Inflation Expectations under Finite Horizon Planning*

Christopher Gust† Edward Herbst‡ David López-Salido§

March 2023

Abstract

We examine the dynamics of inflation expectations in a behavioral New Keynesian (NK) model in which households and firms plan over a finite horizon. With finite horizon planning agents can evaluate a full set of state-contingent paths along which the economy might evolve out to a finite horizon but have limited ability to fully process events beyond that point. We show–analytically and empirically–that the finite horizon planning model can account for key stylized facts about inflation expectations and the predictability of consensus inflation forecast errors, including an initial underreaction and subsequent overreaction of inflation forecasts to new information. We then study the ability of the model to jointly explain survey data on inflation expectations along with the aggregate data on output, inflation, and interest rates. We find that NK-FHP outperforms alternative paradigms of expectations formation, such as those involving information rigidities.

JEL Classification: C11, E52, E70
Keywords: Finite horizon planning, inflation expectations, bayesian estimation.

*The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System.

†Federal Reserve Board. Email: christopher.gust@frb.gov
‡Federal Reserve Board. Email: edward.herbst@frb.gov
§Federal Reserve Board. Email: david.lopez-salido@frb.gov

1 Introduction

Motivated by limitations in the cognitive ability of people to understand and process information, macroeconomists have increasingly begun to incorporate behavioral elements into their models as an alternative to rational expectations. A novel approach in this regard is the finite horizon planning (FHP) framework developed in Woodford (2018) in which agents are boundedly rational, as their ability to evaluate the full set of state-contingent paths along which the economy might evolve is limited to a finite horizon. To highlight the appeal of the approach, Woodford (2018) embeds FHP into a New Keynesian (NK) model and shows that monetary policy does not suffer from a "forward guidance" puzzle in which a credible promise to keep the policy rate unchanged in the distant future produces counterfactually large effects on current inflation and output.

Our previous research, Gust, Herbst, and López-Salido (2022), provides additional evidence that a New Keynesian model with FHP is a compelling framework for understanding aggregate output, inflation, and interest-rate dynamics. Our results suggest that this model is able to generate substantial inflation persistence and realistic costs to an anticipated disinflation announced by a central bank. In addition, the model fits the macroeconomic time series substantially better than other behavioral models as well as the "hybrid" NK model that features rational expectations, habit persistence in consumption, and exogenous price indexation.[1]

While the NK-FHP model has had some success in explaining the macroeconomic time series, it remains an open question how well it accounts for some key stylized facts that have emerged from empirical studies using survey data on inflation expectations. This literature finds that forecast errors are systematically predictable in a way that is difficult to rationalize with macroeconomic models that feature full information, rational expectations (FIRE). This research emphasizes that survey data on expectations help to discriminate across alternative models of expectation formation, and several papers in this literature point to stylized facts that are difficult to reconcile with several popular behavioral models including those emphasizing diagnostic expectations or cognitive discounting.[2] Important contributions to this literature include Coibion and Gorodnichenko (2015) (hereafter CG (2015)), Angeletos, Huo, and Sastry (2020) (hereafter AHS (2020)), and Kohlhas and Walther (2021) hereafter KW (2021).[3] CG (2015) study the correlation between consensus forecast errors and forecast revisions of inflation and find evidence consistent with an *underreaction* of forecasts to revisions. The evidence in KW (2021) also supports the finding of an underreaction of forecasts to revisions but KW (2021) also provide evidence of an *overreaction* of average forecasts to recent data. AHS (2020) study the impulse responses of inflation forecasts from the survey of

[1]The specific behavioral models that we compare to the NK-FHP model are the models of Angeletos and Lian (2018) and Gabaix (2020).

[2]See, for example, Kohlhas and Walther (2021), who point to some evidence that they suggest is challenging to explain with simple formulations of diagnostic expectations. Similarly, Angeletos, Huo, and Sastry (2020) present evidence that seems at odd with simple formulations of cognitive discounting.

[3]We focus on the evidence from these papers because they are directly relevant to the macroeconomic models that we investigate. Another important branch of this literature, including Bordalo, Gennaioli, Ma, and Shleifer (2018), Fuhrer (2018), and Broer and Kohlhas (2018), examines the predictability of forecasts errors of individual forecasters instead of average or consensus forecasts.

professional forecasters (SPF) and find that the average forecast underreacts to shocks initially *but overreacts later on.*

To understand the implications of the FHP model for the predictability regressions and impulse responses of inflation, we begin by focusing on a partial equilibrium model in which FHP firms set prices in a staggered fashion. We characterize the dynamics of aggregate inflation and firms' forecast of inflation analytically. This allows to provide formal conditions under which inflation forecasts under FHP are consistent with the predictability impulse responses of AHS (2020) and predictability regressions of CG (2015) and KW (2021). We find that the version of the FHP model in which price setters learn and update their beliefs about events outside their planning horizons is a key feature necessary to account for this evidence. To provide intuition regarding the role of outside-the-planning-horizon beliefs in the FHP model, it is helpful to first consider the forecasting properties of inflation when price-setters' beliefs for events outside their planning horizon remain fixed and compare it to the case of FIRE. Under FIRE, a price-setting firm with an opportunity to reset its price takes into account the full effects of a persistent shock into the distant future and efficiently incorporates those beliefs into its forecast. Under finite planning without learning, firms' beliefs outside of their planning horizon are fixed, and firms neglect the longer-lasting effects of persistent shocks so that their forecasts systematically underpredict realized inflation.

While an underreaction is consistent with the evidence of CG (2015), the no-learning version of the model fails to account for the longer-run overreaction that AHS (2020) find for inflation forecasts as well as evidence of an overreaction emphasized by KW (2021). The FHP model's predictions can account for this evidence when firms learn adaptively and update their longer-run beliefs (i.e., those outside their planning horizons) about inflation. We show that these outside-the-planning-horizon beliefs depend on past inflation and evolve sluggishly: they do not change much in the short run but can move significantly over time as firms acquire more information on previously observed aggregate inflation. Because a firm's forecast of inflation at shorter horizons also depends on their (slowly evolving) longer-run beliefs, its forecasts inherit this inertia, thus responding sluggishly at first. Accordingly, agents' inflation forecasts undershoot realized inflation initially but overshoot it later on. This later overshoot reflects that firms' inflation forecasts overweigh changes in beliefs about events that occur outside their planning horizon. When these beliefs eventually adjust, they help generate impulse responses in line with the empirical ones in AHS (2020). We formalize these results through three propositions highlighting that there is a wide range of parameter values for which inflation expectations under FHP is qualitatively consistent with the specific patterns of inflation forecast errors emphasized by CG (2015), KW (2021), and AHS (2020). These parameters go beyond the parameter governing the length of firms' planning horizons and include those governing the rate at which firms update their longer-run beliefs and the persistence of shocks.

We then extend the analysis to the dynamic, general equilibrium model studied by Woodford (2018) and find that these results apply in general equilibrium as well. In particular, the NK-FHP model with learning in Woodford (2018) is capable of generating forecast errors in line with the patterns of inflation forecast errors in the empirical literature. Moreover, the general equilibrium version introduces richer inflation dynamics and inflation expectations are influenced by different type of shocks as well as by other aspects of the economy such as the central bank's reaction function for setting interest rates. Accordingly, the parameters governing the shock processes and the central bank's interest-rate reaction function influence the model's implications for the predictability regressions of CG (2015) and KW (2021) and predictability impulse responses of AHS (2020), which, in turn, highlights the importance of examining this evidence using a dynamic, general equilibrium model.

Given the prominent role of structural parameters auxiliary to expectations formation, it is difficult to use only the empirical estimates from the predictability regressions and impulse responses to guide our assessment of the FHP model. Rather, we leverage the approach in our previous research, Gust, Herbst, and López-Salido (2022), and estimate the model employing a Bayesian, full-information likelihood-based approach using U.S. data on output growth, inflation, and nominal interest rates from 1966:Q1 through 2007:Q4, a time period for which there were notable changes in trends in inflation and output growth. This approach allows us to identify all the structural parameters of the model. Armed with these parameter estimates, we compute the model's counterparts to the empirical moments emphasized by CG (2015), KW (2021), and AHS (2020) and evaluate the model using these statistics. An important aspect of our procedure is that it does not directly incorporate information on inflation forecasts or forecast errors. Thus, the information coming from empirical moments emphasized by CG (2015), KW (2021), and AHS (2020) can be viewed as form of external validation of the model. Our main finding using this approach is that the NK-FHP model is able to match the empirical moments in these papers remarkably well, while also explaining fluctuations in U.S. output, inflation, and interest rates better than other alternative models including a NK model incorporating sticky information.

After performing this external validation exercise, following Del Negro and Eusepi (2011) we modify our estimation procedure to include survey data on inflation expectations as an additional observable variable. We find that, overall, the parameter estimates of the model, including the length of agents' planning horizons, change little with the inclusion of inflation expectations in the estimation procedure. This result reflects that the model estimated without the inflation expectations data already fits that data reasonably well. We do find, however, changes in parameter estimates if the objective is to *exclusively* fit the inflation expectations data. In that case, estimates of the length of agents' planning are on the order of four quarters compared to estimates on the order of one quarter when the objective is to jointly fit the macroeconomic time series and inflation expectations data.

Our paper is related to earlier work that uses survey data on inflation expectations in the context of the estimation of DSGE models. Ormeno and Molnar (2015) perform a similar exercise to Del Negro and Eusepi (2011) to show that a NK model incorporating adaptive learning fits the inflation expectations data better than the rational expectations models of Del Negro and Eusepi (2011) though it shows little improvement in fit in terms of macro data. While we also estimate a DSGE model and incorporate survey evidence on inflation expectations into our analysis, our emphasis, unlike those papers, is on evaluating the model's ability to account for the predictability regressions of CG (2015) and KW (2021) and the predictability impulse responses of AHS (2020). Our paper is also related to papers such as Milani (2007) and Slobodyan and Wouters (2012) that estimate models with learning using macro data as well as Eusepi and Preston (2018) and Carvalho, Eusepi, Moench, and Preston (2023), which emphasize learning about long-run trends. A key difference between these papers and the finite-horizon approach used here is that expectation formation in these papers is backward looking while expectation formation under FHP has both a backward-looking and forward-looking component. We find that both components are important in accounting for the predictability results of CG (2015), KW (2021), and AHS (2020).

The rest of the paper proceeds as follows. The next section presents the properties of aggregate inflation when firms' set prices with finite horizon plans and the analytical results regarding the predictability of inflation forecasts under FHP. Section 3 describes the general equilibrium version of the model that we estimate. Section 4 discusses the estimation results of that model, including its fit of the predictability regressions of CG (2015) and KW (2021) and the predictability impulse responses of AHS (2020). It also compares the NK-FHP's empirical performance with alternative models including a NK model emphasizing sticky information. Section 5 concludes.

2 Finite Horizon Planning and Inflation Forecast Predictability

A key finding in the empirical literature using survey data is that inflation forecast errors are systematically predictable. CG (2015) emphasize this predictability by running regressions and showing that median inflation forecast errors in the SPF are correlated with forecast revisions. AHS (2020) show that the impulse response of the median inflation forecast in the SPF under reacts to aggregate shocks in the short run before over reacting later on. Building on CG (2015), KW (2021) find evidence of both an underreaction of the average forecast to new information but a simultaneous overreaction to recent data. In this section, we follow Woodford (2018) and assume firms, setting prices according to Calvo (1983) contracts, have finite planning horizons. We study the inflation forecasts of these firms and derive conditions under which they are consistent with the empirical results of CG (2015), AHS (2020), and KW (2021).

Finite horizon planning. Before discussing the economy's price-setting firms, we first define the expectations operator of an agent who is a finite-horizon planner. As discussed in Woodford (2018), such an agent making decisions at date t can only look forward and formulate plans that take into the model's relationships and all possible realizations of shocks occurring between periods t and $t+k$, where k denotes the length of an agent's planning horizon. Let $\mathbb{E}_t^k$ denote the subjective expectations of a finite-horizon planner. Then, for any endogenous variable in periods $t+k-j$, Z_{t+k-j}, with $j = 0, 1, 2, ..., k$ (i.e., j indexes the number of periods remaining within the planning horizon), the following relationship holds:

$$\mathbb{E}_t^k Z_{t+k-j} = E_t Z_{t+k-j}^j, \tag{1}$$

where E_t denotes the rational expectation (RE) operator conditional on time t information and the variable Z^j_{t+k-j} reflects the subjective expectations of a finite-horizon planner. Because an agent has a limited understanding of events outside of its planning horizons, its subjective expectations differ from rational expectations, and expression (1) provides a mapping between an agent's subjective expectations and rational expectations.

Price setting. The price-setting firms are monopolistic competitors whose prices are staggered á la Calvo (1983). When resetting their price, each firm is assumed to have a finite planning period of length k. As shown in Woodford (2018), under these assumptions, firms' price-setting behavior implies a log-linearized relationship for inflation given by:

$$\pi^j_\tau = \beta E_\tau \pi^{j-1}_{\tau+1} + \kappa y_\tau, \tag{2}$$

where $\tau = t + k - j$ denotes the planning period and $1 < j \leq k$. The variable π^j_τ denotes the (log-linearized) inflation rate implied by firms' plans in period τ. The parameter β is the discount rate, and the parameter κ is a function of the Calvo price-setting parameter, θ_p, and parameters that affect the link between firms' real marginal costs and aggregate output. The variable y_τ represents the (log-linearized) output gap, which is assumed to follow an AR(1) process:

$$y_t = \rho y_{t-1} + e_t \tag{3}$$

with $\rho \geq 0$. For now, we assume that the output gap evolves exogenously. This simplification allows us to derive analytical results, providing a better understanding of the implications that NK-FHP has for the predictability impulse responses of AHS (2020) and the predictability regressions of CG (2015) and KW (2021). Later, we modify this assumption, and allow the output gap to be endogenously determined in the context of the general equilibrium model that we estimate.

Equation (2) reflects the behavior of firms who have the opportunity to change their prices at date t and it holds in each period of those firms' planning horizons except the last period. Iterating forward on the expressions implied by equation (2) results in an expression that determines aggregate inflation:

$$\pi^k_t = \kappa E_t \sum_{i=0}^{k-1} \beta^i y_{t+i} + \beta^k E_t \pi^0_{t+k} \tag{4}$$

where $\pi^k_t = \pi_t$ denotes aggregate inflation (in log deviation from steady state). According to equation (4), aggregate inflation depends on the expected path of the output gap and on the expected inflation rate at the end of firms' planning horizons ($E_t \pi^0_{t+k}$). The NK Phillips curve under rational expectations arises as a special case: as $k \to \infty$, the planning horizon extends over a firm's infinite lifetime and inflation depends on the entire future path of the output gap.

Firms with the opportunity to reset their price at date t make a fully state contingent plan through $t+k$. They use their knowledge of the model's structural equations to do so. However, these firms use continuation value functions to assign value to events outside of their planning horizons (i.e., the longer-run from their viewpoint). These value functions affect the expected inflation rate at the end of firms' planning horizons (π^0_{t+k}). Formally, the (log-linearized) equilibrium condition associated with firms' pricing plans at the end of their horizon is:

$$\pi^0_{t+k} = \kappa y_{t+k} + \beta(1-\theta_p)v_{pt}, \tag{5}$$

where v_{pt} is the (log-linearized) *continuation value* to the plans of firms with the opportunity to reset their prices at date t and $1-\theta_p$ is the fraction of firms that have the opportunity to re-optimize their price at date t.

Learning. While firms are sophisticated in thinking about events within their planning horizon, they are less so when thinking about events further in the future. In this regard, we consider two different situations. In the no learning case, firms' beliefs about longer-run events (i.e., outside their planning horizon) are fixed at their steady state values so that $v_{pt} = 0\ \forall t$. Alternatively, we allow firms to learn and update their beliefs based on *past experience*. In this case, the value function v_{pt} evolves according to:

$$v_{pt+1} = (1-\gamma_p)v_{pt} + \gamma_p v^e_{pt}, \tag{6}$$

where v^e_{pt} is a firm's *new* estimate of its value function. The parameter γ_p determines how much weight they place on that new estimate and satisfies $0 < \gamma_p < 1$. The new estimate of the value function is determined by firms who can re-optimize their prices at time t, as v^e_{pt} is chosen as part of their FHP optimization problem. Woodford (2018) shows that in equilibrium v^e_{pt} satisfies:

$$v^e_{pt} = (1-\theta_p)^{-1}\pi^k_t. \tag{7}$$

According to equation (7), v^e_{pt} depends on aggregate inflation scaled by the fraction of firms who can re-optimize their prices at date t.[4] Combining equations (6) and (7), it follows that firms' beliefs about events outside their planning horizon (i.e., v_{pt}) depends on past realizations of inflation.

Inflation Dynamics. The equilibrium dynamics of inflation can be characterized analytically. Equations (4), (5), and (6) imply that aggregate inflation is the sum of a component that reflects a firm's future beliefs about the output gap over its finite horizon and a component that firm's beliefs about longer-run events outside of its planning horizon:

$$\pi^k_t = A(k)\kappa y_t + \beta^{k+1}(1-\theta_p)v_{pt}, \tag{8}$$

where the parameter $A(k) = \frac{1-(\beta\rho)^{k+1}}{1-\beta\rho}$.[5] Both the parameters affecting the response of inflation to changes in the ouput gap and changes in their longer-run beliefs depend on k, the length of a firm's planning horizon.

The dynamics of inflation under RE correspond to the case in which $k \to \infty$. In that case, a

[4]More specifically, to a first order approximation, Woodford (2018) shows $v^e_{pt} = p^{*k}_t$, where p^{*k}_t denotes the optimal contract price chosen by firms with an opportunity to reset their price at date t. Equation (7) then reflects the equilibrium relationship between aggregate inflation and the contract price: $\pi^k_t = (1-\theta_p)p^{*k}_t$.

[5]The term on the output gap reflects that firms know the process for y_t so that $\mathbb{E}^k_t y_{t+i} = E_t y_{t+i} = \rho^i y_t$.

firm's longer-run beliefs, v_{pt}, become irrelevant and $A(\infty) = \frac{1}{1-\beta\rho}$ so that inflation evolves according to

$$\pi_t^{RE} = \frac{\kappa}{1-\beta\rho} y_t.$$

Inflation dynamics under FHP expectations involves two deviations from the dynamics of inflation under RE. First, since $0 < 1 - (\beta\rho)^{k+1} \leq 1$, inflation in the FHP model is less responsive to fluctuations in the output gap. This muted responsiveness of inflation is a function of ρ and the length of the planning horizon. A shorter planning horizon or a more persistent shock imply a more muted response of inflation to movements in the output gap relative to the RE solution. The second deviation from inflation dynamics under RE is that firm's longer-run beliefs about inflation, as discussed above, depend on past inflation and thus inflation under FHP expectations displays an excess sensitivity to past inflation.

Forecasting. To understand the implications of FHP expectations for forecast predictability, we characterize a firm's one-step ahead forecast for inflation and the associated forecast error. A firm with a planning horizon of length $k > 0$ has a one-step ahead forecast given by:

$$\mathbb{E}_t^k \pi_{t+1} = \rho A(k-1)\kappa y_t + \beta^k (1-\theta_p) v_{pt} = [1-(\beta\rho)^k] \frac{\kappa\rho}{1-\beta\rho} y_t + \beta^k (1-\theta_p) v_{pt}. \tag{9}$$

Like inflation, a firm's one-step ahead forecast is sticky, since a firm's longer-run beliefs about inflation affect $\mathbb{E}_t^k \pi_{t+1}$ and these beliefs depend on lagged inflation.[6] This stickiness diminishes as $k \to \infty$. In that case, expression (9) converges to the forecast under rational expectations: $\mathbb{E}_t^\infty = E_t \pi_{t+1} = \frac{\kappa\rho}{1-\beta\rho} y_t$.

Under FHP expectations, a firm will make systematic forecast errors. To see this, define the one-step ahead forecast error under FHP as $\mathbb{F}_{t+1}^k \equiv \pi_{t+1}^k - \mathbb{E}_t^k \pi_{t+1}$. Using expressions (8) and (9), the one-step ahead forecast error evolves according to:

$$\mathbb{F}_{t+1}^k = \left[\beta^{k+1}\gamma_p A(k) + \rho(\beta\rho)^k\right] \kappa y_t - \beta^k \left[1 - \beta(1-\widetilde{\gamma}_p)\right] (1-\theta_p) v_{pt} + O_{t+1}. \tag{10}$$

where $\widetilde{\gamma}_p = \gamma_p(1-\beta^{k+1})$ and O_{t+1} is an omitted terms that depends on the innovation in the output gap at date $t+1$, e_{t+1}.[7]

Equation (10) is a key equation for determining the forecasting properties of inflation in the FHP model. Firms' forecast errors for inflation are the sum of an unpredictable component (O_{t+1}) and two predictable components. One of those predictable components relates to errors associated with firms underpredicting the responsiveness of inflation to movements in the output gap: In response to an increase changes in the output gap, the inflation forecast error rises because realized inflation responds more than expected inflation. Accordingly, firms' forecasts underreact to changes in the output gap. The other predictable component relates to changes in the value function governing firms' longer-run beliefs. Because $0 < \beta < 1$ and $0 < \widetilde{\gamma}_p < 1$, the forecast error falls in response

[6] As equation (9) highlights, the learning framework in Woodford (2018) uses the "anticipated utility" approach of Kreps (1998) and a firm's forecast of future inflation ignores the fact that v_{pt} will change over time.

[7] To obtain expression (10), note that $A(k) = A(k-1) + (\beta\rho)^k$.

to an increase in the value function, indicating that firms' forecasts overreact to changes in their longer-run beliefs. This overreaction reflects that a firm is closer to the end of its planning horizon when forming expectations of future inflation, making the sensitivity of the one-step ahead forecast to a firm's (continuation) value function greater than that of realized inflation.

2.1 Impulse Response Predictability

Equations (9) and (10) can be used to characterize the impulse responses of inflation forecasts and forecast errors to an innovation in e_t, allowing us to relate the model's implications to the empirical work of AHS (2020). Using data from the SPF, AHS (2020) compute the impulse response of the median respondent's inflation forecast and forecast error from the shock that maximizes the business cycle variation in inflation. Their results are striking, as they show there is a sign switch in the impulse response of the inflation forecast error: it underreacts before overreacting later on. We show that the FHP model is capable of generating this sign switch, and the following proposition established conditions under which it does so.

Proposition 1. *(IRFs of Inflation Forecasts and Forecast Errors). Let* $\frac{\partial \mathbb{E}^k_{t+i}\pi_{t+1+i}}{\partial e_t}$ *and* $\frac{\partial \mathbb{F}^k_{t+1+i}}{\partial e_t}$ *for* $i \geq 0$ *be the impulse response functions for a firm's one-step ahead inflation forecast and forecast error, respectively.*

1. *Without learning:* $\frac{\partial \mathbb{E}^k_{t+i}\pi_{t+1+i}}{\partial e_t} \geq 0$ *and* $\frac{\partial \mathbb{F}_{t+1+i}}{\partial e_t} \geq 0$, $\forall i \geq 0$ *and* $k > 0$.

2. *With learning: If* $\gamma_p \leq \frac{1-\rho}{1-\beta^{k+1}}$, *there is a threshold forecast horizon,* $i^\star$, *such that:*

   (a) $\frac{\partial \mathbb{E}^k_{t+i}\pi_{t+1+i}}{\partial e_t} \geq 0$ *for* $i \geq 0$,

   (b) $\frac{\partial \mathbb{F}_{t+1}}{\partial e_t} > 0$ *and* $\frac{\partial \mathbb{F}_{t+1+i}}{\partial e_t} < 0$ *for* $i \geq i^\star$,

Proof: See the appendix.

Proposition 1 indicates that in the model in which firms do not update their longer-run beliefs about inflation (i.e., no learning), both the inflation forecast and inflation forecast error respond positively to an innovation to the output gap. Accordingly, the impulse responses are characterized by a systematic underreaction — there is no flip in the sign of the impulse response function at any horizon. This underreaction reflects that without learning $v_{pt} = 0 \ \forall t$ so that movements in a firm's inflation forecast reflect only changes in the output gap. And, a firm with finite-planning horizon neglects changes in the output gap that occur outside its planning horizon, implying that its forecast underreacts to such changes. Thus, the FHP model without learning, similar to other behavioral models emphasizing cognitive discounting, can not account for the empirical evidence in AHS (2020).

While incorporating learning into the model may be a necessary condition to account for the evidence in AHS (2020), it is not a sufficient condition. For the forecast error to change signs from an underreaction to an overreaction, Proposition 1 also provides a sufficient condition that puts an upper bound on γ_p, the speed at which firms update their longer-run beliefs using past data. Focusing on $\beta \approx 1$, this condition only imposes a restriction on γ_p for shocks that are highly persistent (i.e., values of ρ close to 1.) For more moderate values of ρ and relatively short planning horizons, this condition is easily satisfied so that all values of γ_p between zero and one are consistent with the result in AHS (2020). Intuitively, if the output-gap shock is not very persistent, the impulse responses of the forecast error in later periods will be mostly determined by changes in a firm's longer-run inflation beliefs. With the effect of the value function more important for the impulse responses in later periods, a firm's forecast will eventually display an overreaction, since firms' forecasts are excessively sensitive to changes in v_{pt}.

Proposition 1 establishes the existence of a threshold horizon at which the impulse response of the inflation forecast error in the FHP model switches signs. Figure 1 illustrates how this threshold horizon depends on the model's structural parameters. The upper panel plots the threshold horizon ($i^\star$) as a function of a firm's planning horizon (k) for two different values of γ_p. The threshold horizon is higher for shorter planning horizons. For instance, when $\gamma_p = 0.3$ and $k = 12$, the impulse response switches from an underreaction to an overreaction after 10 quarters, while when $k = 1$, it takes 16 quarters. A shorter planning horizon has this effect because, all else equal, it makes the magnitude of the underreaction of the forecast to the change in the output gap larger, delaying the eventual overshoot. Figure 1 also highlights that a higher value of γ_p or lower value of ρ results in an earlier overreaction of a firm's inflation forecast. These results reflect that a higher value of γ_p speeds up the learning process so that the model's eventual overreaction occurs earlier. A lower value of ρ has a similar effect, since it implies a smaller and less persistent underreaction of a firm's inflation forecast. Overall, proposition 1 indicates that the sign switch in the impulse response function of the forecast error is a robust, qualitative feature of the FHP model with learning. Later, we conduct a more rigorous empirical evaluation of the FHP model, pinning its parameters down using macroeconomic time series while using the impulse response of AHS (2020) as well as the predictability regressions of CG (2015) and KW (2021) as additional tests regarding the nature of expectation formation embedded in finite horizon planning.

2.2 Inflation Predictability Regressions

To discriminate across alternative models of expectation formation, CG (2015) and KW (2021) emphasize the predictability of forecast errors from regressions using survey data on expectations. CG (2015) regress the median forecast error of inflation on the median forecast revision and show that there is a positive correlation between the forecast error and forecast revision, implying an underreaction of forecasts to new information. KW (2021) also emphasize the underreaction of forecasts to new information but they provide additional evidence that forecasts in survey data also involve an overreaction to recent data. In particular, they regress average forecast errors from survey data on the forecasted variable and show a negative correlation between the forecast error and the forecasted variable, implying an overreaction to recent data. They argue that a wide class of models of expectation formation are unable to account for this simultaneous underreaction to new information implied by the CG (2015) regression and overreaction to recent data implied

Figure 1: Delayed Overreaction of Inflation Forecasts in the FHP Model

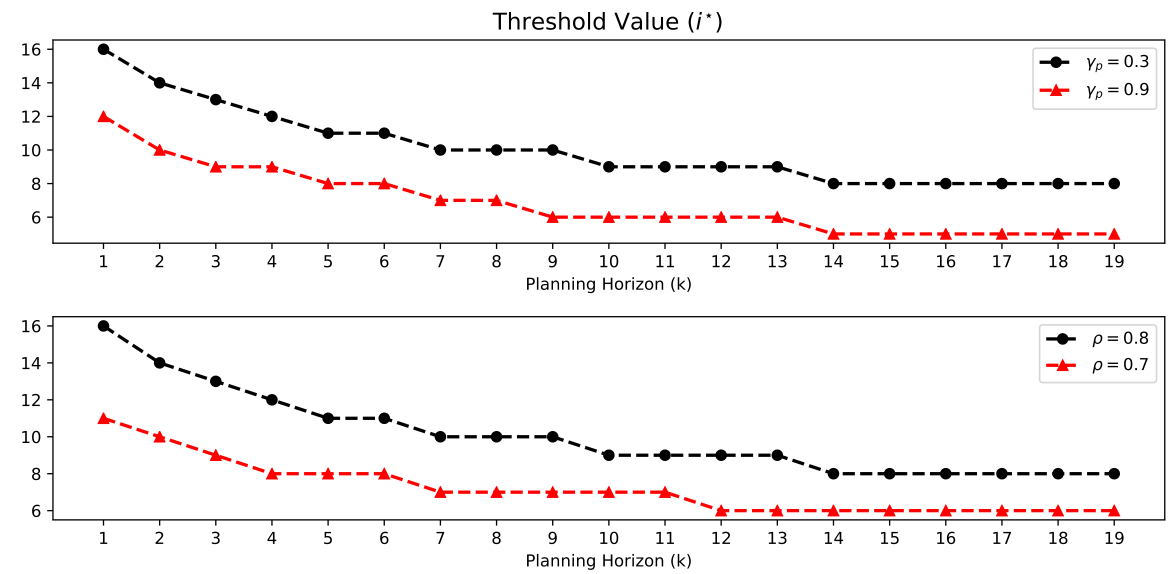

Note: The figure shows the threshold date at which the impulse response of the inflation forecast in the FHP model switches from an underreaction to an overreaction.

by their regression. In this section we investigate the implications of FHP expectations for the predictability regressions of CG (2015) and KW (2021).

To examine the implication of the FHP model for the predictability result of CG (2015), a firm's inflation forecast revision is defined as $\mathbb{R}_t^k = \left[\mathbb{E}_t^k - \mathbb{E}_{t-1}^k\right]\pi_{t+1}$. A firm's forecast at $t-1$ satisfies:

$$\mathbb{E}_{t-1}^k \pi_{t+1} = \rho^2 A(k-2)\kappa y_{t-1} + \beta^{k-1}(1-\theta_p)v_{pt-1} \tag{11}$$

for $k > 1$. At time $t-1$ a firm's expectation for π_{t+1} differs from its expectations at time t because it has less information than at time t. In addition, a firm is looking an extra period ahead and is closer to the end of its planning horizon. Because, it is close to the end of its planning horizon, its forecast of π_{t+1} at time $t-1$ puts more weight on a firm's value function and less weight on the output gap than a firm's one-step ahead forecast. Proposition 2 characterizes the relationship between inflation forecast errors and revisions for a firm with FHP expectations.

Proposition 2. *(Forecast Error and Revision Correlation). Let* $\beta_{CG} = \frac{cov(\mathbb{R}_t^k,\, \mathbb{F}_{t+1}^k)}{var(\mathbb{R}_t^k)}$ *denote the univariate regression coefficient from regressing the one-step ahead forecast error on the forecast revision in the FHP model.*

1. *Without learning: If* $\rho > 0$*, then* $\beta_{CG} > 0$*, for any finite planning horizon* $k > 0$*.*
2. *With learning: If* $\rho = 0$ *and* $\gamma_p < \frac{1-\beta}{1-\beta^{k+1}}$*, then* $\beta_{CG} > 0$*.*

Proof: See the appendix.

Proposition 2 shows that for positively correlated shocks, the FHP model without learning always results in a positive value of β_{CG}, in line with the empirical estimates of CG. Without learning, a persistent increase in the output gap ($\rho > 0$) leads firms with FHP expectations to revise up their forecasts of inflation. From equation (10), it follows that their inflation forecast errors rise persistently, and there is a persistent underreaction of the inflation forecast to new information.

With learning, the dynamics of inflation are richer (and more complex) and the correlation between the forecast error and forecast revision can be either positive or negative depending on the length of a firm's planning horizon (k), how quickly firms update their longer-run beliefs (γ_p), and the persistence of the shock (ρ). When the shocks are uncorrelated, Proposition 2 indicates that the FHP model implies $\beta_{CG} > 0$ if firms do not update their longer-run beliefs too quickly. Specifically, there is an upper bound on γ_p that grows increasingly tight as the length of the planning horizon increases. For instance, for short-horizon planning ($k = 1$) with learning occurring at relatively sluggish rate, (i.e., $\gamma_p < \frac{1}{1+\beta} < 0.5$), the model generates $\beta_{CG} > 0$.

Our final proposition considers the regression statistic of KW (2021). While KW (2021) mainly focus on forecasts of output growth, they show that several survey measures of inflation expectations, including average forecasts of consumer price inflation from the SPF, display a negative correlation between the forecasted variable and survey respondents' forecast errors. Propositon 3 characterizes the relationship between inflation and forecast errors under FHP expectations.

Proposition 3. *(Forecast Error and Inflation Correlation). Let* $\beta_{KW} = \frac{cov(\pi_t^k,\, \mathbb{F}_{t+h}^k)}{var(\pi_t^k)}$ *denote the univariate regression coefficient from regressing the h-step ahead forecast error on inflation in the FHP model.*

1. *Without learning: If* $\rho > 0$*, then* $\beta_{KW} > 0$*, for any finite planning horizon* $k \geq h \geq 1$*.*

2. *With learning: If* $\rho = 0$ *and* $k \geq h > 1$*, then* $\beta_{KW} < 0$ *if and only if:*

$$\frac{\beta^{k+1}(1-\widetilde{\gamma}_p)}{(1-\beta^{k+1})(2-\widetilde{\gamma}_p)}\left[\beta^{-h}(1-\widetilde{\gamma}_p)^{-h} - 1\right] > 1.$$

Proof: See the appendix.

Proposition 3 indicates that the correlation between inflation and the forecast error is positive in the FHP model without learning. However, when the FHP model includes learning, the correlation can be negative for forecasts beyond a quarter (i.e., $h > 1$). For uncorrelated shocks, the condition in proposition 3 indicates that for a fixed value of k, a longer forecast horizon or a larger value of γ are more likely to imply that $\beta_{KW} < 0$. This negative correlation is possible because firms learn and update their longer-run beliefs about events outside their planning horizons by extrapolating from past data on inflation. Such behavior implies that a firm's forecasts can overreact to movements in inflation, and faster learning in which firms' beliefs depend more on recent inflation data exacerbates this overreaction.

Table 1 shows some numerical results for β_{CG} and β_{KW}, varying the persistence of the shock, the length of firms' planning horizons, and the learning speed. The table highlights that $\beta_{CG} > 0$ for firms' planning horizons of 3 to 6 quarters and that β_{CG} is positive for a wide range of values of ρ and γ_p. In particular, when $\rho = 0.9$, β_{CG} is positive even for high values of γ_p. The table also highlights that β_{KW} can be positive or negative depending on how firms update their beliefs about their value functions. When $\gamma_p = 0.9$, the sign of β_{KW} is negative, as firm's beliefs about their value functions responds relatively quickly to past changes in inflation. However, when $\gamma_p = 0.25$, firms' beliefs depend relatively more on inflation in the distant past, and the sign of β_{KW} is positive in those cases. Overall, we conclude that the FHP model can be qualitatively consistent with the empirical evidence of CG (2015) and KW (2021) and below we investigate this question further in the context of an estimated, general equilibrium model.

Table 1: Predictability Regression Results For FHP Model

	$\rho = 0.25$		$\rho = 0.9$	
	β_{CG}	β_{KW}	β_{CG}	β_{KW}
$k = 6$				
$\gamma_p = 0.25$	0.59	0.03	0.52	-0.01
$\gamma_p = 0.9$	0.43	-0.07	0.46	-0.02
$k = 3$				
$\gamma_p = 0.25$	0.59	0.01	0.69	0.0
$\gamma_p = 0.9$	0.44	-0.05	0.58	-0.01

Note: For β_{CG}, entries report population coefficient from a regression of one-step ahead inflation forecast errors on forecast revision. For β_{KW}, entries report population coefficient from a regression of three-quarters ahead inflation forecast errors on inflation. The values of β and κ were set to 0.99 and 0.05, respectively.

3 Dynamic General Equilibrium Model

The previous section establishes that inflation expectations implied by the NK-FHP model are broadly consistent with key stylized facts that have emerged from the empirical literature using survey data on expectations. With this result established, we next turn to investigating whether the NK-FHP model in Woodford (2018) is jointly consistent with the survey data on inflation expectations as well as the fluctuations in output, inflation, and interest rates in U.S. data. To address this question, this section extends the analysis in the previous section to include households with finite planning horizons and monetary policy that is specified to follow an interest rate rule.

As in the previous section, it is assumed that all agents have the same planning horizon of length k. The model's inflation dynamics are determined from similar expressions to those shown in equations (4) and (5) except that now the output gap no longer follows an exogenous process but is endogenously determined. With an endogenously determined output gap, the New Keynesian

Phillips curve becomes:

$$\pi_t^k = \kappa E_t \sum_{i=0}^{k} \beta^i (y_{t+i}^{k-i} - y_{t+i}^{\star}) + \beta^{k+1}(1-\theta_p)v_{pt} \tag{12}$$

where $y_{t+i}^{\star}$ is an exogenous shock to aggregate supply and y_{t+i}^{k-i} is a firm's beliefs about the output gap in period $t+i$ which as discussed below is determined by the level of household expenditures.

Households. There is a large number of identical, infinitely-lived households. Each household makes a consumption\savings decision but like the economy's firms only has the ability to plan k periods ahead. Households also supply their labor services to firms in a perfectly competitive labor market. As shown in the appendix, optimization by households gives rise to a (log-linearized) relationship that relates household expenditures at time t to future interest rates that occur over their planning horizon:

$$y_t^k = -\sigma E_t \sum_{i=0}^{k-1} \left(i_{t+i}^{k-i} - \pi_{t+i+1}^{k-i-1} - r_{t+i}^{\star} \right) + E_t y_{t+k}^0 \tag{13}$$

where y_t^k are a household's demand for expenditures at time t and i_{t+i}^j denotes a household's beliefs about the setting of the policy rate in period $t+i$. The parameter σ is the inverse of a household's relative risk aversion, and the variable $r_t^{\star}$ is an exogenous shock to preferences.[8] This shock as well as the supply shock, $y_t^{\star}$, are assumed to follow AR(1) processes with persistence parameters, ρ_y, for the supply shock, and ρ_r for the demand shock. A household's expenditures at time t, y_t^k, also depend on its plans for expenditures at the end of their planning horizion, y_{t+k}^0, which are given by:

$$y_{t+k}^0 = -\sigma \left(i_{t+k}^0 - r_{t+k}^{\star} \right) + v_{ht} \tag{14}$$

where the variable v_{ht} is the value that household assigns to events that occur outside of their planning horizons and reflects that households, like firms, have a limited ability to understand and evaluate situations that occur in the distant future.

Similar to firms, households update v_{ht} based on past events and do so in a way that is consistent with their optimal finite-horizon plan. In particular, when they decide on their expenditures, y_{t+k}^k, they form a new estimate of their value function, v_{ht}^e, and use it to update their beliefs according to:

$$v_{ht+1} = (1-\gamma)v_{ht} + \gamma v_{ht}^e, \tag{15}$$

where $0 < \gamma < 1$ determines how much weight they place on their new estimate. This new estimate is consistent with household optimization and as shown in the appendix reflects outcomes for both expenditures and inflation:

$$v_{ht}^e = y_t^k + \sigma \pi_t^k \tag{16}$$

[8] As shown in the appendix, this shock affects a household's discount factor and differs from the preference shock used in Woodford (2018).

Substituting equation (16) into equation (15), it follows that v_{ht} depends on past realizations of household expenditures and inflation. Accordingly, a household's longer-run beliefs are determined in a backward-looking manner and because they depend in particular on lagged expenditures, these longer-run beliefs can give rise to persistence in household expenditures.

Trend-Cycle Decomposition. As discussed in Woodford (2018), an interesting feature of the NK-FHP model is that its variables can be decomposed output into a "cyclical" component— reflecting the effect of the model's shocks— and a "trend" component—reflecting changes in household and firm beliefs' about their longer-run continuation values. Specifically, the "trend" components (denoted by $\bar{\pi}_t^j$, $\bar{y}_t^j$, and $\bar{i}_t^j$, respectively, for $j = 0, 1, ..., k$) are defined by abstracting from the effect of shocks in equations (12) and (13). Accordingly, the evolution of these trends can be written as functions of the continuation values of households and firms decisions, v_{ht} and v_{pt}:

$$\bar{\pi}_t^k = \kappa \sum_{i=0}^{k} \beta^i \bar{y}_t^{k-i} + \beta^{k+1}(1 - \theta_p) v_{pt}$$

$$\bar{y}_t^k = -\sigma \left[\sum_{i=0}^{k} \bar{i}_t^{k-i} - \sum_{i=0}^{k-1} \bar{\pi}_t^{k-i-1} \right] + v_{ht}$$

where $\left\{\bar{\pi}_t^k, \bar{y}_t^k\right\}_{j=0}^k$ denote the effect of the continuation value functions on the plans of households and firms. (For these variables, we denote the effects of the v_{ht} and v_{pt} on household and firm plans with only a t subscript since v_{ht} and v_{pt} are fixed at time t.) We use the trend variables to help characterize the model's dynamics and understand the role of household and firms' longer-run beliefs in generating endogenous persistence and influencing inflation expectations.

Monetary Policy. Monetary policy at each date t is specified as an interest-rate rule of the form:

$$i_t^k = \bar{i}_t^k + \phi_\pi(\pi_t - \bar{\pi}_t^k) + \phi_y(y_t - \bar{y}_t^k) + i_t^\star \tag{17}$$

where $i_t^\star$ is an exogenous shock to the rule assumed to follow a first-order autoregressive process with persistence parameter, ρ_i. We assume that the intercept of the policy rule depends on the evolution of the model's trends. In particular, $\bar{i}_\tau^j$ is given by:

$$\bar{i}_\tau^j = \bar{\phi}_\pi \bar{\pi}_\tau^j + \bar{\phi}_y \bar{y}_\tau^j \tag{18}$$

The time-varying intercept in the interest rate rule is intended to capture two aspects of monetary policy. First, it acknowledges that policymakers do not necessarily view the "equilibrium" or longer-run real interest rate as a constant.[9] Second, it also allows for the possibility that policymakers may respond more aggressively to persistent deviations of inflation from their inflation target, as captured by $\bar{\pi}_t^k$, than they do to temporary deviations. In that case, $\bar{\phi}_\pi > \bar{\phi}$, and as shown in our empirical analysis in Gust, Herbst, and López-Salido (2022), we find that such a monetary

[9] This formulation is consistent with policymakers' efforts to inform their decisions distinguishing trend factors—such as demographic or productivity changes—from cyclical variations in output and inflation. It implicitly assumes that policymakers are no better at separating trend from cycle as the private sector.

policy response fits the data substantially better than a rule in which monetary policy responds equi-proportionately to cyclical and trend inflation.

4 Empirical Analysis

In this section, we assess the dynamic, general equilibrium model's ability to to match the impulse response predictability results of AHS (2020) and predictability regressions of CG (2015) and KW (2021). The model and methodology closely follow Gust, Herbst, and López-Salido (2022), where we estimated the NK-FHP model employing a Bayesian, full-information likelihood-based approach using U.S. data on output growth, inflation, and interest rates.[10] Our estimation strategy does not employ information on inflation expectations or the predictability statistics and as discussed below we use predictive checks to evaluate the model's performance on the predictability statistics emphasized by AHS (2020), CG (2015), and KW (2021).

For this assessment, we use plausible parameter configurations from the posterior distribution of the NK-FHP model. The use of a full-information estimation as the basis for assessing these moments is a bit of a departure from the previous literature, which relies on regression analysis or partially identified VAR models. The full information strategy employed here is an attractive approach for a number of reasons. First, as the propositions in Section 2 indicate, the consistency of the NK-FHP model with particular stylized facts of inflation expectations depends not only on parameters directly governing finite horizon planning, but additional structural parameters auxiliary to the planning horizon. The presence of these auxiliary parameters can substantially complicate the estimation of the model. By using the full posterior distribution from the model, we incorporate the mostly likely values of these auxilary parameters into our evaluation. The methodology here also ensures that the assessment of particular moments related to inflation and inflation expectations is done conditional on parameterizations that also can rationalize the realized time series of output growth, inflation, and interest rates. Finally, our use of a fully specified model allows us to examine the behavior of these moments conditional on specific structural shocks. This is in important because—as Section 2 highlights—the persistence of the shock is often critical for determining the values of β_{CG} and β_{KW} and the impulse responses of inflation forecasts to shocks.

The NK-FHP model described in Section 3 is estimated using data on output growth, GDP deflator inflation, and the federal funds rate in the United States from 1967-2007. The priors and computational strategy for eliciting the posterior distribution of the parameters, $p(\theta|Y)$, is extremely similar to Gust, Herbst, and López-Salido (2022). Thus, we relegate most of the estimation details to the appendix.

The NK-FHP model in this section uses a planning horizon $k = 1$. Consistent with the results in Gust, Herbst, and López-Salido (2022), this value maximizes the overall fit—measured using log marginal data densities—of the model in Section 3. That said, our results are broadly consistent for other values of k consistent with a limited planning horizon. Table 2 describes key features of

[10] A minor difference in the model described in Section 3 from the one in Gust, Herbst, and López-Salido (2022) is that the demand shock is specified slightly differently. This change has little effect on the estimation results.

the posterior distribution. The posterior mean estimates of the learning rates, at 0.47 and 0.20 for the households and firms, respectively, indicating that firms weight recent data less in forming their beliefs about events outside their planning horizons than households. Given our focus on inflation expectations, the speed at which price-setters' update their beliefs will be particularly relevant for the predictability statistics that we study. The monetary policy rule has the same features as in Gust, Herbst, and López-Salido (2022): it displays a strong response to trend inflation and cyclical output but essentially no response to trend output. The demand shock and monetary policy shocks are estimated to be highly persistent, while the posterior mean for the AR coefficient for the supply shock is only 0.45. As discussed later, the difference in the persistence of these shocks will be important in determining whether the model is consistent with the evidence regarding inflation forecast predictability.

Table 2: NK-FHP($k = 1$) Model: Selected Posterior Statistics

	Description	Mean	[0, 95
γ	Household learning rate	0.47	[0.30, 0.64]
γ_p	Firm learning rate	0.20	[0.13, 0.29]
κ	Slope of the Phillips curve	0.03	[0.02, 0.04]
σ	Coef. Relative. Risk Aversion	2.68	[1.94, 3.52]
ϕ_π	Int. rule response to $\tilde{\pi}_t$	0.98	[0.72, 1.27]
ϕ_y	Int. rule response to $\tilde{y}_t$	0.89	[0.59, 1.29]
$\overline{\phi}_\pi$	Int. rule response to $\overline{\pi}_t$	1.84	[1.49, 2.23]
$\overline{\phi}_y$	Int. rule response to $\overline{y}_t$	0.13	[0.04, 0.25]
ρ_ξ	AR coeff. for demand shock	0.90	[0.84, 0.96]
ρ_i	AR coeff. for monetary policy shock	0.95	[0.90, 0.98]
ρ_y	AR coeff. for supply shock	0.45	[0.30, 0.60]

Note: The table shows estimates of the posterior means, 5th, and 95th percentiles of the model parameters computed from output of the SMC sampler. See appendix for details.

4.1 Predictive checks for assessing inflation expectation predictability

While the model has been estimated to jointly account for fluctuations in output growth, inflation, and interest rates, the estimation strategy does not use information on inflation expectations or the predictability statistics emphasized by AHS (2020), CG (2015), and KW (2021). Thus, an important additional check on the model's empirical fit is its ability to account for these statistics. To investigate this question, we use the framework of *predictive checks*. These checks involves comparing some statistic or moment from the data to the predictive distribution of that statistic under a given model. Here we use the posterior distribution of the estimated NK-FHP model. Let Y^{DSGE} denote the set of observables used to estimate the DSGE model–output growth, inflation, and interest rates–and let Y be an expanded set of observables which includes SPF consensus inflation expectations data. Formally, let $\mathcal{S}(Y)$ be some statistic of this data, where $\mathcal{S}(Y)$ can be a scalar—like the regression coefficient as in CG (2015)—or a vector—like the impulse responses of AHS (2020). For a given model, $\mathcal{M}$, we can draw from the posterior distribution of parameter estimates to a obtain a simulated counterpart to Y—called $\widetilde{Y}$—and compute the predictive distribution for $\widetilde{Y}$ as

$$p(\widetilde{Y}|\mathcal{M}) = \int p(\widetilde{Y}|\theta, \mathcal{M})p(\theta|Y^{DSGE}, \mathcal{M})d\theta. \tag{19}$$

Using (19), one can compare where the observed statistic $\mathcal{S}(Y)$ lies in the predictive distribution for $S(\widetilde{Y}|\mathcal{M})$. If $\mathcal{S}(Y)$ lies in the tail of the predictive distribution for a particular model, the model is said to be deficient along this dimension of the data. To compute $\mathcal{S}(Y)$ for the predictability IRF of AHS (2020), we estimate a VAR(4) on the observed data and construct impulse responses identifying a shock, as they do, that maximizes the forecast error variance of inflation over the medium term. We follow the same approach to compute the model analogues of these impulse responses from simulated data. For the predictability regression of CG (2015), we run regressions of inflation forecast errors on forecast revisions using survey data and model simulated data. For the predictability regression of KW (2021), we run regressions of inflation forecast errors on lagged inflation. Algorithm 1 provides more details for these posterior predictive checks. Following CG (2015) and others, we use the mean forecast for four-quarter-ahead (GDP deflator) inflation expectations ("Expected Inflation") from the Survey of Professional Forecasters (SPF) as the actual inflation expectations data.[11]

[11] In making their decisions, firms do not need to forecast inflation outside of their planning horizons. Accordingly, we need to make an additional assumption when a firm's forecast horizon exceeds its planning horizon, which is the case since we are forecasting inflation four quarters ahead with $k = 1$. In that case, we assume a firm uses its beliefs at the end of its planning horizon to make its forecast, taking into account its knowledge regarding the persistence of shocks. The appendix provides more details regarding this assumption.

Algorithm 1 Predictive Checks

For $i = 1, \ldots, N$:

1. *Construct* $\widetilde{Y}$. Draw $\theta^i \sim p(\theta|Y^{DSGE})$, and simulate a single trajectory of $\tilde{Y} = \{\Delta y_t, \pi_t, i_t, \mathbb{E}_t^k[\pi_{t+4}^A], \mathbb{E}_{t-1}^k[\pi_{t+4}^A]\}_{t=1}^T$, where $T \in \{168, 3000\}$, either the length of the actual observables ("finite sample") or large sample which eliminates sampling uncertainty ("population"). The variable $\pi_t^A = \pi_t + \pi_{t-1} + \pi_{t-2} + \pi_{t-3}$ is the four-quarter inflation rate.

2. *Construct* $\mathcal{S}_{AHS}(\tilde{Y}|\mathcal{M})$ as the point estimates of the impulse response coefficents of π_t^A and $\mathbb{E}_{t-4}^k[\pi_t^A]$ to an AHS "inflation" shock in a VAR(4) model for $\left[\Delta y_t, \pi_t, i_t, \mathbb{E}_t^k[\pi_{t+4}^A]\right]$. The VAR's inflation shock is identified as the shock that maximizes the variance in inflation over frequencies associated with periods of length 32 to 6 quarters.

3. *Construct* $\mathcal{S}_{CG}(\tilde{Y}|\mathcal{M})$. Estimate the regression model,

$$\pi_t^A - \mathbb{E}_{t-4}^k[\pi_t^A] = \alpha + \beta_{CG}\left(\mathbb{E}_{t-4}^k[\pi_t^A] - \mathbb{E}_{t-5}^k[\pi_t^A]\right) + u_t. \tag{20}$$

   Store the OLS point estimate of $\widehat{\beta}_{CG}$.

4. *Construct* $\mathcal{S}_{KW}(\tilde{Y}|\mathcal{M})$. Estimate the regression model,

$$\pi_t^A - \mathbb{E}_{t-4}^k[\pi_t^A] = \alpha + \beta_{KW}\pi_{t-4}^A + u_t. \tag{21}$$

   Store the OLS point estimate of $\widehat{\beta}_{KW}$.

Notes: We use $N = 200$ draws from the posterior. Further details on the identification of the AHS inflation shock used to construct $\mathcal{S}_{AHS}(\tilde{Y}|\mathcal{M})$ are available in the Appendix.

AHS (2020). The solid black lines in Figure 2 display the impulse responses of inflation, the inflation forecast, and the inflation forecast error for the inflation shock using the identification scheme of AHS on the actual data. The blue lines show the (pointwise) mean estimates using simulated data from the NK-FHP model to compute these impulse responses, and the shaded blue regions correspond to the 90 percent pointwise credible ranges. The mean impulse responses from the NK-FHP model for inflation and the forecast of inflation track those in the data quite well. In both the model and the data, the response of the inflation forecast is more muted than actual inflation on impact so that the forecast error rises on impact. This underreaction of the forecast lasts about two and half years, on average, in the model, and a little longer in the data. This underreaction is then followed by a persistent overreaction in both the model and the data. Accordingly, the NK-FHP model captures well the result in AHS that the inflation forecast underreacts initially and then overreacts later on in response to aggregate shocks.

To understand this result, it is important to realize that the AHS inflation shock is an amalgamation of the model's three structural shocks, in particular the supply and monetary policy shocks. This reflects that these shocks explain most of the variation in inflation at business cycle frequencies (i.e., those with periods between 6 and 32 quarters), the frequency band used to identify the shock. Figure 4 shows the fraction of variance explained by each shock as function of the frequency of fluctions, where these frequencies have been mapped to periods. For short periods, the fluctuations in inflation are mostly explained by supply shocks. As the period grows longer, the importance of the monetary policy shock increases. As shown in Table 2, the supply shock is less persistent than the model's other two shocks. Consistent with discussion in Section 2, the initial underreaction of agents' forecast of inflation is also shorter for the supply shock, given that it is less persistent than the other shocks. Accordingly, the response of inflation to the supply shock in later periods to a greater extent than for the model's other shocks reflects the overreaction of inflation to movements in longer-run beliefs (i.e., the continuation value-functions of price-setting firms). Table 3 highlights this property, as the probability of having the inflation forecast underreact and then overreact within the first 40 quarters after the supply shock is over 90 percent. In contrast, those probabilities are only 52 percent and 68 percent for the monetary policy shock and demand shock, respectively. Moreover, as implied by the mean value of i^*, the sign flip in the forecast error occurs about 10 quarters after the supply shock, on average, which is notably sooner than for the model's other two shocks.

Figure 2: Impulse Response to an AHS "Inflation" Shock

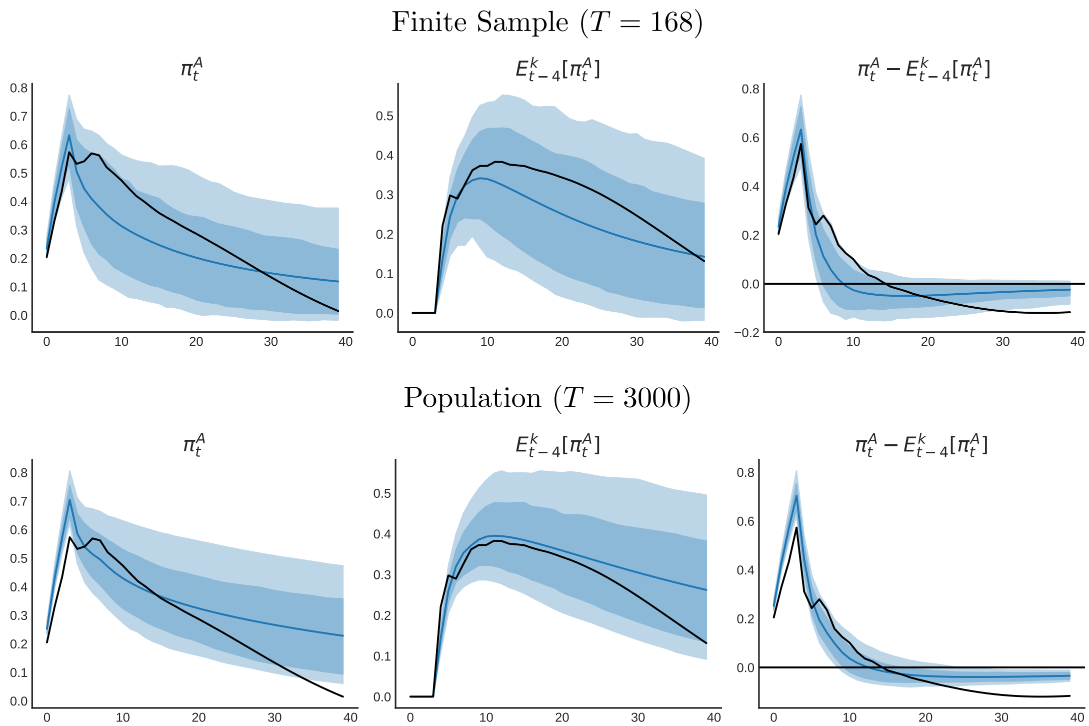

Notes: The figure shows the impulse response of inflation, the inflation forecast, and the inflation forecast error from the AHS-style VAR in the NK-FHP model. The solid blue line denotes the (pointwise) mean across the predictive checks, while the shaded regions denote the ninety (light blue) and sixty-eight (dark blue) percent bands (across the means of the predictive checks). The black line correspond to the impulse responses constructed using the actual data. The top row corresponds to exercise using trajectories of length $T = 168$ and the bottom row $T = 3000$.

Figure 4: Frequency-based Variance Decomposition for Inflation

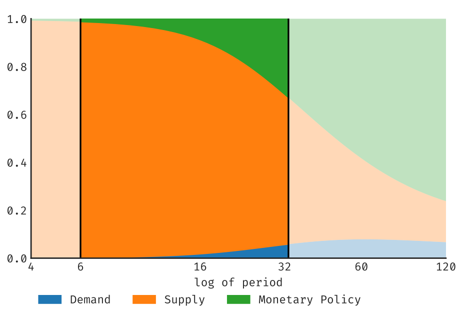

Table 3: AHS (2020) Sign Switch Properties

Shock	$\mathbb{P}(i^* < 40)$	$E[i^*\|i^* < 40]$	$\sqrt{\mathbb{V}[i^*\|i^* < 40]}$
AHS	0.89	10.09	4.49
Supply	0.93	9.52	3.50
Monetary Policy	0.52	30.26	8.43
Demand	0.68	22.83	5.44

Notes: The figure displays the fraction of variance of inflation at particular frequency attributable to demand (blue), supply (orange), and monetary policy (green) for different horizons, evaluated at the posterior mean parameter values. The table displays the "sign switching" properties of the AHS shock and the three structural shocks in the model computed using posterior predictive checks.

CG (2015). Turning to the predictability regression of CG (2015), Figure 5 shows the NK-FHP model's implications for this statistic. The upper left panel compares the distribution of model estimates for $\hat{\beta}_{CG}$ to the point estimate from regressing the inflation forecast error on the forecast revision using the survey data. The point estimate using the survey data is slightly greater than one and is close to the center of the posterior predictive distribution of the NK-FHP model. The positive relationship between inflation forecast errors and forecast revisions reflects the initial underreaction in response to shocks by FHP agents. As shown in Figure 2, because of this underreaction, the impulse reponse of the forecast error increases on impact, reflecting that FHP agents revise up their forecasts but not as much as actual inflation increases. Accordingly, an underreaction is associated with positive co-movement between forecast revisions and forecast errors. While this underreaction of the forecast eventually dissipates and turns into an overreaction, Figure 2 shows that the overreaction that occurs later on is small relative to the earlier underreaction of FHP agents' inflation forecast and thus is relatively less important in affecting $\hat{\beta}_{CG}$.

The model's distribution of $\hat{\beta}_{CG}$ in the upper left panel reflects the average effects of all three of the model's shocks and the remaining panels show these distributions conditional on each shock. Conditional on only demand and monetary policy shocks, the distribution of model estimates lies above the point estimate in the data. As shown in Table 2, these two shocks are highly persistent so that the initial underreaction that occurs in response to these two shocks lasts a long time, generating a stronger positive relationship between FHP agents' forecast errors and forecast revisions. As noted earlier, the supply shock is considerably less persistent than these two shocks so that the initial underreaction is less persistent and the overreaction that occurs later on in response to supply shocks is relatively more important for this shock. Accordingly, the distribution conditional on this shock lies to the left of the model's other two shocks.

Figure 5: Distribution of CG (2015) coefficients

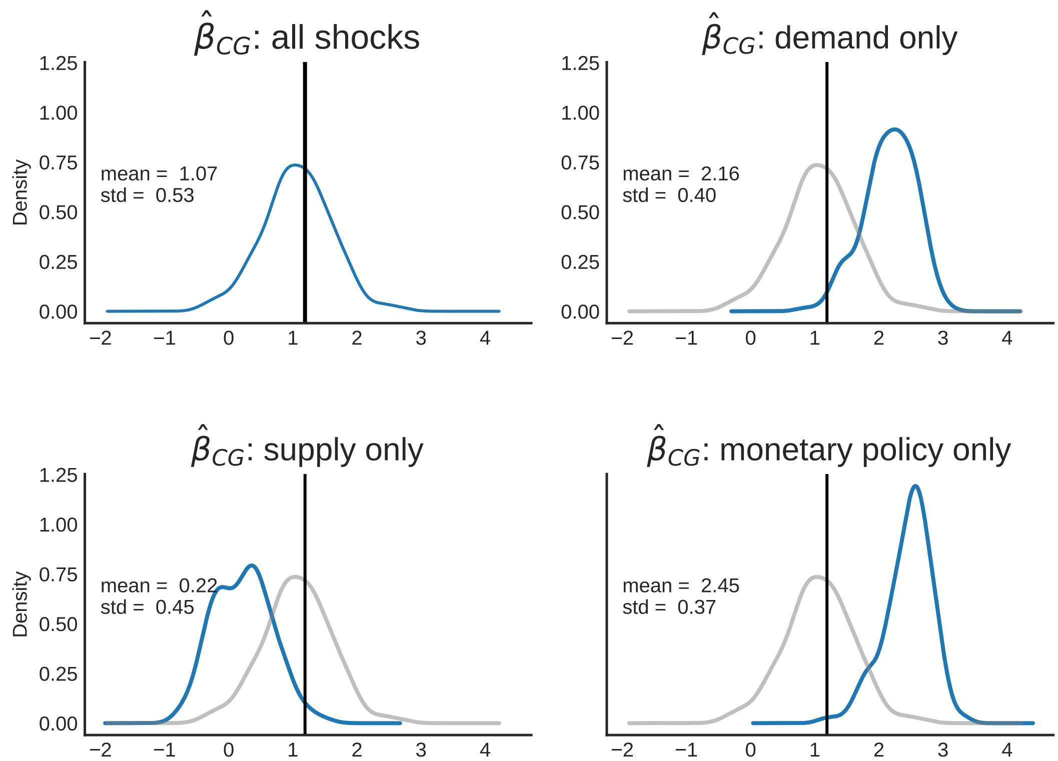

Note: The figure shows the predictive densities for $\hat{\beta}_{CG}$ simulated using all the shocks (upper left panel), and conditional only on demand (upper right panel), supply (bottom left panel), and monetary policy (bottom right panel) shocks. The black vertical line indicates the point estimate using the SPF data.

KW (2021). Figure 6 shows the predictive distributions of $\widehat{\beta}_{KW}$. The conditional on all of the shocks, the mean is about -0.03, and about 70 percent of the simulations feature a $\widehat{\beta}_{KW} < 0$. The posterior mean conditional on supply shocks is about -0.06, while it is about zero for the demand and monetary policy shocks. The mean conditional on supply shock is lower than for the other two shocks, because this lower persistence of this shock implies that the overreaction of inflation forecasts occurs earlier and relatively sooner than for the other shocks, as highlighted in Table 3.

Figure 6: Distribution of KW coefficients

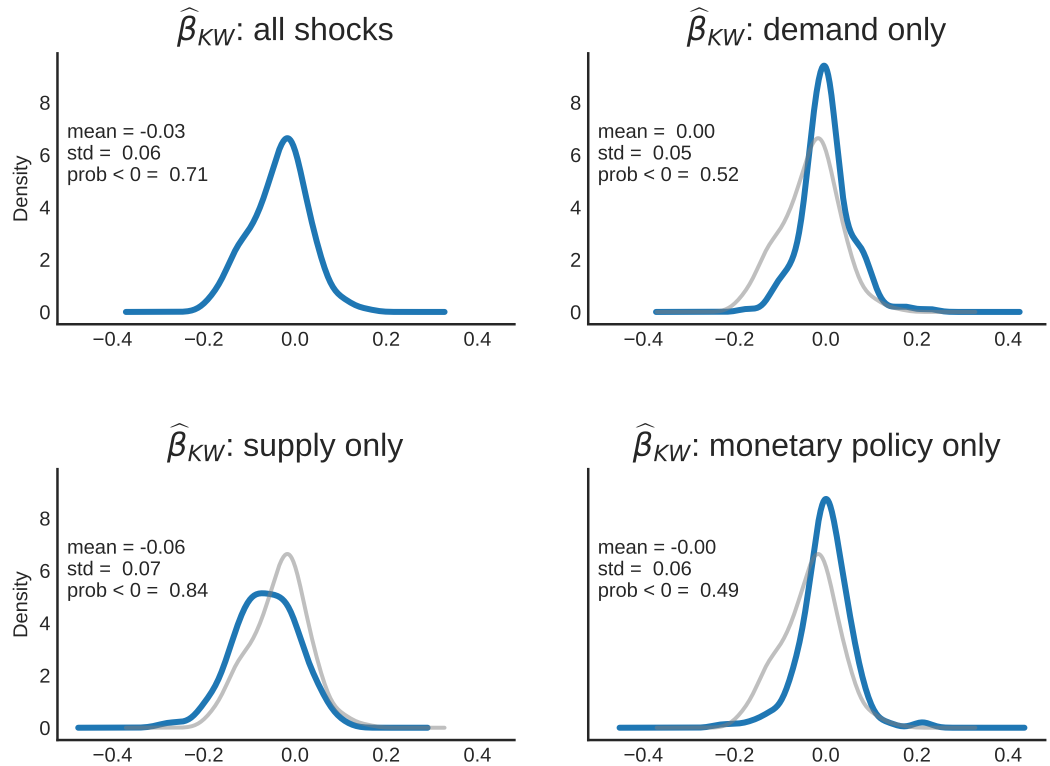

Note: The figure shows the posterior predictive densities for $\hat{\beta}_{KW}$ simulated using all the shocks (upper left panel), and conditional only on demand (upper right panel), supply (bottom left panel), and monetary policy (bottom right panel) shocks. The black vertical line indicates the point estimate using the SPF data.

4.2 Fitting Observed Inflation Expectations

The model used in the previous subsection did not use inflation expectations as an observable. In some sense, this makes the fact that the model can match the inflation expectation predictability statistics more impressive, as these moments are not implicity contained in the likelihood function. That said, it is not obvious that the model can track the time series of inflation expectations and continue to match these moments (and continue to fit output growth, inflation, and interest rates well.) In this subsection, we evaluate the NK-FHP model's ability to fit an additional observable, inflation expectations. We use, as in the previous subsection, the SPF to construct our inflation expectations series.

Recall that $\mathbb{E}_t^k[\pi_{t+h}]$ denotes the expectations of a h-period ahead inflation for an agent with a k-period planning horizon. We link this model variable to observed inflation expectations series using the following measurement equation:

$$\text{Expected Inflation}_t = \pi^A + \mathbb{E}_t^k[\pi_{t+1} + \pi_{t+2} + \pi_{t+3} + \pi_{t+4}] + \eta_t \tag{22}$$

The parameter π^A is the steady state annual inflation rate and $\mathbb{E}_t^k$ denotes the forecast of economic agents with planning horizon of length k. We follow Del Negro and Eusepi (2011) and allow for measurement error, η_t, when including inflation expectations as an observable. The measurement error follows an AR(1) process:

$$\eta_t = \rho_\eta \eta_{t-1} + \epsilon_{\eta,t}, \text{ with } \epsilon_{\eta,t} \overset{iid}{\sim} N\left(0, \sigma_\eta^2\right).$$

As discussed in Del Negro and Eusepi (2011) there are number of reasons why it may be important to include measurement error when adding the SPF measure of inflation expectations to our estimation. One reason is that the information sets of the SPF forecasters and those of the economic agents in the model may not correspond exactly. Indeed, the SPF is produced in the middle of the quarter, while the model-based forecasts are made at the start of every quarter as leading to an information mismatch.[12]

To evaluate how the inflation data affects the NK-FHP model estimates, Table 4 shows the log marginal data density (MDD), a summary measure of fit, for different values of agents' planning horizon, k. The first column shows the MDD using only the macroeconomic data and thus in logs is defined as:

$$\log p(Y) = \log\left(\int p(Y|\theta)p(\theta)d\theta\right).$$

where Y is the observed data consisting of the "standard" macroeconomic observables of output growth, inflation, and interest rates. As shown in this column, a planning horizon in which agents only make state-contingent plans one-quarter ahead ($k = 1$) fits the macroeconomic data better than the models with longer planning horizons.

The second and third columns of the table emphasize the role of the inflation forecast data, which we denote as $E\pi$, as an observable. In particular, these columns show the log predictive data density as well as the MDD inclusive of the inflation forecast data. The log predictive data density provides a measure of the model's fit on the inflation forecast data and is defined as conditional on the standard macreconomic data:

$$\log p(E\pi|Y) = \log\left(\int p(E\pi|Y,\theta)p(\theta|Y)d\theta\right).$$

The log predictive data density allows us to uncouple the model's fit of the survey data of inflation expectations from its fit of the standard macroeconomic data.[13] As shown in the second column,

[12] We follow Del Negro and Eusepi (2011) as well as others by using the current vintage of data for the other observables while the series on SPF inflation forecasts is a real-time measure. This difference leads to an information mismatch between the econometrician and forecaster and hence is another reason for the inclusion of measurement error.

[13] In addition, this object is less sensitive to a researcher's prior distribution, as $p(\theta)$ is replaced by $p(\theta|Y)$.

Table 4: Log MDD Estimates

Model	$\log p(Y)$	$\log p(F\pi\|Y)$	$\log p(Y, F\pi)$
$k = 0$	-720.42	-56.02	-776.44
$k = 1$	-716.54	-49.37	-765.91
$k = 2$	-718.91	-49.53	-768.44
$k = 3$	-721.46	-48.95	-770.41
$k = 4$	-723.52	-48.46	-771.99
$k = 5$	-724.11	-51.47	-775.57

Note: The table shows point estimates of the log MDDs and the log predictive data density computed using the output of the SMC samplers. See appendix for details.

the estimate of the planning horizon changes significantly if we focus on fitting only the survey data on inflation expectations. In that case, the estimate of the planning horizon would include the next four quarters as well as the current quarter ($k = 4$). Such a change has important implications for monetary policy, as it would considerably strengthen the economic effects of forward guidance that policymakers gave about the policy rate.

While focusing exclusively on fitting inflation forecast data implies that the estimates of agents' planning horizons are considerably longer, the third column shows the MDD using both sets of observables, which satisfies:

$$\log p(Y, E\pi) = \log p(Y) + \log p(E\pi|Y).$$

This relationship is highlighted in Table 4, as the values in the third column of the table are the sum of the first two columns. Table 4 indicates that the improved fit of the inflation forecast from the NK-FHP model with $k = 4$ is more than offset by the deterioration in fit, shown in column 1, of the standard macroeconomic data. Because of this deterioration in fit in the standard macroeconomic time series, the estimated planning horizon using this data jointly with the survey data on inflation expectations involves planning only one quarter ahead ($k = 1$).

The inflation expectations data when used jointly as an observable with the other macroeconomic data does not change the estimated planning horizon in the NK-FHP model and we find similar results for the model's other parameters. Table 2 compares the posterior estimates of other structural parameters of the NK-FHP model with $k = 1$ using both sets of observables to their estimates when the observables do not include the inflation expectations data. The posterior distribution of the firm learning rate, γ_p, shifts slightly, with the posterior mean increasing from 0.16 to 0.20. The estimates of the parameters, κ and σ, which determine the sensitivity of inflation to the output gap and the sensitivity of aggregate demand to changes in the policy rate, respectively, are little changed by the inclusion of the survey data on inflation expectations. The same is true for the persistence of the exogenous shocks.

Table 5: NK-FHP($k = 1$) Models: Comparison of Selected Posterior Statistics

	With Expectations Data		Without Expectations Data	
	Mean	[05, 95]	Mean	[0, 95]
		Learning		
γ	0.50	[0.33, 0.67]	0.47	[0.30, 0.64]
γ_p	0.16	[0.13, 0.20]	0.20	[0.13, 0.29]
		Endogenous Propagation		
κ	0.03	[0.02, 0.04]	0.03	[0.02, 0.04]
σ	2.71	[1.95, 3.58]	2.68	[1.94, 3.52]
		Monetary Policy Rule		
ϕ_π	0.96	[0.71, 1.26]	0.98	[0.72, 1.27]
ϕ_y	0.90	[0.60, 1.31]	0.89	[0.59, 1.29]
$\bar{\phi}_\pi$	1.92	[1.56, 2.31]	1.84	[1.49, 2.23]
$\bar{\phi}_y$	0.13	[0.04, 0.25]	0.13	[0.04, 0.25]
		Persistence of Exogenous Processes		
ρ_ξ	0.90	[0.84, 0.96]	0.90	[0.84, 0.96]
ρ_i	0.95	[0.90, 0.98]	0.95	[0.90, 0.98]
ρ_y	0.43	[0.36, 0.50]	0.45	[0.30, 0.60]
ρ_η	0.92	[0.87, 0.97]		

Note: The table shows estimates of the posterior means, 5th, and 95th percentiles of the model parameters computed from output of the SMC sampler. See appendix for details.

4.3 Comparison with Alternative Models of Expectation Formation

In this section, we compare the NK-FHP model's ability to jointly account for the macroeconomic data and the survey data on inflation expectations with an alternative model of inflation expectations formation. In the first, the formation of inflation expectations is imperfect due to the presence of sticky information (SI), as firms' pricing decisions are not always based on current information. Sticky information models are an attractive point of comparison, as a number of researchers have found these models to fit macro time series at least as well as models emphasizing sticky prices.[14] Moreover, as emphasized in CG (2015), sticky information models can successfully account for the correlation between median forecast errors and revisions observed in the SPF. While we mainly focus on the comparison of the NK-FHP model to the SI model, we also compare the performance of these models to the "hybrid" NK model which includes habit persistence in consumption and sticky price contracts that are indexed to lagged inflation. Both the sticky information and hybrid NK models are desribed in detail in the appendix.

In order to compare the overall fit across models, Table 6 displays the log MDDs of the alternative models with and without the inflation forecast series. Table 6 also display the log predictive data densities to assess a model's fit to the inflation forecast series, taking the macroeconomic time

[14]For estimated models with sticky information, see Andrés, López-Salido, and Nelson (2005) and Chung, Herbst, and Kiley (2014).

series as given. Starting with the SI model, the first column of Table 6 shows that the SI model fits the macroeconomic time series about as well as the hybrid-NK model but significantly worse than the NK-FHP model. As discussed in Gust, Herbst, and López-Salido (2022), the improved fit of the NK-FHP model relative to the hybrid-NK model reflects the reduced degree of forward-looking behavior especially in the aggregate demand equation relating the output gap to the policy rate. This reduced degree of forward-looking behavior also accounts for why the NK-FHP model results in an improved fit relative to the SI model in terms of fluctuations in output, inflation, and the policy rate.

Table 6: Log MDD Estimates

Model	$\log p(Y)$	$\log p(F\pi \mid Y)$	$\log p(Y, F\pi)$
Hybrid NK	−736.42	−142.84	−879.42
Sticky Information	−753.73	−50.58	−804.31
FHP ($k = 1$)	−716.54	−49.37	−765.91

Note: The table shows point estimates of the log MDDs and the log predictive data density computed using the output of the SMC samplers. See appendix for details.

The middle column of Table 6 indicates that the SI model fits the inflation forecast data nearly as well as the NK-FHP model. This good fit is also reflected in how well the SI model is able to account for the relationship between inflation forecast errors and forecast revisions in the aggregate data. As indicated in Figure 7, the point estimate of β_{CG} from the data is well within the SI model's credible set for $\widehat{\beta}_{CG}$. The model's somewhat higher values of $\widehat{\beta}_{CG}$ than the data is driven by the estimates of λ, whose median value is close to 0.6. While a positive value of $\widehat{\beta}_{CG}$ is consistent with a forecast that underreacts to new information, the underreaction of the SI inflation forecast can also be seen in Figure 8. As shown there, in the SI model the inflation forecast error rises in response to the AHS "inflation" shock, as realized inflation rises more than an average firm's forecast of inflation. The median impulse response of the forecast error in the SI model does not turn into an overreaction, and instead monotonically converges back to zero.[15] Thus, while the SI model fits the median inflation forecast in the SPF reasonably well, it does not display the sign flip in the impulse response that AHS (2020) document.

[15] In the appendix, we study a partial equilibrium version of the sticky information model and show analytically that there is always an underreaction of the impulse response of the one-step ahead forecast in the SI model.

Figure 7: Distribution of CG coefficient

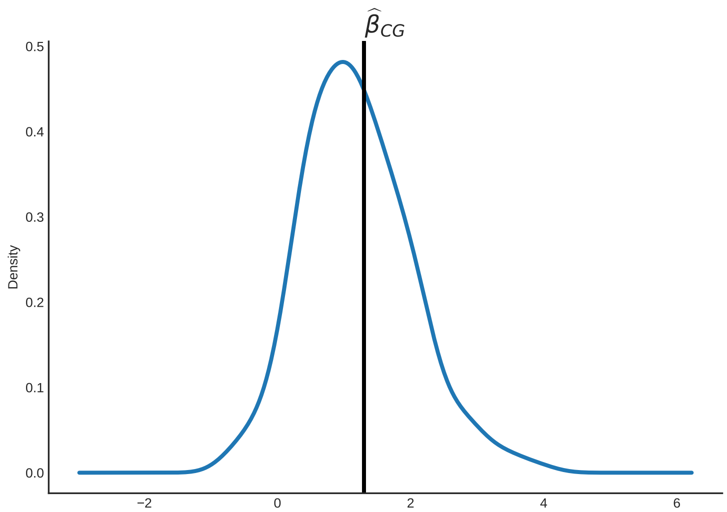

Notes: The figure shows the posterior predictive densities for $\hat{\beta}_{CG}$ simulated for the SI model. The black vertical line indicates the point estimate using the SPF data.

Figure 8: Impulse Response to an AHS "Inflation" Shock

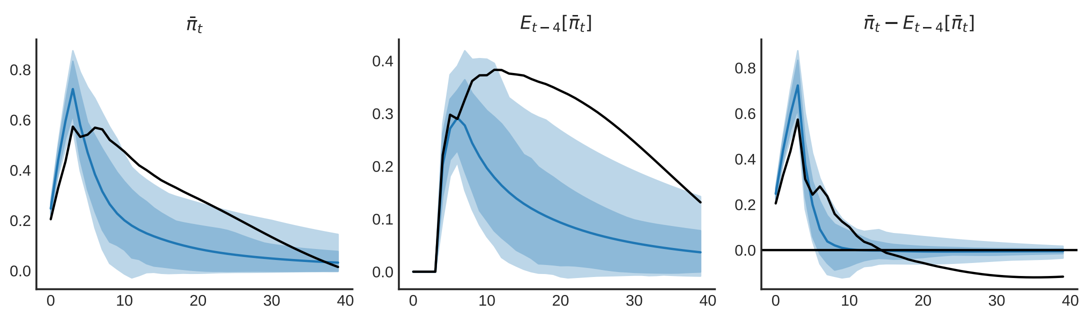

Notes: The figure shows the impulse response of inflation, the inflation forecast, and the inflation forecast error from the AHS-style VAR in the SI model. The solid blue line denotes the (pointwise) mean across the predictive checks, while the shaded regions denote the ninety (light blue) and sixty-eight (dark blue) percent bands (across the means of the predictive checks). This exercise uses trajectories of length $T = 168$.

5 Conclusion

In this paper, we used survey data on inflation expectations as well as aggregate data on output, inflation, and short-term interest rates to estimate and evaluate a NK model featuring FHP. We found that the NK-FHP model can account for the predictability of inflation forecast errors as well as the empirical evidence that the average inflation forecast in the SPF typically underreacts relative to realized inflation but overreacts later on. We also showed that the NK-FHP model can account for survey measures of inflation expectations while also providing a reasonable fit of output, inflation, and interest-rate dynamics over the business cycle. In doing so, we found that planning horizons of about a year fit the survey data on inflation expectations best; however, shorter planning horizons — on the order of a couple of quarters — are best for jointly fitting the inflation survey data and the macroeconomic time series. In addition, we found that the learning that households and firms do to form beliefs about events outside of their planning horizons was crucial to the successful performance of the NK-FHP model in terms of its ability to account for aggregate time series as well as the evidence on forecast error predictability.

The short-planning horizons and the inertia in private sector beliefs about events outside of their planning horizons have important implications for monetary policy. Notably, forward guidance policies are much less effective than when households and firms have lengthy planning horizons and disinflations are much more costly than in the canonical NK model in which households and firms have full information, rational expectations. Given these notable differences, an important avenue of future research is studying optimal monetary policy when households and firms have finite planning horizons to investigate how optimal policy depends on agents' planning horizons as well as the evolution of their beliefs regarding events outside of their planning horizons.

References

Andrés, J., J. D. López-Salido, and E. Nelson (2005): "Sticky-price models and the natural rate hypothesis," *Journal of Monetary Economics*, 52(5), 1025–1053.

Angeletos, G.-M., Z. Huo, and K. A. Sastry (2020): "Imperfect Macroeconomic Expectations: Evidence and Theory," Working Paper 27308, National Bureau of Economic Research.

Angeletos, G.-M., and C. Lian (2018): "Forward Guidance without Common Knowledge," *American Economic Review*, 108(9), 2477–2512.

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer (2018): "Over-reaction in Macroeconomic Expectations," NBER Working Papers 24932, National Bureau of Economic Research, Inc.

Broer, T., and A. Kohlhas (2018): "Forecaster (Mis-)Behavior," Discussion Paper DP12898, CEPR.

Calvo, G. (1983): "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics*, 12(3), 383–398.

Carvalho, C., S. Eusepi, E. Moench, and B. Preston (2023): "Anchored Inflation Expectations," *American Economic Journal: Macroeconomics*, 15(1), 1–47.

Chung, H., E. Herbst, and M. T. Kiley (2014): "Effective Monetary Policy Strategies in New Keynesian Models: A Re-examination," Working Paper 20611, National Bureau of Economic Research.

Coibion, O., and Y. Gorodnichenko (2015): "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts," *American Economic Review*, 105(8), 2644–78.

Del Negro, M., and S. Eusepi (2011): "Fitting observed inflation expectations," *Journal of Economic Dynamics and Control*, 35(12), 2105–2131, Frontiers in Structural Macroeconomic Modeling.

Eusepi, S., and B. Preston (2018): "Fiscal Foundations of Inflation: Imperfect Knowledge," *American Economic Review*, 108(9), 2551–89.

Fuhrer, J. C. (2018): "Intrinsic Expectations Persistence: Evidence from Professional and Household Survey Expectations," Working Paper 18-9, Federal Reserve Bank of Boston.

Gabaix, X. (2020): "A Behavioral New Keynesian Model," *American Economic Review*, 110(8), 2271–2327.

Gust, C., E. Herbst, and D. López-Salido (2022): "Short-term Planning, Monetary Policy, and Macroeconomic Persistence," American Economic Journal: Macroeconomics, forthcoming.

Herbst, E., and F. Schorfheide (2014): "Sequential Monte Carlo Sampling for DSGE Models," *Journal of Applied Econometrics*, 29(7), 1073–1098.

Herbst, E. P., and F. Schorfheide (2015): *Bayesian Estimation of DSGE Models*. Princeton University Press, Princeton, NJ.

Kohlhas, A. N., and A. Walther (2021): "Asymmetric Attention," *American Economic Review*, 111(9), 2879–2925.

Kreps, D. (1998): "Anticipated Utility and Dynamic Choice," in *Frontiers of Research in Economic Theory*, ed. by D. P. Jacobs, E. Kalai, M. I. Kamien, and N. L. Shwartz, pp. 242–274. Cambridge University Press, Cambridge.

Mankiw, N. G., and R. Reis (2002): "Sticky information versus sticky prices: A proposal to replace the new Keynesian Phillips curve," *Quarterly Journal of Economics*, 117(4), 1295–1328.

Milani, F. (2007): "Expectations, Learning and Macroeconomic Persistence," *Journal of Monetary Economics*, 54(7), 2065 – 2082.

Ormeno, A., and K. Molnar (2015): "Using Survey Data of Inflation Expectations in the Estimation of Learning and Rational Expectations Models," *Journal of Money, Credit and Banking*, 47(4), 673–699.

Reis, R. (2009): "Optimal Monetary Policy Rules in an Estimated Sticky-Information Model," *American Economic Journal: Macroeconomics*, 1(2), 1–28.

Slobodyan, S., and R. Wouters (2012): "Learning in a Medium-Scale DSGE Model with Expectations Based on Small Forecasting Models," *American Economic Journal: Macroeconomics*, 4(2), 65–101.

Woodford, M. (2018): "Monetary Policy Analysis when Planning Horizons are Finite," Working Paper 24692, National Bureau of Economic Research.

Appendix for
Inflation Expectations and Macro Dynamics under Finite Horizon Planning

Christopher Gust, Edward Herbst, and David Lopez-Salido

A Finite-Horizon Household Planning

In this section, we describe the optimal finite-horizon plan set by households and derive equations (13) and (16). We focus on the optimal plan chosen by a finite-planning houesehold and abstract from the static labor supply decision that a household makes. A household chooses a state-contingent plan for consumption and bond holdings, $\{C_\tau, B_{\tau+1}\}_{\tau=t}^{t+K}$ to maximize:

$$\mathbf{E}_t^K \left\{ \sum_{\tau=t}^{t+K} \beta^{\tau-t} Q_\tau U(C_\tau) + \beta^{K+1} Q_{t+K+1} V_t(B_{t+K+1}) \right\} \tag{A-1}$$

where $0 < \beta < 1$ and $V_t(B_{t+K+1})$ is the value function a household uses to assign continuation values to its plans over the remainder of its infinite lifetime. As discussed further below, this value function varies over time but is fixed at time t when a household chooses its finite-horizon plans. The value function depends on a household's financial position at the end of its planning period, B_{t+K+1}. However, households are assumed to have limited ability to understand events that occur in the distant future and thus, the value function is not the model consistent one that reflects all possible contingencies that a household may face in the future. The variable Q_τ reflects that the discount factor is stochastic. For $t \le \tau \le t+K+1$, it evolves according to:

$$Q_\tau = \prod_{i=0}^{\tau-t-1} \xi_{t+i}, \tag{A-2}$$

where the variable ξ_t is an exogenous shock that affects a household's rate of time preference between periods t and $t+1$. According to equation (A-2), $Q_t = 1$, $Q_{t+1} = \xi_t$, and Q_{t+K+1} reflects that a household contemplates all possible contingencies of the shocks, $\xi_t, \xi_{t+1}, ...\xi_{t+K}$, that take place over its finite-planning horizon.

A household takes its initial bond holdings, B_t, as given and faces a per-period budget constraint given by:

$$B_{\tau+1} = (1+i_\tau)\left[\frac{B_\tau}{\Pi_\tau} + Y_\tau - C_\tau\right], \tag{A-3}$$

where i_τ is the policy rate, Π_τ is the (gross) inflation rate, and Y_τ denotes a household's disposable income which includes labor income as well as the profits a household receives from ownership of the economy's firms.

The first order conditions from a household's finite horizon plan are given by:

$$\frac{U_{C\tau}}{1+i_\tau} = \beta \xi_\tau \mathbf{E}_\tau^K \frac{U_{C\tau+1}}{\Pi_{\tau+1}} \qquad \text{for } t \le \tau \le t+K-1, \tag{A-4}$$

$$\frac{U_{Ct+K}}{1+i_{t+K}} = \beta \xi_{t+K} V_{Bt}(B_{t+K+1}) \tag{A-5}$$

where $V_{Bt}(B_{t+K+1})$ denotes a household's marginal value function with respect to its financial position at the end of its planning horizon (i.e., $V_{Bt}(B_{t+K+1}) = \frac{\partial V_t(B_{t+K+1})}{\partial B_{t+K+1}}$). A household's marginal utility of consumption, $U_{c\tau}$, satisfies:

$$U_{C\tau} = C_\tau^{\frac{-1}{\sigma}}.$$

We log-linearize a household's first order condtions around a non-stochastic steady state in which aggregate output and consumption satisfy $Y = C = 1$. Also, $\xi = 1$ and the nominal interest rate in steady state satisfies $1 + i = \frac{\Pi}{\beta}$ where Π denotes the (gross) inflation rate. Log-linearizing equation (A-4) implies:

$$\mathbb{E}_\tau^k \{c_\tau - c_{\tau+1} + \sigma [i_\tau - \pi_{\tau+1} - r_\tau^\star]\} = 0 \tag{A-6}$$

for $t \leq \tau \leq t + K - 1$. We use lower case variables to denote the log-linearized variables so that $i_\tau = \log(1 + i_\tau) - \log(1 + i)$ and $\pi_\tau = \log(\Pi_\tau) - \log(\Pi)$. Also, $r_\tau^\star$ is defined as $= r_\tau^\star = -\log(\xi_\tau)$. Using (1), the subjective expectations operator in this expression can be replaced by the rational expectations operator with redefined variables that reflect an agent's subjective expectations:

$$c_{t+K-j}^j = E_t c_{t+K-j+1}^{j-1} - \sigma \left[i_{t+K-j}^j - E_t \pi_{t+K-j+1}^{j-1} - r_{t+K-j}^\star \right] \tag{A-7}$$

for $0 \leq j \leq K$. We also log-linearize a household's terminal condition. This condition requires that we approximate $V_{Bt}(B_t)$. To do so, we log-linearize it around its non-stochastic steady state value of $\frac{1}{\Pi}$ and parameterize it as a linear function whose slope and intercept coefficients can potentially change over time as household's learn and update their longer-run beliefs based on their past experience:

$$V_{Bt}(B_t) \approx -\sigma^{-1} [v_{ht} + \chi_t b_t] \tag{A-8}$$

When making their optimal plan at time t, a household treats v_t and χ_t as fixed and hence their optimal consumption and savings decisions at time t depends on these parameters as well as their initial net asset position, B_t. In expression (A-8), $b_t = \frac{B_t}{\Pi_t}$ and we linearize around b_t, since in steady state $B = 0$. Using equation (A-8), the log-linearized version of equation (A-5) is:

$$c_{t+K}^0 = -\sigma \left[i_{t+K}^0 - r_{t+K}^\star \right] + \left[v_t + \chi_t b_{t+K+1}^0 \right] \tag{A-9}$$

We also linearize a household's budget constraint:

$$b_{t+K-j+1}^j = \beta^{-1} \left[b_{t+K-j}^{j+1} + y_{t+K-j}^j - c_{t+K-j}^j \right] \tag{A-10}$$

where with this notation b_t^{K+1} denotes a household's initial net asset position and $b_{t+K-j+1}^j$ for $0 \leq j \leq K$ denotes a household's plans for the evolution of its net assets.

Households do not know the model-consistent value functions, but they learn adaptively and update their value functions based on observed data. Specifically, a household computes a new estimate of their value function at the same time as choosing its optimal state-contingent plan. This new estimate is consistent with their optimal plan, as it satisfies the envelope condition associated with maximizing equation (A-1):

$$V_{Bt}^E(B_t) = E_t^K \frac{U_{ct}(C_t(B_t))}{\Pi_t}, \tag{A-11}$$

where $V_{Bt}^E(B_t)$ denotes a household's new estimate of its value function. In expression (A-11), $C_t(B_t)$ is a household's optimal consumption decision taking v_t, χ_t, and B_t as given. A household uses $V_{Bt}^E(B_t)$ to form their continuation value function at date $t + 1$ by combining it with their current continuation value function according to:

$$V_{Bt+1}(B_t) = (1 - \gamma) V_{Bt}(B_t) + \gamma V_{Bt}^E(B_t), \tag{A-12}$$

We linearize the functions in expressions (A-11) and (A-12). The latter linearization implies:

$$v_{ht+1} = (1-\gamma)v_{ht} + \gamma v^e_{ht} \tag{A-13}$$
$$\chi_{t+1} = (1-\gamma)\chi_t + \gamma \chi^e_t \tag{A-14}$$

where v^e_{ht} and χ^e_t are the intercept and slope coefficients to the linear approximation to $V^E_{Bt}(B_t)$:

$$V^E_{Bt}(B_t) \approx -\frac{1}{\sigma}\left(v^E_{ht} + \chi^E_t b_t\right) = -\frac{1}{\sigma}\left(c^K_t(b_t; v_{ht}, \chi_t) + \sigma\pi^K_t\right) \tag{A-15}$$

The linearized consumption function, $c^K_t(b_t; v_{ht}, \chi_t)$, depends on the parameters of the value function as well as their initial net asset position. The linearized solution to the optimal consumption function can be determined through recursive substitution using equations (A-7), (A-10), and (A-9). Since it is a linear function, it is convenient to write it in terms of an intercept term and a slope term:

$$c^K_t(b_t; v_{ht}, \chi_t) = c^K_t(0; v_{ht}) + g^K(\chi_t)b_t \tag{A-16}$$

where $c^K_t(0; v_{ht})$ is the intercept term associated with setting $b_t = 0$ and $g^K(\chi_t)$ is given by:

$$g^K(\chi_t) = \frac{\chi_t\beta^{-(K+1)}}{1 + \chi_t\sum_{i=1}^{K+1}\beta^{-i}} = \frac{\chi_t}{\beta^{K+1} + \chi_t\frac{1-\beta^{K+1}}{1-\beta}} \tag{A-17}$$

With optimal consumption defined in this way, equation (A-15) implies that $v^e_t = c^K_t(0; v_{ht}) + \sigma\pi^K_t$ and $\chi^e_t = g^K(\chi_t)$.

In equilibrium, $b_t = 0$, and $y^K_t = c^K_t$ so we can simplify v^e_{ht} further:

$$v^e_{ht} = y^K_t + \sigma\pi^K_t \tag{A-18}$$

which is equation (16) in the main text. In addition, setting $b_t = 0$ in equation (A-9) allows us to determine $y^K_t = c^K_t = c^K_t(0; v_{ht})$ through recursive substitution using equation (A-7):

$$y^K_t = -\sigma E_t\sum_{i=0}^{K}\left(i^{K-i}_{t+i} - r^\star_{t+i}\right) + \sigma E_t\sum_{i=0}^{K-1}\pi^{K-i}_{t+i+1} + v_{ht} \tag{A-19}$$

which is equation (13) in the main text. With a representative household and bonds in fixed supply, the evolution of χ_t is irrelevant for the economy's aggregate dynamics. In particular, only v_{ht} affects aggregate expenditures and thus it is sufficient to use only equations (A-13) and (A-18) to characterize the evolution of longer-run beliefs of households.

B Analytical Results for the FHP Model

In this section, we provide proofs of the paper's three propositions. Proposition 1 characterizes the impulse response functions of inflation forecasts and forecast errors in the FHP model, and the proposition 2 characterizes the FHP model's properties for the predictability regression of CG (2015). Proposition 3 characterizes the impulse response of the average inflation forecast in the SI model.

Proof of Proposition 1. Consider the first part of Proposition 1 that characterizes the impulse response functions of the the model with no learning. In that case, $v_{pt} = 0$ and we need only focus on the effect of changes in the output gap on the impulse responses. The impulse response of the output gap at date $t+i$ is given by:

$$\frac{\partial y_{t+i}}{\partial e_t} = \rho^i \tag{A-20}$$

which reflects that firms know the process for y_t. Using this expression in equation (9) and differentiating it at $t+i$ with respect to e_t implies that the impulse resposse function of agents' one:

$$\frac{\partial \mathbb{E}^k_{t+i}\pi_{t+1+i}}{\partial e_t} = \rho^{i+1}A(k-1)\kappa \tag{A-21}$$

Because $\rho \geq 0$, $0 < \beta < 1$, and $\kappa > 0$, this expression is non-negative for any finite value of $k > 0$. We can also use expression (10) to determine the impulse response of the one-step ahead forecast error. Differentiating this expression at $t+i$ with respect to e_t and using (A-20) yields:

$$\frac{\partial \mathbb{F}_{t+1}}{\partial e_t} = \rho^{i+1}(\beta\rho)^k\kappa \geq 0 \tag{A-22}$$

which completes the proof for the FHP model without learning.

The second part of the proposition characterizes the impulse response function of the model with learning. To characterize the impulse responses of expected inflation and the inflation forecast error, it is convenient to first characterize the impulse response of the value function. Using equation (**??**), the impulse response of the value function at $t+i$ is:

$$\frac{\partial v_{pt+i}}{\partial e_t} = \frac{\gamma_p}{1-\theta}\sum_{j=0}^{i-1}(1-\gamma)^{i-1-j}\frac{\partial \pi^k_{t+j}}{\partial e_t} \tag{A-23}$$

for $i > 0$. We also know that $\frac{\partial v_{pt}}{\partial e_t} = 0$ since the value function is predetermined at time t. We can use the impulse response of inflation to rewrite the impulse response of the value function in terms of the model's parameters. The impulse response of inflation is given by:

$$\frac{\partial \pi^k_{t+i}}{\partial e_t} = A(k)\kappa\rho^i + \beta^{k+1}(1-\theta)\frac{\partial v_{pt+i}}{\partial e_t} \tag{A-24}$$

Using equation (A-24), expression (A-23) can be rewritten as:

$$\frac{\partial v_{pt+i}}{\partial e_t} = (1-\widetilde{\gamma_p})\frac{\partial v_{pt+i-1}}{\partial e_t} + \frac{\gamma_p}{1-\theta_p}A(k)\kappa\rho^{i-1} \tag{A-25}$$

where $\widetilde{\gamma_p} = \gamma_p(1-\beta^{k+1})$. This expression can be rewritten as:

$$\frac{\partial v_{pt+i}}{\partial e_t} = \frac{\gamma_p}{1-\theta_p}A(k)\kappa\sum_{j=0}^{i-1}(1-\widetilde{\gamma_p})^{i-1-j}\rho^j \tag{A-26}$$

which holds for $i > 0$. From this expression, we can see that $\frac{\partial v_{pt+i}}{\partial e_t} \geq 0 \ \forall i$, which implies that the impulse response of expected inflation is always non-negative:

$$\frac{\partial \mathbb{E}^k_{t+i}\pi_{t+1+i}}{\partial e_t} = \rho^{i+1}A(k-1)\kappa + \beta^k(1-\theta)\frac{\partial v_{pt+i}}{\partial e_t} \geq 0. \tag{A-27}$$

For the change in the sign of the impulse response of the one-step ahead inflation forecast error, note that at $i = 0$, $\frac{\partial v_{pt}}{\partial e_t} = 0$ and

$$\frac{\partial \mathbb{F}_{t+1}}{\partial e_t} = \left[\rho(\beta\rho)^k + \beta^{k+1}\gamma_p A(k)\right]\kappa > 0 \tag{A-28}$$

Accordingly, on impact the forecast error rises. For $i > 0$, the impulse response of the forecast error is given by:

$$\frac{\partial \mathbb{F}_{t+1+i}}{\partial e_t} = \left[\rho(\beta\rho)^k + \beta^{k+1}\gamma_p A(k)\right]\kappa\rho^i - \beta^k\left[1 - \beta(1 - \widetilde{\gamma}_p)\right](1 - \theta)\frac{\partial v_{pt+i}}{\partial e_t} \tag{A-29}$$

For the impulse response of the forecast error to be negative at $i > 0$ and $\rho > 0$ requires that:

$$[1 - \beta(1 - \widetilde{\gamma}_p)]\sum_{j=0}^{i-1}\left(\frac{1 - \widetilde{\gamma}_p}{\rho}\right)^{i-1-j} > \left[\beta + \frac{\rho^{k+1}}{\gamma_p A(k)}\right]\rho \tag{A-30}$$

If expression (A-30) holds at response $i^\star$, then it will also hold at $i > i^\star$ given that the sum on left hand side grows over time. If $1 - \widetilde{\gamma}_p > \rho$, then the forecast error is unbounded as $i \to \infty$ and there must exist an $i^\star$ for which expression (A-30) is satisfied. Substituting the expression for $\widetilde{\gamma}_p$ into the condtion, $1 - \widetilde{\gamma}_p > \rho$, yields the expression used in the proposition. Note that if $\rho = 0$, this condition does not apply and expression (A-29) implies $i^\star = 1$.

Proof of Proposition 2. To show that $\beta_{CG} > 0$, it is sufficient to show that the covariance between the inflation forecast error and forecast revision is positive. To show this in the no learning case, note that in the case of no learning equation (10) can be simplified to:

$$\mathbb{F}^k_{t+1} = \rho(\beta\rho)^k\kappa(\rho y_{t-1} + e_t) \tag{A-31}$$

Also, without learning, the forecast revision at date t is given by:

$$\mathbb{R}^k_t = \rho^2\kappa(\beta\rho)^{k-1}y_{t-1} + \rho A(k-1)\kappa e_t \tag{A-32}$$

Using these two expressions, the covariance between forecast errors and revisions is:

$$\text{cov}(\mathbb{R}^k_t, \mathbb{F}^k_{t+1}) = \rho^4(\beta\rho)^{2k-1}\kappa^2\text{var}(y_t) + \rho^2(\beta\rho)^k A(k-1)\kappa^2\text{var}(e_t) \tag{A-33}$$

Expression (A-33) implies that if $\rho > 0$, then $\text{cov}(\mathbb{R}^k_t, \mathbb{F}^k_{t+1}) > 0$, which completes the proof for the no learning case.

With learning, we focus on the case in which $\rho = 0$. In that case, equation (10) can be written as:

$$\mathbb{F}^k_{t+1} = \beta^{k+1}\gamma_p A(k)\kappa e_t - \beta^k\left(1 - \beta(1 - \widetilde{\gamma}_p)\right)(1 - \theta)v_{pt} \tag{A-34}$$

With $\rho = 0$, the forecast revision at time t is given by:

$$\mathbb{R}^k_t = (1 - \theta)\beta^{k-1}(\beta v_{pt} - v_{pt-1}) \tag{A-35}$$

Both v_{pt} and v_{pt-1} are uncorrelated with e_t since they are determined before e_t is realized. This implies that:

$$\text{cov}(\mathbb{R}^k_t, \mathbb{F}^k_{t+1}) = -\beta^{2k-1}\left(1 - \beta(1 - \widetilde{\gamma}_p)\right)(1 - \theta)^2\left[\beta\text{var}(v_{pt}) - \text{cov}(v_{pt}, v_{pt-1})\right] \tag{A-36}$$

With $\rho = 0$, we can write the covariance between the value function at time t and $t - 1$ as:

$$\text{cov}(v_{pt}, v_{pt-1}) = (1 - \widetilde{\gamma}_p)\text{var}(y_t)$$

Substituting this expression into equation (A-36), the covariance between forecast error and revision is:

$$\text{cov}(\mathbb{R}_t^k, \mathbb{F}_{t+1}^k) = -\beta^{2k-1}\left(1-\beta(1-\widetilde{\gamma}_p)\right)(1-\theta)^2\left[\gamma_p(1-\beta^{k+1})-(1-\beta)\right]\text{var}(v_{pt}) \tag{A-37}$$

This expression implies that if $1-\beta > \gamma_p(1-\beta^{k+1})$, then $\text{cov}(\mathbb{R}_t^k, \mathbb{F}_{t+1}^k) > 0$, which is the condition given in Proposition 2.

Proof of Proposition 3. We need to consider a firm's forecast h quarters ahead along with the associated forecast error. For $k \geq h$, a firm's forecast is given by:

$$\mathbb{E}_t^k \pi_{t+h} = \rho^h A(k-h)\kappa y_t + \beta^{k-h+1}(1-\theta_p)v_{pt}. \tag{A-38}$$

A firm's forecast error is given by:

$$\begin{aligned}\mathbb{F}_{t+h}^k = &\left\{\left[\rho^h + \beta^{k+1}\gamma_p B_h(k)\right]A(k) + \rho^h A(k-h)\right\}\kappa y_t + \\ &\left[\beta^{k+1}(1-\widetilde{\gamma}_p) - \beta^{k-h+1}\right](1-\theta_p)v_{pt} + O_{t+h}.\end{aligned} \tag{A-39}$$

where O_{t+h} is an unpredictable component consisting of innovations in the shock from periods $t+1$ to $t+h$. The term $B_h(k)$ satisfies

$$B_h(k) = \sum_{i=0}^{h-1} \rho^{h-1-i}(1-\widetilde{\gamma}_p)^i$$

The first part of the proposition considers the case of no learning. In that case, a firm's forecast error satisfies:

$$\mathbb{F}_{t+h}^k = \frac{(\beta\rho)^{k-h+1}\left[1-(\beta\rho)^h\right]}{1-\beta\rho}\kappa y_t \tag{A-40}$$

Under no learning, equation (8) implies that inflation evolves according to:

$$\pi_t^k = A(k)\kappa y_t \tag{A-41}$$

where $A(k)\kappa > 0$. With $\rho > 0$, the forecast error's coefficient on the output gap is positive for any $h > 1$. Because the coefficient on the output gap for inflation is also positive, the covariance between the forecast error at horizon $h > 1$ and inflation will be positive when $\rho > 0$. Accordingly, under no learning, $\beta_{KW} > 0$ with $\rho > 0$.

The second part of the proposition considers the case of learning when the output-gap shock is *iid*. With $\rho = 0$, the h-step ahead forecast error with learning simplifies to:

$$\mathbb{F}_{t+h}^k = \left[\beta^{k+1}\gamma_p(1-\widetilde{\gamma}_p)^{h-1}\right]A(k)\kappa e_t + \left[\beta^{k+1}(1-\widetilde{\gamma}_p)^h - \beta^{k-h+1}\right](1-\theta_p)v_{pt} + O_{t+h}. \tag{A-42}$$

Using equation (8), the covariance between the forecast error and inflation is given by:

$$\begin{aligned}\text{cov}(\pi_t^k,\ \mathbb{F}_{t+h}^k) = &\left[\beta^{k+1}\gamma_p(1-\widetilde{\gamma}_p)^{h-1}\right](A(k)\kappa)^2\text{var}(e_t) + \\ &+\beta^{k+1}\left[\beta^{k+1}(1-\widetilde{\gamma}_p)^h - \beta^{k-h+1}\right]\text{var}(\widetilde{v}_{pt})\end{aligned} \tag{A-43}$$

where $\widetilde{v}_{pt} = (1-\theta_p)v_{pt}$. Using equation (6), the variance of the value function satisfies:

$$\text{var}(\widetilde{v}_{pt}) = \frac{\gamma_p\kappa^2\text{var}(e_t)}{(1-\beta^{k+1})(2-\widetilde{\gamma}_p)} \tag{A-44}$$

Using this expression in equation (A-43), the covariance between the forecast error and inflation can be rewritten as:

$$\text{cov}(\pi_t^k,\ \mathbb{F}_{t+h}^k) = \beta^{k+1}\gamma_p(1-\widetilde{\gamma}_p)^{h-1}\kappa^2\left\{1-\frac{\beta^{k+1}(1-\widetilde{\gamma}_p)}{(1-\beta^{k+1})(2-\widetilde{\gamma}_p)}\left[\beta^{-h}(1-\widetilde{\gamma}_p)^{-h}-1\right]\right\}\text{var}(e_t)$$

This expression implies that the covariance will be negative if and only if:

$$\frac{\beta^{k+1}(1-\widetilde{\gamma}_p)}{(1-\beta^{k+1})(2-\widetilde{\gamma}_p)}\left[\beta^{-h}(1-\widetilde{\gamma}_p)^{-h}-1\right] > 1.$$

which is the expression shown in proposition 3.

C Computing Inflation Expectations for $h > k$

We need to make an additional assumption about agents' beliefs to compute their forecasts of inflation for forecast horizons h that exceed agents' k-period ahead planning horizons. The assumption that we make in this case is that agents use their beliefs at the end of their planning horizons and do so taking into account their knowledge of the persistence of the shocks.

To understand this assumption, we first consider its implications in the partial equilibrium model. An agent's expectations for inflation at the end of its planning horizon are:

$$E_t^k\pi_{t+k} = \kappa E_t^k y_{t+k} + \beta(1-\theta_p)v_{pt}. \tag{A-45}$$

In the partial equilibrium model, we can use equation (1) and the fact that the output gap follows an exogenous, AR(1) process to write:

$$E_t^k\pi_{t+k} = \kappa E_t y_{t+k} + \beta(1-\theta_p)v_{pt} = \kappa\rho^k y_t + \beta(1-\theta_p)v_{pt} \tag{A-46}$$

Unlike $E_t^k\pi_{t+k}$, an agent's expectations for $h > k$ do not affect their decisions and thus are not needed to solve the NK-FHP model. However, in our empirical exercise to compare the model's implications to the empirical moments in the data, we need to compute $E_t^k\pi_{t+h}$ for $h > k$. We do so assuming agents forecast applying equation (A-45) to periods beyond their planning horizon:

$$E_t^k\pi_{t+h} = \kappa E_t^k y_{t+h} + \beta(1-\theta_p)v_{pt} \tag{A-47}$$

where $h > k$. In this expression, we still need to compute $E_t^k y_{t+h}$ and do so applying equation (1) and the fact that in partial equilibrium the output gap follows an AR(1) process. This implies

$$E_t^k\pi_{t+h} = \kappa E_t y_{t+h} + \beta(1-\theta_p)v_{pt} = \kappa\rho^h y_t + \beta(1-\theta_p)v_{pt} \tag{A-48}$$

for $h > k$. Accordingly, when $h > k$, agents use their beliefs about events outside of their planning horizons to compute $E_t^k\pi_{t+h}$ as well as their knowledge about the persistence of the output gap.

We use the same approach in general equilibrium and compute $E_t^k\pi_{t+h}$ in an analagous manner. The difference, however, is that $E_t^k\pi_{t+h}$ and $E_t^k y_{t+h}$ are simultaneously determined, respond to more shocks, and depend on v_{ht} as well as v_{pt}. Using equation (1), an agent's expectations k-periods ahead in this case are given by:

$$E_t X_{t+k}^0 = A_0^{-1}B_0P^kS_t + A_0^{-1}B_vV_t \tag{A-49}$$

where $X_{t+k} = (y_t, \pi_t)'$ and $S_t = (r_t^\star, y_t^\star, i_t^\star)'$. The vector $V_t = (v_{ht}, v_{pt})'$ and the matrices A_0, B_0, B_v are functions of the model's parameters. The matrix P is a diagonal matrix whose elements along the diagonal consist of the AR(1) coefficients of the three shocks. Equation (A-49) can be used to determine $E_t^k \pi_{t+k}$ and is the analagous expression to equation A-46. We assume that agents apply the same knowledge in making forecasts in which $h > k$ and assume that:

$$E_t^k X_{t+h} = A_0^{-1} B_0 P^h S_t + A_0^{-1} B_v V_t \tag{A-50}$$

for $h > k$, which is the analagous expression to equation (A-48).

D Sticky Information and Hybrid NK Models

In this section of the appendix, we describe the sticky information model and hybrid NK models that we estimate and compare to the NK-FHP models. Under sticky information, price-setting firms do not face costs to adjusting their prices but instead firms infrequently update the set of information upon which their price decisions are based. In particular, following Mankiw and Reis (2002), we assume that price-setters update their information sets in a staggered fashion in which there is a constant probability, $1 - \lambda$, that a firm setting a new price will revise its information set. Accordingly, a fraction, λ, of firms adjust their prices on the basis of previous information.[16] This setup gives rise to a log-linearized Phillips curve of the form:

$$\pi_t = (1 - \lambda)\lambda^{-1} mc_t + \mathbb{E}_{t-1}^{\lambda} [\pi_t + \Delta mc_t], \tag{A-51}$$

where mc_t denotes a firm's real marginal cost and $\mathbb{E}_{t-1}^{\lambda}$ representing the average time $t-1$ forecast across agents. This forecast is a weighted average of past RE forecasts (E_{t-j-1}):

$$\mathbb{E}_{t-1}^{\lambda} = (1 - \lambda) \sum_{j=0}^{\infty} \lambda^j E_{t-j-1}. \tag{A-52}$$

Because the average inflation forecast depends on past expectations of inflation, sticky information induces inertia in inflation with the degree of inertia depending on the information rigidity parameter, λ. Higher values of λ correspond to firms updating their information sets more slowly, which reduces the responsiveness of inflation to marginal cost and increases the importance of past expectations of inflation.

Given the focus of our paper on inflation, we only model price-setting firms as having sticky information. Households are assumed to use current information in their consumption-savings decisions though we still allow for habit persistence in consumption. Accordingly, the (log-linearized) aggregate demand relationship in the model is:

$$[1 + \zeta] y_t = \zeta y_{t-1} + E_t y_{t+1} - \sigma(1 - \zeta) [i_t - E_t \pi_{t+1} - r_t^\star]. \tag{A-53}$$

The presence of habits formation in consumption ($\zeta > 0$) affects the determination of real marginal cost, which satisfies:

$$mc_t = \frac{1}{1 - \zeta} [y_t - \zeta y_{t-1} - y_t^\star].$$

As in the NK model with FHP, $r_t^\star$ and $y_t^\star$ are AR(1) shocks to the equilibrium real rate and aggregate supply, respectively. Finally, in the SI model, monetary policy is specified to follow a Taylor rule:

$$i_t = \phi_\pi \pi_t + \phi_y y_t + i_t^\star. \tag{A-54}$$

[16] Reis (2009) shows how this time-dependent updating of information can arise when firms face a fixed cost to updating their information.

where as in the NK model with FHP, $i_t^\star$ is an AR(1) shock to the monetary policy rule.

We also compare the NK-FHP model and the SI model to the "hybrid" NK model. In the hybrid NK model, prices are sticky and indexed to lagged inflation which implies that aggregate inflation evolves according to:

$$\pi_t = \gamma_b \pi_{t-1} + \gamma_f \mathbb{E}_t \pi_{t+1} + \frac{\kappa}{1-\zeta}\left[y_t - \zeta y_{t-1} - y_t^\star\right], \tag{A-55}$$

In equation (A-55) the parameters $\gamma_b = \frac{1-a}{1+\beta(1-a)}$ and $\gamma_f = \frac{\beta}{1+\beta(1-a)}$ where the parameter $1-a$ determines the extent to which firms index to lagged inflation. For the "hybrid" NK model, the aggregate demand relationship satsfies equation (A-53) and monetary policy is assumed to follow (A-54).

E Analytical Results for the Sticky Information Model

To understand the implications of SI for the predictability IRFs and predictability regressions, we consider a partial equilibrium version of the model in which a firm's marginal cost is exogenous and governed by an AR(1) process:

$$mc_t = \rho_m mc_{t-1} + e_{mt}$$

In that case, we can show analytically that the IRF of the average forecast error across agents to such a shock underreacts relative to realized inflation at each point of the IRF. Moreover, as shown in CG (2015), under SI, there is a positive relationship between the average forecast error and forecast revision:

$$\mathbb{F}_{t+1}^{\lambda} \equiv \pi_{t+1} - E_t^{\lambda}\pi_{t+1} = \frac{\lambda}{1-\lambda}\left(E_t^{\lambda}\pi_{t+1} - E_{t-1}^{\lambda}\pi_{t+1}\right) + \epsilon_{t+1} \tag{A-56}$$

where ϵ_{t+1} is a function of the white noise process, e_{mt+1}. Because e_{mt+1} is unforecastable at date t, βCG, the univariate regression of the SI forecast error on revision satisfies $\beta_{CG} = \frac{\lambda}{1-\lambda}$. Accordingly, in the SI model, this regression coefficient is positive and depends only on the information rigidity parameter, λ. The sticky information model implies a positive relationship between forecast errors and revisions, because only a fraction $1-\lambda$ update their information set to a shock at date t. Accordingly, for a shock that increases marginal cost at date t, the average forecast is not revised up that much, inducing positive co-movement between the average forecast revision and forecast error. The extent of this underreaction of the forecast to the shock depends entirely on the information rigidity parameter, λ, with larger values of λ implying a more sizeable underreaction of the forecast.

While CG (2015) prove this result for the SI model for the predictability regressions that they run, they do not study the implications of sticky information for the preditability impulse responses of AHS (2020). Proposition 4 establishes that the impulse response of the SI inflation forecast to changes in marginal cost underreacts relative to realized inflation at each date of the response. Accordingly, there is no eventual overreaction, as documented by AHS (2020).

Proposition 4. *(Underreaction of IRFs of SI Inflation Forecasts). Let* $\frac{\partial \mathbb{E}_{t+i}\pi_{t+1+i}}{\partial e_{mt}}$ *and* $\frac{\partial \mathbb{E}_{t+i}^{\lambda}\pi_{t+1+i}}{\partial e_{mt}}$ *for* $i \geq 0$ *be the impulse response to an innovation in marginal cost at date* t *for realized inflation and the average inflation forecast across agents in the sticky information model, respectively. Then,* $\frac{\partial \mathbb{E}_{t+i}^{\lambda}\pi_{t+1+i}}{\partial e_{mt}} = (1-\lambda^{i+1})\frac{\partial \mathbb{E}_{t+i}\pi_{t+1+i}}{\partial e_{mt}}$, $\forall i \geq 0$.

Proposition 4 establishes that the impulse response of the average forecast across firms is proportional to the response of realized inflation at each date. Moreover, the response of the average forecast is proportionately smaller than the response of realized inflation at date $t+i$ by a factor, $0 \leq 1-\lambda^{i+1} < 1$ so that there is never an overreaction of the average forecast. The extent of the underreaction depends on λ with higher values implying a slower updating of firms' information sets and a greater underreaction of the response of the average inflation forecast.

Proof of Proposition 4. To prove proposition 4, note that with exogenous marginal cost, the solution to the SI model can be determined analytically. In particular, inflation evolves according to:

$$\pi_t = \sum_{j=0}^{\infty} b_j mc_{t-j} \tag{A-57}$$

With $mc_t = \rho_m mc_{t-1} + e_{mt}$, these coefficients satisfy:

$$b_0 = \frac{1-\lambda}{\lambda} \tag{A-58}$$

and for $j > 0$:

$$b_j = \frac{1-\lambda}{\lambda}\left[\sum_{i=0}^{j-1} \rho_m^{j-i} b_i + \rho_m^{j-1}(\rho_m - 1)\right] \tag{A-59}$$

Using equation (A-58) in equation (A-59) for $j=1$ and repeating this substitution pattern, we can show that for $j > 0$:

$$b_j = \frac{1-\lambda}{\lambda}\left(\frac{\rho_m}{\lambda}\right)^{j-1}\left(\frac{\rho_m}{\lambda} - 1\right) \tag{A-60}$$

Note that for a non-explosive solution to exist, the persistence of the marginal cost shock can not be too large. In particular, the persistence of the shock is bounded by the parameter λ so that $\rho_m < \lambda$.

With this solution in hand, the impulse response of realized inflation one-period ahead as well as the average forecast across firms can also be characterized analytically. The impulse response of realized inflation next period is given by:

$$\frac{\partial E_{t+i}\pi_{t+i+1}}{\partial e_{mt}} = A_{i+1} \tag{A-61}$$

where $A_{i+1} = A_i\rho_m + b_{i+1}$ and $A_1 = (b_0\rho_m + b_1)$. The impulse response of average inflation is given by:

$$\frac{\partial E^{\lambda}_{t+i}\pi_{t+i+1}}{\partial e_{mt}} = (1-\lambda)\sum_{j=0}^{i}\lambda^j \frac{\partial E_{t+i-j}\pi_{t+i+1}}{\partial e_{mt}} \tag{A-62}$$

Note that because we are taking the impulse response at date t with respect to e_{mt}, it is true that:

$$\frac{\partial E_{t+i}\pi_{t+i+1}}{\partial e_{mt}} = \frac{\partial E_t\pi_{t+i+1}}{\partial e_{mt}} \tag{A-63}$$

We can rewrite this expression in expression (A-63) and rewrite the response of the average forecast as:

$$\frac{\partial E^{\lambda}_{t+i}\pi_{t+i+1}}{\partial e_{mt}} = (1-\lambda)\frac{\partial E_t\pi_{t+i+1}}{\partial e_{mt}}\sum_{j=0}^{i}\lambda^j = (1-\lambda^{i+1})A_{i+1} \tag{A-64}$$

Accordingly, the response of the average inflation forecast at each date is proportional to the response of realized inflation, as described in Proposition 4.

In this section, we analyse the FHP model's empirical fit of inflation expectations, with particular attention to the predictability properties described in the previous section. The basis for our empirical analysis is (full information) estimation of dynamic stochastic general equilibrium (DSGE) models using data on inflation expectations in addition to standard macroeconomic observables. This is a methodological departure from the work of CG (2015) and AHS (2020), who examine the predictability of inflation forecasts using a limited information approach. One advantage of our approach is that it allows us to understand the extent to which matching the predictability of inflation expectations is also consistent with overall time series fit of inflation, output, short-term interest rates, and inflation expectations. In this regard, an important benefit of using a completely-specified model is that it also allows to perform historical decompositions and policy analysis. Second, the limited information approach may lack power to discriminate against alternative models of expectation formation. This is because, the estimates associated with the predictability regressions—and impulse response predictability more generally—do not depend only on imperfect expectations parameters, but also on a broader set of structural parameters. In section 2 we mostly emphasized the role of expectation formation in influencing the predictability regressions and predictibility IRFs. However, the model's other structural parameters including the persistence of the shock and the discount factor are also important determinants of these statistics. Accordingly, the full information Bayesian approach offers a viable way to deal with these considerations.

F Estimation of the FHP Model

The solution to the system of equations describing the equilibrium jointly with the observations equations define the measurement and state transition equations of a linear Gaussian state-space system. The state-space representation of a DSGE model yields a likelihood function, $p(Y|\theta)$, where Y is the observed data and θ is a vector comprised of the model's structural parameters. We estimate θ using a Bayesian approach in which the object of interest is the posterior distribution of the parameters θ. The posterior distribution is calculated by combining the likelihood and prior distribution, $p(\theta)$, using Bayes theorem:

$$p(\theta|Y) = \frac{p(Y|\theta)p(\theta)}{p(Y)}.$$

Because we can only characterize the solution to our model numerically, following Herbst and Schorfheide (2014), we use sequential Monte Carlo (SMC) techniques to generate draws from the posterior distribution. Herbst and Schorfheide (2015) provide further details on SMC and Bayesian estimation of DSGE models more generally.

We estimate the FHP model as well as several alternative DSGE models using U.S. data on output growth, inflation, and nominal interest rates from 1966:Q1 through 2007:Q4, a time period for which there were notable changes in trends in inflation and output. The observation equations for the other variables are:[17]

$$\text{Output Growth}_t = \mu^Q + y_t - y_{t-1} \tag{A-65}$$

$$\text{Inflation}_t = \pi^A + 4 \cdot \pi_t \tag{A-66}$$

$$\text{Interest Rate}_t = \pi^A + r^A + 4 \cdot i_t, \tag{A-67}$$

[17]We reparameterize β to be written in terms in the of the annualized steady-state real interest rate: $\beta = 1/(1 + r^A/400)$.

where π^A and r^A are parameters governing a model's steady state inflation rate and real rate, respectively. Also, μ^Q is the growth rate of output, as we view the DSGE models as having been detrended from an economy growing at a constant rate, μ^Q. Thus, we are using the DSGE models to explain low frequency trends in the data but not the average growth rate or inflation rate which are exogenous.

G CG Regressions

In this section, we reproduce the main regression from Coibion and Gorodnichenko (2015). We use data from the Survey of Professional Forecasters (SPF). The SPF is a quarterly panel collecting various economic forecasts from professional forecasters. Our focus will be on forecasts for four-quarter GDP deflator inflation. Specifically, we report OLS estimates of the coefficient β in the regression:

$$\pi_{t+h} - E_t\pi^A_{t+h} = c_{CG} + \beta_{CG}(E_t\pi^A_{t+h} - E_{t-1}\pi^A_{t+h}) + error_t. \tag{A-68}$$

Where $E_t\pi^A_{t+h}$ is the time t consensus (mean) forecast of annual inflation at time $t+h$. The actual inflation π_{t+h} is constructed using the vintage available one year after $t+h$ from the Philadelphia Fed's realtime data set. Table A-1 shows the regression results using data from 1969Q4-2007Q4. (Results are similar for other sample periods and design choices.)

Table A-1: CG Regression Results

$\widehat{c}_{CG}$	0.056
	(0.148)
$\widehat{\beta}_{CG}$	1.30
	(0.50)
n	148
R-squared	0.21

Note: The table shows point estimates and HAC standard errors (parentheses) from the OLS regression of A-68 along with the sample size and adjusted R-squared. The HAC standard errors are Newey-West standard errors with a Bartlett kernel with trunction equal to 4.

H AHS VAR

This section describes the algorithm for computing the "shock" as in AHS. Our starting point is the p-lag vector autoregression for the n dimensional vector y_t:

$$y_t = \Phi_0 + \Phi_1 y_{t-1} + \ldots + \Phi_p y_{t-p} + u_t, \quad u_t \overset{iid}{\sim} N(0, \Sigma).$$

Write the VAR in companion form:

$$\xi_t = F_0 + F_1 \xi_{t-1} + \nu_t,$$

with

$$\xi_t = [y_t', \ldots, y_{t-p-1}']', \quad \nu_t = [u_t', 0, \ldots, 0]', E[\nu_t \nu_t'] = \Omega, \quad \text{and } F_1 = \begin{bmatrix} \Phi_1 & \Phi_2 & \ldots & \Phi_p \\ I & 0 & \ldots & 0 \\ 0 & I & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & 0 \end{bmatrix}.$$

Define the $n \times np$ selection matrix M such that $y_t = M\xi_t$. Consider the variance of y_t over the frequency range $[\omega_0, \omega_1]$:

$$V(\omega_0, \omega_1) = \int_{w_0}^{w_1} M \left(I - F_1 e^{-i\omega}\right)^{-1} \Omega \left(I - F_1 e^{-i\omega}\right)^{-1'} M' d\omega.$$

Consider now identifing a single structural shock ϵ_{1t}, given (estimates for) $\Phi_0, \ldots, \Phi_p$ and Σ. Decompose the covariance matrix Σ as

$$\Sigma = \Sigma_{tr} \left[\alpha_1 \quad \cdots \quad \alpha_n\right],$$

where Σ_{tr} is the lower cholesky factorization of Σ and $\{\alpha_1, \ldots, \alpha_n\}$ is a collection of $n \times 1$ orthonormal vectors (i.e., $\alpha_i \alpha_j' = 1$ if $i = j$ and 0 otherwise.) Identifying the structural shock ϵ_{1t} is equivalent to finding α_1. The variance of innovations attributable to the first structural shock is

$$\Sigma_1(\alpha_1) = \Sigma_{tr} \alpha_1 \alpha_1' \Sigma_{tr}'.$$

Following AHS, we identify α_1 by maximizing contribution of the shock ϵ_{1t} over a particular frequency band. variance of y_t attributable to the first structural shock is given by:

$$S(\omega_0, \omega_1, \alpha_1) = \int_{w_0}^{w_1} M \left(I - F_1 e^{-i\omega}\right)^{-1} \Omega_1(\alpha_1) \left(I - F_1 e^{i\omega}\right)^{-1'`} M' d\omega.$$

where $\Omega_1(\alpha_1)$ is defined analogously to $\Sigma_1(\alpha_1)$. Let i be the index which corresponds to the inflation observable. Then α_1 is such that

$$\alpha_1^* = \text{argmax}_{|\alpha_1|=1} [S(\omega_0, \omega_1, \alpha_1)_{ii} / V(\omega_0, \omega_1)_{ii}].$$

Following AHS, we set the frequencies ω_0 and ω_1 to corresponds to periods of length 32 and 6, respectively. In our computations, the integrals are replaced by sum over 100 grid points.

I Priors, Posteriors, and Selection Figures

I.1 Canonical New Keynesian Model

Table A-2: PRIOR DISTRIBUTION: Canonical New Keynesian Model

Name	Density	Para (1)	Para (2)	Name	Density	Para (1)	Para (2)
r^A	Gamma	2.00	1.00	π^A	Normal	4.00	1.00
γ^Q	Normal	0.50	0.10	κ	Gamma	0.05	0.10
σ	Gamma	2.00	0.50	ζ	Uniform	0.00	1.00
a	Uniform	0.00	1.00	Φ_π	Gamma	1.50	0.25
Φ_y	Gamma	0.25	0.25	σ_ξ	Inv. Gamma	1.00	4.00
σ_y	Inv. Gamma	1.00	4.00	σ_i	Inv. Gamma	1.00	4.00
ρ_ξ	Uniform	0.00	1.00	ρ_i	Uniform	0.00	1.00
ρ_y	Uniform	0.00	1.00	$\sigma_F\pi$	Inv. Gamma	0.10	4.00

Notes: Para (1) and Para (2) correspond to the mean and standard deviation of the Beta, Gamma, and Normal distributions and to the upper and lower bounds of the support for the Uniform distribution. For the Inv. Gamma distribution, Para (1) and Para (2) refer to s and ν, where $p(\sigma|\nu, s) \propto \sigma^{-\nu-1}e^{-\nu s^2/2\sigma^2}$.

Table A-3: POSTERIOR DISTRIBUTION: CANONICAL NEW KEYNESIAN MODEL

	With Expectations Data		Without Expectations Data	
r^A	2.43	[1.68, 3.17]	1.92	[1.10, 2.70]
π^A	3.93	[2.70, 5.27]	4.09	[2.47, 5.72]
γ^Q	0.28	[0.22, 0.34]	0.46	[0.40, 0.53]
κ	0.87	[0.55, 1.31]	0.00	[0.00, 0.00]
σ	1.60	[1.01, 2.30]	1.74	[1.08, 2.52]
ζ	0.26	[0.13, 0.38]	0.87	[0.77, 0.93]
a	0.97	[0.90, 1.00]	0.98	[0.93, 1.00]
Φ_π	3.14	[2.74, 3.59]	1.50	[1.10, 1.91]
Φ_y	0.06	[0.03, 0.09]	0.22	[0.17, 0.29]
σ_ξ	0.28	[0.24, 0.29]	2.56	[1.19, 4.67]
σ_y	0.86	[0.76, 0.91]	0.81	[0.54, 1.20]
σ_i	0.91	[0.78, 1.07]	0.46	[0.38, 0.56]
ρ_ξ	0.83	[0.78, 0.89]	0.55	[0.41, 0.68]
ρ_i	0.48	[0.42, 0.53]	1.00	[0.99, 1.00]
ρ_y	0.99	[0.98, 1.00]	0.96	[0.93, 0.98]
$\sigma_F\pi$	0.10	[0.09, 0.11]		
ρ_F	0.91	[0.85, 0.96]		

Figure A-1: Shock Decomposition: HNK Model

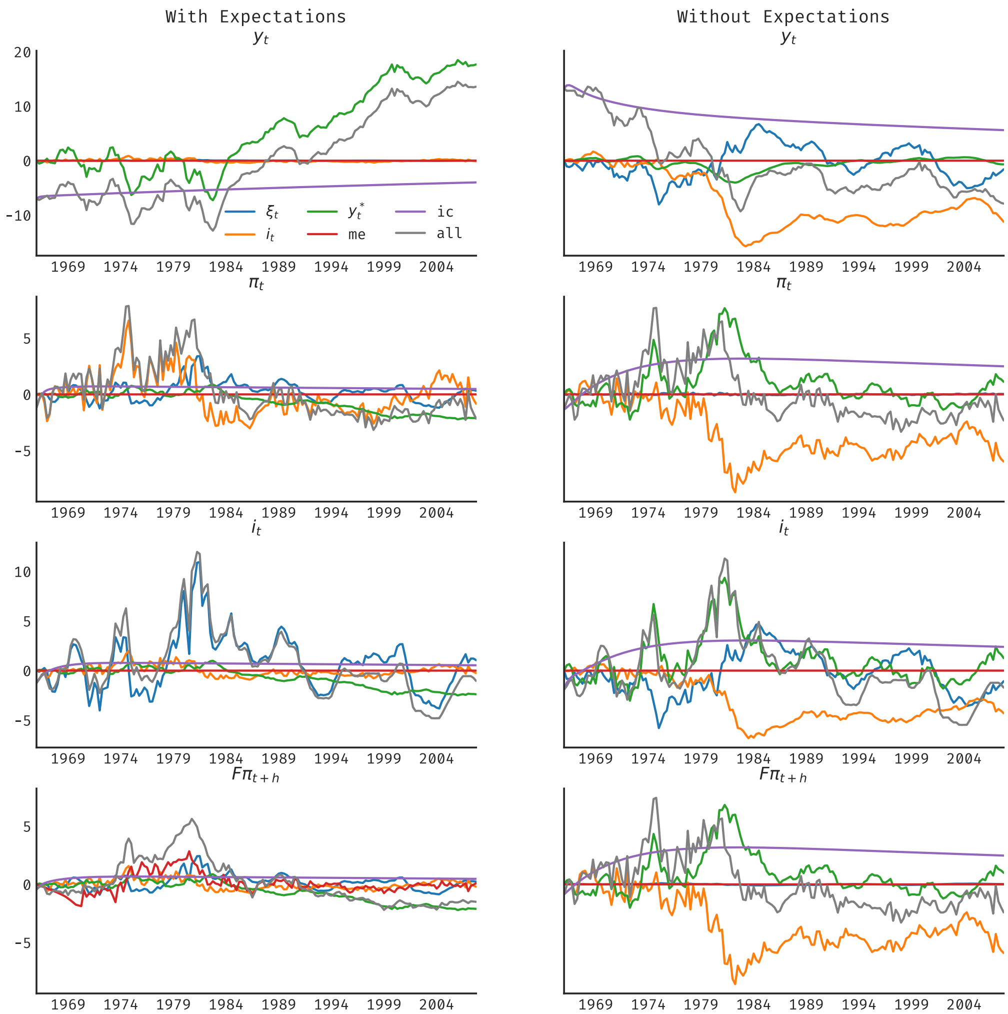

Figure A-2: Impulse Responses: HNK Model

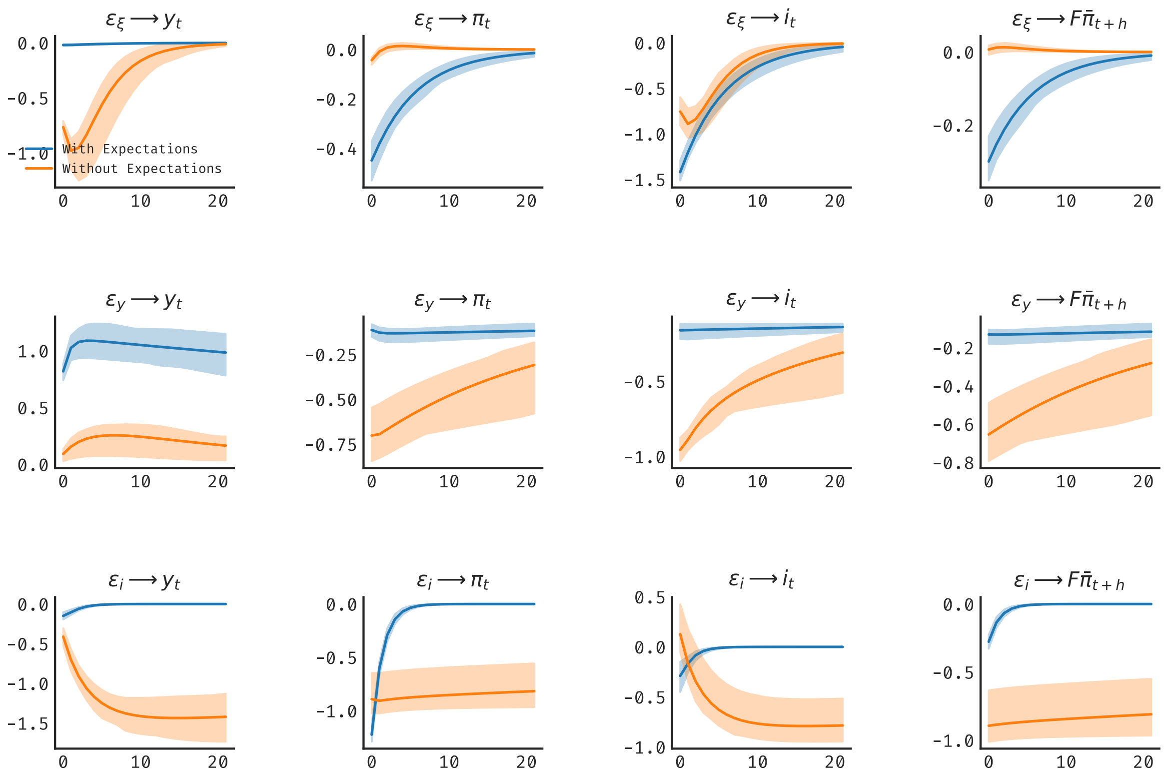

I.2 FHP Model

Table A-4: Prior Distribution: FHP Model

Name	Density	Para (1)	Para (2)	Name	Density	Para (1)	Para (2)
r^A	Gamma	2.00	1.00	π^A	Normal	4.00	1.00
γ^Q	Normal	0.50	0.10	ρ	Uniform	0.00	1.00
κ	Gamma	0.05	0.10	σ	Gamma	2.00	0.50
Φ_π	Gamma	1.50	0.25	Φ_y	Gamma	0.25	0.25
σ_ξ	Inv. Gamma	1.00	4.00	σ_y	Inv. Gamma	1.00	4.00
σ_i	Inv. Gamma	1.00	4.00	ρ_ξ	Uniform	0.00	1.00
ρ_i	Uniform	0.00	1.00	ρ_y	Uniform	0.00	1.00
γ	Uniform	0.00	1.00	γ_f	Uniform	0.00	1.00
$\bar{\phi}_\pi$	Gamma	1.50	0.25	$\bar{\phi}_y$	Gamma	0.25	0.25
$\sigma_F\pi$	Inv. Gamma	0.10	4.00	ρ_F	Uniform	0.00	1.00

Notes: Para (1) and Para (2) correspond to the mean and standard deviation of the Beta, Gamma, and Normal distributions and to the upper and lower bounds of the support for the Uniform distribution. For the Inv. Gamma distribution, Para (1) and Para (2) refer to s and ν, where $p(\sigma|\nu, s) \propto \sigma^{-\nu-1}e^{-\nu s^2/2\sigma^2}$.

Table A-5: Posterior Distribution: FHP Model

	With Expectations Data		Without Expectations Data	
r^A	2.13	[1.07, 3.19]	2.15	[1.11, 3.23]
π^A	3.72	[2.38, 5.10]	3.75	[2.37, 5.18]
γ^Q	0.44	[0.42, 0.46]	0.44	[0.42, 0.46]
ρ	0.50	[0.29, 0.68]	0.47	[0.20, 0.68]
κ	0.02	[0.01, 0.04]	0.03	[0.01, 0.05]
σ	2.67	[1.86, 3.60]	2.75	[1.93, 3.69]
Φ_π	0.85	[0.64, 1.11]	0.88	[0.66, 1.13]
Φ_y	0.64	[0.43, 0.92]	0.63	[0.42, 0.91]
σ_ξ	0.36	[0.31, 0.42]	0.36	[0.31, 0.43]
σ_y	8.88	[5.77, 13.76]	7.72	[4.21, 13.35]
σ_i	0.52	[0.39, 0.71]	0.52	[0.38, 0.71]
ρ_ξ	0.87	[0.79, 0.94]	0.87	[0.79, 0.93]
ρ_i	0.95	[0.91, 0.99]	0.95	[0.91, 0.99]
ρ_y	0.36	[0.30, 0.43]	0.46	[0.31, 0.61]
γ	0.48	[0.32, 0.65]	0.45	[0.30, 0.62]
γ_f	0.20	[0.16, 0.24]	0.22	[0.14, 0.30]
$\bar{\phi}_\pi$	1.88	[1.52, 2.28]	1.86	[1.49, 2.27]
$\bar{\phi}_y$	0.14	[0.04, 0.27]	0.13	[0.03, 0.26]
$\sigma_F\pi$	0.08	[0.07, 0.09]		
ρ_F	0.93	[0.89, 0.98]		

Figure A-3: Shock Decomposition: FHP Model

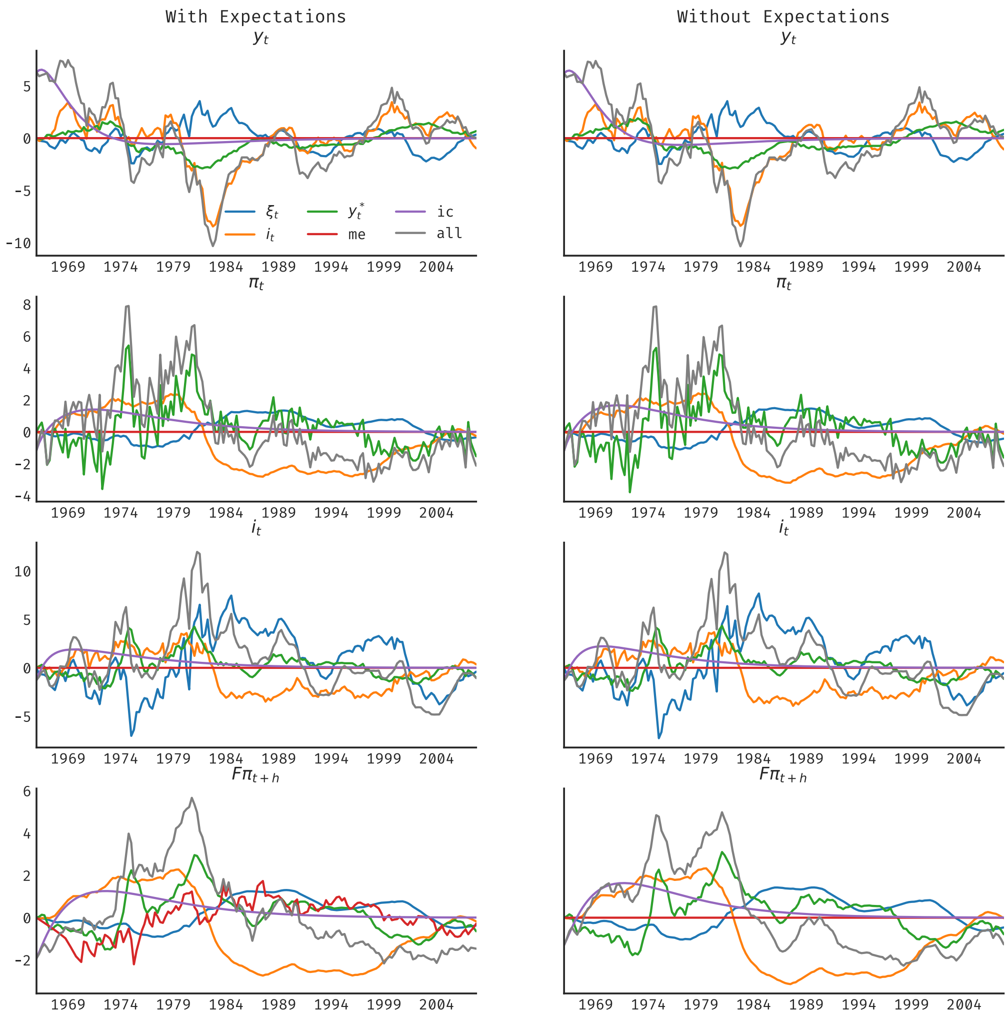

Figure A-4: Impulse Responses: FHP Model