Optimal Policy Without Rational Expectations: A Sufficient Statistic Solution

Jonathan J Adams*

October 21, 2024

WORKING PAPER
Link to Most Current Version

Abstract

How should policymakers respond to mistakes made by agents without rational expectations? I demonstrate in a general setting that the optimal policy is determined by a sufficient statistic: agents' belief distortions. This result is both simple and only semi-structural: in order to calculate policy from the belief distortion, the policymaker does not need to know the whole macroeconomic model. They only need to know how beliefs and policies distort decisions. Crucially, they do not even need to know how expectations are formed; they only need to measure them. Next, I study several examples. In a behavioral RBC model, the optimal policy is to tax capital when agents are overly optimistic about future returns. In a behavioral New Keynesian model, the optimal policy is to raise interest rates when agents misperceive the economy to be running hot.

JEL-Codes: D84, E52, E61, E70
Keywords: Behavioral Expectations, Optimal Policy, Belief Distortions, Information Frictions, Monetary Policy, Fiscal Policy

---

*University of Florida. Website: www.jonathanjadams.com Email: adamsjonathan@ufl.edu. I am grateful for constructive feedback from Marco Airaudo, Philip Barrett, Michał Brzoza-Brzezina, Yu-Ting Chiang, Luigi Iovino, Andre Kurmann, Ricardo Reis, Eugenio Rojas, Mark Wright, Donghoon Yoo, seminar participants at the Kansas City Fed, St. Louis Fed, and Drexel University, and conference participants at the 2024 NTU-AS Symposium on Macroeconomics and Resilience, the 2nd BdF-CEPR-HEC-PSE Paris Conference on the Macroeconomics of Expectations, the 2024 Econometric Society NASM, and the 5th LB-NBP-CEBRA Conference. The behavioral models were solved using the BEET toolkit. All errors are my own.

1 Introduction

When agents do not have full information and rational expectations, macroeconomic outcomes are usually inefficient. How should policymakers respond to this problem? A traditional strategy would be to make assumptions about agents' biases when forming expectations, embed them in a structural macroeconomic model, and then calculate the policy rule that would maximize some welfare criterion. With this approach, the information costs and risk of specification error are high. In this paper, I show that there is a simpler alternative.

I demonstrate that the optimal policy in a general class of models is determined by a sufficient statistic: the belief distortion, i.e. the *ex ante* predictable forecast error made by agents without full information rational expectations (FIRE). If the macroeconomy satisfies a technical condition – *sentiment spanning* – then the optimal policy rule is linear in the belief distortion. The condition is satisfied when the policymaker has "enough" policy tools to offset all the channels through which beliefs distort decisions.

The optimal policy rule is also *semi-structural.* In general, the optimal policy can be implemented *without knowing exactly how agents form expectations.* This is valuable, because while there is strong evidence that agents do not precisely follow rational expectations, there is no consensus on how they *do* form expectations. Moreover, the policymaker does not need to know the full economic model. In order to calculate the policy rule, they only need to be able to measure belief distortions, and to know how belief distortions and policy instruments enter a subset of equilibrium conditions.

If the sentiment spanning condition is not satisfied, the policymaker's problem becomes harder. In this case, the optimal policy rule has two components: the belief distortion component and an economic distortion component. The belief distortion remains a sufficient statistic for the first component. And while the second component requires knowledge of the full macroeconomic model, I prove that it exactly follows the policy rule that would be optimal in the rational expectations version of the model. Thus an optimizing policymaker would need to calculate this second component anyways, and their additional response to non-rational expectations remains summarized by the belief distortion sufficient statistic.

To demonstrate the implications and tractability of this approach, I proceed to calculate optimal policy in several examples. First, I study a behavioral RBC (BRBC) model. The optimal policy is for capital taxes to move proportionately to belief distortions about utility-adjusted returns on investment. If agents mis-forecast future returns too high, or future marginal utility too low, then the capital tax should increase to discourage investment.

Next, I study behavioral New Keynesian models with and without the sentiment spanning condition. Optimal monetary policy is to raise interest rates when agents misperceive the economy as running "too hot": if agents mis-forecast future income too high, or future inflation too high, then monetary policy should contract. Notably, this is *not* what central banks do. Adams and Barrett (2024) show that the reverse is true: after a shock raising these belief distortions, the Federal Reserve *lowers* interest rates.

Lastly, I consider extensions to the baseline environment. First, I consider the possibility that policymakers cannot directly observe the belief distortion, perhaps because surveys measure expectations with error or because the estimates of the rational expectations are misspecified. The paper's main conclusions are robust to these concerns: the policymaker's nowcast replaces the belief distortion in the optimal policy rules. Second, I abstract from standard behavioral expectations and allow the expectation operator to be determined endogenously in a way that depends on the policy. This endogeneity does not change the main conclusions, although the theorem guaranteeing existence and uniqueness of the optimal policy no longer applies.

This paper builds on a large literature studying optimal policy in models without rational expectations. Throughout, I primarily focus on behavioral expectations as the generator of belief distortions. Recent research has adopted a variety of behavioral expectations that fit in this paper's framework in order to study optimal monetary policy, including sticky information (Woodford, 2010a), heterogeneous expectations (Di Bartolomeo, Di Pietro, and Giannini, 2016), heuristics (Hommes, Massaro, and Weber, 2019), cognitive discounting (Gabaix, 2020), and level-k thinking (Iovino and Sergeyev, 2023). But FIRE can be broken by relaxing full information as well, and the theoretical results also apply to some types of information frictions.[1]

The paper is organized as follows. Section 2 lays out the general framework and

[1] A voluminous literature studies optimal policy when agents have incomplete information. Lucas (1972) is seminal, but related recent work on monetary policy includes Adam (2007), Nimark (2008), Lorenzoni (2010), Baeriswyl and Cornand (2010), Paciello and Wiederholt (2014), Angeletos and La'O (2019), Benhima and Blengini (2020), and Angeletos, Iovino, and La'O (2020).

notation. Section 3 derives the optimal policy when the first-best equilibrium is achievable; Section 4 considers the case where it is not. Section 5 derives optimal policies in several examples. Section 6 explores the model extensions. Section 7 concludes.

2 A General Macroeconomic Model with Behavioral Expectations

This section introduces the general class of behavioral models and policies studied in this paper, and defines notation.

2.1 The General Model

Consider a general linear dynamic stochastic macroeconomic model of the following form. $X_t = \begin{pmatrix} X_{t-1}^K \\ X_t^C \end{pmatrix}$ is a $n \times 1$ vector of endogenous variables. n_K of the variables are predetermined state variables X_{t-1}^K, while $n_C = n - n_K$ are control variables X_t^C. Y_t is a vector of exogenous stochastic processes that are realized at time t; $Y_t = Y(L)\omega_t$ is a moving average in the exogenous white noise ω_t. The equilibrium conditions of the model are represented as a single matrix equation:

$$B_{X1}\mathbb{E}_t^b\left[X_{t+1}\right] = B_{X0}X_t + B_Y Y_t + B_G G_t \tag{1}$$

B_{X0}, B_{X1}, and B_Y are all matrices encoding the equilibrium conditions of the model.[2]

$\mathbb{E}_t^b[X_{t+1}]$ denotes a behavioral expectation operator that forecasts X_{t+1} conditional on information available at time t. The information set is the history $\{Y_{t-j}, X_{t-j}, \omega_{t-j}\}_{j=0}^{\infty}$. The superscript b indexes the type of behavioral expectations. When written without a superscript, $\mathbb{E}_t[\cdot]$ indicates the rational expectation. The state variable component of X_{t+1} is known exactly at time t, so the expected vector $\mathbb{E}_t^b[X_{t+1}]$ represents

$$\mathbb{E}_t^b[X_{t+1}] = \begin{pmatrix} X_t^K \\ \mathbb{E}_t^b\left[X_{t+1}^C\right] \end{pmatrix}$$

[2]Appendix C demonstrates that this representation nests more complicated behavioral models where additional objects appear inside the expectation operator.

for any form of expectations.

Adams (2023a) presents the technical details of the general behavioral operator $\mathbb{E}_t^b[\cdot]$, and explores its theoretical properties. The paper works through examples, showing how the operator nests many forms of behavioral expectations, including forward-looking forms such as diagnostic expectations (Bordalo, Gennaioli, and Shleifer, 2018), backward-looking heuristics such as adaptive expectations (Cagan, 1956), and simple information frictions with noisy signals. It allows for heterogeneity such as in sticky information models (Mankiw and Reis, 2002), so long as the average forecast is what enters the linearized equilibrium conditions. The operator does not necessarily allow for information frictions with endogenous noise, because while expectations are endogenous, the expectation operator $\mathbb{E}_t^b[\cdot]$ is assumed to be exogenous to policy. However, Section 6.2 relaxes this restriction, and allows the expectation operator to be endogenously determined; the main conclusions regarding optimal policy still hold.

G_t is the vector of *policy instruments*, such as interest rates, government spending, taxes, etc. The matrix B_G encodes how the policy instruments affect the model behavior.

A **behavioral expectations equilibrium** is defined as stationary time series for X_t, Y_t, and G_t, given a time series of shocks ω_t, such that:

1. X_t, Y_t, and G_t satisfy the equilibrium condition (1)
2. $Y_t = Y(L)\omega_t$
3. X_t and G_t are linear the history of shocks $\{\omega_{t-j}\}_{j=0}^{\infty}$
4. G_t satisfies a policy rule

Of course, this equilibrium definition is incomplete: it requires specifying a policy rule determining G_t.

This framework is very general, but what does it rule out? First, equilibria studied in this paper are stationary, so I cannot consider departures from rationality due to learning, as in Evans and Honkapohja (2003) or Orphanides and Williams (2008), among many others. Second, the behavioral expectation operator does not affect the coefficient matrices $(B_{X0}, B_{X1}, B_Y, B_G)$. This will be true if, for example, the model is linearized around a deterministic steady state, which is unaffected by behavioral expectations. Thus this is a model of dynamic mistakes, rather than steady state biases.

2.2 Additional Assumptions

Throughout the paper, stars denote the rational expectations equilibrium of a model, e.g. X_t^* is the solution to (1) when $\mathbb{E}_t^b$ is the rational expectation $\mathbb{E}_t$.

I assume – except when specified otherwise – that the rational expectation equilibrium X_t^* is the welfare-maximizing equilibrium, by whatever welfare objective the policymaker has in mind. In Section 4, the welfare objective is discussed further.

Next, assume the matrices defining equation (1) satisfy two regularity conditions:

1. *The Blanchard-Kahn condition is satisfied*, so that the rational expectations equilibrium X_t^* is unique.

2. $B_G' B_G$ *is invertible.* This implies that there are no redundant policy instruments that enter the model as a linear combination of other policy instruments. Allowing for additional redundant instruments does not change any conclusions, but complicates notation.

3 Optimal Policy With Sentiment Spanning

This section considers the case where the policymaker has "enough" policy instruments to recover the first-best equilibrium. The technical condition is that the model satisfies "sentiment spanning".

When this condition is satisfied, the policymaker's problem is easy. The optimal policy is given by a sufficient statistic, that does not require knowing either the exact way expectations are formed, nor the full economic model. This sufficient statistic is linear in agents' *belief distortions*, which is the difference between their forecast and the rational expectation. For agents forecasting with behavioral expectation $\mathbb{E}^b$, define the *distortion operator* $\mathbb{D}^b$ by

$$\mathbb{D}_t^b[X_{t+1}] \equiv \mathbb{E}_t^b[X_{t+1}] - \mathbb{E}_t[X_{t+1}]$$

This distortion is crucial to the optimal policy because of the assumption that the first-best equilibrium is a rational expectations equilibrium. In this section, I also assume that there is no policy in the first-best equilibrium, i.e. $G_t^* = 0$. This is a simplification in order to focus on the role of policy to resolve the distortions due to expectations, without the distraction from requiring policy to simultaneously address standard economic frictions. But this assumption is easy to relax, by first solving the traditional optimal policy problem for the rational expectations case, and then secondly considering deviations from this policy in order to resolve the behavioral distortions. The New Keynesian model in Section 5.2 is an example of this approach. Additionally, the more general theoretical results in Section 4 do not impose this assumption.

3.1 Optimal Policy and Belief Distortions

Before examining whether such a policy is feasible, first consider: what equations must a policy satisfy if it were to recover the rational expectations equilibrium X_t^*? Lemma 1 provides the answer.

Lemma 1 *If there is a time series of policy instruments G_t such that the non-rational equilibrium is consistent with the policy-less FIRE equilibrium, then G_t satisfies*

$$B_G G_t = B_{X1} \mathbb{D}_t^b [X_{t+1}] \tag{2}$$

Proof. In the FIRE equilibrium with $G_t = 0$, the endogenous vector X_t^* satisfies the equilibrium conditions

$$B_{X1} \mathbb{E}_t \left[X_{t+1}^* \right] = B_{X0} X_t^* + B_Y Y_t$$

Subtract this equation from the non-rational equilibrium conditions (1):

$$B_{X1} \mathbb{E}_t^b \left[X_{t+1} \right] - B_{X1} \mathbb{E}_t \left[X_{t+1}^* \right] = B_{X0} (X_t - X_t^*) + B_G G_t$$

Next, impose that $X_t = X_t^*$:

$$B_{X1} \mathbb{E}_t^b \left[X_{t+1} \right] - B_{X1} \mathbb{E}_t \left[X_{t+1} \right] = B_G G_t$$

Then rearranging and using the belief distortion operator gives:

$$B_G G_t = B_{X1} \mathbb{D}_t^b [X_{t+1}]$$

■

The proof is straightforward, following from the equilibrium conditions after assuming that X_t matches the first-best equilibrium X_t^* without any policy intervention. The crucial question remains: when does a policy satisfying equation (2) exist?

The next section answers: if the policy instruments satisfy the *sentiment spanning* condition.

3.2 The Sentiment Spanning Condition

What policy instruments are needed to recover the rational expectations equilibrium? Lemma 1 implies that the policy must offset the belief distortions $\mathbb{D}^b[X_{t+1}]$. Fortunately, belief distortions only enter a subset of a model's equilibrium conditions: the forward-looking equations. But, beliefs might distort different equations in different ways. In order to address all of the effects, there must be enough linearly independent policy instruments to offset the belief distortions in the forward-looking equations without creating new distortions in the model's other equations. *Sentiment spanning* is the technical condition that says whether this is possible.

Before stating the condition, additional notation is required. First, subdivide the matrix $B_{X1} \equiv \begin{pmatrix} B_{K1} & B_{C1} \end{pmatrix}$ into coefficients on states and controls. This is to take advantage of the fact that agents know the next period's state variables with certainty:

$$B_{X1}\mathbb{D}_t^b[X_{t+1}] = B_{X1}\mathbb{D}_t^b\left[\begin{pmatrix} 0 \\ X_{t+1}^C \end{pmatrix}\right] = B_{C1}\mathbb{D}_t^b[X_{t+1}^C]$$

Second, define $P_G \equiv B_G(B_G'B_G)^{-1}B_G'$, which is the projection matrix for the space spanned by the columns of B_G.

Condition 1 (Sentiment Spanning) *The macroeconomic model defined in* (1) *is said to satisfy sentiment spanning if*

$$(I - P_G)\, B_{C1} = 0$$

The matrix B_{C1} determines how belief distortions affect the model. The matrix $I - P_G$ projects onto the space orthogonal to the set of policy instruments. Thus Condition 1 says that there is no belief distortion that cannot be offset by some linear combination of policies.

What information does a practitioner have to know in order to evaluate sentiment spanning? They do not need to know the entire model nor how expectations are formed. They only need to know how expectations enter the forward-looking equations (encoded in B_{C1}) and how their policy instruments distort the economy (encoded in B_G, which determines P_G).

When is sentiment spanning satisfied? The policymaker needs at least as many policy instruments (the dimensions of G_t) that span the forward-looking equations (the non-zero rows of B_{C1}) as there are dimensions to the belief distortion $\mathbb{D}_t^b[X_{t+1}^C]$. For example, it is satisfied in the RBC model studied in Section 5.1, because there is one policy instrument and a single forward-looking equation. And it is satisfied in the New Keynesian model studied in Section 5.2, where there are two policy instruments, two forward-looking equations, and each policy affects a different equation.

3.3 Optimal Policy: The Sufficient Statistic

If the sentiment spanning condition is satisfied, then policymakers can recover the rational expectations equilibrium by following a simple policy rule: the measured belief distortion $\mathbb{D}_t^b[X_{t+1}^C]$ is a sufficient statistic for the optimal policy $G_t^\dagger$. Theorem 1 presents this result.

Theorem 1 *If a model satisfies Condition 1, then there exists a policy rule that recovers the FIRE equilibrium, and the policy rule is given by*

$$G_t^\dagger = (B_G' B_G)^{-1} B_G' B_{C1} \mathbb{D}_t^b[X_{t+1}^C] \tag{3}$$

Proof. Let X_t^* denote the equilibrium time series when agents have rational expectations and there is no policy intervention. Per equation (1):

$$B_{X1}\mathbb{E}_t[X_{t+1}^*] = B_{X0}X_t^* + B_Y Y_t$$

Conjecture that when G_t is given by equation (3), the behavioral equilibrium time series satisfies $X_t = X_t^*$. Take the difference between equation (1) under behavioral and rational expectations:

$$B_{X0}(X_t - X_t^*) = B_{X1}\left(\mathbb{E}_t^b[X_{t+1}] - \mathbb{E}_t[X_{t+1}^*]\right) - B_G G_t$$

Imposing the conjecture $X_{t+1} = X^*_{t+1}$ gives

$$B_{X0}(X_t - X^*_t) = B_{X1}\mathbb{D}^b_t[X_{t+1}] - B_G G_t$$

State variables are known exactly under both forms of expectations, so $B_{X1}\mathbb{D}^b_t[X_{t+1}] = B_{C1}\mathbb{D}^b_t[X^C_{t+1}]$:

$$B_{X0}(X_t - X^*_t) = B_{C1}\mathbb{D}^b_t[X^C_{t+1}] - B_G G_t$$

and substituting with equation (3) gives

$$B_{X0}(X_t - X^*_t) = B_{C1}\mathbb{D}^b_t[X^C_{t+1}] - P_G B_{C1}\mathbb{D}^b_t[X^C_{t+1}]$$

$$B_{X0}(X_t - X^*_t) = (I - P_G)B_{C1}\mathbb{D}^b_t[X^C_{t+1}]$$

and Condition 1 implies

$$B_{X0}(X_t - X^*_t) = 0$$

which is consistent with the conjecture. This logic says that if G_t is given by (3), then $X_t = X^*_t$ is a stationary time series that satisfies equation (1) for all t, so it is a behavioral expectations equilibrium.

Lastly, any choice of G_t does not affect the Blanchard-Kahn condition, which is determined by the matrices B_{X0} and B_{X1}. And the policy $G^\dagger_t$ reduces the model equations (1) to the rational expectations case with no policy:

$$G_t = G^\dagger_t \implies B_{X1}\mathbb{E}_t\left[X_{t+1}\right] = B_{X0}X_t + B_Y Y_t$$

so $X_t = X^*_t$ is the unique equilibrium. ■

The proof has two main steps. First, it shows that the sentiment spanning condition implies that it is possible to write a policy $G^\dagger_t$ that delivers the rational expectations equilibrium. Second, it proves that such a policy is feasible: the resulting equilibrium would exist and be unique. This second step is simple in this case, but will be more involved in Section 4 when sentiment spanning fails.

What do policymakers need to know in order to implement the policy $G^\dagger_t$? They need to be able to measure the belief distortion $\mathbb{D}^b_t[X^C_{t+1}]$, and they need to know the same information as is necessary to evaluate sentiment spanning: B_{C1} and B_G, the matrices that determine how expectations enter forward-looking equations and how their policy instruments distort the model. Thus Theorem 1 implies the belief distortion is a sufficient statistic for the optimal policy; in any period, the policymaker need only measure $\mathbb{D}_t^b[X_{t+1}^C]$ in order to choose $G_t^\dagger$.[3]

4 Optimal Policy Without Sentiment Spanning

The last section demonstrated that when a model satisfies the sentiment spanning condition, the optimal policy can recover the rational expectations equilibrium. But what if the condition is not satisfied?

Without sentiment spanning, the policymaker's goal is instead to get as close as possible to the first-best equilibrium, based on some welfare-relevant metric. This optimal policy problem loses the semi-structural property from Theorem 1; the policymaker now needs to know the whole model. But the modeling requirements are no stronger than in the rational expectations case. The policymaker still does not need to know how behavioral expectations are formed: *they only need to measure the belief distortion.*

Moreover, the optimal policy problem is cleanly decomposed into two parts. The first component solves the distortions to beliefs with a similar formula as in Section 3. The second component solves the economic distortion in exactly the same way as under rational expectations. Therefore a policymaker who does not have enough policy instruments to recover the first-best equilibrium does not actually face a harder optimization problem: they solve the problem they would have to solve under rational expectations anyway, and then address the belief distortion using the sufficient statistic.

To solve the optimal policy problem without reaching the first-best equilibrium, it is necessary to specify a welfare function. For simplicity, I focus on a standard problem whereby the policymaker commits to a policy rule that maximizes the unconditional expectation of some welfare function. This is commonly represented as minimizing a quadratic loss function (Rotemberg and Woodford, 1997):

$$\min \mathbb{E}\left[(X_t - X_t^*)' W (X_t - X_t^*)\right] \tag{4}$$

[3] This result contrasts with Woodford (2010b) and Adam and Woodford (2012), who study robust control (max-min) optimal policy responses to belief distortions in New Keynesian models. The source of the difference is that these papers restrict attention to policy rules that are functions of exogenous states, but not belief distortions which are allowed to move for extrinsic reasons. Theorem 1 says that the optimal policy rule is linear in the belief distortions themselves. This is one reason why the robust control optimal policy does not recover the FIRE equilibrium.

The matrix W encodes the policymaker's welfare weights on unconditional covariances. The rational expectations equilibrium X_t^* remains the first-best equilibrium, so it is convenient to write this loss function in terms of deviations from X_t^*. While I maintain the assumption that X_t^* is a rational expectations equilibrium, I drop the simplifying assumption from Section 3 that the optimal policy under rational expectations is zero. Rather, the FIRE optimal policy is now denoted by G_t^*.

Theorem 2 gives the solution to this second-best policy problem. In order to express the solution, define the projection matrix $P_W \equiv B_G \left(B_G' \tilde{W} B_G\right)^{-1} B_G' \tilde{W}$ where $\tilde{W} \equiv \left(B_{X0}^{-1}\right)' W B_{X0}^{-1}$. Sentiment spanning fails, so the matrix B_G is tall; denote its pseudo-inverse by $B_G^+ \equiv (B_G' B_G)^{-1} B_G'$.

Theorem 2 *The constrained-optimal policy rule is*

$$G_t^\dagger = B_G^+ P_W \left(B_{C1} \mathbb{D}_t^b[X_{t+1}^C] + B_{X1} \mathbb{E}_t[X_{t+1} - X_{t+1}^*]\right) + G_t^*$$

Proof. Substitute $\tilde{W}$ into the objective function:

$$\mathbb{E}\left[(X_t - X_t^*)' W (X_t - X_t^*)\right] = \mathbb{E}\left[(X_t - X_t^*)' B_{X0}' \tilde{W} B_{X0} (X_t - X_t^*)\right]$$

$$= \mathbb{E}\left[(X_t - X_t^*)' B_{X0}' C_{\tilde{W}}' C_{\tilde{W}} B_{X0} (X_t - X_t^*)\right]$$

using the Cholesky decomposition $\tilde{W} = C_{\tilde{W}}' C_{\tilde{W}}$. This can be written in terms of the policy using equation (1), which implies $B_{X0} X_t - B_{X0} X_t^* = -B_G \hat{G}_t + B_{X1} E_t^b[X_{t+1}] - B_{X1} E_t[X_{t+1}^*]$ where $\hat{G}_t \equiv G_t - G_t^*$:

$$= \mathbb{E}\Big[(-B_G \hat{G}_t + B_{X1} \mathbb{E}_t^b[X_{t+1}] - B_{X1} \mathbb{E}_t[X_{t+1}^*])' C_{\tilde{W}}'$$
$$C_{\tilde{W}} (-B_G \hat{G}_t + B_{X1} \mathbb{E}_t^b[X_{t+1}] - B_{X1} \mathbb{E}_t[X_{t+1}^*])\Big]$$

$$= \mathbb{E}\Big[(-B_G \hat{G}_t + B_{X1} \mathbb{D}_t^b[X_{t+1}] + B_{X1} \mathbb{E}_t[X_{t+1} - X_{t+1}^*])' C_{\tilde{W}}'$$
$$C_{\tilde{W}} (-B_G \hat{G}_t + B_{X1} \mathbb{D}_t^b[X_{t+1}] + \mathbb{E}_t[X_{t+1} - X_{t+1}^*])\Big]$$

$$= \mathbb{E}\Big[(-C_{\tilde{W}}B_G\hat{G}_t + C_{\tilde{W}}B_{C1}\mathbb{D}_t^b[X_{t+1}^C] + C_{\tilde{W}}B_{X1}\mathbb{E}_t[X_{t+1} - X_{t+1}^*]))'$$
$$(-C_{\tilde{W}}B_G\hat{G}_t + C_{\tilde{W}}B_{C1}\mathbb{D}_t^b[X_{t+1}^C] + C_{\tilde{W}}B_{X1}\mathbb{E}_t[X_{t+1} - X_{t+1}^*])\Big] \quad (5)$$

Thus the objective is to choose $\hat{G}_t$ to minimize equation (5). Written this way, it is a standard least squares minimization problem, which $\hat{G}_t$ solves by projecting $C_{\tilde{W}}B_{C1}\mathbb{D}_t^b[X_{t+1}^C] + C_{\tilde{W}}B_{X1}\mathbb{E}_t[X_{t+1} - X_{t+1}^*]$ onto the space spanned by $C_{\tilde{W}}B_G$:

$$C_{\tilde{W}}B_G\hat{G}_t =$$
$$C_{\tilde{W}}B_G\left((C_{\tilde{W}}B_G)'C_{\tilde{W}}B_G\right)^{-1}(C_{\tilde{W}}B_G)'\left(C_{\tilde{W}}B_{C1}\mathbb{D}_t^b[X_{t+1}^C] + C_{\tilde{W}}B_{X1}\mathbb{E}_t[X_{t+1} - X_{t+1}^*]\right)$$

Left-multiply by $C_{\tilde{W}}^{-1}$ and substitute with $P_W = B_G\left(B_G'\tilde{W}B_G\right)^{-1}B_G'\tilde{W}$:

$$B_G\hat{G}_t = P_W\left(B_{C1}\mathbb{D}_t^b[X_{t+1}^C] + B_{X1}\mathbb{E}_t[X_{t+1} - X_{t+1}^*]\right)$$

Finally, left-multiplying by B_G^+ gives

$$\hat{G}_t = B_G^+P_W\left(B_{C1}\mathbb{D}_t^b[X_{t+1}^C] + B_{X1}\mathbb{E}_t[X_{t+1} - X_{t+1}^*]\right)$$

and substituting in $G_t = \hat{G}_t + G_t^*$ completes the proof. ∎

The proof strategy is to find the optimal policy $G_t^\dagger$ that minimizes the welfare-relevant distance between the equilibrium X_t and the rational expectations counterfactual X_t^*. This is done by projecting the expectation terms onto the space spanned by the policy instruments. When sentiment-spanning is satisfied, this projection is perfect: P_W is the identity, and $X_{t+1} = X_{t+1}^*$. Thus Theorem 1 is a special case of Theorem 2.

The policymaker's problem is now more involved, compared to Section 3. The belief distortion $\mathbb{D}_t^b[X_{t+1}]$ is no longer a sufficient statistic for the optimal policy. Now, the policymaker must also evaluate the rational expectation of the economic distortion $\mathbb{E}_t[X_{t+1} - X_{t+1}^*]$. And minimizing the objective function (4) requires responding to both the belief distortion and the economic distortion.

But this additional concern is already well understood! Let G_t^{RE} denote the policy following the rule that *would* be optimal if agents in the model (1) had rational expectations. Corollary 1 shows that the component of $G_t^\dagger$ that responds to the distortion $\mathbb{E}_t[X_{t+1} - X_{t+1}^*]$ follows exactly the same policy rule as G_t^{RE}.

Corollary 1 *The constrained-optimal optimal policy can be written as*

$$G_t^\dagger = B_G^+ P_W B_{C1} \mathbb{D}_t^b[X_{t+1}^C] + G_t^{RE}$$

Proof. Under rational expectations, there is no belief distortion, i.e. $\mathbb{D}_t^b[X_{t+1}] = 0$. Per Theorem 2, the optimal policy under rational expectations is

$$G_t^{RE} = B_G^+ P_W B_{X1} \mathbb{E}_t[X_{t+1} - X_{t+1}^*] + G_t^* \tag{6}$$

Then, for any behavioral expectation $\mathbb{E}_t^b$, substitute (6) into the expression from Theorem 2. ■

Corollary 1 implies that the *additional* policy challenge introduced by behavioral expectations is resolved by the policy rule $B_G^+ P_W B_{C1} \mathbb{D}_t^b[C_{t+1}]$, for which the belief distortion is still a sufficient statistic. And the information requirements are light, as in Section 3: the exact form of behavioral expectations do not need to be known so long as belief distortions can be measured. However, in order to calculate the sufficient statistic, the policymaker now has to know P_W, which depends on the matrix B_{X0}.

G_t^{RE} follows a rational expectations policy rule. Equation (6) expresses the policy rule as linear in the expectation of economic distortions, plus the FIRE optimal policy G_t^*. With rational expectations, the first term $B_G^+ P_W B_{X1} \mathbb{E}_t[X_{t+1} - X_{t+1}^*]$ is zero, because the G_t^* component ensures $X_t = X_t^*$ under FIRE. What does this term represent without FIRE? It is the constrained optimal policy under FIRE among policy rules that are linear in the expected gap $\mathbb{E}_t[X_{t+1} - X_{t+1}^*]$. The matrix $B_G^+ P_W B_{X1}$ encodes this rational rule.

5 Examples

This section explores the sufficient statistics for optimal policy in several example economies.[4]

In the quantitative exercise that follows, I consider three forms of expectations,

[4]The models are solved using BEET (Adams, 2024), a toolkit for dynamic models with behavioral expectations. Optimal policy in this section is derived explicitly, but the toolkit is also capable of calculating the optimal policy rule automatically.

defined as:

$$\text{Rational Expectations:} \quad \mathbb{E}_t^{RE}[x_{t+1}] = \mathbb{E}_t[x_{t+1}] \tag{7}$$

$$\text{Diagnostic Expectations:} \quad \mathbb{E}_t^{DE}[x_{t+1}] = (1+\theta^{DE})\mathbb{E}_t[x_{t+1}] - \theta^{DE}\mathbb{E}_{t-1}[x_{t+1}] \tag{8}$$

$$\text{Cognitive Discounting:} \quad \mathbb{E}_t^{CD}[x_{t+1}] = \theta^{CD}\mathbb{E}_t[x_{t+1}] \tag{9}$$

Cognitive discounting is useful for resolving a number of puzzles in New Keynesian theory (Gabaix, 2020). If $\theta^{CD} < 1$ (as in the original calibration) forecasts underreact to news, but $\theta^{CD} > 1$ is also allowed, in which case forecasters "overextrapolate" (Angeletos, Huo, and Sastry, 2021). Cognitive discounting is also a convenient choice because it is isomorphic to simple models where agents mis-forecast due to information frictions instead of behavioral biases (Adams, 2023a).

These examples are just a subset of possible behavioral expectations; the main conclusions from Theorems 1 and 2 apply more generally.

5.1 A Behavioral Real Business Cycle (BRBC) Model

This section considers a standard RBC model (Kydland and Prescott, 1982) modified so that households have behavioral expectations. The BRBC model is a convenient initial example for two reasons. First, the rational expectations equilibrium is efficient without any policy intervention. Second, the model has only one forward-looking equation, so only one policy instrument is needed to satisfy sentiment spanning and recover the efficient equilibrium.

The log-linearized model is given by:[5]

$$\text{Euler Equation:} \quad \tau_t = \sigma c_t + \mathbb{E}_t^b[-\sigma c_{t+1} + \overline{R} r_{t+1}] \tag{10}$$

$$\text{Labor Supply:} \quad w_t = \sigma c_t + \eta n_t \tag{11}$$

$$\text{Production Function:} \quad y_t = a_t + \alpha k_{t-1} + (1-\alpha) n_t \tag{12}$$

$$\text{Capital Demand:} \quad r_t = y_t - k_{t-1} \tag{13}$$

$$\text{Labor Demand:} \quad w_t = y_t - n_t \tag{14}$$

$$\text{Resource Constraint:} \quad \overline{Y} y_t = \overline{C} c_t + \overline{K}\left(k_t - (1-\delta)k_{t-1}\right) \tag{15}$$

where uppercase constants denote steady state values, and lowercase variables denote

[5] Appendix B derives the log-linearized BRBC model.

log-linearized deviations. This system of equations uniquely determines the time series for the 6 endogenous variables (consumption c_t, labor l_t, output y_t, capital k_t, real wage w_t and real rental rate r_t), given time series for capital a_t and the policy instrument τ_t. In this setting, τ_t is any policy that affects the economy by creating an "intertemporal wedge". This could be any number of policies that only distort the intertemporal margin (Chari, Kehoe, and McGrattan, 2007), but I will refer to τ_t as an *investment tax*.[6]

This BRBC model is *decentralized*: it includes prices. This is purposeful for two reasons. First, the decentralized and social planner's models are ordinarily isomorphic, *but this is not true with behavioral expectations*, and the social planner's problem is usually considered an analytical convenience rather than representing a real decisionmaker. For example, it may be that:

$$\mathbb{E}_t^b[r_{t+1}] \neq \mathbb{E}_t^b[y_{t+1} - k_t]$$

This is because k_t is known at time t, and for many behavioral expectations, the forecast bias about y_{t+1} is not the same as the forecast bias about r_{t+1}, which has different time series properties. When writing down a behavioral macroeconomic model, the theorist must take a stance on precisely which variables are forecasted by agents in the model.

Second, writing the decentralized problem illustrates one of the advantages of the sufficient statistic approach. The BRBC model has many contemporaneous equations but only one forward-looking equation. And Theorem 1 implies that the policymaker does not need to know all of these equations in order to conduct optimal policy: they do not care what the production function is, or details of the labor market, etc. They only need to know where the policy causes distortions.

In this model, the investment tax τ_t only distorts the Euler equation. Thus sentiment spanning is satisfied, so Theorem 1 implies that the policymaker can recover the efficient equilibrium by setting $B_G G_t = B_{C1}\mathbb{D}_t^b[X_{t+1}^C]$. Appendix B demonstrates how to map the model equations to these matrices. But it is simple to see what the optimal policy must be to offset a belief distortion:

$$\tau_t = \mathbb{D}_t^b[-\sigma c_{t+1} + \overline{R} r_{t+1}]$$

[6] If taken literally, the tax receipts must be returned to households as transfers so as not to distort the budget constraint.

The optimal investment tax responds to two mistakes. If agents are over-optimistic and expect the return on capital investment to be too high ($\mathbb{D}_t^b[\overline{R}r_{t+1}] > 0$) they will build too much capital, so the tax needs to increase to disincentivize investment. Conversely, if agents are over-optimistic about future consumption ($\mathbb{D}_t^b[-\sigma\Delta c_{t+1}] < 0$) then they will build too little capital, so the tax needs to decrease.[7]

To illustrate how this approach works in practice, Figure 1 plots impulse response functions to a positive productivity shock in three example economies. In the first column, agents have rational expectations. After the shock, firms are more productive, so expectations of investment returns rise ($\mathbb{E}_t[\overline{R}r_{t+1}]$, upper left panel). Incomes are higher, so households consume more, and expected consumption rises ($\mathbb{E}_t[c_{t+1}]$, middle left panel). The rational expectations equilibrium is efficient, so the optimal policy is to implement no investment tax.

In the second column, agents have diagnostic expectations, parameterized with $\theta^{DE} = 0.5$. When the shock occurs, agents immediately over-forecast future productivity. Thus they over-forecast investment returns (middle upper panel) and consumption (center panel) relative to the rational expectation. But after one period, their expectations respond rationally to the lagged shock. Without a policy response, when the shock impacts there would not be enough investment; the income effect dominates the substitution effect and households consume would too much ($\mathbb{D}_t^{DE}[\sigma c_{t+1}] > \mathbb{D}_t^{DE}[\overline{R}r_{t+1}]$), so the optimal policy is to subsidize investment for a single period (middle lower panel).

In the third column, agents follow cognitive discounting, parameterized with $\theta^{CD} = 0.5$. Agents' productivity forecasts are now attenuated. Thus they under-forecast investment returns (upper right panel) and consumption (middle right panel) relative to the rational expectation. In contrast to diagnostic expectations, this biased forecasting continues to occur in every period. Without policy, households would perpetually under-consume, because they would expect low consumption in the future; again, the income effect dominates the substitution effect ($-\mathbb{D}_t^{CD}[\sigma c_{t+1}] > -\mathbb{D}_t^{CD}[\overline{R}r_{t+1}]$). Thus there would be too much investment, so the optimal policy is to tax capital, and this tax remains high over the course of the economic boom.

[7] Even though this paper's framework precludes some types of information frictions with endogenous signals, Adams (2023b) studies optimal policy in such a setting, and comes to qualitatively similar conclusions about capital taxation.

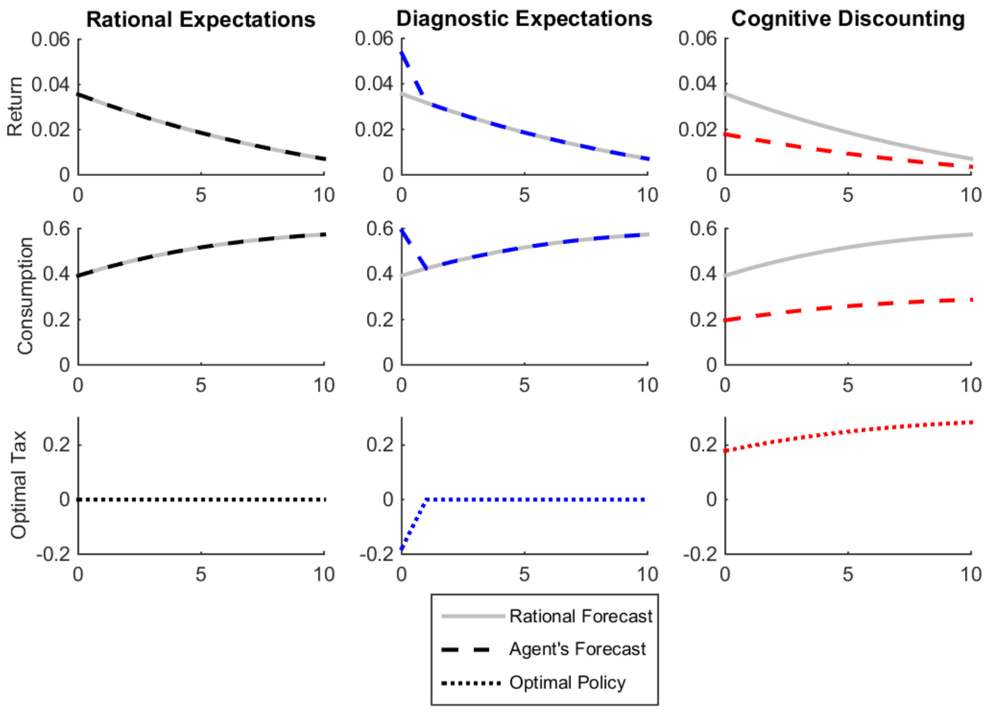

Figure 1: Response of Expectations to a Productivity Shock in Three RBC Models

Each panel plots the impulse response of expectations to a productivity shock. Each column corresponds to a different model of expectations: rational, diagonstic, and cognitive discounting. The first row is forecasts of the one-period-ahead investment return $\overline{R}r_{t+1}$, the second row is consumption, and the third row is the implied optimal investment tax. In all cases, the RBC model is parameterized as follows: $\beta = 0.99$, $\sigma = 1$, $\delta = 0.02$, $\alpha = 0.33$, $\eta = 1$, and productivity follows an AR(1) process with autocorrelation 0.95. The behavioral parameters are $\theta^{DE} = 0.5$ and $\theta^{CD} = 0.5$.

5.2 First-Best Policy in a Behavioral New Keynesian Model

The canonical New Keynesian Model has two forward-looking equations: a Phillips Curve and an Euler equation. Thus two policy variables will be needed to satisfy sentiment spanning.

I modify a New Keynesian model with government spending used by Gnocchi (2013) to study optimal monetary and fiscal policy.[8] When expectations are possibly non-rational, the log-linearized behavioral model is given by:

$$\text{New Keynesian Phillips Curve:} \qquad \pi_t = \kappa y_t - \psi f_t - z_t^{PC} + \beta \mathbb{E}_t^b[\pi_{t+1}] \tag{16}$$

$$\text{Euler Equation:} \qquad \sigma y_t = -i_t - z_t^{EE} + \mathbb{E}_t^b[\sigma y_{t+1} + \pi_{t+1}] \tag{17}$$

where inflation π_t and the output gap y_t are determined by exogenous processes z_t^{PC} and z_t^{EE}, and the policy instruments. The central bank sets the nominal interest rate i_t, while the fiscal authority chooses government spending g_t to set the fiscal gap $f_t = g_t - y_t$.

In this model, the rational expectations solution is not optimal (or unique) in the absence of a policy intervention. Let f_t^* and i_t^* denote the optimal policies under rational expectations, in the style of King (2000). Then decompose the policy instruments into the optimal choices and deviations, denoted with hats:

$$f_t = \hat{f}_t + f_t^* \qquad\qquad i_t = \hat{i}_t + i_t^*$$

Written in the matrix form of equation (1), the model becomes

$$\underbrace{\begin{pmatrix} 1 & -\kappa \\ 0 & \sigma \end{pmatrix}}_{B_{X0}} \underbrace{\begin{pmatrix} \pi_t \\ y_t \end{pmatrix}}_{X_t} + \underbrace{\begin{pmatrix} \psi & 0 \\ 0 & 1 \end{pmatrix}}_{B_G} \underbrace{\begin{pmatrix} \hat{f}_t \\ \hat{i}_t \end{pmatrix}}_{G_t} + \underbrace{\begin{pmatrix} \psi & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{pmatrix}}_{B_Y} \underbrace{\begin{pmatrix} f_t^* \\ i_t^* \\ z_t^{PC} \\ z_t^{EE} \end{pmatrix}}_{Y_t} = \underbrace{\begin{pmatrix} \beta & 0 \\ 1 & \sigma \end{pmatrix}}_{B_{X1}} \mathbb{E}_t^b \left[X_{t+1} \right]$$

[8] As discussed in Section 5.1, different microfoundations and informational assumptions may lead models that are isomorphic under rational expectations to have different log-linear representations under behavioral expectations. This version of the behavioral model is derived in Adams (2023a), with the addition of government spending. It differs from Gabaix (2020) – where agents mis-forecast their incomes but correctly forecast inflation – and from L'Huillier, Singh, and Yoo (2023) where agents mis-forecast the price level instead of the inflation rate.

Written this way, the policies f_t^* and i_t^* are treated as exogenous stochastic processes. This means that the model is technically indeterminate, so assume that some other unmodeled policy choice selects the appropriate welfare-maximizing equilibrium.[9]

This model has two forward looking equations and one policy instrument that affects each: the sentiment spanning condition is satisfied.[10] Therefore, Theorem 1 says a policy following $G_t^\dagger = B_G^{-1} B_{X1} \mathbb{D}_t^b[X_{t+1}]$ can recover the rational expectations equilibrium. This optimal policy is:

$$\hat{f}_t^\dagger = \frac{\beta}{\psi} \mathbb{D}_t^b \left[\pi_{t+1}\right] \qquad \hat{i}_t^\dagger = \mathbb{D}_t^b \left[\pi_{t+1} + \sigma y_{t+1}\right] \tag{18}$$

The optimal interest rate policy says that if agents' forecasts of income are too high ($\mathbb{D}_t^b \left[y_{t+1}\right] > 0$), they will mistakenly consume too much today, so the central bank should raise interest rates to reduce aggregate demand.

However, if agents' forecasts of inflation are too high ($\mathbb{D}_t^b \left[\pi_{t+1}\right] > 0$), this causes two different problems. First, this lowers expectations of the real interest rate, raising aggregate demand, so the central bank needs to respond by raising the nominal rate. Second, the inflation mis-forecast distorts price-setting through the Phillips Curve, which is not directly affected by monetary policy. Expecting high future inflation, firms would like to set prices to raise inflation immediately; this needs to be offset by increasing the fiscal gap $\hat{f}_t$, which lowers inflation in the Gnocchi (2013) model by lowering marginal costs.

Figure 2 demonstrates how straightforward these optimal policies are to implement. The first two panels plot the inflation and real income belief distortions in the data. These belief distortions are calculated by Adams and Barrett (2024), who take household expectations from the Michigan Survey of Consumers and remove a rational expectation estimated from a semi-structural VAR. These two belief distortions are very negatively correlated: households' average expectations of inflation and nominal income have low correlation, so their implied expectations of real income tend to fall when their forecasts of inflation rise.

The remaining panels plot the policies implied by the optimal rule. These are *not*

[9] This could be some off-equilibrium threat by the central bank such as a Taylor rule $i_t = i_t^* + \phi(\pi_t - \pi_t^*)$ with ϕ sufficiently large (King, 2000), an appropriate fiscal policy (Cochrane, 2023), or an arbitrarily small departure from perfect memory (Angeletos and Lian, 2023).

[10] Mathematically: B_G is invertible, so the policy dimensions span the endogenous variable dimensions, i.e. $I - B_G(B_G' B_G)^{-1} B_G' = 0$.

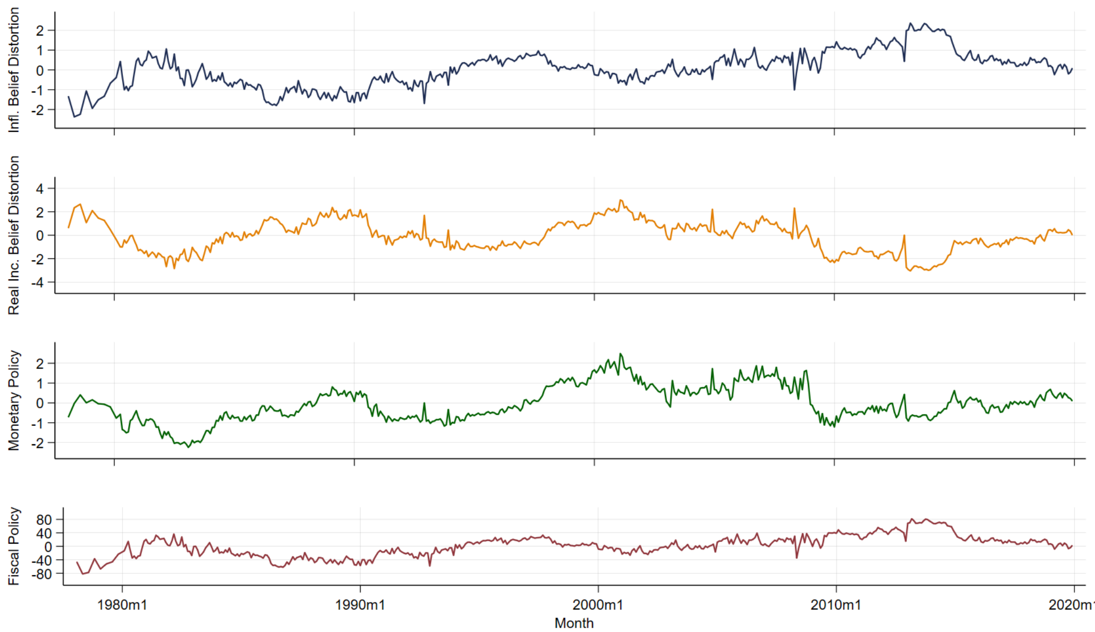

Figure 2: Estimated Belief Distortions and Implied Policies

From top to bottom, the panels plot the inflation belief distortion $\mathbb{D}^b[\pi_{t+1}]$, income belief distortion $\mathbb{D}^b[y_{t+1}]$, nominal interest rate i_t, and fiscal gap f_t. Belief distortions are calculated as in Adams and Barrett (2024). Policies are calculated per equation (18). In all cases, the units are percentage point deviations from trend. The plot truncates before the COVID pandemic, when very large belief distortions dwarf the preceding time series.

counterfactuals: if the policy rule were implemented, the equilibrium belief distortions could change. Rather, these time series demonstrate how tractable the optimal policy is to estimate and calculate.

The third panel (green series) plots the nominal interest rate deviation implied by the optimal policy rule and the measured belief distortions. Equation (18) is parameterized as in Gnocchi (2013): $\beta = 0.99$, $\sigma = 1$, $\psi = 0.03$. When agents mis-forecast the economy as too hot (both belief distortions positive), the prescription for monetary policy is to tighten: for example, this occurs from 2000-2007. When the belief distortions move in different directions, it depends on which one dominates: for example, in the 1982 recession, inflation belief distortions are somewhat positive but income belief distortions are very negative, so monetary policy needs to loosen.

The fourth panel (maroon series) plots the implied optimal fiscal gap. Per equation (18), optimal fiscal policy only depends on the inflation belief distortion, so this series mirrors the blue series. To be clear, these are not counterfactual policies. If the optimal policy rule were actually implemented, all four of these time series will change. To calculate counterfactuals would require taking a stance on the actual form of behavioral expectations in this economy. But following the optimal policy rule only requires measuring the belief distortions.

A consequence of this calibration is that implied fiscal policy is extremely volatile. In this model, government spending affects the Phillips Curve by increasing labor supply through an income effect, and thus reducing marginal costs. But the elasticity of this channel is low, and a policymaker may not be willing or able to stabilize inflation by such large swings in fiscal policy. Therefore the next section considers the case where only monetary policy is available to respond to belief distortions.

5.3 Constrained-Optimal Policy in a Behavioral New Keynesian Model

The sentiment spanning condition was satisfied in the previous section because both monetary and fiscal policy were available to stabilize the economy. What does optimal policy look like if this condition fails and the sufficient statistic from Theorem 1 cannot be applied?

To study this case, suppose that only monetary policy is available to respond to belief distortions. Now, the only policy decision is to choose the interest rate process

i_t.

The policymaker's objective is to choose a rule for interest rates that maximizes the unconditional expectation of consumer welfare. Rotemberg and Woodford (1997) derive the quadratic approximation of welfare w_t in this economy:

$$\mathbb{E}[w_t] = b_\pi Var(\pi_t) + b_y Var(y_t)$$

where $Var(\cdot)$ denotes the unconditional variance, and the positive coefficients b_π and b_y are functions of the model parameters. Mapping this objective to the loss function defined in (4), the policymaker's goal is

$$\min \mathbb{E}\left[\begin{pmatrix} \pi_t - \pi_t^* \\ y_t - y_t^* \end{pmatrix}' \underbrace{\begin{pmatrix} b_\pi & 0 \\ 0 & b_y \end{pmatrix}}_{W} \begin{pmatrix} \pi_t - \pi_t^* \\ y_t - y_t^* \end{pmatrix}\right]$$

which follows from $\pi_t^* = y_t^* = 0$. Proposition 1 gives the solution to this problem, where again i_t^* denotes the FIRE-optimal interest rate.

Proposition 1 *The constrained-optimal monetary policy rule is*

$$i_t^\dagger = \left(1 - \beta\frac{b_\pi \kappa\sigma}{b_\pi\kappa^2 + b_y}\right)\left(\mathbb{D}_t^b[\pi_{t+1}] + \mathbb{E}_t[\pi_{t+1} - \pi_{t+1}^*]\right) + \sigma\left(\mathbb{D}_t^b[y_{t+1}] + \mathbb{E}_t[y_{t+1} - y_{t+1}^*]\right) + i_t^* \quad (19)$$

Proof: Appendix A.1

This policy rule is more complicated than in Section 5.2. However, Corollary 1 says that there is a simpler, more intuitive way of representing it. In the absence of any distortion caused by behavioral expectations, the second best policy can be written as

$$i_t^{RE} = \left(1 - \beta\frac{b_\pi \kappa\sigma}{b_\pi\kappa^2 + b_y}\right)\mathbb{E}_t[\pi_{t+1} - \pi_{t+1}^*] + \sigma\mathbb{E}_t[y_{t+1} - y_{t+1}^*] + i_t^*$$

Therefore the optimal policy is concisely written using its deviation from the rational expectations policy rule:

$$i_t^\dagger - i_t^{RE} = \left(1 - \beta\frac{b_\pi \kappa\sigma}{b_\pi\kappa^2 + b_y}\right)\mathbb{D}_t^b[\pi_{t+1}] + \sigma\mathbb{D}_t^b[y_{t+1}] \quad (20)$$

In contrast to the first best solution (18), this second-best solution includes additional dependence on the inflation belief distortion. The $-\beta\frac{b_\pi\kappa\sigma}{b_\pi\kappa^2+b_y}\mathbb{D}_t^b[\pi_{t+1}]$ term attempts to correct for the distortion to price-setting that fiscal policy would resolve if it were available. This term is more important when prices are especially sticky (κ large) and when inflation is given more weight in the welfare function (b_π large).

In typical applications, the difference between the first-best interest rate rule and the constrained-optimal rule is small. $\beta\frac{b_\pi\kappa\sigma}{b_\pi\kappa^2+b_y}$ is usually close to zero because κ is thought to be small. Thus, the data-implied policy rule nearly exactly follows the time series plotted in Figure 2. However, if the policymaker only cared about stabilizing inflation ($b_y \approx 0$), the coefficient on inflation could become negative; indeed, this is the case for the Gnocchi (2013) calibration, where $\frac{b_\pi}{b_y} \approx 17$. In contrast, if they only cared about stabilizing the real economy ($b_\pi = 0$) they could do so by solely using interest rates to resolve the distortion to the Euler equation; in this case, the coefficients on the belief distortions are the same as in Section 5.2.

To illustrate how the constrained-optimal monetary policy works in practice, consider the effect of "cost-push" shocks to the Phillips curve, i.e. shocks to z_t^{PC}. Mirroring the BRBC analysis, Figure 3 plots impulse response functions to a cost-push shock in three example economies with different types of expectations. Unlike the BRBC model, sentiment spanning is not satisfied: monetary policy alone is insufficient to offset the distortions caused by behavioral expectations, so the rational expectation (gray lines) now vary across the three economies.

In the first column, agents have rational expectations, and the cost-push shock has its usual effect. The shock causes firms to raise prices (upper left panel). The central bank optimally responds by raising nominal rates, engineering a recession (middle left panel). The central bank is willing to tolerate the loss of income because it reduces inflation, although not enough to entirely offset the shock. The third row plots the optimal policy rule's implied response of interest rates. The *rational component* i_t^{RE} resolves the economic distortion as well as possible. When agents are rational, the *belief distortion component* $\hat{i}_t$ (dotted line) is zero.

In the second column, agents have diagnostic expectations, parameterized with $\theta^{DE} = 0.5$. When the shock occurs, agents are immediately too pessimistic: they over-forecast future inflation (middle upper panel) and under-forecast income (center panel) relative to the rational expectation. But this belief distortion only last for one period; their expectations are rational thereafter. The third row plots the optimal

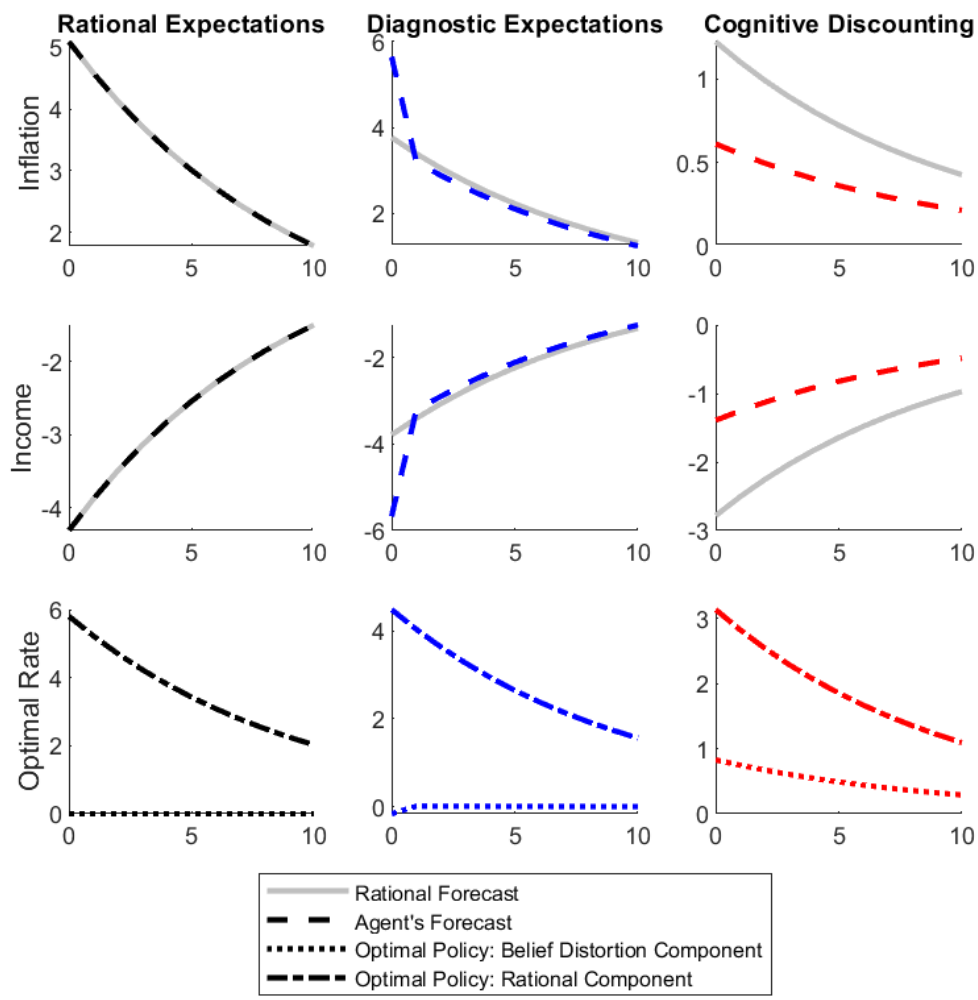

Figure 3: Response of Expectations to a Cost-Push Shock in Three NK Models

Each panel plots the impulse response of expectations to a cost-push shock. Each column corresponds to a different model of expectations: rational, diagnostic, and cognitive discounting. The first row is forecasts of the inflation rate π_{t+1}, the second row is forecasts of income y_{t+1}, and the third row is the implied optimal interest rate $i_t^{\dagger}$. In all cases, the BNK model is parameterized as follows: $\beta = 0.99$, $\sigma = 1$, $\kappa = 0.08$, $\psi = 0.03$, $b_y = b_\pi = 0.5$, and productivity follows an AR(1) process with autocorrelation 0.9. The behavioral parameters are $\theta^{DE} = 0.5$ and $\theta^{CD} = 0.5$.

response of interest rates. As before, the rational component rises to offset the booming inflation, per the standard optimal policy rule. But now, agents have non-zero belief distortions, so the belief distortion component must also move. Agents forecast inflation to be too hot (which should imply a monetary tightening) but also forecast income to be too low (which should imply loosening). These incentives somewhat offset, but the income distortion dominates in the example calibration with equal welfare weights $b_y = b_\pi = 0.5$. The belief distortion component of monetary policy loosens slightly, but only for the one period that forecasts are distorted.

In the third column, agents follow cognitive discounting, parameterized with $\theta^{CD} = 0.5$. Agents' cost-push forecasts now attenuated. Thus they under-forecast inflation (upper right panel) and over-forecast income (middle right panel) relative to the rational expectation. Because this biased forecasting continues to occur in every period, the monetary policy response to the belief distortions continues as well. Again, optimal policy responds primarily to offset the income belief distortion. The attenuated forecasts imply that agents' income belief distortion is positive, so the belief distortion component of monetary policy is contractionary.

These examples demonstrate that the belief distortion component of policy may either offset or reinforce the standard policy response to a shock. The direction depends on whether the behavioral bias amplifies the effect of a shock on expectations – as in the diagnostic expectations case – or attenuates the effect of a shock on expectations – as in the cognitive discounting case. The simplicity of the sufficient statistic is useful because the size of this under- or overreaction is all the policymaker needs to know. This is true even when the expectations are more complicated: Angeletos, Huo, and Sastry (2021) argue that forecasts overreact and underreact at different horizons. The implication for optimal policy is straightforward: policy should also overreact and underreact at those horizons, relative to the rational component.

6 Extensions

This section explores the consequences of relaxing some assumptions made in Section 2's general model. Section 6.1 drops the assumption that the policymaker perfectly observes the belief distortion, and Section 6.2 drops the assumption that the policy rule does not affect expectation formation.

6.1 Measurement Error

Thus far, the policymaker has been able to observe the belief distortion. But a reasonable concern is that the measured belief distortion contains error. This could be due to measurement error in surveyed expectations, or specification error in estimating the rational expectation.

Fortunately, the main conclusions from Section 4 are robust to the existence of measurement error. Theorem 3 states this result formally.

Assume now that the policymaker's observation D_t of the belief distortion is given by

$$D_t = \xi \mathbb{D}_t^b[X_{t+1}] + \upsilon_t$$

with i.i.d. measurement error $\upsilon_t \sim N(0, \sigma_\upsilon^2)$ and constant coefficient matrix ξ.[11] The measurement error may prevent recovery of the FIRE equilibrium even if sentiment spanning were to hold. Therefore it is necessary to choose a welfare function (4) as in Section 4, and I do not assume Condition 1 (Sentiment Spanning) in the results that follow.

Let $\hat{D}_t$ denote the policymaker's nowcast of the belief distortion $\mathbb{D}_t^b[X_{t+1}^C]$ conditional on their measurement and other observables:

$$\hat{D}_t = \mathbb{E}[\mathbb{D}_t^b[X_{t+1}^C] | \Omega_t]$$

where $\Omega_t \equiv \{D_t, G_t, G_t^*, X_t, X_t^*, Y_t, \mathbb{E}_t[X_{t+1}], \mathbb{E}_t[X_{t+1}^*], \Omega_{t-1}\}$ denotes the policymaker's information set at time t.

Theorem 3 states that the optimal policy under measurement error is to use this nowcast $\hat{D}_t$ in place of the true belief distortion from Theorem 2's policy rule.

Theorem 3 *When there is measurement error on the belief distortion, the constrained-optimal policy rule is*

$$G_t^\dagger = B_G^+ P_W \left(B_{C1} \hat{D}_t + B_{X1} \mathbb{E}_t[X_{t+1} - X_{t+1}^*] \right) + G_t^*$$

where $\hat{D}_t$ is the belief distortion nowcast.

[11] The assumption that υ_t is independent of the other exogenous process Y_t implies that if the error is due to estimating the rational expectation, the source is due to the inclusion of erroneous independent predictors. Errors due to misspecification of endogenous regressors are instead encoded in the matrix ξ.

Proof: Appendix A.2

The proof follows the logic from Theorem 2 closely, except it replaces the true belief distortion with the sum of the nowcast and the nowcast error. Orthogonality of the error to the remaining time t observables implies that only the nowcast remains in the least-squares welfare objective.

6.2 Endogenous Expectation Operators

In previous sections, the behavioral expectation operator $\mathbb{E}^b[\cdot]$ was assumed to be given as a primitive of the model. Adams (2023a) demonstrates that this assumption characterizes many forms of behavioral expectations and simple information frictions. If instead the expectation operator were endogenously determined, then changing policy can change how belief distortions behave.[12] How does this possibility affect the optimal policy rule?

The main conclusions of the paper hold: the belief distortion is still a sufficient statistic for optimal policy. But an optimal rule may not be unique or even exist.

6.2.1 Endogenous Expectation Operators: Discussion

In this section, the behavioral expectation $\mathbb{E}_t^b[X_{t+1}; \mathcal{G}]$ and belief distortion $\mathbb{D}_t^b[X_{t+1}; \mathcal{G}]$ are now functions of the policy rule $\mathcal{G}$. Otherwise, I return to the environment from Section 3: Sentiment Spanning holds, and the optimal policy under FIRE is $G_t^* = 0$.

What goes wrong if the expectation operator is endogenous? Not Lemma 1: if there exists a policy G_t that recovers the FIRE equilibrium, it *must* satisfy

$$B_G G_t = B_{X1} \mathbb{D}_t^b[X_{t+1}; \mathcal{G}]$$

When Sentiment Spanning holds, the belief distortions are spanned by the policy vector. The policy rule stated in Theorem 1 (equation 3) still gives the optimal policy:

$$G_t^\dagger = (B_G' B_G)^{-1} B_G' B_{C1} \mathbb{D}_t^b[X_{t+1}^C; \mathcal{G}] \tag{21}$$

This is just a result of matrix multiplication, so it follows from Lemma 1 even when the expectation operator is endogenous.

[12] This would be the case in models such as Lucas (1972). Adams (2023b) studies an optimal taxation problem in such a setting.

But Theorem 1 itself does not necessarily hold, because it states that a unique $G_t^\dagger$ must exist. And this is not true if $\mathbb{E}^b[\cdot;\mathcal{G}]$ is affected by policy; there is no guarantee that there exists a unique fixed point such that equation (21) holds with a rule $\mathcal{G}$ that generates $G_t^\dagger$. Instead, with an endogenous expectation operator, $G_t^\dagger$ is the optimal policy *only if* there exists a policy that recovers FIRE, and it may not be unique.

6.2.2 Endogenous Expectation Operators: Noisy Signals Example

To demonstrate, I proceed by laying out a simple example with incomplete information and learning from endogenous signals. When signal processes are endogenous, it introduces a non-linearity, which can lead to multiplicity or non-existence in models that would otherwise be linear and well-behaved were the signals entirely exogenous. In this example, agents learn from policymakers' decisions, which changes how they form expectations. Equation (3) gives the optimal policy except in extreme cases where no optimal policy exists.

Agents attempt to nowcast a common fundamental $\varphi_t \sim N(0,\sigma_\varphi^2)$. There is a unit measure of agents, indexed by $i \in \mathcal{I}$. Agent i makes an action $x_{i,t}$ based on their nowcast, and perturbed by the policy g_t:

$$x_{i,t} = g_t + \mathbb{E}_{i,t}[\varphi_t]$$

Agents do not observe the aggregates φ_t or g_t when forming their nowcast (even though g_t affects their action $x_{i,t}$ directly). Instead, agent i observes the noisy signals

$$s_{i,t} = \varphi_t + \epsilon_{i,t} \qquad\qquad z_{i,t} = g_t + \nu_{i,t}$$

with i.i.d. noise $\epsilon_{i,t} \sim N(0,\sigma_\epsilon^2)$ and $\nu_{i,t} \sim N(0,\sigma_\nu^2)$ satisfying $\sigma_\epsilon^2 > 0$ and $\sigma_\nu^2 \geq 0$.

The policymaker observes φ_t: in this example, the FIRE nowcast of φ_t is the actual realization. The policymaker cares only about the average agent; their objective is for the average action $x_t \equiv \int_{i\in\mathcal{I}} x_{i,t}di$ to select the FIRE nowcast:

$$x_t^* = \varphi_t$$

If agents did not observe the policy signal $z_{i,t}$, then their forecasting would have a representation as a behavioral expectation operator (Adams, 2023a). But because $z_{i,t}$ is endogenous, the policy choice affects how agents forecast, and the assumptions for Theorem 1 do not hold. Proposition 2 describes exactly how the theorem can fail: the optimal policy rule may not exist, but if it does, it is given by equation (3).

Proposition 2 *The noisy signals model has a unique optimal policy $g_t^\dagger$ – for which the belief distortion is a sufficient statistic – if agents observe the policy with noise ($\sigma_\varphi^2 > 0$):*

$$g_t^\dagger = -(\overline{\mathbb{E}}_t[\varphi_t] - \varphi_t)$$

but no optimal policy exists if agents observe the policy exactly ($\sigma_\varphi^2 = 0$).

Proof: Appendix A.3

If agents observe the policy g_t with noise, then they have two noisy signals and cannot perfectly nowcast the fundamental φ_t. The optimal policy then offsets this error, following the rule implied by Theorem 1.

But Theorem 1 itself does not hold: the optimal policy may not exist. This happens when agents observe the policy g_t exactly. If the policy rule is non-zero, agents can infer the fundamental φ_t, so the policymaker would prefer not to intervene. However, if they choose a policy rule that does not respond to φ_t, then agents mis-forecast, and the policymaker would prefer to intervene. In this example the policymaker can choose a rule that gets arbitrarily close to the optimal action $x_t^* = \varphi$, but cannot achieve it exactly.

7 Conclusion

This paper has demonstrated that policymakers can resolve the distortions due to non-rational expectations with a simple policy rule, for which the belief distortion is a sufficient statistic. This implies that policymakers should devote resources to *measuring belief distortions*.

Progress is already being made on this front. New surveys of expectations have proliferated in recent years, and research on the empirical properties of behavioral expectations is accelerating. With quality measurements of agents' forecasts in hand, the belief distortion only requires estimating the rational expectation. This can be done in a variety of ways. A simple method is to incorporate forecasts in a VAR (Adams and Barrett, 2024), but more robust methods can apply machine learning and microdata (Bianchi, Ludvigson, and Ma, 2022). Alternatively, a structurally-inclined economist might take the rational expectation from a fully specified model.

However, these belief distortion measures are not yet being appropriately applied to policy. For example, Adams and Barrett (2024) show that the Federal Reserve lowers interest rates when inflation belief distortions rise, in contrast to the policy prescribed by Section 5.2. The findings from this paper will allow policymakers to tractably respond to the belief distortions that they now observe.

References

ADAM, K. (2007): "Optimal monetary policy with imperfect common knowledge," *Journal of Monetary Economics*, 54(2), 267–301.

ADAM, K., AND M. WOODFORD (2012): "Robustly optimal monetary policy in a micro-founded New Keynesian model," *Journal of Monetary Economics*, 59(5), 468–487.

ADAMS, J. J. (2023a): "Equilibrium Determinacy With Behavioral Expectations," *University of Florida mimeo.*

ADAMS, J. J. (2023b): "Moderating noise-driven macroeconomic fluctuations under dispersed information," *Journal of Economic Dynamics and Control*, 156, 104752.

——— (2024): "Behavioral Expectations Equilibrium Toolkit," Working Papers 001012, University of Florida, Department of Economics.

ADAMS, J. J., AND P. BARRETT (2024): "Shocks to inflation expectations," *Review of Economic Dynamics*, 54, 101234.

ANGELETOS, G.-M., Z. HUO, AND K. A. SASTRY (2021): "Imperfect Macroeconomic Expectations: Evidence and Theory," *NBER Macroeconomics Annual*, 35, 1–86, Publisher: The University of Chicago Press.

ANGELETOS, G.-M., L. IOVINO, AND J. LA'O (2020): "Learning over the business cycle: Policy implications," *Journal of Economic Theory*, 190, 105115.

ANGELETOS, G.-M., AND J. LA'O (2019): "Optimal Monetary Policy with Informational Frictions," *Journal of Political Economy*, p. 704758.

ANGELETOS, G.-M., AND C. LIAN (2023): "Determinacy without the Taylor Principle," *Journal of Political Economy*, 131(8), 2125–2164, Publisher: The University of Chicago Press.

BAERISWYL, R., AND C. CORNAND (2010): "The signaling role of policy actions," *Journal of Monetary Economics*, 57(6), 682–695.

BENHIMA, K., AND I. BLENGINI (2020): "Optimal Monetary Policy when Information is Market-Generated," *The Economic Journal*, 130(628), 956–975.

BIANCHI, F., S. C. LUDVIGSON, AND S. MA (2022): "Belief Distortions and Macroeconomic Fluctuations," *American Economic Review*, 112(7), 2269–2315.

Bordalo, P., N. Gennaioli, and A. Shleifer (2018): "Diagnostic Expectations and Credit Cycles," *The Journal of Finance*, 73(1), 199–227, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/jofi.12586.

Cagan, P. (1956): "The Monetary Dynamics of Hyperinflation," in *Studies in the Quantity Theory of Money*, ed. by M. Freidman. University of Chicago Press.

Chari, V. V., P. J. Kehoe, and E. R. McGrattan (2007): "Business Cycle Accounting," *Econometrica*, 75(3), 781–836.

Cochrane, J. H. (2023): *The fiscal theory of the price level.* Princeton University Press.

Di Bartolomeo, G., M. Di Pietro, and B. Giannini (2016): "Optimal monetary policy in a New Keynesian model with heterogeneous expectations," *Journal of Economic Dynamics and Control*, 73, 373–387.

Evans, G. W., and S. Honkapohja (2003): "Expectations and the Stability Problem for Optimal Monetary Policies," *The Review of Economic Studies*, 70(4), 807–824.

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

Gnocchi, S. (2013): "Monetary Commitment and Fiscal Discretion: The Optimal Policy Mix," *American Economic Journal: Macroeconomics*, 5(2), 187–216.

Hommes, C., D. Massaro, and M. Weber (2019): "Monetary policy under behavioral expectations: Theory and experiment," *European Economic Review*, 118, 193–212.

Iovino, L., and D. Sergeyev (2023): "Central Bank Balance Sheet Policies Without Rational Expectations," *The Review of Economic Studies*, 90(6), 3119–3152.

King, R. G. (2000): "The new IS-LM model: language, logic, and limits," *FRB Richmond Economic Quarterly*, 86(3), 45–104.

Kydland, F. E., and E. C. Prescott (1982): "Time to Build and Aggregate Fluctuations," *Econometrica*, 50(6), 1345–1370.

Lorenzoni, G. (2010): "Optimal Monetary Policy with Uncertain Fundamentals and Dispersed Information," *The Review of Economic Studies*, 77(1), 305–338.

Lucas, R. E. (1972): "Expectations and the Neutrality of Money," *Journal of economic theory*, 4(2), 103–124.

L'Huillier, J.-P., S. R. Singh, and D. Yoo (2023): "Incorporating Diagnostic Expectations into the New Keynesian Framework," *The Review of Economic Studies*, p. rdad101.

Mankiw, N. G., and R. Reis (2002): "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *The Quarterly Journal of Economics*, 117(4), 1295–1328.

Nimark, K. (2008): "Monetary policy with signal extraction from the bond market," *Journal of Monetary Economics*, 55(8), 1389–1400.

Orphanides, A., and J. C. Williams (2008): "Learning, expectations formation, and the pitfalls of optimal control monetary policy," *Journal of Monetary Economics*, 55, S80–S96.

Paciello, L., and M. Wiederholt (2014): "Exogenous Information, Endogenous Information, and Optimal Monetary Policy," *The Review of Economic Studies*, 81(1), 356–388.

Rotemberg, J. J., and M. Woodford (1997): "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy," *NBER Macroeconomics Annual*, 12, 297–346, Publisher: The University of Chicago Press.

Uhlig, H. (2001): "A Toolkit for Analysing Nonlinear Dynamic Stochastic Models Easily," in *Computational Methods for the Study of Dynamic Economies*, ed. by R. Marimon, and A. Scott, p. 0. Oxford University Press.

Woodford, M. (2010a): "Optimal Monetary Stabilization Policy," in *Handbook of Monetary Economics*, ed. by B. M. Friedman, and M. Woodford, vol. 3, pp. 723–828. Elsevier.

——— (2010b): "Robustly Optimal Monetary Policy with Near-Rational Expectations," *American Economic Review*, 100(1), 274–303.

A Additional Proofs

A.1 Proof of Proposition 1

Proof. Theorem 2 gives the formula $G_t^\dagger = B_G^+ P_W \left(B_{C1}\mathbb{D}_t^b[X_{t+1}^C] + B_{X1}\mathbb{E}_t[X_{t+1} - X_{t+1}^*]\right) + G_t^*$. The coefficient matrices on the belief distortions are derived as follows. The policy instrument is now only $G_t = i_t$, and its coefficient matrix is

$$B_G = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$$

The modified weighting matrix $\tilde{W} = (B_{X0}^{-1})' W B_{X0}^{-1}$ is

$$\tilde{W} = \begin{pmatrix} 1 & 0 \\ -\frac{\kappa}{\sigma} & \frac{1}{\sigma} \end{pmatrix} \begin{pmatrix} b_\pi & 0 \\ 0 & b_y \end{pmatrix} \begin{pmatrix} 1 & -\frac{\kappa}{\sigma} \\ 0 & \frac{1}{\sigma} \end{pmatrix} = \begin{pmatrix} b_\pi & -b_\pi \frac{\kappa}{\sigma} \\ -b_\pi \frac{\kappa}{\sigma} & b_\pi \frac{\kappa^2}{\sigma^2} + b_y \frac{1}{\sigma^2} \end{pmatrix}$$

The projection matrix P_W is

$$P_W = B_G \left(B_G' \tilde{W} B_G \right)^{-1} B_G' \tilde{W}$$

$$= \begin{pmatrix} 0 \\ 1 \end{pmatrix} \left(\begin{pmatrix} 0 & 1 \end{pmatrix} \tilde{W} \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right)^{-1} \begin{pmatrix} 0 & 1 \end{pmatrix} \tilde{W}$$

$$= \begin{pmatrix} 0 \\ 1 \end{pmatrix} \left(b_\pi \frac{\kappa^2}{\sigma^2} + b_y \frac{1}{\sigma^2} \right)^{-1} \begin{pmatrix} -b_\pi \frac{\kappa}{\sigma} & b_\pi \frac{\kappa^2}{\sigma^2} + b_y \frac{1}{\sigma^2} \end{pmatrix}$$

$$= \begin{pmatrix} 0 & 0 \\ \frac{-b_\pi \kappa \sigma}{b_\pi \kappa^2 + b_y} & 1 \end{pmatrix}$$

Then Theorem 2 implies the policy satisfies

$$G_t = (B_G' B_G)^{-1} B_G' P_W B_{X1} \left(\mathbb{D}_t^b[X_{t+1}] + \mathbb{E}_t[X_{t+1} - X_{t+1}^*] \right) + G_t^*$$

Substituting in $B_{X1} = \begin{pmatrix} \beta & 0 \\ 1 & \sigma \end{pmatrix}$ and $(B_G' B_G)^{-1} B_G' = \begin{pmatrix} 0 & 1 \end{pmatrix}$ implies

$$i_t^\dagger = \begin{pmatrix} \frac{-b_\pi \kappa \sigma}{b_\pi \kappa^2 + b_y} & 1 \end{pmatrix} \begin{pmatrix} \beta & 0 \\ 1 & \sigma \end{pmatrix} \left(\mathbb{D}_t^b[X_{t+1}] + \mathbb{E}_t[X_{t+1} - X_{t+1}^*] \right) + i_t^*$$

$$= \left(1 - \beta \frac{b_\pi \kappa \sigma}{b_\pi \kappa^2 + b_y}\right) \left(\mathbb{D}_t^b[\pi_{t+1}] + \mathbb{E}_t[\pi_{t+1} - \pi_{t+1}^*]\right) + \sigma \left(\mathbb{D}_t^b[y_{t+1}] + \mathbb{E}_t[y_{t+1} - y_{t+1}^*]\right) + i_t^*$$

■

A.2 Proof of Theorem 3

Proof. Substitute $B'_{X0} C'_{\tilde{W}} C_{\tilde{W}} B_{X0} = W$ into the objective function:

$$\mathbb{E}\left[(X_t - X_t^*)' W (X_t - X_t^*)\right] = \mathbb{E}\left[(X_t - X_t^*)' B'_{X0} C'_{\tilde{W}} C_{\tilde{W}} B_{X0} (X_t - X_t^*)\right]$$

This can be written in terms of the policy using equation (1), which implies $B_{X0} X_t - B_{X0} X_t^* = -B_G \hat{G}_t + B_{X1} E_t^b[X_{t+1}] - B_{X1} E_t[X_{t+1}^*]$ where $\hat{G}_t \equiv G_t - G_t^*$:

$$= \mathbb{E}\left[(-B_G \hat{G}_t + B_{X1} \mathbb{E}_t^b[X_{t+1}] - B_{X1} \mathbb{E}_t[X_{t+1}^*])' C'_{\tilde{W}} C_{\tilde{W}} (-B_G \hat{G}_t + B_{X1} \mathbb{E}_t^b[X_{t+1}] - B_{X1} \mathbb{E}_t[X_{t+1}^*])\right]$$

Rewriting with the belief distortion gives equation (5):

$$= \mathbb{E}\left[(-C_{\tilde{W}} B_G \hat{G}_t + C_{\tilde{W}} B_{C1} \mathbb{D}_t^b[X_{t+1}^C] + C_{\tilde{W}} B_{X1} \mathbb{E}_t[X_{t+1} - X_{t+1}^*]))' (-C_{\tilde{W}} B_G \hat{G}_t + C_{\tilde{W}} B_{C1} \mathbb{D}_t^b[X_{t+1}^C] + C_{\tilde{W}} B_{X1} \mathbb{E}_t[X_{t+1} - X_{t+1}^*])\right]$$

Next, break the belief distortion into the nowcast $\hat{D}_t$ and error $\mathbb{D}_t^b[X_{t+1}^C] - \hat{D}_t$:

$$= \mathbb{E}\left[(-C_{\tilde{W}} B_G \hat{G}_t + C_{\tilde{W}} B_{C1} (\mathbb{D}_t^b[X_{t+1}^C] - \hat{D}_t) + C_{\tilde{W}} B_{C1} \hat{D}_t + C_{\tilde{W}} B_{X1} \mathbb{E}_t[X_{t+1} - X_{t+1}^*]))' (-C_{\tilde{W}} B_G \hat{G}_t + C_{\tilde{W}} B_{C1} (\mathbb{D}_t^b[X_{t+1}^C] - \hat{D}_t) + C_{\tilde{W}} B_{C1} \hat{D}_t + C_{\tilde{W}} B_{X1} \mathbb{E}_t[X_{t+1} - X_{t+1}^*])\right]$$

The error $\mathbb{D}_t^b[X_{t+1}^C] - \hat{D}_t$ is orthogonal to all other time t variables, so the expression reduces to

$$= \mathbb{E}\left[(-C_{\tilde{W}} B_G \hat{G}_t + C_{\tilde{W}} B_{C1} \hat{D}_t + C_{\tilde{W}} B_{X1} \mathbb{E}_t[X_{t+1} - X_{t+1}^*]))' (-C_{\tilde{W}} B_G \hat{G}_t + C_{\tilde{W}} B_{C1} \hat{D}_t + C_{\tilde{W}} B_{X1} \mathbb{E}_t[X_{t+1} - X_{t+1}^*])\right] \quad (22)$$

Thus the objective is to choose $\hat{G}_t$ to minimize equation (22). $\hat{G}_t$ solves this standard least squares minimization problem by projecting $C_{\tilde{W}} B_{C1} \hat{D}_t + C_{\tilde{W}} B_{X1} \mathbb{E}_t[X_{t+1} - X^*_{t+1}]$ onto the space spanned by $C_{\tilde{W}} B_G$:

$$C_{\tilde{W}} B_G \hat{G}_t = C_{\tilde{W}} B_G \left((C_{\tilde{W}} B_G)' C_{\tilde{W}} B_G\right)^{-1} (C_{\tilde{W}} B_G)' \left(C_{\tilde{W}} B_{C1} \hat{D}_t + C_{\tilde{W}} B_{X1} \mathbb{E}_t[X_{t+1} - X^*_{t+1}]\right)$$

Left-multiply by $C_{\tilde{W}}^{-1}$ and substitute with $P_W = B_G \left(B_G' \tilde{W} B_G\right)^{-1} B_G' \tilde{W}$ and B_G^+:

$$\hat{G}_t = B_G^+ P_W \left(B_{C1} \hat{D}_t + B_{X1} \mathbb{E}_t[X_{t+1} - X^*_{t+1}]\right)$$

and substituting in $G_t = \hat{G}_t + G_t^*$ completes the proof. ■

A.3 Proof of Proposition 2

Proof. Conjecture a policy rule $g_t = \alpha \varphi_t$; α must be solved for. Given α, an agent observes two noisy signals of φ_t:

$$s_{i,t} = \varphi_t + \epsilon_{i,t} \qquad \tilde{z}_{i,t} \equiv \frac{z_{i,t}}{\alpha} = \varphi_t + \frac{\nu_{i,t}}{\alpha}$$

There are two cases to consider: $\sigma_\nu^2 > 0$ and $\sigma_\nu^2 = 0$.

Case 1: Agents observe the policy with noise ($\sigma_\nu^2 > 0$). The nowcast of φ_t given these two noisy signals is

$$\begin{aligned}\mathbb{E}[\varphi_t | s_{i,t}, \tilde{z}_{i,t}] = & \\ & \frac{v(s_{i,t}) cov(\tilde{z}_{i,t}, \varphi_t) - cov(\tilde{z}_{i,t}, s_{i,t}) cov(s_{i,t}, \varphi_t)}{v(\tilde{z}_{i,t}) v(s_{i,t}) - cov(\tilde{z}_{i,t}, s_{i,t})^2} \tilde{z}_{i,t} \\ & \qquad + \frac{v(\tilde{z}_{i,t}) cov(s_{i,t}, \varphi_t) - cov(\tilde{z}_{i,t}, s_{i,t}) cov(\tilde{z}_{i,t}, \varphi_t)}{v(\tilde{z}_{i,t}) v(s_{i,t}) - cov(\tilde{z}_{i,t}, s_{i,t})^2} s_{i,t}\end{aligned}$$

Substituting for the variances gives

$$\begin{aligned}= & \frac{\left(\sigma_\varphi^2 + \sigma_\epsilon^2\right) cov(\tilde{z}_{i,t}, \varphi_t) - cov(\tilde{z}_{i,t}, s_{i,t}) cov(s_{i,t}, \varphi_t)}{\left(\sigma_\varphi^2 + \sigma_\nu^2/\alpha^2\right)\left(\sigma_\varphi^2 + \sigma_\epsilon^2\right) - cov(\tilde{z}_{i,t}, s_{i,t})^2} \tilde{z}_{i,t} \\ & \qquad + \frac{\left(\sigma_\varphi^2 + \sigma_\nu^2/\alpha^2\right) cov(s_{i,t}, \varphi_t) - cov(\tilde{z}_{i,t}, s_{i,t}) cov(\tilde{z}_{i,t}, \varphi_t)}{\left(\sigma_\varphi^2 + \sigma_\nu^2/\alpha^2\right)\left(\sigma_\varphi^2 + \sigma_\epsilon^2\right) - cov(\tilde{z}_{i,t}, s_{i,t})^2} s_{i,t}\end{aligned}$$

and substituting for the covariances gives

$$= \frac{\left(\sigma_\varphi^2 + \sigma_\epsilon^2\right)\sigma_\varphi^2 - \sigma_\varphi^4}{\left(\sigma_\varphi^2 + \sigma_\nu^2/\alpha^2\right)\left(\sigma_\varphi^2 + \sigma_\epsilon^2\right) - \sigma_\varphi^4}\tilde{z}_{i,t} + \frac{\left(\sigma_\varphi^2 + \sigma_\nu^2/\alpha^2\right)\sigma_\varphi^2 - \sigma_\varphi^4}{\left(\sigma_\varphi^2 + \sigma_\nu^2/\alpha^2\right)\left(\sigma_\varphi^2 + \sigma_\epsilon^2\right) - \sigma_\varphi^4}s_{i,t}$$

$$= \frac{\alpha^2\sigma_\epsilon^2\sigma_\varphi^2}{\alpha^2\sigma_\epsilon^2\sigma_\varphi^2 + \sigma_\nu^2\sigma_\epsilon^2 + \sigma_\nu^2\sigma_\varphi^2}\tilde{z}_{i,t} + \frac{\sigma_\nu^2\sigma_\varphi^2}{\alpha^2\sigma_\epsilon^2\sigma_\varphi^2 + \sigma_\nu^2\sigma_\epsilon^2 + \sigma_\nu^2\sigma_\varphi^2}s_{i,t}$$

When $s_{i,t}$ is noisy (σ_ϵ^2 large) the nowcast puts most weight on $\tilde{z}_{i,t}$; when $\tilde{z}_{i,t}$ is noisy (σ_ν^2/α^2 large) the nowcast puts most weight on $s_{i,t}$.

The average nowcast $\overline{\mathbb{E}}_t[\varphi_t] \equiv \int_{i\in\mathcal{I}} \mathbb{E}_{i,t}[\varphi_t]di$ is

$$\overline{\mathbb{E}}_t[\varphi_t] = \frac{\alpha^2\sigma_\epsilon^2\sigma_\varphi^2 + \sigma_\nu^2\sigma_\varphi^2}{\alpha^2\sigma_\epsilon^2\sigma_\varphi^2 + \sigma_\nu^2\sigma_\epsilon^2 + \sigma_\nu^2\sigma_\varphi^2}\varphi_t$$

so the average action is

$$\begin{aligned} x_t &= g_t + \overline{\mathbb{E}}_t[\varphi_t] \qquad (23) \\ &= \alpha\varphi_t + \frac{\alpha^2\sigma_\epsilon^2\sigma_\varphi^2 + \sigma_\nu^2\sigma_\varphi^2}{\alpha^2\sigma_\epsilon^2\sigma_\varphi^2 + \sigma_\nu^2\sigma_\epsilon^2 + \sigma_\nu^2\sigma_\varphi^2}\varphi_t \end{aligned}$$

The policymaker's goal is to choose α such that $x_t = \varphi_t$. This implies that α solves the cubic

$$0 = \alpha^3\sigma_\epsilon^2\sigma_\varphi^2 + \alpha(\sigma_\nu^2\sigma_\epsilon^2 + \sigma_\nu^2\sigma_\varphi^2) - \sigma_\nu^2\sigma_\epsilon^2$$

which has a unique solution. The optimal policy sets $x_t = \varphi_t$, so equation (23) implies it proportional to the belief distortion:

$$g_t = -(\overline{\mathbb{E}}_t[\varphi_t] - \varphi_t)$$

which is this model's special case of the general optimal policy rule in equation (3).

Case 2: Agents observe the policy exactly ($\sigma_\nu^2 = 0$). Then, agents' forecast depends on α: if α is non-zero, observing the policy reveals φ_t. But if $\alpha = 0$, then the policy signal is uninformative and agents nowcast using only the noisy signal $s_{i,t}$: $\mathbb{E}[\varphi_t|s_{i,t}] = \frac{\sigma_\varphi^2}{\sigma_\varphi^2+\sigma_\epsilon^2}s_{i,t}$. Thus the average expectation is

$$\overline{\mathbb{E}}_t[\varphi_t] = \begin{cases} \varphi_t & \alpha \neq 0 \\ \frac{\sigma_\varphi^2}{\sigma_\varphi^2+\sigma_\epsilon^2}\varphi_t & \alpha = 0 \end{cases}$$

and the average action is

$$x_t = \begin{cases} (\alpha+1)\varphi_t & \alpha \neq 0 \\ \frac{\sigma_\varphi^2}{\sigma_\varphi^2+\sigma_\epsilon^2}\varphi_t & \alpha = 0 \end{cases}$$

There is no possible α such that $x_t = \varphi_t$: the optimal policy recovering FIRE does not exist. ■

B Deriving the Behavioral RBC Model

This section derives the behavioral RBC model equations of Section 5.1 from microfoundations.

The representative household's problem is represented by the Bellman equation

$$V(K;A) = \max_{C,N,K'} \frac{C^{1-\sigma}-1}{1-\sigma} - \chi\frac{N^{1+\eta}}{1+\eta} + \beta\mathbb{E}^b\left[V(K';A')|A\right]$$

$$\text{s.t.} \quad RK + WN = C + Q(K' - (1-\delta)K)$$

The household's endogenous state variable is capital K which depreciates by a factor $1-\delta$. The household's budget constraint is real: it earns RK from renting its capital to firms at rate R, and it earns WN form working N hours at wage W; it spends its income to purchase consumption C (the numeraire) and acquire new capital at cost Q. The vector Z includes exogenous state variables, policy, and prices, which atomistic households take as exogenous. $\mathbb{E}^b$ represents the household's behavioral expectation, and primes denote the next period's values.

The household's problem is solved by a labor supply equation

$$\chi N^\eta = WC^{-\sigma} \tag{24}$$

and an Euler equation

$$C^{-\sigma} = \beta\mathbb{E}^b\left[(C')^{-\sigma}(R'+1-\delta)\right] \tag{25}$$

The Euler equation can be derived as usual because $\mathbb{E}^b$ is assumed to be linear: the partial derivative operator passes through it, so that $\frac{\partial}{\partial B'}\mathbb{E}^b\left[V(B';Z')|Z\right] = \mathbb{E}^b\left[\frac{\partial}{\partial B'}V(B';Z')|Z\right]$

Output Y is produced by competitive firms with constant returns to scale production functions. The representative firm produces output Y by

$$Y = AK^{\alpha}N^{1-\alpha} \tag{26}$$

with TFP A exogenously given and the capital share $\alpha \in (0,1)$. The firm hires capital and labor from the household, implying the demand functions:

$$W = (1-\alpha)\frac{Y}{N} \tag{27}$$

$$R = \alpha\frac{Y}{K} \tag{28}$$

The resource constraint for the economy is

$$Y = C + Q(K' - (1-\delta)K) \tag{29}$$

Log-linearizing the equilibrium conditions (24)-(29) is mostly standard, with one exception: the capital price Q is set by the policymaker. Normalize around a steady state of $\overline{Q} = 1$ and denote the log deviation by τ to measure the tax-induced markup over the steady state value. This gives the log-linearized equations from Section 5.1; when written in the matrix form of equation (1), they are:

$$\underbrace{\begin{pmatrix} 0 & -\sigma & 0 & 0 & 0 & 0 \\ 0 & -\sigma & -\eta & 0 & 1 & 0 \\ -\alpha & 0 & -(1-\alpha) & 1 & 0 & 0 \\ 1 & 0 & 0 & -1 & 0 & 1 \\ 0 & 0 & 1 & -1 & 1 & 0 \\ \overline{K}(1-\delta) & -\overline{C} & 0 & \overline{Y} & 0 & 0 \end{pmatrix}}_{B_{X0}} \underbrace{\begin{pmatrix} k_{t-1} \\ c_t \\ n_t \\ y_t \\ w_t \\ r_t \end{pmatrix}}_{X_t} + \underbrace{\begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}}_{B_G} \underbrace{\tau_t}_{G_t} + \underbrace{\begin{pmatrix} 0 \\ 0 \\ -1 \\ 0 \\ 0 \\ 0 \end{pmatrix}}_{B_Y} \underbrace{a_t}_{Y_t}$$

$$= \underbrace{\begin{pmatrix} 0 & -\sigma & 0 & 0 & 0 & \overline{R} \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ -\overline{K} & 0 & 0 & 0 & 0 & 0 \end{pmatrix}}_{B_{X1}} \mathbb{E}_t^b \underbrace{\left[\begin{pmatrix} k_t \\ c_{t+1} \\ n_{t+1} \\ y_{t+1} \\ w_{t+1} \\ r_{t+1} \end{pmatrix}\right]}_{X_t} \tag{30}$$

The relevant steady state values are calculated as follows. Equation (25) in the steady state is

$$1 = \beta\left(\overline{R} + 1 - \delta\right)$$
$$\implies \overline{R} = \frac{1}{\beta} - 1 + \delta$$

I choose χ such that $\overline{N} = 1$, and normalized $\overline{A} = 1$. Then equation (26) implies that steady state output and capital are related by

$$\overline{Y} = \overline{K}^{\alpha}$$

Plugging this into the steady state capital demand (28) gives capital in terms of $\overline{R}$:

$$\overline{K} = \left(\frac{\alpha}{\overline{R}}\right)^{\frac{1}{1-\alpha}}$$

Finally, the steady state resource constraint (29) gives steady state consumption from the preceding values by

$$\overline{C} = \overline{Y} - \delta\overline{K}$$

C Mapping to the General Model

The general model from Section 2 is

$$B_{X1}\mathbb{E}_t^b\left[X_{t+1}\right] = B_{X0}X_t + B_Y Y_t + B_G G_t$$

where the vector $X_t = \begin{pmatrix} X_{t-1}^K \\ X_t^C \end{pmatrix}$ contains pre-determined states in X_{t-1}^K and contemporaneous controls in X_t^C. The behavioral expectation $\mathbb{E}_t^b\left[X_{t+1}\right]$ only applies to the future controls; time t-dated variables are forecasted exactly.

This model nests alternative structures, but some care must be taken to map one to another. Consider instead a behavioral modification of the general model studied in Uhlig (2001):

$$\tilde{\mathbb{E}}_t[Fc_{t+1} + G_1^c c_t + G_1^k k_t + My_{t+1} + N_1 y_t] + G_2^c c_t + G_2^k k_t + Hk_{t-1} + N_2 y_t = TG_t \quad (31)$$

c_t and k_t denote endogenous variables determined at time t, y_t denotes exogenous variables, and G_t denotes the policy instruments. c_t represent controls whose forecasts enter the model, while k_t represent states whose lags enter the model. The behavioral expectations operator $\tilde{\mathbb{E}}_t[\cdot]$ can be applied to some current period variables and exogenous variables in addition to x_{t+1}. This operator has the unique notation $\tilde{\mathbb{E}}$ to clarify that it does not share all the properties of the other behavioral expectations considered in this paper. I assume that when $\tilde{\mathbb{E}}_t[\cdot]$ is applied to a $t+1$ variable, it is equivalent to the usual behavioral expectation operator $\mathbb{E}_t^b[\cdot]$.

In what follows, I show that the alternative model (equation 31) can be rewritten in the general form (equation 1).

Define the new variable z_t:

$$z_t = G_1^c c_t + G_1^k k_{t-1} \tag{32}$$

and Y_t:

$$Y_t = -N_2 y_t - \tilde{\mathbb{E}}_t[M y_{t+1} + N_1 y_t]$$

which allows equation (31) to be rewritten as

$$\tilde{\mathbb{E}}_t[F c_{t+1} + z_{t+1}] + G_2^k k_t = -G_2^c c_t - H k_{t-1} + Y_t + T G_t \tag{33}$$

Combine equations (32) and (33) with matrices:

$$\begin{pmatrix} G_2^k & F & I \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} k_t \\ \mathbb{E}_t^b[c_{t+1}] \\ \mathbb{E}_t^b[z_{t+1}] \end{pmatrix} = \begin{pmatrix} -H & -G_2^c & 0 \\ G_1^k & G_1^c & -I \end{pmatrix} \begin{pmatrix} k_{t-1} \\ c_t \\ z_t \end{pmatrix} + \begin{pmatrix} I \\ 0 \end{pmatrix} Y_t + \begin{pmatrix} T \\ 0 \end{pmatrix} G_t$$

which matches the general form, where $X_t^K = k_t$, $X_t^C = \begin{pmatrix} c_t \\ z_t \end{pmatrix}$, and

$$\begin{pmatrix} G_2^k & F & I \\ 0 & 0 & 0 \end{pmatrix} = B_{X1} \qquad \begin{pmatrix} -H & -G_2^c & 0 \\ G_1^k & G_1^c & -I \end{pmatrix} = B_{X0} \qquad \begin{pmatrix} I \\ 0 \end{pmatrix} = B_Y \qquad \begin{pmatrix} T \\ 0 \end{pmatrix} = B_G$$