Anatomy of Subjective Macroeconomic Expectations*

Jose Lara† and Dalton Rongxuan Zhang‡

Preliminary and Incomplete

Abstract

We develop a structural approach to identify latent belief shocks underlying the cross-sectional subjective macroeconomic expectations system, combining sign restrictions with narrative identification. Departing from existing literature that studies expectations one outcome at a time, we provide a unified anatomy of expectation comovement while remaining agnostic about underlying microfoundations. Contrary to the conventional supply-side view, we show that a pessimism shock—a broad adverse shift in beliefs across macroeconomic outcomes—accounts for the dominant share of covariance in households expectations, extending well beyond output and inflation. We further map these belief shocks to observable heterogeneity, showing that pessimism is concentrated among individuals with low income, unemployment, deteriorating financial conditions, and poor health, who anticipate lower future earnings, higher job-loss and default risk, and plan to dissave.

Keywords: subjective expectations; belief heterogeneity; household expectations; disagreement; narrative and sign restrictions

*We are grateful for comments and guidance from George-Marios Angeletos, Lawrence Christiano, Martin Eichenbaum, Giorgio Primiceri, Matthew Rognlie, Alireza Tahbaz-Salehi. We are grateful for discussion with Aaron Amburgey.

†Northwestern University. Email: joselara@u.northwestern.edu.

‡Northwestern University. Email: r.zhang@u.northwestern.edu

1 Introduction

Core to the rational expectations revolution is the assumption that agents form expectations using a coherent understanding of the economy, so that beliefs are model-consistent, even in the presence of information frictions. A vast literature examines departures from this benchmark by analyzing the behavior of forecast errors and belief revisions, and more recently, how personal experiences shape expectations. However, existing approaches study expectations one outcome at a time (e.g., relating inflation forecast errors to revisions in inflation expectations or personal experiences). By focusing on these "marginal" relationships, the literature is largely silent about how beliefs about different macroeconomic outcomes are *jointly* determined, and therefore about the underlying belief system that governs these expectations.

In this paper, we close this gap by treating expectations as a belief system and develop a structural framework to identify the latent idiosyncratic forces driving joint variation in households expectations. A range of mechanisms-information frictions, heterogeneous subjective macro models, and narrative- or memory-based distortions- can generate disagreement. While they differ in microfoundations, they share a common reduced-form implication: heterogeneity shows up as systematic comovement across expectations, not merely isolated forecast errors. Our approach is deliberately agnostic about the source of dispersion and targets the reduced-form belief shocks implied by these comovement patterns.

We operationalize this idea by adapting SVAR-style identification logic, typically applied to realized macro time series, to the cross-section of beliefs obtained from survey data. In each survey month, households report a vector of expectations across multiple aggregate outcomes. We remove the common component shared within a month (and location) and focus on the remaining within-month disagreement. The covariance structure of this disagreement is our key empirical object: it summarizes how expectations co-move across households and serves as a cross-sectional analogue of SVAR contemporaneous restrictions. As in SVARs, it identifies the structure only up to a rotation, which we discipline using sign restrictions and narrative identification based on salient quasi-exogenous events.

Within this framework, we discipline the rotation by focusing on three latent dimensions of belief dispersion. The first two capture heterogeneity in how households interpret the source of macro fluctuations, which is naturally summarized by the perceived co-movement of real activity and prices: one dimension features negative co-movement ("demand-type" idiosyncratic interpretations) and another features positive co-movement

("supply-type" idiosyncratic interpretations). This structure is motivated by leading models of expectation dispersion, including incomplete information about aggregate shocks and heterogeneous subjective macro models that emphasize demand- versus supply-side channels. Finally, our reduced form evidence, together with existing literature (Cenzon, 2025; Taubinsky et al., 2025), points to an additional "pessimism" dimension: a broad adverse shift in beliefs across many outcomes in the same direction, capturing distortions not tied to a specific equilibrium channel and accounting for a large share of joint comovement in expectations.

We estimate the model using Bayesian methods and follow the narrative-identification strategy in Lagerborg et al. (2023), which leverages mass-shooting events. This helps us identify the mapping from households' expectations to three latent belief shocks (pessimism, demand, and supply). We use this mapping in three complementary exercises. First, we characterize the anatomy of pessimism by relating recovered shocks to demographics. Second, we validate the shock by projecting it onto a broad set of additional economy-wide expectations not used in the identification. Third, we project pessimism onto households expectations about their own outcomes. Taken together, these analyses show that pessimism is not a narrow artifact of the baseline system, but a broad component that organizes belief comovement across domains.

Our headline finding is that the pessimism shock, a broad adverse shift in beliefs across macroeconomic outcomes, accounts for the dominant share of covariance in households expectations, extending well beyond the comovement of output and inflation. Consistent with this interpretation, pessimism is concentrated among more vulnerable respondents and moves a wide set of additional macro expectations in an adverse direction out of sample. It also maps strongly into households idiosyncratic outlooks, worsening expectations about their own income, spending, and job security and raising perceived financial distress and default risk, with larger effects for lower-income households.

Taken together, these findings indicate that belief heterogeneity is not merely idiosyncratic noise scattered across unrelated survey answers. Instead, a small number of latent belief shocks structure how households form expectations across domains. In particular, pessimism emerges as a central force that organizes comovement across expectations in a common adverse direction.

Related Literature Our paper contributes to the large literature that studies how expectations departures from rational expectations by analyzing forecast errors and belief revisions, both at the individual and aggregate level (Coibion and Gorodnichenko, 2015; Kohlhas and Broer, 2023; Coibion et al., 2018, 2020; Bordalo et al., 2020; Angeletos et al.,

2021). The literature typically studies expectations one variable at a time, abstracting from how beliefs about different macroeconomic outcomes are jointly determined. Our paper builds on this literature by shifting attention from marginal forecast errors to the joint distribution of subjective macroeconomic expectations, asking which latent belief shocks govern comovement across macroeconomic forecasts.

Recent work emphasizes that households expectations are shaped not only by aggregate news but also by idiosyncratic experiences. A central finding is that individuals extrapolate from personal histories and salient price or labor-market experiences when forming beliefs about aggregate outcomes, contributing to persistent heterogeneity and systematic biases in reported expectations (Malmendier and Nagel, 2011, 2016; Kuchler and Zafar, 2019a; DAcunto et al., 2021; Cavallo et al., 2017; DAcunto et al., 2022; Cenzon, 2025). Particularly closely related, Taubinsky et al. (2025) document that beliefs about the aggregate economy are excessively sensitive to household-level shocks and propagate across belief domains, consistent with a pessimism channel in which "bad personal news" spills over into broader macro views. We complement this literature by providing a unified framework that separates equilibrium-specific belief dimensions (e.g., demand- versus supply-type interpretations) from broad pessimism-like shifts that move many expectations in the same direction; we show that such latent shock explains a significant share of the variances and covariances across macroeconomic expectations.

Methodologically, our paper brings SVAR-style identification tools—especially sign and narrative restrictions—to cross-sectional survey expectations by exploiting within-month belief covariances to recover interpretable latent "belief shocks." This approach is closely related to the broader SVAR tradition of identifying latent shocks under weak restrictions, including sign-restriction (Uhlig (2005); Rubio-Ramírez et al. (2010)), and the narrative-identification literature that uses external event information (Antolín-Díaz and Rubio-Ramírez, 2018; Romer and Romer, 2004; Ramey, 2011; Mertens and Ravn, 2013). Relative to these contributions, our key departure is that the relevant second moments are cross-sectional rather than time-series: instead of studying the autocovariance structure of realized macro outcomes, we study the covariance structure of subjective expectations within survey month. This shift allows us to study belief comovement directly and to provide an "anatomy" of disagreement—a decomposition into equilibrium-interpretation components (demand versus supply) and a broad pessimism component—while remaining agnostic about the underlying microfoundations. For narrative identification, we follow Lagerborg et al. (2023) that uses mass shooting events in the United States to help discipline the identification of sentiment shocks.

The rest of the paper is structured as follows. Section 2 presents the structural econometric

ric framework we use to analyze households' beliefs. Section 3 describes the SCE data and provides reduced-form evidence on the patterns governing the cross-sectional covariance structure of expectations. Section 4 lays out our identification strategy and explains how we incorporate sign restrictions and narrative evidence. Section 5 reports the structural results and analyzes the implied variance and covariance structure of beliefs. Section 6 studies the anatomy of the identified pessimism shock, relating it to demographics and showing how a one-unit increase shifts both economy-wide and idiosyncratic expectations in an adverse direction. Finally, Section 7 concludes.

2 Structural Econometric Framework

We introduce a structural econometric framework for subjective macroeconomic expectations and outline the identification strategy, which we discuss in detail below. The framework is agnostic about macroeconomic dynamics and does not require specifying a stochastic process for aggregate outcomes or a fully specified equilibrium model. It parallels the Structural Vector Autoregression (SVAR) framework, adapted to individual-level cross-sectional data, with identification based on the cross-sectional covariance structure of expectations.

Environment and structural representation. Fix a survey month t and index respondents by $i \in \{1, \ldots, N_t\}$. Let $B_{i,t} \in \mathbb{R}^G$ denote household i's reported expectations about G aggregate macroeconomic variables over a fixed horizon (e.g., inflation, unemployment, taxes). We represent beliefs as the sum of (i) a component driven by aggregate structural shocks that are common across households and may propagate over time, and (ii) an idiosyncratic component capturing heterogeneous information, interpretations, or belief distortions:

$$B_{i,t} \;=\; \bar{\Gamma}(L)\,\varepsilon_t \;+\; \Gamma\,\varepsilon_{i,t}. \tag{1}$$

Here $\varepsilon_t \in \mathbb{R}^S$ collects aggregate structural shocks, $\bar{\Gamma}(L)$ is a $G \times S$ one-sided lag polynomial mapping the history of aggregate shocks into contemporaneous beliefs, $\varepsilon_{i,t} \in \mathbb{R}^K$ collects idiosyncratic structural shocks, and $\Gamma \in \mathbb{R}^{G \times K}$ maps these idiosyncratic shocks into the cross-sectional dispersion of expectations within month t.

The lag structure in $\bar{\Gamma}(L)$ captures persistence in how common macroeconomic news maps into beliefs, reflecting gradual learning, external and internal shock propagation. By contrast, Γ governs the contemporaneous causal effects of individual-level structural innovations across belief variables. For tractability, we abstract from lagged idiosyncratic

shocks a natural restriction in short rotating panels, where respondents are observed only a handful of times. As a result, the structural shocks should be interpreted not only as innovations to an individuals beliefs, but also as absorbing observed and unobserved individual heterogeneity. We return to this interpretation below. In a separate ongoing work, we investigate persistent idiosyncratic components, however we also note here that most of the persistence in beliefs comes from permanent unobserved heterogeneity. [1]

Our focus is on the structural matrix Γ. The existing literature typically studies the dynamics of average expectations by regressing time series of mean survey forecasts on identified macro shocks (e.g., Bardóczy and Guerreiro (2023); Cai (2024)). In contrast, our goal is to identify the latent shocks that govern cross-sectional differences in beliefs. This distinction matters empirically: cross-sectional dispersion is large and, as we show later, much of its covariance structure is accounted for by a small number of latent shocks loading through Γ. Finally, we impose the following maintained conditions.

Assumption 1 (Moment and rank conditions)**.** *The shocks satisfy* $\mathbb{E}[\varepsilon_t] = 0$ *and* $\mathbb{E}[\varepsilon_{i,t}] = 0$. *The idiosyncratic shocks are cross-sectionally independent and have diagonal covariance,*

$$\mathbb{E}[\varepsilon_{i,t}\varepsilon_{i,t}^\top] \;=\; \Sigma_K, \qquad \Sigma_K = I_K$$

common across i. *The loading matrix* Γ *has full rank* K.

Assumption 1 normalizes scale and rules out redundant shocks while remaining deliberately agnostic about the microfoundations of belief formation. The resulting object can be interpreted as a reduced-form summary of the forces governing belief heterogeneity across individuals and macroeconomic variables. Almost all prominent theories of belief heterogeneity map naturally into our framework, and we leave interpretation through different theoretical lenses to the reader. The following two examples illustrates how such a mapping can be constructed.

Example 1 (Dispersed information with a common subjective model)**.** *Consider an economy in which otherwise homogeneous agents rationally extract information from dispersed signals but hold subjective macroeconomic models, so that their perceived impulse responses may differ from the true ones. Let there be a single i.i.d. aggregate shock* ε_t *each period. Each household observes a private signal*

$$s_{i,t} = \varepsilon_t + \varepsilon_{i,t},$$

[1] If we demean beliefs by State $\times$ Time and individual fixed effects, the autocorrelations of unemployment, inflation, and tax beliefs are not statistically different from zero.

where $\varepsilon_{i,t}$ is mean-zero idiosyncratic noise. In this setting, Γ identifies the agents subjective impulse responses of aggregate outcomes to the aggregate shock ε_t, up to a scale normalization.

In this example, belief heterogeneity arises from dispersed information. Because agents cannot perfectly distinguish idiosyncratic noise from the aggregate shock, their belief responses to idiosyncratic signal innovations reveal their subjective macroeconomic model, up to a scale normalization capturing the subjective Bayesian weight assigned to the idiosyncratic component of the signal.

Example 2 (Perfect Information with heterogeneous subjective model)**.** *Agents perfectly observe the aggregate structural shock, but their perceived impulse responses may differ from the true ones and can be heterogeneous across agents:*

$$B_{i,t}^{type}\Pi_t = IRF\,\varepsilon_t + IRF^{type}\,\varepsilon_{i,t}^{type}.$$

In this case, Γ captures cross-sectional heterogeneity in perceived impulse responses, with agent types indexed by $\varepsilon_{i,t}$.

In this example, Γ identifies heterogeneity in subjective macroeconomic models. This heterogeneity can then be related to observable characteristics—for instance, whether unemployed, lower—income, or less-educated individuals hold systematically different subjective models than others. More generally, Γ may capture distortions arising from attention, memory, or narrative-based information processing, and our framework is designed to accommodate all of these interpretations. Our preferred approach is to "let the data speak" and then identify the theoretical model that best accounts for the observed patterns, in the spirit of the Lucas Program.

Beyond the two environments above, we can also link our object of interest, Γ, to standard objects in workhorse macro models. For instance, in a linearized New Keynesian framework with noisy signals and heterogeneous subjective models we decompose heterogeneous subjective beliefs as a mixture of demand, supply, and pessimism latent components. The supply and demand components mean households have received supply or demand shocks information or have a supply or demand view. In the case of pessimism, it means that households have a pessimistic view of the economy.

Interpretation of idiosyncratic shocks Although we label the components of $\varepsilon_{i,t}$ as shocks, Examples 1 and 2 make clear that they can alternatively be interpreted as capturing either individual private information or agent types underlying heterogeneous subjective

macroeconomic models. These types may be systematically related to observable characteristics such as income, education, or employment status. Because our objective is to remain agnostic about the underlying source of heterogeneity and instead to first "let the data speak" by characterizing the covariance structure of subjective expectations, we do not seek to distinguish between these interpretations at this stage.

Later, we refer to these components as *demand*, *supply*, and *pessimism* shocks. This terminology should be understood broadly. First, dispersed information may lead agents to form different views of the economy—for example, by receiving signals that emphasize demand-side or supply-side forces (such as news about oil prices or aggregate demand conditions). Second, agents may hold intrinsically different beliefs about how the economy operates, including demand-side or supply-side views, or pessimistic narratives in which adverse macroeconomic outcomes tend to move together ("everything bad moves together").

Challenges in identification. To isolate the cross-sectional component, we take cross-sectional means in (1) and difference them out:

$$\bar{B}_t \;\equiv\; \mathbb{E}_i[B_{i,t}] \;=\; \bar{\Gamma}(L)\,\varepsilon_t, \qquad \tilde{B}_{i,t} \;\equiv\; B_{i,t} - \bar{B}_t \;=\; \Gamma\,\varepsilon_{i,t}. \tag{2}$$

The covariance matrix of demeaned beliefs is

$$\Sigma_{\tilde{B}} \;\equiv\; \mathbb{E}\left[\tilde{B}_{i,t}\tilde{B}_{i,t}^{\top}\right] \;=\; \Gamma\,\Gamma^{\top}, \tag{3}$$

which is directly estimable from the cross section. Equation (3) implies an inherent indeterminacy: for any orthogonal matrix $R \in \mathcal{O}(K)$,

$$\Gamma\,\Gamma^{\top} \;=\; (\Gamma R)\,(\Gamma R)^{\top}.$$

Thus, Γ is identified only up to an orthogonal rotation. We discuss our strategy for identifying Γ in Section 4. For now, we focus on characterizing the empirical patterns in the demeaned variables $\tilde{B}_{i,t}$.

3 Reduced Form Evidence

In order to provide further support for our findings and methods, we first present reduced-form evidence using the Survey of Consumer Expectations (SCE) data on expectations and household characteristics.

3.1 SCE Data Description

The SCE is a monthly survey conducted by the New York Fed that aims to measure households' expectations regarding macroeconomic conditions, inflation, spending, and income. It elicits probabilistic expectations and features both aggregate economy-wide and idiosyncratic expectations, as well as detailed household characteristics. Compared to other surveys, it covers a broader range of topics. The SCE consists of a rotating panel of approximately 1,300 households, including both new and returning respondents. Households stay in the survey for up to 12 months. The panel is representative of U.S. demographics and is based on the sampling scheme of the Consumer Confidence Survey. For a broader description of the survey, see Armantier et al. (2017).

In this paper, we focus mainly on the core monthly SCE survey, and our sample starts in 2013M6 and ends in 2024M4. Our main objects of interest are the questions regarding aggregate expectations. Specifically, we focus on the following economy-wide outcomes:

1. What do you expect the rate of inflation to be over the next 12 months?

2. What do you think is the percent chance that 12 months from now the unemployment rate in the U.S. will be higher than it is now?

3. Over the next 12 months, by about what percent do you expect the level of U.S. government debt to increase/decrease?

4. Suppose that 12 months from now, your household income is the same as it is now. By about what percent do you expect your total taxes to have increased/decreased, including federal, state, local income, property, and sales taxes?

We focus on these four expectation measures because they jointly provide a transparent reduced-form map from households' beliefs about aggregate conditions to the outcomes most relevant for perceived economic outlook and policy environment. Expected inflation and the unemployment-risk probability summarize households' views about near-term nominal and real conditions, respectively, while expectations about government debt and tax changes proxy for perceived fiscal policy adjustments. In addition to these economy-wide expectations, the survey also contains information on the probability of an increase in the stock market, changes in interest rates on savings accounts, long-run inflation expectations, and gasoline and food prices. Moreover, the survey also contains information on current and expected household income and spending.

3.2 Reduced-form regressions

In this section, we provide preliminary evidence on the partial associations between households' expectations and show, informally, that a natural pattern emerges: most of the comovements are consistent with pessimism, in the sense that expectations about adverse aggregate outcomes tend to move together. To explore this, we use households' idiosyncratic expectations regarding economy-wide aggregates. Specifically, we estimate the following regression:

$$\mathbb{E}_{i,t}(k) = \alpha + \sum_{m \neq k} \beta_m \, \mathbb{E}_{i,t}(m) + \gamma' Z_{i,t} + FE_{j(i),t} + \epsilon_{i,t}, \tag{4}$$

where $B_{i,t}(k)$ denotes the idiosyncratic expectation elicited from individual i at time t about the economy-wide variable k, and $B_{i,t}(m)$ denotes the idiosyncratic expectation about another aggregate variable $m \neq k$. For instance, when the left-hand-side variable k is inflation, the right-hand-side expectations include all aggregate expectations except inflation. The term $FE_{j(i),t}$ denotes state-by-time (State$\times$Month) fixed effects that control for local business-cycle conditions. Finally, $Z_{i,t}$ is a vector that includes individual i's time-invariant and time-varying demographic characteristics such as age, education, income, tenure in the panel, idiosyncratic employment and financial conditions, marital status, and a numeracy score. The latter measures respondents' ability to solve simple mathematical problems.

Under standard linear-regression logic, the coefficient β_m summarizes the partial correlation between beliefs about k and m, holding fixed demographics and state-by-time conditions. The inclusion of State$\times$Month fixed effects nets out common local business-cycle conditions and other shocks that could simultaneously shift all respondents' expectations in the same state and month. The residual variation left reflects cross-sectional heterogeneity and disagreement in expectations among households facing the same local conditons. Therefore, a positive[2] estimate of β_m indicates that households who are relatively more pessimistic about expectation m (relative to others in the same state-month) also tend to be relatively more pessimistic about expectation k. We hypothesize that a pattern of uniformly "adverse-direction" comovement is consistent with a common pessimism factor shaping household beliefs. In the Structural Identification section, we provide a framework that allows us to interpret in a cleaner way this pessimism factor.

In the baseline estimation, we consider economy-wide expectations regarding unemployment, taxes, inflation, and government debt. Table 1 contains the results. Each column

[2]Note that a higher value in these economy-wide expectations denotes a worse outcome for that variable.

denotes a different dependent variable (the left-hand-side expectation), whereas each row denotes a different right-hand-side expectation that has been standardized. In all specifications, we include the demographic controls $Z_{i,t}$ as well as the State×Month fixed effects.

We find evidence that these conditional comovements exhibit a pattern in which expectations move in a jointly adverse direction. For instance, the coefficient linking government-debt expectations to unemployment expectations implies that, conditional on demographics and local conditions, households who expect higher unemployment also tend to expect higher growth in government debt. Interpreting higher unemployment and higher debt growth as "worse" outcomes, this positive partial association is consistent with pessimistic co-movement in beliefs.

Table 1: Baseline: Reduced-Form Perceived Partial Covariances

	Dep. Var: $\mathbb{E}_{i,t}(k)$ **(Columns)**			
$\mathbb{E}_{i,t}(m)$ (Row)	Gov. Debt	Tax Changes	Inflation	Unemp.
Unemployment	1.2997***	0.3846***	0.5345***	
	(0.0644)	(0.0286)	(0.0349)	
Tax Changes	2.6089***		1.2365***	1.5407***
	(0.0779)		(0.0492)	(0.1140)
Inflation	4.1424***	1.1033***		1.9105***
	(0.0969)	(0.0423)		(0.1229)
Gov Debt		1.3354***	2.3765***	2.6650***
		(0.0405)	(0.0627)	(0.1311)
N	126118	126118	126118	126118
Demographic	yes	yes	yes	yes
Year-Month x State FE	yes	yes	yes	yes
Individual FE	no	no	no	no

Notes: Clustered standard errors by individual in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

In the appendix, Table A1 reports the same reduced-form specification augmented with individual fixed effects, which absorb time-invariant (permanent) differences in respondents' optimism/pessimism. The main finding is unchanged: the adverse-direction comovement in expectations remains present even after subtracting permanent traits. Additionally, Table A2 reports the same regression expanded to a broader set of economy-wide expectations. Overall, with a small number of exceptions, the estimated conditional associations continue to line up in the "pessimistic" direction.

The reduced-form regressions in eq. 4 can be interpreted as a max-share linear index of beliefs. Fix a left-hand-side belief $B_{i,t}(k)$ and let $x_{i,t}$ collect the remaining beliefs $B_{i,t}(m)$,

$m \neq k$, after partialling out demographics and State×Month fixed effects. Then OLS chooses the coefficient vector $\hat{\beta}$ so that the fitted value $\widehat{B}_{i,t}(k) = \hat{\beta}' x_{i,t}$ is the linear combination of the belief system that maximizes explained variation in $B_{i,t}(k)$ (equivalently, it is the best linear predictor of $B_{i,t}(k)$ from the span of the other beliefs). In this sense, the reduced-form "pessimism patterns" summarize how strongly each belief co-moves with the maximally predictive belief index for each $B_{i,t}(k)$. This interpretation strengthens the result that the regression evidence is consistent with a dominant common component that shifts many beliefs in a jointly adverse direction.

To complement this interpretation and to summarize comovement without conditioning on a particular left-hand-side variable, we residualize each belief by State×Month fixed effects and compute principal components of the resulting residual vector. PCA identifies orthogonal linear combinations of beliefs that sequentially maximize the variance explained in the belief system as a whole. In our baseline set of beliefs (unemployment, inflation, tax changes, and government debt), the first principal component[3] explains 43% of the residual cross-sectional variation and loads with the same adverse-direction sign across beliefs, consistent with a common "pessimism" factor. The second component explains 23% and exhibits a rotation in which inflation and unemployment move opposite to (or less strongly with) fiscal variables, a pattern that we interpret as closer to a demand-type dimension. Finally, the fourth component is consistent with a supply factor and it only explains 15%.

Table 2: Baseline: Principal Components

	Component		
Variable	Comp. 1	Comp. 2	Comp. 3
Unemployment	0.32	0.94	-0.01
Inflation	0.56	-0.24	-0.47
Tax Changes	0.50	-0.17	0.84
Gov. Debt	0.58	-0.15	-0.27
Eigenvalue	1.72	0.92	0.76
Variance Share	0.43	0.23	0.19

Income Heterogeneity. Households with different characteristics may have different mental models that map shocks into beliefs about different variables. The baseline reduced-form regression does not capture this; however, the survey allows us to explore heterogeneity in these associations. We focus on how the reduced-form relationships vary with

[3] If we only consider unemployment, inflation and taxes, the first principal component explain 46%.

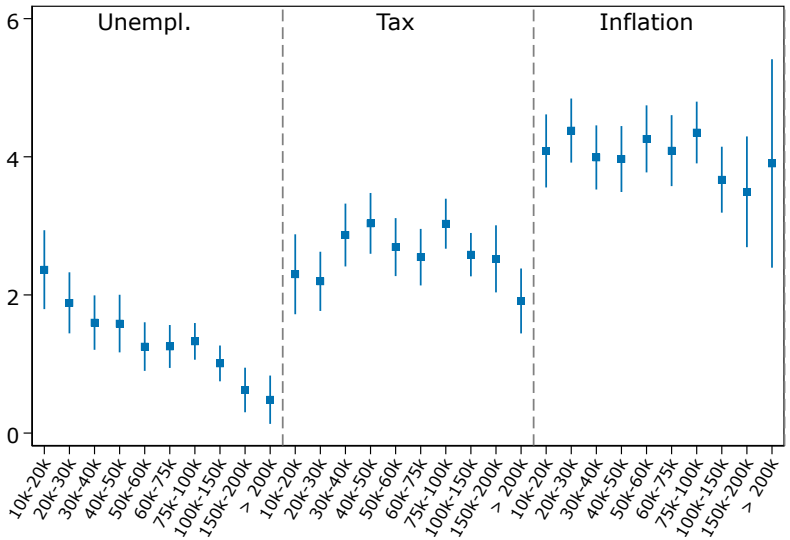
current household income. To be precise, we estimate:

$$B_{i,t}(k) = \alpha + \sum_{m \neq k} \beta_m^{(1)} B_{i,t}(m) + \sum_{m \neq k} \sum_{s=2}^{S} \beta_{m,s}^{(2)} B_{i,t}(m) \times \mathbb{1}\{Income_{i,t} = s\} + Z_{i,t}'\gamma + FE_{j(i),t} + \epsilon_{i,t}, \tag{5}$$

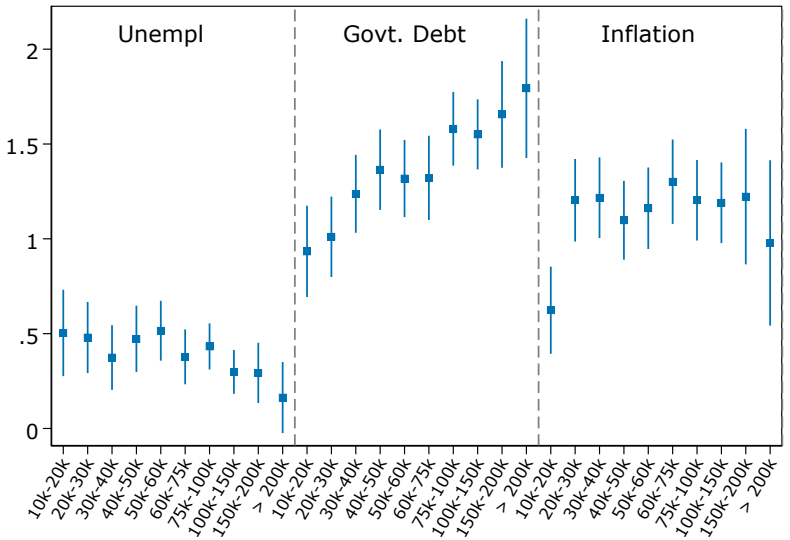
where $Income_{i,t}$ is a categorical variable that assigns household i at time t to one of S income groups, and $\mathbb{1}\{Income_{i,t} = s\}$ is an indicator for income group s. The object of interest is the income-group-specific slope relating expectation m to expectation k: for the reference group it is $\beta_m^{(1)}$, and for group s it is $\beta_m^{(1)} + \beta_{m,s}^{(2)}$.

For our baseline heterogeneity specification, Figure 1 summarizes the results. Each of the four panels fixes the dependent (left-hand-side) expectation, indicated in the panel caption. Within each panel, the dashed vertical lines separate blocks corresponding to different right-hand-side expectations. For each right-hand-side variable m, we plot the income-group-specific effect ($\beta_m^{(1)} + \beta_{m,s}^{(2)}$). Dots report point estimates and the vertical bars report 95% confidence intervals.

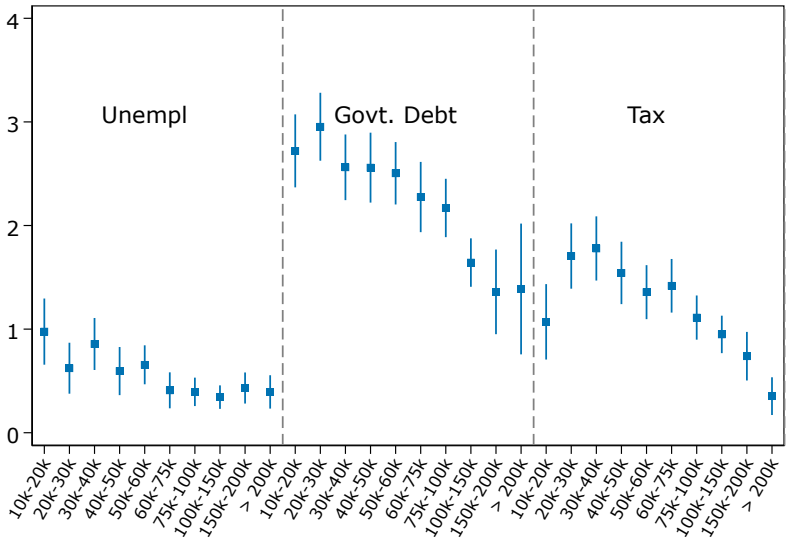
Across all panels, the estimated slopes are predominantly positive for every income group. In this sense, the figure provides reduced-form evidence consistent with our previous evidence regarding a common "pessimism" component. However, the strength of these associations varies systematically with income. In Panel (a), where the left-hand-side variable is expected government debt, the partial association with expected unemployment is strongest among lower-income households and declines with income, while the association with expected inflation remains large and positive across all income groups. In Panel (b), where the left-hand-side variable is expected tax changes, the association with government-debt expectations increases with income, whereas the association with unemployment expectations is positive but smaller and somewhat flatter. Panel (c) shows that when inflation is the left-hand-side variable, the partial associations with government debt and taxes are again positive but tend to attenuate at higher income levels. Finally, Panel (d) indicates that unemployment expectations comove positively with inflation, debt, and tax expectations across the income distribution, with some variation in magnitude across income categories. Overall, this shows household income heterogeneity only shifts the strength of the co-movement rather than the sign.



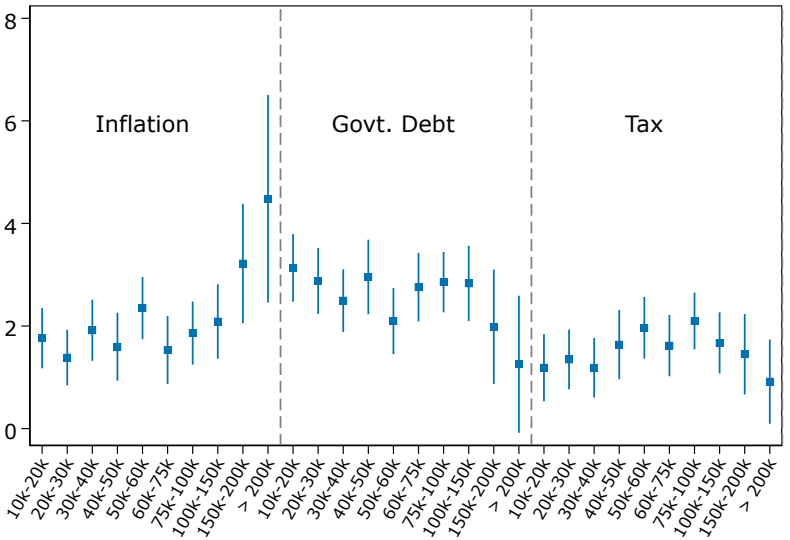
(a) Panel a): LHS Expectation $B_{i,t}(k)$ = Govt. Debt



(b) Panel b): LHS Expectation $B_{i,t}(k)$ = Tax



(c) Panel c): LHS Expectation $B_{i,t}(k)$ = Inflation



(d) Panel d): LHS Expectation $B_{i,t}(k)$ = Unemployment

Figure 1: Heterogeneous Partial Associations by Income

In addition to the estimations described above, we perform a battery of additional robustness checks. Specifically, we consider:

1. Splitting the sample into pre- and post-COVID periods
2. Varying the set of fixed effects and household controls
3. Estimating rolling-window regressions
4. Estimating nonparametric conditional expectation functions
5. Repeating the heterogeneity analysis using education, rather than income

All in all, we confirm that a pessimism component is present in the reduced form specifications.

4 Identification

We use sign and narrative restrictions for identification. Specifically, we postulate that three structural shocks drive agents beliefs: *demand*, *supply*, and *pessimism*. As previewed in Section 2, we adopt this categorization for three reasons. First, under the interpretation that belief dispersion arises from dispersed information, disagreement is naturally driven by heterogeneous private signals about demand- and supply-side conditions, motivating idiosyncratic demand and supply belief shocks. Second, under the interpretation of heterogeneous narratives, existing evidence—such as Coibion et al. (2018)—suggests that agents often hold predominantly "demand-side" or "supply-side" views of the economy, with these views varying across regions, time periods, and economic roles. Third, both our reduced-form evidence and the findings of Taubinsky et al. (2025) point to an additional dimension of belief heterogeneity that can be characterized as *optimism* versus *pessimism*, in which favorable or adverse macroeconomic outcomes tend to move together. Motivated by these considerations, we impose the following identification restrictions.

Sign restrictions. The demand shock is defined by a negative comovement between output and inflation, while the supply shock is defined by a positive comovement between output and inflation. The pessimism shock is defined as a broad shift in beliefs that moves multiple expectations in the same "adverse" direction, without being tied to a particular equilibrium channel.

To discipline the labeling, we impose sign restrictions on a subset of observables. Let (u, π, τ) denote unemployment, inflation, and taxes (all in expectation form, embedded

as components of $B_{i,t}$). We impose that the contemporaneous responses of these belief components to the three shocks satisfy

$$\begin{pmatrix} u \\ \pi \\ \tau \end{pmatrix} = \underbrace{\begin{pmatrix} + & + & + \\ + & - & + \\ - & ? & + \end{pmatrix}}_{\text{sign pattern}} \begin{pmatrix} \varepsilon^{\text{supply}} \\ \varepsilon^{\text{demand}} \\ \varepsilon^{\text{pess}} \end{pmatrix}. \tag{6}$$

The key role of taxes is to help distinguish supply-type interpretations from pessimism. Both a supply shock and a pessimism shock can raise unemployment and inflation expectations, but they differ in their fiscal implications. In particular, a supply shock is expected to lower expected taxes—consistent with a contractionary supply shock that depresses the tax base and necessitates fiscal transfers—whereas pessimism raises expected taxes. The demand shock is restricted to raise unemployment and taxes while lowering inflation.

The restrictions in (6) are intentionally modest, yet sufficient to rule out large classes of rotations. To further strengthen identification, we augment the sign restrictions with narrative restrictions.

Narrative restriction We incorporate a narrative restriction by exploiting the idea that certain salient events can shift pessimism directly ((Lagerborg et al., 2023)). We use mass-shooting exposure as such an event, building on evidence that exposure increases pessimism and worsens subjective assessments of economic conditions. The restriction is explicitly probabilistic rather than deterministic: exposure need not generate a sentiment shock for every treated individual, but instead shifts the distribution of the latent sentiment innovation in a pessimistic direction for a subset of households. Formally, let $T_{i,t} \in \{0, 1\}$ indicate that respondent i is exposed to a mass-shooting event in month t, defined as residing in the same commuting zone as the event and completing the survey after the event within the month. The narrative restriction requires that among treated households who revise unemployment expectations upward, at least 50% experience a negative realization of the sentiment innovation. We operationalize this requirement as a probabilistic inequality that must hold across admissible rotations.

The model is generically set-identified under sign and narrative restrictions. We therefore adopt a Bayesian approach that samples over the space of admissible rotations. The resulting posterior sets and probability bands reflect both sampling uncertainty and partial identification. Appendix **??** documents robustness to alternative probability cutoffs and to wider and narrower exposure windows across months for our narrative identification

scheme.

Table 3: Baseline: Effect of Mass-Shootings on Unemployment Forward Differences

	F_1	F_2	F_3	F_4	F_5	F_6
Treated Respondents	0.1233	1.1530**	1.4858**	1.5928*	1.8683*	1.7020
	(0.5131)	(0.5309)	(0.7417)	(0.8380)	(1.0375)	(1.1883)
Year-Month FE	yes	yes	yes	yes	yes	yes
Commuting Zone FE	yes	yes	yes	yes	yes	yes
N	1.28e+05	1.11e+05	9.63e+04	8.24e+04	6.95e+04	5.73e+04

Notes: Clustered standard errors by commuting zone in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

5 Structural Evidence

Impact Responses We now turn to the results from our structural framework, which are obtained from our identification scheme based on sign and narrative identification. We begin by reporting the contemporaneous impact responses of beliefs about unemployment, inflation, and taxes to the idiosyncratic structural shocks: pessimism, demand, and supply. While the identification restrictions discipline the signs of most responses, the magnitudes are left unrestricted and are therefore estimated from the data. Figure 2 summarizes the impact responses: each subfigure plots the contemporaneous response of the three belief measures to one of the structural shocks. Specifically, for each subplot we report the corresponding column of Γ as well as the 5th–95th percentile of draws that satisfy the sign restrictions.

Three patterns emerge. First, for unemployment beliefs, the largest response (in absolute value) is to a pessimism shock, followed by the demand and supply. Second, for inflation beliefs, the response to a pessimism shock is quantitatively comparable (in absolute value) to the response to a supply shock, while the response to a demand shock is the largest in magnitude. Third, for tax beliefs, the three shocks generate responses of similar absolute magnitude, but with opposite signs: the response is positive for a pessimism shock and negative for both demand and supply shocks.

Variance Decomposition. Next, we discuss the second-moment implications of the structural framework. Figure 3 reports the variance decomposition for each belief: unemployment, inflation, and taxes. The variance decomposition can be obtained from the impact matrix Γ and the covariance of demeaned beliefs so that the share of the variance

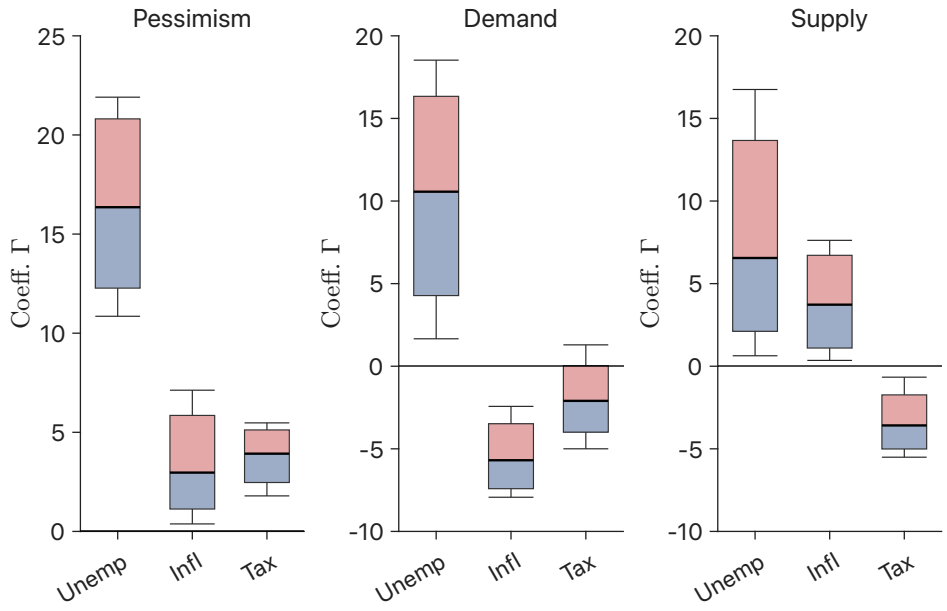


Figure 2: Contemporaneous Responses

of variable i attributable to shock k is as follows:

$$\omega_{i,k} = \frac{\Gamma_{i,k}^2}{\sum_k \Gamma_{i,k}^2}, \text{ since, } \Sigma_{\bar{B},ii} = \sum_k \Gamma_{i,k}^2 \tag{7}$$

Within each panel of Figure 3, we show the contribution of each structural shock to the variance of the corresponding belief variable. By construction, for any given draw and outcome variable, the contributions of the three shocks sum to one. We also report the 5th–95th percentile of draws.

The decompositions highlight clear differences across belief measures. For unemployment beliefs, pessimism shocks account for the largest share of the variance. For inflation beliefs, demand shocks are the dominant driver, followed by supply, with pessimism shocks contributing the smallest share. Finally, for tax beliefs, the three shocks contribute evenly; however, the median contribution of pessimism shocks is slightly larger than that of supply shocks, which in turn exceeds that of demand shocks.

Covariance Decomposition. In addition to decomposing the variance of each belief, we can also decompose the contribution of each structural shock k to the covariance between beliefs i and j. This covariance decomposition tightens the link between our structural and reduced-form evidence: it tells us which shocks account for the cross-sectional comovement in beliefs (the "pessimism" patterns in the reduced form) once we net out common aggregate components through fixed effects. More broadly, it provides a

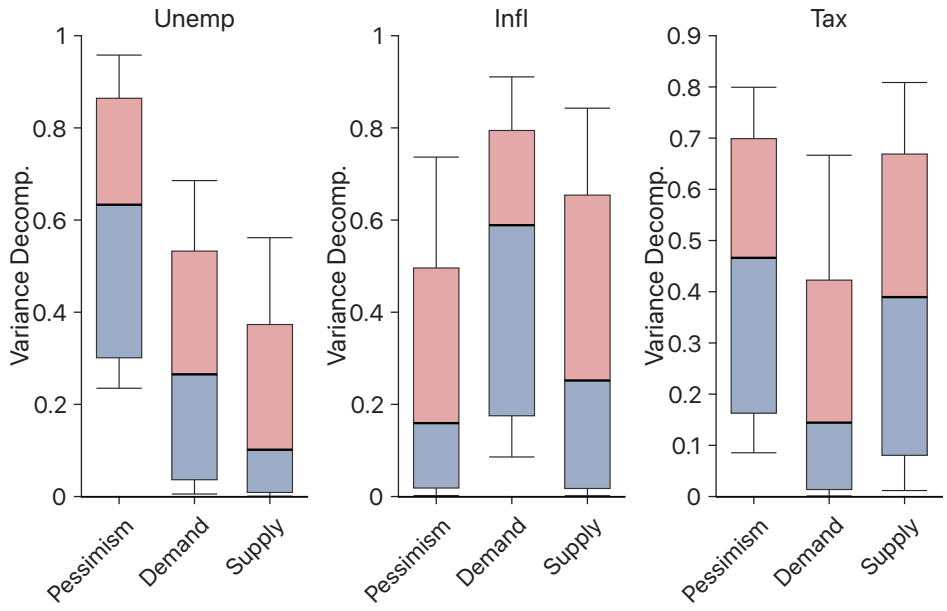


Figure 3: Variance Decomposition

disciplined way of answering the following question: when households disagree about the macroeconomy, what are they disagreeing about? Are belief correlations primarily driven by supply-side narratives, demand-side narratives, or a pessimism component that shifts multiple expectations in the same adverse direction?.

Formally, the share of the covariance between variables i and j attributable to shock k is

$$\omega_{ij,k} = \frac{\Gamma_{i,k}\Gamma_{j,k}}{\sum_{\ell}\Gamma_{i,\ell}\Gamma_{j,\ell}}, \qquad \text{where} \qquad \Sigma_{\bar{B},ij} = \sum_{\ell}\Gamma_{i,\ell}\Gamma_{j,\ell}. \tag{8}$$

By construction, $\sum_k \omega_{ij,k} = 1$ for any pair (i, j), so $\omega_{ij,k}$ provides a transparent accounting of which structural shocks generate the observed covariances. This expression is useful because covariances are the central reduced-form moments behind our "belief system" interpretation: they summarize how expectations move together across households facing the same aggregate environment.

Figure 4 reports the covariance decomposition for the belief system in our framework. Each panel corresponds to a different pairwise covariance, and within each panel we plot the contribution of pessimism, demand, and supply shocks as well as the 5th–95th percentile of draws.

We note the following. First, the sign of an observed covariance can be misleading about the underlying drivers: different shocks or narratives may push two beliefs in opposite directions, so the net covariance can be small or can inherit the sign of a residual component even when other forces are quantitatively important. In other words, a

"supply-like" covariance does not necessarily imply that households primarily hold supply narratives; it may instead reflect offsetting demand and pessimism components whose net effect happens to mimic a supply pattern. Second, the decomposition is informative for interpretation, communication, and forecasting. Policymakers and forecasters routinely track survey expectations as signals of households' perceived macro environment. Knowing whether comovement is driven by demand-type narratives, supply-type narratives, or broad pessimism changes what these survey movements mean: demand-driven disagreement is naturally linked to perceived labor-market strength and spending conditions; supply-driven disagreement reflects perceived cost and price pressures; and pessimism captures a broad reassessment that shifts multiple beliefs jointly. Even if the policy instrument is unchanged, these distinct narratives imply different expectation-management challenges and can generate different implications for how households interpret news about inflation, unemployment, and fiscal policy.

Turning to the results, the covariance between unemployment and inflation beliefs is driven primarily by pessimism and demand shocks, while the contribution of supply shocks is smaller. This finding complements the view emphasized in the literature that households often interpret inflation through a supply-side lens (e.g., Weber et al. (2022)).Our decomposition shows that, at the level of joint belief comovement, supply narratives are not the dominant quantitative driver. Instead, two additional forces matter: (i) a pessimism component that pushes unemployment and inflation beliefs in the same adverse direction across households, and (ii) a demand component that generates systematic comovement through perceived changes in real activity and prices. Importantly, demand-type and pessimism-type contributions can partially offset each other in the net covariance, so that the remaining covariance may appear "supply-like" in sign even when supply is not the leading contributor. This helps reconcile why reduced-form correlations can look consistent with supply narratives even though the underlying belief system contains sizable non-supply components.

The other two panels report the covariance decompositions for unemployment–tax beliefs and inflation–tax beliefs. In both cases, we again find that pessimism and demand shocks account for a large share of the covariance. For unemployment–tax beliefs, a large pessimism share is consistent with households who become more pessimistic, simultaneously expecting weaker labor-market outcomes and greater fiscal strain. For inflation–tax beliefs, the sizable demand contribution suggests that part of the comovement reflects a perceived aggregate "overheating"/"tightening" narrative, while the pessimism component captures a broader adverse reassessment that also raises perceived fiscal burdens. Overall, the covariance decompositions reinforce the central message of the paper: household ex-

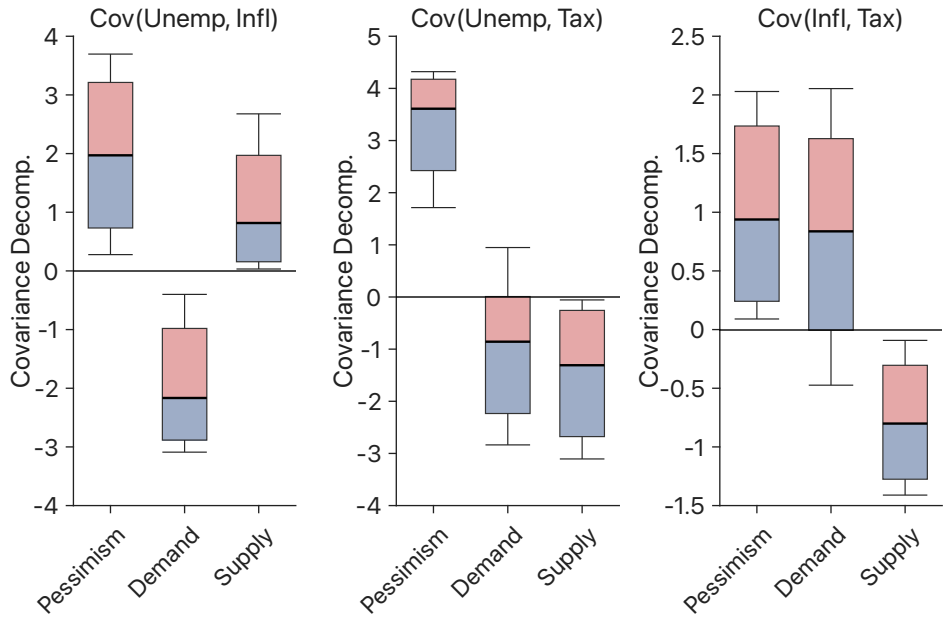


Figure 4: Covariance Decomposition

pectations are shaped by multiple forces beyond supply narratives, and the contributions of demand and pessimism are economically meaningful. Methodologically, the decomposition provides a new way of interpreting reduced-form belief correlations: it maps comovement moments into interpretable latent narratives, clarifying what households are "bundling together" when they update their macroeconomic views.

Permanent Unobserved Heterogeneity in Beliefs To assess whether time-invariant unobserved heterogeneity could affect our structural estimation, we reestimate the framework using beliefs that are demeaned by State$\times$Month fixed effects and respondent fixed effects. The results are reported in Figures B1, B2, and B3, respectively in Appendix B. Overall, our main findings continue to hold. Relative to the baseline, the impact responses change only modestly: tax beliefs become slightly positive for some realizations of demand shocks. The variance decomposition also changes: pessimism becomes more important for inflation and less important for taxes than in the baseline. Finally, the covariance decomposition yields conclusions that are unchanged and, if anything, stronger: pessimism plays a key role in explaining the covariance between unemployment and inflation, as well as between unemployment and taxes. Overall, these results show that controlling for permanent respondent components in beliefs does not alter our conclusions.

6 Anatomy of Expectations

The structural estimation described in the previous section serves two purposes. First, it allows us to recover the impact responses and the variance/covariance decompositions of beliefs to pessimism, demand, and supply shocks. Second, it allows us to construct household-level realizations of these shocks by inverting the mapping from demeaned reduced-form beliefs to the underlying structural shocks using the impact matrix Γ. Mathematically, letting $\tilde{B}_{i,t}$ collect the demeaned beliefs used in the estimation, we obtain

$$\epsilon_{i,t} = \Gamma^{-1}\tilde{B}_{i,t}. \tag{9}$$

The aforementioned step requires an additional condition: not only must Γ have full column rank, but it must be square (the number of belief series equals the number of shocks) so that Γ is invertible. Equivalently, we interpret $\epsilon_{i,t}$ as the unique vector of structural shocks that rationalizes the observed demeaned belief innovations $\tilde{B}_{i,t}$ on impact.

Once we have household-level pessimism, demand, and supply shocks that are constructed from demeaned beliefs about inflation, unemployment, and tax changes, we can study which types of households systematically load on each shock. Put differently, we analyze the anatomy of the recovered structural shocks by relating them to observable characteristics such as income, education, and other demographics. Ee estimate a regression framework in which the dependent variable is the pessimism shock and report the implied conditional means (marginal means) for each demographic category averaging across the rest of the covariates. Importantly, throughout we control for $State \times Time$ fixed effects and controls so that we capture deviations from the average in the households' state and month, and net of additional controls. To be specific we estimate the following regressions and calculate the marginal means for each demographic category and variable

$$Pessimism_{i,t} = \alpha + Z'_{i,t}\gamma + F_{j(i),t} + \epsilon_{i,t}, \tag{10}$$

where the demographic characteristics are embedded in $Z'_{i,t}$.

Figure 5 plots the predicted mean for each demographic category along with 90% and 95% confidence intervals (dark and light navy, respectively). The figure suggests that higher pessimism shocks are concentrated among households in more adverse circumstances. In particular, respondents with lower current income (under \$50k), those who report being financially worse off over the past year, and those who are not in full- or part-time employment exhibit higher predicted pessimism. Higher pessimism is also associated with being under 40 years old, reporting fair or poor health, and having lower

educational attainment (not completing college or postgraduate studies).

In summary, households facing adverse financial, labor-market, or health conditions tend to exhibit larger positive realizations of the pessimism component, consistent with existing evidence that household expectations vary systematically with personal circumstances, experiences, and socioeconomic characteristics (Kuchler and Zafar, 2019b; Coibion et al., 2022; DAcunto et al., 2022).

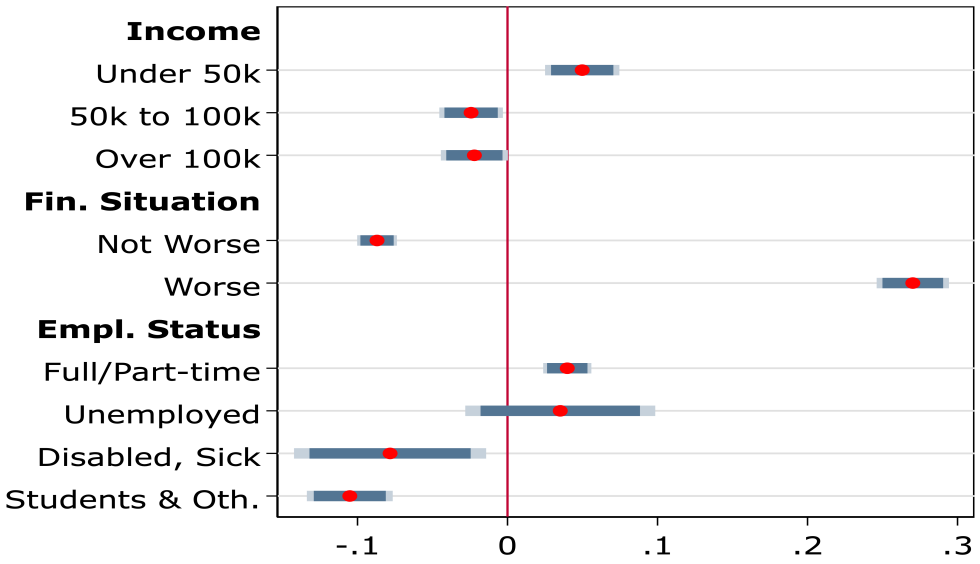
6.1 Projections on Aggregate Expectations

The structural framework only uses a subset of beliefs in the estimation, namely state-time demeaned beliefs regarding inflation, unemployment, and taxes. Nevertheless, the SCE elicits a much broader set of economy-wide expectations. For instance, SCE respondents report beliefs about government debt, interest rates on savings accounts, and a range of prices (e.g., food, housing, gold, and oil), as well as long-run inflation and inflation-uncertainty measures (e.g., the variance and interquartile range) constructed from respondents' density forecasts.

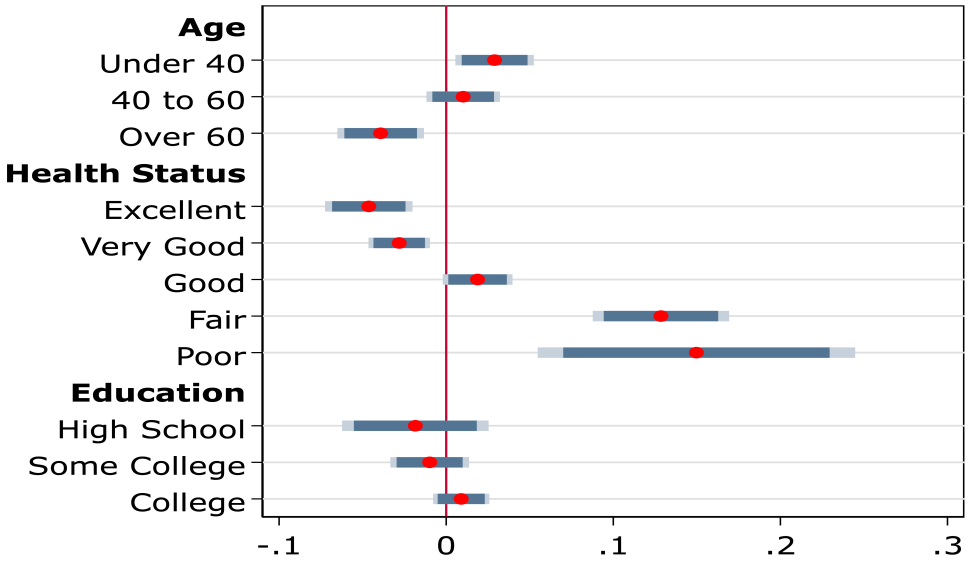
Ideally, we would like to include many of these additional variables directly in the structural estimation. However, doing so is not feasible without imposing substantially stronger identifying restrictions. Expanding the belief system would require specifying the contemporaneous sign patterns of a much larger system. In this paper, we intentionally remain as agnostic as possible about the joint structure of households' beliefs; we do not pursue this approach. Instead, we recover the three structural shocks (pessimism, demand, and supply) from the baseline structural framework and then project these shocks onto the broader set of economy-wide beliefs. This analysis both augments the structural estimation and provides an external validation check by asking how an identified pessimism shock moves other expectations that were not used in the identification. Thus, we estimate the following econometric specification to assess the impact of a pessimism shock on other beliefs:

$$B_{i,t}(k') = \alpha + \beta Pessimism_{i,t} + \gamma Z_{i,t} + FE_{j(i),t} + \epsilon_{i,t}, \tag{11}$$

where $B_{i,t}(k')$ denotes individual i's expectation at time t about economy-wide belief k' that is not included in the baseline structural framework. The covariates are defined analogously to eq.(4). The term $FE_{j(i),t}$ denotes State-by-Time (State$\times$Month) fixed effects that control for local business-cycle conditions, and $Z_{i,t}$ is a vector of individual demographic characteristics. Our coefficient of interest, β, measures the impact response of expectation k' to a one-unit realization of the pessimism shock.



(a) Panel a): By income, financial situation and employment status



(b) Panel b): By age, health status and education

Notes: This Figure plots the mean marginal predicted value by demographic characteristic and the corresponding 90% and 95% CI in dark and light navy blue, respectively.

Figure 5: Anatomy of Pessimism by Demographics

The results from the previous specification are reported in Figure 6. The figure plots the point estimate (red dot) and the corresponding 90% and 95% confidence intervals (dark and light navy, respectively), computed using clustered standard errors at the household level. The pessimism shock is standardized so that it raises unemployment beliefs by 10 percentage points on impact. Consistent with the interpretation of this shock as a broad increase in pessimism, this shock shifts most economy-wide expectations in an adverse direction. In particular, pessimism is associated with higher expected government debt (about 1%) and higher expected taxes (about **0**.6%). It also raises the probability that interest rates on savings accounts will increase (about 2.8 percentage points). Long-run inflation expectations increase by about 1.75%, and inflation uncertainty rises as measured by both the variance and the interquartile range. Households also expect tighter credit conditions following a pessimism shock. Finally, expected prices for gas, food, housing, and gold rise by roughly 0.2% to 0.6%. One exception is the reported probability of a stock-market increase, which rises following a pessimism shock. [4]

Overall, the broad adverse movements across external expectations provide support that our baseline structural estimation isolates a pessimism component in households' belief system, and they are consistent with the reduced-form evidence documented earlier.

6.2 Projections on Idiosyncratic Expectations

In addition to households' economy-wide beliefs, the SCE also asks about a rich set of idiosyncratic expectations about respondents' own outcomes, including their households' savings, income, and spending; job outcomes (e.g., the probability of accepting, losing, or leaving a job); financial conditions (e.g., default a loan and whether respondents expect to be financially worse off); respondents' personal earnings and measures of earnings uncertainty constructed from respondents' density forecasts. We use these additional outcomes to further assess the external validity of our identified pessimism shock. Accordingly, we estimate the following specification:

$$B_{i,t}(d) = \alpha + \beta\, Pessimism_{i,t} + Z'_{i,t}\gamma + FE_{j(i),t} + \epsilon_{i,t}, \tag{12}$$

where $B_{i,t}(d)$ denotes individual i's expectation at time t about idiosyncratic outcome d that is not included in the baseline structural system. The remaining covariates are

[4] In our specification, the pessimism shock increases the probability that the stock market will rise by about 1.5 percentage points. We hypothesize that respondents may answer the stock market question using a different mental model that places greater weight on short-run volatility and mean reversion, in which bad news might already be priced in.

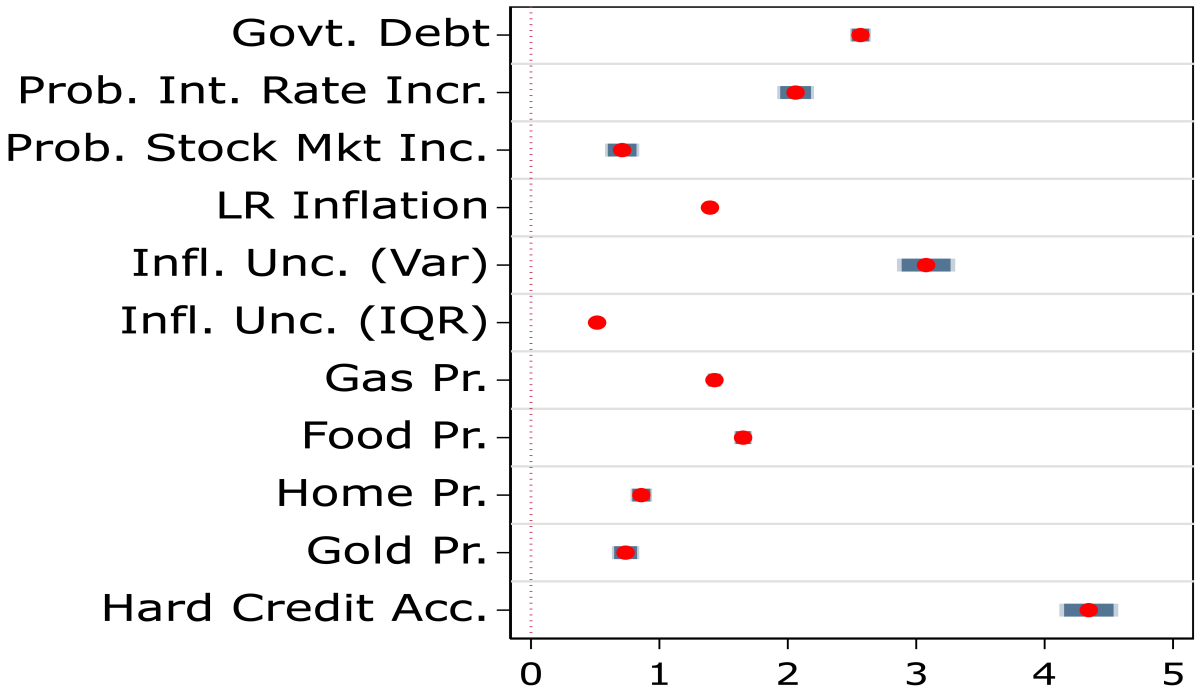


Figure 6: Response of External Economy-Wide Expectations to a Pessimism Shock

defined analogously to equation.(11). As before, the coefficient of interest β measures the impact response of idiosyncratic belief d to a one-unit realization of the pessimism shock, where the shock is normalized to raise economy-wide unemployment beliefs by 10 percentage points on impact. Conceptually, (12) is an auxiliary shock projection in which we recover the structural shocks from the baseline belief system and then project them onto idiosyncratic outcomes not used for identification.

Figure 7 plots the estimated β coefficients (red dots) along with 90% and 95% confidence intervals (dark and light navy, respectively). Consistent with the aggregate projections, the pessimism shock shifts most idiosyncratic expectations in an adverse direction. In particular, greater pessimism is associated with a decrease in expected household savings, spending, and income by about 4% to 5.2%, and with a decline in expected personal earnings of about 0.4%. In addition, the effect on earnings uncertainty is not statistically different from zero. The probability of job outcomes is also affected; the probability of losing or leaving ones job increases by roughly 1 percentage point, whereas the probability of accepting a new job offer is close to zero. Turning to financial outcomes, the probability of defaulting on a loan within the next three months and the probability of being financially worse off over the next year increase by about 0.8 and 2 percentage points, respectively.

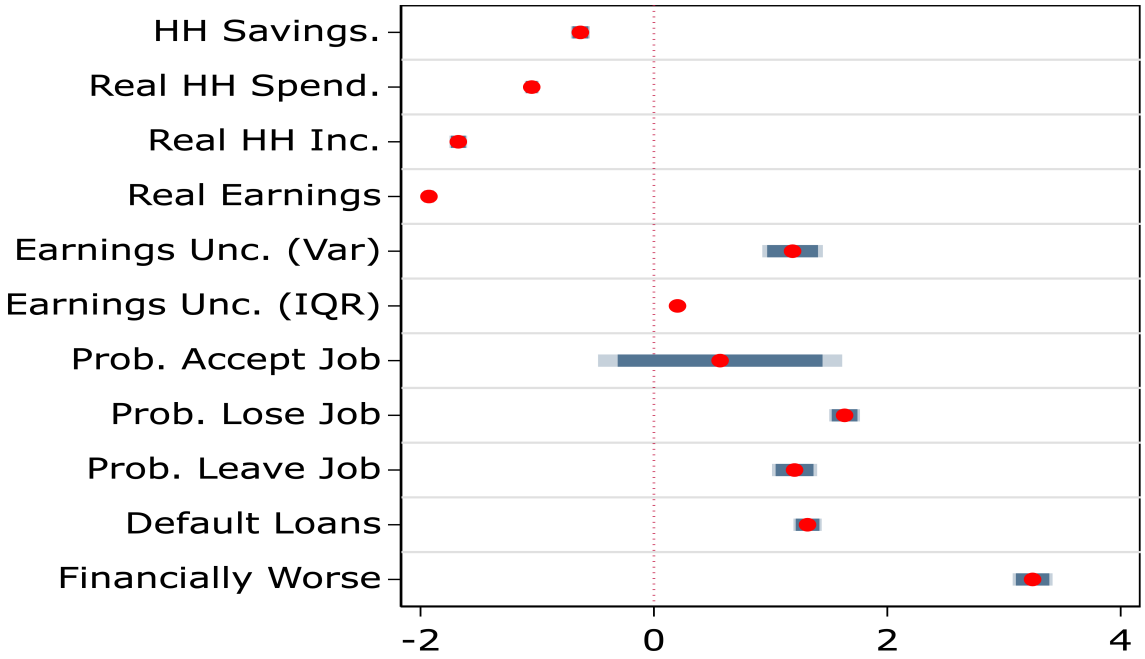


Figure 7: Response of External Idiosyncratic Expectation to Pessimism Shock

Heterogeneity by income. Because beliefs about spending, income, earnings, and savings may vary mechanically with current income, we also interact the pessimism shock with current household-income buckets in equation (12). Figure 8 displays the implied income-group-specific β. We find that for expected household income, spending, and personal earnings, the impact of a pessimism shock is attenuated with income, suggesting that lower-income households exhibit more extreme belief responses. For expected savings, the estimated effects are broadly similar across income groups, except at the top of the income distribution (above \$150k), where savings expectations decline more strongly following a pessimism shock.

Overall, projecting the identified pessimism shock onto idiosyncratic expectations provides additional evidence that this shock captures a broad deterioration in household sentiment, with economically meaningful effects on beliefs about personal income, spending, labor-market risk, and financial situation.

7 Conclusion

This paper studies macroeconomic expectations through the lens of the joint distribution of beliefs. Using the Survey of Consumer Expectations, we document that household



Figure 8: Response of External Idiosyncratic Expectation to Pessimism Shock by Income

expectations about unemployment, inflation, and fiscal variables exhibit systematic cross-sectional comovement that is consistent with a broad "pessimism" component. We then develop a structural framework that identifies a small set of latent belief shocks—demand, supply, and pessimism—from the covariance structure of demeaned expectations, while remaining agnostic about the microfoundations of belief formation. The structural results show that these shocks play distinct roles: pessimism is quantitatively central for unemployment beliefs, demand is the dominant driver of inflation beliefs, and tax beliefs load more evenly across shocks. Covariance decompositions further indicate that the comovement across beliefs is not explained by supply narratives alone, but reflects economically meaningful contributions from demand and pessimism components.

We also provide evidence on the external validity and economic content of the identified pessimism shock. Demographic projections show that higher pessimism realizations are concentrated among more vulnerable households—those with lower income, weaker recent financial conditions, weaker labor-market attachment, worse health, and lower educational attainment. Projections onto additional economy-wide expectations not used for identification show broad adverse movements: pessimism raises expected government debt and taxes beliefs, increases expected interest-rate hikes on savings accounts, lifts long-run inflation expectations and inflation uncertainty, tightens expected credit access, and raises expected prices for several key goods. Finally, projecting pessimism onto

idiosyncratic outcomes reveals negative effects on expected household income, spending, and savings, as well as higher perceived job-loss/exit risk and financial distress, with effects generally stronger for lower-income households. Taken together, these results suggest that a small number of latent belief shocks—and especially a pessimism component—organize belief heterogeneity across households.

References

Angeletos, G.-M., Huo, Z., and Sastry, K. A. (2021). Imperfect macroeconomic expectations: Evidence and theory. *NBER Macroeconomics Annual*, 35(1):1–86.

Antolín-Díaz, J. and Rubio-Ramírez, J. F. (2018). Narrative sign restrictions for svars. *American Economic Review*.

Armantier, O., Topa, G., Van der Klaauw, W., and Zafar, B. (2017). An overview of the survey of consumer expectations. *Economic Policy Review*, (23-2):51–72.

Bardóczy, B. and Guerreiro, J. (2023). Unemployment insurance in macroeconomic stabilization with imperfect expectations. *Working Paper*.

Bordalo, P., Gennaioli, N., Ma, Y., and Shleifer, A. (2020). Overreaction in macroeconomic expectations. *American Economic Review*, 110(9):2748–2782.

Cai, M. (2024). Explaining the macroeconomic inertia puzzle. *Working Paper*.

Cavallo, A., Cruces, G., and Perez-Truglia, R. (2017). Inflation expectations, learning, and supermarket prices: Evidence from survey experiments. *American Economic Journal: Macroeconomics*, 9(3):1–35.

Cenzon, J. (2025). Credit market experiences and macroeconomic expectations: Evidence and theory.

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

Coibion, O., Gorodnichenko, Y., and Kumar, S. (2018). How do firms form their expectations? new survey evidence. *American Economic Review*, 108(9):2671–2713.

Coibion, O., Gorodnichenko, Y., and Ropele, T. (2020). Inflation expectations and firm decisions: New causal evidence. *Quarterly Journal of Economics*, 135(1):165–219.

Coibion, O., Gorodnichenko, Y., and Weber, M. (2022). Monetary policy communications and their effects on household inflation expectations. *Journal of Political Economy*, 130(6):1537–1584.

DAcunto, F., Malmendier, U., Ospina, J., and Weber, M. (2021). Exposure to grocery prices and inflation expectations. *Journal of Political Economy*, 129(5):1615–1639.

DAcunto, F., Malmendier, U., and Weber, M. (2022). What Do the Data Tell Us About Inflation Expectations? *NBER*.

Kohlhas, A. and Broer, T. (2023). Forecaster (mis-)behavior. *Review of Economics and Statistics*.

Kuchler, T. and Zafar, B. (2019a). Personal experiences and expectations about aggregate outcomes. *The Journal of Finance*, 74(5):2491–2542.

Kuchler, T. and Zafar, B. (2019b). Personal Experiences and Expectations about Aggregate Outcomes. *The Journal of Finance*, 74(5):2491–2542.

Lagerborg, A., Pappa, E., and Ravn, M. O. (2023). Sentimental business cycles. *The Review of Economic Studies*, 90(3):1358–1393.

Malmendier, U. and Nagel, S. (2011). Depression babies: Do macroeconomic experiences affect risk taking? *The quarterly journal of economics*, 126(1):373–416.

Malmendier, U. and Nagel, S. (2016). Learning from inflation experiences. *The Quarterly Journal of Economics*, 131(1):53–87.

Mertens, K. and Ravn, M. O. (2013). The dynamic effects of personal and corporate income tax changes in the united states. *American Economic Review*.

Ramey, V. A. (2011). Identifying government spending shocks: It's all in the timing. *Quarterly Journal of Economics*.

Romer, C. D. and Romer, D. H. (2004). A new measure of monetary shocks: Derivation and implications. *American Economic Review*.

Rubio-Ramírez, J. F., Waggoner, D. F., and Zha, T. (2010). Structural vector autoregressions: Theory of identification and algorithms for inference. *Review of Economic Studies*.

Taubinsky, D., Butera, L., Saccarola, M., and Lian, C. (2025). Beliefs about the economy are excessively sensitive to household-level shocks: Evidence from linked survey and administrative data. Technical report, National Bureau of Economic Research.

Uhlig, H. (2005). What are the effects of monetary policy on output? results from an agnostic identification procedure. *Journal of Monetary Economics*.

Weber, M., dAcunto, F., Gorodnichenko, Y., and Coibion, O. (2022). The subjective inflation expectations of households and firms: Measurement, determinants, and implications. *Journal of Economic Perspectives*, 36(3):157–184.

Appendix A: Reduced Form Evidence

Table A1: Baseline: Reduced-Form Perceived Partial Covariances with Indiv. FE.

	Dep. Var: $\mathbb{E}_{i,t}(k)$ **(Columns)**			
$\mathbb{E}_{i,t}(k')$ (Row)	Gov. Debt	Tax Changes	Inflation	Unemp.
Unemployment	0.6996***	0.1540***	0.5247***	
	(0.0463)	(0.0228)	(0.0316)	
Tax Changes	1.2569***		0.5490***	0.5224***
	(0.0554)		(0.0383)	(0.0773)
Inflation	1.4753***	0.5062***		1.6407***
	(0.0714)	(0.0353)		(0.0983)
Gov Debt		0.8341***	1.0616***	1.5744***
		(0.0362)	(0.0517)	(0.1033)
N	124649	124649	124649	124649
Demographic	yes	yes	yes	yes
Year-Month x State FE	yes	yes	yes	yes
Individual FE	yes	yes	yes	yes

Notes: Clustered standard errors by individual in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A2: Reduced-Form Perceived Partial Covariances

	Dep. Var: $\mathbb{E}_{i,t}(k)$ **(Columns)**									
$\mathbb{E}_{i,t}(k')$ (Row)	Govt. D.	Tax	Infl	Unemp	Int. Rate	Stocks	LR Infl	Hard Cred.	Food	Gas
Unemp	1.1812***	0.2976***	0.2546***		5.0489***	2.3079***	0.0868***	0.0661***	0.2245***	0.4110***
	(0.0642)	(0.0291)	(0.0273)		(0.1329)	(0.1188)	(0.0287)	(0.0024)	(0.0261)	(0.0412)
Tax	1.6591***		0.5250***	1.1118***	-0.1763	-0.4684***	0.2966***	0.0152***	0.6930***	0.5010***
	(0.0728)		(0.0382)	(0.1085)	(0.1139)	(0.1004)	(0.0405)	(0.0022)	(0.0313)	(0.0428)
Infl	2.2675***	0.6076***		1.1007***	0.4089***	-0.5000***	4.1869***	0.0090***	1.0592***	-0.0168
	(0.0951)	(0.0428)		(0.1168)	(0.1104)	(0.1072)	(0.0773)	(0.0025)	(0.0407)	(0.0488)
Int. Rate	-0.8809***	-0.0480	0.0962***	5.1336***		8.8324***	-0.0490*	-0.0020	-0.1656***	0.2606***
	(0.0620)	(0.0310)	(0.0260)	(0.1318)		(0.1230)	(0.0284)	(0.0025)	(0.0270)	(0.0464)
Stocks	-0.3237***	-0.1382***	-0.1274***	2.5434***	9.5729***		-0.0364	-0.0530***	-0.1805***	-0.0244
	(0.0627)	(0.0296)	(0.0272)	(0.1308)	(0.1325)		(0.0295)	(0.0025)	(0.0266)	(0.0432)
LR Infl	1.1301***	0.2990***	3.6473***	0.3269***	-0.1815*	-0.1243		0.0107***	0.4443***	0.2591***
	(0.0893)	(0.0410)	(0.0730)	(0.1079)	(0.1051)	(0.1008)		(0.0024)	(0.0386)	(0.0485)
Hard Cred	0.6222***	0.3817***	0.1964***	6.2163***	-0.1875	-4.5204***	0.2665***		0.6673***	-0.2477***
	(0.1257)	(0.0546)	(0.0551)	(0.2265)	(0.2345)	(0.2104)	(0.0614)		(0.0508)	(0.0747)
Gas	0.6767***	0.3913***	-0.0113	1.1991***	0.7478***	-0.0647	0.2007***	-0.0077***	2.6874***	
	(0.0721)	(0.0335)	(0.0329)	(0.1206)	(0.1334)	(0.1144)	(0.0378)	(0.0023)	(0.0382)	
Food	3.2050***	0.9382***	1.2390***	1.1354***	-0.8238***	-0.8286***	0.5966***	0.0359***		4.6586***
	(0.0904)	(0.0421)	(0.0496)	(0.1316)	(0.1339)	(0.1214)	(0.0516)	(0.0027)		(0.0545)
Govt. D		0.8991***	1.0617***	2.3913***	-1.7540***	-0.5947***	0.6075***	0.0134***	1.2829***	0.4696***
		(0.0400)	(0.0495)	(0.1291)	(0.1228)	(0.1152)	(0.0477)	(0.0027)	(0.0376)	(0.0500)
N	125642	125642	125642	125642	125642	125642	125642	125642	125642	125642
Demog.	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
YM x St. FE	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes

Notes: Clustered standard errors by individual in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Appendix B: Permanent Unobserved Heterogeneity

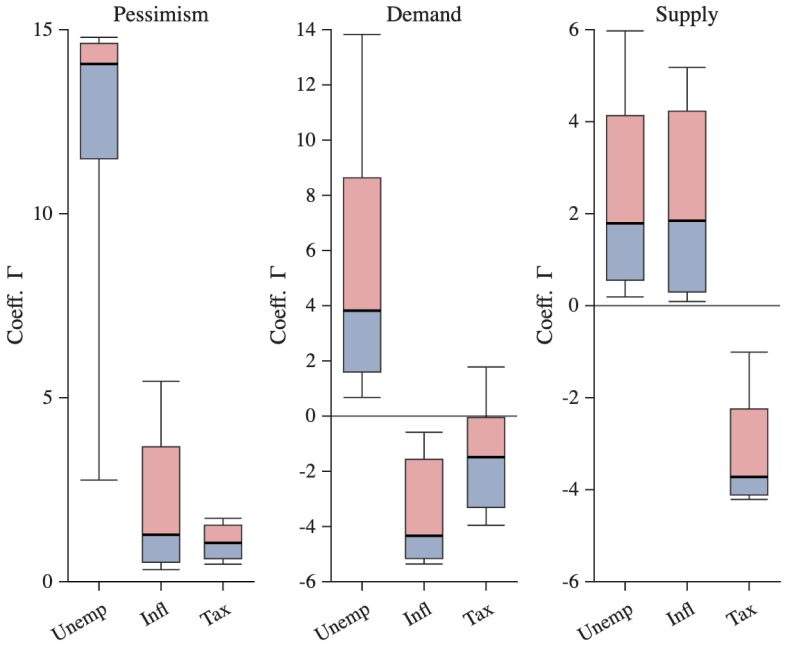


Figure B1: Contemporaneous Responses

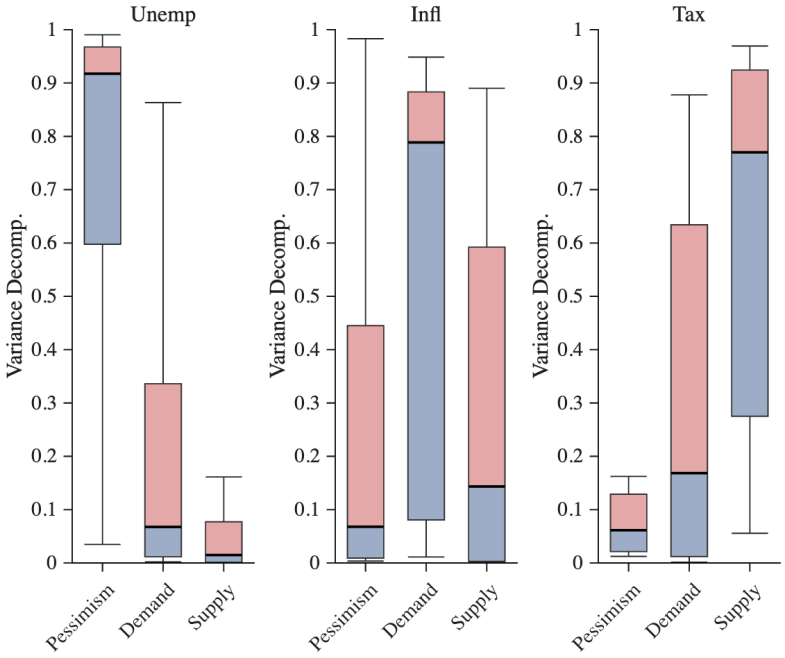


Figure B2: Variance Decomposition

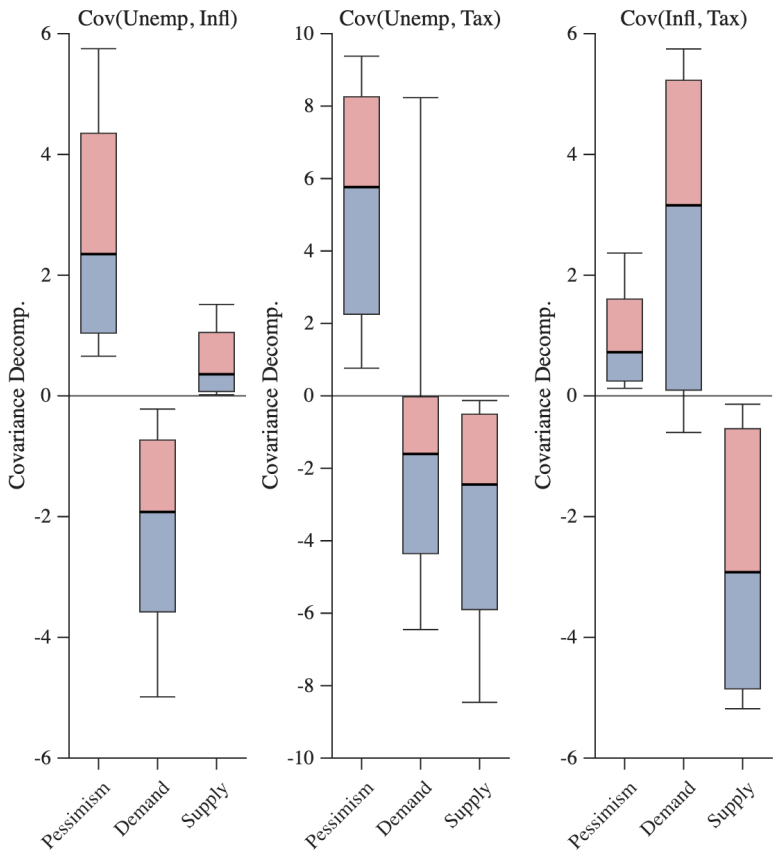


Figure B3: Covariance Decomposition