Anatomy of Subjective Macroeconomic Expectations*

Jose Lara† and Dalton Rongxuan Zhang‡

First Version: Jan 2026
This Version: July 2026

Abstract

We develop a structural approach to identify latent components in households' macroeconomic beliefs. We use the covariance structure of cross-sectional disagreement across macroeconomic expectations, combining sign and narrative restrictions based on mass-shooting exposure. This identifies pessimism, demand, and supply components and provides a unified anatomy of expectation comovement while remaining agnostic about whether it reflects incomplete information or subjective views. Pessimism, a broad adverse shift in beliefs, accounts for a large share of covariance alongside supply-side views. It is concentrated among individuals with deteriorating finances, volatile income, and poor health, who anticipate greater earnings uncertainty and higher job-loss and default risk. The supply component is more common among older individuals with greater lifetime inflation experience; the demand component is more common among higher-income, more-educated individuals. Overall, the results suggest that studying expectation comovements, moving beyond isolated forecast errors, is important for theories of expectation formation.

*We are grateful for guidance from George-Marios Angeletos, Lawrence Christiano, Martin Eichenbaum, Diego Känzig, Giorgio Primiceri, Matthew Rognlie and Alireza Tahbaz-Salehi. We are grateful for comments and discussion with Aaron Amburgey, Michael Cai, Yuriy Gorodnichenko, Jacob Toner Gosselin, Evan Majic, Pooya Molavi, Alessandro Pavan, Ramya Raghavan, Takaaki Sagawa,Kunal Sangani and seminar participants at NBER 2026 Summer Institute, the Econometric Society North American Summer Meeting 2026, the Spring 2026 Midwest Macro Meeting and Northwestern University Macro Lunch.

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

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

1 Introduction

Households form expectations about multiple macroeconomic outcomes at the same time, and these expectations need not vary independently. Instead, a household's beliefs about unemployment, inflation, taxes, and other outcomes may move together in systematic ways. For example, under rational expectations, a perceived demand shock raises expected unemployment and lowers expected inflation. Beliefs may also comove because households receive different information or hold different subjective views of the economy, perhaps shaped by narratives or memories (e.g. people that have experienced high inflation periods in their lifetimes). Although these mechanisms differ in their microfoundations, they share a common implication: belief heterogeneity appears as systematic joint comovement across expectations, rather than only as isolated forecast errors.

In this paper, we explicitly study this joint comovement of expectations by treating expectations as a belief system and developing a structural framework to identify the latent forces driving joint variation in households' expectation. This framework resembles SVAR-style identification but applied to micro data. Our goal is to decompose this joint variation into latent belief components and then study which households hold these views in terms of demographics, how much of the variation and comovement in household expectations the components explain, and how the views map into other expectations. We first remove the component of expectations common to agents in the same month and location. The remaining variation is disagreement across households facing same aggregate and local conditions. The covariance matrix of these residuals records how beliefs move together across agents, such as whether households with high expected unemployment also expect high inflation or high taxes.

To interpret this covariance, we posit a structural relation between latent belief components and observed expectations. Let $\tilde{B}_{it}$ denote the belief vector after removing the component common to households in the same month and location. We write residualized expectations as

$$\tilde{B}_{it} = \Gamma \epsilon_{it},$$

where ϵ_{it} collects household-level scores for the latent belief components, and Γ maps these components into observed beliefs.[1] As in the SVAR identification literature, the covariance matrix identifies Γ only up to an orthogonal rotation. Economic restrictions are therefore needed to distinguish the structural components underlying comovement in the belief system.

[1] Under the normalization $\mathbb{E}[\epsilon_{it}\epsilon_{it}'] = I$, the covariance matrix satisfies $\Sigma_{\tilde{B}} = \Gamma\Gamma'$.

We discipline these rotations and assign them an economic interpretation by adapting SVAR-style identification methods, typically applied to macroeconomic time series, to the cross-section of beliefs observed in the data. Our baseline treats residualized household-month observations as independent Gaussian draws and combines sign restrictions with narrative identification based on salient quasi-exogenous events to identify the latent components.

We focus on three latent dimensions of belief dispersion: demand, supply, and pessimism. The demand component captures views under which expected unemployment rises and expected inflation falls. The supply component captures views under which expected unemployment and inflation rise together, while expected taxes fall. These two components are motivated by leading models of expectation dispersion, including incomplete information about aggregate shocks and heterogeneous subjective views that emphasize demand- or supply-side channels, and are grounded in a New Keynesian framework. The third component is pessimism: a broad adverse shift in beliefs across outcomes, motivated by our reduced-form evidence and recent work on belief spillovers and pessimistic narratives (Cenzon, 2026; Taubinsky et al., 2024).

Supply and pessimism differ on their fiscal implications and we use tax expectations to distinguish between them. Under a supply view, higher expected unemployment is associated with lower expected taxes, consistent with a countercyclical fiscal response and the operation of automatic stabilizers. Under pessimism, by contrast, higher expected taxes form part of a broad adverse shift in beliefs. We impose sign restrictions on the structural mapping that reflect these patterns. To further discipline the pessimism component, we use the narrative event studied by Lagerborg et al. (2023), based on exposure to mass-shooting events. Specifically, we posit that following a mass shooting, at least some exposed households, though not necessarily all households in the affected area, hold more pessimistic views.

The recovered mapping between observed beliefs and latent belief components allows us to provide an anatomy of household expectations. First, we decompose the variance and covariance of expectations across the latent components. Second, we relate the components to observable heterogeneity, specifically we regress household-level component scores on demographics, experiences, and spending patterns. Third, we project these components onto additional economy-wide expectations not used in their identification and onto households' expectations about their own outcomes. Finally, we aggregate the pessimism component into a monthly time-series index.

We find that pessimism accounts for a large part of belief comovement. In the covariance decomposition, pessimism makes the largest positive contribution to the covariance be-

tween unemployment and inflation expectations, so what appears to be a supply-side view, in large part, also reflects broader pessimism. In the variance decomposition, pessimism accounts for the largest share of the variance in unemployment expectations. Pessimism is concentrated among households in fragile economic circumstances: households with volatile incomes, deteriorating financial conditions, poor health, weak labor-market positions, and adverse lifetime labor-market experiences. Higher pessimism is associated with worse expectations, including tighter credit conditions, higher job-loss risk, higher default risk, greater earnings uncertainty, and worse personal financial prospects.

Demand views account for the largest share of the variance in inflation expectations. They are more common among households with higher education and income. Supply views account for part of the positive comovement between expected unemployment and expected inflation, but differ from pessimism because they imply lower expected taxes. They are more common among older households and households with greater lifetime inflation exposure. Importantly, household characteristics associated with pessimism, such as deteriorating financial conditions, health status, weak labor market positions, and volatile income, are not associated with demand or supply views. Table 1 summarizes the resulting anatomy.

Table 1: Anatomy of Expectations: Latent Belief Components

Component	Who holds the view	Other expectations
Pessimism	Poor financial conditions; poor health; volatile income; worse lifetime (aggregate) unemployment experience	Worse macro outlook; higher job-loss default and earnings risk
Demand	Highly educated; higher income	Lower expected taxes
Supply	Older; unemployed or not in labor force; greater lifetime inflation	Higher long-run inflation

These findings show that belief heterogeneity is structured across expectations. A small number of latent belief components organize how households form expectations across domains, with pessimism playing the central role. The broader contribution is to make belief comovement empirically tractable, providing concrete targets for future research on the microfoundations of belief formation and the macroeconomic consequences of heterogeneous expectations.

Related Literature Our paper contributes to the large literature that studies 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; Andre et al., 2022; Ferreira and Pica, 2024; Granziera et al., 2025). The literature typically studies expectations one variable at a time. 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 components govern comovement across macroeconomic forecasts.

Recent work emphasizes that households' expectations are shaped not only by aggregate news but also by idiosyncratic experiences and psychological distortions. 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, 2019; D'Acunto et al., 2021; Cavallo et al., 2017; D'Acunto et al., 2022; Bhandari et al., 2023; Cenzon, 2026). Closely related on the theory side, Bordalo et al. (2026) develop a psychology inspired model of macroeconomic expectations in which salient adverse states can generate pessimistic beliefs across outcomes, while Huo et al. (2024) show that ambiguity aversion under dispersed information can generate both pessimistic bias and heightened sensitivity to signals in forecasts. Closely related on the empirical side, Taubinsky et al. (2024) 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 pessimism explains a large share of the variances and covariances across macroeconomic expectations.

Methodologically, our paper brings SVAR-style identification tools to cross-sectional expectations by exploiting within-month belief covariances to recover interpretable latent belief components. This approach is closely related to the broader SVAR tradition of identifying latent shocks under weak restrictions, including sign restrictions (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) who use 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 describes the SCE data. Section 3 presents the structural econometric framework and provides reduced-form evidence on the patterns governing the cross-sectional covariance structure of expectations. Section 4 lays out our identification strategy, explains how we incorporate sign restrictions and narrative evidence, and describes the independent-observation baseline. Section 5 reports the structural results and analyzes the implied variance and covariance structure of beliefs. Section 6 studies the anatomy of the three components, relating them to demographics and showing how each shifts economy-wide and idiosyncratic expectations. Section 7 constructs an aggregate pessimism index from the estimated components. Finally, Section 8 concludes.

2 Data

We describe the survey data and document the basic variance fact that motivates the structural framework in Section 3: most variation in reported expectations is disagreement across households facing the same aggregate and local conditions.

2.1 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. It also records current and expected household income and spending. Appendix A reproduces the survey questions and defines all variables used in the analysis. Respondents understand the questions well: forecast errors do not vary systematically with panel tenure beyond the first months, and most respondents report non-extreme expectations (Appendix Figures A1 and A2).

One fact organizes the analysis: disagreement dominates the variation in reported expectations. Table 2 decomposes the variance of each belief into the share explained by year-month-by-state fixed effects, which absorb aggregate conditions and local business cycles, and the share left within those cells. The aggregate conditions account for less than 10 percent of the variation in each expectation; more than 90 percent is dispersion across respondents facing the same aggregate environment. This within-cell dispersion is the variation our framework targets, both in the reduced form and in the structural estimation.

Table 2: Variance of Beliefs: Common versus Idiosyncratic Components

	(1) Gov. debt	(2) Taxes	(3) Inflation	(4) Unemployment
Year-month-state FE share	0.096	0.064	0.068	0.064
Within-cell/idiosyncratic share	0.904	0.936	0.932	0.936
Observations	129465	149103	148874	149187

Notes: The first row reports the share of the variance of each expectation explained by year-month-by-state fixed effects. The second row reports the remaining within-cell share.

3 Structural Econometric Framework

We introduce a structural econometric framework for subjective macroeconomic expectations; identification is taken up in Section 4. This section first defines the framework, then documents reduced-form patterns in the cross-sectional covariance structure of expectations, and finally interprets the latent rotations. 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) aggregate belief components that are common across households and may propagate over time, and (ii) an idiosyncratic component capturing heterogeneous information, interpretations, or belief distortions:

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

Here $\varepsilon_t \in \mathbb{R}^G$ collects aggregate belief components, $\bar{\Gamma}_g(L)$ is a $G \times G$ one-sided lag polynomial mapping the history of these aggregate components into contemporaneous beliefs that is common across g groups, $\varepsilon_{i,t} \in \mathbb{R}^G$ collects idiosyncratic latent belief components, and $\Gamma \in \mathbb{R}^{G \times G}$ maps these components into the cross-sectional dispersion of expectations within month t.

The lag structure in $\bar{\Gamma}_g(L)$ captures persistence in how common macroeconomic news maps into beliefs, reflecting gradual learning and the propagation of these aggregate components. By contrast, Γ governs the contemporaneous effects of individual-level belief

components across belief variables. For tractability, we abstract from lagged idiosyncratic components, a natural restriction in short rotating panels, where respondents are observed only a handful of times. As a result, the component scores should be interpreted not only as innovations to an individual's beliefs, but also as absorbing observed and unobserved individual heterogeneity. The estimation in Appendix B makes this explicit by decomposing within-cell belief variation into a persistent household component and a transitory household-month component, akin to clustering.[2]

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 components that govern cross-sectional differences in beliefs. This distinction matters empirically: as Table 2 shows, cross-sectional dispersion accounts for more than 90 percent of belief variation, and, as we show later, much of its covariance structure is accounted for by a small number of latent components loading through Γ. Finally, we impose the following maintained conditions.

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

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

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

Assumption 1 normalizes scale and rules out redundant components 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 for the upcoming sections.

From common components to disagreement. Our focus in our identification framework is Γ. For that we isolate the cross-sectional component of beliefs and take cross-sectional

[2] 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. Including individual fixed effects does not alter our results.

means in (1) and difference them out:

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

The covariance matrix of demeaned beliefs is the key moment in our framework and is given by

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

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. To label these rotations, we need economically meaningful restrictions grounded either in theory or in the data.

3.1 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:

$$B_{i,t}(k) = \alpha + \sum_{m \neq k} \beta_m\, B_{i,t}(m) + FE_{j(i),t} + \epsilon_{i,t}, \tag{4}$$

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.

Under standard linear-regression logic, the coefficient β_m summarizes the partial correlation between beliefs about k and m, holding fixed state-by-time conditions. The inclusion of State$\times$Month fixed effects nets out common local business-cycle conditions and other factors 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 conditions. Therefore, a positive[3] 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. The structural interpretation below gives this factor precise economic content.

In the baseline estimation, we consider economy-wide expectations regarding unemployment, taxes, inflation, and government debt. Table 3 contains the results. Each column denotes a different dependent variable (the left-hand-side expectation), whereas each row denotes a different right-hand-side expectation that has been standardized. All specifications include State×Month fixed effects; the baseline does not additionally condition on demographic controls.

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 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 3: Baseline: Reduced-Form Perceived Partial Covariances

	Dep. Var: $B_{i,t}(k)$ **(Columns)**			
$B_{i,t}(m)$ (Row)	Gov. Debt	Tax Changes	Inflation	Unemp.
Unemployment	1.2393^{***}	0.4083^{***}	0.4313^{***}	
	(0.0610)	(0.0281)	(0.0346)	
Tax Changes	2.6730^{***}		1.2516^{***}	1.7071^{***}
	(0.0743)		(0.0505)	(0.1177)
Inflation	4.5738^{***}	1.0826^{***}		1.5598^{***}
	(0.0917)	(0.0419)		(0.1237)
Gov Debt		1.3821^{***}	2.7342^{***}	2.6792^{***}
		(0.0388)	(0.0634)	(0.1304)
N	128852	128852	128852	128852

Year-Month × State FE. SE clustered by individual.

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

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

demographic controls $Z_{i,t}$, namely age, education, income, tenure in the panel, idiosyncratic employment and financial conditions, marital status, and a numeracy score measuring respondents' ability to solve simple mathematical problems; the adverse-direction comovement pattern is unchanged by this additional conditioning. Table A2 reports the same specification augmented with both demographic and individual fixed effects, which absorb time-invariant (permanent) differences in respondents' optimism/pessimism. The main finding is again unchanged: the adverse-direction comovement in expectations remains present even after subtracting permanent traits. Additionally, Table A6 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 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[4] explains 43% of the residual cross-sectional variation and loads with the same adverse-direction sign on all four beliefs, consistent with a common "pessimism" factor. The second component explains 23% and pairs higher unemployment expectations with lower inflation and fiscal expectations, a demand-type pattern. The remaining components explain smaller shares and mix the fiscal variables. PCA can suggest these patterns, but its orthogonal, variance-ranked components have no economic labels; the structural interpretation in the next sections turns them into interpretable objects.

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

Moreover, the PCA analysis provides a rank test for our data. If every household's beliefs were some reading of demand and supply views, all demeaned belief vectors would lie in a two-dimensional plane and $\Sigma_{\tilde{B}}$ would have at most two nonzero eigenvalues. Table 4 shows instead that the third and fourth eigenvalues account for 19% and 15% of residual variation. This means that over a third of cross-sectional disagreement lies outside any plane and more importantly, the dominant component is the pessimism one and not a purely demand or supply one.

Table 4: 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 affect belief comovement. 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 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 reports the income-specific slopes when inflation expectations are the dependent variable. 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 vertical bars report 95% confidence intervals.

The figure shows that adverse-direction comovement is present within each income

group. Households with higher unemployment, government debt, and tax expectations also report higher inflation expectations, conditional on demographics and State×Month fixed effects. The associations with government debt and taxes decline with income, while the association with unemployment expectations is positive and flatter across the income distribution. Income heterogeneity changes the magnitude of these partial associations; the sign pattern remains stable. Appendix Figures A9–A11 report the remaining dependent variables, and Appendix Figure A12 repeats the exercise by education group.

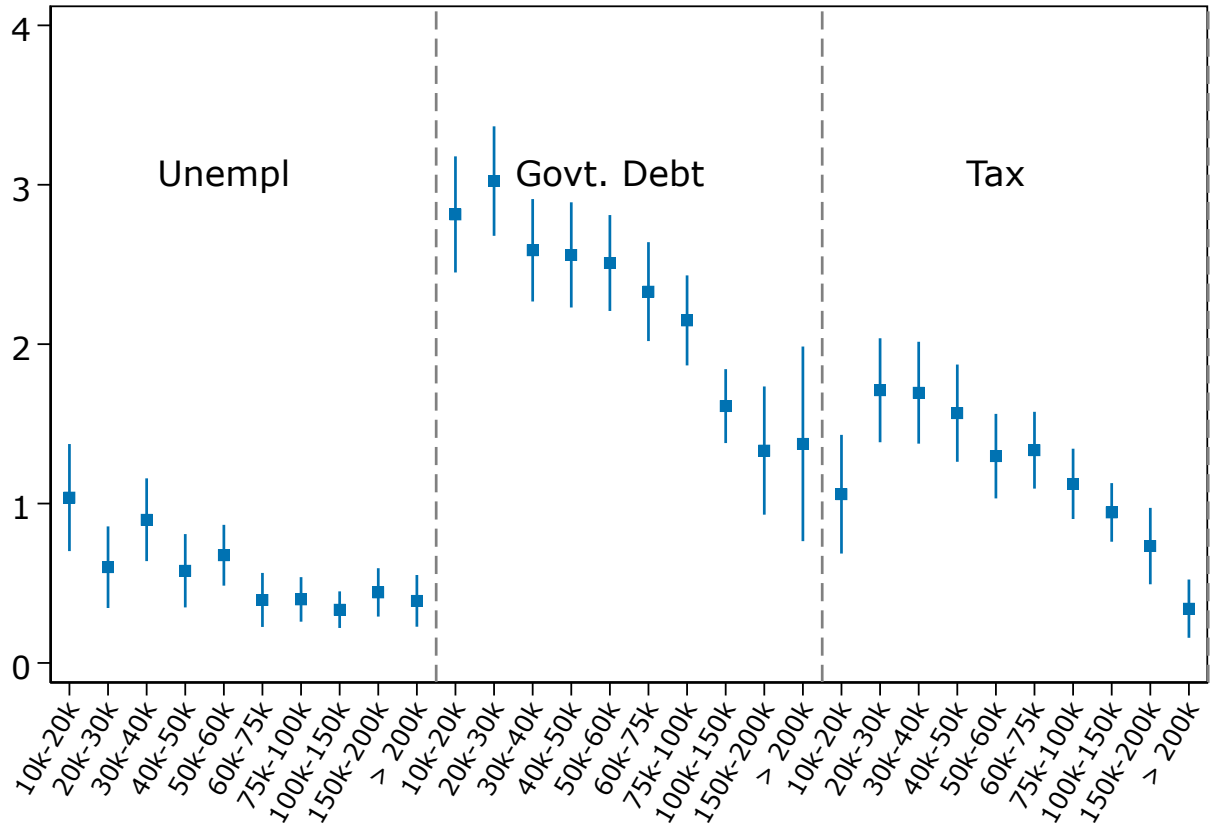


Figure 1: Heterogeneous Partial Associations by Income: Inflation Expectations

Notes: Income-specific slopes from equation (5). The dependent variable is inflation expectations. Dashed vertical lines separate blocks of right-hand-side expectations. Dots report point estimates; vertical bars report 95 percent confidence intervals. Standard errors are clustered by individual.

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

1. Splitting the sample into pre- and post-COVID periods (Tables A3 and A4)
2. Varying the set of fixed effects and household controls

3. Estimating rolling-window regressions (Figure A8)

4. Estimating nonparametric conditional expectation functions

5. Repeating the heterogeneity analysis using education, rather than income

6. Comparing raw and residualized correlation matrices (Table A5)

Overall, the reduced-form evidence supports a pessimism component in the cross-sectional covariance structure of expectations.

3.2 Interpreting the latent components

Our objective is to identify the main latent dimensions that organize cross-sectional disagreement in expectations. In our setting, households do not simply disagree about the level of one variable at a time. Rather, they hold subjective models that imply systematic patterns of comovement across expectations for different macroeconomic outcomes. We focus on a subset of expectations that is especially informative about the structure of disagreement:

$$B_{it} = \begin{bmatrix} \mathbb{E}_{i,t}[\text{Unemployment}] \\ \mathbb{E}_{i,t}[\text{Inflation}] \\ \mathbb{E}_{i,t}[\text{Taxes}] \end{bmatrix}.$$

This subset is useful because it allows us to distinguish broad classes of economic interpretations in a parsimonious way. In particular, the joint behavior of expected unemployment and expected inflation is informative about whether households interpret macroeconomic fluctuations through a demand-side or supply-side lens. Expected taxes provide an additional dimension that helps separate a supply-side interpretation from broader pessimism. Intuitively, supply-type and pessimistic views need not have the same fiscal implications, even if both imply adverse movements in some macroeconomic expectations.

The starting point of our approach is that different information frictions and different subjective views can generate distinct patterns of joint belief comovement. Some households may receive signals that emphasize demand-side forces, while others may receive signals that emphasize supply-side forces. More generally, households may also differ in how they map perceived signals into macroeconomic outcomes, that is, in the subjective views they use to interpret the economy. These two sources of heterogeneity are observationally similar in the reduced form, since both can generate systematic differences in the covariance structure of beliefs across households.

A natural benchmark for organizing these patterns is the New Keynesian framework. In that environment, two canonical forces shape the comovement of inflation and real activity. Demand-side disturbances imply weaker real activity and lower inflation, while supply-side disturbances imply weaker real activity and higher inflation. Translating this logic into our setting implies two natural latent dimensions of belief disagreement: one featuring negative comovement between expected unemployment and expected inflation, which we label the *demand* component, and another featuring positive comovement, which we label the *supply* component. These components capture heterogeneity in how households interpret the source of macroeconomic fluctuations, whether because they receive different information or because they rely on different subjective macroeconomic views.

At the same time, the reduced-form covariance structure points to a broader adverse comovement that is not well described by these two canonical forces alone. In particular, part of the cross-sectional variation in expectations reflects a broad worsening of beliefs across many outcomes in the same direction. This motivates a third component, which we label *pessimism*: a broad adverse shift in beliefs that is not tied to a specific equilibrium channel such as demand or supply, but instead captures a more general deterioration in households' outlook. Taken together, demand, supply, and pessimism constitute the smallest set of economically interpretable components that is consistent both with standard macroeconomic theory and with the empirical patterns in the data in our setting.

To organize the interpretation of these latent components, consider a simple framework in which households observe noisy signals about aggregate conditions and form posterior beliefs about the underlying state of the economy. Let

$$s_{it} = \left(s_{it}^{s},\, s_{it}^{d},\, s_{it}^{\tau}\right)'$$

denote the vector of underlying states, where the elements capture demand, supply, and tax conditions. Household i does not observe s_{it} directly, but instead forms a posterior

$$\tilde{s}_{it} = \left(\tilde{s}_{it}^{s},\, \tilde{s}_{it}^{d},\, \tilde{s}_{it}^{\tau}\right)'.$$

Expectations are then given by

$$B_{it} = H_i \tilde{s}_{it},$$

where H_i is the household-specific mapping or subjective model that maps perceived states into expectations.

This formulation is flexible enough to encompass both information frictions and heterogeneous subjective views. To make this explicit, let $\theta \in \{d, s, p\}$ index the three compo-

nents or latent dimensions: demand, supply, and pessimism. We decompose household i's subjective model mapping as combining these three views

$$H_i = \sum_{\theta \in \{d,s,p\}} h^\theta \left(a^\theta(i)\right)^T,$$

where each h^θ is a vector describing how the θ-component affects next period's beliefs on impact, while $a^\theta(i)$ captures household's i views or how it attributes perceived states into a particular component. Specifically, h^s has a positive entry in the first row, reflecting the idea that supply component has a positive effect on unemployment, a positive effect on inflation and a negative effect on taxes (see Appendix C). As for h^d, it has a positive entry on unemployment beliefs, but a negative one for inflation, leaving taxes unrestricted. Finally h^p has a positive entry in all rows, reflecting the idea that pessimism has a positive effect on all beliefs. This decomposition makes clear that heterogeneity in beliefs can arise through two distinct channels. First, households may differ in the posterior states they form because they receive different noisy signals. Second, even conditional on the same perceived state, they may differ in how they interpret that state because they hold different views about the economy.

The rational expectations NK benchmark is nested in this representation $H_i = H^{RE}$. Under rational expectations, households use the objective state mapping, so that demand and supply signals are interpreted according to their standard macroeconomic content. In that case,

$$a_i^d = e_d = (0,1,0)^T, \qquad a_i^s = e_s = (1,0,0)^T,$$

and

$$H_{RE} = h^s e_s^T + h^d e_d^T.$$

This benchmark implies the sign structure[5]

$$H_{RE} = \begin{bmatrix} + & + & ? \\ + & - & ? \\ - & ? & ? \end{bmatrix},$$

where rows correspond to unemployment, inflation, and taxes, and columns correspond to supply, demand, and tax states. The first two columns summarize the canonical New Keynesian prediction: adverse supply forces raise both unemployment and inflation, whereas adverse demand forces raise unemployment and lower inflation. The negative

[5]Notice that we leave as undefined the signs associated to the third column to simplify the exposition.

entry in the tax row reflects that an adverse supply event depresses the tax base and triggers countercyclical transfers; Appendix C derives this sign structure in a stylized New Keynesian model with a fiscal block.

The same framework also accommodates departures from this benchmark. For example, consider a household with a "purely supply-side view" of the economy, in the sense that it interprets all macroeconomic signals through supply lens. Such a household may effectively use

$$a_i^s = (1, 1, 0)^T,$$

so that both supply and demand signals load only on the supply component. In that case,

$$H_S = h^s(1, 1, 0)^T = \begin{bmatrix} + & + & ? \\ + & + & ? \\ - & - & ? \end{bmatrix}.$$

This example illustrates the core identification challenge. Positive comovement between expected unemployment and expected inflation need not arise solely because a household has received relatively more information about supply events. It may also arise because the household interprets the economy using a supply-side view. For the case of pessimism views, households use

$$a_i^p = (1, 1, 1)^T,$$

so that all signals load into the pessimism component with the same sign. In that case, the pessimism view generates a broad adverse shift in beliefs across all outcomes, consistent with the reduced-form patterns documented above.

More generally, the components we recover from the data should be understood as capturing variation generated by information frictions, subjective views, or both. The theoretical framework provides a direct interpretation of the reduced form in equation (2). Using the decomposition above,

$$\widetilde{B}_{it} = \sum_{\theta \in \{d,s,p\}} h^\theta \underbrace{\left[a^\theta(i)^T \widetilde{s}_{it} - \mathbb{E}_i\left(a^\theta(i)^T \widetilde{s}_{it} \right) \right]}_{\epsilon_{it}^\theta}, \qquad \text{with } h^\theta \text{ equal to the } \theta\text{th column of } \Gamma.$$

This expression clarifies the mapping between theory and data. Each empirical component's score combines heterogeneity in posterior beliefs about the underlying state with heterogeneity in the views used to interpret that state. A supply component in the data may therefore arise because some households receive relatively more supply-related information or because they interpret a broader set of signals through a supply-side lens,

or both. The same logic applies to the demand component. By contrast, the pessimism component captures a more general pessimistic view of the economy, one that generates broad adverse shifts in beliefs rather than movements tied to a particular equilibrium channel.

Our decomposition should be interpreted as identifying latent dimensions of belief disagreement rather than as cleanly separating information from subjective views. This distinction is important. The empirical objects we recover summarize the main reduced-form patterns in how households jointly revise expectations across outcomes, regardless of whether those patterns originate from noisy information, subjective views, or some combination of the two. The framework nests standard rational-expectations models, dispersed- and sticky-information environments, rational inattention, heterogeneous subjective models, and behavioral distortions such as myopia or diagnostic expectations as special cases of the mapping H_i; Appendix D works through three examples and discusses the environments the framework does not cover.

4 Identification and Estimation

The reduced-form covariance matrix of $\widetilde{B}_{it}$ identifies Γ only up to an orthogonal rotation. To pin down that rotation, we impose restrictions with clear economic content, motivated by theory (Section 3.2) and by the reduced-form patterns documented above. We then describe how we estimate the covariance matrix itself, accounting for the rotating-panel structure of the survey.

4.1 Sign and narrative restrictions

Our baseline sign scheme, with rows ordered as unemployment, inflation, and tax expectations and columns as the pessimism, demand, and supply components, is

$$\tilde{B}_{it} = \begin{bmatrix} \mathbb{E}_{i,t}[\text{Unemployment}] \\ \mathbb{E}_{i,t}[\text{Inflation}] \\ \mathbb{E}_{i,t}[\text{Taxes}] \end{bmatrix} = \underbrace{\begin{bmatrix} + & + & + \\ + & - & + \\ + & ? & - \end{bmatrix}}_{\Gamma} \underbrace{\begin{bmatrix} \text{Pessimism} \\ \text{Demand} \\ \text{Supply} \end{bmatrix}}_{\epsilon_{it}}.$$

The first column corresponds to pessimism. We interpret this component as a broad adverse shift in beliefs across outcomes. Accordingly, we require it to raise expected unemployment, expected inflation, and expected taxes. Hence pessimism captures a

diffuse deterioration in outlook that worsens a broad set of macroeconomic expectations in the same direction.

The second column corresponds to the demand component. Here we require expected unemployment to rise and expected inflation to fall. These signs follow directly from the standard demand-side interpretation of macroeconomic fluctuations. We leave the sign on expected taxes unrestricted, reflecting that the fiscal implications of demand-type beliefs are less tightly pinned down by the benchmark theory and may vary across environments (see Appendix C).

The third column corresponds to the supply component. We require this component to raise expected unemployment and expected inflation, consistent with the standard New Keynesian interpretation of adverse supply forces. We additionally impose that it lowers expected taxes. This last restriction is particularly useful empirically because it helps distinguish a supply-side component from broader pessimism: even if both imply adverse movements in some expectations, they differ in their fiscal implications. In particular, an adverse supply component is expected to lower expected taxes, consistent with a contractionary supply shock that depresses the tax base and triggers fiscal transfers, whereas pessimism raises expected taxes.

Several features of these restrictions are worth emphasizing. First, they are intentionally modest. Our goal is not to impose a fully specified structural model on the data, but rather to rule out large classes of observationally equivalent rotations using economically interpretable sign information. This is also the reason why we focus on a subset of all the available aggregate expectations. Second, the restrictions are grounded in both theory and evidence. The signs on the supply and demand components are motivated by the New Keynesian benchmark and by existing evidence on how agents interpret inflation and real activity[6]; Appendix C derives them in a stylized New Keynesian model with a fiscal block. The pessimism component is motivated by the broader adverse comovement that emerges clearly in the reduced form and by recent work emphasizing the role of broad pessimistic narratives in household beliefs (Taubinsky et al., 2024). Third, expected taxes play a central role in the identification strategy. Focusing only on unemployment and inflation would help distinguish demand from supply components, but it would be substantially harder to separate a genuinely supply-side component from broader pessimism. Tax expectations provide the additional variation needed to do so.

[6]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.

Narrative restriction. We strengthen identification of the pessimism component by imposing narrative restrictions based on salient quasi-exogenous events. These events help isolate episodes in which a broad deterioration in sentiment is especially plausible, thereby sharpening the interpretation of the third component. Our approach thus combines the two standard ingredients of modern SVAR identification, sign restrictions and narrative information, but applies them to the cross-section of beliefs rather than to aggregate macroeconomic time series.

We incorporate a narrative restriction by exploiting the idea that certain salient events can shift pessimism directly but have no direct macroeconomic effects (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. 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 an event with four or more fatalities and completing the survey after the event in the same or following month. The narrative restriction requires that, on average, exposed respondents whose unemployment expectations faced an upward change, have a higher pessimism score than non-exposed respondents in the same period. We operationalize this requirement as an inequality on the recovered component scores that must hold across admissible rotations.

Table 5 documents the effect of these events in unemployment beliefs, what we call the first stage behind this restriction. Following Lagerborg et al. (2023) in our context, we regress forward changes in unemployment expectations, $U_{i,t+k} - U_{i,t}$, where t is the month of the event, on the exposure indicator, year-month and commuting-zone fixed effects, and household controls. Exposed respondents revise their unemployment expectations upward relative to non-exposed respondents, and the effect is positive and significant across forward horizons.

Table 5: 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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

4.2 Estimation: independent household-month observations

The model is generically set-identified under sign and narrative restrictions. Our baseline treats the residualized household-month belief vectors as conditionally independent Gaussian draws,

$$\tilde{B}_{it} \overset{\text{iid}}{\sim} N(0, \Sigma_{\tilde{B}}).$$

Let $S = \sum_{i,t} \tilde{B}_{it}\tilde{B}_{it}'$ denote the residual sum of squares after absorbing year-month-by-state cell means. Under a flat prior, the posterior for $\Sigma_{\tilde{B}}$ is inverse-Wishart with scale matrix S and degrees of freedom adjusted for the three belief variables and the absorbed cells. Drawing from this posterior carries sampling uncertainty about the covariance matrix into the structural decomposition.

For each covariance draw $\Sigma_{\tilde{B}}^{(m)}$, we draw a Haar-uniform orthogonal matrix $Q^{(m)}$ and construct the candidate impact matrix

$$\Gamma^{(m)} = \text{chol}(\Sigma_{\tilde{B}}^{(m)})Q^{(m)}.$$

We retain the draw if it satisfies the sign restrictions and the narrative restriction, and continue until 2,000 draws are accepted. The accepted draws reflect both sampling uncertainty in $\Sigma_{\tilde{B}}$ and identification uncertainty across admissible rotations. For each accepted draw we recover household-level component scores

$$\epsilon_{it}^{(m)} = \Gamma^{(m)-1}\tilde{B}_{it},$$

which are orthonormal by construction. Every downstream statistic—the variance and covariance decompositions, anatomy regressions, and projections on outside aggregate and idiosyncratic expectations—is computed separately for each accepted draw. The reported bands therefore propagate posterior uncertainty about Γ into the second-stage estimates. Appendix B replaces the independence assumption with a Bayesian random-effects model that separates persistent household heterogeneity from transitory household-month variation and acts as a clustering device for repeated households' observations; the substantive conclusions are unchanged.

5 Structural Results

We now turn to the results from our structural framework, obtained from the sign and narrative restrictions and the independent-observation covariance model described above.

Throughout, figures report posterior medians together with 68 and 90 percent posterior bands across the 2,000 accepted draws. Appendix B shows that the hierarchical random-effects specification delivers the same qualitative conclusions with wider uncertainty bands.

Impact Responses. We begin with the contemporaneous responses of beliefs about unemployment, inflation, and taxes to the three latent components. While the identification restrictions discipline the signs of most responses, the magnitudes are left unrestricted and are therefore estimated from the data. Figure 2 reports the posterior distribution of each column of Γ: the response of the three beliefs, in their own units, to a one-standard-deviation realization of each component. Unemployment expectations are measured in probability points; inflation and tax expectations in percentage points.

Three patterns emerge. First, unemployment beliefs respond most strongly to pessimism: the median response is about 16 probability points, compared with about 11.5 for demand and 8 for supply. Second, inflation beliefs respond most strongly to the demand component (a median of about -5.5 percentage points), followed by supply and pessimism (both about $+3.5$). Third, tax beliefs respond to the three components with similar absolute magnitudes but opposite signs: the median response is about $+4$ percentage points for pessimism and -3.5 for supply, while the unrestricted demand response is negative at the median (about -2) with a posterior band that includes zero. The tax responses are the fiscal fingerprint that separates the three views: households with a higher pessimism score expect higher taxes, while households with supply-type views expect lower taxes.

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, so that the share of the variance of variable i attributable to component k is

$$\omega_{i,k} = \frac{\Gamma_{i,k}^2}{\sum_k \Gamma_{i,k}^2}, \text{ since } \Sigma_{\tilde{B},ii} = \sum_k \Gamma_{i,k}^2. \tag{6}$$

By construction, for any given draw and outcome variable, the contributions of the three components sum to one.

The decompositions assign clearly different roles to the three components. For unemployment beliefs, pessimism accounts for the largest share of the variance: the median share is about 0.57, against 0.30 for demand and 0.13 for supply. For inflation beliefs, demand is the dominant driver with a median share of about 0.58, while pessimism and

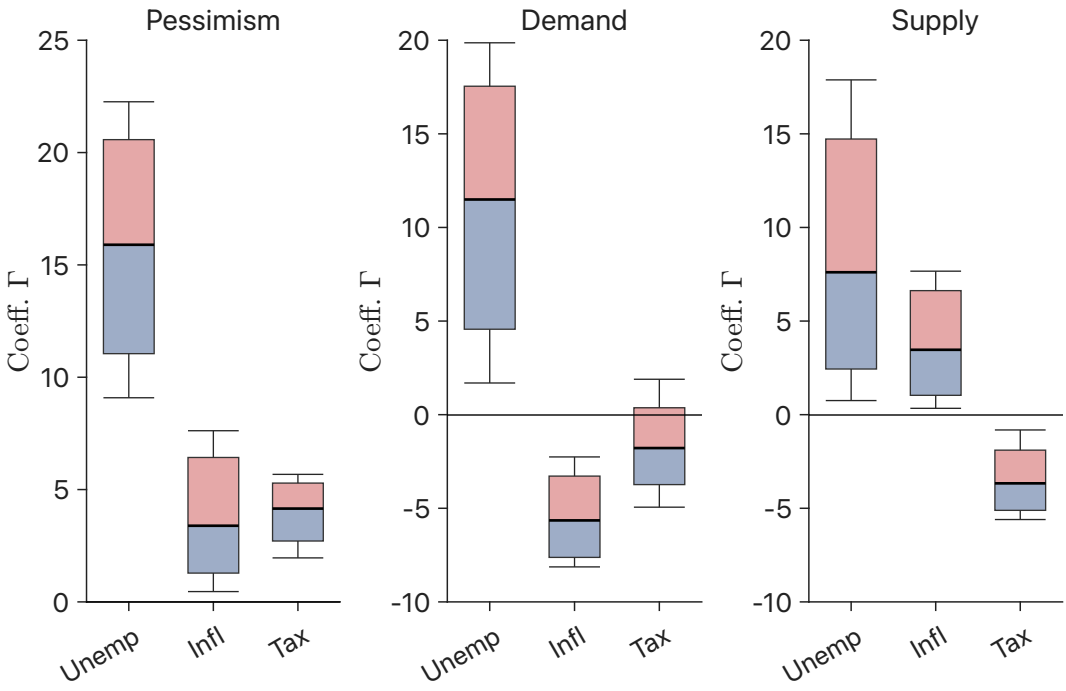


Figure 2: Contemporaneous Responses

Notes: Each panel reports the column of Γ corresponding to one component: the response of unemployment (probability points), inflation, and tax expectations (percentage points) to a one-standard-deviation component. Medians, 68 and 90 percent posterior bands across accepted draws.

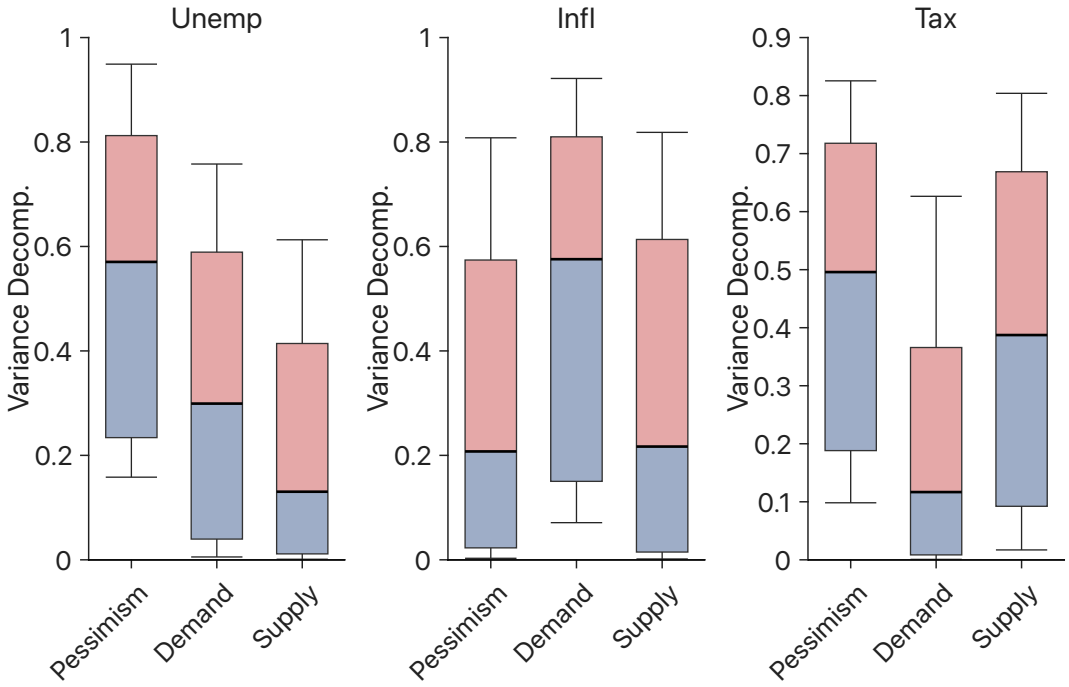


Figure 3: Variance Decomposition

Notes: Share of the within-cell variance of each belief attributable to each component, $\omega_{i,k} = \Gamma_{i,k}^2 / \sum_k \Gamma_{i,k}^2$. Medians, 68 and 90 percent posterior bands across accepted draws.

supply each account for about 0.21. For tax beliefs, pessimism again accounts for the largest median share (0.49), followed by supply (0.39) and demand (0.12).

Covariance Decomposition. In addition to decomposing the variance of each belief, we can also decompose the contribution of each component 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 components 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 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 views, demand-side views, 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 component k is

$$\omega_{ij,k} = \frac{\Gamma_{i,k}\Gamma_{j,k}}{\sum_\ell \Gamma_{i,\ell}\Gamma_{j,\ell}}, \qquad \text{where} \qquad \Sigma_{\tilde{B},ij} = \sum_\ell \Gamma_{i,\ell}\Gamma_{j,\ell}. \tag{7}$$

By construction, $\sum_k \omega_{ij,k} = 1$ for any pair (i, j) draw by draw, so $\omega_{ij,k}$ provides a transparent accounting of which components generate the observed covariances. Because contributions of different components can have opposite signs, individual shares can exceed one; a share above one means that the component contributes more than the entire observed covariance and is partially offset by another component. 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 shares of the pessimism, demand, and supply components.

The unemployment-inflation panel delivers one of the paper's central results. The observed within-cell covariance between unemployment and inflation beliefs is positive: households who expect higher unemployment also expect higher inflation, the type of comovement a pure supply-side view would produce. However, the decomposition shows that supply views generate only a modest part of it. The median contribution share of pessimism is about 2.0, so pessimism generates roughly twice the observed covariance; demand offsets slightly more than the observed covariance (a median share of about -2.1); and supply contributes the remainder (about 0.8). In other words, a "supply-like" covariance does not imply that households primarily hold supply views: it emerges as the net of a large pessimism contribution and an offsetting demand contribution whose sum happens to mimic a supply pattern. 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)): at the level of joint belief comovement.

The other two panels reinforce this reading. For the unemployment-tax covariance, pessimism is the only component with a positive median contribution: only pessimism moves unemployment and tax expectations in the same adverse direction. For the inflation-tax covariance, pessimism and demand contribute positively, while supply contributes negatively, reflecting the negative supply loading on taxes imposed by the fiscal restriction.

This 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 views, supply-type views, 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

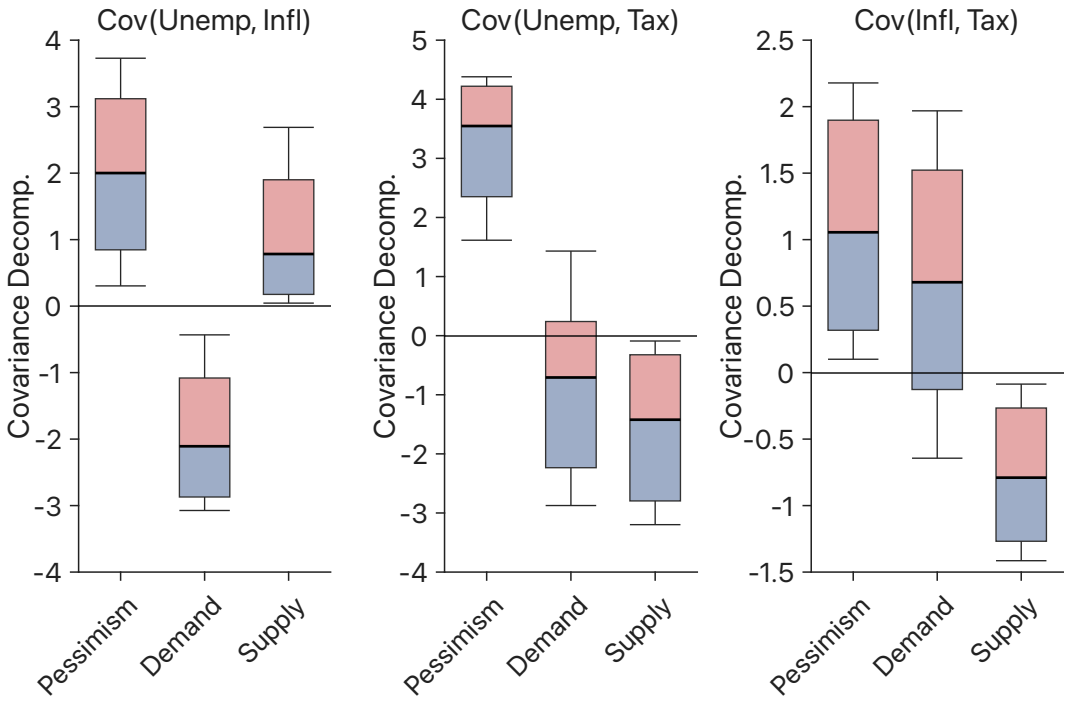


Figure 4: Covariance Decomposition

Notes: Share of each pairwise within-cell belief covariance attributable to each component, $\omega_{ij,k} = \Gamma_{i,k}\Gamma_{j,k} / \sum_{\ell} \Gamma_{i,\ell}\Gamma_{j,\ell}$. Shares sum to one draw by draw; individual shares can exceed one when components offset. Medians, 68 and 90 percent posterior bands across accepted draws.

beliefs jointly. Even if the policy instrument is unchanged, these distinct views imply different expectation-management challenges and can generate different implications for how households interpret news about inflation, unemployment, and fiscal policy.

6 Anatomy of Expectations

The structural estimation serves two purposes. First, it recovers the impact responses and the variance and covariance decompositions reported above. Second, it delivers household-level scores for each component, $\epsilon_{it}^{(m)} = \Gamma^{(m)-1}\tilde{B}_{i,t}$, for every accepted draw m. This section uses the scores to answer two questions: which households hold which views, and what do the views imply for expectations outside the three-variable system used in the identification.

6.1 Which households hold which views?

For each accepted draw we estimate

$$\epsilon_{it}^{k} = \alpha_{\text{state}\times\text{month}} + X_{it}'\beta^{k} + u_{it}^{k}, \qquad k \in \{p, d, s\}, \tag{8}$$

where the score ϵ_{it}^{k} is standardized to mean zero and unit variance and normalized so that the component raises unemployment beliefs, and X_{it} collects household demographics, economic circumstances, experiences, and spending characteristics. We report adjusted predicted means[7] by group for categorical variables and marginal effects for continuous variables. Because equation (8) is estimated draw by draw, the figures show the distribution of these statistics across accepted draws: the median draw, the interquartile range, and the 10th to 90th percentile range.

Pessimism. Figures 5 and 6 report the anatomy of the pessimism score. Pessimism is concentrated among households under economic strain, and the strain shows up in circumstances rather than in the level of income. Respondents who report being financially worse off than a year ago score about 0.27 standard deviations above the state-month average, against -0.08 for everyone else. The gradient in income volatility is equally steep: respondents whose income varies by more than 15 percent score about 0.27, against roughly zero for those with stable incomes. Health matters in the same way: pessimism falls monotonically from poor health (0.16) to excellent health (-0.05). Labor-market positions with high exposure carry high scores: respondents holding three or more jobs (0.19) and the long-term unemployed (0.11) are the most pessimistic groups, while respondents out of the labor force score below average (-0.10). Political alignment matters too: respondents living in counties unaligned with the sitting president score higher (0.04) than aligned

[7] Predicted means are regression-adjusted group averages, holding other covariates at their observed values, and are expressed in standard-deviation units.

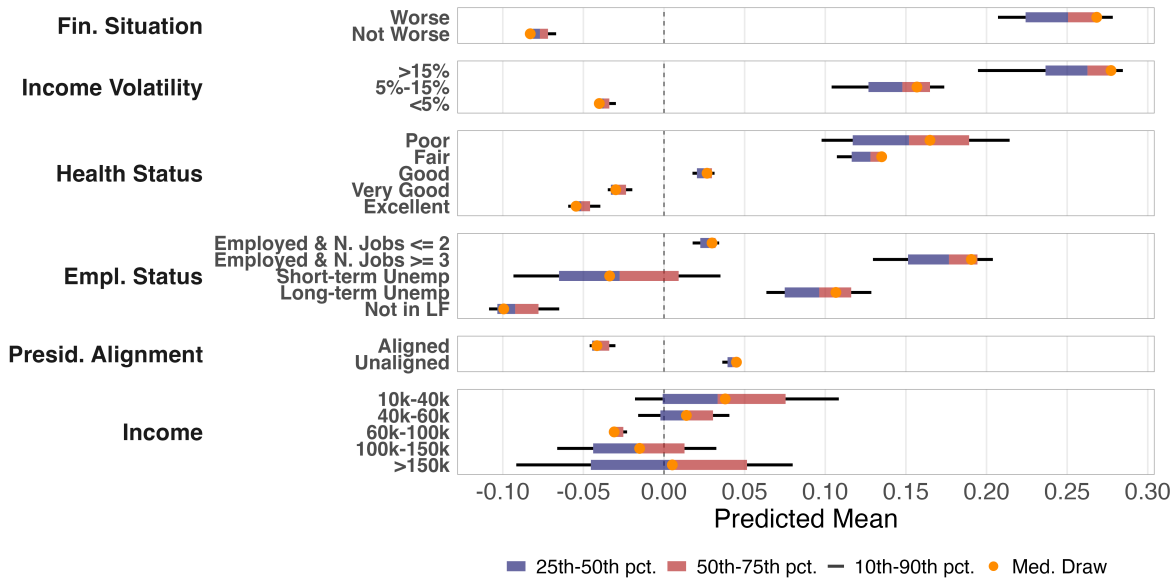
respondents (-0.04). By contrast, current income has no monotone relationship with pessimism: the highest predicted scores appear at both ends of the income distribution, and the middle scores lowest. Older respondents are mildly more pessimistic than younger ones (Figure 7b), and pessimism rises mildly in highly educated people (Figure 7a).

The marginal effects in Figure 6 point to the origins of these views. Lifetime labor-market experience is the strongest single predictor: a respondent whose lifetime average of national unemployment is one point higher scores about 0.44 standard deviations higher on pessimism, while lifetime inflation experience has no comparable effect. Current spending patterns line up with the same interpretation: pessimism rises with the share of spending on medical expenses and mortgages and falls with the share spent on recreation. Pessimism therefore reflects both current difficulties and the imprint of past macroeconomic experiences on how households interpret the economy.

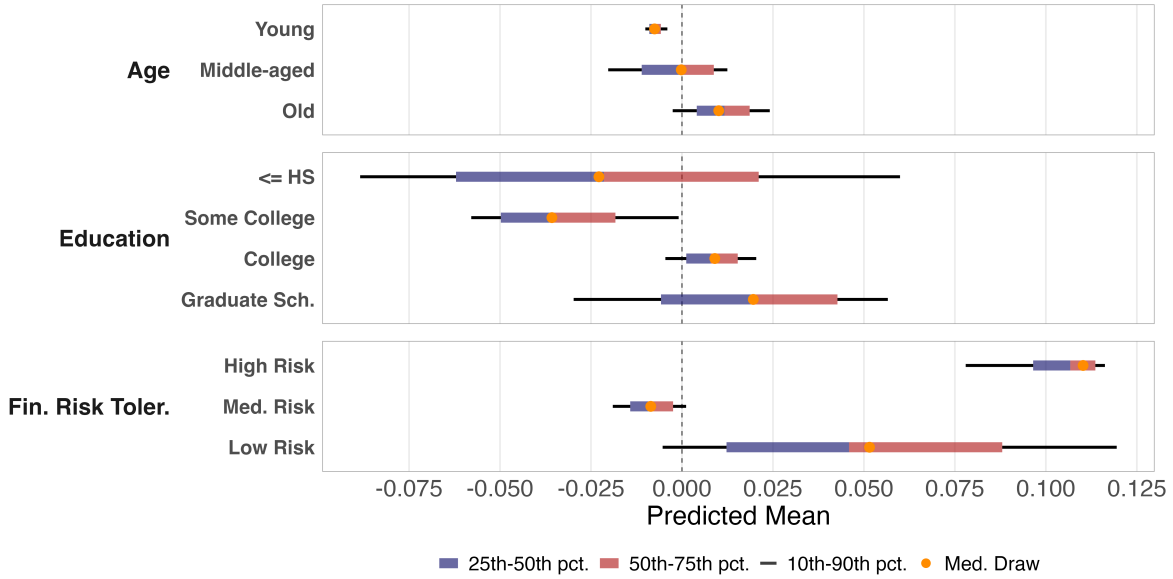
Demand and supply. The demand and supply views have sharply different anatomies, and both differ from pessimism. Figure 7 reports the predicted scores of all three components by education and age; Appendix F contains the corresponding figures for income, employment status, lifetime experiences, income volatility, spending shares, and financial risk tolerance.

Demand views rise steeply with education: predicted scores increase monotonically from -0.17 for respondents with at most a high-school degree to 0.10 for respondents with graduate schooling. They are also more common among middle-aged respondents, higher-income households, respondents willing to take financial risks and households that spend more on recreation. Demand views are thus concentrated in households connected to discretionary consumption, for whom changes in spending, income, and activity are salient.

Supply views show none of these gradients: they are flat in education and rise with age. They are more common among the unemployed and respondents out of the labor force, and they are unrelated to income volatility and to current spending patterns. Their clearest correlate is historical: respondents with worse lifetime inflation experiences load more heavily on the supply view. Where pessimism carries the imprint of lifetime unemployment, the supply view carries the imprint of lifetime inflation, a stark contrast among the two.



(a) Financial situation, income volatility, health, employment status, presidential alignment and income



(b) Age, education and financial risk tolerance

Figure 5: Anatomy of Pessimism: Adjusted Predicted Means by Group

Notes: Adjusted predicted means of the standardized pessimism score by demographic group, from equation (8) estimated for each accepted draw, controlling for State×Month fixed effects and household controls. Orange dots: median draw; boxes: interquartile range; whiskers: 10th to 90th percentile across draws.

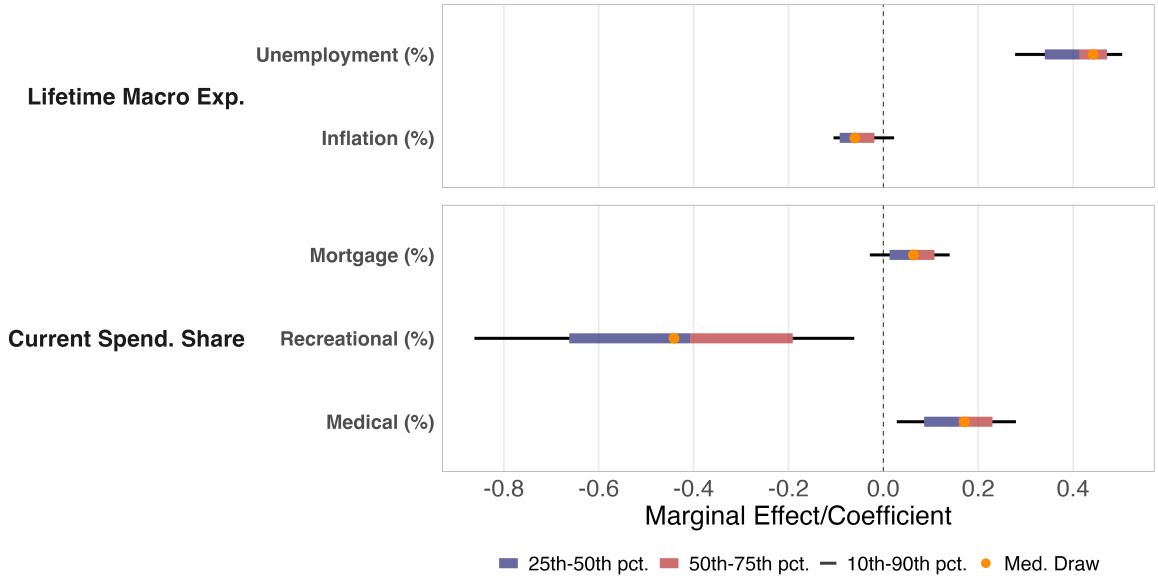


Figure 6: Anatomy of Pessimism: Marginal Effects of Experiences and Spending

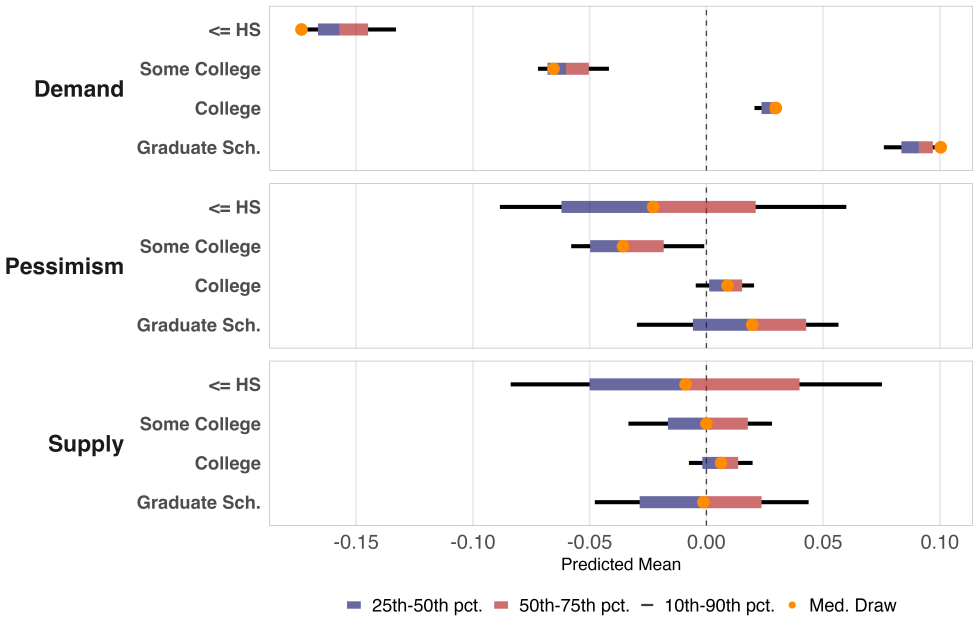
Notes: Marginal effects of continuous covariates on the standardized pessimism score, from equation (8). Lifetime macro experiences are the averages of national unemployment and inflation over the respondent's lifetime. Orange dots: median draw; boxes: interquartile range; whiskers: 10th to 90th percentile across draws.

6.2 Projections on Aggregate Expectations

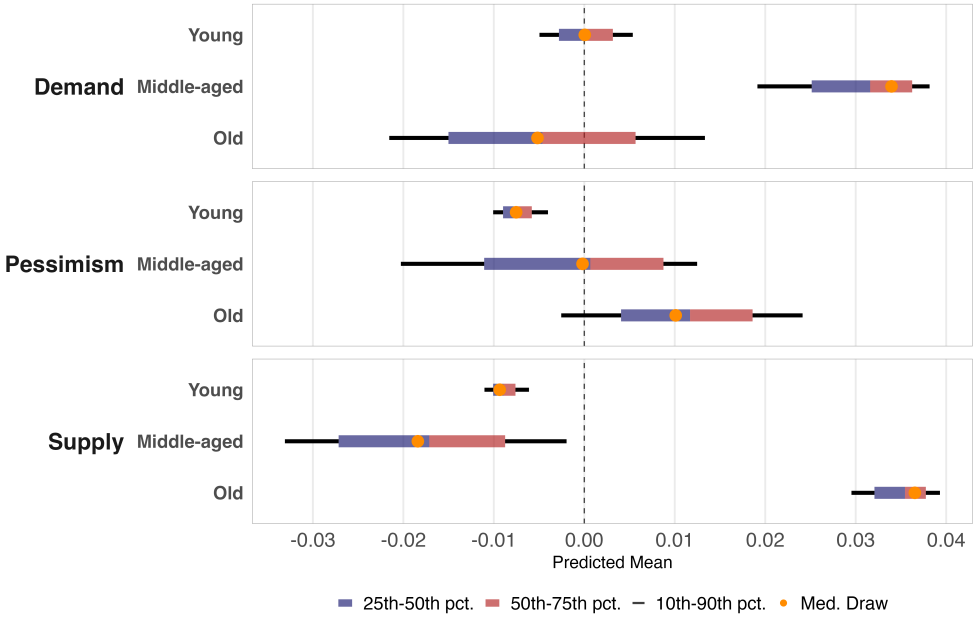
The structural framework uses only three beliefs in the estimation. Nevertheless, the SCE elicits a much broader set of economy-wide expectations: government debt, interest rates on savings accounts, credit access, long-run inflation, inflation uncertainty constructed from density forecasts, and a range of prices (gasoline, food, housing, and gold). Ideally, we would include these variables directly in the structural estimation, but doing so would require specifying the contemporaneous sign patterns of a much larger system. Instead, we recover the component scores from the baseline system and project them onto the broader set of beliefs. This both augments the structural estimation and provides an external validation check: it asks how each identified component moves expectations that were not used in the identification. For each accepted draw we estimate

$$y_{it}^{\ell} = \alpha_{\text{state}\times\text{month}} + \theta_{\ell}^{k}\,\epsilon_{it}^{k} + X_{it}'\beta_{\ell}^{k} + u_{it}^{\ell,k}, \qquad k \in \{p, d, s\}, \tag{9}$$

where y_{it}^{ℓ} is an expectation outside the baseline system, ϵ_{it}^{k} are the component scores, and the score is normalized so that the component raises economy-wide unemployment beliefs by 10 percentage points on impact. The coefficient θ_{ℓ}^{k} measures how expectation ℓ differs for households with a higher score on component k.



(a) By education



(b) By age

Figure 7: Anatomy of Demand, Pessimism, and Supply by Education and Age

Notes: Adjusted predicted means of each standardized component score by group, from equation (8). Orange dots: median draw; boxes: interquartile range; whiskers: 10th to 90th percentile across draws.

Figure 8 reports the results for pessimism. Consistent with a broad adverse shift in beliefs, a pessimism realization that raises unemployment beliefs by 10 percentage points is associated with higher expected government debt growth (about 2 percent), a higher probability that interest rates on savings accounts rise (about 2 percentage points), tighter expected credit access (about 4.7 percentage points), higher long-run inflation expectations (about 1.3 percent), and higher inflation uncertainty under both the variance and interquartile-range measures. Expected prices for gas, food, housing, and gold rise by roughly 0.8 to 1.6 percent. The one exception is the reported probability of a stock-market increase, which weakly rises by about 0.7 percentage points.[8]

The demand and supply projections sharpen the labels (Appendix Figure A18). A higher demand score moves the whole nominal block down: government debt falls by about 2.4 percent, long-run inflation by about 2.5 percent, inflation uncertainty falls, and expected gas, food, and home prices all decline. A higher supply score leaves fiscal beliefs unchanged (the debt bands include zero) and specific prices unaffected, while raising long-run inflation expectations by about 2 percent together with inflation uncertainty. Each view thus has a distinct out-of-sample signature, and the signatures line up with the economic content of the labels: demand views are about perceived slack, supply views are about persistent inflation pressure without government debt implications, and pessimism worsens everything.

6.3 Projections on Idiosyncratic Expectations

The SCE also asks about a rich set of expectations over respondents' own outcomes: household savings, income, and spending; personal earnings and earnings uncertainty constructed from density forecasts; job transitions (the probability of losing a job, leaving a job voluntarily, or accepting a new offer); and financial conditions (the probability of defaulting on a loan and of being financially worse off a year ahead). We estimate equation (9) with these outcomes to assess whether the views recovered from beliefs about the aggregate economy extend to how households see their own futures.

Figure 9 reports the results for pessimism. The same 10-percentage-point pessimism realization is associated with a 3.5-percentage-point higher probability of being financially worse off next year, a 1.4-percentage-point higher probability of defaulting on a loan, and higher perceived job risk: the probability of losing one's job rises by about 1.7 percentage points and of leaving it by about 1.2, while the probability of accepting a new offer is

[8] One explanation for this is 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.

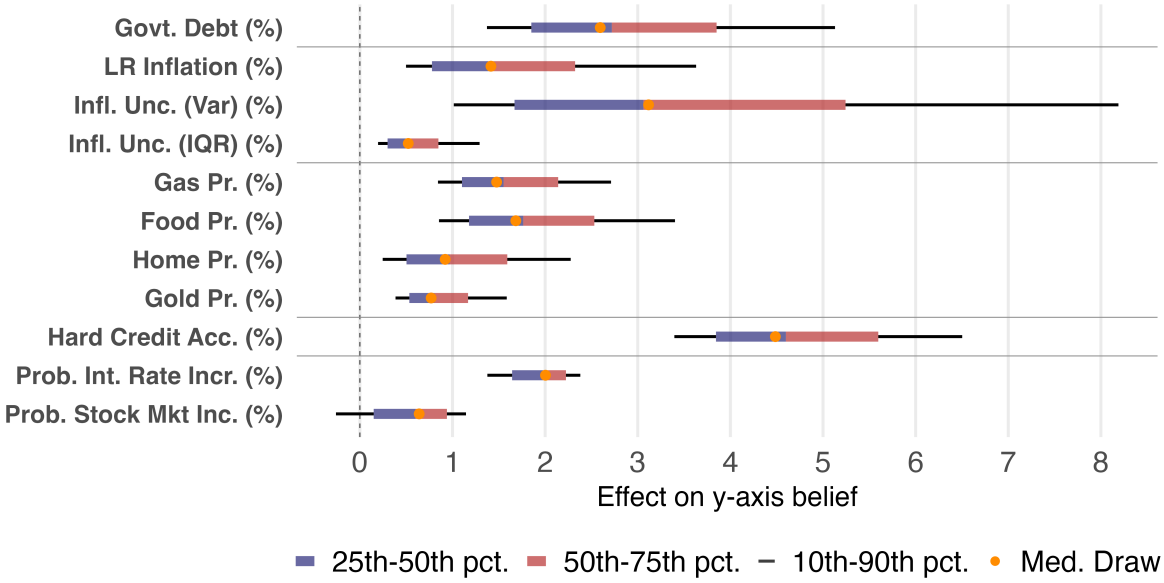


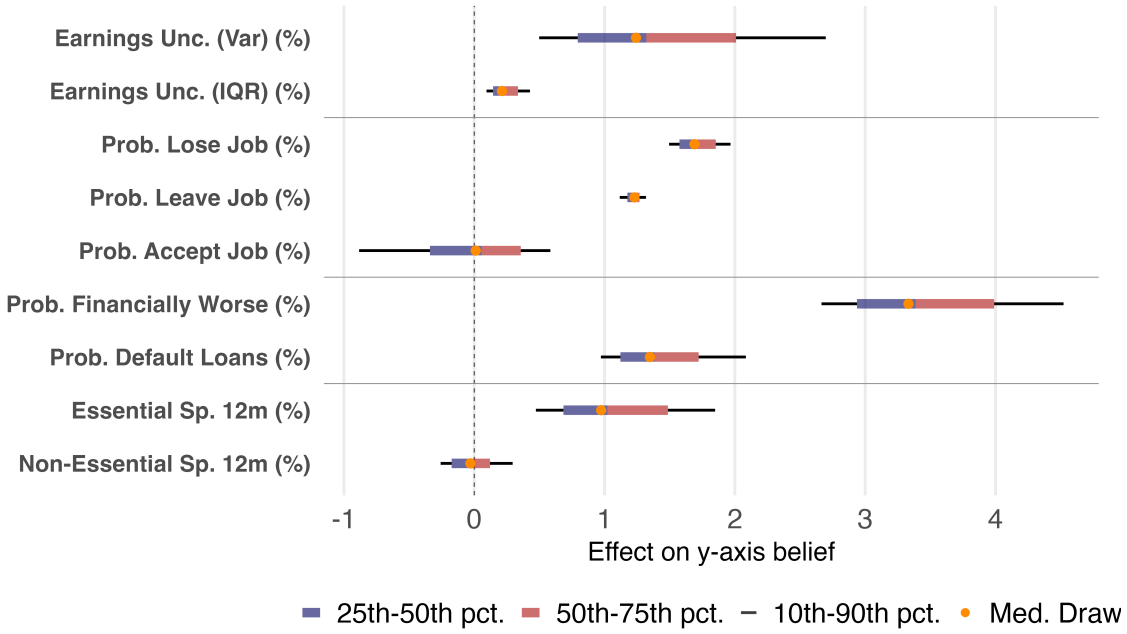
Figure 8: Response of External Economy-Wide Expectations to Pessimism

Notes: Coefficients θ_ℓ^p from equation (9), estimated for each accepted draw. The pessimism score is normalized to raise unemployment beliefs by 10 percentage points. Orange dots: median draw; boxes: interquartile range; whiskers: 10th to 90th percentile across draws.

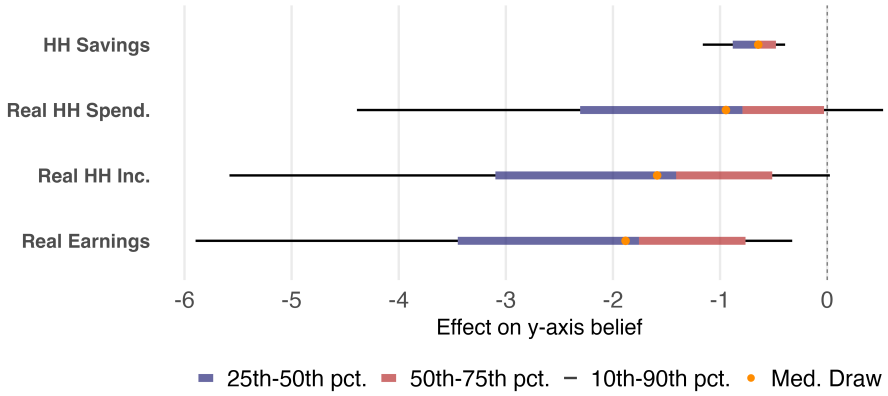
unchanged. Earnings uncertainty rises. Expected household savings, spending, income, and personal earnings all decline, by about 0.6 to 1.5 percent. The composition of expected spending shifts toward necessities: expected essential spending rises by about 1.1 percent while expected non-essential spending is unchanged. Pessimistic households, in short, expect to be squeezed: lower income and earnings, higher default and job-loss risk, and a budget tilted toward essentials.

The demand and supply views differ again here (Appendix Figure A19). A higher demand score lowers earnings uncertainty and expected essential spending while raising all three job-transition probabilities, including the probability of accepting a new offer (about 2 percentage points); demand views are associated with perceived labor-market churn rather than distress. A higher supply score leaves earnings uncertainty and financial-distress probabilities largely unchanged and shows one clear effect: expected non-essential spending falls by about 1.5 percent, consistent with households who expect persistent cost-of-living pressure cutting discretionary items first.

Overall, projecting the identified components onto expectations outside the baseline system supports the structural interpretation. The pessimism component behaves like a broad deterioration in household outlook that crosses the boundary between macroeconomic and personal domains, while the demand and supply components behave like the



(a) Uncertainty, job transitions, financial conditions, and spending composition



(b) Household savings, spending, income, and earnings

Figure 9: Response of Idiosyncratic Expectations to Pessimism

Notes: Coefficients θ_ℓ^p from equation (9), estimated for each accepted draw. The pessimism score is normalized to raise unemployment beliefs by 10 percentage points. Orange dots: median draw; boxes: interquartile range; whiskers: 10th to 90th percentile across draws.

specific equilibrium interpretations their sign restrictions were designed to capture.

7 Aggregate Pessimism Index

The analysis so far uses within-cell variation: the common component of beliefs is removed before estimation, so the identified Γ speaks to disagreement across households. Using the mapping estimated by the baseline model, this section asks what the pessimism direction implies for the aggregate component itself. The exercise treats the average respondent in a state and month "as if" it were an individual respondent: the state-month average belief vector is mapped into the same pessimism, demand, and supply coordinates as individual residual beliefs.[9]

Concretely, for each accepted draw m we compute component coordinates of the state-month average belief vector, $e_{st}^{(m)} = \Gamma^{(m)-1}\bar{B}_{s,t}$, take the pessimism coordinate $P_{s,t}^{(m)}$, and aggregate across states using cell weights,

$$P_t^{(m)} = \sum_s \omega_{s,t}\, P_{s,t}^{(m)}.$$

The result is a monthly national index of common pessimism, with posterior bands inherited from the accepted draws. Figure 10 plots the standardized index.

The index lines up with the narrative record. Pessimism is elevated at the start of the sample in 2013, when uncertainty, international conflicts and fiscal issues dominated the news, declines through the mid-2010s, and reaches its trough in late 2017 and early 2018, after the Tax Cuts and Jobs Act was signed. It rises again in 2019, spikes sharply at the onset of the COVID-19 pandemic in the spring of 2020, and reaches its highest sustained level in 2022 during the inflation surge, debt ceiling discussions and the SVB bank collapse, before easing over 2023 and 2024.

Figure A7, reported in Appendix B.3, compares the index with the University of Michigan's Consumer Sentiment index. The two series are strongly negatively related, with a correlation of -0.80, although the index is constructed from a different survey, different questions, and moments that were never targeted: the pessimism direction comes entirely from the cross-sectional covariance of disagreement. The comovement is tightest

[9] The identification uses cross-sectional second moments, so applying Γ to cell averages is an assumption rather than an implication of the framework: it requires the mapping from views to beliefs to be the same for the common component as for the idiosyncratic one. The external validation in Appendix B, and the fact that the aggregate decomposition is not used in any of the structural results, are the reasons we present the index as a complement rather than as part of the identification.

after 2016, and the 2021–2022 collapse in measured sentiment coincides with the largest sustained rise in the pessimism index.

This correspondence carries a practical lesson for reading sentiment data. Consumer-sentiment indices are widely tracked but hard to interpret: a decline may reflect perceived demand weakness, perceived supply pressure, or a broad loss of confidence, and these call for different policy responses. Our decomposition supplies the missing interpretation. Because the same procedure delivers demand and supply indices alongside pessimism, a fall in sentiment can be decomposed into the component that generates it. The evidence in this section implies that most of what sentiment indices capture is the pessimism dimension: a broad adverse reassessment concentrated among economically fragile households, tied to lived labor-market experience, and spilling over into expectations about taxes, credit, and personal finances.

8 Conclusion

This paper studies macroeconomic expectations through the lens of the joint distribution of beliefs. Using the Survey of Consumer Expectations, we document that more than 90 percent of the variation in households' expectations about unemployment, inflation, and fiscal variables is disagreement across respondents facing the same aggregate conditions, and that this disagreement comoves systematically in an adverse direction. We develop a structural framework that identifies a small set of latent belief components (demand, supply, and pessimism) from the covariance structure of demeaned expectations, combining sign restrictions with a narrative restriction based on mass-shooting exposure. Our baseline estimates the belief covariance under independent household-month observations; a Bayesian hierarchical model that accounts for the rotating-panel design yields similar conclusions in the appendix.

The structural results assign the three components distinct roles. Pessimism dominates the variance of unemployment and tax beliefs and accounts for the bulk of the comovement across beliefs: it contributes roughly twice the observed unemployment-inflation covariance, with demand offsetting the difference, and it is the only component that moves unemployment and tax expectations in the same direction. What looks like a supply-side view of the economy in the raw comovement of expectations largely reflects broad pessimism. The three views are also held by different people: pessimism is concentrated among households with volatile incomes, deteriorating finances, poor health, and adverse lifetime unemployment experiences; demand views among educated, higher-income,

risk-tolerant households; and supply views among older respondents shaped by lifetime inflation experience. Projections onto expectations not used in the identification confirm the labels.

These results change how survey expectations should be read. Belief heterogeneity is not noise scattered across unrelated answers: a small number of latent components organize it, and the dominant one is a broad pessimism that crosses domains, from inflation expectations to perceived default risk. For models of expectation formation, this points toward mechanisms that generate correlated distortions across outcomes, in the spirit of experience effects and belief spillovers, rather than outcome-specific frictions alone. For policymakers, it implies that broad pessimism is concentrated among economically fragile households, whose expectations of higher taxes, tighter credit, and higher job-loss risk are exactly the beliefs that discourage spending and have effects in the real economy.

References

Andre, P., Pizzinelli, C., Roth, C., and Wohlfart, J. (2022). Subjective models of the macroeconomy: Evidence from experts and representative samples. *The Review of Economic Studies*, 89(6):2958–2991.

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*.

Baumeister, C. and Hamilton, J. D. (2015). Sign restrictions, structural vector autoregressions, and useful prior information. *Econometrica*, 83(5):1963–1999.

Bhandari, A., Borovi˘cka, J., and Ho, P. (2023). Survey data and subjective beliefs in business cycle models.

Bordalo, P., Gennaioli, N., Lopez-de Silanes, F., Schröder, S. G., Shleifer, A., and van Rooij, M. (2026). The psychology of macroeconomic expectations. Technical report, National Bureau of Economic Research.

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. (2026). Credit market experiences and macroeconomic expectations: Evidence and theory. Technical report, Bocconi University.

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.

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

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

Ferreira, C. and Pica, S. (2024). Households' subjective expectations: Disagreement, common drivers and reaction to monetary policy. Working Paper 2445, Banco de España.

Granziera, E., Reiche, L., Maffei-Faccioli, N., Weber, M., and Fastbø, T. M. (2025). Five facts about inflation expectations: Evidence from four different groups of agents. *Norges Bank WP*, 7:2025.

Hobert, J. P. and Casella, G. (1996). The effect of improper priors on gibbs sampling in hierarchical linear mixed models. *Journal of the American Statistical Association*, 91(436):1461–1473.

Huo, Z., Pedroni, M., and Pei, G. (2024). Bias and sensitivity under ambiguity. *American Economic Review*, 114(12):4091?4133.

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

Kuchler, T. and Zafar, B. (2019). 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. (2024). Beliefs about the economy are excessively sensitive to household-level shocks: Evidence from linked survey and administrative data. Technical Report 32664, National Bureau of Economic Research. Revised March 2026.

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

Weber, M., d'Acunto, 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.

A Survey Questions, Variable Definitions, and Data Quality

The sample covers June 2013 through April 2024. Depending on the outcome, the belief regressions use between roughly 129,000 and 149,000 household-month observations (Table 2).

A.1 Baseline expectation questions

The four economy-wide expectations used in the reduced form, of which three enter the structural system, are elicited as follows.

- *Inflation:* "What do you expect the rate of inflation to be over the next 12 months?"
- *Unemployment:* "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?"
- *Government debt:* "Over the next 12 months, by about what percent do you expect the level of U.S. government debt to increase/decrease?"
- *Taxes:* "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?"

A.2 Additional expectation questions

The projections in Sections 6.2 and 6.3 use, among others, the following questions.

- *Stock market:* "What do you think is the percent chance that 12 months from now, on average, stock prices in the U.S. stock market will be higher than they are now?"
- *Credit access:* "And looking ahead, do you think that 12 months from now it will generally be harder or easier for people to obtain credit or loans (including credit and retail cards, auto loans, student loans, and mortgages) than it is these days?"
- *Interest rates:* "What do you think is the percent chance that 12 months from now the average interest rate on saving accounts will be higher than it is now?"
- *Prices:* for gasoline, food, housing, and gold, "By about what percent do you expect the price of [item] to have increased/decreased 12 months from now?"

- *Household income:* "Which category represents the total combined pre-tax income of all members of your household (including you) during the past 12 months?"

A.3 Variable definitions

- *Long-run inflation and uncertainty measures:* long-run inflation expectations refer to the three-year-ahead question; inflation and earnings uncertainty (variance and interquartile range) are computed from the subjective density forecasts elicited by the survey.
- *Real tax and income variables:* tax, income, spending, and earnings expectations labeled "real" subtract the respondent's own 12-month inflation expectation.
- *Idiosyncratic job and financial outcomes:* probabilities of losing one's job, leaving it voluntarily, and accepting a new offer within 12 months; probability of missing a debt payment over the next three months; probability of being financially worse off 12 months ahead.
- *Income volatility:* categorical self-reported variability of household income (less than 5 percent, 5 to 15 percent, more than 15 percent).
- *Financial situation:* indicator for reporting being financially worse off than 12 months ago.
- *Presidential alignment:* indicator for whether the respondent's state voted for the sitting president's party in the most recent presidential election.
- *Lifetime macroeconomic experiences:* averages of national unemployment and CPI inflation over the respondent's lifetime, following the experience-effects literature (Malmendier and Nagel, 2011, 2016).
- *Spending shares:* shares of reported monthly household spending on medical expenses, mortgage payments, and recreation; essential and non-essential spending refer to the survey's 12-month-ahead spending-composition questions.
- *Numeracy:* score from the survey's numeracy module, measuring the ability to solve simple probability and interest-rate problems.

A.4 Data quality

Figure A1 plots median forecast errors by panel tenure for inflation and government debt: errors do not vary systematically with tenure beyond the first months in the panel, consistent with modest learning and a good understanding of the questions. Figure A2 reports histograms of reported expectations: most respondents report non-extreme values.

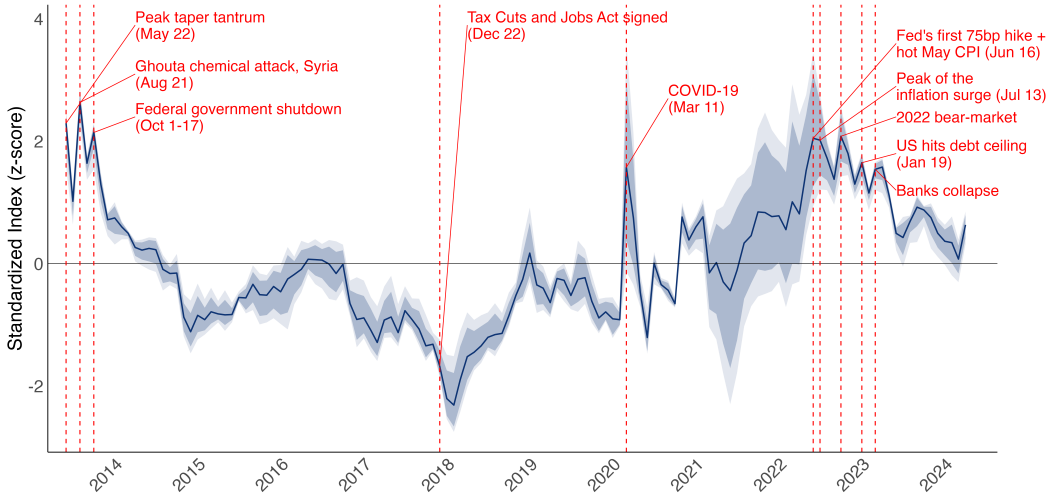
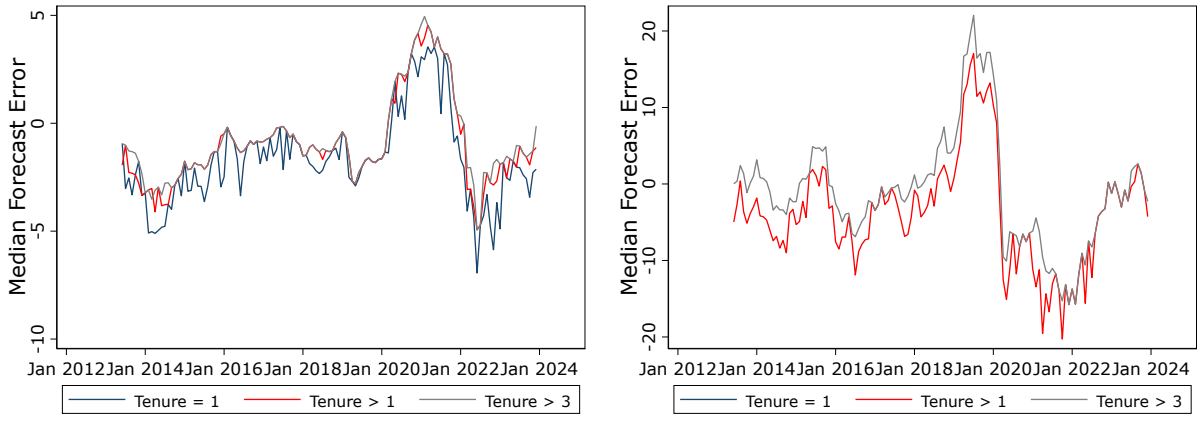


Figure 10: Aggregate Pessimism Index: Baseline Model

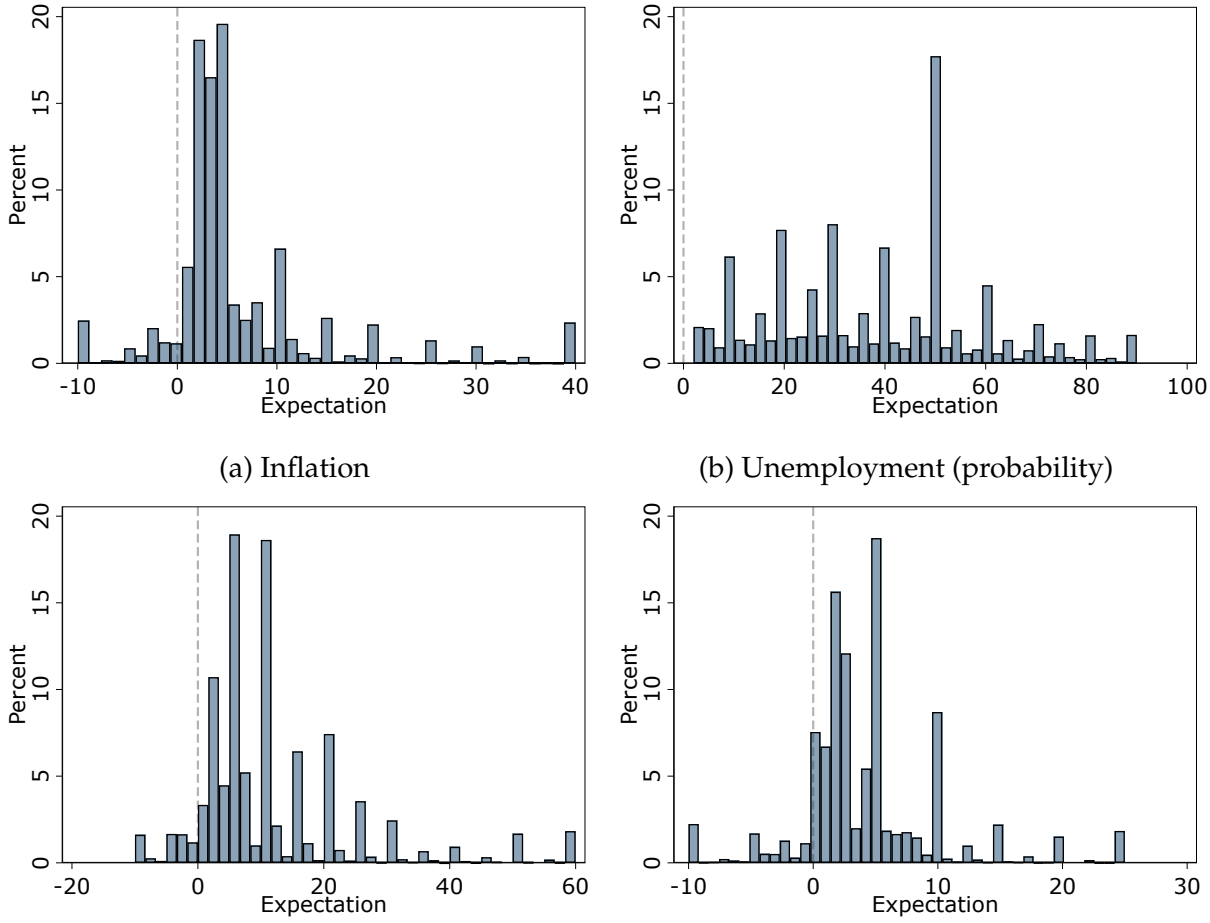
Notes: National monthly pessimism index, constructed by projecting state-month average beliefs onto the baseline-model pessimism coordinate for each accepted draw and aggregating across states with cell weights. Standardized units. Shaded areas: 68 and 90 percent posterior bands across draws.



(a) Inflation

(b) Government debt

Figure A1: Median Forecast Errors by Panel Tenure



(a) Inflation

(b) Unemployment (probability)

(c) Government debt

(d) Tax changes

Figure A2: Histograms of Reported Expectations

B The Hierarchical Random-Effects Covariance Model

This appendix replaces the independent-observation baseline in Section 4.2 with a Bayesian random-effects model that accounts for the SCE's rotating-panel structure. Let $y_{it} \in \mathbb{R}^K$ denote the belief vector of household i in survey month t, with $K = 3$ (unemployment, inflation, and tax expectations), and let $c(i,t)$ denote the year-month-by-state cell of observation (i,t). The model is

$$y_{it} \;=\; \delta_{c(i,t)} \;+\; \alpha_i \;+\; \tau_{it}, \qquad \alpha_i \sim N(0, \Omega_\alpha), \qquad \tau_{it} \sim N(0, \Sigma_\tau), \qquad \alpha_i \perp \tau_{it},$$

where δ_c is a cell fixed effect, α_i captures permanent heterogeneity in beliefs, and τ_{it} captures transitory month-to-month variation around each household's level. The total within-cell covariance is

$$\Sigma_{\tilde{B}} \;=\; \Omega_\alpha + \Sigma_\tau,$$

and is the object passed to the sign-restriction step.

Why random effects. Treating the residualized household-month rows as independent leaves the point estimate of $\Sigma_{\tilde{B}}$ unbiased, but overstates the effective sample size by a design effect $1 + (\bar{k} - 1)\rho$, where $\bar{k}$ is the average number of observations per respondent and ρ the intraclass correlation, so posterior bands are too narrow. A person fixed-effects specification is not an alternative here: with at most twelve observations per respondent, the incidental-parameters problem makes it inconsistent, and it would mechanically remove the permanent component of belief heterogeneity, which is part of the dispersion the framework is designed to explain. The distributional assumption on α_i is what identifies Ω_α and separates it from Σ_τ.

Priors. We use flat priors on both covariance matrices, $p(\Sigma_\tau) \propto 1$ and $p(\Omega_\alpha) \propto 1$. Let T denote the number of household-month observations, C the number of cells, and N the number of unique households. The flat priors imply inverse-Wishart posterior degrees of freedom $\nu_\tau = T - C - K - 1$ and $\nu_\alpha = N - K - 1$. Both priors are improper, so posterior propriety is not automatic in hierarchical models (Hobert and Casella, 1996). In our setting the conditions for propriety are comfortably met: the inverse-Wishart conditionals are proper whenever $\nu_\tau, \nu_\alpha > K - 1$, which holds by orders of magnitude given the panel dimensions, and tens of thousands of households are observed in multiple waves, so Ω_α is identified separately from Σ_τ by the within-household covariance of repeated observations.

Gibbs sampler. The joint posterior of $(\{\alpha_i\}, \Omega_\alpha, \Sigma_\tau)$ admits a Gibbs sampler with conjugate full conditionals. Each iteration cycles through four blocks.

1. *Cell fixed effects.* Given the current $\{\alpha_i\}$, recompute

$$\hat{\delta}_c \;=\; \frac{1}{n_c} \sum_{(i,t)\in c} \big(y_{it} - \alpha_i\big), \qquad y^*_{it} \;=\; y_{it} - \hat{\delta}_{c(i,t)}.$$

   Updating δ_c inside each Gibbs step, rather than fixing it at a pre-sample estimate, is the key correction: the Frisch-Waugh demeaning that is exact under independence no longer removes the fixed effects once α_i enters the model, because the household effects are entangled with the cell means. This block sets δ_c to its conditional posterior mean rather than drawing it from the full conditional $N(\hat{\delta}_c, \Sigma_\tau / n_c)$; the omitted draw noise is of order $1/n_c$ per cell, and the degrees-of-freedom correction $\nu_\tau = T - C - K - 1$ removes the associated $C \times K$ estimated means from the effective sample size.

2. *Household effects.* For each household, α_i has a Gaussian full conditional with variance $V_i = (\Omega_\alpha^{-1} + T_i\,\Sigma_\tau^{-1})^{-1}$ and mean $V_i \Sigma_\tau^{-1} \sum_t y^*_{it}$, where T_i is the number of waves in which household i is observed. The posterior mean is a shrinkage estimator: households observed few times are shrunk strongly toward zero.

3. *Transitory covariance.* With $\hat{\tau}_{it} = y^*_{it} - \alpha_i$,

$$\Sigma_\tau \mid \cdot \;\sim\; \mathrm{IW}\Big(\textstyle\sum_{i,t} \hat{\tau}_{it}\hat{\tau}_{it}',\; T - C - K - 1\Big).$$

4. *Persistent covariance.*

$$\Omega_\alpha \mid \cdot \;\sim\; \mathrm{IW}\Big(\textstyle\sum_{i} \alpha_i \alpha_i',\; N - K - 1\Big).$$

Sign-restriction step. After a burn-in period, every fifth Gibbs draw we form $\Sigma_{\tilde{B}}^{(m)} = \Omega_\alpha^{(m)} + \Sigma_\tau^{(m)}$, draw a Haar-uniform orthogonal matrix $Q^{(m)}$, and set $\Gamma^{(m)} = \mathrm{chol}(\Sigma_{\tilde{B}}^{(m)})\, Q^{(m)}$. The draw is kept if $\Gamma^{(m)}$ satisfies the sign restrictions and if the recovered scores satisfy the narrative restriction,

$$\mathbb{E}\big[e^{(m)}_{it,\text{pess}} \mid \text{exposed}\big] \;\geq\; \mathbb{E}\big[e^{(m)}_{it,\text{pess}} \mid \text{non-exposed}\big], \qquad e^{(m)}_{it} = \Gamma^{(m)-1} y^*_{it}.$$

Sampling continues until 2,000 draws are accepted. The scores are computed from the total within-cell residuals $y^*_{it} = \alpha_i + \tau_{it}$ rather than from the transitory part alone: $\text{Var}(y^*_{it}) = \Sigma_{\tilde{B}}$, so the scores are orthonormal, $\text{Var}(e^{(m)}_{it}) = I$, and they carry the permanent belief heterogeneity that the anatomy regressions explain with household characteristics. One implication is worth stating explicitly: because the scores contain $\Gamma^{-1}\alpha_i$, they are serially dependent within household, with $\text{Corr}(e_{it}, e_{is}) = \Gamma^{-1}\Omega_\alpha(\Gamma^{-1})'$ for $t \neq s$. The identity covariance in Assumption 1 is a contemporaneous normalization; it does not assert independence of the scores over time.

What the accepted draws represent. The algorithm attempts one rotation per thinned Gibbs draw and keeps the pair $(\Sigma_{\tilde{B}}, Q)$ when the restrictions hold. The accepted draws are therefore distributed as the posterior of $\Sigma_{\tilde{B}}$ tilted by the acceptance probability of each covariance draw, combined with the uniform distribution over each draw's admissible rotation set. This is the standard convention in the sign-restriction literature (Uhlig, 2005; Rubio-Ramírez et al., 2010). Antolín-Díaz and Rubio-Ramírez (2018) show that narrative restrictions make the tilt informative and propose importance weights, proportional to the inverse of each covariance draw's acceptance probability, computed by attempting many rotations per draw; applying their reweighting in our setting is a direct extension of the algorithm above.

What is and is not assumed. The Gaussian assumptions serve two distinct purposes that are useful to separate. For point estimates, they are inessential: $\Sigma_{\tilde{B}}$ is a second moment, its posterior mode coincides with a method-of-moments estimator, and the variance and covariance decompositions inherit consistency without normality. For the posterior bands, the Gaussian likelihood matters: reported survey beliefs are heaped at focal values and have heavy tails, and under excess kurtosis a Gaussian likelihood understates the sampling variability of covariance estimates. The bands should therefore be read as model-based credible sets rather than distribution-free confidence sets. The same reading applies to the second-stage figures: the ribbons report the distribution of per-draw point estimates across the 2,000 accepted draws, which propagates the posterior uncertainty about $\Sigma_{\tilde{B}}$ and the rotation into every downstream statistic, and does not additionally include the within-draw sampling error of the second-stage regressions.

Relation to the independence baseline. The model nests the baseline that treats rows as independent as the limit $\Omega_\alpha \to 0$. Ignoring the panel structure leaves the point estimate of $\Sigma_{\tilde{B}}$ unbiased but overstates the effective sample size by the design effect $1 + (\bar{k} - 1)\rho$,

where $\bar{k}$ is the average number of waves per respondent and ρ the intraclass correlation of residual beliefs. In the data, the point estimates of $\Sigma_{\tilde{\beta}}$ and Γ are essentially unchanged relative to that baseline, while the posterior bands widen to reflect within-household clustering.

B.1 Hierarchical-model structural results

Figures A3–A5 report the core structural results under the hierarchical covariance model. The point estimates remain close to the independent-observation baseline, while the posterior bands widen to reflect persistent within-household correlation. The substantive conclusions are unchanged: pessimism has the largest unemployment loading, accounts for the largest share of unemployment and tax-belief variance, and is the dominant positive contributor to comovement between unemployment and inflation expectations.

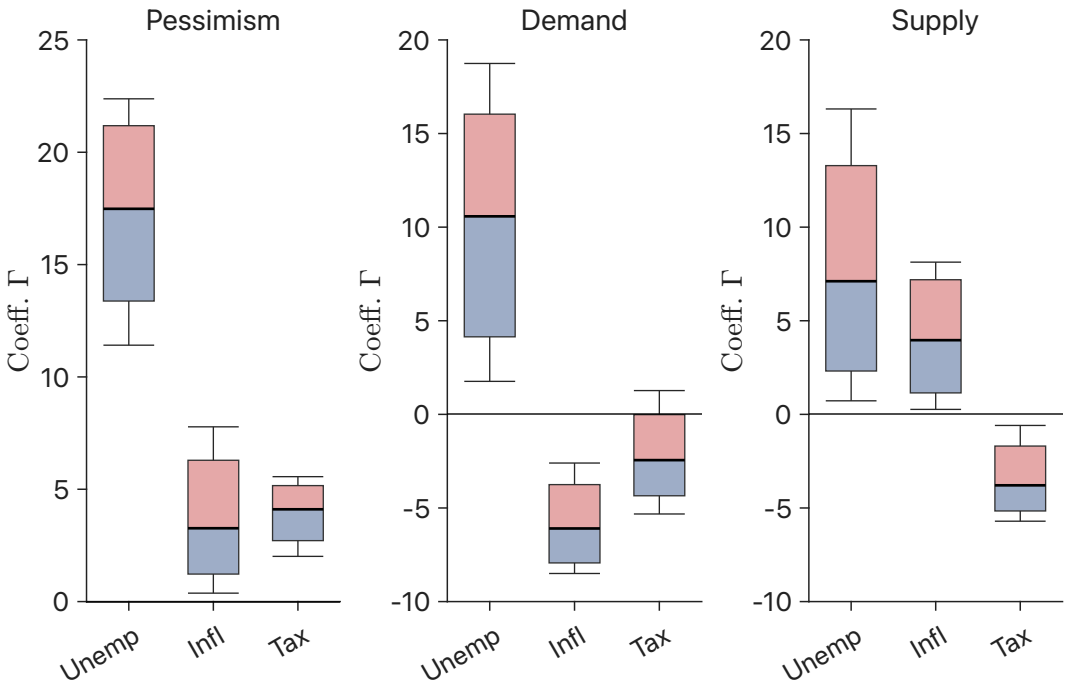


Figure A3: Contemporaneous Responses: Hierarchical Model

Notes: Each panel reports the column of Γ corresponding to one component under the hierarchical random-effects covariance model. Medians, 68 and 90 percent posterior bands across accepted draws.

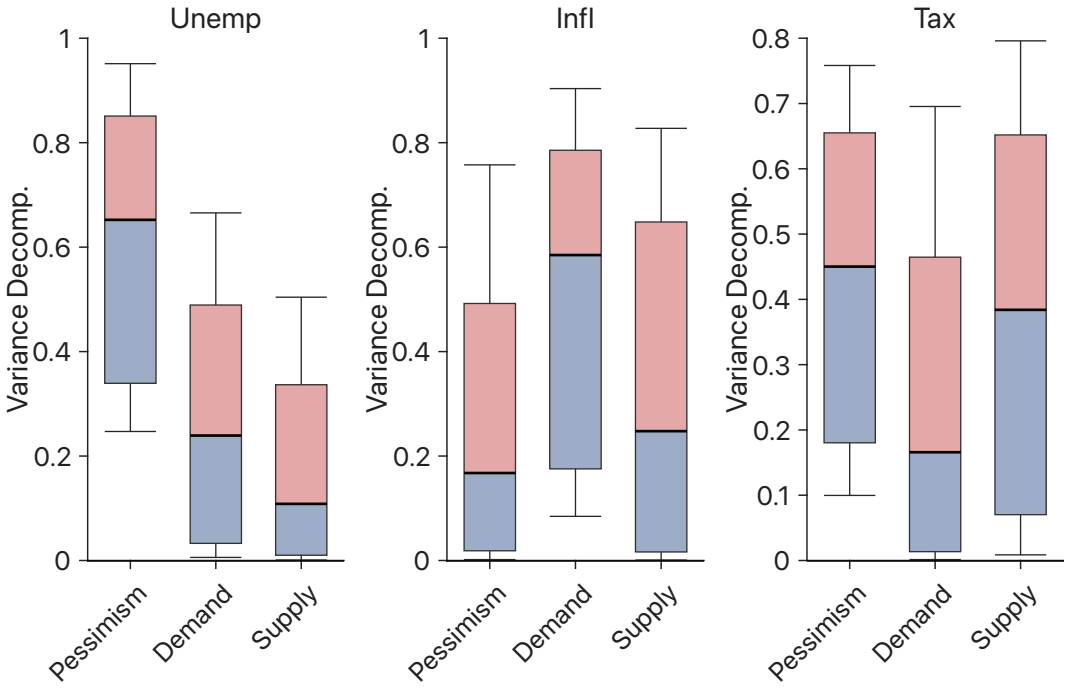


Figure A4: Variance Decomposition: Hierarchical Model

Notes: Share of the within-cell variance of each belief attributable to each component under the hierarchical random-effects covariance model. Medians, 68 and 90 percent posterior bands across accepted draws.

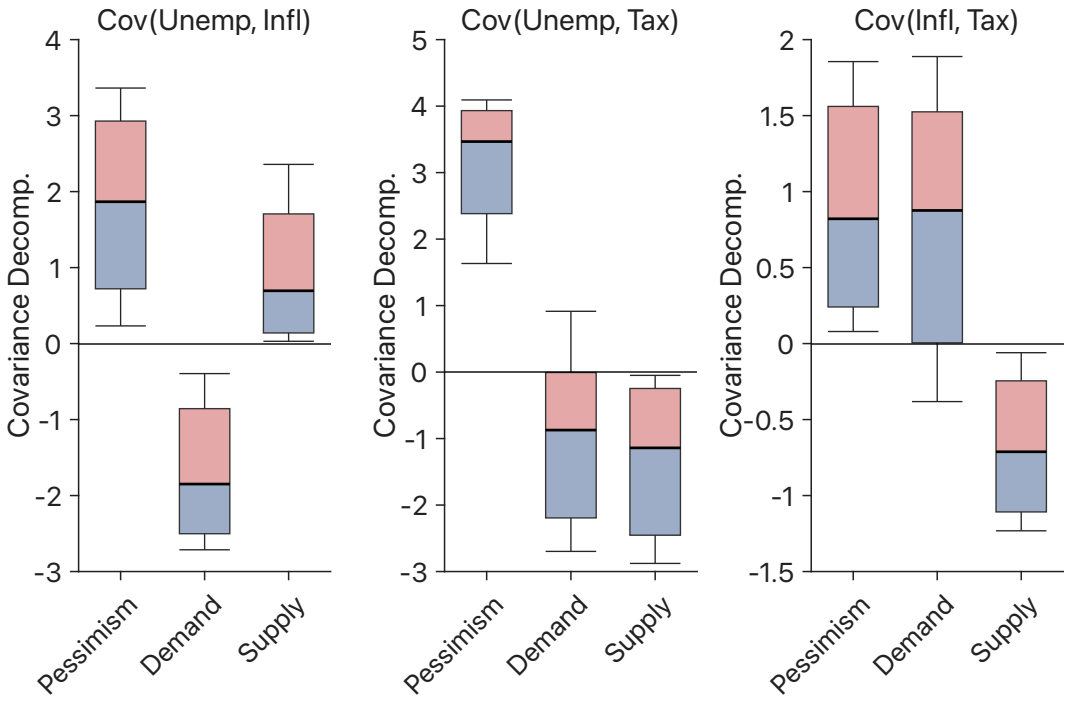


Figure A5: Covariance Decomposition: Hierarchical Model

Notes: Share of each pairwise within-cell belief covariance attributable to each component under the hierarchical random-effects covariance model. Medians, 68 and 90 percent posterior bands across accepted draws.

B.2 Robustness to alternative specifications

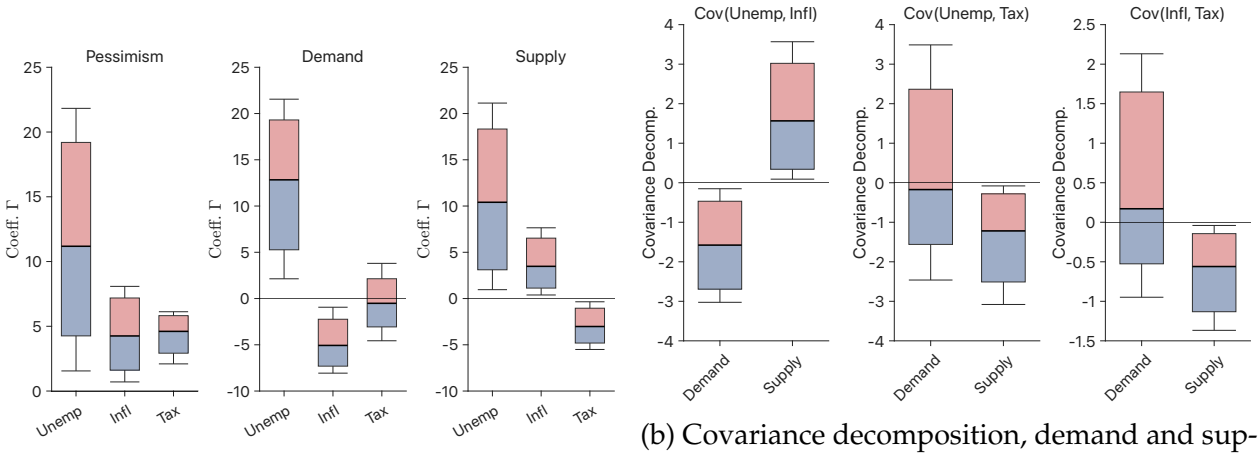
Because the model is set-identified, we assess the robustness of the hierarchical-model results to two alternative specifications, both estimated under the baseline (independent-observations) covariance model. The first removes the narrative restriction while keeping the sign restrictions, which isolates the identifying content of the mass-shooting event. The second imposes the demand and supply sign columns and leaves the remaining column unrestricted, which asks how the covariance matrix is interpreted when the pessimism component is omitted. Figure A6 reports the two exercises.

In the first specification, estimated under the baseline covariance model without the narrative restriction, the sign restrictions continue to select the same qualitative impact matrix, while the posterior bands widen. The 68 percent band for the response of unemployment beliefs to pessimism widens from roughly $[13, 21]$ probability points in the hierarchical specification to roughly $[4, 19]$, and the bands for the variance and covariance shares widen as well. The narrative restriction therefore carries identifying content. It removes admissible rotations in which the pessimism component is too small to match the higher pessimism realizations of respondents exposed to mass shootings. This tightening addresses a standard concern in sign-restricted models: posterior bands can partly reflect the uniform prior over rotations rather than information in the data (Baumeister and Hamilton, 2015). The comparison between the sign-only specification and the hierarchical specification measures the information added by the narrative event.

The second specification shows why disciplining pessimism matters for the economic interpretation of the covariance matrix. We impose the demand and supply sign columns under the baseline covariance model and leave the remaining column unrestricted. In the unemployment-inflation covariance, demand contributes negatively because it raises expected unemployment and lowers expected inflation, while supply contributes positively because it raises both expectations. When the pessimism column is unrestricted, the model assigns a larger share of this positive covariance to supply: the median supply share is about 1.5, compared with about 0.7 in the hierarchical three-component specification.

The fiscal covariances discipline this interpretation. The observed unemployment-tax covariance is positive, yet both demand and supply have non-positive median contributions in this specification, about -0.2 for demand and about -1.2 for supply: supply cannot account for this moment because it raises expected unemployment, and if anything both components work against the observed comovement. The unrestricted column must therefore account not only for the full covariance but also for offsetting the negative contributions of demand and supply. The inflation-tax covariance points in the same

direction. These moments require a component that moves expected unemployment, expected inflation, and expected taxes in a common adverse direction. The three-component specification labels this residual pessimism, preserving the intended interpretation of demand and supply.



(a) Impact responses, sign restrictions only

(b) Covariance decomposition, demand and supply only

Figure A6: Robustness to Alternative Specifications

Notes: Panel (a) estimates the baseline (independent-observations) covariance model without the narrative restriction, keeping the sign restrictions. Panel (b) imposes only the demand and supply sign columns under the baseline covariance model and leaves the third column unrestricted; the plotted shares are those of the demand and supply components, so the residual component absorbs the remainder. Medians, 68 and 90 percent posterior bands.

B.3 External Validation of the Pessimism Index

Section 7 constructs a national monthly pessimism index by projecting state-month average beliefs onto the baseline-model pessimism coordinate. Figure A7 compares that index with the University of Michigan's Consumer Sentiment index.

C Sign Restrictions in a New Keynesian Model with a Fiscal Block

This appendix derives the sign pattern of H_{RE} used in Section 3.2 from a stylized static New Keynesian environment with a fiscal block. The purpose is to make explicit the equilibrium logic behind the sign restrictions, in particular the restriction that adverse supply forces lower expected taxes while pessimism raises them.

Let d denote an adverse demand disturbance, s an adverse supply (cost-push) disturbance, and s^τ an exogenous tax disturbance, all mean zero. The equilibrium block is the static version of the three-equation model, with an Okun mapping from the output gap x to unemployment u:

$$\pi = \kappa x + s, \qquad x = -\sigma\, i - d, \qquad i = \phi_\pi \pi, \qquad u - \bar{u} = -\omega x,$$

with $\kappa, \sigma, \omega > 0$ and $\phi_\pi \geq 0$. Substituting the policy rule and the IS relation into the Phillips curve and solving,

$$\pi = \frac{s - \kappa d}{\Delta}, \qquad x = -\frac{\sigma\phi_\pi s + d}{\Delta}, \qquad u - \bar{u} = \omega\, \frac{\sigma\phi_\pi s + d}{\Delta}, \qquad \Delta \equiv 1 + \kappa\sigma\phi_\pi > 0.$$

Both adverse disturbances raise unemployment. Inflation rises with the supply disturbance and falls with the demand disturbance. These are the standard comparative statics of the three-equation model under an interest-rate rule that does not fully offset either disturbance.

The fiscal block combines automatic stabilizers with the effect of inflation on the real debt burden. Expected tax obligations at a given household income respond to the state of the economy through two channels,

$$\tau = -\chi\,(u - \bar{u}) + \eta\, b + s^\tau, \qquad b = b_0 - \mu\, \pi, \qquad \chi, \eta, \mu > 0.$$

The first channel is the stabilizer response: a deterioration in activity lowers expected

taxes over the next year, because recessions trigger tax credits, transfer expansions, and delayed consolidation rather than contemporaneous tax increases. The second channel is the perceived financing burden b of outstanding nominal debt, which future taxes must service and which inflation erodes.

The two channels reinforce each other for supply disturbances and oppose each other for demand disturbances. Substituting the equilibrium solutions into the fiscal rule gives the tax responses

$$\frac{\partial \tau}{\partial s} = -\frac{\chi\omega\sigma\phi_\pi + \eta\mu}{\Delta} < 0, \qquad \frac{\partial \tau}{\partial d} = \frac{\eta\mu\kappa - \chi\omega}{\Delta} \lesseqgtr 0.$$

An adverse supply disturbance raises u and raises π: expected taxes fall through the stabilizer and fall again because inflation erodes the real debt burden, so the supply-tax entry is negative for all parameter values. An adverse demand disturbance raises u but lowers π: the stabilizer pushes expected taxes down while the heavier real debt burden pushes them up, and the sign turns on $\eta\mu\kappa$ against $\chi\omega$. A household with rational expectations who assigns posterior means $(\tilde{s}^d, \tilde{s}^s, \tilde{s}^\tau)$ to the three disturbances then expects

$$H_{RE} = \frac{1}{\Delta}\begin{bmatrix} \omega\sigma\phi_\pi & \omega & 0 \\ 1 & -\kappa & 0 \\ -(\chi\omega\sigma\phi_\pi + \eta\mu) & \eta\mu\kappa - \chi\omega & \Delta \end{bmatrix} \sim \begin{bmatrix} + & + & 0 \\ + & - & 0 \\ - & ? & + \end{bmatrix},$$

with columns ordered as supply, demand, and tax states and rows as unemployment, inflation, and taxes. The first two columns reproduce the sign pattern imposed on the supply and demand components, including the unrestricted demand-tax entry. The pessimism component has no counterpart in this equilibrium mapping: it is a belief distortion that shifts perceived outcomes adversely across the board, including the perceived fiscal burden, and its positive loading on expected taxes is what distinguishes it from the equilibrium response to an adverse supply disturbance.

D Encompassing Representation: Examples

This appendix illustrates the flexibility of the representation $B_{i,t} = H_i \tilde{s}_{i,t}$ introduced in Section 3.2, in which $\tilde{s}_{i,t}$ denotes the household's perceived state and H_i the household-specific mapping from perceived states to observable beliefs. The setup encompasses a wide class of environments in which disagreement arises either because households observe different signals or because they interpret similar information through differ-

ent subjective models. It cleanly nests standard rational-expectations New Keynesian models, noisy-information and dispersed-information models, sticky-information environments, and rational-inattention formulations, since in all of these cases beliefs can be represented as functions of a low-dimensional perceived state. It also accommodates heterogeneous subjective models in which households differ in the causal narratives they apply to the economy, and behavioral distortions such as myopia, diagnostic expectations, or experience-based belief formation, provided these mechanisms operate through the perceived state $\tilde{s}_{i,t}$, the mapping H_i, or both.

Example 1 (Dispersed information with a common subjective model)**.** *Suppose there is a vector of aggregate states s_t, and household i observes a noisy private signal*

$$x_{i,t} = s_t + v_{i,t},$$

where $v_{i,t}$ is mean-zero idiosyncratic noise. Households share a common mapping $H_i = H$, but differ in the posterior mean they form about the underlying state. Let $\tilde{s}_{i,t} = \mathbb{E}_i[s_t \mid x_{i,t}]$ denote household i's posterior. Beliefs then satisfy

$$B_{i,t} = H\tilde{s}_{i,t}.$$

In this environment, cross-sectional disagreement arises because households receive different signals and therefore hold different posterior beliefs about the same aggregate state, even though they share the same subjective model of how states map into observable outcomes.

In this case, the framework captures a standard dispersed-information mechanism. Heterogeneity operates through $\tilde{s}_{i,t}$, while the mapping H is common across households. The estimated latent components therefore reflect differential exposure to, or differential extraction of, information about the underlying aggregate state rather than disagreement about the structure of the economy itself.

Example 2 (Common information with heterogeneous subjective models)**.** *Suppose all households observe the same perceived state $\tilde{s}_t$, but differ in the mapping they use to translate this state into beliefs:*

$$B_{i,t} = H_i\tilde{s}_t.$$

A convenient representation is

$$H_i = \sum_{\theta \in \{d,s,p\}} \lambda_\theta(i)\, h^\theta(a^\theta(i))',$$

where h^θ is the impact-response of beliefs to view θ, $a^\theta(i)$ determines which combination of states household i interprets as evidence for that view, and $\lambda_\theta(i)$ is the weight household i assigns to it. Notice that in the previous formulation λ was absorbed into a^θ. In this environment, households may differ because some interpret the same information through a more demand-side view, others through a more supply-side view, and others through a broader pessimistic view.

Here, disagreement arises from heterogeneity in these views rather than from heterogeneity in information. The framework thus encompasses environments in which households share the same information set but differ in the causal view they apply to the economy. This interpretation is particularly natural when one wishes to relate estimated components to observable characteristics such as income, employment status, education, or financial fragility.

Example 3 (Myopia as a distortion of the benchmark mapping)**.** *Suppose the benchmark rational-expectations mapping is H^{RE}, obtained from a solved New Keynesian model. Household i's beliefs are instead given by*

$$B_{i,t} = H^{RE} M_i \tilde{s}_{i,t},$$

where M_i is a household-specific distortion matrix and recall $B_{i,t}$ contains next period's beliefs. In the simplest case,

$$M_i = \mathrm{diag}(m_i^s, m_i^p, m_i^\tau), \qquad 0 \le m_i^k \le 1,$$

so that each m_i^k governs the degree to which household i attends to, or reacts to, the corresponding underlying state. Lower values of m_i^k capture attenuated responses relative to the benchmark rational-expectations mapping.

This formulation nests Gabaix-style myopia in a particularly transparent way. Relative to the benchmark model, households may underreact to some dimensions of the aggregate state while responding more strongly to others. In this sense, myopia can be interpreted as a structured restriction on H_i, rather than as a fundamentally different object. More generally, the same logic can be extended to other behavioral distortions, including selective attention, diagnostic expectations, or experience-based updating, provided these operate through distortions in the perceived state $\tilde{s}_{i,t}$, the mapping H_i, or both.

Taken together, these examples illustrate that the framework is intentionally agnostic about the ultimate source of belief heterogeneity. The same reduced-form object may reflect dispersed information, heterogeneous subjective models, or structured distortions of a benchmark rational-expectations mapping. Our empirical strategy therefore aims to recover the low-dimensional directions organizing disagreement in the data, and only subsequently interpret these directions through the lens of particular theoretical environments.

At the same time, it is important to distinguish representation from identification of these specific models. Although the framework can nest a broad range of models of expectation formation, the baseline empirical strategy does not, by itself, sharply distinguish among them without imposing further assumptions. An estimated supply-type component, for instance, may reflect differential exposure to supply-related information, a supply-side view, or an attention distortion of the kind emphasized by behavioral myopia models. Accordingly, the framework is best understood as providing a common language for organizing cross-sectional disagreement, rather than as uniquely identifying one specific theory of expectations absent additional structure or further empirical restrictions.

By contrast, this framework is less naturally suited to environments in which nonlinearities, regime changes, tail risk, evolving laws of motion, or higher moments are the central primitives. Rare-disaster models are the clearest example. In those settings, beliefs typically depend on time-varying disaster probabilities, extreme downside states, or shifts in uncertainty that are not well summarized by the low-dimensional linear state space considered here.

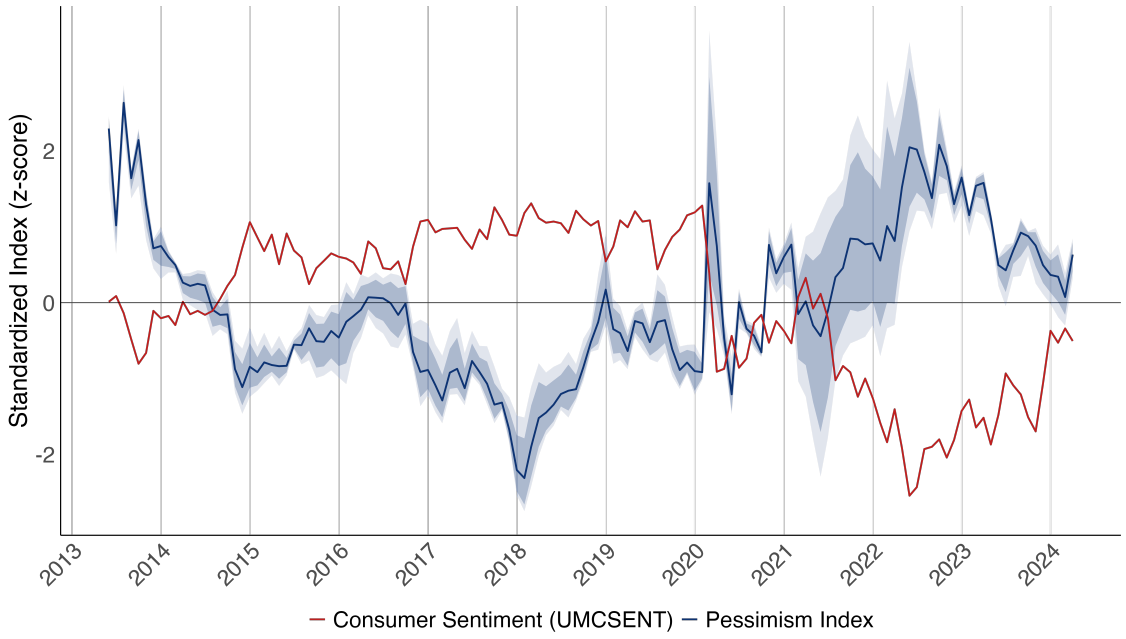


Figure A7: Pessimism Index versus Consumer Sentiment: Baseline Model

Notes: Standardized national baseline-model pessimism index (blue, with 68 and 90 percent posterior bands) and the standardized University of Michigan Consumer Sentiment index (red). The correlation between the two series is -0.76.

E Additional Reduced-Form Evidence

Table A1: Reduced-Form Perceived Partial Covariances with Demographic Controls

	Dep. Var: $B_{it}(k)$ **(Columns)**			
$B_{it}(m)$ (Row)	Gov. Debt	Tax Changes	Inflation	Unemp.
Unemployment	1.2383^{***}	0.3716^{***}	0.5351^{***}	
	(0.0628)	(0.0289)	(0.0353)	
Tax Changes	2.5415^{***}		1.2083^{***}	1.5071^{***}
	(0.0751)		(0.0496)	(0.1167)
Inflation	4.0402^{***}	1.0880^{***}		1.9541^{***}
	(0.0936)	(0.0430)		(0.1268)
Gov Debt		1.3483^{***}	2.3802^{***}	2.6644^{***}
		(0.0401)	(0.0626)	(0.1337)
N	120609	120609	120609	120609
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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A2: Reduced-Form Perceived Partial Covariances with Demographic and Indiv. FE.

	Dep. Var: $B_{it}(k)$ **(Columns)**			
$B_{it}(k')$ (Row)	Gov. Debt	Tax Changes	Inflation	Unemp.
Unemployment	0.6672^{***}	0.1505^{***}	0.5210^{***}	
	(0.0457)	(0.0233)	(0.0320)	
Tax Changes	1.2296^{***}		0.5277^{***}	0.5177^{***}
	(0.0543)		(0.0392)	(0.0801)
Inflation	1.4177^{***}	0.4893^{***}		1.6622^{***}
	(0.0709)	(0.0364)		(0.1015)
Gov Debt		0.8398^{***}	1.0441^{***}	1.5676^{***}
		(0.0365)	(0.0524)	(0.1065)
N	119186	119186	119186	119186
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 A3: Reduced-Form Perceived Partial Covariances: Pre-COVID Sample

	Dep. Var: $B_{it}(k)$ **(Columns)**			
$B_{it}(m)$ (Row)	Gov. Debt	Tax Changes	Inflation	Unemp.
Unemployment	1.3674^{***}	0.4042^{***}	0.4499^{***}	
	(0.0784)	(0.0381)	(0.0422)	
Tax Changes	2.2271^{***}		1.0007^{***}	1.6099^{***}
	(0.0955)		(0.0601)	(0.1510)
Inflation	4.1831^{***}	1.0480^{***}		1.8768^{***}
	(0.1298)	(0.0602)		(0.1712)
Gov Debt		1.3082^{***}	2.3460^{***}	3.1994^{***}
		(0.0562)	(0.0838)	(0.1833)
N	69716	69716	69716	69716
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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A4: Reduced-Form Perceived Partial Covariances: Post-COVID Sample

	Dep. Var: $B_{it}(k)$ **(Columns)**			
$B_{it}(m)$ (Row)	Gov. Debt	Tax Changes	Inflation	Unemp.
Unemployment	1.0713^{***}	0.3199^{***}	0.6656^{***}	
	(0.0992)	(0.0436)	(0.0583)	
Tax Changes	2.9160^{***}		1.4785^{***}	1.3252^{***}
	(0.1171)		(0.0808)	(0.1802)
Inflation	3.8447^{***}	1.1295^{***}		2.1064^{***}
	(0.1329)	(0.0606)		(0.1834)
Gov Debt		1.3774^{***}	2.3773^{***}	2.0963^{***}
		(0.0562)	(0.0921)	(0.1919)
N	50893	50893	50893	50893
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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A5: Correlation Matrices of Beliefs: Raw and Residualized

Panel A: Raw beliefs

	Gov. Debt	Taxes	Unemployment	Inflation
Gov. Debt	1.000	0.304	0.155	0.402
Taxes	0.304	1.000	0.123	0.270
Unemp.	0.155	0.123	1.000	0.121
Inflation	0.402	0.270	0.121	1.000

Panel B: Residualized beliefs

	Gov. Debt	Taxes	Unemployment	Inflation
GovDebt	1.000	0.295	0.160	0.389
Taxes	0.295	1.000	0.123	0.259
Unemp	0.160	0.123	1.000	0.127
Inflation	0.389	0.259	0.127	1.000

Notes: Residualized beliefs partial out state-by-month fixed effects.

Table A6: Reduced-Form Perceived Partial Covariances: All Aggregate Expectations

	Dep. Var: $B_{it}(k)$ **(Columns)**									
$B_{it}(k')$ (Row)	Govt. D.	Tax	Infl	Unemp	Int. Rate	Stocks	LR Infl	Hard Cred.	Food	Gas
Unemp.	1.1289***	0.2841***	0.2529***		5.0521***	2.3086***	0.0868***	0.0660***	0.2252***	0.4080***
	(0.0620)	(0.0294)	(0.0276)		(0.1357)	(0.1216)	(0.0291)	(0.0024)	(0.0264)	(0.0417)
Tax	1.6189***		0.5150***	1.0747***	-0.1807	-0.4840***	0.2698***	0.0157***	0.6930***	0.5128***
	(0.0700)		(0.0384)	(0.1112)	(0.1172)	(0.1035)	(0.0404)	(0.0022)	(0.0319)	(0.0433)
Infl	2.2064***	0.6025***		1.1192***	0.3765***	-0.4651***	4.1973***	0.0099***	1.0474***	-0.0119
	(0.0923)	(0.0435)		(0.1207)	(0.1148)	(0.1114)	(0.0793)	(0.0026)	(0.0416)	(0.0503)
Int. Rate	-0.8606***	-0.0485	0.0864***	5.1309***		8.8400***	-0.0454	-0.0013	-0.1599***	0.2516***
	(0.0608)	(0.0314)	(0.0264)	(0.1343)		(0.1257)	(0.0289)	(0.0026)	(0.0273)	(0.0472)
Stocks	-0.3141***	-0.1408***	-0.1157***	2.5407***	9.5792***		-0.0274	-0.0544***	-0.1832***	-0.0230
	(0.0614)	(0.0300)	(0.0275)	(0.1337)	(0.1353)		(0.0299)	(0.0025)	(0.0268)	(0.0441)
LR Infl	1.1187***	0.2745***	3.6498***	0.3341***	-0.1722	-0.0957		0.0107***	0.4394***	0.2874***
	(0.0865)	(0.0413)	(0.0742)	(0.1117)	(0.1094)	(0.1045)		(0.0025)	(0.0395)	(0.0502)
Hard Cred.	0.5725***	0.3917***	0.2102***	6.2300***	-0.1245	-4.6588***	0.2627***		0.6607***	-0.2672***
	(0.1215)	(0.0554)	(0.0556)	(0.2317)	(0.2399)	(0.2149)	(0.0624)		(0.0512)	(0.0755)
Gas	0.6118***	0.4010***	-0.0080	1.2072***	0.7331***	-0.0620	0.2210***	-0.0084***	2.6758***	
	(0.0685)	(0.0341)	(0.0336)	(0.1239)	(0.1377)	(0.1186)	(0.0389)	(0.0024)	(0.0390)	
Food	3.1258***	0.9431***	1.2184***	1.1595***	-0.8104***	-0.8571***	0.5878***	0.0361***		4.6562***
	(0.0883)	(0.0430)	(0.0502)	(0.1353)	(0.1381)	(0.1248)	(0.0527)	(0.0028)		(0.0559)
Govt D.		0.9100***	1.0602***	2.4005***	-1.8020***	-0.6068***	0.6181***	0.0129***	1.2911***	0.4397***
		(0.0396)	(0.0496)	(0.1303)	(0.1266)	(0.1184)	(0.0473)	(0.0027)	(0.0380)	(0.0492)
N	120165	120165	120165	120165	120165	120165	120165	120165	120165	120165
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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

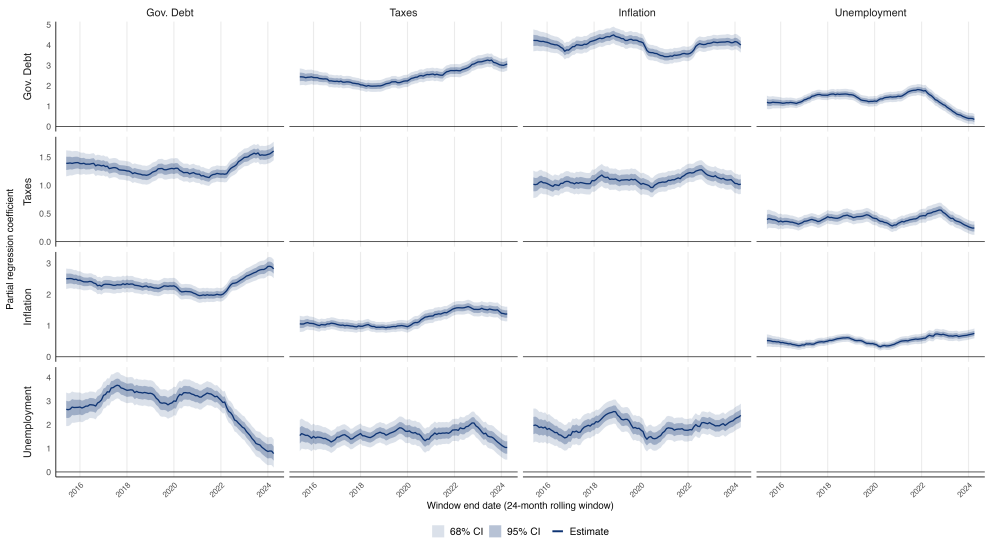


Figure A8: Rolling-Window Estimates of the Partial Associations

Notes: 24-month rolling-window estimates of equation (4). Rows correspond to the dependent variable.

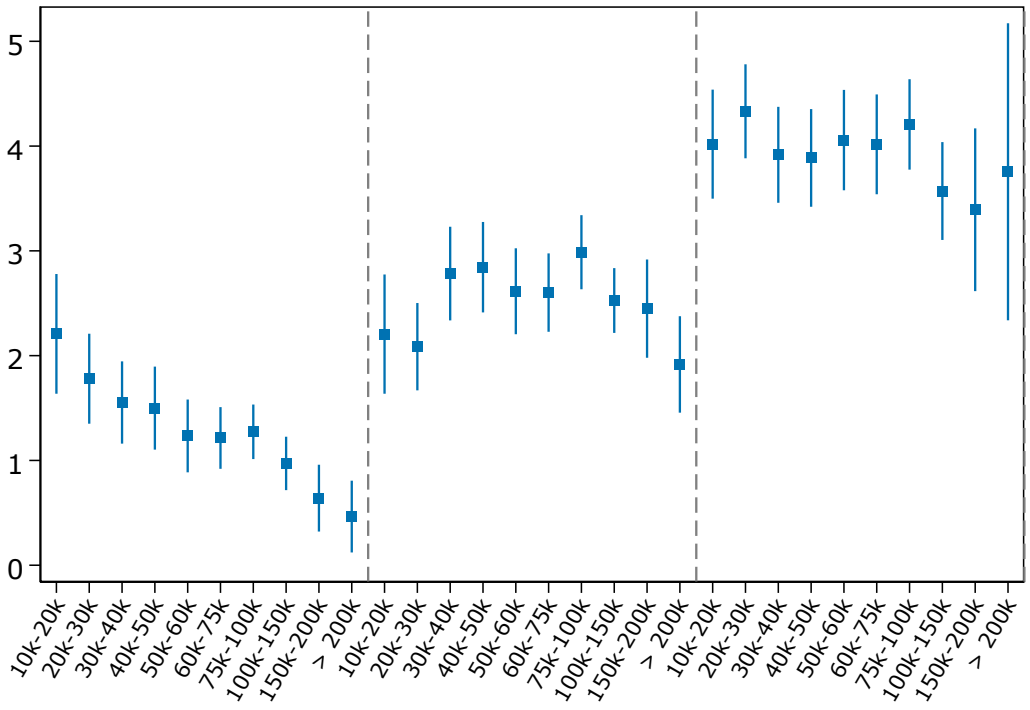


Figure A9: Heterogeneous Partial Associations by Income: Government Debt Expectations

Notes: Income-specific slopes from equation (5). The dependent variable is expected government debt growth. Dashed vertical lines separate blocks of right-hand-side expectations. Dots report point estimates; vertical bars report 95 percent confidence intervals.

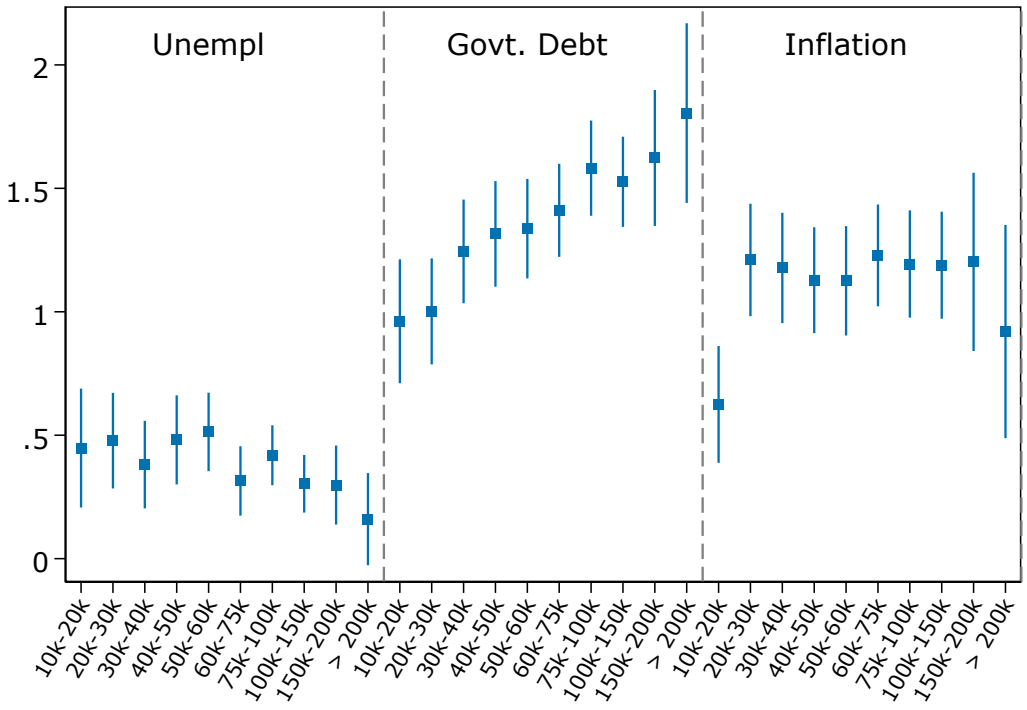


Figure A10: Heterogeneous Partial Associations by Income: Tax Expectations

Notes: Income-specific slopes from equation (5). The dependent variable is expected tax changes. Dashed vertical lines separate blocks of right-hand-side expectations. Dots report point estimates; vertical bars report 95 percent confidence intervals.

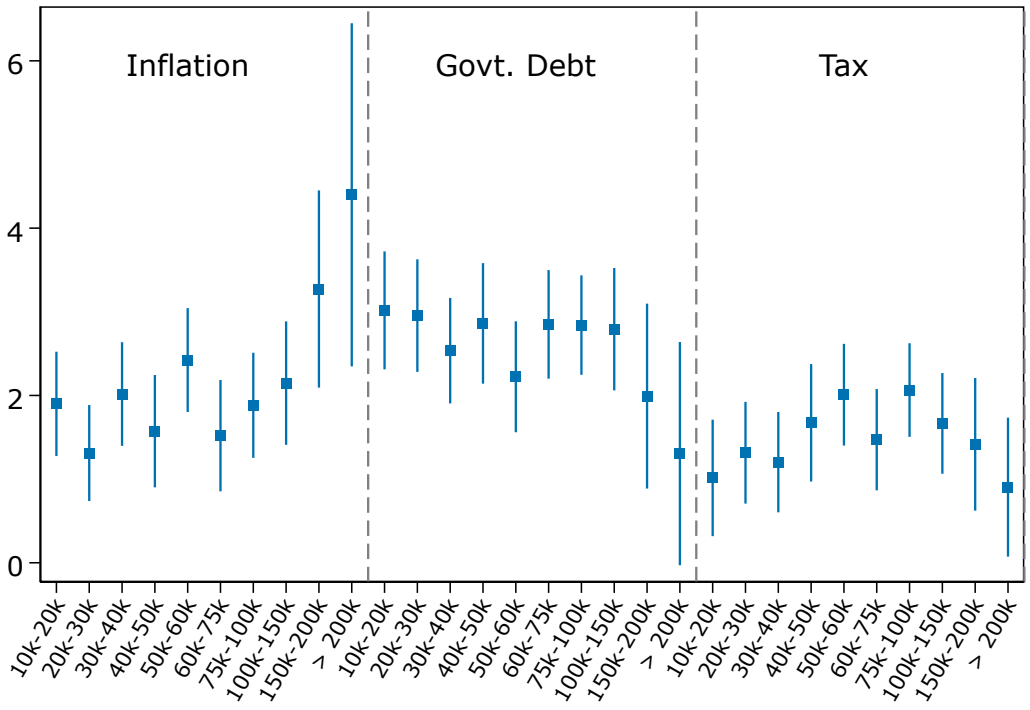
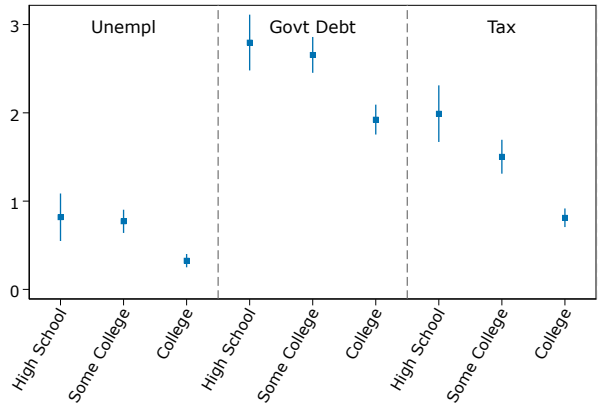
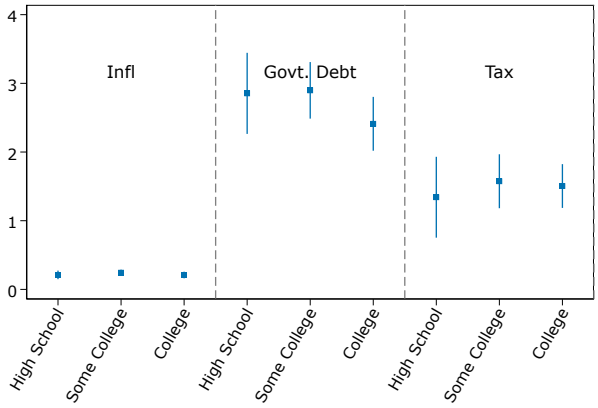


Figure A11: Heterogeneous Partial Associations by Income: Unemployment Expectations

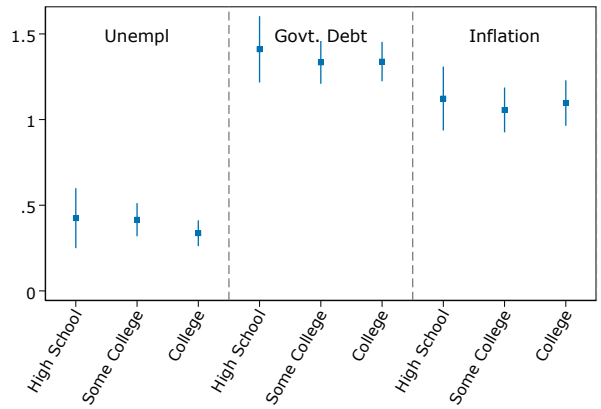
Notes: Income-specific slopes from equation (5). The dependent variable is the expected probability that unemployment rises over the next 12 months. Dashed vertical lines separate blocks of right-hand-side expectations. Dots report point estimates; vertical bars report 95 percent confidence intervals.



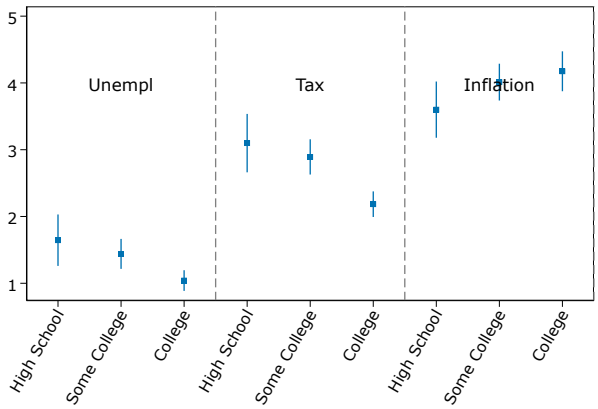
(a) Inflation as dependent variable



(b) Unemployment as dependent variable

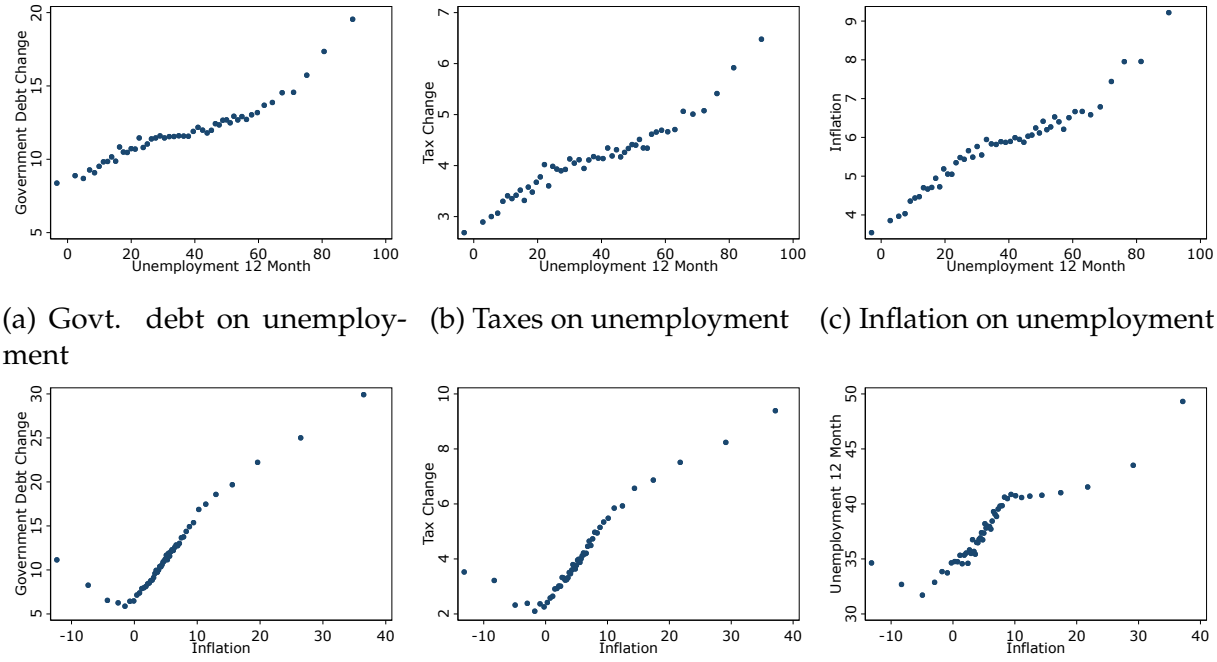


(c) Taxes as dependent variable



(d) Government debt as dependent variable

Figure A12: Heterogeneous Partial Associations by Education



(a) Govt. debt on unemployment (b) Taxes on unemployment (c) Inflation on unemployment

(d) Govt. debt on inflation (e) Taxes on inflation (f) Unemployment on inflation

Figure A13: Nonparametric Conditional Expectation Functions

Notes: Binned scatter plots controlling for year-month-by-state fixed effects and demographics. Adverse-direction comovement holds throughout, with some nonlinearity at the extremes.

F Structural Robustness and Additional Anatomy

F.1 Anatomy of demand and supply views

Figures A14 to A17 complete the anatomy of Section 6. Three patterns stand out. First, the gradients that define pessimism are absent for the other views: demand and supply scores are unrelated to financial deterioration, income volatility and presidential alignment, the three strongest pessimism correlates. Second, the experience channel splits cleanly by view: lifetime unemployment experience predicts pessimism while lifetime inflation experience predicts supply views. Third, demand views carry the signature of discretionary consumption: they rise with income, education, recreation spending, and self-reported willingness to take financial risk, while supply views show none of these patterns and are concentrated among the unemployed, old, and respondents out of the labor force.

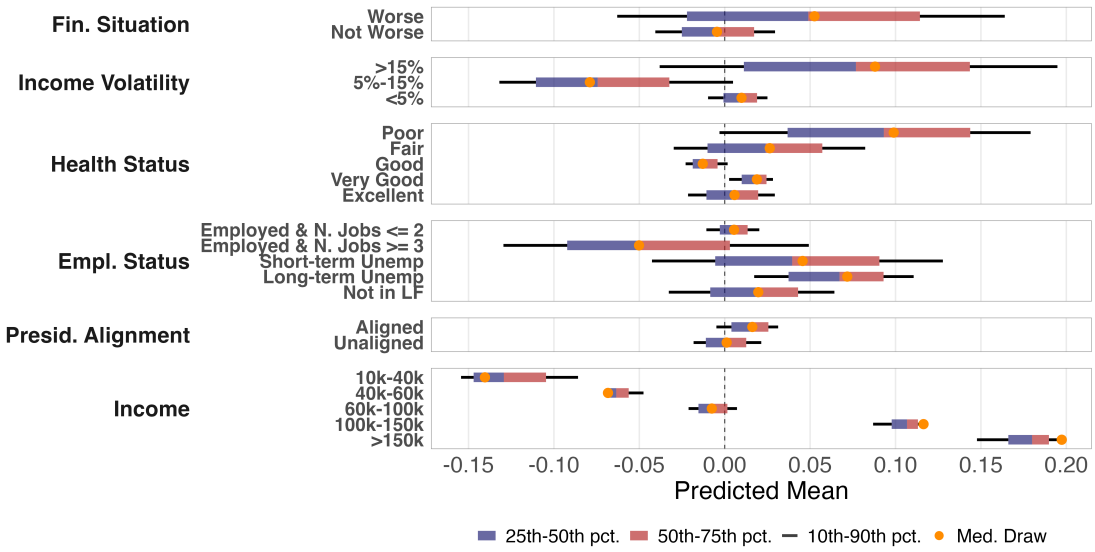


Figure A14: Anatomy of Demand: Adjusted Predicted Means by Group

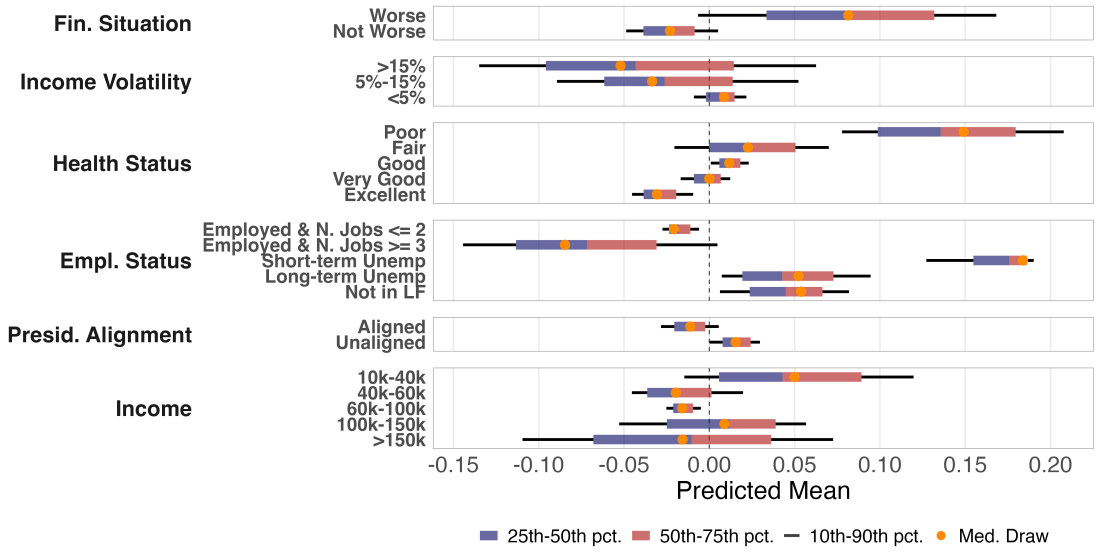
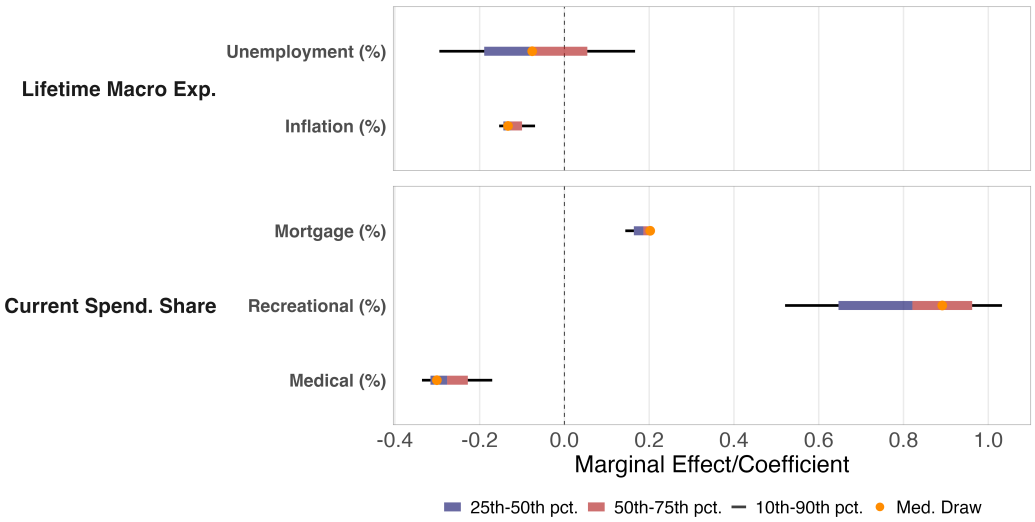
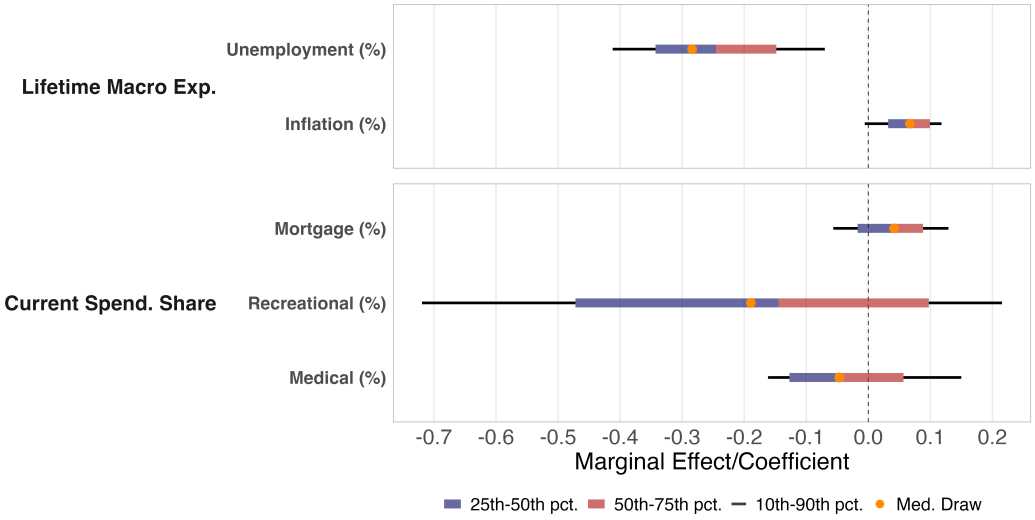


Figure A15: Anatomy of Supply: Adjusted Predicted Means by Group

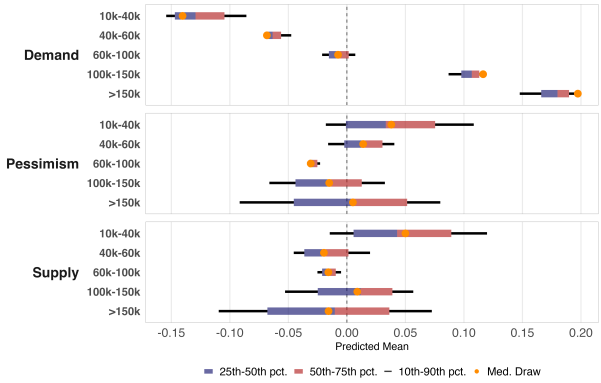


(a) Demand

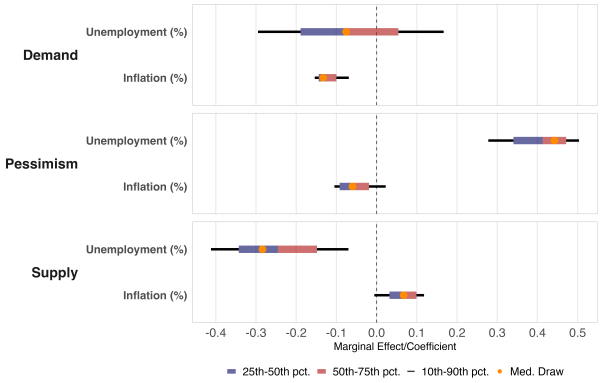


(b) Supply

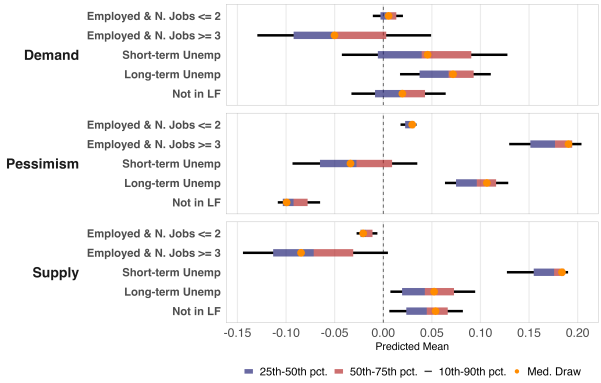
Figure A16: Anatomy of Demand and Supply: Marginal Effects of Experiences and Spending



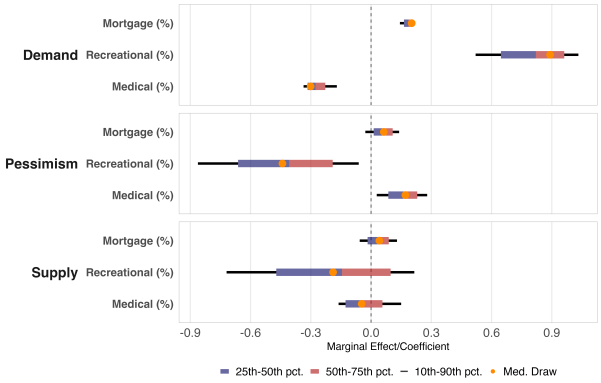
(a) Income



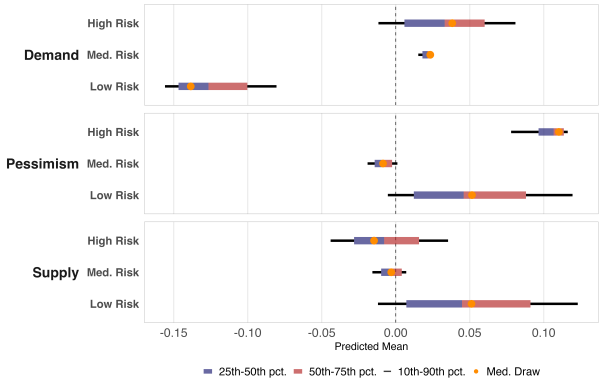
(b) Lifetime macro experiences



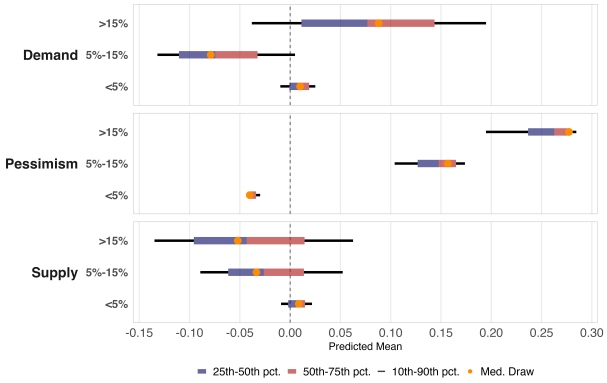
(c) Employment status



(d) Spending shares



(e) Financial risk willingness

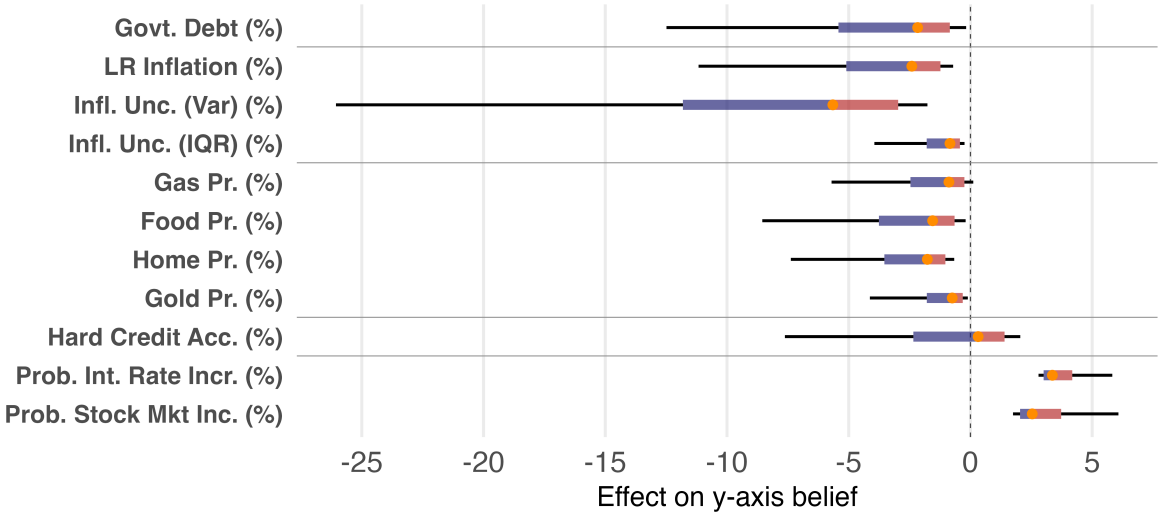


(f) Income volatility

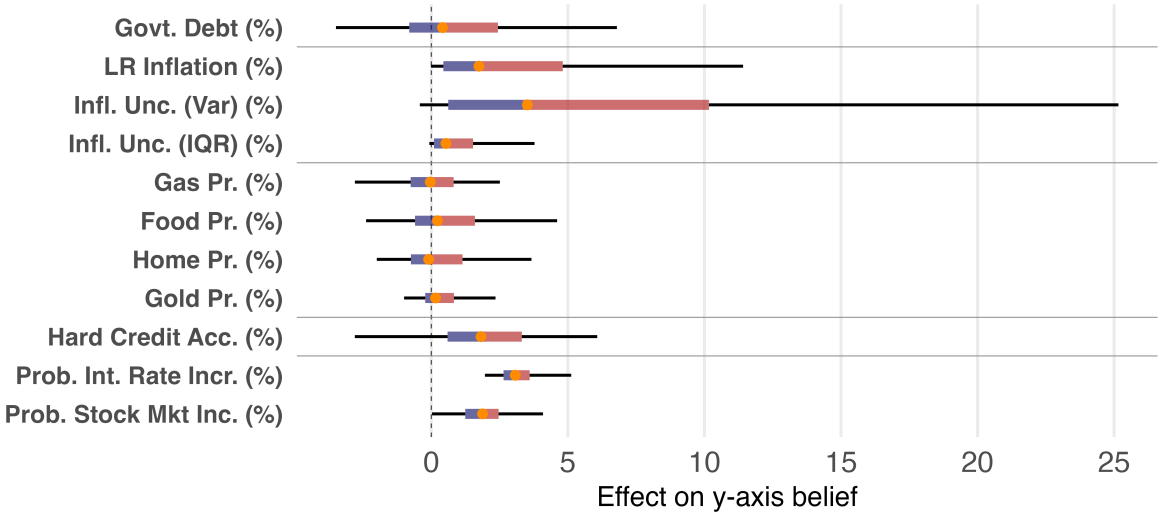
Figure A17: Anatomy of Demand, Pessimism, and Supply: Additional Characteristics

F.2 Projections of demand and supply views

Figures A18 and A19 report the projections summarized in Sections 6.2 and 6.3. On the aggregate side, the demand view lowers every nominal and fiscal expectation while the supply view moves only the inflation block, with tax, debt, and price bands that include zero. On the idiosyncratic side, the demand view lowers earnings uncertainty and raises all three job-transition probabilities, a churn pattern rather than a distress pattern, while the supply view's one clear effect is a decline in expected non-essential spending.



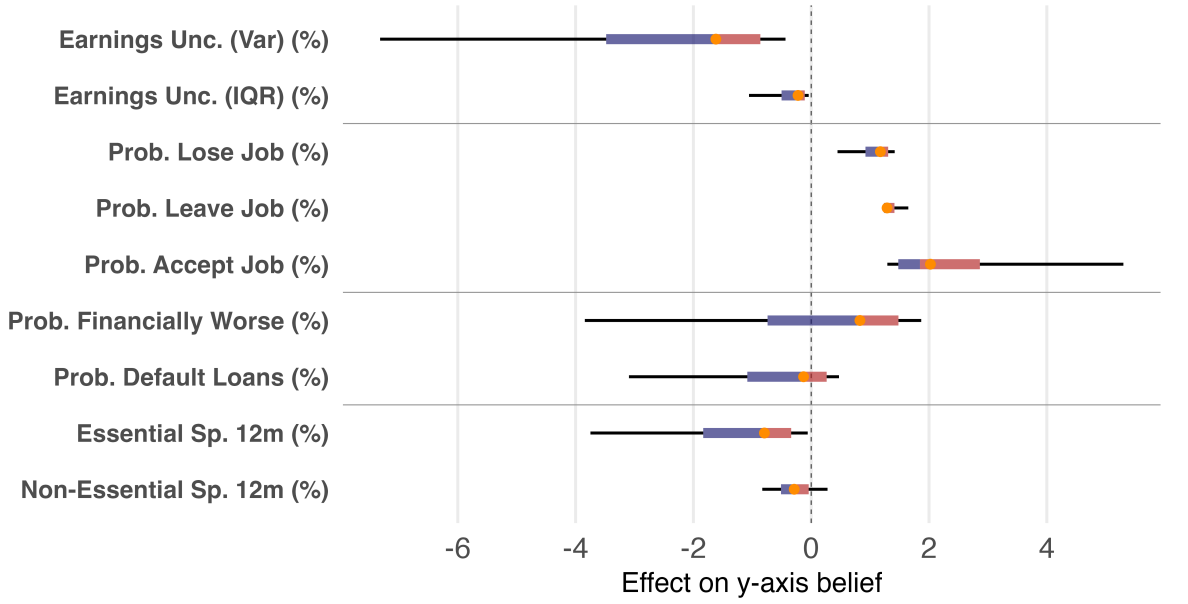
(a) Demand



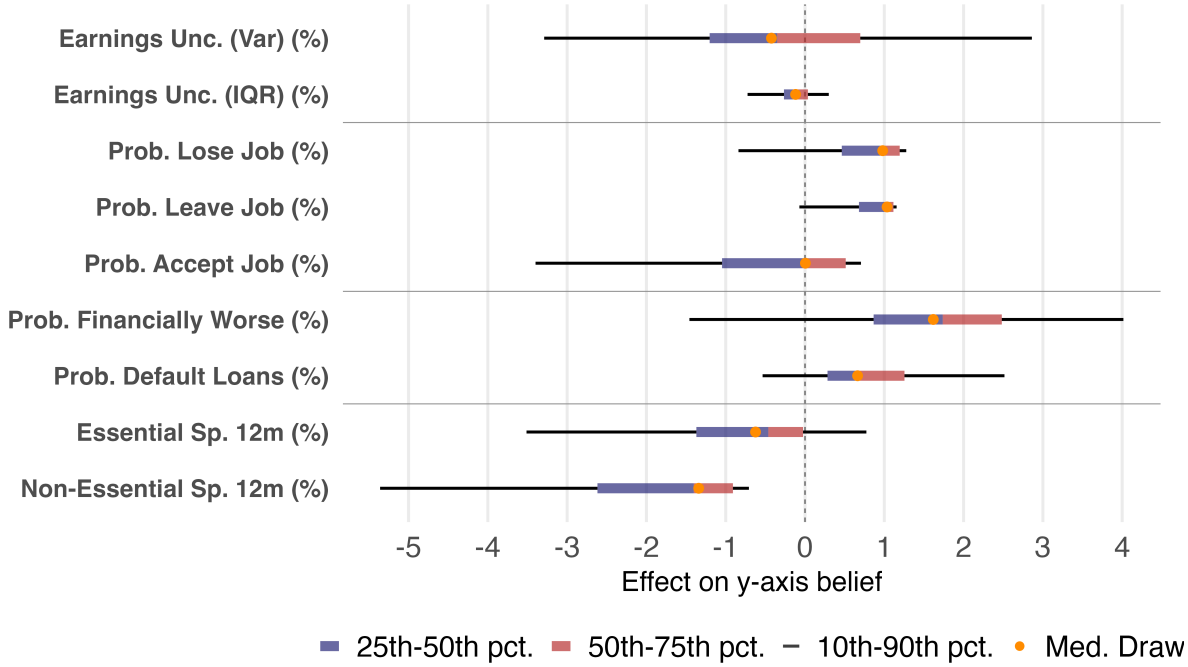
(b) Supply

Figure A18: Response of External Economy-Wide Expectations to Demand and Supply

Notes: Coefficients θ_ℓ^k from equation (9), estimated for each accepted draw. Each score is normalized to raise unemployment beliefs by 10 percentage points. Orange dots: median draw; boxes: interquartile range; whiskers: 10th to 90th percentile across draws.



(a) Demand



(b) Supply

Figure A19: Response of Idiosyncratic Expectations to Demand and Supply

Notes: Coefficients θ_ℓ^k from equation (9), estimated for each accepted draw. Each score is normalized to raise unemployment beliefs by 10 percentage points. Orange dots: median draw; boxes: interquartile range; whiskers: 10th to 90th percentile across draws.