_{12}\mathrm{mc}_{it} = \tilde{\rho}^{\mathrm{ann}}\tilde{\mathcal{K}}^{\mathrm{ann}}\Delta_{12}\mathrm{mc}_{it},$$

and so the quarterly parameters can be determined by matching coefficients. The parameter α^{FIRE} is constant across quarterly and annual calibrations.

B.6 Fit of the calibrated beliefs model to sectoral time series of beliefs

We now fit beliefs from our annual, calibrated beliefs model, truncating after one year back. For the FIRE expectation of future costs we use the Michigan forecast for CPI over the next year. The fit is shown in Figure B.13.

B.7 Fit of the calibrated beliefs model to Fact 5

Figure B.14 shows, in blue, the path of beliefs in response to the change in the oil price, as seen in the main text. In black we add the prediction from the calibrated model. This matches the key dynamics, with quickly exposed firms' beliefs jumping on impact, and slowly exposed firms' beliefs rising gradually over time. By contrast, under FIRE (orange) all firms' beliefs should jump on impact, and for the firms where pass-through occurs within the year, their forecast should start declining immediately.

B.8 Learning over a longer horizon

As we want to extend our analysis to before the start of the BIE sample, we assume costs can be decomposed into two components

$$\Delta_{12}\mathrm{mc}_{i,t} = \Delta_{12}\mathrm{mc}_{t}^{\mathrm{agg}} + \Delta_{12}\mathrm{mc}_{i,t}^{idio} \tag{69}$$

Figure B.13: Model predicted time series for sectoral cost beliefs closely match the data

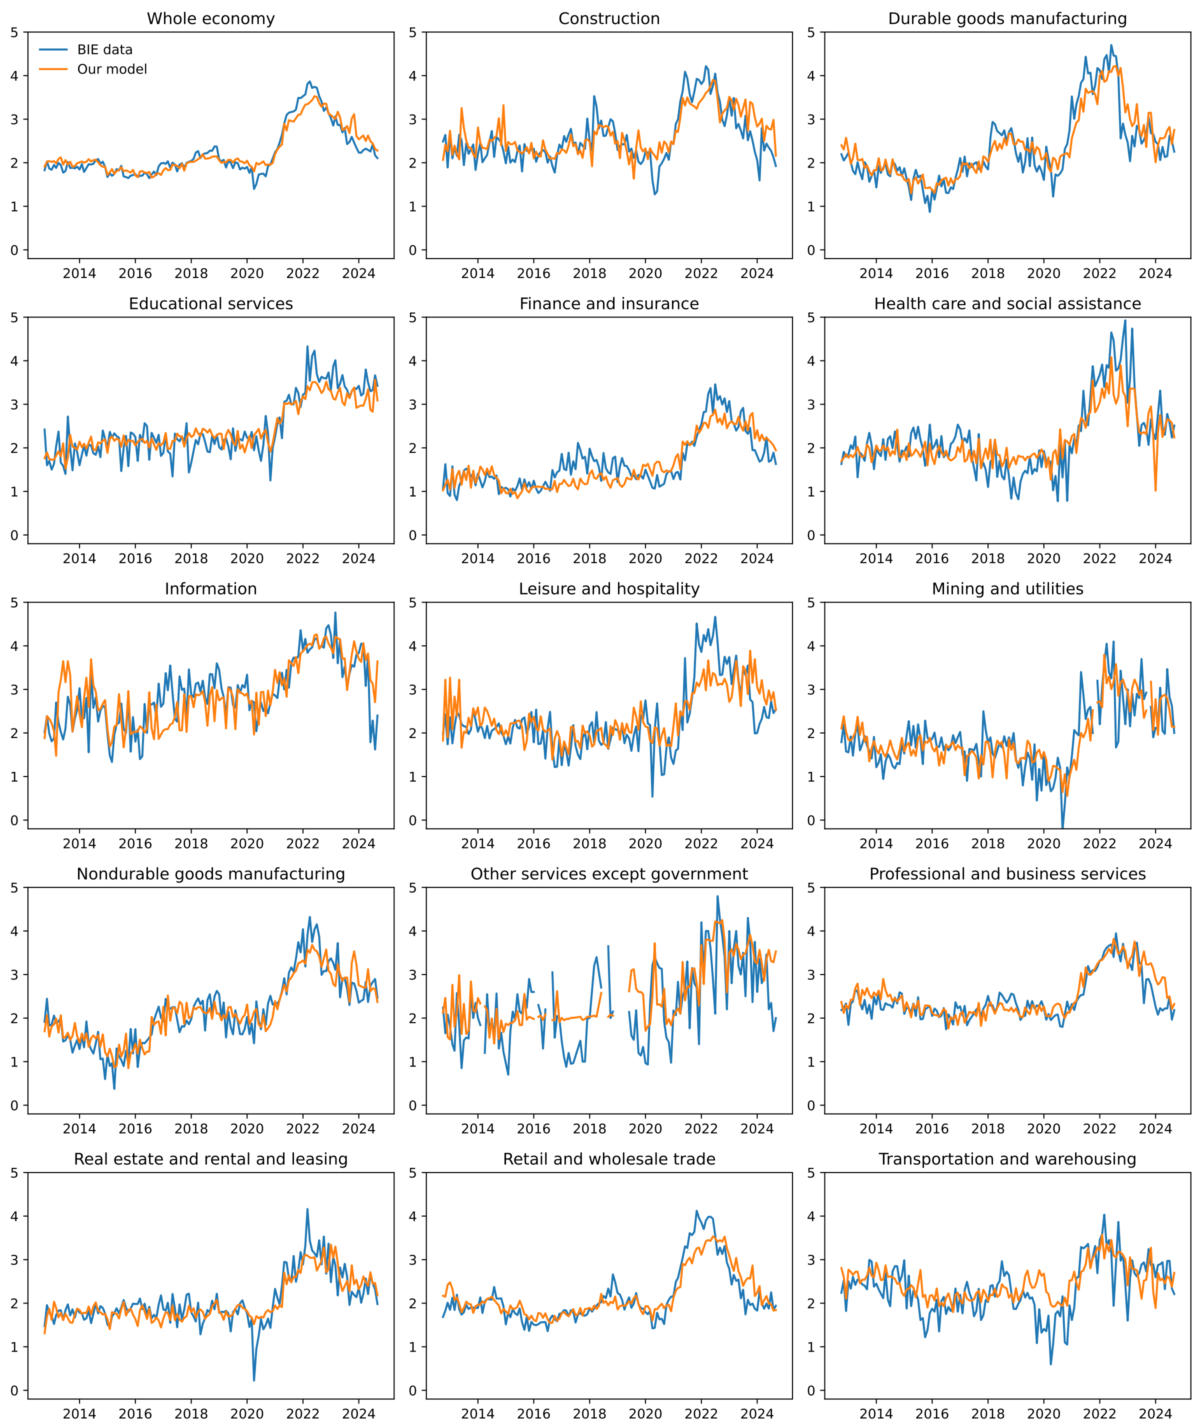

Figure B.14: Comparing model fits: Response of belief to an oil price change

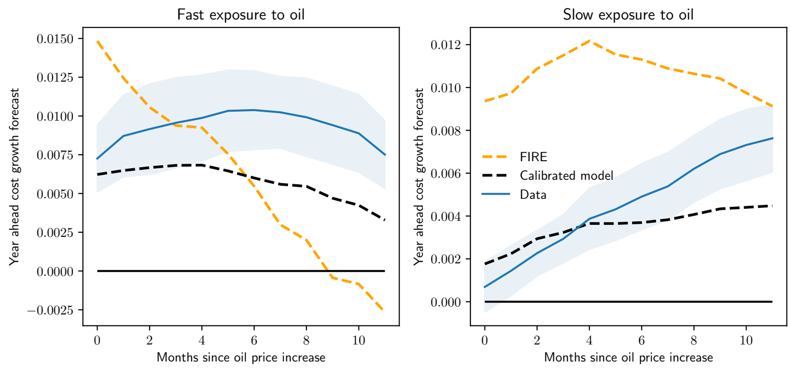

where the idiosyncratic component follows an AR(1): $\Delta_{12}\text{mc}_{i,t}^{idio} = \rho^{\text{idio}}\Delta_{12}\text{mc}_{i,t-12}^{idio} + e_{i,t}$ and $e_{i,t} \overset{\text{iid}}{\sim} \mathcal{N}\left(0, \sigma_e^2\right)$. We know that over the BIE sample period, the year-on-year percent change in the log GDP deflator (denoted $\Delta_{12}p_t^{\text{GDP}}$) provides a very close approximation to the average of the BIE cost growth series (see Figure A.2). So we use this series to capture aggregate costs. To estimate the idiosyncratic cost process, we subtract this aggregate measure from overall costs and estimate ρ^{idio} and σ_e over the sample period, at 0.08 and 1.41, respectively. We can then simulate costs in earlier periods as

$$\Delta_{12}\text{mc}_{i,t}^{\text{sim}} = \Delta_{12}p_t^{\text{GDP}} + \Delta_{12}\text{mc}_{i,t}^{idio,sim} \tag{70}$$

where $\Delta_{12}\text{mc}_{i,t}^{idio,sim}$ is drawn from an AR(1) process with the calibrated persistence and shock variance. This allows us to extend a cost series back to 1949, under the key assumptions that (1) the average costs in the BIE also closely tracked the GDP deflator before 2011, and (2) the 'idiosyncratic' component of costs followed a similar process before 2011 as after it.

B.9 Overreaction in the long run

In the main text, we show that the one-year ahead forecast error,

$$\text{FE}_{i,t,t+12} \equiv \Delta_{12}\text{mc}_{i,t+12} - \mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right] \tag{71}$$

is predictable from cost growth over the last year. In this section, we ask if this overreaction can be explained entirely as a result of over-extrapolating permanent shocks. Decompose cost growth into a permanent component, and deviations from this,

$$\Delta_{12}\text{mc}_{i,t} = \Delta_{12}\text{mc}_{i,t}^{\text{Permanent}} + \Delta_{12}\text{mc}_{i,t}^{\text{Transitory}} \tag{72}$$

Then, we can likewise decompose the forecast error

$$\mathrm{FE}_{i,t,t+12} = \mathrm{FE}^{\mathrm{SR}}_{i,t,t+12} + \mathrm{FE}^{\mathrm{LR}}_{i,t,t+12} \tag{73}$$

where

$$\mathrm{FE}^{\mathrm{LR}}_{i,t,t+12} = \Delta_{12}\mathrm{mc}^{\mathrm{Permanent}}_{i,t+12} - \mathbb{E}_{i,t}\left[\Delta_{12}\mathrm{mc}^{\mathrm{Permanent}}_{i,t}\right]$$

and

$$\mathrm{FE}^{\mathrm{SR}}_{i,t,t+12} = \Delta_{12}\mathrm{mc}^{\mathrm{Transitory}}_{i,t+12} - \mathbb{E}_{i,t}\left[\Delta_{12}\mathrm{mc}^{\mathrm{Transitory}}_{i,t}\right]$$

On a consistent sample, it is then clear that the regression coefficient of the overall one-year ahead forecast error on $\Delta_{12}\mathrm{mc}_{i,t}$ equals the sum of the coefficients from regressing the long-run and the short-run forecast errors on that same regressor. We can then ask: how much of the overreaction to own costs is about perceived permanent cost dynamics, versus short-run dynamics around a permanent level.

To implement this in the data, we make use of a BIE survey question asking firms to forecast their unit cost growth five- to ten-years ahead. We treat this as equal to the belief of the permanent component (i.e. we assume the transitory component is not relevant at this horizon). Finally, we need to know the actual movements in the permanent component. This is not trivial to map to the data, as the horizon is not clear (five or ten years), and we have relatively few observations a full ten years apart. And recall from Table 2 that we estimated an annual persistence of overall own cost growth around 0.1. Here we assume that conditional on sector time fixed effects, own cost movements do not influence the true five to ten-year ahead forecasts. Note that if there is in fact some positive long-run predictive power, this will bias our analysis towards suggesting that long-run overreaction drives our result.

Our findings are shown in Table B.9. We run this on a consistent sample.[91] We see that cost growth over the past year significantly predicts an overreaction both about the long-run, and deviations from the long run.

Table B.9: Firms over-react to cost growth at a long and short horizon

	$\mathrm{FE}_{i,t,t+12}$	$\mathrm{FE}^{\mathrm{SR}}_{i,t,t+12}$	$\mathrm{FE}^{\mathrm{LR}}_{i,t,t+12}$
$\Delta_{12}\mathrm{mc}_{i,t}$	-0.23***	-0.10***	-0.13***
	(0.03)	(0.03)	(0.02)
Fixed effects	Firm and Sector-Date	Firm and Sector-Date	Firm and Sector-Date
Observations	6228	6228	6228

Note: $^{*} < 0.10$, $^{**} < 0.05$, $^{***} < 0.01$. Standard errors clustered at the firm level.

[91] The sample is much smaller as firms are asked to forecast long-run costs only quarterly. However, note that the overall coefficient of the forecast error on own costs, -0.23, is very close to that in Table 2, -0.22.

B.10 Historical importance of CPI in predicting own costs

The results are shown in Table B.10.

Table B.10: Estimating predictability of future costs with future CPI using (23)

	Year-on-year growth in GDP deflator, $\pi^{PGDP}_{t,t-12}$
Const	0.27*
	(0.15)
$\pi^{PGDP}_{t-12,t-24}$	0.17**
	(0.07)
$\pi^{CPI}_{t,t-12}$	0.68***
	(0.07)
Observations	308
R-squared	0.90

Note: *p<0.10, **p<0.05, ***p<0.01. HAC standard errors with maximum lag at 12 quarters.

B.11 Other models of beliefs

Table B.11: Which models can match these facts?

Model	**Fact 1** CPI	**Fact 2** Overreact	**Fact 3** Stability	**Fact 4** Inertia	**Fact 5** Oil
FIRE	✓	×	×	×	×
LIRE Mankiw and Reis (2002), Maćkowiak and Wiederholt (2009)	✓	×	✓	×	×
Cognitive discounting Gabaix (2020)	✓	×	✓	?	×
Diagnostic expectations Bordalo et al. (2018)	✓	✓	×	×	×
Delayed overshooting Angeletos et al. (2021)	✓	✓	✓	✓	×

Mankiw and Reis (2002) In this model, every period each firm has a probability θ^{MR} of learning about the current state of the economy, while other firms continue to base their pricing on information from the last time they were informed. So,

$$\mathbb{E}_{i,t}^{\text{MR}}\left[\text{mc}_{t+\tau}\right] = \begin{cases} \mathbb{E}_t\left[\text{mc}_{t+\tau}\right] & \text{with probability } 1-\theta^{\text{MR}} \\ \mathbb{E}_{i,t-1}^{\text{MR}}\left[\text{mc}_{t+\tau}\right] & \text{with probability } \theta^{\text{MR}} \end{cases}$$

Carroll et al. (2020) In the Mankiw-Reis model, firms may not be up-to-date on macroeconomic information; but they might also not know their current costs. In the spirit of a sticky expectations model, we allow firms to know their current costs perfectly, even as their forecasts for future costs might be based on out-of-date information.

$$\mathbb{E}_{i,t}^{\text{SE}}\left[\text{mc}_{t+\tau}\right] = \begin{cases} \text{mc}_t & \text{if } \tau = 0 \\ \mathbb{E}_t\left[\text{mc}_{t+\tau}\right] & \text{with probability } 1-\theta^{\text{SE}}, \text{ if } \tau > 0 \\ \mathbb{E}_{i,t-1}^{\text{SE}}\left[\text{mc}_{t+\tau}\right] & \text{with probability } \theta^{\text{SE}}, \text{ if } \tau > 0 \end{cases}$$

Gabaix (2020) In the cognitive discounting model, we suppose that firms attenuate the FIRE forecast for their future costs towards some default. In order to allow the level of nominal costs to permanently move, we follow Gabaix (2020) in applying the myopia to deviations of costs from current prices,[92]

$$\mathbb{E}_{i,t}^{\text{CD}}\left[\text{mc}_{i,t+\tau} - p_t\right] = \left(m^{\text{CD}}\right)^{\tau} \mathbb{E}_t\left[\text{mc}_{i,t+\tau} - p_t\right]$$

Bordalo et al. (2018) We follow the formulation of diagnostic expectations that

$$\mathbb{E}_{i,t}^{\text{DE}}\left[\text{mc}_{i,t+\tau}\right] = \mathbb{E}_t\left[\text{mc}_{i,t+\tau}\right] + \theta^{\text{DE}}\left(\mathbb{E}_t\left[\text{mc}_{i,t+\tau}\right] - \mathbb{E}_{t-1}\left[\text{mc}_{i,t+\tau}\right]\right)$$

The central idea is that firms overreact to surprises about their marginal costs.

B.12 CPI and own-cost beliefs in Maćkowiak and Wiederholt (2009)

Consider the transitory shock case in Section IV of Maćkowiak and Wiederholt (2009) that is solved analytically. In this case, real output is

$$y_t = \frac{1-\alpha^{\text{MW}}}{\alpha^{\text{MW}}} p_t$$

where $\alpha^{\text{MW}} \in [0,1]$. α^{MW} is greater than zero when firms pay any attention to aggregate conditions, and less than one when the solution deviates from FIRE.

Consider the flexible wage setting where real marginal costs are increasing in consumption

[92] This leads to the same NKPC as derived in Gabaix (2020).

and hours worked. Then

$$\text{mc}_{i,t} = p_t + \sigma y_t = \left(1 + \sigma \frac{1 - \alpha^{\text{MW}}}{\alpha^{\text{MW}}}\right) p_t$$

So then, beliefs about nominal marginal cost growth and CPI over the next year satisfy

$$\mathbb{E}_{i,t}\left[\Delta \text{mc}_{i,t+1}\right] = \left(1 + \sigma \frac{1 - \alpha^{\text{MW}}}{\alpha^{\text{MW}}}\right) \mathbb{E}_{i,t}\left[\pi_{t+1}\right]$$

So the pass-through is weakly greater than one-to-one.

The intuition is more general than this analytic case. In Maćkowiak and Wiederholt (2009), the only aggregate shocks drive the price level and real output is the same direction. Therefore, revisions in beliefs about the aggregate price level lead to revisions, in the same direction, about aggregate real output. So in any setting where real costs are increasing in real output, this implies that nominal cost beliefs move more than one-to-one in CPI beliefs.

C Additional Theorems and Proofs

C.1 Lemma for Theorem 3

Lemma 3. *Decompose nominal marginal costs into the real marginal cost and the price level. Then, applying* $1 - \beta\theta_p \mathbb{E}_t F$ (41) *becomes*

$$\begin{aligned}
\pi_t = {} & \frac{1-\theta_p}{\theta_p}\left(\left(1 - \alpha^{FIRE}\right) + \alpha^{FIRE}\left(1 - \beta\theta_p\right)\right) mc_t^{real} \\
& - \left(1 - \alpha^{FIRE}\right)\left(1 - \theta_p\right)\beta\mathbb{E}_t\left[mc_{t+1}^{real}\right] \\
& + \left(1 - \alpha^{FIRE}\right)\frac{\beta\tilde{\rho}\left(1-\theta_p\right)}{1 - \beta\theta_p\tilde{\rho}}\left(b_t - \beta\theta_p\mathbb{E}_t\left[b_{t+1}\right]\right) \\
& + \left(\theta_p + \alpha^{FIRE}\left(1 - \theta_p\right)\right)\beta\mathbb{E}_t\left[\pi_{t+1}\right]
\end{aligned} \tag{74}$$

And (37) *becomes*

$$b_t = \tilde{\mathcal{K}}\Delta mc_t^{real} + \tilde{\mathcal{K}}\pi_t + \left(1 - \tilde{\mathcal{K}}\right)\tilde{\rho} b_{t-1} \tag{75}$$

For any ζ, *we can write*

$$\begin{aligned}
\mathbb{E}_t\left[b_{t+1}\right] = {} & \zeta\mathbb{E}_t\left[b_{t+1}\right] \\
& + (1-\zeta)\tilde{\mathcal{K}}\mathbb{E}_t\left[\Delta mc_{t+1}^{real}\right] + (1-\zeta)\tilde{\mathcal{K}}\mathbb{E}_t\left[\pi_{t+1}\right] + (1-\zeta)\left(1 - \tilde{\mathcal{K}}\right)\tilde{\rho} b_t
\end{aligned}$$

Define

$$\delta \equiv \alpha^{FIRE} + \left(1 - \alpha^{FIRE}\right) \theta_p \left[1 - (1-\zeta)\mathcal{M}\right]$$

and

$$\omega_b \equiv \left(1 - \alpha^{FIRE}\right) \frac{\beta\tilde{\rho}\left(1-\theta_p\right)}{1-\beta\theta_p\tilde{\rho}} \left[1 - (1-\zeta)\beta\theta_p\left(1-\tilde{\mathcal{K}}\right)\tilde{\rho}\right]$$

where $\mathcal{M} \equiv \frac{\left(1-\theta_p\right)\beta\tilde{\rho}\tilde{\mathcal{K}}}{1-\beta\theta_p\tilde{\rho}}$. *Then we can combine these to get*

$$\begin{aligned}\pi_t - \omega_b b_t = &\left[\frac{1-\theta_p}{\theta_p} + \beta\left(\theta_p - \delta\right)\right] mc_t^{real} \\ &- (1-\delta)\beta\mathbb{E}_t\left[mc_{t+1}^{real}\right] \\ &+ \beta\delta\left(\mathbb{E}_t\left[\pi_{t+1}\right] - \frac{\left(1-\alpha^{FIRE}\right)\frac{\beta\tilde{\rho}\left(1-\theta_p\right)}{1-\beta\theta_p\tilde{\rho}}\theta_p\zeta}{\delta}\mathbb{E}_t\left[b_{t+1}\right]\right)\end{aligned} \tag{76}$$

Notice that this still holds for any ζ. *We choose* ζ *such that*

$$\omega_b = \left(1-\alpha^{FIRE}\right)\frac{\left(1-\theta_p\right)\beta\tilde{\rho}\theta_p\zeta}{\delta\left(1-\beta\theta_p\tilde{\rho}\right)}$$

Given this choice, we immediately recover (42).

There are generally two solutions to the implied fixed point for δ. *We next show that one always lies in* $(\theta_p, 1)$ *and the other is greater than* $1/\beta$. *The fixed point is*

$$f(\delta) = 0$$

where

$$\begin{aligned}f(\delta) = &\beta\left(1-\tilde{\mathcal{K}}\right)\tilde{\rho}\cdot\delta^2 \\ &+ \left\{\mathcal{M}\left(1-\alpha^{FIRE}\right) - \beta\left(1-\tilde{\mathcal{K}}\right)\tilde{\rho}\left(\left(1-\alpha^{FIRE}\right)\theta_p + \alpha^{FIRE}\right) - 1\right\}\cdot\delta \\ &+ \left\{\alpha^{FIRE} + \theta_p(1-\mathcal{M})\left(1-\alpha^{FIRE}\right)\right\}\end{aligned}$$

Then, notice that

$$f\left(\theta_p\right) = \alpha^{FIRE}\left(1-\theta_p\right)\left(1-\beta\tilde{\rho}\theta_p\left(1-\theta_p\right)\left(1-\tilde{\mathcal{K}}\right)\right) > 0$$

but

$$f(1) = \left(1-\alpha^{FIRE}\right)\left(1-\theta_p\right)\left[\mathcal{M} - \left(1-\beta\left(1-\tilde{\mathcal{K}}\right)\tilde{\rho}\right)\right] < 0$$

since $\mathcal{M} < 1 - \beta\left(1 - \tilde{\mathcal{K}}\right)\tilde{\rho}$. This implies one root lies in $\left(\theta_p, 1\right)$. Next, notice that

$$f\left(\beta^{-1}\right) = \left(1 - \left(1 - \tilde{\mathcal{K}}\right)\tilde{\rho}\right)\left[\left(1 - \alpha^{FIRE}\right)\theta_p + \alpha^{FIRE} - \beta^{-1}\right] + \left(1 - \alpha^{FIRE}\right)\mathcal{M}\left(\beta^{-1} - \theta_p\right)$$

Since $\mathcal{M}\left(\beta^{-1} - \theta_p\right) < \left(1 - \left(1 - \tilde{\mathcal{K}}\right)\tilde{\rho}\right)\left(1 - \theta_p\right)$,

$$f\left(\beta^{-1}\right) < \left(1 - \left(1 - \tilde{\mathcal{K}}\right)\tilde{\rho}\right)\left(1 - \beta^{-1}\right) < 0$$

Therefore, the second root is above β^{-1}.

C.2 Derivation of the Friedman-Phelps Phillips curve

Suppose that inflation depends on current real marginal costs and expected inflation,

$$\pi_t = \kappa^{\text{AE}}\text{mc}_t^{\text{real}} + \alpha^{\text{AE}}\pi_t^e \tag{77}$$

Take the (discrete time) expectations formulation in Cagan (1956), and used by Phelps (1967),

$$\pi_t^e = (1 - \gamma^{\text{AE}})\pi_{t-1}^e + \gamma^{\text{AE}}\pi_t \tag{78}$$

Combining (77) and (78),

$$\pi_t = \kappa^{\text{AE}}\text{mc}_t^{\text{real}} + \alpha^{\text{AE}}\gamma^{\text{AE}}\sum_{\ell=0}^{\infty}(1 - \gamma^{\text{AE}})^{\ell}\pi_{t-\ell} \tag{79}$$

C.3 Adaptive Learning Comparative Statics

Lemma 4. *The AEPC slope increases with price flexibility:*

$$\frac{d\kappa^{AL}}{d\theta_P} < 0.$$

Proof. First note that

$$\frac{d\kappa^{AL}}{d\theta_P} = \left(\frac{1-\theta_P}{\theta_P} + \mathcal{M}\right)\frac{d}{d\theta_P}\left(\frac{1}{1-\mathcal{M}}\right) + \frac{1}{1-\mathcal{M}}\left(-\frac{1}{\theta_P^2} + \frac{d\mathcal{M}}{d\theta_P}\right),$$

and

$$\frac{d}{d\theta_P}\left(\frac{1}{1-\mathcal{M}}\right) = \frac{1}{(1-\mathcal{M})^2}\frac{d\mathcal{M}}{d\theta_P}.$$

On the interior of the parameter space, $\mathcal{M} \in (0,1)$, and so $\frac{d\kappa^{AL}}{d\theta_P} < 0$ when $\frac{d\mathcal{M}}{d\theta_P} < 0$. We have

$$\frac{d\mathcal{M}}{d\theta_P} = -\frac{\beta\tilde{\rho}\tilde{\mathcal{K}}(1-\beta\tilde{\rho})}{(1-\beta\theta_P\tilde{\rho})^2} < 0.$$

$\square$

Lemma 5. *The adaptive Phillips curve under our beliefs becomes more inertial as prices become more flexible. In particular, the cumulative weight of past inflation up to any lag, L, satisfies*

$$\frac{d}{d\theta_p}\left[\omega_b^{AL}\sum_{\ell=1}^{L}\left(\lambda^{AL}\right)^{\ell}\right]\leq 0 \tag{80}$$

Proof. Define $\Omega_L = \omega_b^{\mathrm{AL}}\sum_{\ell=1}^{L}\left(\lambda^{\mathrm{AL}}\right)^{\ell}$. Then

$$\Omega_L = \Omega_\infty\left[1-\left(\lambda^{\mathrm{AL}}\right)^{L}\right]$$

We start by showing that $d\Omega_\infty/d\theta_p \leq 0$, i.e. the total weight of past inflation increases in price flexibility. First note that

$$\Omega_\infty = \frac{\omega_b^{\mathrm{AL}}\lambda^{\mathrm{AL}}}{1-\lambda^{\mathrm{AL}}} = \frac{\beta\tilde{\rho}\tilde{\mathcal{K}}}{1-\left(1-\tilde{\mathcal{K}}\right)\tilde{\rho}}\cdot\left[\frac{\left(1-\tilde{\mathcal{K}}\right)\tilde{\rho}}{\left(1-\beta\theta_p\tilde{\rho}\right)-\left(1-\theta_p\right)\beta\tilde{\rho}\tilde{\mathcal{K}}}-\frac{\theta_p}{\left(1-\beta\theta_p\tilde{\rho}\right)}\right]$$

Taking the derivative, to get $d\Omega_\infty/d\theta_p < 0$, we need that

$$\sqrt{\beta}\tilde{\rho}\left(1-\tilde{\mathcal{K}}\right)\left(1-\beta\theta_p\tilde{\rho}\right) < \left(1-\beta\theta_p\tilde{\rho}\right)-\left(1-\theta_p\right)\beta\tilde{\rho}\tilde{\mathcal{K}}$$

This is implied by the stronger inequality dropping the $\sqrt{\beta}$, which gives

$$0 < \left(1-\beta\theta_p\tilde{\rho}\right)\left(1-\tilde{\rho}\right)+\tilde{\rho}\tilde{\mathcal{K}}\left(\left(1-\beta\right)+\beta\theta_p\left(1-\tilde{\rho}\right)\right)$$

We next show that this holds to any lag horizon. When $\Omega_\infty \geq 0$, a sufficient condition for this to hold is

$$\frac{d}{d\theta_p}\left[1-\left(\lambda^{\mathrm{AL}}\right)^{L}\right]\leq 0 \iff \frac{d\lambda^{\mathrm{AL}}}{d\theta_p}\geq 0$$

since $L\left(\lambda^{\mathrm{AL}}\right)^{L-1}\geq 0$. So this follows immediately from $\frac{d\lambda^{\mathrm{AL}}}{d\theta_p}\geq 0$.

If $\Omega_\infty < 0$, then use that

$$\frac{d\Omega_L}{d\theta_p} = \frac{d\Omega_\infty}{d\theta_p}\left[1-\left(\lambda^{\mathrm{AL}}\right)^{L}\right]-\Omega_\infty L\left(\lambda^{\mathrm{AL}}\right)^{L-1}\frac{d\lambda^{\mathrm{AL}}}{d\theta_p}$$

and $\left(\lambda^{\mathrm{AL}}\right)^{L-1}L \leq \frac{1-\left(\lambda^{\mathrm{AL}}\right)^{L}}{1-\lambda^{\mathrm{AL}}}$ to get

$$\frac{d\Omega_L}{d\theta_p}\leq\frac{1-\left(\lambda^{\mathrm{AL}}\right)^{L}}{1-\lambda^{\mathrm{AL}}}\frac{d\Omega_1}{d\theta_p}$$

Therefore if the condition holds for $L=1$, it holds for all L. From $\Omega_1 = \frac{\left(1-\tilde{\mathcal{K}}\right)\tilde{\rho}\frac{\mathcal{M}}{1-\mathcal{M}}-\mathcal{M}\theta_p}{\left(1-\theta_p\right)+\mathcal{M}\theta_p}$, and

$\mathcal{M} \equiv \frac{(1-\theta_p)\beta\tilde{\rho}\tilde{\mathcal{K}}}{(1-\beta\theta_p\tilde{\rho})}$, we can show that $d\Omega_1/d\theta_p < 0$. $\square$

C.4 Adaptive Learning Phillips curve with observed permanent cost shocks

Recall that in Section B.3, we characterized beliefs when firms perceive cost growth as driven by a permanent, a persistent, and a transitory component, and they are aware of the persistent level of cost growth. Using (67) and (68), we derive the Phillips curve

$$\begin{aligned}\hat{\pi}_t = &\frac{1}{1-\mathcal{M}}\left[\mathcal{M}+\frac{(1-\theta_p)}{\theta_p}\right]\hat{\mathrm{mc}}_t^{\mathrm{real}} - \frac{1}{1-\mathcal{M}}\left[\mathcal{M}+\frac{(1-\theta_p)}{\theta_p}\left(1-\tilde{\mathcal{K}}\right)\tilde{\rho}\right]\hat{\mathrm{mc}}_{t-1}^{\mathrm{real}} \\ &+\mathcal{B}\hat{\pi}_{t-1} \\ &+\frac{\mathcal{M}}{1-\mathcal{M}}\frac{\theta_p(1-\beta)}{(1-\theta_p)(1-\beta\theta_p)}\Delta\pi_t^{\mathrm{LR}}\end{aligned} \tag{81}$$

where $\hat{\pi}_t \equiv \pi_t - \pi_t^{\mathrm{LR}}$, $\hat{\mathrm{mc}}_t^{\mathrm{real}} \equiv \mathrm{mc}_t^{\mathrm{real}} - \mathrm{mc}_t^{\mathrm{real,\,LR}}$, $\pi_t^{\mathrm{LR}} \equiv \mathbb{E}_t\left[\pi_\infty\right]$, and $\mathrm{mc}_t^{\mathrm{real,\,LR}} \equiv \mathbb{E}_t\left[\mathrm{mc}_\infty^{\mathrm{real,\,LR}}\right]$. Also $\mathcal{M} \equiv \frac{(1-\theta_p)\beta\tilde{\rho}\tilde{\mathcal{K}}}{1-\beta\theta_p\tilde{\rho}}$ and $\mathcal{B} \equiv \frac{(1-\tilde{\mathcal{K}})\tilde{\rho}}{1-\mathcal{M}}$.

When there are no permanent shocks, $\Delta\pi_t^{\mathrm{LR}} = 0$, (81) is identical to (43) except that it characterizes deviations from the long run. Any permanent real marginal cost shocks affect the long-run level of inflation the same in our model as under FIRE, i.e. $\pi_t^{\mathrm{LR}} = \frac{(1-\theta_p)(1-\beta\theta_p)}{\theta_p(1-\beta)}\mathrm{mc}_t^{\mathrm{real,\,LR}}$. This implies that the long-run trade-off between inflation and output for the policy maker is the same as in the standard New Keynesian model. However, in our model, such a shock generates additional short-run inflation.

C.5 Forward Guidance with Myopic Households

Since firms are not themselves very forward-looking in our calibration, their responsiveness to future real interest rates is relatively more dependent on how much costs move today. This suggests that when it comes to news about the future, other agents' expectations may be more important in our model. This intuition is born out in Theorem A1, which suggests an interesting interaction between household myopia and Adaptive Learning:

Theorem A1. *Consider the IS equation but with myopia (see Gabaix (2020))*

$$y_t = -eis \cdot r_t + m^{hh}\mathbb{E}_t y_{t+1}$$

Suppose that at date 0, the central bank announces an interest rate cut at date h that will fully revert at date h + 1. The effect on inflation at date 0, under FIRE and Adaptive Learning (AL), respectively, are

$$-\frac{d\pi_0^{FIRE}}{dr_h} = eis \cdot \varphi\kappa^{FIRE} \cdot \left(m^{hh}\right)^h \sum_{\tau=0}^{h}\left(\frac{\beta}{m^{hh}}\right)^\tau, \quad -\frac{d\pi_0^{AL}}{dr_h} = eis \cdot \varphi\kappa^{AL} \cdot \left(m^{hh}\right)^h.$$

Proof. Iterating the IS equation forward and imposing the boundary condition $\lim_{H\to\infty}(m^{\text{hh}})^H y_H = 0$, we have $y_0 = -eis\sum_{\tau=0}^{\infty}(m^{hh})^\tau r_\tau$. Therefore, $\frac{dy_0}{dr_h} = -\text{eis}\cdot(m^{\text{hh}})^h$. Similarly, $\frac{dy_\tau}{dr_h} = -\text{eis}\cdot(m^{\text{hh}})^{h-\tau}$ for all $\tau \in \{0,1,...,h\}$. The FIRE and AL NKPCs can be solved forward exactly as in Theorem 5. Taking the derivative and plugging in the output effects just derived yields the result. $\square$

C.6 Optimal Monetary Policy

We now turn to optimal monetary policy. To consider the canonical setting, we allow wages to be flexible. As labor is the only input to production, real marginal costs are proportional to output:

$$\text{mc}_t^{\text{real}} = \varphi y_t \tag{82}$$

where $\varphi \equiv \frac{1}{\text{frisch}} + \frac{1}{\text{eis}}$. We break divine coincidence in our model by introducing cost-push shocks. In particular, we microfound these as markup shocks exactly as in Galí (2015) Chapter 5.2, so that the log-linearized optimal reset price is

$$p_{i,t}^* = \left(1-\beta\theta_p\right)\sum_{\tau=0}^{\infty}\left(\beta\theta_p\right)^\tau \mathbb{E}_{i,t}\left[\mu_{i,t+\tau} + \text{mc}_{i,t+\tau}\right] \tag{83}$$

where $\mu_{i,t+\tau}$ are time-varying desired markups. We assume that firms in our model hold FIRE beliefs about their own desired markups. We also introduce a standard quadratic loss function for the central bank

$$\mathcal{L}_t = \mathbb{E}_t\sum_{\tau=0}^{\infty}\beta^\tau\left[\pi_{t+\tau}^2 + \vartheta y_{t+\tau}^2\right] \tag{84}$$

Under full-information rational expectations (FIRE), this gives rise to the familiar results. Under commitment, the optimal policy response to a given shock $\{\mu_t\}$ satisfies

$$\mathbb{E}_0\left[\pi_t + \frac{\vartheta}{\kappa^{\text{FIRE}}}\left(y_t - y_{t-1}\right)\right] = 0 \tag{85}$$

where $\kappa^{\text{FIRE}} \equiv \frac{\left(1-\theta_p\right)\left(1-\beta\theta_p\right)}{\theta_p}\varphi$.[93] And under discretion, optimal policy induces

$$\pi_t + \frac{\vartheta}{\kappa^{\text{FIRE}}} y_t = 0 \tag{86}$$

In Theorem A2, we derive optimal monetary policy under Adaptive Learning.

Theorem A2. *Under purely Adaptive Learning beliefs ($\alpha^{FIRE} = 0$), optimal monetary policy is the same*

[93] This definition of κ^{FIRE} differs from elsewhere in the text by including φ. We've done this here to condense notation.

under commitment and discretion, and is given by

$$\left(\frac{1-\theta_p}{\theta_p}+\mathcal{M}\right)\pi_t+(1-\mathcal{M})\frac{\vartheta}{\varphi}y_t$$
$$=\beta\mathbb{E}_t\left[\left(\frac{1-\theta_p}{\theta_p}\left(1-\tilde{\mathcal{K}}\right)\tilde{\rho}+\mathcal{M}\right)\pi_{t+1}+\left(1-\tilde{\mathcal{K}}\right)\tilde{\rho}\frac{\vartheta}{\varphi}y_{t+1}\right] \tag{87}$$

Proof. To get the NKPC with markups in the Adaptive Learning model, we combine the law of motion for prices, $p_t=\theta_p p_{t-1}+\left(1-\theta_p\right)p_t^*$, the optimal reset price (83), and the belief model in (12) and (13). This gives

$$\pi_t=\frac{1}{1-\mathcal{M}}\left\{\left[\frac{1-\theta_p}{\theta_p}+\mathcal{M}\right]\mathrm{mc}_t^{\mathrm{real}}-\left[\frac{1-\theta_p}{\theta_p}\left(1-\tilde{\mathcal{K}}\right)\tilde{\rho}+\mathcal{M}\right]\mathrm{mc}_{t-1}^{\mathrm{real}}\right\}$$
$$+\frac{\left(1-\tilde{\mathcal{K}}\right)\tilde{\rho}}{(1-\mathcal{M})}\pi_{t-1} \tag{88}$$
$$+\frac{1}{1-\mathcal{M}}\frac{\left(1-\theta_p\right)\left(1-\beta\theta_p\right)}{\theta_p}\left(m_t-\left(1-\tilde{\mathcal{K}}\right)\tilde{\rho}m_{t-1}\right)$$

where $\mathcal{M}\equiv\frac{(1-\theta_p)\beta\tilde{\rho}\tilde{\mathcal{K}}}{1-\beta\theta_p\tilde{\rho}}$ and $m_t=\sum_{\tau=0}^{\infty}\left(\beta\theta_p\right)^{\tau}\mathbb{E}_t\left[\mu_{i,t+\tau}\right]$. Given the objective (84) and this NKPC, the first order conditions are

$$\lambda_t=\pi_t+\beta\frac{\left(1-\tilde{\mathcal{K}}\right)\tilde{\rho}}{(1-\mathcal{M})}\mathbb{E}_0\lambda_{t+1}$$

$$\lambda_t=-\frac{1-\mathcal{M}}{\frac{1-\theta_p}{\theta_p}+\mathcal{M}}\frac{\vartheta}{\varphi}y_t+\beta\frac{\frac{1-\theta_p}{\theta_p}\left(1-\tilde{\mathcal{K}}\right)\tilde{\rho}+\mathcal{M}}{\frac{1-\theta_p}{\theta_p}+\mathcal{M}}\mathbb{E}_0\lambda_{t+1}$$

where λ_t is the lagrangian multiplier on the NKPC. Combining these, we get (87). □

D Model extensions

D.1 Comparison across belief models

In Figure D.15 we show that the price level converges to the same point under all the belief models.

D.2 Jacobian of the New Keynesian Phillips Curve across belief models

In Figure D.16 we plot the Jacobians of the NKPC in various models.

D.3 The Adaptive Learning Phillips curve with learning about a permanent cost component

We begin by asking if our model implies a different inflation-output menu for policy makers, in the spirit of Samuelson and Solow (1960). Suppose that firms believe cost shocks are permanent

Figure D.15: Percent price level response to 1pp permanent change in nominal marginal costs at date s

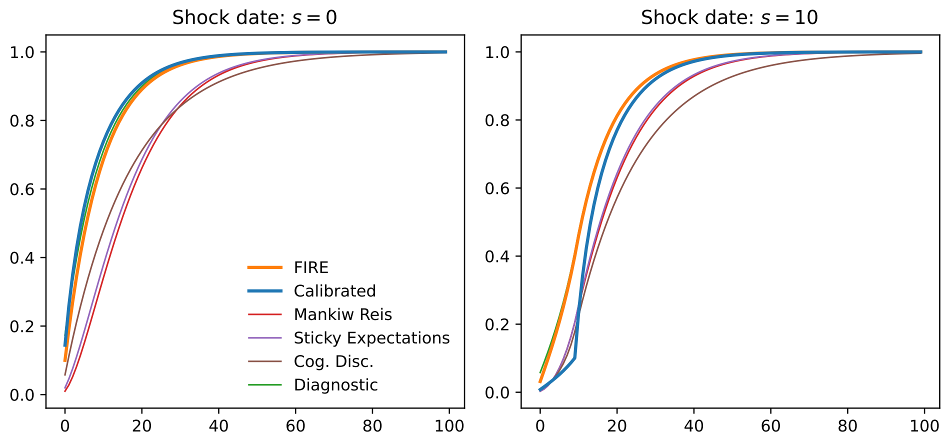

Note: Impulse response of prices to a permanent rise in nominal marginal costs under six models. 'Calibrated' is our model with α^{FIRE}, $\tilde{\rho}$ and $\tilde{\mathcal{K}}$ calibrated as in Section 2.2

or transitory, so $\tilde{\rho} = 1$. It is then easy to see that under Adaptive Learning, to maintain marginal costs (unemployment) permanently high (low) requires steady state inflation at

$$\pi = \frac{\kappa^{\text{FIRE}}}{1-\beta} \text{mc}^{\text{real}} \tag{89}$$

where $\kappa^{\text{FIRE}} \equiv \frac{(1-\theta_p)(1-\beta\theta_p)}{\theta_p}$. This is identical to the trade-off faced under FIRE. Furthermore, the result holds in a model with a transitory, a persistent, and a permanent component. Intuitively, the firms eventually learn the new steady state cost inflation rate.

One immediate implication of this result is that

$$\left(\frac{\pi}{\text{mc}^{\text{real}}}\right)_{\text{LR}}^{\text{AL}} = \frac{\left(1-\lambda^{\text{AL}}\right)\kappa^{\text{AL}}}{1-\left(1+\omega_b^{\text{AL}}\right)\lambda^{\text{AL}}} = \frac{\kappa^{\text{FIRE}}}{1-\beta} = \left(\frac{\pi}{\text{mc}^{\text{real}}}\right)_{\text{LR}}^{\text{FIRE}} \tag{90}$$

That is, the sum of the weights on past inflation $\Omega^{\text{AL}} \equiv \omega_b^{\text{AL}} \sum_{\ell=1}^{\infty} \left(\lambda^{\text{AL}}\right)^{\ell}$ must satisfy $\frac{\kappa^{\text{AL}}}{\kappa^{\text{FIRE}}} = \frac{1-\Omega^{\text{AL}}}{1-\beta}$. Typically this suggests $\Omega^{\text{AL}} \gg 0$. However, the effect on inflation today from a one off real marginal cost shock yesterday, $\frac{\partial \pi_1}{\partial \text{mc}_0^{\text{real}}} / \frac{\partial \pi_0}{\partial \text{mc}_0^{\text{real}}} = \omega_b^{\text{AL}} \lambda^{\text{AL}}$ may remain small.

D.4 Effect of demand and supply shocks on inflation under a Taylor rule

Here we take our simple model from Section 4.1, and again feed demand and supply shocks into models with various firm beliefs. Only now, instead of a real rate rule, we suppose that the nominal interest rate is set according to

$$i_t = \phi_\pi \pi_t + i_t^{\text{shock}}$$

with $\phi_\pi = 1.5$. The results are shown in Figure D.17

Figure D.16: Jacobian of the NKPC across belief models

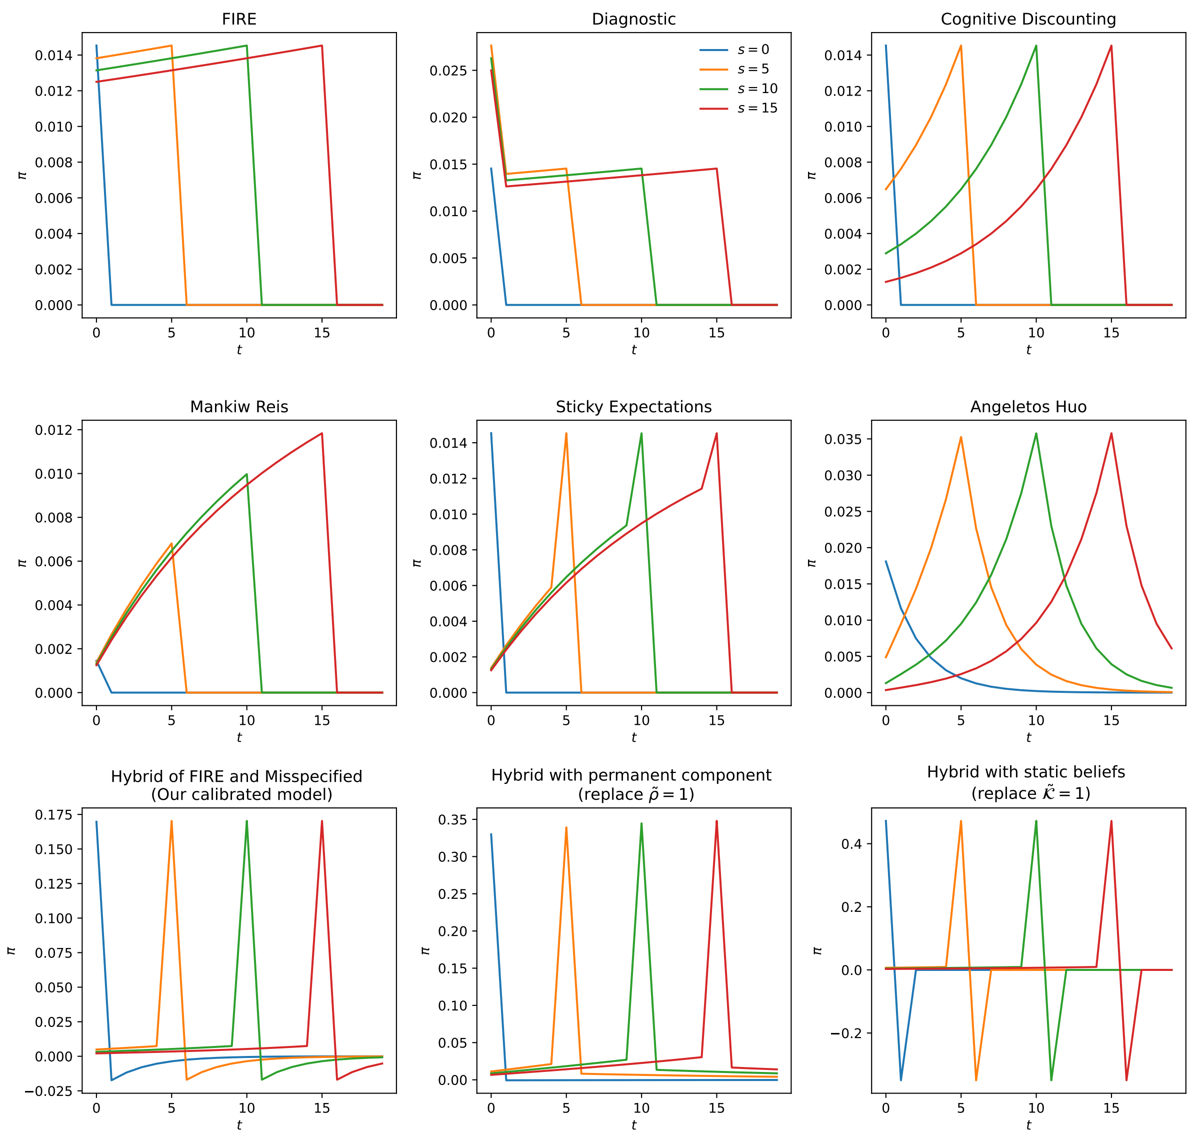

Figure D.17: Inflation response to supply and demand shocks under Taylor Rule

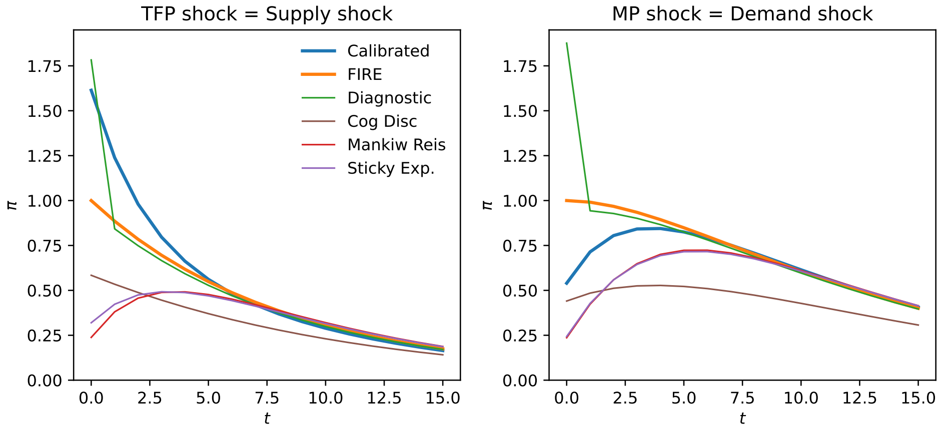

Note: The response of inflation to a monetary policy shock (left) and to a TFP shock (right) under different models of firm beliefs. Impulse responses normalized to 1 under FIRE. Our model, calibrated in Section 2.2 above, is shown in blue.

D.5 Demand and supply shocks under rational expectations

In this section, we show that this prediction can also be consistent with a model of limited-information rational expectations (LIRE) in which firms observe only their own current and past costs but are otherwise rational.

More formally, we consider an economy with only the two shocks above (TFP and monetary policy) and where firms have Adaptive Learning (AL) beliefs except that $\tilde{\rho}$ and $\tilde{\mathcal{K}}$ are chosen optimally. By restricting beliefs in this way, we do not quite recover LIRE, since the true form of costs may not have the AR(1) + noise form posited for AL beliefs. We solve the fixed point as follows: for given $\tilde{\rho}$ and $\tilde{\mathcal{K}}$, we can simulate the equilibrium under both shocks. We normalize the shock variances so they each explain an equal share of the variation in inflation. The simulation generates observational data on cost growth. We suppose that firms use the first three autocovariances of marginal cost growth to infer $\tilde{\rho}$, σ_ε, and σ_η.[94] From these we can also construct the Kalman gain $\tilde{\mathcal{K}}$. We iterate until the believed $\tilde{\rho}$ and $\tilde{\mathcal{K}}$ sustain an economy in which these are the values generated by the simulated data.

The results are shown in Figure D.18. We find the same directional asymmetry in shock response as in our calibrated model, suggesting that the core intuition of our results holds up in such a model. However, it is important to note that the similarity is not compelling evidence that firms are nearly rational, since the simple model presented above (with only its two aggregate shocks) is not sufficiently rich to capture the true range of shocks a fully rational firm would consider.[95]

[94] The firm could use other moments of the data, and this would generate different results to the extent that the believed cost process is incorrect. One alternative would be to model the firm as choosing the parameters that maximise expected profits under the true probability distribution.

[95] One notable example is that we have not included highly transitory (perhaps idiosyncratic) shocks here. If these make up a large share of the variation firms face, then under LIRE, firms may overreact to these and underreact to both our aggregate shocks. Our finding that firms overreact to own costs already demonstrates that LIRE fails, and is quantitatively important to match how firms respond to aggregate shocks.

Figure D.18: Inflation response to supply and demand shocks

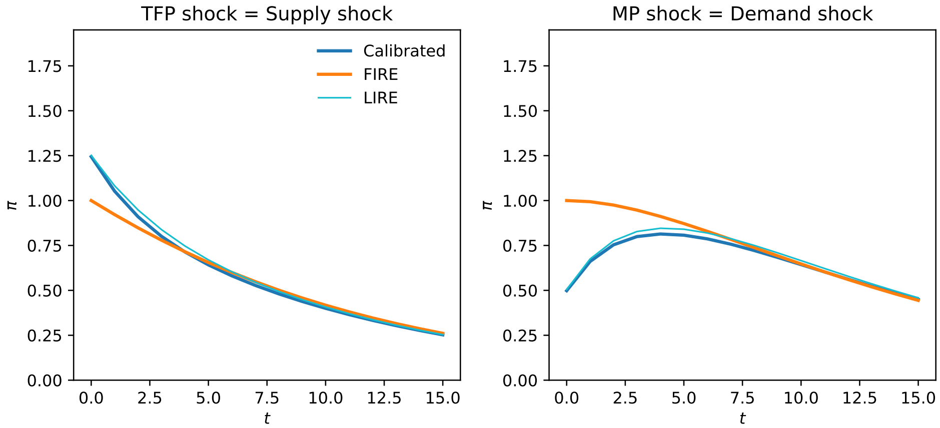

Note: The response of inflation to a monetary policy shock (left) and to a TFP shock (right) under different models of firm beliefs. Impulse responses normalized to 1 under FIRE. Our model, calibrated in Section 2.2 above, is shown in blue.

Figure D.19: Forward guidance is much less powerful at long horizons

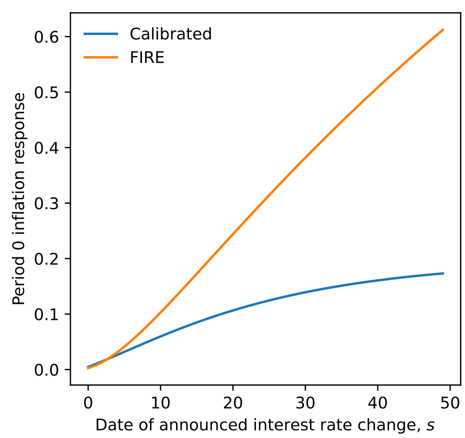

Note: Impact inflation, π_0, in response to date s monetary policy shock announced at date 0.

D.6 Forward guidance

Figure D.19 shows the effect of forward guidance on inflation today under our calibrated beliefs, and in a model with sticky prices and wages. The model is calibrated as in Section 4.1.

D.7 Demand and supply shocks in an estimated model

In this section, we follow Smets and Wouters (2007). To match our pricing model, and since we do not model firms' forecasts of output, we must deviate from this model in two ways. First, we set the degree of indexation to past inflation of prices to zero. Second, we set the curvature of the Kimball goods market aggregator to zero. More importantly, we deviate by estimating the remaining parameters under our fully calibrated belief model. The resulting parameters are shown in Table D.12. For comparison, we also show the parameters values estimated under FIRE (though still without indexed prices or Kimball curvature).[96]

Parameter	FIRE	Calibrated (expanded markup)
Price and wage setting		
Calvo prices, ξ_p	0.77	0.78
Fixed cost / gross markup, Φ	1.66	1.66
Calvo wages, ξ_w	0.76	0.77
Wage indexation, ι_w	0.57	0.51
Preferences and technology		
Risk aversion, σ_c	1.25	1.25
Habit, h	0.81	0.81
Frisch elasticity, σ_l	2.46	2.49
Investment adjustment cost, φ	6.34	6.56
Capacity-utilization cost, ψ	0.38	0.41
Capital share, α	0.20	0.20
Gov.-spending feedback on TFP, ρ_{ga}	0.59	0.59
Monetary policy rule		
Inflation feedback, r_π	1.90	1.86
Interest-rate smoothing, ρ	0.87	0.88
Output feedback, r_y	0.12	0.13
Output-growth feedback, $r_{\Delta y}$	0.12	0.13
Steady state		
SS inflation, $\bar{\pi}$	0.64	0.64
Discount rate, $100(\beta^{-1}-1)$	0.11	0.11
Trend growth, $\bar{\gamma}$	0.51	0.51
SS hours, $\bar{l}$	1.38	1.38
Shock persistence		
TFP shock, ρ_a	0.98	0.98
Monetary-policy shock, ρ_r	0.21	0.17

Table D.12: Estimated parameters under the calibrated cost-beliefs model with the expanded (equation-consistent) price-markup Phillips curve and, for comparison, under FIRE. Both variants set price indexation and Kimball curvature to zero. All seven shock standard deviations are estimated but omitted here; only the TFP and monetary-policy shock persistences are shown.

Under these estimated parameters, we then ask how TFP and monetary policy shocks propagate into inflation. For reference, the estimated persistences of these two shocks are also shown in Table D.12. We show the impulse responses under our calibrated beliefs and under FIRE in Figure D.20.

[96]To be clear, however, only the parameters estimated under our belief model are used to produce the IRFs below.

Figure D.20: Inflation response to estimated TFP and Monetary Policy shocks

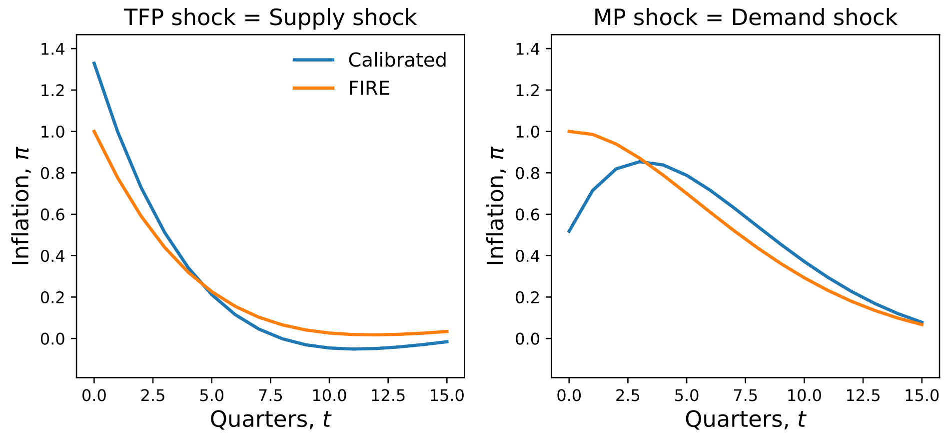

Note: The response of inflation to an estimated TFP shock (left) and to an estimated monetary policy shock (right) under different models of firm beliefs. Impulse responses normalized to 1 under FIRE. Our model, calibrated in Section 2.2 above, is shown in blue.

E Pricing regressions

E.1 Extension of Hazell et al. (2022)

In this section, we estimate an adaptive Phillips curve using state-level inflation data and instruments from HHNS. HHNS estimate the NKPC slope from

$$\pi_{i,t}^{N} = -\kappa_{\text{unemp}}^{\text{FIRE}} \sum_{j=0}^{T} \beta^{j} u_{i,t+j} + \text{FE}_{i,t} + \text{Controls}_{i,t} + \varepsilon_{i,t}, \tag{HHNS eq.17}$$

where $\pi_{i,t}^{N}$ is non-tradable goods inflation in state i at time t and $u_{i,t+j}$ is the unemployment rate in state i at time $t+j$. Our regressions use the same fixed effects and controls as HHNS. HHNS also estimate a myopic variant of their NKPC (in which future unemployment terms are omitted),

$$\pi_{i,t}^{N} = -\kappa_{\text{unemp}}^{\text{myopic}} u_{i,t-4} + \text{FE}_{i,t} + \text{Controls}_{i,t} + \varepsilon_{i,t} \tag{HHNS eq.19}$$

To estimate our adaptive Phillips curve (43), we follow the timing convention of (HHNS eq.19) but allow a role for lagged inflation,

$$\pi_{i,t}^{N} = -\kappa_{\text{unemp}}^{\text{AL}} u_{i,t-4} + \omega_{b}^{\text{AL}} \pi_{i,t-4}^{N} + \text{FE}_{i,t} + \text{Controls}_{i,t} + \varepsilon_{i,t} \tag{91}$$

The results are shown in Table E.13. Under the null of no forward-lookingness, the Phillips curve is an order of magnitude steeper; estimates of $\kappa_{\text{unemp}}^{\text{myopic}}$ and $\kappa_{\text{unemp}}^{\text{AL}}$ are both near 0.08, more than ten times larger than the estimate of $\kappa_{\text{unemp}}^{\text{FIRE}}$. We can translate our estimate for $\kappa_{\text{unemp}}^{\text{AL}}$ into an estimate for κ^{AL}, the slope of the cost-based adaptive Phillips curve (43), using the result from Gagliardone et al. (2023) that a one percentage point decrease in unemployment increases real

Table E.13: Estimating an Adaptive Expectations Phillips Curve

Estimates of $\kappa^{\text{FIRE}}_{\text{unemp}}$ from equation (HHNS eq.17)			
$\kappa^{\text{FIRE}}_{\text{unemp}}$	0.0003	0.0062	0.0062
	(0.0019)	(0.0028)	(0.0025)
Panel B: Estimates of $\kappa^{\text{myopic}}_{\text{unemp}}$ from equation (HHNS eq.19)			
$\kappa^{\text{myopic}}_{\text{unemp}}$	0.004	0.028	0.085
	(0.007)	(0.014)	(0.032)
Panel C: Estimates of adaptive expectations PC params			
$\kappa^{\text{AL}}_{\text{unemp}}$	0.01	0.03	0.08
	(0.01)	(0.01)	(0.03)
ω^{AL}_{b}	0.251	0.097	0.034
	(0.049)	(0.039)	(0.037)
State effects	✓	✓	✓
Time effects		✓	✓
Instruments			Tradeable demand

marginal costs by 0.25 percent.[97] Using this mapping, the cost-based adaptive Phillips curve slope from our calibrated model, $\kappa^{\text{AL}} \approx 0.5$, corresponds to an unemployment-based adaptive Phillips curve slope of about 0.12, in line with the preferred IV estimate of $\kappa^{\text{AL}}_{\text{unemp}}$.

Further, Table E.13 shows evidence of limited inertia. The largest estimate of ω^{AL}_{b} is just 0.25, and under the preferred IV specification, the estimate is small and statistically insignificant. In Table E.14, we also show that inertia remains limited when including additional lags of inflation in the estimation of (91).

In Table E.14 we show the regression result when allowing for a two year lag of inflation. And in Table E.15 we show the result when splitting the sample into pre- and post-1990.

[97] This mapping accounts for wage stickiness in addition to the usual factors, such as the Frisch elasticity of labor supply and elasticity of intertemporal substitution, which cause cost-based and traditional NKPC slopes to differ.

Table E.14: Estimating an Adaptive Expectations Phillips Curve

Estimates of $\kappa^{\text{FIRE}}_{\text{unemp}}$ from equation (HHNS eq.17)				
$\kappa^{\text{FIRE}}_{\text{unemp}}$	-0.0037	0.0003	0.0062	0.0062
	(0.0013)	(0.0019)	(0.0028)	(0.0025)
Panel B: Estimates of adaptive expectations PC params				
$\kappa^{\text{AL}}_{\text{unemp}}$	0.007	0.014	0.022	0.072
	(0.005)	(0.006)	(0.0128)	(0.032)
$\omega^{\text{AL}}_{b,Lag4}$	0.313	0.217	0.0713	0.027
	(0.040)	(0.042)	(0.032)	(0.034)
$\omega^{\text{AL}}_{b,Lag8}$	0.232	0.133	0.083	0.074
	(0.029)	(0.030)	(0.038)	(0.038)
State effects		✓	✓	✓
Time effects			✓	✓
Instruments				Tradeable demand

Table E.15: Estimating an Adaptive Expectations Phillips Curve

Estimates of $\kappa^{\text{FIRE}}_{\text{unemp}}$ from equation (HHNS eq.17)				
$\kappa^{\text{FIRE}}_{\text{unemp}}$	0.011	0.011	0.005	0.005
	(0.008)	(0.006)	(0.004)	(0.003)
Panel B: Estimates of adaptive expectations PC params				
$\kappa^{\text{AL}}_{\text{unemp}}$	0.044	0.092	0.023	0.086
	(0.026)	(0.053)	(0.015)	(0.041)
ω^{AL}_{b}	0.155	0.179	-0.025	-0.036
	(0.061)	(0.078)	(0.031)	(0.033)
Sample	Pre 1990	Pre 1990	Post 1990	Post 1990
State effects	✓	✓	✓	✓
Time effects	✓	✓	✓	✓
Instruments		Tradeable demand		Tradeable demand