Beliefs About Government Debt Valuation and Sustainability

Ricardo Delao Wenhao Li

August 2026

Abstract

What ultimately supports the market value of U.S. government debt? We measure *valuation frameworks*: the perceived channels that back debt value. We survey bond investors, registered voters, and economics or finance graduate-degree holders, asking them to allocate debt value across six sources derived from open-ended answers and the literature. Global safe-asset demand and permanent rollover together account for almost half of perceived value in every group, so debt value heavily relies on investor absorption rather than on conventional fiscal fundamentals. Future primary surpluses, the textbook debt backing, receive about 16% even among the graduate-trained; roughly one-third say they provide little or no support. Allocations concentrate within respondents but disperse across them, forming belief archetypes that vary little with fiscal knowledge. A model shows why belief elicitation matters: identical debt quantities and price schedules can conceal differences in saleable public backing and in investors' willingness to hold government debt under stress. These differences imply different degrees of debt sustainability. Concern about debt sustainability is widespread, and all three groups put the ten-year crisis odds near 50%. However, among respondents expressing concern, 90% of voters report that it is not decisive in their voting decisions, and 72% of investors report no concrete portfolio change.

Keywords: Government debt valuation; Fiscal sustainability; Belief heterogeneity; Safe-asset demand; Rollover beliefs; Survey experiment

*Delao: Marshall School of Business, University of Southern California, rdelao@marshall.usc.edu; Li: Marshall School of Business, University of Southern California and NBER, liwenhao@marshall.usc.edu. This project was stimulated by Gregory Mankiw's 2025 Martin Feldstein Lecture, "The Fiscal Future," at the NBER Summer Institute and by discussions that followed. We thank Jim Poterba for encouragement and helpful discussions, and Francesca Bastianello, Cary Frydman, John Graham, David Hirshleifer, Gerald Hoberg, Mete Kilic, John Matsusaka, Tracie Majors, Miao Ben Zhang, and Tom Chang for comments. We gratefully acknowledge financial support from the Institute for Outlier Research in Business (iORB) at the USC Marshall School of Business.

1. Introduction

Opinion is the soul of it; and this is affected by appearances as well as realities.

– Alexander Hamilton, *First Report on the Public Credit*, January 1790

What backs the value of government debt? Hamilton's 1790 answer, drafted as the United States moved to fund its government debt, placed beliefs alongside realities. Every generation since has re-answered the question. Ricardo (1951) formalized neutrality and then rejected it on belief grounds;[1] Barro (1974) recast debt as deferred taxation under rational expectations; the fiscal theory of the price level (Leeper 1991; Sims 1994; Woodford 2001; Cochrane 2023) pinned debt to the present value of surpluses; and the safe-asset and rational-bubble literatures restored non-surplus valuation channels through convenience services, retrading motives, and low-rate debt issuance (Tirole 1985; Krishnamurthy and Vissing-Jorgensen 2012; Nagel 2016; Caballero and Farhi 2018; Brunnermeier et al. 2024).

The core literature debate remains unresolved and is sharpened by the public-debt valuation puzzle in Jiang et al. (2024b): are debt values anchored mainly by future surpluses or by non-surplus channels? Aggregate evidence alone cannot reveal which theory of value people have in mind, because the same facts, including high debt values, low yields, and stable rollover, can be rationalized by several mechanisms: expected surpluses, safe-asset demand, permanent debt rollover, or Fed purchases. What is missing is direct evidence on how people allocate perceived backing across the channels that could support outstanding debt.

We answer this question by measuring valuation frameworks directly. An open-ended survey first asks respondents, without a menu, what ultimately backs the value of U.S. government debt. Five sources arise in these answers; we add financial repression from the government-debt literature to form a six-source taxonomy: primary surpluses, government assets, Fed purchases, global safe-asset demand, financial repression, and permanent debt rollover. The main survey presents this taxonomy to 985 respondents screened for government-bond ownership, 1,001 respondents screened for U.S. voter registration, and 247 holders of a graduate or doctoral degree in Economics or Finance.[2] Respondents identify the sources providing the greatest and the least support, explain their choices, and allocate 100 points across the sources they named as providing the greatest support and an "All others" row containing the remaining sources. All three samples then report

[1] The "Ricardian Equivalence" doctrine (debt as deferred taxation under infinite horizons and operative bequests) bearing Ricardo's name is one Ricardo himself rejected. In the *Essay on the Funding System*, he wrote that the taxpayer "deludes himself with the belief that he is as rich as before" under loan finance, precisely the belief-based rejection that the modern theorem assumes away.

[2] We refer to this third sample as the Economics/Finance sample, which captures respondents with formal training and better knowledge of economics and finance.

beliefs about debt sustainability, including perceived crisis risk, which we link to portfolio and voting choices.

The first finding is that perceived backing is overwhelmingly non-fiscal. Global safe-asset demand is the largest single source of support in all three samples, with benchmark means of 29.5 points out of 100 among bond investors, 25.7 among voters, and 26.0 in the Economics/Finance sample. Expected future primary surpluses, the textbook anchor, never exceed 16.0 points. Together, safe-asset demand and permanent rollover, two channels whose value depends strongly on investors' continued willingness to absorb Treasury supply, receive 43 to 46 points in every sample. This is close to half of perceived debt value and roughly three times the primary-surplus share. Their combined share never falls below 40 points under any admissible allocation of the survey's aggregate residual. Respondent counts yield the same picture: only 14.9% of bond investors, and a similar share of voters, give primary surpluses their largest allocation, while roughly one-third of each sample say that surpluses provide little or no support.

Formal training does not restore the fiscal benchmark's prominence. Across all three samples, the upper endpoint of the 95% confidence interval for the mean primary-surplus share never exceeds 18.8 points—less than one-fifth of the textbook benchmark of full primary-surplus backing.

The second finding is that this distribution changes little with fiscal knowledge and does not mirror mechanism frequency in national newspaper coverage. Bond investors with accurate debt/GDP estimates assign only 4.4 points more to primary surpluses than less accurate investors, with little change in the non-fiscal shares. In a 2016–2025 corpus of debt-related articles from three national newspapers, primary surpluses account for 59.4% of valuation-relevant coverage, whereas permanent rollover accounts for only 4.5%. This pattern stands in sharp contrast to our survey results across all three samples.

The result is not specific to the structured allocation exercise. In the separate open-ended survey, respondents name the same broad channels, and safe-asset demand is prominent in all three samples. A best/worst elicitation gives safe-asset demand strongly positive net support everywhere, while primary surpluses receive negative net support in all three samples. Source-specific rationales use language aligned with the channel the respondent selected. The pattern survives changes in sample, elicitation format, and response scale.

We build a parsimonious model to show why this heterogeneity matters: normal-time prices cannot reveal which sources of support would survive fiscal stress. An investor who would withdraw once repayment is doubted and one who would stay can hold the same bond at the same normal-time price. The price schedule therefore identifies total voluntary demand but cannot reveal its division among safety demand, sticky rollover, and vigilante rollover, or how much survives stress. The government budget similarly identifies the combined flow from primary balances and

public resources while concealing its composition. In a debt crisis, future surpluses are difficult to pledge because they depend on future government policy and cannot readily be transferred to creditors. Real public assets can be sold, and central-bank capital can be deployed without issuing a new fiscal promise. Economies with the same debt and normal-time prices can therefore have different refinancing capacities under stress. The survey measures the backing respondents attribute to these channels and how those beliefs respond to stress.

The third finding is that perceived crisis risk is high and broadly consistent across three separately elicited measures. Across all three samples, the average stated probability of a U.S. debt crisis within ten years is near 50%. A similar share selects a most likely crisis date within ten years, and a similar share reports sustainability headroom (maximum sustainable debt/GDP minus current debt/GDP) small enough to be exhausted by the CBO debt path within ten years. The measures also align within respondents. In every sample, respondents who select a later crisis date report a lower ten-year probability on average, fewer than 10% give answers that violate basic probability bounds, and respondents with more sustainability headroom tend to select later crisis dates. Taken together, the responses are mostly consistent across questions, samples, and respondents rather than arbitrary numerical answers.

However, high perceived risk translates weakly into concrete action. Among respondents expressing concern, 91.2% of voters say it is not decisive in their voting decisions, and 72.0% of investors report no concrete portfolio change. Even among the most concerned voters, the decisive share reaches only 18.6%. Respondents who place the most backing on primary surpluses are no more likely to act than other respondents, and behavior does not differ significantly across the six dominant valuation types.

The experiment confirms that beliefs translate weakly into intended action. Among investors, showing the current debt level and the CBO's long-run projection raises stated ten-year crisis probability by 14.9 percentage points, yet raises the share planning to reduce Treasury holdings by an imprecise 4.2 points. Information that materially changes perceived crisis risk therefore need not produce a comparable intended portfolio response.

Related literature. Our paper contributes to the literature on government debt valuation by measuring which sources people believe support outstanding debt. The surplus present-value view equates debt with discounted future primary surpluses (Barro 1979; Lucas and Stokey 1983; Bohn 1998). The fiscal theory of the price level embeds this relation in an equilibrium linking debt, surpluses, and the price level (Sargent et al. 1981; Leeper 1991; Sims 1994; Woodford 2001; Cochrane 2023), with Bassetto (2002) providing its equilibrium-selection content. Other theories allow debt value to reflect rational bubbles and low-rate debt valuation when $r < g$ (Tirole 1985; Blanchard 2019), safe-asset services (Krishnamurthy and Vissing-Jorgensen 2012; Nagel

2016; Caballero and Farhi 2018; Brunnermeier et al. 2024), self-fulfilling rollover (Calvo 1988; Cole and Kehoe 2000), or financial repression (Reinhart and Sbrancia 2015; Reis 2021). Central-bank support can also draw on finite capital beyond Treasury-controlled assets, including foreign reserves, residual seigniorage capacity, and the fiscal value of the banking franchise (Buiter 2008; Bassetto and Messer 2013; Reis 2015; Li and Merkel 2026). Hall and Sargent (2011) decompose postwar U.S. debt dynamics; Jiang et al. (2024b) show that risk-adjusted primary surpluses are far too small to account for the market value of U.S. debt; and the Jiang et al. (2024a) p-theory allows fiscal and bubble valuations to coexist without identifying their mixing proportions. Our identification result shows why aggregate prices alone cannot recover this mixture. We therefore measure the perceived mixture directly.

We also contribute to research on Treasury demand and sovereign rollover by measuring the mechanisms investors believe sustain debt value. Treasuries command value from liquidity, information insensitivity, preferred-habitat demand, retrading services, and insurance in bad states (Gorton and Pennacchi 1990; Krishnamurthy and Vissing-Jorgensen 2012; Greenwood and Vayanos 2014; Gorton 2017; Caballero and Farhi 2018; Vayanos and Vila 2021; Krishnamurthy and Li 2023; Brunnermeier et al. 2024). This literature estimates how far yields must move for the market to absorb a given supply and how intermediary balance sheets constrain that absorption (Vayanos and Vila 2021; Krishnamurthy and Li 2023; Du et al. 2023; Jansen et al. 2026). A parallel literature shows how expectations of continued refinancing can sustain debt until those expectations collapse (Eaton and Gersovitz 1981; Calvo 1988; Cole and Kehoe 2000; Arellano 2008). Our survey instead asks which sources investors believe support government debt. Respondents describe rollover as the government's capacity to continue borrowing and refinance outstanding debt, the continued-market-access mechanism in rational-bubble and low-rate debt-valuation models (Tirole 1985; Blanchard 2019). These mechanisms need not remain stable under stress. Our historical vignettes test whether beliefs respond to evidence that safe-asset status or rollover access can fail.

Our evidence complements survey research on fiscal and inflation beliefs by eliciting a perceived mechanism rather than a quantity or forecast. Survey evidence documents systematic gaps in how people understand tax policy and public finance (Stantcheva 2021; Roth et al. 2022; Bianchi et al. 2025). Roth et al. (2022) show that many U.S. respondents underestimate debt relative to GDP and that correcting this belief changes views about fiscal sustainability and government spending more than views about taxation. Research on inflation expectations shows that firms hold dispersed macroeconomic beliefs, devote limited attention to inflation, and change decisions when their beliefs change (Coibion et al. 2018; Weber et al. 2022). Household inflation expectations also reflect lifetime inflation experiences (Malmendier and Nagel 2016). Information treatments similarly move household inflation expectations and subsequent spending (Coibion et al. 2022). Our design separately measures revisions to targeted valuation channels, perceived crisis risk, and

intended behavior. The results distinguish responsiveness to information from revisions to the model through which that information is interpreted.

Our paper elicits valuation frameworks and subjective beliefs to analyze competing theories of government debt value. Households and experts hold heterogeneous causal narratives and incomplete models of equilibrium pricing (Andre et al. 2022, 2026a,b). Text-based methods recover narratives from business news and link them to asset prices (Bybee et al. 2023). Related work documents sparse models among financial analysts and heterogeneous models of climate-risk pricing (Bastianello et al. 2025; Bauer et al. 2024), while Haaland et al. (2024) systematize open-ended questions for measuring subjective models, narratives, attention, and recall. Jiang et al. (2025) link elicited recall to administrative trading records and find that what investors remember explains more of the cross-investor variation in beliefs than what they actually experienced. Elicited beliefs and motives also inform portfolio choice, trading, and models of business cycles, asset prices, monetary-policy transmission, and market design (Giglio et al. 2021; Liu et al. 2022; Bhandari et al. 2025; Adam et al. 2017; Roth et al. 2023; Kapor et al. 2020). Normal-time prices reveal total voluntary demand but not its persistence under stress. Our point allocation measures the perceived composition of backing, while open-ended responses validate the categories. Together, these measures highlight the difference between possessing different information and using a different valuation model.

Finally, our results add a belief-side condition to political-economy accounts of delayed fiscal adjustment. These accounts emphasize distributional conflict, institutional fragmentation, and electoral incentives (Alesina and Drazen 1991; Drazen and Grilli 1993; Persson and Tabellini 2000; Yared 2019). Most respondents in the screened voter sample do not value debt through a surplus-conditioned framework, and concern about sustainability rarely becomes a decisive stated voting constraint. We do not establish that the first fact causes the second. The evidence describes coexistence rather than realized market or electoral discipline.

The remainder of the paper is structured as follows. Section 2 builds the six-source taxonomy from an open-ended survey and describes the samples, the instrument, and the information treatments. Section 3 reports the perceived value decomposition of debt, the belief archetypes behind it, and the evidence that its dispersion is not an information gap. Section 4 develops the model showing why normal-time prices cannot identify stress-resilient backing. Section 5 turns to sustainability beliefs, their weak transmission to voting and portfolio decisions, and the effects of the information treatments. Section 6 concludes.

2. From Open-Ended Beliefs to Structured Valuation Frameworks

To measure valuation frameworks without imposing them, we begin with an unprompted survey asking what ultimately supports the value of U.S. government debt. We code the answers into economic channels, retain vague and non-substantive answers as separate categories, and use the resulting taxonomy in the main survey's six-source allocation task. The menu uses respondents' language but assigns each source to a distinct channel in the government-debt literature. Figure 1 summarizes the two-stage design. Panel A shows how open-ended responses generate the six-source taxonomy. Panel B shows how new, screened samples move through the structured elicitation, randomized information treatments, and post-treatment outcomes. Appendix Table A1 summarizes the full instrument, and Appendix B documents the open-text mapping.

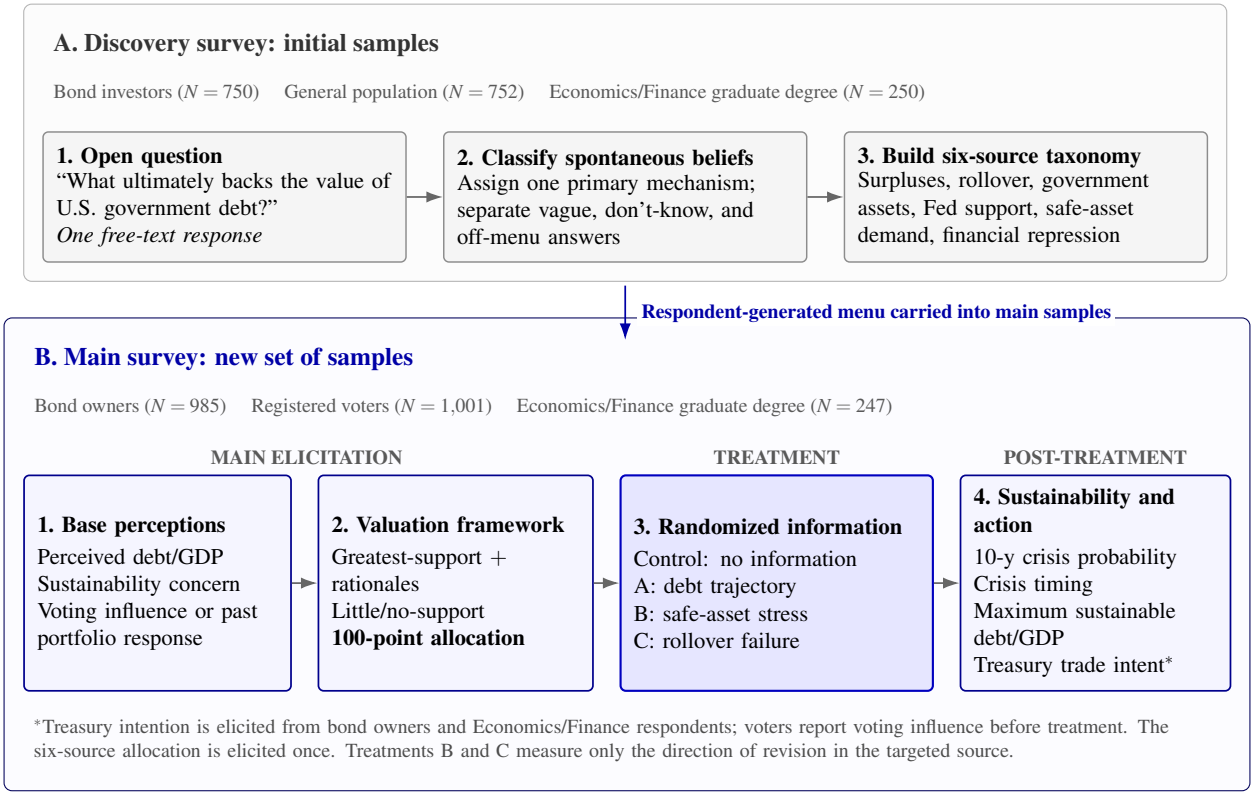


Figure 1. Survey design. A separate open-ended survey recovers spontaneous explanations of debt value and generates five of the six sources used in the structured taxonomy; financial repression is added from the government-debt literature. New, screened samples then complete the common main-survey sequence.

2.1. From spontaneous beliefs to six valuation sources

Panel A of Figure 1 shows the discovery stage. We asked three initial, unscreened samples a single free-text question: what ultimately backs the value of U.S. government debt? Respondents saw no menu of mechanisms. The samples comprise bond investors ($N = 750$), the general population ($N = 752$), and respondents whose highest completed degree is a graduate or doctoral degree in Economics or Finance ($N = 250$). Because respondents who had never considered the question were retained, the survey recovers both the theories available before researcher prompting and the share with no substantive view.

We assign each answer one primary economic mechanism and retain vague, don't-know, and off-menu explanations as separate categories. When an answer names several mechanisms, we code the one receiving the greatest emphasis. If emphasis is unclear, an explicit mechanism takes precedence over a generic explanation. For example, taxing power is coded as primary surpluses, while unspecified economic strength remains a generic explanation. The rules were fixed before the main analysis and applied independently by ChatGPT, Claude, and Gemini. The complete ChatGPT pass supplies the reported classification; the other two provide reliability checks. All three classifiers agree exactly on 81.7% of answers. Appendix B gives the full coding rules, representative answers, disambiguation cases, and reliability results.

Figure 2 shows that five channels arise without prompting in all three populations: expected future primary surpluses, permanent rollover, government real assets, Fed purchases, and global safe-asset demand. Bond investors mention rollover particularly often, respondents with graduate Economics/Finance training frequently name both surpluses and safe-asset demand, and the general-population sample supplies more vague or don't-know answers. The exact population comparisons are reported in Appendix Table A8.

Open text establishes that these channels are cognitively available before researcher prompting, but it cannot by itself provide a comparable decomposition: answers often use overlapping language, name multiple mechanisms, and differ in specificity. The main survey therefore translates the taxonomy into a common structured allocation. The resulting six-source menu carries forward the five spontaneous channels and adds financial repression from the government-debt literature, the only source not clearly identified in the open-ended responses.

The open-ended evidence also motivates the awareness screen used for the main samples. Respondents who had never considered the question were much more likely to give vague or don't-know answers, while safe-asset and rollover explanations appeared on both sides of the screen. Appendix Table A5 reports these comparisons.[3] The main survey checks the resulting beliefs using greatest-support selections, little/no-support selections, source-specific rationales,

[3] Section 3.4 shows that prior awareness cannot explain the main valuation results.

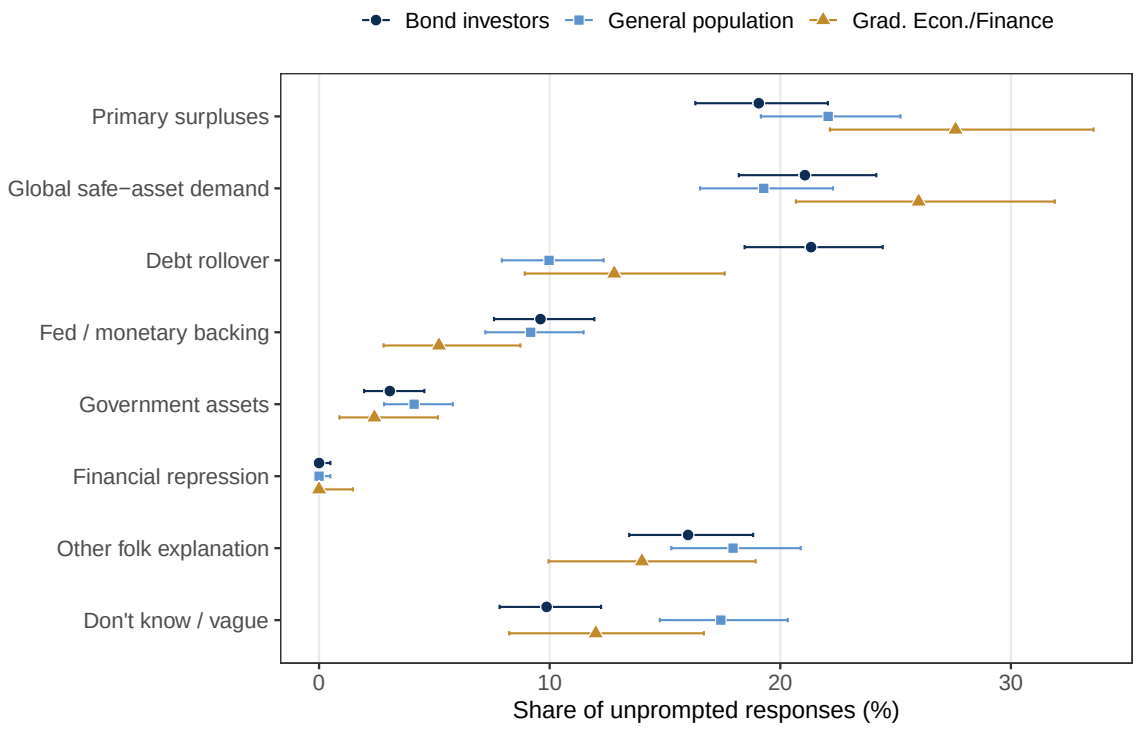


Figure 2. Unprompted debt-backing explanations across three populations. Each free-text response is assigned one primary category. Points report sample shares, and bars are exact 95% binomial confidence intervals. "Other folk explanation" combines generic-economy, production, labor, stock-market, and instrument-only answers. Primary coding covers all responses; two additional independent model passes provide reliability checks for all three samples. Bond investors $N = 750$; general population $N = 752$; Economics/Finance benchmark $N = 250$.

and a point allocation. Sections 3.2 and 3.3 report the validation results.

2.2. Main samples and survey sequence

Panel B of Figure 1 identifies the three main samples, which differ in their exposure to government-debt markets and in their economic training but complete the same survey core. Appendix D.1 gives fielding dates and additional recruitment details.

The **bond-investor sample** consists of $N = 985$ U.S. residents recruited through Prolific's prescreener for government-bond ownership. The screen establishes self-reported ownership. It records neither position size nor trading activity. The sample represents individuals with direct bond-market exposure.

The **voter/general-population sample** consists of $N = 1{,}001$ U.S. residents recruited through Prolific's registered-voter prescreener. It measures the political channel through which fiscal beliefs may affect voting. Its core instrument matches that of the investor survey, except that the

opening behavior item asks how debt sustainability has affected voting; investors instead describe past portfolio decisions.

The **Economics/Finance sample** contains $N = 247$ respondents whose highest completed degree is a graduate or doctoral degree in Economics or Finance. It provides a formal-training benchmark for the two main samples.

Table 1 reports the demographic composition of the full investor and voter samples against U.S. adult benchmarks, with every category defined to match Prolific's self-reported fields. The Economics/Finance sample is omitted from the comparison because its eligibility screen fixes education. The voter/general-population sample tracks U.S. adults reasonably closely on race and sex but, like most online panels, is younger and more educated. The bond-investor sample, recruited for direct bond-market exposure, is more male and more likely to be full-time employed. The main analyses are unweighted because identification rests on within- and cross-respondent variation in beliefs rather than on population totals. Appendix F verifies that reweighting the voter sample to U.S. adult margins leaves the value decomposition of debt and the source-valence ranking essentially unchanged.

All three samples complete the valuation, treatment, and sustainability stages. The bond-investor and Economics/Finance samples receive the final Treasury portfolio-intention item. Voters report the prior influence of debt concern on voting before treatment.

2.3. Questionnaire structure

The four boxes in Panel B correspond to three substantive modules, with a randomized-treatment block between Modules 2 and 3. The order is fixed across respondents.

Module 1: Baseline fiscal perceptions, concern, and behavior. The first main-survey box collects respondents' perceptions of the current federal debt-to-GDP ratio (Question 2) and concern about debt sustainability on a four-point scale (not at all, slightly, moderately, very concerned; Question 3). Respondents reporting any concern then receive a sample-specific behavior item. Bond investors provide an open-ended account of portfolio decisions linked to sustainability concerns (Question 4). Voters instead answer a four-point item asking how much debt concern has affected their votes for Congress or the President, from no influence to "decisive factor: I have voted differently because of a candidate's debt policy" (Question 4), and briefly explain their answer. Both items are retrospective and pre-treatment. Our cross-sample comparison pairs the voters' decisive category with a concrete portfolio change coded from the investors' open-ended responses. Because one item is closed-ended and the other is coded from open text, their levels should not be interpreted as responses to the same scale. The voter scale measures the full gradient

Table 1. Demographic composition: main samples versus U.S. adults

	Bond investors ($N = 985$)	Voters ($N = 1001$)	U.S. adults
Panel A: Sex			
Male	57.4%	44.4%	48.5%
Female	42.6%	55.6%	51.5%
Panel B: Age group			
Under 35	38.4%	38.4%	28%
35–49	37.2%	39.1%	24%
50–64	19.1%	17.1%	25%
65+	5.4%	5.5%	23%
Panel C: Race / ethnicity			
White	64.6%	69.6%	61%
Black	18.2%	14.6%	13%
Asian	6.2%	5.5%	7%
Mixed / other	11.1%	10.3%	19%
Panel D: Employment status			
Full-time employed	59.9%	–	50%
Part-time employed	16.4%	–	10%
Not in paid work	7.6%	–	37%
Unemployed (job seeking)	9.0%	–	2%
Other	7.0%	–	
Panel E: Education			
Less than bachelor's degree	–	40.2%	62%
Bachelor's degree	–	41.7%	23%
Graduate or professional degree	–	18.2%	15%

Notes. Demographic composition of the two main samples versus the U.S. adult population, with every category defined to match Prolific's self-reported fields. *Panel C:* Prolific records a single combined race/ethnicity item (*Ethnicity simplified*) with no separate Hispanic-origin question, so the comparison to the two-question Census classification is approximate; the U.S. column reports non-Hispanic White, Black, and Asian, with Hispanic and multiracial respondents combined in "Mixed / other." *Panel D:* Prolific's point-in-time *Employment status* is compared to CPS labor-force status (civilian noninstitutional population 16+, 2023 annual averages); "unemployed" is self-reported job-seeking and need not coincide with the CPS active-search definition, and "Other" covers respondents transitioning between jobs. *Panel E:* the bachelor's-degree boundary is used because Prolific records the highest *completed* qualification and has no "some college, no degree" category. The voter survey does not elicit employment, and the bond survey does not elicit education (shown as "–"). The voter sample tracks the U.S. adult population closely on sex and race; like other online panels, both samples skew younger and more educated, a tilt the reported analyses address by post-stratifying the voter sample on age, sex, and education. U.S. benchmarks: Census ACS 2019–2023 (sex, age, race); CPS 2023 annual averages (employment); ACS 2022, ages 25+ (education).

of political salience while anchoring the decisive option behaviorally. The debt/GDP benchmark is based on official statistics at the time of fielding and is not revealed during the survey.

Module 2: Beliefs about debt valuation. The second box carries the six-source taxonomy constructed in Panel A into the main samples. The model in Section 4 later gives primary surpluses, government assets, and central-bank resources terms of their own in a debt valuation equation, while rollover, safe-asset demand, and financial repression enter through the demand block. The survey names the policy action, Fed purchases; the model represents the underlying economic source as central-bank capital, broadly including foreign reserves, residual seigniorage capacity, the fiscal value of the banking franchise, and institutional credibility. Li and Merkel (2026) microfound the banking-franchise case: QE-created reserves displace bank-dependent lending and reduce taxable bank profits.

The elicitation proceeds in three steps. First, respondents select up to two sources they regard as providing the *greatest support* to the current market value of U.S. government debt (Question 5); the order of the six source choices is randomized across respondents. For each selected source, they provide a brief open-ended rationale in their own words (Question 6a–f). Second, from the sources they did not select, respondents identify those that provide *little or no support* (Question 7). The answer choices are dynamically populated with each respondent's unselected options from the first step and are again shown in random order. Third, respondents distribute 100 points across the selected top sources and an aggregate "All others" category (Question 8). This sequence anchors beliefs qualitatively before respondents assign quantitative weights. Both selection screens randomize independently, and no source is pinned to a fixed position or offered as a catch-all option, so the source rankings reported in Section 3 cannot be an artifact of the order in which the six sources appear. The open-ended rationales test whether selected sources are accompanied by substantive, source-specific reasoning and characterize the language behind each valuation archetype (Section 3.3). Appendix A.2 documents the pilot waves and the wording and sequencing changes they motivated.

Reconstructing the six-source allocation. The allocation screen records points for each selected top source and a single aggregate "All others" row. Our benchmark reconstruction takes the preceding qualitative answer at face value: it assigns zero to the sources a respondent marked as providing little or no support and divides the residual equally among the remaining non-top sources. The residual being reallocated is small, averaging 12–14 of the 100 points, with the other 86–88 assigned directly to named top sources (Appendix Table A10). Safe-asset demand exceeds primary surpluses in all three samples even when the residual is allocated arbitrarily across non-top sources; Appendix Table A11 reports the sharp bounds and shows which finer rankings and zero

shares remain reconstruction-dependent.

Module 3: Beliefs about debt sustainability. The final box contains three post-treatment sustainability questions. Respondents state the probability of a U.S. debt crisis within the next ten years (Question 11), where a crisis is defined as a sudden and severe drop in debt value leading to default or very high inflation. They select the most likely crisis horizon, including "Never" (Question 12), and report the highest debt-to-GDP ratio the U.S. can sustain without a crisis (Question 13). The last item recalls the respondent's earlier debt/GDP estimate except in Treatment A, where it recalls the official 100% value shown in the treatment. Bond investors additionally report whether they plan to adjust their U.S. government bond or bond-fund holdings over the next 12 months (Question 14). This forward-looking item complements the pre-treatment account of past portfolio decisions and lets us test whether information changes planned behavior.

2.4. Information treatments

The third main-survey box shows the four mutually exclusive groups between the valuation and sustainability modules: Control, debt trajectory, safe-asset stress, and rollover failure. This subsection describes the experimental design; Section 3.5 reports the first stage for the mechanism vignettes, and Section 5.3 reports effects on sustainability beliefs.

Control and Treatment A: Debt trajectory. Control-group respondents receive no information between Modules 2 and 3. Respondents in Treatment A are shown the historical path of U.S. federal debt held by the public and the CBO's long-run projection under current policy (Appendix Figure A1). The current level, approximately 100% of GDP, is highlighted. Because Treatment A is administered before Question 13, it affects the anchor used in the personalized crisis-threshold question: trajectory-treated respondents see "Earlier we showed you that U.S. government debt held by the public is currently about **100% of GDP**," while respondents in the other groups see their own earlier estimate from Question 2.

Treatments B and C: Belief-archetype vignettes. Treatments B and C are assigned across the full analysis sample, irrespective of respondents' earlier source selections. We can therefore estimate how source-specific historical context changes views of the targeted source and later sustainability beliefs in the same population used for the main descriptive analysis.

Treatment B: Safe-asset stress. The fielded vignette describes the decline of U.K. safe-asset status: it states that U.K. government debt accounted for about 50% of global foreign reserves

in the 1950s and about 10% by the 1970s as Britain's relative economic, fiscal, and geopolitical position weakened (Treatment B in Appendix A). The historical series behind this episode measures sterling's share of global foreign-exchange reserves, not the share of reserves invested in U.K. government debt (Schenk 2009). The follow-up asks whether the vignette changes the role respondents assign to global safe-asset demand (smaller, unchanged, or larger) and elicits a rationale. These responses distinguish fragility readings from U.S.-exceptionalism readings. Because the vignette contains only U.K. evidence and no U.S.-specific anchor, it presents erosion of safe-asset status without portraying the United States as the declining issuer.

Treatment C: Rollover failure. The fielded vignette presents 1976 U.K. debt-market turmoil as a rollover-fragility episode and states that an investor pullback from new government debt caused long-term government-bond prices to fall by about 20% within months (Treatment C in Appendix A). The follow-up first asks whether the vignette changes the respondent's view of debt rollover as a source of debt value (smaller, unchanged, or larger), and then elicits an open-ended rationale.

Both vignettes are bundled historical narratives. Treatment B links the decline of sterling's reserve share to U.K. government debt; Treatment C combines the 1976 episode, a quantitative bond-price statement, and a rollover interpretation. The randomized comparisons therefore identify responses to each package as fielded, not to any one historical component in isolation.

The valuation allocation is elicited only once. Each mechanism-treatment follow-up records whether the targeted source now plays a smaller, unchanged, or larger role and then asks for a rationale; it does not elicit a new six-source allocation. Crisis probability and timing are asked identically across all four groups. The debt-limit question is identical for Control and Treatments B and C but uses the official 100% reference in Treatment A. Cross-group differences are therefore causal effects of the randomized information packages.

Randomization is balanced on six pre-treatment measures in all three samples, with the smallest Holm-adjusted p-values of 0.335 among bond investors, 0.841 among voters, and 0.221 in the Economics/Finance sample. Appendix Table A16 reports the group means and the joint balance tests, adjusted for collection wave in the two multiwave samples, and Appendix Table A2 summarizes the four groups and the outcomes each is meant to move.

2.5. Screening, data quality, and weights

Screening. The two panels in Figure 1 use different respondents and screening rules. The initial open-ended samples retain respondents with no prior view, but the main survey screens them out. Before entering the main survey, respondents report whether they have ever thought about what supports the market value of U.S. government debt (Question 1), with options "Yes, a lot," "Yes,

somewhat," and "No." Those answering "No" do not proceed. Because the screen precedes the preamble, eligibility reflects prior awareness. Throughout the paper, "bond investors" and "voters" refer to these screened Prolific samples, not to the investing public or electorate generally. The screen-out rates, 14.2% of bond investors and 24.5% of voters, are close to, though above, the shares unable to name a source in the unscreened open-ended survey (10% and 17%; Figure 2). This comparison does not eliminate selection, so the results describe respondents with at least some prior awareness.

Exclusions. Screened-out respondents are routed past the instrument and appear in the raw export as completed responses that are blank on every substantive item, so removing blank responses removes exactly this group and no respondent who passed the screen. Among the remaining eligible respondents, we exclude those who (i) do not complete the survey (`Finished` $= 0$); (ii) submit more than once, in which case we keep the longest-duration completed submission per Prolific ID; (iii) enter an implausible debt/GDP estimate above 500% of GDP, a threshold applied uniformly across all three samples;[4] or (iv) have a Prolific submission status other than approved (returned, timed out, or absent from the Prolific export). These criteria remove 12.5% of eligible bond investors (141 of 1,126) and 10.5% of eligible voters (117 of 1,118). Incomplete responses account for most of the attrition, 120 and 108, respectively, so the three criteria that require a judgment call (duplicate submissions, the debt/GDP cap, and Prolific status) together remove only 1.9% and 0.8% of eligible respondents. In the Economics/Finance sample, the debt/GDP cap is the only binding criterion, removing 3 of 250. Manual review of all open-text fields in the bond-investor and voter samples flagged no responses as likely AI-generated, so no observations are excluded on that basis.

Collection waves. The bond-investor and voter samples were each collected in two waves. The second bond-investor wave began about five weeks after the initial wave, while the two voter waves were fielded roughly eight weeks apart. We pool the waves within each sample and include a collection-wave fixed effect in every randomized-treatment specification. The safe-asset-versus-surplus comparison holds in both waves.

Weights. All statistics and regressions in the main text are unweighted. The main analysis concerns within- and cross-respondent belief differences, but demographic composition can affect population-level aggregates, chief among them the mean value-share decomposition (Section 3.1) and the source-valence ranking (Section 3.2). As a robustness check, Appendix F reconstructs

[4]The 500% cap retains all responses consistent with any serious fiscal scenario while excluding clear data-entry errors. Results are robust to alternative thresholds.

both for the voter/general-population sample under post-stratification weights, raking it to the age and sex margins for U.S. adults and the available education margin for adults aged 25 and older (the demographic fields Prolific records; Census region is not in the Prolific export and is omitted). These weights are not intended to recover U.S. population parameters; they test whether the descriptive results are sensitive to the observed demographic composition of the retained voter sample. The reweighted value shares stay within 1.1 percentage points of the unweighted shares, and the valence ordering is preserved exactly, so both descriptive conclusions are stable under this demographic calibration. Raking cannot correct for selection on unobserved characteristics.

3. Beliefs about Debt Valuation

With the taxonomy, samples, and allocation rule defined, we now measure how respondents allocate support for the current market value of U.S. government debt. Safe-asset demand and permanent rollover together account for close to half of perceived value, and the ranking changes little across samples. Best/worst selections and open-ended rationales validate the allocations. We then examine formal training, fiscal knowledge, prior awareness, newspaper topic salience, and targeted mechanism treatments.

3.1. Valuation frameworks

Section 2.3 describes the three-step elicitation and the benchmark reconstruction. Here we use the resulting six-source vector of perceived support shares. Figure 3 plots the mean share by source for bond investors ($N = 985$), voters ($N = 1{,}001$), and Economics/Finance graduate-degree holders ($N = 247$). Appendix Table A8 reports the three-population benchmark means and formal comparisons; Appendix Table A9 gives additional investor and voter descriptives.

Three patterns stand out. First, under the benchmark reconstruction, global safe-asset demand is the single largest source in all three samples: a mean of 29.5 points among bond investors, 25.7 among voters, and 26.0 in the Economics/Finance sample. The safe-asset mean stays above the primary-surplus mean even when the residual allocation is left arbitrary (Appendix Table A11). The ranking also survives dropping the point allocation altogether. On the primary-minus-little contrast, the share naming a source as providing the greatest support minus the share dismissing it as providing little or none, safe-asset demand scores $+45$ percentage points among bond investors, $+34$ among voters, and $+35$ in the Economics/Finance sample (Section 3.2 reports all six sources). In the separate unprompted survey, where respondents saw no menu, safe-asset demand is again among the most frequent answers for sources of debt value.

Second, future primary surpluses remain a minority component even in the Economics/Finance sample. The benchmark means are 14.7, 14.3, and 16.0 points across investors, voters, and Economics/Finance graduate-degree holders. Global safe-asset demand and permanent rollover together receive 46.4, 43.5, and 44.8 points, close to half of perceived value and roughly three times the surplus share. These two channels depend on investors continuing to absorb Treasury supply. The survey separates them by stated motivation, while the model shows that stress capacity also depends on which demand persists. The allocation and persistence measures therefore do not identify structural market depth. Roughly one-third in each sample say that primary surpluses provide little or no support.

Third, value decompositions vary widely across respondents and are highly concentrated. Between 27 and 53 percent of respondents give a source no weight, either by marking it as providing little or no support or by leaving the "All others" row at zero after naming their top sources. Among bond investors, four of the six sources have a median of zero (Appendix Table A9).[5] Each mean therefore comes from the minority who give that source weight. A respondent's largest allocation averages about 60 of the 100 points. The aggregate means pool respondents who place dominant weight on different sources; they do not describe a representative diversified framework. The archetype analysis below summarizes these respondent-level configurations.

Valuation archetypes. To summarize these distinct respondent-level configurations, we cluster the 1,986 pooled benchmark vectors of the bond-investor and voter samples using k-means. We choose $k = 5$ because almost no respondent chooses financial repression as dominant. Five archetypes emerge:

1. **Primary-surplus believers** assign the largest share of value to expected future primary surpluses. They comprise 361 respondents, 18.2% of the pooled sample.

2. **Rollover optimists** place dominant weight on the government's ability to issue new debt to retire old debt indefinitely. They comprise 444 respondents, 22.4% of the pooled sample, and are proportionally more common among voters (23.9%) than among bond investors (20.8%).

3. **Government-asset believers** attribute value to the government's real assets and collateral. They comprise 319 respondents, with a centroid government-assets share of roughly 60 points.

[5]Both routes to a zero are respondent choices rather than imputation. Between 44 and 73 percent of the zeros, depending on the source and sample, come from the explicit little/no-support marking, and the remainder from a zero "All others" allocation. A source a respondent named among their own top sources is essentially never zero: at most five respondents per source in either sample, in each case from a blank allocation box.

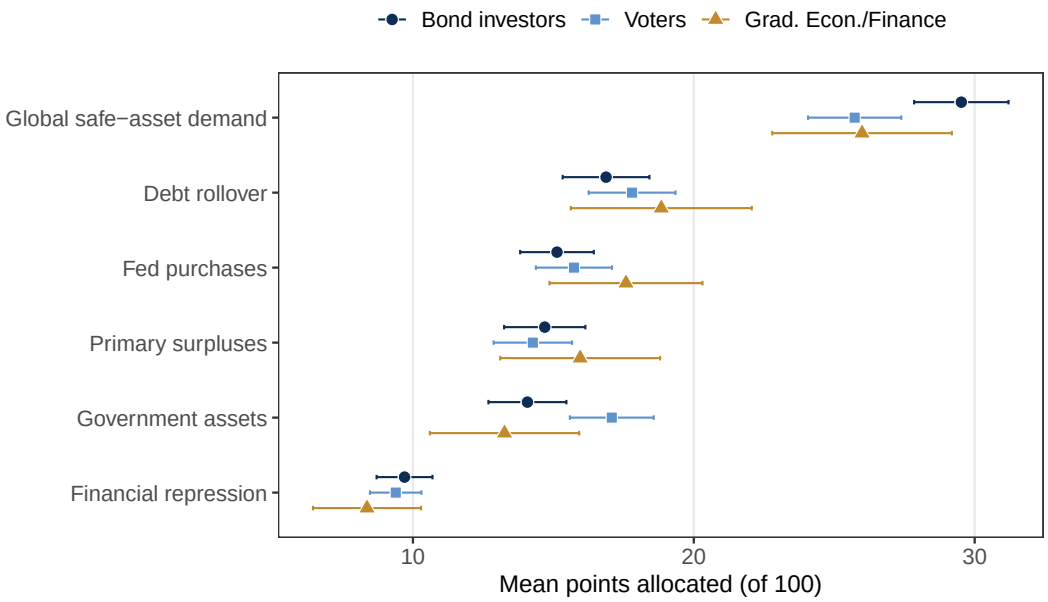


Figure 3. Mean value shares under the benchmark reconstruction, across bond investors, voters, and Economics/Finance graduate-degree holders. The benchmark assigns zero to little/no-support sources and divides the "All others" residual equally among the remaining non-top sources. Error bars show 95% confidence intervals.

4. **Fed-backstop believers** attribute value to Fed purchases. They comprise 320 respondents, with a centroid Fed-purchases share of roughly 52 points.

5. **Safe-asset believers** emphasize the safety and liquidity services that Treasuries provide to global financial markets. They are the most prevalent type overall (542 respondents, 27.3% of the pooled sample).

Figure 4 shows the centroid profiles. All five archetypes appear in both main samples and among the Economics/Finance respondents assigned to the nearest centroid. Appendix Table A13 reports each archetype's prevalence by sample alongside its defining sources.

Because clustering requires choosing k, we also summarize each respondent by the source receiving the most points. This k-free measure gives the same substantive result: only about 15% of either main sample is surplus-dominant, while safe-asset demand is the most common dominant source (Appendix Table A14 and Appendix Figure A2). Both summaries show concentrated, source-specific frameworks dominated by non-fiscal channels.

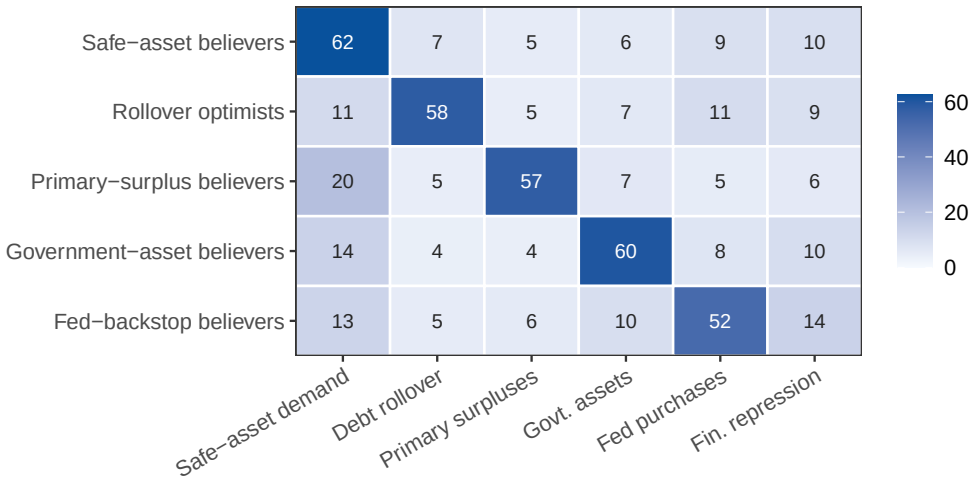


Figure 4. Belief-archetype centroids. Each row is one of the five k-means clusters estimated on the benchmark value shares over the pooled bond-investor and voter samples; each cell is the cluster's mean share of the source in that column. Archetypes are named for the source that dominates their centroid.

3.2. Valence and contestedness of value sources

Mean support shares cannot distinguish broad agreement from a distribution in which a committed minority gives a source substantial weight while many others reject it. The valuation module therefore also asks each respondent to name the sources providing the *greatest* support (Question 5) and those providing *little or no* support (Question 7), a most/least (best/worst) elicitation (Louviere et al. 2015).[6] These selections reproduce the safe-asset lead without using the point allocation (Appendix Table A15). Combined with the concentrated allocations and respondent-level archetypes in Section 3.1, they show structured positive and negative judgments about particular sources. They also recover the *balance of opinion* source by source, revealing consensus and division.

We focus on two summaries. The first is the primary-minus-little contrast introduced above, a net-favorability score of the kind business surveys such as the ISM and PMI report as a diffusion index (the "Diffusion" column of Appendix Table A15). The second is *contestedness*, $2\min(P,Q)/(P+Q) \in [0,1]$, which is near zero when opinion is one-sided and near one when a source is named greatest and little about equally often. It provides a transparent summary of disagreement about a source, in the spirit of research on group differences in choices (Gentzkow et al. 2019), belief disagreement in finance (Meeuwis et al. 2022), and the polarization-measurement tradition (Esteban and Ray 1994). Appendix Table A15 also reports two alternative summaries.[7]

[6] Survey elicitation is widely used to study economic beliefs and concerns (Stantcheva 2021; Ferrario and Stantcheva 2022; Stantcheva 2023). Respondents hold heterogeneous subjective models of how the economy works (Andre et al. 2022), and belief heterogeneity (Giglio et al. 2021) and disagreement (Meeuwis et al. 2022) shape portfolio choice.

[7] The log-odds measure is $\log(\%\text{prim.}/\%\text{little})$, with a percentile bootstrap interval over respondents. The model-

The three-sample result is sharp. Global safe-asset demand is the only source with strongly positive net support in every sample, with a primary-minus-little contrast between +34 and +45 percentage points. Fed purchases are a distant second, positive everywhere but never above +19. Primary surpluses, the textbook fiscal anchor, are net negative in all three samples, as is financial repression. Debt rollover is negative among investors and voters and roughly balanced in the Economics/Finance sample. Government assets are the one source whose sign differs across samples, marginally positive among voters and negative in the other two. Within each sample, the ranking separates a safe-asset consensus from mechanisms that divide or repel respondents; across samples, global safe-asset demand and Fed purchases retain the same top-two positions. Figure 5 plots the contrast for all six sources, and Appendix Table A15 reports it alongside the log-odds, contestedness, and latent-valence statistics for investors and voters.[8]

The Economics/Finance sample therefore reinforces the interpretation of the value shares: formal training changes neither the identity of the consensus source nor the divided status of fiscal backing. The sources that split the sample (surpluses, government assets, and rollover) are the ones whose weights separate the belief archetypes, whereas the consensus safe-asset channel is shared across them.

3.3. Internal coherence of valuation frameworks

We next ask whether respondents' open-ended reasoning aligns with the valuation source they select. The rationales from Question 6a–f explain in respondents' own words why a selected source provides the greatest support. The separate backing survey shows that the channels arise without a menu, and Appendix B gives representative mappings from raw responses to the six-source taxonomy. The main-survey rationales provide respondent-level qualitative validation.

Across the main archetypes, respondents offer, in their own language, coherent rationales that closely match the corresponding economic channels. Rollover optimists appeal to the government's ability to "always borrow more," the absence of any creditor capable of forcing repayment, and the historical success of refinancing. One respondent explained that the government can repay old debt by issuing new debt; another wrote that rollover has "seemed to work for the last 50+ years" and expected no near-term change. These explanations express the continued-market-access mechanism emphasized in formal accounts of permanent debt rollover (Tirole 1985; Blanchard 2019). Safe-asset believers emphasize global demand for "something safe," the dollar's reserve-

based latent-valence measure is the centered source fixed effect from an ordered logit of the three-level response (little $<$ neutral $<$ greatest) on source.

[8]The voter valence ranking is robust to post-stratification reweighting of the voter sample to U.S. adult margins on age, sex, and education: the contrasts shift by at most 2.3 percentage points, and the ordering of sources is preserved (Appendix F).

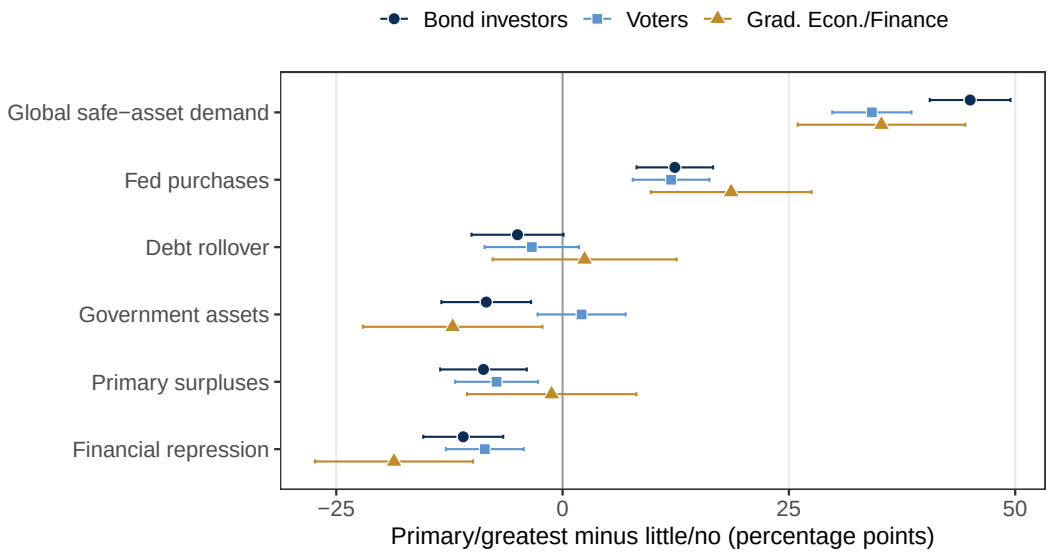


Figure 5. Net support for debt-backing sources across the three populations. For each source, the contrast is the share selecting it as providing the greatest support minus the share selecting it as providing little or no support. Positive values indicate net support. Error bars show 95% confidence intervals. Bond investors $N = 984$; voters $N = 1{,}001$; Economics/Finance benchmark $N = 247$.

currency role, the unmatched liquidity of the Treasury market, and the absence of an alternative sovereign safe haven, echoing formal accounts of safe-asset services (Brunnermeier et al. 2024). Primary-surplus believers reason by analogy to household budgets: debt must eventually be repaid from income, so future tax revenue must exceed non-interest spending. Fed-backstop believers describe an institutional put under which the Fed can "step in and buy debt directly" when private demand is weak.

Appendix Table A17 provides the systematic check: 95–100% of source rationales are substantive, their median length is 18–27 words, and 48–86% contain language specific to the selected channel.

The rationales state compact theories of value in institutional and heuristic terms. They express the mechanisms in the six-source taxonomy and support interpreting the selections as valuation frameworks.

3.4. Why the distribution is not an information gap

Formal training. If non-surplus frameworks mainly reflect limited economic knowledge, they should recede with formal training. The Economics/Finance comparison provides this test. Its ranking closely resembles those of the bond-investor and voter samples: safe-asset demand remains the largest source, safe-asset demand and rollover together receive close to half of perceived value,

and primary surpluses remain a minority component. Five of the six benchmark source shares, including primary surpluses, are statistically indistinguishable between the Economics/Finance sample and either comparison sample after within-source adjustment. Formal Economics/Finance training has little relation to the composition of perceived government-debt backing.

Fiscal knowledge. The survey elicits one benchmarked fiscal perception before the information treatments: the respondent's estimate of current U.S. debt held by the public as a share of GDP, denoted $\hat{b}_i$. We compare this estimate to the contemporaneous benchmark used in the survey, $b^{\text{data}} = 100\%$, and define the debt/GDP perception error

$$\Delta_i^b = \hat{b}_i - b^{\text{data}}.$$

Misperceptions are large. Respondents *underestimate* the current debt-to-GDP ratio: against an actual value near 100%, the median perceived ratio is 60% (mean 66%) in the bond-investor main survey, and perceptions are widely dispersed, ranging from near zero to several times GDP. Even bond investors thus have large factual uncertainty about the aggregate fiscal position. We use the accuracy of this debt/GDP perception as our operational measure of fiscal knowledge.

To summarize accuracy, we construct the fiscal knowledge index

$$\text{FKI}_i = -\left|\hat{b}_i - b^{\text{data}}\right|, \tag{1}$$

so that higher values of FKI_i correspond to smaller absolute debt/GDP errors, with zero denoting an exactly calibrated response.

We estimate the relation between fiscal knowledge and each valuation share:

$$\theta_k^i = \alpha_k + \beta_k \text{FKI}_i + X_i'\gamma_k + \varepsilon_k^i, \tag{2}$$

for each source k, where θ_k^i is respondent i's benchmark share of source k and X_i includes demographics, sample indicators, and political affiliation. The only economically meaningful relation is with the primary-surplus share. More-accurate bond investors (absolute debt/GDP error within one standard deviation; $N = 493$) assign primary surpluses a mean share of 16.9%, compared with 12.5% among less-accurate investors ($N = 492$), a 4.4-point gap (Table 2). The government-assets share falls by an offsetting amount. Safe-asset demand rises by 2.6 points and rollover falls by 2.2 points.

Fiscal knowledge raises the weight on primary surpluses but leaves the non-fiscal frameworks largely unchanged. Table 2 reports the full split for all six reconstructed shares. Across the benchmark and three alternative residual rules, defined in Appendix Table A12, the largest R^2 from

regressing any reconstructed share on the continuous knowledge measure ranges from 0.012 to 0.013 among investors and remains about 0.003 among voters. Fiscal knowledge is more strongly related to sustainability assessments in Section 5.

Table 2. Fiscal knowledge and value-share allocations, main surveys

	Bond investors			Voters		
	Knows less ($N = 492$)	Knows more ($N = 493$)	Diff. (m.−l.)	Knows less ($N = 534$)	Knows more ($N = 467$)	Diff. (m.−l.)
Value-share decomposition θ_k (points of 100)						
Global safe-asset demand	28.2	30.8	+2.6	25.2	26.3	+1.0
Primary surpluses	12.5	16.9	+4.4	13.0	15.7	+2.7
Debt rollover	18.0	15.8	-2.2	17.7	17.9	+0.2
Fed purchases	15.4	14.8	-0.6	16.0	15.5	-0.5
Government assets	15.5	12.6	-2.9	18.0	16.1	-1.9
Financial repression	10.3	9.1	-1.3	10.1	8.5	-1.6

Notes. Both main surveys split on fiscal knowledge: *knows more* if perceived current debt/GDP is within one cross-sectional SD of the true value (100%), giving $N = 493$ vs. $N = 492$ (bond) and $N = 467$ vs. $N = 534$ (voters). Both the value-share allocation and the knowledge item are pre-treatment, so the split is a baseline cross-section (associational, not randomized). The table reports mean value shares by group and the difference.

Prior awareness. A third possibility is that the distribution reflects how much thought respondents had given the question before the survey. The screen admits only respondents who had considered it, but the retained samples still divide between those who had thought about it "a lot" and "somewhat" (Question 1). This split barely moves the value shares. Safe-asset demand receives 29.3 and 29.6 points among the two investor groups and 23.8 and 26.2 among the two voter groups; primary surpluses receive 16.2 and 13.9 points among investors and 12.3 and 14.8 among voters, with the sign of the gap reversing between samples (Appendix Table A7).

The same split moves other outcomes sharply. The share of respondents reporting themselves as very concerned about sustainability runs from 27.8 to 52.5 percent across the two investor groups and from 30.3 to 57.5 percent across the two voter groups, perceived current debt/GDP is about 11 and 8 points higher in the more-considered groups, and the action margins rise from 27.6 to 36.5 percent for planned Treasury reductions and from 6.5 to 17.3 percent for decisive debt votes. Prior attention to the question therefore shifts concern, factual accuracy, and stated action, but not the composition of perceived backing.

Newspaper topic salience. Respondents may reproduce the valuation mechanisms most prominent in news coverage. We assess this explanation using ProQuest records from the *Wall Street*

Journal, the *New York Times*, and the *Washington Post* over 2016–2025. After we enforce the date window, the corpus contains 16,591 debt-related articles. We submit each full article to `gpt-4o-mini` under a common rubric; the LLM identifies 5,654 articles as valuation relevant and assigns one primary mechanism to each. The comparison measures article frequency. Figure 6 compares the survey allocations with normalized mechanism frequencies. Appendix Table A6 reports the pooled mechanism distribution and newspaper-level breakdown.

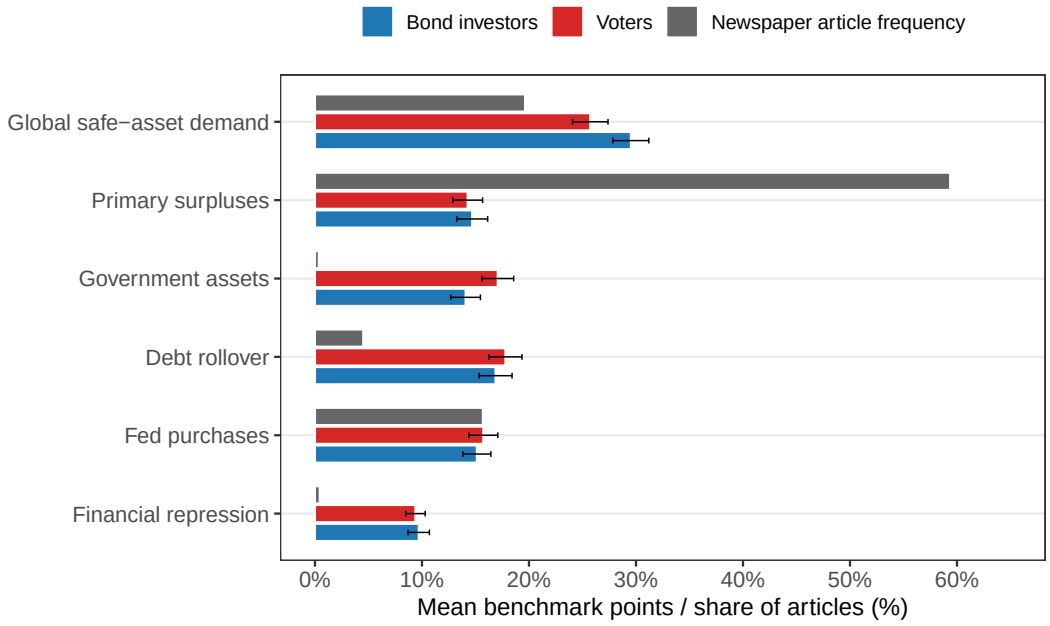


Figure 6. Respondent valuation allocations and newspaper mechanism frequency. Survey bars report mean points out of 100 for bond investors ($N = 985$) and voters ($N = 1{,}001$). Newspaper bars report the percentage of 5,583 valuation-relevant articles assigned to each of the six common mechanisms; 71 relevant "Other" articles are omitted. Survey error bars are 95% confidence intervals. No sampling interval is shown for the article frequencies. The newspaper series is a topic-frequency measure, not a measure of respondent exposure.

The two measures are considerably different: the survey bars are mean points under the benchmark reconstruction, whereas the newspaper bars are the percentage of relevant articles assigned to each mechanism. If beliefs simply mirrored topic frequency, the two mechanism mixes should resemble one another. They do not. Primary surpluses account for 59.4% of the newspaper distribution but only 14.7 and 14.3 benchmark points among bond investors and voters. Conversely, rollover accounts for only 4.5% of newspaper articles but receives 16.9 and 17.8 points. Government assets and financial repression together account for less than 1% of classified articles, yet receive 23.8 benchmark points among bond investors and 26.5 among voters. Appendix Table A11 shows that the comparison is not driven by the benchmark residual allocation. The mismatch rejects a literal topic-frequency account of the elicited beliefs.

Formal training and fiscal knowledge move the valuation mix little, and the mix does not mirror mechanism frequency in newspaper coverage. This pattern is consistent with evidence that elicited beliefs track what respondents retrieve rather than the information available to them (Jiang et al. 2025). The next question is whether respondents revise these frameworks when confronted with evidence that a supporting mechanism can fail.

3.5. Mechanism stress and heterogeneous updating

Treatments B and C (Section 2.4) test how respondents interpret the failures of two leading non-fiscal channels. Treatment B describes the erosion of U.K. safe-asset status, targeting *global safe-asset demand*; Treatment C describes a U.K. buyers' strike that disrupted *debt rollover*. Appendix A gives their complete wording and follow-up questions. Immediately afterward, respondents report whether the targeted source now plays a smaller, unchanged, or larger role in supporting U.S. debt value and explain their answer. Both vignettes are assigned across the full sample, so the estimates are population-average responses rather than within-archetype treatment effects.

Figure 7 shows the first-stage response distribution. A majority report no change in all six cells. In the two main samples, the rollover vignette produces net downweighting of 17 percentage points among bond investors and 18 among voters, compared with only 5 and 2 points after the safe-asset vignette. The Economics/Finance sample moves little after either vignette: 6 points for rollover and 5 for safe-asset demand. A historical rollover failure thus produces modest revision in the main samples, whereas another country's loss of safe-asset status barely shifts average beliefs.

The mechanism treatments produce responses in both directions. Some respondents infer that the targeted channel is more fragile; others infer that its failure makes the channel more important; and a further group exempts the United States from the historical analogy. Appendix Table A18 codes the rationales into these readings, and Appendix Table A19 shows that these readings line up with the stated direction. This heterogeneity is consistent with the small average first-stage shifts and the sample-specific downstream effects.

3.6. Implications for debt-valuation theories

These responses reveal how respondents interpret the channels, not which debt-valuation theory is correct. The survey identifies the mechanisms respondents believe support Treasury value; it does not adjudicate among the underlying theories.

The fiscal benchmark ties the real value of nominal government liabilities to expected discounted primary surpluses, with the monetary-fiscal regime determining how that restriction is

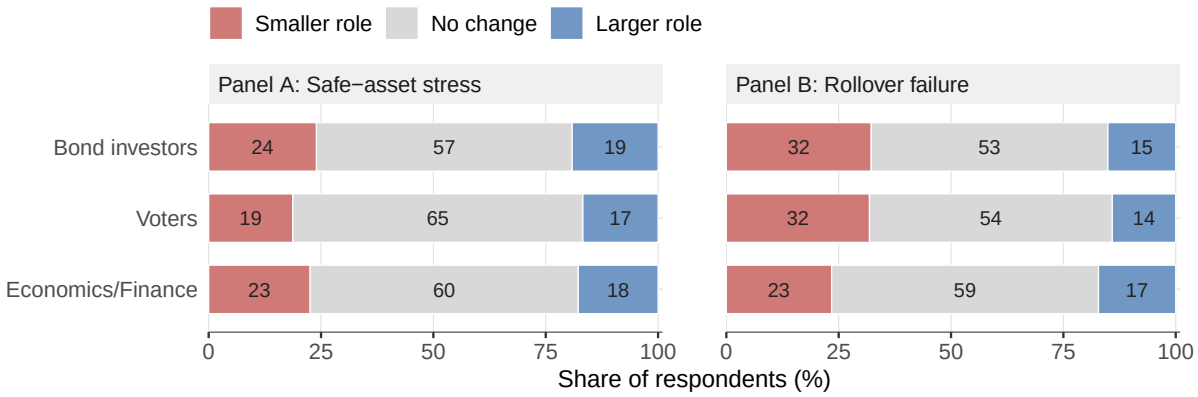


Figure 7. Directional updating after mechanism-stress vignettes. Panel A reports responses to the safe-asset-stress vignette (Treatment B), and Panel B responses to the rollover-failure vignette (Treatment C). Bars report the share of treated respondents who say the targeted source now plays a smaller, unchanged, or larger role in supporting U.S. government debt value; segment labels are percentages and each row sums to 100. Net downweighting, quoted in the text, is the smaller-role share minus the larger-role share, the difference between the two colored end segments. The item measures directional revision in the targeted source; the 100-point valuation allocation was not re-elicited.

satisfied (Woodford 2001; Cochrane 2023). Jiang et al. (2024b) document the empirical challenge: the risk-adjusted present value of U.S. primary surpluses falls far short of the market value of Treasury debt. Our result concerns beliefs rather than that valuation gap. Primary surpluses receive only 14.7, 14.3, and 16.0 points among bond investors, voters, and the Economics/Finance sample, and roughly one-third in each sample say that surpluses provide little or no support. This evidence does not test the fiscal valuation restriction. It shows that most respondents do not use surplus backing as their dominant account of Treasury value.

Respondents instead give the largest mean share to safe-asset demand. Existing work shows that Treasuries command a premium because they provide safety and liquidity (Krishnamurthy and Vissing-Jorgensen 2012; Nagel 2016); retrading can also create service flows beyond the bond's cash flows (Brunnermeier et al. 2024). Respondents distinguish these services from rollover: safe-asset demand supports value because investors prize the properties of the claim, whereas rollover supports value because investors expect future issuance to be absorbed. Under the benchmark allocation, safe-asset demand receives 29.5 versus 16.9 points for rollover among bond investors and 25.7 versus 17.8 among voters.

Rollover itself covers theories that should not be conflated. Blanchard (2019) shows that debt may be rolled over without later tax increases when the safe interest rate remains below the growth rate. Calvo (1988) derives multiple debt-service equilibria from repayment expectations, while Cole and Kehoe (2000) show how a failed bond auction can make pessimism self-fulfilling. The survey cannot choose among these models. Its contribution is to show that respondents invoke continued market access as a separate and quantitatively important source of value.

4. Why Prices Do Not Reveal Stress-Resilient Backing

The valuation evidence separates safe-asset demand from rollover, and primary surpluses from other public resources. This section shows why normal-time prices cannot recover either distinction.

An investor who would withdraw if repayment came into question and one who would continue lending hold the same bond at the same price while neither expects default. Normal-time prices reveal how much demand absorbs Treasury supply but not how much would survive a loss of safe-asset status or rollover confidence. That distinction determines refinancing capacity under stress.

We formalize this measurement limit by extending Li and Merkel (2026). Voluntary demand has three components: safety demand, sticky rollover, and vigilante rollover. All three generate the same normal-time price schedule, but only sticky rollover survives the joint stress experiment. Compelled holdings, fiscal assets, and Federal Reserve capital complete the mapping to the six survey categories. The main result concerns investor demand: normal-time prices identify its total depth but not its stress resilience.

4.1. A minimal environment

Dates are indexed by $t \in \{0, 1, \ldots\}$. A Lucas tree in fixed unit supply pays $Y > 0$ each period, and one-period government bonds promise one unit of output. All quantities are real, so a failure to repay means either outright default or inflation that erodes the real payment.

Let B_t be gross government debt maturing at t, B_t^{CB} the part held by the central bank, and $\tilde{B}_t \equiv B_t - B_t^{CB}$ the privately held remainder; $\delta_t \in \{0, 1\}$ is the realized nonpayment on debt maturing at t. Netting out the central bank is not a normalization. Consolidating the two budgets cancels the central bank's bond purchases and receipts against the Treasury's, so central-bank Treasury holdings never appear as backing and only its real resources do. The fiscal authority and central bank hold nonnegative tree shares $\alpha_t^f, \alpha_t^{\text{Fed}} \geq 0$, with $\alpha_t \equiv \alpha_t^f + \alpha_t^{\text{Fed}} \leq \bar{\alpha} < 1$; private investors

hold the remainder.[9] With constant primary balance $\bar{T}$, the no-default consolidated budget is

$$\tilde{B}_t = P^B\left(\tilde{B}_{t+1}\right)\tilde{B}_{t+1} + \bar{T} + \alpha_t Y - P_t^Y \Delta\alpha_{t+1}, \qquad \Delta\alpha_{t+1} \equiv \alpha_{t+1} - \alpha_t. \tag{3}$$

Appendix E derives this equation from the separate fiscal and central-bank budgets.

Three competitive investor blocks $j \in \mathscr{J} \equiv \{\text{Safe}, \text{Sticky}, \text{Vig}\}$ discount at rate $\beta \in (0,1)$ and evaluate

$$\mathbb{E}_0^j \sum_{t=0}^{\infty} \beta^t \left[c_t^j + v^j\left((1-\delta_t^j) b_t^j\right)\right]. \tag{4}$$

Here $b_t^j \geq 0$ is the block's holding of bonds that mature at t, bought one period earlier at $t-1$, and δ_t^j is the view block j holds when it buys: $\delta_t^j = 1$ if block j expects nonpayment at t and $\delta_t^j = 0$ if it expects payment, so $\mathbb{E}_{t-1}^j[1-\delta_t] = 1-\delta_t^j$. Perceptions govern demand; the common realized δ_t governs repayment. The linear consumption good is an unrestricted numeraire, which pins the risk-free bond price at β and the stationary tree price at $P^Y = \beta Y/(1-\beta)$.

Bonds provide non-pecuniary services through

$$v^j(b) = \chi b\left(1 - \frac{b}{2b_0^j}\right), \tag{5}$$

where $\chi \in (0,1)$ is common across blocks and $b_0^j > 0$ is block j's demand depth. This reduced form produces the downward-sloping Treasury demand used in the convenience-yield literature (Krishnamurthy and Vissing-Jorgensen 2012; Nagel 2016; Vayanos and Vila 2021; Krishnamurthy and Li 2023). The blocks differ only in depth and in whether they participate under an imposed stress state.

Total normal-time voluntary depth is

$$b_0^V = b_0^{\text{Safe}} + b_0^{\text{Sticky}} + b_0^{\text{Vig}}. \tag{6}$$

For an imposed pre-default state $\delta = (\delta^{\text{Safe}}, \delta^{\text{Sticky}}, \delta^{\text{Vig}})$, the three blocks' views of the same

[9] The public-resource tree generates a current flow and has continuation value that can be mobilized under stress. For the fiscal authority, $\alpha_t^f Y$ represents income from land, natural resources, and government-owned enterprises. For the central bank, α_t^{Fed} represents support backed by economic central-bank capital, not accounting equity or balance-sheet size. Following Li and Merkel (2026), this broad object can include foreign reserves, residual seigniorage capacity, the fiscal value of the banking franchise, and institutional credibility. Fed purchases are the policy action, while α_t^{Fed} is the underlying economic source that supports it. The acquired Treasuries do not themselves survive consolidation as backing. The tree is literal for tradable reserves and a reduced-form fiscal equivalent for the other sources. Li and Merkel microfound the banking-franchise case: QE-created reserves use scarce bank balance-sheet capacity, displace bank-dependent lending, and reduce taxable bank profits, thereby depleting α_t^{Fed}.

payment date, surviving voluntary depth is

$$\hat{b}_0^V(\delta) = \sum_{j \in \mathscr{J}} (1 - \delta^j) b_0^j. \tag{7}$$

Normal times have $\delta = (0,0,0)$. Safety failure sets $\delta^{\text{Safe}} = 1$; vigilante withdrawal sets $\delta^{\text{Vig}} = 1$. Sticky rollover carries over the last realized outcome, so it stays active before a realized default, and $\delta^{\text{Sticky}} = 0$ in every stress experiment considered here.

Policy also compels regulated institutions to hold a fixed quantity $\bar{b}^{\text{Rep}} \geq 0$. Compelled holdings have no demand schedule or failure trigger. We maintain $\bar{b}^{\text{Rep}} < (1+\chi) b_0^{\text{Sticky}} / \chi$, so voluntary investors remain marginal even when only sticky demand survives.

Bond market clearing allocates net privately held debt across compelled holdings and the three voluntary blocks:

$$\tilde{B}_{t+1} = \bar{b}^{\text{Rep}} + \sum_{j \in \mathscr{J}} b_{t+1}^j. \tag{8}$$

The tree market clears with private investors holding $1 - \alpha_t$.

4.2. What prices reveal and stress requires

These participation states matter under stress but collapse into a single demand aggregate in normal times. Investor optimality and market clearing (8) give

$$P^B\left(\tilde{B}_{t+1}; \delta_{t+1}\right) = \beta \left[1 + \chi \left(1 - \frac{\tilde{B}_{t+1} - \bar{b}^{\text{Rep}}}{\hat{b}_0^V(\delta_{t+1})}\right)\right]. \tag{9}$$

In normal times,

$$P^B\left(\tilde{B}_{t+1}\right) = \beta \left[1 + \chi \left(1 - \frac{\tilde{B}_{t+1} - \bar{b}^{\text{Rep}}}{b_0^V}\right)\right]. \tag{10}$$

The entire normal-time schedule contains the three voluntary depths only through their sum b_0^V. It therefore cannot distinguish demand that would withdraw from demand that would remain.

Let $\Phi(b; \bar{b}^{\text{Rep}})$ be the most revenue the government can raise when voluntary depth b survives:

$$\Phi(b; \bar{b}^{\text{Rep}}) = \frac{\beta}{4\chi b} \left[(1+\chi) b + \chi \bar{b}^{\text{Rep}}\right]^2. \tag{11}$$

On the maintained interior domain, Φ is strictly increasing in surviving depth. It measures a one-period funding limit and does not specify an equilibrium default rule.

Public resources differ from future primary surpluses when fiscal trust is impaired. Future

surpluses remain promises of the government and can be monetized only through the stressed bond market. Fiscal assets can instead be sold, and central-bank capital can be deployed without issuing a new fiscal promise. Let $\rho \in (0,1]$ be the fraction of normal continuation value that can be mobilized under stress. For fiscal assets, it captures saleability and any fire-sale discount; for central-bank capital, it is a reduced-form measure of immediate deployability. Using a common ρ is a benchmark restriction; component-specific rates leave the result unchanged whenever the added public resource has positive mobilization value. Deployable public resources are

$$D(\rho) \equiv \bar{T} + \alpha Y + \rho \alpha P^Y, \tag{12}$$

so even limited deployability leaves additional capacity whenever public resources retain positive mobilizable value. Full mobilization, $\rho = 1$, gives the earlier benchmark $\bar{T} + \alpha(Y + P^Y)$. We restrict attention to parameters satisfying $\Phi(b_0^{\text{Sticky}}; \bar{b}^{\text{Rep}}) + D(\rho) > 0$, so every state-contingent capacity is positive. Stress refinancing capacity in state δ is

$$B^{\text{cap}}(\delta) = \Phi\big(\hat{b}_0^V(\delta); \bar{b}^{\text{Rep}}\big) + D(\rho). \tag{13}$$

The two isolated stress states are

$$B^{\text{cap}}_{\text{Safe}} = \Phi\Big(b_0^{\text{Sticky}} + b_0^{\text{Vig}}; \bar{b}^{\text{Rep}}\Big) + D(\rho), \tag{14}$$

$$B^{\text{cap}}_{\text{Vig}} = \Phi\Big(b_0^{\text{Sticky}} + b_0^{\text{Safe}}; \bar{b}^{\text{Rep}}\Big) + D(\rho). \tag{15}$$

Joint failure leaves only sticky voluntary depth:

$$B^{\text{cap}}_{\text{Joint}} \equiv \min_{\delta^{\text{Safe}}, \delta^{\text{Vig}} \in \{0,1\}} B^{\text{cap}}(\delta) = \Phi\Big(b_0^{\text{Sticky}}; \bar{b}^{\text{Rep}}\Big) + D(\rho). \tag{16}$$

4.3. Normal-time prices do not identify stress-resilient backing

Normal-time prices and stress capacity depend on different aggregates. Consider a normal-time steady state with constant $\tilde{B}^* > \bar{b}^{\text{Rep}}$, α, and $\bar{T}$. Give an observer β, Y, ρ, $\bar{b}^{\text{Rep}}$, $\tilde{B}^*$, and the full normal-time price schedule, including its right-hand limit at $\bar{b}^{\text{Rep}}$.

Proposition 1 (Normal-time prices do not identify stress-resilient backing)**.** *The normal equilibrium identifies two aggregates but not the compositions entering stress refinancing capacity:*

(i) Investor demand. *The price schedule identifies total voluntary depth b_0^V, but not its division into safety demand, sticky rollover, and vigilante rollover. Sum-preserving reallocations leave the full normal-time schedule unchanged but alter which demand survives stress.*

(ii) Public resources, conditional on the baseline market structure. *The government budget identifies normal fiscal flow $\bar{s} \equiv \bar{T} + \alpha Y$, but not its division between the primary balance and income generated by fiscal assets and Federal Reserve capital. If those resources retain any mobilizable continuation value ($\rho > 0$) while claims on future taxes can be issued only through the capped bond market, that division changes $D(\rho)$ and hence every stress capacity.*

For the demand result, fix b_0^V and shift a small amount ε from either safety demand or vigilante rollover to sticky rollover. Then

$$\frac{dP^B(\tilde{B})}{d\varepsilon} = 0 \quad \text{for every } \tilde{B}, \tag{17}$$

$$\frac{dB^{\text{cap}}_{\text{Joint}}}{d\varepsilon} = \frac{\partial \Phi}{\partial b}\left(b_0^{\text{Sticky}}; \bar{b}^{\text{Rep}}\right) > 0. \tag{18}$$

The result is an off-path identification limit: normal prices reveal market depth but not participation after a common stress trigger.

The public-resource result rests on pledgeability. Holding $\bar{s} = \bar{T} + \alpha Y$ fixed while replacing surplus flow with public-resource income raises stress resources at rate

$$\left.\frac{\partial D(\rho)}{\partial(\alpha Y)}\right|_{\bar{s}} = \rho \frac{\beta}{1-\beta} > 0. \tag{19}$$

The current flow is unchanged, but the underlying public resource has mobilizable continuation value. An equivalent claim on future taxes has no value outside the bond market once fiscal promises are no longer trusted. Limited deployability reduces the gain in proportion to ρ, but some gain remains whenever $\rho > 0$. Longer debt maturity or the ability to pledge future taxes would weaken this one-period asymmetry. The result isolates a mechanism; it is not an estimate of U.S. fiscal capacity.

4.4. How the survey disciplines the hidden compositions

The hidden compositions have direct survey counterparts. The appendix derives a normal-time accounting allocation that mirrors the six survey categories:

$$1 = \theta_S + \theta_{\text{Asset}} + \theta_{\text{Fed}} + \theta_{\text{Safe}} + \theta_{\text{Roll}} + \theta_{\text{Rep}}. \tag{20}$$

Survey shares are perceived counterparts to these categories and do not measure structural market depth.

The survey contributes three readings that normal prices omit. First, the allocation measures

perceived exposure to public resources, safety demand, rollover, and repression. Mapping those shares into holder quantities or mobilizable resources requires external data and assumptions.

Second, Treatments B and C probe perceived state dependence. Treatment B targets the safety-failure margin in $B^{\text{cap}}_{\text{Safe}}$, while Treatment C targets vigilante withdrawal in $B^{\text{cap}}_{\text{Vig}}$. Responses and rationales reveal whether respondents view the targeted channel as fragile, but neither treatment measures withdrawn depth or recreates joint failure.

Third, reported and intended portfolio adjustments provide evidence about persistence. Treasury reductions are consistent with vigilante withdrawal. Inaction after perceived risk rises is consistent with sticky rollover, mandates, investment horizons, hedging motives, or a signal too weak to trigger adjustment. These responses describe an extensive margin and contain no auction quantities.

Allocations, vignettes, and portfolio responses measure perceived exposure, state dependence, and stated persistence. They restrict the composition hidden by normal prices without structurally identifying stress refinancing capacity. The framework does not solve a default game or predict which stress state will occur. Its purpose is to show why normal-time prices cannot substitute for these measurements.

5. Beliefs about Debt Sustainability

Section 3 measures why respondents believe government debt retains value. Section 4 shows why those value shares do not reveal how much support would survive stress. We now ask where respondents place the debt limit, how sustainability concern relates to their self-reported past voting and portfolio responses, and whether fiscal information affects their future portfolio intentions. The survey elicits the probability of a U.S. debt crisis within ten years, its most likely horizon, and the highest debt-to-GDP ratio the United States can sustain without a crisis. A crisis is defined as a sudden and severe drop in government debt value that leads to default or very high inflation.

The central empirical result is that high perceived fiscal risk translates weakly into stated action. Crisis assessments are high and mostly consistent across elicitations, and concern is widespread. Yet debt concern is rarely decisive for reported past voting, and most investors report no concrete past portfolio adjustment. The debt-trajectory experiment sharpens this gap prospectively: fiscal information raises perceived crisis risk without producing a comparable change in planned Treasury holdings over the next 12 months. A ten-year risk need not outweigh other considerations in a voting decision or a 12-month portfolio plan.

5.1. High perceived crisis risk

Figure 8 shows that perceived crisis risk is high in every sample. The mean subjective probability of a U.S. debt crisis within ten years is 46.2% among bond investors, 54.2% among voters, and 41.2% among Economics/Finance respondents. Voters report the highest risk and Economics/Finance respondents the lowest; concern produces the same ordering. The modal timing category lies within ten years in every sample, and 52.4% of bond investors, 65.4% of voters, and 48.6% of Economics/Finance respondents select a horizon within that period. These timing-based shares are close to the corresponding mean crisis probabilities and preserve the same cross-sample ordering. The probability and timing items therefore give mostly consistent aggregate readings.

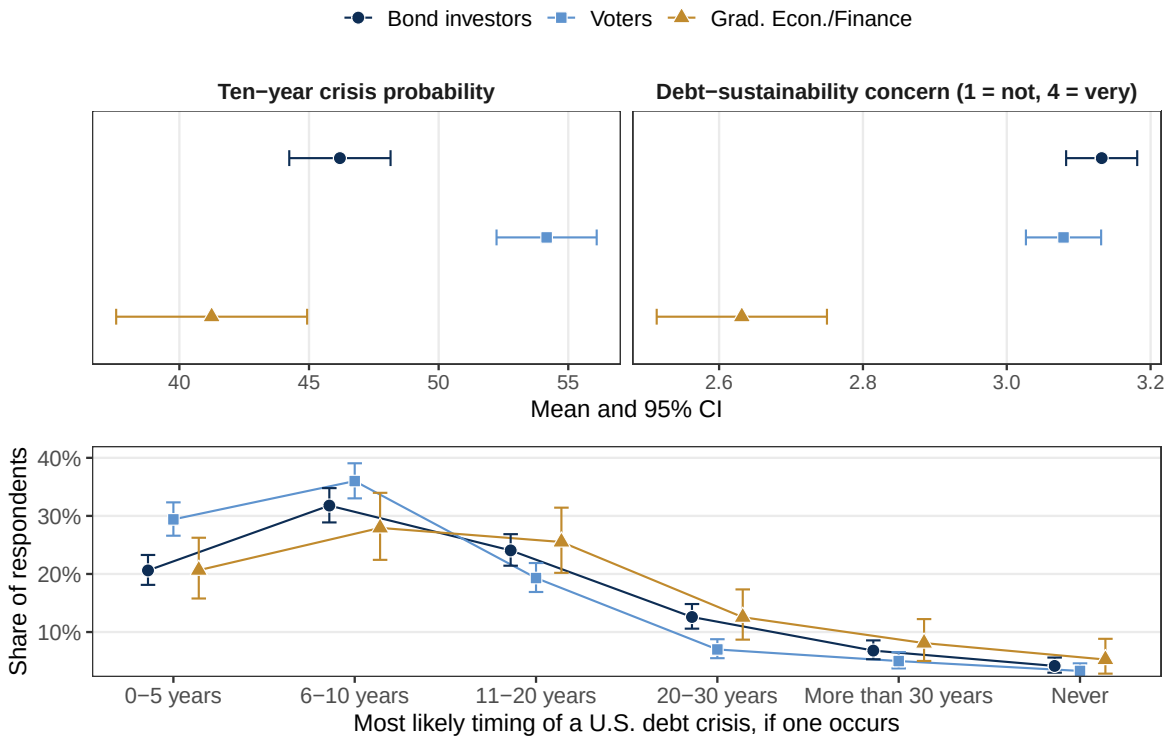


Figure 8. Perceived crisis risk across bond investors, voters, and Economics/Finance graduate-degree holders. The upper panels report the sample mean of the subjective probability of a U.S. debt crisis within ten years and the four-point debt-sustainability concern rating (1 = not concerned, 4 = very concerned). The lower panel reports the distribution of responses to the expected-timing question. Error bars show 95% confidence intervals.

Figure 9 provides a second view of these risk assessments by reporting perceived current debt/GDP, the reported ceiling, and headroom. The reported ceiling, $b^i_{\max}$, is a subjective belief about the debt ratio the United States can sustain, not a structural estimate of the model's joint-stress capacity $B^{\text{cap}}_{\text{Joint}}$. Mean perceived current debt/GDP is 66.1% among bond investors, 59.6%

among voters, and 81.1% among Economics/Finance respondents, against an actual level of about 100%. All three groups therefore understate current debt, although formal training narrows the error. This pattern is consistent with existing survey evidence on public debt (Roth et al. 2022; Grigoli and Sandri 2024; Bianchi et al. 2025).

For a common current-debt reference, consider respondents assigned to the debt-trajectory treatment, whose ceiling question reminds them that debt/GDP is approximately 100%. Their mean reported ceiling is 119.5% among bond investors, 101.2% among voters, and 115.5% among Economics/Finance respondents. Across the full samples, using the reference shown in each respondent's ceiling question, mean headroom is 29.9 percentage points among bond investors, 9.5 points among voters, and 35.1 points among Economics/Finance respondents. This ranking is the inverse of the crisis-probability ranking: voters report the least headroom and the highest crisis probability, bond investors occupy the middle, and Economics/Finance respondents report the most headroom and the lowest crisis probability.

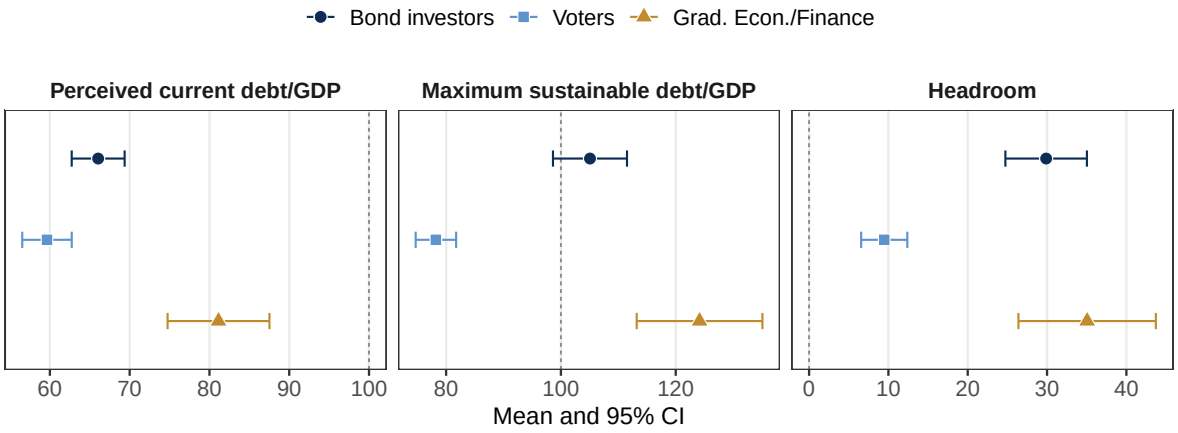


Figure 9. Perceived debt/GDP position across bond investors, voters, and Economics/Finance graduate-degree holders: perceived current debt/GDP, the perceived sustainability ceiling, and headroom. Points report sample means; error bars show 95% confidence intervals. In the left and center panels, the dashed line marks the actual current debt/GDP of about 100 percent. Headroom in the right panel is the maximum sustainable ratio minus the perceived current ratio, except for debt-trajectory respondents, whose supplied reference is 100 percent.

Figure 10 makes the consistency across the three sustainability measures explicit. Panel A compares the probability and timing measures described above. For Panel B, we map each reported ceiling among respondents shown the CBO path to the first horizon at which projected debt/GDP reaches it. The resulting shares assigned a horizon within ten years differ from the directly reported timing shares by only 4.9, 0.8, and 1.6 percentage points among bond investors, voters, and Economics/Finance respondents, respectively. The agreement across these separately

elicited questions provides a direct check on measurement consistency.[10]

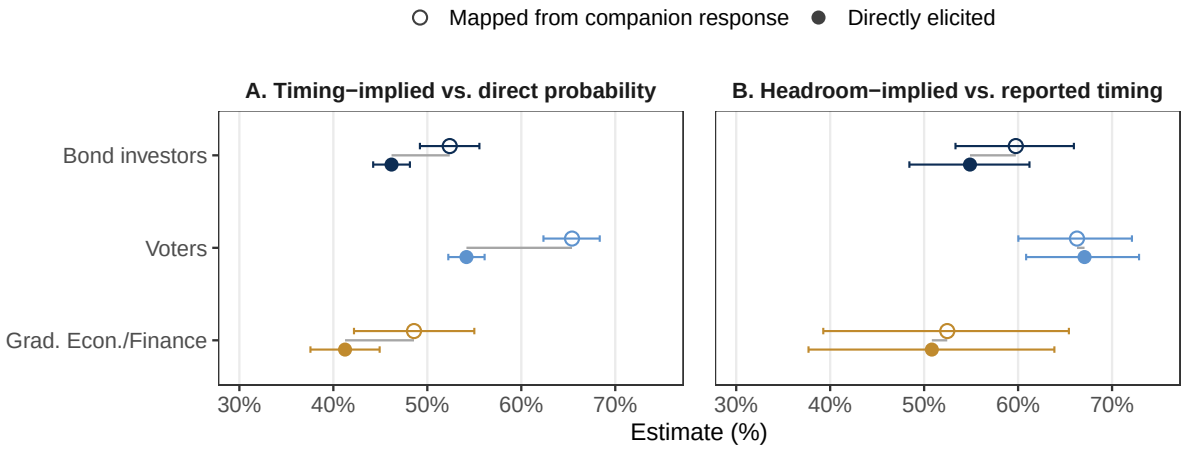


Figure 10. Comparison of aggregate sustainability measures. Panel A compares the share selecting a most likely crisis horizon of at most ten years with the mean reported ten-year crisis probability, using all valid responses to each item. Panel B uses debt-trajectory-group respondents with valid ceiling and timing responses. It codes an implied crisis by 2036 when the stated ceiling is at or below the CBO's projected 2036 debt/GDP ratio and compares the resulting share with the share directly selecting a horizon of at most ten years. Open markers are mapped measures; filled markers summarize directly elicited crisis responses. Lines connect estimates within each sample; whiskers show 95% confidence intervals.

Valuation frameworks and sustainability assessments nevertheless remain distinct. Appendix Table A24 relates each respondent's allocation across the six sources of value to that respondent's ten-year crisis probability, headroom, and expected crisis timing. A larger safe-asset share is associated with a lower crisis probability and a later expected crisis in both main samples, but with greater headroom only among voters. The rollover share is not statistically associated with any of the three outcomes in either sample, and the other source shares show no common pattern across samples. Thus, the value shares do not consistently account for respondent-to-respondent differences in perceived crisis risk or remaining fiscal room.

The model's identification result in Section 4.3 provides a natural interpretation of this weak relationship. Value shares measure the perceived composition of support, whereas stress resilience in the model also depends on the total amount of support and how much persists or remains deployable under stress. The allocation omits those margins, so respondents' separate crisis assessments need not vary tightly with their value shares.

[10]Respondent-level diagnostics point the same way: nearer reported horizons are associated with higher ten-year probabilities, few answers violate the probability bounds implied by the selected horizon, and ceiling-implied and directly reported ten-year classifications agree more often than chance. Appendix Tables A22 and A23 report the full diagnostics.

5.2. Widespread concern translates weakly into self-reported past action

Figure 11 compares respondents' reports of how debt concern affected their past choices. Concern is widespread: before treatment, 80.0% of bond investors and 75.5% of voters report being at least moderately concerned about debt sustainability (Appendix Table A25). Yet its voter panel shows that, among voters expressing some concern, only 8.8% say debt has been decisive for a past vote, even though 66.6% say it has been at least one factor. Among the most concerned voters, the decisive share rises to 18.6% (Appendix Table A25). Debt enters the political issue bundle but rarely determines the reported voting choice.

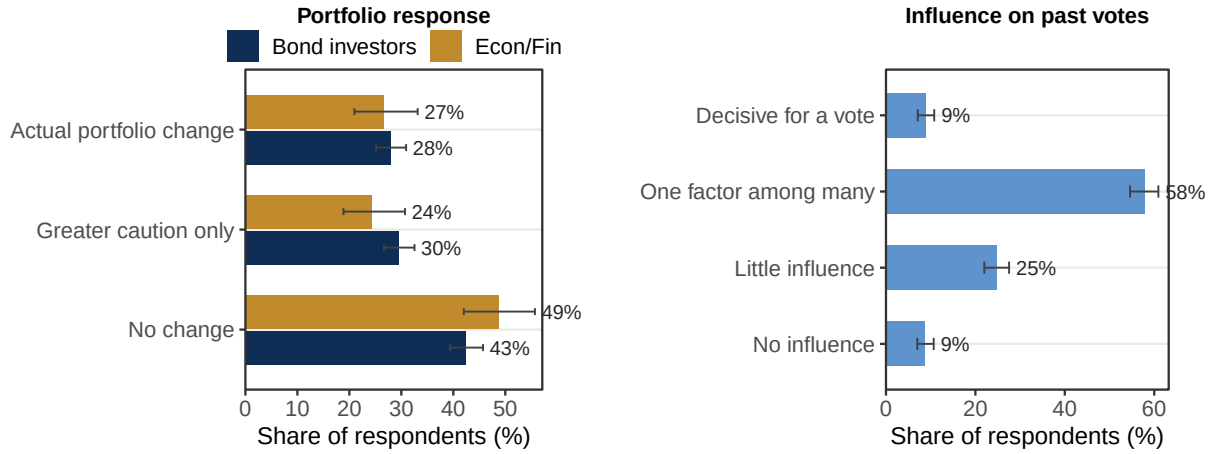


Figure 11. Self-reported past responses to debt concern. The left panel classifies open-ended accounts of whether debt concern has affected past investment or portfolio decisions among bond investors (both collection waves, $N = 955$) and the Economics/Finance sample ($N = 217$). The right panel reports voters' stated influence of debt concern on past votes ($N = 964$); respondents reporting no concern skip this item by design. Bars report sample shares; whiskers show exact 95% binomial confidence intervals.

Retrospective investor reports show limited response relative to widespread concern. Among the $N = 955$ matched open-ended answers about past investment or portfolio decisions, 42.5% report no portfolio change, 29.5% describe greater caution or monitoring without a reallocation, and 28.0% report a past portfolio adjustment. The Economics/Finance sample is similar: 26.7% report a past portfolio adjustment. Overall, more than one-quarter describe some reallocation, while nearly three-quarters describe either no change or greater attention without a reallocation. The voter item is closed-ended while the investor outcome is coded from open text, so their levels are not directly comparable.

The open-ended responses also reveal how bond investors reporting an adjustment changed their portfolios. This category includes partial reallocations and does not require a complete exit from an asset class. Among these respondents, 36% cut or avoid Treasuries or shorten duration, 21% diversify into foreign or non-dollar assets, 19% shift toward equities or real estate, 17% hold

more cash or certificates of deposit, 12% add inflation hedges such as TIPS or commodities, and 10% invest less or reduce risk. Smaller groups mention gold, crypto, or other hedges. These categories overlap because one response can report several changes. The responses therefore identify the direction of adjustment, not its size.

Figure 12 shows two distinct gradients among bond investors. Fiscal knowledge is most closely associated with the valuation framework: higher-knowledge respondents are 11 percentage points more likely to name primary surpluses among the greatest sources of support, but report only a modestly lower ten-year crisis probability and are no more likely to plan to reduce Treasury holdings. Concern about debt sustainability displays a different pattern. From the least to the most concerned group, mean crisis probability rises from 31 to 56 percent and the share planning to reduce Treasury holdings rises from 7 to 42 percent, while the share naming primary surpluses exhibits only a weak gradient. Appendix Table A25 confirms in both main samples that concern is more closely associated with crisis probability and headroom than with the six value shares. Thus, knowledge primarily distinguishes views about what supports debt value, whereas concern is much more closely associated with perceived crisis risk and intended action. Even among the most concerned investors, however, a majority does not plan to reduce Treasury holdings.

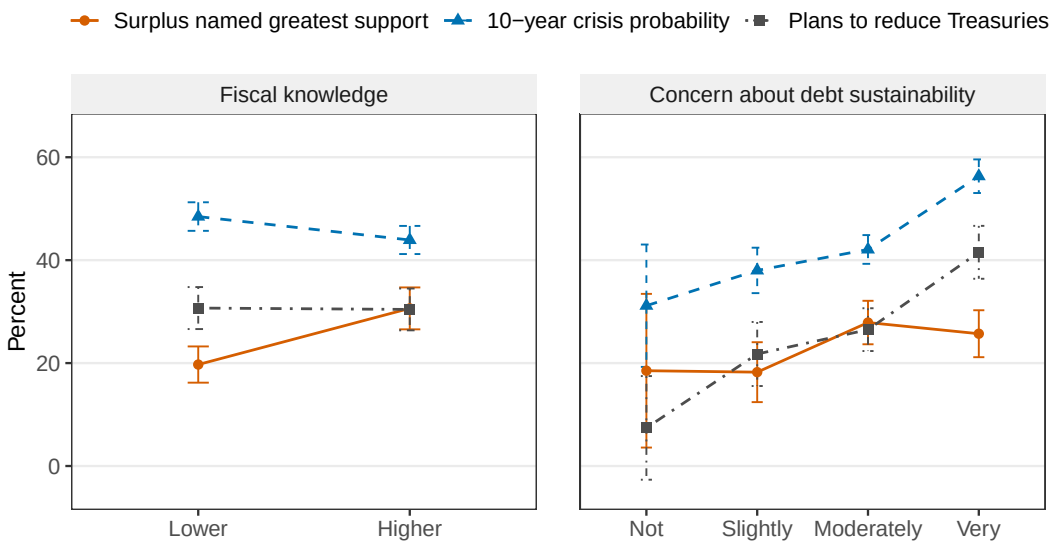


Figure 12. Valuation beliefs, crisis beliefs, and prospective portfolio intentions among bond investors ($N = 985$). The series report the share naming primary surpluses among the greatest-support sources, the mean ten-year crisis probability, and the share planning to reduce Treasury holdings over the next 12 months. Points are cross-sectional means; bars show 95% confidence intervals. Higher fiscal knowledge means that the respondent's debt/GDP estimate is within one sample standard deviation of the 100% benchmark; concern uses the survey's four response categories.

5.3. Fiscal information moves risk more than intended holdings

Table 3 reports how the randomized treatments affect perceived crisis risk and headroom. The debt-trajectory treatment shows respondents current debt/GDP of approximately 100% and the CBO's long-run projection. Because mean perceived debt is below 100% in every sample, the treatment moves the average starting point upward. It raises the stated ten-year crisis probability by 14.9 percentage points among bond investors and 6.9 points among voters, while reducing headroom relative to the reference shown by 19.5 and 10.6 points. Respondents also raise their reported ceiling, but by less than the treatment-induced increase in that reference. Appendix Table A20 provides the full decomposition, while Appendix Table A21 reports the noisier most-likely-horizon results.

Table 3. Information-treatment effects on perceived crisis risk and headroom

Group (vs. Control)	Ten-year crisis probability (pp)	Perceived headroom (pp)
Panel A. Bond investors		
A: Debt trajectory	14.87*** (2.73)	-19.55*** (6.55)
B: Safe-asset vignette	2.89 (2.61)	-1.95 (6.98)
C: Rollover vignette	4.75* (2.65)	-16.22** (6.56)
N	985	985
Control mean	41.0	38.0
Panel B. Voters		
A: Debt trajectory	6.88** (2.69)	-10.63** (4.58)
B: Safe-asset vignette	-5.60** (2.59)	-0.95 (3.52)
C: Rollover vignette	-2.97 (2.74)	2.18 (3.95)
N	1001	1001
Control mean	54.8	11.9

Notes. OLS of each outcome on the three information-treatment dummies (A debt/GDP trajectory, B safe-asset-stress vignette, C rollover-failure vignette; omitted category = Control) with pre-treatment controls (perceived debt/GDP, concern, the fiscal-knowledge index, age, and the safe-asset value share) and collection-wave fixed effects; assignment is randomized over the full sample, so controls enter for precision only. The first outcome is the respondent's stated probability of a U.S. debt crisis within ten years, measured from 0 to 100. Headroom is constructed as the reported maximum sustainable debt/GDP ratio minus the reference recalled in the ceiling question: 100 for Treatment A and the respondent's pre-treatment debt/GDP estimate for Control and Treatments B and C. Both outcomes are in percentage points. Because Treatment A changes both the information shown and the headroom reference, its headroom coefficient is the effect of that bundled design, not an average treatment effect on a uniformly defined outcome. Appendix Table A20 reports the ceiling and reference components and a uniform-headroom alternative. HC1 standard errors are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 4 reports the corresponding effects on planned Treasury-holding adjustments. Relative to a 28.6% control mean, the debt-trajectory treatment raises the share of bond investors planning to reduce Treasury or bond-fund holdings by an imprecise 4.2 percentage points and barely changes

planned increases. The belief and intention outcomes are elicited in the same experiment: information raises investors' crisis probability by 14.9 points without a comparable change in intended Treasury exposure.

Table 4. Treatment effects on planned Treasury-holding adjustments

Group (vs. Control)	Plan to reduce	Plan to increase
A: Debt trajectory	4.24 (4.07)	0.52 (2.63)
B: Safe-asset vignette	6.27 (4.08)	1.36 (2.74)
C: Rollover vignette	0.04 (4.04)	-0.82 (2.59)
N	985	985
Control mean	28.6	9.7

Notes. The outcomes are indicators for planning to reduce or increase Treasury or bond-fund holdings over the next 12 months. Each cell reports a treatment-effect estimate in percentage points relative to the omitted control group, with HC1 standard errors in parentheses. The bond-investor regressions include the same pre-treatment controls and collection-wave fixed effect as Table 3.

The remaining rows of Tables 3 and 4 show no common downstream response to the mechanism vignettes. The safe-asset vignette lowers voter crisis probability but does not move investor sustainability beliefs; the rollover vignette raises investor crisis probability and reduces investor headroom but has no corresponding voter effect. Neither vignette significantly changes planned holdings. Section 3.5 shows why a common direction should not be expected: respondents interpret the same historical episode as evidence of fragility, renewed salience, or U.S. exceptionalism. Fiscal information can therefore change perceived proximity to a crisis without producing a comparable change in investors' prospective portfolio intentions.

6. Conclusion

This paper measures what screened samples of government-bond owners and registered voters, together with an Economics/Finance graduate-degree sample, believe backs the value of U.S. government debt. We organize these beliefs into six sources and elicit their perceived value shares.

First, safe-asset demand is the largest source in all three samples, while primary surpluses receive only 14.3–16.0 points under the benchmark reconstruction. Roughly one-third in each sample say that primary surpluses provide little or no support. The safe-asset-versus-surplus comparison survives alternative allocations of the aggregate non-top residual, and the rankings barely move with formal training. Respondents disagree with the textbook fiscal benchmark about what backs Treasury value.

Second, respondents concentrate weight on dominant sources, and the resulting patterns reflect

real heterogeneity rather than menu choice. The same archetypes recur across the bond-investor and voter samples, and open-ended answers recover the broad categories in respondents' own language. Fiscal knowledge raises the surplus share by only a few points. The mechanism vignettes generate heterogeneous updating: rollover failure produces modest net downweighting in the two main samples, while respondents interpret safe-asset stress as either evidence of fragility or an improvement in the relative position of U.S. Treasuries. The trained sample moves little after either vignette.

Third, market prices cannot recover this composition. The bond price schedule and the stress refinancing capacity aggregate the same primitives differently. Two economies with identical debt and prices can therefore differ in the demand that survives stress and the public backing that remains pledgeable. Future surpluses cannot be sold separately from the government's promise, whereas real assets can be sold and central-bank capital can be deployed without issuing a new fiscal promise. Demand that withdraws likewise supports normal-time prices but not stress funding. The survey restricts both hidden compositions without structurally identifying refinancing capacity.

Finally, high perceived fiscal risk translates weakly into stated action. Across all three samples, respondents report substantial ten-year crisis risk, yet most voters say debt concern was not decisive and most investors report no concrete portfolio change. The randomized debt-trajectory treatment sharpens this weak transmission: it raises investors' perceived crisis probability without producing a comparable change in planned Treasury holdings. Economics/Finance respondents are more accurate about debt/GDP and more optimistic about sustainability than voters and, on several outcomes, bond investors, but their planned-reduction rate is statistically indistinguishable from investors'. Formal training improves knowledge and shifts sustainability assessments more than it changes valuation frameworks or intended action. These outcomes measure stated beliefs and intentions rather than realized trades, votes, or prices.

The central implication is that relevant fiscal beliefs are not summarized by a single assessment of debt or crisis risk. They also contain a model of what makes Treasuries valuable, which shapes how respondents interpret fiscal information. Respondents can observe the same fiscal facts yet draw different conclusions about sustainability; when they see Treasury value as supported by non-fiscal channels, even sizable revisions in perceived risk need not change stated portfolio or voting intentions.

References

Klaus Adam, Albert Marcet, and Johannes Beutel. Stock price booms and expected capital gains. *American Economic Review*, 107(8):2352–2408, 2017.

Alberto Alesina and Allan Drazen. Why are stabilizations delayed? *American Economic Review*, 81(5): 1170–1188, 1991.

Peter Andre, Carlo Pizzinelli, Christopher Roth, and Johannes Wohlfart. Subjective models of the macroeconomy: Evidence from experts and representative samples. *Review of Economic Studies*, 89 (6):2958–2991, 2022.

Peter Andre, Ingar Haaland, Christopher Roth, Mirko Wiederholt, and Johannes Wohlfart. Narratives about the macroeconomy. *Review of Economic Studies*, 2026a. doi: 10.1093/restud/rdag014. Advance online publication.

Peter Andre, Philipp Schirmer, and Johannes Wohlfart. Mental models of the stock market. *Quarterly Journal of Economics*, 2026b. doi: 10.1093/qje/qjag039. Advance online publication.

Cristina Arellano. Default risk and income fluctuations in emerging economies. *American Economic Review*, 98(3):690–712, 2008.

Robert J. Barro. Are government bonds net wealth? *Journal of Political Economy*, 82(6):1095–1117, 1974.

Robert J Barro. On the determination of the public debt. *Journal of political Economy*, 87(5, Part 1): 940–971, 1979.

Marco Bassetto. A game-theoretic view of the fiscal theory of the price level. *Econometrica*, 70(6):2167–2195, 2002.

Marco Bassetto and Todd Messer. Fiscal consequences of paying interest on reserves. *Fiscal Studies*, 34(4): 413–436, 2013.

Francesca Bastianello, Paul H. Décaire, and Marius Guenzel. Mental models and financial forecasts. Working paper, November 2025. `https://ssrn.com/abstract=5004839`, 2025.

Rob Bauer, Katrin Gödker, Paul Smeets, and Florian Zimmermann. Mental models in financial markets: How do experts reason about the pricing of climate risk? IZA Discussion Paper 17030, IZA Institute of Labor Economics, 2024.

Anmol Bhandari, Jaroslav Borovička, and Paul Ho. Survey data and subjective beliefs in business cycle models. *Review of Economic Studies*, 92(3):1375–1437, 2025.

Francesco Bianchi, Era Dabla-Norris, and Salma Khalid. Perceptions of public debt and policy expectations: Evidence from cross country surveys. NBER Working Paper 34382, National Bureau of Economic Research, 2025. Also IMF Working Paper WP/25/197.

Olivier Blanchard. Public debt and low interest rates. *American Economic Review*, 109(4):1197–1229, 2019.

Henning Bohn. The behavior of US public debt and deficits. *the Quarterly Journal of economics*, 113(3): 949–963, 1998.

Markus K. Brunnermeier, Sebastian Merkel, and Yuliy Sannikov. Safe assets. *Journal of Political Economy*, 132(11):3603–3657, 2024.

Willem H. Buiter. Can central banks go broke? Technical report, CEPR Discussion Paper No. 6827, 2008.

Leland Bybee, Bryan T. Kelly, and Yinan Su. Narrative asset pricing: Interpretable systematic risk factors from news text. *The Review of Financial Studies*, 36(12):4759–4787, 2023. doi: 10.1093/rfs/hhad042.

Ricardo J Caballero and Emmanuel Farhi. The safety trap. *Review of Economic Studies*, 85(1):223–274, 2018.

Ricardo J. Caballero and Alp Simsek. A model of fickle capital flows and retrenchment. *Journal of Political Economy*, 128(6):2288–2328, 2020.

Guillermo A. Calvo. Servicing the public debt: The role of expectations. *American Economic Review*, 78 (4):647–661, 1988.

John H Cochrane. *The Fiscal Theory of the Price Level*. Princeton University Press, 2023.

Olivier Coibion, Yuriy Gorodnichenko, and Saten Kumar. How do firms form their expectations? new survey evidence. *American Economic Review*, 108(9):2671–2713, 2018. doi: 10.1257/aer.20151299.

Olivier Coibion, Yuriy Gorodnichenko, and Michael Weber. Monetary policy communications and their effects on household inflation expectations. *Journal of Political Economy*, 130(6):1537–1584, 2022. doi: 10.1086/718982.

Harold L. Cole and Timothy J. Kehoe. Self-fulfilling debt crises. *Review of Economic Studies*, 67(1):91–116, 2000. doi: 10.1111/1467-937X.00123.

Congressional Budget Office. The long-term budget outlook data: 2026 to 2056, 2026. Publication 62044, February 2026. `https://www.cbo.gov/publication/62044`.

Allan Drazen and Vittorio Grilli. The benefit of crises for economic reforms. *American Economic Review*, 83(3):598–607, 1993.

Wenxin Du, Benjamin Hébert, and Wenhao Li. Intermediary balance sheets and the treasury yield curve. *Journal of Financial Economics*, 150(3):103722, 2023.

Jonathan Eaton and Mark Gersovitz. Debt with potential repudiation: Theoretical and empirical analysis. *Review of Economic Studies*, 48(2):289–309, 1981.

Joan-María Esteban and Debraj Ray. On the measurement of polarization. *Econometrica*, 62(4):819–851, 1994.

Beatrice Ferrario and Stefanie Stantcheva. Eliciting people's first-order concerns: Text analysis of open-ended survey questions. *AEA Papers and Proceedings*, 112:163–169, 2022.

Matthew Gentzkow, Jesse M. Shapiro, and Matt Taddy. Measuring group differences in high-dimensional choices: Method and application to congressional speech. *Econometrica*, 87(4):1307–1340, 2019.

Stefano Giglio, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus. Five facts about beliefs and portfolios. *American Economic Review*, 111(5):1481–1522, 2021.

Gary B. Gorton. The history and economics of safe assets. *Annual Review of Economics*, 9:547–586, 2017.

Gary B. Gorton and George Pennacchi. Financial intermediaries and liquidity creation. *Journal of Finance*, 45(1):49–71, 1990.

Robin Greenwood and Dimitri Vayanos. Bond supply and excess bond returns. *The Review of Financial Studies*, 27(3):663–713, 2014.

Francesco Grigoli and Damiano Sandri. Public debt and household inflation expectations. *Journal of International Economics*, 152:104003, 2024. doi: 10.1016/j.jinteco.2024.104003.

Ingar K. Haaland, Christopher Roth, Stefanie Stantcheva, and Johannes Wohlfart. Understanding economic behavior using open-ended survey data. NBER Working Paper 32421, National Bureau of Economic Research, 2024. Earlier version circulated as "Measuring What Is Top of Mind".

George J. Hall and Thomas J. Sargent. Interest rate risk and other determinants of post-WWII US government debt/GDP dynamics. *American Economic Journal: Macroeconomics*, 3(3):192–214, 2011.

Kristy A. E. Jansen, Wenhao Li, and Lukas Schmid. Granular treasury demand with arbitrageurs. NBER Working Paper No. 33243, 2026.

Wei Jiang, Thomas J Sargent, Neng Wang, and Jinqiang Yang. A p theory of government debt and taxes. *The Journal of Finance*, 2024a.

Zhengyang Jiang, Hanno Lustig, Stijn Van Nieuwerburgh, and Mindy Z Xiaolan. The U.S. public debt valuation puzzle. *Econometrica*, 92(4):1309–1347, 2024b.

Zhengyang Jiang, Hongqi Liu, Cameron Peng, and Hongjun Yan. Investor memory and biased beliefs: Evidence from the field. *Quarterly Journal of Economics*, 140(4):2749–2804, 2025.

Adam J. Kapor, Christopher A. Neilson, and Seth D. Zimmerman. Heterogeneous beliefs and school choice mechanisms. *American Economic Review*, 110(5):1274–1315, 2020.

Arvind Krishnamurthy and Wenhao Li. The demand for money, near-money, and treasury bonds. *The Review of Financial Studies*, 36(5):2091–2130, 2023.

Arvind Krishnamurthy and Annette Vissing-Jorgensen. The aggregate demand for treasury debt. *Journal of Political Economy*, 120(2):233–267, 2012.

Eric M. Leeper. Equilibria under 'active' and 'passive' monetary and fiscal policies. *Journal of Monetary Economics*, 27(1):129–147, 1991.

Wenhao Li and Sebastian Merkel. Quantitative easing and government debt sustainability. NBER Working Paper 35421, 2026.

Hongqi Liu, Cameron Peng, Wei A. Xiong, and Wei Xiong. Taming the bias zoo. *Journal of Financial Economics*, 143:716–741, 2022.

Jordan J. Louviere, Terry N. Flynn, and A. A. J. Marley. *Best-Worst Scaling: Theory, Methods and Applications*. Cambridge University Press, 2015.

Robert E Lucas and Nancy L Stokey. Optimal fiscal and monetary policy in an economy without capital. *Journal of monetary Economics*, 12(1):55–93, 1983.

Ulrike Malmendier and Stefan Nagel. Learning from inflation experiences. *Quarterly Journal of Economics*, 131(1):53–87, 2016. doi: 10.1093/qje/qjv037.

Maarten Meeuwis, Jonathan A. Parker, Antoinette Schoar, and Duncan Simester. Belief disagreement and portfolio choice. *The Journal of Finance*, 77(6):3191–3247, 2022.

Stefan Nagel. The liquidity premium of near-money assets. *The Quarterly Journal of Economics*, 131(4): 1927–1971, 2016.

Torsten Persson and Guido Tabellini. *Political Economics: Explaining Economic Policy*. MIT Press, Cambridge, MA, 2000.

Carmen M Reinhart and M Belen Sbrancia. The liquidation of government debt. *Economic Policy*, 30(82): 291–333, 2015.

Ricardo Reis. Different types of central bank insolvency and the central role of seignorage. Technical report, NBER Working Paper No. 21226, 2015.

Ricardo Reis. The constraint on public debt when $r < g$ but $g < m$. Technical report, BIS Working Paper No. 939, 2021.

David Ricardo. Essay on the funding system. In Piero Sraffa, editor, *The Works and Correspondence of David Ricardo, Vol. IV: Pamphlets and Papers 1815–1823*. Cambridge University Press, Cambridge, 1951. Originally published 1820.

Christopher Roth, Sonja Settele, and Johannes Wohlfart. Beliefs about public debt and the demand for government spending. *Journal of Econometrics*, 231(1):165–187, 2022. doi: 10.1016/j.jeconom.2020.09.011.

Christopher Roth, Mirko Wiederholt, and Johannes Wohlfart. The effects of monetary policy: Theory with measured expectations. CESifo Working Paper No. 10216, 2023.

Thomas J Sargent, Neil Wallace, et al. Some unpleasant monetarist arithmetic. *Federal reserve bank of minneapolis quarterly review*, 5(3):1–17, 1981.

Catherine R. Schenk. The retirement of sterling as a reserve currency after 1945: Lessons for the US dollar? Working paper, University of Glasgow, May 2009. URL `https://www.cirje.e.u-tokyo.ac.jp/research/workshops/history/history_paper2009/history1109.pdf`.

Christopher A Sims. A simple model for study of the determination of the price level and the interaction of monetary and fiscal policy. *Economic Theory*, 4(3):381–399, 1994.

Stefanie Stantcheva. Understanding tax policy: How do people reason? *The Quarterly Journal of Economics*, 136(4):2309–2369, 2021.

Stefanie Stantcheva. How to run surveys: A guide to creating your own identifying variation and revealing the invisible. *Annual Review of Economics*, 15:205–234, 2023.

Jean Tirole. Asset bubbles and overlapping generations. *Econometrica: Journal of the Econometric Society*, pages 1499–1528, 1985.

Dimitri Vayanos and Jean-Luc Vila. A preferred-habitat model of the term structure of interest rates. *Econometrica*, 89(1):77–112, 2021.

Michael Weber, Francesco D'Acunto, Yuriy Gorodnichenko, and Olivier Coibion. The subjective inflation expectations of households and firms: Measurement, determinants, and implications. *Journal of Economic Perspectives*, 36(3):157–184, 2022. doi: 10.1257/jep.36.3.157.

Michael Woodford. Fiscal requirements for price stability. *Journal of Money, Credit and Banking*, 33(3): 669–728, 2001.

Pierre Yared. Rising government debt: Causes and solutions for a decades-old trend. *Journal of Economic Perspectives*, 33(2):115–140, 2019.

Internet Appendix of "Beliefs About Government Debt Valuation and Sustainability"

Ricardo Delao Wenhao Li

A. Survey Instrument Summary

This appendix documents the order, constructs, sample branches, and randomization in the main survey. Most items are described at the construct level. The two mechanism vignettes are quoted verbatim, and the debt-trajectory display appears in Figure A1. Question numbers match those cited throughout the paper and are identical across the three samples.

A.1. Overview

The instrument consists of a screen, three substantive modules, and a treatment block between Modules 2 and 3. Respondents first complete an awareness screen. Eligible respondents then report factual perceptions and debt-sustainability concern, complete the debt-valuation module, receive randomized information where applicable, and finally answer sustainability and behavior questions. Table A1 summarizes the construct map.

A.2. Instrument development and piloting

Instrument development began in September 2025, with successive questionnaire versions and student feedback rounds during the fall. Before the June 2026 main surveys, we fielded more than a dozen Prolific pilot waves across general-population and bond-investor samples. The pilots led to three changes. First, short labels such as "safe-asset demand" and "rollover" were sometimes read as generic confidence statements, so the final menu pairs each source with a plain-language description. Second, some respondents wanted to reject a source, so the final instrument asks which sources provide little or no support before the point allocation. Third, some pilot respondents had never formed a view about debt value, motivating the prior-thought screen. These revisions preserve the economic distinctions while making the elicitation understandable.

A.3. Screening and opening measures

Question 1. The first item measures whether the respondent had previously considered what ultimately supports the market value of U.S. government debt.

Table A1. Construct-level summary of the main survey instrument

Instrument block	Construct measured	Description
Screening	Prior awareness of the debt-valuation question	Identifies respondents who have previously thought about what supports the market value of U.S. government debt.
Module 1	Fiscal perceptions and concern	Elicits perceived current debt/GDP, concern about debt sustainability, and whether that concern has affected voting or portfolio behavior.
Module 2	Debt-valuation framework	Elicits qualitative greatest-support sources, open-ended rationales for selected sources, qualitative little/no-support sources, and a 100-point allocation across the selected top sources plus an aggregate "All others" residual.
Treatment block	Information effects on targeted-source updating and sustainability assessments	Randomizes respondents into a debt-trajectory information group or a source-specific historical vignette group about safe-asset status or rollover fragility.
Module 3	Sustainability beliefs and planned behavior	Elicits ten-year crisis probability, the most likely crisis-timing category (including "Never"), maximum sustainable debt/GDP, and planned bond-portfolio adjustment for bond investors.

Notes. Blocks are administered in the order listed, and the item-by-item wording follows in this appendix. The screen precedes the preamble, so eligibility reflects prior awareness. The treatment block sits between Modules 2 and 3. It can move sustainability beliefs, while the valuation allocation has already been elicited and is not repeated.

Question 2. Respondents provide a numeric estimate of the current debt held by the public as a share of GDP. This item serves two roles. First, it measures factual perceptions about the fiscal position. Second, for respondents not assigned to the debt-trajectory information group, it provides the personalized anchor used later in the maximum-sustainable-debt item.

Question 3. Respondents report their level of concern about U.S. government debt sustainability on an ordered categorical scale.

Question 4. Respondents who report any debt-sustainability concern answer a sample-specific behavior item. In the voter sample, the item measures how much debt concern has affected voting decisions, using an ordered scale from no influence to decisive influence, followed by a brief explanation of why debt concern was decisive, was one factor among many, or had little to no influence. In the bond-investor sample, the item asks for an open-ended account of any portfolio or investment decisions affected by debt-sustainability concern. The item is skipped for respondents with no concern.

For the retrospective cross-sample comparison, decisive voting is the voter action margin. Investor explanations are coded as no portfolio change, greater caution or monitoring without a concrete change, or a concrete portfolio change. The last of these is the investor action margin. These outcomes align in timing and behavioral meaning, but not in elicitation format: voters select an ordered category, whereas investor outcomes are coded from open text.

A.4. Beliefs about debt valuation

The valuation module measures how respondents decompose the current market value of U.S. government debt across six sources: expected future primary surpluses, permanent debt rollover, government real assets, Fed purchases, global safe-asset demand, and financial repression. Section 2 describes the construction of this menu from open-ended responses and the government-debt literature; Appendix B reports representative mappings from spontaneous theories to these categories.

Question 5. Respondents first identify the source or sources they view as providing the greatest support for the current market value of U.S. government debt. The six sources correspond to the valuation channels analyzed throughout the paper: expected future primary surpluses, permanent debt rollover, government real assets, Fed purchases, global safe-asset demand, and financial repression. The Qualtrics instrument randomizes the display order of all six source choices independently for each respondent. No source is held in a fixed position, and the screen offers no "other," "none," or "don't know" option, so every respondent chooses among the same six substantive sources.

Question 6a–f. For each source selected in Question 5, respondents provide a brief open-ended rationale. These source-conditional rationales are used to assess whether the selected source reflects substantive reasoning and to characterize the language used within each valuation framework.

Question 7. Respondents next identify which remaining sources they view as providing little or no support for current debt value. The choices carry forward the sources left unselected in Question 5 and are randomized again, independently of the order used on the previous screen. This best/worst structure is used to measure both positive support and source-level disagreement.

Question 8. Respondents then complete a constrained allocation task. The selected greatest-support sources appear as separate rows, together with one "All others" row that is always displayed last and combines every non-top source. Because the displayed rows are the respondent's own selections carried forward from Question 5, this screen is not randomized further. Respondents distribute 100 points across those displayed rows to reflect the relative contribution of each source to the current market value of U.S. government debt. The survey therefore records source-specific points for the selected top sources and one aggregate residual for all remaining sources; Section 2.3 describes the benchmark reconstruction, and Appendix Table A11 reports sharp bounds that leave the residual allocation unrestricted.

A.5. Information treatments

The survey randomizes respondents into four mutually exclusive groups: a control group, a debt-trajectory information group, a safe-asset-stress vignette group, and a rollover-fragility vignette group.

Question 9. *Treatment A: debt trajectory.* Respondents assigned to the debt-trajectory group are shown Figure A1; respondents in the other groups do not see this display. Because the treatment appears before Question 13, it also determines whether the later maximum-sustainable-debt prompt is anchored to the official debt/GDP value or to the respondent's own earlier estimate.

Question 10. *Treatments B and C: source-specific historical vignettes.* Respondents assigned to Treatment B see the safe-asset-stress vignette, and those assigned to Treatment C see the rollover-fragility vignette. Assignment is independent of earlier source selections in Question 5. Both screens appear below as fielded, together with the directional follow-up and the open-ended rationale prompt.

Treatment B: safe-asset stress.

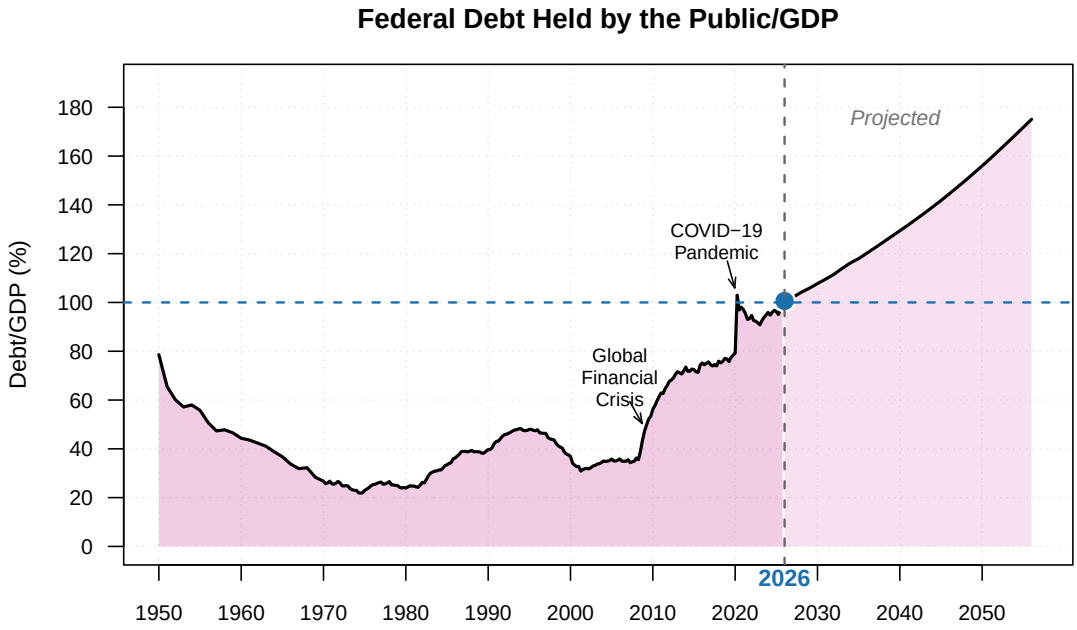


Figure A1. Debt-trajectory information shown in Treatment A. Respondents in the debt-trajectory group see the historical path of U.S. federal debt held by the public together with the CBO's long-run projection under current policy (Congressional Budget Office 2026); respondents in the control and vignette groups do not. The display highlights the current level of about 100 percent of GDP, which then serves as the reference in the maximum-sustainable-debt question for this group, whereas respondents in the other groups are referred back to their own earlier estimates.

Treatment B. *According to economic historians, U.K. government debt used to play the role of the* ***world's leading safe asset****, accounting for about 50% global foreign reserves in the 1950s, but that* ***declined*** *to about 10% by the 1970s, as Britain's relative economic, fiscal, and geopolitical position weakened.*

After reading this information, consider the role of **global safe-asset demand** (investors worldwide valuing U.S. government debt for its safety and liquidity) in supporting U.S. government debt value. Which of the following best describes how your view has changed?

a. I now think global safe-asset demand plays a **smaller** role in supporting U.S. government debt value than I previously thought.

b. This information does **not** change my view of the role of global safe-asset demand.

c. I now think global safe-asset demand plays a **larger** role in supporting U.S. government debt value than I previously thought.

Follow-up [not separately numbered]: In one or two sentences, what was your reasoning for this answer?

Treatment C: rollover fragility.

Treatment C. *Historically,* ***rolling over*** *government debt* ***can be fragile****. In 1976, the UK experienced a debt market turmoil in which investors pulled back from buying new government debt, causing long-term government bond prices to fall by about 20% within months.*

After reading this information, consider the role of **debt rollover** (the government issuing new debt to pay off old debt indefinitely) in supporting U.S. government debt value. Which of the following best describes how your view has changed?

a. I now think debt rollover plays a **smaller** role in supporting U.S. government debt value than I previously thought.
b. This information does **not** change my view of the role of debt rollover.
c. I now think debt rollover plays a **larger** role in supporting U.S. government debt value than I previously thought.

Follow-up [not separately numbered]: In one or two sentences, what was your reasoning for this answer?

A.6. Beliefs about debt sustainability

Question 11. Respondents report a subjective probability that the United States experiences a government-debt crisis within the next ten years. The survey defines a debt crisis as a severe decline in debt value associated with default or very high inflation.

Question 12. Respondents select the most likely timing of a U.S. government-debt crisis from six fielded options: 0–5, 6–10, 11–20, 20–30, more than 30 years, or "Never." The "Never" option makes the response unconditional on a crisis occurring. The six options cover every horizon, as required by the probability bounds in Section 5.1; the third and fourth options share the boundary year 20.

Question 13. Respondents report the highest debt-to-GDP ratio they believe the United States can sustain without a debt crisis. For trajectory-treated respondents, the prompt references the

Table A2. Treatment assignment and measured outcomes

Treatment component	Assignment	Outcomes used in the paper
Debt trajectory	Randomized over the full sample	Crisis probability, crisis timing, maximum sustainable debt/GDP, and planned portfolio adjustment for bond investors.
Safe-asset stress	Randomized over the full sample	Direction of updating about safe-asset demand, open-ended update rationale, crisis probability, crisis timing, maximum sustainable debt/GDP, and planned portfolio adjustment for bond investors.
Rollover fragility	Randomized over the full sample	Direction of updating about rollover, open-ended update rationale, crisis probability, crisis timing, maximum sustainable debt/GDP, and planned portfolio adjustment for bond investors.

Notes. Assignment is a single four-way randomization over these three groups and a control group that sees no information between Modules 2 and 3. Covariate balance across the four groups is reported in Table A16. The directional updating item is the only post-treatment belief measure about the targeted source; the 100-point valuation allocation is not re-elicited, so a continuous reallocation cannot be observed.

official current debt/GDP value shown in the treatment. For respondents in the control group and in Treatments B and C, the prompt references their own earlier estimate from Question 2. The resulting variable is the subjective debt limit $b^i_{\max}$ analyzed in Section 5.

Question 14. Bond-investor respondents report whether, based on their debt-sustainability views, they plan to adjust their holdings of U.S. government bonds or bond funds over the next 12 months. This item is not shown to the voter sample.

B. Mapping Spontaneous Theories of Debt Value

This appendix maps free-text answers in the open-ended backing survey onto the categories reported in Section 2.1. It defines the six main-survey sources, reports answers outside those sources, resolves ambiguous phrases, and assesses agreement across three model passes. The mapping shows how respondents' language generates the main-survey menu.

B.1. Coding principle

The open-ended backing survey asks respondents to state, without a menu, what ultimately backs the value of U.S. government debt. We use these answers to recover the spontaneous theories that respondents bring to the debt-valuation question before they see the six-source allocation task in the main survey.

For the main comparison, each response is assigned one primary mechanism. When a response names several mechanisms, we code the one receiving the most emphasis. If emphasis is unclear, an explicit mechanism takes precedence over a vague category. For example, the government's power to tax is coded as primary surpluses; generic economic strength remains a separate category. Responses naming no mechanism are recorded as folk, instrument-only, or vague/don't-know answers.

Table A3 shows how the substantive codes correspond to the six valuation sources used in the main survey. The examples are representative response fragments from the open-ended backing survey, with respondent identifiers removed.

B.2. Off-menu and residual answers

The open-ended survey also produces answers that are informative but do not fit cleanly into one of the six valuation sources. We keep these categories separate because folding them into the six sources would overstate how many formal mechanisms respondents supplied on their own. Table A4 summarizes the residual coding.

B.3. Prior awareness, concern, and spontaneous backing theories

The open-ended backing survey also records two pre-question variables: whether respondents had previously thought about what gives value to U.S. government debt, and how concerned they were about debt sustainability. Table A5 reports how the spontaneous backing theories vary with these two variables across all 1,752 responses. The awareness gradient holds separately in each population: respondents who would pass the later screen are more likely to name one of the six mechanisms and less likely to give a vague answer. This is the pattern that motivates the screen in the main survey; the table reports the corresponding magnitudes.

B.4. Disambiguation rules

Several phrases require consistent treatment. First, references to "the economy" or "GDP" are not coded as primary surpluses unless the respondent explicitly names taxes, revenue, or spending

Table A3. Mapping spontaneous theories to the six valuation sources

Main-survey source	Coding rule in open text	Representative response fragments
Expected future primary surpluses	Explicit taxes, tax revenue, taxpayers, the power to tax, spending cuts, or other fiscal extraction.	"ability to collect taxes"; "income taxes"; "American tax payers"; "raising revenue and paying its obligations."
Permanent debt rollover	Continued borrowing, refinancing, promise to repay, full faith and credit, creditworthiness, or trust that the government will keep paying, when not framed mainly as foreign or global demand.	"full faith"; "the word of the government and the credibility to repay it"; "the fact that the government has never not held to their word."
Government real assets	Government land, property, natural resources, oil, commodities, or the country itself described as collateral or backing.	"oil"; "gold and property assets"; "resources to back their debt"; "Is America its self its ultimate backing?"
Fed purchases	Federal Reserve purchases, money creation, monetary sovereignty, printing money, currency issuance, or gold/gold-reserve language used as a monetary backing theory.	"the Federal Reserve"; "print money"; "the U.S.'s own currency"; "gold reserves."
Global safe-asset demand	Reserve-currency status, global demand for dollars or Treasuries, foreign holders, safety and liquidity, geopolitical power, or institutional trust.	"the dollar being the reserve currency of the world"; "strength and safety of US debt"; "other countries holding the dollar"
Financial repression	Captive or regulated demand, especially banks or other institutions being required or pressured to hold government debt.	No response clearly stated this mechanism in the final primary classification.

Notes. Rules used to map free-text answers in the separate open-ended backing survey onto the six sources the main survey names on its menu, so the two elicitations can be compared. Respondents saw no menu, and each answer receives one primary code; answers falling outside these six are coded in Table A4. Fragments are verbatim and lightly truncated. Every response is classified independently by three models, with reliability reported in Appendix B.5.

Table A4. Off-menu responses in the open-ended backing survey

Residual code	Coding rule	Representative response fragments
Generic economy / GDP	The economy, GDP, national wealth, or economic strength without an explicit tax or asset-backing mechanism.	"the overall economic output of the country"; "backed by the country with the strongest economy."
Production / exports	Trade, exports, production, or goods and services sold abroad.	"trade is the primary source"; "its exports"; "oil, technology, automobiles, steel and other commodities."
Labor / citizens	Workers, citizens, labor, or the American people as the source of value without an explicit tax mechanism.	"The American worker"; "The American people."
Stocks / equity markets	The respondent identifies the stock market or public equities as backing government debt.	"public stocks and securities"; "the stock market."
Instrument-only	The answer names the debt instrument but not the backing mechanism.	"Treasuries"; "Bonds and treasuries"; "Treasury bonds."
Don't know / vague	The respondent gives no substantive theory or says they do not know.	"I have no idea"; "I'm not sure"; "I don't know enough about financial markets."

Notes. Codes for answers in the separate open-ended backing survey that fall outside the six menu sources of Table A3. They remain separate categories, so the open-ended and menu-based distributions are not mechanically aligned. Fragments are verbatim and lightly truncated. Classification reliability is reported in Appendix B.5.

Table A5. Spontaneous backing theories by prior awareness and concern

	Full sample ($N = 1752$)	Would pass screen ($N = 1424$)	Would fail screen ($N = 328$)
Panel A: By prior awareness			
Theory mechanism	70.0	73.2	56.4
Folk / residual answer	16.6	16.6	16.2
Don't know / vague	13.4	10.2	27.4
Primary surpluses	21.6	23.6	12.8
Global safe-asset demand	21.0	21.2	20.1
Permanent rollover	15.2	16.1	11.6
Fed / monetary backing	8.8	8.7	9.1
Government assets	3.4	3.6	2.7
Financial repression	0.0	0.0	0.0
	Not concerned ($N = 117$)	Slightly concerned ($N = 451$)	Moderately / very concerned ($N = 1184$)
Panel B: By sustainability concern			
Theory mechanism	61.5	65.4	72.6
Folk / residual answer	13.7	18.0	16.3
Don't know / vague	24.8	16.6	11.1
Primary surpluses	12.8	19.1	23.4
Global safe-asset demand	26.5	20.0	20.9
Permanent rollover	12.8	14.2	15.9
Fed / monetary backing	5.1	8.2	9.4
Government assets	4.3	4.0	3.1
Financial repression	0.0	0.0	0.0

Notes. Open-ended debt-value-backing survey: 750 bond investors, 752 general-population respondents, and 250 Economics/Finance respondents. Entries are percentages within each column. Panel A splits respondents by whether they would have passed the later main-survey screen: "would pass screen" includes respondents who answered that they had thought about what gives value to U.S. government debt either "a lot" or "somewhat"; "would fail screen" includes respondents who answered "No." Panel B splits respondents by their stated concern about U.S. debt sustainability. The theory rows are the six categories used in the main survey; folk / residual answers include generic economy, production, labor, stock-market, and instrument-only responses; don't know / vague answers name no substantive backing mechanism.

adjustment. The economy is the tax base, not the fiscal mechanism. Second, trust or full-faith language is coded as rollover when it refers to the government's promise or ability to keep paying, but as safe-asset demand when the response emphasizes foreign buyers, global reserve status, or international confidence. Third, gold language is coded as Federal Reserve or monetary backing unless the response explicitly emphasizes government-owned property or natural resources as collateral. Fourth, references to banks as holders are coded as financial repression only when the answer points to banks or regulated institutions as a captive source of demand; generic buyers or investors are coded as safe-asset demand.

These rules preserve the paper's distinction between what the two surveys measure. The open-ended survey measures which theories are cognitively available without researcher prompting. The main survey then presents a common six-source menu, records qualitative source selections and rationales, and elicits quantitative weights for selected top sources plus an aggregate residual. Section 2.3 documents the benchmark reconstruction, and Appendix Table A11 gives sharp bounds that leave the residual allocation unrestricted.

B.5. Coding reliability

The primary LLM-assisted pass, run with ChatGPT, classifies all 1,752 responses: 750 bond investors, 752 general-population respondents, and 250 respondents in the Economics/Finance benchmark. Claude and Gemini provide independent passes over the same responses, so every record is classified three times. All three models assign exactly the same category to 1,432 responses (81.7%). Pairwise exact agreement ranges from 85.2% to 89.6%; after collapsing categories to theory, folk explanation, and vague/don't know, agreement ranges from 91.3% to 94.1%. The paper uses the complete primary pass, while these independent passes assess whether the substantive distribution depends on one classifier.

One human-adjudication step follows the independent passes, and it concerns financial repression alone. In total, 8 responses were tagged as financial repression by at least one model: ChatGPT tagged 8, Claude 6, and Gemini 6, and every such tag fell on the same 8 responses. Each response invokes banks or banking without stating that regulation or compulsion creates captive demand; examples include "its banking," "money in banks," and "Maybe banking." Applying the disambiguation rule in Appendix B.4, we reclassify 6 as don't-know, 1 as safe-asset demand, and 1 as Fed purchases. Each change is recorded with its original tag and reason, and the changes are applied to the analysis outputs while the raw model files are left untouched, so the agreement statistics reported here are computed on unadjudicated codes. No model tagged any response beyond these, so the absence of spontaneous financial-repression reasoning is a property of the responses rather than of one classifier or of the review step.

Agreement is highest in the Economics/Finance sample. All three models agree exactly on 216 of the 250 Economics/Finance responses (86.4%), against 81.3% for bond investors and 80.6% for voters; coarse agreement ranges from 92.4% to 96.8%. Respondents with formal training write answers that name a mechanism explicitly, leaving the classifier less to adjudicate. The Economics/Finance shares in Table A8 are therefore the least classifier-dependent numbers in the open-ended analysis.

C. Newspaper Topic Salience

This appendix documents the newspaper comparison used in Section 3.4. It describes the corpus and classification rubric and reports the resulting mechanism frequencies.

C.1. Corpus construction

We construct the corpus from ProQuest records for the *Wall Street Journal*, the *New York Times*, and the *Washington Post* from January 1, 2016 through December 31, 2025. The search runs over every article these three newspapers published in that window, and we retain every record it returns for the following terms:

> "national debt" OR "federal debt" OR "public debt" OR "government debt" OR "debt-to-GDP" OR "debt sustainability" OR "Treasury yield*" OR "bond yield*" OR "sovereign debt" OR "Treasuries" OR "federal borrowing" OR "interest on debt" OR "long-term yields".

C.2. LLM-based relevance and mechanism classification

The broad search, which returns 16,591 debt-related articles, also includes articles that mention debt or Treasury markets without necessarily discussing what supports debt value, so a second step separates the two. We send each full article to an LLM through the OpenAI API under a common classification rubric. The LLM's first task is to identify articles that connect fiscal policy, market demand, monetary policy, regulation, or refinancing conditions to the price, yield, sustainability, rollover risk, or expected repayment value of U.S. federal debt. The rubric excludes passing debt mentions, routine auction calendars, and articles concerned only with consequences running from debt or yields to other outcomes.

For each relevant article, the LLM's second task is to assign one primary mechanism: primary surpluses, global safe-asset demand, Fed purchases, debt rollover, financial repression, government assets, or "Other." When an article contains several mechanisms, the rubric instructs the LLM to

record the one receiving the main explanatory emphasis. The categories use the same economic definitions as the survey. Safe-asset demand includes safety, liquidity, reserve-currency demand, and flight to quality. Fed purchases require Federal Reserve purchases, sales, holdings, reinvestment, QE, QT, or balance-sheet policy. Primary surpluses cover taxes, spending, deficits, fiscal capacity, and repayment. Rollover covers refinancing maturing debt. Financial repression requires policy-induced demand through institutions or regulation. Government assets cover public assets that could support repayment. Classification runs on `gpt-4o-mini` at temperature zero.

After classification, the corpus contains 5,654 valid articles relevant to U.S. federal debt valuation. Every relevant row carries a nonempty supporting quotation. Table A6 reports the pooled mechanism distribution and its breakdown by newspaper. Of the 5,654 relevant articles, 71 are classified as "Other." Primary surpluses rank first among the six common mechanisms in all three newspapers.

Table A6. Mechanism distribution in valuation-relevant newspaper articles, pooled and by source. Normalized shares are divided by the number of valuation-relevant articles assigned to one of the six mechanisms, pooled in the "All" column and within each newspaper in the final three columns. WSJ denotes the *Wall Street Journal*, NYT the *New York Times*, and WaPo the *Washington Post*. The "Other" row reports its article count but is excluded from all four normalized-share columns.

Mechanism	Articles	Normalized share (%)			
		All	WSJ	NYT	WaPo
Primary surpluses	3,314	59.4	47.5	68.7	82.7
Global safe-asset demand	1,096	19.6	26.9	12.0	8.3
Fed purchases	876	15.7	20.4	13.5	4.4
Debt rollover	252	4.5	4.3	5.2	4.2
Financial repression	25	0.4	0.6	0.3	0.3
Government assets	20	0.4	0.4	0.4	0.2
Other	71	–	–	–	–

D. Additional Empirical Evidence

This appendix reports supporting evidence on sample construction, the valuation module, randomization, open-text responses, and sustainability beliefs. It also reports the audit, sharp bounds, and sensitivity analysis for the "All others" reconstruction used to calculate value shares.

D.1. Samples and waves

Detailed recruitment and sample composition. The bond-investor sample pools a June 2026 wave ($N = 491$) and a July 2026 wave ($N = 494$). U.S. residents were recruited from Prolific's panel using its investment prescreener, and recruitment was restricted to respondents who report having invested in "Government bonds." The voter/general-population sample also pools June and July 2026 waves. Respondents were recruited from Prolific's U.S. general-public panel using its voter-registration prescreener, and the matched Prolific field confirms registered-voter status for every cleaned respondent.

The Economics/Finance sample was recruited in July 2026. Of its $N = 247$ respondents, 230 hold a master's-level graduate degree and 17 a doctorate; 128 report Economics and 119 Finance as the subject of their highest completed degree. The sample measures completed formal training; occupation was not an eligibility condition.

Multiple comparisons. When the paper compares samples on several related quantities, it reports Holm-adjusted p-values and states the comparison family at each use. Holm's step-down procedure controls the probability of any false rejection within the family and remains valid when comparisons are dependent because the same respondents enter several of them. Pairwise comparisons use Fisher's exact test for categorical shares and Welch's t-test for means unless otherwise stated.

Prior awareness within the retained samples. The opening screen removes respondents who had never thought about what supports the market value of U.S. government debt, but the retained samples still differ in how much they had thought about it. Table A7 splits each sample by the "a lot" and "somewhat" answers to Question 1 and reports the value shares, fiscal perceptions, concern, crisis probability, and action margins used in the main text. Section 3.4 draws on this split.

D.2. Valuation evidence and the residual reconstruction

Table A8 compares the Economics/Finance sample with the two main samples across structured sources, open-text categories, and sustainability outcomes. Table A9 gives investor and voter descriptives. The next three tables audit the aggregate "All others" row, which the survey records as one number and the reconstruction divides across sources, and report what is identified when that row is allocated without imposing the benchmark rule.

Table A7. Opening awareness screen and main outcomes

	Bond investors		Voters	
	Thought a lot	Thought somewhat	Thought a lot	Thought somewhat
N	326	659	212	789
Mean concern, 1–4	3.4	3.0	3.5	3.0
Moderately / very concerned, %	89.9	75.1	92.0	71.1
Very concerned, %	52.5	27.8	57.5	30.3
Perceived current debt/GDP, %	73.8	62.2	66.1	57.9
Primary-surplus share θ, of 100	16.2	13.9	12.3	14.8
Safe-asset share θ, of 100	29.3	29.6	23.8	26.2
Rollover share θ, of 100	15.2	17.7	19.1	17.5
P(crisis within 10y), %	45.1	46.7	56.8	53.5
Plan reduce / decisive vote, %	36.5	27.6	17.3	6.5

Notes. Main-survey respondents all passed the initial screen, which excludes respondents who say they had never thought about what supports the market value of U.S. government debt. The table splits the retained samples by whether they had thought about the question "a lot" or "somewhat." Value shares are benchmark reconstructed shares out of 100 points. "Plan reduce / decisive vote" is the share of bond investors planning to reduce U.S. government bond or bond-fund holdings over the next twelve months, and the share of voters saying debt sustainability was decisive for voting. The split is a baseline cross-section and is not randomized.

Table A8. The Economics/Finance sample across valuation and sustainability outcomes

	Bond investors	Voters / gen. population	Grad. Econ./Finance	$p(E=B)$	$p(E=V/G)$
Panel A: Mean structured value share (main survey)					
Global safe-asset demand	29.5	25.7	26.0	0.110	0.886
Debt rollover	16.9	17.8	18.8	0.838	0.838
Fed purchases	15.1	15.7	17.6	0.332	0.463
Government assets	14.1	17.1	13.3	0.593	0.028
Financial repression	9.7	9.4	8.4	0.678	0.691
Primary surpluses	14.7	14.3	16.0	0.888	0.888
Panel B: Unprompted primary category (separate open-text survey, percent)					
Primary surpluses	19.1	22.1	27.6	0.017	0.169
Global safe-asset demand	21.1	19.3	26.0	0.229	0.076
Debt rollover	21.3	10.0	12.8	0.006	0.237
Fed / monetary backing	9.6	9.2	5.2	0.106	0.106
Government assets	3.1	4.1	2.4	0.748	0.748
Financial repression	0.0	0.0	0.0	1.000	1.000
Other folk explanation	16.0	18.0	14.0	0.672	0.519
Don't know / vague	9.9	17.4	12.0	0.340	0.093
Panel C: Knowledge and sustainability beliefs (main survey)					
Absolute debt/GDP error	46.4	50.1	41.0	0.088	0.002
Concern (1–4)	3.13	3.08	2.63	< 0.001	< 0.001
Ten-year crisis probability	46.2	54.2	41.2	0.020	< 0.001
Maximum sustainable debt/GDP	105.1	78.2	124.2	0.003	< 0.001
Headroom	29.9	9.5	35.1	0.313	< 0.001

Notes. Panels A and C use the main-survey samples ($N = 985$, 1,001, and 247); their second column is the registered-voter sample. Panel B uses the separate unprompted open-text samples ($N = 750$, 752, and 250); its second column is the general-population sample. The last two columns compare the Economics/Finance sample with the bond-investor and second comparison samples shown in each panel. For each structured source, open-text category, or outcome, the three pairwise population comparisons are Holm adjusted. Open-text categories use the final primary classification from the completed coding pass; no response in any sample was classified primarily as financial repression. Comparisons are descriptive and do not identify a causal effect of education or training.

Table A9. Perceived value-share decomposition, main surveys

	Bond investors ($N = 985$)			Voters ($N = 1001$)		
Source	Mean	Median	% zero	Mean	Median	% zero
Global safe-asset demand	29.5	30	27	25.7	15	31
Debt rollover	16.9	0	51	17.8	0	51
Fed purchases	15.1	5	39	15.7	5.7	40
Primary surpluses	14.7	0	53	14.3	0	51
Government assets	14.1	0	53	17.1	4	45
Financial repression	9.7	0	51	9.4	0	50

Notes. Benchmark reconstructed value shares θ_k^i from Question 8. The survey records points for each selected top source and one aggregate "All others" residual. The benchmark sets little/no-support sources to zero and divides the residual equally among the remaining non-top sources; the two voter responses with no remaining neutral source use all non-top sources. Shares sum to 100 within respondent. Columns report the cross-sectional mean, median, and benchmark percentage at zero. Rows are ordered by the bond-investor mean. Tables A10–A12 report the reconstruction audit, sharp bounds that do not impose this allocation rule, and sensitivity to three alternative rules.

Reconstruction of the "All others" residual. The allocation screen shows each selected top source separately but combines every non-top source into one "All others" row. Table A10 quantifies that residual. It is positive for 54–62 percent of respondents but averages only 12–14 of the 100 points; the remaining 86–88 points are assigned directly to named top sources. The benchmark reconstruction preserves the 100-point total for every respondent. Two voter responses assign positive points to "All others" after marking every non-top source as little or no support; for those two cases only, the residual is divided across all non-top sources so that the observed 100-point total is preserved.

Table A10. Audit of the aggregate "All others" residual

Sample	N	Residual > 0 (%)	Mean	Median	P75	Non-top (mean N)	Sum $= 100$ (%)
			Residual points				
Bond investors	985	54.6	12.3	7.0	20.0	4.1	100.0
Voters	1001	54.1	13.5	6.0	21.0	4.1	100.0
Grad. Econ./Finance	247	61.9	13.0	10.0	20.0	4.1	100.0

Notes. "Residual" is the number of the respondent's 100 points assigned to the single "All others" row. "Non-top" is the number of sources not selected as greatest support. "Sum = 100" is the percentage of respondents whose six benchmark reconstructed shares sum to 100. The benchmark assigns zero to sources marked little or no support and divides the residual equally among the remaining non-top sources. If no such source remains, it divides the residual among all non-top sources.

Table A11 separates conclusions identified by the instrument from conclusions that require a residual-allocation rule. For each source, the lower mean bound assigns none of any respondent's residual to that source; the upper bound assigns all of it whenever the source is non-top. The bounds are sharp because residual allocations are respondent-specific and otherwise unrestricted. The zero-share bounds are constructed analogously. The final column bounds the mean safe-

asset share minus each other source's share. Safe-asset demand exceeds primary surpluses in all three samples throughout the identified set. It also exceeds every source throughout the investor identified set. The voter bounds do not rule out rollover or government assets exceeding safe-asset demand, and the Economics/Finance bounds do not rule out rollover exceeding it. The qualitative valence and unprompted results therefore provide independent evidence for safe-asset demand's broader prominence.

Table A11. Sharp bounds under unrestricted allocation of the non-top residual

Source	Bench.	Equal-all	Mean bounds	Bench. zero	Zero bounds	Safe minus source
	Mean points			Percent		Mean bounds
Bond investors						
Global safe-asset demand	29.5	29.3	[28.0, 33.7]	27.1	[21.1, 42.4]	–
Debt rollover	16.9	17.2	[15.2, 23.7]	50.7	[32.0, 69.7]	[4.3, 18.5]
Fed purchases	15.1	14.5	[12.3, 21.5]	38.9	[30.1, 70.6]	[6.5, 21.3]
Government assets	14.1	14.5	[12.3, 21.6]	52.7	[31.4, 72.7]	[6.4, 21.4]
Financial repression	9.7	9.6	[7.2, 17.2]	50.8	[34.6, 80.1]	[10.8, 26.5]
Primary surpluses	14.7	14.9	[12.6, 22.1]	53.0	[35.1, 75.2]	[5.9, 21.0]
Voters						
Global safe-asset demand	25.7	25.4	[23.6, 31.3]	31.2	[24.6, 52.4]	–
Debt rollover	17.8	18.5	[16.4, 25.3]	50.6	[31.0, 66.6]	[-1.8, 14.8]
Fed purchases	15.7	15.2	[13.0, 22.8]	40.2	[32.0, 70.0]	[0.8, 18.3]
Government assets	17.1	17.3	[14.9, 24.9]	45.3	[27.7, 68.6]	[-1.4, 16.3]
Financial repression	9.4	9.3	[6.8, 17.6]	50.4	[36.6, 80.0]	[6.0, 24.5]
Primary surpluses	14.3	14.4	[11.9, 22.8]	51.3	[34.9, 76.4]	[0.8, 19.4]
Grad. Econ./Finance						
Global safe-asset demand	26.0	25.9	[24.2, 31.1]	29.6	[18.2, 49.0]	–
Debt rollover	18.8	19.1	[17.1, 25.7]	45.3	[26.7, 66.0]	[-1.4, 14.0]
Fed purchases	17.6	17.3	[15.3, 23.9]	35.6	[23.5, 64.4]	[0.4, 15.8]
Government assets	13.3	13.4	[11.0, 21.0]	50.6	[25.1, 74.9]	[3.3, 20.1]
Financial repression	8.4	8.5	[5.9, 16.9]	52.6	[32.0, 84.2]	[7.3, 25.2]
Primary surpluses	16.0	15.8	[13.5, 23.2]	44.1	[30.4, 72.9]	[1.0, 17.6]

Notes. "Bench." uses the benchmark reconstruction. "Equal-all" divides the residual equally among all non-top sources, including sources marked little or no support. Mean and zero-share bounds allow the residual to be allocated arbitrarily and non-negatively across all non-top sources. The last column gives sharp bounds on the mean safe-asset share minus the row source; a positive lower bound establishes that safe-asset demand has the larger mean under every admissible allocation. The zero-share bounds are correspondingly wide, because the allocation is permitted to place weight on sources the respondent marked little or no support.

Table A12 reports three disciplined alternatives to the benchmark, giving four rules in total. "Equal across all non-top" ignores the little/no distinction. "Qualitative weighted" gives each neutral source twice the residual weight of each little/no-support source. "Pooled top-frequency weighted" divides each respondent's residual in proportion to the source's pooled probability of being selected as greatest support. Safe-asset demand ranks first, and primary surpluses remain near 14–16 points under every rule. At least 98.9 percent of investor, 99.2 percent of voter, and 99.2 percent of Economics/Finance dominant-source assignments agree with the benchmark; non-fiscal dominant shares remain 83–86 percent. Re-estimating the pooled $k = 5$ partition under the three alternatives preserves the same five centroid-dominant sources, with adjusted Rand indices of 0.914–0.982 relative to the benchmark.

Table A12. Central results under alternative residual-allocation rules

Residual rule	Safe mean	Safe + rollover Points	Surplus mean	Surplus zero (%)	Safe rank	Non-fiscal (%)	Type agreement (%)
Bond investors							
Benchmark	29.5	46.4	14.7	53.0	1	85.1	100.0
Equal across all non-top	29.3	46.5	14.9	35.1	1	85.0	98.9
Qualitative weighted	29.4	46.5	14.8	35.1	1	85.1	99.4
Pooled top-frequency weighted	30.1	47.5	14.5	35.1	1	85.1	99.0
Voters							
Benchmark	25.7	43.5	14.3	51.3	1	85.7	100.0
Equal across all non-top	25.4	43.9	14.4	34.9	1	85.6	99.2
Qualitative weighted	25.5	43.7	14.3	34.9	1	85.6	99.5
Pooled top-frequency weighted	26.4	45.1	14.0	34.9	1	85.6	99.5
Grad. Econ./Finance							
Benchmark	26.0	44.8	16.0	44.1	1	83.4	100.0
Equal across all non-top	25.9	45.0	15.8	30.4	1	83.4	99.2
Qualitative weighted	25.9	45.0	15.8	30.4	1	83.4	99.6
Pooled top-frequency weighted	26.9	46.2	15.4	30.4	1	83.4	99.2

Notes. "Safe + rollover" is the combined mean share of global safe-asset demand and permanent debt rollover, the two channels whose support comes from investors absorbing supply. "Surplus zero" is the share allocating zero to primary surpluses under each rule; it falls under the three alternatives because they do not use the little/no-support item. "Safe rank" ranks the six sample means. "Non-fiscal" is the share whose dominant reconstructed source is not primary surpluses. "Type agreement" is the agreement of the rule-specific dominant source with the benchmark. Across these rules, the largest R^2 from regressing any reconstructed share on fiscal knowledge is 0.013 for investors and 0.003 for voters. Treatment-A effects change by at most 0.11 percentage points because the safe-asset share enters only as a precision control. The unprompted distributions, qualitative valence rankings, newspaper frequencies, crisis-risk levels, stated behavior, and vignette directional responses do not use the reconstruction and are mechanically unaffected.

The reconstruction's conclusions also hold when the residual rules are combined with the paper's other robustness checks. Combining voter raking weights with the four rules leaves safe-asset demand first, at 24–25 points, and primary surpluses at 14–15 points.

Five of the six benchmark source shares do not differ between the Economics/Finance sample and either comparison sample after the within-source Holm adjustment. Voters assign more weight to government assets than the Economics/Finance sample does (17.1 versus 13.3 points; adjusted $p = 0.028$), and this difference survives all four residual rules. The primary-surplus comparisons remain insignificant under every rule. The evidence supports a common safe-asset-versus-surplus pattern across samples, not equality of every source mean.

The remaining results in this subsection describe respondent-level allocations. Table A13 reports the prevalence and centroid profile of the five belief archetypes, and Table A14 and Figure A2 report each respondent's largest allocation and its lead over the runner-up.

The best/worst selections provide a reconstruction-free check on the allocation results. Table A15 reports the corresponding source rankings and measures of contestedness.

Table A13. Belief archetypes: prevalence and profile

Archetype	% Bond	% Voters	% Econ./Fin.	% Pooled	Dominant source(s) (centroid share)
Safe-asset believers	29.7%	24.9%	25.5%	27.3%	Safe-asset demand (62)
Rollover optimists	20.8%	23.9%	23.9%	22.4%	Debt rollover (58)
Primary-surplus believers	19.1%	17.3%	20.6%	18.2%	Primary surpluses (57), Safe-asset demand (20)
Government-asset believers	14.5%	17.6%	11.3%	16.1%	Govt. assets (60)
Fed-backstop believers	15.8%	16.4%	18.6%	16.1%	Fed purchases (52)

Notes. Belief archetypes from a deterministic $k = 5$ k-means partition of the pooled bond-investor and voter samples on the six benchmark reconstructed shares (centered shares, `nstart=50`; identical assignment reused across all analyses). Economics/Finance respondents are assigned out of sample to the nearest main-sample centroid, so the smaller comparison sample does not redefine the taxonomy. Columns report each archetype's share within the bond, voter, and Economics/Finance samples and within the pooled main samples ($N_{\text{bond}} = 985$, $N_{\text{voters}} = 1{,}001$, pooled $N = 1{,}986$). Pooled head counts are 542 safe-asset, 444 rollover, 361 primary-surplus, 319 government-asset, and 320 Fed-backstop. The final column gives each cluster's defining source(s), with centroid shares in parentheses. The same five centroid-dominant sources recur under the alternative residual rules in Table A12.

Table A14. Dominant valuation source: prevalence and decisiveness

Dominant source	% Bond	% Voters	% Pooled	Mean top share	% decisive
Safe-asset demand	36.6%	31.5%	34.0%	59	75%
Debt rollover	19.1%	21.9%	20.5%	59	81%
Fed purchases	11.6%	12.3%	11.9%	61	76%
Govt. assets	13.3%	16.5%	14.9%	62	85%
Fin. repression	4.5%	3.6%	4.0%	56	74%
Primary surpluses	14.9%	14.3%	14.6%	61	89%

Notes. Each respondent is assigned to the single value source receiving their largest allocated share (the argmax of the six shares); this differs from the k-means archetypes of Table A13, which cluster the full share vector. $N_{\text{bond}} = 985$, $N_{\text{voters}} = 1{,}001$, pooled $N = 1{,}986$. "Mean top share" is the average winning share within the sample; "% decisive" is the share whose winning source leads the runner-up by at least 10 points (pooled). A primarily fiscal view remains a minority: 85% of bond investors and 86% of voters lead with a non-fiscal channel.

Table A15. Source valence and contestedness from primary-versus-little selections

Source	% prim.	% little	Diffusion	Log-odds [95% CI]	Contest	Valence
Panel A: Bond investors						
Global safe-asset demand	58	13	+45	1.50 [1.32, 1.71]	0.37	1.08
Fed purchases	30	17	+12	0.54 [0.35, 0.73]	0.74	0.17
Debt rollover	31	36	−5	−0.15 [−0.31, −0.00]	0.92	-0.24
Government assets	27	36	−8	−0.27 [−0.43, −0.11]	0.87	-0.32
Primary surpluses	25	34	−9	−0.30 [−0.46, −0.14]	0.85	-0.33
Financial repression	20	31	−11	−0.44 [−0.62, −0.26]	0.78	-0.37
Panel B: Voters						
Global safe-asset demand	48	14	+34	1.26 [1.07, 1.46]	0.44	0.74
Fed purchases	30	18	+12	0.51 [0.33, 0.70]	0.75	0.16
Debt rollover	34	37	−3	−0.10 [−0.25, 0.05]	0.95	-0.21
Government assets	32	30	+2	0.07 [−0.10, 0.22]	0.97	-0.07
Primary surpluses	24	31	−7	−0.27 [−0.44, −0.10]	0.87	-0.30
Financial repression	20	29	−9	−0.36 [−0.54, −0.16]	0.82	-0.32

Notes. For each source, "% prim." is the share selecting it as providing the *greatest* support and "% little" the share selecting it as providing *little or no* support; both items allow multiple selections, so columns need not sum to 100. "Diffusion" = %prim. − %little is the primary-minus-little contrast, a net-favorability score of the kind business surveys report as a diffusion index. Log-odds = $\log(\%\text{prim.}/\%\text{little})$ is reported with a 95% percentile bootstrap interval over respondents ($B = 2000$). Contestedness = $2\min(P,Q)/(P+Q) \in [0,1]$ is the balance of opinion among the engaged (0 consensus, 1 evenly split for/against). "Valence" is the centered source fixed effect from an ordered logit of the three-level response (little < neutral < greatest) on source, a model-based measure of latent appeal. Rows are ordered by the bond-investor contrast.

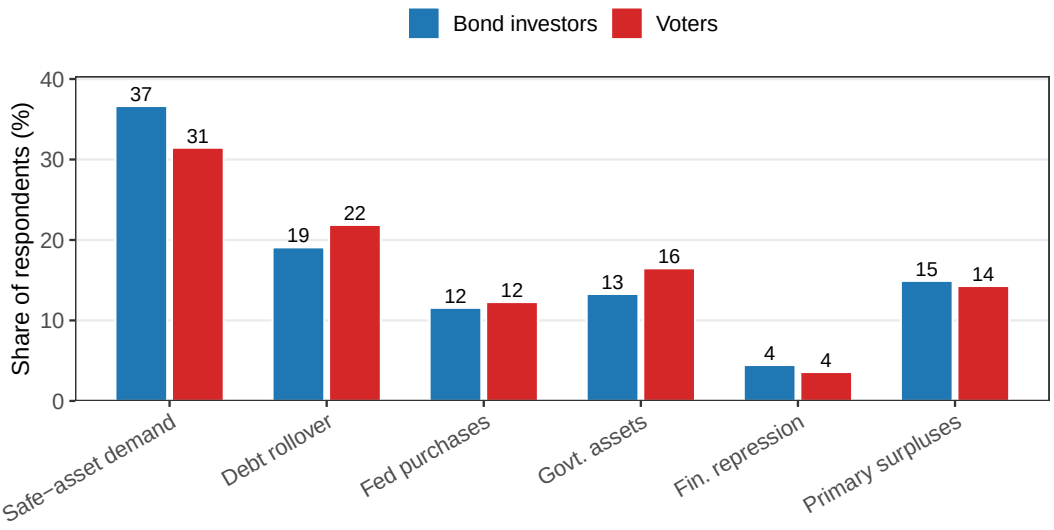


Figure A2. Distribution of the dominant valuation source by sample, where a respondent's dominant source is the one receiving their largest benchmark value share. Bars give the share of respondents in each sample, so they sum to 100 within sample. Ties at the top are rare and are reported in Table A14, which also shows that the assignment is nearly unchanged under the alternative residual-allocation rules.

D.3. Randomization balance

Table A16 reports pre-treatment balance across the four randomized groups. Group means are similar in all three samples, and every Holm-adjusted balance-test p-value exceeds 0.2, providing no evidence of imbalance in the observed covariates.

D.4. Open-text response analyses

This subsection uses the open text to check whether respondents' closed-form answers are accompanied by meaningful, channel-specific reasoning. Two sets of responses are coded: the rationales respondents write for each source they name as providing the greatest support, and the rationales they write after a mechanism vignette.

Rationales for primary debt-value sources. After selecting the sources providing the greatest support for U.S. government debt, respondents wrote a short rationale for each choice. Table A17 summarizes all 4,211 source-conditional rationales across the three samples. Almost all contain substantive answers. We flag a rationale as source-aligned if it contains language associated with the selected channel: taxes, revenues, repayment, or fiscal balance for primary surpluses; borrowing, refinancing, or issuing new debt for rollover; Federal Reserve or money-creation language for Fed purchases; safety, global demand, liquidity, dollar, or reserve-currency language

Table A16. Pre-treatment covariate balance across randomized information groups

Variable	N	Control	A	B	C	Holm p
Panel A. Bond investors						
Perceived debt/GDP	985	63.7	63.5	66.4	70.8	1.000
Debt concern (1–4)	985	3.2	3.1	3.1	3.1	1.000
Fiscal accuracy ($-$|error|)	985	-47.0	-49.1	-43.6	-45.7	1.000
Age	985	39.9	42.1	39.0	40.7	0.335
Safe-asset value share	985	29.5	29.6	29.8	29.2	1.000
Female (%)	977	44.4	40.7	40.2	45.1	1.000
Panel B. Voters						
Perceived debt/GDP	1001	58.1	63.5	56.8	60.2	1.000
Debt concern (1–4)	1001	3.1	3.0	3.2	3.0	1.000
Fiscal accuracy ($-$|error|)	1001	-48.9	-54.4	-48.2	-48.8	1.000
Age	1001	40.5	39.4	40.6	39.8	1.000
Safe-asset value share	1001	25.8	27.5	27.1	22.5	0.841
Female (%)	994	55.2	57.0	55.8	54.5	1.000
Panel C. Economics/Finance						
Perceived debt/GDP	247	79.9	67.8	93.3	83.3	0.239
Debt concern (1–4)	247	2.7	2.6	2.7	2.6	1.000
Fiscal accuracy ($-$|error|)	247	-35.5	-52.3	-37.3	-39.1	0.221
Age	247	36.1	34.7	38.0	35.5	1.000
Safe-asset value share	247	29.5	22.0	24.9	27.6	1.000
Female (%)	246	36.7	37.7	35.5	41.3	1.000

Notes. Entries are raw group means. Treatment A is the debt-trajectory information group, B the safe-asset-stress vignette, and C the rollover-failure vignette. Female is reported in percent; fiscal accuracy equals minus the absolute error in perceived debt/GDP relative to 100, so values nearer zero indicate greater accuracy. The last column reports the Holm-adjusted p-value from a joint test of equality across the four groups. The tests include collection-wave fixed effects for the bond-investor and voter samples; the Economics/Finance sample was collected once. Adjustment is across the six balance variables within each sample.

for safe-asset demand; land, resources, assets, or collateral language for government assets; and regulatory, mandated, captive-institution, or bank-holding language for financial repression. Under this deliberately imprecise keyword screen, alignment ranges from 48% to 86% across source-sample cells, and the median rationale is 18–27 words. The source selections are accompanied by substantive, channel-specific explanations.

Table A17. Open-text rationales for greatest-support source selections

Source	*N*	% subst.	% aligned	Med. words
Bond investors				
Global safe-asset demand	572	100%	79%	24
Debt rollover	300	100%	52%	22
Fed purchases	293	100%	66%	20
Government assets	270	100%	81%	21
Financial repression	196	99%	60%	21
Primary surpluses	248	100%	83%	27
Voters				
Global safe-asset demand	476	100%	71%	21
Debt rollover	337	99%	48%	21
Fed purchases	299	99%	60%	19
Government assets	316	99%	78%	20
Financial repression	199	100%	51%	19
Primary surpluses	237	100%	74%	21
Economics/Finance				
Global safe-asset demand	126	100%	86%	21
Debt rollover	84	99%	55%	19
Fed purchases	88	98%	68%	18
Government assets	63	98%	75%	18
Financial repression	41	95%	59%	22
Primary surpluses	66	98%	74%	21

Notes. Source-conditional open-text rationales in the bond-investor, voter, and Economics/Finance surveys. *N* is the number of nonempty rationales for each selected greatest-support source. "% subst." flags responses that are not blank, nonsensical, or too short to classify. "% aligned" flags substantive responses containing source-specific language for the selected channel. The alignment rule is intentionally conservative and keyword-based, so non-aligned responses should not be read as invalid; they often use general institutional language.

The keyword screen is intentionally limited. In particular, language such as "safe" can describe ordinary portfolio safety or a distinctive global role for Treasuries, or it can merely repeat the offered label. We do not use this screen to assign a finer semantic taxonomy or to claim that every source selection has the intended model interpretation.

Why the mechanism vignettes move beliefs only weakly. The source-specific vignettes were designed to stress the key assumption behind two non-fiscal channels: rollover and safe-asset demand. The average first-stage shift is small (Figure 7). We code the open-text rationales from all survey waves to examine the interpretations behind that average. Table A18 assigns each available explanation to one of six channels. A *fragility* reading accepts the intended message that rollover access or safe-asset status can fail. An *exceptionalism* reading argues that the U.S. is different from the historical case, especially because of reserve-currency status or institutional dominance. A

reinforced reading treats the example as evidence that the targeted mechanism is more important. A *prior-unchanged* reading says the vignette confirms what the respondent already believed. A *low-importance* reading says other channels matter more, and *noise* collects non-substantive answers.

The coding explains the muted treatment effects. Bond-investor rollover rationales lean toward the intended fragility reading, but safe-asset rationales split almost evenly between fragility and U.S.-exceptionalism. Voter rationales contain more prior-unchanged and noise responses than investor rationales do. Economics/Finance rationales contain less noise but the same fragility-versus-exceptionalism split. The directional choices line up with the text channels (Table A19): fragility and low-importance rationales mostly accompany "smaller role" responses, reinforced-importance rationales accompany "larger role" responses, and prior-unchanged rationales accompany "no change." Many respondents interpret the same historical information through a prior framework that exempts the United States or makes the targeted mechanism look more central.

Table A18. Open-text interpretation of mechanism vignettes

Sample	Vignette	*N*	Fragility	Exceptionalism	Reinforced	Prior unchanged	Low importance	Noise
Bond investors	Rollover	121	39%	9%	16%	12%	12%	12%
Bond investors	Safe asset	119	29%	29%	15%	12%	2%	13%
Voters	Rollover	123	33%	7%	10%	20%	8%	21%
Voters	Safe asset	123	16%	24%	15%	15%	9%	21%
Economics/Finance	Rollover	31	32%	13%	10%	23%	19%	3%
Economics/Finance	Safe asset	31	23%	35%	16%	16%	10%	0%

Notes. Open-text rationales after the rollover and safe-asset mechanism vignettes, restricted to respondents in the current cleaned samples and coded by dominant reasoning channel under a common six-category rubric. The coding covers all 240 nonempty bond-investor, 246 nonempty voter, and 62 nonempty Economics/Finance rationales. "Reinforced" means the respondent interpreted the vignette as making the targeted mechanism more important. "Exceptionalism" means the respondent exempted the United States from the historical analogy. The categories describe respondents' interpretations; they are not randomized treatment moderators.

Table A19. Text channel and stated direction after mechanism vignettes

Text channel	*N*	Smaller	No change	Larger
Fragility	160	62%	32%	6%
Exceptionalism	99	2%	87%	11%
Reinforced	75	1%	25%	73%
Prior unchanged	84	0%	100%	0%
Low importance	47	55%	45%	0%
Noise	83	19%	60%	20%

Notes. The 548 rationales in Table A18, pooled across samples, waves, and vignette groups. Rows are text-channel codes; columns show the respondent's closed-form directional response about whether the targeted source now plays a smaller, unchanged, or larger role. Percentages are row percentages.

D.5. Sustainability evidence

Table A20 separates the Treatment-A information effect from the change in the reference value induced by the same treatment. Table A21 reports the treatment effect across alternative codings of the most-likely-horizon response. The remaining results document respondent-level consistency across sustainability measures, their relation to valuation frameworks, and the gap between concern and intended action.

Table A20. Decomposing the Treatment-A headroom contrast

Treatment-A contrast	Bond investors	Voters
Reported maximum sustainable debt/GDP	17.07** (7.24)	29.47*** (4.67)
Reference used in headroom	36.62*** (2.63)	40.10*** (2.31)
Group-specific headroom	-19.55*** (6.55)	-10.63** (4.58)
Uniform headroom: $b_{\max}-$ prior estimate	17.07** (7.24)	29.47*** (4.67)

Notes. Each cell reports the Treatment-A coefficient relative to Control using the same controls and collection-wave fixed effects as Table 3, with HC1 standard errors in parentheses. The reported maximum sustainable debt/GDP ratio, $b_{\max}$, is the directly elicited ceiling. The reference used in the group-specific headroom index is 100 for Treatment A and the respondent's pre-treatment debt/GDP estimate for Control and Treatments B and C. Hence the third row equals the first row minus the second. The final row instead defines headroom uniformly as $b_{\max}$ minus the respondent's pre-treatment estimate for every group. Because that estimate also enters as a regression control, its Treatment-A coefficient equals the coefficient on $b_{\max}$. The sign reversal establishes that the negative group-specific headroom contrast is not an average treatment effect on a uniformly defined outcome. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A21. Treatment-A effects across crisis-timing measures

Timing outcome	Bond investors	Voters
Horizon within 5 years	0.46 (3.56)	0.15 (4.08)
Horizon within 10 years	2.38 (4.29)	3.14 (4.23)
Horizon within 20 years	9.22 (3.76)	3.36 (3.23)
Horizon within 30 years	7.65** (2.73)	-1.51 (2.36)
Finite horizon (not Never)	2.48 (1.78)	-1.96 (1.42)
Log representative years, finite only	-0.09 (0.08)	-0.06 (0.08)
Log representative years, Never = 1000	-0.19 (0.11)	0.03 (0.10)
Log ordered category	-0.06 (0.05)	-0.02 (0.05)

Notes. Each cell reports the Treatment-A coefficient relative to Control using the same controls and collection-wave fixed effects as Table 3, with HC1 standard errors in parentheses. The first five outcomes are indicators scaled to percentage points. The final three use the ordered timing response: representative years assign 2.5, 8, 15.5, 25, and 50 years to the five finite categories; the first log specification drops "Never," the second assigns it 1000 years, and the last logs the category number itself. Stars use Holm-adjusted p-values across the eight specifications within each sample. The investor response is concentrated at the 20- and 30-year cutoffs; no specification is statistically significant among voters. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The ceiling-implied horizon coefficients in the debt-trajectory group are 0.28 for bond investors, 0.23 for voters, and 0.10 for Economics/Finance respondents; the last estimate is imprecise

Table A22. Individual-level alignment across sustainability measures

	Bond investors	Voters	Econ./Finance	All samples
Panel A. Dependent variable: reported horizon within ten years (100/0)				
Ten-year crisis probability (per 10 pp)	9.18***	7.27***	8.41***	8.25***
	(0.37)	(0.41)	(0.87)	(0.26)
Observations	984	1000	247	2231
R^2	0.334	0.236	0.290	0.292
Panel B. Dependent variable: reported horizon category, 1–5 (debt-trajectory group)				
Ceiling-implied horizon category	0.277***	0.226***	0.104	0.233***
	(0.050)	(0.057)	(0.097)	(0.035)
Observations	246	249	61	556
R^2	0.164	0.088	0.090	0.120
Panel C. Dependent variable: reported horizon category, 1–5 (all groups)				
Headroom-implied horizon category	0.239***	0.231***	0.194***	0.233***
	(0.028)	(0.032)	(0.055)	(0.020)
Observations	984	1000	247	2231
R^2	0.108	0.092	0.089	0.108
Age and sex controls	Yes	Yes	Yes	Yes
Treatment-group fixed effects (Panels A and C)	Yes	Yes	Yes	Yes
Collection-wave controls	Yes	Yes	–	Yes
Education controls	–	Yes	Yes	–
Sample fixed effects	–	–	–	Yes

Notes. Each panel compares a pair of items in the units that the pair shares. Panel A is a linear-probability regression: the outcome equals 100 if the respondent names 0–5 or 6–10 years as the most likely crisis horizon, and the stated ten-year probability is in 10-percentage-point units. Panels B and C regress the reported horizon on a horizon implied by the respondent's own debt ceiling, both coded 1 to 5 on the survey's categories with "more than 30 years" and "Never" pooled. Panel B maps the ceiling level onto the CBO projection, using the debt-trajectory group alone because only that group was anchored to the official figure of 100 percent of GDP and shown the projection. Panel C converts headroom, defined against whichever reference the respondent saw, into the years the projection needs to accumulate it; the two constructions coincide in the trajectory group. Exact agreement would imply a unit slope, but that is a benchmark rather than a null worth testing, since measurement error pushes the slope below one on its own; stars mark departures from zero. HC1 standard errors in parentheses; the coefficients are descriptive associations. Table A23 reports the matching agreement statistics. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

in the 61-respondent cell. Applying the headroom mapping to all randomized groups gives corresponding coefficients of 0.24, 0.23, and 0.19, all positive at the one percent level. This broader mapping assumes the CBO accumulation rate for respondents who did not see the projection and uses their perceived current debt/GDP. Kendall's τ_b is 0.29, 0.26, and 0.17. The slope estimates do not differ significantly across samples or between the debt-trajectory group and the remaining randomized groups ($p = 0.81$ and $p = 0.80$).

Table A23. Respondent-level consistency across the sustainability items

	Bond investors	Voters	Econ./Finance
Panel A. Bounds implied by the modal-horizon item			
Respondents	985	1000	247
Violating the implied bound (%)	8.3	9.8	9.3
Modal horizon within ten years (%)	7.0	6.9	9.2
Modal horizon beyond ten years (%)	9.8	15.3	9.4
Panel B. Ceiling vs. reported horizon, within-ten-years cut			
Respondents (trajectory group)	246	249	61
Difference in sample shares (pp)	4.9	-0.8	1.6
Respondents whose two answers disagree (%)	35.8	32.9	37.7
Ceiling implies within ten only (%)	20.3	16.1	19.7
Reported within ten only (%)	15.4	16.9	18.0
Observed agreement (%)	64.2	67.1	62.3
Chance agreement (%)	51.0	55.6	50.0
Controlled LPM coefficient (pp)	27.1***	26.0***	17.9
Panel C. Ceiling mapped into the survey's own horizon categories			
Exact category match (%)	29.3	32.5	14.8
Match within one category (%)	69.1	70.7	59.0
Chance-expected exact match (%)	19.8	23.9	16.2
Kendall's τ_b	0.29	0.26	0.17
Spearman's ρ	0.34	0.29	0.21

Notes. In Panel A, a reported most-likely horizon within ten years puts a floor of 16.7 percent on the stated ten-year probability, while a later horizon puts a cap of 66.7 percent on it. The indented rows are violation rates within each horizon category, not shares of the sample. Pairwise comparisons of the overall violation rate use Fisher's exact test, with the three sample comparisons adjusted jointly by the Holm method; the smallest adjusted p-value is 0.82. Panels B and C use debt-trajectory-group respondents with valid ceiling and horizon responses. Panel B codes the ceiling as implying a crisis within ten years if it is at or below the CBO's projected 2036 ratio. Chance agreement is what independence with the observed marginals would give, so the gap between the observed and chance rows is the agreement attributable to the respondent rather than to the marginals. The controlled LPM includes age, sex, collection-wave controls where applicable, and education where collected. Panel C compares the full categories, taking the implied category from the first projection year that reaches the stated ceiling and pooling "more than 30 years" with "Never." The Economics/Finance sample has 61 trajectory-group respondents, so its estimates are imprecise.

The modal-horizon bounds are violated by 8.3% of bond investors, 9.8% of voters, and 9.3% of Economics/Finance respondents; no pairwise difference is statistically significant. At the ten-year cutoff, observed agreement between the ceiling-implied and reported classifications is 64.2%,

67.1%, and 62.3%, compared with chance benchmarks of 51.0%, 55.6%, and 50.0%. The controlled indicator regressions give coefficients of 27.1, 26.0, and 17.9 percentage points. Disagreement runs in both directions: the ceiling alone implies a crisis within ten years for 20.3%, 16.1%, and 19.7%, while the reported horizon alone places it within ten years for 15.4%, 16.9%, and 18.0%.

Table A24. Value shares and sustainability beliefs

Value share / control	Crisis prob.	Headroom	log-timing
Panel A. Bond investors ($N = 977$)			
Safe-asset demand	-0.184*** (0.056)	0.013 (0.126)	0.004* (0.002)
Debt rollover	-0.034 (0.058)	-0.031 (0.128)	0.003 (0.003)
Govt. assets	0.067 (0.061)	-0.416*** (0.127)	-0.000 (0.003)
Fin. repression	0.106 (0.072)	0.000 (0.191)	-0.001 (0.003)
Primary surpluses	-0.084 (0.057)	-0.088 (0.154)	0.002 (0.002)
Perceived debt/GDP	-0.039** (0.018)	0.408*** (0.103)	0.001 (0.001)
Adj. R^2	0.065	0.107	0.035
N	977	977	977
Panel B. Voters ($N = 994$)			
Safe-asset demand	-0.188*** (0.053)	0.235*** (0.086)	0.005** (0.002)
Debt rollover	-0.089 (0.057)	0.038 (0.092)	0.003 (0.002)
Govt. assets	-0.054 (0.056)	-0.055 (0.102)	0.002 (0.002)
Fin. repression	-0.017 (0.081)	-0.106 (0.103)	-0.005* (0.003)
Primary surpluses	-0.147** (0.057)	0.105 (0.089)	0.001 (0.002)
Perceived debt/GDP	0.042* (0.024)	-0.063 (0.063)	0.000 (0.001)
Adj. R^2	0.032	0.040	0.030
N	994	994	993

Notes. OLS of each sustainability outcome on the benchmark reconstructed shares, perceived current debt/GDP, and demographic controls (age, sex), with HC1 heteroskedasticity-robust standard errors in parentheses. Because the six shares sum to 100 within respondent, the regression includes five of the six sources and omits *Fed purchases* as the base; each reported coefficient is the association from moving one benchmark point from the Fed-purchases base into that source. Outcomes: *Crisis prob.* is the stated probability of a debt crisis within ten years (0–100); *Headroom* is the maximum sustainable debt/GDP less the anchor (100 for the trajectory group, otherwise the respondent's perceived level), in percentage points; *log-timing* is the log of representative years assigned to the six timing categories (2.5, 8, 15.5, 25, 50, and 1000). Because the headroom column uses the group-specific reference, it should be read together with the anchor audit in Table A20. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A25. Concern, behavior, and the action gap

	Bond investors	Voters
Panel A. Does concern (vs. knowledge) predict the behavioral action?		
Action = plan reduce / decisive vote (% yes)	30.6	8.8
Action $\sim$ concern: LPM slope	0.108***	0.089***
R^2	0.034	0.056
Action $\sim$ knowledge (FKI): LPM slope	0.0000	0.0001
R^2	0.000	0.000
Panel B. The action gap (% taking the action, by concern level)		
A little concerned (2)	21.8	2.4
Moderately concerned (3)	26.5	3.3
Very concerned (4)	41.5	18.6
Worried (concern ≥ 3, % of sample)	80.0	75.5
of whom *inert* at top concern (%)	58.5	81.4
Panel C. Concern and the two elicited belief dimensions		
Concern $\sim$ crisis prob. + headroom: R^2	0.058	0.035
Concern $\sim \theta$ shares only: R^2	0.007	0.004
incremental R^2 from adding θ	0.007	0.004
joint test θ block $= 0$ (p)	0.209	0.534

Notes. Concern is the four-point worry item, where 4 means very concerned. Panels A and B use respondents with an observed action: all 985 investors and the 964 voters who report some concern and therefore receive the voting item. For investors the behavioral action is planning to reduce Treasury holdings; for voters it is debt sustainability being a decisive vote factor. Panel A reports linear-probability slopes with HC1 robust standard errors from separate regressions on concern and on the fiscal-knowledge index FKI $= -|\text{perceived} - 100|$. Panel B reports action rates by concern. Panel C reports cross-sectional R^2 values from regressions relating concern to crisis probability and headroom, to the value shares θ, and to both sets of beliefs. Adding θ to the sustainability-belief regression raises R^2 negligibly, and the block is jointly insignificant; this comparison is an association, not a mediation or causal ordering. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A26. Intended action by dominant valuation source

	Bond: plan to reduce		Voters: decisive vote	
Dominant source	%	*N*	%	*N*
Safe-asset demand	28.0	361	8.2	304
Debt rollover	33.5	188	8.7	206
Fed purchases	31.6	114	11.9	118
Govt. assets	34.4	131	9.8	163
Fin. repression	36.4	44	5.7	35
Primary surpluses	27.2	147	7.2	138
All respondents	30.6	985	8.8	964
χ^2 test (action vs. type), p	.504		.770	

Notes. Intended action by dominant valuation source. For bond investors the action is planning to *reduce* Treasury holdings over the next twelve months; for voters it is debt sustainability being a *decisive* vote factor. Columns give the percentage taking the action and the cell *N*. The χ^2 test is of action against dominant source. Dominant-source differences in intended action are weak, consistent with Table A25, where concern is more strongly associated with action than the value-share vector is. Repression-dominant cells are small (bond 44, voters 35) and should be read descriptively.

E. Model Derivations

This appendix derives investor demand, the normal-time price schedule, the consolidated public budget, the six-category accounting allocation, and stress refinancing capacity. It characterizes what normal-time prices identify and proves Proposition 1.

E.1. Demand and Pricing

Each voluntary block $j \in \mathscr{J}$ has the objective

$$\mathbb{E}_0^j \sum_{t=0}^{\infty} \beta^t \left[c_t^j + v^j \Big((1-\delta_t^j) b_t^j \Big) \right] \tag{4}$$

and trades bonds and the tree subject to the flow budget

$$c_t^j + P_t^B b_{t+1}^j + P_t^Y \alpha_{t+1}^j = y_t^j + (1-\delta_t) b_t^j + (P_t^Y + Y) \alpha_t^j. \tag{A1}$$

Here $y_t^j > 0$ is the block's endowment and α_t^j its tree share, which matters only through market clearing for the tree. The resource flow depends on the common realized-default indicator δ_t; the objective depends on block j's perceived state δ_t^j. The unrestricted linear numeraire pins the risk-free price at β. It also gives the tree first-order condition $P_t^Y = \beta(P_{t+1}^Y + Y)$, whose stationary solution is $P^Y = \beta Y/(1-\beta)$. The private sector is therefore willing to absorb the residual tree supply at this price, so tree market clearing is $\sum_j \alpha_t^j = 1 - \alpha_t$, where the unsuperscripted α_t is the consolidated public holding of the main text.

The budget fixes the timing. The block buys b_{t+1}^j at date t and is repaid $(1-\delta_{t+1}) b_{t+1}^j$ at $t+1$, and it already holds its view of that payment at t, so $\mathbb{E}_t^j[1-\delta_{t+1}] = 1 - \delta_{t+1}^j$. This is the main text's $\mathbb{E}_{t-1}^j[1-\delta_t] = 1-\delta_t^j$ moved forward one period. Perceived survival therefore determines both expected repayment and the bond-service payoff in the polar experiment. Safe investors set $\delta_{t+1}^{\text{Safe}} = 1$ when the claim is no longer regarded as safe; vigilantes set $\delta_{t+1}^{\text{Vig}} = 1$ when they expect continued placement to fail. Sticky rollover obeys $\delta_{t+1}^{\text{Sticky}} = \delta_t$, so it remains active throughout every pre-default scenario. The blocks differ in their perceived states but face the same realized sovereign default.

Let λ_t^j be the multiplier on block j's date-t budget. The consumption first-order condition gives $\lambda_t^j = \beta^t$. Taking the block's date-t subjective expectation of next period's repayment and bond services in the bond Kuhn–Tucker condition, then dividing by β^t, gives

$$P_t^B \;\geq\; \beta(1-\delta_{t+1}^j) \left[1 + v^j \Big((1-\delta_{t+1}^j) b_{t+1}^j \Big)' \right]. \tag{A2}$$

With $v^j(b)' = \chi(1 - b/b_0^j)$, a block with $\delta_{t+1}^j = 0$ follows the common intensive-margin schedule, whereas a block with $\delta_{t+1}^j = 1$ chooses zero at any positive price. Thus block demand is

$$b_{t+1}^j = (1 - \delta_{t+1}^j) b_0^j \max\left\{1 + \chi^{-1}\left(1 - \frac{P_t^B}{\beta}\right), 0\right\}, \tag{A3}$$

so δ_{t+1}^j directly determines the block's participation. The three voluntary blocks share the same intensive-margin demand curve; their different responses enter through δ_{t+1}^j, while b_0^j measures their depths. The regulated block instead holds the fixed aggregate quantity $\bar{b}^{\text{Rep}}$ imposed by policy; it has no first-order condition or perceived failure state.

The intensive margin is the standard reduced form used to estimate Treasury demand from price responses (Krishnamurthy and Vissing-Jorgensen 2012; Nagel 2016; Vayanos and Vila 2021; Krishnamurthy and Li 2023). The participation margin supplies the distinction needed here. Safe-asset demand can disappear when information-insensitive debt loses safe status (Gorton and Pennacchi 1990; Gorton 2017; Brunnermeier et al. 2024); vigilante rollover captures withdrawal when continued placement is doubted (Calvo 1988; Cole and Kehoe 2000); and sticky rollover corresponds to investors who continue funding despite the same signal (Caballero and Simsek 2020). Normal-time prices discipline the common intensive margin. The stress experiment asks which blocks remain on the participation margin.

Equilibrium price. Substituting (A3) into the market-clearing condition (8) gives

$$\tilde{B}_{t+1} = \bar{b}^{\text{Rep}} + \left[\sum_{j \in \mathscr{J}} (1 - \delta_{t+1}^j) b_0^j\right] \max\left\{1 + \chi^{-1}\left(1 - \frac{P_t^B}{\beta}\right), 0\right\}, \tag{A4}$$

where the consolidated budget derived below splits gross debt B_{t+1} into the central-bank part B_{t+1}^{CB} and the privately held stock $\tilde{B}_{t+1}$. The bracketed sum is voluntary depth $\hat{b}_0^V(\delta_{t+1})$ from (7). Because $\tilde{B}_{t+1} > \bar{b}^{\text{Rep}}$, the positive-voluntary-demand branch applies:

$$\hat{b}_0^V(\delta_{t+1})\left(1 + \chi^{-1}\left(1 - \frac{P_t^B}{\beta}\right)\right) = \tilde{B}_{t+1} - \bar{b}^{\text{Rep}}. \tag{A5}$$

Solving for P^B gives equation (9). Setting $\delta_{t+1} = (0,0,0)$ gives $\hat{b}_0^V = b_0^V$ and equation (10). In joint stress, $\delta_{t+1}^{\text{Safe}} = \delta_{t+1}^{\text{Vig}} = 1$ and $\delta_{t+1}^{\text{Sticky}} = 0$, so only sticky voluntary depth and compelled holdings remain.

Demand and market clearing thus deliver the first identification result: the normal price schedule reveals total voluntary depth but not the depths attached to different stress responses. To connect that demand result to government debt value, the next step derives the consolidated public

budget.

E.2. Consolidated Budget and Normal-Time Valuation

Write T_t^{CB} for the central bank's remittance to the Treasury. The fiscal authority's flow budget, written with $B_{t+1} = \tilde{B}_{t+1} + B_{t+1}^{CB}$, is

$$(1-\delta_t)B_t = P_t^B \tilde{B}_{t+1} + P_t^B B_{t+1}^{CB} + \bar{T} + T_t^{CB} + \alpha_t^f Y + P_t^Y(\alpha_t^f - \alpha_{t+1}^f). \tag{A6}$$

The central bank's budget is

$$P_t^Y \alpha_{t+1}^{\text{Fed}} + P_t^B B_{t+1}^{CB} + T_t^{CB} = P_t^Y \alpha_t^{\text{Fed}} + \alpha_t^{\text{Fed}} Y + (1-\delta_t)B_t^{CB}. \tag{A7}$$

After solving the central-bank constraint for T_t^{CB},

$$T_t^{CB} = P_t^Y(\alpha_t^{\text{Fed}} - \alpha_{t+1}^{\text{Fed}}) + \alpha_t^{\text{Fed}} Y + (1-\delta_t)B_t^{CB} - P_t^B B_{t+1}^{CB}, \tag{A8}$$

substituting into the fiscal constraint, and using $B_t = \tilde{B}_t + B_t^{CB}$, the central-bank bond terms $(1-\delta_t)B_t^{CB}$ and $P_t^B B_{t+1}^{CB}$ appear on both sides and cancel, leaving

$$\begin{aligned}(1-\delta_t)\tilde{B}_t = {} & P_t^B \tilde{B}_{t+1} + \bar{T} \\ & + P_t^Y(\alpha_t^{\text{Fed}} - \alpha_{t+1}^{\text{Fed}}) + \alpha_t^{\text{Fed}} Y + \alpha_t^f Y + P_t^Y(\alpha_t^f - \alpha_{t+1}^f).\end{aligned} \tag{A9}$$

Collecting the public-resource terms with $\alpha_t \equiv \alpha_t^{\text{Fed}} + \alpha_t^f$,

$$(1-\delta_t)\tilde{B}_t = P_t^B \tilde{B}_{t+1} + \bar{T} + \alpha_t Y - P_t^Y(\alpha_{t+1} - \alpha_t). \tag{A10}$$

In a period without default, this is equation (3). The cancellation is why central-bank Treasury holdings do not themselves appear as backing; only central-bank capital, represented by α_t^{Fed}, survives consolidation. Following Li and Merkel (2026), this is a broad economic capacity that can include foreign reserves, residual seigniorage capacity, the fiscal value of the banking franchise, and institutional credibility. Their banking microfoundation maps QE into this object: reserves use scarce bank balance-sheet capacity, displace bank-dependent lending, and reduce taxable bank profits.

Present-value decomposition. Along a normal-time path, write the bond price as $P^B = \beta + \mathrm{CY}_{t+1}$ with

$$\mathrm{CY}_{t+1} = \beta\chi\left(1 - \frac{\tilde{B}_{t+1} - \bar{b}^{\mathrm{Rep}}}{b_0^V}\right) \tag{A11}$$

from (10). Substituting into (3) and rearranging,

$$\tilde{B}_t = \bar{T} + \alpha_t Y - P^Y \Delta\alpha_{t+1} + \beta\tilde{B}_{t+1} + \mathrm{CY}_{t+1}\tilde{B}_{t+1}, \tag{A12}$$

where $P^Y = \beta Y/(1-\beta)$ is constant in equilibrium. For an integer horizon $H \geq 0$, apply (A12) recursively to eliminate $\tilde{B}_{t+1}$, then $\tilde{B}_{t+2}$, and so on up to $\tilde{B}_{t+H}$:

$$\tilde{B}_t = \sum_{s=0}^{H} \beta^s \left[\bar{T} + \alpha_{t+s} Y - P^Y \Delta\alpha_{t+s+1} + \mathrm{CY}_{t+s+1}\tilde{B}_{t+s+1}\right] + \beta^{H+1}\tilde{B}_{t+H+1}. \tag{A13}$$

Impose $\lim_{H\to\infty} \beta^{H+1}\tilde{B}_{t+H+1} = 0$. Bounded debt is sufficient because $\beta \in (0,1)$. The component public-resource holdings are nonnegative and bounded above, since $\alpha_t \leq \bar{\alpha}$, so the component present-value series converge. Taking the limit and grouping by economic channel, using $\alpha_t = \alpha_t^{\mathrm{Fed}} + \alpha_t^f$, gives

$$\begin{aligned}\tilde{B}_t = {} & \underbrace{\sum_{s=0}^{\infty} \beta^s \bar{T}}_{V_t^S} + \underbrace{\sum_{s=0}^{\infty} \beta^s \alpha_{t+s}^f Y - P^Y \sum_{s=0}^{\infty} \beta^s \Delta\alpha_{t+s+1}^f}_{V_t^{\mathrm{Asset}}} + \underbrace{\sum_{s=0}^{\infty} \beta^s \alpha_{t+s}^{\mathrm{Fed}} Y - P^Y \sum_{s=0}^{\infty} \beta^s \Delta\alpha_{t+s+1}^{\mathrm{Fed}}}_{V_t^{\mathrm{Fed}}} \\ & + \underbrace{\sum_{s=0}^{\infty} \beta^s \mathrm{CY}_{t+s+1}\tilde{B}_{t+s+1}}_{V_t^{\mathrm{Dem}}}.\end{aligned} \tag{A14}$$

For each voluntary block, define its holder allocation of the common demand value by

$$\theta_j \equiv \frac{1}{\tilde{B}_t} \sum_{s=0}^{\infty} \beta^s \mathrm{CY}_{t+s+1} b_{t+s+1}^j, \qquad j \in \{\mathrm{Safe}, \mathrm{Sticky}, \mathrm{Vig}\}, \tag{A15}$$

and define

$$\theta_{\mathrm{Rep}} \equiv \frac{1}{\tilde{B}_t} \sum_{s=0}^{\infty} \beta^s \mathrm{CY}_{t+s+1} \bar{b}^{\mathrm{Rep}}. \tag{A16}$$

Market clearing implies that these terms sum to $V_t^{\mathrm{Dem}}/\tilde{B}_t$. Dividing (A14) by debt then gives the six-category allocation in (20), with $\theta_{\mathrm{Roll}} \equiv \theta_{\mathrm{Sticky}} + \theta_{\mathrm{Vig}}$.

The valuation identity explains the current market value of debt. Locating the economy relative to refinancing stress also requires the normal-time debt map, which fixes the steady states shared by observationally equivalent demand and public-resource compositions.

E.3. Normal-Time Values and Debt Dynamics

At the stable steady state $\tilde{B}_{t+s} = \tilde{B}^*$ and $\alpha_{t+s} = \alpha = \alpha^{\text{Fed}} + \alpha^f$, so

$$V^S = \frac{\bar{T}}{1-\beta}, \tag{A17}$$

$$V^{\text{Asset}} = \frac{\alpha^f Y}{1-\beta} = \alpha^f (Y + P^Y), \tag{A18}$$

$$V^{\text{Fed}} = \frac{\alpha^{\text{Fed}} Y}{1-\beta} = \alpha^{\text{Fed}} (Y + P^Y). \tag{A19}$$

Summing the three public components gives steady-state public backing:

$$V^S + V^{\text{Asset}} + V^{\text{Fed}} = \frac{\bar{T} + \alpha Y}{1-\beta}, \tag{A20}$$

which depends on the public-resource composition only through the flow $\bar{T} + \alpha Y$.

V^{Dem}: At the steady state $\text{CY}^* = \beta\chi[1 - (\tilde{B}^* - \bar{b}^{\text{Rep}})/b_0^V]$ is constant, so

$$V^{\text{Dem}} = \sum_{s=0}^{\infty} \beta^s \text{CY}^* \tilde{B}^* = \frac{\text{CY}^*}{1-\beta} \tilde{B}^* = \frac{\beta\chi}{1-\beta} \left(1 - \frac{\tilde{B}^* - \bar{b}^{\text{Rep}}}{b_0^V}\right) \tilde{B}^*, \tag{A21}$$

which depends on safety, sticky-rollover, and vigilante-rollover depths only through their sum. This is the demand-identification premise of Proposition 1: every voluntary composition consistent with b_0^V delivers the same value for a fixed compelled quantity.

To close the normal-time dynamics, suppose the fiscal authority selects the smaller face-value issuance whenever the budget admits two quantities that raise the required resources. Equivalently, it remains on the high-price, rising branch of the issuance-revenue curve. This is an issuance rule, not an implication of investor demand.

Lemma 1 (Debt dynamics)**.** *Let* $\kappa \equiv \beta(1+\chi) - 1 + \beta\chi\bar{b}^{Rep}/b_0^V$ *and assume*

$$0 < \kappa \leq 1, \qquad -\frac{\kappa^2 b_0^V}{4\beta\chi} < \bar{s} < 0, \qquad \bar{b}^{Rep} < \frac{b_0^V}{2\beta\chi} \left(\kappa - \sqrt{\kappa^2 + \frac{4\beta\chi\bar{s}}{b_0^V}}\right). \tag{A22}$$

Then funding supply and funding need cross exactly twice on the branch with positive voluntary demand, at $\bar{b}^{Rep} < B_1^* < B_2^*$*. The smaller steady state is stable and the larger is unstable. From any feasible initial level below* B_2^**, debt converges to* B_1^**. From any level above* B_2^**, debt rises while refinancing remains feasible.*

Proof. Hold the asset stock fixed, so $\Delta\alpha_{t+1} = 0$, and write $\bar{s} \equiv \bar{T} + \alpha Y$. Substituting (10) into the budget (3) with $\tilde{B}_t = \tilde{B}_{t+1} = B^*$ gives

$$\frac{\beta\chi}{b_0^V}(B^*)^2 - \kappa B^* - \bar{s} = 0, \qquad \kappa = \beta(1+\chi) - 1 + \frac{\beta\chi\bar{b}^{\text{Rep}}}{b_0^V}, \tag{A23}$$

With $\Delta \equiv \kappa^2 + 4\beta\chi\bar{s}/b_0^V$, the roots are

$$B_{1,2}^* = \frac{b_0^V}{2\beta\chi}\left(\kappa \mp \sqrt{\Delta}\right). \tag{A24}$$

The assumed bounds on $\bar{s}$ give $0 < \Delta < \kappa^2$, so both roots are real, distinct, and positive; the final condition in (A22) places both on the positive-voluntary-demand branch. Issuance revenue is $P^B(B)B = (1+\kappa)B - \beta\chi B^2/b_0^V$, with slope $1 \pm \sqrt{\Delta}$ at the two roots. Under the maintained issuance rule, the debt map is the smaller solution of $P^B(B_{t+1})B_{t+1} = B_t - \bar{s}$, which lies on the rising branch of issuance revenue, so the map is increasing with slope the reciprocal of the revenue slope and meets the 45° line exactly at the roots. At B_1^* the slope is $1/(1+\sqrt{\Delta}) < 1$; the map sits above the line below B_1^* and below it between the roots, so debt converges to B_1^* monotonically from every feasible level below B_2^*. Because $\kappa \leq 1$ and $\bar{s} < 0$ give $\Delta < 1$, the larger root also lies on the rising branch, the map's slope there is $1/(1-\sqrt{\Delta}) > 1$, and the map sits above the line beyond B_2^*, so debt rises while refinancing remains feasible from every level above it. This is Lemma 3 of Li and Merkel (2026) with b_0^V in place of b_0, $\bar{s}$ in place of $\bar{T}$, and the compelled-holdings term added to κ.

Compelled holdings raise κ and therefore B_2^*, so a captive block widens the safe region. Public-resource income and primary surpluses enter only through $\bar{s}$ and therefore have the same root-specific effects. Increasing either contribution to $\bar{s}$ lowers B_1^* and raises B_2^*, since $\partial B_1^*/\partial\bar{s} = -1/\sqrt{\Delta}$ and $\partial B_2^*/\partial\bar{s} = 1/\sqrt{\Delta}$. □

The smaller root governs normal convergence; the larger root marks the point beyond which debt rises even while auctions continue to clear. Under the maintained issuance rule, we therefore call $[\bar{b}^{\text{Rep}}, B_2^*)$ the safe region. Let $B^{\text{flow}} \equiv \bar{s} + \Phi(b_0^V; \bar{b}^{\text{Rep}})$ denote the largest current debt that normal issuance can refinance, since $\Phi(b_0^V; \bar{b}^{\text{Rep}})$ is the maximum of issuance revenue $P^B(B)B$ over B. Debt rises on the conditional no-return region $[B_2^*, B^{\text{flow}}]$. One-off public-resource mobilization can extend current funding capacity beyond B^{flow}, but the model does not specify recursive dynamics afterward.

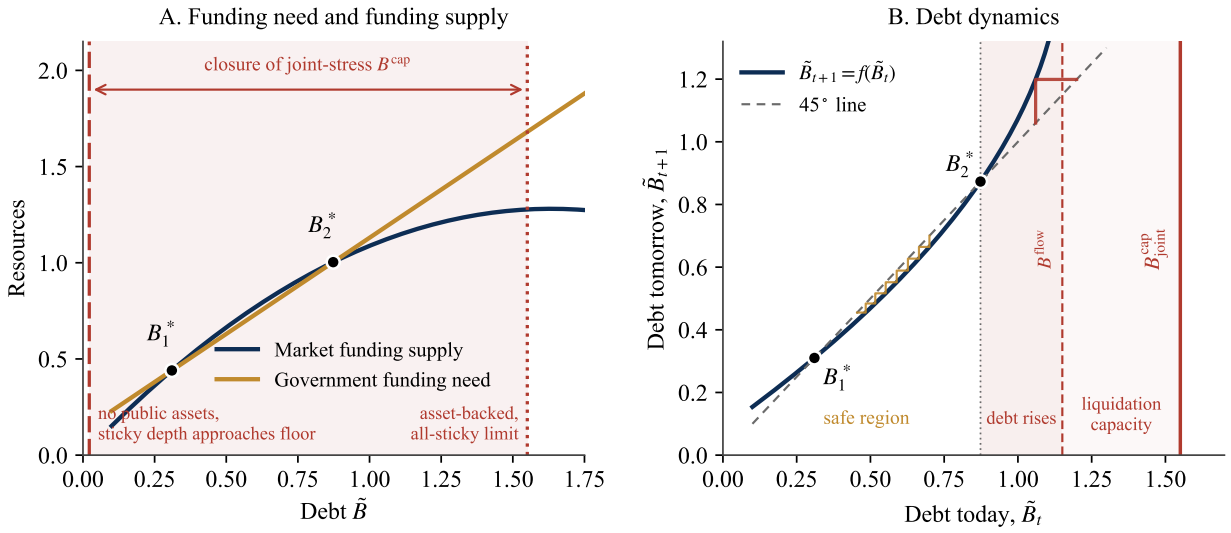


Figure A3. **Steady states, no-return region, and unidentified joint-stress funding limits.** Panel A plots normal-time funding need and supply, whose crossings are B_1^* and B_2^*. Its shaded span is the closure of the one-period joint-stress funding limits generated by hidden demand and public-resource compositions consistent with the same normal-time prices and quantities. Panel B shows the debt map under the maintained smaller-issuance rule. Debt rises between B_2^* and the normal refinancing limit B^{flow}. The interval from B^{flow} to the most optimistic joint-stress limit is one-period liquidation capacity, not a recursive debt path. The economy repays normally at B_1^*. A joint-stress limit below B_1^* means that, under that hidden composition, imposing the joint trigger at the current debt would cause immediate refinancing failure; it does not imply default in the observed normal state. The shaded span varies the hidden demand split holding b_0^V fixed.

Figure A3 separates the objects fixed by normal-time observations from those required for stress capacity. The full price schedule reveals total voluntary depth b_0^V; observed debt and its price imply $\bar{s} = \tilde{B}^*[1 - P^B(\tilde{B}^*)]$; and holdings data reveal $\bar{b}^{\text{Rep}}$. These objects determine the funding-supply curve, funding-need line, and both steady states. The joint-failure object $B^{\text{cap}}_{\text{Joint}}$ in (16) also depends on sticky depth and net public resources.

Because each demand block is strictly positive and the interior restriction on sticky depth is strict, the shaded closed interval includes two limiting compositions that are not attained. Its interior is the set of conditional one-period joint-stress funding limits; the endpoints are its infimum and supremum. At the illustrative parameters, demand composition contributes 1.13 of the closure's total width of 1.53, compared with 0.40 from the public-resource split. Normal-state refinancing remains feasible at B_1^* for every composition in this exercise. Values $B^{\text{cap}}_{\text{Joint}} < B_1^*$ instead identify hidden compositions under which the joint trigger would cause immediate refinancing failure. Conditioning on survival of joint stress would add the restriction $B^{\text{cap}}_{\text{Joint}} \geq B_1^*$.

The illustrative parameters satisfy the bounds in (A22), so both steady states lie on the positive-voluntary-demand branch characterized by Lemma 1.

Normal dynamics therefore pin down the path observed before stress without pinning down the resources available once participation changes. We now derive that state-contingent object.

E.4. State-Contingent Refinancing Capacity

The valuation identity capitalizes resources that support indefinite rollover. Stress refinancing capacity asks how much maturing debt can be covered in one period after participation changes. The calculation conditions on perceived failure states and does not solve a recursive default equilibrium. Write $r = \bar{b}^{\text{Rep}}$ and let $b > 0$ denote surviving voluntary depth. The experiment is restricted to the positive-price portion of the positive-voluntary-demand branch: net privately placed issuance is $\tilde{B}_{t+1} = q + r$ with $q > 0$ absorbed voluntarily. From (9),

$$P^B(q;b) = \beta\left(1 + \chi - \frac{\chi q}{b}\right), \tag{A25}$$

so total issuance revenue is

$$\mathscr{R}(q;b,r) = \beta(q+r)\left(1 + \chi - \frac{\chi q}{b}\right). \tag{A26}$$

This is a strictly concave quadratic. Its stationary point is

$$q^*(b,r) = \frac{b(1+\chi) - \chi r}{2\chi}. \tag{A27}$$

For $b > \chi r/(1+\chi)$, the stationary point lies in the maintained domain. The bound on compelled holdings assumed in Section 4.1,

$$\bar{b}^{\text{Rep}} < \frac{1+\chi}{\chi}\, b_0^{\text{Sticky}}, \tag{A28}$$

imposes this inequality at the smallest possible surviving depth, $b = b_0^{\text{Sticky}}$, and therefore at every larger depth. The unique maximizer lies on the maintained branch. Substitution gives

$$\Phi(b;r) \equiv \max_{\substack{q>0 \\ P^B(q;b)>0}} \mathscr{R}(q;b,r) = \frac{\beta}{4\chi b}\left[(1+\chi)b + \chi r\right]^2, \tag{A29}$$

which is (11). Differentiating,

$$\frac{\partial \Phi(b;r)}{\partial b} = \frac{\beta}{4\chi}\left[(1+\chi)^2 - \frac{\chi^2 r^2}{b^2}\right] > 0 \tag{A30}$$

on the maintained domain, so Φ is strictly increasing in surviving depth. The price at the optimizer is

$$P^B(q^*;b) = \frac{\beta}{2}\left(1+\chi+\frac{\chi r}{b}\right) > 0. \tag{A31}$$

At a date with maturing net debt $\tilde{B}_t$, allow the stress-mobilization value $P^{Y,S} = \rho P^Y$ to differ from its normal-time counterpart. The consolidated budget becomes

$$\tilde{B}_t = P_t^B \tilde{B}_{t+1} + \bar{T} + \alpha_t Y - P^{Y,S}(\alpha_{t+1} - \alpha_t). \tag{A32}$$

The current primary balance and public-resource income together supply $\bar{T} + \alpha Y$. Future primary surpluses cannot be pledged outside the stressed bond market: they remain promises of a government in which trust has been impaired. Fiscal assets can instead be sold, and central-bank capital can be deployed without issuing a new fiscal promise. Let $\rho \in (0,1]$ be the fraction of normal continuation value mobilized under stress. For fiscal assets it captures saleability and any fire-sale discount; for central-bank capital it is a reduced-form measure of immediate deployability. Mobilization adds $\rho\alpha P^Y$, so public resources are $D(\rho) = \bar{T} + \alpha Y + \rho\alpha P^Y$ as in (12). Conditional on a stress pattern with surviving voluntary depth b, maximum debt service within the maintained branch is $D(\rho) + \Phi(b;r)$.

Safety failure alone leaves surviving depth $b_0^{\text{Sticky}} + b_0^{\text{Vig}}$, while vigilante withdrawal alone leaves $b_0^{\text{Sticky}} + b_0^{\text{Safe}}$. Substituting each into $D(\rho) + \Phi(b;r)$ gives (14) and (15). Since Φ is strictly increasing, the isolated capacities are ranked by the depths that survive:

$$\operatorname{sign}\left(B_{\text{Vig}}^{\text{cap}} - B_{\text{Safe}}^{\text{cap}}\right) = \operatorname{sign}\left(b_0^{\text{Safe}} - b_0^{\text{Vig}}\right). \tag{A33}$$

Taking the minimum over these two isolated states yields

$$\min\left\{B_{\text{Safe}}^{\text{cap}}, B_{\text{Vig}}^{\text{cap}}\right\} = \Phi\left(b_0^{\text{Sticky}} + \min\{b_0^{\text{Safe}}, b_0^{\text{Vig}}\}; r\right) + D(\rho), \tag{A34}$$

which is conditional on at most one fragile margin failing.

For every combination of safety and vigilante failure,

$$\hat{b}_0^V(\delta) = b_0^{\text{Sticky}} + (1-\delta^{\text{Safe}})b_0^{\text{Safe}} + (1-\delta^{\text{Vig}})b_0^{\text{Vig}} \geq b_0^{\text{Sticky}}. \tag{A35}$$

Joint failure attains the lower bound. Strict monotonicity of Φ therefore gives

$$\min_{\delta^{\text{Safe}},\, \delta^{\text{Vig}} \in \{0,1\}} \left\{\Phi\left(\hat{b}_0^V(\delta); r\right) + D(\rho)\right\} = \Phi(b_0^{\text{Sticky}}; r) + D(\rho), \tag{A36}$$

which is (16).

These calculations condition on sticky rollover continuing to participate and compelled holdings remaining fixed through the period. Each isolated capacity keeps the other fragile block active; joint stress allows both fragile blocks to withdraw. The calculation does not select a stress state in a recursive equilibrium. A self-fulfilling default region would require a default policy and a beliefs-to-participation fixed point, which are outside this benchmark.

The capacity formula separates market funding, Φ, from deployable public resources, $D(\rho)$. Demand composition governs the former. The public-resource split governs the latter, and the valuation identity provides a direct way to express that split in the same units as the survey allocations.

E.5. Public-Resource Composition and the Survey Shares

The mapping is conditional on one-period debt, positive stress-mobilization value for fiscal assets and central-bank capital, and the restriction that claims on future taxes can be issued only through the capped bond market. Under those assumptions, public-resource composition affects stress capacity even when normal-time public backing is unchanged.

Using $P^Y = \beta Y/(1-\beta)$, write

$$D(\rho) = \bar{T} + \alpha Y + \rho\,\alpha P^Y = \bar{T} + \left(1 + \rho\,\frac{\beta}{1-\beta}\right)\alpha Y. \tag{A37}$$

The marginal values are

$$\frac{\partial V^S}{\partial \bar{T}} = \frac{1}{1-\beta} > \frac{\partial D(\rho)}{\partial \bar{T}} = 1, \tag{A38}$$

$$\frac{\partial (V^{\text{Asset}} + V^{\text{Fed}})}{\partial(\alpha Y)} = \frac{1}{1-\beta}, \qquad \frac{\partial D(\rho)}{\partial(\alpha Y)} = 1 + \rho\,\frac{\beta}{1-\beta}. \tag{A39}$$

Thus a unit of either flow contributes $1/(1-\beta)$ to normal-time value. In one-period stress, a unit of current primary surplus contributes one unit. A unit of public-resource income contributes that same current unit plus mobilizable continuation value. Holding the observed flow $\bar{s} = \bar{T} + \alpha Y$ fixed therefore gives

$$\left.\frac{\partial D(\rho)}{\partial(\alpha Y)}\right|_{\bar{s}} = \rho\,\frac{\beta}{1-\beta} > 0. \tag{A40}$$

Any positive mobilization value preserves the ranking; a lower ρ reduces the magnitude of this derivative linearly.

If the accounting components in (20) are supplied externally, a constant primary surplus and

constant public holding imply

$$\bar{T} = (1-\beta)\theta_S \tilde{B}, \qquad \alpha Y = (1-\beta)(\theta_{\text{Asset}} + \theta_{\text{Fed}})\tilde{B}. \tag{A41}$$

Substitution gives

$$D(\rho) = [(1-\beta)\theta_S + (1-\beta+\rho\beta)(\theta_{\text{Asset}} + \theta_{\text{Fed}})]\,\tilde{B}. \tag{A42}$$

At full mobilization ($\rho = 1$) and $\beta = 0.8$, a surplus-only composition deploys 20.0% of debt. Applying the formula to the elicited shares gives 35.7% for voters, 32.1% for investors, and 34.0% for the Economics/Finance sample, or 1.6–1.8 times the surplus-only benchmark. Excluding the Federal Reserve share gives 16.5–19.9%, or 0.8–1.0 times the benchmark.

The discount factor here is the illustrative value used in Figure A3, chosen so that the figure's public-resource multiplier $1/(1-\beta)$ stays on the scale of its debt axis, and not an estimated annual rate: linear consumption utility ties β to the risk-free rate through $R = \beta^{-1}$, so 0.8 would correspond to an implausibly high one. The comparison is sensitive to that choice in one direction only. A surplus share enters $D(\rho)$ at $1-\beta$ while an asset or central-bank share enters at one, so raising β shrinks the benchmark far faster than it shrinks the elicited compositions. At $\beta = 0.96$ the surplus-only benchmark falls to 4.0% and the elicited compositions deploy 7.4–8.3 times it. The illustrative value is therefore the conservative one for the comparison drawn here.

These calculations illustrate the mapping and do not estimate resources the United States could mobilize in a crisis. A lower mobilizable fraction reduces the public-resource contribution through ρ. Longer-maturity debt reduces the amount that must be refinanced immediately, while restored commitment or an enforceable way to sell claims on future taxes weakens the pledgeability distinction itself.

E.6. Proof of Proposition 1

Write $r = \bar{b}^{\text{Rep}}$.

Part (i): investor demand. Equation (10) contains the three voluntary depths only through b_0^V. Because the observer knows β and r and observes the schedule's right-hand limit at r,

$$\chi = \frac{\lim_{\tilde{B}\downarrow r} P^B(\tilde{B})}{\beta} - 1, \tag{A43}$$

$$b_0^V = -\frac{\beta\chi}{P^B(\tilde{B})'}. \tag{A44}$$

Every sufficiently small reallocation among positive depths that preserves their sum leaves the full schedule and (A21) unchanged.

By (13) and (A30), capacities are strictly increasing in surviving depth. Joint failure uniquely attains the minimum b_0^{Sticky}. Shift $0 < \varepsilon < b_0^j$ from either fragile block j to sticky rollover. Joint-failure depth rises by ε; between the two isolated states, the state in which j fails gains ε and the other is unchanged. Hence joint-stress capacity rises strictly and the smaller isolated capacity weakly.

At fixed sticky depth and fixed fragile total $b_0^{\text{Safe}} + b_0^{\text{Vig}}$,

$$\min\{b_0^{\text{Safe}}, b_0^{\text{Vig}}\} \leq \frac{b_0^{\text{Safe}} + b_0^{\text{Vig}}}{2}, \tag{A45}$$

with equality only at $b_0^{\text{Safe}} = b_0^{\text{Vig}}$. Strict monotonicity of Φ proves that the smaller isolated capacity is uniquely maximized there, while (16) proves the invariance of joint-stress capacity.

These reallocations leave the proposition's aggregate normal-time observables unchanged but change joint-stress capacity. Swapping unequal positive safety and vigilante depths also preserves their sum, reverses the sign in (A33), and therefore swaps the isolated-capacity ranking.

Part (ii): public resources. Setting $\Delta\alpha = 0$ in (3) gives

$$\bar{T} + \alpha Y = \tilde{B}^* \left[1 - P^B(\tilde{B}^*)\right]. \tag{A46}$$

Any feasible reallocation with $d\bar{T} = -d(\alpha Y)$ preserves this equation and the total public-backing value $(\bar{T} + \alpha Y)/(1-\beta)$. Forward iteration therefore adds no restriction on the split.

Let $\bar{s} \equiv \bar{T} + \alpha Y$. Since $P^Y = \beta Y/(1-\beta)$,

$$D(\rho) = \bar{s} + \rho \frac{\beta}{1-\beta} \alpha Y. \tag{A47}$$

Holding $\bar{s}$, the three depths, r, β, χ, Y, and $\rho > 0$ fixed, every capacity therefore rises at rate $\rho\beta/(1-\beta) > 0$ for every feasible public-resource reallocation. The reallocation preserves normal fiscal flow but changes deployable resources. Thus normal equilibrium does not identify stress refinancing capacity. □

F. Post-Stratification Robustness

This appendix tests the sensitivity of two population-level summaries to the demographic composition of the voter sample. Both are stable under reweighting.

The main text reports unweighted statistics. Two of its summaries, however, are population-level aggregates of the voter/general-population sample, and a demographic tilt could in principle move both: the benchmark reconstructed mean value shares (Figure 3 and Table A9) and the source-valence ranking (Figure 5 and Table A15). Because that sample, like other online panels, skews younger and more educated than the U.S. adult population (Table 1), we recompute both under post-stratification weights. Random assignment provides internal validity for the treatment comparisons without demographic weighting, so the exercise here addresses only whether the two descriptive aggregates move when the observed sample margins are aligned with the population.

We use iterative proportional fitting to rake the voter sample to the age, sex, and education margins in Table 1. The age and sex targets describe U.S. adults; the available education target covers adults aged 25 and older. Census region is unavailable in the Prolific export. Of the 1,001 voters, 994 have complete age, sex, and education fields. Respondents aged 65 and older are 5.5% of the sample against a 23% benchmark, so the weighted estimates rely heavily on those few respondents.

Table A27 compares the unweighted and reweighted benchmark mean shares. No share moves by more than 1.1 percentage points, and the ordering of the channels is unchanged. Global safe-asset demand remains the single largest benchmark source, primary surpluses remain a minor component near 15%, and the non-surplus channels continue to account for the bulk of benchmark perceived value. This weighting exercise is separate from the unrestricted residual-allocation bounds in Appendix Table A11.

Table A28 repeats the exercise for the valence ranking. The primary-minus-little contrast that orders the sources (the share naming a source as a primary source of support minus the share dismissing it) is a difference of two population proportions. Individual contrasts shift by at most 2.3 percentage points, and the ordering of the sources is preserved exactly. Global safe-asset demand remains, by a wide margin, the consensus source, while financial repression and primary surpluses remain the channels with the most negative net support. Neither the value decomposition nor the valence ranking depends on the voter sample's demographic tilt, which is why we report unweighted statistics throughout the main text.

Table A27. Mean value-share decomposition: sensitivity to post-stratification weights (voters)

Source	Unweighted (%)	Weighted (%)	Δ (pp)
Global safe-asset demand	25.6	24.5	−1.1
Debt rollover	17.9	17.5	−0.3
Government assets	17.1	17.4	+0.3
Fed purchases	15.7	16.6	+0.9
Primary surpluses	14.3	14.6	+0.3
Financial repression	9.4	9.3	−0.1

Notes. Robustness of the benchmark reconstructed mean value shares (Figure 3 and Table A9) to demographic reweighting, voter/general-population sample. "Unweighted" is the cross-sectional benchmark mean; "Weighted" is the same mean after raking the sample to the age, sex, and education benchmarks in Table 1 by iterative proportional fitting. Age and sex refer to U.S. adults; the education benchmark refers to ages 25 and older. Census region is not in the Prolific export and is omitted. Both columns are computed on the 994 respondents with non-missing raking variables. Δ is weighted minus unweighted, in percentage points. Bond investors are the bond-market sample and are not reweighted.

Table A28. Source-valence ranking: sensitivity to post-stratification weights (voters)

	% primary		% little		Diffusion	
Source	Unwt.	Wt.	Unwt.	Wt.	Unwt.	Wt.
Global safe-asset demand	48	45	14	14	+34	+31
Fed purchases	30	31	18	18	+12	+13
Government assets	32	31	30	31	+2	+0
Debt rollover	34	33	37	35	−3	−2
Primary surpluses	24	25	31	30	−7	−5
Financial repression	20	19	29	29	−9	−10

Notes. Robustness of the source-valence ranking (Table A15) to demographic reweighting, voter/general-population sample. "% primary" is the share naming a source as providing the *greatest* support, and "% little" the share naming it as providing *little or no* support; their difference is the primary-minus-little contrast, reported in the "Diffusion" columns, and the rows are ordered by the unweighted contrast. "Unwt." columns are the cross-sectional values; "Wt." columns use the same raking weights as Table A27: age and sex targets refer to U.S. adults, while the education target refers to adults aged 25 and older. Both sets of columns are computed on the 994 respondents with non-missing raking variables. Bond investors are the bond-market sample and are not reweighted.