Finance and Economics Discussion Series
Divisions of Research & Statistics and Monetary Affairs
Federal Reserve Board, Washington, D.C.

Short-term Planning, Monetary Policy, and Macroeconomic Persistence

Christopher Gust, Edward Herbst, David López-Salido

2020-003

Please cite this paper as:
Gust, Christopher, Edward Herbst, David López-Salido (2020). "Short-term Planning, Monetary Policy, and Macroeconomic Persistence," Finance and Economics Discussion Series 2020-003. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2020.003.

Short-term Planning, Monetary Policy, and Macroeconomic Persistence

Christopher Gust*　　Edward Herbst†　　David López-Salido‡

December 2019

Abstract

This paper uses aggregate data to estimate and evaluate a behavioral New Keynesian (NK) model in which households and firms plan over a finite horizon. The finite-horizon (FH) model outperforms rational expectations versions of the NK model commonly used in empirical applications as well as other behavioral NK models. The better fit of the FH model reflects that it can induce slow-moving trends in key endogenous variables which deliver substantial persistence in output and inflation dynamics. In the FH model, households and firms are forward-looking in thinking about events over their planning horizon but are backward looking regarding events beyond that point. This gives rise to persistence without resorting to additional features such as habit persistence and price contracts indexed to lagged inflation. The parameter estimates imply that the planning horizons of most households and firms are less than two years which considerably dampens the effects of expected future changes of monetary policy on the macroeconomy.

JEL Classification: C11, E52, E70
Keywords: Finite-horizon planning, learning, monetary policy, New Keynesian model, Bayesian estimation.

Special thanks to Mike Woodford and George-Marios Angeletos for very useful comments and suggestions. We also thank Marty Eichenbaum, Larry Christiano, Jordi Galí, Ben Johannsen, Mark Gertler and participants at the NBER-SI 2019 for comments. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of anyone else associated with the Federal Reserve System.

*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

Macroeconomists have long understood the important role that expectations play in determining the effects of monetary policy. Although it is common to analyze the effects of monetary policy in models in which agents form expectations rationally, the role of this assumption has come under increased scrutiny. Motivated by experimental evidence on human judgement and limits to cognitive abilities, a growing literature has moved away from this assumption and developed behavioral macroeconomic models to examine the effects of monetary policy.[1] This literature has highlighted several advantages of these models, including that the effects of changes in future policy rates on the macroeconomy are more realistic in these models. More specifically, advocates of behavioral macro models point to a "forward guidance puzzle" in New Keynesian (NK) models with rational agents because a credible promise to keep the policy rate unchanged in the distant future produces counterfactual large effects on current inflation and output. In contrast, NK models in which agents' expectations are consistent with behavioral evidence do not display such a puzzle.[2] From this standpoint, behavioral macro models are a promising alternative to those with rational expectations. Nonetheless, it remains an open question whether these models can be developed into empirically-realistic ones capable of providing guidance to monetary policy on a broad range of issues.

In this paper, we take a step towards addressing this question by estimating several New Keynesian (NK) models with behavioral features and assessing their ability to account for fluctuations in inflation, output, and interest rates in the United States. Our analysis suggests that the finite-horizon (FH) approach developed in a recent contribution by Woodford (2018) is a promising framework for explaining aggregate data and analyzing monetary policy. A chief advantage of the FH approach that we identify in the aggregate data is its ability to deliver an explanation for the persistence observed in aggregate output and prices. This ability hinges in the behavioral assumptions that underlie the way households and firms plan about their future decisions. In particular, in this model households and firms are forward-looking in thinking about events over their planning horizon but are backward looking regarding events beyond that point. This generates sluggishness in their consumption and price decisions that translate into output and price persistence without resorting to additional features such as habit persistence and price contracts indexed to lagged inflation.

As argued in Schorfheide (2013), one of the key challenges in developing empirically-realistic

[1] Recent contributions include Gabaix (2018), Garcia-Schmidt and Woodford (2019) and Farhi and Werning (2018), and Angeletos and Lian (2018). For a broader discussion of behavioral macroeconomics, see De Grauwe (2012). This literature is closely related to earlier work in models with boundedly rational agents; see, for example, Sargent (1993) and Evans and Honkapohja (2001).

[2] See Del Negro, Giannoni, and Patterson (2012) and McKay, Nakamura, and Steinsson (2016) for a discussion of the forward guidance puzzle. While the behavioral NK literature has emphasized the importance of incorporating boundedly rational agents into monetary models, others have emphasized the assumption that households and firms may not view promises about future rates as perfectly credible. In an estimated model, Gust, Herbst, and Lopez-Salido (2018), for example, show that imperfect credibility could have been an important reason why the Federal Reserve's forward guidance was less effective than otherwise.

macroeconomic models is to explain the substantial low frequency variation in macroeconomic data and produce accurate inference about business cycle fluctuations. A number of researchers have attempted to address this issue by incorporating shock processes into models to capture movements in trends; however, under this approach most of the persistence observed in aggregate data remains largely exogenous.[3] In contrast, in the finite-horizon approach of Woodford (2018), cyclical fluctuations as well as low frequency variation in aggregate prices and quantities depend on agents' planning horizon. In a nutshell, the key contribution of this paper is to empirically evaluate this feature of the model. To understand how Woodford (2018)'s FH model gives rise to a theory by which the cycle contributes to slow-moving trends, it is useful to review key features of his framework.

The backbone of the model is still New Keynesian, as monopolistically-competitive firms set prices in a staggered fashion and households make intertemporal consumption and savings decisions. Households and firms, as in a standard NK model, are infinitely-lived and need to look into the far distant future to make their current decisions. In doing so, households and firms are still quite sophisticated in that their current decisions involve making forecasts and fully-state contingent plans over their finite planning horizons. But, beyond their planning horizon, households and firms are less sophisticated, as a key assumption of the FH approach is that households and firms are boundedly rational regarding the continuation values of their plans over their infinite lifetimes. In particular, instead of viewing the value functions that govern these continuation values as fully state-contingent—as it would be if their expectations were fully rational, these continuation values are assumed to be coarser in their state dependence. Moreover, households and firms do not fully use the relationships in the model to infer these continuation values; instead, they update them, based on past data that they observe. In this way, households' and firms' learn about their continuation values as the economy evolves.[4]

Because of this decision-making process, households and firms who are forward-looking in thinking about events over their planning horizon are also backward looking in thinking about events beyond that point. If the planning horizons of households and firms becomes very long, the dynamics of the FH model mimic those of a standard NK model so that the backward-looking behavior of households and firms becomes irrelevant. However, if a significant fraction of agents have short planning horizons, the dynamics of the model are notably different from those of a more standard NK model. In particular, changes in future policy rates are not as effective in influencing current output, and future changes in output also have a much smaller effect on current inflation.

Most importantly, when a material fraction of agents have short-planning horizons, the model is capable of generating persistence endogenously through the way agents update their beliefs about the continuation values to their plans. Because of this feature, the model's equilibrium conditions can be decomposed into a cyclical component governed by agents' forward looking behavior and a

[3] See Canova (2014) for a discussion of the issue and approaches in which the trends are modeled exogenously and independently from the structural model used to explain business cycle fluctuations.

[4] Woodford (2018) motivates such decision making based on the complex intertemporal choices made by sophisticated artificial intelligence programs.

trend component explained by the way agents update the continuation values to their plans. Because agents update their beliefs in a backward-looking manner, the model is capable of generating substantial persistence in output and inflation. For instance, in line with empirical evidence, we show that the model is capable of generating substantial inflation persistence and a hump-shaped response of output following a monetary policy shock. Notably, it does so in the absence of incorporating habit persistence in consumption, price-indexation contracts tied to lagged inflation, or adjustment costs to investment.[5]

In our empirical analysis, we employ Bayesian methods to estimate the FH model as well as other behavioral macro models using U.S. quarterly data on output growth, inflation, and interest rates from 1966 until 2007. Besides comparing the FH model's performance to other behavioral macro models, we also compare its performance to a hybrid NK model that incorporates habit persistence and price-indexation contracts tied to lagged inflation. Because there is notable low frequency variation in the variables over the sample period that we estimate, we also compare the FH model's ability to fit the data relative to a NK model in which we introduce exogenous and separate processes to approximate the trends in output, inflation, and nominal short-term interest rates.

Regarding the estimation of the FH model, we find that we can reject parameterizations in which there is a considerable fraction of agents with long planning horizons including the standard NK model in which agents are purely forward looking. Our mean estimates suggest that about 50 percent of households and firms have planning horizons that include only the current quarter, 25 percent have planning horizons of two quarters, and only a small fraction have a planning horizon beyond 2 years. Thus, our estimates imply that there is a substantial degree of short term planning. Our evidence is also consistent with agents updating their value functions slowly in response to recently observed data so that the model's implied trends also adjust slowly. We show that because of this feature the model can account for the substantial changes that occurred to trend inflation and trend interest rates in the 1970s and 1980s. Interestingly, as a form of external validation, we show that the model's measure of trend inflation displays similar movements to a measure of longer-term inflation expectations coming from the Survey of Professional Forecasters.

We also show that the FH model fits the observed dynamics of output, inflation, and interest rates better than the hybrid NK model. This reflects both the endogenous persistence generated by agents' learning about their value functions as well as the reduced degree of forward-looking behavior associated with short-term planning horizons. Because of this model's ability to generate slow moving trends, its "goodness of fit" measure is substantially better than that of behavioral macro models of Angeletos and Lian (2018) and Gabaix (2018). Surprisingly, the FH model fit is nearly as good as the fit of a NK model that incorporates exogenous and separate trends in output, inflation, and interest rates. This is true despite the considerable flexibilty of the model with exogenous trends relative to the FH model. Finally, we show that the heterogeneity in plan-

[5]These mechanisms are often described as forms of generating "intrinsic persistence" in output and inflation (e.g., Smets and Wouters (2007) and Christiano, Eichenbaum, and Evans (2005)). Sims (1998) is an important earlier contribution to the discussion of issues related to modeling persistence in the context of macroeconomic models.

ning horizons embodied in the FH model is not necessary to track the aggregate variables in our estimation. Indeed, a simplified version of the model, where all agents have a one-period planning horizon, slightly outperforms the version of the model in which agents have heterogenous planning horizons. Overall, we view these results as suggesting that the FH framework offers a parsimonious and fruitful way for understanding monetary policy and business cycle dynamics.

The rest of the paper is structured as follows. The next section describes the FH model of Woodford (2018) paying particular attention to the role of monetary policy and the model's trend-cycle decomposition. Section 3 analyzes the dynamic properties of the model further and shows that the model is capable of generating realistic dynamics following a monetary policy shock. Section 4 discusses the data and methodology we use to estimate the model, while Section 5 presents the estimation results including the role of demand, supply, and monetary shocks in accounting for trend and cyclical movements in output and inflation. Section 6 compares the fit of the FH model to the other models that we estimate. Section 7 concludes and offers directions for further research.

2 An NK Model with Finite-Horizon Planning

We now present a description of the key structural relationships of the finite-horizon model that we estimate. The derivation of these expressions can be found in Woodford (2018).

To help motivate the finite-horizon approach, it is helpful to first review the structural relationships from the canonical NK model.[6] In that model, aggregate output y_t and inflation π_t (expressed in log-deviations from steady state) evolve according to the following expressions:

$$\begin{aligned} y_t - \xi_t &= E_t[y_{t+1} - \xi_{t+1}] - \sigma\left[i_t - E_t(\pi_{t+1})\right] && (1) \\ \pi_t &= \beta E_t[\pi_{t+1}] + \kappa(y_t - y_t^*) && (2) \end{aligned}$$

where E_t denotes the model-consistent expectations operator conditional on available information at time t, ξ_t is a demand or preference shock and y_t^* is exogenous and captures the effects of supply shocks. The parameters β, σ, and κ are the discount factor, the inverse of the household's relative risk aversion, and the slope of the inflation equation with respect to aggregate output. The parameter κ itself is a function of structural model parameters including the parameter governing the frequency of price adjustment and the elasticity of output to labor in a firm's production function. To close the model, a central bank is assumed to follow an interest-rate (i_t) policy rule:

$$i_t = \phi_\pi \pi_t + \phi_y y_t + i_t^*, \tag{3}$$

where $\phi_\pi > 0$, $\phi_y > 0$, and i_t^* as an exogenous monetary policy surprise. These three equations can be used to characterize the equilibrium for output, inflation and the short-term interest rate in the canonical NK model.

The finite-horizon model in Woodford (2018) maintains two key ingredients of the canonical

[6] See Woodford (2003) or Galí (2008) for the derivations of the canonical NK model.

model. In particular, monopolistically-competitive firms set prices in a staggered fashion according to Calvo (1983) contracts and households make intertemporal choices regarding consumption and savings. However, the finite-horizon approach departs from the assumption that households and firms formulate complete state-contingent plans over an infinite-horizon. Instead, infinitely-lived households and firms make state-contingent plans over a *fixed* $k-$*period horizon* taking their infinite-horizon continuation values as given. While households and firms are sophisticated about their plans over this fixed horizon, they are less sophisticated in thinking about the continuation value to their plans. In particular, Woodford (2018) assumes that agents are not able to use their model environment to correctly deduce their value functions and how they differ across each possible state. Instead, the value function is coarser in its state dependence. Agents update their beliefs about their value functions as they gain information about them as the economy evolves.

This assumption introduces a form of bounded rationality in which agents choose a plan at date t over the next k periods but only implement the date t part of the plan. To make their decisions about date t variables, households and firms take into account the state contingencies that could arise over the next k periods, working backwards from their current beliefs about their value-functions.[7] In period $t+1$, an agent will not continue with the plan originally chosen at time t but will choose a new plan and base their time $t+1$ decisions on that revised plan. An agent will also not necessarily use the same value-function that she used at date t, as an agent may update her value function for decisions at date $t+1$.

The model allows for heterogeneity over the horizons with which firms and households make their plans. In the presence of this heterogeneity, Woodford (2018) is able to derive a log-linear approximation to the finite-horizon model whose aggregate variables evolve in a manner resembling the equilibrium conditions of the canonical NK model. In particular, aggregate output and inflation satisfy:

$$y_t - \xi_t - \overline{y}_t \quad = \quad \rho E_t[y_{t+1} - \xi_{t+1} - \overline{y}_{t+1}] - \sigma\left[i_t - \overline{i}_t - \rho E_t(\pi_{t+1} - \overline{\pi}_{t+1})\right] \tag{4}$$

$$\pi_t - \overline{\pi}_t \quad = \quad \beta\rho E_t[\pi_{t+1} - \overline{\pi}_{t+1}] + \kappa(y_t - y_t^* - \overline{y}_t) \tag{5}$$

Two elements stand out about aggregate dynamics of the finite-horizon model. First, there is an additional parameter, $0 < \rho < 1$, in front of the expected future values for output and inflation. Second, aggregate output and inflation are written in deviations from endogenously-determined “trends”; the trends vary over time and are represented by a “bar” over a variable.

Finally, monetary policy responds to the deviation of inflation and output from their trends and allows for a time-varying intercept ($\overline{i}_t$):

$$i_t - \overline{i}_t \quad = \quad \phi_\pi(\pi_t - \overline{\pi}_t) + \phi_y(y_t - \overline{y}_t) + i_t^* \tag{6}$$

We discuss each of these elements in more detail below.

[7] Woodford (2018) motivates this approach based on sophisticated, artificial intelligence programs constructed to play games like chess and go.

2.1 Microeconomic Heterogeneity and Short-term Planning

The parameter ρ is an aggregate parameter reflecting that planning horizons differ across households and firms. To understand this, let ω_j and $\widetilde{\omega}_j$ be the fraction of households and firms, respectively, that have planning horizon j for $j = 0, 1, 2, ...$; the sequences of $\omega's$ satisfy $\sum_j \omega_j = \sum_j \widetilde{\omega}_j = 1$. The parameter ρ also satisfies $\omega_j = \widetilde{\omega}_j = (1-\rho)\rho^j$ where $0 < \rho < 1$. Aggregate spending and inflation are themselves the sum of spending and pricing decisions over the heterogeneous households and firms. As a result, $y_t = \sum_j (1-\rho)\rho^j y_t^j$ and $\pi_t = \sum_j (1-\rho)\rho^j \pi_t^j$, where y_t^j denotes the amount of spending of a household with planning horizon j and π_t^j denotes the inflation rate set by a firm with planning horizon j.[8]

The parameter ρ governs the distribution of planning horizons agents have in the economy and has important implications for aggregate dynamics. A relatively low value of ρ implies that the fraction of agents with a short planning horizon is relatively high. And, as a consequence, the dynamics characterizing aggregate output and inflation are less "forward-looking" than in the canonical model. When $\rho \to 1$, there are an increasing number of households and firms with long planning horizons and the aggregate dynamics become like those of the canonical model in which agents have rational expectations. Because of its prominent role in affecting the cyclical component of aggregate dynamics, one aim of this paper is to estimate the value of ρ and see how much short-term planning by households and firms is necessary to explain the observed persistence in output, inflation, and interest rates.

Woodford (2018) also emphasizes the important role that $\rho < 1$ plays in overcoming the forward guidance puzzle inherent in the canonical NK model – i.e., the powerful effects on current output and inflation of credible promises about future interest rates. To understand how this works, equation (4) can be rewritten as:

$$\widetilde{y}_t = -\sigma \sum_{s=0}^{\infty} \rho^s E_t(\widetilde{i}_{t+s} - \widetilde{\pi}_{t+s+1}) - \sigma(1-\rho) \sum_{s=0}^{\infty} \rho^s E_t \widetilde{\pi}_{t+s+1} \tag{7}$$

where the symbol " $\sim$ " represents the cyclical component of the variables (i.e., the value of the variable in deviation from its trend: $\widetilde{x}_t = x_t - \overline{x}_t$).[9] The expression above differs in two important ways from the aggregate output equation in the canonical NK. Current (cyclical) output depends on the "discounted future" path of the (cyclical) short-term real rates, and the geometric weights of future cyclical rates on cyclical output are a function of the parameter ρ. In particular, the effect on cyclical output from a change in the cyclical real rate in period $t+s$ is given by $-\sigma\rho^s$. With $\rho < 1$, a near-term change in the real rate has a larger effect on cyclical output than a longer-run change. In contrast, in the canonical NK model in which $\rho = 1$, there is no difference in the effect

[8]With an infinite number of household and firm types, the existence of the equilibrium requires that these (infinite) sums converge.

[9]All cyclical variables are defined in this way except for cyclical output, which is defined as $\widetilde{y}_t = y_t - \xi_t - \overline{y}_t$.

of a near-term change and one in the far future.[10]

The second term on the right-hand side of expression (7) reflects that, as noted in Woodford (2018), the Fisher equation does not hold in the short run. For the cyclical variables, short-run planning horizons introduces a form of "money illusion" in which higher expected inflation relative to trend, holding the (cyclical) real rates constant, reduces cyclical output.[11] However, the Fisher equation does hold in the long run. In particular, once the response of the trends is incorporated into the analysis, a permanent increase in inflation leads to a permanently higher nominal rate, leaving the level of output unchanged in the long run. The next section describes how the model's trends are determined.

2.2 A Theory-Based Trend-Cycle Decomposition

Expressions (4) and (5) describe the evolution of aggregate output and inflation in deviation from trend. The trend variables in equations (4)-(6) themselves are in deviation from nonstochastic steady state so that these trends can reflect very low frequency movements in these variables that are not associated with a change in the model's steady state. Instead, these time-varying trends reflect changes in agents' beliefs about the longer-run continuation values of their plans. Because agents update their continuation values based on observation of past data, this updating can induce persistent trends in output, inflation, and the short-term interest rate. We now turn to discussing how the finite-horizon approach leads to such a trend-cycle decomposition.

To understand how agents in the model parse trend from cycle, it is necessary to describe the value functions of households and firms. In the case of a household, its value function depends on its wealth or asset position. The derivative of the value function with respect to wealth is a key determinant of their optimal decisions. This derivative determines the marginal (continuation) value to a household of holding a particular amount of wealth. Unlike under rational expectations, this function is not fully state-contingent and it is assumed that agents can not deduce it using the relationships of the model. Instead, they update the parameters governing the marginal value of wealth based on past experience. More specifically, in log-linear form the marginal value of wealth consists of an intercept term and a slope coefficient on household wealth. Under constant-gain learning, Woodford (2018) proves that only the intercept-term needs to be updated, as the slope coefficient can be shown to converge to a constant. Constant-gain learning implies that a household

[10] A similar property holds for equation (5) which can be rewritten as:

$$\widetilde{\pi}_t = \kappa \sum_{s=0}^{\infty} (\beta\rho)^s \, E_t \widetilde{y}_{t+s}$$

where $\beta\rho < \beta$. Accordingly, in the FH horizon, the effects of future changes in the output gap on cyclical inflation can in principle be much smaller than in the canonical NK model.

[11] This effect is related to the one discussed in Modigliani and Cohn (1979). In their model, agents do not distinguish correctly between real and nominal rates of return and mistakenly attribute a decrease (increase) in inflation to a decline (increase) in real rates. See also Brunnermeier and Julliard (2008).

updates the intercept-term of her marginal value of wealth according to:

$$v_{t+1} = (1-\gamma)v_t + \gamma v_t^{est}, \tag{8}$$

where v_t denotes the (log-linearized) intercept of the marginal value of wealth at date t. Also, the constant-gain parameter, γ, satisfies $0 < \gamma < 1$ and determines how much weight a household put on the current estimate of this intercept in the updating step. The variable v_t^{est} denotes an updated (estimate for the) intercept that a household computes based on information acquired at date t. Through recursive substitution of expression (8), one can see that v_t, is a weighted sum of all the past values of v_t^{est} with distant past values getting more weight the larger is $(1-\gamma)$. As shown in Woodford (2018), up to first order, this new continuation value has an intercept term satisfying:

$$v_t^{est} = y_t - \xi_t + \sigma\pi_t. \tag{9}$$

Thus, the updated intercept term depends on current spending and current inflation as well as the shock to preferences. Combining this expression with equation (8) yields an expression in which the marginal value of wealth depends on all past values of y_t, ξ_t, and π_t.

We turn now to the discussion of how firms update their value functions. Each period only a fraction $1-\alpha$ of firms have the opportunity to reset its (relative) price. Accordingly, a firm that has the opportunity to do so at date t maximizes its expected discounted stream of profits taking into account that it may not have the opportunity to re-optimize its price in future periods. A firm that can re-optimize its price at date t only plans ahead for a finite number of periods and evaluates possible situations beyond that point with its value function.

As was the case for households, this value function is not fully state-contingent and it is assumed that firms can not deduce it using the relationships of the model. Instead, the firm updates it based on past experience. Specifically, a firm's first order conditions for its optimal price depend on the derivative of a firm's value-function with respect to its relative price, or in other words the marginal continuation value associated with its price. The (marginal) continuation value for a firm is a function that in its log-linear form has an intercept term and a slope coefficient with respect to a firm's relative price. Similar to households, it is sufficient only to update the intercept term, $\tilde{v}_t$, which evolves according to:

$$\tilde{v}_{t+1} = (1-\tilde{\gamma})\tilde{v}_t + \tilde{\gamma}\tilde{v}_t^{est}, \tag{10}$$

where $\tilde{\gamma}$ is the constant-gain learning parameter and $\tilde{v}^{est}$ is a new estimate of a firm's marginal continuation value.

Taking its current continuation value as given, a firm's objective function can be differentiated with respect to a firm's price to determine a new estimate of the marginal continuation value-function. Woodford (2018) shows that a first order approximation of a firm's marginal continuation value satisfies:

$$\tilde{v}_t^{est} = (1-\alpha)^{-1}\pi_t, \tag{11}$$

so that the new estimate depends on the average duration of a price contract, $(1-\alpha)^{-1}$, as well as the average inflation rate.

Because the (longer-run) continuation values of households and firms reflect averages of past values of spending and inflation, they can induce slow moving trends in these aggregate variables. More concretely, the finite-horizon model implies that output and inflation can be decomposed into a cyclical component (denoted using a tilde) and trend component (denoted using a bar) so that $y_t = \widetilde{y}_t + \overline{y}_t$ and $\pi_t = \widetilde{\pi}_t + \overline{\pi}_t$. The trend components represent how the spending and pricing decisions are affected by v_t and $\tilde{v}_t$, while the cyclical component represents these decisions in the absence of any changes in v_t and $\tilde{v}_t$, respectively. Because a household is (still) forward-looking, its plan for spending in future periods as well as the plan's continuation value matter for its current spending decision. Similarly, a firm with the opportunity to reset its price is forward looking so that $\tilde{v}_t$ matters for that decision.

Formally, averaging across the different household types, Woodford (2018) shows that the effect of v_t on aggregate spending is given by:

$$\bar{y}_t = \frac{-\sigma}{1-\rho}\left(\bar{\imath}_t - \rho\bar{\pi}_t\right) + v_t, \tag{12}$$

where $\bar{\imath}_t$ is the trend interest rate discussed further below. Similarly, averaging across firms with different planning horizons, the effect of $\tilde{v}_t$ on average price inflation is given by:

$$\bar{\pi}_t = \frac{\kappa}{1-\beta\rho}\bar{y}_t + \frac{(1-\rho)(1-\alpha)\beta}{1-\beta\rho}\tilde{v}_t. \tag{13}$$

Holding fixed the trend interest rate, equations (12) and (13) relate trend output and inflation to the longer-run continuation values of households and firms. These continuation values, as reflected in v_t and $\tilde{v}_t$, in turn depend on the entire past history of aggregate spending and inflation and thus the model is capable of generating substantial persistence in output and inflation trends. Importantly, as indicated by the presence of $\bar{\imath}_t$ in equation (12), trend output and inflation depend on agents' views about monetary policy, which we now specify.

2.3 Monetary Policy

Given our focus on the empirical performance of this model, we extend Woodford (2018) by allowing for the monetary policy rule to have a time-varying intercept that depends on agents' beliefs about longer-run trends.[12] In particular, the intercept term in equation (6) is specified as:

$$\bar{\imath}_t = \overline{\phi}_\pi \overline{\pi}_t + \overline{\phi}_y \overline{y}_t. \tag{14}$$

[12] We also estimated the model with a policy rule that allowed the policy rate to depend on its lag instead of using equation (3). The fit of the model deteriorated noticeably relative to including a time-varying intercept that varies with agents' perceptions of longer-run inflation.

We view this time-varying intercept as capturing two aspects of monetary policy. First, monetary policymakers do not view the 'equilibrium' or longer-run real interest rate as a constant. Instead, they see it as time-varying and view agents' beliefs about trend output as useful in determining its value. Second, by allowing the time-varying intercept to depend on agent's beliefs about trend inflation, we allow for the possibility that monetary policy responds more aggressively to deviations from the inflation target than agents believe will persist (i.e., $\overline{\pi}_t$) than to deviations from the inflation target that those deviations they view as less persistent (i.e., $\tilde{\pi}_t$). In that case, $\overline{\phi}_\pi > \phi_\pi$, and in our empirical analysis, we evaluate whether such a response is a better characterization of monetary policy than the case in which monetary policy responds equiproportionately so that $\overline{\phi}_\pi = \phi_\pi$.

In our empirical analysis, we assess the performance of the model both when the intercept is time-varying according to equation (14) and when $\phi_\pi = \overline{\phi}_\pi$ and $\phi_y = \overline{\phi}_y$ so that the intercept is a constant. In the latter case, the rule is the same as the one in Woodford (2018) and is given by equation (3), where the policy rate simply depends on the deviations of aggregate output and inflation from their steady state.

3 Short-Term Planning and Macroeconomic Persistence

In this section, we investigate the model's trend-cycle decomposition more thoroughly and show how it induces persistent movements in output and inflation following a monetary policy shock. We begin by showing that in the finite-horizon model, cyclical fluctuations are independent from the trend. However, the trends depend on the cycle and thus on monetary policy.

3.1 Trend-Cycle Decomposition and Monetary Policy

To see that the cycle is independent of the trend, note that equations (4)-(6) are block recursive when we express output, inflation, and the policy rate as deviations from trends. Specifically, after substituting out the policy rate deviation using the interest-rate rule, the remaining two equations yield:

$$\tilde{x}_t = \rho M \cdot E_t[\tilde{x}_{t+1}] + N \cdot u_t, \tag{15}$$

where $\tilde{x}_t = (\tilde{y}_t - \xi_t, \tilde{\pi}_t)'$ and $u_t = (i_t^* + \phi_y \xi_t, \xi_t - y_t^*)'$ Also, M and N are 2-by-2 matrices whose elements depend on the model's structural parameters including the rule parameters, ϕ_π and ϕ_y. (The appendix shows the elements of M and N as a function of the model's parameters.) This system can be used to solve for the cyclical variables, $\tilde{x}_t$, as a function of the economy's shocks, u_t, independently of the trends for output, inflation, or the policy rate. As a result, the cyclical variables do not depend on the long-run response of monetary policy to the trends (i.e., $\bar{\phi}_\pi$ and $\bar{\phi}_y$).

The trends, however, depend on the cycle. To see that, expressions (12) and (13) can be used to solve for $\overline{y}_t$ and $\overline{\pi}_t$ as a function of v_t and $\tilde{v}_t$:

$$\overline{x}_t = (1-\rho)\Theta V_t, \tag{16}$$

where we have substituted out $\bar{i}_t$ using equation (14), $\overline{x}_t = (\overline{y}_t, \overline{\pi}_t)'$, and $V_t = (v_t, \tilde{v}_t)'$. The 2-by-2 matrix, Θ, is shown in the appendix and depends on structural model parameters that include $\bar{\phi}_\pi$ and $\bar{\phi}_y$. Thus, since monetary policy affects agents' longer-run continuation values, the trends for output and inflation depend on how monetary policy reacts to their movements.

To express the trends, $\overline{x}_t$, as a function of the cycle, it is convenient to rewrite the laws of motion for the intercepts of the marginal value-function as:

$$V_t = (I-\Gamma)V_{t-1} + \Gamma\Phi x_{t-1}, \tag{17}$$

where $x_t = (y_t - \xi_t, \pi_t)'$, and Γ and Φ are 2-by-2 matrices shown in the appendix. Importantly, they do not depend on the monetary policy rule parameters. Combining expression (17) with equation (16) yields:

$$\overline{x}_t = \Lambda\overline{x}_{t-1} + (1-\rho)\gamma Q x_{t-1}, \tag{18}$$

where $\Lambda = \Theta(I-\Gamma)\Theta^{-1}$ and $Q = \Theta\Gamma\Phi$ are also 2-by-2 matrices shown in the appendix. Using $\widetilde{x}_t = x_t - \overline{x}_t$, we can rewrite this expression so that the trends for output and inflation are a function of the past cyclical values for these variables:

$$\overline{x}_t = \left[\Lambda + (1-\rho)\gamma Q\right]\overline{x}_{t-1} + (1-\rho)\gamma Q\tilde{x}_{t-1} \tag{19}$$

The aggregate equilibrium consists of the forward-looking system given by expression (15) characterizing the cycle and a backward-looking system given by expression (19) characterizing the trends. Because the cycle is independent of agents' beliefs about the trends, one can determine the cycle by solving the system in expression (15) for $\tilde{y}_t$ and $\tilde{\pi}_t$ and then using these values to determine the trends using expression (19).

Discussion. So far, our analysis of the model's trend-cycle decomposition has followed Woodford (2018). Here we extend the analysis. First, while Woodford (2018) shows that the stability of the trends depends on $0 < \gamma < 1$ and $0 < \tilde{\gamma} < 1$, we show that a modified Taylor principle is necessary for stability of the forward-looking system. For the stability of the system given by expression (15), the standard Taylor principle needs to be modified. As shown in the appendix, the modified Taylor principle for the FH model is:

$$\left(\frac{1-\rho\beta}{\kappa}\right)\phi_y + \phi_\pi > \rho. \tag{20}$$

Accordingly, the canonical model is a special case in which $\rho \to 1$, and in general the Taylor principle is relaxed relative to the canonical model when agents have finite horizons (i.e. $\rho < 1$). Moreover, the Taylor principle depends on how policy responds in the short run and not on how policy responds to fluctuations in trends.

Second, in the appendix, we provide analytical expressions for the matrices, Λ and Q, allowing for a better understanding of the model's trend-cycle decomposition. Thus, from expression (19), it follows that the impact the cycle has on trend inflation depends on the planning horizon of agents, the speed at which they update their value functions, and how responsive policy is to movements in trend variables. As ρ increases toward one, agents have long planning horizons and the trends no longer depend on the cycle. In fact, the trends become constants at their steady state values and the model's cyclical dynamics mimic those of the canonical NK model.

Third, monetary policy has important implications for the dynamics of the trends. With $\gamma = \tilde{\gamma}$, households and firms update their value-functions at the same rate, the analysis simplifies considerably. As shown in the appendix, the feedback matrix Λ becomes a scalar, $1-\gamma$, and the matrix Q is independent of γ. Thus, from equation (19) follows that if agents update their value functions more quickly (i.e., the value of γ approaches one), then both trends become more responsive to cycles. From expression (19) it also follows that the "long-run monetary policy response coefficients" affect the persistence of these trends as well as the pass-through of the cycle through the matrix Q. To get some insights on the trend and cycle dynamics, the appendix shows that the matrix Q simplifies to:

$$Q = \frac{1}{\Delta}\begin{pmatrix} 1-\beta\rho & \sigma(1-\beta\overline{\phi}_\pi) \\ \kappa & \kappa\sigma + (1-\rho+\sigma\overline{\phi}_y)\beta \end{pmatrix} \tag{21}$$

where $\Delta = (1-\beta\rho)(1-\rho+\sigma\overline{\phi}_y) + \kappa\sigma(\overline{\phi}_\pi - \rho)$.

Our estimates imply that monetary policy responds aggressively to movements in trend inflation. In that case, the matrix Q implies that trend inflation becomes less sensitive to movements in cyclical inflation or output. Moreover, trend output also becomes less responsive to movements in cyclical output; however, trend output falls more in response to a cyclical increase in inflation for larger values of $\bar{\phi}_\pi$ for values of $\bar{\phi}_\pi > 1$.

These results highlight that the FH approach gives rise to a theory through which trend and cycle can be correlated. In the context of macroeconomic models, this idea has been considered in reduced-form econometric analysis since Nelson and Plosser (1982). In statistical models, allowing for such a correlation can make identification difficult without stark assumptions (i.e., independence of trend and cycle). In the finite-horizon approach, theoretical restrictions from the model preclude confounding of trend and cycle. In fact, the model's trend-cycle decomposition can be directly related to monetary policy and to assumptions about household and firm behavior. Moreover, the finite-horizon approach allows one to decompose the cycle and trend into structural shocks. However, it remains an open question how well such an approach can explain aggregate data. This is the key question that we investigate in our empirical analysis.

3.2 Dynamic Responses to a Monetary Policy Shock

An important feature of the model is its ability to generate endogenous persistence without any need for habit persistence or the indexation of inflation to past values of inflation. To illustrate this property, we examine the impulse responses to a shock that affects the monetary policy rule. This shock is assumed to follow an AR(1) process:

$$i_t^* = \rho_{i^*} i_{t-1}^* + \epsilon_{i,t} \tag{22}$$

We examine a policy tightening for three different parameterizations. In the first, $\rho = 1.0$, which corresponds to the forward-looking, canonical NK model in which the responses of the aggregate and cyclical variables are the same, and the model's trend corresponds to the nonstochastic steady state. In Figure 1, the canonical NK model's impulse responses are labelled "Forward". In the second and third parameterizations of the model, we set $\rho = 0.5$ which corresponds to 50 percent of households and firms doing their planning within the existing quarter, 25 percent of them doing it in two quarters, and only a small fraction – less than 0.5 percent – of households and firms having a planning horizon of two years or more. The second parameterization, labelled "Large gain" in Figure 1, sets $\gamma = 0.5$, which implies that households and firms put a relatively large weight on current observations in updating their value functions. The third parameterization, labelled "Small gain", is the same as the second one except that $\gamma = 0.05$. This value implies that current observations get a relatively small weight in the updating of agents' value functions.[13]

Figure 1 displays the impulse responses of output, y_t, inflation, π_t, and the short-term interest rates, i_t to a unit increase in $\epsilon_{i,t}$ at date 0. (All variables are expressed in deviation from their values in the nonstochastic steady state.) The first row in the figure corresponds to the responses of the aggregate variables, the second row to the trend responses, and the third row to the cyclical responses. As shown in the first row of the figure, a policy tightening results in an immediate fall in output of a little more than 2 percent and a 15 basis point fall in inflation in the canonical model (green lines). Thereafter, the responses of output and inflation converge back monotonically to their steady state values. This monotonic convergence entirely reflects the persistence of the shock. The middle and lower panels of the figure confirm that in the canonical model, there is no difference between the trends and steady state values of the model so that the aggregate and cyclical responses are the same.

The blue lines, labelled "Large Gain," in Figure 1 show the impulse responses in the finite horizon model in which agents heavily weigh recent data in updating their value functions. As in the canonical NK model, aggregate output and inflation fall on impact; however, the fall is dampened substantially. Moreover, output and inflation display hump-shaped dynamics despite the lack of indexation or habit persistence in consumption. While output reaches its peak decline after about a year, it takes substantially longer for inflation to reach its peak decline. As shown in the middle panel, these hump-shaped dynamics are driven by the gradual adjustment of the trends. The trend values for output and inflation fall in response to the policy tightening, reflecting that the policy shock persistently lower aggregate output and inflation. For output this return back to trend is relatively quick with a slight overshoot (not shown). However, the inflation trend returns

[13] For these three cases, we set the remaining parameters as follows: $\beta = 0.995$, $\sigma = 1$, $\kappa = 0.01$, $\phi_\pi = 1.5$, $\phi_y = \frac{0.5}{4}$, and $\rho_{i^*} = 0.85$.

back to its steady state very gradually as agents with finite horizons only come to realize slowly over time that the policy tightening will have a persistent effect on inflation.

The orange lines, labelled "Small Gain," show a similar parameterization except that agents update their value function even more slowly. In this case, the responses of the output and inflation trends is smaller and even more drawn out over time. Because of the dampened response of trend output, the response of aggregate output is no longer hump-shaped, as the aggregate effect is driven primarily by the monotonic cyclical response shown in the bottom left panel. In contrast, the aggregate inflation response is both dampened and more persistent. In sum, the finite horizon model is capable of generating substantial persistence in inflation and hump-shaped output responses following a monetary policy shock. Such dynamics are in line with empirical work examining the effects of monetary policy shocks on the macroeconomy.[14]

4 Estimation

4.1 Data and Methodology

We estimate several variants of the model 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.[15] The observation equations for the model are:[16]

$$\begin{aligned} \text{Output Growth}_t &= \mu^Q + y_t - y_{t-1} &(23)\\ \text{Inflation}_t &= \pi^A + 4 \cdot \pi_t &(24)\\ \text{Interest Rate}_t &= \pi^A + r^A + 4 \cdot i_t, &(25) \end{aligned}$$

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

The solution to the system of equations (15) and (19) jointly with these observations equations define the measurement and state transition equations of a linear Gaussian state-space system. The state-space representation of the DSGE model has 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 the interest is the posterior distribution of the parameters θ. The posterior distribution is calculated by combining the likelihood and prior

[14] See, for instance, Christiano, Eichenbaum, and Evans (2005) and the references therein.

[15] The appendix details the construction of this data.

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

Figure 1: Impulse Responses to an Unexpected Monetary Tightening

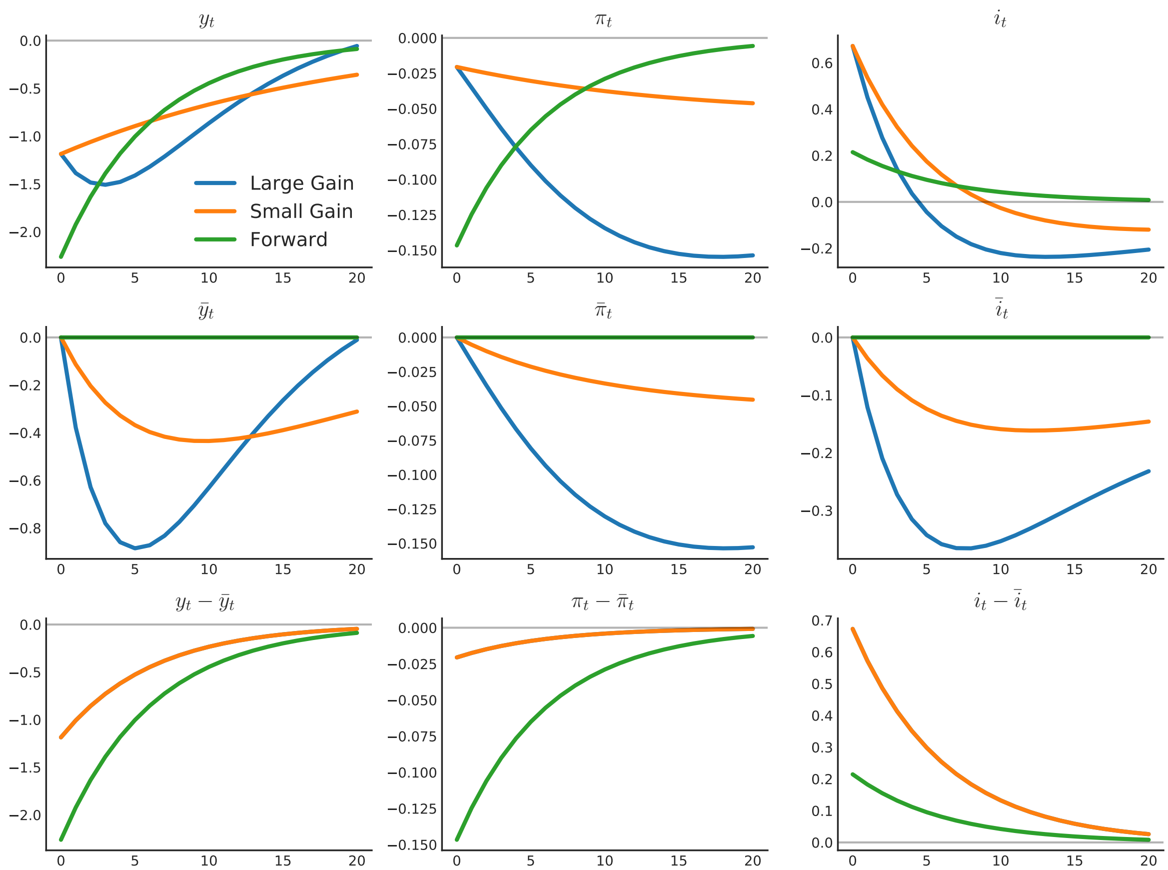

NOTE: The figure shows impulse responses to a monetary policy shock. In the Forward model (green lines), agents have infinite planning horizons ($\rho = 1.0$), and two in the two remaining models, agents have finite planning horizons ($\rho = 0.5$). The first of these models, Large Gain (blue lines), agents update their value function quickly, ($\gamma = 0.5$); in the second one, Small Gain (green lines), agents update their value function slowly ($\gamma = 0.05$).

distribution, $p(\theta)$, using Bayes theorem:

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

The prior distribution for the model's parameters is generated by a set of independent distributions for each of the structural parameters that are estimated. These distributions are listed in Table 1. For the shocks, we assume they follow AR(1) processes and use relatively uninformative priors regarding the coefficients governing these processes. Specifically, the monetary policy shock follows the AR(1) process given by equation (22) and the processes for the other two shocks are given by:

$$\xi_t = \rho_\xi \xi_{t-1} + \epsilon_{\xi,t} \tag{26}$$

$$y_t^* = \rho_{y^*} y_{t-1}^* + \epsilon_{y^*,t}. \tag{27}$$

The prior for each of the AR(1) coefficients is assumed to be uniform over the unit interval, while each of the priors for the standard deviations of shocks is assumed to be an inverse gamma distribution with 4 degrees of freedom.

The priors for the gain parameters, γ and $\tilde{\gamma}$, in the household's and firm's learning problems are also assumed to follow uniform distributions over the unit interval. Similarly, we assume that the prior distribution for the parameter governing the length of agents' planning horizons, ρ, is also a uniform distribution over the unit interval. The prior for r^A and π^A are chosen to be consistent with a 2% average real interest rate and 4% average rate of inflation. The prior of the slope of the Phillips curve, κ, is consistent with moderate-to-low pass through of output to inflation.[17] The prior for σ, the coefficient associated with degree of intertemporal substitution, follows a Gamma distribution with a mean of 2 and standard deviation of 0.5, and hence encompasses the log preferences frequently used in the literature. The prior distributions of the coefficients of the monetary policy rule, ϕ_π and ϕ_γ, are consistent with a monetary authority that responds strongly to inflation and moderately to the output gap and encompasses the parameterization in Taylor (1993).

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. The appendix provides information about the tuning parameters used to estimate the model as well as convergence diagnostics associated with the SMC algorithm.

4.2 Models

Table 2 displays the models that we estimate. These models differ in the restrictions on the parameters governing the length of the horizon, the parameters governing how quickly firms and households update their value functions, and the parameters in the reaction function for monetary policy.

The first model, referred to as "Forward" in Table 2, corresponds to the canonical New Keynesian model with three shocks, purely forward looking agents, and a Taylor-type rule for monetary policy. It is consistent with setting $\rho = 1$. Because the trends in this model are simply constants,

[17] The parameter κ is a reduced form parameter that is related to the fraction of firms that have an opportunity to reset their price, $1 - \alpha$, a parameter governing the elasticity of substitution for each price-setter's demand, θ, the elasticity of production to labor input, $\frac{1}{\phi}$, and the Frisch labor supply elasticity, ν. The mean value of our prior for κ is 0.05, which implies an $\alpha \approx \frac{1}{3}$ with $\nu = 1$, $\theta = 10$, and $\phi = 1.56$. Thus, the mean of the prior for κ is consistent with an average duration of a firm's price contract that is under one year.

Table 1: Prior Distributions

Parameter	Distribution		
	Type	Par(1)	Par(2)
r^A	Gamma	2	1
π^A	Normal	4	1
μ^Q	Normal	0.5	0.1
$(\rho, \gamma, \widetilde{\gamma})$	Uniform	0	1
σ	Gamma	2	0.5
κ	Gamma	0.05	0.1
ϕ_π	Gamma	1.5	0.25
ϕ_y	Gamma	0.25	0.25
$(\sigma_\xi, \sigma_{y^*}, \sigma_{i^*})$	Inv. Gamma	1	4
$(\rho_\xi, \rho_{y^*}, \rho_{i^*})$	Uniform	0	1

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

Table 2: Key Parameters of the Estimated Models

Model	Parameters		
Type	Estimated	Fixed	Not identified
Forward	ϕ_π, ϕ_y	$\rho = 1$	$\gamma, \widetilde{\gamma}, \overline{\phi}_\pi, \overline{\phi}_y$
Exog. Trends	AR(1) trends	$\rho = 1$	$\gamma, \widetilde{\gamma}, \overline{\phi}_\pi, \overline{\phi}_y$
FH-baseline	$\rho, \gamma, \phi_\pi, \phi_y$	$\gamma = \widetilde{\gamma}, \overline{\phi}_\pi = \phi_\pi, \overline{\phi}_y = \phi_y$	-
FH-$\widetilde{\gamma}$	$\rho, \gamma, \widetilde{\gamma}, \phi_\pi, \phi_y$	$\overline{\phi}_\pi = \phi_\pi, \overline{\phi}_y = \phi_y$	-
FH-$\overline{\phi}$	$\rho, \gamma, \phi_\pi, \phi_y, \overline{\phi}_\pi, \overline{\phi}_y$	$\gamma = \widetilde{\gamma}$	-

NOTE: This table presents the key parameters of the different estimated models.

we also consider a version of this model, "Exog. Trends," which allows for stochastic trends as in Canova (2014) and Schorfheide (2013). Specifically, with $\rho = 1$, we augment the model with three more shocks that allow the trends for output, inflation, and the nominal interest rate to evolve exogenously:[18]

$$\overline{y}_t = \rho_{\overline{y}} \overline{y}_{t-1} + \epsilon_{\overline{y},t} \tag{28}$$

$$\overline{\pi}_t = \rho_{\overline{\pi}} \overline{\pi}_{t-1} + \epsilon_{\overline{\pi},t} \tag{29}$$

$$\overline{i}_t = \rho_{\overline{i}} \overline{i}_{t-1} + \epsilon_{\overline{i},t}. \tag{30}$$

[18]The prior for the additional parameters is the product of six independent priors, with each autoregressive coefficient prior following a Beta distribution with mean 0.95 and standard deviation 0.05. The prior standard deviation of the innovations to these trends is and Inverse Gamma distribution with $s = 0.1$ and $\nu = 6$. This joint distribution is informative; it is consistent with the view that these trends are very persistent and that the magnitude of their innovations are small relative to the shocks of the model.

The remaining models in Table 2 are all different versions of the FH model. The first, referred to as "FH-baseline", estimates ρ and γ but assumes that the constant gain parameter, γ, is the same across households and firms. In addition, in this baseline version, the intercept term in the central bank's reaction function responds to trends in inflation and output in the same manner as it does to short-run cyclical fluctuations (i.e., $\bar{\phi}_\pi = \phi_\pi, \bar{\phi}_y = \phi_y$). The second variant of the FH model, referred to as "FH-$\widetilde{\gamma}$", allows for firms and households to learn about their value function at different rates so that γ and $\widetilde{\gamma}$ may differ. The third variant of the FH model, referred to as "FH-$\overline{\phi}$", allows for the parameters governing the policy response to trends to differ from those governing the cyclical response of policy.

5 Results

5.1 Parameter Estimates

Table 3 displays the means and standard deviations from the posterior distribution of the estimated parameters. The results suggest that incorporating finite horizon planning into an otherwise canonical NK model is helpful in accounting for movements in U.S. output growth, inflation and interest rates over the 1966-2007 period. In particular, the estimates of ρ in the FH versions of the model are all substantially less than one. Such estimates are consistent with, but not identical to, the recent evidence in Gabaix (2018), who estimates that the values for discounting future output and inflation are around 0.75. In comparison, these mean estimates shown in Table 3 are closer to 0.5. As discussed earlier, a value of $\rho = 0.5$ substantially reduces the degree of forward-looking behavior and as a result dampens the responsiveness of output to interest rate changes and inflation to changes in the cyclical position of the economy. For example, using $\beta = 0.995$, in the canonical NK model, the effect on current inflation of a (constant) of a 1 percentage point increase in the output gap over eight consecutive quarters is $\kappa\frac{1-\beta^9}{1-\beta}\overline{y} \approx 9\kappa\overline{y}$. In contrast, in the FH-baseline model with $\rho = 0.5$, this response is given by $\kappa\frac{1-(\beta\rho)^9}{1-\beta\rho}\overline{y} \approx \kappa\overline{y}$ and is about 9 times smaller.

The estimates also suggest that the slow updating of agents' value functions is helpful in explaining aggregate data. In particular, for all three FH models, the posterior distributions for γ are concentrated at low values, with means around 0.1. For the "FH-$\widetilde{\gamma}$" model, the posterior distribution of $\widetilde{\gamma}$, with a mean of 0.31, is similarly consistent with slow updating. Thus, households and firms both update their value functions relatively slowly to the new data that they observe, imparting considerable persistence into trend components. As a result of this sluggishness, the supply shock is much less persistent in the FH versions of the model than in the canonical NK model. In particular, the mean estimate of ρ_{y^*} is near one in the canonical NK model and close to 0.5 in the FH-baseline model.

Figure 2 provides additional information about the posterior distribution for ρ and γ derived from the FH-$\bar{\phi}$ model. The grey dots represent draws from the prior distribution while the blue dots represent draws from the posterior distribution. As indicated by the much smaller blue region than the grey region, there is substantial information about the values ρ and γ in the data. In particular,

Table 3: Posterior Distributions

	Forward		Exog. Trends		FH-baseline		FH-$\bar{\phi}$		FH-$\tilde{\gamma}$	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD	Mean	SD
r^A	2.25	0.60	2.06	0.76	2.51	0.37	2.39	0.30	2.55	0.46
π^A	3.76	0.76	3.88	0.86	3.98	1.00	3.80	0.91	3.96	0.99
μ^Q	0.40	0.06	0.43	0.03	0.45	0.01	0.45	0.02	0.44	0.01
ρ					0.50	0.13	0.46	0.14	0.69	0.12
γ					0.14	0.03	0.11	0.02	0.06	0.04
$\tilde{\gamma}$									0.31	0.09
σ	0.45	0.33	1.75	0.46	3.57	0.62	3.72	0.65	3.15	0.60
κ	0.31	0.12	0.00	0.00	0.04	0.01	0.03	0.01	0.01	0.01
ϕ_π	2.14	0.33	1.57	0.26	1.07	0.13	0.94	0.15	1.01	0.15
ϕ_y	0.10	0.29	0.86	0.19	0.79	0.16	0.75	0.16	0.93	0.20
$\bar{\phi}_\pi$							2.09	0.26		
$\bar{\phi}_y$							0.05	0.05		
ρ_ξ	0.93	0.05	0.83	0.07	0.98	0.02	0.97	0.02	0.93	0.04
ρ_{y^*}	0.99	0.02	0.90	0.22	0.53	0.09	0.57	0.08	0.31	0.11
ρ_{i^*}	0.71	0.09	0.97	0.01	0.97	0.01	0.97	0.02	0.97	0.01
σ_ξ	1.11	0.60	2.44	1.21	2.17	0.42	2.08	0.39	2.62	0.61
σ_{y^*}	1.18	1.60	1.58	0.83	5.93	2.20	5.99	1.94	17.35	8.87
σ_{i^*}	0.63	0.10	0.70	0.15	0.67	0.12	0.58	0.11	0.77	0.16
Log MDD	-753.63	0.17	-712.52	0.13	-725.69	0.08	-714.59	0.10	-724.57	0.06

NOTE: The table shows estimates of the posterior means and standard deviations of the model parameters computed using 10 runs of the SMC sampler. The mean and standard deviation of the log MDD is computed across the 10 runs of the SMC algorithm. See the appendix for details.

while the prior contains many draws of ρ near one, there are essentially zero posterior draws greater than 0.75. This is substantial evidence against models in which ρ is high including the canonical NK model in which $\rho = 1$. The data are also very informative about γ which determines how quickly the finite-horizon households and firms update their value functions. The posterior distribution for γ lies almost entirely between 0.05 and 0.2, which implies that agent's update their value functions slowly and that trends in inflation, output, and the interest rate are highly persistent.

The estimated coefficients of the monetary policy rule imply that the policy rate is less responsive to cyclical movements in inflation and the output gap in the FH versions of the model than in the canonical NK model. For example, the responsiveness of the policy rate to inflation deviations is about 1.5 in the canonical model and in the Exog. Trends model compared to a value close to 1 in the FH-baseline.

Another important feature of the estimated policy rule is that the data prefers rule coefficients that differ significantly in the short run from those in the long run. In the FH-$\bar{\phi}$ version of the model, the coefficient on trend inflation deviations is near 2 while the coefficient on trend output deviations is close to zero. Hence, the monetary policy rule responds more aggressively to stabilize deviations of trend inflation from the steady state inflation rate than it does to short-run inflation

Figure 2: Joint Posterior Distribution of Parameters ρ and γ

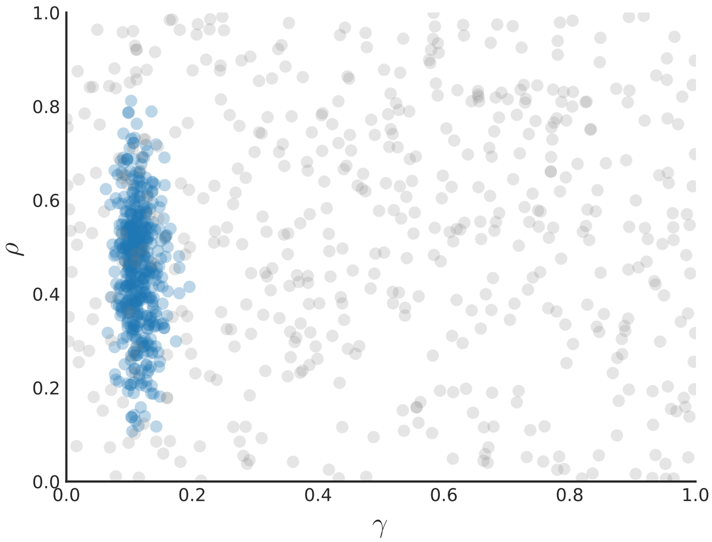

NOTE: The grey dots represent draws from the prior distribution of (ρ, γ) while the blue dots represent draws from the posterior distribution of (ρ, γ) from the FH-$\bar{\phi}$ model.

deviations from trend. In addition, policy responds aggressively to short-run deviations of output from trend but very little to the deviation of trend output from steady state.

5.2 Model Fit

The last row of Table 3 shows, for each model, an estimate of the log marginal data density (MDD), defined as:

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

This quantity provides a measure of overall model fit, and an estimate of it is computed as a by-product of the SMC algorithm used to the estimate the posterior of the model. The MDD of the canonical New Keynesian model is less than MDDs of the FH models by about 20 to 30 log points. This indicates substantial evidence in favor the FH models: In a strict application of Bayesian calculus, a researcher with equal prior odds on the canonical New Keynesian model FH-$\bar{\phi}$ models, would end up with posteriors odds of about eight trillion to one in favor of the FH-$\bar{\phi}$ model. Comparing the FH models, the data moderately favors the FH-$\bar{\phi}$ version of the model which allows for monetary policy to respond more aggressively to deviations in trend inflation than to short-run deviations of inflation from trend. More surprisingly, the FH-$\bar{\phi}$ is competitive with the exogenous trends model. The mean estimates of the MDDs are seperated by only two log points, with the exogenous trends model fitting slightly better. This only small difference in model fit emerges despite the fact that the exogenous trends model is substantially more flexible–it has 3 additional shocks to fit the data–than the FH-$\bar{\phi}$ model. Accordingly, the fit of the FH-$\bar{\phi}$ model is impressive in light of its parsimony. More broadly, the estimates in Table 3 suggest that finite horizons, slow learning about the observed trends, and an aggressive policy response to trend inflation are all important in accounting for movements in inflation, output, and interest rates.

Figure 3 compares the fit of the FH model relative to the canonical NK model over an expanding sample. Specifically, the figure plots

$$\Delta_t = \log \hat{p}_{\text{FH}-\bar{\phi}}(Y_{1:t}) - \log \hat{p}_{\text{Forward}}(Y_{1:t}), \tag{31}$$

where $Y_{1:t}$ is the observables through period t and $\log \hat{p}_{\mathcal{M}}(Y_{1:t})$ is an estimate of the log MDD for model $\mathcal{M}$ for the subsample of Y that ends in period t. Thus, Δ_t measures the cumulative difference in the estimates of the log MDD for the FH-$\bar{\phi}$ from the canonical NK model. The figure shows that the data begins systematically preferring the FH-$\bar{\phi}$ model beginning in 1979. For the canonical NK model, this period is difficult to rationalize, since it must capture the upward inflation trend in the 1970s and its subsequent reversal in the 1980s through large and persistent shocks. In contrast, the FH-$\bar{\phi}$ model embeds persistence into trend inflation that makes it easier to fit the Great Inflation episode. Although the relative fit of the canonical NK model improves somewhat during the Volcker disinflation, as inflation moves back toward the model's mean estimate for π^A of 4 percent, it continues to fit much worse than the FH-$\bar{\phi}$ for the remainder of the sample. This better fit of the FH-$\bar{\phi}$ model reflects that this model does a relatively good job capturing the secular decline in inflation, as inflation moves and remains well below 4 percent over the latter part of the sample.

5.3 Estimated Effects of a Monetary Policy Shock

As discussed earlier, empirical evidence from the VAR literature has emphasized that following a monetary policy shock, there is considerable persistence in the price response and a delayed response in output. Figure 4 plots the 90-percent pointwise credible bands for impulse responses of output, inflation, and the short-term interest rate to a one standard deviation increase in $\epsilon_{i,t}$ from the FH-$\bar{\phi}$ model. There is a persistent fall in output following a tightening in monetary policy with the decline in output after one year on par with the initial fall. This response in part reflects the hump shaped pattern in trend output, which falls slowly over the next year or so before recovering. As shown earlier, the responses from the estimated model for inflation are highly persistent. Inflation only drops slightly on impact and its response grows over time as agents revise down their estimates of the trend. Overall, however, its response is small.

Figure 3: Difference in Log MDD over time

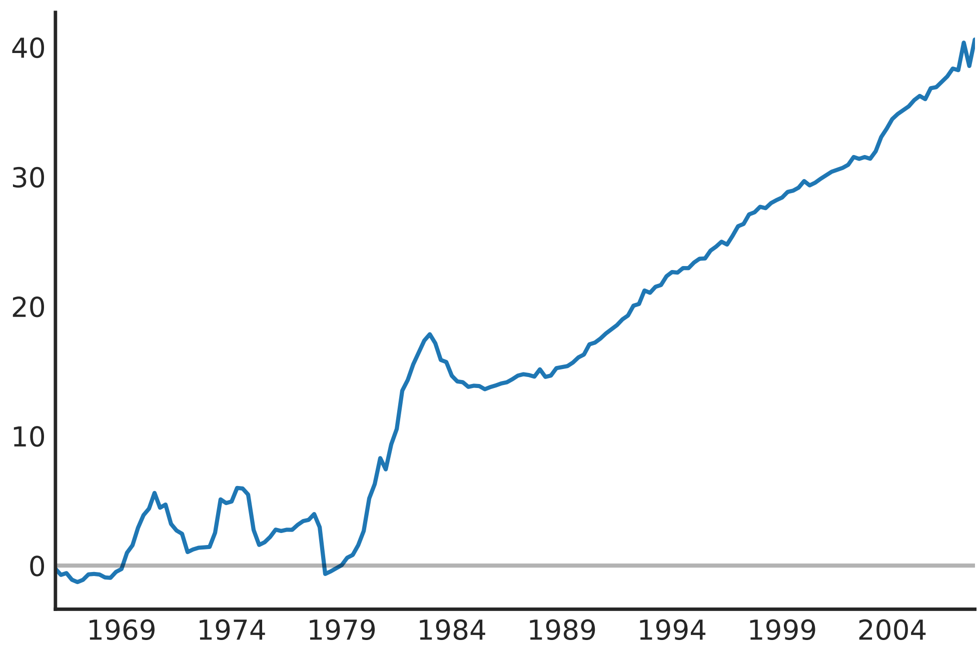

NOTE The figure displays Δ_t defined in equation (31).

5.4 Estimated Trend-Cycle Decomposition

Figure 5 decomposes observed inflation into its trend and cyclical components. The top panel displays the smoothed estimates from the FH-$\bar{\phi}$ model of trend inflation in the top panel. Trend inflation, according to the model, rose sharply during the 1970s, declined during the 1980s, and then remained relatively constant from 1990 to 2007. The middle panel shows that the model's measure of the deviation of inflation from trend displays little persistence with the possible exception of the early 1980s when inflation remained below trend for a couple of years. Moreover, as the middle panel suggests, the model's estimate of $\pi_t - \bar{\pi}_t$ implies that the volatility of inflation relative to trend declined during the period of the Great Moderation. The bottom panel of Figure 5 compares the FH-$\bar{\phi}$ model's trend inflation estimates to an estimate of longer-run inflation expectations computed using survey data.[19] Although the model uses the GDP deflator to compute trend inflation and the survey-based measure is for the CPI, the two series display a similar pattern: both measures fall sharply during the Volcker disinflation and then stabilized in the 1990s at a level well below

[19]This time series is available in the public FRB/US dataset. Starting in 1991, the variable corresponds to the Survey of Professional Forecasters median estimate of 10-year average inflation expectations. Prior to 1991, the variable is constructed using additional surveys along the lines of Kozicki and Tinsley (2001).

Figure 4: Impulse Responses to a Monetary Policy Tightening

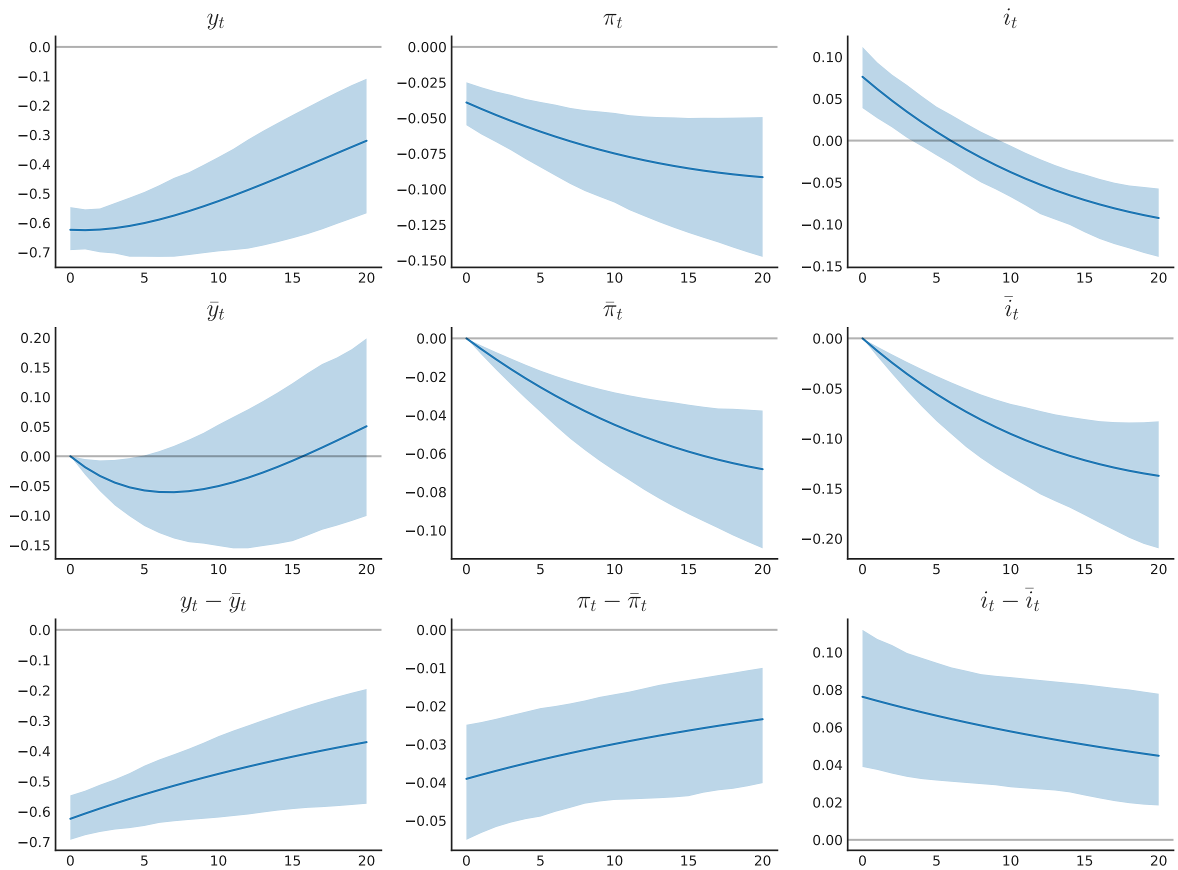

NOTE: This figure plots the posterior mean and the 90-percent pointwise credible bands for impulse responses of model variables to a one standard deviation increase in $\epsilon_{i,t}$ for the FH-$\bar{\phi}$ model using 250 draws from the posterior distribution.

their respective measures in the early 1980s.

The top panel of Figure 6 displays the smoothed estimates from the FH-$\bar{\phi}$ model of the trend interest rate. The trend interest rate follows the same pattern as the model's trend inflation series: rising substantially in the 1970s, falling sharply in the 1980s, and then recovering in the 1990s. The fact that the movements in the trend interest rate is so similar to those for trend inflation in the FH-$\bar{\phi}$ model is not too surprising, since the estimates of that model imply that the trend interest rate is driven almost entirely by trend inflation rather than the trend in output. The lower panel displays FH-$\bar{\phi}$ model's estimates of the deviation of the interest rate from trend. The estimates suggest that monetary policy responded by cutting rates aggressively well below trend during the recessions in late 1960s and mid-1970s. In both the recessions of 1981-82 and in 2001, $i_t - \bar{i}_t$ also fell but from relatively elevated levels.

The top panel of Figure 7 displays the smoothed estimates of the output gap, measured as the

Figure 5: Trend-Cycle Decomposition: Inflation

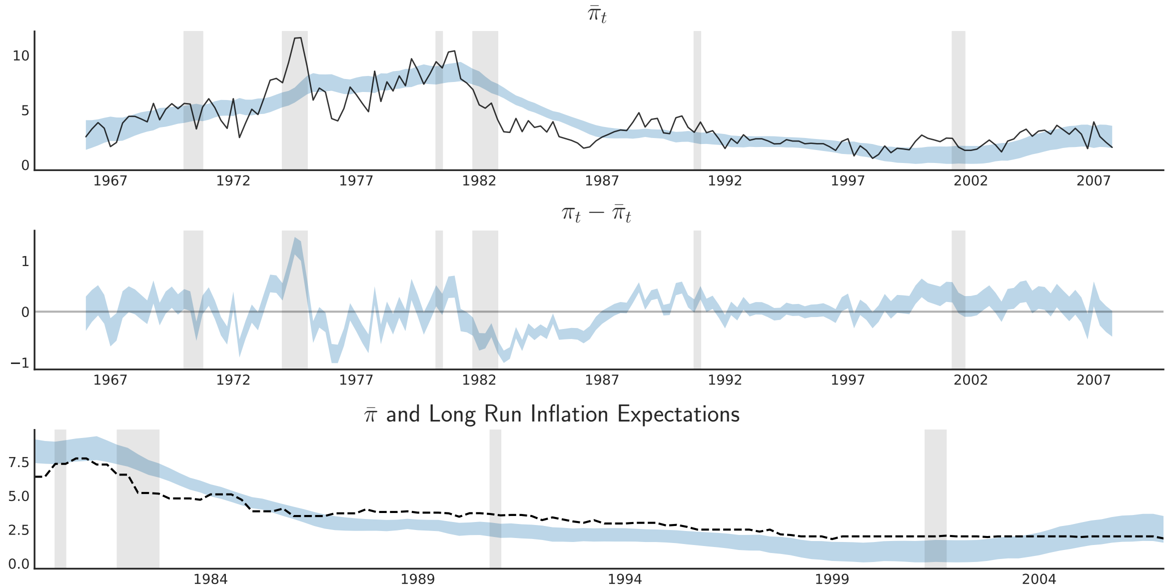

NOTE: The top panel of this figure shows the time series of 90 percent pointwise credible interval for the smoothed mean of $\bar{\pi}_t$, annualized and adjusted by π^A (shaded region), as well as observed inflation (solid line.) The middle panel shows the time series of 90 percent pointwise credible interval for the smoothed mean of $\pi_t - \bar{\pi}_t$ (shaded region). The bottom panel shows the time series of 90 percent pointwise credible interval for the smoothed mean of $\bar{\pi}_t$, annualized and adjusted by π^A (shaded region) along with an estimate of long-run inflation expectations (dashed line) constructed from survey data.

deviation of output relative to trend from the FH-$\bar{\phi}$ model. As shown there, the model's estimate of the output gap falls sharply during NBER recession dates. For example, in both of the recessions in the mid-1970s and in 1981-82, the estimate of $y_t - \bar{y}_t$ falls more than 2 percentage points. In contrast, as shown in the middle panel, the model's estimate of the trend moves much less during NBER recessions. Trend output, for instance, declines slightly during the severe recession in the mid-1970s but this decline is small relative to the fall in the model's cyclical measure for output. In addition, the level of trend output is unchanged or even increases a bit during other NBER recessions. The top panel of the figure also compares the smoothed estimates of the output gap to the output gap measured published by the CBO. The model's estimate of the output gap and the CBO measure have a correlation of about 0.65. The two measures differ notably in terms of how they saw the cyclical position of the economy in the mid to late 1970s and during the Great Moderation. While the CBO measure saw a significant improvement in the cyclical position of the economy following the recession in the mid-1970s, the model-based measure shows little improvement following that recession. In addition, the CBO measure indicates that output was below potential for most of the 1990-2007 period, while the model-based measure of output was closer to trend through much of that period.

Figure 6: Trend-Cycle Decomposition: Short-term Interest Rate

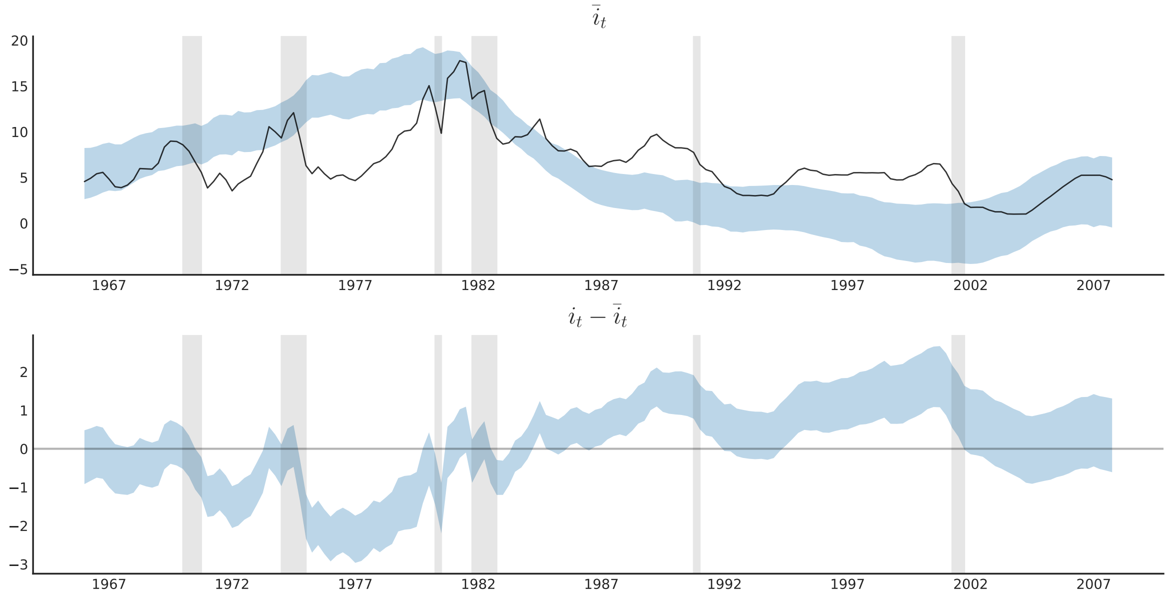

NOTE: The top panel of this figure shows the time series of 90 percent pointwise credible interval for the smoothed mean of $\bar{i}_t$, annualized and adjusted by $\pi^A + r^A$ (shaded region), as well as the observed federal funds rate (solid line.) The bottom panel shows the time series of 90 percent pointwise credible interval for the smoothed mean of $i_t - \bar{i}_t$ (shaded region).

5.5 Estimated Shocks

Figure 8 displays the smoothed estimates of the demand, supply, and monetary policy shocks over our sample period.[20] As shown in the first panel, the estimated demand shock, ξ_t, is highly autocorrelated as the mean estimate of its AR(1) coefficient is 0.97. This shock captures a variety of autonomous factors that affect aggregate spending including changes in household preferences, financial shocks, and changes in fiscal policy. Our reading of the literature is that several of these factors were at play over our sample period. For example, the demand shock was low in the 1970s and fell sharply during the 1973-1975 recession; this is consistent with the large reductions in defense spending that followed the end of the Vietnam War, as discussed in Ramey (2011). In addition, Mishkin (1977) emphasizes an unfavorable shift in household balance sheet positions that contributed to weakness in consumer expenditures during the 1973-1975 recession. The model's estimated demand shock picks up notably in the mid-1980s, which is consistent with the evidence in Ramey (2011) regarding expansionary fiscal policy during the Reagan Presidency. The demand shock is also high in the early 2000s, possibly reflecting the runup in household wealth associated with the boom in housing that occurred at the time.

The model's supply shock, y_t^* captures a variety of forces that affect aggregate supply including

[20] The *iid* innovations underlying these shocks are shown in the appendix. As shown there, the model is consistent with there being very infrequent large shocks. See Blanchard and Watson (1986), for example, find a similar result).

Figure 7: Trend-Cycle Decomposition: Output

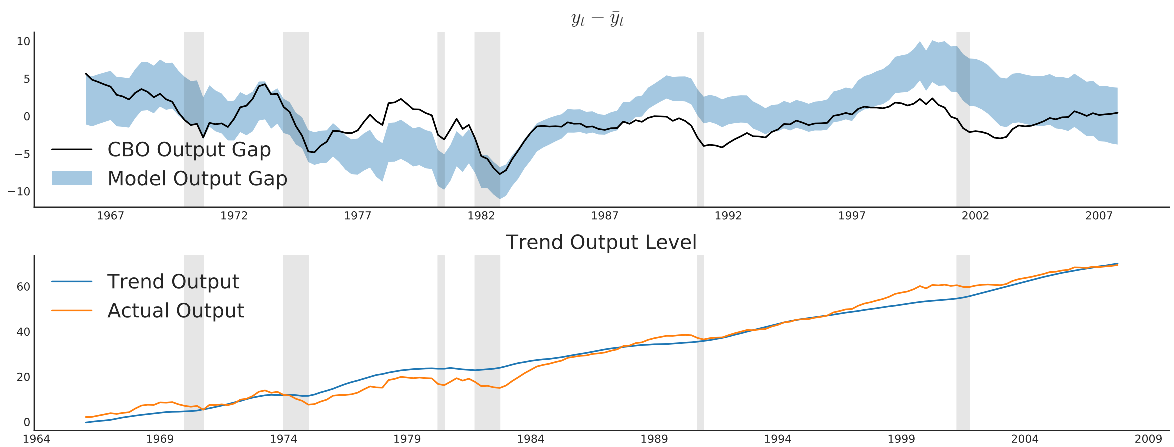

NOTE: The top panel of the figure shows the time series of 90 percent pointwise bands of the cyclical position of output, $y_t - \bar{y}_t$, for the FH-$\bar{\phi}$ model, as well as the CBO estimate of the output gap. The bottom panel shows the level of actual output (orange line) as well as the smoothed mean estimates of the trend level of output (blue line) for the FH-$\bar{\phi}$ model inclusive of trend growth (μ^Q).

technological and regulatory changes, shocks to oil prices, and other factors that affect firms' costs to production. As shown in the middle panel, the estimated supply shock is positively correlated but less so than the demand shock, as there are a number of infrequent but sharp changes in supply. Two of those sharp changes occur in 1973-1974 and in 1979-80 when there were large disruptions in world oil supply. These large contractions in the model's supply shocks are important in contributing to the sharp contraction in output and concurrent rise in inflation that occurred during those episodes. Finally, the supply shock is relatively high in the late 1990s and early 2000s, a time period which has been identified with a sustained increase in total factor productivity growth by for example, Fernald (2016).

The lower panel of Figure 8 displays the exogenous shock to the monetary policy rule, i_t^*. This shock is highly persistent as the mean estimate of the AR(1) coefficient for this shock is about 0.97. Consistent with the evidence in Romer and Romer (2004), there are some notable departures of monetary policy from the systematic portion of the rule. First, early in the sample – the period from the mid-1960s until 1974 – our estimates imply that monetary policy was overly accommodative relative to the rule. Second, our estimates imply that monetary policy tightened considerably relative to the systematic portion of the rule in the early 1980s, a period commonly referred to as the "Volcker disinflation." Finally, our estimates imply that monetary policy was tightening in the years leading up to the Great Recession but remained relatively accommodative in the years leading up to that event.

Figure 8: Estimated Shocks

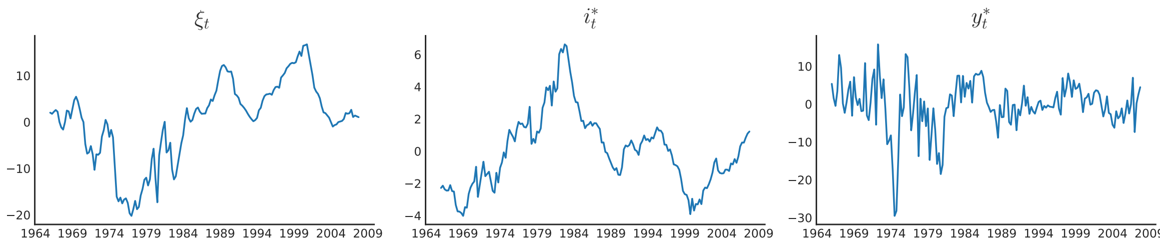

NOTE: The panels of the figure shows the smoothed estimates of the shocks from the FH-$\overline{\phi}$ model.

5.6 Shock Decomposition

Figure 9 shows the contribution of each of the shocks to the estimates of the cyclical and trend fluctuations in output and inflation.[21] Because of the model's property that the cycle affects the trend, the three shocks affect not only cyclical fluctuation in output and inflation but also their trends. As discussed earlier, the monetary policy rule also plays an important role in influencing the model's trend-cycle dynamics, and it is useful to reproduce the rule using the mean parameter estimates:

$$\begin{aligned} i_t - \bar{i}_t &= \phi_\pi(\pi_t - \overline{\pi}_t) + \phi_y(y_t - \overline{y}_t) + i_t^* \\ \bar{i}_t &= \overline{\phi}_\pi \overline{\pi}_t \end{aligned}$$

where in line with those estimates we have imposed that $\overline{\phi}_y = 0$. In addition, it is useful to note that our estimated rule implies an aggressive response to cyclical output, as $\phi_y \approx 0.75$ and a more aggressive response to trend inflation ($\overline{\phi}_\pi \approx 2$) than cyclical inflation ($\phi_\pi \approx 1$).

As shown in Figure 9, the severe upward spiral in inflation during the 1970s is manifested mainly in the model's estimate of trend inflation. In the early part of the decade, the upward pressure on trend inflation reflects accommodative policy stemming from non-systematic deviations from the rule. However, by the middle of the decade the predominant force pushing up trend inflation is weakness in aggregate demand that was accommodated by the systematic policy response. Because these demand shocks lowered $y_t - \overline{y}_t$ and $\pi_t - \overline{\pi}_t$ sharply in 1973-74, monetary policy eased significantly in response. However, the easing was overly aggressive in the sense that it generated a substantial increase in agents' expectations of longer-run or trend inflation. As trend inflation rose in the mid-to-late 1970s, monetary policy tightened in accordance with the intercept portion of the rule but too slowly to stave off the persistent increase in inflation. Accordingly, from the model's perspective monetary policy was too easy in the early half of the decades and "fell behind the curve" in reacting to upward pressure on inflation stemming from agents' expectations of

[21] These contributions are constructed using the smoothed shocks under the mean parameter estimates from the version of the model in which the short and long-run coefficients of the rule differ (i.e., FH-$\overline{\phi}$).

Figure 9: Trend and Cycle of Output and Inflation: Historical Counterfactuals

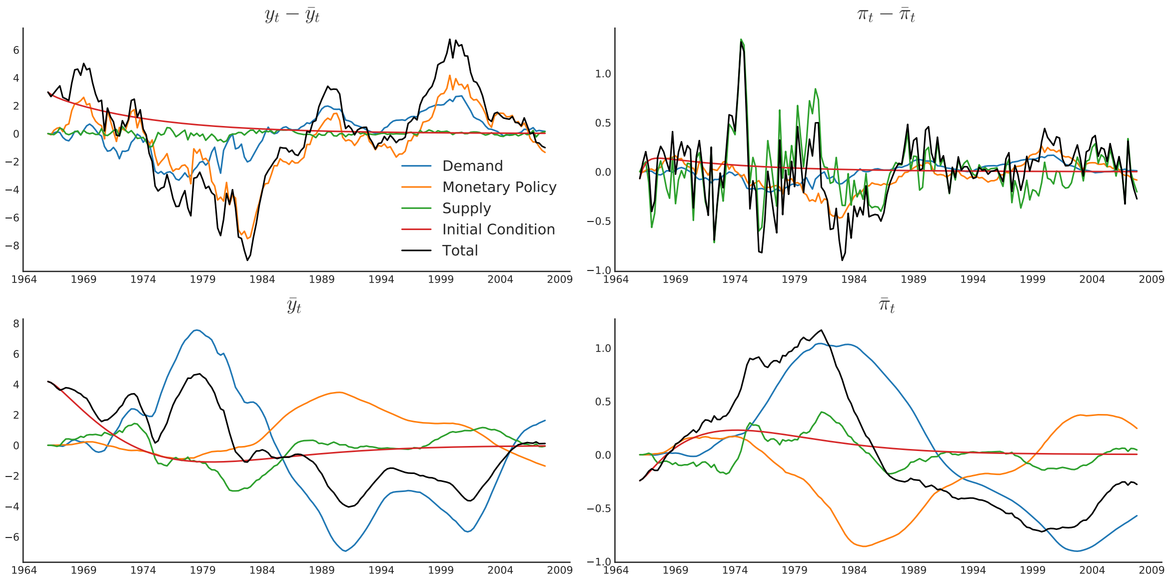

NOTE: Figure shows the time series of the smoothed means of the trend and cycle of output and inflation for the FH-$\bar{\phi}$ model(black lines). The figure also shows the contribution to these trajectories of demand shocks (blue lines), supply shocks (green lines), monetary policy shocks (orange lines), and initial conditions (red lines) for the FH-$\bar{\phi}$ model.

longer-run inflation.

The model's estimates also imply that the subsequent disinflation in the 1980s is largely driven by a substantial and persistent tightening of monetary policy. Tight monetary policy in the early 1980s, both because trend inflation remains high and more importantly because of substantial non-systematic policy shocks, puts downward pressure on cyclical inflation and is largely responsible for the significant downturn in output in 1981-82. With the large, non-systematic tightening of policy persisting through the first half of the 1980s, cyclical inflation remains low and eventually agents' expectations for longer-run inflation begin to decline as well. While monetary policy shocks play an important role in the decline in trend inflation, the demand shock also contributes to bringing trend inflation down.

While the model's estimates imply that demand and monetary shocks played an important role in driving trend inflation over our sample period, supply shocks are relegated to a lesser role. However, these shocks, which are estimated to be less persistent than demand or monetary shocks, play a prominent role in driving cyclical movements in inflation. In particular, the supply shock largely explains the jumps in cyclical inflation in 1973-1974 and 1979-80 when there were disruptions to world oil supply.

5.7 Aggregate Data and the Role of Heterogeneity

The underlying model used so far allows for heterogeneity in the length of the planning horizons across households and firms. But, since our focus is on explaining aggregate fluctuations, a legitimate question to ask is how important is this heterogeneity in explaining aggregate data. To address this question, Table 4 compares the marginal data density of the FH-$\overline{\phi}$ model—in which households and firms are heterogeneous in their planning horizons—with versions of the model in which all households and firms have the same planning horizon. In particular, all households solve an identical problem, with a k-period planning horizon; and all firms solve an identical problem, also with a k-period planning horizon.

As shown there, the marginal data density is slightly lower for the representative agent version of the model in which planning horizon includes the current and next quarter ($k = 1$) than the marginal data density of the FH-$\overline{\phi}$ model, indicating a slightly better fit using representative agents. Table 4 also shows that the fit of the representative-agent model deteriorates as the planning horizon increases. Overall, relative to the heterogeneous planning horizon version of the model, these results suggest that representative-agent versions of the model remain useful in explaining aggregate fluctuations. However, this better fits may reflect that the FH-$\overline{\phi}$ model uses a relatively rigid distribution of agents (i.e. exponential distribution) whose main appeal is in deriving the elegant expressions for aggregate output and inflation shown in equations (4) and (5).

6 Comparison with other Behavioral NK Models

So far, we have shown how the finite-horizon model can do nearly as well in accounting for inflation and output dynamics over the Great Inflation and Volcker disinflation periods than a NK model that incorporates stochastic trends. In this section, we compare the finite-horizon model's performance to other ways of incorporating behavioral features into the NK model. In addition, we compare the model's performance to the hybrid NK model which includes habit persistence and inflation indexation in order to generate persistent movements in output and inflation.

The model with finite horizons is closely related to two recent extensions of the NK model. The first is discussed in Gabaix (2018), who departs from rational expectations by assuming agents' beliefs are distorted when forecasting future variables. Angeletos and Lian (2018) also extend the NK

Table 4: Log Marginal Data Density Estimates for Single Agent Models

	mean	std
FH-$\bar{\phi}$ Het. Agent	-714.59	0.10
FH-$\bar{\phi}$ Rep. Agent ($k = 0$)	-718.12	0.11
FH-$\bar{\phi}$ Rep. Agent ($k = 1$)	-710.57	0.41
FH-$\bar{\phi}$ Rep. Agent ($k = 2$)	-714.90	0.80
FH-$\bar{\phi}$ Rep. Agent ($k = 3$)	-722.97	1.60
FH-$\bar{\phi}$ Rep. Agent ($k = 4$)	-725.57	0.72

model so that strategic interactions between agents affect expectations of future variables. Though the microfoundations differ from the finite-horizon approach discussed here, these recent extensions give rise to similar expressions characterizing linearized aggregate dynamics. In Angeletos and Lian (2018), the linearized expressions for output and inflation are given by:

$$y_t = \rho E_t y_{t+1} - \sigma(i_t - \lambda E_t \pi_{t+1} - r_t^n) \quad (32)$$

$$\pi_t = \beta \rho_f E_t \pi_{t+1} + \kappa y_t + u_t \quad (33)$$

where the parameters ρ, λ, and $\rho_f \in [0,1]$. The expressions determining aggregate output and inflation in Gabaix (2018) are very similar except that $\lambda = 1$.

The expressions (32) and (33) are similar to those determining aggregate inflation and output in the finite-horizon approach; however, in the finite-horizon model, ρ, ρ_f, and λ are constrained to be the same. A more important difference is that the variables in the finite-horizon model are expressed in deviation from trends which are determined endogenously as agents update their value functions. In contrast, this feature is absent from Angeletos and Lian (2018) and the variables are expressed as a deviation from their nonstochastic steady state.

Table 5 compares our measure of model fit, the log marginal data density, for the Angeletos and Lian (2018) to the alternative estimated versions of the finite horizon model discussed earlier.[22] The log marginal data density is about 4 points higher in the version of the finite horizon model in which the monetary policy reaction function is the same in the short and long run (labelled "FH") and about 15 points higher when the policy reaction function differs in the long and short run (labelled FH-$\bar{\phi}$). Accordingly, the fit of the finite-horizon models is better than the model in Angeletos and Lian (2018) and when the policy reaction function is allowed to differ in the short and long run in the finite-horizon model, the fit is substantially better. This improved fit reflects the endogenous persistence the finite-horizon approach can generate through slow moving trends for output, inflation, and interest rates. Similar results would apply to the model in Gabaix (2018), since the aggregate dynamics of that model (up to a log-linear approximation) are a special case of Angeletos and Lian (2018) with $\lambda = 1$.

It is interesting to compare the FH model to the hybrid NK model, since both generate endogenous persistence but through different mechanisms. The hybrid NK model introduces persistence into output and inflation by introducing habit formation in household preferences and indexation to past inflation in the price contracts of firms, and these features have been used extensively in empirical applications in the literature.[23] In the hybrid NK model, the log-linear aggregate dynamics

[22] For the Angeletos and Lian (2018) model, we use uniform priors on $[0,1]$ for ρ_f and λ. The Appendix shows the posterior distributions of the parameters of this model.

[23] The underlying preferences for the households are $\log[C_t - \nu C_{t-1}^a] - H_t$, where H_t are hours worked and the parameter ν captures the presence of (external) habits reflecting the influence of "aggregate" past consumption on current utility. Firms set prices in a staggered way (à la Calvo) and price contracts are indexed to past aggregate inflation. The indexation parameter is $1 - a$.

Table 5: Overall Fit of Alternative Models

	Mean	Std. Dev.
Angeletos-Lian	-729.22	0.12
Exog. Trends	-712.52	0.13
FH-$\bar{\phi}$	-714.59	0.10
FH-$\tilde{\gamma}$	-724.57	0.06
FH-baseline	-725.69	0.09
Forward	-753.63	0.18
Hybrid NK	-734.78	0.31

NOTE: Means and standard deviations are over 10 runs of each algorithm.

(around the non-stochastic steady state) for output and inflation are given by:

$$\begin{aligned}
(1+2\alpha)y_t &= \alpha y_{t-1} + (1+\alpha)E_t y_{t+1} - E_t\left[i_t - \pi_{t+1} - \xi_t\right] \\
[1+\beta(1-a)]\pi_t &= (1-a)\pi_{t-1} + \beta E_t \pi_{t+1} \\
&\quad + \kappa(1+\alpha)y_t - \kappa\alpha y_{t-1} + y_t^* \qquad (34)
\end{aligned}$$

where $\alpha = \frac{\nu}{1-\nu}$, ν is the habit-formation parameter in the households' preferences, β is the households' discount factor, and $1-a$ is the indexation to past inflation of the Calvo's price contracts of firms.[24]

Table 5 shows that the three versions of the finite-horizon model that we estimate all fit the observed dynamics of output, inflation, and the interest rate better than the hybrid NK model. This better fit reflects both the endogenous persistence generated by agents' learning about their value functions as well as the reduced degree of forward-looking behavior associated with $\rho < 1$. Overall, the results in this section suggest the finite-horizon approach with agents' learning about their value functions is a parsimonious and fruitful way to fit movements in longer-run trends and aggregate business cycle dynamics.

7 Conclusion

In this paper, we used aggregate data to estimate and evaluate a behavioral New Keynesian (NK) model in which households and firms have finite horizons. Our parameter estimates implied that most households and firms have planning horizons under two years, and we could reject parameterizations of the model in which agents had long planning horizons such as the canonical NK model with rational expectations. Our parameter estimates also implied that households and firms update their beliefs about their value functions slowly. These slowly evolving beliefs allowed the model

[24]For the Hybrid NK model, we use uniform distribituion for the prior on a and a Beta distribution with mean 0.7 and standard deviation 0.15 for the prior on ν. The Appendix shows the posterior distributions of the parameters of this model.

to generate endogenous persistence that helped it explain persistent trends observed in inflation, output, and interest rates in the United States over the 1966-2007 period. We also showed that the FH model outperformed other behavioral NK models as well as rational expectations versions of the NK model commonly used in empirical applications. Overall, our empirical analysis suggests that the FH model is a promising framework for explaining aggregate data and analyzing monetary policy.

Our paper provides estimates of important parameters of the FH model that can be used to study the heterogeneity of households and firms that underlies the (aggregate) model. Recent studies in NK models has emphasized the importance of heterogeneity (e.g., Kaplan, Moll, and Violante (2018)), and with different planning horizons across both individual households and firms, the FH approach naturally gives rise to disperse beliefs about expected inflation and output. In future work, it would be interesting to investigate whether these disperse beliefs across households and firms are consistent with surveys of households and firms.

References

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

Blanchard, O., and M. Watson (1986): "Are Business Cycles All Alike?," in *The American Business Cycle: Continuity and Change*, pp. 123–180. National Bureau of Economic Research, Inc.

Brunnermeier, M. K., and C. Julliard (2008): "Money Illusion and Housing Frenzies," *The Review of Financial Studies*, 21(1), 135–180.

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

Canova, F. (2014): "Bridging DSGE Models and the Raw Data," *Journal of Monetary Economics*, 67(C), 1–15.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (2005): "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 113(1), 1–45.

De Grauwe, P. (2012): *Lectures on Behavioral Macroeconomics*. Princeton University Press, Princeton, NJ.

Del Negro, M., M. Giannoni, and C. Patterson (2012): "The Forward Guidance Puzzle," Staff Reports 574, Federal Reserve Bank of New York.

Evans, G. W., and S. Honkapohja (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, NJ.

Farhi, E., and I. Werning (2018): "Monetary Policy, Bounded Rationality, and Incomplete Markets," Revise and resubmit, American Economic Review.

Fernald, J. G. (2016): "Reassessing Longer-Run U.S. Growth: How Low?," Working Paper Series 2016-18, Federal Reserve Bank of San Francisco.

Gabaix, X. (2018): "A Behavioral New Keynesian Model [Working Paper]," Revise and resubmit, American Economic Review.

Galí, J. (2008): *Monetary Policy, Inflation, and the Business Cycle*. Princeton University Press, Princeton, NJ.

Garcia-Schmidt, M., and M. Woodford (2019): "Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis," *American Economic Review*, 109(1), 86–120.

Gust, C. J., E. Herbst, and J. D. Lopez-Salido (2018): "Forward Guidance with Bayesian Learning and Estimation," Finance and Economics Discussion Series 2018-072, Board of Governors of the Federal Reserve System (US).

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.

Kaplan, G., B. Moll, and G. L. Violante (2018): "Monetary Policy According to HANK," *American Economic Review*, 108(3), 697–743.

Kozicki, S., and P. Tinsley (2001): "Term structure views of monetary policy under alternative models of agent expectations," *Journal of Economic Dynamics and Control*, 25(1), 149 – 184, Computational Methods in Economic Dynamics and Finance.

McKay, A., E. Nakamura, and J. Steinsson (2016): "The Power of Forward Guidance Revisited," *American Economic Review*, 106(10), 3133–58.

Mishkin, F. (1977): "What Depressed the Consumer? The Household Balance Sheet and the 1973-75 Recession," *Brookings Papers on Economic Activity*, 8(1), 123–174.

Modigliani, F., and R. A. Cohn (1979): "Inflation, Rational Valuation and the Market," *Financial Analysts Journal*, 35(2), 24–44.

Nelson, C. R., and C. R. Plosser (1982): "Trends and Random Walks in Macroeconmic Time Series," *Journal of Monetary Economics*, 10(2), 139 – 162.

Ramey, V. A. (2011): "Identifying Government Spending Shocks: It's all in the Timing," *The Quarterly Journal of Economics*, 126(1), 1–50.

Romer, C. D., and D. H. Romer (2004): "Choosing the Federal Reserve Chair: Lessons from History," *Journal of Economic Perspectives*, 18(1), 129–162.

Sargent, T. J. (1993): *Bounded Rationality in Macroeconomics*. Clarendon Press; Oxford University Press, Oxford, England.

Schorfheide, F. (2013): "Estimation and Evaluation of DSGE Models: Progress and Challenges," in *Advances in Economics and Econometrics: Tenth World Congress*, vol. 3, p. 184. Cambridge University Press.

Sims, C. A. (1998): "Stickiness," *Carnegie-Rochester Conference Series on Public Policy*, 49, 317 – 356.

Smets, F., and R. Wouters (2007): "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 97(3), 586–606.

Taylor, J. B. (1993): "Discretion versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy*, 39, 195 – 214.

Woodford, M. (2003): *Interest and Prices Foundations of a Theory of Monetary Policy*. Princeton University Press, Princeton, NJ.

——— (2018): "Monetary Policy Analysis When Planning Horizons Are Finite," Columbia University, New York, mimeo.

Appendices

A Model Dynamics

In this section of the appendix, we provide some of the details that help characterize the dynamics of the finite-horizon model.

A.1 The Cycle and the Taylor Principle

The system determining the cycle is:

$$\tilde{x}_t = \rho M \cdot E_t[\tilde{x}_{t+1}] + N \cdot u_t, \tag{A-1}$$

where the matrices $M = \frac{1}{\delta}\begin{pmatrix} 1 & \sigma(1-\beta\phi_\pi) \\ \kappa & \kappa\sigma + \beta(1+\sigma\phi_y) \end{pmatrix}$ and $N = \frac{1}{\delta}\begin{pmatrix} -\sigma & -\sigma\kappa\phi_\pi \\ -\kappa\sigma & \kappa(1+\sigma\phi_y) \end{pmatrix}$, with $\delta = 1 + \sigma(\phi_y + \kappa\sigma_\pi)$. To determine the Taylor principle for the FH model, rewrite the system (A-1) as

$$E_t[\tilde{x}_{t+1}] = A[\tilde{x}_t] + Bu_t,$$

where the relevant matrix A is given by

$$A = \begin{pmatrix} (\beta\rho)^{-1} & -\kappa(\beta\rho)^{-1} \\ \sigma(\phi_\pi - \beta^{-1}) & 1 + \sigma(\phi_y + \kappa\beta^{-1}) \end{pmatrix}.$$

The equilibrium is determinate if and only if the matrix A has both eigenvalues outside the unit circle (i.e., with modulus larger than one). Invoking proposition (C.1) in Woodford (2003), this condition is satisfied if and only if

$$\det(A) - tr(A) > -1.$$

This condition implies:

$$\left(\frac{1-\beta\rho}{\kappa}\right)\phi_y + \phi_\pi > \rho.$$

A.2 Trend-Cycle Decomposition

In this section, we report the matrices that determine the evolution of the model's trends. The evolution equations of v_t and $\tilde{v}_t$ are given by:

$$V_{t+1} = (I - \Gamma)V_t + \Gamma\Phi x_t, \tag{A-2}$$

where $V_t' = \begin{pmatrix} v_t & \widetilde{v}_t \end{pmatrix}$, and the matrices $\Gamma = \begin{pmatrix} \gamma & 0 \\ 0 & \widetilde{\gamma} \end{pmatrix}$ and $\Phi = \begin{pmatrix} 1 & \sigma \\ 0 & \frac{1}{(1-\alpha)} \end{pmatrix}$. The trends can be written in terms of V_t as: :

$$\overline{x}_t = (1-\rho)\Theta V_t, \tag{A-3}$$

where the matrix of coefficients $\Theta = \frac{1}{\Delta}\begin{pmatrix} 1-\beta\rho & -\sigma(\overline{\phi}_\pi - \rho)(1-\alpha)\beta \\ \kappa & (1-\rho+\sigma\overline{\phi}_y)(1-\alpha)\beta \end{pmatrix}$ and $\Delta = (1-\beta\rho)(1-\rho+\sigma\overline{\phi}_y) + \kappa\sigma(\overline{\phi}_{\ pi} - \rho)$.

Combining expression (A-2) with expression (A-3) yields:

$$\overline{x}_t = \Lambda\overline{x}_{t-1} + (1-\rho)\gamma Q x_{t-1},$$

where $\Lambda = \Theta(I-\Gamma)\Theta^{-1}$ and $(1-\rho)\gamma Q = \Theta\Gamma\Phi$. After some algebra these matrices can be written as:

$$\Lambda = \frac{1}{\Delta}\begin{pmatrix} (1-\gamma)(1-\beta\rho)(1-\rho+\sigma\overline{\phi}_y)+(1-\widetilde{\gamma})\frac{(\overline{\phi}_\pi-\rho)}{(\sigma\kappa)^{-1}} & \sigma(1-\beta\rho)(\overline{\phi}_\pi-\rho)(\widetilde{\gamma}-\gamma) \\ (\widetilde{\gamma}-\gamma)\kappa(1-\rho+\sigma\overline{\phi}_y) & (1-\widetilde{\gamma})(1-\beta\rho)(1-\rho+\sigma\overline{\phi}_y)+(1-\gamma)\frac{(\overline{\phi}_\pi-\rho)}{(\sigma\kappa)^{-1}} \end{pmatrix}$$

$$Q = \frac{1}{\Delta}\begin{pmatrix} (1-\beta\rho) & \sigma(1-\beta\rho)-\frac{\widetilde{\gamma}}{\gamma}\sigma(\overline{\phi}_\pi-\rho)\beta \\ \kappa & \kappa\sigma+\frac{\widetilde{\gamma}}{\gamma}(1-\rho+\sigma\overline{\phi}_y)\beta \end{pmatrix}.$$

When $\gamma = \widetilde{\gamma}$, the system simplifies to:

$$\overline{x}_t = (1-\gamma)\overline{x}_{t-1} + (1-\rho)\gamma Q x_{t-1},$$

with $Q = \frac{1}{\Delta}\begin{pmatrix} 1-\beta\rho & \sigma(1-\beta\overline{\phi}_\pi) \\ \kappa & \kappa\sigma+(1-\rho+\sigma\overline{\phi}_y)\beta \end{pmatrix}$. Note that in this case the feedback of $\overline{x}_t$ on its lag can be characterized by the scalar, $1-\gamma$, and that Q is independent of γ. Finally, Q can be simplified further if $\overline{\phi}_y = 0$: $Q = \frac{1}{\Delta}\begin{pmatrix} 1-\beta\rho & \sigma(1-\beta\overline{\phi}_\pi) \\ \kappa & \kappa\sigma+(1-\rho)\beta \end{pmatrix}$, with $\Delta = (1-\beta\rho)(1-\rho)+\kappa\sigma(\overline{\phi}_\pi-\rho) > 0$ if $\overline{\phi}_\pi > \rho$.

B Data

The data used in the estimation is constructed as follows.

1. **Per Capita Real Output Growth.** Take the level of real gross domestic product, (`FRED` mnemonic "GDPC1"), call it GDP_t. Take the quarterly average of the Civilian Non-institutional Population (`FRED` mnemonic "CNP16OV" / BLS series "LNS10000000"), call it POP_t. Then,

$$\begin{aligned} &\text{Per Capita Real Output Growth} \\ &= 100\left[\ln\left(\frac{GDP_t}{POP_t}\right) - \ln\left(\frac{GDP_{t-1}}{POP_{t-1}}\right)\right]. \end{aligned}$$

2. **Annualized Inflation.** Take the GDP deflator, (`FRED` mnemonic "GDPDEF"), call it $PGDP_t$. Then,

$$\text{Annualized Inflation} = 400\ln\left(\frac{PGDP_t}{PGDP_{t-1}}\right).$$

3. **Federal Funds Rate.** Take the effective federal funds rate (`FRED` mnemonic "FEDFUNDS"), call it FFR_t. Then,

$$\text{Federal Funds Rate} = FFR_t.$$

The figures in the paper include two additional series, the CBO estimate of the Output Gap and longer-run inflation expectations. These data are constructed as follows.

1. **CBO Output Gap.** The CBO's estimate of the level of Potential GDP, (`FRED` mnemonic "GDPPOT"), call it POT_t.

$$\text{CBO Output Gap}_t = 100\ln\left(\frac{GDP_t}{POT_t}\right).$$

2. **Longer-run Inflation Expectations.** An estimate of historical inflation expectations can be found in the public FRB/US dataset. The variable is called PTR_t. Then,

$$\text{Longer-run Inflation Expectations} = PTR_t$$

C Posterior Sampler: Details and Additional Results

For most of the models in the paper, our estimation follows Herbst and Schorfheide (2014) with the following hyperparameters: $N_{part} = 16,000, N_\phi = 500, \lambda = 2.1, N_{blocks} = 3, N_{intmh} = 1$. We run each sampler $N_{run} = 10$ times, and pool the draws from the runs, yielding a posterior distribution with $160,000$ draws. There are two exceptions: for the Forward and the FH-$\tilde{\gamma}$ models, we use $N_{part} = 25,000$, $N_\phi = 2000$, and $N_{blocks} = 6$ because of bimodalities in the posterior.

We assess the convergence and efficiency of our algorithm by analyzing the variation of the estimate of the sample mean across the N_{run} runs of the algorithm. This variance serves as an estimate of the CLT variance associated with the SMC-based estimate of the sample mean (as the number of particles becomes large). Call this estimated variance $\text{VAR}[\bar{\theta}]$ for any parameter θ. We also construct a measure of efficiency of the sampler based on the following idea: Suppose we were able to compute M i.i.d. draws from the marginal posterior distribution for θ. The variance of the mean, $\bar{\theta}$, of these draws would be given by

$$\mathbb{V}[\bar{\theta}] = \frac{\mathbb{V}[\theta]}{M},$$

where $\mathbb{V}[\theta]$ is the posterior variance of θ. We define the number of effective draws as:

$$\text{number of effective draws} = \frac{\hat{\mathbb{V}}[\theta]}{\text{VAR}[\bar{\theta}]},$$

where the hat indicates that we are using our estimated posterior variance. Such a measure indicates this (in)efficiency of the sampler, relative to hypothetical i.i.d. draws. Tables A-1 through A-7 display the estimated mean and 5th and 95th percentiles of the posteriors, in addition to the standard deviation of the mean across the N_{runs} runs and N_{eff}, the number of effective draws for each of the estimated models.

In general the SMC-based estimates of the posterior mean are relatively precise. The parameter σ_{y^*}, whose posterior mean lies in the tail of its prior distribution for many models, typically has the noisiest estimates. Across models, the Forward model is the most difficult to estimate, owing to a bimodality in σ. However, this bimodality does not affect the stability of the estimate model fit (log MDD), as each model has about the same density height.

Table A-1: Posterior Distribution of the Forward Model

	Mean	Std(Mean)	Q05	Q95	Neff
r^A	2.25	0.01	1.27	3.26	7060.56
π^A	3.76	0.01	2.55	5.03	6140.92
μ^Q	0.40	0.00	0.31	0.50	1764.48
σ	0.45	0.02	0.23	1.30	410.38
κ	0.31	0.01	0.01	0.49	539.78
ϕ_π	2.14	0.01	1.42	2.61	619.23
ϕ_y	0.10	0.01	0.00	0.95	393.38
ρ_ξ	0.93	0.00	0.80	0.97	348.04
ρ_{y^*}	0.99	0.00	0.95	1.00	491.94
ρ_{i^*}	0.71	0.00	0.62	0.96	524.75
σ_ξ	1.11	0.04	0.62	2.39	288.46
σ_{y^*}	1.18	0.06	0.81	2.44	783.37
σ_{i^*}	0.63	0.00	0.51	0.82	726.86

NOTE: The table displays the mean, 5th, and 95th percentile of the posterior distribution of the Forward model, as well as the standard deviation of the posterior mean across 10 runs of the sampler.

Table A-2: Posterior Distribution of the Exogenous Trends Model

	Mean	Std(Mean)	Q05	Q95	Neff
r^A	2.06	0.01	0.89	3.36	5815.27
π^A	3.88	0.03	2.46	5.30	942.33
μ^Q	0.43	0.00	0.38	0.47	2611.56
σ	1.75	0.01	1.08	2.57	1790.34
κ	0.00	0.00	0.00	0.00	2098.67
ϕ_π	1.57	0.00	1.17	2.01	3247.53
ϕ_y	0.86	0.00	0.60	1.21	2092.40
ρ_ξ	0.83	0.00	0.70	0.92	959.68
ρ_{y^*}	0.90	0.01	0.29	1.00	426.79
ρ_{i^*}	0.97	0.00	0.95	0.99	1767.18
σ_ξ	2.44	0.07	1.02	4.67	279.73
σ_{y^*}	1.58	0.03	0.75	3.15	807.21
σ_{i^*}	0.70	0.00	0.50	0.98	1895.24
$\rho_{\bar{\pi}}$	0.78	0.01	0.58	0.95	286.80
$\rho_{\bar{i}}$	0.96	0.00	0.90	0.99	1965.97
$\rho_{\bar{y}}$	0.95	0.00	0.86	1.00	3038.06
$\sigma_{\bar{\pi}}$	0.23	0.00	0.19	0.27	468.25
$\sigma_{\bar{i}}$	0.12	0.00	0.07	0.19	3742.10
$\sigma_{\bar{y}}$	0.12	0.00	0.07	0.19	1254.13

NOTE: The table displays the mean, 5th, and 95th percentile of the posterior distribution of the Exog. Trends model, as well as the standard deviation of the posterior mean across 10 runs of the sampler.

Table A-3: Posterior Distribution of the FH Model

	Mean	Std(Mean)	Q05	Q95	Neff
r^A	2.51	0.01	1.85	3.07	5251.54
π^A	3.98	0.01	2.34	5.62	6468.59
μ^Q	0.45	0.00	0.43	0.47	5544.45
ρ	0.50	0.01	0.27	0.71	639.00
γ	0.14	0.00	0.09	0.19	2192.39
σ	3.57	0.01	2.59	4.64	3790.00
κ	0.04	0.00	0.02	0.06	692.01
ϕ_π	1.07	0.00	0.89	1.30	6865.39
ϕ_y	0.79	0.01	0.57	1.07	926.28
ρ_ξ	0.98	0.00	0.94	1.00	4893.74
ρ_{y^*}	0.53	0.00	0.39	0.67	1533.31
ρ_{i^*}	0.97	0.00	0.95	0.99	5370.33
σ_ξ	2.17	0.02	1.62	2.94	515.89
σ_{y^*}	5.93	0.12	3.24	10.12	335.06
σ_{i^*}	0.67	0.00	0.51	0.89	1090.33

NOTE: The table displays the mean, 5th, and 95th percentile of the posterior distribution of the FH model, as well as the standard deviation of the posterior mean across 10 runs of the sampler.

Table A-4: Posterior Distribution of the FH-$\tilde{\gamma}$ Model

	Mean	Std(Mean)	Q05	Q95	Neff
r^A	2.55	0.00	1.76	3.24	12107.09
π^A	3.96	0.01	2.34	5.59	19699.67
μ^Q	0.44	0.00	0.42	0.46	9307.94
ρ	0.69	0.00	0.47	0.85	915.59
γ	0.06	0.00	0.01	0.14	3165.49
$\tilde{\gamma}$	0.31	0.00	0.16	0.46	1650.56
σ	3.15	0.01	2.24	4.21	1719.39
κ	0.01	0.00	0.01	0.03	790.72
ϕ_π	1.01	0.00	0.78	1.29	15744.44
ϕ_y	0.93	0.00	0.65	1.30	3400.78
ρ_ξ	0.93	0.00	0.86	0.99	7124.87
ρ_{y^*}	0.31	0.00	0.14	0.51	1662.95
ρ_{i^*}	0.97	0.00	0.95	0.99	8313.81
σ_ξ	2.62	0.01	1.85	3.76	2125.10
σ_{y^*}	17.35	0.42	5.96	34.94	446.74
σ_{i^*}	0.77	0.00	0.56	1.06	4066.76

NOTE: The table displays the mean, 5th, and 95th percentile of the posterior distribution of the FH-$\tilde{\gamma}$ model, as well as the standard deviation of the posterior mean across 10 runs of the sampler.

Table A-5: Posterior Distribution of the FH-$\bar{\phi}$ Model

	Mean	Std(Mean)	Q05	Q95	Neff
r^A	2.39	0.01	1.88	2.84	2243.90
π^A	3.80	0.01	2.33	5.33	3716.65
μ^Q	0.45	0.00	0.42	0.48	8951.67
ρ	0.46	0.01	0.22	0.68	525.70
γ	0.11	0.00	0.08	0.15	7959.75
σ	3.72	0.02	2.70	4.84	1001.96
κ	0.03	0.00	0.02	0.06	556.04
ϕ_π	0.94	0.00	0.71	1.20	9018.26
ϕ_y	0.75	0.00	0.53	1.03	1697.69
$\bar{\phi}_\pi$	2.09	0.00	1.68	2.52	3116.96
$\bar{\phi}_y$	0.05	0.00	0.00	0.16	4517.02
ρ_ξ	0.97	0.00	0.93	0.99	4036.74
ρ_{y^*}	0.57	0.00	0.45	0.70	1171.34
ρ_{i^*}	0.97	0.00	0.94	0.99	5604.30
σ_ξ	2.08	0.01	1.56	2.78	775.48
σ_{y^*}	5.99	0.11	3.64	9.61	308.57
σ_{i^*}	0.58	0.00	0.43	0.78	1970.05

NOTE: The table displays the mean, 5th, and 95th percentile of the posterior distribution of the FH-$\bar{\phi}$ model, as well as the standard deviation of the posterior mean across 10 runs of the sampler.

Table A-6: Posterior Distribution of the Angeletos-Lian Model

	Mean	Std(Mean)	Q05	Q95	Neff
r^A	1.82	0.02	0.69	3.25	1039.92
π^A	4.03	0.01	2.67	5.37	10840.71
μ^Q	0.41	0.00	0.38	0.45	5028.95
ρ	0.76	0.01	0.44	0.96	206.74
ρ_f	0.86	0.01	0.24	1.00	317.40
λ	0.08	0.00	0.01	0.22	6899.07
σ	1.88	0.01	1.11	2.79	3745.82
κ	0.03	0.00	0.01	0.12	350.62
ϕ_π	1.45	0.01	1.06	1.89	431.32
ϕ_y	0.51	0.01	0.25	0.86	215.16
ρ_ξ	0.87	0.00	0.81	0.94	2129.60
ρ_{y^*}	0.97	0.00	0.93	0.99	415.24
ρ_{i^*}	0.98	0.00	0.96	1.00	474.52
σ_ξ	0.36	0.00	0.30	0.44	231.83
σ_{y^*}	1.43	0.04	0.77	2.60	312.13
σ_{i^*}	0.56	0.00	0.44	0.76	530.02

NOTE: The table displays the mean, 5th, and 95th percentile of the posterior distribution of the Angeletos-Lian model, as well as the standard deviation of the posterior mean across 10 runs of the sampler.

Table A-7: Posterior Distribution of the Habit Model

	Mean	Std(Mean)	Q05	Q95	Neff
r^A	1.91	0.01	1.08	2.74	1678.01
π^A	4.03	0.03	2.41	5.65	1458.24
μ^Q	0.46	0.00	0.37	0.56	1381.77
ν	0.87	0.01	0.78	0.93	86.60
a	0.98	0.00	0.94	1.00	6597.47
σ	1.71	0.02	1.05	2.51	879.31
κ	0.00	0.00	0.00	0.00	196.18
ϕ_π	1.65	0.01	1.27	2.06	1926.67
ϕ_y	0.23	0.00	0.17	0.28	137.91
ρ_ξ	0.52	0.01	0.38	0.66	81.64
ρ_{y^*}	0.99	0.00	0.97	1.00	325.29
ρ_{i^*}	0.99	0.00	0.98	1.00	352.65
σ_ξ	2.62	0.19	1.25	4.76	35.45
σ_{y^*}	1.78	0.08	0.95	3.25	101.93
σ_{i^*}	0.49	0.00	0.41	0.59	1932.00

NOTE: The table displays the mean, 5th, and 95th percentile of the posterior distribution of the habit model, as well as the standard deviation of the posterior mean across 10 runs of the sampler.

Figure A-1: Estimated Innovations

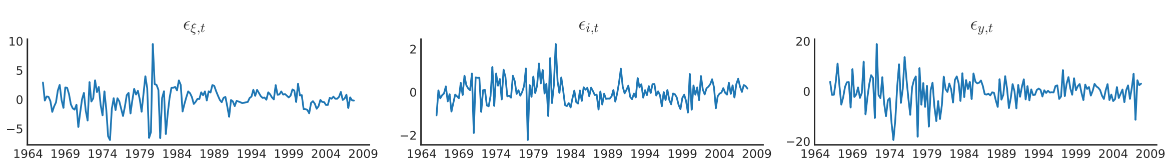

NOTE: The figure shows the time series of the posterior mean smoothed innovations for the FH-$\bar{\phi}$ model.