Bank of England

A flexible deviation from FIRE in the sequence space

Staff Working Paper No. 1,197
Juky 2026

Jamie Lenney and Biagio Rosso

Bank of England

Staff Working Paper No. 1,197

A flexible deviation from FIRE in the sequence space

Jamie Lenney[(1)] and Biagio Rosso[(2)]

Abstract

This paper proposes a flexible approach to departing from full information rational expectations (FIRE) in DSGE models using the sequence space framework. We implement a reduced-form behavioural expectations process in which agents can simultaneously overreact or underreact to current economic conditions and underreact to news, and we derive an associated behavioural expectations solver deployable to a wide class of DSGE and HANK models. The approach nests several different expectations models as special cases while remaining agnostic on the precise source of belief frictions. We apply it to a medium-scale two-asset HANK model and jointly estimate the model's dynamic and behavioural parameters on US business cycle data, including inflation expectation data. Behavioural expectations quantitatively improve the empirical fit of the model and the qualitative properties of its impulse response functions.

Key words: Behavioural expectations, HANK models, sequence space, monetary policy, information rigidity, asymmetric attention, macroeconomic transmission, business cycle estimation, DSGE.

JEL classification: D84, E31, E52, E70.

---

(1) Bank of England. Email: jamie.lenney@bankofengland.co.uk
(2) University of Cambridge. Email: br421@cam.ac.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England or its committees. We are grateful for comments and support from Adrien Auclert, Daniel Albuquerque, Bence Bardoczy, Lukasz Rachel, Ricardo Reis and Vincent Sterk. We also thank participants to the Bank of England internal seminar series, 2026 Scottish Economic Conference, and 20th Annual Dynare Conference.

The Bank's working paper series can be found at www.bankofengland.co.uk/working-paper/staff-working-papers

Bank of England, Threadneedle Street, London, EC2R 8AH
Email: enquiries@bankofengland.co.uk

ISSN 1749-9135 (on-line)

1 Introduction

Can we depart from full information rational expectations (FIRE) in a flexible and standardised manner that is consistent with empirical evidence from macro, micro, and behavioural economics? This question has renewed focus given that the recent heterogeneous agent New Keynesian (HANK) literature emphasises the need to match both microeconomic evidence and macroeconomic dynamics, particularly hump-shaped impulse responses to monetary policy shocks, simultaneously. While Auclert et al. (2020) propose sticky expectations as a solution to generate persistence in macro humps, this approach tends to reduce amplification — a feature highlighted by Albuquerque et al. (2025) when trying to match large impulse responses. Further, the approach ignores microeconomic evidence that agents not only update their information and expectations infrequently, but in forming their expectations also tend to extrapolate from current conditions.

Taking stock, this paper proposes a semi-structural departure from rational expectations based on the micro-founded empirical framework of Kohlhas and Walther (2021). Their reduced-form specification captures how actual expectations deviate from rational expectations through two channels: extrapolation from current conditions and underreaction to news. Our main contribution is to flexibly implement this behavioural expectations process in the sequence space framework for the solution and estimation of DSGE and HANK models developed by Auclert et al. (2021), allowing for tractable incorporation of non-rational expectations into complex, quantitatively rich models. In so doing, we develop a fast and flexible algorithm that solves for these alternative expectations processes in seconds as a function of the rational expectations solution.[1] As we show when laying out our theoretical framework, our implementation of the process building on Kohlhas and Walther (2021) is best understood as a generalisation of sticky-expectations and informational rigidity models to a richer case, such that whenever agents infrequently update their information set they incur a further bias from over- or under-extrapolating from current macroeconomic conditions.

Our framework delivers four key advantages. First, it offers a clean mapping from empirically estimated expectation parameters to quantitative macro models. Second, it allows agents to simultaneously overreact or underreact to current outcomes while underreacting to news about future economic conditions — the joint pattern documented in survey data. We further show, for a general class of hump-shaped impulse responses, that the synergistic interaction between stickiness and over-reactive extrapolation is necessary and sufficient to generate dynamic overshooting — the pattern documented by Angeletos et al. (2021) as critical for expectations frictions to drive endogenous amplification of macroeconomic shocks. Third, it nests as special cases the leading alternative models: FIRE, sticky information, noisy information, asymmetric attention, cognitive discounting, and extrapolation, enabling direct model comparisons. Fourth, it is internally rational in the sense of Adam and Marcet (2011): agents optimise given their subjective beliefs, expectations converge to rational expectations in the long run, and a law of iterated expectations holds so that forecast errors are orthogonal to agents' own projections. Putting this together, by including just two extra parameters we can open up the estimation of a DSGE model in the sequence space to capture a wide range of behavioural frictions in a tractable manner, that a FIRE or purely sticky model could not capture.

We demonstrate the usefulness, speed, and flexibility of our approach by applying it to the two-asset HANK model of Auclert et al. (2020). The authors in that paper advocate for sticky expectations in order to produce empirically consistent hump-shaped impulse responses of macroeconomic aggregates while maintaining empirically realistic marginal propensities to consume. Their IRF matching approach calls for a considerable degree of stickiness in household expectations.[2] We estimate this model more generally by allowing for departures from rational expectations outside of the household sector, and jointly estimate the model's dynamic and behavioural parameters on US business cycle data using a full-information Bayesian maximum likelihood approach. We also incorporate survey expectations data into the estimation explicitly. We compare the properties of the model estimated

[1] Repository link: `https://github.com/s0840389/Flexibile_Deviations_SSJ/tree/main`

[2] Kaplan et al. (2018) discuss how household consumption problems depend on a larger set of inputs — wages, taxes, and so on — than the RANK equivalent, where the direct real interest rate channel dominates. This makes a HANK model a more natural setting in which to apply our algorithm, since expectations over a broader array of prices matter more for general equilibrium outcomes.

under: (1) FIRE; (2) calibrated behavioural parameters based on the empirical evidence of Kohlhas and Walther (2021); (3) estimated behavioural parameters using US business cycle data and survey expectations data; and (4) estimated behavioural parameters using US business cycle data without survey expectations data. To the best of our knowledge, we are the first to estimate a HANK model jointly on macroeconomic and survey expectations data and to do so with a behavioural expectations process.

Approaches (2), (3) and (4) all deliver statistically significant improvements over FIRE in terms of model likelihood and the qualitative properties of the model's impulse responses. Notably approach (1) and (2) estimate the same number of parameters. When comparing root mean squared errors (RMSE) and autocorrelation functions, however, FIRE still performs well outside of the first few horizons. This is consistent with the main contribution of the behavioural expectations process lying in the generation of more realistic impulse responses to shocks, particularly in the first few periods after a shock. Overall across model validation exercises approach (3) that estimates the behavioural parameters jointly with the other dynamic and shock parameters and includes survey expectations data performs best.

Approaches (3) and (4) estimate the behavioural parameters jointly with the other dynamic and shock parameters. When survey expectations data are included (approach (3)), the estimated behavioural parameters are consistent with the empirical evidence of Kohlhas and Walther (2021) of simultaneous overreaction to current conditions and underreaction to news. Both approaches, however, call for a considerably greater degree of stickiness than the micro empirical evidence suggests, in line with the stickiness parameters estimated in macro studies such as Auclert et al. (2020) and Albuquerque et al. (2025). Despite this, approaches (3) and (4) capture reasonably well the dynamic response of household and professional inflation forecast errors relative to the data, indicating that the higher stickiness parameters are not inconsistent with the final combination of model, data and dynamic empirical forecast errors.

In terms of the other dynamic and shock parameters, we do not find large differences between the FIRE and behavioural specifications. Perhaps disappointingly the behavioural models do not significantly reduce the model's reliance on exogenous shocks to explain the data. We also note that behavioural frictions do not map to reductions in other real and nominal frictions. This suggests that the behavioural frictions are not simply substitutes for other frictions in the model but rather a complementary source of amplification and persistence, with the determination of dynamic and behavioural parameters reflecting the deep interaction between the model's structure, its impulse responses, and the data.

Approaches (2) and (3), which embed stickiness and extrapolation, generate amplified, delayed, and persistent impulse responses relative to FIRE. The extrapolation channel increases the peak response of GDP and inflation to a monetary policy shock significantly relative to the purely sticky case.[3] As most HANK models are not fully estimated on business cycle data or indeed with the behavioural frictions that allow it to best match broader macrodynamics, in a final exercise we show that the impulse response properties of the model estimated on US business cycle data (1978–2018) are more consistent with the estimated empirical impulse responses to monetary policy based on the Romer and Romer (2004) series than the more amplified, state-of-the-art impulse responses of Bauer and Swanson (2023).

1.1 Related Literature

This paper relates to literatures spanning behavioural macroeconomics, particularly the role of expectations and informational frictions, heterogeneous agent business cycle models, and expectation formation. Our main theoretical and technical contribution is to flexibly embed the general properties of these literatures into a mainstream macroeconomic solution method.

A growing body of empirical evidence documents systematic deviations from rational expectations in survey data and provides suggested micro-foundations for those deviations. Coibion and Gorodnichenko (2015) find evidence of information rigidity in the expectations formation process akin to

[3] By a factor of approximately 2–3 in the calibrated case of approach (2) and 1–2 in the estimated case of approach (3).

the sticky expectations of Mankiw and Reis (2003), showing that forecast errors are predictable using past information. Kohlhas and Walther (2021) document that households, firms, and professional forecasters simultaneously underreact to news and overreact to current conditions. Bordalo et al. (2018) propose diagnostic expectations where agents overweight recent news. Adam et al. (2024) studies the household response to house price movements and documents initial underreaction followed by overreaction, explaining this through extrapolative expectations. Importantly, Angeletos et al. (2021) document a new stylised fact that unifies much of this seemingly conflicting evidence: in response to the main shocks driving the business cycle, expectations initially underreact but subsequently overshoot actual outcomes. They show that this pattern of dynamic overshooting speaks in favour of models combining dispersed, noisy information with overextrapolation. We formally show that the developed expectations formation framework (i) nests under-reactive expectations models, and (ii) contrary to the latter is able to help sustain dynamic overshooting on the part of agents—a critical feature informing the map from aggregate shock processes to equilibrium outputs through the SSJ implementation.

Our approach is implemented in the sequence space using the methods of Auclert et al. (2021). Auclert et al. (2020) and Bardóczy and Guerreiro (2024) show how this solution method can be extended to accommodate departures from FIRE: Auclert et al. (2020) implement sticky expectations, while Bardóczy and Guerreiro (2024) provide a general treatment of how the sequence-space method can be applied under deviations from FIRE, combining noisy information, diagnostic expectations, and long-term memory to study the expectations channel in the design of unemployment insurance. As in their work, our framework maintains the internal rationality of Adam and Marcet (2011). We build on their generalisation by mapping it to the reduced-form specification of Kohlhas and Walther (2021), which we argue provides a specific, empirically grounded, yet sufficiently flexible departure from rational expectations that approximates the leading alternative models as special cases. We further derive the behavioural block Jacobian in closed form directly from the fake-news matrix of Auclert et al. (2021), yielding a fast and modular solver that delivers the behavioural solution block by block.

Angeletos et al. (2025) study the transition from RANK to HANK models without the assumption of full information rational expectations, and similarly try to distil multiple expectations processes into a framework that captures the essence of each. Their work highlights how the interaction between household heterogeneity and departures from FIRE shapes the transmission of aggregate shocks, emphasising that the quantitative implications of HANK models depend critically on the assumed expectations process. Our paper complements their analysis by providing an estimated quantitative framework that disciplines the expectations process with survey data, and by showing how the specific combination of sticky expectations and extrapolation affects both the amplification and persistence of impulse responses in HANK models.

Relating more broadly to the HANK literature, Auclert et al. (2020) argue that departures from rational expectations are key to simultaneously matching the macro and micro evidence in the calibration of HANK models. They rely upon sticky expectations in their application, which we compare to both quantitatively in terms of model estimation against both empirical IRFs and US business cycle data and theoretically, as the proposed behavioural expectations model robustly nests and extends their focus on macro-inattention by incorporating extrapolation from current conditions as a source of further bias, above and beyond the Calvo rigidity, whenever agents update their information sets. Albuquerque et al. (2025) directly leverages the methodology outlined in this paper in order to match the simultaneous response of house prices and rents to a monetary policy shock in a HANK model, where they find sticky expectations insufficient to generate the requisite amplification and hump-shaped dynamics, and instead require a combination of sticky expectations and extrapolation from current conditions to match the empirical evidence.

Finally, this paper contributes to a growing literature that uses survey expectations data as observables in the Bayesian estimation of DSGE models. Pioneering work by Del Negro and Eusepi (2011) showed that including SPF inflation expectations among the observables provides strong discriminating power between alternative models of expectations formation, significantly improving model fit and helping to pin down the behaviour of latent expectation processes. Del Negro and Schorfheide

(2013) demonstrated that incorporating survey forecasts into DSGE model estimation improves out-of-sample forecasting performance. Slobodyan and Wouters (2012) estimate a medium-scale DSGE model with adaptive learning and find that the beliefs implied by learning are more closely related to survey evidence on inflation expectations than those under rational expectations, while also reducing the estimated persistence of exogenous shock processes. Our contribution to this literature is twofold. First, we estimate a HANK model in the sequence space—rather than the state-space representative agent models that dominate the existing literature on estimation with survey observables. Second, we allow for differences in expectations formation between households and firms—this directly leverages for identification purposes the modular nature of our framework and SSJ solver, allowing for asymmetric departures from the FIRE benchmark within different blocks of the DSGE model.

2 Generalised Expectations Framework

Our framework captures the key insight from Kohlhas and Walther (2021) that agents may simultaneously over or underreact to current conditions while underreacting to news about the future. We do so using a small number of reduced-form parameters. We start from the reduced-form regression equations of Kohlhas and Walther (2021) that express the forecast errors of agents about an input x, k periods ahead, as a function of observations of x today and news about future x.[4]

$$x_{t+k} - f_{it}x_{t+k} = \alpha_i + \gamma x_t + \delta(\bar{f}_t x_{t+k} - \bar{f}_{t-1}x_{t+k}) + \varepsilon_{i,t|t+k} \tag{1}$$

The parameter γ captures extrapolation from current conditions (the "Extrapolation" effect), while δ captures the reaction to news (the "News" effect). In empirical applications statistically significant differences from zero in either γ or δ are a rejection of the rational expectations hypothesis. Studies such as Coibion and Gorodnichenko (2015) consistently find that δ is statistically significant and greater than zero even for professional forecasters, indicating a tendency to underreact to news about future economic conditions. Kohlhas and Walther (2021) further demonstrate that not only is $\delta > 0$ consistently across surveys and agent types, but simultaneously γ is consistently and statistically significantly less than zero, indicating that agents overreact to x_t today while underreacting to news about future x_{t+k}. Coibion and Gorodnichenko (2015) show that for the case $\gamma = 0$, the reduced form of equation 1 is consistent with either sticky expectations, where agents update their forecasts infrequently, or with a noisy information setup, where agents observe x with measurement error. For the more general case $\gamma \neq 0$, Kohlhas and Walther (2021) show that equation 1 is consistent with asymmetric attention, whereby agents pay different levels of attention to different components of news about x.

By remaining agnostic about the precise micro-foundation, our reduced-form approach is flexible enough to nest several prominent models from the literature. Beyond the sticky information and noisy information frameworks discussed by Coibion and Gorodnichenko (2015), the specification accommodates Bordalo et al. (2018)'s diagnostic expectations — where agents overweight states made more likely by recent data — through γ, which reflects the overweighting of salient current information. The framework is also consistent with extrapolative expectations, where agents project recent trends into the future. Where γ and δ are allowed to vary by input x, it can also approximate rational inattention models, where agents optimally choose to remain uninformed about certain aspects of the economy due to information processing costs. Finally, by allowing δ to be time-varying we can implement cognitive discounting, where less weight is given to news further ahead in time. This flexibility allows the specification to remain agnostic about the precise source of departures from rational expectations, while capturing the key asymmetric patterns in attention and updating that characterise different expectations processes.

[4] As is typical in this literature, news about future x is proxied empirically by the average forecast $\bar{f}$ of x in the current and previous period.

2.1 Implementation in the Sequence Space and the Behavioural DSGE Solver

We now describe how to map the reduced-form equation 1 into the sequence space so that we can apply the framework of Auclert et al. (2021) under this more general expectations process. For ease of exposition we focus on the case of constant δ and γ, which can be relaxed as shown in Appendix A.5. These parameters can also vary by input x as required or supported by the empirical evidence.

2.1.1 Behavioural Forecast and Forecast Mappings

The first step is to derive a forecasting equation for x. We take expectations[5] over equation 1, assuming no correlation of $\varepsilon_{i,t|t+k}$ across agents. Rearranging yields the average k-period-ahead forecast of x:

$$\bar{f}_t x_{t+k} = c + \frac{1}{1+\delta}\left(\delta \bar{f}_{t-1} x_{t+k} + E_t^*[x_{t+k}] - \gamma x_t\right) \tag{2}$$

where $\bar{f}_t x_{t+k} = E[f_{it} x_{t+k}]$ and E^* is the expectations operator conditional on the full information set for x. Since in the sequence space this information set coincides with the time-zero one, this is equivalent to perfect foresight, and we suppress the operator in the derivations below. We additionally assume: (1) agents have unbiased expectations in the steady state, i.e. $c = \frac{\gamma}{1+\delta} x_{ss}$; (2) for some large T, agents expect a return[6] to the steady state x_{ss}; and (3) a law of iterated expectations holds for agents, so that they do not expect to make forecast errors. Together these assumptions ensure agents are internally rational in the sense of Adam and Marcet (2011), as set out in appendix A.2. Agents optimise given subjective beliefs that are unbiased and stationary in the long run ((1)–(2)) and subjectively consistent in that agents do not forecast their own errors (3). The friction therefore resides entirely in how the subjective measure is formed—through infrequent updating and extrapolation—rather than in any failure of agents to act optimally given it. In particular, assumption (3) holds with respect to each agent's own information set and subjective measure, and is fully consistent with forecast errors that are predictable under the objective DGP: the over- and under-reaction induced by γ and δ are visible to the econometrician but, by construction, orthogonal to agents' own projections. Under these assumptions it can be shown in appendix A.1 that, at a first-order perturbation around the steady state, the law of motion for the behavioural forecast is:

$$\bar{f}_t dx_{t+k} = \frac{1}{1+\delta}\left(\delta \bar{f}_{t-1} dx_{t+k} + dx_{t+k} - \gamma dx_t\right).$$

This can be rewritten as:

$$\bar{f}_t dx_{t+k} = \overbrace{\frac{\delta}{1+\delta}}^{\text{Stickiness: } P(\Omega_t = \Omega_{t-1})} \bar{f}_{t-1} dx_{t+k} + \overbrace{\frac{1}{1+\delta}}^{\text{Pr. update}} \left(dx_{t+k} - \gamma dx_t\right).$$

The forecast at time t is a weighted average of the previous period's forecast — arising from Calvo-style stickiness with weight $\delta/(1+\delta)$ on information updates — and the innovation made by the fraction $1/(1+\delta)$ of agents who update at the start of this period. Contrary to the purely sticky case, updating agents still make a forecast error through extrapolation, $-\gamma dx_t$, from the economic conditions prevailing when they update. We can express the same dynamics in terms of a law of motion for the forecast error on aggregate input x at some fixed future period $t+k$. The forecast error is:

$$h_t^{t+k} = dx_{t+k} - \bar{f}_t dx_{t+k}.$$

Using the law of motion for the forecast, this implies, for all t, k:

$$h_t^{t+k} = \left(\frac{\delta}{1+\delta}\right) dx_{t+k} + \frac{1}{1+\delta}(\gamma dx_t - \delta \bar{f}_{t-1} dx_{t+k})$$

[5] The natural application is to average over all households, though the same logic applies to distinct groups as long as expectation differences are permanent. This is consistent with the block-level aggregation in Auclert et al. (2021), which combines policy function choices with distributional dynamics.

[6] This implies that equation 2 holds for $k < T$ and that the model must converge to x_{ss} in general equilibrium.

$$h_t^{t+k} = \left(\frac{\delta}{1+\delta}\right) dx_{t+k} - \frac{\delta}{1+\delta} dx_{t+k} + \frac{1}{1+\delta}(\gamma dx_t + \delta h_{t-1}^{t+k})$$

$$\underbrace{h_t^{t+k}}_{\text{posterior forecast error}} = \frac{\delta}{1+\delta} \underbrace{h_{t-1}^{t+k}}_{\text{prior forecast error}} + \frac{1}{1+\delta} \underbrace{\gamma dx_t}_{\text{extrapolation}} \tag{3}$$

The forecast error at time t is a weighted average of the prior forecast error and extrapolation from current conditions (γdx_t), with weight $\delta/(1+\delta)$. The forecast error therefore evolves as a weighted average of past forecast errors and current extrapolation — a structure analogous to a filtering algorithm from the econometrician's perspective.

Whether extrapolation or the prior forecast error dominates depends on the degree of stickiness: as $\delta \to \infty$, the forecast error law of motion approaches a unit root process, and extrapolation from current conditions ceases to matter. More formally, modelling the updating process as i.i.d. Bernoulli trials with success probability $(1-\theta)$, where

$$\delta = (1-\theta) \sum_{s=0}^{\infty} \theta^s s = \frac{\theta}{1-\theta}$$

is the expected number of periods before agents update,[7] the law of motion for the forecast error can equivalently be written as:

$$h_t^{t+k} = \theta h_{t-1}^{t+k} + (1-\theta)\gamma dx_t. \tag{4}$$

The two behavioural channels interact in complementary ways. The news-revision parameter δ sustains inertia in forecast revisions over time, while extrapolation (with $\gamma < 0$, i.e. overreaction to current conditions) makes the forecast error depend pro-cyclically on current conditions. This interaction is key to the model's ability to generate stronger endogenous amplification at similar levels of inattention, while preserving the hump-shaped dynamics familiar from inattentive-agent models. Finally, for bounded positive δ (so that θ lies within the unit circle), the forecast error converges to zero along any non-stochastic path where aggregate shocks have settled to zero (see Assumption 4 in Appendix A.2).

2.1.2 Backward solution and behavioural forecast mapping

To derive the general equilibrium SSJ solver, as in Bardóczy and Guerreiro (2024), we construct a mapping $\mathcal{E}_\tau : d\mathbf{X} \to X^{e,\tau} = \Lambda_\tau d\mathbf{X}$ from the space of true aggregate input sequences $d\mathbf{X}$ to the space of time-τ expectations on the path of inputs $d\mathbf{X}^{e,\tau}$. Starting from the steady state, we trace how the forecast of $x_k = x_{ss} + dx_k$ evolves as we move from period 0 (when shocks first hit) to period k:[8]

$$\bar{f}_0 x_k = x_{ss} + \frac{1}{1+\delta} dx_k - \frac{\gamma}{1+\delta} dx_0$$

$$\bar{f}_1 x_k = x_{ss} + \left(\frac{1}{1+\delta} + \frac{\delta}{(1+\delta)^2}\right) dx_k - \frac{\gamma}{1+\delta} dx_1 - \frac{\gamma\delta}{(1+\delta)^2} dx_0$$

$$\bar{f}_t x_k = x_{ss} + \sum_{j=0}^{t} \left(\frac{\delta}{1+\delta}\right)^{t-j} \frac{1}{1+\delta} dx_k - \gamma \sum_{j=0}^{t} \left(\frac{\delta}{1+\delta}\right)^{t-j} \frac{1}{1+\delta} dx_j \tag{5}$$

For $t < k$, the first sum represents the news effect, which grows over time and asymptotically converges to perfect foresight for $\delta > 0$.[9] With bounded δ, so that there is always a positive probability of updating, almost-sure asymptotic rationality is equivalent to the set of never-updating households

[7] The number of periods s elapsed before a single household updates, and under a law of large numbers the fraction of agents who last updated s periods ago, follow the geometric distribution with failure probability θ.

[8] This is shown in Appendix A.1.

[9] Noting that for any t:

$$\frac{1}{1+\delta} \sum_{j=0}^{t} \left(\frac{\delta}{1+\delta}\right)^{t-j} = \frac{1}{1+\delta} \sum_{j=0}^{t} \left(\frac{\delta}{1+\delta}\right)^{j},$$

having measure zero.[10] The second sum captures the extrapolation effect: agents extrapolate from observed realisations of x up to period k, with greatest weight on the most recent observations. Forecasts are therefore linearly decomposable as a function of the history of x and news about future x.

This backward solution for the linearised forecast forms the basis for the mapping from actual aggregate input paths to time-varying expectations used to construct the sequence-space DSGE solver. Letting $d\bar{f}_t x_k = \bar{f}_t x_k - x_{ss}$, equation (5) can be rewritten as:

$$d\bar{f}_t x_k = \frac{1}{1+\delta}\sum_{j=0}^{t}\left(\frac{\delta}{1+\delta}\right)^{t-j}(dx_k - \gamma dx_j)$$

$$d\bar{f}_t x_k = (1-\theta)\sum_{j=0}^{t}\theta^{t-j}(dx_k - \gamma dx_j) \tag{6}$$

where we have used the identification of δ with the macro-inattention parameter θ from above. Similarly to Auclert et al. (2020), the average behavioural forecast can be viewed as the cross-sectional average of forecasts from agents who differ in the time they last updated, where θ^{t-j} is both the probability of not having updated for the j periods preceding t, and — under a law of large numbers — the fraction of agents who last updated at $t-j$. Relative to pure sticky expectations, agents in the behavioural framework who update their information sets incur a further bias from pro-cyclical extrapolation from current conditions at the time of updating, $-\gamma x_k$. Suppressing this extra bias, as noted in Coibion and Gorodnichenko (2015), yields the nested purely sticky case.

2.1.3 Generalisation of the Sequence-Space Jacobian to Linear Forecast Mappings

While the original SSJ framework relies on a FIRE assumption for constructing the Jacobian J of any given forward-looking agent block – such to identify inputs path with the perfect foresight equilibrium path – as Bardóczy and Guerreiro (2024) point out this can be viewed as a restriction of a more general case in which forecasts themselves enter as inputs in the forward looking agents' problems in an augmented space of sequences $(\mathbf{X}, \mathbf{Z}, \mathbf{U}) \to (\mathbf{X}, \mathbf{X}^e, \mathbf{Z}, \mathbf{U})$.[11] In theory, given the derived laws of motion for forecasts over the sequence space, one could thus always explicitly solve for an augmented Jacobian with respect to both realised inputs and forecasts $\hat{J} : (d\mathbf{X}, d\mathbf{X}^e) \to d\mathbf{Y}$. However, this strategy would prove soon computationally very expensive, particularly so when departures from FIRE involve multiple blocks, include heterogeneous agents one, or are time-varying – all of which are supported in the proposed framework. A brute force approach here would substantially slow down the solution and consequently hamper scalability to structural estimation through modern Bayesian methods.

Expectations revisions as iterative Fake-News shocks. We propose a solver that generalises the SSJ approach of Auclert et al. (2021) to accommodate departures from FIRE, operating under the assumptions of (i) internal rationality and (ii) the existence of a linear approximation $d\mathbf{X}^{e,h} = \Lambda_h d\mathbf{X}$ to the expectations process around the steady state mapping realised inputs $d\mathbf{X}$ to expectations $d\mathbf{X}^{e,t}$ in force at time t. The solver is based on the insight that, under internal rationality and particularly

it follows from the Neumann Series Lemma that

$$\lim_{t\to\infty}\frac{1}{1+\delta}\sum_{j=0}^{t}\left(\frac{\delta}{1+\delta}\right)^{t-j} = \frac{1}{1+\delta}\sum_{j=0}^{\infty}\left(\frac{\delta}{1+\delta}\right)^{j} = \frac{1}{1+\delta}\left(1-\frac{\delta}{1+\delta}\right)^{-1} = 1.$$

[10] This follows from the first Borel–Cantelli Lemma:

$$\sum_{j}^{\infty}P(A_j) < \infty \to P(\liminf_{j} A_j) \leq P(\limsup_{j} A_j) = 0,$$

where A_j is the set of households who fail to update for j consecutive periods and P is the associated probability measure.

[11] In Appendix A.3 we relate our derivation which builds on the Fake News matrix to the generalised derivation of Bardóczy and Guerreiro (2024).

the assumed law of iterated expectations, the behavioural Jacobian $\hat{J}$ can be expressed as a function of the Fake-News matrix[12] with respect to the realised or perfect-foresight inputs only. The key object is the **forecast unlearning** matrix $\mathcal{P}_h$, which computes the impact of the iterative unlearning of prior forecasts as new expectations are formed at time h. The core economic logic is threefold, exploiting the fact that because the fake-news matrix is a partial equilibrium object, it does not matter whether the shocks it maps are realised shocks or shocks to expectations:

1. Under internal rationality, at any time t shocks to future expectations formed at $s \geq t$ are unanticipated and therefore have no impact at t.

2. For the same reason, shocks to expectations at time t unanticipated at $t-1$ are equivalent to a new set of time-zero shocks. Because the end of period distribution is a pre-determined variable when these hit, their impact is through the adjustment of optimal policies only – that is, the first row of the Fake News matrix.

3. At time t, once off shocks to expectations at time $t-s$ for s $s = 1, ..., t$ has no impact on optimal policies at time t. However, they have an impact on current aggregates via the adjustment caused to the end of period distribution at $t-s$ – forecast unlearning is the same as a fake-news shock: expectations on shocks announced at $t-s$ are retracted or unlearnt at $t-s+1$. Because at time $t-s$ shocks to expectations at that time are like time-zero shocks, then the impact of unlearning them at $t-s+1$ is captured by the $(t-s)-th$ row of the fake-news matrix.

The overall intuition can be stated as follows: iterative forecasts revisions between successive periods $t-1, t$ are equivalent to iteratively unlearning the previously held forecast at the end of the previous period and learning a new forecast at the beginning of the current period. The impact can thus be decomposed into two channels: (i) the unlearning of prior forecasts, that is retracting at the beginning of time t expectations news previously announced at time $t-1$ – this is essentially a fake-news shock in Auclert et al. (2021) which propagates through the distributional dynamics; and (ii) the impact of newly formed forecasts, which propagate through the first row of the fake-news matrix as if they were time-zero shocks. The fake-news matrix therefore spans the mapping from shocks to expectations to the impact on aggregate outcomes.

Generalised SSJ solver. To build intuition more formally, consider the following stylised deviation from FIRE: agents make forecast errors in period zero but converge to FIRE immediately afterwards – the set of forecast matrices is thus $\{\Lambda_0, \mathbf{I}, ...\}$. With the initial distribution pre-determined on impact and internally rational expectations, agents do not anticipate reversion to FIRE in the next period and the response of aggregates is through the adjustment in optimal policies only: $dY_0 = \mathcal{Y}_0 = \mathcal{F}_0 \Lambda_0 d\mathbf{X}$ as in Auclert et al. (2020). By contrast, at the beginning of $t = 1$, the aggregate response is:

$$\begin{aligned} dY_1 = \frac{\partial dY_1}{\partial X_0^e} + \frac{\partial dY_1}{\partial X_1^e} &= \mathcal{F}_{1,0} dX_0^{e,0} + \mathcal{F}_{1,:} dX_{s\geq 1}^{e,0} + \mathcal{F}_{0,:-1} dX_{s\geq 1}^{e,1} \\ &= \underbrace{\mathcal{F}_{1,:}}_{d\mathbf{D}_1 | d\mathbf{X}^{e,0}} \Lambda_0 d\mathbf{X} + \underbrace{\begin{bmatrix} 0 & \mathcal{F}_{0,:-1} \end{bmatrix}}_{d\mathbf{y}_1 | d\mathbf{X}^{e,1}} d\mathbf{X}, \end{aligned}$$

Forecast revision between periods 0 and 1 is thus a mixture of two shocks: an unanticipated and hence time-zero like shock to optimal policies, and a fake-news shock to the beginning of period distribution equivalent to the impact of unlearning or retracting the previously held forecast. Suppose reversion to FIRE now occurs instead at $t = 2$, then:

[12] Recall from Auclert et al. (2020) that for some scalar-valued shock $dx \to 0$ of the s-th period input with all else kept the same, the (t,s) element of the fake-news matrix $\mathcal{F}$ is defined as (eq. 18):

$$\mathcal{F}_{t,s} dx = dY_{t,s} - dY_{t,s-1} = (\mathbf{y}_{ss}' \Lambda_{ss}')(d\mathbf{D}_{t,s} - d\mathbf{D}_{t-1,s-1}) = \mathbf{y_{ss}} (\Lambda_{ss}')^{t-1} d\mathbf{D}_{1,s}$$

where $d\mathbf{D}_{1,s}$ is the change in the end-of-period distribution at $t = 1$ from a shock to the s-th period input. Hence the (t, s)-th element captures the impact on the aggregate outcome at time t of learning at time zero a shock the s-th period input and unlearning it at the beginning of period 1.

$$d\mathbf{Y}_2 = \mathcal{F}_{2,:}\Lambda_0 d\mathbf{X} + \mathcal{F}_{1,:}\Lambda_1 d\mathbf{X} + [0 \quad 0 \quad F_{0,:-2}]d\mathbf{X}$$

Similarly to the above, the second two terms capture respectively the current impact of unlearning at time $t = 2$ the forecast formed at $t = 1$ and the impact of the newly formed forecast at $t = 2$. The newly appeared first term instead captures the sluggish propagation through the distribution of unlearning the time zero forecasts at $t = 1$. Carrying this forth this delayed reversion algorithm, in general the aggregate response at time t when expectations revert to FIRE in that period can be expressed as:

$$dY_t = \sum_{j=0} \mathcal{F}_{t-j,:}\Lambda_j d\mathbf{X} + [0'_t \quad F_{0,:}]d\mathbf{X}$$

Where the first term accumulates the impact of prior fake-news shocks – unlearning the one formed at $t-1$ plus the propagation of iteratively unlearning the $(t-s)$ forecast at $(t-s+1)$ – and the second term encodes the shock to optimal policies from the unanticipated forecast formed at t.

Because at any time t the impact of outcomes today only depends on expectations formed from time zero through to now, and not on future expectations under internal rationality, the above stylised case with sequentially delayed reversion to FIRE applies even when forecasts at t do not revert to the FIRE one – i.e. $d\mathbf{X}^{e,t} = \Lambda_t d\mathbf{X}$ with $\Lambda_t \neq \mathbf{I}$ at t. The above delayed reversion algorithm therefore recursively generates the requisite mapping $d\mathbf{X}^{e,0}, ..., d\mathbf{X}^{e,T} \rightarrow d\mathbf{Y}_t$ – stacking the expansions and writing them in compact form via the matrix $\mathcal{P}_h$ and differentiating with respect to $d\mathbf{X}$ yields the generalised Jacobian $\hat{J}$ as a function of the fake-news matrix and the forecast matrices Λ_h:

$$\hat{J} = \sum_{h \geq 0} \mathcal{P}_h \Lambda_h. \tag{7}$$

where

$$\mathcal{P}_h = \begin{pmatrix} 0 & 0'_h \\ 0_h & \mathcal{F} \end{pmatrix}$$

is the forecast-unlearning matrix that slides the fake-news matrix $\mathcal{F}$ along the main-diagonal to the appropriate period h. The above solver breaks down the mapping from expectations to aggregate response into two components: (1) applied iteratively as we move forward in time (along the main diagonal), the first non-zero row of $\mathcal{P}_t$ maps impact of newly formed forecasts at time t, which propagate as if they were time-zero shocks; and (2) for any fixed t, each row $h < t$ of $\mathcal{P}_t$ captures the impact of prior forecasts formed at h that are unlearned at $h+1$, which propagate through the distribution as fake-news shocks.[13]

2.2 Generalised SSJ Solver with Behavioural Forecast Mapping

The forecast matrices Λ_h follow directly from the backward solution in equation 5. Each Λ_k has four distinct regions. The top-left region is the identity matrix $\mathbf{I}_{h+1}$, reflecting the fact that agents observe current and past inputs perfectly. Below it is $\mathbf{A}_h$, which captures the extrapolation effect: how each observed x (along the columns) affects forecasts of future X (down the rows), and persists but dampens ($\delta > 0$) through time. Each entry reflects the pass-through $-\gamma$ at the time agents last updated, weighted by $(\delta/(1+\delta))^{h-j}$, the probability they have not updated since. The bottom-right region is the news effect, a diagonal matrix $\mathbf{B}_h$ which captures how as we move forward in time from the shock, agents update their forecasts towards the FIRE value. $\mathbf{B}_h$ collects along its main diagonal the news-revision coefficients on the forecast input at time h, corresponding to the combined probability of agents incorporating news on the forecast input at times $t = 0, \ldots, h$.

[13]See Appendix A.3 for more details on the derivation of the generalised Jacobian $\hat{J}$.

$$\Lambda_0 = \begin{pmatrix} 1 & 0 & 0 & 0 & \cdots & 0 \\ \frac{-\gamma}{1+\delta} & \frac{1}{1+\delta} & 0 & 0 & \cdots & 0 \\ \frac{-\gamma}{1+\delta} & 0 & \frac{1}{1+\delta} & 0 & \cdots & 0 \\ \frac{-\gamma}{1+\delta} & 0 & 0 & \frac{1}{1+\delta} & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \frac{-\gamma}{1+\delta} & 0 & 0 & 0 & \cdots & \frac{1}{1+\delta} \end{pmatrix}$$

$$\Lambda_1 = \begin{pmatrix} 1 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & 0 & \cdots & 0 \\ \frac{-\gamma\delta}{(1+\delta)^2} & \frac{-\gamma}{1+\delta} & \frac{1}{1+\delta} + \frac{\delta}{(1+\delta)^2} & 0 & \cdots & 0 \\ \frac{-\gamma\delta}{(1+\delta)^2} & \frac{-\gamma}{1+\delta} & 0 & \frac{1}{1+\delta} + \frac{\delta}{(1+\delta)^2} & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \frac{-\gamma\delta}{(1+\delta)^2} & \frac{-\gamma}{1+\delta} & 0 & 0 & \cdots & \frac{1}{1+\delta} + \frac{\delta}{(1+\delta)^2} \end{pmatrix}$$

$$\Lambda_2 = \begin{pmatrix} 1 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 1 & 0 & \cdots & 0 \\ \frac{-\gamma\delta^2}{(1+\delta)^3} & \frac{-\gamma\delta}{(1+\delta)^2} & \frac{-\gamma}{1+\delta} & \frac{1}{1+\delta} + \frac{\delta}{(1+\delta)^2} + \frac{\delta^2}{(1+\delta)^3} & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \frac{-\gamma\delta^2}{(1+\delta)^3} & \frac{-\gamma\delta}{(1+\delta)^2} & \frac{-\gamma}{1+\delta} & 0 & \cdots & \frac{1}{1+\delta} + \frac{\delta}{(1+\delta)^2} + \frac{\delta^2}{(1+\delta)^3} \end{pmatrix}$$

For arbitrary h, the pattern generalises to:

$$\Lambda_h = \begin{pmatrix} \mathbf{I}_{h+1} & \mathbf{0}_{(h+1)\times(T-(h+1))} \\ \mathbf{A}_h & \mathbf{B}_h \end{pmatrix} \tag{8}$$

where the block matrices $\mathbf{A}_h$ and $\mathbf{B}_h$ are defined by:

$$\mathbf{A}_{h(i,j)} = \frac{-\gamma\delta^{h-j}}{(1+\delta)^{h-j+1}} \quad \text{for } j = 0, 1, \ldots, h$$

$$\mathbf{B}_{h(i,j)} = \begin{cases} \displaystyle\sum_{l=0}^{h} \frac{\delta^l}{(1+\delta)^{l+1}} & i = j \\ 0 & i \neq j \end{cases}$$

Asymptotic rationality requires that the forecast matrix satisfies $\lim_{h\to\infty} \Lambda_h = \mathbf{I}$, so that the forecast eventually converges to the perfect foresight (and hence FIRE) case. Based on an analogous argument to that in footnote (7), it is straightforward to see that the news-revision matrix $\mathbf{B}_h$ converges to an identity matrix:

$$\lim_{k\to\infty} \mathbf{B}_k = \mathbf{I},$$

while the extrapolation block matrix $\mathbf{A}$ converges to a matrix with zero entries:

$$\lim_{k\to\infty} \mathbf{A}_k = \mathbf{0},$$

which taken together with the definition of Λ_h imply the limit condition is satisfied.

2.3 General Equilibrium Solver

We now assemble the partial-equilibrium objects — the forecast mappings Λ_h and the behavioural Jacobian $\hat{\mathbf{J}}$ — into a general-equilibrium solution. The central observation is that behavioural expectations enter the solver *only* through the Jacobians of blocks that contain forward-looking expectations, and that these Jacobians enter the general-equilibrium computation as nothing more than the building blocks of a chain rule. We can therefore replace the FIRE Jacobian of any forward-looking block with its behavioural counterpart and then solve for general equilibrium by exactly the same steps as under FIRE.

A DSGE model can be written as a root-finding problem in the space of length-T sequences of endogenous aggregate unknowns $\mathbf{U} \in \mathbb{R}^{n\times T}$ and aggregate shocks $\mathbf{Z} \in \mathbb{R}^{m\times T}$:

$$\mathbf{H}(\mathbf{U}, \mathbf{Z}) = \mathbf{0},$$

where $\mathbf{H}$ denotes the target block (the equilibrium conditions that must hold in general equilibrium). As shown in Auclert et al. (2021), a first-order sequence-space solution is obtained by linearising around the steady state,

$$\mathbf{H}_U\, d\mathbf{U} + \mathbf{H}_Z\, d\mathbf{Z} = 0 \quad \Rightarrow \quad d\mathbf{U} = -\mathbf{H}_U^{-1}\mathbf{H}_Z\, d\mathbf{Z},$$

where $\mathbf{H}_U \equiv \partial\mathbf{H}/\partial\mathbf{U}$ and $\mathbf{H}_Z \equiv \partial\mathbf{H}/\partial\mathbf{Z}$ are the derivatives of the target conditions evaluated at the steady state. The $T \times T$ matrix

$$\mathbf{G} = -\mathbf{H}_U^{-1}\mathbf{H}_Z$$

is the general-equilibrium Jacobian, whose columns collect the linearised impulse response functions of aggregate unknowns to each aggregate shock.

To compute $\mathbf{G}$, Auclert et al. (2021) represent the model as a directed acyclic graph (DAG) of blocks, each mapping a set of inputs to a set of outputs and summarised by its *partial* Jacobians: for an output o and a direct input m of the same block, the partial (block) Jacobian $\mathbf{J}^{o,m}$ is the $T \times T$ matrix mapping dm into do holding all other inputs of that block fixed. The *total* Jacobian $\mathbf{J}^{o,i}$ of an output o with respect to a shock or unknown $i \in \mathbf{Z} \cup \mathbf{U}$ is then assembled by forward accumulation along the DAG: initialising $\mathbf{J}^{i,i} = \mathbf{I}$ and proceeding in topological order over blocks b,

$$\mathbf{J}^{o,i} = \sum_{m\in\mathcal{I}_b} \mathbf{J}^{o,m}\, \mathbf{J}^{m,i},$$

where $\mathcal{I}_b$ denotes the inputs of block b. This is simply the chain rule applied block by block: the factor $\mathbf{J}^{o,m}$ is the partial Jacobian of the current block, while $\mathbf{J}^{m,i}$ is the already-accumulated total Jacobian of its input m with respect to i (equal to the identity, and hence absent, when $m = i$). Collecting the total Jacobians of the target block delivers $\mathbf{H}_U = \mathbf{J}^{\mathbf{H},\mathbf{U}}$ and $\mathbf{H}_Z = \mathbf{J}^{\mathbf{H},\mathbf{Z}}$, and hence $\mathbf{G} = -\mathbf{H}_U^{-1}\mathbf{H}_Z$; the impulse responses of any other output are recovered by the same recursion, $\mathbf{G}^{o,Z} = \sum_{m\in\mathcal{I}_b} \mathbf{J}^{o,m}\, \mathbf{G}^{m,Z}$.

Within this representation, a forward-looking block is simply a node whose partial Jacobian $\mathbf{J}^{o,m}$ with respect to an input m embeds an expectation of the future path of m. Under FIRE that expectation is perfect foresight; under our framework it is the behavioural forecast generated by the mappings Λ_h. The partial Jacobian of any such block is therefore replaced by its behavioural counterpart $\hat{\mathbf{J}}^{o,m}$, constructed exactly as in Section 2.1. This construction transfers to any forward-looking block — not only the heterogeneous-agents block for which it is derived — because it relies only on the law of iterated expectations and the resulting fact that forecast revisions arrive as unanticipated, re-dated time-zero shocks, a property of the forecast mapping Λ_h rather than of any particular block. Because forward accumulation only ever multiplies and sums these partial Jacobians, substituting $\hat{\mathbf{J}}^{o,m}$ for $\mathbf{J}^{o,m}$ in any subset of blocks leaves every subsequent step — the total Jacobians, $\mathbf{H}_U$, $\mathbf{H}_Z$, and $\mathbf{G}$ — structurally unchanged. The general principle is that wherever a solver $\mathbf{G}(\mathbf{J})$ applies under FIRE, the behavioural solver $\mathbf{G}(\hat{\mathbf{J}})$ is obtained by swapping $\mathbf{J}$ for $\hat{\mathbf{J}}$ in those blocks where agents form expectations about future aggregates, and leaving all other blocks untouched.

Crucially, nothing in this argument is specific to the heterogeneous-agents block: the same swap applies to *any* forward-looking block, since each is just a node of the DAG carrying its own partial Jacobian. Whether that Jacobian is computed under perfect foresight or under behavioural expectations is a property of the block in isolation and does not interact with the composition rules that deliver $\mathbf{G}$. This block-local character of the substitution is also what lets us assign different behavioural parameters — different degrees of stickiness δ (or θ) and extrapolation γ — to different blocks of the model. Each block's $\hat{\mathbf{J}}$ is built from forecast mappings Λ_h that use that block's own parameters, so expectations can differ across sectors simply by constructing the corresponding blocks' Jacobians with distinct (δ, γ), while the general-equilibrium assembly proceeds unchanged.

3 Extrapolation, Overreaction, and Dynamic Overshooting

A key stylised fact documented by Angeletos et al. (2021) is that expectations initially underreact to aggregate shocks before subsequently overshooting actual outturns, and that this dynamic overshooting is critical for expectations frictions to drive amplification rather than merely introducing inertia. In this section we formalise the conditions under which dynamic overshooting arises in our behavioural expectations framework. We show formally that for a broad class of impulse responses extrapolation is a necessary condition for dynamic overshooting, and establish conditions under which it is also sufficient. [14] Pure stickiness ($\delta > 0, \gamma = 0$) produces forecasts error that vanish asymptotically – as the set of sequentially non-updating agents shrinks in measure to zero – as the average forecast approaches the perfect foresight path from below. While stickyness warrants under-reaction on impact, sluggish incorporation of news is never enough on its own to overshoot the outturn, making extrapolation necessary. In terms of sufficiency we require extrapolation that is sufficiently strong to eventually dominate the decaying stickiness component but not strong enough so as to dominate on impact. While the formal necessity results rely on some specific assumptions, the established conditions remain natural diagnostics in richer general equilibrium settings.

3.0.1 Formal Analysis of Dynamic Overshooting

We work with a generalised, parametrically hump-shaped impulse response of the form:

$$dx_t \;=\; x(t) \;=\; (1 + bt)\,\rho^t, \qquad \rho \in (0,1),\; b \geq 0, \tag{9}$$

normalised so that $dx_0 = 1$. This is a non-negative, single-peaked response: treating t as continuous, $x'(t) = \rho^t\big[b + (1+bt)\ln\rho\big]$ vanishes at

$$t^* = \frac{1}{|\ln\rho|} - \frac{1}{b}, \tag{10}$$

so x rises to an interior peak at t^* when $b > |\ln\rho|$ and declines monotonically from impact otherwise; the AR(1) response is the special case $b = 0$. In all cases $dx_t > 0$ for every t and $dx_t \to 0$.

Consider the k-period-ahead forecast error under behavioural expectations at a first-order perturbation, $h_{t,t+k} = dx_{t+k} - dx_t^e$, with the behavioural expectation formed from time zero,

$$dx_t^e = (1-\theta)\sum_{j=0}^{t}\theta^{t-j}\big(dx_{t+k} - \gamma\,dx_j\big), \tag{11}$$

where $\theta \in (0,1)$ captures inattention and $\gamma < 0$ captures extrapolation bias. Substituting and using the identity $1 - (1-\theta)\sum_{j=0}^{t}\theta^{t-j} = \theta^{t+1}$, the forecast error decomposes as

$$h_{t,t+k} = \underbrace{\theta^{t+1}\,dx_{t+k}}_{h^u_{t,t+k}} \;-\; \underbrace{|\gamma|(1-\theta)\sum_{j=0}^{t}\theta^{t-j}\,dx_j}_{h^o_{t,t+k}}, \tag{12}$$

[14] Specifically a class of linear hump-shaped IRFs $\{x(t) = (1+bt)\rho^t\}$ that do not overshoot (oscillate). We show in footnote 16 that oscillatory IRFs ($b < 0$) can generate a pattern consistent with dynamic overshooting, mimicking over-reactive extrapolation.

where $h^u_{t,t+k}$ is the stickiness (under-reaction) component and $h^o_{t,t+k}$ is the stickiness–extrapolation interaction. Under the functional form (9) these are, explicitly,

$$h^u_{t,t+k} = \theta^{t+1}\big(1+b(t+k)\big)\rho^{t+k}, \qquad h^o_{t,t+k} = |\gamma|(1-\theta)\,S_t, \quad S_t := \sum_{j=0}^{t} \theta^{t-j}(1+bj)\rho^j. \tag{13}$$

where S_t is the history of the economy through time t from which agents extrapolated at the time of their last update $j < t$, weighted by the relative probability of having last updated at that time. Note that $h^o_{t,t+k}$ depends on t but not on the forecast horizon k. We establish three results in turn: the necessity of extrapolation, an on-impact undershooting condition, and the eventual domination of extrapolation.

Proposition 1 (Necessity)**.** *Under (9), if $\gamma \geq 0$ then $h_{t,t+k} > 0$ for all finite t, k, so the forecast error exhibits no dynamic overshooting. Dynamic overshooting therefore requires $\gamma < 0$.*

Proof. If $\gamma = 0$ the interaction term vanishes and $h_{t,t+k} = h^u_{t,t+k} = \theta^{t+1}(1+b(t+k))\rho^{t+k}$, a product of strictly positive factors, hence > 0 for all finite t, k and $\to 0$ as $t \to \infty$. If $\gamma > 0$ the interaction term enters (12) with a positive sign, so $h_{t,t+k} > h^u_{t,t+k} > 0$. In neither case does the error cross zero from above, so there is no undershooting-then-overshooting pattern. □

Pure informational stickiness thus generates a forecast error that peaks on (or near) impact and decays to zero from above without ever overshooting. With $\gamma < 0$, dynamic overshooting requires that stickiness dominate at least on impact, so that the error begins in the undershooting (positive) region.

Proposition 2 (On-impact undershooting)**.** *Under (9) with $\gamma < 0$, the forecast error undershoots on impact, $h_{0,k} > 0$, if and only if*

$$|\gamma| \;<\; \frac{\theta}{1-\theta}\,\frac{dx_k}{dx_0} \;=\; \frac{\theta}{1-\theta}\,(1+bk)\rho^k. \tag{14}$$

Proof. Evaluating (12) at $t = 0$, where $\sum_{j=0}^{0} \theta^{-j} dx_j = dx_0$, gives $h_{0,k} = \theta\, dx_k - (1-\theta)|\gamma|\, dx_0$. Rearranging $h_{0,k} > 0$ yields (14), and $dx_0 = 1$, $dx_k = (1+bk)\rho^k$ give the explicit form. □

Extrapolation must therefore not be too strong relative to the persistence of the response at the forecast horizon: a longer horizon k (smaller dx_k) tightens the bound, requiring weaker extrapolation to preserve on-impact undershooting. The complementary requirement is that the interaction component eventually overtake the stickiness component, driving the error negative. Under the functional form this holds *for any* $|\gamma| > 0$, with no lower bound required.

Proposition 3 (Eventual domination)**.** *Under (9), for any $|\gamma| > 0$,*

$$h^o_{t,t+k} \;\longrightarrow\; m > 0 \qquad \textit{as} \qquad t \to \infty, \tag{15}$$

where m is bounded, and since $h^u_{t,t+k} \longrightarrow 0$ as $t \to \infty$, there exists a finite horizon $T(k)$ such that $h_{t,t+k} < 0$ for all $t \geq T(k)$.

Proof. The key observation is that S_t is a convolution of two convergent absolutely summable sequences: the geometric decay from $a_k = \theta^k$ and the dampened growth from $b_j = (1+bj)\rho^j$, with $k = t - j$. By the Cauchy product theorem,[15] the convolution S_t thus converges to a finite limit as $t \to \infty$. Since $h^o_{t,t+k} = |\gamma|(1-\theta)S_t$, it follows that $h^o_{t,t+k}$ converges to a finite limit as well. The limit is a positive constant. In particular, we have that:

$$\lim_{t\to\infty} \sum_{j=0}^{t} \theta^{t-j} = \frac{1}{1-\theta}$$

[15] Or convolution theorem.

$$\lim_{t\to\infty}\sum_{j=0}^{t}(1+bj)\rho^j = \lim_{t\to\infty}\sum_{j=0}^{t}\rho^j + \lim_{t\to\infty} b\rho\frac{d}{d\rho}\sum_{j=0}^{t}\rho^j = \frac{1-\rho+b\rho}{(1-\rho)^2}$$

where the second equality uses the chain rule to differentiate the geometric series. Therefore, we have that:

$$h^o_{t,t+k} \longrightarrow |\gamma|\frac{1-\rho+b\rho}{(1-\rho)^2} > 0 \qquad \text{as} \qquad t\to\infty.$$

□

Figure 1: Dynamic overshooting with a hump-shaped impulse response

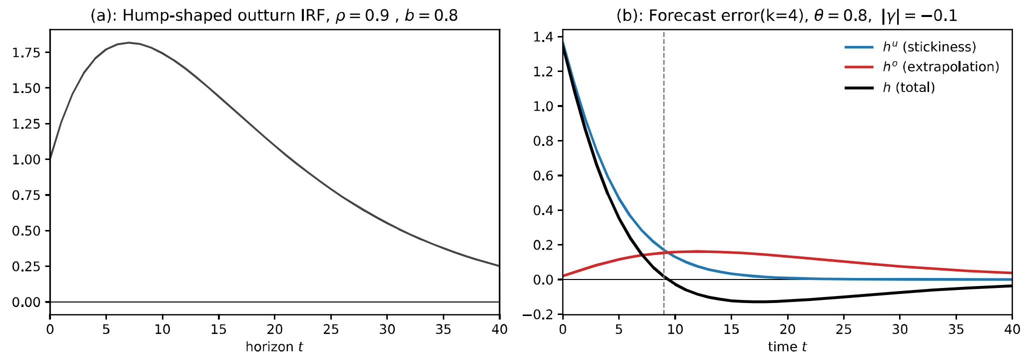

Note: The impulse response is of the form $x(t) = (1+bt)\rho^t$. The forecast error in panel (b) initially undershoots as the stickiness component (blue line) dominates, then overshoots as the accumulated extrapolation component (red line) eventually dominates.

Figure B.6 provides a visual illustration of propositions (1)-(3) against numerical experiments. The mechanism is a comparison of decay rates. The stickiness component measures the gap between the sticky forecast and the specific future outturn dx_{t+k}; this gap is discounted *twice*—once by inattention, θ^{t+1}, and once by the response's own decay, ρ^{t+k}—and therefore vanishes at the fast compound rate $\theta\rho$. The extrapolation component is different: it is a stock of geometrically-weighted memory of the entire realised history, $S_t = \sum_{j\le t}\theta^{t-j}dx_j$. As long as the agent keeps observing the (fading) response, that stock keeps being replenished, and it erodes only at the slower of the two rates that govern it—the memory weights θ or the observations ρ—that is, at $\max(\theta,\rho)$. Since $\theta\rho < \max(\theta,\rho)$ strictly, the slowly-eroding memory stock inevitably overtakes the doubly-discounted forecast gap.

On impact the memory stock is negligible—only dx_0 has been accumulated—while the stickiness gap is at its largest, so stickiness dominates and the error is positive; at long horizons the ranking of the two decay rates takes over and the error turns negative. The strength of extrapolation $|\gamma|$ scales the memory stock and so sets when the crossover occurs—a smaller $|\gamma|$ lowers the entire h^o path and delays $T(k)$—and how deep the overshoot runs. The hump parameter b enters only through the polynomial factor, which is asymptotically irrelevant to this exponential race, though it shifts weight onto h^o during the response upswing and thereby pushes the extrapolation peak, and the zero-crossing, to later interior horizons.

Figure 1 illustrates the mechanism. The stickiness component $h^u_{t,t+k}$ peaks on (or near) impact and decays. The extrapolation component $h^o_{t,t+k}$ starts small, rises through the response upswing, peaks at an interior horizon $\tilde{t}$, then decays. The total forecast error starts positive (stickiness dominates) and crosses zero when accumulated extrapolation overtakes the decaying stickiness, generating the characteristic undershooting-then-overshooting pattern.[16]

[16] As noted, the necessity of extrapolation rests on the assumption of a non-oscillatory IRF. If the outturn IRF dx_t is oscillatory, the assumption of non-negativity is violated and with it convergence to zero of the sticky component of the forecast error from above. We thus point out that oscillatory IRFs – generated by $b < 0$ in the above family of processes – can warrant a degree of dynamic overshooting even in the absence of extrapolation. Formally, under the assumption

4 Application

As an application of the developed framework, we study the two-asset HANK model of Auclert et al. (2020), which emphasised the importance of departures from rational expectations in HANK models and demonstrated how to implement sticky expectations in the sequence space. This is a natural model to build on given its focus on sticky expectations. Our application also allows us to assess the validity of the parameter estimates in that paper by estimating all parameters jointly on business cycle data, rather than their two-stage approach using IRF matching. Furthermore, as discussed by Kaplan et al. (2018), in HANK models households place greater weight on expectations of a wider array of input prices, whereas in RANK models the real interest rate dominates in determining consumption. We first replicate the model's steady state and confirm we can replicate its macrodynamics (see Table 1 and Figure B.1). We then estimate the model's dynamic, behavioural, and shock process parameters using a Bayesian full-information maximum likelihood approach on quarterly US business cycle data, including time series on inflation expectations from the Survey of Professional Forecasters and the Michigan Survey.

4.1 Two-Asset HANK Model

The two-asset heterogeneous agent New Keynesian (HANK) model of Auclert et al. (2020) was designed to simultaneously match micro-level consumption dynamics and macro-level impulse responses. The model features rich household heterogeneity replicating key features of the US distribution of disposable income and illiquid asset wealth, nominal rigidities in price and wage setting, an aggregate investment sector with time-to-build and adjustment costs, long-term government debt to replicate US maturity structures, and a quarterly calibration targeting both key US aggregates and the micro-structure of intertemporal marginal propensities to consume (iMPCs). The following section outlines the main features of the model and its steady-state solution.[17]

Heterogeneous households. Households differ in two main respects. First, there is ex-post heterogeneity in exogenous idiosyncratic labour productivity $z_{i,t}$, assumed to follow a Markov process with transition matrix Π, and in liquid asset positions $a_{i,t}$, as a consequence of uninsurable idiosyncratic risk and ad hoc borrowing constraints $a_{i,t} \geq -\bar{\phi}$ with $\bar{\phi} = 0$. Second, there is ex-ante heterogeneity in illiquid savings. Specifically, agents permanently belong to one of six non-overlapping groups g with measure μ_g satisfying $\sum \mu_g = 1$. Groups differ in average labour productivity $\bar{e}_g$, discount factor β_g, and steady-state illiquid asset holdings $\bar{a}_g$. Agents within the same group are identical in these respects and differ only ex-post along the idiosyncratic productivity and liquid asset dimensions. As shown in Auclert et al. (2020), this structure yields a tractable two-asset HANK model that replicates empirical inequality in the illiquid asset distribution while remaining computationally inexpensive — a key requirement for the estimation procedures in our application, which require repeated recomputation of the model Jacobians.

$\gamma = 0$ so that $h_{t,t+k} = h^u_{t,t+k}$, under $b < 0$ one can always write the sticky forecast error as:

$$h^u_{t,t+k} = \underbrace{\tilde{h}^u_{t,t+k}}_{\text{Pure AR component}} - \underbrace{|b(t+k)|\theta^{t+1}\rho^{t+k}}_{\text{Hump contribution}}$$

where $\tilde{h}^u_{t,t+k} = \theta^{t+1}\rho^{t+k}$ is the sticky forecast error associated to the nested pure AR component. Comparing with (12), the above representation makes it clear that a sufficiently strong hump-parameter b can mimic a pattern in forecast error similar to what obtains under over-reactive extrapolation. Contrary to the latter, this is not the result of a the stock of procyclical biases from staggered updates weighing in on the average forecast – instead it reflects sluggish incorporation into agents information set of news of a (dampened) downward trend. Based on the above representation, it is clear that for a given forecast horizon $h^u_{t,t+k} > 0$ as long as

$$t < t^* = \frac{1}{|b|} - k.$$

Clearly, as $b \to 0^-$ which retrieves the nested pure AR case, the $t^* \to \infty$, implying no dynamic overshooting occurs at any bounded time as expected. See Figure B.7 in the Appendix.

[17]For a more detailed overview, we refer the reader to Auclert et al. (2020) and Appendix B.1.

In each period, a group-g household with preferences $f(c,n) = u(c) - v(n)$, separable over consumption and labour, observes its productivity realisation $\bar{e}_g z_t$ and, given tax rate τ_t, its real disposable labour income $(1-\tau_t)w_t N_t \bar{e}_g z_t$ from supplying labour $n_t = N_t$ inelastically to meet union demand (allocated uniformly across households), the value of its liquid asset portfolio $(1+r_t^p)a$, and the distribution from the illiquid portfolio $d_{g,t}$ received at the start of the period. The household chooses consumption and its end-of-period liquid asset position to maximise the expected present value of its utility stream, subject to budget and borrowing constraints. Recursively, this amounts to the solution the Bellman equation:

$$V_{g,t}(a,x,z;\mathbf{X}_t) = u(c_t) - v(N_t) + \beta_g \mathbb{E}_{z'|z} V_{t+1}(a',x',z';\mathbf{X}_{t+1})$$
$$c + a' = (1-\tau_t)w_t N_t \bar{e}_g z_t + (1+r_t^p)a + d_{g,t}$$
$$a' \geq 0$$
$$x' = (1+r_t^{ill})x - d_{g,t}$$

where $\mathbf{X}_t = [w_t, r_t^p, \tau_t, N_t, d_{g,t}]'$ is the relevant aggregate state vector for group-g households, and the distribution from the illiquid portfolio is governed by the rule (for $\chi \to 0^+$), r_t^p and r_t^{ill} are the return on liquid assets held against the financial intermediary as outlined below:

$$d_{g,t} = \frac{r_{ss}^{ill}}{1+r_{ss}^{ill}}(1+r_t^{ill})x + \chi[(1+r_t^{ill})x - (1+r_{ss}^{ill})\bar{x}_g].$$

As in Auclert et al. (2020), for the quantitative estimation exercises we further allow for stickiness in attention to the portfolio value $(1+r_t^{ill})x$, as outlined in Appendix B.2.1. At the non-stochastic steady state the distribution policy consumes the return on the stationary illiquid portfolio (the annuity value of gross returns):

$$d_g = r_{ss}^{ill}\bar{x},$$

which combined with the distribution rule implies:

$$x = \bar{x}_g.$$

Agents therefore save when the illiquid portfolio value is low relative to steady state and dissave when it is high, with illiquid holdings converging to the group-specific mean $\bar{x}_g$. Since the distribution policy and illiquid asset position are invariant over the household idiosyncratic state-space conditional on belonging to the same group, the latter can be treated as an aggregate state variable chosen by a representative agent for each group – at any time t for a household in group g we hence let the end of period position $x' = X_{g,t}$.

Mutual funds and financial investment. Financial investment is carried out by a perfectly competitive mutual fund that issues liquid and illiquid accounts to households. Its liabilities are household liquid and illiquid account balances borrowed at rates r_t^p and r_t^{ill}; its assets include equities in the capital-building and intermediate firm (with aggregate supply normalised to one), government bonds, and central-bank reserves (in zero net supply). The intermediary faces linear costs on liquid accounts, which causes the liquid return to lie strictly below the illiquid return in equilibrium. The fund chooses its portfolio composition to maximise expected profits subject to the balance sheet constraint,[18] yielding a set of asset-pricing conditions and a Fisher equation describing the pass-through from monetary policy to asset prices.

Firms. The production side features nominal rigidities in both price setting by intermediate producers (with Cobb-Douglas technology using labour and capital) and wage setting by a representative union that sells labour services to a competitive labour packer. This combination of sticky prices and sticky wages captures the sluggish adjustment of the price level and the feedback between price and wage inflation. The investment sector consists of capital-building firms subject to time-to-build

[18]Equivalent by construction to asset market clearing.

constraints and Rotemberg-style quadratic adjustment costs, generating realistic and gradual investment dynamics. In particular, these frictions imply that investment responds slowly to a time-zero monetary policy shock even under FIRE, so that its income effects on household budgets are sustained over time, acting as a source of amplification. As Auclert et al. (2020) note, hump-shaped investment and income dynamics are not sufficient on their own to match the hump-shaped response to monetary policy shocks; departures from FIRE are also needed.

Fiscal and monetary policy. The government issues long-term debt with a maturity structure matching US Treasury data, allowing bond prices to respond meaningfully to monetary policy shocks. Tax revenue is collected according to a simple fiscal rule under which the tax rate adjusts gradually to return real government debt to its steady-state value. Monetary policy follows a standard Taylor rule responding to inflation and output.

Calibration. The quarterly calibration follows the strategy in Auclert et al. (2020). It targets aggregate moments of the US economy — wealth, rates of return, and bond maturity structures — and, via between-group differences in illiquid asset holdings and discount factors, the micro-structure of consumption responses to income shocks (iMPCs) and the illiquid asset distribution. This positions the model to address the dual challenge of matching micro-jumps and macro-humps simultaneously. The resulting steady state is reported in Table 1.

Table 1: Model Calibration and Estimation

Panel A: Calibrated Parameters					
Parameter		Value	Parameter		Value
ν	EIS	1	α_h	Labour share	0.76
ς	Frisch	0.5	δ_K	Depreciation (p.a.)	0.053
r	Real interest rate (p.a.)	0.050	μ_p	Price markup	1.06
K/Y	Capital to GDP (p.a.)	2.23	L/Y	Liquid assets to GDP (p.a.)	0.23
ξ	Intermediation spread (p.a.)	0.065	G/Y	Spending-to-GDP	0.16
qB/Y	Gov. bonds to GDP (p.a.)	0.42	$\frac{1+r}{1+r-\delta}$	Debt maturity (a.)	5
ψ	Tax response to debt (p.a.)	0.1	ϕ_π	Taylor rule coefficient	1.5

Panel B: Household Heterogeneity						
Group g	1	2	3	4	5	6
Description	Bottom 50%	Next 20%	Next 10%	Next 10%	Next 5%	Top 5%
Population share (μ_g)	0.500	0.200	0.100	0.100	0.050	0.050
Illiquid asset share	2.7%	7.0%	7.0%	13.0%	12.2%	58.0%
Labor income share	26.7%	18.3%	10.8%	14.4%	11.0%	18.8%
Discount factor β_g (p.a.)	0.905	0.919	0.933	0.946	0.950	0.975

4.2 Business Cycle Estimation

We estimate the model's dynamic parameters on US business cycle data. As outlined in the introduction, this allows us to explore the quantitative implications of the proposed expectations framework for structural estimation in a state-of-the-art medium-scale HANK model. A set of secondary exercises — documented in the following sections — complement the analysis by examining model performance along three dimensions: (i) the impulse responses to monetary policy shocks; (ii) the dynamic pattern of forecast errors; and (iii) broader model fit through root mean squared errors (RMSE) and autocorrelation functions (ACF) relative to the data.

We conduct a full-information maximum likelihood estimation of the model's dynamic, shock, and behavioural parameters using nine quarterly macroeconomic time series over 1978–2018: GDP, consumption, business investment, hours worked, real wages, value-added inflation, the federal funds rate, professional forecaster inflation expectations, and household inflation expectations. The first seven series are those used in Smets and Wouters (2007) and Auclert et al. (2020), transformed as in Smets and Wouters (2007) and described in the appendix (see Figure B.2). The inflation expectation series are transformed into time series of forecast errors by subtracting the actual year on year inflation outturn from the forecasted value one year prior. These series are also linearly detrended.[19]

To estimate the model we include seven exogenous shock processes (following Auclert et al. (2020)) and two measurement error terms for the two inflation expectation series, with prior means for the measurement errors set to the bootstrapped quarterly standard deviation of each survey's quarterly mean. The seven shock processes are all AR(1), with the price and wage markup series additionally including an MA(1) term.[20] These exogenous processes add sixteen shock process parameters to the estimation. We additionally estimate four dynamic parameters governing real and nominal rigidities: the investment adjustment cost parameter ψ, Taylor rule inertia ρ_m, and price and wage update probabilities ζ_p and ζ_w. Priors for these parameters are taken from Auclert et al. (2020), which in turn largely follow Smets and Wouters (2007). In Auclert et al. (2020) these parameters were estimated by IRF matching before estimating the shock parameters in a second step on the business cycle data. We instead estimate all parameters jointly in a single full-information maximum likelihood step, which is computationally more demanding but allows the data to inform all parameters simultaneously, and allows the inflation expectation series to discipline the behavioural parameters.[21]

Finally, we introduce four behavioural parameters $\{\theta_{hh}, \theta_{prof}, \gamma_{hh}, \gamma_{prof}\}$ capturing household and firm information stickiness and extrapolation tendencies.[22] [23] As in Auclert et al. (2020), these information frictions apply equally to all household inputs $\{w, r^p, \tau, N, d_g\}$; for firms and financial markets they apply to the relevant inputs in the forward-looking equations for price inflation, investment, and dividends. To help identify the behavioural parameters, we include separate equations for four-quarter-ahead inflation forecasts by households and firms, incorporating inattention and extrapolation with respect to current and future inflation.

All estimation exercises are conducted in a Bayesian framework targeting the mode of the posterior density. Table 2 summarises the parameters, priors, posterior modes, and standard errors across four specifications:

1. Full information rational expectations (FIRE): $\theta_{hh} = \theta_{prof} = \gamma_{hh} = \gamma_{prof} = 0$.

2. Calibrated behavioural expectations (BE): behavioural parameters fixed in line with the empirical values reported in Kohlhas and Walther (2021) for households and firms separately, with $\theta_{hh} = 0.396$, $\theta_{prof} = 0.302$, $\gamma_{hh} = -0.096$, and $\gamma_{prof} = -0.148$.

3. Estimated behavioural expectations (Est. BE): all parameters, including θ_{hh}, θ_{prof}, γ_{hh}, and γ_{prof}, estimated jointly on the nine observables including the two inflation expectation series.

4. Estimated behavioural expectations without survey data (Est. BE ex. $\mathbb{E}$): as in (3) but excluding the inflation expectation series.

4.2.1 Estimation Results

Table 2 reports the posterior modes for each specification, with Hessian-based standard errors in brackets. All three behavioural specifications statistically significantly raise the macroeconomic log-likelihood relative to the FIRE benchmark. Of particular note is that specification (2) outperforms the FIRE benchmark despite estimating the same number of free parameters.

[19] Transforming the expectations series into series of forecast errors improves identification of the behavioural parameters by ensuring consistent de-trending with actual inflation outturns and reducing the correlation of these series with the other inflation series observables the model is estimated on.

[20] The seven shock processes are: a discount factor shock, monetary policy shock, government spending shock, TFP shock, risk premium shock, and price and wage markup shocks.

[21] Specifically, for each model we run a Nelder-Mead algorithm for 1500 iterations at a time until the log likelihood improvement over the iteration batch falls to less than 1.0. We then conduct a local search using the L-BFGS-B algorithm (scipy) until local convergence of the log likelihood of the model.

[22] For γ_{hh} and γ_{prof}, we apply a quarterly decay rate of 0.86 starting from quarter 20 after a shock. This limits the horizon over which extrapolation is active to ten years (beginning the decay from five years), and was introduced to prevent long-run cycles when extrapolation is applied to both household and firm inputs. It did not materially affect the parameter estimates but significantly reduced the time required for impulse responses to converge, which was of practical importance for the estimation.

[23] The stickiness parameter θ maps to the δ parameterisation of the previous section via $\theta = 1 - \frac{1}{1+\delta}$.

Table 2: Parameter Estimates

Parameter	Prior	(1) RE	(2) Fixed BE	(3) Est. BE	(4) Est. BE (ex. $\mathbb{E}$)
Shock Processes					
ρ_C	Beta(0.5,0.2)	0.950 [0.026]	0.957 [0.023]	0.968 [0.023]	0.965 [0.019]
ρ_r	Beta(0.5,0.2)	0.069 [0.044]	0.192 [0.071]	0.150 [0.067]	0.172 [0.069]
ρ_g	Beta(0.5,0.2)	0.954 [0.035]	0.933 [0.028]	0.940 [0.025]	0.927 [0.028]
ρ_z	Beta(0.5,0.2)	0.988 [0.007]	0.974 [0.045]	0.932 [0.024]	0.927 [0.021]
ρ_{rp}	Beta(0.5,0.2)	0.900 [0.035]	0.917 [0.027]	0.817 [0.053]	0.925 [0.042]
ρ_p	Beta(0.5,0.2)	0.877 [0.022]	0.093 [0.068]	0.123 [0.074]	0.152 [0.097]
ρ_w	Beta(0.5,0.2)	0.897 [0.038]	0.887 [0.057]	0.992 [0.002]	0.792 [0.066]
σ_C	Inv-Gamma(0.1,2)	0.117 [0.021]	0.112 [0.021]	0.095 [0.019]	0.096 [0.016]
σ_r	Inv-Gamma(0.1,2)	0.235 [0.015]	0.226 [0.014]	0.200 [0.012]	0.203 [0.011]
σ_g	Inv-Gamma(0.1,2)	0.303 [0.017]	0.304 [0.017]	0.286 [0.016]	0.304 [0.017]
σ_z	Inv-Gamma(0.1,2)	0.530 [0.029]	0.530 [0.030]	0.532 [0.031]	0.533 [0.030]
σ_{rp}	Inv-Gamma(0.1,2)	0.215 [0.055]	0.218 [0.047]	2.240 [0.765]	0.810 [0.402]
σ_p	Inv-Gamma(0.1,2)	0.275 [0.017]	0.080 [0.005]	0.159 [0.020]	0.195 [0.059]
σ_w	Inv-Gamma(0.1,2)	0.186 [0.019]	0.113 [0.007]	0.188 [0.021]	0.226 [0.019]
θ_p	Beta(0.5,0.2)	0.984 [0.012]	0.404 [0.070]	0.692 [0.099]	0.461 [0.101]
θ_w	Beta(0.5,0.2)	0.602 [0.100]	0.715 [0.082]	0.916 [0.036]	0.602 [0.095]
Dynamic Parameters					
ψ	Normal(10,4)	0.423 [0.063]	0.905 [0.123]	9.371 [3.765]	10.184 [3.101]
ρ_m	Beta(0.5,0.2)	0.639 [0.030]	0.650 [0.037]	0.855 [0.029]	0.821 [0.030]
ζ_p	Beta(0.7,0.2)	0.821 [0.017]	0.904 [0.012]	0.939 [0.013]	0.950 [0.012]
ζ_w	Beta(0.7,0.2)	0.875 [0.024]	0.923 [0.013]	0.946 [0.017]	0.957 [0.011]
Behavioural Parameters					
θ_{hh}	Beta(0.5,0.2)	0	0.396	0.862 [0.030]	0.886 [0.024]
θ_{prof}	Beta(0.5,0.2)	0	0.302	0.906 [0.029]	0.462 [0.204]
γ_{hh}	Normal(0.0,0.3)	0	-0.096	-0.054 [0.056]	0.002 [0.036]
γ_{prof}	Normal(0.0,0.3)	0	-0.148	-0.127 [0.082]	-0.022 [0.031]
Log-likelihood		183.7	206.2	250.9	268.6

Note: Table reports parameter modes with Hessian-based standard errors in brackets for four specifications: (1) rational expectations (RE) with $\theta = \gamma = 0$ fixed; (2) behavioural expectations fixed at empirical values from Kohlhas and Walther (2021); (3) behavioural expectations estimated jointly with all parameters using nine observables including inflation expectations; (4) as in (3) but without inflation expectations data. Behavioural parameters in columns (1) and (2) are not estimated and have no standard errors reported. Log-likelihood values are reported for each specification, with higher values indicating better fit. The log-likelihood is based on the seven macroeconomic time series in each case, including for the model (3) disciplined against expectations data. Laplace-approximated log Bayes factors provide very strong evidence in favour of all three behavioural specifications over the RE benchmark, with log Bayes factors of 25.42, 55.02, and 85.42 for models (2), (3), and (4) respectively ($2 \times \log \mathrm{BF} \gg 10$ in all cases).

The behavioural estimates that include expectations data call for a high degree of informational stickiness relative to the rational expectations benchmark and, more relevantly, relative to the calibrated model. Household stickiness rises from the empirical value of $\theta_{hh} = 0.396$ to 0.862 and 0.886 in columns (3) and (4), while professional stickiness rises from 0.302 to 0.906 in column (3). The comparison of columns (3) and (4) isolates the role of the survey data: including the expectations series pushes professional inattention near its upper bound ($\theta_{prof} = 0.906[0.029]$), implying an average inattention spell of approximately 10 quarters in the cross-section of firms. This is plausibly needed to reproduce the persistence and initial underreaction observed in the survey forecast errors. Dropping the expectations series leaves professional stickiness considerably more moderate ($\theta_{prof} = 0.462[0.204]$), implying on average information sets turn over every 1.85 quarters, consistent with the lower bound of the estimate in Coibion and Gorodnichenko (2015).

This suggests some tension between the degree of stickiness implied by aggregate time series and that implied by survey data. Several explanations, net of model misspecification issues, may contribute. Forecast errors series backed out from survey expectations data may contain informative components generating inertia and large initial underreaction that the current specification absorbs by construction into the inattention parameter — this could be the case, for example, in the presence of a persistent expectations shock process or cognitive discounting. Additionally, the constant inattention parameter currently imposed as a restriction is effectively forced to reconcile both initially large under-reaction with the persistence properties of the forecast error series in the time domain, which as shown in the dynamic overshooting section must correlate positively with a constant inattention rate. Both, together with time-varying inattention rates as derived in Appendix A.5, constitute, however, feasible refinements for future work. At the same time, the higher estimated stickiness is associated with a significant improvement in overall model fit according to both full-information estimation likelihood metrics and with reference to partial-information critera such as the dynamic response of forecast errors, which we rationalise as reflecting the complex nonlinear interaction between expectations formation and aggregate dynamics at market clearing. Different data de-trending strategies or, the inclusion of additional observables or further variance of coefficients between prices may also help to reconcile the tension between the two data sources. This is a topic beyond the scope of the current paper, and we leave it to future work.

At the same time, the specifications with full-information estimated behavioural parameters (3)-(4) offer some favourable evidence on the role of the extrapolation channel in the US business cycle, particularly on the part of professional forecasters. To the best of our knowledge, this constitutes the first such finding within a full-information DSGE and HANK estimation exercise, highlighting the power of the proposed framework in bridging the gap between microeconomic evidence on deviations from FIRE and their role in estimated business cycle models. The comparison between estimation exercises (3) and (4) merits some attention. Linking to Slobodyan and Wouters (2012) in the context of the larger-scale estimation exercise we proposed, the above results similarly document that integrating expectations data in the structural estimation exercise offers an informative wedge to better identify and discipline the parameter estimates. Interpretation-wise, the recovered extrapolation parameter in (3) – with point-estimates for households and professional forecasters respectively $\gamma_{hh} = -0.054[0.056]$ and $\gamma_{prof} = -0.127[0.082]$ – are qualitatively consistent with priors from the empirical microeconomics literature (Kohlhas and Walther, 2021), suggesting over-reactive extrapolation with firms' sluggish forecast revisions exhibiting pro-cyclical properties with respect to macroeconomic and outturn conditions. As outlined in the dynamic overshooting section and illustrated below through IRF decomposition, the friction introduced by extrapolation plays is key for sustaining endogenous amplification and replicating empirical forecast error patterns relative to purely sticky expectations and FIRE.

Turning to the shock process and dynamic parameters, there are no consistent differences across specifications. The shock processes remain highly persistent under all behavioural specifications, implying that the additional behavioural frictions do not reduce the model's reliance on exogenous persistence. This is broadly consistent with our analysis of dynamic overshooting: generating sufficient weight for history in agents' forecasts requires shock processes that are persistent relative to the probability of updating beliefs, so it is perhaps unsurprising that persistence remains similarly high across specifications. The behavioural models also call for higher price and investment adjustment cost frictions, suggesting these do not act as substitutes for expectations frictions — contrary to what one might expect a priori. This reflects the complex interaction between time series volatility, behavioural frictions, and market clearing: underreaction may necessitate larger price moves to clear markets, calling in turn for higher adjustment costs to prevent excessive volatility in investment and the price level.

4.3 Extrapolation Properties

Figure 2 isolates the contribution of extrapolation to the transmission of monetary policy shocks. We re-solve the model at the estimated posterior modes from Table 2 but set the extrapolation parameters γ_{hh} and γ_{prof} to zero for both households and firms. The blue areas trace out the impulse responses under this counterfactual — a model with the same estimated degree of stickiness but no extrapolation — while the red shaded area represents the additional contribution of extrapolation. The FIRE impulse response is shown as a benchmark.

Two things stand out. First, in both panels the impulse responses of key real variables are hump-shaped relative to the FIRE benchmark: the stickiness channel slows the incorporation of news into agents' information sets, generating the delayed peak responses familiar from inattentive-agent models. Second, in both the calibrated and estimated case, the overreactive extrapolation component drives stronger amplification of macro-humps in both real and nominal variables relative to the nested sticky benchmark. This is a direct result of the fact that agents undergoing staggered updating of their information sets also extrapolate pro-cyclically from current conditions, overshooting the outturn and rational expectations paths, and amplifying endogenous fluctuations in response to the underlying shock. This is the dynamic overshooting mechanism formalised in Section 3, and it is precisely the channel through which over-reactive extrapolation can alleviate the amplification concerns raised by Albuquerque et al. (2025) relative to purely sticky expectations.

In relative terms, the model with calibrated behavioural parameters displays higher frequency cycles – i.e. an earlier occurring hump – than the estimated case. At the same time, the amplification induced by extrapolation relative to the sticky counterfactual is also more sizable. This is a direct consequence of two issues: as we noted, the estimated stickiness is larger than in the calibrated case. This immediately implies lower frequency in the information cycle, with the average agent taking longer before revising their forecast. Secondly, the estimated over-reaction particularly for households is weaker in the estimated models than in the calibrated case. The result of this is that the overall pro-cyclical bias in forecast revisions is lower,[24], dampening the extent of dynamic overshooting. These two qualitatively different contributions of extrapolation illustrate the flexibility of the framework: by remaining agnostic on the sign and magnitude of γ, the approach can accommodate a continuum of amplifying and dampening deviations from the rational expectations, spanning amplitude and frequency properties of hump-shaped IRFs, while the developed solver provide an off-the-shelf way for taking arbitrary DSGE and HANK models to the data to discriminate between them.

4.4 Dynamic Pattern of Forecast Errors

Figure 3 assesses the model-implied forecast error impulse responses against their empirical counterparts across the four estimated specifications. We construct the dynamic forecast error path following Adam et al. (2024) which plots the response of four quarter-ahead inflation forecast errors to a lagged change in inflation.[25] The model implied version is constructed by simulating the models for 5000 periods under the modal estimated shock standard deviations. We then estimate the local projection on the generated times series from the simulation. The FIRE model produces a flat forecast error at zero by construction. The empirical forecast error IRF is consistent with dynamic overshooting: the forecast error jumps on impact and subsequently crosses zero from above, indicating initial underreaction

[24] Equivalently, at similar inattention rates, it takes a more sustained expansionary phase to sustain forecast revisions of the same magnitude

[25] Adam et al. (2024) conduct a similar exercise when comparing their behavioural model of house prices to house price expectations in the Michigan survey.

Figure 2: Extrapolation contribution to a monetary policy shock

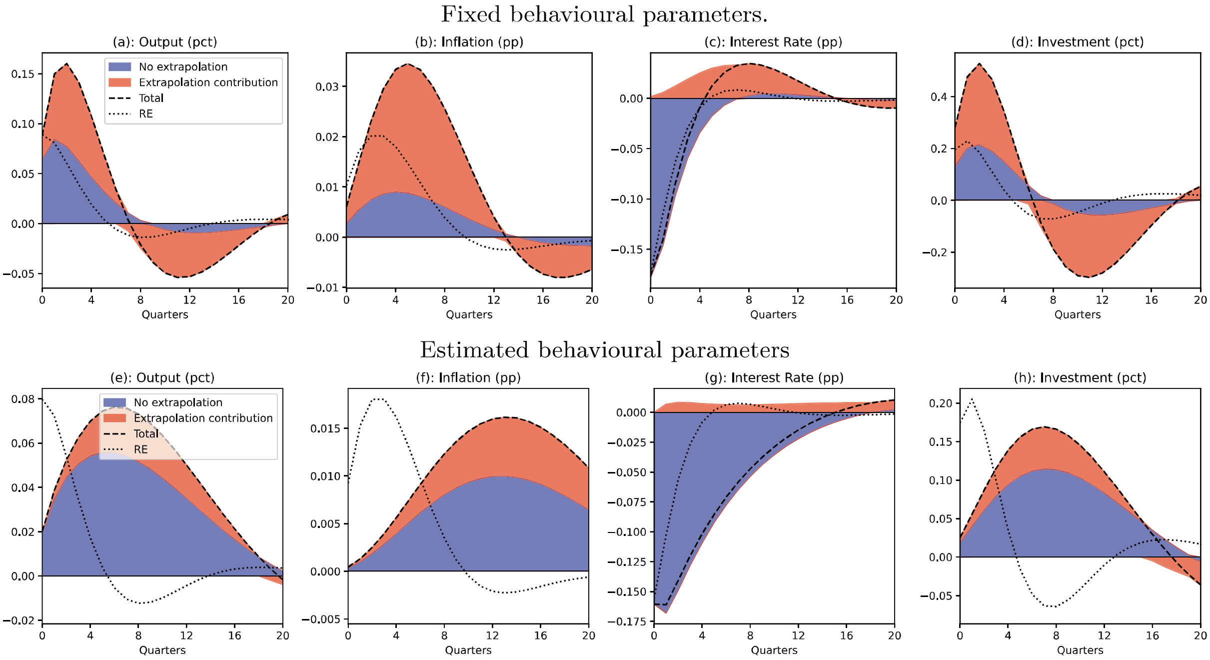

Note: The figure isolates the contribution of extrapolation by re-solving the model with $\gamma_{hh} = \gamma_{prof} = 0$ for both firms and households, otherwise setting all parameters at the posterior mode from Table 2. The top panel reports results for column (2) (BE calibrated at empirical values from Kohlhas and Walther (2021)); the bottom panel for column (3) (BE estimated jointly with all parameters). The blue shaded area shows the impulse responses without extrapolation; the red shaded area represents the additional contribution of extrapolation; and the black dashed line shows the total BE response. The black dotted line plots the FIRE impulse response from column (1). The shock is a 25bp contractionary monetary policy shock. For decompositions of other shocks, see Appendix B.4.2.

Figure 3: Dynamic response of forecast errors to an inflation innovation

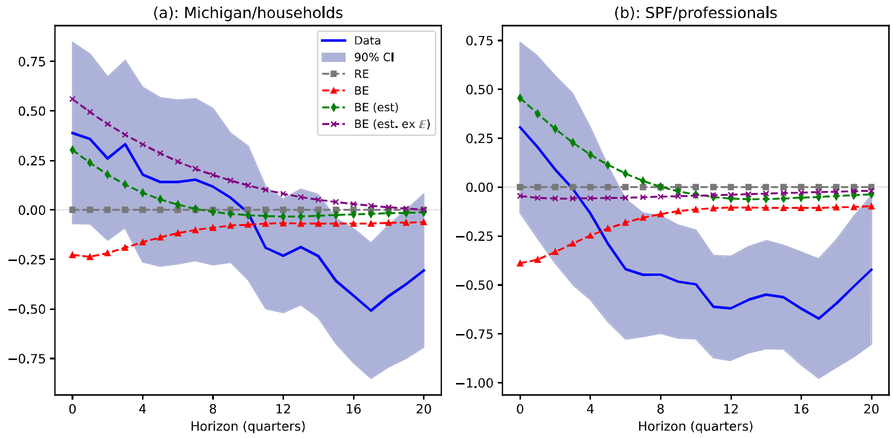

Note: The figure plots the response of the four-quarter-ahead inflation forecast error — the difference between the forecast and the outturn — to a unit inflation innovation, for Michigan households (panel a) and SPF professionals (panel b). The empirical series are estimated by local projections following Adam et al. (2024). Specifically, for each horizon we estimate $FE_{t+h} = \beta_h \pi_{t-1} + \varepsilon_{t+h}$ where $FE_t = \sum_{k=0}^{3} \pi_{t+k} - f_t^{(4)}$ is the four-quarter cumulative inflation outturn minus the survey forecast made at t, and π_{t-1} is lagged quarterly inflation. Standard errors are Newey-West with shaded 90% confidence intervals reported. Model series are computed directly from the sequence-space Jacobian as a variance-share weighted average of forecast-error impulse responses across all seven structural shocks, normalised per unit impact response of inflation.

followed by overshooting of the respective outturns.

Among the specifications, the estimated model featuring high stickiness and over-reactive extrapolation among firms and households (column (3)) achieves the closest match to the empirical forecast errors, particularly for households. This is a direct manifestation of the stickiness–extrapolation interaction formalised in Section 3, and stands in contrast to the monotone convergence from below that a purely sticky, time-dependent model would produce – for comparison, this is plausibly the case in model (4) that combines high stickiness with counter-cyclical or under-reactive extrapolation on the part of households. The better fit follows from the relatively protracted undershooting phase in specification (3): the relatively strong stickiness implies more sluggish updating which in turn implies that the average forecast produced by staggered revisions extrapolates from a longer history. Combined with over-reactive extrapolation, such accumulating pro-cyclical bias in forecasts eventually gains sufficient momentum to temporarily offset convergence to perfect foresight and over-predict the outturn. It is interesting – and at first sight counter-intuitive – that the calibrated case (2) fails to deliver a similar or stronger result. Despite the stronger extrapolation channel for firms and especially households, here overshooting of outturn occurs on impact rather than display the characteristic under-reaction and interior crossing documented in the empirical local projection and literature Angeletos et al. (2021). Upon closer inspection through the formal framework developed for dynamic overshooting, this is what one would expect given the relatively low inattention calibrated on the basis of micro estimates. The reason is twofold. First, in line with the above narrative, the much higher frequency of information updating (around 1.85 quarters) implies that the average forecast is driven in disproportionately larger measure by the revisions of recent updaters. In other words, while extrapolation from current conditions is stronger, the more rapid turnover of information sets effectively offsets the pass-through to average forecast from the history of staggered information updates, by reducing the influence and weight of further away segments of the history in the prevailing forecast. Secondly, in line with the outlined constraints on the persistence of the forecast processes relative to the inattention parameter, this persistence requirements becomes more binding in the calibrated model due to the fact agents learn more rapidly. Viewed through the lenses of condition the dynamic overshooting framework, this is precisely the case in which the relatively low inattention rates tighten the bound on the magnitude of over-reactions, making calibrated extrapolation relatively too strong.

4.5 Broader Model Fit

This section assesses the fit of the estimated models to the broader set of business cycle moments using three complementary yardsticks: root mean squared forecast errors (RMSEs) across horizons, autocorrelation functions (ACFs), and the match to empirically estimated impulse responses to monetary policy shocks from Romer and Romer (2004) and Bauer and Swanson (2023). Together these yardsticks speak to a pattern that runs through the results: behavioural expectations generate a clear improvement in the short-run dynamics and shape of impulse responses while FIRE continues to perform well on unconditional second moments at longer horizons.

Figure 4 summarises the in-sample forecast performance. The fixed empirical-BE and the estimated-BE specifications tend to reduce RMSEs over the first few quarters relative to the FIRE benchmark. The estimated BE specifications generally outperform the FIRE benchmark across horizons, though notably FIRE performs best on the inflation series beyond the first few quarters.

Figure 4: Model forecast errors

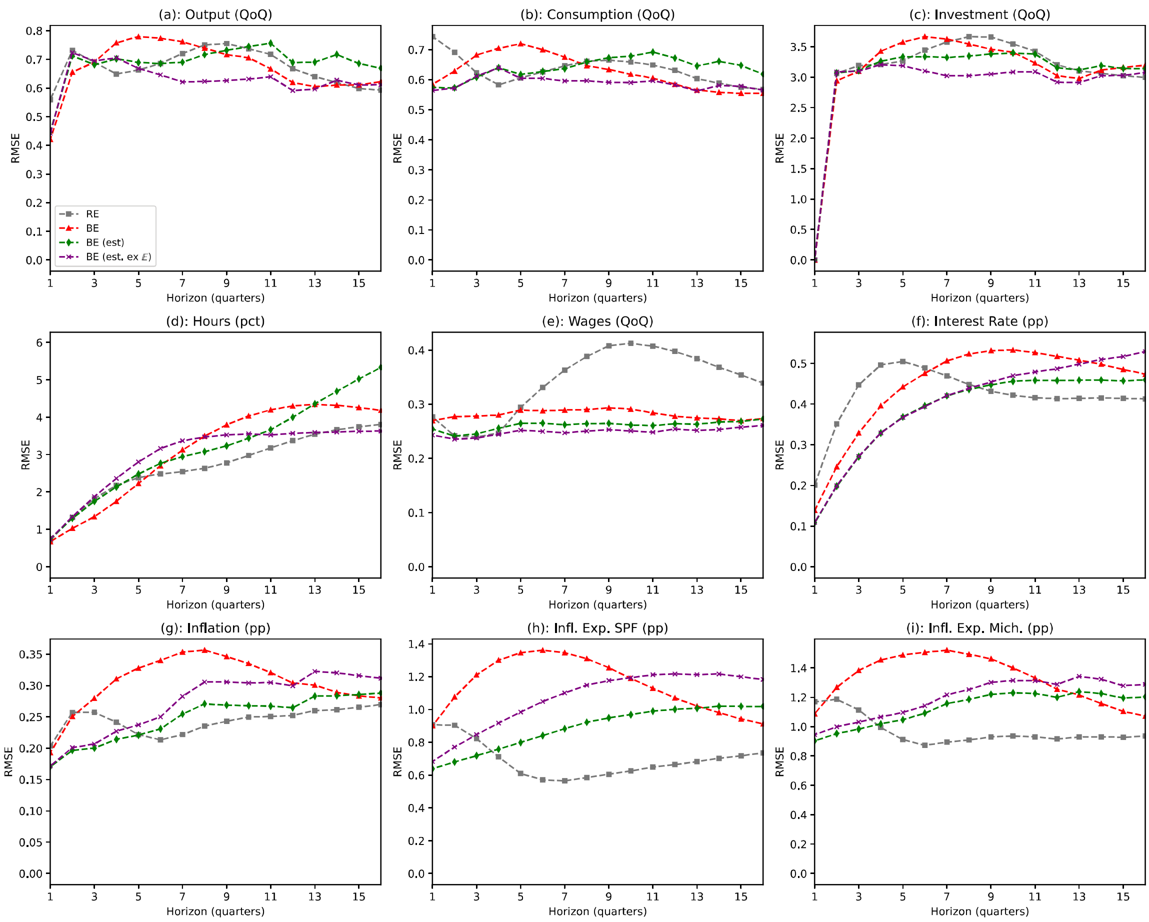

Note: Figure reports the in-sample root mean squared error (RMSE) of the model's one-quarter-ahead forecasts for the macroeconomic time series used in estimation, for each of the four specifications in Table 2. RMSEs are calculated over the estimation sample 1978–2018 with the first 40 quarters used as a burn-in period.

Figure 5 compares model-implied and empirical ACFs. This is complemented by the analysis of the spectral densities presented in Figure B.5 in Appendix B.4.2. Generally the estimated behavioural specifications create too much persistence relative to the data, whereas the FIRE or calibrated behavioural specification (which is closer to FIRE) matches the empirical ACFs more closely. The exception are hours and wages which are much more persistent in the data and therefore more in line with the estimated behavioural specifications. A common feature across all specifications is the relatively higher frequency cycle characterising in-sample dynamics of the interest rate. One possible explanation is that the benchmark model does not capture time-varying aspects of the monetary policy rule. Further, whether enabling deviations from FIRE on the part of the monetary authority itself would improve the fit of the interest rate series is an interesting question for future work.

Figure 5: Model autocorrelation functions (ACFs)

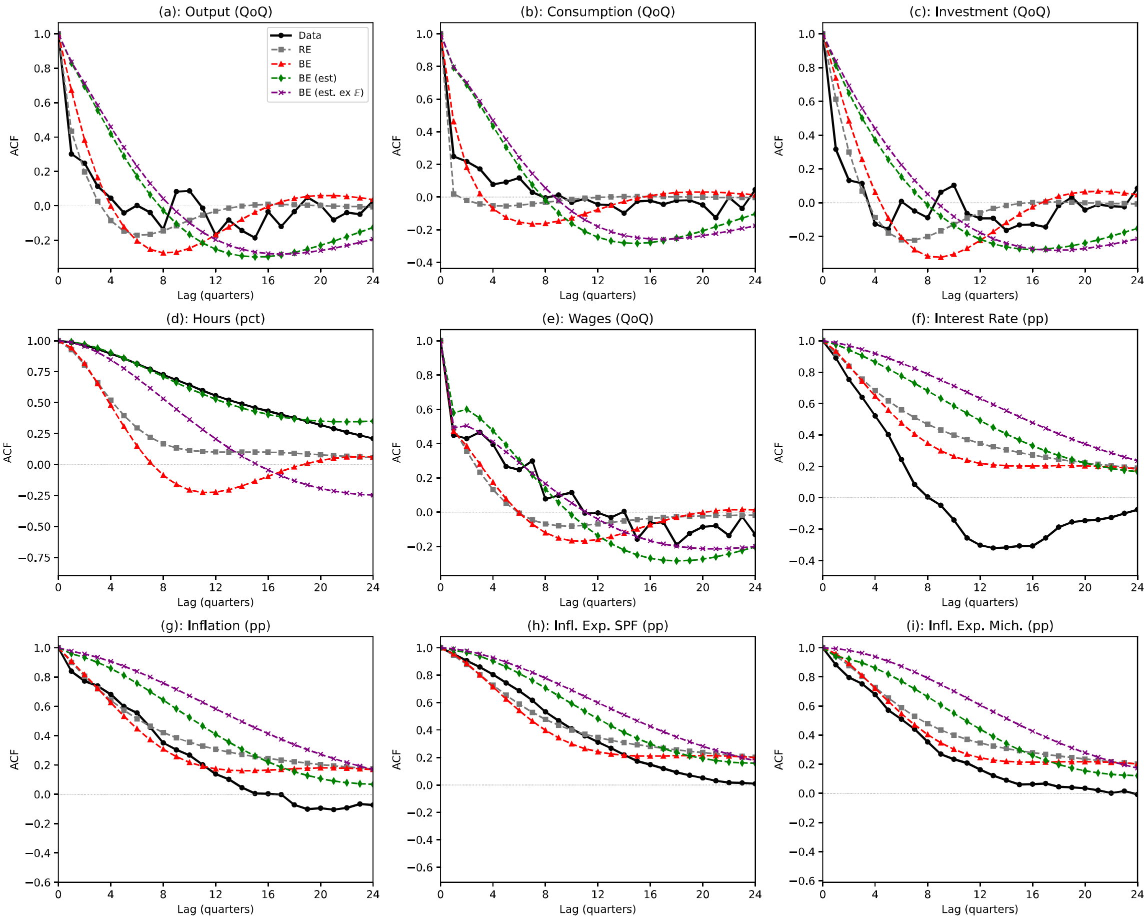

Note: Figure reports model-implied autocorrelation functions for the macroeconomic time series used in estimation, for each of the four specifications in Table 2, compared to empirical autocorrelation functions calculated on the data over the estimation sample 1978–2018.

Finally previous papers in the literature have adopted an IRF matching strategy to estimate dynamic parameters, whereas our approach does not. It is therefore of interest to compare the indirectly estimated impulse responses to a monetary shock to the empirical estimates. Figure 6 documents this. The estimated-BE models (columns (3)–(4) in Table 2) provide the closest visual match to the Romer and Romer empirical IRFs. Spanning the IRF evidence, when estimated on business cycle data the IRFs sit far closer to the empirical IRFs based on the Romer and Romer shock series than the IRFs estimated on the Bauer and Swanson shock series. Even the calibrated behavioural model which contains significantly greater extrapolation struggles to breach the 90 percent confidence interval.

Figure 6: Impulse responses to a monetary policy shock

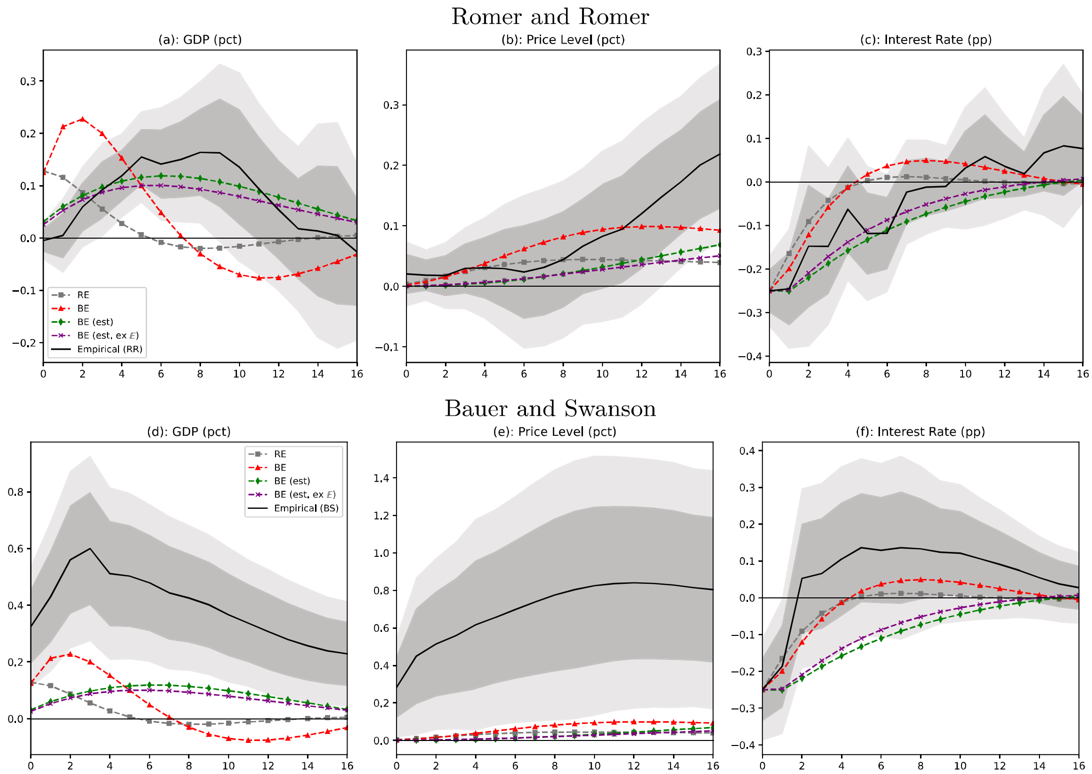

Note: Figure reports model-implied impulse response functions to a 25bp expansionary monetary policy shock for each of the four specifications in Table 2, compared to empirical impulse responses estimated using the Romer and Romer (2004) shock series as in Auclert et al. (2020) (panels (a)–(c)) and the high-frequency identified shock series of Bauer and Swanson (2023) (panels (d)–(f)). Monthly GDP has been interpolated from monthly industrial production data.

5 Conclusion

This paper develops a parsimonious, computationally efficient framework for departing from full information rational expectations in the sequence space. Building on Kohlhas and Walther (2021), the behavioural expectations process captures the two empirically dominant deviations from FIRE—extrapolation from current conditions and underreaction to news—through a two-parameter law of motion that is internally rational in the sense of Adam and Marcet (2011), nests leading alternative expectation models as special cases, and delivers the sequence-space Jacobian in closed form as a function of the rational-expectations solution. Applying the framework to the two-asset HANK model of Auclert et al. (2020) and estimating it on US business cycle data with SPF and Michigan Survey inflation expectations, we find that behavioural specifications outperform FIRE in model likelihood, the qualitative properties of model impulse responses and the dynamic pattern of forecast errors.

References

Adam, Klaus and Albert Marcet, "Internal rationality, imperfect market knowledge and asset prices," *Journal of Economic Theory*, 2011, *146* (3), 1224–1252.

_ , Oliver Pfäuti, and Timo Reinelt, "Subjective housing price expectations, falling natural rates, and the optimal inflation target," *Journal of Monetary Economics*, 2024, p. 103647.

Albuquerque, Diogo, Thomas Lazarowicz, and Jamie Lenney, "Monetary transmission through the housing sector," 2025. Working paper.

Angeletos, George-Marios, Joao Guerreiro, and Dalton Rongxuan Zhang, "From RANK to HANK, without FIRE," Technical Report, National Bureau of Economic Research 2025.

_ , 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," Technical Report, National Bureau of Economic Research 2020.

Bardóczy, Bence and Joao Guerreiro, "Unemployment insurance in macroeconomic stabilization with imperfect expectations," *Manuscript, April*, 2024, *5.*

Bauer, Michael D and Eric T Swanson, "A reassessment of monetary policy surprises and high-frequency identification," *NBER Macroeconomics Annual*, 2023, *37*, 87–155.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, "Diagnostic expectations and credit cycles," *Journal of Finance*, 2018, *73* (1), 199–227.

Coibion, Olivier and Yuri Gorodnichenko, "Information rigidity and the expectations formation process: A simple framework and new facts," *American Economic Review*, 2015, *105* (8), 2644–2678.

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

Kohlhas, Alexandre N and Ansgar Walther, "Asymmetric attention," *American Economic Review*, 2021, *111* (9), 2879–2925.

Mankiw, N Gregory and Ricardo Reis, "Sticky information: A model of monetary nonneutrality and structural slumps," *Knowledge, information, and expectations in modern macroeconomics: In honor of Edmund S. Phelps*, 2003, pp. 64–86.

Negro, Marco Del and Frank Schorfheide, "DSGE model-based forecasting," in "Handbook of economic forecasting," Vol. 2, Elsevier, 2013, pp. 57–140.

_ and Stefano Eusepi, "Fitting observed inflation expectations," *Journal of Economic Dynamics and control*, 2011, *35* (12), 2105–2131.

Romer, Christina D and David H Romer, "A new measure of monetary shocks: Derivation and implications," *American economic review*, 2004, *94* (4), 1055–1084.

Slobodyan, Sergey and Raf Wouters, "Learning in an estimated medium-scale DSGE model," *Journal of Economic Dynamics and Control*, 2012, *36* (1), 26–46.

Smets, Frank and Rafael Wouters, "Shocks and frictions in US business cycles: A Bayesian DSGE approach," *American economic review*, 2007, *97* (3), 586–606.

A Further Methodological Details

A.1 Law of Motion for Behavioural Forecasts around the Steady State

Derivation of the recursive forecast law of motion. This appendix shows how the average cross-sectional forecast for an aggregate input at future period $t+k$ evolves over time around the steady-state path, establishing the law of motion used in equation (5).

Recall that the average forecast formed at time t on an aggregate input k periods ahead is given by equation (2):

$$\bar{f}_t x_{t+k} = c + \frac{1}{1+\delta}\left(\delta \bar{f}_{t-1} x_{t+k} + x_{t+k} - \gamma x_t\right)$$

where we suppress the full-information expectations operator E_t^* since, under the sequence-space solution concept, there is no aggregate uncertainty and all aggregate sequences are part of the time-zero information set.

Performing a first-order perturbation around the steady state, writing $x_r = x_{ss} + dx_r$ for any period r:

$$\bar{f}_t(x_{ss} + dx_{t+k}) = c + \frac{1}{1+\delta}\left(\delta \bar{f}_{t-1}(x_{ss} + dx_{t+k}) + x_{ss} + dx_{t+k} - \gamma(x_{ss} + dx_t)\right).$$

Using $\bar{f}_t x_{ss} = x_{ss}$ (agents have unbiased steady-state expectations) together with the no-bias assumption (Assumption 1 in Appendix A.2), which pins down the constant as $c = \frac{\gamma}{1+\delta} x_{ss}$, the steady-state terms cancel to give:

$$x_{ss} + \bar{f}_t dx_{t+k} = x_{ss} + \frac{1}{1+\delta}\left(\delta \bar{f}_{t-1} dx_{t+k} + dx_{t+k} - \gamma\, dx_t\right).$$

Subtracting x_{ss} from both sides yields the linearised law of motion:

$$\bar{f}_t\, dx_{t+k} = \frac{1}{1+\delta}\left(\delta \bar{f}_{t-1} dx_{t+k} + dx_{t+k} - \gamma\, dx_t\right). \tag{A.1}$$

This is a first-order linear difference equation in the sequence of forecasts $\{\bar{f}_t dx_{t+k}\}_{t\geq 0}$, expressing how the forecast of x_{t+k} at time t depends on: the previous period's forecast of the same future period (formed at $t-1$, when the horizon was $k+1$ periods away); the full-information value dx_{t+k}; and extrapolation from current conditions dx_t.

Solving equation (A.1) backwards, applying the boundary condition $\bar{f}_{-1} dx_{t+k} = 0$ (which follows from the assumption that agents start from the non-stochastic steady state at the origin of the transition, so that no deviations from the steady-state forecast are in place before the shock), yields equation (5) in the main text:

$$\bar{f}_t\, dx_{t+k} = \sum_{j=0}^{t}\left(\frac{\delta}{1+\delta}\right)^{t-j} \frac{1}{1+\delta}\, dx_{t+k} - \gamma \sum_{j=0}^{t}\left(\frac{\delta}{1+\delta}\right)^{t-j} \frac{1}{1+\delta}\, dx_j.$$

The first sum converges to dx_{t+k} as $t \to \infty$ (agents eventually attain perfect foresight), and the second captures the accumulated extrapolation from observed conditions, with exponentially declining weights on older observations.

A.2 Technical Assumptions for Internal Rationality

The behavioural expectations framework rests on four assumptions that together ensure agents are internally rational in the sense of Adam and Marcet (2011): they optimise given their subjective beliefs, their expectations converge to rational expectations in the long run, and they do not forecast their own forecast errors.

1. **No bias along the steady-state path.** Let α_i be the individual-specific intercept in the Kohlhas-Walther regression equation (1), and let P denote the cross-sectional distribution of

agents. The average forecast is unbiased along the steady-state path if and only if the average intercept satisfies:

$$\int \alpha_i \, P(di) = \gamma \, x_{ss}.$$

This pins down the constant in equation (2) as $c = \frac{\gamma}{1+\delta} x_{ss}$, ensuring that in the steady state agents' average forecast equals the steady-state value of the input: $\bar{f}_{ss} x_{ss} = x_{ss}$.

2. **Law of iterated expectations.** For all $t \geq s$, agents do not forecast their own future forecast errors:

$$\bar{f}_s \left[\bar{f}_t x_{t+k} \right] = \bar{f}_s x_{t+k}.$$

Equivalently, future forecast errors are orthogonal to the information set at any earlier time $s \leq t$:

$$\bar{f}_s \left[x_{t+k} - \bar{f}_t x_{t+k} \right] = 0.$$

This holds under the assumption that information sets lie in a natural filtration, so that agents cannot anticipate revisions they have not yet made. This assumption is critical for the derivation in Appendix A.3: because forecast updates between consecutive periods are unanticipated, they hit the economy as unanticipated time-zero shocks, which is precisely what allows the block Jacobians $\tilde{J}^\tau$ to be constructed from the FIRE Jacobian.

3. **Asymptotic convergence to the steady state.** For large enough T, agents expect the forecast variable to return to its steady-state value:

$$\exists \, T \in \mathbb{N} : \forall \, t \leq T, \quad \bar{f}_t x_T = x_{ss}.$$

This terminal condition, used in Bardóczy and Guerreiro (2024), ensures the sequence-space transition has a well-defined endpoint and that the linearised path is bounded.

4. **Positive news-revision coefficient.** The parameter $\delta > 0$, which implies $\theta = \delta/(1+\delta) \in (0,1)$. This is sufficient to ensure that the law of motion (A.1) is stable — the eigenvalue of the difference equation is $\delta/(1+\delta) < 1$ — and that the forecast error converges to zero along any non-stochastic path on which aggregate shocks have settled to zero (as noted in equation (4)). In the limit $\delta \to 0$ agents update fully each period (FIRE), while as $\delta \to \infty$ agents almost never update, recovering the purely sticky expectations case of Auclert et al. (2020).

A.3 The Heterogeneous Block Jacobian under Behavioural Expectations

Purpose and logical structure. This appendix derives the behavioural Jacobian $\hat{J}$ — which maps true aggregate input sequences $d\mathbf{X}$ to aggregate outputs $d\mathbf{Y}$ under the Kohlhas-Walther expectations process — from first principles. The derivation proceeds in four steps. First, we write down the full heterogeneous agents Jacobian $\tilde{J}$ with respect to both true inputs and all time-varying forecast sequences, partitioned into column blocks. Second, we use the structure of the recursive dynamic programming (DP) problem to express each block in terms of the FIRE Jacobian J and Fake-News matrix $\mathcal{F}$ from Auclert et al. (2021). Third, we use the linearity of the forecast mapping Λ_τ to collapse the full expanded Jacobian back to $\hat{J}$ with respect to true inputs only, recovering the formula in equation (**??**). Fourth, we reformulate the result entirely in terms of the Fake-News matrix $\mathcal{F}$ rather than J directly, which is the computationally efficient representation since $\mathcal{F}$ is a more primitive object: the Jacobian is defined by recursive sums over its diagonals, $J_{t,s} = \sum_{r=0}^{\min\{s,t\}} F_{t-r,s-r}$.

Throughout, we build on the extension of Auclert et al. (2021) to non-FIRE expectations environments developed in Bardóczy and Guerreiro (2024), and show how the specific structure of the Kohlhas-Walther forecast mapping Λ_τ yields closed-form entries for $\hat{J}$ in terms of $\mathcal{F}$ and Λ_τ.

Setup. Consider any scalar endogenous aggregate output $Y_{i,t} \in \mathbb{R}_+$ of the heterogeneous agents block. In the DAG formulation of the sequence-space equilibrium (Auclert et al., 2021; Bardóczy and Guerreiro, 2024), this satisfies:

$$Y_{i,t} = h(\mathbf{X}, \mathbf{X}^{e,0}, \ldots, \mathbf{X}^{e,t})$$

where $h : \mathcal{X} \to \mathbb{R}_+$ maps the set of length-T sequences of aggregate inputs $\mathcal{X} \subset \mathbb{R}^{n_i \times T}$ into the reals. The set of inputs includes both the true equilibrium aggregate input sequences $\mathbf{X}$ and the family of expectations $\{\mathbf{X}^{e,\tau}\}_{\tau=0,\ldots,T}$, where τ denotes the time at which expectations are formed and s the period being forecast. The equality above embeds the assumption that the aggregate output at time t depends only on expectations formed up to and including time t: agents cannot make use of information not yet in their information set. Note that under FIRE, $\mathbf{X}^{e,\tau} = \mathbf{X}$ for all τ, and the mapping collapses to the standard one in Auclert et al. (2021).

We make three assumptions throughout. (A1) There is a time-varying linear forecast mapping from true inputs to agent expectations: for each τ there exists a $T \times T$ matrix Λ_τ such that

$$d\mathbf{X}^{e,\tau} = \Lambda_\tau \, d\mathbf{X},$$

given by the Kohlhas-Walther law of motion derived in equation (5). (A2) The steady-state expectation of any input equals its steady-state value. (A3) A law of iterated expectations holds, so agents do not forecast their own future forecast errors.

At a first-order perturbation around the non-stochastic steady state, for all $t = 0, \ldots, T$:

$$dY_{i,t} = \sum_{s=0}^{T} \frac{\partial Y_{i,t}}{\partial X_s} dX_s + \sum_{\tau=0}^{t} \sum_{s=0}^{T} \frac{\partial Y_{i,t}}{\partial X_s^{e,\tau}} dX_s^{e,\tau}. \tag{A.2}$$

This is the key object. We now pin down the two groups of partial derivatives in turn.

Step 1: Partial derivatives with respect to true inputs — the block $\tilde{J}^*$. Consider the partial equilibrium response $dY_{i,t}^s$ when only the true aggregate input at time s is shocked ($dX_s \neq 0$) and all expectations are held at their steady-state values ($dX_r^{e,\tau} = 0$ for all τ, r).

If $s > t$: by the structure of the recursive DP problem, with expectations at their steady-state values the value function satisfies $v_t = \mathcal{V}(v, X_{ss})$ for all $t < s$, implying no deviation from the steady-state policy. Because there is no policy deviation prior to s, the beginning-of-period distribution at s is also at steady state, and therefore $dY_{i,t}^s = 0$ for $s > t$.

If $s = t$: since there are no policy deviations prior to t, the distribution is at steady state at t. A first-order perturbation of the aggregation rule gives $dY_{i,t}^s = dy_{i,t}^{s\,\prime} D_{ss}$, which by stationarity of the DP problem equals the first element of the FIRE Jacobian:

$$dY_{i,t}^s\big|_{t=s} = \mathcal{Y}_0 \, dX = J_{0,0} \, dX.$$

If $t > s$: optimal policies have reverted to steady state, so only the distribution propagates the shock. The linearised Chapman-Kolmogorov equation gives $dD_t^s = (\Lambda_{ss}')^{t-s-1} dD_{s+1}^s$, and using the expectations vector $\mathcal{E}_r = \Lambda_{ss}^r y_{ss}$ from Auclert et al. (2021):

$$dY_{i,t}^s = \mathcal{E}_{t-s-1}' \, \mathcal{D}_0 \, dX.$$

Combining, and using the change of variable $\tau = t - s$, these results coincide with Proposition 1 of Auclert et al. (2021):

$$\frac{\partial Y_{i,t}}{\partial X_s} = \begin{cases} J_{t-s,0} = \mathcal{F}_{t-s,0} & t \geq s \\ 0 & t < s \end{cases} \tag{A.3}$$

The block $\tilde{J}^*$ is therefore lower triangular with entries given by the first column of the FIRE Jacobian (equivalently, the first column of the fake-news matrix $\mathcal{F}$):

$$\tilde{J}^* = \begin{pmatrix} J_{0,0} & 0 & 0 & \cdots & 0 \\ J_{1,0} & J_{0,0} & 0 & \cdots & 0 \\ J_{2,0} & J_{1,0} & J_{0,0} & \cdots & 0 \\ \vdots & & & \ddots & \vdots \\ J_{T,0} & J_{T-1,0} & J_{T-2,0} & \cdots & J_{0,0} \end{pmatrix}$$

Note that $\tilde{J}^*$ is equivalent to the forecast-error matrix in Bardóczy and Guerreiro (2024): both give the IRF to an unanticipated shock to the time-s input. The equivalence is a direct consequence of assumption (A3) — because agents do not predict their future forecast errors, a shock to the true input dX_s generates a forecast error of the same magnitude that hits the economy at time s exactly as an unanticipated time-zero shock would.

Step 2: Partial derivatives with respect to forecast inputs — the blocks $\tilde{J}^\tau$. Now consider the partial equilibrium response when only the time-τ expectation of the period-r input is shocked, $dX_r^{e,\tau} \neq 0$, with all true inputs and other expectations at steady state.

If $t < \tau$: by assumption (A3), agents at time t cannot respond to expectations not yet formed. Hence $\frac{\partial Y_{i,t}}{\partial X_r^{e,\tau}} = 0$ for all $t < \tau$, and the first τ rows of $\tilde{J}^\tau$ are zero.

If $r \leq \tau$: current or past inputs are observed perfectly (the identity block of Λ_τ), so a shock to $X_r^{e,\tau}$ for $r \leq \tau$ has no effect on agents who already observed X_r directly. Hence $\frac{\partial Y_{i,t}}{\partial X_r^{e,\tau}} = 0$ for $r \leq \tau$, and the first $\tau + 1$ columns of $\tilde{J}^\tau$ are zero.

If $t = \tau$ and $r > \tau$: from the DP problem, the response is identical to a future shock under FIRE, since the time-τ expectation of a future input acts exactly like a time-zero expectation of the same future input in a re-dated economy. By stationarity:

$$\frac{\partial Y_{i,t}}{\partial X_r^{e,\tau}} = J_{0,r-\tau} = \mathcal{F}_{0,r-t}.$$

If $t > \tau$ and $r > \tau$: optimal policies have reverted to steady state, so only the distribution propagates. By an argument analogous to Step 1, using the timing shift $t' = t - \tau$, $r' = r - \tau$:

$$\frac{\partial Y_{i,t}}{\partial X_r^{e,\tau}} = \mathcal{E}_{t-\tau-1}\, \mathcal{D}_{r-\tau}\, dX = F_{t-\tau,r-\tau} = J_{t-\tau,r-\tau} - J_{t-\tau-1,r-\tau-1}.$$

Combining, the block $\tilde{J}^\tau$ has the structure: zero rows for $t < \tau$, zero columns for $r \leq \tau$, and remaining entries equal to elements of the FIRE fake-news matrix:

$$\frac{\partial Y_{i,t}}{\partial X_r^{e,\tau}} = \mathcal{F}_{t-\tau,r-\tau} \quad \text{for } t \geq \tau,\ r > \tau.$$

Writing out $\tilde{J}^0$ and $\tilde{J}^1$ explicitly:

$$\tilde{J}^0 = \begin{pmatrix} 0 & J_{0,1} & J_{0,2} & \cdots & J_{0,T} \\ 0 & F_{1,1} & F_{1,2} & \cdots & F_{1,T} \\ 0 & F_{2,1} & F_{2,2} & \cdots & F_{2,T} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & F_{T,1} & F_{T,2} & \cdots & F_{T,T} \end{pmatrix}, \quad \tilde{J}^1 = \begin{pmatrix} 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & J_{0,1} & \cdots & J_{0,T-1} \\ 0 & 0 & F_{1,1} & \cdots & F_{1,T-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & F_{T-1,1} & \cdots & F_{T-1,T-1} \end{pmatrix}$$

where $F_{t,s} = J_{t,s} - J_{t-1,s-1}$ denotes the (t, s) element of the fake-news matrix.

Step 3: Collapsing to the behavioural Jacobian $\hat{J}$. Partitioning $\tilde{J} = [\tilde{J}^* \quad \tilde{J}^0 \quad \cdots \quad \tilde{J}^T]$ and using the linear forecast mapping $d\mathbf{X}^{e,\tau} = \Lambda_\tau\, d\mathbf{X}$, the condition $d\mathbf{Y}_i = \hat{J}\, d\mathbf{X}$ requires:

$$\hat{J}\, d\mathbf{X} = \tilde{J}^*\, d\mathbf{X} + \sum_{\tau=0}^{T} \tilde{J}^\tau \Lambda_\tau\, d\mathbf{X}$$

and therefore:

$$\hat{J} = \tilde{J}^* + \sum_{\tau=0}^{T} \tilde{J}^\tau \Lambda_\tau. \tag{A.4}$$

This confirms equation (**??**), which can equivalently be written as $\hat{J} = J\Lambda_0 + \sum_h R_h(\Lambda_h - \Lambda_{h-1})$; the two representations are algebraically equivalent via the fake-news decomposition.

Step 4: Representation in terms of the Fake-News matrix. Substituting the results of Steps 1 and 2 into equation (A.2), and using $\mathcal{F}_{t,s} = J_{t,s} - J_{t-1,s-1}$ together with $J_{t,s} = \sum_{r=0}^{\min\{s,t\}} F_{t-r,s-r}$:

$$dY_{i,t} = \sum_{s=0}^{t} \mathcal{F}_{t-s,0}\, dX_s + \sum_{\tau=0}^{t-1} \sum_{s=\tau+1}^{T} \mathcal{F}_{t-\tau,s-\tau}\, dX_s^{e,\tau} + \sum_{s=t+1}^{T} \mathcal{F}_{0,s-t}\, dX_s^{e,t}. \tag{A.5}$$

The three terms have a clean interpretation. The first captures the direct effect of true shocks at times $s \leq t$, weighted by the first column of $\mathcal{F}$. The second captures the effect of expectation revisions at all intermediate times $\tau < t$, as agents update their forecasts when new information arrives. The third captures the effect of expectations formed at time t about future inputs $s > t$, weighted by the first row of $\mathcal{F}$.

Substituting the forecast mapping $d\mathbf{X}^{e,\tau} = \Lambda_\tau\, d\mathbf{X}$, the (t,r) element of $\hat{J}$ follows by an undetermined coefficients argument:

$$\hat{J}_{t,r} = \mathcal{F}_{t-r,0} + \sum_{\tau=0}^{t-1} \sum_{s=\tau+1}^{T} \mathcal{F}_{t-\tau,s-\tau}\, \Lambda_{\tau,s,r} + \sum_{s=t+1}^{T} \mathcal{F}_{0,s-t}\, \Lambda_{t,s,r} \tag{A.6}$$

where $\Lambda_{\tau,s,r}$ denotes the (s,r) element of the forecast matrix Λ_τ, given by:

$$\Lambda_{\tau,s,r} = \frac{\partial \bar{f}_\tau X_s}{\partial X_r} = \begin{cases} 1 & s \leq \tau,\ r = s \\ \displaystyle\sum_{l=0}^{\tau} \frac{\delta^l}{(1+\delta)^{l+1}} & s > \tau,\ r = s \\ -\gamma \left(\dfrac{\delta}{1+\delta}\right)^{\tau-r} \dfrac{1}{1+\delta} & s > \tau,\ r \leq \tau \\ 0 & \text{otherwise} \end{cases} \tag{A.7}$$

The four cases correspond directly to the four regions of Λ_τ identified in Section 8: (i) the identity block (agents observe current and past inputs perfectly); (ii) the news-revision diagonal $\mathbf{B}_\tau$ (cumulative probability of incorporating news on future inputs); (iii) the extrapolation block $\mathbf{A}_\tau$ (pass-through from observed conditions at each update); and (iv) zero elsewhere.

Equation (A.6) is the central result of this appendix. It shows that the behavioural Jacobian $\hat{J}$ can be computed entirely from the FIRE fake-news matrix $\mathcal{F}$ and the parameters (δ, γ) of the Kohlhas-Walther process, without ever computing J explicitly. Since $\mathcal{F}$ is a more primitive object than J — the Jacobian being defined by recursive column sums of $\mathcal{F}$ — this representation is both computationally efficient and conceptually transparent.

Relationship to Bardóczy and Guerreiro (2024). The derivation above follows the procedure of Bardóczy and Guerreiro (2024) closely. Our contribution relative to their general treatment is threefold: we specialise to the Kohlhas-Walther forecast mapping Λ_τ, we provide explicit closed-form entries for each block $\tilde{J}^\tau$ via equation (A.7), and we show the result can be stated directly in terms of $\mathcal{F}$ rather than J, which is the representation used in our computational implementation.

A.4 Equivalence of the generalised solver with Bardóczy and Guerreiro (2024)

In this section we show equivalence between the proposed generalised solver and the alternative implementation in Bardóczy and Guerreiro (2024) derived in the last section. The key idea is twofold:

1. The iterative revision of forecasts $d\mathbf{X}^{e,t} - d\mathbf{X}^{e,t-1}$ that hits as a time-zero shock in Bardóczy and Guerreiro (2024) is equivalent to iteratively unlearning the previous period forecast and learning a new one, which is the device underlying our more efficient fake-news matrix based implementation.

2. The sequence-space fake-news matrix and Jacobian are related by the recursion proved in Auclert et al. (2020). This implies that a similar recursive relationship holds between our **forecast unlearning** matrix $\mathcal{P}_h$ and the forecast revision matrix in Bardóczy and Guerreiro (2024).

The proof of equivalence is in two parts. First we prove that any Jacobian satisfying the Bardóczy and Guerreiro (2024) solver also satisfies ours (necessity). Then we prove the reverse (sufficiency). This establishes that the two solvers span the same space of solutions $\mathcal{J} \subset \mathcal{M}_T$.

We first provide the following R-P recursion lemma, linking the *forecast revision matrix* R_h in Bardóczy and Guerreiro (2024)and the *forecast-unlearning matrix* $\mathcal{P}_h$ in our framework:

R-P recursion. The forecast-revision matrix sequence R_h and the forecast-unlearning matrix sequence $\mathcal{P}_h$, satisfy for all $0 \leq h \leq T$:

$$R_{h+1} = R_h - \mathcal{P}_h.$$

Further, given a fixed horizon or length T of the transition sequence around the nonstochastic steady state, $R_T = \mathcal{P}_T$.

Proof. By definition of the shifter matrix, it follows that for all $h \geq 0$:

$$R_h - R_{h+1} = \begin{pmatrix} 0 & 0_h' \\ 0_h & J \end{pmatrix} - \begin{pmatrix} 0 & 0_{h+1}' \\ 0_{h+1} & J \end{pmatrix} = \begin{pmatrix} 0 & 0_h' \\ 0_h & \mathcal{F} \end{pmatrix} = \mathcal{P}_h \tag{A.8}$$

Where the second equality uses the fact that as shown in Auclert et al. (2021)

$$J_{t,s} - J_{t-1,s-1} = \mathcal{F}_{t,s},$$

noting that by construction of the shifter matrix the operation is subtracting from each element the first Jacobian the element located to its upper left along the diagonal, and that the first row and column of the Jacobian and fake-news matrix coincide by construction, as in Auclert et al. (2021). The result $R_T = \mathcal{P}_T$ follows immediately from the latter observation.[26]

Equation (6) is key for the proofs of the two following corollaries, showing equivalence of the solution sets defined by the two solvers, by establishing a key link between the Jacobian-shifter matrix and the fake-news shifter matrix, or the *forecast-revision* matrix R_h and the *forecast-unlearning* matrix $\mathcal{P}_h$:

$$R_{h+1} = R_h - \mathcal{P}_h.$$

The economic intuition behind the above identity was expounded earlier. That is, the impact of tomorrow's forecast revision from tomorrow onwards can be backed out from the impact of the forecast revision today from today onwards and the impact tomorrow of fully unlearning the shocks expected today. This is because, as noticed to be true for a generalised sequence-space environments with time-varying expectations as inputs alongside the true aggregate input sequence:

- The shock to expectations between yesterday and today is equivalent to the combination of the newly expected shocks and fully unlearning (i.e. negative shocks of the same magnitude) expectations formed yesterday.

[26] And equal to the displacement in optimal policies from a shock to the input in the same period $\mathcal{Y}_0/dx$, while the distribution is fixed at the steady state as in Auclert et al. (2021).

- The impact of unlearning shocks between yesterday and today, as shown in Auclert et al. (2021), amounts to the impact driven by end of-period distributional adjustments as optimal policies converge back to the steady state.

Note that a further immediate implications of Lemma 1 is that that when $\mathbf{I} = \Lambda_h = \Lambda_{h+1}$ for all h, i.e. with expectations constant at the perfect foresight/FIRE ones, the above collapses to the FIRE solver in Auclert et al. (2021). This is the case because, using Lemma 1:

$$\sum_h \mathcal{P}_h \mathbf{I} = \sum_h \mathcal{P}_h = R_0 - R_1 + R_1 - R_2 + R_2 - ... - R_{T-1} + R_{T-1} - R_T + \mathcal{P}_T$$

$$\mathcal{P}_h = R_0 - R_T + \mathcal{P}_T = J$$

The economic intuition within the developed *forecast-unlearning* framework is that, because there is perfect foresight or FIRE, it is exactly as if the initially (correct) expectations are unlearned and reformed renewed at each point in time. However, the impacts of this shock combination – equivalent to maintaining expectations constant – offset each other fully, so that agents only end up responding to the news at time zero. That a restriction to FIRE replicates the same solver as in Auclert et al. (2021) is a shared property with the one in Bardóczy and Guerreiro (2024), which is not really surprising because as we now go on to show the two are fully equivalent.[27]

A.4.1 Necessity corollary

Having fleshed out the inductive logic sustaining the solver above and Lemma 1, we now prove the necessity corollary. In fact, proving this shows also that the solver can be backed-out directly by manipulating the Bardoczy-Guerreiro generalised solver. Recall this is for $h \geq 1$:

$$\hat{J} = J\Lambda_0 + \sum_h R_h(\Lambda_h - \Lambda_{h-1})$$

Noting that the forecast revision matrix at time zero $R_0 = J$, we can accordingly rewrite the solver for $h \geq 1$ as:

$$\hat{J} = R_0\Lambda_0 + \sum_h R_h(\Lambda_h - \Lambda_{h-1})$$

$$\hat{J} = (R_0 - R_1)\Lambda_0 + (R_1 - R_2)\Lambda_1 + (R_2 - R_3)\Lambda_2 + ... + R_T\Lambda_T.$$

The *efficient* fake-news matrix SSJ solver for linear deviations from FIRE is thus:

$$\hat{J} = \sum_{h\geq 0} \mathcal{P}_h\Lambda_h \qquad \mathcal{P}_h = \begin{pmatrix} 0 & 0_h' \\ 0_h & \mathcal{F} \end{pmatrix} \tag{A.10}$$

where we have made use of the definition of the forecast-unlearning matrix $\mathcal{P}_h$ and the result in Lemma 1 in order to substitute in $R_h - R_{h+1} = \mathcal{P}_h$ for $h \leq T-1$ and $R_T = \mathcal{P}_T$.

A.4.2 Sufficiency corollary

The above shows that if some SSJ Jacobian, i.e. a candidate solver, satisfies the equation in Bardóczy and Guerreiro (2024), then it is also an efficient fake-news generalised solver. Here we prove the reverse is true as well, so that the classes of solutions to DSGE/HANK models in the sequence-space defined

[27] An alternative proof of consistency and convergence to the Auclert et al. (2021) SSJ-solver under FIRE is in fact using the next two corollaries establishing equivalence to the Bardóczy and Guerreiro (2024), which obviously converges to the FIRE case as its expression makes clear:

$$\hat{J} = J\Lambda_0 + \sum_h R_h(\Lambda_h - \Lambda_{h-1}) = J \tag{A.9}$$

for $\Lambda_0 = \Lambda_h = \mathbf{I}$ for all h.

by the Bardóczy and Guerreiro (2024) SSJ solver and the fake-news efficient solver are *coextensive*. To show this, consider equation (6) which defines the efficient fake-news solver, noting again that $\mathcal{P}_0 = \mathcal{F}$:

$$\hat{J} = \sum_{h \geq 0} \mathcal{P}_h \Lambda_h$$

$$\hat{J} = \mathcal{F}\Lambda_0 + \sum_{h \geq 1}^{T-1} (R_h - R_{h+1})\Lambda_h + R_T \Lambda_T$$

Where we used again the observation the fact that $\mathcal{P}_h = R_h - R_{h+1}$, and that $\mathcal{P}_T = R_T$ in the last period. Then the above can be re-arranged as:

$$\hat{J} = \mathcal{F}\Lambda_0 + R_1\Lambda_1 - R_2\Lambda_1 + R_2\Lambda_2 - R_3\Lambda_2 + ... + R_{T-1}\Lambda_T - R_T\Lambda_{T-1} + R_T\Lambda_T$$

$$\hat{J} = \mathcal{F}\Lambda_0 + R_1\Lambda_1 + \sum_{h \geq 2}^{T} R_h(\Lambda_h - \Lambda_{h-1})$$

Adding $R_1(\Lambda_0 - \Lambda_0) = 0$ to the right-hand side, we have that the efficient fake-news SSJ-solver in (6) satisfies:

$$\hat{J} = (\mathcal{F} + R_1)\Lambda_0 + \sum_{h \geq 1}^{T} R_h(\Lambda_h - \Lambda_{h-1})$$

However, note that

$$\mathcal{F} + R_1 = J,$$

which follows from the more general result in equation (6) or Lemma 1:

$$R_{h+1} = R_h - P_h,$$

with $h = 0$, for which by construction of the forecast-revision and forecast-unlearning matrices $R_0 = J$ and $\mathcal{P}_0 = \mathcal{F}$, implying the above equation can be written as

$$\hat{J} = J\Lambda_0 + \sum_{h \geq 1} R_h(\Lambda_h - \Lambda_{h-1}).$$

Hence if a solver $\hat{J}$ is an efficient fake-news solver, it is also a Bardóczy and Guerreiro (2024) generalised SSJ-solver. Because we proved that the reverse of the conditional is also true, the two solvers are fully equivalent.

A.5 Generalisation to Time-Varying Coefficients γ_k, δ_k

The Kohlhas-Walther reduced form equation for the evolution of the forecast with time-varying extrapolation and news-revision coefficients is:

$$x_{t+k} - f_{it}x_{t+k} = \alpha_{i,k} + \gamma_k x_t + \delta_k(\bar{f}_t x_{t+k} - \bar{f}_{t-1}x_{t+k}) + \varepsilon_{i,t|t+k}.$$

Relaxing the framework presented in the main text, the above equation allows for the strength of the extrapolation and news-revision effects to vary in function of the time elapsed between the time at which the forecast is formed and the timing of the forecast variable. This also implies that the strength of extrapolation and news-revision is the same at times t and r when respectively forecasting x_{t+k} and x_{r+k}. Again, let $\bar{f}_t x_{t+k} := \int_j f_{jt}x_{t+k}dP(j)$ denote the average expectation in the cross-section. Integrating over the measure of households and re-arranging:

$$x_{t+k} - \bar{f}_t x_{t+k} = \alpha_{i,k} + \gamma_k x_t + \delta_k(\bar{f}_t x_{t+k} - \bar{f}_{t-1}x_{t+k}).$$

$$\bar{f}_t x_{t+k} = c_k + \frac{1}{1+\delta_k}(\delta_k \bar{f}_{t-1}x_{t+k} + x_{t+k} - \gamma_k x_t),$$

where $c_k = -\frac{1}{1+\delta_k}\int_j \alpha_{j,k} + \epsilon_{j,t|t+k}P(dj) = -\frac{1}{1+\delta_k}\int_j \alpha_{j,k}P(dj)$. We impose again a no-bias assumption on the constant term c_k for all $k \in N$, so that $\bar{f}_t x_{ss} = \bar{f}_{t-1}x_{ss} = \bar{f}_{ss}x_{ss} = x_{ss}$:

$$c_k = \frac{\gamma_k}{1+\delta_k}x_{ss}.$$

Note that, using the definition of c_k the no-bias assumption implies a restriction on the intercept or household fixed effects, $\forall k$:

$$\gamma_k x_{ss} + \int_j \alpha_{j,k}P(dj) = 0.$$

One interpretation of the above is that the bias of the agents when forming individual forecasts exactly offsets the extrapolation effect at the steady state. The reason is that, as the news revision channel ceases to play a role at the steady state (the public information set becomes stationary), then an unbiased aggregate forecast requires that agents adjust their idiosyncratic bias so as to offset aggregate extrapolation. For an aggregate input with a positive-valued steady state, if we require $\gamma_k \leq 0$, consistent with the micro-founded model of extrapolation from asymmetric attention and empirical estimates in Kohlhas and Walther (2021), then the no-bias assumption implies

$$\int_j \alpha_{j,k}P(dj) < 0.$$

With this in mind, the generalisation of the Kohlhas-Walther mapping to time-varying extrapolation and news revision parameters proceeds in the same way as for the case with time-invariant ones. At a first order perturbation around the non-stochastic steady state, the forecast at time t of an aggregate input k periods ahead satisfies:

$$\bar{f}_t(x_{ss} + dx_{t+k}) = c_k + \frac{1}{1+\delta_k}\left(\delta_k \bar{f}_{t-1}(x_{ss} + dx_{t+k}) + x_{ss} + dx_{t+k} - \gamma_k(x_{ss} + dx_t)\right)$$

Using the no-bias assumption pinning down each c_k:

$$\bar{f}_t(x_{ss} + dx_{t+k}) = x_{ss} + \frac{1}{1+\delta_k}\left(\delta_k \bar{f}_{t-1}dx_{t+k} + dx_{t+k} - \gamma_k dx_t\right)$$

$$\bar{f}_t dx_{t+k} = \frac{1}{1+\delta_k}\left(\delta_k \bar{f}_{t-1}dx_{t+k} + dx_{t+k} - \gamma_k dx_t\right)$$

This yields, like previously, a first order difference equation in the forecasts formed at times $t = 0...T$. Solving the equation backwards, together with the usual boundary conditions:

$$\bar{f}_t dx_{t+k} = \frac{1}{1+\delta_k}\left[\delta_k \frac{1}{1+\delta_{k-1}}\left(\delta_{k-1}\bar{f}_{t-1}dx_{t+k} + dx_{t+k} - \gamma_{k+1}dx_{t-1}\right) + dx_{t+k} - \gamma_k dx_t\right]$$

$$\bar{f}_t dx_{t+k} = \sum_{j=0}^{t}\left(\prod_{s=j}^{t-1}\frac{\delta_{k+1-(t-s)}}{1+\delta_{k-(t-s)}}\right)\frac{1}{1+\delta_k}dx_{t+k} - \sum_{j=0}^{t}\gamma_{k+(t-j)}\left(\prod_{s=j}^{t-1}\frac{\delta_{k+1-(t-s)}}{1+\delta_{k-(t-s)}}\right)\frac{1}{1+\delta_k}dx_j.$$

B Two Asset HANK Model

B.1 Model Equations

This appendix collects the equilibrium conditions of the model.

B.2 Households

There are six groups of households $g \in \{1, \dots, 6\}$ with population shares ω_g and group-specific discount factors β_g. Within each group, households face idiosyncratic productivity risk z governed by a Markov chain with time-invariant transition matrix Π. Households choose consumption c and liquid asset holdings a' subject to a borrowing constraint $a' \geq \underline{a}$.

Budget constraint. For a household of type g with productivity z and liquid assets a:

$$c + a' = (1 + r^p)a + (1 - \tau)w\, h\, e_g\, z + d_g + T^{gov} \tag{B.1}$$

where r^p is the return on liquid assets, τ is the labour income tax rate, w is the real wage, h is hours worked (determined by the union, see below), e_g is the group-specific income share, d_g is the distribution from the illiquid portfolio to group g, and T^{gov} is a lump-sum government transfer.

Euler equation (liquid asset). The household's optimality condition is:

$$c^{-1/v} \geq \tilde{\beta}_g\, \mathbb{E}_t \left[(1 + r^p_{t+1})\, c_{t+1}^{-1/v} \right], \quad \text{with equality if } a' > \underline{a} \tag{B.2}$$

where v is the elasticity of intertemporal substitution and $\tilde{\beta}_g = \beta_g \cdot e^{\eta_{C,t}}$ incorporates the discount factor shock.

Aggregation. Aggregate consumption C_t, liquid asset demand A_t, disposable income Y_t^S, and the marginal rates of substitution MRSC_t, MRSH_t are obtained by integrating over the stationary distribution:

$$C_t = \sum_g \frac{\omega_g}{\mu_g} \int c_g(z, a)\, \mathrm{d}\mu_g(z, a), \qquad A_t = \sum_g \frac{\omega_g}{\mu_g} \int a'_g(z, a)\, \mathrm{d}\mu_g(z, a) \tag{B.3}$$

where μ_g is the within-group measure.

B.2.1 Illiquid Asset and Portfolio Inattention

Illiquid asset law of motion. For each group g, illiquid assets are subject to a law of motion:

$$X_t^g = (1 + r_t^{ill})X_{t-1}^g - d_t^g \tag{B.4}$$

where X_t^g is group g aggregate illiquid wealth, r_t^{ill} is the return on illiquid assets, and d_t^g is the aggregate distribution from illiquid to liquid accounts.

Distribution policy under sticky expectations. The distribution policy for a household ex-ante belonging in group g follows a simple rule Auclert et al. (2020):

$$d_t^g = \frac{\bar{r}^{ill}}{1 + \bar{r}^{ill}}\, \tilde{X^g}_t + \chi \left(\tilde{X^g}_t - (1 + \bar{r}^{ill})\bar{X^g} \right) \tag{B.5}$$

where $\tilde{X^g}_t = \mathbb{E}(X_t^g | \Omega_t^g)$ denotes the expected illiquid wealth (defined below) conditional on the g-th group information set Ω_t^g at time t and $\chi \to 0^+$ governs the sensitivity of distributions to perceived deviations from the group-specific steady state illiquid wealth position X_g.

Sticky portfolio expectations. Average portfolio expectations are subject to sticky information sets, which are assumed to update infrequently in a time-dependent fashion with probability $(1-\theta)$ per period. Under a LLN:

$$\tilde{X^g}_t = \int_{\mathbb{R}_+} \tilde{X}_t^g(j)\mu(dj) = (1-\theta)(1+r_t^{ill})X_{t-1}^g + \theta\, \tilde{X^g}_{t-1} \tag{B.6}$$

where $X_t^g(j)$ is the expectation of a household j in group g who last updated their information set j periods ago and μ is the probability mass function $\mu(s) = \theta^s$. By a similar argument to that presented in footnote (8), it is clear that asymptotically expectations on the illiquid portfolio value converge to the group-specific steady state. When $\theta = 0$, households have full-information rational expectations over their portfolio: $\tilde{X}_t = (1 + r_t^{ill})X_{t-1}$.

Aggregate consistency. Total illiquid wealth X_t in the economy satisfies an aggregate consistency condition at the end of each period:

$$X_t = \sum_{\mathcal{G}} X_t^g \mu_g \tag{B.7}$$

B.2.2 Production

Final good firm (production function). Output is produced with a Cobb–Douglas technology:

$$Y_t = Z_t\, K_{t-1}^{1-\alpha_h}\, (N_E \cdot h_t)^{\alpha_h} \tag{B.8}$$

where $Z_t = e^{\eta_{z,t}}$ is total factor productivity, K_{t-1} is the capital stock, N_E is aggregate labour efficiency units, and h_t is hours worked. Rearranging for hours:

$$h_t = \frac{1}{N_E}\left(\frac{Y_t}{Z_t\, K_{t-1}^{1-\alpha_h}}\right)^{1/\alpha_h} \tag{B.9}$$

Marginal cost.

$$mc_t = \frac{w_t}{\alpha_h\, Y_t/(h_t\, N_E)} \tag{B.10}$$

Rental rate of capital.

$$r_t^k = mc_t\,(1-\alpha_h)\,\frac{Y_t}{K_{t-1}} \tag{B.11}$$

Dividends of the intermediate firm.

$$\mathrm{Div}_t = (1 - mc_t)\, Y_t \tag{B.12}$$

B.2.3 Capital Firm

Capital accumulation.

$$K_t = I_{t-1} + (1-\delta)K_{t-1} \tag{B.13}$$

Investment Euler equation (Tobin's Q). With convex adjustment costs parameterised by ψ:

$$Q_t = 1 + \frac{\psi}{2}\left(\frac{I_t}{I_{t-1}} - 1\right)^2 + \psi\left(\frac{I_t}{I_{t-1}} - 1\right)\frac{I_t}{I_{t-1}} - \frac{\psi}{1 + r_{t+1} + \eta_{rp,t+1}}\left(\frac{I_{t+1}}{I_t} - 1\right)\left(\frac{I_{t+1}}{I_t}\right)^2 \tag{B.14}$$

Asset price of capital.

$$Q_t = \frac{1}{1 + r_{t+1} + \eta_{rp,t+1}}\left[r_{t+2}^k + (1-\delta)Q_{t+1}\right] \tag{B.15}$$

Capital firm dividends.

$$\text{Div}_t^K = r_t^k K_{t-1} - I_{t-1} - \frac{\psi}{2}\left(\frac{I_{t-1}}{I_{t-2}} - 1\right)^2 I_{t-1} \tag{B.16}$$

Measured investment.

$$\text{Inv}_t = I_{t-1} \tag{B.17}$$

B.2.4 Nominal Block

Price Phillips curve. With Calvo parameter ζ_p, Kimball super-elasticity ν_p, and demand elasticity ϵ^p:

$$\pi_t = \pi_{t-1} + \frac{(1-\zeta_p)\left(1 - \frac{\zeta_p}{1+\bar{r}}\right)}{\zeta_p} \cdot \frac{\epsilon^p}{\epsilon^p + \nu_p - 1}\left(mc_t - \frac{\epsilon^p - 1}{\epsilon^p}\right) + \frac{1}{1+r_t}\left(\pi_{t+1} - \pi_t\right) + \eta_{p,t} \tag{B.18}$$

Wage Phillips curve. With Calvo parameter ζ_w, labour supply elasticity (Frisch) φ, and union markup ϵ:

$$\pi_t^w = \pi_{t-1} + \frac{(1-\zeta_w)(1-\bar{\beta}\zeta_w)}{\zeta_w} \cdot \frac{\epsilon}{\epsilon + \nu_w - 1}\left(\phi_l \frac{\text{MRSH}_t}{\text{MRSC}_t} - \frac{\epsilon - 1}{\epsilon}\right) + \bar{\beta}\left(\pi_{t+1}^w - \pi_t\right) + \eta_{w,t} \tag{B.19}$$

where ϕ_l is a steady-state labour disutility scaling parameter and $\bar{\beta}$ is the average discount factor.

In the above we further set the steady state wage markup $\frac{\epsilon}{\epsilon-1} = \frac{\epsilon^p}{\epsilon^p-1}$ so that it matches the price markup, and similarly for the Kimball super-elasticities $v_w = v_p = 10$. This is the same as in the original model of Auclert et al. (2020).

Real wage evolution.

$$\ln\left(w_t(1+\pi_t)\right) - \ln w_{t-1} = \pi_t^w \tag{B.20}$$

Taylor rule.

$$i_t = \rho_m\, i_{t-1} + (1-\rho_m)\left[r^* + \phi_\pi\, \pi_t + \phi_y \frac{Y_t - \bar{Y}}{\bar{Y}}\right] + \eta_{r,t} \tag{B.21}$$

Throughout, we set the Taylor rule parameter $\phi_y = 0$ so that the central bank does not respond to deviations of output from its steady state value.

Fisher equation.

$$1 + r_t = \frac{1+i_t}{1+\pi_{t+1}} \tag{B.22}$$

B.2.5 Asset Pricing and Returns

Return on liquid assets.

$$r_t^p = r_{t-1} - \xi \tag{B.23}$$

where ξ is the financial intermediation spread.

Government bond price. With maturity structure governed by decay parameter δ_b:

$$q_t^b = \frac{1 + \delta_b\, q_{t+1}^b}{1+r_t} \tag{B.24}$$

Equity prices. No-arbitrage conditions for the intermediate firm and capital firm:

$$p_t = \frac{p_{t+1} + \mathrm{Div}_{t+1}}{1 + r_t + \eta_{rp,t}} \tag{B.25}$$

$$p_t^K = \frac{p_{t+1}^K + \mathrm{Div}_{t+1}^K}{1 + r_t + \eta_{rp,t}} \tag{B.26}$$

B.2.6 Fiscal Policy

Government spending.

$$G_t = \bar{G} + \bar{Y}\,\eta_{g,t} \tag{B.27}$$

Tax rule.

$$\tau_t = \bar{\tau} + \gamma_\tau\,\bar{q}^b\,\frac{B_{t-1} - \bar{B}}{\bar{Y}} \tag{B.28}$$

Government budget constraint.

$$q_t^b B_t = G_t + (1 + \delta_b q_t^b)B_{t-1} + T^{gov} - \tau_t\left(Y_t - r_t^k K_{t-1} - \mathrm{Div}_t\right) \tag{B.29}$$

B.2.7 Market Clearing

Asset market.

$$A_t + X_t = q_t^b B_t + p_t + p_t^K \tag{B.30}$$

Goods market.

$$Y_t = C_t + G_t + I_{t-1} + \frac{\psi}{2}\left(\frac{I_{t-1}}{I_{t-2}} - 1\right)^2 I_{t-1} + \xi\,A_{t-1} \tag{B.31}$$

Funds market.

$$(1 + r_t^p + \xi)A_{t-1} + (1 + r_t^{ill})X_{t-1} = p_t + \mathrm{Div}_t + p_t^K + \mathrm{Div}_t^K + (1 + \delta_b q_t^b)B_{t-1} \tag{B.32}$$

B.2.8 Exogenous Shock Processes

Five of the seven structural shocks follow AR(1) processes:

$$\eta_{C,t} = \rho_C\,\eta_{C,t-1} + \varepsilon_t^C \qquad \text{(discount factor)} \tag{B.33}$$

$$\eta_{rp,t} = \rho_{rp}\,\eta_{rp,t-1} + \varepsilon_t^{rp} \qquad \text{(risk premium)} \tag{B.34}$$

$$\eta_{z,t} = \rho_z\,\eta_{z,t-1} + \varepsilon_t^z \qquad \text{(TFP)} \tag{B.35}$$

$$\eta_{r,t} = \rho_r\,\eta_{r,t-1} + \varepsilon_t^r \qquad \text{(monetary policy)} \tag{B.36}$$

$$\eta_{g,t} = \rho_g\,\eta_{g,t-1} + \varepsilon_t^g \qquad \text{(government spending)} \tag{B.37}$$

The price and wage markup shocks follow ARMA(1,1) processes as in Smets and Wouters (2007):

$$\eta_{p,t} = \rho_p\,\eta_{p,t-1} + \varepsilon_t^p - \theta_p\,\varepsilon_{t-1}^p \qquad \text{(price markup)} \tag{B.38}$$

$$\eta_{w,t} = \rho_w\,\eta_{w,t-1} + \varepsilon_t^w - \theta_w\,\varepsilon_{t-1}^w \qquad \text{(wage markup)} \tag{B.39}$$

All innovations $\varepsilon_t^j \sim \mathcal{N}(0, \sigma_j^2)$ are mutually independent.

B.3 Model Dynamics

Figure B.1: Replicating Auclert et al. (2020)

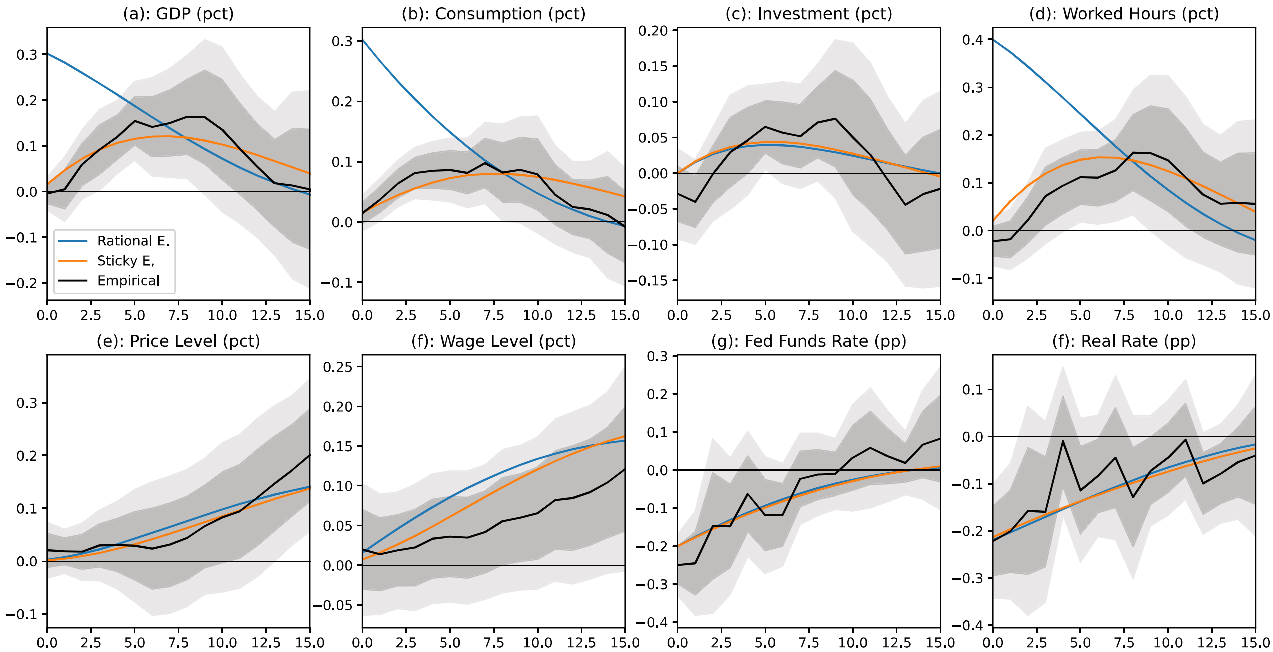

Note: This figure replicates the impulse response to a monetary policy shock under rational and sticky expectations as in Auclert et al. (2020). The sticky impulse responses were IRF matched to the empirical evidence of a local projection to a Romer and Romer (2004) monetary policy shock.

B.4 Business Cycle Estimation

B.4.1 Data and Estimation

This appendix describes the construction of the dataset used to estimate the model. The sample runs from 1978 Q4 to 2018 Q4. The first seven observables follow the transformations of Smets and Wouters (2007) closely; two additional series capture realised inflation forecast errors for professionals and households.

All macroeconomic aggregates are quarterly and sourced from FRED. The raw series are: nominal GDP, nominal personal consumption expenditures, nominal business fixed investment, compensation per hour (nonfarm business sector), an index of aggregate hours (average weekly hours $\times$ civilian employment), the GDP deflator, civilian non-institutional population, and the effective federal funds rate. Real quantities are deflated by the GDP deflator throughout.

The forecast error series are constructed from two survey sources:

- **Survey of Professional Forecasters (SPF):** individual-level microdata from the Philadelphia Fed. The four-quarter-ahead forecast of GDP deflator inflation is computed as $100 \times (X_{t+4}/X_t - 1)$ from each respondent's level forecasts and averaged across forecasters each quarter to obtain $\widehat{\pi}_t^{\text{SPF}}$.
- **Michigan Survey of Consumers:** individual-level microdata on expected price changes over the next 12 months (`px1`), aggregated to a quarterly frequency using survey sampling weights to obtain $\widehat{\pi}_t^{\text{Mich}}$.

The realised four-quarter inflation ending at t is $\Pi_t \equiv \pi_t + \pi_{t-1} + \pi_{t-2} + \pi_{t-3}$, where $\pi_t = 100 \ln(P_t/P_{t-1})$ is log quarterly GDP-deflator inflation. The forecast errors are then

$$\text{FE}_t^{\text{SPF}} = \Pi_t - \widehat{\pi}_{t-3}^{\text{SPF}}, \qquad \text{FE}_t^{\text{Mich}} = \Pi_t - \widehat{\pi}_{t-3}^{\text{Mich}},$$

Figure B.2: Business cycle data

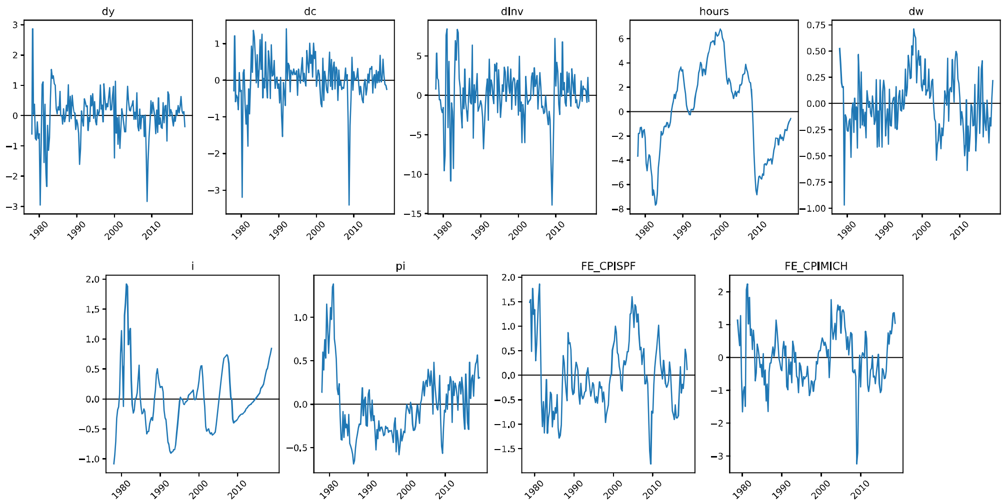

Note: Figure plots the transformed data used to estimate the Two Asset Hank Model. The sample range spans 1978 to 2018.

comparing outturn inflation over the four quarters ending at t to the expectation formed three quarters earlier. Because three lags of the survey series are required, the effective sample begins in 1978 Q4.

Table B.1 summarises the nine observables. All series are linearly detrended (OLS residuals from a regression on a linear time trend) rather than simply demeaned; this accounts for low-frequency movements in trend growth and the secular decline in nominal interest rates over the sample.

Variable	**Transformation**
Δy_t	$100 \times \left[\ln(GDP_t/P_t/N_t) - \ln(GDP_{t-1}/P_{t-1}/N_{t-1})\right]$, linearly detrended
Δc_t	$100 \times \left[\ln(C_t/P_t/N_t) - \ln(C_{t-1}/P_{t-1}/N_{t-1})\right]$, linearly detrended
Δi_t	$100 \times \left[\ln(I_t/P_t/N_t) - \ln(I_{t-1}/P_{t-1}/N_{t-1})\right]$, linearly detrended
h_t	$100 \times \ln(\bar{h}_t \cdot E_t/N_t)$, linearly detrended
Δw_t	$100 \times \left[\ln(W_t/P_t) - \ln(W_{t-1}/P_{t-1})\right]$, linearly detrended
i_t	Federal funds rate divided by 4 (quarterly), linearly detrended
π_t	$100 \times \ln(P_t/P_{t-1})$, linearly detrended
$\mathrm{FE}_t^{\mathrm{SPF}}$	$\Pi_t - \widehat{\pi}_{t-3}^{\mathrm{SPF}}$: realised four-quarter GDP-deflator inflation minus the SPF mean four-quarter-ahead forecast form
$\mathrm{FE}_t^{\mathrm{Mich}}$	$\Pi_t - \widehat{\pi}_{t-3}^{\mathrm{Mich}}$: realised four-quarter GDP-deflator inflation minus the Michigan Survey weighted-mean expected

Table B.1: Observable variables and their transformations. P_t denotes the GDP deflator, N_t is civilian non-institutional population (normalised to an index), E_t is civilian employment (normalised), $\bar{h}_t$ is average weekly hours (normalised), W_t is nominal compensation per hour, and $\Pi_t \equiv \sum_{j=0}^{3} \pi_{t-j}$ is realised four-quarter inflation.

Figure B.3: Business cycle structural decomposition

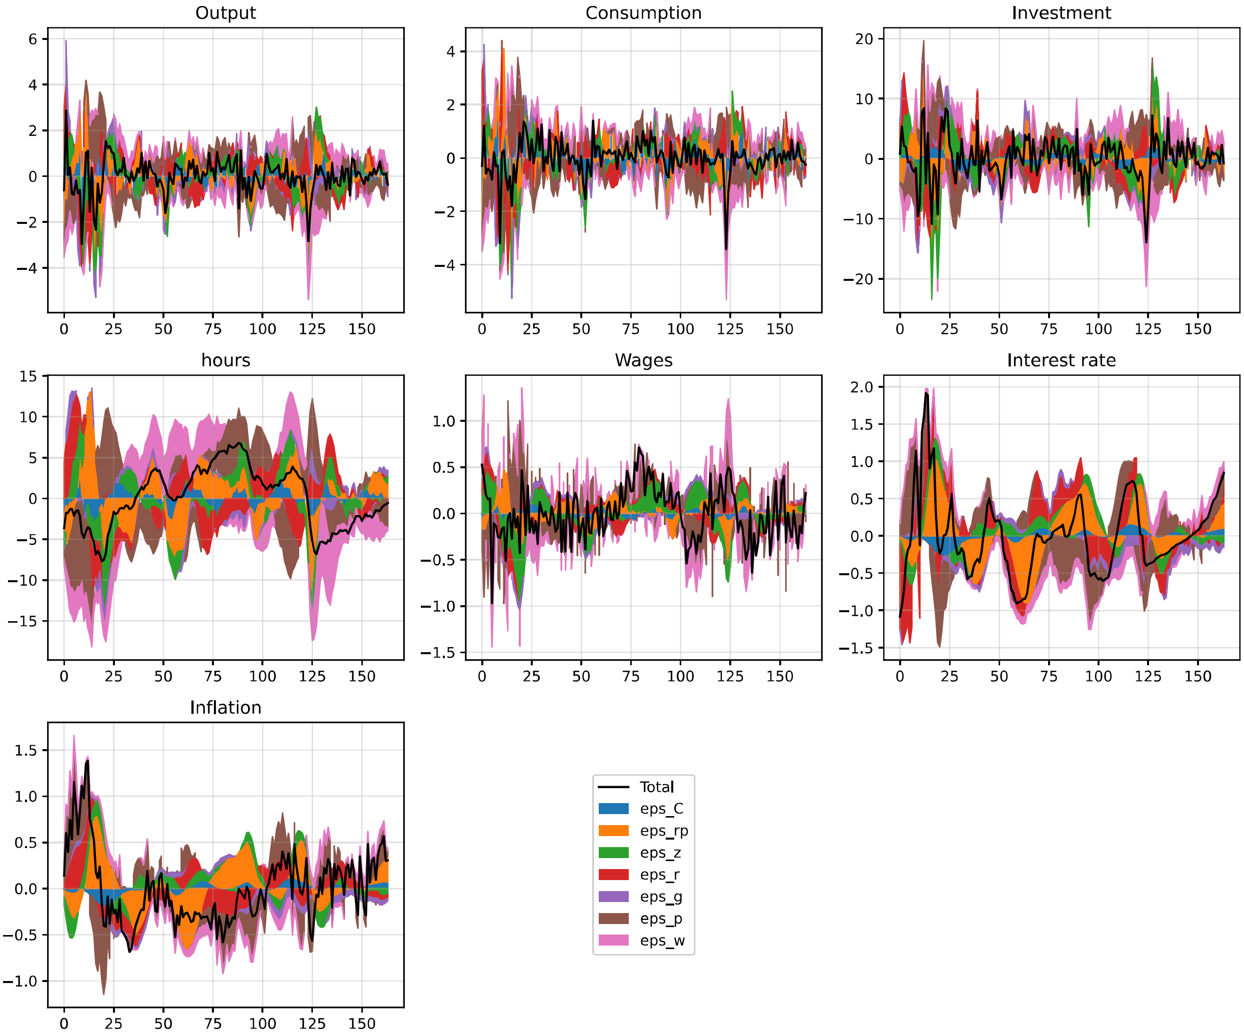

Note: Figure plots the structural decomposition of the aggregate time series on the components driven by each of the seven structural shocks for model (3). Innovations to the shocks are obtained by filtering against business cycle data using the MA representation of the process at the estimated parameters.

B.4.2 Further Results

Figure B.4: Other Shocks

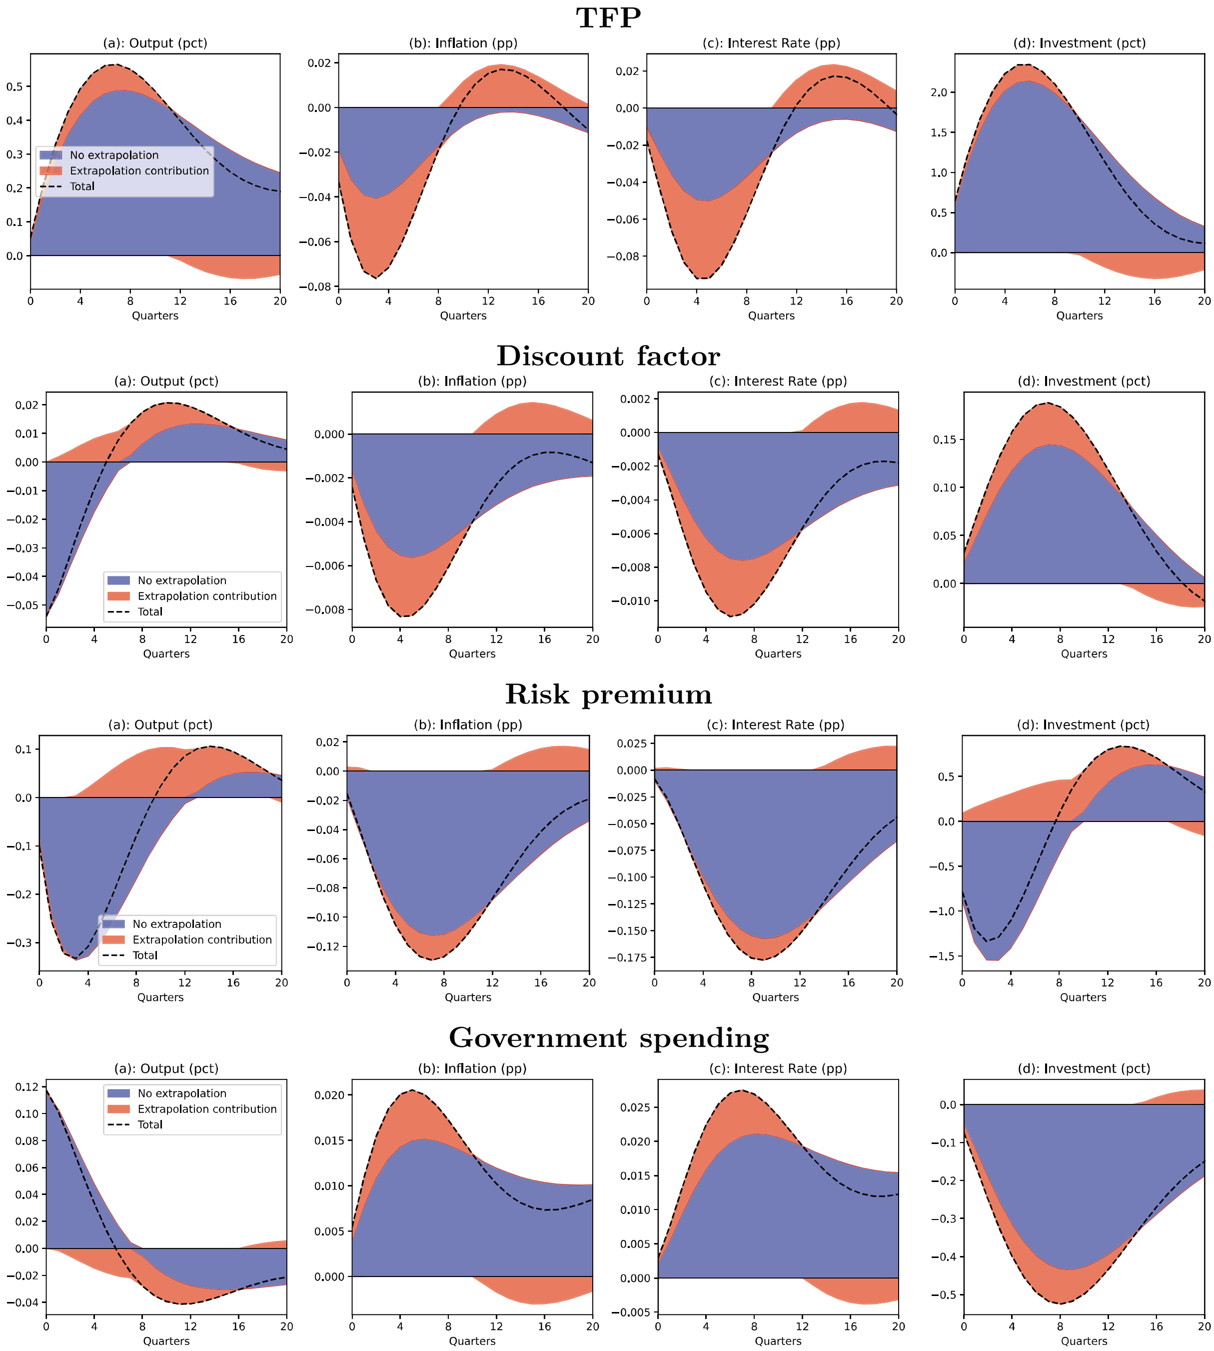

Price markup

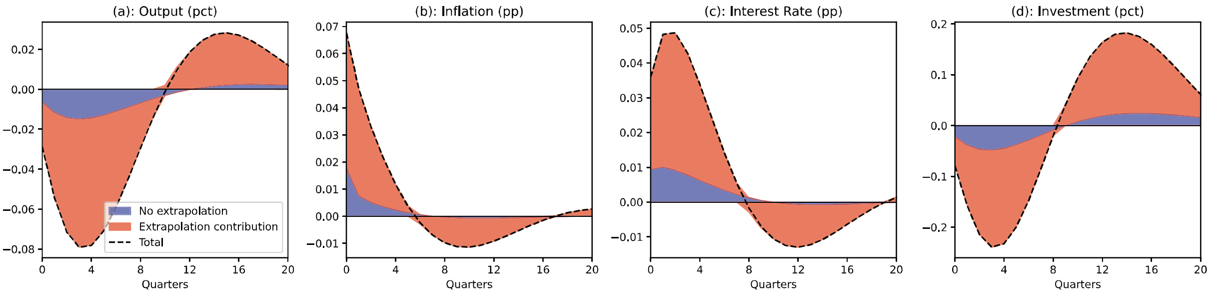

Wage markup

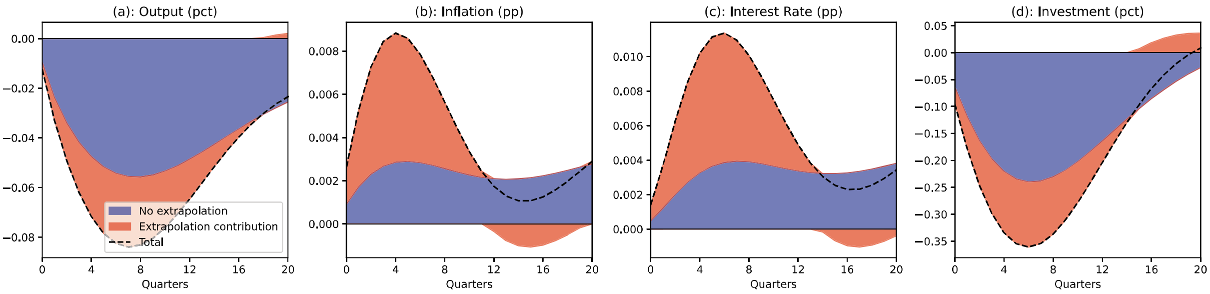

Note: Impulse response to a one standard deviation shock decomposed as described in Figure 2.

Figure B.5: Model autospectral density functions (ASDFs)

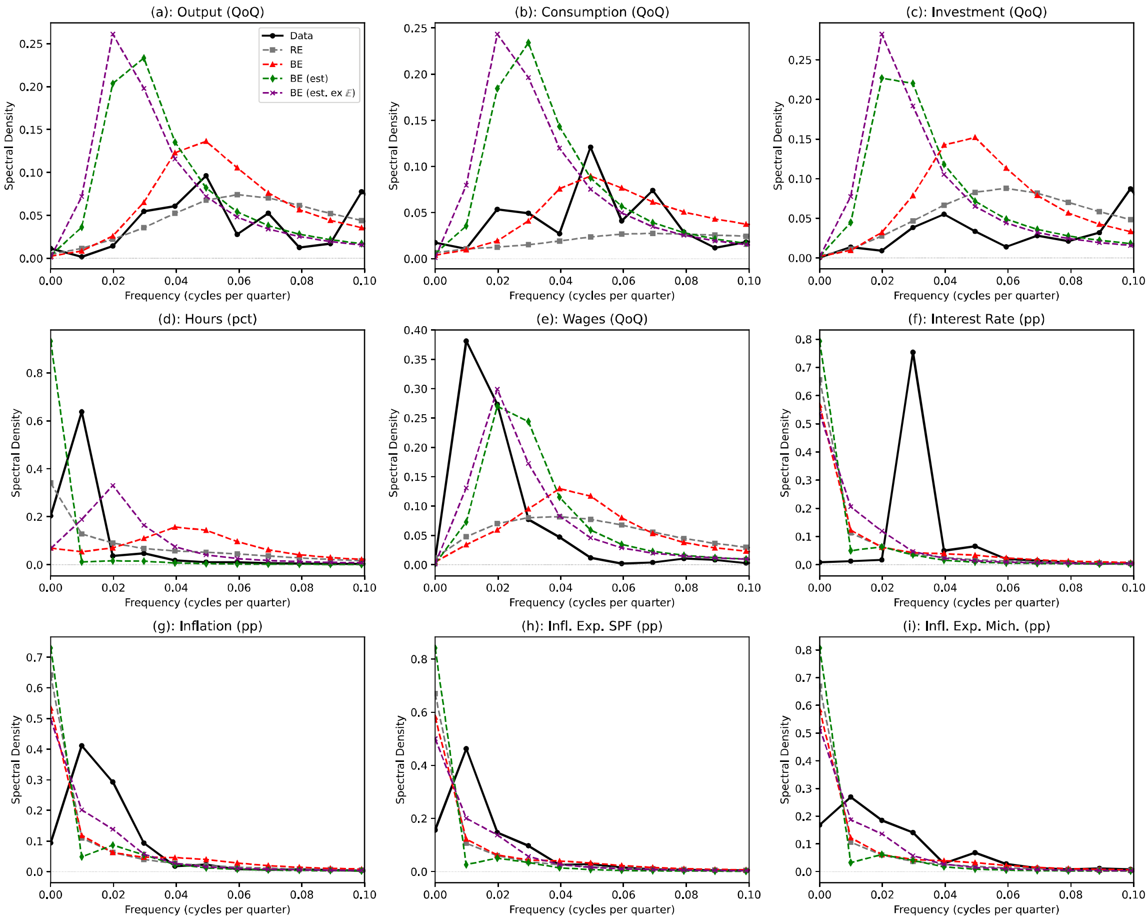

Note: Figure reports the autospectral density functions for the macroeconomic time series used in estimation, decomposing the time-domain signal into constituent frequency components, for each of the four specifications in Table 2. They are compared to their empirical equivalent for the same series, employing the ACFs estimated above. For comparability, the power spectra for each process and model are normalised by the total power carried by the signal over time. For all outputs, cross-model variation for frequencies $\lambda \in [0.1, 0.5]$ cycles per quarter is negligible, hence only the lower frequencies through to 0.1 cycles per quarters are shown. [28]

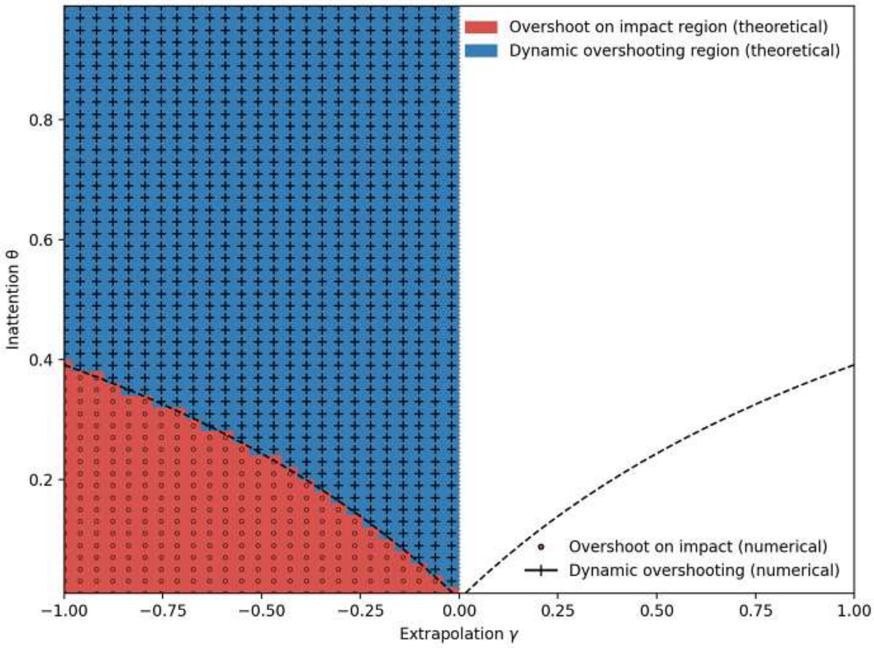

Figure B.6: Necessity & sufficiency of over-reactive extrapolation for dynamic overshooting: numerical test of theoretical conditions

Note: The figure verifies the theoretical results in propositions (1-3). The figure plots the results from a set of numerical experiments on one-period ahead forecast errors against the IRF in figure 1 for a meshgrid of values of the extrapolation parameter *gamma* and inattention parameter $\theta = \delta/1 + \delta$. The white region corresponds to no dynamic overshooting, while the blue and red regions respectively map to dynamic overshooting (interior timing) and on-impact overshooting. The dashed line plots the theoretical boundary between the two regions derived in proposition (2). Numerical tests align with the theoretical results, confirming the necessity of over-reactive extrapolation for dynamic overshooting and the sufficiency of the conditions in proposition (2) for interior timing: (i) without over-reactive extrapolation no overshooting occurs, (ii) over-reactive extrapolation must not be too large relative to the inattention parameter, and (iii) larger inattention loosens the upper boundary on extrapolation consistent with dynamic overshooting.

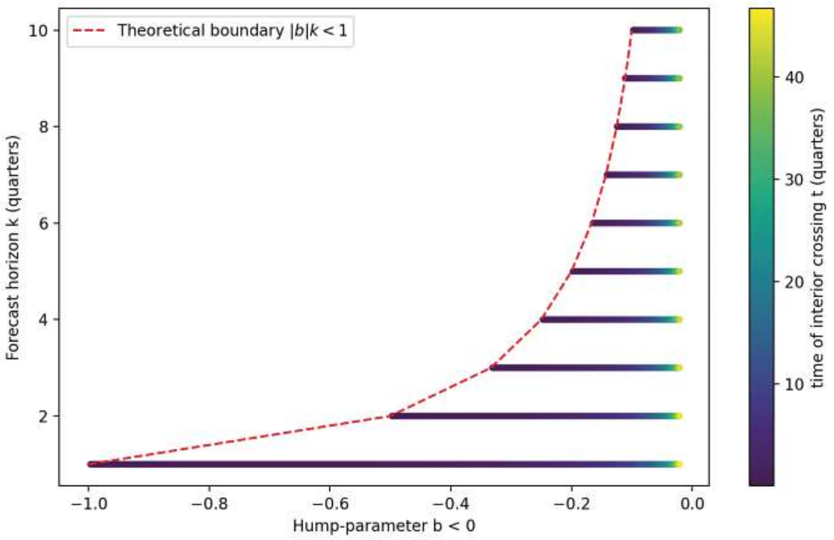

Figure B.7: Interior dynamic overshooting times.

Note: The figure shows the existence and timing of interior overshoot under purely sticky expectations for the case with an oscillatory hump-shaped IRF ($b < 0$). The dashed line plots the theoretical interior time derived for such case in footnote (16).