How Firms Form Beliefs and the Implications for Inflation

Robert Minton* and Hugo Monnery†

February 24, 2026

Click here for latest version

Abstract

Firms' forecasts of their own future marginal costs are central to the propagation of shocks into inflation. Using survey data on US firms, we establish five novel facts about these forecasts: they (1) are driven by idiosyncratic cost shocks, (2) are overly sensitive to the firm's realized costs (a failure of rational expectations), (3) over-react consistently across time, sectors, and firm size, (4) incorporate cost movements slowly over time, and (5) under-react to aggregate shocks until costs move. We estimate a model of firms' beliefs that is able to match these facts and embed it in a New Keynesian model. The New Keynesian Phillips curve becomes steeper and less forward-looking. We confirm our Phillips curve predictions in external pricing data. Supply shocks are made more inflationary, because they hit costs quickly, leading firms' beliefs to over-react, while demand shocks are made less inflationary, as firms fail to anticipate future wage pressure. We show that optimal monetary policy is less reliant on forward guidance and commitment, since it is difficult to move firms' beliefs without first moving their costs. In contrast, current interest rates remain powerful and may have lasting effects.

*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 Tilman Graff, Helene Hall, Laura Nicolae, and Jieying Zhang. 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, and the Federal Reserve Bank of Boston. 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.

It would be foolish, in forming our expectations, to attach great weight to matters which are very uncertain. 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.

— Keynes (1936)

1 Introduction

Firms' forecasts of their own future 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 increasing study of 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 make use of the Federal Reserve Bank of Atlanta's Business Inflation Expectations (BIE) survey, which asks firms to report their cost growth over the past year and to forecast their cost growth over the next year. We show that firms over-react to their past costs: a failure of rational expectations. Firms also under-react to other information, such as public signals. Figure 1 plots the average of the firms' forecasts for cost growth over the next year alongside a best fit using reported cost growth over the last year. The close fit suggests that extrapolation from costs *might* explain almost the entire forecast. However, it does not prove firms rely on their own costs when forecasting. Nor does it tell us how well they use other information.

To establish the outsized role of realized costs in forecasting future costs, we leverage the panel structure of our survey data. We document five facts about firms' own-cost forecasts.

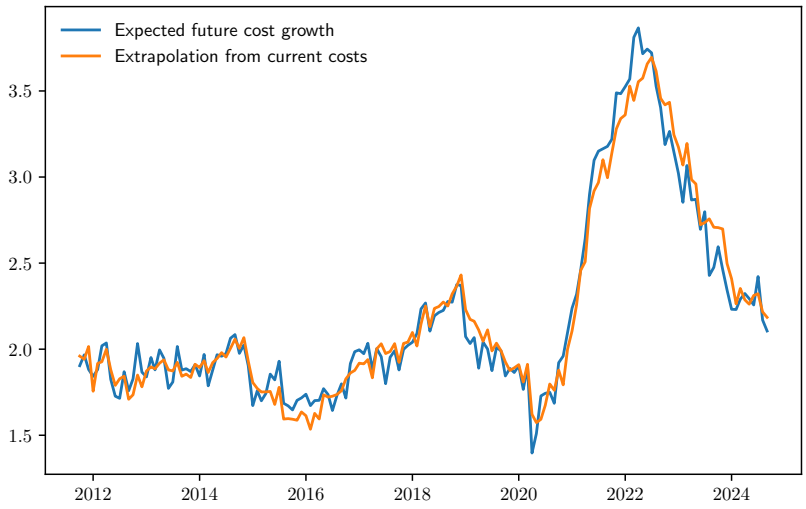
Fact 1 Firms' beliefs about future costs are highly correlated with recent cost growth. This finding holds even when we isolate the role of idiosyncratic cost movements.

Fact 2 Firms' beliefs over-react to their recent cost growth. Because firms' forecast errors are predictable from information provided to us by firms, this finding reflects a failure of rational expectations rather than a lack of information.

Fact 3 This over-reaction to recent cost growth is ubiquitous, holding across sectors, firm sizes, and over time.

Fact 4 There is substantial inertia in belief formation. Forecasts depend not only on cost growth from the current month, but also from a year ago.

Figure 1: Average beliefs are well explained by constant extrapolation from current costs



Note: 'Expected future cost growth' is the cross-sectional equal-weighted average at any date of firms' forecasts for their cost growth over the next year. 'Extrapolation from current costs' is created as average beliefs over time plus a deviation equal to a constant, 0.6, times the deviation of costs from their average over time.

Fact 5 Beliefs respond to aggregate shocks mostly when firms' costs are affected. They under-react while the shock is public information but has not yet passed into costs.

We present a simple model of firms' beliefs, which nests full-information rational expectations (FIRE) but also allows 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, our model can match the five facts established above.

We then calibrate this model, taking advantage of the panel structure of the BIE survey data to use estimated firm-specific loadings on macroeconomic shocks. Equation (19) below characterizes how firm i's forecast of their own cost growth τ periods ahead, $\mathbb{E}_{i,t}\left[\Delta \text{mc}_{i,t+\tau}\right]$, depends on cost growth ℓ periods in the past, $\Delta \text{mc}_{i,t-\ell}$, and the full-information rational expectation of their future costs, $\mathbb{E}_t\left[\Delta \text{mc}_{i,t+\tau}\right]$,

$$\mathbb{E}_{i,t}\left[\Delta \text{mc}_{i,t+\tau}\right] = \underbrace{0.25 \times \mathbb{E}_t\left[\Delta \text{mc}_{i,t+\tau}\right]}_{\text{FIRE component of beliefs}} + \underbrace{0.16 \times 0.82^{\tau} \times \sum_{\ell=0}^{\infty} 0.64^{\ell} \Delta \text{mc}_{i,t-\ell}}_{\text{Adaptive Learning}}$$

The 0.25 weight on FIRE suggests that firms do learn from information outside their costs, but the extent is limited. There is also a significant weight placed on Adaptive Learning. Firms believe 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 the behavior we have documented: why firms are too responsive

to their own costs and insufficiently responsive to other news. First, we show that firms' over-reaction may result from using an average cost persistence over a much longer horizon than our sample, 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 slowly — but at a rate consistent with the learning literature — then we can rationalize the over-reaction observed in our sample.[1]

Second, we explore why firms under-react to *public* information. We ran an information treatment in the May 2025 BIE survey, in which we inform all firms of the average Blue Chip forecast for CPI inflation over the next year. Asking firms to make forecasts before and after the treatment, we observe the effect on their beliefs about (i) CPI inflation over the next year and (ii) their own cost growth over the next year. If the information induces a firm to increase its CPI forecast by 1 percentage point, this leads to only a 0.15 percentage point increase in their own-cost forecast. This disconnect between firms' own-cost and aggregate beliefs reaffirms the importance of understanding own-cost forecasts and further suggests the reason that firms learn so little from public signals may be a failure of rational expectations (rather than purely a lack of information). This finding also suggests it is not possible to treat firms' beliefs about the nominal and real components of their costs as separate; when firms learn that CPI is higher than they thought, they revise down their real cost beliefs.[2]

We next explore the implications of Adaptive Learning for pricing decisions. 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 in 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. We show that in our model, sufficiently persistent cost shocks are under-reacted to, while sufficiently transitory cost shocks are over-reacted 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 marginal costs 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 find that cost growth beliefs contribute positively to inflation, acting like a cost-push term. Despite the backward lookingness of cost growth beliefs, our pa-

[1] This explanation would suggest that over a much longer sample, or over other time periods, we may not find unconditional over-reaction, as in **Fact 2**. However, many of the qualitative implications of Adaptive Learning beliefs remain even when firm's extrapolate correctly on average from their costs; see Section 5.1.1.

[2] This contrasts with a reduced-form reading of the FIRE NKPC. It also differs from the implications of a simple rational inattention model in which firms see the price level and their own real costs as independent signals, see Maćkowiak and Wiederholt (2009).

rameters suggest little inertia in the Phillips curve: there are limited effects on inflation in the future from one-off shocks to real marginal costs today.[3] We show that the degree of inertia in the Phillips curve is increasing in price *flexibility*, via Adaptive Learning behavior. 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 in inflation data. First we estimate an Adaptive Expectations Phillips curve, using the state-level inflation data and estimation methodology from Hazell, Herreño, Nakamura and Steinsson (2022). We find a steep slope and limited inertia; moreover, the limited inertia we find is driven by the pre-1990 period, when the frequency of price adjustment was higher. Second, we consider industry-level pricing and cost data from the Bureau of Economic Analysis and the Bureau of Labor Statistics, showing that inflation is more responsive to cost movements and less responsive to expected cost movements than we would expect under FIRE.

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 over-react unconditionally to costs.[6]

We then analyze optimal monetary policy. Forward guidance is weaker in our model, particularly at longer horizons, as firms respond less today to anticipated future costs. Households' forecasts play a more pivotal role in forward guidance in our model, as they act as a mediator between news about the future and firms' current costs, which are passed through strongly into inflation. 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".

[3] 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.

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

[5] This stands in contrast to many other models of beliefs. For example, Angeletos, Huo and Sastry (2021) present a model of belief formation that is very similar to ours, except that firms learn from observing each shock directly, rather than learning about every shock only via own costs. This assumption implies a different NKPC for each shock, and does not generate the asymmetry result.

[6] In a model of LIRE in which firms observe only their own costs, firms will initially under-react to the most persistent shocks, and over-react to the least persistent. However, in this model, beliefs must be correct on average, in contrast to **Fact 2** above. Therefore, it is still quantitatively important to use actual beliefs to calibrated the model.

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.

Even under discretion, optimal policy differs in our model from FIRE. As our NKPC is steeper, current interest rates are more powerful and can affect inflation as quickly as they can affect costs. Additionally, in the New Keynesian model under FIRE, the central bank does not have an incentive to act in anticipation of future inflation, since a one-off monetary policy shock today does not have lasting effects.[7] Under Adaptive Learning, early action by the central bank can be optimal. The conditions for this are the same as those governing whether the Phillips curve is inertial. As such, while we find little support for the benefits of such early reactions 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 the pandemic.

Literature review 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 over-reaction, suggesting a failure of rational expectations. Broer and Kohlhas (2022) found the same result of over-reaction in forecast revisions; but they additionally show that, while some public information is over-reacted to, other public information is under-reacted 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 under-react and then over-react. Compared to this literature, we focus on firms' beliefs. But we too find evidence for both over-reaction and under-reaction, 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 under-reaction; but if costs peak on impact and then sharply decline, there will be initial over-reaction.

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.[8] Information treatments have also been used to test the causal effect of inflation expectations on firm decisions, with mixed findings.[9] There has also been significant interest in the higher order moments of firm beliefs.[10] [11] In contrast to this literature, we focus on firms' beliefs about their own nominal marginal costs,

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

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

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

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

[11] There are also some papers looking at other firm outcomes, for example Bloom, Codreanu and Fletcher (2025) and

the primitive object for setting prices in the New Keynesian model. We empirically demonstrate that firms' beliefs about their own future costs differ from their beliefs about CPI.

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, while Born, Enders, Menkhoff, Müller and Niemann (2022) argue that expectations about the aggregate economy over-react to micro news and under-react to macro news. Both these papers, while suggesting firms react to local information, do not isolate what that information is.[12] Dovern, Müller and Wohlrabe (2023) show German firms extrapolate from local conditions to expected aggregate growth rates. 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 behavioral 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. This role for information frictions is also present in more recent papers.[13] Gabaix (2020) spells out the implications of one particular form of myopia – cognitive discounting – in the New Keynesian model.[14] 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 over-reaction 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 over-reaction. We compare the predictions of our behavioral theory to those from these alternative models and demonstrate that 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 a la Evans and Honkapohja (2001). For a recent overview, see Eusepi and Preston (2018).[15] 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.

Ma, Ropele, Sraer and Thesmar (2020) find systematic mistakes when firms forecast sales growth.

[12] Sector specific prices could reflect variation in costs, but also many other factors. Beliefs in turn may also affect prices, or other forces may lead to sector-specific responses of both prices and beliefs.

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

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

[15] 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.

Meyer and Sheng (2025) highlight the BIE's focus on firms' own cost forecasts, and show that surveying firms about their own costs provides meaningful data: average own cost growth forecasts match aggregate statistics, and average forecasts of future cost growth perform well compared to other forecasters. In other work in the BIE, Meyer, Prescott and Sheng (2022) argue that, in the early part of the Covid period, firms saw the Covid pandemic mostly as a negative demand shock, lowered their CPI expectations, but had stable long run beliefs.

2 How do firms forecast their costs?

In this section, we introduce our data on firms' own cost forecasts and present reduced-form evidence that (1) firms' beliefs react to their idiosyncratic costs, (2) firms over-react to their own costs, (3) firms' beliefs depend on own costs in a stable way across sectors, firm sizes, and over time, (4) firms learn from their costs slowly over time, and (5) firms under-react to aggregate shocks until their costs move.

2.1 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 all participants' responses 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.[16]

The key variables of interest in our study are firms' realized and expected 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. 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

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

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

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

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

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

As an external check on firms' responses, and as previously argued in Meyer and Sheng (2025), we show in Appendix A.2 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.[17] 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.3, 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 $t+12$

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

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.

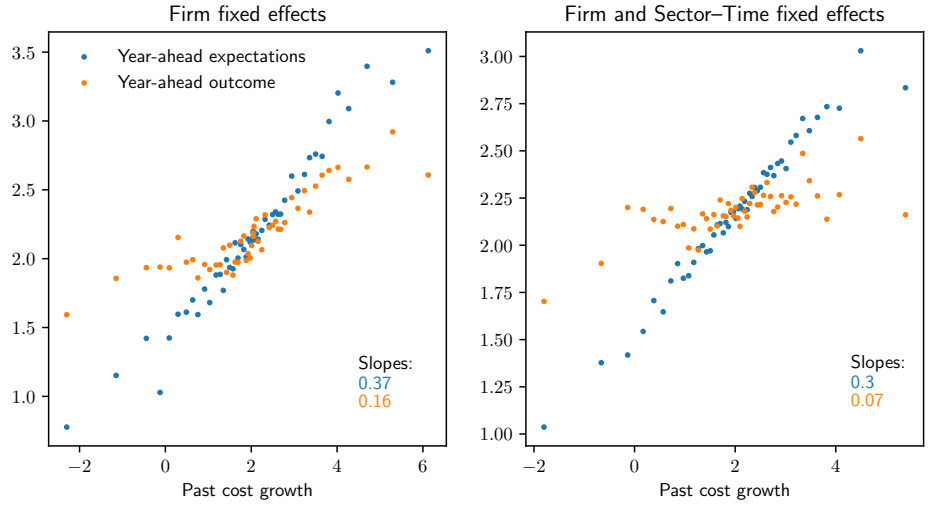
Facts 1&2: Firms (over) react to their own costs

We start by relating firms' cost growth over the last year to their expectations for their cost growth over the next year. We can visualize this relationship in a binned scatter plot, shown by the blue dots in Figure 2. We see a strong, positive, and linear relationship between costs and forecasts at the firm level (controlling for firm fixed effects).[18] This relationship persists when we also control

[17] In Appendix A.4 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.

[18] 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.5 we show that firms appear to hold fairly accurate beliefs on average over time. Cattaneo, Crump, Farrell and Feng (2024) suggest jointly computing binned

Figure 2: Cost Growth is Associated with Expected and Realized Year-ahead Cost Growth



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

for sector-time fixed effects. This suggests that firms believe that their nominal cost growth is persistent. In the orange dots, we test 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 costs are more persistent than they truly are.

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 there may be substantial measurement error given the binning of the survey responses.[19]

We next turn to linear regressions, to explore these relationships and their robustness. Our specifications are

$$\mathbb{E}_{it}[\Delta_{12}\text{mc}_{i,t+12}] = \alpha_i + \text{Controls}_{it} + \beta_{\text{expectations}}\Delta_{12}\text{mc}_{it} + u_{i,t}^{\text{expectations}} \tag{4}$$

$$\Delta_{12}\text{mc}_{i,t+12} = \alpha_i + \text{Controls}_{it} + \beta_{\text{outcome}}\Delta_{12}\text{mc}_{it} + u_{i,t}^{\text{outcome}} \tag{5}$$

$$\text{FE}_{i,t,t+12} = \alpha_i + \text{Controls}_{it} + \beta_{\text{FE}}\Delta_{12}\text{mc}_{it} + u_{i,t}^{\text{FE}} \tag{6}$$

where α_i is a firm fixed effect and the u_{it} are error terms. Table 1 shows the OLS estimates of

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.6 we show that our conclusions hold when following their procedure too.

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

$\beta_{\text{expectations}}$, 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.7). Standard errors are clustered at the firm level.

Table 1: Firms (over) react to their own costs

	Outcomes		
	$\mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right]$	$\Delta_{12}\text{mc}_{i,t+12}$	$\text{FE}_{i,t,t+12}$
Controls			
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: *p<0.10, **p<0.05, ***p<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 1 shows the regression coefficient for expectations of year-ahead cost growth. No matter the controls included, firms strongly believe that cost growth over the last year predicts cost growth over the next year. This includes accounting for aggregate components of costs; for sector loadings on aggregates or sectoral shocks; and even for the first four principal components of cost growth, which together explain 71% of cost growth. We conclude that firms' cost forecasts depend on idiosyncratic cost movements. Quantitatively, these results say that a firm whose costs typically rise by 2% per year, but whose costs were 3% over the last year, would thereby increase their forecast from 2% to 2.3%.

The second column shows the regression coefficient for realized year-ahead cost growth. We find cost growth to be persistent when there are no controls, and to remain so after controlling for time fixed effects, albeit less so. As we isolate more idiosyncratic variation, firms' cost growth becomes predictive of a decline in year-ahead cost growth, suggesting slight mean reversion.

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 seem to over-react 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.

CPI beliefs and long-term expectations. In Appendix A.8, we show that, as in Albagli et al. (2022), firms also use their past costs to forecast CPI growth over the next year. However, in Appendix A.9, we show that firms still overreact to their own cost growth conditional on their CPI beliefs, suggesting firms form own cost beliefs differently from the CPI beliefs more commonly studied in the literature. Our survey experiment in Section 3.4 further demonstrates the disconnect between CPI and own cost beliefs.

In addition, we show that firms use their past costs to forecast their annualized cost growth five to ten years ahead. As Table 1 shows that cost growth is not very persistent even at a one-year horizon, its perceived persistence at a five to ten year horizon is likely further evidence of overreaction.

Robustness. In Appendix A.10, 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.[20] In Appendix A.10, 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%.[21] 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.9, we include two measures of sales and a measure of profits from the BIE survey and continue to find our result. Appendix A.11 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

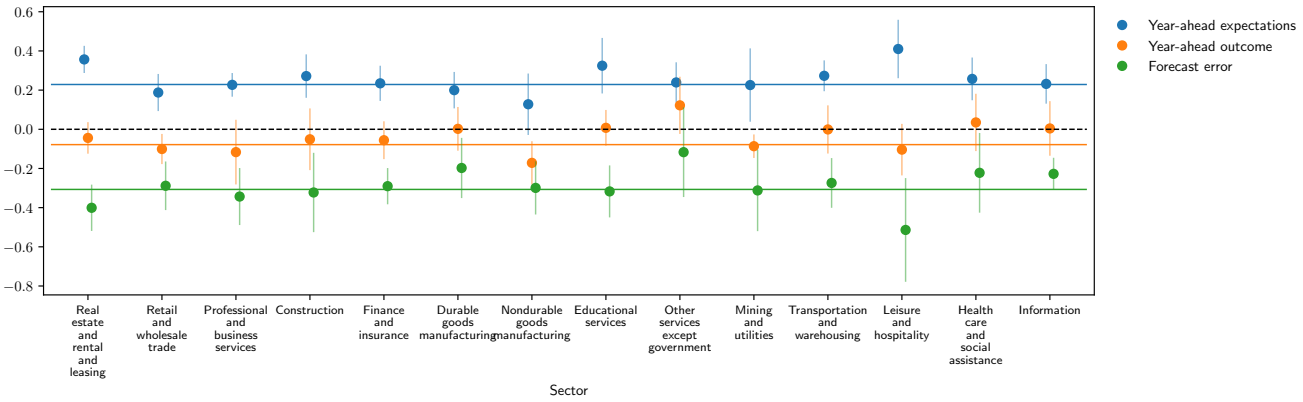
Does this dependence of firms' forecasts on their own costs vary with observables? An answer helps us understand why firms make these mistakes, but also guides our modeling. For example, if firms only over-reacted to their costs in the post-Covid period, it would suggest the over-reaction was state contingent. Instead, we find remarkable stability in how firms use their costs to inform their beliefs.

In Figure 3, 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

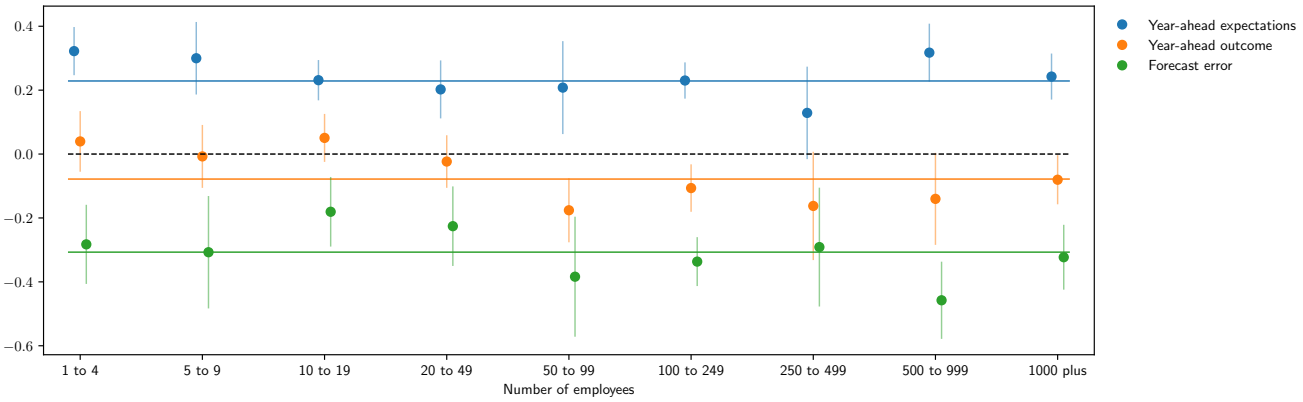
[20] Note also that measurement error poses an even larger issue for regressions of forecast errors on forecast revisions—a common regression run in this literature—than it does for us here, since those regressions contain the current forecast on the left-hand side (negatively) and on the right-hand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

[21] 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. Firms are asked to assign probabilities and report realizations of the same variable over the same intervals.

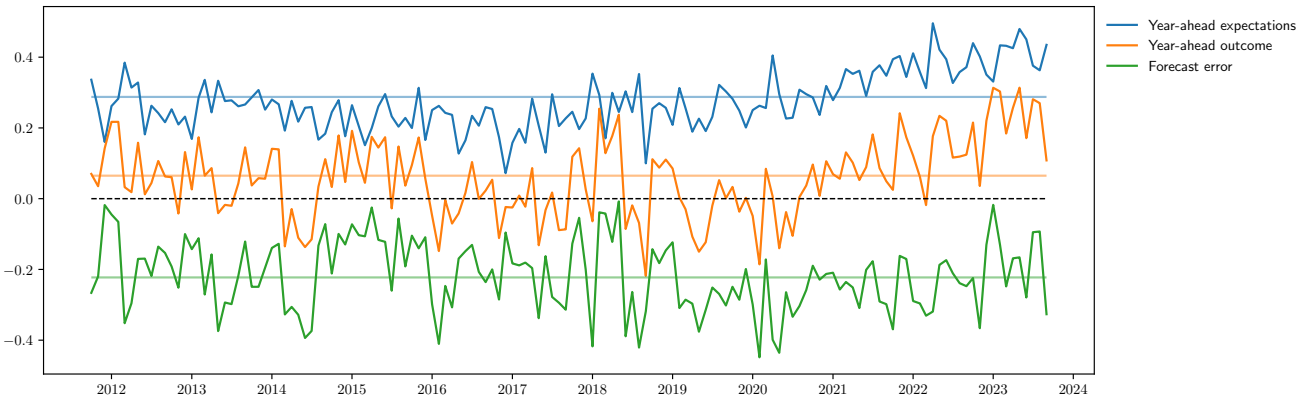
Figure 3: Heterogeneity analysis



(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).

find that beliefs are positively related to costs, and that these costs are over-reacted 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 relatively stable.

Further results on heterogeneity. This evidence suggests that over-reaction to costs is a widespread feature in the data. 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.12, we show that there is little variation in over-reaction to own costs along this margin.

While **Fact 2** showed overreaction to costs on average, Appendix A.13 argues that firms with more persistent cost growth overreact less. This heterogeneity suggests too much stability in firms' extrapolation behavior. 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.[22]

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[23]

$$\mathbb{E}_{it}[\Delta_{12}\text{mc}_{i,t+12}] = \alpha_i + \text{Controls}_{it} + \sum_{\ell=0}^{L} \beta_\ell^{\text{expectations}} \Delta\text{mc}_{i,t-\ell} + u_{i,t} \tag{7}$$

We cannot run the regression equation (7) since we do not actually observe month-on-month cost changes. However, we can transform (7) into a regression on observables

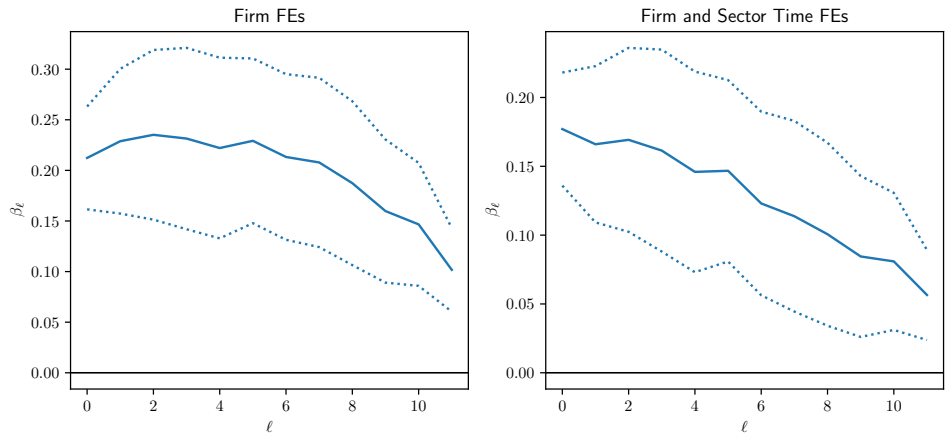
$$\mathbb{E}_{it}[\Delta_{12}\text{mc}_{i,t+12}] - \mathbb{E}_{it-12}[\Delta_{12}\text{mc}_{i,t}] = \Delta_{12}\text{Controls}_{it} + \sum_{\ell=0}^{L} \beta_\ell^{\text{expectations}} \Delta\left(\Delta_{12}\text{mc}_{i,t-\ell}\right) + \Delta_{12} u_{i,t}$$

In Figure 4, 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. So

[22] 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 5.1.1.

[23] 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 4: Firm costs take time to drive beliefs



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

firms do not simply extrapolate from their most recent monthly cost increase, but instead react more to cost increases that have been persistent. This suggests firms may be learning about an underlying persistent component of cost growth.

This result is highly suggestive that firms are over-reacting not just to costs today, but also to past costs. In Appendix A.14, we formally show that firms over-react to their past costs.

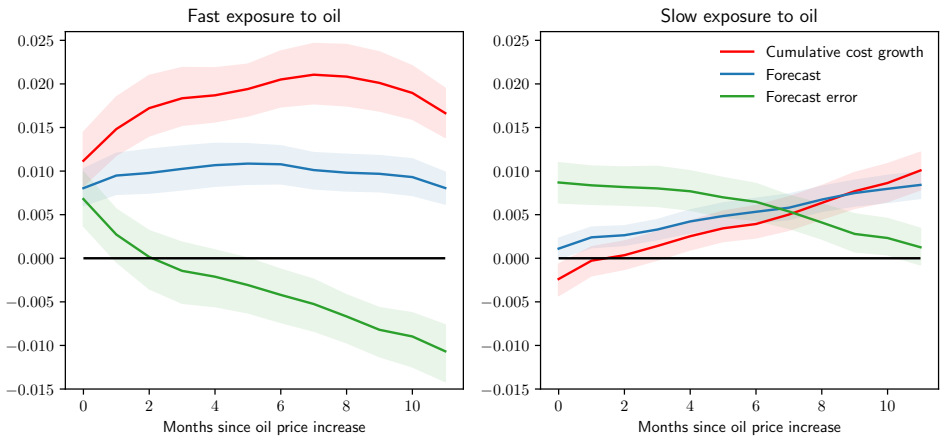
Fact 5: Beliefs under-react until costs move

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 over-react to some public signals and under-react to others. Furthermore, firms' beliefs may over-react and under-react 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.[24] This suggests a particular time profile to learning about aggregate shocks. There is initial under-reaction while news about the shock is contained only in public signals, followed by a move toward over-reaction 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 over-reaction" is possible, so is over-reaction with no delay, or beliefs that under-react at all dates.

We look at how firms react to a change in the Brent oil price. As shown in Minton and Wheaton (2023), persistent oil price movements filter through the supply chain gradually and have substan-

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

Figure 5: Under-reaction and Over-reaction 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.

tial 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 5. While these lines differ by 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 over-reaction. However, for firms only slowly exposed to oil, they under-react 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.

3 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 2. We first outline why existing models in the literature cannot jointly match the facts we established above. Table 2 summarizes these findings.

In **Fact 1** we demonstrated a prominent role for learning from own-costs. This fact is not necessarily inconsistent with FIRE. However, **Fact 2** establishes over-reaction to own-costs on average, a finding of forecast error predictability that does violate FIRE. Since (reported) own-costs are within the firm's information set, this finding is also inconsistent with models of limited-information ra-

Table 2: Which models can match these facts?

Model	**Fact 1** React	**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)	✓	✓	✓	✓	✗

tional expectations (LIRE) – e.g., Mankiw and Reis (2002) and Maćkowiak and Wiederholt (2009). Likewise, models with cognitive discounting in beliefs, such as the New Keynesian model of Gabaix (2020), deviate from FIRE by attenuating forecasts back towards steady state, which would lead to less, not more, extrapolation from own costs.

By contrast, "diagnostic expectations", as in Bordalo et al. (2018), can generate the over-reaction we see. But the only mistake this model predicts is over-reaction to surprises (relative to a FIRE forecast made in the previous period), and so it cannot explain that firms are insufficiently sensitive to their true cost persistence, as in **Fact 3**.[25] Furthermore, Bordalo et al. (2018) predicts only over-reaction to surprises relative to FIRE last period.[26] 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 over-reaction while allowing for inertia in beliefs. In response to an aggregate shock, their model generates initial under-reaction as firms only slowly process the news, followed by an over-reaction, 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 argue 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 4.1.1 below we show quantitatively how our model's predictions differ from those of other theories.

[25] See in particular Section A.13 for further details.

[26] 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 over-reactive extrapolation) is combined with a slower moving belief component ("sentiment").

3.1 Adaptive Learning

We propose a parsimonious model of firms' beliefs that matches the key features we found above. Firms use their own costs as a signal to learn about a persistent underlying process. We also nest full-information rational expectations (FIRE) as a special case of our model, allowing firms to forecast using other information.

In particular, when a firm makes a forecast, we place a weight α^{FIRE} on the FIRE forecast.[27] And we place a weight $1-\alpha^{\text{FIRE}}$ on a behavioural 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) above, our model features firms extrapolating from costs in a constant way; we call these "Adaptive Learning" (AL) beliefs.[28] These beliefs are formed as follows. Suppose that a firm believes that its nominal cost growth has a persistent and transitory component following the process

$$\Delta \text{mc}_{i,t} = \Delta \tilde{\text{mc}}^{p}_{i,t} + \tilde{\eta}_{i,t}, \qquad \tilde{\eta}_{i,t} \overset{\text{iid}}{\sim} \mathcal{N}\left(0, \tilde{\sigma}^2_{\eta}\right) \tag{8}$$

$$\Delta \tilde{\text{mc}}^{p}_{i,t} = \tilde{\rho} \Delta \tilde{\text{mc}}^{p}_{i,t-1} + \tilde{\varepsilon}_{i,t}, \qquad \tilde{\varepsilon}_{i,t} \overset{\text{iid}}{\sim} \mathcal{N}\left(0, \tilde{\sigma}^2_{\varepsilon}\right) \tag{9}$$

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.[29] Given the perceived law of motion implied by (8) and (9), the firm forecasts their costs as follows

$$\tilde{\mathbb{E}}_{i,t}\left[\Delta \text{mc}_{i,t+\tau}\right] = \tilde{\rho}^{\tau} \tilde{\mathbb{E}}_{i,t}\left[\Delta \tilde{\text{mc}}^{p}_{i,t}\right] \tag{10}$$

$$\tilde{\mathbb{E}}_{i,t}\left[\Delta \tilde{\text{mc}}^{p}_{i,t}\right] = \tilde{\mathcal{K}} \Delta \text{mc}_{i,t} + \left(1-\tilde{\mathcal{K}}\right) \tilde{\rho} \tilde{\mathbb{E}}_{i,t-1}\left[\Delta \tilde{\text{mc}}^{p}_{i,t-1}\right] \tag{11}$$

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,

$$\mathbb{E}_{i,t}\left[\Delta \text{mc}_{i,t+\tau}\right] = \alpha^{\text{FIRE}} \mathbb{E}_{t}\left[\Delta \text{mc}_{i,t+\tau}\right] + \left(1-\alpha^{\text{FIRE}}\right) \tilde{\mathbb{E}}_{i,t}\left[\Delta \text{mc}_{i,t+\tau}\right] \tag{12}$$

To calibrate firms' beliefs we therefore need to estimate three parameters: α^{FIRE}, $\tilde{\rho}$, and $\tilde{\mathcal{K}}$. α^{FIRE}

[27] 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.

[28] In Appendix B.2, we allow firms to choose a firm-specific $\tilde{\rho}_i$, allowing for heterogeneity as well as endogenising the choice of $\tilde{\rho}_i$. Our findings there are approximately captured by the model presented in the main text.

[29] 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.

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 fact, this simple model is very similar to that in Angeletos, Huo and Sastry (2021) (AHS) with two key differences. 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. 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.[30] The existence of a separate Phillips curve per shock is exactly why AHS *can* find delayed overshooting in response to all shocks.

3.2 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 calibration that aligns best with the format of our data is therefore at an annual frequency. Casting the model from Section 3.1 in annual time and combining equations (10), (11), and (12), beliefs depend on FIRE beliefs and cost growth according to

$$\begin{aligned}\mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right] &= \alpha^{\text{FIRE}}\mathbb{E}_t\left[\Delta_{12}\text{mc}_{i,t+12}\right] \\ &\quad + \left(1-\alpha^{\text{FIRE}}\right)\tilde{\rho}^{\text{ann}}\tilde{\mathcal{K}}^{\text{ann}}\sum_{\ell=0}^{\infty}\left[\left(1-\tilde{\mathcal{K}}^{\text{ann}}\right)\tilde{\rho}^{\text{ann}}\right]^{\ell}\Delta_{12}\text{mc}_{i,t-12\ell}. \end{aligned} \tag{13}$$

Note that 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 $t+12$:

$$\Delta_{12}\text{mc}_{i,t+12} = \mathbb{E}_t[\Delta_{12}\text{mc}_{i,t+12}] + \epsilon^{\text{FIRE}}_{i,t+1,t+12}. \tag{14}$$

Plugging the rational expectations relationship (14) into (13) and absorbing most lags of cost growth into the error term, we have the regression equation

$$\begin{aligned}\mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right] &= \alpha^{\text{FIRE}}\Delta_{12}\text{mc}_{i,t+12} + \left(1-\alpha^{\text{FIRE}}\right)\tilde{\rho}^{\text{ann}}\tilde{\mathcal{K}}^{\text{ann}}\Delta_{12}\text{mc}_{i,t} \\ &\quad + \left(1-\alpha^{\text{FIRE}}\right)\tilde{\rho}^{\text{ann}}\tilde{\mathcal{K}}^{\text{ann}}\left[\left(1-\tilde{\mathcal{K}}^{\text{ann}}\right)\tilde{\rho}^{\text{ann}}\right]\Delta_{12}\text{mc}_{i,t-12} + \epsilon_{it}, \end{aligned} \tag{15}$$

where

$$\epsilon_{it} = -\alpha^{\text{FIRE}}\epsilon^{\text{FIRE}}_{i,t+1,t+12} + \left(1-\alpha^{\text{FIRE}}\right)\tilde{\rho}^{\text{ann}}\tilde{\mathcal{K}}^{\text{ann}}\sum_{\ell=2}^{\infty}\left[\left(1-\tilde{\mathcal{K}}^{\text{ann}}\right)\tilde{\rho}^{\text{ann}}\right]^{\ell}\Delta_{12}\text{mc}_{i,t-12\ell}. \tag{16}$$

[30] Additionally, if the shock is not an AR(1), the model becomes more complicated.

Endogeneity. 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 $t+12$. So we need an instrument for $\Delta_{12}\text{mc}_{i,t+12}$. 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, $\epsilon^{\text{FIRE}}_{i,t+1,t+12}$.

Further, marginal cost growth may be autocorrelated, meaning regressors may also be correlated with the lagged cost components of ϵ_{it}. However, Table 1 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.

Instruments and Relevance. Following the above requirements for the exclusion restriction, we need a series from periods $t-\ell$, with $\ell \geq 0$, that meaningfully affects $\Delta_{12}\text{mc}_{i,t+12}$. 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. Our baseline calibration instruments future cost growth with average current cost growth in the BIE. In addition, 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 and (2) Nakamura and Steinsson (2018) monetary policy shocks.

Using cross-sectional instruments helps us to deal with weak instruments issues common to this setting.[31] 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.[32] We also show our results are robust to using sectoral, instead of firm-level, loadings.

Regression equations. As in our reduced-form 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 will include time fixed effects to isolate how our time series instruments with firm-level loadings induce differential responses of costs across firms. As a result, we do not need good predictions of future aggregates for instrument relevance.

The final form of our calibration second stage is therefore

$$\mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right] = \mu_i + \lambda_t + \beta_1\Delta_{12}\text{mc}_{i,t} + \beta_2\Delta_{12}\text{mc}_{i,t-12} + \beta_3\Delta_{12}\text{mc}_{i,t+12} + \epsilon_{it}, \tag{17}$$

[31] 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.

[32] 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)).

with $\beta_1 = \left(1-\alpha^{\text{FIRE}}\right)\tilde{\rho}^{\text{ann}}\tilde{\mathcal{K}}^{\text{ann}}$, $\beta_2 = \left(1-\alpha^{\text{FIRE}}\right)\tilde{\rho}^{\text{ann}}\tilde{\mathcal{K}}^{\text{ann}}\left(1-\tilde{\mathcal{K}}^{\text{ann}}\right)\tilde{\rho}^{\text{ann}}$, and $\beta_3 = \alpha^{FIRE}$. The parameters are exactly identified in a regression of beliefs on these three marginal cost growth objects, though they need to be computed from the regression coefficients. Unlike the other two parameters, the parameter α^{FIRE} can be read directly from the β_3.

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 $\Delta_{12}\text{mc}_{i,t-12}$.

Parameter estimates are constructed from the regression coefficients as determined above, and standard errors are constructed via the delta method. We use Driscoll and Kraay (1998) standard errors with bandwidth 12.[33] The results are shown in Table 3.

Results Discussion. 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. 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 $\tilde{\mathcal{K}}^{ann} < 1$.

While we explore the implications of these results more formally when we turn to pricing and optimal monetary policy, it is clear that these beliefs entail less rational anticipation than FIRE and more embedding of historical cost changes into beliefs. For example, even when the rational expectation predicts cost growth to return to normal levels in the next year, recent and historical cost growth above the norm will lead firms to expect elevated cost growth in the future.

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.[34] Instead, we emphasize 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.

[33] Recall that we have monthly data on twelve-month changes, which typically induces a mechanical MA(12) error process.

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

Table 3: Annual Calibration

Panel A: 2SLS estimates

			$\mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right]$	
Regression Estimates	$\Delta_{12}\text{mc}_{i,t}$	0.236***	0.243***	0.236***
		(0.022)	(0.021)	(0.020)
	$\Delta_{12}\text{mc}_{i,t-12}$	0.040***	0.034**	0.039***
		(0.015)	(0.015)	(0.015)
	$\Delta_{12}\text{mc}_{i,t+12}$	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

			$\mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right]$	
Regression Estimates	$\Delta_{12}\text{mc}_{i,t}$	0.296***	0.292***	0.278***
		(0.030)	(0.031)	(0.025)
	$\Delta_{12}\text{mc}_{i,t-12}$	0.022	0.026	0.042*
		(0.017)	(0.018)	(0.022)
	$\Delta_{12}\text{mc}_{i,t+12}$	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.

Time fixed effects. Because our calibration regression includes time fixed effects, we identify α^{FIRE} using cross-sectional variation in costs and cost beliefs.[35] 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.[36]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.

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 — over-reacting to micro news and under-reacting to macro news (Kumar et al. (2015), Andrade et al. (2022), Born et al. (2022)). Second, our experiment in Section 3.4 below shows that the average perceived relevance of CPI to a firms' costs (in absolute, not relative, terms) is low; and the α^{FIRE} estimate from this experiment aligns with our calibration.

Robustness. Recall our motivating Figure 1 in the introduction, where we showed the potentially strong explanatory power of current costs for beliefs in the time series. 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.

In Appendix Table B.5, 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.6, 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.

Quarterly Calibration. 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

[35]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.

[36]We showed this for CPI in Appendix A.8, and this echoes similar results in Andrade et al. (2022) and Albagli et al. (2022).

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.

Result 1. *Firm i's forecast of its own future nominal marginal cost growth is given by the equation*

$$\mathbb{E}_{i,t}\left[\Delta mc_{i,t+\tau}\right] = 0.25 \times \mathbb{E}_t\left[\Delta mc_{i,t+\tau}\right] + 0.16 \times 0.82^{\tau} \times \sum_{\ell=0}^{\infty} 0.64^{\ell} \Delta mc_{i,t-\ell} \tag{19}$$

3.3 Why do firms over-react 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 3.1 but with $\tilde{\sigma}_\eta = 0$ and cast in annual time. So the perceived law of motion for costs is

$$\Delta_{12}\text{mc}_{i,t} = a_i + \tilde{\rho}^{\text{ann}}_{i,t} \Delta_{12}\text{mc}_{i,t-12} + \tilde{\varepsilon}_{i,t} \tag{20}$$

We suppose that each firm updates the perceived cost persistence, $\tilde{\rho}^{\text{ann}}_{i,t}$, as new data comes in, under constant gain learning at rate γ.[37]

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.[38] Given this simulated data and a gain parameter, γ, we can compute what the learning model predicts for $\tilde{\rho}^{\text{ann}}_{i,t}$ at each point in time in (20).

Over our sample period in the BIE, as shown in Table 1, we estimated a constant $\tilde{\rho}^{\text{ann}}$ 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

[37] 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 $y_t = \Delta_{12}\text{mc}_{i,t}$.

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

growth, proxied by the GDP deflator. This very slow learning is not out of line with estimates in the learning literature.[39]

3.4 Why do firms under-react to public signals?

In this section, we argue that a major reason that firms fail to incorporate public signals into their beliefs is a failure of rational expectations (RE), not purely a failure of full-information (FI). It is not simply that firms do not know about these signals, but that they do not know what to do with them. Given **Fact 2**, this suggests failures of RE both in how firms learn from their own costs and also from more aggregate information.

To establish this, we run a Special Question in the Atlanta Fed's Business Inflation Expectations (BIE) survey. This approach also provides us with a plausible upper bound on α^{FIRE} that independently corroborates our estimate above. Finally, this exercise re-asserts the importance of studying firms' own-cost beliefs, as we find a substantial disconnect between these beliefs and CPI beliefs.

The survey question. In the May 2025 BIE survey, we first asked firms to report

- their expectations for their own cost growth over the next year, $\mathbb{E}_{i,t}^{\text{pre}}\left[\Delta \text{mc}_{i,t+12}\right]$
- their expectation for CPI inflation over the next year, $\mathbb{E}_{i,t}^{\text{pre}}\left[\pi_{t+12}^{\text{CPI}}\right]$

We then tell *all* firms the latest Blue Chip forecast (from April 2025) for inflation over the next year. We then ask firms, in light of this, to report

- their updated forecast for their own costs, $\mathbb{E}_{i,t}^{\text{post}}\left[\Delta \text{mc}_{i,t+12}\right]$
- their updated forecast for CPI, $\mathbb{E}_{i,t}^{\text{post}}\left[\pi_{t+12}^{\text{CPI}}\right]$

The exact wording of the questions is shown in Appendix B.9. The own-cost responses are given by assigning weight to bins (as is standard in the BIE survey, and as discussed in Section 2.1), whereas we elicit point estimates for the CPI forecasts. To prevent a mechanical bias against RE, we ex-post truncate the CPI estimate to the range $[-2\%, +6\%]$; our results will also demonstrate robustness to this truncation assumption.

Own-cost beliefs are disconnected from CPI beliefs. Denote the change in CPI beliefs, from before to after the information treatment, by

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

[39] 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.

Table 4: 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
$d\mathbb{E}_{i,t}\pi^{\text{CPI}}_{t+12}$	0.15^{*}	0.09^{*}	
	(0.08)	(0.05)	
$\mathbb{1}\{\text{Changed CPI forecast}\}$			0.27^{***}
			(0.05)
Observations	349	349	349

Note: *p<0.10, **p<0.05, ***p<0.01. 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 4. We find a slope coefficient of 0.15: a 1 percentage point increase in the forecast for CPI over the next year induces only a 0.15 percentage point increase in the forecast for own-cost growth over the next year. 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.

We noted above that the literature has focused on firms' CPI beliefs, and we showed that own-cost beliefs have dynamics that are not captured by CPI beliefs. Here we reinforce that point: if some news (e.g. an information treatment) is found to affect CPI beliefs substantially, this may still have almost no effect on a firm's forecast of its own costs. In a standard pricing model, this would imply very limited effects on pricing, and so inflation.[40]

Failure of rational expectations. We now transform the estimate above into a quantitative measure of the deviation from rational expectations. Suppose that year-on-year cost growth is given by

$$\Delta_{12}\text{mc}_{i,t} = \rho_{\text{idio}}\Delta_{12}\text{mc}_{i,t-12} + \psi_{\text{CPI}}\Delta_{12}p^{\text{CPI}}_{t} + \varepsilon_{i,t} \tag{21}$$

[40] In a model with strategic complementarity in price setting, beliefs beyond a firm's own costs are relevant; but it is still important to understand how much own-cost beliefs are moved as distinct from beliefs about competitors' prices.

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

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

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 (21). As this is essentially a time series regression, and the BIE sample only begins in October 2011, we instead take the integral over (21) 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 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 $\psi_{\text{CPI}} = 0.67$, suggesting $\alpha^{\text{RE}} \approx 0.22$.

We can also use this framework to decompose our α^{FIRE} estimate from Section 3.2. Suppose there is some news about CPI inflation over the next year (with no effect on current costs), denoted s^{CPI}_t. 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,

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

where $\alpha^{\text{FI}} \equiv \frac{d\mathbb{E}_{i,t}\left[\Delta_{12}p^{\text{CPI}}_{t+12}\right]}{ds^{\text{CPI}}_t} \Big/ \frac{d\mathbb{E}^{\text{FIRE}}_t\left[\Delta_{12}p^{\text{CPI}}_{t+12}\right]}{ds^{\text{CPI}}_t}$ captures how much of the news is incorporated into CPI beliefs, compared to under FIRE.[41] 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.

4 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 passthrough 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, first following the methodology in Hazell et al. (2022)

[41] 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}).

and second using sectoral price and cost data.

4.1 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*

$$p_t = \theta_p p_{t-1} + \left(1-\theta_p\right)\alpha^{FIRE} p_t^{*,FIRE} + \left(1-\theta_p\right)\left(1-\alpha^{FIRE}\right) p_t^{*,AL} \tag{22}$$

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

$$p_t^{*,FIRE} = \left(1-\beta\theta_p\right)\sum_{\tau=0}^{\infty}\left(\beta\theta_p\right)^{\tau}\mathbb{E}_t\left[mc_{t+\tau}\right], \tag{23}$$

and the Adaptive Learning optimal reset price is

$$p_t^{*,AL} = mc_t + \frac{\beta\theta\tilde{\rho}\tilde{\mathcal{K}}}{1-\beta\theta\tilde{\rho}}\sum_{\ell=0}^{\infty}\left[\left(1-\tilde{\mathcal{K}}\right)\tilde{\rho}\right]^{\ell}\Delta mc_{t-\ell}. \tag{24}$$

Proof. Under Calvo updating, the log-linearized aggregate price evolves as $p_t = \left(1-\theta_p\right)\int p_{i,t}^{*}\mathrm{d}i + \theta_p p_{t-1}$. Firm i's optimal reset price is given by $p_{i,t}^{*} = \left(1-\beta\theta\right)\sum_{\tau=0}^{\infty}\left(\beta\theta\right)^{\tau}\mathbb{E}_{i,t}\left[\mathrm{mc}_{i,t+\tau}\right]$. Substituting in the marginal cost beliefs in (10), (11), and (12), 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 over-react (relative to FIRE) to transitory shocks and under-react 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,*

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

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

$$\tilde{\rho} > \frac{\rho}{\left(1-\beta\theta_p\rho\right)\tilde{\mathcal{K}} + \rho\beta\theta_p} \tag{25}$$

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

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

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

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

Proof. The proof to Theorem 2 is in Appendix C.2 below. □

Theorem 2 shows that, under our model of beliefs, macroeconomic shocks may be either under- or over-reacted 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 under-react. 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.[42]

4.1.1 Comparing to other models of beliefs

We now introduce four alternative theories of expectation formation and compare the models' predictions for pricing.

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, Crawley, Slacalek, Tokuoka and White (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

[42] Setting $\beta = 0.99$ and $\theta_p = 0.9$, and using the calibration for $\tilde{\rho}$ and $\tilde{\mathcal{K}}$ from Section 3.2, Theorem 2 implies that the on impact response is equal under FIRE and our calibrated model when $\rho = 0.42$.

perfectly, even as their forecasts for future costs might be based on out-of-date information.

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

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,[43]

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

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

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

The central idea is that firms over-react to surprises about their marginal costs.

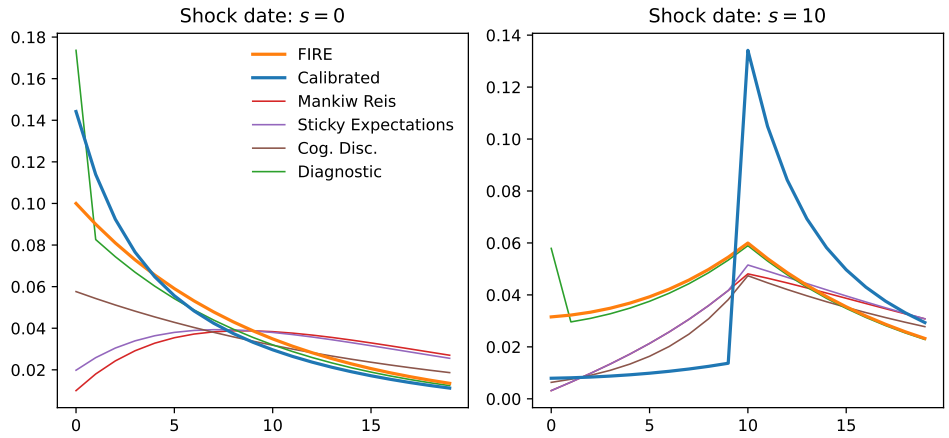
Comparing our model to other behavioral models. To calibrate these behavioral models at quarterly frequency, we set $\theta^{\text{MR}} = \theta^{\text{SE}} = m^{\text{CD}} = 0.9$ and $\theta^{\text{DE}} = 1$. For our model we use the quarterly calibration of α^{FIRE}, $\tilde{\rho}$ and $\tilde{\mathcal{K}}$ from Section 3.2. In Figure 6, 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. In Appendix D.1, we show the long run of the price level under all the beliefs models is indeed the same. Relative to FIRE, our model generates faster passthrough when the shock is immediate, as do diagnostic expectations, whereas the other three ("dampening") models suggest very slow passthrough to prices. By contrast, when the shock is in the future, our model implies the least anticipation, followed by strong passthrough 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 over-react. However, we also agree with the "dampening" theories that prior to costs moving, there will be little reaction.

4.2 The New Keynesian Phillips curve

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

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

Figure 6: 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 3.2

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

$$\pi_t = \frac{(1-\beta\delta\theta_P)\left(1-\frac{\theta_P}{\delta}\right)}{\theta_P} \cdot \sum_{\tau=0}^{\infty} (\beta\delta)^\tau \, \mathbb{E}_t\left[mc_{t+\tau}^{real}\right] + \frac{1-\delta}{\delta} \cdot mc_t^{real} + \omega_b \cdot b_t \tag{26}$$

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

$$b_t = \tilde{\mathcal{K}}\Delta mc_t + \left(1-\tilde{\mathcal{K}}\right)\tilde{\rho} b_{t-1} \tag{27}$$

and the weight placed on this belief term is

$$\omega_b \equiv \left(1-\alpha^{FIRE}\right)\frac{(1-\theta_P)\,\beta\tilde{\rho}\theta_P\zeta}{\delta\left(1-\beta\theta_P\tilde{\rho}\right)} \tag{28}$$

The additional discounting of the future is captured by the term

$$\delta \equiv \alpha^{FIRE} + \left(1-\alpha^{FIRE}\right)\theta_P\left(1-\frac{(1-\theta_p)\,\beta\tilde{\rho}\tilde{\mathcal{K}}}{\left(1-\beta\theta_p\tilde{\rho}\right)}(1-\zeta)\right) \tag{29}$$

We see that both ω_b and δ depend on ζ, which is a more complicated term that satisfies $\zeta = \frac{\left(1-\beta\tilde{\rho}\left(1-\tilde{\mathcal{K}}\right)\theta_P\right)\delta}{\left(1-\beta\tilde{\rho}\left(1-\tilde{\mathcal{K}}\right)\delta\right)\theta_P}$, giving a fixed point.

The proof to Theorem 3 is in Appendix C.3 below. 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 stan-

dard FIRE NKPC but with additional discounting, δ, given by (29).[44] Since $\delta \in [0,1]$, sufficiently distant shocks are always less important relative to FIRE.

(2) **There is higher inflation passthrough 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.

(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 (27), and so depends positively on the history of average nominal marginal cost growth.[45] The weight on this backward looking term is ω_b, and this is always 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 passthrough on impact. In our calibration, there are limited lasting effects of real marginal costs today on future inflation.

4.2.1 The Adaptive Expectations Phillips curve

The NKPC derived in Theorem 3 has a slightly non-standard form, in that the beliefs term depends on *nominal* costs. We can more fully solve out the NKPC in the special case where firms are fully Adaptive Learners, i.e. $\alpha^{\text{FIRE}} = 0$. In this case, we can write our NKPC as a function of current real marginal costs and the history of inflation,

$$\pi_t = \kappa^{\text{AL}} \text{mc}_t^{\text{real}} + \omega_b^{\text{AL}} \sum_{\ell=1}^{\infty} \left(\lambda^{\text{AL}}\right)^{\ell} \pi_{t-\ell} \tag{30}$$

where $\kappa^{\text{AL}} \equiv \frac{1}{1-\mathcal{M}} \left(\frac{1-\theta_P}{\theta_P} + \mathcal{M}\right)$, $\omega_b^{\text{AL}} \equiv \frac{\frac{\mathcal{M}}{\theta_P}\left(\frac{(1-\tilde{\mathcal{K}})\tilde{\rho}}{1-\mathcal{M}} - \theta_P\right)}{\frac{1-\theta_P}{\theta_P}(1-\tilde{\mathcal{K}})\tilde{\rho}+\mathcal{M}}$, and $\lambda^{\text{AL}} \equiv \frac{\frac{1-\theta_P}{\theta_P}(1-\tilde{\mathcal{K}})\tilde{\rho}+\mathcal{M}}{\frac{1-\theta_P}{\theta_P}+\mathcal{M}}$, and $\mathcal{M} \equiv \frac{(1-\theta_p)\beta\tilde{\rho}\tilde{\mathcal{K}}}{(1-\beta\theta_p\tilde{\rho})}$ is a multiplier term. This Phillips curve is derived with only a persistent and transitory cost component (though it nests the case of a permanent and transitory cost component). In Appendix C.5, we show that if firms are informed about permanent shocks, (30) characterizes the path of inflation around the long-run level of inflation.[46]

The functional form of (30) matches the "Friedman-Phelps" adaptive-expectations Phillips curve, derived in Appendix C.4. We provide a modern microfoundation for this Phillips curve:

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

[45] 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.

[46] In Appendix D.3, we show that if firms learn about permanent cost shock 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.

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

In particular, we can show that as the frequency of price adjustment rises (θ_P falls), the Phillips curve becomes steeper (κ^{AL} rises) and more inertial ($\omega_b^{\text{AL}} \sum_{\ell=1}^{\infty} \left(\lambda^{\text{AL}}\right)^{\ell}$ rises).

In the post-2010s United States, we have estimates for the parameters. Using our calibration from Section 3.2, $\beta = 0.99$ 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 passthrough than under FIRE ($\kappa^{\text{FIRE}} \approx 0.1$). We also predict no inflation inertia.

4.3 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. Each of the following subsections assesses a subset of these predictions in pricing data. In Section 4.3.1, we 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. Next, we use the HHNS framework to estimate an Adaptive Expectations Phillips curve and find little inertia.

The degree of forward-lookingness in the NKPC is central to interpreting the results in HHNS. Section 4.3.2 proposes a new approach to estimating forward-lookingness directly in sectoral price and cost data. We show that primitive sectoral NKPCs exhibit very little forward-lookingness and are steeper than under FIRE, both of which are suggested by our model.

4.3.1 Estimating an Adaptive Expectations Phillips curve using Hazell, Herreño, Nakamura and Steinsson (2022)

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 (30). 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: 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 substantial myopia, their results are consistent with those in this paper. The next subsection provides further empirical evidence on limited forward-lookingness – beyond our low estimate of α^{FIRE} in survey data.[47]

4.3.2 Estimating the Phillips curve in sectoral data

In this section, we empirically test whether firms' prices respond more to contemporaneous marginal costs and less to expected future marginal costs than they should under FIRE. Using data from the BEA and the BLS, we can construct quarterly changes in marginal costs net of changes in TFP. We denote these observed changes in marginal cost as

$$\Delta\widetilde{mc}_{it} \equiv \Delta mc_{it} + \Delta a_{it} = s_{iLt}\Delta w_t + (1 - s_{iLt})\Delta p_{int,t},$$

where s_{iLt} is the labor share of gross output in industry i in year t, $(1 - s_{iLt})$ is the intermediate input share, Δw_t is the change in the log wage, and $\Delta p_{int,t}$ is the change in the industry's log intermediate input price index. All information except the wage is obtained from the BEA, and the wage is obtained from the BLS's employment cost index.[48] Our sample period covers 2005 Q1 (when the BEA starts publishing this quarterly data) through 2023 Q4.[49]

To derive our estimating equation and determine the relevant assumptions for our instruments, we start with the null hypothesis that the Calvo model under FIRE pricing is correct at the industry level. We then use a primitive Phillips curve at the sectoral level from Rubbo (2023). Defining

$$\delta_i = \delta(\theta_{ip}, \beta) = \frac{(1 - \theta_{ip})(1 - \beta\theta_{ip})}{1 - \beta\theta_{ip}(1 - \theta_{ip})},$$

we have

$$\pi_{it} = \delta_i(mc_{it} - p_{i,t-1}) + \beta(1 - \delta_i)\mathbb{E}_t\pi_{i,t+1}.$$

[47] Forward-lookingness is difficult to test in HHNS's framework because it would require an additional instrument meeting F-stat requirements for predicting state-level recessions at least one quarter in advance.

[48] While the BEA publishes a price index for value added by industry, this series is a residual constructed from the measured price indices for gross output from which we will construct sectoral inflation and the industry's measured intermediate input price index.

[49] The limit on the end of the sample comes from the labor shares, which are published annually.

As in our beliefs calibration, we will use the fact that realized inflation next period is the rational expectation plus shocks that occurred since period t:

$$\pi_{i,t+1} = \mathbb{E}_t \pi_{i,t+1} + \epsilon^{FIRE}_{i,t+1}.$$

Further, we will incorporate that we only observe changes in marginal cost net of changes in TFP, requiring us to add TFP to the error term and include industry fixed effects (since we observe changes rather than levels). The altered equation is

$$\begin{aligned}\pi_{it} &= \alpha_i + \delta_i(\widetilde{mc}_{it} - p_{i,t-1}) + \beta(1-\delta_i)\pi_{i,t+1} + \epsilon_{it} \\ \epsilon_{it} &= -\delta_i a_{it} - \beta(1-\delta_i)\epsilon^{FIRE}_{i,t+1}\end{aligned}$$

Now, we will run this regression with regressors incorporating measured industry heterogeneity in frequencies of price adjustment from Pasten, Schoenle and Weber (2020). We set $\beta = 0.99$ quarterly. Formally, we will run

$$\pi_{it} = \alpha_i + b_1 x_{i,Impact,t} + b_2 x_{i,Future,t} + \epsilon_{it},$$

with $x_{i,Impact,t} = \delta_i(\widetilde{mc}_{it} - p_{i,t-1})$ and $x_{i,Future,t} = \beta(1-\delta_i)\pi_{i,t+1}$. If the Calvo model under FIRE is correct, we should find $b_1 = b_2 = 1$.

Because contemporaneous TFP enters the error term, and because future inflation is correlated with at least the future shock component of the error term, we require instruments for both regressors and cannot run OLS. We will use the Känzig (2021) shocks from periods t and $t-1$ as the instruments. This automatically handles the FIRE component of the error term, since the shock is known in period t. We also believe high frequency movements in oil futures around OPEC announcements are unlikely to be correlated with domestic sectoral TFP. To ensure this is the case, we also exclude two sectors from our sample, (1) oil and gas extraction and (2) petroleum and coal products. We assess whether our instruments are weak by using the Sanderson and Windmeijer (2016) conditional F-statistic for each endogenous variable.

The results are shown in Table 5. We use Driscoll Kraay standard errors with bandwidth three quarters. We can reject both $b_1 = 1$ and $b_2 = 1$, which would be consistent with the Calvo model under FIRE. Instead, prices appear to be much less forward-looking than under FIRE, and contemporaneous pass-through is higher. Both of these rejections of the Calvo model under FIRE beliefs are consistent with our alternative model of beliefs.

5 Aggregate implications of Adaptive Learning

5.1 Supply and Demand shocks in a Simple New Keynesian model

We start with 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

Table 5: Inflation and FIRE NKPC

	Dependent variable:
	$\pi_{i,t}$
$\beta(1-\delta_i)\mathbb{E}_t\pi_{i,t+1}$	-0.249
	(0.344)
$\delta_i(mc_{i,t}-p_{i,t-1})$	1.960^{***}
	(0.377)
Model	2SLS
Instruments	(DK_t, DK_{t-1})
Conditional F-stat $\beta(1-\delta_i)\mathbb{E}_t\pi_{i,t+1}$	16.68
Conditional F-stat $\delta_i(mc_{i,t}-p_{i,t-1})$	19.26
Fixed Effects	Industry
Observations	4,736

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$.

sticky wages and assuming a constant real rate monetary rule.

Consumption, and hence output, follows a standard Euler equation

$$y_t = -\text{eis} \cdot r_t + \mathbb{E}_t y_{t+1} \tag{31}$$

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

$$r_t = i_t^{\text{shock}} \tag{32}$$

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 to a Calvo wage setting friction,

$$\pi_t^w = \kappa^w \left(\frac{1}{\text{frisch}} n_t + \frac{1}{\text{eis}} y_t - w_t^{\text{real}} \right) + \beta \mathbb{E}_t \left[\pi_{t+1}^w \right] \tag{33}$$

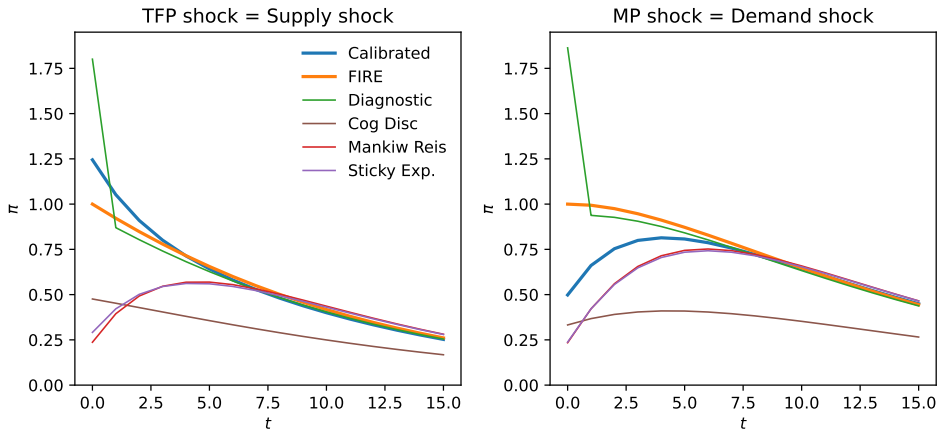
where $y_t = a_t + n_t$. Nominal marginal costs are

$$\text{mc}_t = w_t - a_t \tag{34}$$

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

$$\pi_t = \frac{1-\theta_p}{\theta_p} \left(p_t^* - p_t \right) \tag{35}$$

Figure 7: 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 3.2 above, is shown in blue.

where the optimal reset price is

$$p_t^* = (1 - \beta\theta_p) \sum_{\tau=0}^{\infty} (\beta\theta_p)^\tau \mathbb{E}_t^f [\text{mc}_{t+\tau}] \tag{36}$$

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 equate the slopes of the wage and price Phillips curves, $\kappa_p = \kappa_w = 0.9$.

In Figure 7 we show the impulse response of inflation under our calibrated model of firm beliefs (in blue), and under FIRE (in orange).[50] In our model, firms "over-react" 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 over-estimate the persistence. On the other hand, firms under-react 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 7 we also show the impulse responses under the other models of firm beliefs we

[50] 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.

introduced in Section 4.1.1 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 passthrough in many cases. In Appendix D.5, we show how a general production network has essentially the same structure. 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 F below, we also show that our results hold for supply and demand shocks as estimated in the data.

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.

5.1.1 A model with limited-information constrained rational expectations

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. We now 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 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.

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.[51] 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.

[51] 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.

The results are shown in Figure D.19. 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.[52]

5.2 How interest rates drive inflation and the forward guidance puzzle

We saw in Figure 7 that monetary policy operates differently in our model than under FIRE. In Theorem 4, we specifically focus on the effects of forward guidance. To simplify, we compare FIRE to a model with fully Adaptive Learning, and we make wages flexible.

Theorem 4. *Consider the standard New Keynesian model set out in Section 5.1, and with flexible wages. Suppose that at date* 0*, the central bank announces that at date* h*, there will be an interest rate cut. The effect on inflation at date* 0*, under FIRE, is then*

$$-\frac{d\pi_0^{FIRE}}{dr_h} = \left(1 + \frac{eis}{frisch}\right) \frac{1-\beta\theta_p}{\theta_p} \left(1 - \beta^{h+1}\right)$$

And under Adaptive Learning,

$$-\frac{d\pi_0^{AL}}{dr_h} = \left(1 + \frac{eis}{frisch}\right) \frac{1}{1-\mathcal{M}} \left[\frac{1-\theta_p}{\theta_p} + \mathcal{M}\right]$$

where $\mathcal{M} \equiv \frac{(1-\theta_p)\beta\rho\tilde{\mathcal{K}}}{1-\beta\theta_p\rho}$*. So the power of forward guidance does not grow with the horizon under fully Adaptive Learning.*

Since the contemporaneous slope of the NKPC is steeper in our model than under FIRE, we find that monetary policy is more effective at stimulating inflation over short horizons under Adaptive Learning. On the other hand, we find that distant rate decreases have a much smaller effect on inflation today than under FIRE. In Figure D.20, we quantitatively show the effects of forward guidance when wages are also sticky, and under our fully calibrated beliefs model. In this case, we see that the forward guidance puzzle is still significantly dampened, with policy more effect under our beliefs than under FIRE up to four quarters out, but weaker thereafter.

These results are despite the fact that households remain completely rational and forward-looking in our model.[53] Since firms are not themselves very forward-looking, their responsiveness to future real interest rates is relatively more dependent on how much costs move today. This

[52] 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 over-react to these and under-react to both our aggregate shocks. Our finding that firms over-react to own costs already demonstrates that LIRE fails, and is quantitatively important to match how firms respond to aggregate shocks.

[53] In particular, given log utility, a 1pp decrease in r_s leads c_t to rise by 1pp for all $t \leq s$.

suggests that when it comes to news about the future, other agents' expectations may be more important in our model. This intuition is borne out in Theorem 5.

Theorem 5. *Consider a standard household Euler equation, but with myopia (see Gabaix (2020))*

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

And suppose a one-to-one relationship between demand and real marginal costs. Under fully adaptive learning, the on impact effect of an announced future change in real interest rates is

$$\frac{d\pi_0^{AL}}{dr_t} = -\frac{eis}{1-\mathcal{M}}\left(\frac{1-\theta}{\theta}+\mathcal{M}\right)\left(m^{hh}\right)^t$$

where $\mathcal{M} \equiv \frac{\left(1-\theta_p\right)\beta\tilde{\rho}\tilde{\mathcal{K}}}{\left(1-\beta\theta_p\tilde{\rho}\right)}$. *Whereas under FIRE this is*

$$\frac{d\pi_0^{FIRE}}{dr_t} = -eis\frac{(1-\beta\theta)(1-\theta)}{\theta}\frac{\left(m^{hh}\right)^{t+1}-\beta^{t+1}}{m^{hh}-\beta}$$

Household expectations are less important under FIRE, in the sense that the proportionate effect of making households less forward-looking on the inflation response is larger when firms are adaptive,

$$\frac{d\log\frac{d\pi_0}{dr_t}}{d\log\left(m^{hh}\right)} > \frac{d\log\frac{d\pi_0^{FIRE}}{dr_t}}{d\log\left(m^{hh}\right)}$$

Theorem 5 formalizes that adaptive learning beliefs act as an amplifier for other agents' beliefs. For example, myopia of households is even more detrimental to the power of forward guidance in our model than under FIRE. Taken together with Theorem 4, the results could suggest a pivotal role for (highly forward-looking) financial markets in mapping news about future interest rates into contemporaneous inflation — but only insofar as they can affect firms' costs today.

5.3 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:

$$\text{mc}_t^{\text{real}} = \varphi y_t \tag{37}$$

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

$$p_{i,t}^* = \left(1-\beta\theta_p\right)\sum_{\tau=0}^{\infty}\left(\beta\theta_p\right)^\tau \mathbb{E}_{i,t}\left[\mu_{i,t+\tau} + \text{mc}_{i,t+\tau}\right] \tag{38}$$

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

$$\mathcal{L}_t = \mathbb{E}_t \sum_{\tau=0}^{\infty} \beta^\tau \left[\pi_{t+\tau}^2 + \vartheta y_{t+\tau}^2\right] \tag{39}$$

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

$$\mathbb{E}_0 \left[\pi_t + \frac{\vartheta}{\kappa^{\text{FIRE}}} \left(y_t - y_{t-1}\right)\right] = 0 \tag{40}$$

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

$$\pi_t + \frac{\vartheta}{\kappa^{\text{FIRE}}} y_t = 0 \tag{41}$$

In Theorem 6, we derive optimal monetary policy under Adaptive Learning.

Theorem 6. *Under purely Adaptive Learning beliefs ($\alpha^{FIRE} = 0$), optimal monetary policy is the same under commitment and discretion, and is given by*

$$\begin{aligned}
&\left(\frac{1-\theta_p}{\theta_p} + \mathcal{M}\right)\pi_t + \left(1-\mathcal{M}\right)\frac{\vartheta}{\varphi} y_t \\
&\quad = \beta \mathbb{E}_t \left[\left(\frac{1-\theta_p}{\theta_p}\left(1-\tilde{\mathcal{K}}\right)\tilde{\rho} + \mathcal{M}\right)\pi_{t+1} + \left(1-\tilde{\mathcal{K}}\right)\tilde{\rho}\frac{\vartheta}{\varphi} y_{t+1}\right]
\end{aligned} \tag{42}$$

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 (38), and the belief model in (10) and (11). This gives

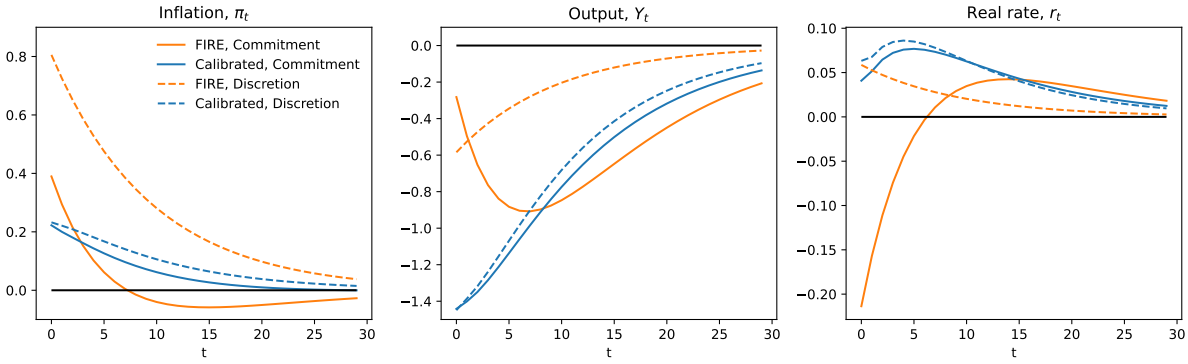
$$\begin{aligned}
\pi_t = &\frac{1}{1-\mathcal{M}}\left\{\left[\frac{1-\theta_p}{\theta_p} + \mathcal{M}\right]\text{mc}_t^{\text{real}} - \left[\frac{1-\theta_p}{\theta_p}\left(1-\tilde{\mathcal{K}}\right)\tilde{\rho} + \mathcal{M}\right]\text{mc}_{t-1}^{\text{real}}\right\} \\
&+ \frac{\left(1-\tilde{\mathcal{K}}\right)\tilde{\rho}}{\left(1-\mathcal{M}\right)}\pi_{t-1} \\
&+ \frac{1}{1-\mathcal{M}}\frac{\left(1-\theta_p\right)\left(1-\beta\theta_p\right)}{\theta_p}\left(m_t - \left(1-\tilde{\mathcal{K}}\right)\tilde{\rho} m_{t-1}\right)
\end{aligned} \tag{43}$$

where $\mathcal{M} \equiv \frac{\left(1-\theta_p\right)\beta\rho\tilde{\mathcal{K}}}{1-\beta\theta_p\rho}$ 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 (39) and this NKPC, the first order conditions are

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

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

Figure 8: 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.

where λ_t is the lagrangian multiplier on the NKPC. Combining these, we get (42). □

Since firms are not forward-looking, central bank promises (for the output gap) do not affect inflation expectations today, and the optimal policy does not vary whether solved under commitment or discretion.

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 8, 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 indeed see that under commitment, the output response under FIRE (orange, solid) is much more delayed than that under our Calibrated model (blue, solid).[54] Turning to discretion makes the response under FIRE much smaller are more frontloaded (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.

The results above for monetary policy are driven first and foremost by a lack of forward-lookingness by firms.[55] We now examine the differences in optimal monetary policy under discre-

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

[55] Note that the results on commitment are not straightforwardly the same as those on forward guidance. In particular, forward guidance asks how a given change in interest rates affects output and inflation, whereas the optimal policy problem asks how a given change in output affects inflation. While fully Adaptive Learning beliefs eliminate any returns to commitment, they do not eliminate forward guidance, and may amplify or dampen it at varying horizons.

tion to focus on other implications of our beliefs model. In particular, firms are not simply myopic, but instead extrapolate from current and past costs, and this too alters the effects of interest rates on inflation.

Theorem 7. *Consider the problem of the central bank under discretion. Under FIRE, if monetary policy induces an output gap it has only a contemporaneous effect on inflation, with magnitude κ (the slope of the NKPC)*

$$\frac{\partial \pi_t}{\partial y_s} = \kappa^{FIRE} 1_{t=s} \tag{44}$$

Under Adaptive Learning, however,

$$\frac{\partial \pi_t}{\partial y_s} = \begin{cases} 0 & if\ s > t \\ \kappa^{AL} & if\ s = t \\ \kappa^{Lag} \left(\frac{(1-\tilde{\mathcal{K}})\tilde{\rho}}{1-\mathcal{M}} \right)^{t-s} & if\ s < t \end{cases} \tag{45}$$

where $\kappa^{AL} \equiv \frac{\varphi}{1-\mathcal{M}} \left(\frac{1-\theta_P}{\theta_P} + \mathcal{M} \right)$ *and* $\kappa^{Lag} = \frac{\varphi \mathcal{M}}{(1-\mathcal{M})\theta_P} \left(1 - \frac{(1-\mathcal{M})\theta_P}{(1-\tilde{\mathcal{K}})\tilde{\rho}} \right)$. *So*

(i) The contemporaneous effect of a monetary policy induced recession is larger, as the effective slope of the NKPC is steeper, $\kappa^{AL} > \kappa^{FIRE}$

(ii) And inducing a recession today helps to fight inflation tomorrow if and only if beliefs are sufficiently backward looking

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

The first implication from Theorem 7 is that monetary policy is more powerful in the short run. This may lead to a more or less strong reaction of interest rates by the central bank, depending on the loss function.

The second implication is that the central bank may lean against anticipated future inflation by hiking rates today. 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 4, in microfounding this mechanism, we uncover a countervailing force. Overall, we found little evidence for inertia, suggesting the central bank should not lean against future inflation for this reason.[56] However, we showed above that there was some inflation inertia pre 1990 in the US, when prices were more flexible. Our model suggests monetary policy would have had more lasting effects in this period. Further, Montag and Villar Vallenas (2025), Cavallo, Lippi and Miyahara (2024), and Gagliardone, Gertler, Lenzu and

[56] However, the fact that monetary policy operates with lags provides an independent reason to act early.

Tielens (2025) argue that the frequency of price adjustment rose during Covid. Our model suggests this may have led shocks in this period to have longer lasting effects, but likewise allowed monetary policy to do the same.

6 Conclusion

We provide novel empirical results on a key object for macroeconomic propagation: firms' forecasts of their own future costs. We show that firms are overly reliant on their own past costs when making these forecasts, leading them to over-react to shocks that affect costs quickly and to under-react to shocks that affect costs slowly. In terms of policy, forward guidance and commitment are less powerful, though short term interest rates remain effective.

References

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

__ , **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, 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.

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.

Christiano, Lawrence J, Martin Eichenbaum, and Charles L Evans, "Nominal rigidities and the dynamic effects of a shock to monetary policy," *Journal of political Economy*, 2005, *113* (1), 1–45.

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.

Erceg, Christopher J, Dale W Henderson, and Andrew T Levin, "Optimal monetary policy with staggered wage and price contracts," *Journal of monetary Economics*, 2000, *46* (2), 281–313.

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.

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

Gagliardone, Luca, 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.

Känzig, Diego R, "The macroeconomic effects of oil supply news: Evidence from OPEC announcements," *American Economic Review*, 2021, *111* (4), 1092–1125.

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.

Pasten, Ernesto, Raphael Schoenle, and Michael Weber, "The propagation of monetary policy shocks in a heterogeneous production economy," *Journal of Monetary Economics*, 2020, *116*, 1–22.

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

Rubbo, Elisa, "Networks, Phillips curves, and monetary policy," *Econometrica*, 2023, *91* (4), 1417–1455.

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

Sanderson, Eleanor and Frank Windmeijer, "A weak instrument F-test in linear IV models with multiple endogenous variables," *Journal of econometrics*, 2016, *190* (2), 212–221.

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.

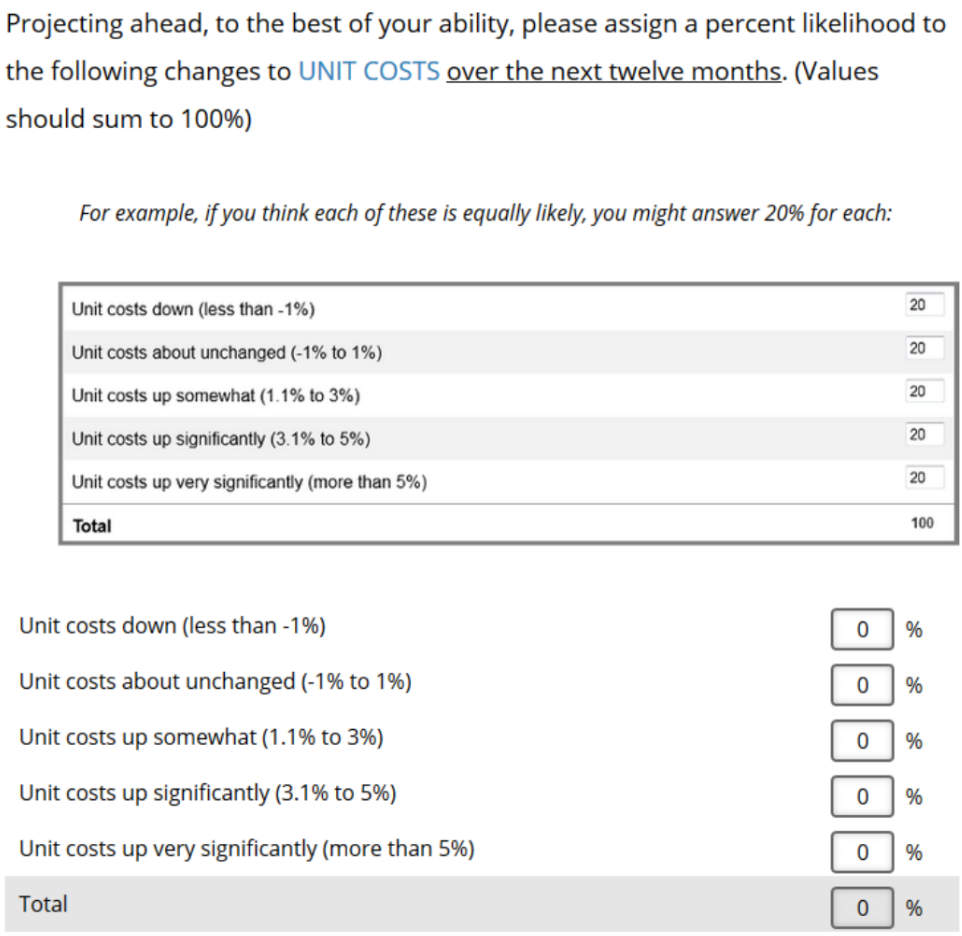
Weber, Michael, Francesco D'Acunto, Yuriy Gorodnichenko, and Olivier Coibion, "The subjective inflation expectations of households and firms: Measurement, determinants, and implications," *Journal of Economic Perspectives*, 2022, *36* (3), 157–184.

Wu, Jing Cynthia and Fan Dora Xia, "Measuring the macroeconomic impact of monetary policy at the zero lower bound," *Journal of Money, Credit and Banking*, 2016, *48* (2-3), 253–291.

A Robustness and extensions in the BIE data

A.1 Example question from the BIE

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

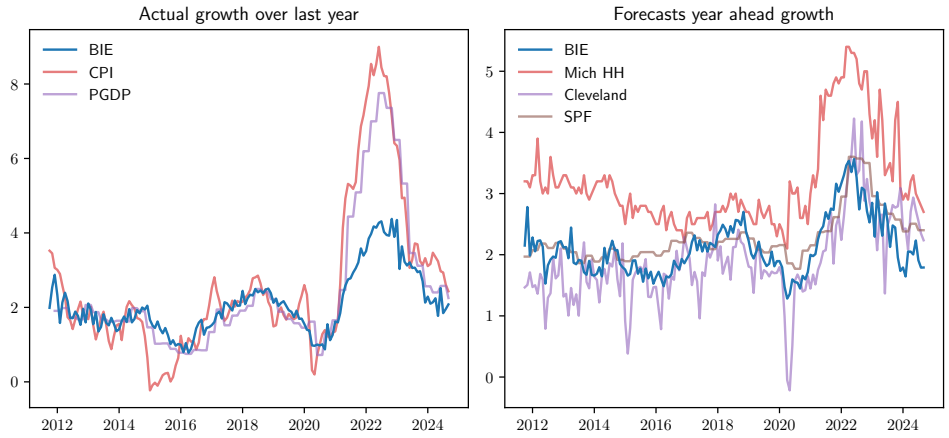
Unit costs down (less than -1%)	0	%
Unit costs about unchanged (-1% to 1%)	0	%
Unit costs up somewhat (1.1% to 3%)	0	%
Unit costs up significantly (3.1% to 5%)	0	%
Unit costs up very significantly (more than 5%)	0	%
Total	0	%

Source: FRBA Business Inflation Expectations (BIE) Survey.

A.2 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 and GDP deflator. 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, the Federal Reserve Bank of Philadelphia's Sur-

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 $\mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right]$ (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.

vey of Professional Forecasters, and the Federal Reserve Bank of Cleveland's inflation expectation measure.

A.3 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.[57] A typical question framing was

> 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%.[58] We regress expected price growth onto expected future cost growth, controlling for past price and

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

[58] 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 passthrough, if we drop any price responses outside the range $[-2\%, +6\%]$.

cost growth[59]

$$\mathbb{E}_{i,t}\left[\Delta_{12} p_{i,t+12}\right] \sim \alpha_i + \text{Controls}_{i,t} + \beta^{\text{future}} \mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right] + \beta^{\text{past}} \Delta_{12}\text{mc}_{i,t} + \delta \Delta_{12} p_{i,t} + u_{i,t} \quad (46)$$

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 passthrough 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 $[-2\%, +6\%]$.

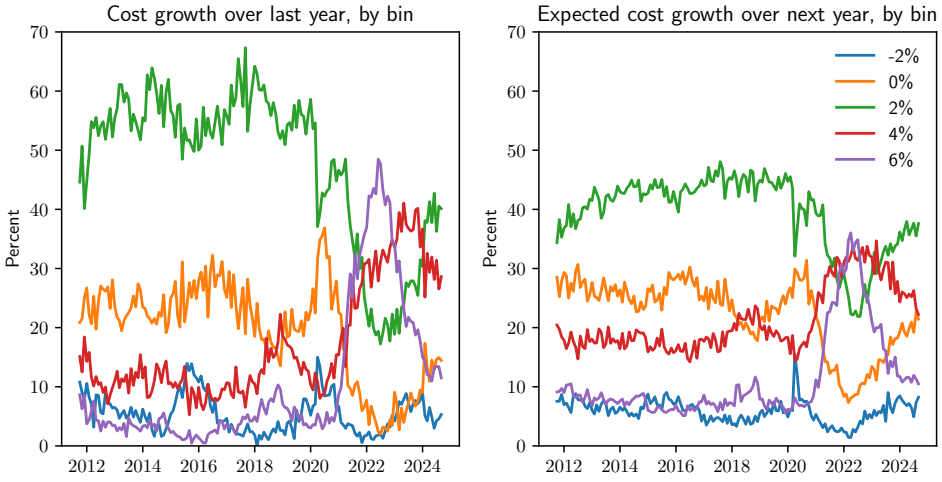
A.4 Proportion of cost responses in each bin over time

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%.[60]

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

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

Figure A.3: Average beliefs and average costs over time for each firm

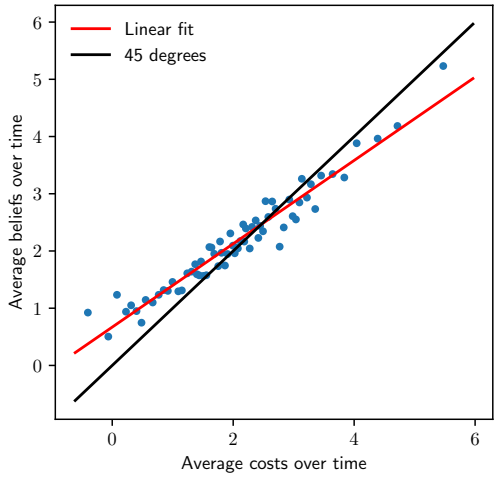


Note: Left hand side breaks down past cost growth response into five bins. Right hand side is average share of probability weight placed on the given bin when forecasting costs.

A.5 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.

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.6 Binscatter

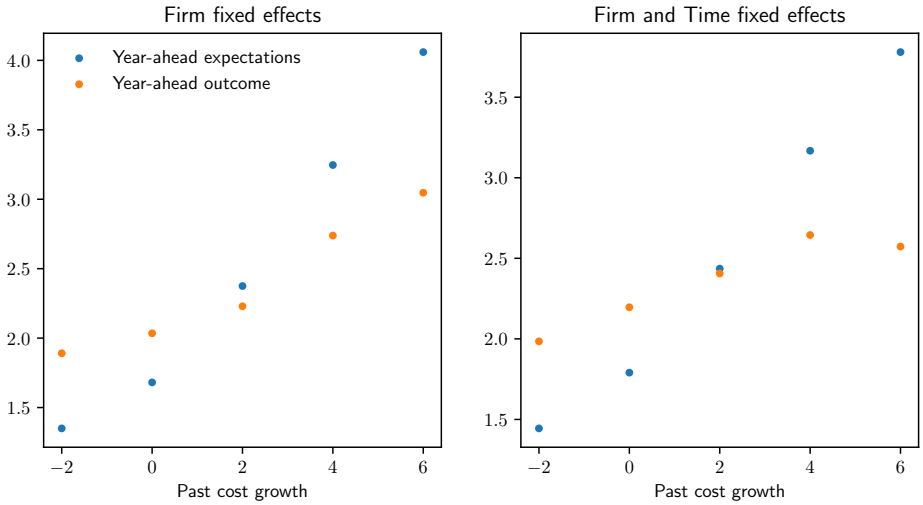
Following Cattaneo et al. (2024), we jointly estimate the mean at each of our bins and the fixed effect controls. For example, with beliefs as the dependent variable and with only firm fixed effects, we run

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

The point evaluated at j is then $\widehat{\alpha_i} + \hat{\beta}_j$. 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.

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



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

A.7 Principal component analysis (PCA) of cost growth

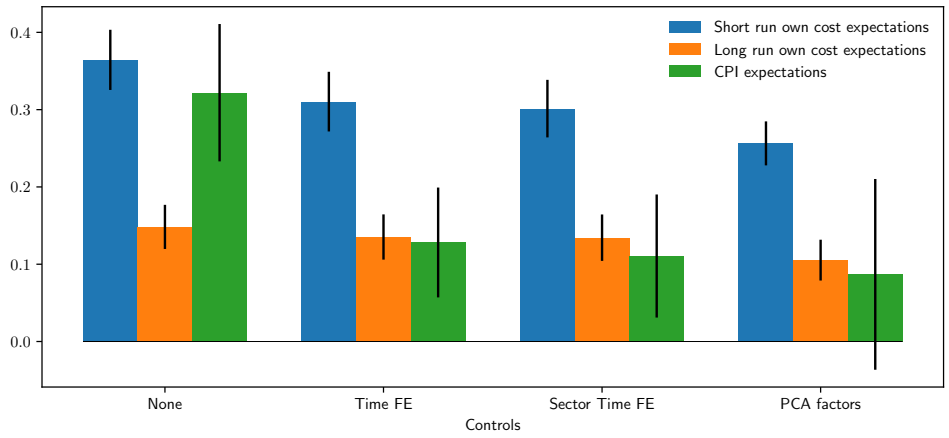
To run the PCA we first need to create a balanced panel. Whenever an observation is missing, we fit its value from firm and time fixed effects estimated through the panel, so

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

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

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

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



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

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.8 Other belief outcomes

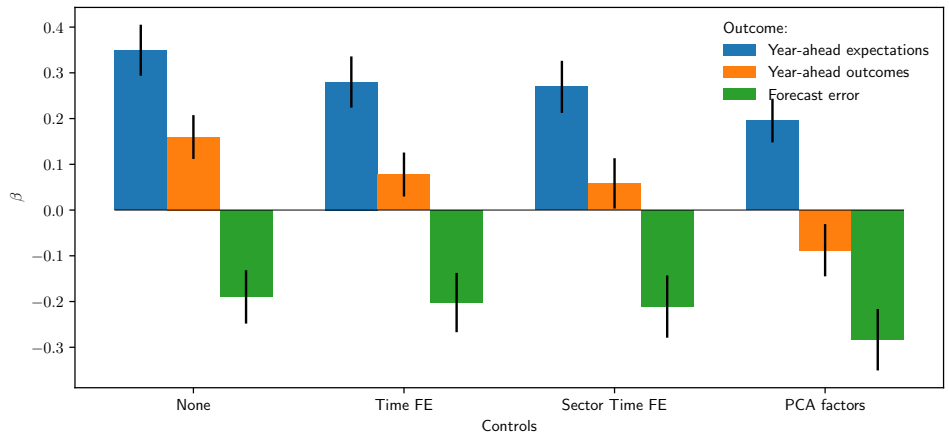
In Figure A.6, we show the coefficients from regressing other beliefs onto costs. In particular, firms appear to rely on movements in their own costs also when forecasting the very long run (their own cost growth five to ten years ahead) and forecasting aggregates (CPI inflation over the next year). 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 3.4.

A.9 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, though the results are similar with any subset of these three controls.

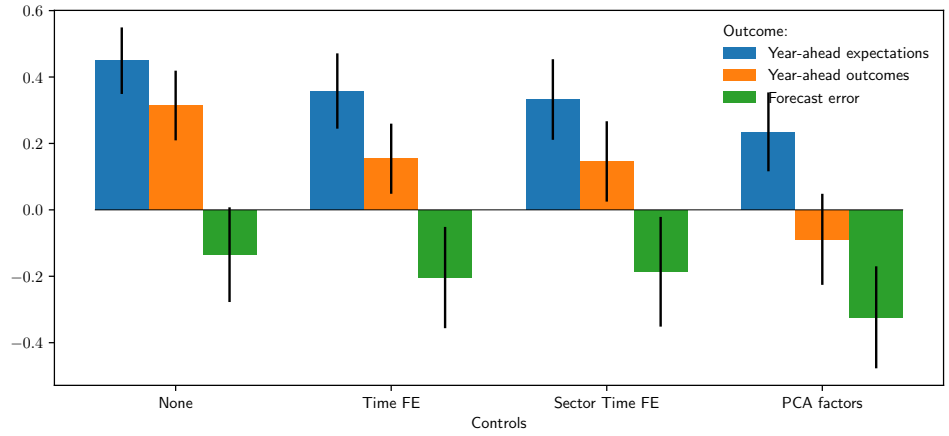
Next, we show that our effect is not undone by controlling for CPI beliefs. We show the results in Figure A.8. So although we do find a role for firms learning about CPI via their own costs in Section A.8 below – a result suggested by the existing literature – we note that understanding this channel 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.

Figure A.7: Including sales and profits as controls



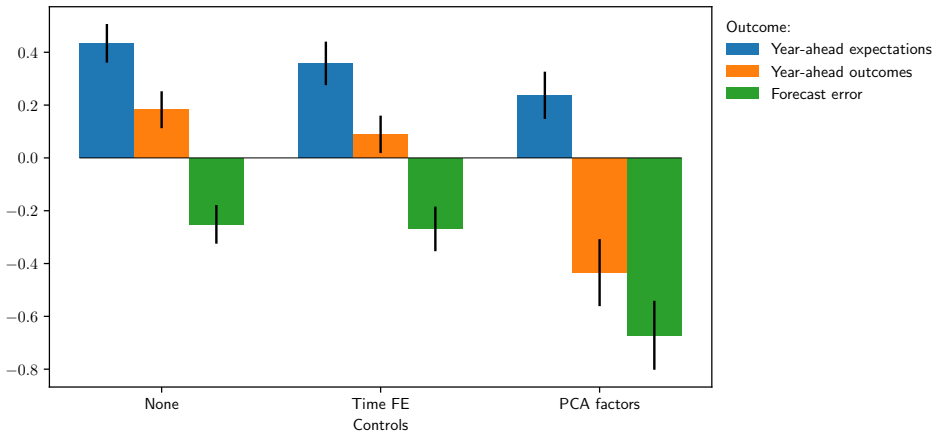
Note: Black lines are 95% confidence intervals, standard errors clustered at firm level.

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.

A.10 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.

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

$$\mathbb{E}_{it}[\Delta_{12}\text{mc}_{i,t+12}] = \alpha_i + \text{Controls}_{it} + \beta_{\text{expectations}}\widehat{\Delta_{12}\text{mc}_{it}} + u_{i,t}^{\text{expectations}} \tag{47}$$
$$\Delta_{12}\text{mc}_{i,t+12} = \alpha_i + \text{Controls}_{it} + \beta_{\text{outcome}}\widehat{\Delta_{12}\text{mc}_{it}} + u_{i,t}^{\text{outcome}} \tag{48}$$
$$\text{FE}_{i,t,t+12} = \alpha_i + \text{Controls}_{it} + \beta_{\text{FE}}\widehat{\Delta_{12}\text{mc}_{it}} + u_{i,t}^{\text{FE}} \tag{49}$$

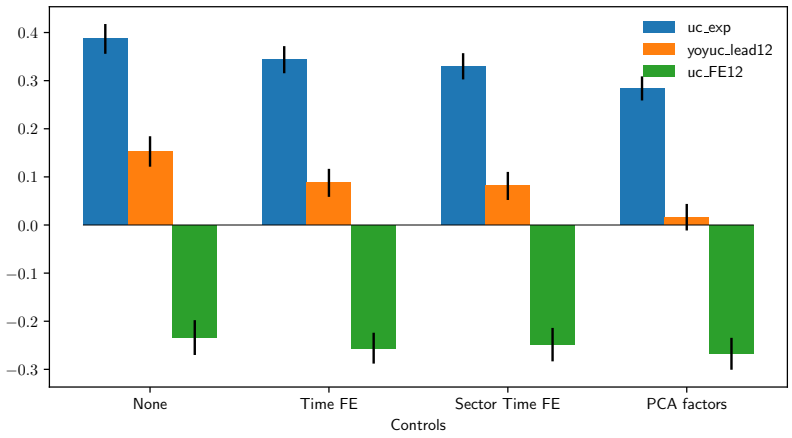
where the instrumented costs are fitted from the first-stage regression

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

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.4, we know the truncation was far less

Figure A.10: Robustness to truncation of the data



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

binding in the earlier period.

A.11 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

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

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.

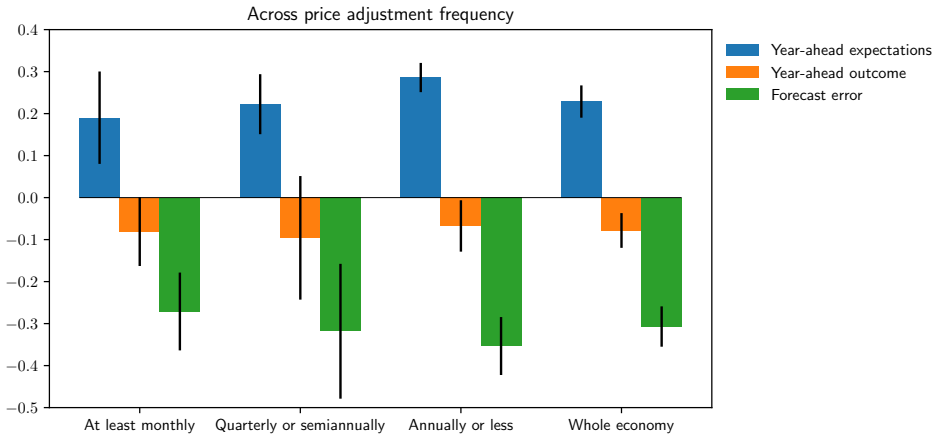
A.12 Heterogeneity by frequency of price adjustment

Firms who adjust prices less often may have more incentive to forecast costs well. In Figure A.11 we show the 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.[61] We create three groups for frequency of price adjustment: "At least monthly", "Quarterly or semiannually", and "Annually of 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

[61] 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.

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.

the last response.[62] 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.[63]

A.13 Do firms who should learn more from costs do so?

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

$$\Delta_{12}\text{mc}_{i,t} = \alpha_i + \rho_i \Delta_{12}\text{mc}_{i,t-12} + u_{i,t} \tag{50}$$

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

$$\mathbb{E}_{it}[\Delta_{12}\text{mc}_{i,t+12}] = \alpha_i + \tilde{\rho}_i \Delta_{12}\text{mc}_{i,t} + u_{i,t} \tag{51}$$

to estimate firm i's perceived cost persistence.

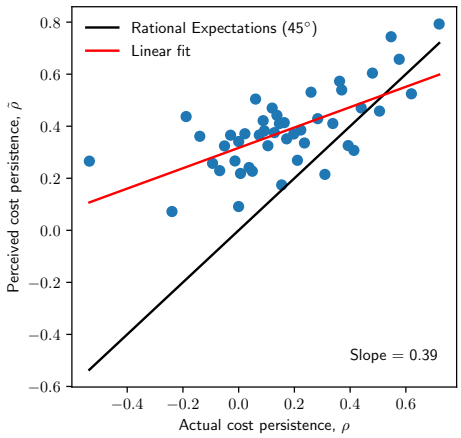
In Figure A.12, we show a scatter plot of $\tilde{\rho}_i$ against ρ_i, including only firms with at least 30 observations.[64] A 45 degree *slope* would reflect each firm correctly judging how much more persistent their costs are than the average. If firms also correctly judged the *level* of cost persistence, we would recover the 45 degree line passing through the origin. Instead, we find a much shallower 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.

[62] 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.

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

[64] 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.

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



Note: Binned scatter plot of $\hat{\bar{\rho}}_i$ and $\hat{\rho}_i$, recovered from regression equations (50) and (51) 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.

A.14 Firms over-react to costs strictly in the past

Here, we run the regression

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

This differs from (6) only in using the cost growth from last month. The results are shown in Table A.2. We continue to find significantly negative estimates for β_{FE}, telling us that firms do not only over-react in period t to surprises (relative to FIRE) in that same period.

Table A.2: Firms over-react to cost growth from last month

	Forecast error, $\Delta_{12}\text{mc}_{i,t+12} - \mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right]$			
β_{FE}	-0.16***	-0.15***	-0.15***	-0.22***
	(0.02)	(0.02)	(0.02)	(0.02)
$\text{Controls}_{i,t}$	None	Date FEs	Sector-Date FEs	PCA controls
Observations	17401	17401	17401	17401

Note: *p<0.10, **p<0.05, ***p<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

$$\mathbb{E}_t\left[\Delta \text{mc}_{i,t+\tau}\right] \sim \mathcal{N}\left(\tilde{\mathbb{E}}^{\text{AL}}_{i,t}\left[\Delta \text{mc}_{i,t+\tau}\right], \sigma^2_{\text{AL}}\right) \tag{52}$$

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,

$$s^{\text{FIRE}}_{i,t} = \mathbb{E}_t\left[\Delta \text{mc}_{i,t+\tau}\right] + \nu_{i,t}, \qquad \nu_{i,t} \overset{\text{i.i.d.}}{\sim} \mathcal{N}\left(0, \sigma^2_\nu\right) \tag{53}$$

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

$$\mathbb{E}_{i,t}\left[\Delta \text{mc}_{i,t+\tau}\right] = \alpha^{\text{FIRE}}\mathbb{E}_t\left[\Delta \text{mc}_{i,t+\tau}\right] + \left(1-\alpha^{\text{FIRE}}\right)\tilde{\mathbb{E}}^{\text{AL}}_{i,t}\left[\Delta \text{mc}_{i,t+\tau}\right] + \alpha^{\text{FIRE}}\nu_{i,t} \tag{54}$$

where $\alpha^{\text{FIRE}} \equiv \frac{\sigma^2_{\text{AL}}}{\sigma^2_{\text{AL}}+\sigma^2_\nu}$. 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^2_{AL}, 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 3.2, 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

$$\mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right] = \tilde{\rho}^{\text{ann}}_i \Delta_{12}\text{mc}_{i,t} \tag{55}$$

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

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

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},

$$\tilde{\rho}_i^{\text{ann}} = \left(1 - \alpha^{\text{RE}}\right)\overline{\rho}^{\text{ann}} + \alpha^{\text{RE}}\rho_i^{\text{ann}} \tag{56}$$

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 (55) 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.3. 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.4, 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 3.2 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.3: 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: *p<0.10, **p<0.05, ***p<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

$$\Delta\text{mc}_{i,t} = \Delta\tilde{\text{mc}}_{i,t}^{\text{Permanent}} + \Delta\tilde{\text{mc}}_{i,t}^{\text{Persistent}} + \Delta\tilde{\text{mc}}_{i,t}^{\text{Transitory}} \tag{57}$$

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

where the permanent component is given by

$$\Delta\tilde{\text{mc}}_{i,t}^{\text{Permanent}} = \Delta\tilde{\text{mc}}_{i,t-1}^{\text{Permanent}} + \tilde{\varepsilon}_{i,t}^{\text{Permanent}}, \qquad \tilde{\varepsilon}_{i,t}^{\text{Permanent}} \overset{\text{iid}}{\sim} \mathcal{N}\left(0, \tilde{\sigma}^2_{\text{Permanent}}\right) \tag{58}$$

the persistent component is given by

$$\Delta\tilde{\text{mc}}_{i,t}^{\text{Persistent}} = \tilde{\rho}\Delta\tilde{\text{mc}}_{i,t-1}^{\text{Persistent}} + \tilde{\varepsilon}_{i,t}^{\text{Persistent}}, \qquad \tilde{\varepsilon}_{i,t}^{\text{Persistent}} \overset{\text{iid}}{\sim} \mathcal{N}\left(0, \tilde{\sigma}^2_{\text{Persistent}}\right) \tag{59}$$

and the transitory component is driven by

$$\Delta\tilde{\text{mc}}_{i,t}^{\text{Transitory}} = \tilde{\varepsilon}_{i,t}^{\text{Transitory}}, \qquad \tilde{\varepsilon}_{i,t}^{\text{Transitory}} \overset{\text{iid}}{\sim} \mathcal{N}\left(0, \tilde{\sigma}^2_{\text{Transitory}}\right) \tag{60}$$

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

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

where the firm observes

$$\Delta\widehat{\tilde{\text{mc}}}_{i,t} \equiv \Delta\text{mc}_{i,t} - \Delta\tilde{\text{mc}}_{i,t}^{\text{Permanent}} \tag{62}$$

and given (59) and (60). The two key beliefs are then

$$\mathbb{E}_{i,t}\left[\Delta\tilde{\text{mc}}_{i,t}^{\text{Permanent}}\right] = \Delta\tilde{\text{mc}}_{i,t}^{\text{Permanent}} \tag{63}$$

and

$$\mathbb{E}_{i,t}\left[\Delta\tilde{\text{mc}}_{i,t}^{\text{Persistent}}\right] = \tilde{\mathcal{K}}\Delta\widehat{\text{mc}}_{i,t} + \left(1-\tilde{\mathcal{K}}\right)\tilde{\rho}\mathbb{E}_{i,t-1}\left[\Delta\tilde{\text{mc}}_{i,t-1}^{\text{Persistent}}\right] \tag{64}$$

B.4 Robustness to other instruments

The results using sector loadings are shown in Table B.5. The results using firm-level loadings and a cutoff of at least 72 observations per firm are shown in Table B.6. The results when accounting for the two-year lag of own costs are shown in Table B.7.

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 write the year ahead cost forecast in the quarterly model as

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

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

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

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

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

This is exactly analogous to the annual equation

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

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.

Table B.5: 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.

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

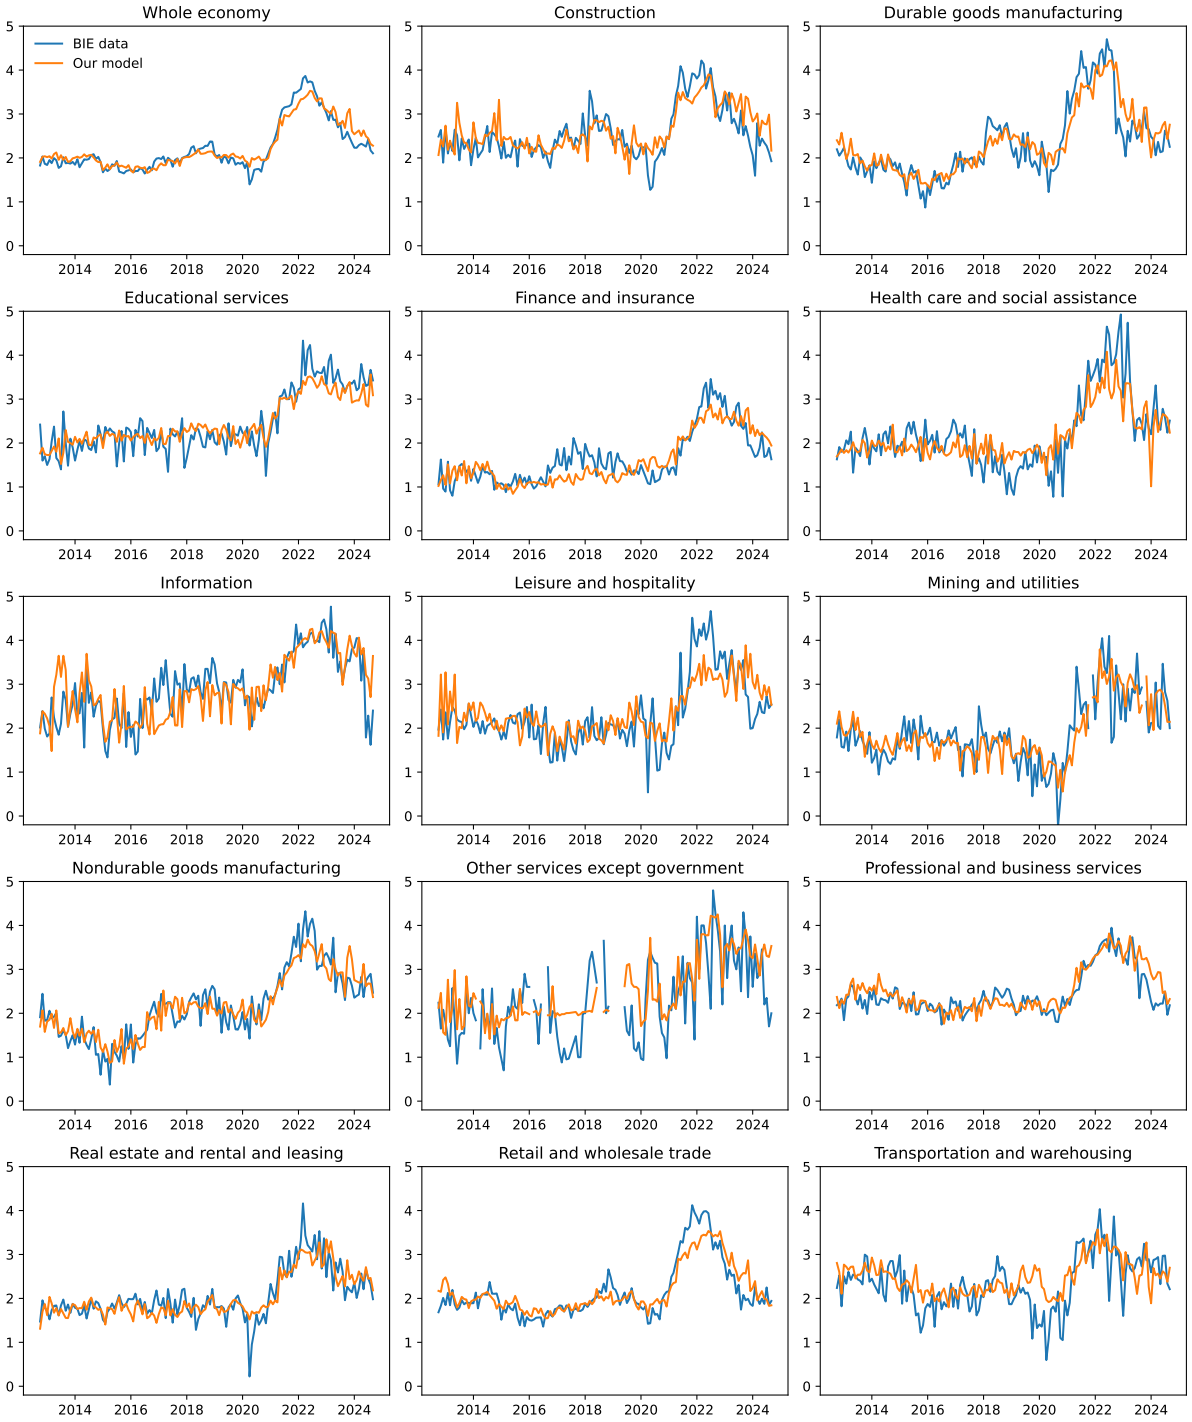


Table B.6: Annual Calibration with loadings on firms with at least 72 observations

		$\mathbb{E}_{i,t}\left[\Delta_{12}\text{mc}_{i,t+12}\right]$		
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.

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 passthrough 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

$$\Delta_{12}\text{mc}_{i,t} = \Delta_{12}\text{mc}_t^{\text{agg}} + \Delta_{12}\text{mc}_{i,t}^{idio} \tag{65}$$

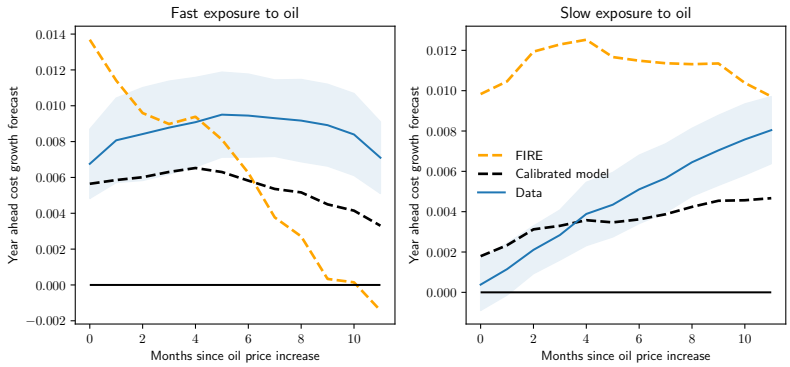
where the idiosyncratic component follows an AR(1): $\Delta_{12}\text{mc}_{i,t}^{idio} = \rho^{\text{idio}}\Delta_{12}\text{mc}_{i,t-12}^{idio} + e_{i,t}$ and $e_{i,t} \overset{\text{iid}}{\sim} \mathcal{N}\left(0, \sigma_e^2\right)$. We know that over the BIE sample period, the year-on-year percent 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

Table B.7: 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.

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



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

$$\Delta_{12}\text{mc}_{i,t}^{\text{sim}} = \Delta_{12}p_t^{\text{GDP}} + \Delta_{12}\text{mc}_{i,t}^{idio,sim} \tag{66}$$

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 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?

B.10 Historical importance of CPI in predicting own costs

The results are shown in Table B.8.

Table B.8: Estimating predictability of future costs with future CPI using (21)

	Year-on-year growth in GDP deflator, $\pi^{\text{PGDP}}_{t,t-12}$
Const	0.28
	(0.16)
$\pi^{\text{PGDP}}_{t-12,t-24}$	0.17**
	(0.07)
$\pi^{\text{CPI}}_{t,t-12}$	0.67***
	(0.07)
Observations	305
R-squared	0.90

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

C Additional Theorems and Proofs

C.1 Proof of Theorem 1

Proof. Under rational expectations, (**??**) becomes

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

Under Adaptive Learning, (**??**) becomes

$$\begin{aligned}
p_{i,t}^{*,\text{AL}} &= \text{mc}_{i,t} + \sum_{\tau=1}^{\infty}(\beta\theta)^{\tau}\,\mathbb{E}_t^{\text{AL}}\left[\Delta\text{mc}_{i,t+\tau}\right] \\
&= \text{mc}_{i,t} + \frac{\beta\theta\tilde{\rho}}{1-\beta\theta\tilde{\rho}}\mathbb{E}_t^{\text{AL}}\left[\Delta\text{mc}_{i,t}^{P}\right] \\
&= \text{mc}_{i,t} + \frac{\beta\theta\tilde{\rho}\tilde{\mathcal{K}}}{1-\beta\theta\tilde{\rho}}\sum_{\ell=0}^{\infty}\left[\left(1-\tilde{\mathcal{K}}\right)\tilde{\rho}\right]^{\ell}\Delta\text{mc}_{i,t-\ell}
\end{aligned}$$

Integrating over i, the average optimal reset price under Adaptive Learning is

$$p_t^{*,\text{AL}} = \text{mc}_t + \frac{\beta\theta\tilde{\rho}\tilde{\mathcal{K}}}{1-\beta\theta\tilde{\rho}}\sum_{\ell=0}^{\infty}\left[\left(1-\tilde{\mathcal{K}}\right)\tilde{\rho}\right]^{\ell}\Delta\text{mc}_{t-\ell}$$

□

C.2 Proof of Theorem 2

Proof. From (**??**) and (23), the period 0 response under FIRE is

$$p_0^{\text{FIRE}} = \frac{(1-\theta_p)}{(1-\beta\theta_p\rho)} \tag{67}$$

From (22), (23), and (24), the period 0 price in our model is

$$p_0 = (1-\theta_p)\left[\frac{\alpha^{\text{FIRE}}}{(1-\beta\theta_p\rho)} + \left(1-\alpha^{\text{FIRE}}\right)\left(1+\frac{\beta\theta_p\tilde{\rho}\tilde{\mathcal{K}}}{1-\beta\theta_p\tilde{\rho}}\right)\right] \tag{68}$$

So the response is larger in our model, $p_0 > p_0^{\text{FIRE}}$, if and only if

$$\tilde{\rho} > \frac{\rho}{(1-\beta\theta_p\rho)\,\tilde{\mathcal{K}} + \beta\theta_p\rho} \tag{69}$$

From (22), the difference between our price and the FIRE price dynamically satisfies

$$p_t - p_t^{\text{FIRE}} = \theta_p\left(p_{t-1} - p_{t-1}^{\text{FIRE}}\right) + (1-\theta_p)\left(1-\alpha^{\text{FIRE}}\right)\left(p_t^{*,\text{AL}} - p_t^{*,\text{FIRE}}\right) \tag{70}$$

And from (23) the FIRE optimal reset price is

$$p_t^{*,\text{FIRE}} = \frac{1}{1-\rho} + \left(\frac{\beta\theta_p\rho}{1-\beta\theta_p\rho} - \frac{\rho}{1-\rho}\right)\rho^t \tag{71}$$

In the special case where $\tilde{\mathcal{K}} = 1$, (24) implies

$$p_t^{*,\text{AL}} = \text{mc}_t + \frac{\beta\theta_p\tilde{\rho}}{1-\beta\theta_p\tilde{\rho}}\Delta\text{mc}_t = \frac{1}{1-\rho} + \left(\frac{\beta\theta_p\tilde{\rho}}{1-\beta\theta_p\tilde{\rho}} - \frac{\rho}{1-\rho}\right)\rho^t \tag{72}$$

So, then

$$p_t^{*,\text{AL}} - p_t^{*,\text{FIRE}} = (\tilde{\rho}-\rho)\,\frac{\beta\theta_p}{(1-\beta\theta_p\tilde{\rho})\,(1-\beta\theta_p\rho)}\rho^t \tag{73}$$

Combining (70) and (73),

$$p_t - p_t^{\text{FIRE}} = \left(1-\alpha^{\text{FIRE}}\right)\frac{(1-\theta_p)\,\beta\theta_p}{(1-\beta\theta_p\tilde{\rho})\,(1-\beta\theta_p\rho)}\frac{\rho^{t+1}-\theta_p^{t+1}}{\rho-\theta_p}\,(\tilde{\rho}-\rho) \tag{74}$$

In the special case where $\tilde{\rho} = \rho$, (24) implies

$$p_t^{*,\text{AL}} = \frac{1}{1-\rho} + \left(\frac{\beta\theta_p\rho}{1-\beta\theta_p\rho}\left(1-\left(1-\tilde{\mathcal{K}}\right)^{t+1}\right) - \frac{\rho}{1-\rho}\right)\rho^t \tag{75}$$

So,

$$p_t^{*,\text{AL}} - p_t^{*,\text{FIRE}} = -\frac{\beta\theta_p\rho\,(1-\tilde{\mathcal{K}})}{1-\beta\theta_p\rho}\left(\left(1-\tilde{\mathcal{K}}\right)\rho\right)^t \tag{76}$$

So

$$p_t - p_t^{\text{FIRE}} = -\left(1-\alpha^{\text{FIRE}}\right)\frac{(1-\theta_p)\,\beta\theta_p\rho}{1-\beta\theta_p\rho}\frac{\left(\left(1-\tilde{\mathcal{K}}\right)\rho\right)^{t+1}-\theta_p^{t+1}}{\left(1-\tilde{\mathcal{K}}\right)\rho-\theta_p}\left(1-\tilde{\mathcal{K}}\right) \tag{77}$$

□

C.3 Proof of Theorem 3

Proof. From the usual linearized pricing equations (**??**) and (**??**), inflation can be written as

$$\pi_t = \frac{1-\theta_P}{\theta_P}\int(1-\beta\theta_P)\sum_{\tau=0}^{\infty}(\beta\theta_P)^\tau\,\mathbb{E}_{i,t}\left[\text{mc}_{i,t+\tau}\right]\text{d}i - \frac{1-\theta_P}{\theta_P}p_t$$

Under the belief formation equations (10) and (12), this becomes

$$\pi_t = \alpha^{\text{FIRE}} \frac{(1-\theta_P)}{\theta_P} \left[(1-\beta\theta_P) \sum_{\tau=0}^{\infty} (\beta\theta_P)^\tau \mathbb{E}_t \left[\text{mc}_{t+\tau}^{\text{real}} \right] + \sum_{\tau=1}^{\infty} (\beta\theta_P)^\tau \mathbb{E}_t \left[\pi_{t+\tau} \right] \right] + \left(1-\alpha^{\text{FIRE}}\right) \frac{(1-\theta_P)}{\theta_P} \left[\text{mc}_t^{\text{real}} + \frac{\beta\theta_P\tilde{\rho}}{1-\beta\theta_P\tilde{\rho}} b_t \right]$$

where average costs are $\text{mc}_{t+\tau} \equiv \int \text{mc}_{i,t+\tau} \text{d}i$ and the average belief in the own-cost persistent component is $b_t \equiv \int \mathbb{E}_{i,t}^{\text{AL}} \left[\Delta \text{mc}_{i,t}^p \right] \text{d}i$. Applying $\mathbb{E}_t (1-\beta\theta_P F)$ to both sides, where F is the forward operator, we get the recursion

$$\begin{aligned} \pi_t = {} & \alpha^{\text{FIRE}} \frac{(1-\theta_P)}{\theta_P} (1-\beta\theta_P) \text{mc}_t^{\text{real}} \\ & + \left(1-\alpha^{\text{FIRE}}\right) \frac{(1-\theta_P)}{\theta_P} \left(\text{mc}_t^{\text{real}} - \beta\theta_P \mathbb{E}_t \left[\text{mc}_{t+1}^{\text{real}} \right] \right) \\ & + \left(1-\alpha^{\text{FIRE}}\right) \frac{(1-\theta_P)}{\theta_P} \frac{\beta\theta_P\tilde{\rho}}{1-\beta\theta_P\tilde{\rho}} \left(b_t - \beta\theta_P \mathbb{E}_t \left[b_{t+1} \right] \right) \\ & + \left(\theta_P + \alpha^{\text{FIRE}} (1-\theta_P) \right) \beta \mathbb{E}_t \left[\pi_{t+1} \right] \end{aligned}$$

Then using (11) to decompose $\mathbb{E}_t \left[b_{t+1} \right] = \zeta \mathbb{E}_t \left[b_{t+1} \right] + (1-\zeta) \tilde{\mathcal{K}} \mathbb{E}_t \left[\Delta \text{mc}_{t+1} \right] + (1-\zeta) \left(1-\tilde{\mathcal{K}}\right) \tilde{\rho} b_t$, we can write

$$\begin{aligned} (\pi_t - \omega_b b_t) = {} & \left(\frac{(1-\theta_P)(1-\beta\theta_P)}{\theta_P} + \left(1-\alpha^{\text{FIRE}}\right) \beta \left[(1-\theta_P) + (1-\zeta)\theta_P \mathcal{M} \right] \right) \text{mc}_t^{\text{real}} \\ & - \left(1-\alpha^{\text{FIRE}}\right) \left(\frac{(1-\theta_P)}{\theta_P} + (1-\zeta)\mathcal{M} \right) \beta\theta_P \mathbb{E}_t \left[\text{mc}_{t+1}^{\text{real}} \right] \\ & + \beta\delta \mathbb{E}_t \left(\pi_{t+1} - \omega_b b_{t+1} \right) \end{aligned}$$

where $\mathcal{M} \equiv \frac{(1-\theta_P)\beta\tilde{\rho}\tilde{\mathcal{K}}}{1-\beta\theta_P\tilde{\rho}}$, $\delta \equiv \alpha^{\text{FIRE}} + \left(1-\alpha^{\text{FIRE}}\right)\theta_P \left(1-\mathcal{M}(1-\zeta)\right)$, $\omega_b \equiv \frac{\left(1-\alpha^{\text{FIRE}}\right)\frac{(1-\theta_P)\beta\tilde{\rho}\theta_P\zeta}{1-\beta\theta_P\tilde{\rho}}}{\delta}$, and $\zeta \equiv \frac{\left(1-\beta\tilde{\rho}\left(1-\tilde{\mathcal{K}}\right)\theta_P\right)\delta}{\left(1-\beta\tilde{\rho}\left(1-\tilde{\mathcal{K}}\right)\delta\right)\theta_P}$. This defines a quadratic in δ, and we choose the root in $[0,1]$; we show this always exists in Lemma xx below. Given $\delta < 1$, we can iterate the above expression forward to get

$$\begin{aligned} \pi_t = {} & \left[\frac{(1-\theta_P)(1-\beta\theta_P)}{\theta_P} - \frac{(1-\delta)(1-\beta\delta)}{\delta} \right] \sum_{\tau=0}^{\infty} (\beta\delta)^\tau \mathbb{E}_t \left[\text{mc}_{t+\tau}^{\text{real}} \right] \\ & + \left(1-\alpha^{\text{FIRE}}\right) \left(\frac{(1-\theta_P)}{\theta_P} + (1-\zeta)\mathcal{M} \right) \frac{\theta_P}{\delta} \text{mc}_t^{\text{real}} \\ & + \omega_b b_t \end{aligned}$$

using $\lim_{T\to\infty} (\beta m)^T \mathbb{E}_t \left(\pi_{t+T} - \omega_b b_{t+T} \right) = 0$. This simplifies to (26). □

C.4 Derivation of the Friedman-Phelps Phillips curve

Suppose that inflation depends on current real costs and expected inflation, with some coefficient,

$$\pi_t = \kappa^{\text{AE}} \text{mc}_t^{\text{real}} + \alpha^{\text{AE}} \pi_t^e \tag{78}$$

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

$$\pi_t^e = (1 - \gamma^{\text{AE}}) \pi_{t-1}^e + \gamma^{\text{AE}} \pi_t \tag{79}$$

Combining (78) and (79),

$$\pi_t = \kappa^{\text{AE}} \text{mc}_t^{\text{real}} + \alpha^{\text{AE}} \gamma^{\text{AE}} \sum_{\ell=0}^{\infty} (1 - \gamma^{\text{AE}})^{\ell} \pi_{t-\ell} \tag{80}$$

C.5 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 (63) and (64), we derive the Phillips curve

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

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

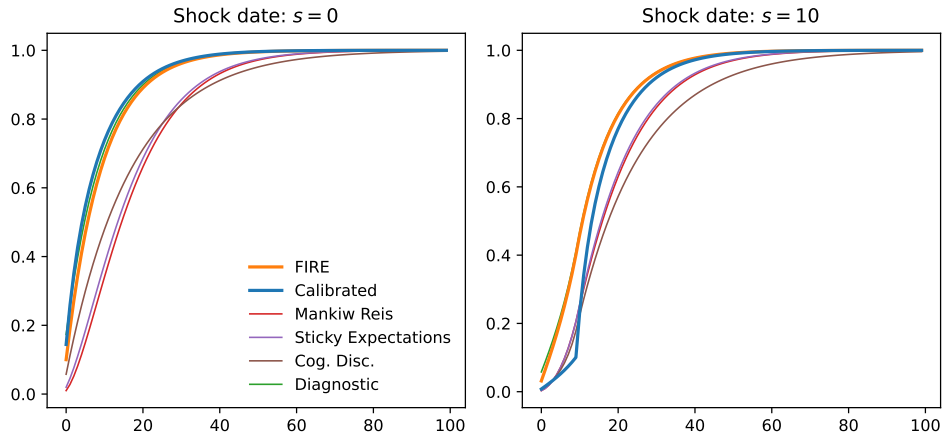
When there are no permanent shocks, $\Delta \pi_t^{\text{LR}} = 0$, (81) is identical to (30) except that it characterizes deviations from the long run. Any permanent real marginal cost shocks affect the long-run level of inflation the same in our model as under FIRE, i.e. $\pi_t^{\text{LR}} = \frac{(1-\theta_p)(1-\beta\theta_p)}{\theta_p(1-\beta)} \text{mc}_t^{\text{real, LR}}$. 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.

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.

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 3.2

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

In Figure D.16 we plot the Jacobians of the NKPC is 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 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

$$\pi = \frac{\kappa^{\text{FIRE}}}{1-\beta}\text{mc}^{\text{real}} \tag{82}$$

where $\kappa^{\text{FIRE}} \equiv \frac{(1-\theta_P)\left(1-\beta\theta_p\right)}{\theta_P}$. 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

$$\left(\frac{\pi}{\text{mc}^{\text{real}}}\right)^{\text{AL}}_{\text{LR}} = \frac{\left(1-\lambda^{\text{AL}}\right)\kappa^{\text{AL}}}{1-\left(1+\omega_b^{\text{AL}}\right)\lambda^{\text{AL}}} = \frac{\kappa^{\text{FIRE}}}{1-\beta} = \left(\frac{\pi}{\text{mc}^{\text{real}}}\right)^{\text{FIRE}}_{\text{LR}} \tag{83}$$

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.

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

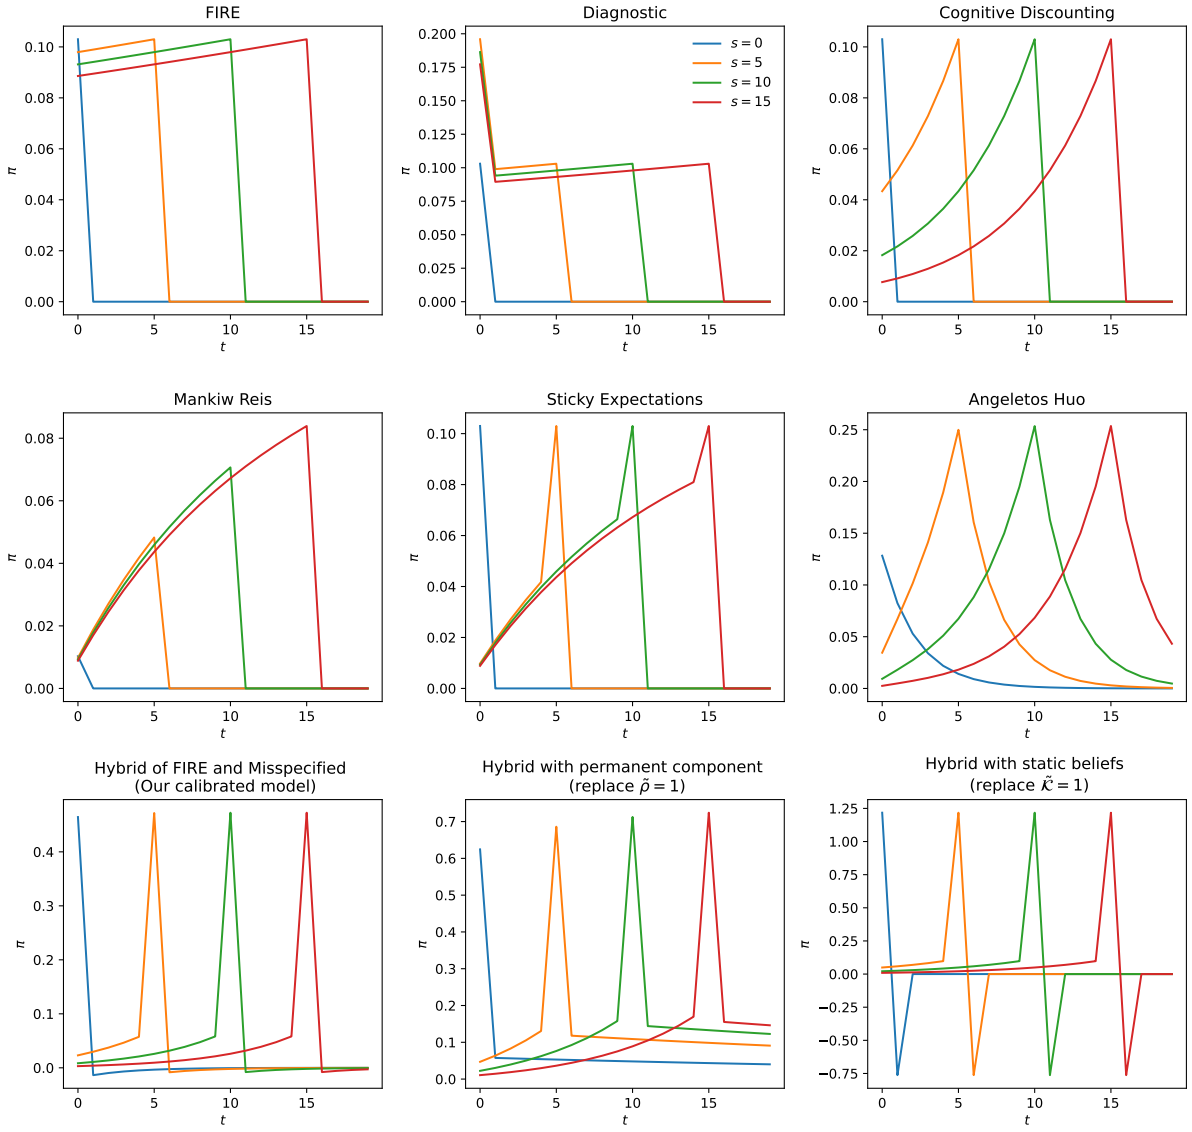
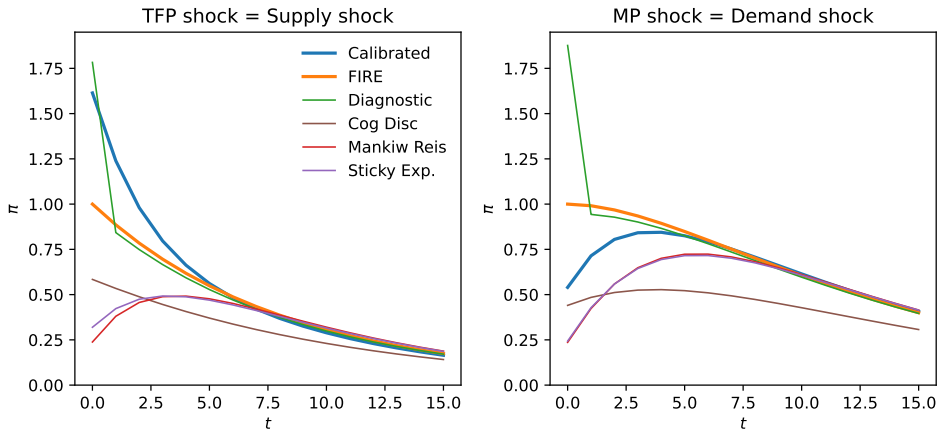


Figure D.17: Inflation response to supply and demand shocks under Taylor Rule



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 3.2 above, is shown in blue.

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

Here we take our simple model from Section 5.1, 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

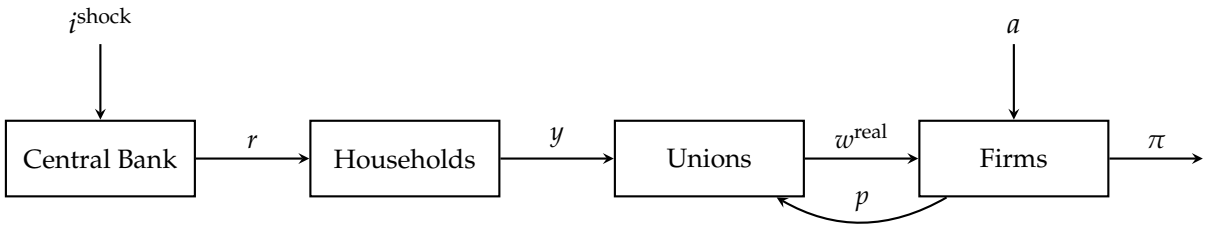
D.5 Demand versus supply logic carries over to production network

In Figure D.18a we represent our New Keynesian model graphically. The key difference between supply and demand shocks in our model is that demand shocks pass through the union block, slowing down transmission. In Figure D.18b, we see that exactly the same intuition holds for a general production network.

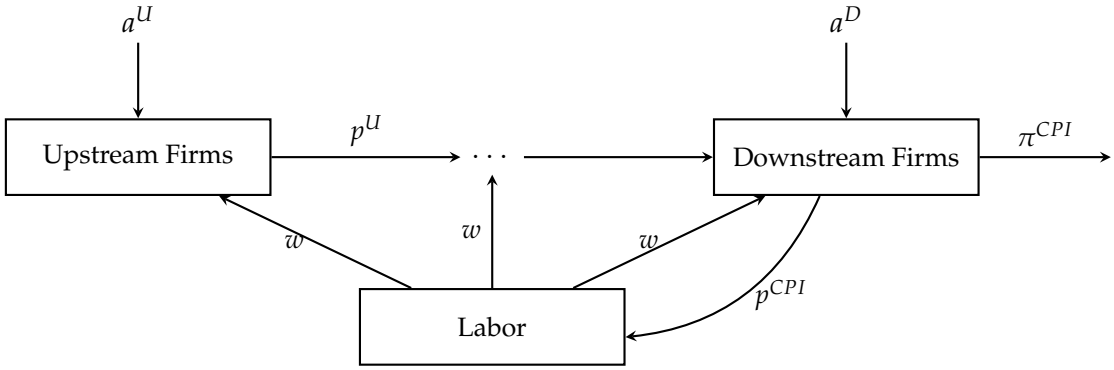
D.6 Demand and supply shocks under rational expectations

We show the response of inflation under limited information (constrained) rational expectations in Figure D.19.

Figure D.18: Upstream and Downstream shocks behave like Demand and Supply shocks



(a) New Keynesian model



(b) Production network

D.7 Forward guidance

E Pricing regressions

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

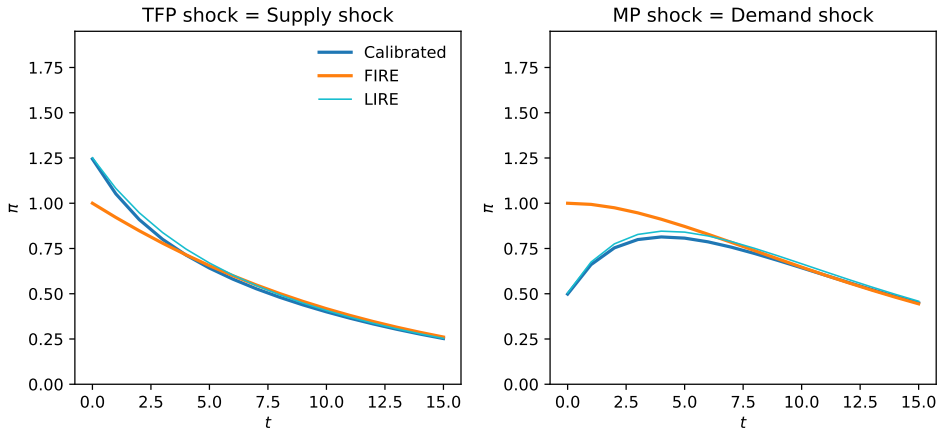
In this section, we estimate as adaptive Phillips curve using state-level inflation data and instruments from HHNS. HHNS estimate the NKPC slope from

$$\pi^N_{i,t} = -\kappa^{\text{FIRE}}_{\text{unemp}} \sum_{j=0}^{T} \beta^j u_{i,t+j} + \text{FE}_{i,t} + \text{Controls}_{i,t} + \varepsilon_{i,t}, \qquad \text{(HHNS eq.17)}$$

where $\pi^N_{i,t}$ 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),

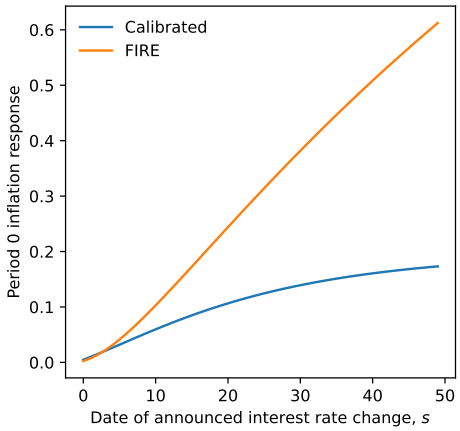
$$\pi^N_{i,t} = -\kappa^{\text{myopic}}_{\text{unemp}} u_{i,t-4} + \text{FE}_{i,t} + \text{Controls}_{i,t} + \varepsilon_{i,t} \qquad \text{(HHNS eq.19)}$$

Figure D.19: 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 3.2 above, is shown in blue.

Figure D.20: Forward guidance is much less powerful in our model than under FIRE



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

Table E.9: 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)	(0.038)	(0.037)
State effects	✓	✓	✓
Time effects		✓	✓
Instruments			Tradeable demand

To estimate our adaptive Phillips curve (30), we follow the timing convention of (HHNS eq.19) but allow a role for lagged inflation,

$$\pi^{N}_{i,t} = -\kappa^{\text{AL}}_{\text{unemp}} u_{i,t-4} + \omega^{\text{AL}}_{b} \pi^{N}_{i,t-4} + \text{FE}_{i,t} + \text{Controls}_{i,t} + \varepsilon_{i,t} \tag{84}$$

The results are shown in Table E.9. Under the null of no forward-lookingness, the Phillips curve is an order of magnitude steeper; estimates of $\kappa^{\text{myopic}}_{\text{unemp}}$ and $\kappa^{\text{AL}}_{\text{unemp}}$ are both near 0.08, more than ten times larger than the estimate of $\kappa^{\text{FIRE}}_{\text{unemp}}$. We can translate our estimate for $\kappa^{\text{AL}}_{\text{unemp}}$ into an estimate for κ^{AL}, the slope of the cost-based adaptive Phillips curve (30), using the result from Gagliardone et al. (2023) that a one percentage point decrease in unemployment increases real marginal costs by 0.25 percent.[65] 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.9 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.10, we also show that inertia remains limited when including additional lags of inflation in the estimation of (84).

In Table E.10 we show the regression result when allowing for a two year lag of inflation.

In Table E.11 we show the result when splitting the sample into pre- and post-1990.

[65] 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.10: Estimating an Adaptive Expectations Phillips Curve

Estimates of κ from equation (HHNS eq.17)				
κ^{FIRE}	-0.0037	0.0003	0.0062	0.0062
	(0.0013)	(0.0019)	(0.0028)	(0.0025)
Estimates of adaptive expectations PC params				
κ^{A}	0.027	0.055	0.087	0.289
	(0.022)	(0.022)	(0.0513)	(0.130)
$\omega^{A}_{b,\text{Lag4}}$	0.313	0.217	0.0713	0.027
	(0.040)	(0.042)	(0.032)	(0.034)
$\omega^{A}_{b,\text{Lag8}}$	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.11: Estimating an Adaptive Expectations Phillips Curve

Estimates of κ from equation (HHNS eq.17)				
κ^{FIRE}	0.011	0.011	0.005	0.005
	(0.008)	(0.006)	(0.004)	(0.003)
Estimates of adaptive expectations PC params				
κ^{A}	0.177	0.369	0.094	0.345
	(0.102)	(0.213)	(0.059)	(0.165)
ω^{A}_{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

F Quantitative macroeconomic model

We now show that our results hold up in an estimated, quantitative model. We embed our beliefs model in a medium scale New Keynesian model, following Christiano, Eichenbaum and Evans (2005) and Smets and Wouters (2007).

F.1 GE model

Households. We consider a representative household, with FIRE expectations, facing internal additive habits. The household solves

$$\begin{aligned} \max_{C_t} \quad & \mathbb{E}_t \sum_{\tau=0}^{\infty} \beta^\tau Z_{t+\tau} u\left(C_{t+\tau} - \varphi_h C_{t+\tau-1}\right) \\ s.t. \quad & C_t + \mathcal{A}_t = w_t^{\text{real}} N_t + \text{Profit}_t - T_t + (1 + r_{t-1})\mathcal{A}_{t-1} \end{aligned}$$

where the utility function is $u(C) = \frac{C^{1-\text{eis}^{-1}}}{1-\text{eis}^{-1}}$. Z_t is a discount factor shock following an AR(1) process,

$$\log Z_t = \rho_z \log Z_{t-1} + \varepsilon_t^z$$

Goods production. The firm production function is

$$Y_{i,t} = A_t K_{i,t}^{\alpha} L_{i,t}^{1-\alpha}$$

Labor and capital are both sourced from a common market, at prices w_t^{real} and r_t^K, respectively. So firm i's period t real marginal cost is

$$\Psi_{i,t}^{\text{real}} = \frac{1}{A_t} \left(\frac{r_t^K}{\alpha}\right)^{\alpha} \left(\frac{w_t^{\text{real}}}{1-\alpha}\right)^{1-\alpha}$$

where TFP follows the AR(1),

$$\log A_t = \rho_a \log A_{t-1} + \varepsilon_t^a$$

And the aggregate factor demands are

$$\begin{aligned} K_t &= \alpha \frac{\Psi_t^{\text{real}}}{r_t^K} Y_t \\ L_t &= (1-\alpha) \frac{\Psi_t^{\text{real}}}{w_t^{\text{real}}} Y_t \end{aligned}$$

Unions. Unions set wages subject to a Calvo friction as in Erceg, Henderson and Levin (2000),

$$\pi_t^W = \kappa^W \left(\frac{\mu_{W,t} N_t^{1/\text{frisch}} C_t^{1/\text{eis}}}{w_t^{\text{real}}} - 1 \right) + \beta \mathbb{E}_t \pi_{t+1}^W$$

where $\kappa^W \equiv \frac{(1-\beta\theta_W)(1-\theta_W)}{\theta_W(1+\varepsilon_W/\text{frisch})}$.

Investment. Capital is produced by firms subject to a quadratic adjustment cost. They solve the problem

$$\max_{I_t^K} \mathbb{E}_t \sum_{\tau=0}^{\infty} \frac{1}{\prod_{j=0}^{\tau-1}(1+r_{t+j})} \left[r_{t+\tau}^K K_{t+\tau} - I_{t+\tau}^K \left(1 + \frac{\varphi_K}{2} \left(\frac{I_{t+\tau}^K}{I_{t+\tau-1}^K} - 1 \right)^2 \right) \right]$$

subject to

$$K_{t+1} = (1-\delta_K) K_t + A_t^{\text{inv}} I_t^K$$

where there are shocks to the investment technology,

$$\log A_t^{\text{inv}} = \rho_{\text{inv}} \log A_{t-1}^{\text{inv}} + \varepsilon_t^{\text{inv}}$$

Central bank. The central bank sets the nominal interest rate according to

$$1 + i_t = (1+i_{t-1})^{\rho_m} \cdot \left((1+r_{ss}) \cdot (1+\pi_t)^{\phi_\pi} \right)^{1-\rho_m} \cdot (1+i_t^{\text{shock}})$$

Here r_{ss} is the steady state real interest rate, ρ_m captures monetary policy inertia, and ϕ_π is the Taylor rule coefficient on inflation. The monetary policy shock term follows the process

$$\log(1+i_t^{\text{shock}}) = \rho_m \log(1+i_{t-1}^{\text{shock}}) + \varepsilon_t^m$$

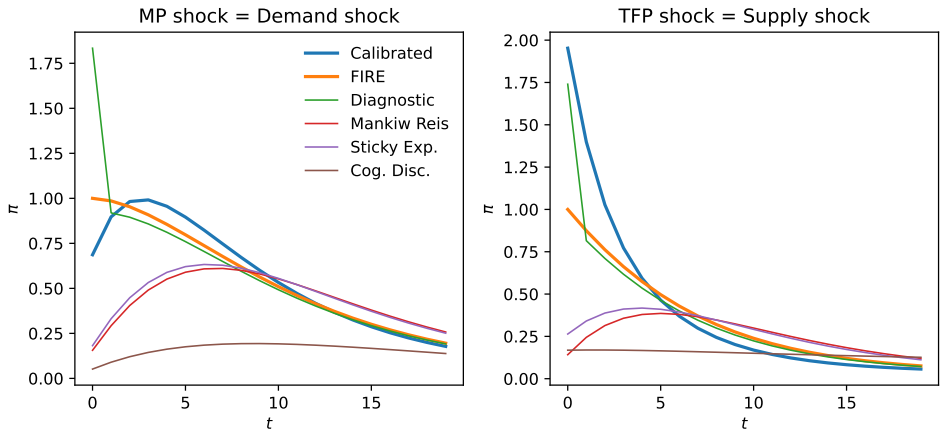
We also have the Fisher equation,

$$1 + r_t = \frac{1+i_t}{1+\mathbb{E}_t \pi_{t+1}}$$

Government. Each period, real government spending is G_t, and this is financed out of same-period lump-sum taxation: $T_t = G_t$.[66] We assume

$$\log G_t = (1-\rho_g) \log G_{ss} + \rho_g \log G_{t-1} + \varepsilon_t^g$$

Figure F.21: Inflation response to estimated TFP and Monetary Policy shocks



Note: The response of inflation to an estimated monetary policy shock (left) and to an estimated TFP shock (right) under different models of firm beliefs. Impulse responses normalized to 1 under FIRE. Our model, calibrated in Section 3.2 above, is shown in blue.

F.2 The response of inflation to estimated shocks

We use Bayesian estimation to pin down the remaining parameters in the model We also estimate the variance and persistence terms for our seven shocks: monetary policy (σ_m, ρ_m), total factor productivity (σ_a, ρ_a), discount rate (σ_z, ρ_z), government spending (σ_g, ρ_g), investment $(\sigma_{\text{inv}}, \rho_{\text{inv}})$, price markup (σ_p, ρ_p), and wage markup (σ_w, ρ_w). We match HP filtered data on GDP, consumption, investment, hours, real wages, CPI inflation, and the nominal interest rate.[67] The full calibration is shown in Table xx.

We can now plot the impulse response of inflation, under different models of firm beliefs, to a monetary policy and a TFP shock in our model. The key difference to Section 5.1 is that the persistence of these shocks is estimated, and the shocks are mediated through our more realistic quantitative model. The results are shown in Figure F.21, and we see that the key insights remain: under our beliefs, the TFP shock is amplified and the monetary policy shock is dampened.

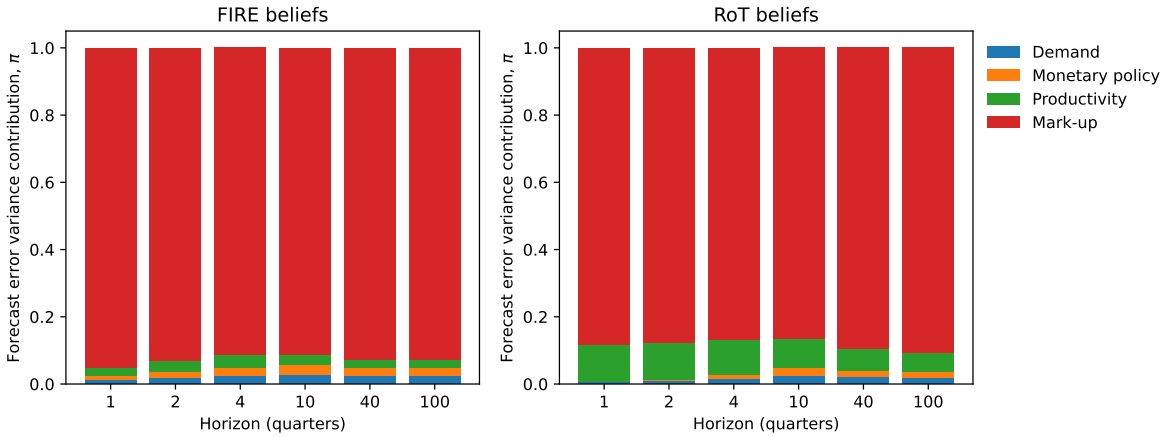
F.3 What has driven inflation historically?

In this section, we replicate the forecast error variance decomposition exercise in Smets and Wouters (2007), focusing on inflation. We compare the decomposition under the Smets and Wouters model with FIRE beliefs against the Smets and Wouters model with (calibrated) Adaptive Learning beliefs. The results are shown in Figure F.22. We see that the one-year-ahead contribution of productivity shocks rises from just 4% under FIRE to 11% under Adaptive Learning, with even larger increases at shorter horizons.

[66] Given Ricardian equivalence holds in our model, the fiscal rule does not matter.

[67] We use data up to the end of 2019, so we avoid the Covid period. To properly account for interest rates during the ZLB period, we use the shadow Federal Funds Rate from Wu and Xia (2016).

Figure F.22: Forecast error variance decomposition of inflation



We estimate the model with seven shocks, just as in Smets and Wouters (2007). To make the decomposition more readable, we group the wage and price markup contributions; and we group investment, exogenous spending, and risk premium shocks into 'Demand'. We fully estimate the model in both cases, and the only change under AL beliefs is to edit the NKPC equation, substituting in (as exogenous parameters) our calibrated values for $\bar{\rho}$ and $\bar{\mathcal{K}}$.