Spend What You Learn: Heterogeneous Agents, Sentiments, and Aggregate Demand

Adrien Auclert* Hugo Monnery† Matthew Rognlie‡ Ludwig Straub§

March 2026

Abstract

We integrate dispersed information into a heterogeneous-agent New Keynesian model in which households cannot fully distinguish idiosyncratic from aggregate income, creating a two-way feedback between household sentiment about future income and aggregate demand. We develop a portable methodology—based on a *forecast matrix*—to embed various types of learning into arbitrary sequence-space models. Our model yields three key results: delayed overreaction of beliefs to shocks, a "sentiment inertia" feedback that amplifies hump-shaped responses, and a "sentiment accelerator" that generates persistently slow recoveries. These results depend crucially on the interaction between household heterogeneity and dispersed information.

*Stanford University, CEPR and NBER. Email: aauclert@stanford.edu.
†Harvard University. Email: hmonnery@g.harvard.edu.
‡Northwestern University and NBER. Email: matthew.rognlie@northwestern.edu.
§Harvard University, CEPR and NBER. Email: ludwigstraub@fas.harvard.edu.

1 Introduction

In his General Theory, Keynes (1936) famously observed that much of economic activity is influenced by "sentiments" — waves of optimism and pessimism that sweep through the economy. This idea has been enormously influential in popular discussions of recessions and recoveries, and has inspired a large theoretical literature (Lorenzoni 2009, Angeletos and La'o 2013, Angeletos, Collard and Dellas 2018). At the core of this literature is a link between sentiments, interpreted as optimistic expectations about future income, and aggregate demand. The strength of this link relies crucially on two elements: first, how households form their income expectations, and second, how large marginal propensities to consume out of future income are.

In this paper, we develop a new model that is designed to improve both elements and give a new perspective on how sentiments influence, and are influenced by, business cycles. Our model integrates learning and dispersed information (Lucas 1972, Nimark 2014, Angeletos and Huo 2021) into a modern heterogeneous-agent New-Keynesian model (Werning 2015, Oh and Reis 2012, Kaplan, Moll and Violante 2018, McKay, Nakamura and Steinsson 2016, Auclert, Rognlie and Straub 2024). In doing so, we develop a new and fully portable methodology to integrate learning in any model written in the sequence space (Auclert, Bardóczy, Rognlie and Straub 2021a).

Households in the model cannot fully distinguish idiosyncratic from aggregate income (Adams and Rojas 2024) and learn about future income from past income observations. This creates a tight link between sentiments and aggregate demand: optimistic sentiment about future income arises naturally from the persistent component of the idiosyncratic income process; and it passes through to demand via realistically high intertemporal MPCs (Auclert et al. 2024).

Our model produces four main findings. First, it predicts that beliefs first undershoot, then overshoot, in response to monetary policy shocks. This idea of "delayed overreaction" can be found in theoretical work by Barberis, Shleifer and Vishny (1998) and in empirical work by Bordalo, Gennaioli, Ma and Shleifer (2020b) and Angeletos, Huo and Sastry (2021). We extend previous empirical results in section 2 below.

Second, the model generates a "sentiment inertia" effect. Because households in our model learn from a slow-moving aggregate variable—aggregate income—the consumption response is even more hump-shaped than in regular models with dispersed information about exogenous $AR(1)$ variables (Angeletos and Huo 2021). In our model, dispersed information slows down learning about the monetary shock, which slows down the consumption and hence output responses, which in turn further slows learning from output.

Third, we find that our model of sentiments predicts slow recoveries from recessions. Weak output leads to pessimistic sentiment, which persists somewhat into the future, begetting weak consumption and demand in the future. In analogy to the "investment accelerator" (Samuelson 1939), we term this phenomenon the "sentiment accelerator".

Fourth, we find a significant role for the interaction between heterogeneity and dispersed information in causing these results. Without dispersed information, none of these results hold. Without heterogeneity and realistic intertemporal MPCs, the results either do not hold, or are

significantly weaker.

Facts. We begin our paper with a set of five basic facts that motivate and set the stage for our analysis. We construct a usable "household sentiment" index by extracting the principal component from all 13 questions about future economic activity or economic well-being in the University of Michigan's Survey of Consumers (Curtin 2019). We begin with a basic fact, namely that sentiment is highly procyclical (Fact 1), correlating tightly with GDP growth. This raises the question of (a) whether sentiments affect the real economy and (b) whether the real economy affects sentiments.

On the former, we find evidence suggestive of a transmission from sentiments to the real economy. We show that our broad sentiment index is, in fact, tightly correlated to expectations about future income, both over time, as well as in the cross-section across individual survey respondents (Fact 2). Likewise, sentiments are also tightly correlated with the intention to buy durable goods, again over time as well as in the cross-section (Fact 3). Combining these two facts suggests a connection between optimistic income expectations and consumer spending. We confirm this connection by showing data from Colarieti, Mei and Stantcheva (2024) indicating that households very much respond to higher income anticipated over the course of the next year, and generally have high MPCs out of income (Fact 4).

On the latter, to better understand how sentiments are formed, we consider identified monetary policy shocks a la Romer and Romer (2004). We construct forecast errors for the two questions in the Michigan survey for which this is possible. We show that both suggest initial underreaction of sentiment to economic activity, followed by overreaction (Fact 5), in line with the findings of Angeletos et al. (2021). Eventually, forecast errors go to zero. We confirm our facts where we can (Facts 1–3) with data from the European Union's Business and Consumer Surveys (BCS, European Commission 2024).

Model. We develop a new model that is consistent with all five facts. The model is based on a conventional heterogeneous-agent (HA) model, in which household income is partly driven by idiosyncratic income shocks, and partly driven by aggregate income shocks. Traditional models with full information rational expectations (FIRE) assume households can precisely distinguish idiosyncratic from aggregate income fluctuations. Our model, instead, assumes that some of the idiosyncratic risk cannot be observed independently of aggregate risk. This means that households form expectations about their own future income, as well as aggregate future income, based on income observations that are only noisy indicators of true aggregate income. In line with the literature (Woodford 2003, Angeletos and La'o 2013, Angeletos and Huo 2021), we call this form of expectation formation dispersed information rational expectations (DIRE).

Methodology. Models of learning, including our DIRE model, are typically hard to solve, and especially so in the context of a heterogeneous-agent model. We develop a new methodology

that can easily embed many existing forms of learning as alternatives to FIRE into macroeconomic models. Our method works in the sequence space (Auclert et al. 2021a) and can be readily combined with any model written and solved to first order in the sequence space. Aside from DIRE, our methodology allows for least-squares learning (Evans and Honkapohja 2001, Marcet and Nicolini 2003, Eusepi and Preston 2011, Cogley and Sargent 2005), bounded rationality of the type in Krusell and Smith (1998), learning using local projections, structural reinforcement learning (Yang, Wang, Schaab and Moll 2025), natural expectations (Fuster, Laibson and Mendel 2010), Bayesian learning with and without unknown parameters (Kozicki and Tinsley 2001, Monnery 2025, Farmer, Nakamura and Steinsson 2024), and a form of learning from memory (Bordalo, Gennaioli and Shleifer 2020a) and diagnostic expectations (Bordalo, Gennaioli and Shleifer 2018). Our method for solving DIRE models is quite different from existing methods, which work in the state space or frequency domain, such as Rondina and Walker (2009, 2021), Huo and Takayama (2025).

The key insight behind our method is that all these forms of belief formation can be reduced to a simple idea: Namely that the average expectation about the first-order movement of an aggregate variable, $\overline{\mathbb{E}}_t[Y_{t+h}]$, can be expressed as a linear combination of past realizations of the variable,

$$\overline{\mathbb{E}}_t[Y_{t+h}] = \sum_{s=0}^{\infty} F_{s,h} Y_{t-s} \tag{1}$$

To give two examples, if agents believe Y_t follows an AR(1) with persistence ρ, $F_{0,h}$ would simply equal ρ^h, with $F_{s,h} = 0$ elsewhere. If agents instead see a noisy signal $Y_t + \psi_{it}$, with ψ_{it} orthogonal to Y_t but possibly autocorrelated itself, $F_{s,h}$ is given as the solution to a straightforward linear system of equations

$$\sum_{s=0}^{\infty} F_{s,h} \lambda_{r-s} = \lambda_{r+h} \qquad \forall r, h \geq 0 \tag{2}$$

where $\lambda_m \equiv \text{Cov}(Y_{t+m}, Y_t) + \text{Cov}(\psi_{it+m}, \psi_{it})$. We collect $(F_{s,h})$ in a matrix we call **forecast matrix** $\mathbf{F} \equiv (F_{s,h})$. The approach naturally generalizes to having multiple aggregate variables on the right hand side of (1).

We show how the forecast matrix $\mathbf{F}$ can be used to turn the original FIRE sequence-space Jacobians of our model into the correct Jacobians with learning. Thus, with these modified Jacobians in hand, our model can be solved with standard sequence-space techniques (Auclert et al. 2021a). If $\mathbf{F}$ depends on endogenous variables (e.g. the autocovariances of Y_t), we show how these can be easily computed in an outer fixed point.

Results. Equipped with a standard calibration, we show that our HA model with DIRE is in line with our five motivational facts. In particular, it features a two-way interaction between sentiments and economic activity. Sentiments drive aggregate demand, as optimistic expectations of future income translate to consumer spending due to realistically high MPCs out of future income. At the same time, past aggregate income drives sentiments via (1), first too little (initial

underreaction) then too much (delayed overreaction).

This interaction between sentiments and aggregate demand has both a short- and a long-run effect. In the short run, it leads to very inertial sentiment. Consider, for example, an accommodative monetary policy shock. Since sentiment takes time to build up, aggregate output and income are initially weaker, and this in turn dampens sentiment early on, since it depends on past aggregate income. This interaction amplifies the existing tendency of DIRE to generate hump-shaped responses to monetary shocks (see, e.g., Angeletos and Huo 2021), which emphasize slow learning about whether the shock occurred.

In the long run, the interaction between sentiment and economic activity leads to what we call the "sentiment accelerator". Here, the interaction plays out dynamically. Consider, for example, a negative demand shock causing a recession, associated with poor economic activity and pessimistic sentiment. In our HA model with DIRE, *both* of these variables are persistent on their own: poor economic activity persists because of intertemporal MPCs, as low income today spills over into low spending going forward; pessimistic sentiment is persistent as it crucially depends on past income realizations. The fact that both economic activity and sentiment are persistent and affect one another leads the total persistence to be greater than the sum of its parts. Our model generates significantly slower recoveries than an equivalent representative-agent or two-agent model would with DIRE, or an HA model would with FIRE.

Having spelled out the conceptual insights from our model, we move on to quantifying the two-way interaction between sentiments and economic activity. We do so by conducting a Bayesian estimation of our model with historical time series data of U.S. business cycles. We specifically include our sentiment series from the University of Michigan's Survey of Consumers among the series we use to discipline our model.

Extensions. We consider two extensions to our work. First, we show that DIRE is not the only way to match the facts and generate the conceptual results mentioned above. This can also be done with dispersed information coupled with diagnostic expectations (Bordalo et al. 2020b, Bardóczy and Guerreiro 2024). Second, building on the idea of the forecast matrix in (1), we develop a general methodology that can accommodate *all* possible information frictions to first order, allowing straightforward efficient solution in the sequence space.

Related Literature. Our paper is most related to three key strands of the literature on deviations from FIRE models. First, it is related to the methodology of simulating models with dispersed information, and more generally, with different kinds of learning. The closest work here is a small set of papers proposing alternative ways of solving models with deviations from FIRE in the sequence space. The first such paper is Auclert, Rognlie and Straub (2020), which considers a specific set of deviations from FIRE, such as sticky expectations (Mankiw and Reis 2002, Carroll 2003), and cognitive discounting (Gabaix 2020). The technology used in that paper cannot be used to solve DIRE models and is generally quite distinct from the one here. In Appendix E we show

how both approaches can be nested in a single, most general, approach. This general approach is similar to that in Bardóczy and Guerreiro (2024), which Lenney and Rosso (2025) builds on. Neither paper, however, studies models with dispersed information about endogenous variables or learning.[1]

In solving DIRE models, our paper is related to prior methods that do not work in the sequence space, including Rondina and Walker (2009, 2021), Nimark (2014), and Huo and Takayama (2025). This literature has found various relatively complex ways of exactly or approximately representing the hierarchy of beliefs implied by DIRE models in the state space. Our approach circumvents much of this complexity in the sequence space, making it just as efficient to solve models with DIRE as with FIRE, even with large-scale heterogeneity, without need for a Kalman filter or any explicit description of higher-order beliefs.

The second strand of literature our paper is most related to is the one on sentiments. Lorenzoni (2009) models exogenous sentiment as a driver of aggregate demand in a representative-agent economy. Sentiments have an effect in his model as they shift *permanent* income expectations. Our model focuses on persistent, but not permanent, income expectations, which matter precisely in the presence of high MPCs. Angeletos and La'o (2013) study the effects of exogenous sentiment shocks, which are best understood as jointly driving aggregate demand and aggregate supply in their model. Angeletos et al. (2018) quantify the effects of exogenous sentiment shocks in a representative-agent DSGE model.

More generally, DIRE models have also been studied to generate myopia (Angeletos and Lian 2022) and inertia (Angeletos and Huo 2021). Angeletos et al. (2021) show how delayed overreaction can be generated by combining dispersed information from iid noise with natural expectations (Fuster et al. 2010). Our paper shows that delayed overreaction emerges naturally from assuming heterogeneous households cannot distinguish a fraction of their idiosyncratic income from aggregate income.

The third strand our paper is closely related to is the literature on deviations from FIRE in HA models. The first such model is the one by Farhi and Werning (2019), which considers level-k thinking and forward guidance. Following this paper, there is Auclert et al. (2020), Pfäuti and Seyrich (2024), Guerreiro (2024), Bardóczy and Guerreiro (2024), and Yang et al. (2025), among many others. Adams and Rojas (2024) is notable here, as it shares with our work the assumption that households cannot cleanly distinguish idiosyncratic from aggregate income. However, unlike our paper, Adams and Rojas (2024) focuses on partial equilibrium consumption volatility.

2 Five facts about household sentiment

In this section, we make use of two surveys: the University of Michigan's Survey of Consumers, and the household component of the European Union's Business and Consumer Surveys (BCS). We define sentiment as the first principal component of multiple forward-looking survey ques-

[1] See Overhage (2026) for a different approach to implement dispersed information in the sequence space.

tions put to households. We focus on questions about economic activity—in particular, we exclude questions about prices. The exact survey questions used in the construction of our sentiment measure are listed in appendix A.1.

To construct sentiment, we take the N question series, $i = 1, \ldots, N$, and standardize each

$$\tilde{Q}_{i,t} = \frac{Q_{i,t} - \mu_i}{\sigma_i}$$

where $\mu_i = \frac{1}{T}\sum_{t=1}^{T} Q_{i,t}$ and $\sigma_i = \sqrt{\frac{1}{T}\sum_{t=1}^{T}\left(Q_{i,t} - \frac{1}{T}\sum_{t=1}^{T} Q_{i,t}\right)^2}$. Where data are missing, we first impute the median for that survey series. We run a principal components analysis on all the questions included, and define sentiment as the first principal component,

$$\mathcal{S}_t = \sum_{i=1}^{N} w_i^{(1)} \tilde{Q}_{i,t}$$

where the weights $w_i^{(1)}$ are the elements of the eigenvector corresponding to the largest eigenvalue of the covariance matrix of the standardized series. The sign of sentiment is normalized to have a positive weight on expected personal financial situation. The resulting weights in each case are listed in appendix A.1.

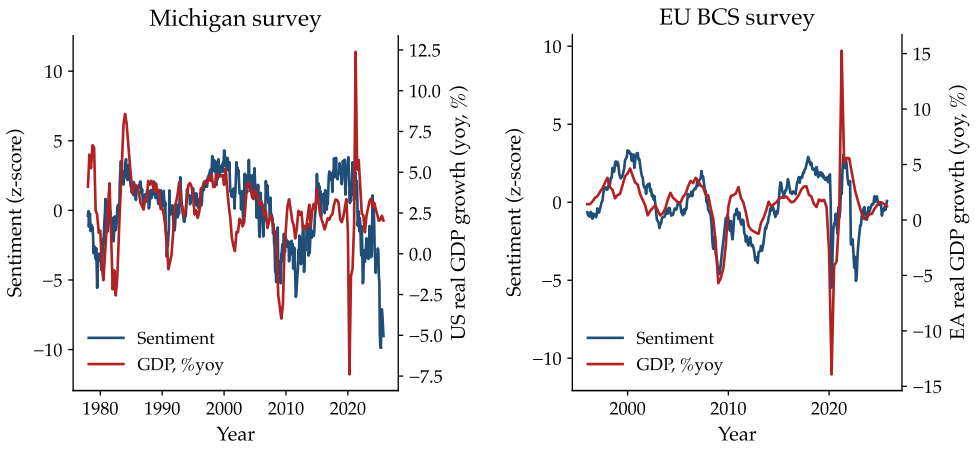
2.1 Household sentiment is very cyclical

In figure 1, we compare our sentiment measure across these two surveys to real GDP growth (in the US and Euro Area, respectively).[2]

As the left panel of figure 1 shows, the correlation between our sentiment index and GDP growth is remarkably tight in the United States, both at business-cycle and higher frequencies, consistent with Bhandari, Borovicka and Ho (2019). The right panel confirms that this pattern is not specific to the US: Euro Area sentiment moves closely with Euro Area GDP growth. Sentiment drops sharply during recessions and recovers during expansions, suggesting a strong link between household perceptions and economic activity.

[2]To match the monthly timing of the consumer surveys, we first interpolate GDP to monthly frequency.

Figure 1: Sentiment is procyclical



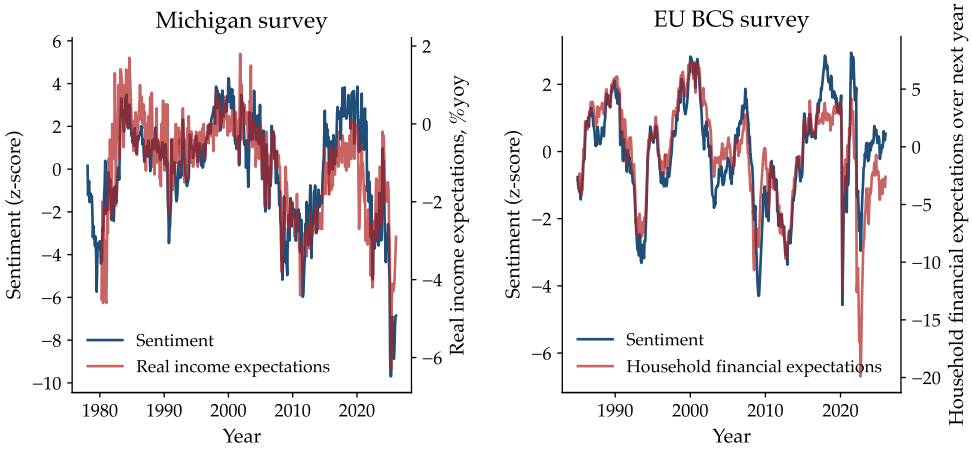
Note. Left panel: normalized Michigan sentiment index and year-on-year real GDP growth. Right panel: normalized EU BCS sentiment index and year-on-year real GDP growth.

2.2 Sentiment is highly correlated with income expectations

Our measure of sentiment speaks to households' overall perception of economic activity going forward. In the left plot of figure 2, we use quantitative responses to the Michigan survey of households' own expected nominal income growth over the next year, and expected inflation over the next year, to create a quantitative real income growth forecast. We see that our sentiment measure is very strongly correlated with these income expectations. To make the comparison more meaningful, for this figure, we computed sentiment without using the nominal income growth responses (and in any case we did not include the inflation forecast in our sentiment measure). Recall that the variables used in constructing sentiment are standardized, but this plot suggests a mapping of around two sentiment units to one percent real income growth expectation.

As the EU BCS does not have a quantitative income growth question, on the right plot of figure 2, we compare sentiment to the question "How do you expect the financial position of your household to change over the next 12 months?" Again, for this figure, we exclude that question from construction of the sentiment measure. We again see a strong co-movement.

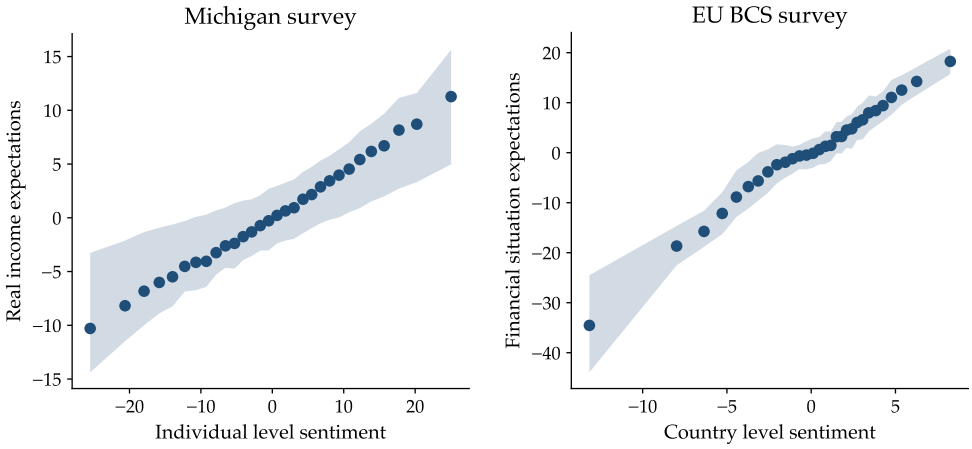
Figure 2: Sentiment and income expectations co-move



Note. Left panel: normalized Michigan sentiment index (excl. the real income expectations question) and real income expectations from Michigan survey. Right panel: normalized EU BCS sentiment index (excl. the household financial expectations question) and household financial expectations from EU BCS survey.

We also show this pattern in the cross section, in figure 3, with the shaded band representing the interquartile range. This is constructed using the same principal component weights as above, applied at the individual level, but where the sentiment measure does not include the income expectation question. We include fixed effects, parsing out aggregate comovement of sentiment and income expectations. In the case of the Michigan survey, following Bachmann, Berg and Sims (2015) we drop observations where the absolute one year ahead inflation expectation exceeds 20%. We also drop observations with absolute income change expectations exceeding 50%.

Figure 3: Sentiment and income expectations co-move in the cross section

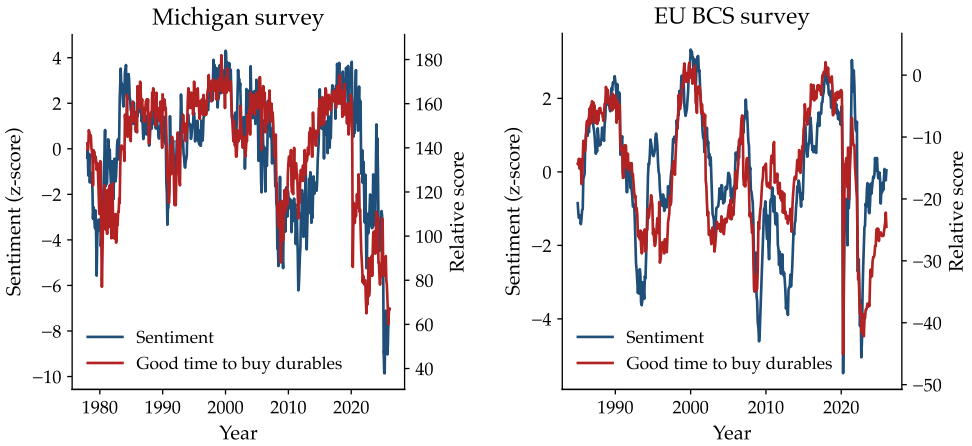


Note. Figure uses the principal component weights from the sentiment index to construct individual-level sentiment (x axis) (excl. income / financial situation expectations) and plots it against Michigan real income expectations (left panel) and EU BCS financial situation expectations (right panel). Binned scatter plot.

2.3 Sentiment correlated with good time to make purchases

Sentiment matters because it drives economic decisions. In the next subsection, we will show evidence for high MPC agents, whose consumption decisions ought to depend on their near term income forecasts. Here, we show in the survey data that our measure of sentiment co-moves with household perception that it is a good time to spend. In the left hand plot, we compare sentiment against the response to the question, in the Michigan survey, "... do you think now is a good or bad time for people to buy major household items?". In the right hand plot, we compare Euro Area sentiment against the question "... do you think that now it is the right moment for people to make major purchases such as furniture, electrical/electronic devices, etc.?" Neither question was itself used in constructing the sentiment measure.

Figure 4: Sentiment and purchasing intention



Note. Left panel: normalized Michigan sentiment index and intention to buy durable goods from Michigan survey. Right panel: normalized EU BCS sentiment index and intention to buy durable goods from EU BCS survey.

Next, we show this relationship holds in the cross-section. We construct a respondent-level sentiment measure in the Michigan survey, consistently applying the principal component loadings estimated through the macro series. To do so, for each individual j, we first normalize their responses with the same mean and standard deviation used to normalize the aggregated survey question

$$\tilde{Q}_{j,i,t} = \frac{Q_{j,i,t} - \mu_i}{\sigma_i}$$

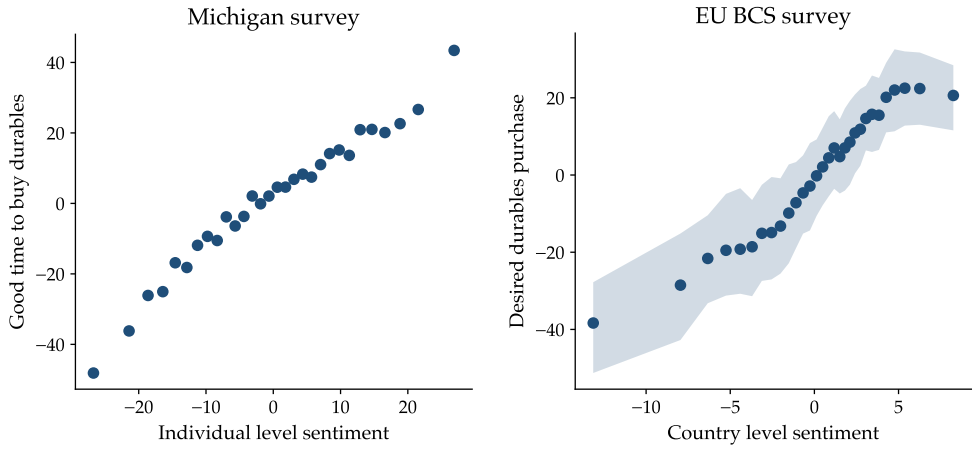
And then construct individual sentiment as

$$\mathcal{S}_{j,t} = \sum_{i=1}^{N} w_i^{(1)} \tilde{Q}_{j,i,t}$$

We show a binned scatter plot of the durable good survey question, $Q_{j,\text{Durables},t}$ against sentiment, $\mathcal{S}_{j,t}$, controlling for a time fixed effect, on the left of Figure 5. We proceed analogously for the BCS. We construct a country-level sentiment measure using the same loadings estimated for the Euro

Area above, and applying them to the country level survey series. We show a binned scatter plot of the durable goods series against this measure of sentiment, controlling for a time fixed effect. In both cases, there is a clear positive relationship.

Figure 5: Sentiment and purchasing intention in the cross-section

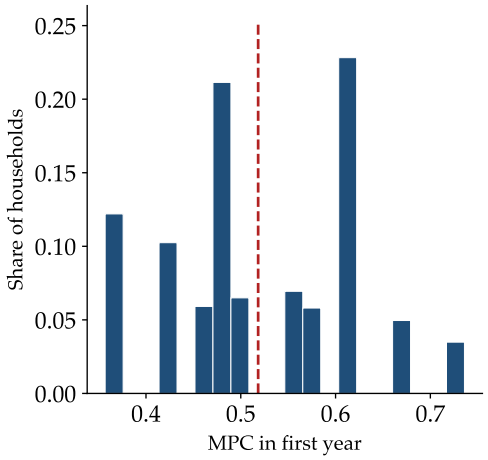


Note. Figure uses the principal component weights from the sentiment index to construct individual-level sentiment (x axis) and plots it against intention to buy durables in Michigan survey (left panel) and EU BCS (right panel). Binned scatter plot.

2.4 High MPCs

Two-agent models feature a share of hand-to-mouth agents, whose marginal propensity to consume is 100%, and a share of permanent-income agents, whose consumption responds to transitory income shocks at only approximately the rate of interest. Richer heterogeneous agent models, as we will consider in this paper, feature intermediate MPC agents. These are essential to the central role of sentiment in business cycles, as these agents uniquely respond to short- and medium-horizon income forecasts. In Figure 6, we plot the empirical distribution, across liquidity cells, of MPCs in the first year after a lottery win found in Fagereng, Holm and Natvik (2021).

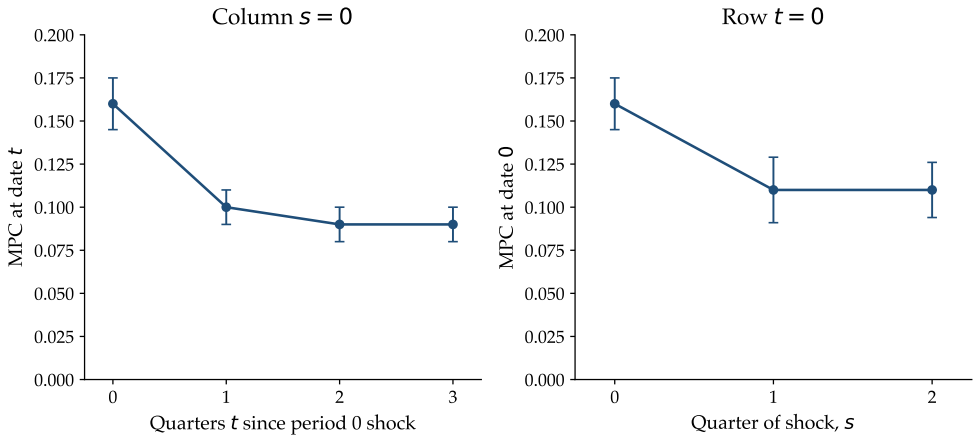
Figure 6: Distribution of MPCs



Note. Distribution of estimated MPCs from spending out of unexpected lottery earnings in Fagereng et al. (2021).

In figure 7, we plot two dimensions of the intertemporal MPCs from Colarieti et al. (2024). On the left, we show how the consumption response to a period 0 income shock evolves over time. There is a decline after the first period, but MPCs remain elevated. On the right, we show the period 0 response to news about income zero, one, or two quarters out. In Figure 9, we compare the iMPCs from a heterogeneous agent (HA) model to those from representative agent and two-agent models, with the HA iMPCS matching these patterns much more closely.

Figure 7: iMPCs



Note. Left panel: dynamic spending response to a one-time unanticipated income increase at date 0. Right panel: Date-0 spending response to an expected one-time income shock to arrive at date s. Both from survey data collected by Colarieti et al. (2024), see their figure 3A.

2.5 Delayed overreaction in household expectations

We next ask how sentiment evolves in response to a monetary policy shock. Since expectations should rationally respond to such shocks, we focus on forecast errors to understand whether sen-

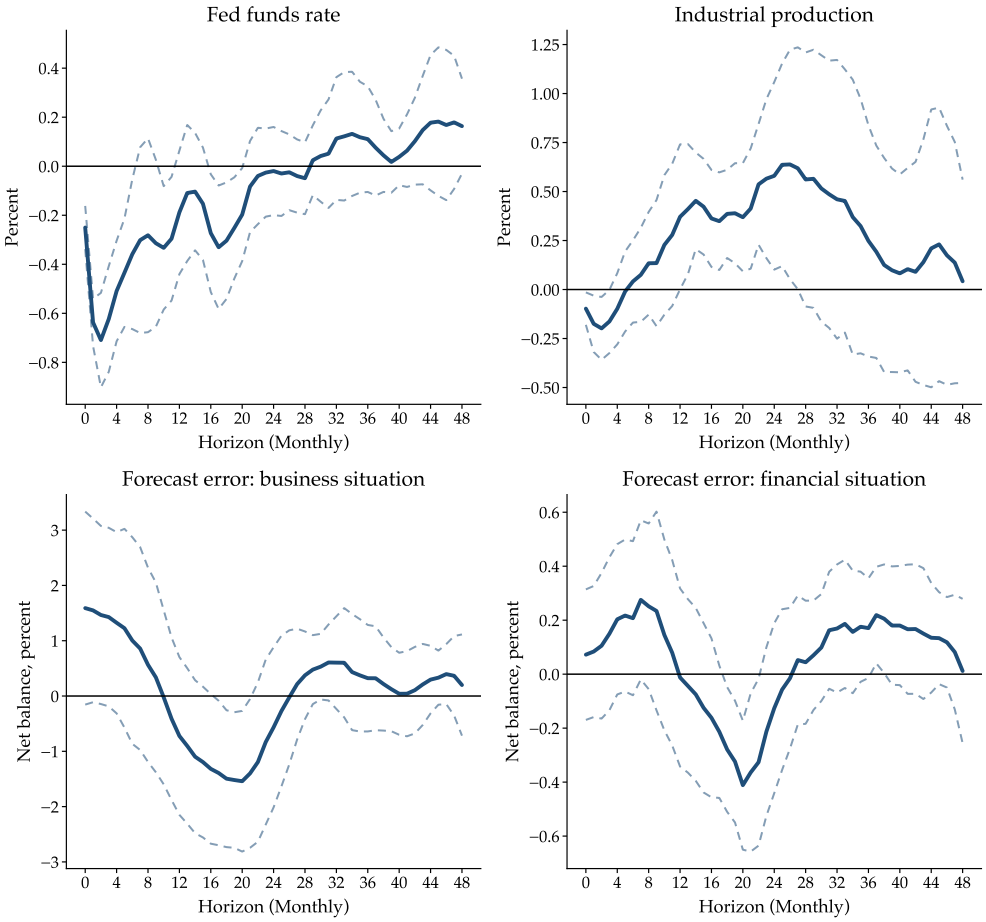
timent moves too much or too little. We use the Romer-Romer shocks, and, since we need a sufficiently long time series, we focus on the Michigan survey. Many questions in the Michigan survey ask only about the future, or only about the past; to construct forecast error-like objects, we restrict ourselves to questions with both a future and past dimension. This leaves us with two pairs of questions. Households are asked to forecast their personal financial situation over the next year, and to report it over the past year. They are also asked analogous questions but about overall business conditions. We define the forecast error here as the difference between the report made one year in the future about the change over the past year, and the forecast about the forthcoming year made today. Thus, a positive forecast error at time t suggests that, on average, households at t were too pessimistic about the change between t and $t+12$, when compared to what they report for that same period in $t+12$. As the data, particularly on personal financial situation, is noisy, we first apply a bandpass filter. We retain standard business cycle frequency fluctuations, between 6 and 32 quarters. We set a truncation length of 12.

We run local projections at monthly frequency: for each outcome, Y, at horizon h, we run

$$Y_{t+h} = \beta_h^Y \varepsilon_t^{\text{RR}} + \Gamma_h^Y X_t + u_{t,h}^Y \tag{3}$$

$\varepsilon_t^{\text{RR}}$ is the Romer-Romer series. The controls X_t include one and two-period lagged industrial production, unemployment, CPI, a commodity price index, the fed funds rate, and the Romer-Romer series, as in Ramey (2016). Finally, $u_{t,h}^Y$ is a regression error. We plot the estimated coefficients $\{\beta_h^Y\}$ in Figure 8. We compute Newey and West (1987) HAC-robust standard errors, and plot 95% confidence intervals. An accommodative monetary policy boosts industrial production. We find that in response, households first under-react and then over-react to the shock. This aligns with the finding in Angeletos et al. (2021) of delayed overshooting in the Survey of Professional Forecasters, but here in the case of households.

Figure 8: Delayed over-reaction to a monetary policy shock



Note. Panels show impulse responses from local projections onto Romer-Romer shocks, estimated as in (3). Top left: Effective federal funds rate. Top right: Industrial production. Bottom row: Forecast errors (1y-backcast 12 months from now minus current 1y-forecast) for question on business situation and financial situation in Michigan survey.

3 Heterogeneous agents with learning about income

We base our analysis on a standard heterogeneous-agent model, which we introduce in this section. We show how different types of learning can be integrated into the model and how they modify the aggregate consumption response. Throughout this section, we focus on a partial-equilibrium setup in which households face a constant interest rate and an exogenous shock to aggregate income.

3.1 Model setup

There is a continuum of households, labeled by $i \in [0, 1]$. Each household i has a strictly concave flow utility function $u(c_{it})$ over its consumption c_{it} whose expected present discounted value it

maximizes at each point in time τ

$$\max \mathbb{E}_{i\tau}\left[\sum_{t=\tau}^{\infty}\beta^{t-\tau}u(c_{it})\right]$$

We allow for an arbitrary process of expectations operators $\mathbb{E}_{i\tau}$ here to nest the various different learning models that we analyze below. Household i faces a constant interest rate r, a standard budget constraint

$$c_{it} + a_{it} = (1+r)\,a_{it-1} + y_{it} \tag{4}$$

as well as a borrowing constraint, requiring that its asset position a_{it} at the end of period t be non-negative,

$$a_{it} \geq 0 \tag{5}$$

Household i's income at date t, y_{it}, is the product of three distinct components:

$$y_{it} = e_{it}\psi_{it}Y_t$$

The third component Y_t is aggregate income in the economy. We assume $Y_t = Y = 1$ in the steady state. We assume that first-order deviations from the steady state dY_t follow some $MA(\infty)$ process,

$$dY_t = \sum_{s=0}^{\infty}\hat{Y}_s\epsilon_{t-s}$$

with MA coefficients $\hat{Y}_s$ and Gaussian innovations $\epsilon_t \sim \mathcal{N}(0,\sigma^2)$. The other two components capture idiosyncratic income shocks. We assume that e_{it} is a Markov chain with mean 1, $\mathbb{E}e_{it} = 1$. ψ_{it} follows an arbitrary stationary process. Both e_{it} and ψ_{it} are iid across households, and are independent of one another. The main distinction we draw at this point between risk emanating from e_{it} compared to that from ψ_{it} is that we assume the latter scales in the amount of aggregate risk σ, that is, $\log\psi_{it} = \sigma\log\tilde{\psi}_{it}$ where $\tilde{\psi}_{it}$ is independent of σ.

The reason we allow for two sources of idiosyncratic shocks is that it gives us the ability to nest different informational assumptions on the two later on. Specifically, in some of the models we consider below, e_{it} will capture idiosyncratic risk known to households, while movements in ψ_{it} will capture firm-level or sector-specific shocks that may be hard for households to distinguish from Y_t.

3.2 Solution under FIRE

Next, we analyze the behavior of aggregate consumption under different assumptions on the expectations operators $\mathbb{E}_{i\tau}$. First, we assume each household has full information about past and current realizations of $\{e_{it}, \psi_{it}, \epsilon_t\}$ and forms rational expectations. Henceforth, we refer to these as "FIRE" expectations.

Proposition 1. *Under full information and rational expectations, aggregate consumption follows, to first order in aggregate risk σ,*

$$dC_t = \sum_{s=0}^{\infty} \hat{C}_s \epsilon_{t-s}$$

where the impulse response $\hat{C}_s$ is determined as

$$\underbrace{\begin{pmatrix} \hat{C}_0 \\ \hat{C}_1 \\ \hat{C}_2 \\ \vdots \end{pmatrix}}_{\equiv \hat{\mathbf{C}}} = \mathbf{M} \cdot \underbrace{\begin{pmatrix} \hat{Y}_0 \\ \hat{Y}_1 \\ \hat{Y}_2 \\ \vdots \end{pmatrix}}_{\equiv \hat{\mathbf{Y}}}$$

where $\mathbf{M}$ is the intertemporal MPC (iMPC) matrix of the heterogeneous-agent model.

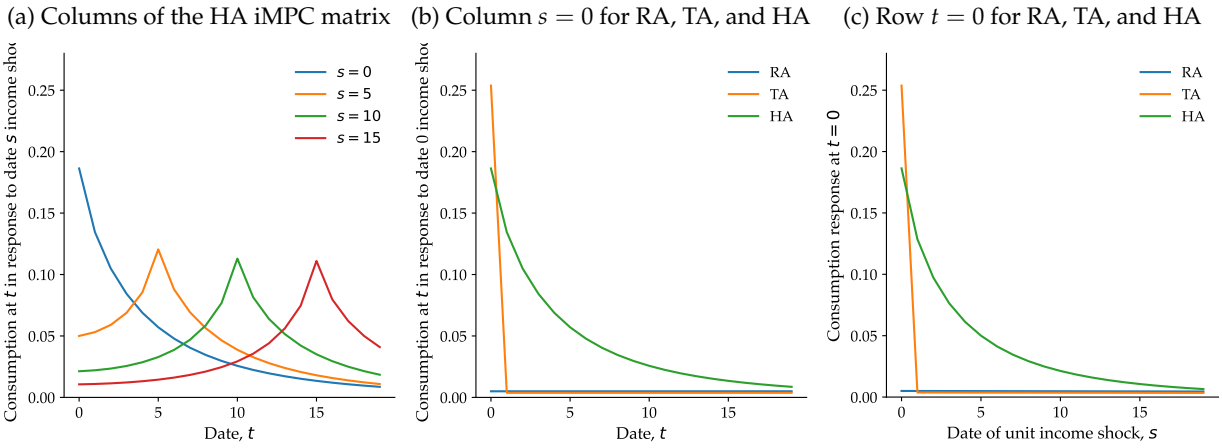
Proposition 1 shows that the first-order impulse response of consumption to the aggregate income shock is a simple linear function of the MA coefficients of income, $\hat{\mathbf{Y}} = (\hat{Y}_s)$. The infinite matrix mapping income to consumption is the iMPC matrix $\mathbf{M}$ (Auclert et al. 2024), whose (t, s) entry is defined as the derivative of the perfect-foresight consumption response t periods after learning of an anticipated increase in income in s periods. Mathematically,

$$M_{t,s} \equiv \frac{\partial C_t}{\partial \epsilon_0}\Big|_{\sigma=0, \hat{Y}_r = \delta_{rs}} \tag{6}$$

That is, for a given s, $M_{t,s}$ is the dynamic consumption response to a one-time perfect-foresight shock of greater income in period s, without any future aggregate risk, or risk in ψ_{it}, $\sigma = 0$. Intertemporal MPCs can be estimated in the data (e.g. Fagereng et al. 2021; Colarieti et al. 2024), and are informative about fiscal policy (Auclert et al. 2024, Angeletos, Lian and Wolf 2024, Sundram 2025), monetary policy (Auclert et al. 2020), exchange rate shocks (Auclert, Rognlie, Souchier and Straub 2021b), and energy shocks (Auclert, Monnery, Rognlie and Straub 2023), among others.

Panel (a) of figure 9 plots a few columns $M_{t,s}$ as function of t, for $s \in \{0, 5, 10, 15\}$. The plot uses the calibration we introduce in section 4.2 below. Consumption responses in heterogeneous-agent models are generally largest in the period of income receipt, that is, when $t = s$. Panel (b) of figure 9 compares the first column $M_{t,0}$ in our heterogeneous-agent (HA) economy to the first columns of analogously computed iMPC matrices of a representative-agent (RA) and a two-agent economy (TA). The RA model predicts an unrealistically small and flat consumption response. The TA model predicts a high spending response in period 0, but then falls even below the response of the RA model. Panel (c) of figure 9 compares the first *row* $M_{0,s}$ between the RA, TA, and HA models. The first row captures how households respond to anticipated future income changes. A mirror image to the first columns emerges: in the HA model, households respond significantly to anticipated income increases over the next few quarters, while in the RA and TA models, households do not respond very much to future income over the next few quarters. This pattern of responses will

Figure 9: Intertemporal MPCs



be crucial to understanding the differences in how optimism and pessimism about future income map into consumption in the three different models.

3.3 Introducing learning

Full-information rational expectations imply that the average income expectation in the economy is determined by past innovations to aggregate income, that is,

$$d \int \mathbb{E}_{it}[y_{it+h}]\, di = \sum_{s=0}^{\infty} \hat{Y}_{s+h} \epsilon_{t-s} \tag{7}$$

At each date t, the average household expects their income to move in line with the anticipated effects, $\hat{Y}_{s+h}$, of past income innovations, ϵ_{t-s}, for all s. This is a fairly complicated problem to solve: It requires keeping track of past innovations, as well as full knowledge of the anticipated effects of the income innovations on income.

We now consider an alternative approach based on households learning from past *income*, instead of income innovations, about future income. We show below that this approach nests many common formulations of learning. Specifically, we assume that there exists an infinite matrix $\mathbf{F} \equiv (F_{s,h})$ such that the average income expectation is given by

$$d \int \mathbb{E}_{it}[y_{it+h}]\, di = \sum_{s=0}^{\infty} F_{s,h} dY_{t-s} \tag{8}$$

The main difference in (8) relative to (7) is that the average income expectation here is an explicit function of past income dY_{t-s}, not past income innovations, ϵ_{t-s}. The way past income influences the average income expectation depends on the matrix $\mathbf{F}$, which we call the *forecast matrix*: its entry $F_{s,h}$ exactly captures how an increase in income that happened s periods in the past influences the belief about income in h periods.

We do not imagine **F** to be a primitive of the model. Instead, as we illustrate below, several common ways of modeling learning imply an equation like (8) for *some* forecast matrix **F**. The process of arriving at a forecast matrix may involve least-squares regressions on past data, extrapolation from recent news, or inference from noisy signals. Once the matrix is set, however, it defines a clear mapping from past Y's to average expectations about future incomes.

Strictly speaking, FIRE expectations can also be written in the form (8), provided that $\hat{Y}_0 \neq 0$, by defining $F_{s,h}$ recursively in s

$$F_{s,h} = \frac{\hat{Y}_{s+h}}{\hat{Y}_0} - \sum_{j=0}^{s-1} F_{j,h} \frac{\hat{Y}_{s-j}}{\hat{Y}_0}$$

Many of the examples we give below, however, do not allow agents to learn sufficiently well to achieve FIRE expectations. We expand this approach in section 5 to the case of income following a $VMA(\infty)$ process, with many aggregate shocks and a multi-dimensional innovation series $\{\epsilon_t\}$. In that case, FIRE expectations are generically no longer nested in (8).

3.4 Examples

Many learning models admit a representation as in (8), with a forecast matrix **F**. Among them are: least-squares learning (Evans and Honkapohja 2001, Krusell and Smith 1998), local projection learning, learning from dispersed information (Angeletos and La'o 2013, Angeletos and Huo 2021), Bayesian learning (Kozicki and Tinsley 2001, Farmer et al. 2024, Monnery 2025), cognitive discounting (Gabaix 2020), learning from memory of past shocks (Bordalo et al. 2020a), and dispersed information with diagnostic expectations (Bordalo et al. 2020b, Bardóczy and Guerreiro 2024).

We focus on the first three learning models here and show how each pins down a specific forecast matrix **F**. We discuss the remaining learning models in appendix B.1. Our approach does not nest sticky information, or learning directly about innovations ϵ_t. For those, see Auclert et al. (2020). Appendix E presents a unified framework that can handle arbitrary first-order deviations from FIRE, nesting the approaches we focus on in this paper as well as those in Auclert et al. (2020).

3.4.1 Least-squares learning (LSL)

Our first example is *least-squares learning* (LSL). With this type of learning, each household fully observes aggregate income Y_t, and models it as an $AR(p)$ process,

$$dY_t = \sum_{j=1}^{p} a_j dY_{t-j} + u_t$$

with estimated autoregressive coefficients $a_1, \ldots, a_p$. If dY_t is stationary, the roots of the lag polynomial $a(L) \equiv 1 - \sum_{j=1}^{p} a_j L^j$ will all lie outside the unit circle. Thus, there exists an MA represen-

tation, of the estimated $AR(p)$ given by

$$dY_t = a(L)^{-1}u_t = b(L)u_t = \sum_{j=0}^{\infty} b_j u_{t-j}$$

for some coefficients b_j. Then, we can write

$$\mathbb{E}_t\left[dY_{t+h}\right] = \sum_{s=0}^{\infty}\left\{b_{h+s} - \sum_{k=1}^{\min\{p,s\}} b_{h+s-k}a_k\right\} dY_{t-s}$$

Since the idiosyncratic component of income averages to one, $\mathbb{E}_{it}\left[e_{it}\psi_{it}\right] = 1$, we have that

$$d\int \mathbb{E}_{it}\left[y_{it+h}\right] di = \sum_{s=0}^{\infty} F_{s,h}^{LSL} dY_{t-s}$$

with forecast matrix

$$F_{s,h}^{LSL} = b_{h+s} - \sum_{k=1}^{\min\{p,s\}} b_{h+s-k}a_k$$

Thus, the estimated set of coefficients $a_0, \ldots, a_p$ determines b_j and thus the entire forecast matrix $\mathbf{F}^{LSL}$.

Least-squares learning is most commonly used with an $AR(1)$, that is, $p = 1$. In that case, $a_1 = \rho \in (0,1)$, the persistence parameter, and $b_j = \rho^j$. Thus,

$$F_{s,h}^{AR1} = \rho^{h+s} - 1_{\{s>0\}}\rho^{h+s} = \begin{cases} \rho^h & s = 0 \\ 0 & s > 0 \end{cases}$$

Written in matrix form, the $AR(1)$ forecast matrix is

$$\mathbf{F} = \begin{pmatrix} 1 & \rho & \rho^2 & \rho^3 & \cdots \\ 0 & 0 & 0 & 0 & \cdots \\ 0 & 0 & 0 & 0 & \cdots \\ 0 & 0 & 0 & 0 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}$$

or in other words,

$$d\int \mathbb{E}_{it}\left[y_{it+h}\right] di = \rho^h dY_t$$

Agents simply take the most recent aggregate income observation and assume it will persist with some persistence parameter ρ. We think of the persistence parameter ρ as generally being estimated on a large historical sample of income observations here. This approach generalizes to $VAR(p)$'s.

This type of learning has been studied extensively (see, e.g., Evans and Honkapohja 2001,

for a textbook treatment). The special case of $AR(1)$ learning was popularized by Krusell and Smith (1998) to approximate equilibria in heterogeneous-agent models. We apply this methodology specifically to the Krusell and Smith (1998) model in appendix B.5.

3.4.2 Local projection learning (LPL)

A natural alternative to modeling dY_t as an $AR(p)$ is to run local projections of future aggregate income dY_{t+h} on current and past income, $dY_t, \ldots, dY_{t-m+1}$. We label this "local projection learning" (LPL). The coefficient on lag $s \in \{0, \ldots, m-1\}$ of income is exactly equal to $F^{LPL}_{s,h}$, with entries for $s \geq m$ equal to zero.

The simplest case is that of a single regressor, $m = 1$. In that case, the first row is identical to the autocorrelation function of dY_t, $\lambda_h \equiv Cov(dY_{t+h}, dY_t)/Var(dY_t)$,

$$\mathbf{F}^{LPL} = \begin{pmatrix} 1 & \lambda_1 & \lambda_2 & \lambda_3 & \cdots \\ 0 & 0 & 0 & 0 & \cdots \\ 0 & 0 & 0 & 0 & \cdots \\ 0 & 0 & 0 & 0 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}$$

As we show in appendix B.2, local projection learning is first-order identical to *structural reinforcement learning*, as defined in Yang et al. (2025). In particular, when agents use an $AR(1)$ specification under LSL, the resulting forecast matrix is identical to the one from LPL with $m = 1$: that is, $AR(1)$ least-squares learning is equivalent to structural reinforcement learning to first order.

3.4.3 Dispersed information rational expectations (DIRE)

The third example we consider here is dispersed-information rational expectations (DIRE). In this model, households do not observe Y_t directly in period t. Instead, they observe $z_{it} \equiv \psi_{it} Y_t$, that is, a noisy version of aggregate income. This captures the idea that a household's labor income can change due to a variety of factors, including sectoral, firm-level, or foreign business conditions, which are not always distinguishable for households from changes in aggregate income Y_t.[3] Dispersed information models of this type have been widely studied (e.g., Lorenzoni 2009, Angeletos and La'o 2013, Angeletos and Huo 2021).

To obtain how households update their beliefs in response to signals, agents can again employ a local projection, estimating, for each h,

$$dy_{it+h} = \sum_{s=0}^{\infty} F_{s,h} dz_{it-s} + u_{it+h}$$

[3]DIRE models are often written in terms of noisy signals, where a household earns Y_t but only "sees" a noisy signal $\log Y_t + \log \psi_{it}$. This works well in RA economies, but in HA models, it is critical that households know their current income and asset position to avoid violating the borrowing constraint. This is why, aside from being realistic, a "noisy income" setup in which ψ_{it} influences actual income is methodologically preferable.

that is, they project their own future income on past signals received. This approach is equivalent to the Kalman filter, but convenient for our purposes, since the coefficients are directly determining the forecast matrix $\mathbf{F}$. The standard moment condition associated with this local projection gives us

$$\mathbb{E}\left[\left(dz_{it+h} - \sum_{s=0}^{\infty} F_{s,h} dz_{it-s}\right) dz_{it-r}\right] = 0$$

where we replaced dy_{it+h} by dz_{it+h}, ignoring the idiosyncratic component as it averages to zero. Rearranging the moment condition, we find a simple linear system of equations,

$$\sum_{s=0}^{\infty} F_{s,h} \lambda_{r-s} = \lambda_{r+h} \qquad \forall r, h \geq 0 \tag{9}$$

where $\lambda_m \equiv Cov\left(dY_{t+m}, dY_t\right) + Cov\left(d\psi_{it+m}, d\psi_{it}\right)$ is the sum of the autocovariances of aggregate income Y_t and sectoral (or firm-level) income ψ_{it}. We show in appendix B.3 how this approach can be generalized to handle several signals and multiple shocks.

To give an example, consider an AR(1) aggregate shock, $dY_t = \rho dY_{t-1} + \epsilon_t$ with cross-sectional variance $\frac{\sigma^2}{1-\rho^2}$, and iid sectoral shocks $d\psi_{it} \sim \mathcal{N}(0, \frac{\sigma^2}{1-\rho^2} \cdot \sigma_\psi^2)$ where $\sigma_\psi > 0$ is the standard deviation relative to that of the aggregate innovation dY_t. Then,

$$\lambda_m = \frac{\sigma^2}{1-\rho^2}\left(\rho^m + \sigma_\psi^2 1_{\{m=0\}}\right)$$

Substituting λ_m into (9), we find

$$\sum_{s=0}^{\infty} F_{s,h}\left(\rho^{|r-s|} + \sigma_\psi^2 1_{\{r=s\}}\right) = \rho^{r+h} + \sigma_\psi^2 1_{\{r+h=0\}}$$

For $h = 0$, it is easy to see that this equation is solved by $F_{s,h} = 1_{\{s=0\}}$. For $h > 0$, we can guess and verify that

$$F_{s,h} = \rho^h K\left(\rho\left(1-K\right)\right)^s$$

where $K \in (0,1)$, the Kalman gain, solves

$$\sigma_\psi^2 \frac{\rho^2}{1-\rho^2} K^2 + \left(1 + \sigma_\psi^2\right) K - 1 = 0 \tag{10}$$

In matrix notation,

$$\mathbf{F}^{DIRE} = \begin{pmatrix} 1 & \rho K & \rho^2 K & \rho^3 K & \cdots \\ 0 & \rho K \rho\left(1-K\right) & \rho^2 K \rho\left(1-K\right) & \rho^3 K \rho\left(1-K\right) & \cdots \\ 0 & \rho K\left(\rho\left(1-K\right)\right)^2 & \rho^2 K\left(\rho\left(1-K\right)\right)^2 & \rho^3 K\left(\rho\left(1-K\right)\right)^2 & \cdots \\ 0 & \rho K\left(\rho\left(1-K\right)\right)^3 & \rho^2 K\left(\rho\left(1-K\right)\right)^3 & \rho^3 K\left(\rho\left(1-K\right)\right)^3 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}$$

The intuition for $\mathbf{F}^{DIRE}$ is straightforward: As a shock to aggregate income hits, expectations about future income increase, but not to their FIRE level of ρ, the true persistence, but instead to ρK, where K is a measure of the informativeness of observed income $\psi_{it} Y_t$ for aggregate income. If there was no sectoral income component, $\sigma_\psi^2 = 0$, the Kalman gain solving (10) comes out at $K = 1$, and the economy is back at FIRE beliefs. With $\sigma_\psi^2 > 0$, $K < 1$, and agents first need to learn of the presence of the shock. If there is no further shock (rows $s > 0$ of $\mathbf{F}^{DIRE}$), forecasts fall to zero with powers of $\rho(1 - K)$. Having iid sectoral shocks is one way to slow down the adjustment to a shock.

3.5 Solution under learning

These three examples show how several distinct learning models give rise to forecast matrices $\mathbf{F}$ of different shapes. Our next proposition shows that $\mathbf{F}$ is, in fact, a sufficient statistic for how learning affects aggregate behavior.

Proposition 2. *With learning as in* (8)*, consumption responds as*

$$\hat{\mathbf{C}} = \tilde{\mathbf{M}} \cdot \hat{\mathbf{Y}} \tag{11}$$

where $\tilde{\mathbf{M}}$ *is the lower triangular Toeplitz matrix*[4]

$$\tilde{M}_{t+j,j} = \sum_{k=0}^{t} \sum_{s=k}^{\infty} M_{t-k,s-k} \left(F_{k,s-k} - F_{k-1,s-k+1} \right) \qquad \forall t, j \geq 0 \tag{12}$$

This result, which we formally prove in appendix B.4, shows that no matter how complex the underlying learning model is, if it can be represented with a forecast matrix $\mathbf{F}$, the aggregate implications of the learning model can be obtained by simply "modifying" the iMPC matrix $\mathbf{M}$, to a new one, $\tilde{\mathbf{M}}$. The only information needed for this modification are the original FIRE iMPC matrix $\mathbf{M}$ as well as the forecast matrix $\mathbf{F}$. This also means that any two learning models that share the same, or a similar, forecast matrix, will have the same, or similar, aggregate implications.

The formula (12) for $\tilde{\mathbf{M}}$ is quite intuitive: $\tilde{\mathbf{M}}$ is lower triangular Toeplitz, that is, it has all zeros above the diagonal, and in each column j, from the diagonal down, it has the exact same entries, $\tilde{M}_{t,0}$. These entries correspond to the consumption response to a one-off income shock, holding all beliefs and the forecast matrix fixed. On impact, the consumption response is given by

$$\tilde{M}_{0,0} = \sum_{s=0}^{\infty} M_{0,s} F_{0,s}$$

which is precisely the consumption response of the average agent believing that income will increase by $F_{0,s}$, s periods in the future. As time goes by, however, the agent slowly comes to realize that there will not be a sustained increase in aggregate income. This leads to revisions in the

[4]The convention here is such that any occurrence of F with a negative index is assumed to be equal to zero.

agent's forecast. Specifically, k periods after the shock, the agent's forecast about date s income is revised by $F_{k,s-k} - F_{k-1,s-k+1}$. This forecast revision is a "news shock" to the agent, and is responded to by adjusting consumption in all periods $t \geq k$, in line with the intertemporal MPC $M_{t-k,s-k}$. This explains the formula in (12).

Of course the actual income shock that hits households, $\hat{\mathbf{Y}}$, is not just a one-off shock. Proposition 2 shows, however, that once the forecast matrix $\mathbf{F}$ is found, the consumption response is simply the sum of a sequence of responses to one-off shocks,

$$\hat{C}_t = \sum_{s=0}^{t} \tilde{M}_{t-s,0} \hat{Y}_s$$

The key is that this expression is taking the forecast matrix as given; in practice, in many examples, the stochastic process for aggregate income influences the forecast matrix: e.g. if one analyzes least squares learning, the first step may be to derive the correct estimate of the autocorrelation of dY_t. Once that is found, $\mathbf{F}$ can be calculated. After that, (11) is used to compute the response of consumption in the least-squares learning model.

Everything we derived here carries over to other blocks in sequence-space models, whether these are other household blocks (like RA or TA blocks), or other blocks with heterogeneity, such as heterogeneous-firm blocks. In the general case, $\mathbf{M}$ is simply the sequence-space Jacobian of the output whose response one would like to analyze with learning, with respect to the input that needs to be learned about. In appendix E, we generalize proposition 2 to *arbitrary* first-order deviation from FIRE.

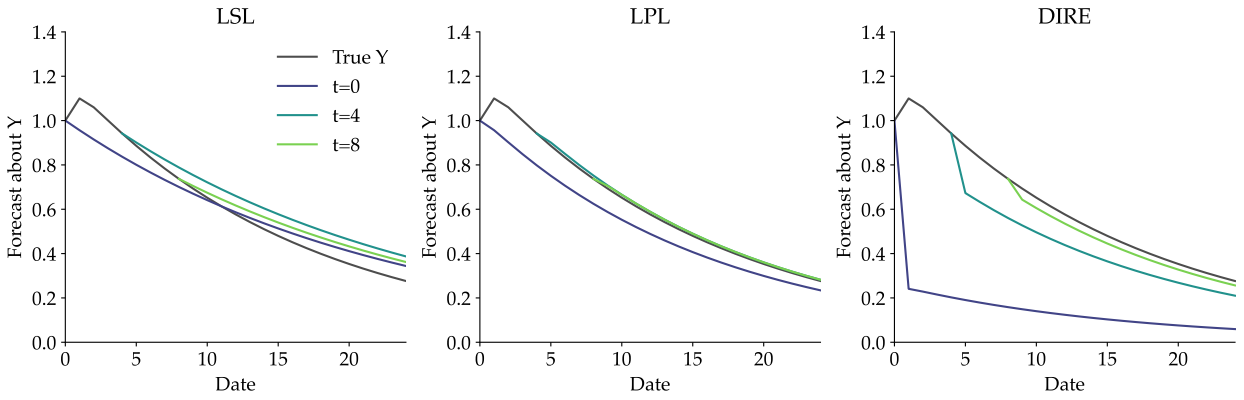
3.6 Numerical examples

We demonstrate proposition 2 with a numerical example. We assume aggregate income, dY_t, follows an exogenous $AR(2)$ process, with parameterization $dY_t = 1.1\, dY_{t-1} - 0.15\, dY_{t-2} + \epsilon_t$. The heterogeneous-agent model is calibrated as in section 4.2 and assumed to be in the steady state initially. We then compare the responses of the economy across the three learning models we outlined in section 3.4: (a) least-squares learning (LSL), where agents approximate dY_t with an estimated $AR(1)$ process; (b) local projection learning (LPL) where agents run local projections of dY_{t+h} on contemporaneous output dY_t; (c) dispersed-information rational expectations (DIRE) with iid sectoral income $d\psi_{it}$.

Figure 10 plots the correct $AR(2)$ income path (black), $\hat{Y}_s$, as well as the associated forecasts. The $AR(1)$ LSL forecasts initially undershoot, but eventually overshoot somewhat. The LPL forecasts also undershoot at first and then approximate $\hat{Y}_s$ fairly well. The DIRE forecasts are the slowest to respond to the shock, as any movement in $\psi_{it} Y_t$ is initially confused with the iid sectoral income. Only over time do agents learn that aggregate income moved.

Figure 11 computes the aggregate consumption response according to the three learning models and compares it with the FIRE solution. In the example, the LSL and LPL consumption responses are slightly hump-shaped, due to their initial undershooting. The DIRE consumption

Figure 10: Beliefs in response to an AR2 aggregate income shock



Note: We plot beliefs at each date t about future income in response to the AR2 aggregate income shock: $Y_t = 1.1Y_{t-1} - 0.15Y_{t-2} + \varepsilon_t^Y$. Under DIRE, the signal is $Y_{i,t} = Y_t + \varepsilon_{i,t}$, where $\varepsilon_{i,t} \sim \mathcal{N}\left(0, \sigma_\psi^2\right)$ and $\varepsilon_t^Y \sim \mathcal{N}\left(0, \sigma_Y^2\right)$. We set $\sigma_\psi = 0.15$, to match the overall calibrated sectoral income variance in section 4.2 below. We set $\sigma_Y = 0.04$ so the variance in $Y_{i,t}$ is driven equally by the aggregate and sectoral components.

response is quite inertial, due to agents' beliefs moving very slowly.

4 Heterogeneity and learning in general equilibrium

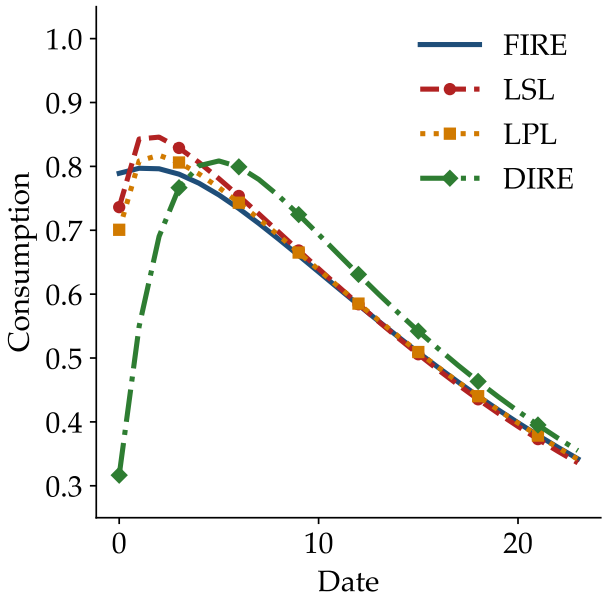
In the previous section, we analyzed three different learning models in partial equilibrium. Next, we move to a general equilibrium setting. Our model will be a conventional heterogeneous-agent New-Keynesian (HANK) model, endowed with dispersed-information rational expectations (DIRE). We interpret the beliefs formed by agents in response to past shocks as capturing their endogenous sentiments.

We argue that, in general equilibrium, the interaction between heterogeneity and sentiments gives rise to three important insights: First, sentiment is inertial – it takes time to build up, especially if output is slow to move itself. This causes a feedback mechanism that helps generate a more pronounced hump shape in monetary transmission than previous work (e.g. Angeletos and Huo 2021). Second, sentiments prolong recoveries from recessions, as pessimistic sentiment feeds into a weak economy, and vice-versa. Third, sentiment shocks themselves can move consumption and thus influence the business cycle. We show that these channels are significantly dampened with representative-agent or two-agent models, compared to a heterogeneous-agent model.

4.1 Setup

We set up the model as follows.

Figure 11: Consumption in responses to an AR2 aggregate income shock



Note: We plot the partial equilibrium consumption response to an exogenous AR2 aggregate income shock: $Y_t = 1.1Y_{t-1} - 0.15Y_{t-2} + \varepsilon_t^Y$.

Households. The household side is very close to the one introduced in section 3.1. The utility function now includes a disutility term from labor supply,

$$\max \mathbb{E}_{i\tau} \left[\sum_{t=\tau}^{\infty} \beta^{t-\tau} \left\{ u(c_{it}) - v(n_{it}) \right\} \right]$$

Labor supply n_{it} is effectively chosen by a continuum of unions to be specified below, and will be equal to

$$n_{it} = \psi_{it} N_t$$

where, as before, ψ_{it} is a stochastic process that is iid across agents. Per-period income y_{it} is now given by

$$y_{it} = (1 - \tau_t)\, e_{it} n_{it} w_t$$

where τ_t is the tax rate on labor income and w_t is the real wage. We will see below that $w_t = 1$ at all times. In the DIRE version of the model, we assume that households can perfectly observe e_{it}, but cannot distinguish the other terms entering y_{it}, that is,

$$y_{it} = e_{it} \times (1 - \tau_t)\, n_{it} w_t$$

where each factor is separately observable. In particular, households cannot distinguish ψ_{it} and N_t as drivers of n_{it}. The budget and borrowing constraints (4) and (5) are otherwise unchanged.

Production. There is a representative firm, converting labor one-for-one into final goods, $Y_t = N_t$. The goods market is perfectly competitive, so the real wage is equal to 1, $w_t = 1$.

Monetary policy. For this section, we assume that the central bank implements an exogenous real interest rate r_t between periods t and $t+1$. As usual, we can think of this as a Taylor rule policy of the sort $i_t = r_t + \phi_\pi \mathbb{E}_t \pi_{t+1}$, with $\phi_\pi \searrow 1$.

Fiscal policy. There is a government issuing real government debt B_t, purchasing G goods each period, and setting tax rate τ_t. Its budget constraint is

$$G + (1 + r_{t-1}) B_{t-1} = B_t + \tau_t Y_t$$

In case the real interest rate r_t fluctuates, we assume that τ_t adjusts to keep debt repayment $(1 + r_t) B_t$ equal to its steady state value $(1 + r)B$, that is, τ_t is given by

$$\tau_t = \frac{1}{Y_t} \left\{ G + (1 + r) B - \frac{(1 + r) B}{(1 + r_t)} \right\}$$

The behavior is intuitive: a lower real interest rate r_t relaxes the government's budget constraint, and allows for a reduced tax rate.

Nominal rigidity. We build on the wage rigidity formulations in Erceg, Henderson and Levin (2000) for an RA model and Auclert et al. (2024) for a HA model. There is a continuum of monopolistically competitive unions, and each household stands ready to supply each union with an infinitesimal amount of labor. Due to union-specific shocks, we assume that household i works $n_{it} = \psi_{it} N_t$ in period t, which can be above or below the average hours worked in the economy. The wage setting behavior by unions translates into a Phillips curve, which we derive in appendix C.1. Given that our simulations in this section will assume a real interest rate rule, the details of the Phillips curve will not matter for our results here. We will discuss the Phillips curve in detail in section 5.

4.2 Calibration

We model $\log e_{it}$ as a discretized process that has a $AR(1)$ component and an iid component. The AR(1) component has persistence ρ_e and cross-sectional standard deviation $\sigma_{e,pers}$. The iid component has standard deviation $\sigma_{e,iid}$. We assume $\log \psi_{it}$ follows the exact same process,

$$\log \psi_{it} = \log \psi_{it}^P + \varepsilon_{i,t}^{\psi,\text{iid}}$$

$$\log \psi_{it}^P = \rho_e \log \psi_{it-1}^P + \varepsilon_{i,t}^{\psi,\text{pers}} \tag{13}$$

Parameter	Symbol	Value	Target
GDP	Y	1	Normalization
Gov. spending / GDP	G	0.2	20% of GDP
Gov. debt / GDP	B	0.69	Annual MPC of 0.51
Real interest rate	r	0.005	2% real rate p.a.
Consumption / GDP	C	0.8	1 - gov. spending share
Labor income tax rate	τ	0.204	Implied by gov. budget
Discount factor	β	0.968	Match real interest rate

Table 1: Calibration

except that its overall cross-sectional standard deviation is given by $\sigma_\psi \sigma$. We choose $\rho_e = 0.985$ quarterly, $\sigma_{e,pers} = 0.873$, $\sigma_{e,iid} = 0.291$, very much in line with the literature estimating income processes.[5] We set the relative size of sectoral shocks $\log \psi_{it}$ equal to $\sigma_\psi = 0.3$. We assume $u(c) = \log c$ and calibrate β to match an annual MPC of 0.51 at a steady state interest rate of 2% annually. This annual MPC estimate is equal to the one found by Fagereng et al. (2021) and close to those typically estimated in the literature (see Jappelli and Pistaferri 2014 for a survey). We normalize steady state output to 1, $Y = N = 1$, and assume government spending is 20% of GDP. Government debt is equal to asset demand by households at the calibrated β, and comes out to be 19% of annual GDP.[6] We do not have to take a stance on the disutility from labor supply $v(n)$ here, as it only enters the Phillips curve, and inflation is irrelevant with a real interest rate rule (see Auclert et al. 2024 for details). We specify it in section 5 below. Table 1 summarizes our calibrated parameters.

Next, we take this heterogeneous-agent model and analyze three distinct experiments with it, always comparing outcomes with different kinds of learning to the outcome under FIRE. The experiments are a monetary policy shock, a demand shock, and a sentiment shock. At the end, we show how all three results crucially hinge on heterogeneity and high intertemporal MPCs.

4.3 Monetary policy: Sentiment inertia

We begin by considering an AR(1) monetary policy shock in the model,

$$r_t = \rho_r r_{t-1} + \varepsilon_t^r, \qquad \varepsilon_t^r \sim \mathcal{N}\left(0, \sigma_r^2\right)$$

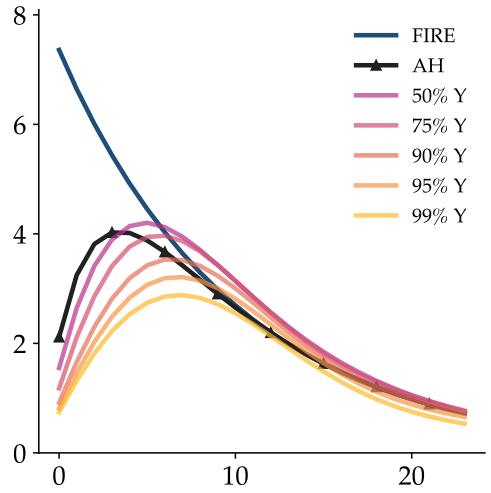
We set the persistence to $\rho_r = 0.9$ similar to that in Romer and Romer (2004). We set the standard deviation of the monetary policy shock to 100bp, $\sigma_r = 0.01$.

We assume households learn about interest rates similar to how they learn about their income. Specifically, under LSL and LPL, households run $AR(1)$ and LP regressions to predict interest rates. Since r_t is an $AR(1)$ itself, this means households will be able to perfectly forecast interest rates. Under DIRE, we follow Angeletos and Huo (2021) and assume households observe a noisy

[5] See, e.g., Flodén and Lindé (2001), Guvenen (2009).

[6] While this is clearly too low, we note that it is more important for this section to match iMPCs than it is to match total wealth. The quantitative model in the next section will have a reasonable amount of household wealth.

Figure 12: Learning about a monetary policy shock



Note. This figure shows the consumption response to an accommodative monetary policy shock. FIRE denotes the full information rational expectations response. "AH" stands for Angeletos-Huo, and denotes a version of our model in which 100% of the information about the monetary shocks stems from dispersed information about interest rates. "50% Y" stands for a model in which 50% of the information stems instead from observing a noisy version of aggregate income, $\psi_{it}Y_t$. Similarly for the remaining lines. x-axis: Quarters.

signal of the interest rate,

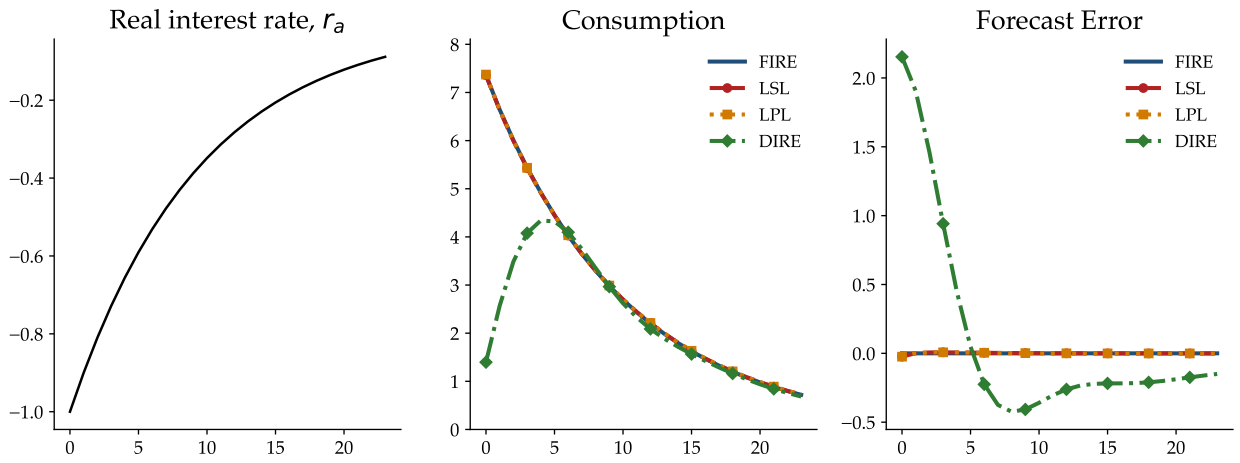
$$s^r_{it} = r_t + \eta^r_{i,t}, \qquad \eta^r_{i,t} \sim \mathcal{N}\left(0, \sigma^2_{\eta_r}\right)$$

Thus, in that model, households combine their information from s^r_{it} as well as from $\psi_{it}Y_t$ to learn about the monetary policy shock. To parameterize the information content of each of these sources, σ_{η_r} and σ_ψ, we imagine a single household in the FIRE version of the model, obtaining information from s^r_{it} and ψ_{it}. We target a joint informativeness of both signals, such that, upon impact of a monetary shock ϵ^r_t, the household expectation of $\mathbb{E}_t[r_{t+1}]$ is a fraction 30% of the true interest rate shock r_{t+1}. Moreover, we parameterize by χ the share of information coming from learning from output, $\psi_{it}Y_t$. The remainder, share $1-\chi$, comes from the exogenous signal about interest rates.

Figure 12 plots the consumption response of the economy under alternative assumptions about χ. With $\chi = 0$, and thus no learning from output, we are in the world of Angeletos and Huo (2021), with learning entirely from an exogenous AR(1) process. DIRE generates a hump shaped response as households are initially unsure about whether a monetary shock in fact happened. As we increase χ, and push it closer to 100%, households learn more and more from their income. But income responds slowly to the shock, as it is hump-shaped, slowing down their learning. But slow learning implies an even slower response of output to the monetary shock, and so on. This feedback mechanism causes household sentiment to be very inertial, amplifying the hump-shaped output response to a monetary shock.

In figure 13, we return to the calibrated process for ψ_{it} from section 4.2, and calibrate the signal noise to set $\chi = 0.9$. We compare this DIRE model to the two alternative learning models, LSL and LPL, as well as the FIRE solution. We see that LSL, LPL, and FIRE all nearly coincide. This

Figure 13: Monetary policy shock



Note. $AR(1)$ real interest rate shock with quarterly persistence of 0.90 in the HA economy. FIRE = full-information rational expectations. LSL = least squares learning (households estimate an $AR(1)$). LPL = local projection learning (households form beliefs using a local projection of future output on current output). DIRE = dispersed information rational expectations (households learn about output through their own income, partly obscured by idiosyncratic shocks). Forecast error computed as the difference between actual realized output 4 quarters from now and the current forecast.

is because the output response in the FIRE model is almost equal to an $AR(1)$, which LSL and LPL can both learn perfectly. The DIRE response first undershoots (positive forecast errors), then overshoots, mirroring our fact 3 in section 2.5.

4.4 Recessions: Sentiment accelerator

Next, we consider a demand shock, which we model as a negative shock to the discount factor, β_t. We model the shock as an $AR(1)$,

$$\log \beta_t = \rho_\beta \log \beta_{t-1} + \epsilon_t^\beta$$

with preferences given by

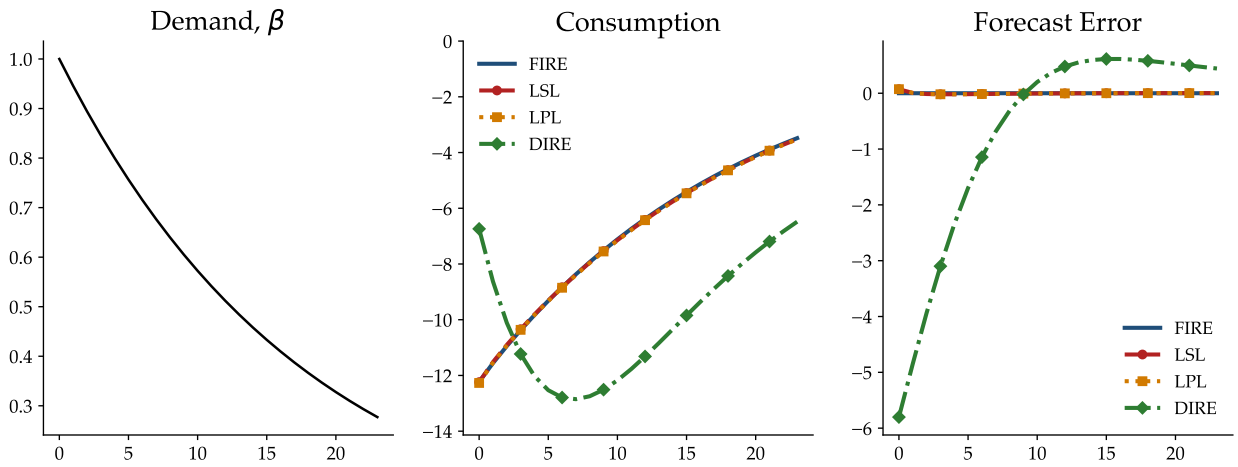
$$\max \mathbb{E}_{i\tau} \left[\sum_{t=\tau}^{\infty} \prod_{s=\tau}^{t-1} \beta_s \left\{ u(c_{it}) - v(n_{it}) \right\} \right]$$

Every agent perfectly observes the discount factor shock.

Figure 14 shows the impulse responses a reduction in demand, that is, $\epsilon_t^\beta > 0$. Consumption falls across the board, but again, since output in the FIRE economy roughly inherits the $AR(1)$ structure, LSL and LPL are very close to FIRE. The DIRE economy not only underreacts initially; it is also quite a bit more persistent as time goes by.

This difference in persistence is especially strong if the demand shock itself is somewhat persistent. We show this in figure 15, where we plot the half life of the output response as function of the half life of the demand shock. As the demand shock gets more and more persistent, the

Figure 14: Demand shock



Note. Discount factor shock, annual persistence 0.75 in the HA economy. Legend just like in figure 13. Forecast error computed as the difference between actual realized output 4 quarters from now and the current forecast.

DIRE economy generates more and more additional persistence, leading to a slower and slower recovery from the demand shock.

This is driven by an important dynamic interaction between sentiments and aggregate demand we call the "sentiment accelerator", similar to the "investment accelerator" Samuelson (1939). Imagine a recession caused by a negative demand shock is over, but households are still pessimistic. This induces them to spend less. But less spending leads to lower income, and lower income prevents a fast normalization of household sentiments.

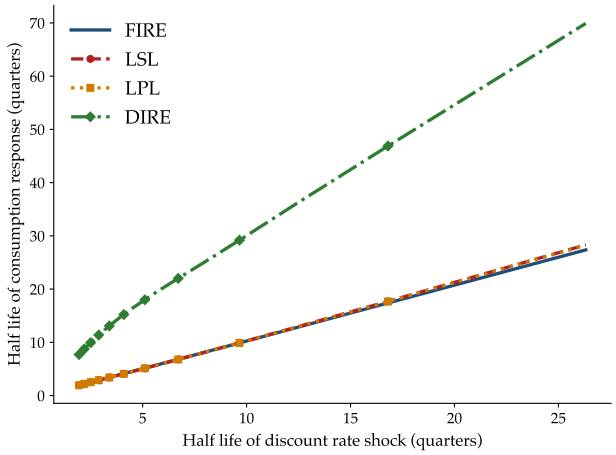
4.5 Leaning against sentiment

The increased persistence of demand shocks, due to the sentiment accelerator, suggests a more important role for the timing of monetary policy. If the central bank leaves a demand shock to play out, sentiment builds over time, and the central bank faces a worse trade off at the later date (even setting aside the costs borne up to that point). To see this, first define a simple central bank loss function only in the output gap,

$$\mathcal{L}_t = \sum_{\tau=0}^{\infty} \beta^\tau y_{t+\tau}^2$$

We then shock the economy with a discount rate shock, just as in section 4.4. If the central bank does nothing, the present value losses at t are then $\mathcal{L}_t^{\text{passive}}$. Now suppose the central bank becomes active only in period t, and at that points commits to an $AR(1)$ monetary policy shock, as in section 4.3. Under the optimally sized monetary intervention, the losses are $\mathcal{L}_t^{\text{active}}$. Then the percent of losses that can be averted by first intervention in period t, evaluated in period t, is $100\left(1 - \frac{\mathcal{L}_t^{\text{active}}}{\mathcal{L}_t^{\text{passive}}}\right)$. DIRE beliefs in this model are solved for households observing income, subject

Figure 15: Inertia amplifies persistence of demand



Note. Discount factor shock in the HA economy. Plot varies half-life $-\log 2/\log \rho_\beta$ in quarters of the discount factor shock, and for each computes the half-life of the consumption response, i.e. finds date t^* such that $dC_{t^*} = 0.5dC_0$. Legend just like in figure 13.

to the $\psi_{i,t}$ process and also driven by discount rate and monetary policy shocks, and observing a monetary policy signal with noise as calibrated in figure 13. We plot the gains from intervention at time t in figure 16.

4.6 Sentiment shocks

Consumer sentiment is not always slow-moving. In the following exercise, we build on the existing literature on sentiment shocks (Lorenzoni 2009, Angeletos and La'o 2013, Angeletos et al. 2018 among others) and simultaneously shock households' beliefs about the $AR(1)$ component of ψ_{it}: all households believe income will increase, fueling optimism, while no fundamental shock has happened. We introduce a signal

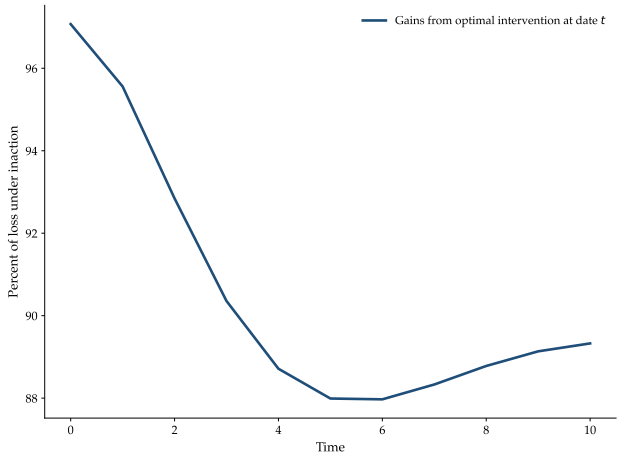
$$s_{it}^{\psi} = \varepsilon_{i,t}^{\psi,\text{pers}} + \eta_{i,t}^{\psi}, \qquad \eta_{i,t}^{\psi} \sim \mathcal{N}\left(0, \sigma_{\eta_\psi}^2\right)$$

where $\varepsilon_{i,t}^{\psi,\text{pers}}$ is a shock to the persistent component of ψ_{it} income, defined in (13). We set σ_{η_ψ} equal to the standard deviation of $\varepsilon_{i,t}^{\psi,\text{pers}}$. We label a one-off, ex-post common shock to $\eta_{i,t}^{\psi}$ a sentiment shock. This does not affect any household's belief about their income today. Figure 17 shows the impulse response to the sentiment shock. A one-time shock to sentiment causes a strong increase in consumption. This increase sets off the "sentiment accelerator", prevents a fast normalization of the optimistic sentiment, and propagates the shock into the future.

4.7 Role of heterogeneity

Each one of the three experiments above requires non-trivial intertemporal MPCs, both such that optimistic or pessimistic sentiment translates into actual spending decisions, and such that current income changes translate into future spending decisions.

Figure 16: Gains from early intervention



Note. Plot considers $AR(1)$ demand shock and asks, at each quarter t after the impact of the shock: What % of the discounted loss $\sum_{\tau=0}^{\infty} \beta^{\tau} y_{t+\tau}^2$ can be averted by using monetary policy from date t onwards. Greater number = bigger gain from intervening. Monetary policy assumed to be $AR(1)$ as in figure 13 but optimally chosen scale.

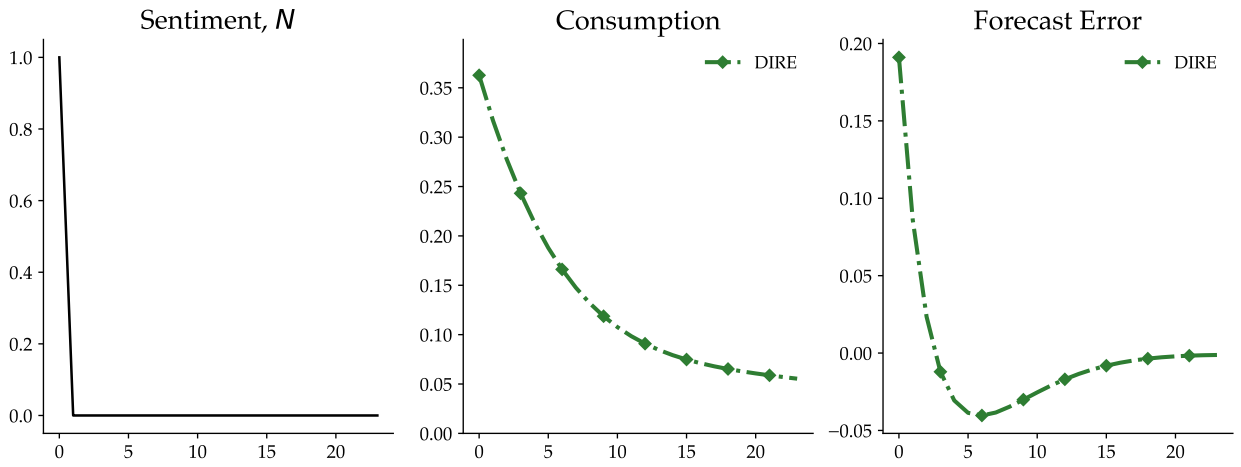
We show this in figure 18, which computes the consumption responses in each of the three experiments for a representative-agent (RA), two-agent (TA), and our baseline HA model. All models are calibrated in exactly the same way and have the exact same information structures as the HA model.

Panel a compares the response to the monetary shock from section 4.3. With DIRE, all three generate a hump shape, though it is most pronounced in the HA model. Our mechanism for sentiment inertia is also present in the RA and TA models, since learning from the endogenously hump-shaped aggregate income response is also slow in those models. However, in those models, almost all learning is about the future path of real interest rates, as this is the most important variable to determine household behavior. In the HA economy, instead, the learning about future output is a lot more important. As households learn that a monetary policy shock actually happened, they expect greater incomes and spend in anticipation. This leads to the more pronounced hump shape.

Panel b considers the demand shock from section 4.4. Quite strikingly, the RA and TA models do not generate much of a sentiment accelerator. This is precisely for the reasons we mentioned above: The accelerator stems from the interaction between persistence in sentiment and the persistence in spending via intertemporal MPCs (iMPCs). The RA and TA models lack any persistence in iMPCs (figure 9) and therefore do not generate a meaningful sentiment accelerator. We show how this interaction matters analytically in the context of a bonds-in-utility model in appendix C.2

Panel c revisits the sentiment shock from section 4.6. In RA and TA, iMPCs out of anticipated income are small, so optimistic sentiment barely moves spending. This is different in the HA model, which is why shocks to sentiment have a large effect on spending here.

Figure 17: Sentiment shock



Note. Shock to belief about persistent component of ψ_{it}. Forecast error computed as the difference between actual realized output 4 quarters from now and the current forecast.

5 Quantitative model *[incomplete]*

In this section, we bring our model of heterogeneous agents with dispersed-information rational expectations (DIRE) to U.S. quarterly data. The quantitative model enriches the general-equilibrium framework of the previous section along several dimensions. On the supply side, we introduce sticky prices via a New Keynesian Phillips curve and an inertial Taylor rule for monetary policy. On the household side, the information structure is enriched to include learning from multiple signals. We estimate the model using Bayesian methods on five quarterly U.S. time series—output, consumption, hours, the policy rate, and our measure of household sentiment—over the period 1966Q1–2019Q4.

Our key finding is that sentiment plays a quantitatively important role in business cycles, through both exogenous sentiment shocks and the endogenous amplification and propagation of other shocks via the learning channel. Exogenous sentiment shocks account for XX% of the variance of de-trended GDP. Endogenous sentiment movement—that is, the additional GDP fluctuations generated by DIRE relative to FIRE for the same non-sentiment shocks—accounts for a further XX% of the variance of GDP.

5.1 Model

The household block is the same as in section 4.2, with the same income process, budget constraint, and borrowing constraint. The household objective is still

$$\max \mathbb{E}_{i\tau}\left[\sum_{t=\tau}^{\infty}\beta^{t-\tau}\left\{u(c_{it})-v(n_{it})\right\}\right]$$

where $u(c) = \log c$ and the disutility of labor takes the form

$$v(n) = \frac{\chi}{1+\phi^{-1}} n^{1+\phi^{-1}} \tag{14}$$

Here, ϕ is the Frisch elasticity of labor supply and $\chi > 0$ a scaling parameter. We embed our heterogeneous household block in a general-equilibrium environment similar to the one in section 4, with a few modifications, which we now describe.

Production and prices. We now allow for a productivity shifter A_t, that is, output is produced using the technology

$$Y_t = A_t \cdot N_t$$

We continue assuming perfect competition and flexible prices in the goods market.

Wage setting. We continue assuming nominal wage rigidities, with a union-based microfoundation. As derived in appendix C.1 in great detail, we derive the following wage Phillips curve,

$$\pi_t^w(1+\pi_t^w) = \kappa \int u'(c_{it}) \left(\frac{v'(n_{it})}{u'(c_{it})\,(1-\tau_t)\,A_t} - 1 \right) di + \beta \mathbb{E}_t \left[\pi_{t+1}^w \,(1+\pi_{t+1}^w) \right]$$

where $n_{it} = \psi_{it} N_t$. To keep things simple, we assume unions have rational expectations.

Monetary policy. The central bank sets the nominal policy rate i_t according to an inertial Taylor rule,

$$1 + i_t = \left(1 + i_t^{\text{shock}}\right)^{1-\rho_m} (1+i_{t-1})^{\rho_m} \,(1+\pi_t)^{(1-\rho_m)\phi_\pi} \tag{15}$$

where ρ_m governs interest-rate smoothing, ϕ_π is the inflation response coefficient, and i_t^{shock} is an exogenous monetary policy shifter. The real interest rate r_t is defined as the ex ante real rate,

$$1 + r_t = \frac{1+i_t}{1+\mathbb{E}_t \pi_{t+1}}$$

Fiscal policy. The government purchases G_t goods each period, issues nominal debt, and levies a proportional tax on labor income. Taxes adjust to keep debt repayment $(1+r_t)B_t$ at its steady-state value. Aggregate post-tax income is $Z_t \equiv Y_t - T_t$, where T_t denotes total taxes.

Sentiment. We define model sentiment as the average household expectation of their post-tax income four quarters ahead,

$$\text{Sentiment}_t \equiv \int \mathbb{E}_{it} \left[y_{i,t+4} \right] di \tag{16}$$

Under FIRE, this reduces to the rational-expectations forecast of Z_{t+4}. Under DIRE, it reflects the average of Bayesian forecasts of future income across households, incorporating the learning dynamics discussed in the previous sections.

5.2 Data

We construct quarterly time series for five observables over the sample period 1966Q1–2019Q4. Real output Y_t is measured as nominal GDP divided by the GDP deflator. Real consumption C_t is nominal personal consumption expenditures (PCE) divided by the same GDP deflator to make it comparable to GDP. Hours N_t are nonfarm business sector hours. The policy rate i_t is the quarterly average of the effective federal funds rate, divided by 400 to express it in quarterly units comparable to the model. For output, consumption, and hours we take logs.

For sentiment, we use the first principal component of the forward-looking questions in the Michigan Survey of Consumers, as constructed in section 2.2. To express the sentiment index in model-consistent units, we scale it by $\frac{1}{200}$, reflecting our finding that a one-unit movement in the index corresponds to approximately a half percentage point shift in one-year-ahead income growth expectations.

All five series are detrended using the Hodrick–Prescott filter with smoothing parameter $\lambda =$ 10,000, as an intermediate between the regular quarterly smoothing parameter of $\lambda =$ 1,600 and linear de-trending. The cyclical component of consumption is further rescaled by steady-state consumption to express it as a share of output, consistent with the model's units.

5.3 Calibration

All non-shock parameters are calibrated prior to estimation. The calibration is summarized in Table 2. Household preferences feature log utility and a Frisch elasticity of labor supply of $\phi = 0.5$. The elasticity of substitution across intermediate varieties is $\xi = 5$, implying a 25% steady-state markup. The slope of the wage Phillips curve is given by $\kappa = 0.01$. The Taylor rule parameters are $\phi_\pi = 1.5$ and $\rho_m = 0.89$.

The income process is calibrated as in section 4.2, with quarterly persistence $\rho_e = 0.985$, total idiosyncratic income standard deviation of 0.92, a share of known idiosyncratic income of 90%, and a transitory share of 10%. We normalize steady-state output and TFP to 1. Government spending is 20% of GDP. The steady-state real interest rate is 2% annually, which we match by varying the discount factor β. The level of government debt is calibrated to match a first-year MPC of 0.51, in line with the evidence in Fagereng et al. (2021).

5.4 Estimation strategy

Shock processes. We allow for five structural disturbances, each following an $AR(1)$ process with innovation standard deviation σ_j and persistence ρ_j:

a) Government spending, G_t;

b) Total factor productivity, A_t;

c) Monetary policy, i_t^{shock}, the exogenous component of the Taylor rule;

Parameter	Symbol	Value	Target / Source
Frisch elasticity	ϕ	0.5	Standard
Discount factor	β	solved	First-year MPC = 0.51
Elasticity of substitution	ξ	5	25% markup
Slope of wage Phillips curve	κ	0.01	
Steady-state TFP	A	1	Normalization
Inflation coefficient	ϕ_π	1.5	Standard
Interest-rate smoothing	ρ_m	0.89	Standard
Government spending	G/Y	0.2	Standard
Steady-state real rate	r	0.02	Standard
Steady-state output	Y	1	Normalization
Persistence	ρ_e	0.985	0.94 annual persistence
Total idiosyncratic s.d.	σ_y	0.92	U.S. data

Table 2: Calibrated parameters for the quantitative model

d) Discount factor (demand), β_t;

e) Sentiment signal noise, η_t^ψ, an exogenous common innovation to the sector signal.

The first four shocks are standard in the DSGE literature. The fifth shock—sentiment signal noise—deserves further discussion. In our model, households observe a sector signal that they perceive to be driven entirely by idiosyncratic persistent income shocks and iid signal noise. When η_t^ψ is positive, all households simultaneously receive a signal that is more favorable than warranted by fundamentals, generating a coordinated wave of optimism. This is the "sentiment shock" analyzed in section 4.6, now disciplined by the data.

Estimated parameters. There are ten estimated parameters: the innovation standard deviations $(\sigma_G, \sigma_A, \sigma_i, \sigma_\beta, \sigma_\eta)$ and the $AR(1)$ persistence coefficients $(\rho_G, \rho_A, \rho_i, \rho_\beta, \rho_\eta)$ for each of the five shocks. All other parameters are calibrated as described above.

Priors. We place independent priors on each parameter. Each innovation standard deviation σ_j follows an inverse-gamma distribution with shape and scale parameters 0.1 and 2. Each persistence parameter ρ_j follows a beta distribution with mean 0.5 and standard deviation 0.2. Table 3 summarizes the prior specification and reports the posterior mode under both the FIRE and HA-DIRE models.

Likelihood and posterior mode. The linearized model implies a mapping from the five structural shocks to the five observables. For given shock parameters $\theta = (\sigma_1, \ldots, \sigma_5, \rho_1, \ldots, \rho_5)$, we compute the model-implied autocovariance of the observables under the assumption that the structural innovations are iid standard normal. The log-likelihood $\log p(\mathbf{Y} \mid \theta)$ is evaluated as the multivariate normal density given this covariance structure. The posterior mode is obtained

Parameter	Symbol	Prior	Prior param.	FIRE	HA-DIRE
Gov. spending	σ_G	Inv-Gamma	(0.1, 2)	XX	XX
TFP	σ_A	Inv-Gamma	(0.1, 2)	XX	XX
Monetary policy	σ_i	Inv-Gamma	(0.1, 2)	XX	XX
Demand	σ_β	Inv-Gamma	(0.1, 2)	XX	XX
Sentiment noise	σ_η	Inv-Gamma	(0.1, 2)	XX	XX
Gov. spending	ρ_G	Beta	(0.5, 0.2)	XX	XX
TFP	ρ_A	Beta	(0.5, 0.2)	XX	XX
Monetary policy	ρ_i	Beta	(0.5, 0.2)	XX	XX
Demand	ρ_β	Beta	(0.5, 0.2)	XX	XX
Sentiment noise	ρ_η	Beta	(0.5, 0.2)	XX	XX

Table 3: Estimated shock parameters: priors and posterior modes

Note. Each shock follows an $AR(1)$ with innovation standard deviation σ_j and persistence ρ_j. Posterior modes are obtained by maximizing the log posterior using Powell's method.

by maximizing $\log p(\mathbf{Y} \mid \theta) + \log p(\theta)$ using Powell's derivative-free method with box constraints $\sigma_j \in [10^{-4}, 1]$ and $\rho_j \in [0, 0.99]$.

[Rest of this section: TBD]

6 Conclusion

This paper has developed a new framework for understanding the role of sentiment in business cycles. By integrating learning and dispersed information into a heterogeneous-agent New-Keynesian model, we provide a tight link between sentiments—interpreted as expectations about future income—and aggregate demand. Households in our model cannot fully distinguish idiosyncratic from aggregate income and learn about future income from past income realizations. Combined with realistically high intertemporal MPCs, this gives rise to three key mechanisms: delayed overreaction of beliefs to aggregate shocks; a "sentiment inertia" effect, in which the interaction between slow-moving beliefs and slow-moving aggregate income generates hump-shaped consumption and output responses; and a "sentiment accelerator," whereby pessimistic sentiment and weak economic activity reinforce one another, prolonging recoveries from recessions. These mechanisms rely crucially on the interaction between heterogeneity and dispersed information; neither ingredient alone is sufficient.

Our paper also contributes a portable methodology for embedding learning into sequence-space models. The key insight is that many common forms of belief formation—including dispersed-information rational expectations, least-squares learning, diagnostic expectations, and others—can be summarized by a forecast matrix $\mathbf{F}$ that maps past realizations of aggregate variables into average expectations of their future values. This matrix can be used to transform the Jacobians of the household block, allowing any model solved in the sequence space to be augmented with learning at minimal additional cost. We hope this framework proves useful for future work exploring the macroeconomic consequences of non-rational expectations.

References

Adams, Jonathan J and Eugenio Rojas, "Household signal extraction and slow recoveries," Technical Report, Working Paper 2024.

Angeletos, George-Marios and Chen Lian, "Confidence and the propagation of demand shocks," *The Review of Economic Studies*, 2022, *89* (3), 1085–1119.

__ and Jennifer La'o, "Sentiments," *Econometrica*, 2013, *81* (2), 739–779.

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

__ , Chen Lian, and Christian K Wolf, "Can deficits finance themselves?," *Econometrica*, 2024, *92* (5), 1351–1390.

__ , Fabrice Collard, and Harris Dellas, "Quantifying confidence," *Econometrica*, 2018, *86* (5), 1689–1726.

__ , 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.

__ , Hugo Monnery, Matthew Rognlie, and Ludwig Straub, "Managing an energy shock: Fiscal and monetary policy," Technical Report, NBER Working Paper 31543 2023.

__ , Matthew Rognlie, and Ludwig Straub, "Micro jumps, macro humps: Monetary policy and business cycles in an estimated HANK model," Technical Report, NBER Working Paper 26647 2020.

__ , __ , and __ , "The intertemporal keynesian cross," *Journal of Political Economy*, 2024, *132* (12), 4068–4121.

__ , __ , Martin Souchier, and Ludwig Straub, "Exchange rates and monetary policy with heterogeneous agents: Sizing up the real income channel," Technical Report, NBER Working Paper 28872 2021.

Bachmann, Rüdiger, Tim O. Berg, and Eric R. Sims, "Inflation Expectations and Readiness to Spend: Cross-Sectional Evidence," *American Economic Journal: Economic Policy*, 2015, *7* (1), 1–35.

Barberis, Nicholas, Andrei Shleifer, and Robert Vishny, "A model of investor sentiment," *Journal of Financial Economics*, 1998, *49* (3), 307–343.

Bardóczy, Bence and Joao Guerreiro, "Expectation formation in general equilibrium," Technical Report, Working Paper 2024.

Bhandari, Anmol, Jaroslav Borovicka, and Paul Ho, "Survey Data and Subjective Beliefs in Business Cycle Models, Working Paper 19-14," 2019.

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

__ , __ , and __ , "Memory, attention, and choice," *Quarterly Journal of Economics*, 2020, *135* (3), 1399–1442.

__ , __ , Yueran Ma, and Andrei Shleifer, "Overreaction in macroeconomic expectations," *American Economic Review*, 2020, *110* (9), 2748–2782.

Carroll, Christopher D, "Macroeconomic expectations of households and professional forecasters," *Quarterly Journal of Economics*, 2003, *118* (1), 269–298.

Cogley, Timothy and Thomas J Sargent, "The Conquest of U.S. Inflation: Learning and Robustness to Model Uncertainty," *Review of Economic Dynamics*, 2005, *8* (2), 528–563.

Colarieti, Roberto, Pierfrancesco Mei, and Stefanie Stantcheva, "The how and why of household reactions to income shocks," Technical Report, National Bureau of Economic Research 2024.

Curtin, Richard, "Consumer Expectations: Micro Foundations and Macro Impact," Technical Report, University of Michigan, Survey Research Center 2019.

Erceg, Christopher J, Dale W Henderson, and Andrew T Levin, "Optimal monetary policy with staggered wage and price contracts," *Journal of monetary Economics*, 2000, *46* (2), 281–313.

European Commission, "The Joint Harmonised EU Programme of Business and Consumer Surveys: User Guide," Technical Report, Directorate-General for Economic and Financial Affairs 2024.

Eusepi, Stefano and Bruce Preston, "Expectations, learning, and business cycle fluctuations," *American Economic Review*, 2011, *101* (6), 2844–2872.

Evans, George W and Seppo Honkapohja, *Learning and Expectations in Macroeconomics*, Princeton University Press, 2001.

Fagereng, Andreas, Martin B Holm, and Gisle J Natvik, "MPC heterogeneity and household balance sheets," *American Economic Journal: Macroeconomics*, 2021, *13* (4), 1–54.

Farhi, Emmanuel and Iván Werning, "Monetary policy, bounded rationality, and incomplete markets," *American Economic Review*, 2019, *109* (11), 3887–3928.

Farmer, Leland E, Emi Nakamura, and Jón Steinsson, "Learning about the long run," *Journal of Political Economy*, 2024, *132* (10), 3334–3377.

Flodén, Martin and Jesper Lindé, "Idiosyncratic risk in the United States and Sweden: Is there a role for government insurance?," *Review of Economic Dynamics*, 2001, *4* (2), 406–437.

Fuster, Andreas, David Laibson, and Brock Mendel, "Natural expectations and macroeconomic fluctuations," *Journal of Economic Perspectives*, 2010, *24* (4), 67–84.

Gabaix, Xavier, "A behavioral New Keynesian model," *American Economic Review*, 2020, *110* (8), 2271–2327.

Guerreiro, Joao, "Expectation formation and heterogeneous-agent economies," Technical Report, Job Market Paper 2024.

Guvenen, Fatih, "An empirical investigation of labor income processes," *Review of Economic Dynamics*, 2009, *12* (1), 58–79.

Huo, Zhen and Naoki Takayama, "Rational expectations models with Higher-Order beliefs," *Review of Economic Studies*, 2025, *92* (5), 3138–3173.

Jappelli, Tullio and Luigi Pistaferri, "Fiscal policy and MPC heterogeneity," *American Economic Journal: Macroeconomics*, 2014, *6* (4), 107–136.

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

Keynes, John Maynard, *The General Theory of Employment, Interest and Money*, Macmillan, 1936.

Kozicki, Sharon and Peter A Tinsley, "Shifting endpoints in the term structure of interest rates," *Journal of Monetary Economics*, 2001, *47* (3), 613–652.

Krusell, Per and Anthony A Smith, "Income and wealth heterogeneity in the macroeconomy," *Journal of Political Economy*, 1998, *106* (5), 867–896.

Lenney, Jamie and Federico Rosso, "Learning in sequence-space models," Technical Report, Working Paper 2025.

Lorenzoni, Guido, "A theory of demand shocks," *American economic review*, 2009, *99* (5), 2050–2084.

Lucas, Robert E, "Expectations and the neutrality of money," *Journal of Economic Theory*, 1972, *4* (2), 103–124.

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

Marcet, Albert and Juan P Nicolini, "Recurrent hyperinflations and learning," *American Economic Review*, 2003, *93* (5), 1476–1498.

McKay, Alisdair, Emi Nakamura, and Jón Steinsson, "The power of forward guidance revisited," *American Economic Review*, 2016, *106* (10), 3133–3158.

Monnery, Hugo, "Adaptive learning in heterogeneous-agent models," Technical Report, Working Paper 2025.

Newey, Whitney K and Kenneth D West, "Hypothesis testing with efficient method of moments estimation," *International Economic Review*, 1987, pp. 777–787.

Nimark, Kristoffer P, "Man-bites-dog business cycles," *American Economic Review*, 2014, *104* (8), 2320–2367.

Oh, Hyunseung and Ricardo Reis, "Targeted transfers and the fiscal response to the great recession," *Journal of Monetary Economics*, 2012, *59*, S50–S64.

Overhage, Tobias, "Dispersed information in the sequence space," Technical Report, Working Paper 2026.

Pfäuti, Oliver and Fabian Seyrich, "A behavioral heterogeneous agent New Keynesian model," Technical Report, Working Paper 2024.

Ramey, Valerie A, "Macroeconomic shocks and their propagation," *Handbook of macroeconomics*, 2016, *2*, 71–162.

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

Rondina, Giacomo and Todd B Walker, "Information equilibria in dynamic economies with dispersed information," Technical Report, Working Paper 2009.

__ and __ , "Confounding dynamics," *Journal of Economic Theory*, 2021, *196*, 105287.

Samuelson, Paul A, "Interactions between the multiplier analysis and the principle of acceleration," *The Review of Economics and statistics*, 1939, *21* (2), 75–78.

Sundram, Jacob, "Fiscal policy in small open economies: The international intertemporal Keynesian cross," Technical Report, Working Paper 2025.

Werning, Iván, "Incomplete markets and aggregate demand," Technical Report, NBER Working Paper 21448 2015.

Woodford, Michael, *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press, 2003.

Yang, Yucheng, Chiyuan Wang, Andreas Schaab, and Benjamin Moll, "Structural Reinforcement Learning for Heterogeneous Agent Macroeconomics," *arXiv preprint arXiv:2512.18892*, 2025.