Unemployment Insurance in Macroeconomic Stabilization with Imperfect Expectations*

Bence Bardóczy[†] Joao Guerreiro[‡]

August 2026

Abstract

Automatic stabilizers can respond to a recession without a new policy decision, but their effects on aggregate demand may still only arrive with delay. We study the transmission lag for unemployment insurance extensions in a heterogeneous-agent New Keynesian model that closely matches the incidence of unemployment risk and the response of household spending to job loss and UI extensions. More generous UI stimulates demand largely by relaxing precautionary saving, so households must anticipate future unemployment risk and benefit duration. To quantify this channel, we work directly with survey expectations, using a measured history of forecast errors and revisions, rather than committing to a specific model of belief formation. Our estimated model implies a significant *Expectations-Driven Efficacy Lag*. The direct effect of UI extensions on consumption peaks only after UI duration has begun to recede. In general equilibrium, the consumption multiplier is about 0.6 on impact, less than half its value under full-information rational expectations, but the policy becomes more effective after the first year of the recession.

Keywords: information frictions, bounded rationality, expectations surveys, unemployment insurance, heterogeneous agent New Keynesian models, stabilization policy.

*For helpful comments, we thank Hassan Afrouzi, George-Marios Angeletos, Adrien Auclert, Borağan Aruoba, Marty Eichenbaum, Chen Lian, Kurt Mitman, Camilo Morales-Jiménez, Matthew Rognlie, Luminita Stevens, and Christian Wolf, as well as seminar participants at Barcelona Summer Forum 2023, North American Summer Meeting 2023, Oslo Macro Conference 2023, SITE 2023: Frontiers of Macroeconomic Research, Católica Lisbon, Dartmouth, Federal Reserve Board, Johns Hopkins, Maryland, Toronto. We are very grateful to Riccardo Bianchi-Vimercati who contributed to an early version of this project. Harrison Snell provided excellent research assistance.

[†]BI Norwegian Business School. Email: bence.bardoczy@bi.no.

[‡]UCLA and NBER. Email: jguerreiro@econ.ucla.edu.

1 Introduction

The classic case for automatic stabilizers is fundamentally about lags. Friedman (1948) distinguished three lags in stabilization policy: the time to recognize the need for action, the time between recognition and action, and the time between action and its effects. His proposed fiscal framework used tax and transfer schedules, explicitly including unemployment insurance, which would vary automatically over the cycle. A well-designed rule can compress the first two lags because it does not require policymakers to diagnose each downturn and legislate a new response. It cannot eliminate the third.

We revisit the third lag in the context of unemployment insurance (UI).[1] Temporary extensions of UI duration are among the most commonly used fiscal-policy instruments in U.S. recessions. During the Great Recession, maximum benefit duration rose from 26 to 99 weeks; in 2020, benefits were again extended. A modern quantitative literature has also identified potentially powerful stabilization channels for social insurance. McKay and Reis (2016) show that tax-and-transfer programs can affect aggregate fluctuations through redistribution and social insurance, rather than only by smoothing disposable income. Kekre (2021) shows that more generous UI can raise aggregate demand through redistribution and precautionary saving, and estimates a contemporaneous output multiplier around one for the Great Recession extensions.

This paper studies a neglected limit on this stabilization power which we call the *Expectations-Driven Efficacy Lag*. We first establish quantitatively that the aggregate-demand effect of UI extensions operates mainly through precautionary saving, rather than only through transfers to households that are currently unemployed. This mechanism makes beliefs central. Most households remain employed even in a severe recession, so an extension raises their spending only insofar as it reduces the income risk they perceive from a possible future job loss. A policy can therefore be well timed in a statutory sense and still transmit slowly due to household expectations. Our question is whether measured expectations create such a gap between the timing of the policy and the timing of its effects.

Figure 1 previews the empirical fact that is central for this mechanism: unemployment expectations initially underreact to a downturn and remain too pessimistic after unemployment has turned around, a pattern that also appears conditional on an identified business-cycle shock. We quantify the resulting expectations-driven efficacy lag in a heterogeneous-agent New Keynesian (HANK) model with search and matching frictions, incomplete markets, and endogenous UI duration. To this end, we also discuss a general method to solve dynamic models under arbitrary beliefs. The method represents decisions using the initial forecast and the subsequent history of forecast revisions, objects that can be measured directly in expectations surveys. We therefore use Survey of Professional Forecasters (SPF) forecasts wherever they are available, without first selecting a parametric model from the "wilderness of bounded rationality."

[1] The first two lags may not even be quantitatively large in current UI policy. Chodorow-Reich et al. (2022) compare the ad hoc extensions enacted in the 2001 and 2007 to 2009 recessions with common automatic-trigger proposals. They find that the historical extensions compare favorably ex post.

Figure 1: Consensus Forecast of the Unemployment Rate

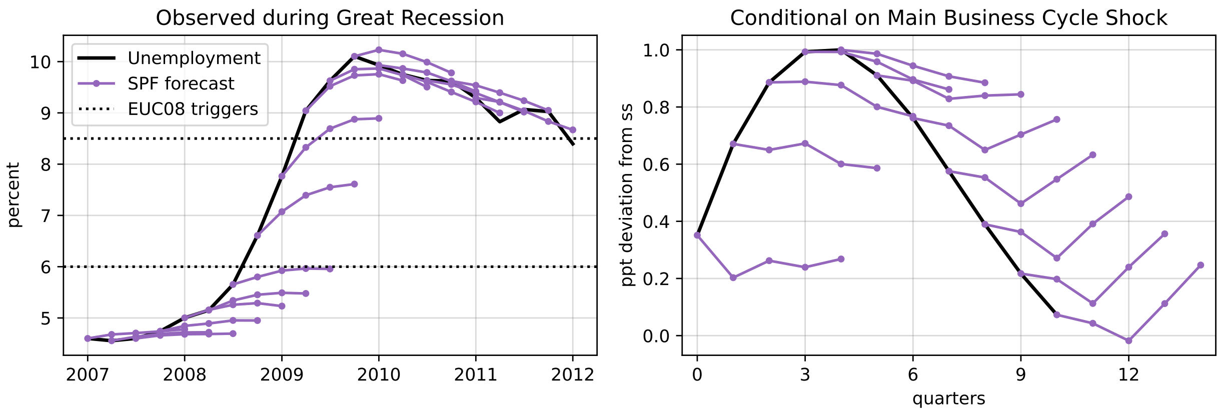

Notes: On both panels, the black line is based on the seasonally adjusted civilian unemployment rate from the U.S. Bureau of Labor Statistics, and the purple line is based on the median unemployment forecast from the Survey of Professional Forecasters. On the left panel, we plot the level of the unemployment rate and forecasts between 2007Q1 and 2012Q1. On the right panel, we plot the impulse responses of these variables to the main business-cycle shock targeting unemployment from Angeletos et al. (2020). We estimate the impulse responses with a classical structural VAR, identifying the main business-cycle shock following Angeletos et al. (2020) (see Section 5.2 for details). We scale the shock so that the peak response of unemployment is 1 percentage point.

Our main finding is that imperfect expectations change the timing of stabilization, not simply its magnitude. In partial equilibrium, the consumption response under full-information rational expectations (FIRE) is largest on impact and then declines. Under estimated beliefs, the impact response is less than half as large, builds for several quarters, and peaks only after UI duration, which tracks unemployment in this exercise, has begun to recede. In general equilibrium, the cumulative consumption multiplier of the extension is about 0.6 on impact under estimated beliefs, compared with about 1.2 under FIRE. Thus, the cumulative spending effect remains positive over the horizon we report, but substantially less of it is delivered at the onset of the recession. This result provides a data-driven microfoundation for part of Friedman's third lag: the statutory rule has moved, while the private spending response has not yet fully arrived.

Environment. Households face idiosyncratic productivity risk, stochastic employment transitions, and borrowing constraints. Unemployed workers receive UI benefits whose expected duration is indexed to the aggregate unemployment rate. This state-contingent rule isolates the automatic-stabilizer logic and is calibrated to the change in benefit duration under the Emergency Unemployment Compensation program of 2008. The model also features capital accumulation with investment adjustment costs, sticky prices and wages, long-term government debt, and a financial intermediary that intermediates between household savings and firm equity.

The model is designed to provide a quantitative account of incomplete markets, unemployment dynamics, and countercyclical income risk. This structure is essential for our application: UI works not only through transfers to unemployed households, but also through precautionary saving. We discipline the household block using average marginal propensities to consume (MPCs), differences in MPCs across employed and unemployed households, and income dynamics around

UI receipt and benefit expiration from Ganong and Noel (2019). The model then accounts well for untargeted evidence on consumption through an unemployment spell, the incidence of UI receipt, liquid net worth, and consumption shares across the wealth distribution.

Dynamics with arbitrary expectations and the power of UI. Building on the Sequence-Space Jacobian (SSJ) framework of Auclert et al. (2021), and its extension to departures from FIRE in Auclert et al. (2020), we show that linearized impulse responses under non-rational expectations can be constructed from two additional objects: the initial forecast and subsequent forecast revisions. The associated responses can be computed directly from FIRE Jacobians, preserving tractability and speed in rich heterogeneous-agent environments. Because the initial forecasts and revisions can be measured in the SPF, the framework makes the timing of Friedman's transmission lag an empirical object rather than an assumed delay.

We estimate the remaining structural parameters by matching the empirical impulse responses of unemployment, consumption, investment, inflation, and the nominal interest rate to an identified aggregate shock (as in Christiano et al., 2005 and Auclert et al., 2020).

Using the estimated model, we first show that beliefs also account for the timing of the recession itself. The initial decline in consumption is smaller under estimated beliefs than under FIRE because households underforecast the deterioration in unemployment and income. As beliefs become too pessimistic, consumption falls more strongly. The trough-to-impact ratio of the consumption response is therefore nearly three times as large under estimated beliefs as under FIRE. This exercise validates the mechanism that subsequently governs the policy experiment: measured expectations turn a front-loaded consumption response into a delayed and hump-shaped one.

Policy counterfactuals and the efficacy lag. For counterfactual policy experiments, we estimate a structural model that combines noisy information with long-memory diagnostic expectations, building on Bordalo et al. (2020) and Bianchi et al. (2021). This model is able to reproduce the initial underreaction and delayed overreaction observed in the SPF. We compare the economy with endogenous UI duration to a counterfactual in which UI duration remains constant. Under FIRE, households immediately anticipate the extension, and the general-equilibrium cumulative consumption multiplier is about 1.2 on impact. Under estimated beliefs, the impact multiplier is about 0.6. The direct effect of UI on spending is also temporal: the FIRE consumption response is front-loaded, whereas the response under estimated beliefs is initially muted, overtakes the FIRE response after a few quarters, and peaks after UI duration does. Measured by household spending per dollar of UI, the policy remains effective over the horizon we report, but its efficacy relative to FIRE catches up only two quarters after the recession peak. Our results establish that imperfect expectations reallocate stimulus away from the onset of a recession and toward its later stages.

Relationship to the literature. Our first contribution is to connect modern quantitative work on automatic stabilizers to the policy-lag problem in Friedman (1948). McKay and Reis (2016) measure the role of automatic tax-and-transfer programs in U.S. business cycles, while Kekre (2021) studies UI extensions in a model with incomplete markets, search frictions, and nominal rigidities. Christiano et al. (2016) show that nominal rigidities and constraints on monetary policy adjustment can reverse the contractionary effects of UI extensions in Krusell et al. (2010), Nakajima (2012), and Mitman and Rabinovich (2015, 2019). These papers identify the mechanisms that make UI a potentially powerful stabilizer.[2] Our contribution is to show that the most powerful mechanism, the reduction in precautionary saving, is also the mechanism most exposed to an expectations-driven transmission lag.

The idea that unemployment expectations matter for consumption and business cycles goes back to Carroll (1992). Ravn and Sterk (2017) show that endogenous countercyclical income risk via job uncertainty can significantly amplify recessions. Bilbiie et al. (2022) show that cyclical income risk and precautionary saving can substantially amplify business cycles, and Fernandes and Rigato (2022) studies UI when present bias reduces the responsiveness of precautionary saving to extensions. Recent micro evidence directly ties unemployment risk and beliefs to self-insurance: Juelsrud and Wold (2025) show that higher unemployment risk raises safe saving and accounts for much of its recessionary increase, while Hartmann and Leth-Petersen (2024) show that subjective unemployment risk predicts subsequent unemployment and induces both liquid saving and UI enrollment. On spending, Dias da Silva et al. (2025) find that anticipating job loss mutes the consumption drop when displacement occurs, and Fagereng et al. (2024) estimate sizable consumption losses after unemployment that vary with household liquidity and debt. Using pandemic administrative data, Ganong et al. (2024) find that expanded UI had large spending effects but small job-finding effects. These findings support the micro mechanism at the center of our paper; our contribution is to show how the aggregate dynamics of unemployment expectations govern when that mechanism makes UI effective.

Our second contribution is to the literature analyzing macroeconomic shocks and policies without FIRE and using survey data to discipline expectations. Closely related, Bianchi-Vimercati et al. (2024) show that bounded rationality weakens government-spending stimulus while leaving some tax-based stimulus comparatively effective, and that communication is central to transmission. Our model builds on the recent literature with HANK models[3] without FIRE including Farhi and Werning (2019), Farhi et al. (2020), Auclert et al. (2020), Pappa et al. (2023), Dobrew et al. (2023), Pfäuti and Seyrich (2022), Morales-Jiménez and Stevens (2025), and Guerreiro (2022). These papers specify parametric models of bounded rationality. We instead provide a method that

[2] Recent quasi-experimental evidence emphasizes the distinction between the net employment effect of UI extensions and their aggregate-demand channel. Acosta et al. (2023) exploit cross-state variation in optional Extended Benefits trigger rules and find that a 13-week extension raises unemployment when baseline potential benefit duration is short, but has little effect when duration is already long. We abstract from endogenous search effort and isolate the effects on aggregate demand via precautionary savings.

[3] Auclert et al. (2025) survey the large HANK policy literature, highlighting direct and indirect transmission, cyclical income risk, behavioral frictions, and sequence-space methods, the ingredients our paper combines in its study of UI.

can quantify deviations from FIRE using expectations surveys directly. This allows us to remain agnostic about the model of belief formation wherever the relevant forecasts are observed.

Key to our analysis is how beliefs deviate from FIRE. Coibion and Gorodnichenko (2012, 2015) find evidence of underreaction, consistent with rational inattention or information rigidities as in Sims (2003), Woodford (2001), Carroll (2003), Mankiw and Reis (2002), and Gabaix (2020). Bordalo et al. (2020) find overreaction, consistent with diagnostic expectations and overextrapolation as in Bordalo et al. (2018). Complementary experimental evidence from Afrouzi et al. (2023) shows that overreaction is stronger for less persistent processes and at longer forecast horizons, providing another microfoundation for horizon-dependent forecast distortions. Angeletos et al. (2021) reconcile these findings by documenting initial underreaction followed by delayed overreaction. Our framework accommodates this non-monotonic pattern and lets the survey evidence determine when the transmission of UI arrives. We show that this empirical pattern of expectations is responsible for the delayed hump-shaped response of consumption to a recession and for the delayed consumption response to UI extensions.[4]

Outline. The structure of the paper is as follows. In Section 2, we introduce an illustrative model to emphasize the importance of expectations in determining the economy's response to UI extensions. In Section 3, we present the quantitative model. In Section 4, we discuss the model's micro incidence and calibration. In Section 5, we present our framework for dynamic responses under imperfect expectations and describe the estimation. In Section 6, we discuss the estimation results. In Section 7, we quantify the consequences for UI policy. The final section concludes.

2 Illustrative model

We start with an analytical demonstration that imperfect expectations interfere with the power of unemployment insurance (UI) extensions to stabilize business cycles. We consider a simple two-period environment. We analyze how equilibrium output at time 0 depends on households' expectations and the implementation of UI. Appendix A contains detailed derivations and proofs.

2.1 Setup

Consider a two-period model, $t = 0, 1$. The economy is populated by a measure-one continuum of households, a representative firm, and a government. The sequence of events within the two periods is the same. First, the representative firm randomly hires a fraction of households. Second, production takes place and households make a consumption-saving decision. Nominal wages are fully rigid, so there is no wage inflation.

[4] In recent work, Guerreiro et al. (2026) show that a similar hump-shaped response arises in the context of stimulus checks, and argue that the hump can be explained by the same expectations mechanism.

Firm. A competitive firm produces a final good Y_t from labor N_t according to the production function $Y_t = N_t$. In equilibrium, firm optimality implies that the real wage is constant, $w_t = 1$, and profits are zero. This result implies that firms keep prices constant, so there is no price inflation.

Households. In period t, a fraction $N_t \in [0, 1]$ of households is employed. The remaining $1 - N_t$ households are unemployed. The probability that an individual household is employed is the same for all workers and equal to the employment rate N_t. Employed workers earn real wage $w_t = 1$. Unemployed workers receive real benefits $b_t \in (0, 1)$, financed by a lump-sum tax τ_t levied on all households. Once their employment status for the current period ($e_t \in \{0, 1\}$) is determined, households choose consumption c_t and savings a_t in a non-contingent bond with real return r to maximize their anticipated lifetime utility $u(c_0) + \beta\mathbb{E}[u(c_1)]$ subject to period budget constraints

$$c_t + a_t = (1 + r)a_{t-1} + e_t w_t + (1 - e_t)b_t - \tau_t \tag{1}$$

and borrowing constraints $a_t \geq 0$ for $t = 0, 1$. Let the felicity function $u(\cdot)$ be smooth, increasing, concave, and have a positive third derivative, so households are prudent in the sense of Kimball (1990).

At time 0, households may not have perfect foresight of the endogenous variables N_1, b_1, τ_1, and hence even of their own consumption $E[c_1]$. Let $E[N_1]$, $E[b_1]$, and $E[\tau_1]$ denote the beliefs for each variable. We assume that all households have the same beliefs and, since we focus on the first-order response of this economy around a deterministic equilibrium, certainty equivalence holds.

The household Euler equation connects employment uncertainty to current spending:

$$\begin{aligned} u'(c_0) \geq \beta(1+r)\Big[&E[N_1]\, u'\big(1 - E[\tau_1] + (1+r)a_0\big) \\ &+ (1 - E[N_1])\, u'\big(E[b_1] - E[\tau_1] + (1+r)a_0\big)\Big]. \end{aligned} \tag{2}$$

As is standard in models at the zero liquidity limit (Werning, 2015), we assume that at least one Euler equation holds with equality.

Government and monetary policy. The government runs a balanced budget

$$\tau_t = (1 - N_t)b_t. \tag{3}$$

Monetary policy controls the interest rate r and the nominal GDP target for time $t = 1$, $P_1 C_1 = M_1$.

Equilibrium. Goods market clearing implies that $Y_t = C_t$ and asset markets clear.

Beliefs. For the purposes of this section, we impose a simple model of beliefs, based on the idea of belief distortion about future deviations from steady state. Formally, we assume that, for a given variable x_1, households' beliefs are given by

$$E[dx_1] = \lambda dx_1 \tag{4}$$

where λ is a cognitive bias and $dx_1 = x_1 - \overline{x}_1$ denotes the deviation of x_1 from its unshocked value $\overline{x}_1$. Note that $\lambda = 1$ corresponds to full-information and rational expectations (FIRE). If $\lambda < 1$, then beliefs under-react relative to FIRE. If $\lambda > 1$, then beliefs over-react relative to FIRE. As we discuss next, whether beliefs under-react or over-react with respect to the FIRE benchmark is essential in addressing our questions.

Main result. We use this model to illustrate the importance of expectations for the power of UI extensions to stabilize the economy. Consider a period-1 perturbation that raises output and UI generosity by $dY_1 \geq 0$ and $db_1 > 0$, respectively. Under FIRE, the associated period-0 output response is dY_0^{FIRE}. Proposition 1 characterizes how the period-0 response, dY_0, differs from this benchmark when beliefs deviate from FIRE.

Proposition 1. *Consider a period-1 perturbation with changes (dM_1, db_1). Then*

$$dY_0 = dY_0^{FIRE} - \frac{1-\lambda}{1-\text{MPC}_0} \left\{ \text{MPC}_1 \cdot dY_1 + M_{1,b} \cdot db_1 \right\}. \tag{5}$$

Proposition 1 decomposes the effect of imperfect expectations on the transmission of unemployment insurance into two channels, both scaled by the degree of cognitive bias $(1 - \lambda)$ and amplified by the Keynesian multiplier.

The first channel operates through the *direct demand effect of UI*. Under FIRE, households that anticipate more generous future benefits reduce their precautionary savings today, boosting current aggregate demand. When households under-react to the policy ($\lambda < 1$), they underestimate the insurance they will receive if unemployed, and therefore cut their precautionary savings by less than they would under rational expectations. This muted response is captured by $M_{1,b}$, which measures the sensitivity of current consumption to expected future benefits, a quantity closely related to the precautionary savings motive.

The second channel operates through *general-equilibrium feedback*. A UI extension stimulates output and reduces unemployment. Under rational expectations, households anticipate this improvement and raise their consumption accordingly. When expectations are imperfect, households underestimate the increase in future output and the associated decline in unemployment risk, leading them to consume less today than they otherwise would.

Both channels are proportional to the deviation from rational expectations, $1 - \lambda$. When $\lambda = 1$, expectations are rational and the two terms vanish: the policy transmits exactly as under the full-information benchmark. When $\lambda < 1$, households under-react to the future, and both channels reduce the effectiveness of the policy. When $\lambda > 1$, households over-react, and the policy is more powerful than rational expectations would predict. Finally, in general equilibrium, both channels are amplified by the contemporaneous Keynesian multiplier $\frac{1}{1-\mathrm{MPC}_0}$, as any initial shortfall (or surplus) in demand propagates through the usual income-spending feedback loop.

3 Quantitative Model

Next we present a dynamic general-equilibrium model that is suitable for a quantitative evaluation of unemployment benefit extensions with imperfect expectations. It is a New Keynesian model with heterogeneous households and search and matching frictions that is motivated by Kekre (2021) and Auclert et al. (2020). The key distinguishing feature of our model is that it allows for flexible deviations of household expectations from FIRE.

3.1 Households

There is a unit mass of households, who are heterogeneous along four dimensions: labor productivity $z_{i,t}$, liquid assets $a_{i,t}$, employment status $e_{i,t}$, and discount factor β_i. Employment status takes three values: $e_{i,t} = E$ for employment, $e_{i,t} = U$ for unemployment with UI benefits, and $e_{i,t} = X$ for unemployment after benefit exhaustion. The discount factor is permanently fixed at one of three values, $\beta_i \in \{\beta_1, \beta_2, \beta_3\}$, with an equal mass $\mu_k = 1/3$ of households at each.

Time is discrete and each period (a quarter) unfolds in three stages. First, household i draws labor productivity $z_{i,t}$ from an exogenous income distribution $\Pi_z(z_{i,t}|z_{i,t-1})$. Second, labor market transitions take place. Employed workers lose their job with probability $s_{i,t}$; unemployed workers, including those who just separated, find a job with endogenous probability f_t; newly-unemployed workers qualify for UI benefits with probability π^{get}; and ongoing UI recipients lose eligibility with probability π_t^{lose}. The latter probability is the policy instrument for UI extensions. These labor market transitions may be summarized by a single transition matrix for $e_{i,t} \in \{E, U, X\}$:

$$\begin{array}{c c} & \begin{array}{ccc} E_t & U_t & X_t \end{array} \\ \begin{array}{c} E_{t-1} \\ U_{t-1} \\ X_{t-1} \end{array} & \begin{pmatrix} 1 - s_i(1-f_t) & \pi^{get} s_i (1-f_t) & (1-\pi^{get}) s_i (1-f_t) \\ f_t & (1-\pi_t^{lose})(1-f_t) & \pi_t^{lose}(1-f_t) \\ f_t & 0 & 1-f_t \end{pmatrix} \end{array} \tag{6}$$

Following Kekre (2021), we let the separation rate of household i depend on his discount factor

$$s_i = s \exp\left(\Delta_\beta(\beta_i - \beta)\right),$$

which ensures a positive separation rate and allows the model to match the observed differences in wealth and marginal propensities to consume between employed and unemployed households.

In the final stage, the household chooses consumption $c_{i,t}$ and liquid assets $a_{i,t}$ (risk-free deposits at a financial intermediary) to maximize expected lifetime utility subject to a budget constraint and a borrowing limit. The associated Bellman equation is

$$\begin{aligned} V_t(e_{i,t}, z_{i,t}, a_{i,t-1}) = &\max_{c_{i,t}, a_{i,t}} u(c_{i,t}) + \beta_i E_t[V_{t+1}(e_{i,t+1}, z_{i,t+1}, a_{i,t})] \\ &\text{s.t. } c_{i,t} + a_{i,t} = (1 + r^a_{t-1})a_{i,t-1} + (1 - \tau_t)y_t(e_{i,t}, z_{i,t})^{1-\lambda} \\ &\qquad a_{i,t} \geq \underline{a}, \end{aligned} \tag{7}$$

where pre-tax labor income depends on the common wage rate per efficiency unit, w_t, and on replacement rates b_1 and b_2:

$$y_t(e_{i,t}, z_{i,t}) = \begin{cases} w_t z_{i,t}, & \text{if } e_{i,t} = E, \\ b_1 w_t z_{i,t}, & \text{if } e_{i,t} = U, \\ b_2 w_t z_{i,t}, & \text{if } e_{i,t} = X. \end{cases} \tag{8}$$

As is usual, we index unemployment benefits to current productivity instead of productivity at the time of job loss, thereby economizing on the state space. The replacement rate after UI exhaustion, b_2, captures other sources of income, e.g., spousal income as in Bardóczy (2020).

Expectations. $E_t[\cdot]$ denotes households' subjective expectations; $\mathbb{E}_t[\cdot]$ denotes the econometrician's/modeler's objective, complete-information expectations. As in Adam and Marcet (2011), we assume that individuals are internally rational but may hold incorrect beliefs about variables external to them. These variables include the real interest rate, tax rate, aggregate wage, and the job-finding probability. We impose the following restrictions on $E_t[\cdot]$:

Assumption 1. *The expectations operator $E_t[\cdot]$ satisfies the following properties.*

1. ***Observe contemporaneous variables:*** *People observe contemporaneous aggregate variables, i.e., $E_t[x_t] = x_t$ for any variable x_t.*

2. ***Correct expectations of idiosyncratic productivity shocks:*** *People forecast future idiosyncratic productivity shocks correctly, i.e., they use the correct transition probabilities Π_z.*

3. ***Long-run convergence:*** *For any variable x_t that converges to a steady state x, we assume that $\lim_{t\to\infty} E_t[x_{t+h}] = x$.*

4. ***Law of iterated expectations:*** *For any variable x_{t+h}, we assume that $E_t[E_{t+1}[x_{t+h}]] = E_t[x_{t+h}]$.*[5]

The first assumption ensures that individuals know their current income and other relevant budget-constraint variables, so their individual problem is well defined and they do not violate borrowing constraints. The second assumption narrows our focus on boundedly rational expectations to macroeconomic outcomes by imposing correct beliefs about idiosyncratic shocks, i.e.,

[5]Note that this does not imply that $\mathbb{E}_t[E_{t+1}[x_{t+h}]]$ is equal to $E_t[x_{t+h}]$. In fact, it is a feature of our analysis that people make forecast errors that are predictable from the point of view of the econometrician.

to focus on the implications of lack of understanding of the aggregate implications of the policy, we assume that people understand the idiosyncratic component of their income but not the aggregate component. The third assumption implies that, as the economy returns to its steady state, individuals eventually stop making forecast errors. Finally, the fourth assumption means that individuals do not expect to make forecast errors, even if they make systematic errors under the objective probability distribution, i.e., their forecast errors are predictable from the point of view of the econometrician, but households do not expect to make such errors.[6] These properties are shared by a variety of models of beliefs, including full-information rational expectations, noisy-information rational expectations, sticky expectations, diagnostic expectations, and cognitive discounting.

Taken together, these assumptions imply that the steady state of our HANK economy coincides exactly with that which is obtained under full-information and rational expectations. It follows that imperfect expectations affect the transition paths in our economy, but do not alter the long-run steady state.

3.2 Financial intermediary

All assets in the economy are held by a representative financial intermediary. The assets are three: shares in firm equity v_t^E, long-term nominal government bonds B_t, and short-term nominal reserves M_t. The liabilities of the financial intermediary are net worth N_t^{FI} and short-term deposits A_t from households. Thus the balance sheet, in date-t real terms, is

$$p_t v_t^E + q_t^B \frac{B_t}{P_t} + \frac{M_t}{P_t} = N_t^{FI} + A_t \tag{9}$$

where P_t is the price level, p_t is the equity price, q_t^B is price of long nominal bonds, and the price of reserves is 1. Going forward, let $\pi_t = P_t/P_{t-1} - 1$ denote the inflation rate.

The nominal return on these assets are the following. One share of equity purchased in period $t-1$ yields dividend stream $\{P_{t+s}d_{t+s}\}$ for all $s \geq 0$. One government bond purchased in period $t-1$ pays a coupon δ_B^s in period $t+s$ for all $s \geq 0$. One unit of reserves purchased in period $t-1$ pays $(1+i_{t-1})$ in period t. Finally, the intermediary pays dividends d_t^{FI} in period t. This implies that net worth is

$$N_t^{FI} = \underbrace{(d_t + p_t)v_{t-1}^E}_{\text{gross return on equity}} + \underbrace{\frac{1+\delta_B q_t^B}{1+\pi_t}\frac{B_{t-1}}{P_{t-1}} + \frac{1+i_{t-1}}{1+\pi_t}\frac{M_{t-1}}{P_{t-1}}}_{\text{gross return on nominal assets}} - (1+r_{t-1}^a)A_{t-1} - d_t^{FI} \tag{10}$$

For simplicity, we assume that dividends depend on net worth via a simple function

$$d_t^{FI} = d_{ss}^{FI} + \phi_N \left(N_t^{FI} - N_{ss}^{FI} \right). \tag{11}$$

[6]This assumption is violated in models where people are sophisticated about their future forecast biases, as in Lian (2023) or in the sophisticated diagnostic expectations model in Bianchi et al. (2021).

The dividends are paid to a representative owner of the financial intermediary, who simply consumes their dividends in each period.

The financial intermediary is risk neutral. Optimality implies the following asset pricing equations

$$1 + r_t^a = \mathbb{E}_t \left[\frac{d_{t+1} + p_{t+1}}{p_t} \right] = \mathbb{E}_t \left[\frac{1 + \delta_B q_{t+1}^B}{q_t^B (1 + \pi_{t+1})} \right] = \mathbb{E}_t \left[\frac{1 + i_t}{1 + \pi_{t+1}} \right] \equiv 1 + r_t \tag{12}$$

where we defined r_t as the economy-wide ex-ante real interest rate.

3.3 Firms

The aggregate supply block is standard, and affected by three types of frictions: nominal rigidities, investment adjustment costs, and search and matching frictions. To ease the exposition, we delegate these frictions to separate firms: retailers, capital producers, labor agencies.[7]

Retailers. There is a unit mass of retailers indexed by j who engage in monopolistic competition. They produce differentiated goods using a Cobb-Douglas production function with common productivity $y_{jt} = \Theta_t k_{jt}^{\alpha} l_{jt}^{1-\alpha}$. Firms hire capital k_{jt} and labor l_{jt} on spot markets at prices r_t^k and r_t^l and pay a fixed cost Ξ. They also set the price of their product, p_{jt}, subject to a demand curve with constant elasticity ϵ and a quadratic price adjustment cost à la Rotemberg (1982). We allow for price indexation, so the adjustment cost is paid on price changes relative to a fraction ι_p of last period's price change. The firms' objective is to maximize the present value of their future profits. The Bellman equation is

$$J_t^R(p_{jt-1}, p_{jt-2}) = \max_{k_{jt}, l_{jt}, y_{jt}, p_{jt}} \left\{ \frac{p_{jt}}{P_t} y_{jt} - r_t^l l_{jt} - r_t^k k_{jt} - \Psi_{jt}^p - \Xi + \mathbb{E}_t \left[\frac{J_{t+1}^R(p_{jt}, p_{jt-1})}{1 + r_t} \right] \right\}$$

$$\text{s.t. } y_{jt} = \Theta_t k_{jt}^{\alpha} l_{jt}^{1-\alpha}$$

$$y_{jt} = \left(\frac{p_{jt}}{P_t} \right)^{-\epsilon} Y_t$$

$$\Psi_{jt}^p = \frac{\epsilon}{2\kappa_p} \left[\log \left(\frac{p_{jt}}{p_{jt-1}} \right) - \iota_p \log \left(\frac{p_{jt-1}}{p_{jt-2}} \right) \right]^2 Y_t.$$

In a symmetric equilibrium, all firms choose the same level of output, capital, and labor. So, they have the same marginal cost:

$$mc_t = \frac{1}{\Theta_t} \left(\frac{r_t^k}{\alpha} \right)^{\alpha} \left(\frac{r_t^l}{1 - \alpha} \right)^{1-\alpha} \tag{13}$$

and set the same prices according to the Phillips curve

$$\pi_t - \iota_p \pi_{t-1} = \kappa_p \left(mc_t - \frac{\epsilon - 1}{\epsilon} \right) + \frac{1}{1 + r_t} \mathbb{E}_t \left[\frac{Y_{t+1}}{Y_t} \left(\pi_{t+1} - \iota_p \pi_t \right) \right]. \tag{14}$$

[7] Because these firms interact with each other on competitive markets, this formulation is equivalent to treating them as a single firm.

Capital producer. A representative firm owns the capital stock and rents it to retailers at rate r_t^k. Its Bellman equation is

$$J_t^K(K_{t-1}, I_{t-1}) = \max_{K_t, I_t} \left\{ r_t^k K_{t-1} - I_t + \mathbb{E}_t \left[\frac{J_{t+1}^K(K_t, I_t)}{1+r_t} \right] \right\}$$
$$\text{s.t. } K_t = (1-\delta)K_{t-1} + \mu_t \left[1 - S\left(\frac{I_t}{I_{t-1}} \right) \right] I_t \tag{15}$$

where $\delta \in (0,1)$ is the depreciation rate, I_t is investment, μ_t is the marginal efficiency of investment (MEI) as in Justiniano et al. (2010), and $S(\bullet)$ is a convex function that satisfies $S(1) = S'(1) = 0$. In our impulse response matching exercise, we assume that the model economy is perturbed by the MEI shock. Angeletos et al. (2020) argue that this typical shock used in business cycle analysis captures the central features of the main business cycle shock.

Defining Tobin's Q as the marginal value of capital at the end of period t, investment dynamics is characterized by

$$Q_t = \mathbb{E}_t \left[\frac{r_{t+1}^k + Q_{t+1}(1-\delta)}{1+r_t} \right] \tag{16}$$

$$1 = Q_t \mu_t \left[1 - S\left(\frac{I_t}{I_{t-1}} \right) - \left(\frac{I_t}{I_{t-1}} \right) S'\left(\frac{I_t}{I_{t-1}} \right) \right] + \mathbb{E}_t \left[\frac{\mu_{t+1} Q_{t+1}}{1+r_t} \left(\frac{I_{t+1}}{I_t} \right)^2 S'\left(\frac{I_{t+1}}{I_t} \right) \right] \tag{17}$$

Labor agency. A representative firm hires workers on a frictional labor market and rents homogeneous labor services to retailers at rate r_t^l. The agency posts vacancies v_t, each of which is filled with probability q_t. Following Christiano et al. (2016), we assume a two-tiered cost of hiring. The firm pays κ_v to create a vacancy and then κ_h for each vacancy it fills. Incumbent workers have an average productivity of $Z_t = \mathbb{E}[z_{i,t} | e_{i,t} = E]$ and separate exogenously with probability $s_t = \mathbb{E}[s_{i,t} | e_{i,t-1} = E]$. The Bellman equation is

$$J_t^L(N_{t-1}) = \max_{N_t, v_t} \left\{ (r_t^l - w_t) Z_t N_t - (\kappa_v + \kappa_h q_t) v_t + \mathbb{E}_t \left[\frac{J_{t+1}^L(N_t)}{1+r_t} \right] \right\}$$
$$\text{s.t. } N_t = (1-s_t) N_{t-1} + q_t v_t. \tag{18}$$

Solving this problem yields a standard job creation curve, which equates the cost and benefit of hiring the marginal worker:

$$\frac{\kappa_v}{q_t} + \kappa_h = (r_t^l - w_t) Z_t + \mathbb{E}_t \left[\frac{1 - s_{t+1}}{1+r_t} \left(\frac{\kappa_v}{q_{t+1}} + \kappa_h \right) \right]. \tag{19}$$

3.4 Government policies

Fiscal policy. The fiscal authority issues long-term nominal bonds, collects income taxes, and provides unemployment benefits. The government budget constraint is

$$G + UI_t + \frac{(1+\delta_B q_t^B)}{1+\pi_t}\frac{B_{t-1}}{P_{t-1}} = T_t + q_t^B \frac{B_t}{P_t} \tag{20}$$

where UI_t denotes unemployment insurance payments, and $T_t = \mathbb{E}[y_{i,t} - (1-\tau_t)y_{i,t}^{1-\lambda}]$ denotes tax revenues. Government spending G is constant. The income tax rate τ_t is chosen according to the fiscal rule

$$T_t - T_{ss} = \phi_B q_{ss}^B \left(\frac{B_{t-1}}{P_{t-1}} - \frac{B_{ss}}{P_{ss}}\right) \tag{21}$$

We introduce state-dependent UI extensions via a simple rule that links the expected UI duration, $1/\pi_t^{lose}$, to fluctuations in the aggregate unemployment rate via an elasticity $\zeta_b > 0$:

$$\frac{1}{\pi_t^{lose}} - \frac{1}{\pi_{ss}^{lose}} = \zeta_b(U_t - U_{ss}), \tag{22}$$

where $U_t \equiv 1 - N_t$ denote the aggregate unemployment rate.

Automaticity and transmission. Equation (22) is the precise sense in which UI is an automatic stabilizer in the model: conditional on unemployment, benefit duration changes without a new policy decision. Households observe current aggregates, but the precautionary-saving channel depends on the expected future path of unemployment and therefore on expected future eligibility. The rule removes the policymaker's recognition and action problem by construction; it does not remove the private transmission lag.

Monetary policy. The monetary authority sets the short-term nominal interest rate according to

$$i_t = \rho_R i_{t-1} + (1-\rho_R)\left(i_{ss} + \phi_\pi \pi_t\right), \tag{23}$$

where $\phi_\pi \geq 0$ is the Taylor coefficient and $\rho_R \geq 0$ is the interest rate smoothing parameter.

3.5 Equilibrium

Wage setting. We assume that real wages are determined by a simple wage rule that allows for sluggishness in wage adjustments:

$$\frac{w_t}{w_{ss}} = \left(\frac{r_t^l}{r_{ss}^l}\right)^{1-\rho_w}\left(\frac{w_{t-1}}{w_{ss}}\right)^{\rho_w}. \tag{24}$$

Christiano et al. (2016) show that this type of wage rule approximates well the dynamics of more complex bargaining problems.[8]

Matching. New matches are formed on the labor market according to a Cobb-Douglas matching function

$$M(JS_t, v_t) = A_m (JS_t)^{\ell} v_t^{1-\ell} \tag{25}$$

where the mass of job seekers equals the mass of unemployed workers from last period plus the mass of newly separated workers

$$JS_t = 1 - N_{t-1} + s_t N_{t-1} \tag{26}$$

Let $\theta_t \equiv v_t / JS_t$ denote labor market tightness. Job finding and vacancy filling probabilities are

$$f_t = A_m \theta_t^{1-\ell} \qquad \text{and} \qquad q_t = \frac{f_t}{\theta_t}. \tag{27}$$

Market clearing. Factor market clearing requires

$$L_t = \int l_{jt} dj \tag{28}$$

$$K_{t-1} = \int k_{jt} dj \tag{29}$$

The notation reflects that capital is predetermined from the perspective of capital producers but not from the perspective of retailers. Aggregate dividends are given by

$$d_t = d_t^R + d_t^K + d_t^L = Y_t - w_t L_t - I_t - \Psi_t^p - (\kappa_v + \kappa_h q_t) v_t - \Xi \tag{30}$$

Firm equity is then priced according to (12). Asset market clearing corresponds to the balance sheet of the financial intermediary (9), imposing that the intermediary holds all shares $v_t^E = 1$, and nominal reserves are zero $M_t = 0$. Nominal reserves are in zero net supply; their purpose is to deliver a Fisher equation in (12). Goods market clearing requires that the final good is used for household consumption, financial intermediary consumption, investment (including adjustment costs), government spending, price adjustment costs, hiring costs, and the fixed cost.

$$Y_t = C_t + I_t + G_t + \Psi_t^p + (\kappa_v + \kappa_h q_t) v_t + \Xi + d_t^{FI} \tag{31}$$

[8]In our setting, this assumption simplifies the complex problem of modeling wage determination through a bargaining protocol when agents have heterogeneous and non-rational expectations. Gertler et al. (2020) and Hazell and Taska (2025) also estimate wage rules empirically.

4 The Micro Incidence of Unemployment and UI Benefits

Before turning to the dynamic response of the economy to aggregate shocks, it is useful to consider the steady-state solution of our model. This will help us understand the key micro features that we need to match in order to deliver a realistic quantitative evaluation of UI extensions. We assume that expectations are correct in the steady state because all variables are constant. Both our empirical evidence and the model used have the property that expectations converge to the steady state in the long run (see Figure 3).

Calibration of steady state. Our calibration strategy is to choose parameters to match key micro moments in the data. The goal is to provide a quantitatively accurate description of the incidence of unemployment and unemployment insurance, which is essential for a realistic evaluation of the effects of UI extensions. We also target some key macro moments to ensure that the model delivers a realistic description of the macroeconomic environment in which UI extensions operate. The targeted moments are summarized in the table 1.

Table 1: Targeted Moments

Moment	Model	Data	Source
Mean quarterly MPC	0.21	0.21	Kekre (2021)
Annual MPC unemp. - emp.	0.25	0.25	Kekre (2021)
Unemployment rate	0.05	0.05	Kekre (2021)
UI recipient share	0.39	0.39	Kekre (2021)
HH income w UI / pre job loss	0.76	0.76	Kekre (2021)
HH income after UI / pre job loss	0.55	0.55	Kekre (2021)
Reg. coeff. unemployment risk on wealth (%)	-0.02	-0.02	Kekre (2021)

Notes: This table reports the moments targeted in the steady-state calibration of the model. "Mean quarterly MPC" is the average marginal propensity to consume out of a transitory income shock, measured at a quarterly frequency. "Annual MPC unemp. - emp." is the difference in annualized marginal propensities to consume between unemployed and employed households. "Unemployment rate" is the fraction of the labor force that is unemployed. "UI recipient share" is the fraction of unemployed workers currently receiving unemployment insurance benefits. "HH income w UI / pre job loss" is the ratio of household income while receiving UI benefits to household income prior to job loss. "HH income after UI / pre job loss" is the ratio of household income after UI benefit exhaustion to household income prior to job loss. "Reg. coeff. unemployment risk on wealth" is the coefficient from a regression of the annual job-loss probability on log liquid wealth, expressed in percentage points.

Household parameters. Households have CRRA utility, $u(c) = c^{1-\sigma^{-1}}/(1-\sigma^{-1})$, and an EIS of $\sigma = 0.5$. We set the borrowing limit to $\underline{a} = 0$. Our Markov process for labor productivity $(\mathcal{G}_z, \Pi_z)$ is the discrete-time equivalent of the process estimated by Kaplan et al. (2018). We set the income tax progressivity parameter $\lambda = 0.181$ as in Heathcote et al. (2017). Mean productivity is normalized to 1. We choose the probability of losing benefits, $\pi^{lose} = 0.5$, so that expected benefit duration is 6 months. Benefits are a constant fraction of pre-displacement income, so the replacement rates map one for one into the income losses during and after UI eligibility. We set $b_1 = 0.76$ and $b_2 = 0.55$ to match the evidence reviewed by Kekre (2021).

Labor market. We set the job finding rate to $f = 0.6$ and the vacancy filling rate to $q = 0.7$ quarterly. We assume that κ_h accounts for 94% of total search cost, leaving 6% for vacancy posting cost per hire κ_v/q, and set total hiring cost to 7% of the quarterly wage of an average worker. These cost targets, together with the steady-state job-creation condition, determine κ_v, κ_h, and the steady-state wage; the implementation does not introduce a separate bargaining-power parameter.

Aggregate parameters. We assume that the annual economy-wide real interest rate is $r = 2\%$. We calibrate total factor productivity, Θ, to normalize output to $Y = 1$. We set government debt, B/P, to 46% of annual output, and choose the coupon, δ_B, to match the average duration of U.S. government debt of 5 years. We set government spending, G, to 16% of output, which leads to a tax rate of $\tau = 0.29$. We set the depreciation rate to $\delta_K = 0.083/4$ quarterly and calibrate the capital share α to match a quarterly capital-to-output ratio of 8.92. This implies that the steady-state labor share is 62%. The fixed cost Ξ is calibrated to make total wealth $p + q_b B/P$ equal to 382% of annual output.

Internally calibrated parameters. This leaves five parameters to be calibrated jointly: the average discount factor β_2, the dispersion in discount factors $\beta_3 - \beta_2 = \beta_2 - \beta_1$, the scale of the separation rate s, the semi-elasticity of the separation rate with respect to the discount factor Δ_β, and the probability of qualifying for unemployment benefits upon job loss π^{get}. We target five moments in Table 1: the unemployment rate, the average quarterly MPC, the difference between the average MPC of unemployed and employed workers, the share of unemployed workers receiving benefits, and the semi-elasticity of the annual job-loss probability with respect to liquid wealth, all from Kekre (2021).

The unemployment rate pins down the average separation rate: together with the job-finding rate, an unemployment rate of 5% requires employed workers to separate at 7.9% per quarter on average, which the scale s delivers. The two MPC moments are informative about the level and the dispersion of discount factors. The wealth gradient of job-loss risk separates that dispersion from Δ_β, since patient households both accumulate more wealth and, through Δ_β, face lower job-loss risk. The UI recipient share pins down π^{get}. The model matches all five moments exactly, as well as the two income-loss ratios assigned above.

Appendix Table 6 reports every fixed and internally calibrated steady-state parameter, including units, sources, and the primary identifying target.

Untargeted moments. Table 2 shows that the model does well on key untargeted moments. Most importantly, it provides an almost perfect match to the average consumption drop upon job loss and upon benefit exhaustion. The main miss is the level of liquid wealth: households hold almost twice as much relative to income as in the data, although the difference between employed and unemployed households is close to its empirical counterpart. We view the level as acceptable. What counts as a liquid asset is a matter of convention, and the definition behind our target, from

Kaplan et al. (2014), sits at the narrow end of the range; it excludes, e.g., directly held stocks and mutual funds. Finally, the model broadly reproduces the distribution of consumption across wealth quintiles.

Table 2: Untargeted Micro Moments

Moment	Model	Data	Source
HH consumption w UI / pre job loss	0.91	0.92	Ganong and Noel (2019)
HH consumption after UI / pre job loss	0.77	0.76	Ganong and Noel (2019)
Mean liquid net worth / monthly HH income	7.3	3.7	Kekre (2021)
Liquid net worth unemp. - emp.	-3.0	-2.7	Kekre (2021)
Aggregate consumption shares by net worth			
Quintile 5 (highest)	0.40	0.37	Krueger et al. (2016)
Quintile 4	0.20	0.22	Krueger et al. (2016)
Quintile 3	0.15	0.17	Krueger et al. (2016)
Quintile 2	0.14	0.12	Krueger et al. (2016)
Quintile 1 (lowest)	0.09	0.11	Krueger et al. (2016)

Notes: This table reports untargeted moments from the steady-state model alongside their empirical counterparts. "HH consumption w UI / pre job loss" is the ratio of household consumption while receiving UI benefits to consumption prior to job loss. "HH consumption after UI / pre job loss" is the ratio of household consumption after UI benefit exhaustion to consumption prior to job loss. "Mean liquid net worth / monthly HH income" is the ratio of average household liquid net worth to average monthly household income. "Liquid net worth unemp. - emp." is the difference in the ratio of liquid net worth to monthly income between unemployed and employed households. "Aggregate consumption shares by net worth" reports the share of total consumption accounted for by each quintile of the liquid net worth distribution.

Figure 2 elaborates on the expected income and consumption responses of a typical unemployed worker who stays unemployed until a given quarter. The figure displays a good match to the empirical patterns in Ganong and Noel (2019). The key findings is that the model reproduces the large drops in consumption upon job loss and upon benefit exhaustion. Interestingly, the model accounts well for these empirical patters without relying on bounded rationality. Instead, the model delivers this result via a combination of liquidity constraints and heterogeneity in the exposure to unemployment risk.

5 The Dynamic Response to Shocks with Imperfect Expectations

The previous section shows that the model delivers a realistic description of the micro incidence of unemployment and UI benefits in the steady state. We now evaluate its dynamic response to aggregate shocks and ask how the effects of UI extensions change when we replace FIRE with imperfect expectations. To do so, we estimate the remaining transition parameters by matching the model's impulse responses to their empirical counterparts while using observed expectations and forecast errors directly.

We begin by laying out the general framework for dynamic models with imperfect expectations. We then describe the estimation strategy and its results before evaluating UI extensions under alternative assumptions about expectations. By characterizing responses to an initial forecast

Figure 2: Income and Consumption Dynamics Through an Unemployment Spell

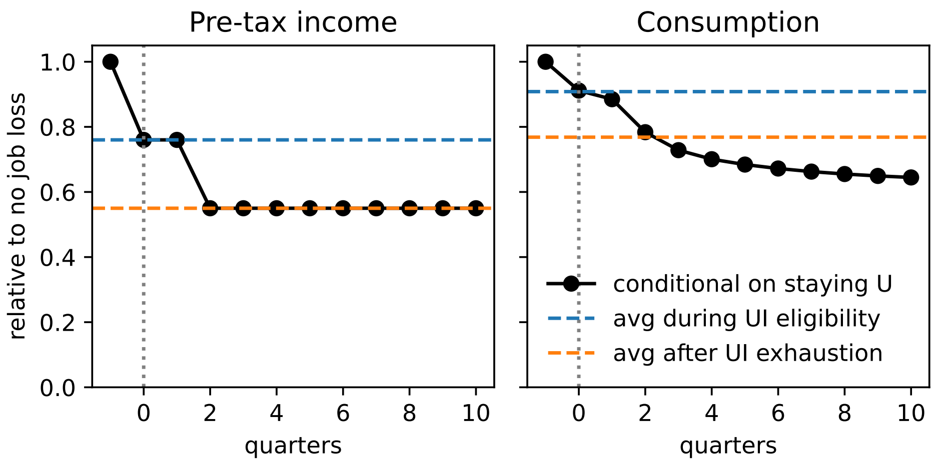

Notes: This figure plots model-implied income and consumption dynamics for an unemployed worker who remains unemployed through each horizon. Income and consumption are expressed relative to their pre-job-loss levels. The empirical moments used for comparison are from Ganong and Noel (2019).

and a sequence of forecast revisions, the framework turns Friedman's third lag from an unspecified delay into a sequence of empirically measured revisions and their effects on current decisions.

5.1 A Framework for Dynamic Models with Imperfect Expectations

The framework has two components. First, a model of how a dynamic block maps expectations into decisions and aggregate outcomes. Second, a model of how expectations are formed from observations. In Appendix C, we present further details and show how to map many popular models of bounded rationality and information frictions into our framework.

Consider a forward-looking block of the economy that produces an aggregate outcome Y_t over periods $t = 0, 1, \dots, T-1$. The block could be a household block, a firm block, a financial block, or any other part of the model for which one can compute sequence-space Jacobians. Let $\mathbf{Y}$ denote the path of this outcome. For ease of exposition, let every object (parameters, initial and terminal conditions, and other exogenous paths) that matters for the block be fixed and known except the path of a single univariate input $\mathbf{X}$. The extension to multiple time-varying inputs and multiple outputs is straightforward. Auclert et al. (2021) cast such dynamic problems under FIRE as a mapping between sequences

$$\mathbf{Y} = f(\mathbf{X}) \tag{32}$$

Their SSJ method computes the Jacobian $\mathcal{J}$ and then computes impulse responses to any shock $d\mathbf{X}$ via matrix multiplication, $d\mathbf{Y} = \mathcal{J} \cdot d\mathbf{X}$.[9]

[9]The Jacobian is computed at a baseline path $\bar{\mathbf{X}}$, typically a constant path corresponding to the steady state $\bar{\mathbf{Y}} = f(\bar{\mathbf{X}})$. So the shock $d\mathbf{X} = \mathbf{X} - \bar{\mathbf{X}}$ and the impulse response $d\mathbf{Y} = \mathbf{Y} - \bar{\mathbf{Y}}$ are both deviations from the baseline path.

We build on recent contributions by Auclert et al. (2020) who show how to extend the SSJ method to accommodate a wide range of expectations beyond FIRE. We retain certainty equivalence, so only the mean expectation matters. However, the mean expectation may not be correct and may evolve over time. In period 0, agents expect the path $E_0[d\mathbf{X}]$; in period 1, they expect a path $E_1[d\mathbf{X}]$; and so on. Each vector $E_\tau[d\mathbf{X}] = \begin{bmatrix} E_\tau[dX_0] & E_\tau[dX_1] & \ldots & E_\tau[dX_t] & \ldots \end{bmatrix}'$ captures the beliefs that agents hold at time τ about the variable X at every date. The only restrictions on expectations are those in Assumption 1: observing contemporaneous variables[10], long-run convergence, and the law of iterated expectations. Given this setup, relaxing FIRE implies that we have to keep track of the entire history of expectations, $E_t[\mathbf{X}]$ for all t. Formally,

$$\mathbf{Y} = g\left(\mathbf{X}, \{E_t[\mathbf{X}]\}_t\right). \tag{33}$$

Conceptually, if we could compute all the Jacobians of $g(\bullet)$, we could compute linearized impulse responses. But the domain of $g(\bullet)$ is a much larger space than the domain of $f(\bullet)$, since it involves an infinite path of expectations, which are themselves infinite paths. We leverage the insights in Auclert et al. (2020) to show that we can compute the impulse response to any shock $d\mathbf{X}$ and any history of expectations $\{E_t[d\mathbf{X}]\}_t$ using only the FIRE Jacobian $\mathcal{J}$ and some simple manipulations of it. In Proposition 2, we show that, for general expectations, we need to describe the response of the output path, not just to the initial forecast $E_0[d\mathbf{X}]$, but also to forecast revisions $E_h[d\mathbf{X}] - E_{h-1}[d\mathbf{X}]$ for all h. Furthermore, we show that the Jacobians that capture the responses to forecast revisions can be readily computed from the FIRE Jacobian $\mathcal{J}$, at virtually no additional cost.[11]

Proposition 2. *Suppose that expectations satisfy Assumption 1, and truncate all paths at a common horizon T. Then, the linearized impulse response $d\boldsymbol{Y} \in \mathbb{R}^T$ to arbitrary shocks $d\boldsymbol{X} \in \mathbb{R}^T$ and beliefs $E_t[d\boldsymbol{X}] \in \mathbb{R}^T$ is given by*

$$d\boldsymbol{Y} = \mathcal{J} \underbrace{E_0[d\boldsymbol{X}]}_{\textit{initial forecast}} + \sum_{h=1}^{T-1} \mathcal{R}_h \underbrace{\left(E_h[d\boldsymbol{X}] - E_{h-1}[d\boldsymbol{X}]\right)}_{\textit{forecast revision}}. \tag{34}$$

*For $h = 1, \ldots, T-1$, let $\mathcal{J}^{(T-h)} = [\mathcal{J}_{ts}]_{t,s=0}^{T-h-1}$ denote the leading $(T-h) \times (T-h)$ submatrix of $\mathcal{J}$. The **forecast-revision Jacobian** is the $T \times T$ matrix*

$$\mathcal{R}_h = \begin{bmatrix} \boldsymbol{0}_{h\times h} & \boldsymbol{0}_{h\times(T-h)} \\ \boldsymbol{0}_{(T-h)\times h} & \mathcal{J}^{(T-h)} \end{bmatrix}, \qquad (\mathcal{R}_h)_{ts} = \mathbf{1}\{t \geq h,\ s \geq h\}\mathcal{J}_{t-h,s-h}.$$

[10] This assumption ensures that agents do not violate any constraints.

[11] This representation is not specific to consumption. If the block is the household sector, one natural output is aggregate consumption, $Y_t = C_t = \int c_{i,t} di$, and this is the case we emphasize in the quantitative exercises below. But the same logic applies to any other aggregate outcome produced by a dynamic block. The paper focuses on imperfect expectations among households, so we apply the framework to the household block while keeping the notation in this subsection general.

To see the intuition for the representation in Proposition 2, consider the following. Suppose that, at date 0, agents hold an incorrect belief about the variable X at some future date s, so that

$$E_0[dX_s] \neq dX_s,$$

and suppose further that this mistaken belief persists until a forecast revision arrives at some date $h \leq s$, that is,

$$E_h[dX_s] - E_{h-1}[dX_s] \neq 0.$$

Then the resulting response of the output at any date $t \leq s$ can be written as the sum of two components, $dY_t = dY_t^{\text{anticipated}} + dY_t^{\text{revision}}$, where the first term captures the effect of the belief held at date 0, while the second term captures the additional adjustment induced when beliefs are revised at date h.

This decomposition is useful because the two objects correspond exactly to the two Jacobians that appear in our characterization. The matrix $\mathcal{J}$ governs the response to beliefs that are already in place and therefore anticipated from the outset. By contrast, the matrix $\mathcal{R}_h$ governs the response to belief changes that are realized unexpectedly at date h. The key point is that arbitrary belief dynamics can be understood by combining these two building blocks.

To see this, first isolate the effect of the initial mistaken belief $E_0[dX_s]$. Suppose, for the moment, that there are no subsequent forecast revisions, so agents continue to hold that belief until date s and then observe the truth thereafter. In that case, the change in expectations is fully anticipated from date 0 onward, and the outcome responds exactly as in the perfect-foresight sequence-space Jacobian:

$$dY_t^{\text{anticipated}} = \mathcal{J}_{t,s}\, E_0[dX_s], \qquad \text{for all } t \geq 0.$$

Intuitively, this is just the standard response to a known future movement in the input: agents internalize the mistaken forecast immediately and adjust their plans accordingly.

Next consider the effect of the forecast revision at date h. At that date, agents learn that their earlier view of the input at date s was incorrect, and they update beliefs by an amount $E_h[dX_s] - E_{h-1}[dX_s]$. Since this revision is not anticipated before h, it has no effect on the outcome at earlier dates. But once the revision arrives, agents react as if hit by a new MIT shock at date h. Therefore,

$$dY_t^{\text{revision}} = \begin{cases} 0, & t < h, \\[2ex] \mathcal{J}_{t-h,\,s-h}\ \left(E_h[dX_s] - E_{h-1}[dX_s]\right), & t \geq h. \end{cases}$$

Equivalently, the forecast-revision Jacobian $\mathcal{R}_h$ can be viewed as a time-shifted version of the original Jacobian $\mathcal{J}$: it describes how an unanticipated change in beliefs propagates once it is realized.

Taken together, these two pieces clarify the economic content of our representation. The matrix $\mathcal{J}$ captures how the output responds to anticipated paths of the perceived input, while the matrices $\mathcal{R}_h$ capture how the output responds when beliefs are revised unexpectedly over time. By stacking these responses across horizons s and revision dates h, one obtains the generalized sequence-space representation used in the paper. This representation makes transparent how the original Jacobians can be repurposed to handle rich belief dynamics without introducing any new conceptual ingredients.

Remark 1 (Heterogeneous Beliefs)**.** *Proposition 2 applies to models in which individuals have heterogeneous beliefs, as long as the distribution of beliefs is uncorrelated with other idiosyncratic characteristics in the cross-section. In this case, we can redefine $E_t[X]$ as the cross-sectional average expectation, and the framework still applies.*[12]

Remark 2 (Implementing Models of Beliefs)**.** *Proposition 2 provides a general representation of the impulse response to arbitrary beliefs, i.e., in the* temporary equilibrium*. This general methodology allows us to estimate the model without taking a stance on the specific model by which people form expectations. However, we also note that this representation allows us to compute the IRF to any shocks and under any model of beliefs, which is useful for counterfactual analysis. In Appendix C.4, we show how to map many popular models of bounded rationality and information frictions into our framework. In many models of belief formation, expectations take the form of a forecast matrix Λ_t times the change in the variable $d\boldsymbol{X}$:*

$$E_t[d\boldsymbol{X}] = \Lambda_t \cdot d\boldsymbol{X},$$

for some matrix Λ_t. When this is the case, the representation in Proposition 2 can be simplified to

$$d\boldsymbol{Y} = \underbrace{\left[\mathcal{J}\Lambda_0 + \sum_{h\geq 1} \mathcal{R}_h(\Lambda_h - \Lambda_{h-1})\right]}_{\equiv \tilde{\mathcal{J}}} \cdot d\boldsymbol{X},$$

where $\tilde{\mathcal{J}}$ can be interpreted as the boundedly-rational Jacobian that captures the response to shocks under the given model of beliefs.

5.2 Estimation

The framework in the previous section allows us to compute impulse responses to any shock and any history of expectations. To use this framework, we need to estimate the parameters that govern the model's dynamics.

We estimate the model in two steps, similarly to Christiano et al. (2005) and Auclert et al. (2020). We first set some parameters exogenously and then estimate the remaining parameters by matching the model's impulse responses to empirical impulse responses.

[12] It is also possible to use this framework to allow for meaningful belief disagreement as long as beliefs are with permanent individual characteristics. In this case, one has to set up a heterogeneous-agent block for each permanent type, and apply the propositions type by type. Guerreiro (2022) follows this approach in his study of disagreements over the business cycle.

The exogenous parameters are those related to the steady state, which we have already pinned down in the previous section. Additionally, we set ζ_b, the semi-elasticity of average unemployment duration with respect to the unemployment rate, to approximate the Emergency Unemployment Compensation (EUC08) program. The unemployment rate in 2007Q1 was 4.6%, close to the steady state value of 4.5% in our model. Unemployment rate peaked at 10.1% in 2009Q4. During the same time, unemployment benefit duration was raised from 26 weeks to 99 weeks (in states with unemployment rate above 8.5%). So, we calibrate the semi-elasticity to $\zeta_b = (99-26)/13/(0.101-0.046) \approx 102$. That is, a one percentage point increase in the unemployment rate triggers 1.02 quarter increase in average UI duration.

We now describe our estimation methodology.

Estimation: IRF matching. We estimate the model by matching the impulse response functions obtained in our model to their empirical counterparts obtained with a standard business-cycle shock. We generate the empirical impulse responses by following the empirical strategy in Angeletos et al. (2020) (and replicated also in Angeletos et al., 2021). We estimate a multivariate structural VAR with 12 variables, plus the forecast variables of interest. We then apply their classical inference procedure and replicate the identification of the main business-cycle shock in Angeletos et al. (2020). This shock is constructed to account for most business-cycle fluctuations in the unemployment rate.

We generate impulse response functions not only for outcomes, but also for expectations of the relevant variables at several horizons. These IRFs are the key empirical target that we use to estimate the parameters governing the model's dynamics. The estimation sample begins in 1981Q3, when the SPF data become available. We target the first twelve quarters of the empirical IRFs for unemployment, consumption, investment, inflation, and the three-month Treasury bill rate. Figure E.1 shows the full set of estimated impulse response functions.

To perturb the economy from its steady-state level, we consider a single shock to the marginal efficiency of investment (MEI) which follows an AR(1) process

$$\mu_t = \rho_\mu \mu_{t-1}.$$

The process for this shock has two free parameters: the persistence ρ_μ and the initial level of the shock μ_0. Angeletos et al. (2020) show that the MBC and MEI shocks are closely related. Using our model, we solve the impulse responses of variables in our model to this shock for a given set of parameters.

Importantly, our methodology in equation (34) allows us to directly use the estimated impulse responses for expectations to solve the model. We focus on the implications of imperfect expectations for aggregate demand and so assume that all other economic agents have full information and rational expectations.[13]

[13] It is easy to extend the estimation to incorporate imperfect expectations into the decision problems of other economic agents.

We recover the model-implied impulse responses $\text{IRF}_{v,h}(\Omega)$, where Ω denotes the parameters in Table 4, $v \in \mathcal{V} \equiv \{u, c, \text{inv}, \pi, i\}$ indexes the five targeted variables, and $h = 0, \ldots, 11$ indexes quarters. The implemented minimum-distance criterion uses diagonal weights:

$$\widehat{\Omega} = \arg\min_{\Omega} \sum_{v \in \mathcal{V}} \sum_{h=0}^{11} \left[w_{v,h} \left(\text{IRF}_{v,h}(\Omega) - \widehat{\text{IRF}}_{v,h} \right) \right]^2, \qquad w_{v,h} \equiv \frac{1}{\widehat{\text{IRF}}_{v,h}^{\text{upper}} - \widehat{\text{IRF}}_{v,h}^{\text{lower}}}, \tag{35}$$

where the weights $w_{v,h}$ are inversely proportional to the standard errors of the empirical IRFs.

Expectations: source, data limitations, and solution. We use data from the Survey of Professional Forecasters (SPF). From this dataset, we use one- to four-quarter-ahead unemployment rate, interest rate, and inflation forecasts. Figure 3 Panel A plots the estimated impulse responses of the unemployment rate and the forecasts of the unemployment rate at different horizons to the MBC shock. Figure 3 Panel B plots the analogous impulse responses of the real interest rate and its forecasts. The figure shows that there is an initial underreaction in the forecasts, followed by a delayed overreaction, as pointed out by Angeletos et al. (2021).

In practice, however, the exercise described in the previous section cannot be fully implemented due to the unavailability of all necessary expectations data. In this dimension, we confront two limitations: (1) we do not have survey data on expectations for all relevant variables, and do not observe expectations of tax rates or job finding rates, for example; and (2) even for the variables for which we do have survey data, we observe expectations only for a finite number of future horizons (four quarters), not the infinite number of horizons that would be required to solve the model.

We solve both of these issues by imposing a parametric model of beliefs to generate beliefs of variables for which expectations data are lacking and extrapolate unemployment forecasts beyond the fourth quarter horizon. We estimate a parametric model of beliefs that is flexible enough to capture the empirical patterns of observed expectations and use this model to fill the gaps in the data.

Filling the gaps with a model of beliefs: how to capture delayed overreaction? As discussed before, a key pattern in the data is the initial underreaction followed by delayed overreaction in expectations. This pattern is crucial for our analysis, as it implies that the effects of shocks on expectations can be non-monotonic over time, which in turn can have important implications for the dynamic response of the economy to shocks. Therefore, it is important that the model of beliefs we use to fill the gaps in the data is able to capture this pattern.

To capture the estimated pattern of initial underreaction followed by delayed overreaction, we combine noisy information with diagnostic expectations and long memory. In doing so, we build on Bordalo et al. (2020) (who combined noisy information with standard diagnostic expectations) and on Bianchi et al. (2021) (who introduced diagnostic expectations with long memory).[14]

[14]In Appendix D.2, we discuss the merits of this model of beliefs relative to other popular models in the literature.

In this model, people receive noisy signals about the state of the economy and update their beliefs according to Bayes' rule on a distorted probability distribution that overweights states that become relatively more likely given the new signal, as in Bordalo et al. (2018).

As we discuss in Appendix C.4, the noisy-information and long-memory diagnostic expectations model implies that the time t average expectation to a deterministic shock takes the following form:

$$\overline{E}_t[dX_{t+h}] = \left[(1+\theta)\xi_t - \theta \sum_{j=1}^{t} \alpha_j \xi_{t-j}\right] dX_{t+h}, \tag{36}$$

where $\overline{E}_t[dX_{t+h}]$ denotes the average expectation, $\xi_t \equiv \frac{t+1}{\tau_\epsilon/\tau_\nu + t+1}$ denotes the precision of information accumulated up to time t, and we use the convention $\sum_{j=1}^{0}(\cdot) = 0$. This model features several parameters: θ denotes the degree of belief overreaction, $\alpha_j \geq 0$ for $j \geq 1$ denote the memory weights and satisfy $\sum_{j=1}^{\infty} \alpha_j = 1$, and $\tau \equiv \tau_\epsilon/\tau_\nu$ denotes the relative precision of priors to the precision of the noisy signals.

We define the cognitive distortion coefficient as the elasticity of average beliefs with respect to the econometrician's forecast:

$$\lambda_t \equiv \frac{\overline{E}_t[dX_{t+h}]}{dX_{t+h}}. \tag{37}$$

Underreaction of expectations corresponds to $\lambda_t < 1$, while overreaction corresponds to $\lambda_t > 1$. The success of this model in capturing the observed pattern of expectations hinges on the fact that λ_t can be non-monotonic over time, as the initial underreaction (when ξ_t is small) can be followed by a delayed overreaction (as ξ_t increases and the diagnostic component of beliefs becomes more important). The long-memory component allows for the overreaction to persist over time, albeit with a declining intensity, which is crucial in capturing the observed pattern of expectations.

The combination of features in this model of belief formation provides an intuitive mechanism for the observed pattern of expectations. Information arrives slowly over time, through the accumulation of noisy signals. This slow information accumulation generates the initial underreaction. Over time, as information accumulates, the diagnostic component of beliefs generates a delayed overreaction. The standard model of diagnostic expectations without long memory can only capture one period of overreaction to new information. The long-memory component allows the overreaction to persist over time, albeit with declining intensity. Provided that the overreaction persists sufficiently long, the model can capture the delayed overreaction as people continue to overreact to the accumulation of signals they receive.

Here, we discuss relative to models that are nested within this framework, see Remark 3.

Table 3: Belief Parameters

Parameter	Description	Value
θ	Diagnostic expectation param	4.966
τ	Noisy information param	16.877
α	Long memory param 1	7.878
β	Long memory param 2	42.881

Notes: This table reports the calibrated parameters of the noisy-information and long-memory diagnostic expectations model. The parameter θ controls diagnostic overreaction, τ controls the relative precision of priors to noisy signals, and α and β parameterize the Beta-binomial memory weights used for $\{\alpha_j\}_{j\geq 1}$.

We assume that the α_j are determined by a Beta-binomial distribution with parameters α and β. This assumption implies that we have four parameters to calibrate in this model: θ, τ, α, and β. We calibrate these parameters so that the beliefs they imply for unemployment rate forecasts line up with those that we observe in the data. However, note that, in solving the model, we actually use directly the observed unemployment rate forecasts and not the ones implied by this model. The calibrated parameters are found in Table 3, and the model's empirical match can be seen in Figure 3. Overall, the model provides a good fit to the observed patterns of expectations. So, we use this model to fill the gaps in the data, but we use the observed expectations directly whenever possible.

Figure 3: Calibration of Belief Parameters

(a) Panel A: Unemployment forecast

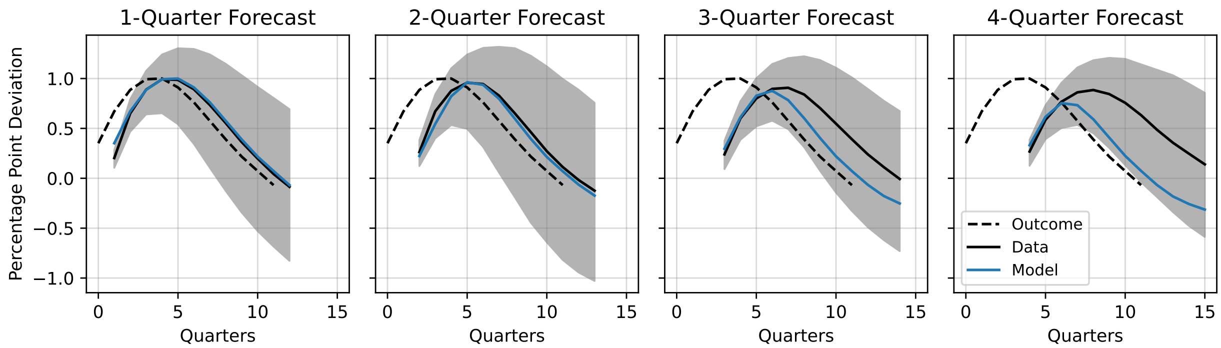

(b) Panel B: Real rate forecast

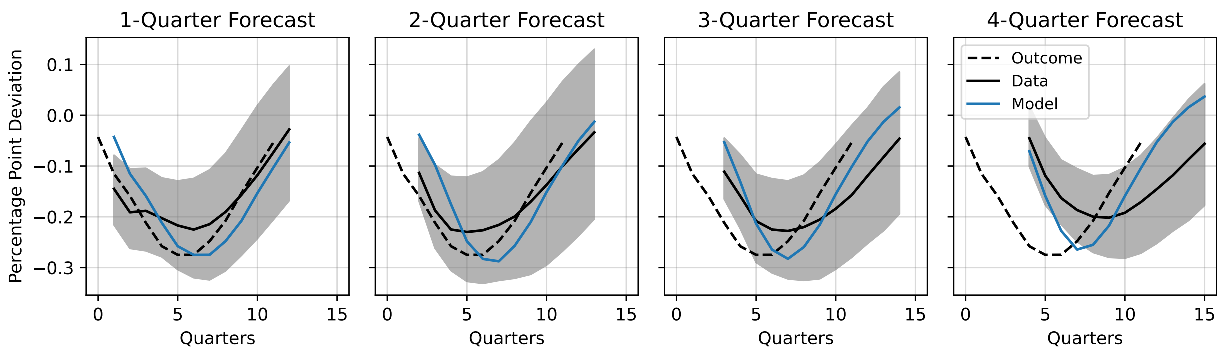

Notes: The blue lines plot the forecasts implied by the estimated model of beliefs. The black lines plot the forecasts observed in the data. The shaded areas plot the 68% confidence interval around the empirical point estimates.

A comparison to nested models. We now briefly comment on how our model of beliefs compares to some standard models in the literature. This discussion helps understand the importance of the combination of features in our model in capturing the observed pattern of expectations.

Remark 3 (Nested models)**.** *This model nests four known models as special cases:*

(1) The model nests the standard diagnostic expectations of Bordalo et al. (2018), when $\tau = 0$ and $\alpha_1 = 1$. In this case, the cognitive distortion is given by

$$\lambda_t = \begin{cases} 1+\theta, & \text{if } t = 0 \\ 1, & \text{if } t > 0, \end{cases}$$

so $\lambda_t \geq 1$ for all t, which implies that beliefs always overreact, but only at the initial date. Thus, this baseline model cannot capture the empirical pattern.

(2) The model nests the long-memory diagnostic expectations model in Bianchi et al. (2021), when $\tau = 0$ but allowing for the memory weights to assign mass to further away expectations. In this case, the cognitive

distortion is given by

$$\lambda_t = 1 + \theta \left(1 - \sum_{j=1}^{t} \alpha_j\right) \geq 1.$$

Note that long memory delivers more persistent overreaction as λ_t falls smoothly over time. However, it can only capture overreaction of expectations.

(3) The model nests the noisy-information and rational expectations model as in Appendix H of Angeletos and Huo (2021), when $\theta = 0$ but $\tau > 0$. In this case, the cognitive distortion is given by

$$\lambda_t = \frac{t+1}{\tau + t + 1} < 1,$$

so beliefs always underreact, and the model cannot capture the empirical pattern.

(4) The model nests the standard noisy-information and diagnostic expectations model used in Bordalo et al. (2020), when $\theta > 0$ and $\tau > 0$ but $\alpha_1 = 1$. In this case, the cognitive distortion is given by

$$\lambda_t = (1+\theta)\frac{t+1}{\tau + t + 1} - \theta \frac{t}{\tau + t}.$$

In this model, $\lambda_t > 1 \Leftrightarrow t < \theta - \tau$, and otherwise $\lambda_t \leq 1$. Intuitively, this change in relative reaction occurs when signals are very uninformative (τ is small), and there is substantial overreaction (θ is large), so that $\theta - \tau > 0$. Otherwise, the model can only capture underreaction. This model can generate periods of underreaction and overreaction, but the timing is reversed: it can only produce an initial overreaction followed by delayed underreaction, which is the opposite of what we see in the data.

Our model is best understood as extending this final model to allow for long memory, which turns out to be essential in capturing the pattern of initial underreaction followed by delayed overreaction that can be seen in Figure 3.[15]

6 Estimation Results

In this section, we present the main results of our estimation exercise. Furthermore, we compare the implied impulse responses in our model to the benchmarks of perfect and no anticipation of future changes. This exercise allows us to understand the implications of the patterns of imperfect expectations observed in the data.

We estimate the remaining parameters that are relevant for the transition dynamics in our economy. These parameters are as follows. The monetary policy parameters are the Taylor-rule coefficient on inflation, ϕ_π, and Taylor-rule inertia, ρ_R. The investment adjustment cost is ψ, and the wage-rule parameter is ρ_w. The nominal-rigidity parameters are price indexation, ι_p, and the slope of the Phillips curve, κ_p. The marginal response of intermediary payouts to current net worth is ϕ_N. Finally, μ_0 and ρ_μ control the scale and persistence of the MEI shock. The elasticity of the tax rate to debt, ϕ_B, is not well identified by the targeted impulse responses, so we fix it to one

[15] See Appendix D.2 for a discussion.

rather than estimate it. The estimated values for the parameters that affect the transition dynamics in our economy can be found in Table 4. Fitting the estimated inflation and nominal rate responses requires a steep Phillips curve and a relatively large inflation coefficient in the Taylor rule. Here κ_p is the slope with respect to the level deviation of real marginal cost; the corresponding elasticity is $\kappa_p \, mc_{ss} = 0.13$, still steep relative to conventional estimates. The rest of the estimated parameters are typical values.

Table 4: Estimated Parameters

Parameter	Description	Value
Core model parameters		
ϕ_π	Taylor rule coef on inflation	3.868
ρ_R	Taylor rule inertia	0.747
ψ	Investment adjustment cost	5.415
ρ_w	Wage rule parameter	0.363
ι_p	Price indexation	0.248
κ_p	Phillips curve slope	0.156
ϕ_N	Marginal payout response	0.135
MEI shock process		
μ_0	Scale of MEI shock	0.127
ρ_μ	Persistence of MEI shock	0.639

Notes: This table reports the estimated transition-dynamics parameters used in the quantitative model.

The model's match to the targeted impulse response functions is shown by the blue lines in Figure 4 for the unemployment rate, consumption, investment, inflation rate, and the nominal interest rate. The black line shows the associated empirical impulse responses, and the shaded region plots the 68% confidence interval around the empirical point estimates. The model provides a good fit to its targeted empirical counterparts.

Figure 4: Targeted Impulse Responses: Outcomes

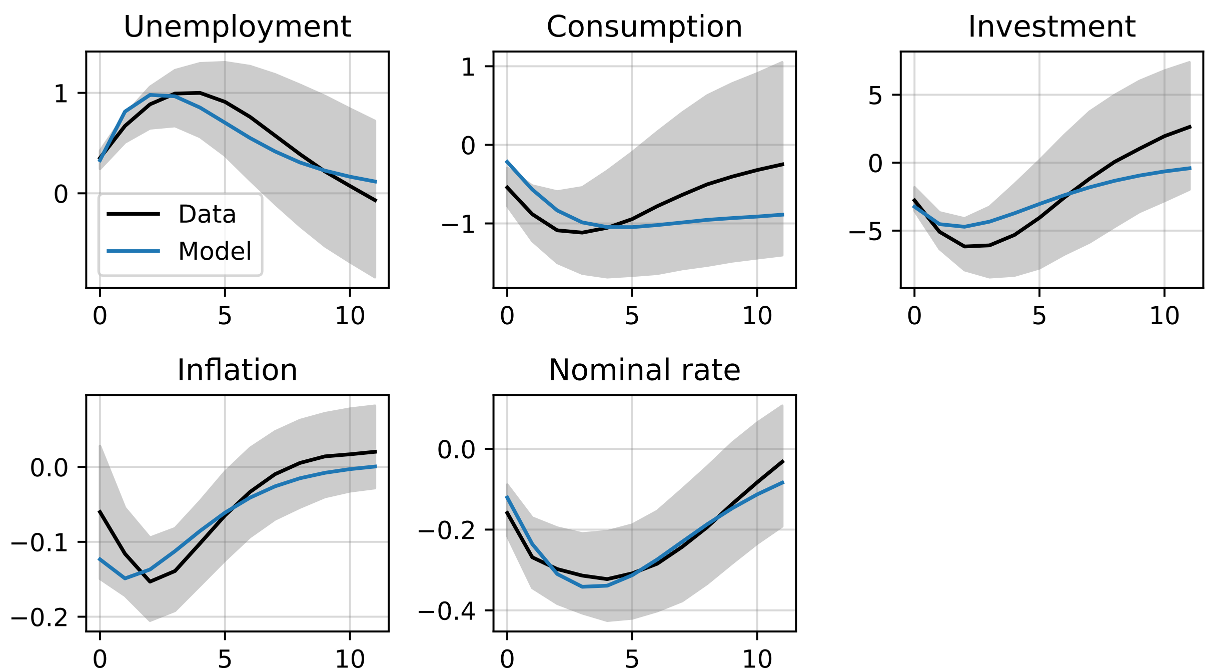

Notes: This figure compares the model-implied impulse responses with the empirical impulse responses targeted in the estimation. Blue lines denote the estimated model, black lines denote the empirical responses, and shaded areas denote 68% confidence intervals around the empirical point estimates.

The determination of hump-shaped responses. The model's ability to match the empirical impulse responses is a key feature of our estimation. We now argue that the success in matching these empirical patterns is largely driven by its ability to capture the patterns of expectations observed in the data. To see this, we contrast the model's consumption response with its counterpart under FIRE. We summarize the comparison using two statistics: (1) the initial consumption response, dC_0, and (2) the trough-to-impact magnitude ratio, $|\min_t dC_t|/|dC_0|$. The first statistic captures the impact response, while the second captures the hump-shaped nature of the contraction.

The results of this comparison are shown in Figure 5 and summarized in Table 5. Relative to FIRE, the estimated model delivers a much smaller initial decline in consumption and a trough-to-impact ratio nearly three times as large. The FIRE response is therefore substantially more front-loaded and less hump-shaped.[16] The impact decline under FIRE is more than twice as large as under estimated beliefs, while its trough-to-impact ratio is about 35% as large. These comparisons show that the expectations observed in the data are crucial for the dynamic consumption response to shocks.

[16]Consumption remains slightly hump-shaped under FIRE in this model because slow-moving unemployment also contributes to the response.

Table 5: IRF Statistics in Estimated Model vs. FIRE

Statistic	Definition	Specification	
		Estimated	FIRE
Initial response	dC_0	-0.28	-0.64
Trough-to-impact ratio	$\frac{\lvert \min_t dC_t \rvert}{\lvert dC_0 \rvert}$	3.87	1.43

Notes: The table reports the impact response and the ratio of the largest contraction in consumption to the impact contraction. Values are generated by the same model solution plotted in Figure 5.

The mechanism underlying this result is intimately related to the empirical patterns of expectations. Initially, people's expectations underreact, so they underforecast the future decline in income and increase in unemployment. They therefore do not adjust their consumption plans much on impact, which delivers a smaller initial decline in consumption. Over time, however, people's expectations turn more pessimistic and eventually overreact. They then overforecast the future decline in income and increase in unemployment, so they adjust consumption by more than the realized paths would justify. This mechanism delivers a larger trough-to-impact ratio.

Figure 5: Consumption Response: Estimated Model vs. FIRE

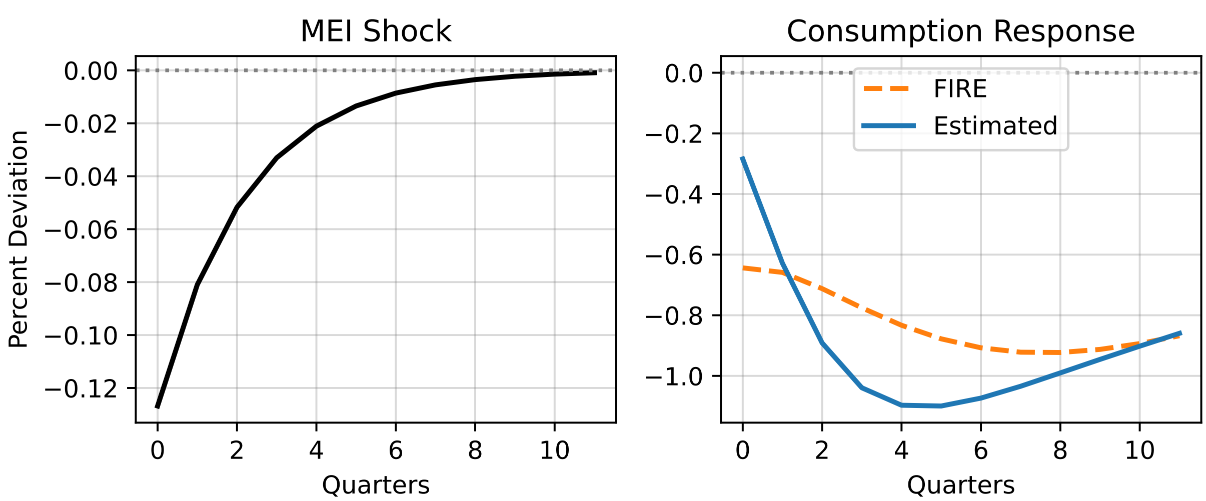

Notes: The left panel plots the marginal-efficiency-of-investment shock; the right panel plots the general-equilibrium consumption impulse response from impact through quarter 10, under the estimated belief model and under FIRE.

7 The Consequences of Imperfect Expectations for UI Policy

In this section, we assess the efficacy of UI extensions on stimulating demand with imperfect expectations. The goal is to quantify the role of imperfect anticipation of endogenous UI extensions

Figure 6: Direct Impact of UI Duration and Consumption

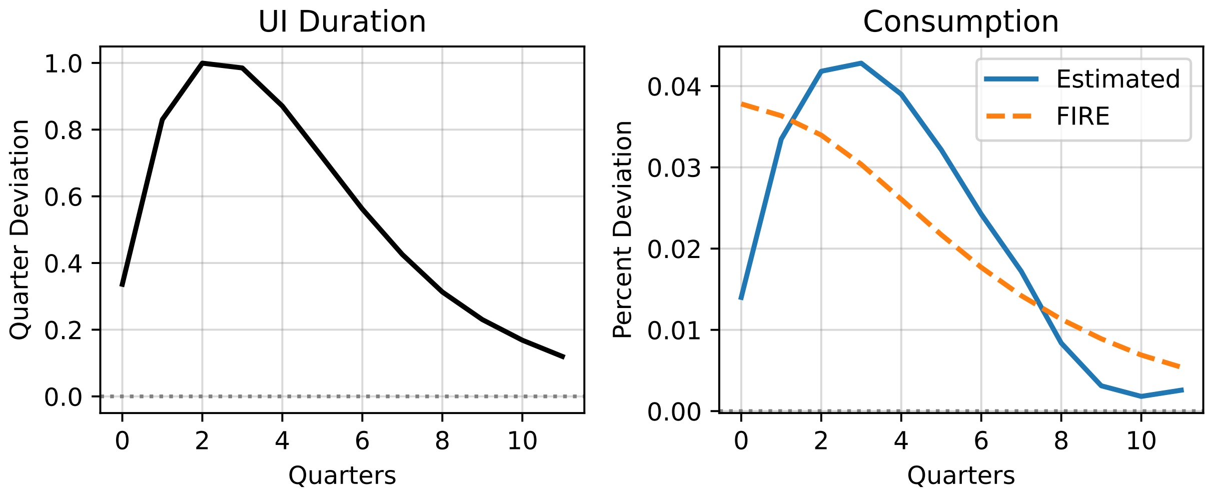

Notes: This figure reports the direct impact on aggregate consumption of the UI-duration path implied by the empirical recessionary shock. The exercise feeds the UI-duration path into the calibrated household block and compares the response under estimated beliefs with the response under FIRE.

in affecting aggregate demand. As we have discussed before, we do so by working directly with the empirical response of expectations, avoiding the need to choose a particular model of belief formation. We show that the direct effect of UI extensions on the distribution of income is less important than their indirect effect on precautionary saving. This implies that the power of UI extensions to boost aggregate demand is diminished if households do not anticipate them. We show this result in partial equilibrium (using only the calibrated household block to compute the direct effect), as well as in general equilibrium (using the full estimated model).

7.1 The Direct Effect of UI Extensions on Aggregate Demand

UI extensions can boost aggregate demand by two channels. First, directly, by raising the income of unemployed households who get to keep their benefits thanks to the extension. Second, indirectly, by reducing the precautionary savings of employed households facing the risk of job loss, and of unemployed households facing the risk of losing benefits. Our first goal is to establish that the precautionary saving channel is quantitatively relevant.

To establish the importance of the precautionary savings channel, we perform a partial equilibrium exercise in which we feed the path of UI duration implied by the empirical responses to our calibrated household block and compute the resulting response of aggregate consumption to the UI extension. In doing so, we isolate the effect of imperfect anticipation of UI extensions on aggregate demand, abstracting from any general equilibrium effects that may arise in the full model. So, this exercise isolates the direct impact of UI extensions on aggregate demand, both via disposable incomes and via precautionary savings.

Figure 6 plots the path of UI duration in the left panel and the resulting response of aggregate

Figure 7: Decomposition of the Partial-Equilibrium Consumption Response

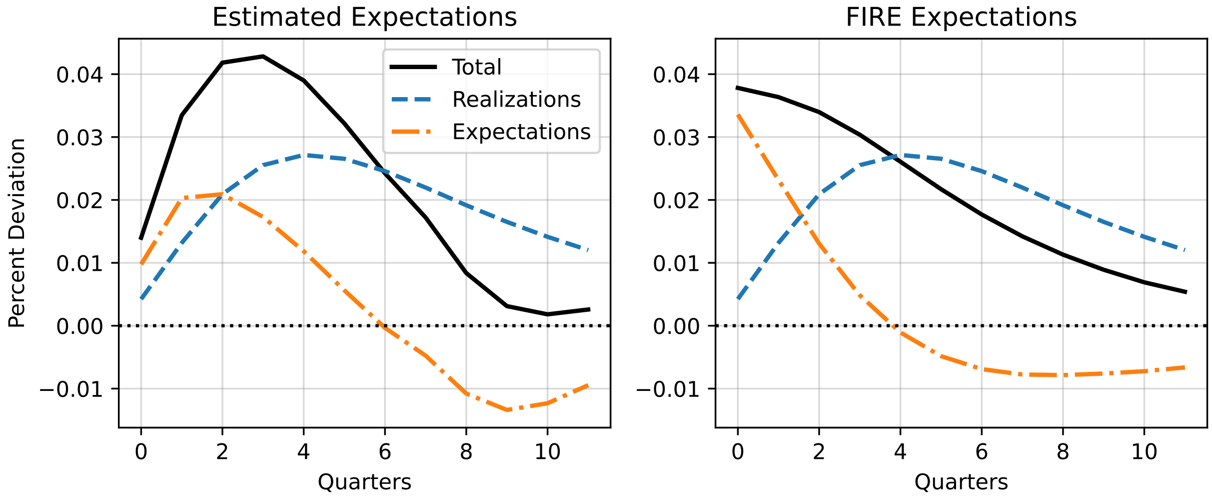

Notes: This figure decomposes the partial-equilibrium consumption response to a UI extension into the response to realized UI transfers and the additional response due to anticipation of future UI generosity. The decomposition is shown under estimated beliefs and under FIRE.

consumption to this UI extension for our estimated model (blue) and the FIRE model (dotted orange). The patterns of responses are markedly different across the two models. In the model with FIRE, people anticipate the UI extension from the outset, leading to a large immediate reduction in precautionary savings and a sharp increase in consumption. Instead, the response under the estimated beliefs is more muted on impact, but it becomes larger than the response under FIRE after a few quarters and, in particular, after the peak of the recession. The reason for this pattern is that, with the estimated beliefs, people do not fully anticipate the UI extension on impact, but they eventually overreact to it as their beliefs become more pessimistic than what would be justified by the actual path of UI duration.

To see these effects more clearly, we leverage our results in Proposition 2 to decompose the response of aggregate consumption into two terms. First, we consider the effect of realized UI extensions that are never anticipated, i.e., such that $E_t[d\text{UI duration}_{t+s}] = 0$ for all $s \geq 1$. This myopic scenario isolates the effect of the direct transfer of income to unemployed households, without any change in precautionary savings. Second, we compute the additional impact of expectations of UI extensions on consumption, by netting out the myopic response from the full response under the estimated beliefs. This second term captures the effect of anticipation of UI extensions on consumption through the precautionary savings channel.

The left panel in Figure 7 shows this decomposition of the response of aggregate consumption under the estimated beliefs. The right panel shows the same decomposition under FIRE. The effect of realizations is common across both models because it does not depend on expectations. The expectations effects, however, differ markedly. Full anticipation under FIRE leads to a large increase in consumption on impact that then falls rapidly. Thus, under FIRE, expectations generate a sharp and front-loaded consumption response.

In contrast, the empirically slow update of expectations implies that the initial response of consumption is much more muted and closer to that implied by the myopic scenario. However, as expectations update over time, this term grows in importance, leading to a backloaded and hump-shaped response of consumption. The overreaction of expectations in the later periods also implies that the consumption response eventually becomes larger than under FIRE.

In this exercise, UI duration is proportional to unemployment and therefore provides a natural date for the peak of the downturn. Figure 6 shows that duration peaks in quarter 2 and then begins to recede, whereas the consumption response under estimated beliefs peaks in quarter 3. Thus, within the partial-equilibrium exercise, the policy's demand effect is strongest one quarter after UI duration has peaked. The same policy under FIRE is strongest on impact. This comparison provides evidence of an expectations-driven efficacy lag. To better understand the implications of this lag for the effectiveness of UI extensions, we next turn to a general-equilibrium analysis.

7.2 Quantifying the Spending Effectiveness of UI Extensions in GE

The results in the previous section allow us to understand the importance of anticipating UI extensions in determining their efficacy in stimulating consumption. Furthermore, they allow us to compare the implications of the empirical patterns of beliefs relative to FIRE. However, those results do not allow us to quantify the power of UI extensions in general equilibrium. For that purpose, we need to compare the impulse responses obtained in the previous section with those obtained in a counterfactual economy assuming no extensions, i.e., $\zeta_b = 0$.

For the purposes of computing a counterfactual response, we must take into account how beliefs differ in this counterfactual economy relative to our baseline analysis. This means that we must specify a model for belief formation, and can no longer rely on the estimated beliefs. We assume that beliefs are determined by the noisy-information and long-memory diagnostic expectations model that we estimation to perform the extrapolation exercise in Section 5.2. This choice is in line with our discussion in that section.

We evaluate the efficacy of UI extensions by computing the cumulative consumption multipliers of the policy under the two belief specifications. The cumulative consumption multiplier is a standard measure of the effectiveness of fiscal stimulus, and it is defined as the present value of the additional consumption generated per dollar of additional UI spending (see, e.g., Ramey and Zubairy 2018). To define the cumulative multiplier, let ΔC_t and $\Delta \mathrm{UI}_t$ denote change in consumption and UI transfers at date t relative to the counterfactual economy without extensions. Then, the present-value cumulative multiplier is defined as

$$\mathcal{M}_t = \frac{\sum_{s=0}^{t}(1+r)^{-s}\,\Delta C_s}{\sum_{s=0}^{t}(1+r)^{-s}\,\Delta \mathrm{UI}_s}, \tag{38}$$

We report this object from impact through quarter 10, the horizon displayed in Figure 8. The left

Figure 8: Cumulative Consumption Multipliers of UI Extensions

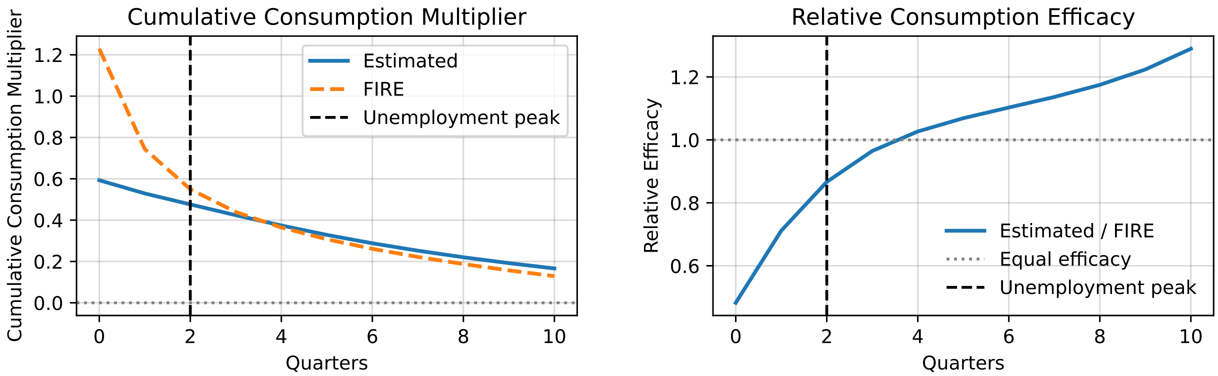

Notes: The left panel plots the cumulative general-equilibrium consumption multiplier of the UI extension, $\mathcal{M}_t$, under FIRE and estimated beliefs. The right panel plots the ratio of the estimated-beliefs multiplier to the FIRE multiplier. The multiplier is the discounted cumulative consumption response divided by the discounted cumulative change in UI transfers, each measured relative to the counterfactual economy without extensions. Quarter 0 is the impact quarter.

panel plots the cumulative multiplier under FIRE and estimated beliefs. The right panel plots the ratio of the estimated-beliefs multiplier to the FIRE multiplier.

The figure shows that spending effectiveness is significantly diminished on impact when households do not fully anticipate the extension. The estimated-beliefs multiplier is less than half the FIRE multiplier on impact and catches up in quarter 4, two quarters after unemployment peaks.

The spending effect depends sharply on how households form expectations. Under FIRE, households immediately anticipate the more generous UI and adjust their spending on impact, so the multiplier is largest at $t = 0$ (about 1.2) and declines thereafter as the extension is unwound and taxes rise to finance it. Under estimated beliefs, initial underreaction mutes this impact response: the multiplier on impact is more than halved (about 0.6). By quarter 4, the multiplier under estimated beliefs has caught up to the FIRE multiplier.

The cumulative consumption multiplier remains positive over the horizon shown. Its low impact value nevertheless implies that automaticity does not guarantee timely spending effectiveness. Estimated beliefs shift spending support away from impact, and the multiplier catches up to FIRE only after unemployment has peaked. This finding is consistent with the partial-equilibrium results and highlights the importance of expectations for the timing of household spending.

8 Conclusion

Economists have long emphasized the benefits of linking UI duration to aggregate economic conditions (see, e.g., Chodorow-Reich and Coglianese, 2019, Eichenbaum, 2019, Mitchell and Husak, 2021). The attraction of such a rule is immediate: it can respond to a downturn without waiting for policymakers to recognize the recession and choose a new instrument. Our results show that automaticity does not guarantee timely policy efficacy.

We study UI extensions in a heterogeneous-agent New Keynesian model with search and matching frictions and a household block disciplined by the micro incidence of unemployment and benefit expiration. We develop a framework that solves the model under arbitrary beliefs using initial forecasts and subsequent forecast revisions. This lets us use survey expectations directly and makes the timing of private adjustment an empirical object.

The measured beliefs imply initial underreaction followed by delayed overreaction. For UI, this pattern matters because the largest demand channel is the reduction in precautionary saving, which operates through expected future insurance. Measured by additional household spending per dollar of UI, the cumulative effect remains positive under estimated beliefs in the baseline, but its impact multiplier is roughly half the FIRE benchmark. In partial equilibrium, the consumption response becomes strongest one to two quarters after the recession has peaked. The lesson is close to Friedman's original warning: rules can largely solve policy-decision lags without solving the lag between policy action and private behavior. Expectations determine not only how much spending support UI provides, but whether that support arrives when it is most valuable.

References

Acosta, Miguel, Andreas I Mueller, Emi Nakamura, and Jón Steinsson, "Macroeconomic Effects of UI Extensions at Short and Long Durations," Technical Report, National Bureau of Economic Research 2023.

Adam, Klaus and Albert Marcet, "Internal Rationality, Imperfect Market Knowledge and Asset Prices," *Journal of Economic Theory*, 2011, *146* (3), 1224–1252.

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

Angeletos, George-Marios and Karthik A Sastry, "Managing Expectations: Instruments Versus Targets," *The Quarterly Journal of Economics*, 2021, *136* (4), 2467–2532.

_ **and Zhen Huo**, "Myopia and Anchoring," *American Economic Review*, 2021, *111* (4), 1166–1200.

_ **, Fabrice Collard, and Harris Dellas**, "Business-Cycle Anatomy," *American Economic Review*, 2020, *110* (10), 3030–70.

_ **, Zhen Huo, and Karthik A Sastry**, "Imperfect Macroeconomic Expectations: Evidence and Theory," *NBER Macroeconomics Annual*, 2021, *35* (1), 1–86.

Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub, "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models," *Econometrica*, 2021, *89* (5), 2375–2408.

— **, Matthew Rognlie, and Ludwig Straub**, "Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model," Working Paper 26647, National Bureau of Economic Research 2020.

— **, — , and —** , "Fiscal and Monetary Policy with Heterogeneous Agents," *Annual Review of Economics*, 2025, *17*, 539–562.

Bardóczy, Bence, "Spousal Insurance and the Amplification of Business Cycles," Job Market Paper, Northwestern University 2020.

Bianchi, Francesco, Cosmin L Ilut, and Hikaru Saijo, "Diagnostic Business Cycles," Working Paper 28604, National Bureau of Economic Research 2021.

Bianchi-Vimercati, Riccardo, Martin S. Eichenbaum, and Joao Guerreiro, "Fiscal Stimulus with Imperfect Expectations: Spending vs. Tax Policy," *Journal of Economic Theory*, 2024, *217*, 105814.

Bilbiie, Florin, Giorgio E Primiceri, and Andrea Tambalotti, "Inequality and Business Cycles," Manuscript 2022.

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

— **, — , Yueran Ma, and Andrei Shleifer**, "Overreaction in Macroeconomic Expectations," *American Economic Review*, 2020, *110* (9), 2748–82.

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

Carroll, Christopher D, "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence," *Brookings Papers on Economic Activity*, 1992, *1992* (2), 61–156.

— , "Macroeconomic Expectations of Households and Professional Forecasters," *Quarterly Journal of Economics*, 2003, *118* (1), 269–298.

— **, Edmund Crawley, Jiri Slacalek, Kiichi Tokuoka, and Matthew N White**, "Sticky Expectations and Consumption Dynamics," Working Paper 224377, National Bureau of Economic Research 2018.

Chodorow-Reich, Gabriel and John M Coglianese, "Unemployment Insurance and Macroeconomic Stabilization," *Unemployment Insurance and Macroeconomic Stabilization." In Recession Ready, ed. Heather Boushey, Ryan Nunn, and Jay Shambaugh*, 2019.

— **, Peter Ganong, and Jonathan Gruber**, "Should We Have Automatic Triggers for Unemployment Benefit Duration and How Costly Would They Be?," 2022. Manuscript.

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

—, **Martin S Eichenbaum, and Mathias Trabandt**, "Unemployment and Business Cycles," *Econometrica*, 2016, *84* (4), 1523–1569.

Coibion, Olivier and Yuriy Gorodnichenko, "What Can Survey Forecasts Tell us About Information Rigidities?," *Journal of Political Economy*, 2012, *120* (1), 116–159.

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

da Silva, António Dias, Desislava Rusinova, and Marco Weißler, "Consumption Effects of Job Loss Expectations: New Evidence for the Euro Area," *European Economic Review*, 2025, *179*, 105080.

Dobrew, Michael, Rafael Gerke, Sebastian Giesen, and Joost Röttger, "Make-Up Strategies with Incomplete Markets and Bounded Rationality," Bundesbank Discussion Paper 2023.

Eichenbaum, Martin S, "Rethinking Fiscal Policy in an Era of Low Interest Rates," *Economic Policy Group Monetary Authority of Singapore*, 2019, *90.*

Fagereng, Andreas, Helene Onshuus, and Kjersti Næss Torstensen, "The Consumption Expenditure Response to Unemployment: Evidence from Norwegian Households," *Journal of Monetary Economics*, 2024, *146*, 103578.

Farhi, Emmanuel and Iván Werning, "Monetary policy, Bounded Rationality, and Incomplete Markets," *American Economic Review*, 2019, *109* (11), 3887–3928.

—, **Mikel C. Petri, and Iván Werning**, "The Fiscal Multiplier Puzzle: Liquidity Traps, Bounded Rationality, and Incomplete Markets," Technical Report, mimeo 2020.

Fernandes, Adriano and Rodolfo Rigato, "Precautionary Savings and Stabilization Policy in a Present-Biased Economy," 2022. Job market paper.

Friedman, Milton, "A Monetary and Fiscal Framework for Economic Stability," *American Economic Review*, 1948, *38* (3), 245–264.

—, "A Theory of the Consumption Function," in "A Theory of the Consumption Function," Princeton University Press, 1957, pp. 1–6.

Gabaix, Xavier, "A Sparsity-Based Model of Bounded Rationality," *The Quarterly Journal of Economics*, 2014, *129* (4), 1661–1710.

—, "Behavioral Macroeconomics via Sparse Dynamic Programming," Technical Report, National Bureau of Economic Research 2016.

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

Ganong, Peter and Pascal Noel, "Consumer Spending During Unemployment: Positive and Normative Implications," *American Economic Review*, 2019, *109* (7), 2383–2424.

_ **, Fiona Greig, Pascal Noel, Daniel M. Sullivan, and Joseph S. Vavra**, "Spending and Job-Finding Impacts of Expanded Unemployment Benefits: Evidence from Administrative Micro Data," *American Economic Review*, 2024, *114* (9), 2898–2939.

Gertler, Mark, Christopher Huckfeldt, and Antonella Trigari, "Unemployment Fluctuations, Match Quality, and the Wage Cyclicality of New Hires," *The Review of Economic Studies*, 07 2020, *87* (4), 1876–1914.

Guerreiro, Joao, "Belief Disagreement and Business Cycles," Job Market Paper, Northwestern University 2022.

_ **, Jonathon Hazell, Diego R. Kanzig, and Ed Manuel**, "The Macroeconomic Effect of Stimulus Checks: Evidence from Postwar Veterans' Transfers," Working Paper 2026.

Hartmann, Ida Maria and Søren Leth-Petersen, "Subjective Unemployment Expectations and (Self-)Insurance," *Labour Economics*, 2024, *90*, 102579.

Hazell, Jonathon and Bledi Taska, "Downward Rigidity in the Wage for New Hires," *American Economic Review*, December 2025, *115* (12), 4183–4217.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L Violante, "Optimal Tax Progressivity: An Analytical Framework," *The Quarterly Journal of Economics*, 2017, *132* (4), 1693–1754.

Juelsrud, Ragnar E. and Ella Getz Wold, "The Importance of Unemployment Risk for Individual Savings," *Management Science*, 2025, *71* (8), 6746–6769.

Justiniano, Alejandro, Giorgio E Primiceri, and Andrea Tambalotti, "Investment Shocks and Business Cycles," *Journal of Monetary Economics*, 2010, *57* (2), 132–145.

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

_ **, Giovanni L Violante, and Justin Weidner**, "The Wealthy Hand-to-Mouth," *Brookings Papers on Economic Activity*, 2014, *2014* (1), 77–138.

Kekre, Rohan, "Unemployment Insurance in Macroeconomic Stabilization," Working Paper 29505, National Bureau of Economic Research 2021.

Kimball, Miles S., "Precautionary Saving in the Small and in the Large," *Econometrica*, 1990, *58* (1), 53–73.

Krueger, Dirk, Kurt Mitman, and Fabrizio Perri, "Macroeconomics and Household Heterogeneity," in "Handbook of Macroeconomics," Vol. 2, Elsevier, 2016, pp. 843–921.

Krusell, Per, Toshihiko Mukoyama, and Ayşegül Şahin, "Labour-Market Matching with Precautionary Savings and Aggregate Fluctuations," *Review of Economic Studies*, 2010, *77* (4), 1477–1507.

Lian, Chen, "Mistakes in Future Consumption, High MPCs Now," *American Economic Review: Insights*, 2023, *5* (4), 563–581.

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

McKay, Alisdair and Ricardo Reis, "The Role of Automatic Stabilizers in the US Business Cycle," *Econometrica*, 2016, *84* (1), 141–194.

Mitchell, David and Corey Husak, "How to Replace COVID Relief Deadlines with Automatic 'Triggers' that Meet the Needs of the US Economy," *The Washington Center for Equitable Growth*, 2021.

Mitman, Kurt and Stanislav Rabinovich, "Optimal Unemployment Insurance in an Equilibrium Business-Cycle Model," *Journal of Monetary Economics*, 2015, *71*, 99–118.

__ **and** __ , "Do Unemployment Benefit Extensions Explain the Emergence of Jobless Recoveries?," CEPR Discussion Paper No. DP13760 2019.

Morales-Jiménez, Camilo and Luminita Stevens, "Price Rigidities in U.S. Business Cycles," Unpublished Working Paper March 2025.

Nakajima, Makoto, "A Quantitative Analysis of Unemployment Benefit Extensions," *Journal of Monetary Economics*, 2012, *59* (7), 686–702.

Pappa, Evi, Morten O. Ravn, and Vincent Sterk, "Expectations and Incomplete Markets," in "Handbook of Economic Expectations," Elsevier, 2023, pp. 569–611.

Pfäuti, Oliver and Fabian Seyrich, "A Behavioral Heterogeneous Agent New Keynesian Model," Discussion Paper, DIW Berlin 2022.

Ramey, Valerie A and Sarah Zubairy, "Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data," *Journal of political economy*, 2018, *126* (2), 850–901.

Ravn, Morten O and Vincent Sterk, "Job Uncertainty and Deep Recessions," *Journal of Monetary Economics*, 2017, *90*, 125–141.

Rotemberg, Julio J, "Sticky Prices in the United States," *Journal of Political Economy*, 1982, *90* (6), 1187–1211.

Sims, Christopher A, "Implications of Rational Inattention," *Journal of Monetary Economics*, 2003, *50* (3), 665–690.

Werning, Iván, "Incomplete Markets and Aggregate Demand," Working Paper 21448, National Bureau of Economic Research 2015.

Woodford, Michael, "Imperfect Common Knowledge and the Effects of Monetary Policy," Working Paper 8673, National Bureau of Economic Research December 2001.

_ , "Monetary Policy Analysis when Planning Horizons are Finite," Working Paper 24692, National Bureau of Economic Research June 2018.

A Appendix to Section 2

A.1 Deriving the equilibrium

We solve the model backwards, starting at the end of period 1. At this point, households observe the relevant equilibrium outcomes $\{b_1, \tau_1\}$ as well as their employment status $e_1 \in \{E, U\}$. Their problem is

$$\begin{aligned} V_1(e_1, a_0) = \max_{c_1, a_1} u(c_1) \qquad \text{s.t. } & c_1 + a_1 = (1+r)a_0 + \mathbf{1}_{\{e_1=E\}} + \mathbf{1}_{\{e_1=U\}} \cdot b_1 - \tau_1 \\ & a_1 \geq 0 \end{aligned} \tag{A.1}$$

where we already imposed $w_1 = 1$. Clearly, the optimal decision is to consume the entire cash on hand

$$a_1(e_1, a_0) = 0 \tag{A.2}$$

$$c_1(e_1, a_0) = (1+r)a_0 + \mathbf{1}_{\{e_1=E\}} + \mathbf{1}_{\{e_1=U\}} \cdot b_1 - \tau_1 \tag{A.3}$$

Since assets are in zero net supply and borrowing is not allowed, all workers have zero assets in equilibrium, $a_0 = 0$. Given the policies $\{b_1, M_1\}$, the time-1 equilibrium can be computed recursively as

$$N_1 = M_1 \tag{A.4}$$

$$\tau_1 = (1 - N_1)b_1 \tag{A.5}$$

$$c_1(E) = 1 - \tau_1 \tag{A.6}$$

$$c_1(U) = b_1 - \tau_1 \tag{A.7}$$

Note that time-1 equilibrium is independent of what happens in period 0, including the beliefs $\{E[N_1], E[b_1], E[\tau_1]\}$ that households hold in period 0.

Now turn to period 0. Combining the consumption policy function (A.3) with the fact that the probability of employment is iid, the expected continuation value at the end of period 0 can be written as

$$V_1^e(a_0) = E[N_1] \cdot u\Big((1+r)a_0 + 1 - E[\tau_1]\Big) + (1 - E[N_1]) \cdot u\Big((1+r)a_0 + E[b_1] - E[\tau_1]\Big) \tag{A.8}$$

At time $t = 0$, households solve

$$\begin{aligned} \max_{c_0, a_0} u(c_0) + \beta V_1^e(a_0) \qquad \text{s.t. } & c_0 + a_0 = \mathbf{1}_{\{e_0=E\}} + \mathbf{1}_{\{e_0=U\}} \cdot b_0 - \tau_0 \\ & a_0 \geq 0 \end{aligned} \tag{A.9}$$

Taking FOCs yields the Euler equation (2) in the main text

$$u'(c_0) \geq \beta (V_1^e)'(a_0) \tag{A.10}$$

$$= \beta(1+r)\Big[E[N_1] \cdot u'\Big(1 - E[\tau_1] + (1+r)a_0\Big) + (1 - E[N_1]) \cdot u'\Big(E[b_1] - E[\tau_1] + (1+r)a_0\Big)\Big] \tag{A.11}$$

Note that the right-hand side does not depend on time-0 employment status. Since there cannot be any saving in equilibrium ($a_0 = 0$), households consume their income in period 0, which is higher for employed households. Given that $u'(\bullet)$ is decreasing, this implies that either both types of households are borrowing constrained or only unemployed workers are constrained. We assume that β and r are such that only unemployed households are constrained.

Then, given beliefs $\{E[N_1], E[\tau_1], E[b_1]\}$ and policies $\{b_0, r\}$, the time-0 equilibrium can be computed as follows

$$u'\Big(c_0(E)\Big) = \beta(1+r)\Big[E[N_1] \cdot u'\Big(1 - E[\tau_1]\Big) + (1 - E[N_1]) \cdot u'\Big(E[b_1] - E[\tau_1]\Big)\Big] \tag{A.12}$$

$$\tau_0 = 1 - c_0(E) \tag{A.13}$$

$$c_0(U) = b_0 - \tau_0 \tag{A.14}$$

$$N_0 = 1 - \frac{\tau_0}{b_0} \tag{A.15}$$

$$M_0 = N_0 \tag{A.16}$$

A.2 Proof of proposition 1

Proof. Consider an infinitesimal period-1 perturbation (dY_1, db_1) with $dY_1 \geq 0$ and $db_1 > 0$. Differentiating the period-1 equilibrium from (A.4) to (A.7) yields

$$dN_1 = dY_1 = dM_1 \tag{A.17}$$

$$d\tau_1 = (1 - Y_1)db_1 - dY_1 \cdot b_1 \tag{A.18}$$

$$dc_1(E) = -d\tau_1 = dY_1 \cdot b_1 - (1 - Y_1)db_1 \tag{A.19}$$

$$dc_1(U) = db_1 - d\tau_1 = dY_1 \cdot b_1 + Y_1 \cdot db_1 \tag{A.20}$$

Given the belief rule in equation (4), the perceived deviations of period-1 output and UI generosity are

$$dE[Y_1] = \lambda dY_1 \qquad \text{and} \qquad dE[b_1] = \lambda db_1. \tag{A.21}$$

Linearizing the government budget constraint and period-1 policies gives

$$dE[c_1(E)] = b_1 dE[Y_1] - (1 - Y_1)dE[b_1], \tag{A.22}$$

$$dE[c_1(U)] = b_1 dE[Y_1] + Y_1 dE[b_1]. \tag{A.23}$$

From equation (A.12), together with $c_0(E) = 1 - \tau_0 = 1 - (1 - Y_0)b_0$ and $db_0 = 0$, we can write

$$\begin{aligned}
dc_0(E) &= -d\tau_0 = dY_0 \cdot b_0 \\
u''(c_0(E)) \cdot dc_0(E) &= \beta(1+r)\Big[dE[Y_1] \cdot (u'(c_1(E)) - u'(c_1(U))) + Y_1 \cdot u''(c_1(E))dE[c_1(E)] \\
&\quad + (1 - Y_1) \cdot u''(c_1(U))dE[c_1(U)]\Big]
\end{aligned}$$

Collecting terms, define

$$\mathrm{MPC}_0 \equiv 1 - b_0, \tag{A.24}$$

$$M_{1,b} \equiv \beta(1+r)Y_1(1-Y_1)\left[-\frac{u''(c_1(E))}{u''(c_0(E))} + \frac{u''(c_1(U))}{u''(c_0(E))}\right], \tag{A.25}$$

$$\mathrm{MPC}_1 \equiv \beta(1+r)\left[\frac{u'(c_1(E)) - u'(c_1(U))}{u''(c_0(E))} + Y_1 b_1 \frac{u''(c_1(E))}{u''(c_0(E))} + (1-Y_1)b_1 \frac{u''(c_1(U))}{u''(c_0(E))}\right]. \tag{A.26}$$

Concavity implies $\mathrm{MPC}_1 > 0$, and prudence together with $c_1(E) > c_1(U)$ implies $M_{1,b} > 0$. These definitions yield

$$dY_0 = \frac{1}{1 - \mathrm{MPC}_0}\left[\mathrm{MPC}_1\, dE[Y_1] + M_{1,b} dE[b_1]\right]. \tag{A.27}$$

Using $dE[Y_1] = dY_1 - (1-\lambda)dY_1$ and $dE[b_1] = db_1 - (1-\lambda)db_1$, we obtain

$$dY_0 = dY_0^{\mathrm{FIRE}} - \frac{1-\lambda}{1 - \mathrm{MPC}_0}\left[\mathrm{MPC}_1\, dY_1 + M_{1,b} db_1\right], \tag{A.28}$$

where

$$dY_0^{\mathrm{FIRE}} = \frac{1}{1 - \mathrm{MPC}_0}\left[\mathrm{MPC}_1\, dY_1 + M_{1,b} db_1\right]. \tag{A.29}$$

□

B Appendix to Section 3

B.1 Financial intermediary

The quantitative model uses the reduced-form intermediary block stated in equations (9) to (12). Its balance sheet holds equity and long-term government bonds against household deposits and net worth. With equity supply normalized to $v_t^E = 1$ and reserves in zero net supply, the implemented net-worth and payout equations are

$$N_t^{FI} = d_t + p_t + \frac{1 + \delta_B q_t^B}{1 + \pi_t}\frac{B_{t-1}}{P_{t-1}} - (1 + r_{t-1}^a)A_{t-1} - d_t^{FI}, \tag{B.1}$$

$$d_t^{FI} = d_{ss}^{FI} + \phi_N\left(N_t^{FI} - N_{ss}^{FI}\right). \tag{B.2}$$

Thus payouts reduce current net worth and respond in levels to contemporaneous net worth, exactly as in the quantitative implementation. The balance-sheet identity, these two equations, and the no-arbitrage conditions jointly close the asset markets; the model imposes no additional intermediary optimization problem beyond these conditions.

B.2 Steady-state calibration details

Table 6 records the complete steady-state calibration used for the reported v7 results. Parameters labeled "calibrated" are chosen jointly in the square five-moment system described in Section 4. Parameters labeled "assigned" are fixed directly before solving the steady state. The remaining entries are equilibrium objects implied by those choices and the stated normalizations.

Table 6: Complete Steady-State Calibration

Parameter	Value, unit, or role	Source, target, or derivation
Households and unemployment risk		
σ	0.500 (EIS)	Assigned CRRA preference parameter.
$\underline{a}$	0 (asset units)	Assigned no-borrowing constraint.
$(\mathcal{G}_z, \Pi_z)$	Data files (quarterly process)	Discrete-time version of the productivity process estimated by Kaplan et al. (2018); mean productivity is normalized to one.
λ	0.181 (tax progressivity)	Assigned from Heathcote et al. (2017).
β_2	0.9223 (quarterly)	Calibrated; primarily identified by the mean quarterly MPC.
$\beta_3 - \beta_2$	0.0504 (discount-factor spread)	Calibrated jointly; primarily identified by the unemployed-minus-employed annual MPC difference. The three types are equally weighted and symmetric around β_2.
s	0.0703 (quarterly hazard scale)	Calibrated; the resulting average employed-worker separation probability is 0.0789 and matches 5% unemployment given $f = 0.6$.
Δ_β	-12.5078 (hazard semi-elasticity)	Calibrated in $s_i = s\exp[\Delta_\beta(\beta_i - \beta_2)]$; identified by the wealth gradient of annual job-loss risk.
Δ_z	0 (hazard semi-elasticity)	Assigned. The baseline excludes a separate persistent-productivity gradient in job-loss risk.
π^{get}	0.520 (probability)	Calibrated to the 0.39 UI-recipient share among unemployed workers.
π^{lose}	0.500 (quarterly probability)	Assigned to an expected baseline eligibility duration of two quarters.
b_1	0.760 (pre-displacement income share)	Assigned directly to the income ratio while receiving UI in Table 1.
b_2	0.550 (pre-displacement income share)	Assigned directly to the income ratio after UI exhaustion in Table 1.
Asset grid	$[0, 2500]$, 100 points	Fixed implementation grid; the upper endpoint is nonbinding in the calibrated steady state.
Labor market		
f	0.600 (quarterly probability)	Assigned job-finding rate.

Continued on next page

Table 6 continued

Parameter	Value, unit, or role	Source, target, or derivation
q	0.700 (quarterly probability)	Assigned vacancy-filling rate.
ℓ	0.500 (matching elasticity)	Assigned Cobb-Douglas matching elasticity.
$\kappa_v/(\kappa_v + q\kappa_h)$	0.060 (cost share)	Assigned vacancy-posting share of total hiring cost.
Hiring cost	0.070 (quarterly wage share)	Assigned total hiring cost per filled job; the steady-state job-creation condition implies $\kappa_v = 0.00193$ and $\kappa_h = 0.04321$.
A_m	0.6481 (matching efficiency)	Implied by $f = A_m\theta^{1-\ell}$ and $\theta = f/q$.
Aggregate production, policy, and assets		
r	0.005 (quarterly real rate)	Assigned to 2% at an annual rate. The intermediary spread is $\xi = 0$, so $r^a = r$.
ϵ	7.000 (demand elasticity)	Assigned retailer elasticity, implying steady-state marginal cost $6/7$.
Y	1.000 (output units)	Normalization.
G/Y	0.160 (output share)	Assigned author calibration.
$q^B B/(4Y)$	0.460 (annual-output share)	Assigned government-debt target.
Bond duration	20 quarters	Assigned five-year duration target; implies coupon-decay parameter $\delta_B = 0.95475$.
δ_K	0.02075 (quarterly rate)	Assigned to 8.3% at an annual rate.
K/Y	8.920 quarters	Assigned capital-output target.
α	0.2680 (capital elasticity)	Implied by the capital-output target and steady-state user cost; the resulting labor share is approximately 62%.
Total wealth$/(4Y)$	3.820 (annual-output multiple)	Assigned aggregate balance-sheet target.
Θ	0.5776 (TFP units)	Implied by the normalization $Y = 1$.
Ξ	0.1205 (output units)	Implied fixed cost that reconciles profits with the aggregate wealth target.
τ	0.2877 (tax-rate level)	Implied by the steady-state government budget.

B.3 Retailers

For the price-setting problem, define indexed inflation for retailer j as

$$\widetilde{\pi}_{jt} \equiv \log\left(\frac{p_{jt}}{p_{jt-1}}\right) - \iota_p \log\left(\frac{p_{jt-1}}{p_{jt-2}}\right). \tag{B.3}$$

The retailer's problem is

$$J_t^R(p_{jt-1}, p_{jt-2}) = \max_{k_{jt}, l_{jt}, y_{jt}, p_{jt}} \left\{ \frac{p_{jt}}{P_t} y_{jt} - r_t^k k_{jt} - r_t^l l_{jt} - \frac{\epsilon}{2\kappa_p} \widetilde{\pi}_{jt}^2 Y_t - \Xi \right. \tag{B.4}$$

$$\left. + \mathbb{E}_t \left[\frac{J_{t+1}^R(p_{jt}, p_{jt-1})}{1 + r_t} \right] \right\}, \tag{B.5}$$

subject to

$$y_{jt} = \Theta_t k_{jt}^{\alpha} l_{jt}^{1-\alpha} = \left(\frac{p_{jt}}{P_t}\right)^{-\epsilon} Y_t. \tag{B.6}$$

Let mc_t denote real marginal cost. The static input conditions and symmetry imply

$$r_t^l = mc_t(1-\alpha)\frac{Y_t}{L_t}, \tag{B.7}$$

$$r_t^k = mc_t\alpha\frac{Y_t}{K_{t-1}}, \tag{B.8}$$

$$Y_t = \Theta_t K_{t-1}^{\alpha} L_t^{1-\alpha}. \tag{B.9}$$

Linearizing the price-setting first-order condition around the zero-inflation steady state, where $\overline{mc} = (\epsilon - 1)/\epsilon$, gives

$$\widetilde{\pi}_t = \kappa_p(mc_t - \overline{mc}) + \frac{1}{1+r_t}\mathbb{E}_t\left[\frac{Y_{t+1}}{Y_t}\widetilde{\pi}_{t+1}\right], \qquad \widetilde{\pi}_t \equiv \pi_t - \iota_p\pi_{t-1}. \tag{B.10}$$

This is the Phillips curve used in the quantitative implementation. Aggregate price-adjustment costs and retailer dividends are

$$\Psi_t^p = \frac{\epsilon}{2\kappa_p}\widetilde{\pi}_t^2 Y_t, \tag{B.11}$$

$$d_t^R = Y_t - r_t^l L_t - r_t^k K_{t-1} - \Psi_t^p - \Xi. \tag{B.12}$$

B.4 Capital producer

The capital producer solves

$$\begin{aligned} J_t^K(K_{t-1}, I_{t-1}) = \max_{K_t, I_t} &\left\{ r_t^k K_{t-1} - I_t + \mathbb{E}_t\left[\frac{J_{t+1}^K(K_t, I_t)}{1+r_t}\right]\right\} \\ \text{s.t.} \qquad & K_t = (1-\delta)K_{t-1} + \mu_t\left[1 - S\left(\frac{I_t}{I_{t-1}}\right)\right] I_t. \end{aligned} \tag{B.13}$$

Let Q_t be the multiplier on the capital law of motion and define $x_t \equiv I_t/I_{t-1}$. The first-order conditions are

$$Q_t = \mathbb{E}_t\left[\frac{r_{t+1}^k + (1-\delta)Q_{t+1}}{1+r_t}\right], \tag{B.14}$$

$$1 = Q_t\mu_t\left[1 - S(x_t) - x_t S'(x_t)\right] + \mathbb{E}_t\left[\frac{\mu_{t+1}Q_{t+1}}{1+r_t}x_{t+1}^2 S'(x_{t+1})\right]. \tag{B.15}$$

Capital-producer dividends are

$$d_t^K = r_t^k K_{t-1} - I_t. \tag{B.16}$$

For concreteness, the adjustment-cost function is quadratic:

$$S\left(\frac{I_t}{I_{t-1}}\right) = \frac{\psi}{2}\left(\frac{I_t}{I_{t-1}} - 1\right)^2 \tag{B.17}$$

$$S'\left(\frac{I_t}{I_{t-1}}\right) = \psi\left(\frac{I_t}{I_{t-1}} - 1\right) \tag{B.18}$$

B.5 Labor agency

The labor agency solves

$$J_t^L(N_{t-1}) = \max_{N_t, v_t}\left\{(r_t^l - w_t)Z_t N_t - (\kappa_v + \kappa_h q_t)v_t + \mathbb{E}_t\left[\frac{J_{t+1}^L(N_t)}{1+r_t}\right]\right\} \tag{B.19}$$
$$\text{s.t.} \qquad N_t = (1-s_t)N_{t-1} + q_t v_t.$$

The value of a filled job, J_t, equals its recruiting cost,

$$J_t = \frac{\kappa_v}{q_t} + \kappa_h, \tag{B.20}$$

and satisfies the job-creation condition

$$J_t = (r_t^l - w_t)Z_t + \mathbb{E}_t\left[\frac{1-s_{t+1}}{1+r_t}J_{t+1}\right]. \tag{B.21}$$

Labor-agency dividends are

$$d_t^L = (r_t^l - w_t)Z_t N_t - (\kappa_v + \kappa_h q_t)v_t. \tag{B.22}$$

B.6 Household expectations

The aggregate variables that households must forecast to make their consumption-saving decision are $\{f_t, \pi_t^{lose}, w_t, r_t^a, \tau_t\}$. We have expectations data on the unemployment rate, which is closely related to the job-finding rate f_t and the UI expiration probability π_t^{lose}. We construct expectations for $\{f_t, \pi_t^{lose}\}$ by assuming that households have first-order knowledge of the structural relationships between these variables and the unemployment rate.

In matrix notation, we look for Jacobians $\mathcal{J}^{f,U}$ and $\mathcal{J}^{\pi^{lose},U}$ such that

$$d\mathbf{f}^{e,t} = \mathcal{J}^{f,U} d\mathbf{U}^{e,t} \qquad \text{and} \qquad d(\pi_t^{lose})^{e,t} = \mathcal{J}^{\pi^{lose},U} d\mathbf{U}^{e,t} \tag{B.23}$$

First, the relationship between the unemployment rate and UI expiration rate is given by the policy rule

$$\frac{1}{\pi_t^{lose}} - \frac{1}{\pi_{ss}^{lose}} = \zeta_b(U_t - U_{ss}) \tag{B.24}$$

Differentiating the policy rule at the steady state yields

$$d\pi_t^{lose} = -\left(\pi_{ss}^{lose}\right)^2 \zeta_b dU_t \tag{B.25}$$

which shows that $\mathcal{J}^{\pi^{lose},U}$ is a diagonal matrix.

Second, the relationship between the employment rate and job-finding rate is given by the law of motion

$$\begin{aligned} U_t &= (1-f_t)U_{t-1} + s(1-f_t)(1-U_{t-1}) & \text{(B.26)} \\ &= (1-f_t)\left[s + (1-s)U_{t-1}\right]. & \text{(B.27)} \end{aligned}$$

Differentiating this law of motion yields

$$df_t = \frac{(1-f)(1-s)dU_{t-1} - dU_t}{s + (1-s)U} \tag{B.28}$$

which shows that $\mathcal{J}^{f,U}$ is a lower bidiagonal matrix.

B.7 Household aggregation and equilibrium closure

Let Γ_t denote the post-transition joint distribution of household states (a, z, e, β). Household policies and this distribution imply

$$C_t = \int c_t(a,z,e,\beta)\, d\Gamma_t, \qquad A_t = \int a_t(a,z,e,\beta)\, d\Gamma_t, \tag{B.29}$$

$$N_t = \int \mathbf{1}\{e=E\}\, d\Gamma_t, \qquad L_t = \int \mathbf{1}\{e=E\} z\, d\Gamma_t, \tag{B.30}$$

$$UI_t = \int \mathbf{1}\{e=U\} b_1 w_t z\, d\Gamma_t. \tag{B.31}$$

The law of motion for Γ_t combines the productivity process with the employment-transition matrix in Section 3 and the endogenous asset policy. Factor markets clear at $L_t = \int l_{jt} dj$ and $K_{t-1} = \int k_{jt} dj$. The intermediary holds all firm equity, nominal reserves are in zero net supply, and household deposits equal aggregate liquid assets. Finally, the government budget constraint and the goods-market condition in Section 3 close the fiscal and resource sides of the model. These aggregation equations make explicit the distinction between raw employment N_t and effective labor input L_t.

C Appendix to Section 5.1

C.1 Additional Results

Proposition 3 handles the special case of non-rational but time-invariant expectations $d\mathbf{X}^e \neq d\mathbf{X}$. An example of this is level-k thinking. The total response $d\mathbf{Y}$ is the sum of two effects. First, the response to the expected part of the shock. Second, the responses to the forecast errors that the agent observes along the way. The key new object is the forecast-error Jacobian, $\mathcal{E}$, that captures the second effect. Column s of $\mathcal{E}$ can be interpreted as the impulse response to the forecast error in dX_s which the agent learns in period s. Constructing $\mathcal{E}$ is straightforward. It is a lower diagonal matrix whose columns are shifted versions of the first column of $\mathcal{J}$. The intuition is that observing a forecast error in period t is equivalent to observing an unexpected shock in period 0. The formal proof is in appendix C.2.

Proposition 3. *Assuming beliefs are constant over time, $X^{e,t}_{t+h} = X^{e}_{t+h}$ for all $t, h > 0$, the linearized impulse response $d\mathbf{Y}$ to arbitrary shocks $d\mathbf{X}$ and $d\mathbf{X}^e$ is given by*

$$d\mathbf{Y} = \mathcal{J} \underbrace{d\mathbf{X}^e}_{\textit{forecast}} + \mathcal{E} \underbrace{(d\mathbf{X} - d\mathbf{X}^e)}_{\textit{forecast error}} \tag{C.1}$$

*where the **forecast-error Jacobian** $\mathcal{E}$ is given by*

$$\mathcal{E} = \begin{bmatrix} \mathcal{J}_{0,0} & 0 & \dots & 0 \\ \mathcal{J}_{1,0} & \mathcal{J}_{0,0} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \mathcal{J}_{t,0} & \mathcal{J}_{t-1,0} & \dots & \mathcal{J}_{0,0} \end{bmatrix} \tag{C.2}$$

C.2 Proof of Proposition 3

For convenience, we restate the proposition.

Proposition. *Assuming beliefs are constant over time, $X^{e,t}_{t+h} = X^{e}_{t+h}$ for all $t, h > 0$, the linearized impulse response $d\mathbf{Y}$ to arbitrary shocks $d\mathbf{X}$ and $d\mathbf{X}^e$ is given by*

$$d\mathbf{Y} = \mathcal{J} \underbrace{d\mathbf{X}^e}_{\textit{forecast}} + \mathcal{E} \underbrace{(d\mathbf{X} - d\mathbf{X}^e)}_{\textit{forecast error}} \tag{C.3}$$

*where the **forecast-error Jacobian** $\mathcal{E}$ is given by*

$$\mathcal{E} = \begin{bmatrix} \mathcal{J}_{0,0} & 0 & \dots & 0 \\ \mathcal{J}_{1,0} & \mathcal{J}_{0,0} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \mathcal{J}_{t,0} & \mathcal{J}_{t-1,0} & \dots & \mathcal{J}_{0,0} \end{bmatrix} \tag{C.4}$$

We consider a generic representation of a heterogeneous-agent problem. Following Auclert et al. (2021), we can write this representation as a mapping from aggregate inputs $\mathbf{X}$ to a time path for aggregate outputs $\mathbf{Y}$. To allow for generic beliefs, we extend their framework by allowing X^e_t to denote the expectations that individuals hold about the variable X_t. For now, we assume that beliefs are constant over time. As we will see below, the behavior of individuals is determined by their expectations $\mathbf{X}^e$ and the realizations $\mathbf{X}$. We assume that the distribution is discretized on n_g points, i.e., D_t is the $n_g \times 1$ matrix summarizing the distribution of agents. Let $y_t = y\left(v^e_{t+1}, X_t\right)$ be the $n_g \times 1$ matrix of individual outcomes, given the expected continuation value v^e_{t+1} and the current value for the input X_t. The heterogeneous-agent problem can be written as follows:

$$v_t = v\left(v^e_{t+1}, X_t\right) \tag{C.5}$$

$$v^e_{t+1} = v\left(v^e_{t+2}, X^e_{t+1}\right) \tag{C.6}$$

$$D_{t+1} = \Lambda\left(v^e_{t+1}, X_t\right)' D_t \tag{C.7}$$

$$Y_t = y\left(v^e_{t+1}, X_t\right)' D_t \tag{C.8}$$

for $t \geq 0$.

We describe the first-order response of Y_t to shocks to inputs X_t and expectations X^e_t. In doing so, we consider the behavior of the heterogeneous-agent block around the steady state with correct beliefs. We write (Y, v, v^e, D) as the steady state for their respective variables when expectations are correct $X^e_t = X_t = X$. This immediately implies that $v = v^e$. For notational convenience, let $\Lambda \equiv \Lambda\left(v^e, X\right)$ denote the steady-state transition matrix. In what follows, we consider transitions of time length T that satisfy that the inputs and expectations converge back to the original steady state after T periods, i.e., $X^e_t = X_t = X$ for $t \geq T$, which also implies that $v^e_t = v_t = v$ for $t \geq T$, i.e., we assume that both the input and expectations converge back to the original steady state. We assume that the block starts from its steady-state distribution $D_0 = D$.

This heterogeneous-agent problem defines a $T \times 1$ vector of stacked outputs

$$\mathbf{Y} = h\left(\mathbf{X}, \mathbf{X}^e\right),$$

where $\mathbf{Y} = \begin{bmatrix} Y_0 & Y_1 & \ldots \end{bmatrix}'$, $\mathbf{X} = \begin{bmatrix} X_0 & X_1 & \ldots \end{bmatrix}'$, and $\mathbf{X}^e = \begin{bmatrix} X^e_0 & X^e_1 & \ldots \end{bmatrix}'$. If all functions are differentiable in $\mathbf{X}$ and $\mathbf{X}^e$, then h is also differentiable. Our goal is to characterize the Jacobian $\mathcal{J}$ of the function h with respect to X and X^e, evaluated at the steady state. To do this, we start at the steady state and perturb, in turn, the time-s input by dX_s or its time-s expectation by dX^e_s. This procedure allows us to disentangle the consequences of changes in expectations from changes in realizations.

Response to input dX_s Consider a change to input X at time s, dX_s, with $dX_t = 0$ for all $t \neq s$ and $dX^e_{t+1} = 0$ for all t. Since expectations do not change,

$$v^e_{t+1} = v^e = v,$$

for all t. Furthermore, it follows that, for all $t \neq s$, $y_t \equiv y\left(v^e_{t+1}, X_t\right) = y\left(v, X\right) = y$ and $\Lambda_t \equiv \Lambda\left(v^e_{t+1}, X_t\right) = \Lambda\left(v, X\right) = \Lambda$, so $dy_t = 0$ and $d\Lambda_t = 0$.

Using the chain rule, we can decompose the change in output at time t by the change in choices

y_t and the change in the distribution D_t:

$$dY_t = dy_t' D + y dD_t$$

and

$$dD_{t+1} = d\Lambda_t' D + \Lambda' dD_t.$$

Using these expressions and the results above, it follows that, prior to date s, the output and distribution remain unchanged, i.e., $dY_t = 0$ and $dD_{t+1} = 0$ for all $t < s$. Furthermore, for $t = s$, we find that $dD_s = 0$, so

$$dD_{s+1} = d\Lambda_s' D, \quad \text{and} \quad dY_s = dy_s' D,$$

After date s, there is no other change in behavior, so we only need to account for the change in the distribution, i.e., for $t > s$ we find that

$$dY_t = y dD_t$$

where

$$dD_t = \left(\Lambda'\right)^{t-(s+1)} dD_{s+1}.$$

Finally, note that dy_s and $d\Lambda_s$ do not depend on the date s. It follows that

$$\frac{\partial Y_t}{\partial X_s} = \begin{cases} 0 & \text{if } t < s \\ \frac{\partial Y_{t-s}}{\partial X_0}, & \text{if } t \geq s \end{cases} = \begin{cases} 0 & \text{if } t < s \\ \mathcal{J}_{t-s,0}, & \text{if } t \geq s, \end{cases} \tag{C.9}$$

where $\mathcal{J}$ denotes the FIRE Jacobian derived in Auclert et al. (2021).

Response to expectation dX_s^e We now consider the response to a shock in expectations that is not realized, i.e., $dX_s^e \neq 0$ but $dX_s = 0$. Also $dX_t = dX_t^e = 0$ for $t \neq s$.

Note that at and after date s the value functions, transition matrices, and decisions return to their steady state values, i.e., $v_t^e = v_t = v$, $\Lambda_t = \Lambda$ and $y_t = y$, for all $t \geq s$. It follows that, for $t > s$,

$$dY_t = y dD_t$$

and

$$dD_{t+1} = \Lambda' dD_t$$

where $dD_t = \Lambda' dD_{t-1} = (\Lambda')^{t+1-(s+1)} dD_s$ for $t \geq s+1$.

From the perspective of $t < s$, individuals believe that the shock will realize. So, their response is the same as in FIRE, i.e., $dY_t = \mathcal{J}_{t,s}$. For $t = s$, we find that $v_{s+1}^e = v$ and since $X_s = X$, then $y_s = y$ and $\Lambda_s = \Lambda$.

To connect the responses after date s to the FIRE Jacobians, let the superscript $*$ denote the variables in case $dX_s = dX_s^e$, i.e., when both the input and the expectation change by the same amount. Note that since $dD_s = dD_s^*$ and $dD_{s+1}^* = \Lambda' dD_s^* + (d\Lambda_s^*)' D$ then

$$dD_{s+1} = \Lambda' dD_s = \Lambda' dD_s^* = dD_{s+1}^* - (d\Lambda_s^*)' D$$

and since $dY_s^* = y' dD_s^* + (dy_s^*)' D$ then

$$dY_s = y' dD_s = y' dD_s^* = dY_s^* - (dy_s^*)' D.$$

Let dD_s^0 and $d\Lambda_s^0$ denote the responses to an unanticipated change, i.e., $dX_s \neq 0$ with $dX_s^e = 0$ as we computed above. Then, since $d\Lambda_s^* = d\Lambda_s^0$ and $dD_{s+1}^0 = (d\Lambda_s^0)' D$

$$dD_{s+1} = dD_{s+1}^* - (d\Lambda_s^*)' D = dD_{s+1}^* - dD_{s+1}^0$$

and since $dy_s^* = dy_s^0$ then

$$dY_s = dY_s^* - (dy_s^*)' D = dY_s^* - dY_s^0.$$

Finally, for $t > s$, we also find that $v_{t+1}^e = v$ and $X_t = X$ so that decisions and transitions do not change. As a result,

$$dD_t = \left(\Lambda'\right)^{t-(s+1)} dD_{s+1} = \left(\Lambda'\right)^{t-(s+1)} \left(dD_{s+1}^* - dD_{s+1}^0\right) = dD_t^* - dD_t^0$$

$$dY_t = y' dD_t = y' dD_t^* - y' dD_t^0 = dY_t^* - dY_t^0.$$

These results mean that we can decompose an expectation shock as follows. Prior to date s, the response is exactly identical to what would be obtained under a perfect-foresight (FIRE) transition to a shock dX_s, since the only effects come from anticipation. From date s on, we can decompose the response to an unrealized expectation shock as the difference between the FIRE response and the response to an unanticipated time-s shock of the opposite sign. It follows that we can write:

$$\frac{dY_t}{dX_s^e} = \begin{cases} \mathcal{J}_{t,s} & \text{if } t < s \\ \mathcal{J}_{t,s} - \mathcal{J}_{t-s,0} & \text{if } t \geq s \end{cases} \tag{C.10}$$

Putting it together Define

$$\mathcal{E} \equiv \begin{bmatrix} \mathcal{J}_{0,0} & 0 & 0 & ... \\ \mathcal{J}_{1,0} & \mathcal{J}_{0,0} & 0 & ... \\ \mathcal{J}_{2,0} & \mathcal{J}_{1,0} & \mathcal{J}_{0,0} & ... \\ ... & ... & ... & ... \end{bmatrix}, \tag{C.11}$$

then we can summarize these results in the following expressions

$$d\mathbf{Y} = (\mathcal{J} - \mathcal{E}) \cdot d\mathbf{X}^e + \mathcal{E} \cdot d\mathbf{X} = \mathcal{J} \cdot d\mathbf{X}^e + \mathcal{E} \cdot \underbrace{(d\mathbf{X} - d\mathbf{X}^e)}_{\text{Forecast Error}}. \tag{C.12}$$

C.3 Proof of proposition 2

For convenience we restate the proposition.

Proposition. *Assume that expectations satisfy Assumption 1, and truncate all paths at a common horizon T. The linearized impulse response $d\boldsymbol{Y} \in \mathbb{R}^T$ to arbitrary shocks $d\boldsymbol{X} \in \mathbb{R}^T$ and beliefs $E_t[d\boldsymbol{X}] \in \mathbb{R}^T$ is given by*

$$d\boldsymbol{Y} = \mathcal{J} \underbrace{E_0[d\boldsymbol{X}]}_{\text{initial forecast}} + \sum_{h=1}^{T-1} \mathcal{R}_h \underbrace{\left(E_h[d\boldsymbol{X}] - E_{h-1}[d\boldsymbol{X}]\right)}_{\text{forecast revision}}. \tag{C.13}$$

*For $h = 1, \ldots, T-1$, define $\mathcal{J}^{(T-h)} = [\mathcal{J}_{ts}]_{t,s=0}^{T-h-1}$. The **forecast-revision Jacobian** is*

$$\mathcal{R}_h = \begin{bmatrix} \boldsymbol{0}_{h\times h} & \boldsymbol{0}_{h\times(T-h)} \\ \boldsymbol{0}_{(T-h)\times h} & \mathcal{J}^{(T-h)} \end{bmatrix}, \qquad (\mathcal{R}_h)_{ts} = \mathbf{1}\{t \geq h,\ s \geq h\}\mathcal{J}_{t-h,s-h}.$$

Let $X_s^{e,t}$ denote the agent's beliefs at each point at time t for the value of the input at time s. We can extend the previous representation as follows:

$$v_t = v\left(v_{t+1}^{e,t}, X_t\right) \tag{C.14}$$

$$v_s^{e,t} = v\left(v_{s+1}^{e,t}, X_s^{e,t}\right) \quad \text{for } s \geq t+1 \tag{C.15}$$

$$D_{t+1} = \Lambda\left(v_{t+1}^{e,t}, X_t\right)' D_t \tag{C.16}$$

$$Y_t = y\left(v_{t+1}^{e,t}, X_t\right)' D_t, \tag{C.17}$$

where $v_{T+1}^{e,t} = v$ and $X_{T+1}^{e,t} = X$ for all t. We still denote the full-information and rational-expectations Jacobian by $\mathcal{J}$.

The response to an unanticipated date s shock, dX_s, is still the same as in the previous case with time-invariant expectations

$$\frac{dY_t}{dX_s} = \begin{cases} 0 & \text{if } t < s \\ \mathcal{J}_{t-s,0} & \text{if } t \geq s \end{cases}$$

So, we only need to extend the previous analysis for shocks to expectations. When expectations at time τ about the input at time s change by $dX_s^{e,\tau}$, outcomes do not change for $t < \tau$. For $t \geq \tau$, the response depends only on the distance from the revision date and the forecast horizon:

$$\frac{\partial Y_t}{\partial X_s^{e,\tau}} = \frac{\partial Y_{t-\tau}}{\partial X_{s-\tau}^{e,0}}. \tag{C.18}$$

Using this fact, we need only consider changes in time-0 expectations.

Note that if we shock the time-0 expectations for time $s > 0$, $dX_s^{e,0}$, then $v_s^{e,t} = v$, $y_t = y$, $\Lambda_t = \Lambda$ for $t > 0$, i.e., only time 0 decisions change. What happens at time 0? Let the superscript $*$ denote the response to a perfect foresight shock, i.e., under FIRE $dX_s = dX_s^{e,t}$ for all t. Then:

$$dY_0 = dy_0' D = \left(dy_0^*\right)' D = dY_0^*$$

$$dD_1 = \left(d\Lambda_0^*\right)' D = dD_1^*$$

For $t > 0$, since decisions revert to steady state, then we only need to account for the change in the distribution:

$$dY_t = y' dD_t$$

$$dD_{t+1} = \Lambda' dD_t = \left(\Lambda'\right)^t dD_1 = \left(\Lambda'\right)^t dD_1^*$$

Following Auclert et al. (2021), the perfect foresight (FIRE) responses are given by

$$dD_t^* = \left(d\Lambda_{t-1}^*\right)' D + \Lambda' dD_{t-1}^* = \sum_{m=0}^{t-2} \left(\Lambda'\right)^m \left(d\Lambda_{t-1-m}^*\right)' D + \left(\Lambda'\right)^{t-1} dD_1^*$$

$$dY_t^* = \left(dy_t^*\right)' D + y' dD_t^*$$

Note that the first equation implies that

$$dD_t^* = \sum_{m=0}^{t-2} \left(\Lambda'\right)^m \left(d\Lambda_{t-1-m}^*\right)' D + \left(\Lambda'\right)^{t-1} dD_1^* = \sum_{m=0}^{t-2} \left(\Lambda'\right)^m \left(d\Lambda_{t-1-m}^*\right)' D + dD_t$$

It is useful to put superscripts for the time of the shock, s. For example, define the FIRE response as:

$$dD_t^{*,s} = \sum_{m=0}^{t-2} \left(\Lambda'\right)^m \left(d\Lambda_{t-1-m}^{*,s}\right)' D + \left(\Lambda'\right)^{t-1} dD_1^{*,s}$$

$$dY_t^{*,s} = \left(dy_t^{*,s}\right)' D + y' dD_t^{*,s}$$

So, the response to the time-0 expectation shock for the time-s input is given by:

$$dD_t^s = dD_t^{*,s} - \sum_{m=0}^{t-2} \left(\Lambda'\right)^m \left(d\Lambda_{t-1-m}^{*,s}\right)' D$$

$$dY_t^s = y' dD_t^s = \underbrace{y' dD_t^{*,s}}_{=dY_t^{*,s} - \left(dy_t^{*,s}\right)' D} + y' \left(dD_t^s - dD_t^{*,s}\right) = dY_t^{*,s} - \left(dy_t^{*,s}\right)' D + y' \left(dD_t^s - dD_t^{*,s}\right).$$

Note furthermore, that

$$\begin{aligned}
dD_t^{*,s-1} &= \left(d\Lambda_{t-1}^{*,s-1}\right)' D + \Lambda' dD_{t-1}^{*,s-1} \\
&= \left(d\Lambda_{t-1}^{*,s-1}\right)' D + \Lambda' \left(d\Lambda_{t-1}^{*,s-1}\right)' D + \left(\Lambda'\right)^2 dD_{t-2}^{*,s-1} \\
&= \sum_{m=0}^{t-1} \left(\Lambda'\right)^m \left(d\Lambda_{t-1-m}^{*,s-1}\right)' D + \left(\Lambda'\right)^t \underbrace{dD_0^{*,s-1}}_{=0} \\
&= \sum_{m=0}^{t-1} \left(\Lambda'\right)^m \left(d\Lambda_{t-1-m}^{*,s-1}\right)' D
\end{aligned}$$

and now using the fact that

$$d\Lambda_{t-1-m}^{*,s-1} = d\Lambda_{t-m}^{*,s}$$

we can write

$$dD_t^{*,s-1} = \sum_{m=0}^{t-1} \left(\Lambda'\right)^m \left(d\Lambda_{t-m}^{*,s}\right)' D \quad \text{and } dD_{t-1}^{*,s-1} = \sum_{m=0}^{t-2} \left(\Lambda'\right)^m \left(d\Lambda_{t-1-m}^{*,s}\right)' D$$

It follows that:

$$dD_t^s = dD_t^{*,s} - dD_{t-1}^{*,s-1}.$$

Finally, replacing dD_t^s in

$$dY_t^s = dY_t^{*,s} - \left(dy_t^{*,s}\right)' D + y' \left(dD_t^s - dD_t^{*,s}\right)$$

we obtain

$$dY_t^s = dY_t^{*,s} - \left(dy_t^{*,s}\right)' D - y' dD_{t-1}^{*,s-1}.$$

Finally, since $dy_t^{*,s} = dy_{t-1}^{*,s-1}$ then

$$dY_t^s = dY_t^{*,s} - \left(\left(dy_{t-1}^{*,s-1}\right)' D + y' dD_{t-1}^{*,s-1}\right) = dY_t^{*,s} - dY_{t-1}^{*,s-1}.$$

It follows that:

$$\frac{\partial Y_t}{\partial X_s^{e,0}} = \mathcal{J}_{t,s} - \mathcal{J}_{t-1,s-1}. \tag{C.19}$$

Putting everything together We have found that

$$\frac{\partial Y_t}{\partial X_s} = \begin{cases} 0 & \text{if } t < s \\ \mathcal{J}_{t-s,0} & \text{if } t \geq s \end{cases}$$

and

$$\frac{\partial Y_t}{\partial X_s^{e,\tau}} = \begin{cases} 0 & \text{if } t < \tau \text{ or } s \le \tau \\ \mathcal{J}_{t-\tau,s-\tau} - \mathcal{J}_{t-\tau-1,s-\tau-1} & \text{if } t > \tau \text{ and } s > \tau \\ \mathcal{J}_{0,s-t} & \text{if } t = \tau \text{ and } s > \tau = t. \end{cases}$$

Putting everything together we can write

$$\begin{aligned}
dY_t &= \sum_{s=0}^{t} \mathcal{J}_{t-s,0} \cdot dX_s + \sum_{\tau=0}^{t-1} \sum_{s=\tau+1}^{\infty} \left(\mathcal{J}_{t-\tau,s-\tau} - \mathcal{J}_{t-\tau-1,s-\tau-1}\right) \cdot dX_s^{e,\tau} + \sum_{s=t+1}^{\infty} \mathcal{J}_{0,s-t} \cdot dX_s^{e,t} \\
\Leftrightarrow dY_t &= \sum_{s=0}^{t} \mathcal{J}_{t-s,0} dX_s + \sum_{\tau=0}^{t-1} \sum_{s=\tau+1}^{\infty} \mathcal{J}_{t-\tau,s-\tau} dX_s^{e,\tau} - \sum_{\tau=0}^{t-1} \sum_{s=\tau+1}^{\infty} \mathcal{J}_{t-\tau-1,s-\tau-1} dX_s^{e,\tau} + \sum_{s=t+1}^{\infty} \mathcal{J}_{0,s-t} dX_s^{e,t} \\
\Leftrightarrow dY_t &= \sum_{\tau=0}^{t} \mathcal{J}_{t-\tau,0} dX_\tau^{e,\tau} + \sum_{\tau=0}^{t-1} \sum_{s=\tau+1}^{\infty} \mathcal{J}_{t-\tau,s-\tau} dX_s^{e,\tau} - \sum_{\tau=0}^{t-1} \sum_{s=\tau+1}^{\infty} \mathcal{J}_{t-\tau-1,s-\tau-1} dX_s^{e,\tau} + \sum_{s=t+1}^{\infty} \mathcal{J}_{0,s-t} dX_s^{e,t} \\
\Leftrightarrow dY_t &= \mathcal{J}_{0,0} dX_t^{e,t} + \sum_{\tau=0}^{t-1} \sum_{s=\tau}^{\infty} \mathcal{J}_{t-\tau,s-\tau} dX_s^{e,\tau} - \sum_{\tau=0}^{t-1} \sum_{s=\tau+1}^{\infty} \mathcal{J}_{t-\tau-1,s-\tau-1} dX_s^{e,\tau} + \sum_{s=t+1}^{\infty} \mathcal{J}_{0,s-t} dX_s^{e,t} \\
\Leftrightarrow dY_t &= \sum_{\tau=0}^{t-1} \sum_{s=\tau}^{\infty} \mathcal{J}_{t-\tau,s-\tau} dX_s^{e,\tau} - \sum_{\tau=0}^{t-1} \sum_{s=\tau+1}^{\infty} \mathcal{J}_{t-\tau-1,s-\tau-1} dX_s^{e,\tau} + \sum_{s=t}^{\infty} \mathcal{J}_{0,s-t} dX_s^{e,t} \\
\Leftrightarrow dY_t &= \sum_{\tau=0}^{t} \sum_{s=\tau}^{\infty} \mathcal{J}_{t-\tau,s-\tau} dX_s^{e,\tau} - \sum_{\tau=0}^{t-1} \sum_{s=\tau+1}^{\infty} \mathcal{J}_{t-\tau-1,s-\tau-1} dX_s^{e,\tau} \\
\Leftrightarrow dY_t &= \sum_{s=0}^{\infty} \mathcal{J}_{t,s} dX_s^{e,0} + \sum_{\tau=1}^{t} \sum_{s=\tau}^{\infty} \mathcal{J}_{t-\tau,s-\tau} dX_s^{e,\tau} - \sum_{\tau=0}^{t-1} \sum_{s=\tau+1}^{\infty} \mathcal{J}_{t-\tau-1,s-\tau-1} dX_s^{e,\tau} \\
\Leftrightarrow dY_t &= \sum_{s=0}^{\infty} \mathcal{J}_{t,s} dX_s^{e,0} + \sum_{\tau=1}^{t} \sum_{s=\tau}^{\infty} \mathcal{J}_{t-\tau,s-\tau} dX_s^{e,\tau} - \sum_{\hat{\tau}=1}^{t} \sum_{s=\hat{\tau}}^{\infty} \mathcal{J}_{t-(\hat{\tau}-1)-1,s-(\hat{\tau}-1)-1} dX_s^{e,\hat{\tau}-1} \\
\Leftrightarrow dY_t &= \sum_{s=0}^{\infty} \mathcal{J}_{t,s} dX_s^{e,0} + \sum_{\tau=1}^{t} \sum_{s=\tau}^{\infty} \mathcal{J}_{t-\tau,s-\tau} dX_s^{e,\tau} - \sum_{\tau=1}^{t} \sum_{s=\tau}^{\infty} \mathcal{J}_{t-\tau,s-\tau} dX_s^{e,\tau-1} \\
\Leftrightarrow dY_t &= \sum_{\tau=1}^{t} \sum_{s=\tau}^{\infty} \mathcal{J}_{t-\tau,s-\tau} \left(dX_s^{e,\tau} - dX_s^{e,\tau-1}\right) + \sum_{s=0}^{\infty} \mathcal{J}_{t,s} dX_s^{e,0}
\end{aligned}$$

from where equation (34) follows immediately.

C.4 Special cases

Throughout, we maintain the following notation. $E_t[dX_{t+h}]$ denotes the agent's time-t expectations about the variable at horizon h. $\mathbb{E}_t[dX_{t+h}]$ denotes the full-information and rational expectation for the same variable.

C.4.1 Shallow reasoning

(Angeletos and Sastry, 2021). $E_t[dX_{t+h}] = \lambda \cdot dX_{t+h}$.

$$\Lambda_0 = \begin{bmatrix} 1 & 0 & 0 & 0 & \dots \\ 0 & \lambda & 0 & 0 & \dots \\ 0 & 0 & \lambda & 0 & \dots \\ 0 & 0 & 0 & \lambda & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix} \quad \Lambda_1 = \begin{bmatrix} 1 & 0 & 0 & 0 & \dots \\ 0 & 1 & 0 & 0 & \dots \\ 0 & 0 & \lambda & 0 & \dots \\ 0 & 0 & 0 & \lambda & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix} \tag{C.20}$$

This diagonal example also covers scalar sparsity. More generally, horizon-specific salience parameters $m_h \in [0, 1]$ imply $E_t[dX_{t+h}] = m_h dX_{t+h}$, so the first $t + 1$ diagonal entries of Λ_t equal one and its remaining diagonal entries are $m_1, m_2, \dots$.

C.4.2 Finite planning horizon

(Woodford, 2018). With planning horizon H, agents observe realized dates and attend fully to inputs no more than H periods ahead:

$$E_t[dX_s] = \begin{cases} dX_s, & s \leq t + H, \\ 0, & s > t + H. \end{cases}$$

Thus, Λ_t is diagonal, with ones through diagonal entry $t + H$ and zeros thereafter.

C.4.3 Cognitive discounting

(Gabaix, 2020). $E_t[dX_{t+h}] = \lambda^h \cdot dX_{t+h}$.

$$\Lambda_0 = \begin{bmatrix} 1 & 0 & 0 & 0 & \dots \\ 0 & \lambda & 0 & 0 & \dots \\ 0 & 0 & \lambda^2 & 0 & \dots \\ 0 & 0 & 0 & \lambda^3 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix} \quad \Lambda_1 = \begin{bmatrix} 1 & 0 & 0 & 0 & \dots \\ 0 & 1 & 0 & 0 & \dots \\ 0 & 0 & \lambda & 0 & \dots \\ 0 & 0 & 0 & \lambda^2 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix} \tag{C.21}$$

C.4.4 Sticky expectations

(Mankiw and Reis, 2002, Carroll et al., 2018). A date-0 shock ϵ causes a sequence of disturbances $\{dX_t\}$. At each date $t \geq 0$, some agents learn about ϵ and deduce $\{dX_{t+h}\}$ for all $h \geq 0$. The probability of learning ϵ is $1 - \lambda$ for every agent who hasn't learned it already. Thus the share of ignorant agents at date t is λ^{t+1}. They believe that the disturbances observed so far were special events, and don't expect any disturbances in the future. This setup implies that average expectations are $E_t[dX_{t+h}] = (1-\lambda^{t+1}) \cdot dX_{t+h}$.

$$\Lambda_0 = \begin{bmatrix} 1 & 0 & 0 & 0 & \dots \\ 0 & 1-\lambda & 0 & 0 & \dots \\ 0 & 0 & 1-\lambda & 0 & \dots \\ 0 & 0 & 0 & 1-\lambda & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix} \qquad \Lambda_1 = \begin{bmatrix} 1 & 0 & 0 & 0 & \dots \\ 0 & 1 & 0 & 0 & \dots \\ 0 & 0 & 1-\lambda^2 & 0 & \dots \\ 0 & 0 & 0 & 1-\lambda^2 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix} \tag{C.22}$$

C.4.5 Noisy information and rational expectations

(Angeletos and Huo, 2021). A date-0 shock ϵ causes a sequence of disturbances $\{dX_t\}$ according to an MA process

$$dX_t = M_t \epsilon \tag{C.23}$$

Suppose that agents know the MA coefficients, M_t, but they don't observe ϵ. Their prior is that ϵ is distributed $\mathcal{N}(0, 1/\tau_\epsilon)$. At each date $t \geq 0$, agents receive independent private signals $\epsilon + \nu_t$, where $\nu_t \sim \mathcal{N}(0, 1/\tau_\nu)$. Bayesian updating implies that the average posterior belief is

$$E_t[\epsilon] = \frac{t+1}{\tau_\epsilon/\tau_\nu + t + 1}\epsilon \tag{C.24}$$

Then, the average expectation of dX_{t+h} at date t is

$$E_t[dX_{t+h}] = M_{t+h} E_t[\epsilon] = M_{t+h} \underbrace{\left(\frac{t+1}{\tau_\epsilon/\tau_\nu + t + 1}\right)}_{\lambda_t} \epsilon = \lambda_t dX_{t+h} \tag{C.25}$$

Thus the associated Λ_t matrices are

$$\Lambda_0 = \begin{bmatrix} 1 & 0 & 0 & 0 & \dots \\ 0 & \lambda_0 & 0 & 0 & \dots \\ 0 & 0 & \lambda_0 & 0 & \dots \\ 0 & 0 & 0 & \lambda_0 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix} \quad \Lambda_1 = \begin{bmatrix} 1 & 0 & 0 & 0 & \dots \\ 0 & 1 & 0 & 0 & \dots \\ 0 & 0 & \lambda_1 & 0 & \dots \\ 0 & 0 & 0 & \lambda_1 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix} \tag{C.26}$$

C.4.6 Extrapolation

Geometric extrapolation. $E_t[dX_{t+h}] = \lambda^h dX_t$. First example of non-diagonal Λ matrices.

$$\Lambda_0 = \begin{bmatrix} 1 & 0 & 0 & 0 & \dots \\ \lambda & 0 & 0 & 0 & \dots \\ \lambda^2 & 0 & 0 & 0 & \dots \\ \lambda^3 & 0 & 0 & 0 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix} \qquad \Lambda_1 = \begin{bmatrix} 1 & 0 & 0 & 0 & \dots \\ 0 & 1 & 0 & 0 & \dots \\ 0 & \lambda & 0 & 0 & \dots \\ 0 & \lambda^2 & 0 & 0 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix} \tag{C.27}$$

C.4.7 Adaptive expectations

(Cagan, 1956, Friedman, 1957). Current and past realizations are observed, so $E_t[dX_s] = dX_s$ for $s \leq t$. For $h \geq 1$, let

$$E_t[dX_{t+h}] = \kappa \sum_{\tau=0}^{t} \lambda^{h+\tau} dX_{t-\tau},$$

where $\kappa > 0$ scales the geometrically weighted history. The associated matrices begin

$$\Lambda_0 = \begin{bmatrix} 1 & 0 & 0 & 0 & \dots \\ \kappa\lambda & 0 & 0 & 0 & \dots \\ \kappa\lambda^2 & 0 & 0 & 0 & \dots \\ \kappa\lambda^3 & 0 & 0 & 0 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix} \qquad \Lambda_1 = \begin{bmatrix} 1 & 0 & 0 & 0 & \dots \\ 0 & 1 & 0 & 0 & \dots \\ \kappa\lambda^2 & \kappa\lambda & 0 & 0 & \dots \\ \kappa\lambda^3 & \kappa\lambda^2 & 0 & 0 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix} \tag{C.28}$$

C.4.8 Diagnostic expectations

(Bordalo et al., 2018, Bianchi et al., 2021). Let $E_t^r[dX_{t+h}]$ denote a reference expectation for the variable h periods ahead. Then, the diagnostic expectation with parameter θ is given by:

$$E_t[dX_{t+h}] = \mathbb{E}_t[dX_{t+h}] + \theta\left(\mathbb{E}_t[dX_{t+h}] - E_t^r[dX_{t+h}]\right).$$

Bordalo et al. (2018) assume that $E_t^r[dX_{t+h}] = \mathbb{E}_{t-1}[dX_{t+h}]$. In this case,

$$\Lambda_0 = \begin{bmatrix} 1 & 0 & 0 & 0 & \dots \\ 0 & 1+\theta & 0 & 0 & \dots \\ 0 & 0 & 1+\theta & 0 & \dots \\ 0 & 0 & 0 & 1+\theta & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}, \quad \Lambda_t = \begin{bmatrix} 1 & 0 & 0 & 0 & \dots \\ 0 & 1 & 0 & 0 & \dots \\ 0 & 0 & 1 & 0 & \dots \\ 0 & 0 & 0 & 1 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix} \tag{C.29}$$

for $t \geq 1$.

Bianchi et al. (2021) develop a generalization of this framework to allow for long memory, which assumes that $E_t^r[dX_{t+h}] = \sum_{j=1}^{\infty} \alpha_j \mathbb{E}_{t-j}[dX_{t+h}]$. With this assumption, we find

$$\Lambda_0 = \begin{bmatrix} 1 & 0 & 0 & 0 & \dots \\ 0 & 1+\theta & 0 & 0 & \dots \\ 0 & 0 & 1+\theta & 0 & \dots \\ 0 & 0 & 0 & 1+\theta & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}, \quad \Lambda_1 = \begin{bmatrix} 1 & 0 & 0 & 0 & \dots \\ 0 & 1 & 0 & 0 & \dots \\ 0 & 0 & 1+\theta(1-\alpha_1) & 0 & \dots \\ 0 & 0 & 0 & 1+\theta(1-\alpha_1) & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}, \tag{C.30}$$

and, for any t,

$$\Lambda_t = \begin{bmatrix} 1 & 0 & 0 & 0 & \dots \\ \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & 1+\theta(1-\sum_{j=1}^{t}\alpha_j) & 0 & \dots \\ 0 & 0 & 0 & 1+\theta(1-\sum_{j=1}^{t}\alpha_j) & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}, \tag{C.31}$$

C.4.9 Noisy information and diagnostic expectations

(Bordalo et al., 2020) As in noisy information and rational expectations, the agent observes a signal $\epsilon + \nu_t$, where $\nu_t \sim \mathcal{N}(0, 1/\tau_\nu)$. However, the forecaster then overweighs representative states by using the distorted posterior

$$f^\theta(\epsilon|S_t^i) = f(\epsilon|S_t^i) R_t^i(\epsilon)^\theta \frac{1}{Z_t} \tag{C.32}$$

Bordalo et al. (2020) assume that $R_t^i(\epsilon) = f(\epsilon|S_t^i)/f(\epsilon|S_{t-1}^i \cup \{\mathbb{E}_{i,t-1}[\epsilon]\})$. This assumption implies that the mean of the distorted posterior is given by:

$$E_{i,t}^\theta[\epsilon] = \mathbb{E}_{i,t}[\epsilon] + \theta\left(\mathbb{E}_{i,t}[\epsilon] - \mathbb{E}_{i,t-1}[\epsilon]\right) \tag{C.33}$$

where $\mathbb{E}_{i,t}[\epsilon]$ denotes the time-t rational expectation with information set S_t^i. It follows that the average expectation is given by

$$\overline{E}_t^\theta[\epsilon] = \underbrace{\left[\frac{\left(\frac{t+1+\theta}{t+1}\right)\tau_\epsilon/\tau_\nu + t}{\tau_\epsilon/\tau_\nu + t}\right] \frac{t+1}{\tau_\epsilon/\tau_\nu + t + 1}}_{\equiv \lambda_t} \epsilon \tag{C.34}$$

Thus the associated Λ_t matrices are

$$\Lambda_0 = \begin{bmatrix} 1 & 0 & 0 & 0 & \dots \\ 0 & \lambda_0 & 0 & 0 & \dots \\ 0 & 0 & \lambda_0 & 0 & \dots \\ 0 & 0 & 0 & \lambda_0 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}, \quad \Lambda_1 = \begin{bmatrix} 1 & 0 & 0 & 0 & \dots \\ 0 & 1 & 0 & 0 & \dots \\ 0 & 0 & \lambda_1 & 0 & \dots \\ 0 & 0 & 0 & \lambda_1 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}. \tag{C.35}$$

Analogously to Bianchi et al. (2021), we can extend this model to include long-memory as follows. Assume that $R_t^\theta(\epsilon) = f(\epsilon|S_t^i)/f^*(\epsilon|S_t^i)$, where $\epsilon \sim_{f^*|S_t^i} \mathcal{N}(E_t^r[\epsilon], \tau_\epsilon + (t+1)\tau_\nu)$ and $E_t^r[\epsilon] = \sum_{j=1}^{t} \alpha_j \mathbb{E}_{i,t-j}[\epsilon]$. It follows that

$$E_{i,t}^\theta[\epsilon] = \mathbb{E}_{i,t}[\epsilon] + \theta\left(\mathbb{E}_{i,t}[\epsilon] - E_{i,t}^r[\epsilon]\right). \tag{C.36}$$

As a result, the average expectation is given by

$$\overline{E}_t^{\theta}[\epsilon] = \left[(1+\theta)\frac{t+1}{\tau_\epsilon/\tau_\nu + t + 1} - \theta \sum_{j=1}^{t} \alpha_j \left(\frac{t+1-j}{\tau_\epsilon/\tau_\nu + t + 1 - j}\right)\right] \epsilon, \tag{C.37}$$

and defining now $\lambda_t \equiv (1+\theta)\frac{t+1}{\tau_\epsilon/\tau_\nu+t+1} - \theta \sum_{j=1}^{t} \alpha_j \left(\frac{t+1-j}{\tau_\epsilon/\tau_\nu+t+1-j}\right)$ we obtain the analogous Λ_t matrices as above.

D Appendix to Sections 5, 6, and 7

D.1 A flexible model of expectations

Propositions 3 and 2 enable us to compute linearized impulse responses to any shock $d\mathbf{X}$ given the path of expectations $\{dE_t[\mathbf{X}]\}_t$ conditional on the same shock.[17] In some cases, the response of expectations may be estimated directly. We do so in section 5.2 with respect to unemployment. Another route is to impose a model of expectation formation. In this subsection, we present a tractable yet flexible specification that nests many popular models including: (1) sticky expectations (Mankiw and Reis, 2002, Carroll et al., 2018), (2) noisy-information and rational expectations (Angeletos and Huo, 2021), (3) cognitive discounting (Gabaix, 2020), (4) sparsity (Gabaix, 2014, 2016, Guerreiro, 2022), (5) shallow reasoning (Angeletos and Sastry, 2021), (6) finite planning horizons (Woodford, 2018), (7) adaptive expectations (Cagan, 1956, Friedman, 1957), (8) diagnostic expectations (Bordalo et al., 2018, Bianchi et al., 2021), (9) noisy-information diagnostic expectations (Bordalo et al., 2020), among others. Appendix C.4 discusses how to map each of these models into our framework.

Propositions 3 and 2 deal with linear mappings in sequence space. So it is natural for us to model expectations in the same way. Given a single time-varying input $d\mathbf{X}$, the most general linear sequence-space model of expectations is a sequence of matrices $\Lambda_t \in \mathbb{R}^{T\times T}$ that map realized outcomes $d\mathbf{X} \in \mathbb{R}^T$ into time-t expectations $dE_t[\mathbf{X}] \in \mathbb{R}^T$ according to

$$dE_t[\mathbf{X}] = \Lambda_t d\mathbf{X} \tag{D.1}$$

We maintain the assumption that expectations of current and past realizations of $d\mathbf{X}$ are correct. Because dates are indexed from 0, this implies that the first $t+1$ rows of Λ_t reproduce the first $t+1$ realizations:

$$\Lambda_t = \begin{bmatrix} \mathbf{I}_{(t+1)\times(t+1)} & \mathbf{0}_{(t+1)\times(T-t-1)} \\ A_t & B_t \end{bmatrix} \tag{D.2}$$

where $A_t \in \mathbb{R}^{(T-t-1)\times(t+1)}$ and $B_t \in \mathbb{R}^{(T-t-1)\times(T-t-1)}$ govern forecasts of future dates. Equation (D.1) looks simple but it can capture rich theories of expectation formation. It can account for

[17] Recall that deviations $d\mathbf{X}$ and $\{dE_t[\mathbf{X}]\}_t$ are all relative to the paths around which one wishes to compute the Jacobian.

an understanding of the data generating process as well as for updating of priors in light of new observations. In our expository environment, $d\mathbf{X}$ is the only input the agent has to form expectations about. So it's not restrictive to assume that $dE_t[\mathbf{X}]$ depends only on the realized path of $d\mathbf{X}$ itself. In richer environments with multiple inputs, one may introduce additional linear terms, allowing the agent to think about cross-equation restrictions directly. For example, expectations of unemployment benefits and income taxes may be related via an understanding of the government budget constraint.

Corollaries 1 and 2 substitute (D.1) into propositions 3 and 2. Their bracketed terms are the Jacobians of the non-FIRE decision map g. These Jacobians account for both the direct effect of $d\mathbf{X}$ on $d\mathbf{Y}$ and its indirect effect through $\{E_t[\mathbf{X}]\}_t$. They remain $T \times T$ matrices, just like the FIRE Jacobians of f. Thus, the remaining SSJ machinery of Auclert et al. (2021) applies without modification. The framework therefore accommodates rich heterogeneity under general deviations from FIRE.

Corollary 1. *Consider the setup of proposition 3 with FIRE Jacobian $\mathcal{J}$, and forecast-error Jacobian $\mathcal{E}$. Let the constant expectations be $d\boldsymbol{X}^e = \Lambda d\boldsymbol{X}$, according to* (D.1). *The linearized impulse response $d\boldsymbol{Y}$ to an arbitrary shock $d\boldsymbol{X}$ is*

$$d\boldsymbol{Y} = \Big[(\mathcal{J} - \mathcal{E})\,\Lambda + \mathcal{E} \Big] d\boldsymbol{X} \tag{D.3}$$

Corollary 2. *Consider the setup of proposition 2 with FIRE Jacobian $\mathcal{J}$ and forecast-revision Jacobians $\mathcal{R}_\tau$. Let expectations be $dE_t[\boldsymbol{X}] = \Lambda_t d\boldsymbol{X}$, according to* (D.1). *The linearized impulse response $d\boldsymbol{Y}$ to an arbitrary shock $d\boldsymbol{X}$ is*

$$d\boldsymbol{Y} = \left[\mathcal{J}\Lambda_0 + \sum_{\tau \geq 1} \mathcal{R}_\tau \left(\Lambda_\tau - \Lambda_{\tau-1}\right) \right] d\boldsymbol{X} \tag{D.4}$$

Given a model for beliefs, the matrices given in equations (D.3) or (D.4) fully summarize the response of the heterogeneous agent block to the path $d\mathbf{X}$, taking into account forecast errors and revisions. These matrices are easy and fast to compute after obtaining the FIRE Jacobians $\mathcal{J}$.

D.2 Fitting a model of beliefs

Figure D.1: Cognitive Distortion in Parametric Belief Models

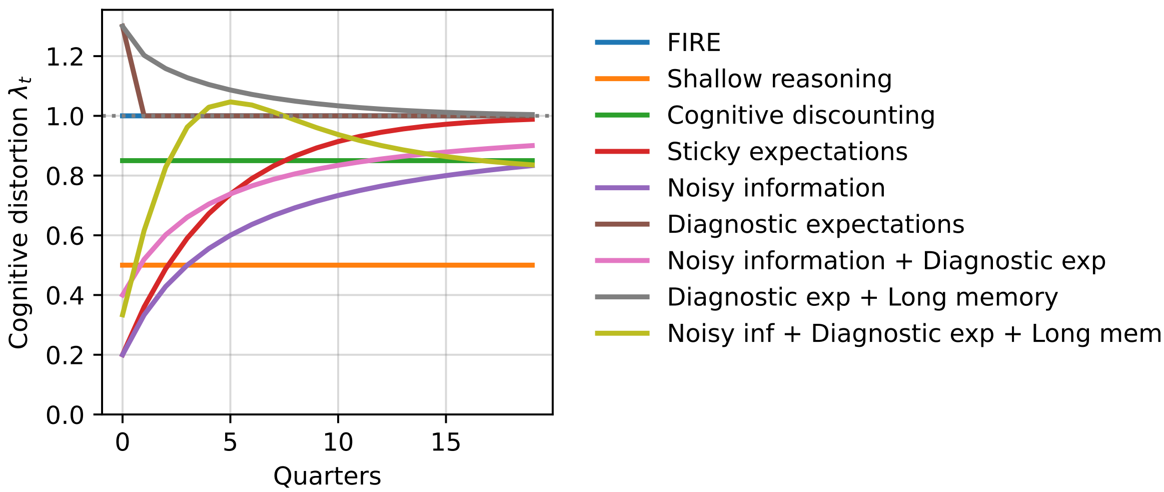

Notes: This figure plots the cognitive distortion coefficient $\lambda_t \equiv \overline{E}_t[dX_{t+h}]/dX_{t+h}$ implied by each parametric belief model, as a function of the belief-formation date t (for cognitive discounting we report the one-step-ahead, $h = 1$, coefficient). Values below one indicate underreaction and values above one overreaction relative to the full-information rational-expectations benchmark $\lambda_t = 1$. All models use illustrative calibrations, except the noisy-information long-memory diagnostic model, which uses our estimated calibration (Table 3). It is the only model that can generate initial underreaction followed by delayed overreaction.

Figure D.2: Estimated Memory Weights

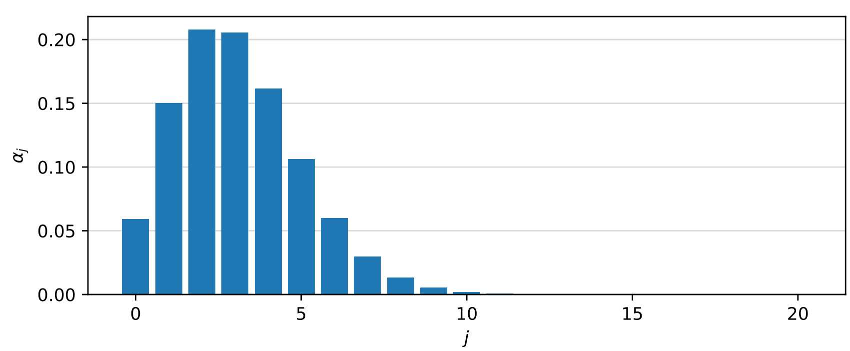

Notes: This figure plots the estimated memory weights $\{\alpha_j\}_{j\geq 1}$ of the noisy-information long-memory diagnostic expectations model under our classical calibration. The weights follow a Beta-binomial distribution with the parameters reported in Table 3, and govern how past forecast revisions enter the diagnostic component of beliefs.

E Additional Figures for Sections 6 and 7

Figure E.1: Classically Estimated IRFs to the Main Business-Cycle Shock

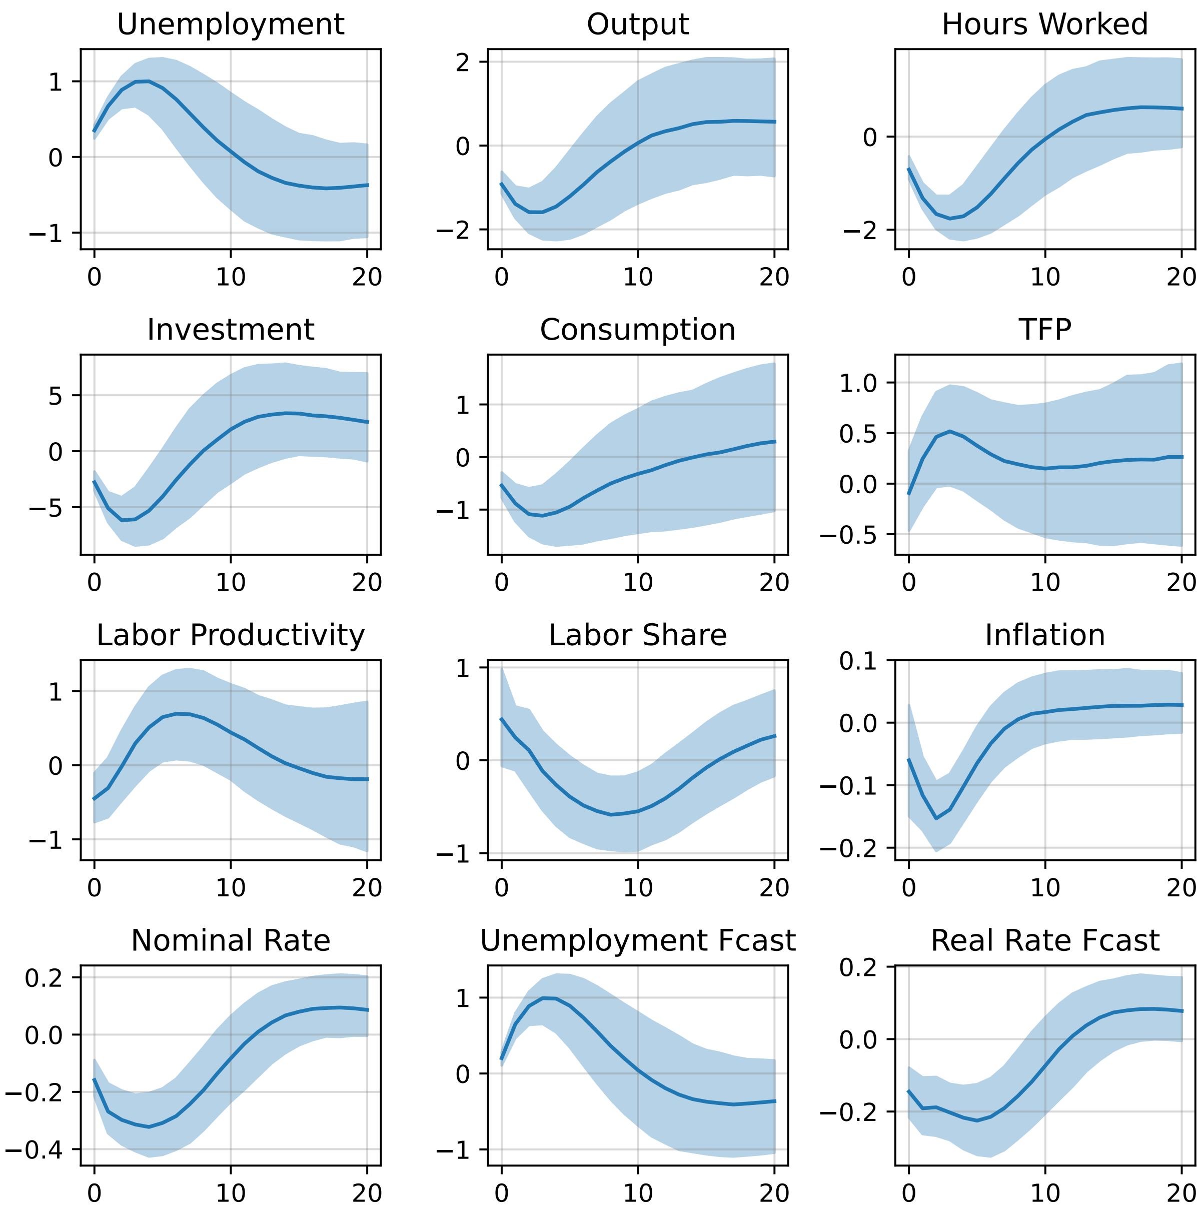

Notes: This figure reports the classically estimated impulse responses to the main business-cycle shock used in the empirical analysis. These responses provide the empirical moments for the model's dynamic estimation.