The Role of Stock-Flow Reasoning in Understanding the Social Security Trust Fund

Megan E. Weber, University of California, Los Angeles

Hal E. Hershfield, University of California, Los Angeles

Stephen A. Spiller, University of California, Los Angeles

Suzanne B. Shu, Cornell University & NBER

The research reported herein was performed pursuant to grant RDR18000003 from the US Social Security Administration (SSA) funded as part of the Retirement and Disability Research Consortium. The opinions and conclusions expressed are solely those of the authors and do not represent the opinions or policy of SSA, any agency of the Federal Government, or NBER. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of the contents of this report. Reference herein to any specific commercial product, process or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply endorsement, recommendation or favoring by the United States Government or any agency thereof.

Abstract

The financial future of Social Security's trust funds is an important policy topic with significant implications for members of the public who pay taxes and expect to receive benefits in retirement. The funds were created to hold and invest surplus tax revenue not used to pay out benefits, but in recent years, Social Security has started to use this money to fulfill benefits obligations. The funds are projected to become depleted in 2035, at which point benefits payments will have to be reduced. In this research, we draw from the literature on stock-flow reasoning errors and inconsistencies to explore how communication about the trust funds impacts understanding of the situation. In Studies 1 and 2 we randomly assign participants to see information about the trust funds over time presented as a stock (i.e., balance) or in terms of flows (i.e., tax revenue and benefits payments), finding that those who see the stock presentation are significantly more likely to expect benefits to go away completely after depletion. In a third study, we show that explicitly prompting participants to reflect on the continuity of the inflows (via payroll taxes) significantly reduces this common misunderstanding even further. Applying the theoretical lens of stock-flow reasoning, results of this research highlight a key aspect of communications about the trust funds that may contribute to – or be used to remedy – the widespread misconception that benefits will cease when the funds are depleted.

Introduction

For American workers anticipating receiving Social Security retirement (OASI) benefits during their retirement, the solvency of the Social Security system is a relevant and pressing concern. Solvency of the system generally focuses on the status of the Social Security trust funds. The trust funds represent the accumulated surplus that remains from payroll tax income paid into the system by current workers, minus benefits that are paid out to current beneficiaries. Separate projections exist for the OASI (retirement benefits) trust fund versus the DI (disability insurance) trust fund, but the combined OASDI funds are the focus of most reports. Current projections are that both the OASI trust fund and the combined trust funds will be depleted by 2035 (Social Security and Medicare Boards of Trustees, 2022).

Media coverage of Social Security projections intensifies each year when the Board of Trustees releases its Annual Report documenting updates to the expected depletion date for the trust funds. Communicating the implications of these projections is critical but prone to inducing misleading inferences. Coverage that asks, "Will Social Security Run Out of Money?" (Paul, 2022) or highlights a "worst-case scenario" (Werschkul, 2021), and politicians who declare that Social Security will go "bankrupt" (Kiely, 2022), can make the actual impact on retirement benefits appear more severe than trust fund insolvency would actually imply. Because media headlines often emphasize the balance running to zero, non-experts may be likely to arrive at the erroneous conclusion that Social Security benefits will dry up too.

In this project, we argue that the faulty reasoning about the relationship between the trust funds balance and the end of benefits may result from a "stock-flow" reasoning error. After reviewing background literature on the Social Security trust funds, prior research on workers' understanding of the funds, and an overview of research on stock-flow reasoning errors and

inconsistencies, we present two studies that vary the presentation format of information about the trust funds and examine how these different communications impact understanding, focusing specifically on how stock vs. flow information influences perceptions of what will happen to benefits as a result of depletion. Indeed, we find that compared to respondents presented with flow information (about the funds' income and costs), respondents presented with stock information (about the balance of the funds) are more likely to erroneously believe that when the trust fund balance is depleted, benefits will cease altogether. Finally, in a third study, we test an intervention that directly prompts participants to reflect on whether the trust funds will continue to receive inflows (i.e., income collected through payroll taxes), finding that explicitly drawing attention to the continuation of these inflows can further reduce this misconception.

History of the Social Security Trust Fund

When President Franklin D. Roosevelt established the Social Security program in 1935, extra funds (i.e., the difference between worker payroll contributions and paid benefits) were held in an "Old Age Reserve Account". This was replaced by the modern OASI trust fund in 1940, which, unlike the Reserve Account, allowed the excess funds to be temporarily utilized by Treasury for other governmental purposes until they were needed for benefit payouts. The trust fund is overseen by the Board of Trustees, who release an annual report on the health of the fund with a (typically) 75-year projection of taxes and benefits. Even in the beginning years of the fund, the Trustees were concerned about whether the trust fund would be sufficient as the number of eligible workers increased and as life expectancy improved (see Hines & Taylor, 2005, and Pattison, 2015, for histories). One key measure, regularly reported in the Annual Report, is the number of workers per beneficiary; this number has fallen to an all-time low of 2.7 workers per beneficiary in 2020 and is projected to fall lower still (Social Security and Medicare

Boards of Trustees, 2022, page 63). Another key measure from the report, and often the focus of media headlines, is the year in which the trust fund will be depleted, currently estimated at 2035 for the combined trust funds in the 2022 Report. It is this metric that generates the most sensational headlines about the trust fund.

Over the years, a variety of reform measures have been proposed to extend the date at which the Trust Fund is expected to become insolvent. Progress was made in 1983 when payroll tax rates were increased and the full retirement age was also increased (National Commission on Social Security Reform, 1982), at which time solvency was predicted to last through 2060. However, as life expectancy has increased, the date for depletion has moved earlier; by 2004, the insolvency date was projected to be 2044, and by 2022, it is projected to be 2035. Additional reforms have been proposed in recent decades, including the possible introduction of personal accounts (e.g., Kotlikoff & Sachs, 1998) and additional tax increases and benefit cuts (Diamond & Orszag, 2003) but no additional action has yet been taken by Congress. Nevertheless, Hines and Taylor (2005) note that since the early days of Social Security trust fund, there has been a "spirit of optimism that long-term issues confronting the program can be addressed successfully" (p. 8).

One feature of the trust fund worth noting is that its balance is kept on the books at the Treasury Department and can be used as a source of liquidity to reduce the amount of borrowing that needs to be done by Treasury to cover other budget items. Some researchers have argued that this use of off-budget surplus may lead to government overspending (Smetters, 2004). However, this should have negligible impact on the health of the trust fund itself. In this paper, we set aside the use of trust fund's role within the Treasury unified budget and avoid communicating the complicated details of how Treasury's use of the trust fund can affect

national borrowing, saving, and interest rates (as noted in Holahan & Schug, 2000). Instead, we focus only on the dynamics of the system within the boundaries of Social Security itself and treat the trust fund as a Social Security-specific reserve (not unlike its original 1935 "reserve account" design).

Communication about Social Security Trust Fund Status

The concern that media communications about trust fund depletion may lead to incorrect inferences about the end of retirement benefits has led to a small number of prior projects to better communicate the workings of the trust fund to future beneficiaries. One project, by Holahan and Schug (2000), focuses on building a training tool that explains the system to students of economics. They write that, "worries about the solvency of the trust fund often are the result of failure to understand how it works and serve as a distraction from the key factor in the solvency of Social Security system." Their teaching materials introduce flow charts to track how money flows between payroll taxes from workers, Treasury, and retired beneficiaries. They explain the basic concept of payroll tax income minus benefits payouts to retirees, which forms basis for their surplus and deficit definitions, but they expand their model to include how the funds are available to the Treasury for use in government operations and/or investment in the credit markets, and Treasury's use of bonds to track debts and reserves. Since trust fund money can affect national savings and investments through Treasury actions, national interest rates and cost of capital are also affected more generally. Their approach highlights the impact of the trust fund on the larger system of the economy, and the uses of both private and public investments to build a robust economy. However, they do not report any results of empirical testing to show how successful their approach is at communicating the workings of the trust fund to student

audiences, and thus it's unclear whether the tool can be used to reduce incorrect inferences about the trust fund among a broader population.

More directly relevant to the current project is a recent effort by Quinby and Wettstein (2021) which considers the impacts of a poor understanding of the trust fund by looking at how variations in newspaper headlines can affect investing and benefit claiming intentions. The authors test four types of headlines for an otherwise identical article about the projected insolvency of the trust fund and its impact on benefits. In their control condition, participants are simply told that the Social Security Administration has a "long-term financing shortfall". The authors then test three treatment conditions that use more specific headline language. One condition highlights the year in which the trust fund will be "depleted", similar to most current media headlines; a second condition increases the sensationalism by reporting the depletion date as the date the Social Security trust fund reaches "insolvency"; and the third treatment condition includes information in the headline that revenue will still cover three-fourths of benefits after 2034. Compared to the control, the treated respondents expressed an increased desire to claim OASI benefits earlier and also reported more realistic expectations (rather than extreme predictions of 0% or 100%) of monthly benefit size. However, they report no impact on personal savings intentions. As Quinby and Wettstein (2021) summarize, "shifting the media narrative around the trust fund to highlight ongoing revenues could improve the public's understanding of actuarial projections" (p. 3). We attempt to take this approach a step farther by incorporating existing behavioral science research on understanding of stock-flow models into our design of better communications around trust fund projections.

Stock-Flow Reasoning Problems

As previewed above, a key difficulty in understanding the Social Security trust funds may arise from what researchers have termed a stock-flow reasoning error. The cumulative amount of a resource, like the balance of money in a checking account or the amount of water in a reservoir, is a "stock." The changes in the amount of a resource over some defined period, like deposits to and withdrawals from a checking account or water flowing into and out of a reservoir, are the "flows." Given an initial value of the stock, there is a one-to-one correspondence between the stock and the net flow: the stock is the integral of the net flow, and the net flow is the derivative of the stock. As a result, given either a time series of the stock or a time series of the flows (with a starting or ending value of the stock), the information content is calculably the same. But calculus is difficult. As a result, even though the calculable information is the same, people do not respond to the two representations in the same way.

In particular, such stock-flow accumulation processes lead to two related types of problems. First, formal mathematical transformations between stocks and flows are difficult and error-prone tasks even for highly-educated people (e.g., Booth Sweeney & Sterman, 2000; Brunstein, Gonzalez, & Kanter, 2010; Cronin, Gonzalez, & Sterman, 2009; Sterman & Booth Sweeney, 2007). This has sometimes been called the *stock-flow failure*. Because the task is so difficult (i.e., integrating a series of flow values or differentiating a series of stock values), people are prone to rely on a faulty *correlation heuristic*: they wrongly expect the stock trend will tend to match the flow trend. For example, a constant flow into a reservoir followed by a constant flow out of a reservoir ought to lead to a linear increase in the level followed by a linear decrease in the level. Instead, use of the correlation heuristic may lead people to infer that the level of the reservoir suddenly drops when the direction of flow shifts from in to out. Except in rare cases (e.g., no net flows at all, or exponential growth), the two will typically not show the

same pattern. Using such a heuristic can lead to both quantitative and qualitative mistakes, leading to a number of serious errors like violating the conservation of mass in physical systems.

Second, presentation of stocks versus flows can lead to qualitatively different evaluations and forecasts – creating *stock-flow inconsistencies*. Whereas the research cited above indicates that people often cannot successfully translate from one format (e.g., stock) to the other (flow), research on stock-flow inconsistencies addresses how evaluations about the past and forecasts about the future depend on the presentation format even when no translation is necessary. In other circumstances, people may act *as if* they were able to successfully carry out complex calculations, even if the likelihood of successfully conducting such calculations is unlikely. These findings regarding stock-flow inconsistencies suggest reasoning about stock-flow problems is not a case of such successful as-if reasoning but instead reflects sensitivity to the ways in which the same data are presented.

Such stock-flow inconsistencies hold for personal finances (Goldstein, Hershfield, & Benartzi, 2016), evaluations of national employment (Spiller, Reinholtz, & Maglio, 2020), and risk evaluations given COVID test data (Reinholtz, Maglio, & Spiller, 2021). As a concrete example, consider employment in the United States in 2009 (as tested in Spiller et al., 2020). During this time, the number of employed people was decreasing from one month to the next, but at a slowing rate. When one considers the number of employed people, the flow is increasing (from a large negative number to a small negative number). But when one considers the stock, the stock is decreasing (due to the negative flow). As a result, when shown the flow and asked about the economy, a majority of respondents indicated that the economy was getting better. When shown the stock and asked about the economy, a majority of respondents indicated that the economy was getting worse. Presenting the same data in different ways led to qualitatively

different evaluations. There were corresponding effects on forecasts of what respondents anticipated would happen next.

The literature on stock-flow reasoning failures and inconsistencies shows them to be remarkably robust and replicable. They apply across a wide variety of domains, including: atmospheric accumulation of carbon dioxide (Sterman & Booth Sweeney, 2007; Sterman, 2008); water accumulating in a bathtub (Booth Sweeney & Sterman, 2000; Cronin et al., 2009); weight (Brunstein et al., 2010); people in a store (Cronin et al., 2009; Brunstein et al., 2010); distance between cars (Cronin et al., 2009); product inventory in warehouses (Booth Sweeney & Sterman 2000; Spiller et al., 2020); national employment (Spiller et al., 2020); COVID cases (Reinholtz et al., 2021; Villanova 2022; Harman et al., 2021; Padilla et al., 2022); blood glucose levels and other medical measurements (Brunstein et al., 2010); and perhaps most relevantly for the current investigation, corporate and personal cash flows (Booth Sweeney & Sterman 2000; Newell et al., 2016; Spiller et al., 2020). Such stock-flow failures and inconsistencies can be taken as a "stylized fact," even among highly educated participants (e.g., medical students and MIT graduate students; Booth Sweeney & Sterman 2000; Brunstein et al., 2010; Cronin et al., 2009; Sterman & Booth Sweeney 2007). Moreover, they hold across multiple presentation formats, including scatterplots, line charts, barcharts, tables, and verbal descriptions (Cronin et al., 2009; Newell et al., 2016; Spiller et al., 2020).

Understanding Social Security Trust Funds is a Stock-Flow Reasoning Problem

Together, these problems (stock-flow failures regarding translating between formats and stock-flow inconsistencies regarding evaluations and forecasts) have meaningful implications for public understanding of the Social Security Trust Funds. Accumulation and decumulation in the trust funds is a textbook stock-flow reasoning problem. Given the direct applicability of the

domain and the generalizability of the findings above, it would be highly surprising if people *did* fully understand the implications of accumulation for proper interpretation of what happens when the trust fund is depleted. Yet to our knowledge, the robust finding on such problems with understanding accumulation has not been brought to bear on this critical topic.

The time course of the trust fund is illustrated as both a stock and as flows in current SSA communications[1], in which some depictions focus on the balance of money in the trust fund (i.e., the stock, which may present a dire picture to consumers in that it shows depletion by 2035) and some show the changes in inflows and outflows over time (i.e., the flows, which may present a more optimistic view by demonstrating that there will still be taxes paid into the system that can be used to pay for benefits obligations). Reasoning errors like the correlation heuristic and effects on evaluations described above may lead people to wrongly infer that zero balance of the trust funds implies zero outflow from the trust funds.

Reinforcing this problem, and potentially providing initial evidence for confusing stocks and flows, is media characterization that does not unambiguously distinguish the balance of the trust funds from the flows of the trust funds. For example, on December 1, 2020 *CNBC* ran a headline "How Social Security invests its money – and why it may run out of cash really soon"; on August 31, 2021 the *New York Times* tweeted: "Social Security will be depleted in 2033, a year earlier than previously projected…"; and on June 2, 2022 CNBC ran a headline "Social Security fund will be able to pay benefits one year longer than expected, Treasury says." Without distinguishing the stock (the trust funds balance) from the flows (taxes collected and benefits paid out), these headlines may suggest to readers that the Social Security system itself (rather

[1] E.g., https://www.ssa.gov/oact/TRSUM/index.html

than just the trust funds) may be depleted and suggest the headline writer did not find the distinction sufficiently important to draw attention to.

Indeed, such headlines can affect beliefs about future benefits. Even when the content of the article is the same, changing the headline from "Social Security Fund Headed toward Insolvency in 2034, Trustees Find" to "Revenues Projected to Cover Only 75 Percent of Scheduled Social Security Benefits After 2034" affected respondents' beliefs about future benefits (Quinby & Wettstein, 2021, described in more detail above). This perhaps contributes to widespread uncertainty and misunderstanding of the long-term prospects of Social Security retirement benefits. For instance, a recent Pew Research poll (2019) found that 42% of survey respondents doubt they will receive retirement benefits at all, with younger generations showing more pessimism than those closer to retirement (Parker, Morin, & Horowitz, 2019). In a survey by the National Academy of Social Insurance, 28% of respondents reported that after the trust fund was depleted, "Social Security would be unable to pay benefits at all" (Walker, Reno, & Bethell, 2014). A study by Luttmer and Samwick (2018) found that about three-quarters of respondents were either "not too confident" or "not at all confident" that Social Security would be able to pay them the benefits amount they were supposed to receive and that people saw about a 1 in 6 chance that they would not receive anything.

Research on stock-flow reasoning problems both suggests why people may hold such misconceptions and provides some potential tools to reduce the extent of mistaken inferences regarding the Social Security trust fund. The literature on such stock-flow problems has found that reasoning errors arising from stock vs. flow presentations are stubbornly persistent, suggesting education is unlikely to be a feasible, scalable solution. Instead, we propose that alternative presentations, such as deemphasizing the trust funds balance and instead emphasizing

the projection of paid amounts may help to lead to more accurate judgments and inferences based on communications. That is, rather than attempting to solve the stock-flow reasoning problem through education, it may be possible to describe the trust funds using alternative representations or to increase the salience of certain features of the workings of the funds to enhance understanding.

Specifically, in the present research we propose experimentally varying the presentation format and observing how key metrics vary depending on the presentation format. In particular, alternative presentations such as deemphasizing the trust funds balance and instead emphasizing the projection of paid benefits may lead to more accurate judgments and inferences based on communications. We expect that presenting relative inflows and outflows may lead to better-calibrated benefit expectations, which are important for beneficiaries, whereas presenting stocks may lead to better-calibrated forecasts of when the trust fund will reach zero, an important milestone for policymakers. Furthermore, we test an additional intervention that specifically prompts participants to think about whether the trust funds will continue to receive income, hypothesizing that such reflection will further reduce misunderstanding about what happens to benefits by increasing the salience of the fact that inflows will not stop when the funds are depleted. This intervention further supports the proposed role of stock-flow reasoning errors.

Studies

We conducted one pilot study and three full studies to begin testing these research questions. Our first pilot study (N = 403 adults from Amazon Mechanical Turk (AMT)) served as a test for our stimuli and key questions used in later studies. This pilot began with an open-ended question where we asked participants to write in anything they knew (without consulting other materials) about the financial future of Social Security and what might happen to benefits.

Similar to prior work (e.g., Parker et al., 2019), preliminary coding of these responses indicated that 42% of participants believed Social Security would run out of money or benefits would completely cease to exist in the future.

Building on this, Study 1 was designed to test our main hypotheses using graphs that we designed to cleanly manipulate whether participants saw the OASDI trust funds' balance or income and costs over time (both in dollars). We then conducted Study 2 in order to test closer replicas of SSA visuals about the trust funds that used different metrics (i.e., not dollar amounts) and to see whether a third type of graph that showed payable benefits might further improve understanding. Finally, Study 3 tested whether drawing participants' attention to the fact that payroll taxes would still be collected after depletion would further reduce misperceptions about benefits going away after depletion, providing both a potential intervention as well as evidence regarding process. The methods and results for each study are described in the sections that follow. All studies were preregistered on AsPredicted.org prior to data collection[2].

Study 1

In Study 1 we sought to test the effect of stock vs. flow stimuli based on data and presentations of those data used in the 2022 Trustees Reports and related communications. We focused on measures of objective understanding.

Materials and Procedure

Based on a power analysis relying on data from the pilot study, we recruited 1,001 participants from AMT to take this survey. First, participants read the following brief description about the OASDI trust funds:

> *The Social Security Administration uses accounts called "trust funds" to store income that's collected through Social Security taxes and then eventually uses that money to pay*

[2] Study 1: https://aspredicted.org/RRY_76Y; Study 2: https://aspredicted.org/MBF_SVL; Study 3: https://aspredicted.org/5X4_6GF

out benefits. The OASI Trust Fund pays retirement and survivors benefits and the DI Trust Fund pays disability benefits, but the two are often referred to together as the OASDI Trust Funds. At the end of the year in 2021, the OASDI trust funds held $2.85 trillion in total.

Social Security uses the OASDI trust funds to make benefits payments and pay administrative expenses. The trust funds receive income through two sources: tax revenue collected from workers and interest that comes from the investment of the money in US Government securities. In 2021, Social Security's total income from both of these sources was $56 billion lower than its total costs (benefits payments plus administrative expenses). This was the first time in many years that total income was lower than total costs. Social Security predicts that in future years, total income will continue to be lower than total costs. Because of this continued projected deficit, the trust funds balance is projected to reach $0 at some time in 2035.

On the same page as the above description, participants were randomly assigned to see an accompanying graph that showed either the balance (*stock* condition) or total income and expenditures (*flows* condition) of the trust funds for the period 1994 through 2034. Historical data were shown on both graphs with solid lines, while projected data were represented with a dotted line, and both graphs reported the respective metrics in trillions of dollars[3]. These graphs are shown in Figure 1 below.

Figure 1. Study 1 graph stimuli

Stock condition graph and description
The chart below provides some more information on the situation of the OASDI trust funds. The chart depicts the historical and projected OASDI trust fund balance for the period 1994 through 2034, in trillions of dollars (scaled to the current dollar). The solid line shows the historical trust fund balance, and the dotted line shows the projected balance.

[3] The description was based on the 2022 Trustees Report, and the data were taken from the 2022 Supplemental Single-Year Tables (specifically, Table VI.G8 available at https://www.ssa.gov/oact/TR/2022/lr6g8.html).

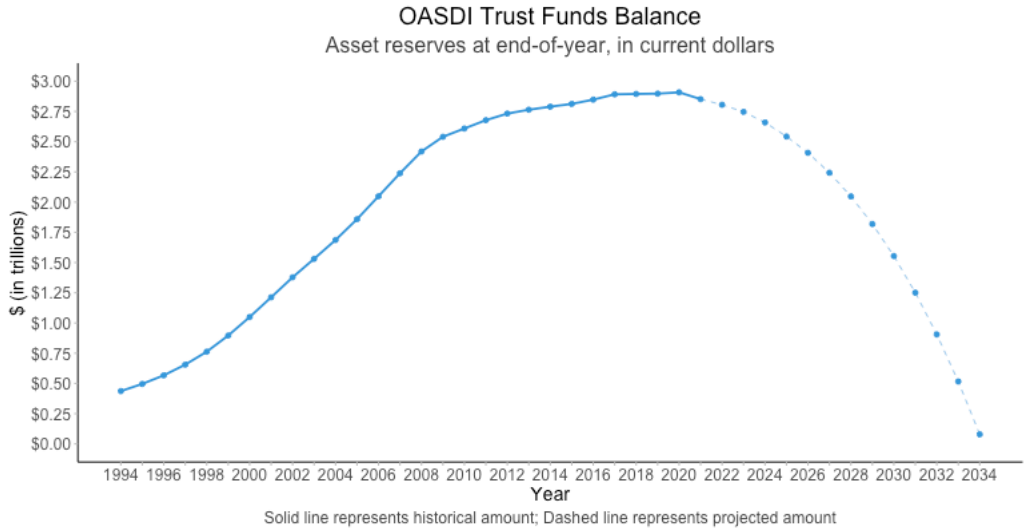

Flows condition graph

The chart below provides some more information on the situation of the OASDI trust funds. The chart depicts the historical and projected OASDI total income (from payroll taxes and interest) and expenditures for the period 1994 through 2034, in trillions of dollars (scaled to the current dollar). The solid lines show the historical trust funds income and expenditures, and the dotted lines show the projected income and expenditures.

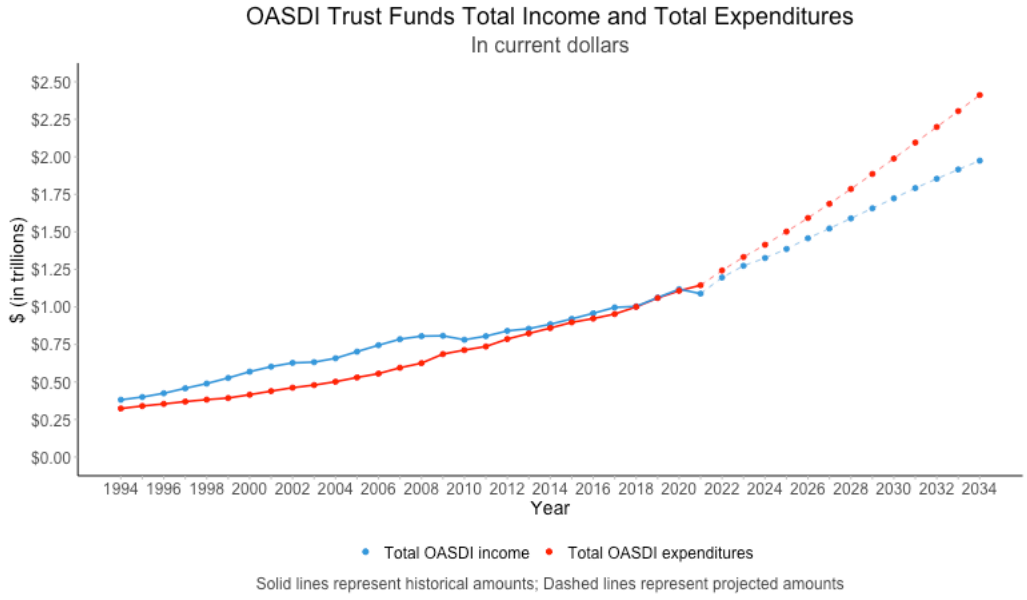

Next, we asked participants four key objective understanding questions: (1) when total costs did/will begin to exceed income; (2) when the trust funds did/will become depleted; (3)

what will happen to benefits if trust funds are depleted; and (4) after depletion, what the monthly retirement benefits amount would be for someone expecting $1,000/month. The order of the first two questions (when costs began to exceed income and when the trust funds would become depleted) was randomized across participants, and we control for this factor in statistical tests reported below. Importantly, whereas both key dates were explicitly mentioned in the text of the description, what would happen to benefits was not discussed, meaning that answering these latter questions required participants to make inferences that went beyond the data that was explicitly provided. After this section, participants responded to questions about their own eligibility for benefits and provided demographic information. The final page of the survey provided links to further resources about the trust funds and benefits calculations. Key measures and manipulations are detailed in Appendix 1.

Analyses

As specified in our preregistration, we coded answers to the first three objective understanding questions according to accuracy (see Table 1 for the breakdown of correct vs. incorrect answers to each question)[4]. While only slightly more than half of participants (56%) correctly identified within 1 year that costs began to exceed income in 2021, about three-fourths (76%) accurately reported within 1 year that the trust funds would become depleted in 2035. Accuracy was much lower when participants were asked about what would happen to benefits: only about a third (33%) chose the correct option stating that benefits would still be paid out in smaller amounts. However, as specified in our preregistration, analyses for this question about

[4] For Studies 1, 2, and 3, missing answers due to a participant failing to answer a question are marked as missing but are included in the base for calculating the proportion of correct answers. These participants are excluded from regression analyses.

what would happen to benefits are actually focused on the proportion of people (60%) who chose the incorrect option reflecting the common belief that benefits would go away completely.

Table 1. Study 1 overall accuracy for objective understanding questions

Question (correct answer)	Correct		Incorrect	
	%	n	%	n
1. When did costs begin to exceed total income? (2021)	56%	559	44%	442
2. When will the trust funds become depleted? (2035)	76%	759	24%	241
3. What will happen to benefits? (still paid out, but smaller)	33%	335	67%	666

To analyze differences in accuracy by condition, we conducted three separate logistic regressions using contrast-coded predictors (-1, 1) for condition (and include another contrast-coded predictor for order of the first two questions as well as an interaction term between the two factors). There was no significant difference across conditions in accuracy in answers to the question about when costs began to exceed income ($b = 0.00$, $z = -0.01$, $p > .99$), though a larger proportion of participants in the stock condition (80%) correctly identified when the funds were projected to become depleted, compared to the flow condition (72%; $b = 0.21$, $z = 2.87$, $p = .004$). Most interestingly, those in the stock condition (64%) were more likely to incorrectly answer that benefits would completely go away as a result of depletion, compared to the flow condition (56%; $b = 0.17$, $z = 2.62$, $p = .009$; see Table 2 for the full breakdown of answers by condition).

Table 2. Study 1 answers to benefits amount question, by condition

Answer	Stock condition		Flow condition	
	%	n	%	n
Benefits go away completely	64%	326	56%	276
Paid, smaller amount (correct)	31%	157	36%	178
Paid, the same amount	4%	22	7%	36
Paid, larger amount	1%	3	1%	3

Following our preregistration, we analyzed the fourth question (about the benefits amount someone expecting $1,000/month would get after depletion) as a continuous variable. Since this question was only asked of those who indicated in the prior question that benefits would be smaller or larger, we treat those who indicated that benefits would go away completely as giving an answer of $0 to this question and those who indicated that benefits would stay the same as giving an answer of $1,000[5]. The average amount given by those who answered this question was $641.40 (*SD* = $237.79), but looking at the whole sample (i.e., using our imputed values), the mean is $276.07 (*SD* = $375.45). Using the same predictors as specified above, we conducted a regression to analyze differences across condition. As expected, the average benefits amount is significantly lower for those in the stock condition (M_{stock} = $239.42, SD_{stock} = $356.10) compared to those in the flow condition (M_{flows} = $313.77, SD_{flows} = $391.18; b = −37.00, $t(996) = -3.13$, $p = .002$). To put these numbers into context, based on the projections in the 2022 Trustees Report, for every $1,000 in scheduled benefits, trust funds income post-depletion would be sufficient to pay $750 to $800, suggesting both groups likely tended to underestimate the amount.

As mentioned in the prior section, we randomized the order of the first two questions that participants answered such that some participants saw the question about when depletion would happen first while others saw the question about when costs began to exceed income first. While this factor was included to counterbalance stimuli and was not of theoretical interest, results indicated that question order did have a significant impact on accuracy for the question about when costs began to exceed income ($b = 0.22$, $z = 2.92$, $p < .001$). Specifically, those who saw this question first were less likely to answer it correctly (51%) than those who saw the question

[5] Here and in future studies, responses greater than $2000 were excluded as more than likely indicating misunderstanding or inattentive responses.

about the depletion date first (61%). We did not have a specific theoretical expectation for this result but return to it in the General Discussion.

Discussion

Using comparable stock and flows graphs based on Social Security data, the results from this study show that different ways of showing trust funds information over time can significantly impact important facets of understanding. Specifically, we found that those shown a stock graph that displayed the trust fund balance over time (making particularly salient the decline to zero around 2035) were more accurate in their understanding of when depletion would happen. This is in line with the existing literature on stock-flow reasoning – and notable given that this date was explicitly provided in the text description for both conditions. However, while the stock condition boosted accuracy on this question, those who saw the stock graph were *less* accurate in translating this information into an understanding of what impact trust fund depletion would have on benefits.

Study 2

Study 2 was designed for two main purposes. First, we sought to investigate whether our main findings would replicate with materials more closely based on those typically included in Trustees Reports (as our stimuli in Study 1 were generated to ensure formal equivalence across conditions). Second, we added a third "enhanced flows" condition, also based on existing SSA communications, that showed the same information on income and costs plus information on "payable benefits." We included this condition to explore whether this additional information on payable benefits would further enhance understanding by making it clear when and how benefits would be impacted.

Materials and Procedure

We recruited 1,503 participants from AMT to participate in this study. The structure of the survey was very similar to that of Study 1: we showed participants a description and graph related to the trust funds, asked a number of objective understanding questions, and finished with questions about eligibility and demographics. The key changes in this study came in the content of the description and the manipulations. Participants were randomly assigned to see one of three graphs displaying information about the combined OASDI trust funds for the period 2000-2050: a *stock* graph that showed the trust fund ratios (balance as a percentage of projected costs for the ensuing year), a *plain flows* graph that showed non-interest income and cost (scheduled benefits) as percentages of taxable payroll, and an *enhanced flows* graph that built on the plain flows graph by including a line for expenditures (payable benefits). We designed these graphs to mimic graphs from the Trustees Report as closely as possible, though we did constrain the x-axis to show the same date range across conditions in order to be consistent (see Figure 2 for our graphs and the SSA graphs we based them on). Since these stimuli were closely based on the graphs included in Trustees Reports, which use data on non-interest income (rather than total income, as we used in Study 1), we updated the description to include key dates and other information for this metric (specifically, the date for when costs began to exceed non-interest income is 2010). Appendix 2 provides further details on the materials used in this study.

Figure 2. Stimuli for Study 2, with comparison to SSA graphs

Study 2 Stimuli

Stock condition

The chart below provides some more information on the situation of the OASDI trust funds. The chart depicts the historical and projected OASDI trust fund ratio for the period 2000 through 2050. The "trust fund ratio" is the value of trust fund asset reserves at the start of a year expressed as a percentage of the projected costs for the ensuing year.

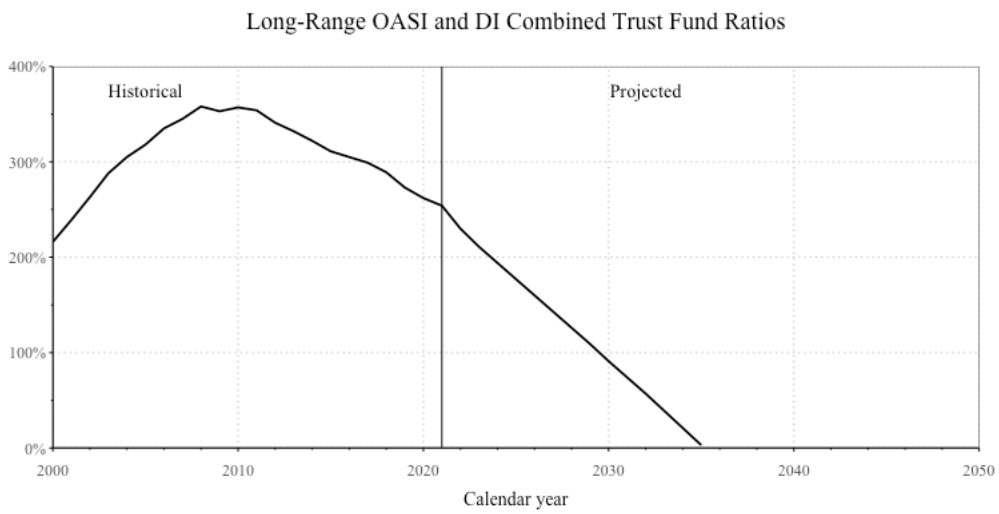

Plain flows condition

The chart below provides some more information on the situation of the OASDI trust funds. The chart depicts the historical and projected year-by-year relationship between OASDI income (excluding interest) and cost (including scheduled benefits) for the period 2000 through 2050. The figure shows all values as percentages of taxable payroll.

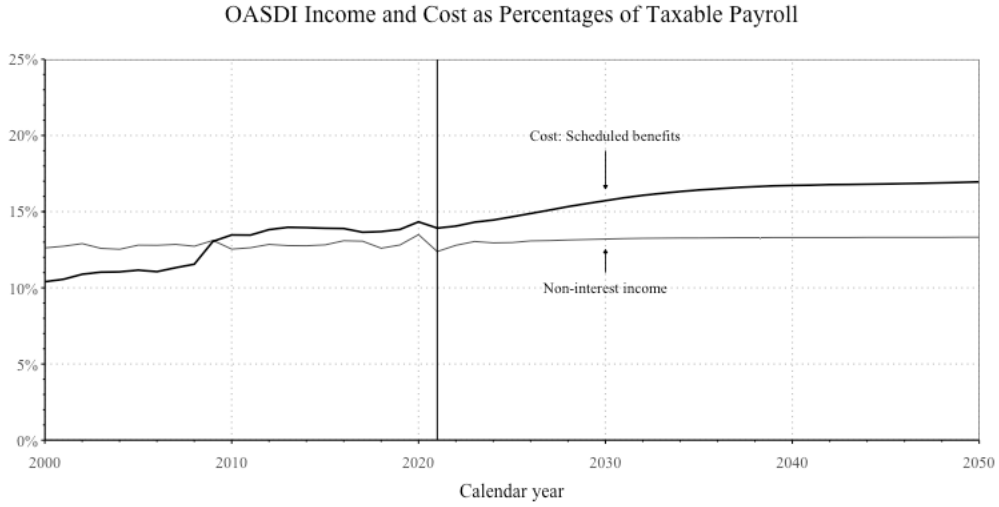

Enhanced flows condition

The chart below provides some more information on the situation of the OASDI trust funds. The chart depicts the historical and projected year-by-year relationship between OASDI income (excluding interest), cost (including scheduled benefits), and expenditures (including payable benefits) for the period 2000 through 2050. The figure shows all values as percentages of taxable payroll.

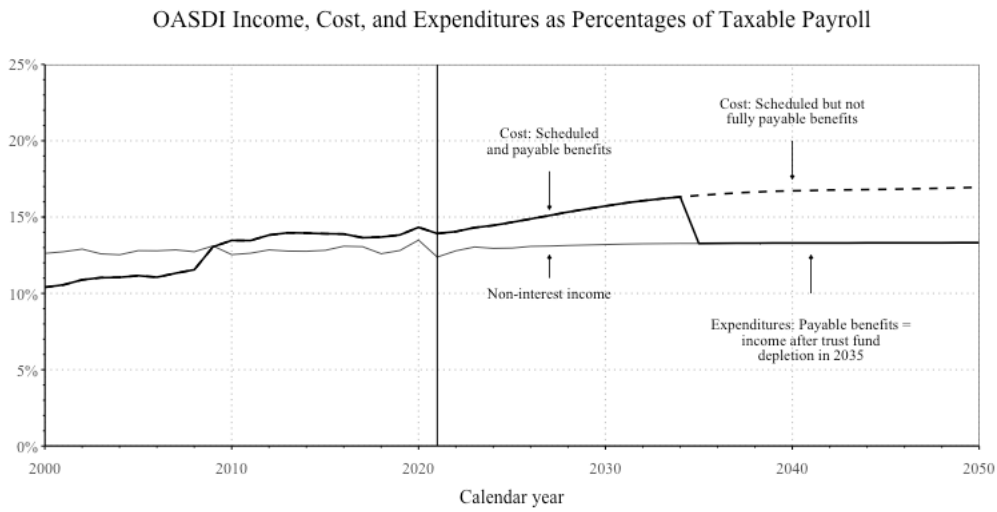

SSA Graphs

(source: https://www.ssa.gov/oact/TR/2022/II_D_project.html#105057)

Stock graph

(note: we only show the intermediate projections line, labeled as II)

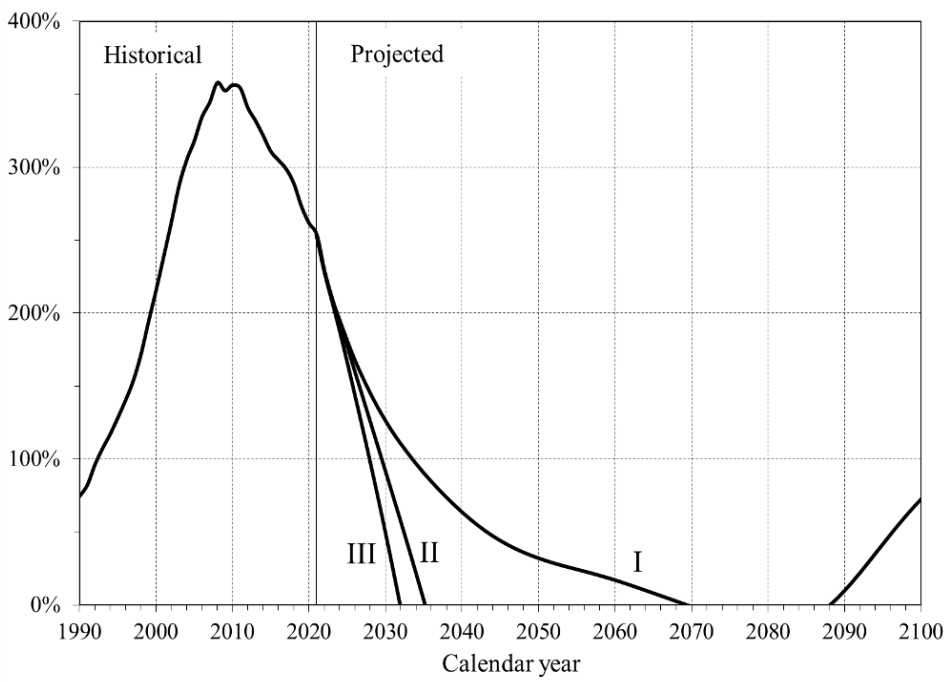

Flows graph

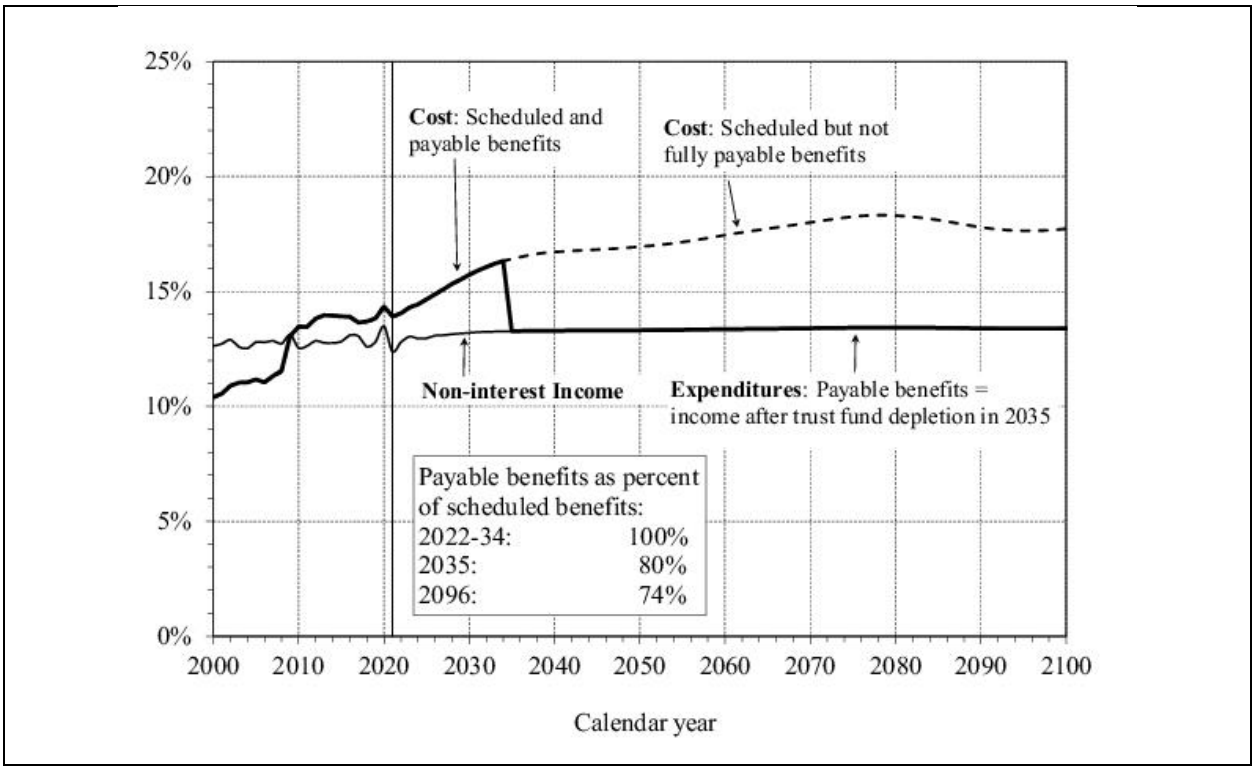

Analyses

As in Study 1, we preregistered that we would recode the objective understanding questions for accuracy (and focus on proportions of incorrect answers about benefits going away). Table 3 below summarizes these results. Across all conditions, only 45% of participants correctly identified when costs began to exceed income and 73% were accurate in reporting when depletion of the funds would occur. Similar to Study 1, around one-third (35%) understood that depletion would lead to a reduction in benefits, while the majority (57%) chose the answer reflecting the misunderstanding that benefits would go away completely.

Table 3. Study 2 overall accuracy for objective understanding questions

Question (correct answer)	Correct		Incorrect	
	%	n	%	n
1. When did costs begin to exceed non-interest income? (2010)	45%	680	54%	819
2. When will the trust funds become depleted? (2035)	73%	1092	27%	410
3. What will happen to benefits? (still paid out, but smaller)	35%	529	65%	974

Based on our design, we created two dummy coded predictors to capture the effect of condition (stock vs. plain flows: 1 if stock condition, 0 otherwise; enhanced vs. plain flows: 1 if enhanced flows condition, 0 otherwise; leaving the plain flows condition as the reference group). As in the previous study, we randomized the order of the two questions about when costs began to exceed income and when depletion would happen and included a contrast-coded predictor (-1, 1) for order and all interactions in the analyses reported here. As in Study 1, there were no significant differences across conditions for accuracy on the question about when costs began to exceed (non-interest) income (stock vs. plain flows: $b = 0.00$, $z = 0.00$, $p > .99$; enhanced vs. plain flows: $b = -0.16$, $z = -1.21$, $p = .22$)[6]. In contrast with Study 1, however, there were also no significant differences for the question about when depletion would happen (stock vs. plain flows: $b = 0.15$, $z = 1.08$, $p = .28$; enhanced vs. plain flows: $b = 0.12$, $z = 0.87$, $p = .38$).

Our main results about understanding of what would happen to benefits were replicated such that those in the stock condition were more likely to choose the wrong answer about benefits going away completely than those in the plain flows condition ($b = 0.29$, $z = 2.22$, $p = .027$). Contrary to our expectations, however, the enhanced flows condition did not provide a further reduction in inaccuracy on this question, as the difference between the plain flows and enhanced flows condition was negligible and not significant ($b = –0.01$, $z = –0.04$, $p = .96$). Table 4 below shows the full breakdown of answers to this question by experimental condition.

We also analyzed data for the question about the expected benefits amount after depletion using the same method as Study 1. Among those who answered this question, the average amount entered was \$626.49 (SD = \$240.16); looking at the whole sample the average was

[6] There was, however, a significant interaction between the dummy code for stock vs. plain flows and question order. The interaction term is difficult to interpret in this context, and we do not have a specific theoretical explanation for it.

$300.91 (*SD* = $387.40). Replicating Study 1, those in the stock condition (M_{stock} = $264.18, SD_{stock} = $373.23) thought benefits would be significantly lower than those in the plain flows condition ($M_{plain\ flows}$ = $317.88, $SD_{plain\ flows}$ = $390.25; $b = -52.66$, $t(1495) = -2.15$, $p = .032$). However, the difference between the plain flows and enhanced flows condition was not significant ($M_{enhanced\ flows}$ = $320.53, $SD_{enhanced\ flows}$ = $396.45; $b = 3.79$, $t(1495) = 0.15$, $p = .88$), providing further support for the conclusion that the mere inclusion of the "payable benefits" line did not necessarily improve understanding.

Table 4. Study 2 answers to benefits amount question, by condition

Answer	Stock condition		Plain flows condition		Enhanced flows condition	
	%	n	%	n	%	n
Benefits go away completely	61%	307	54%	272	54%	272
Paid, smaller amount (correct)	30%	151	37%	188	38%	190
Paid, the same amount	8%	41	7%	37	7%	36
Paid, larger amount	0%	1	1%	5	1%	3

Discussion

In contrast to the prior study, neither question about dates showed any significant differences, which may not be surprising given that the description shown to all participants mentioned these specific dates (although this was also the case in Study 1). More importantly, though, this study provides further evidence of the "stickiness" of the misconception that when the trust funds become depleted, benefits will fully disappear. Building on the previous study, we expected and found that the flows condition mitigates this misunderstanding relative to the stock condition – though more than half of participants across all conditions still think this will be the case. Contrary to our expectations, the results from this study also suggest that including an

explicit line corresponding to "payable benefits" in the enhanced flow condition does not provide any additional benefit for understanding.

Study 3

While the flows manipulations used thus far significantly decreased the extent to which participants thought benefits would go away completely, levels of misunderstanding remained consistently high – over half of participants in the flows conditions in Studies 1 and 2 still chose the incorrect option. Study 3 was designed to test a stronger manipulation that encourages participants to reason through what would happen to the inflows and outflows when the trust funds become depleted.

Materials and Procedure

We recruited 1,001 participants from AMT for this study. The beginning of the survey was identical to previous studies, where participants read a description of the trust funds accompanied by a stock or flow diagram (the same description and graphs as in Study 1) and then answered two questions (in random order) about when costs began to exceed income and when depletion was projected to happen. The key change in this study was the addition of two self-reflection questions: first a yes/no question that asked if participants thought Social Security would continue to collect payroll taxes after depletion and second an open-ended question that asked participants about what they thought Social Security would do with the money if payroll taxes continued to be collected. These questions did not provide participants with any new information. Instead, they merely asked them to consider for themselves the consequences of whatever they already knew. Crucially, participants were randomly assigned to answer these questions as an intervention that came either before or after the questions about what would happen to benefits as a result of depletion. Thus, this study utilized a 2 (presentation: *stock* vs.

flow) by 2 (intervention: *before* vs. *after*) design. The remainder of the survey was the same as Studies 1 and 2 (see Appendix 3 for full measures and manipulations).

Analyses

Overall accuracy for each question (using the same criteria for accuracy as in Study 1) is reported in Table 5 below. Just over half of participants (56%) correctly identified the date when total costs began to exceed total income, and about three-fourths (77%) of participants gave a correct answer when asked about when the trust funds would be depleted. Across all conditions, accuracy about what would happen to benefits after depletion was still very low, with only 38% of participants choosing the correct answer that benefits would still be paid out in smaller amounts.

Table 5. Study 3 overall accuracy for objective understanding questions

Question (correct answer)	Correct		Incorrect	
	%	n	%	n
1. When did costs begin to exceed non-interest income? (2021)	56%	560	44%	441
2. When will the trust funds become depleted? (2035)	77%	766	23%	233
3. What will happen to benefits? (still paid out, but smaller)	38%	381	62%	620

We used regression analyses with contrast-coded predictors (-1, 1) to test the impact of presentation condition and intervention condition. As in previous studies, we also randomized the order of the first two objective knowledge questions about key dates, so all models include a contrast-coded predictor for order (-1, 1) that controls for this factor (we discuss findings related to this at the end of this section). Our models reported here also include all two-way interactions between factors as well as the three-way interaction; however, since none of these interactions are significant, we do not report on them below. Following our preregistration, we focus on stock-flow presentation effects only for the two date questions and both presentation and intervention effects for the questions about what happens to benefits.

As in previous studies, we did not see a significant main effect of presentation on accuracy in identifying when costs began to exceed income ($b = 0.04$, $z = 0.61$, $p = .54$). We did, however, replicate Study 1's finding that significantly more participants in the stock condition (79%) correctly identified the projected depletion date, compared to the flow condition (74%; $b = 0.17$, $z = 2.21$, $p = .027$). (Given that this held in Studies 1 and 3 but not 2, it is possible it relates to properties of the stimuli used in Studies 1 and 3 but not Study 2.) The effect of intervention condition was non-significant for both of these dependent variables (income vs. costs: $b = 0.06$, $z = 0.90$, $p = .37$; depletion: $b = 0.03$, $z = 0.35$, $p = .73$), as expected given that everyone responded to these questions before the intervention.

Table 6 shows the breakdown of answers to the multiple-choice question about what happens to benefits for each experimental condition. Analyses focused on the incorrect answer (i.e., that benefits will go away completely) reveal that the main effect of presentation condition was not significant ($b = 0.04$, $z = 0.59$, $p = .56$; see discussion for further investigation on this lack of an effect). However, the main effect of the intervention condition is significant and quite large: participants who were asked to reflect on what would happen to the inflow of payroll taxes before this question were significantly less likely to choose the incorrect answer option (43%) compared to those who saw the intervention questions after answering (66%; $b = 0.48$, $z = 7.29$, $p < .001$). Additionally, we conducted a secondary (and non-preregistered) analysis focused on the correct answer, finding that participants who answered the intervention questions before were also more likely to indicate that benefits would still be paid out at a smaller amount ($b = -0.36$, $z = -5.40$, $p < .001$).

Turning to the benefits amount question, the overall pattern of results is similar. Among those who answered the question, the average amount written in was \$592.33 (SD = \$260.25);

using imputed values based on our preregistered approach, the overall average was $290.60 (*SD* = $371.87). The main effect of stock-flow presentation was not significant ($b = -4.01$, $t(992) = -0.34$, $p = .73$), but the effect of the intervention condition was significant such that those who saw the intervention questions before answering estimated significantly higher benefits amounts after trust fund depletion (M_{Before} = $355.23, SD_{Before} = $384.60), compared to those who saw the intervention questions after answering (M_{After} = $228.51, SD_{After} = $348.47; $b = -64.04$, $t(992) = -5.49$, $p < .001$). While still below the levels of what Social Security projections suggest could still be paid out (about $750-800), note that the mean amount in the intervention before condition is closer to realistic projections than in the flows conditions from prior studies.

Table 6. Study 3 answers to benefits amount question, by condition

Answer	Intervention before				Intervention after			
	Stock condition		Flows condition		Stock condition		Flows condition	
	%	n	%	n	%	n	%	n
Benefits go away completely	44%	111	43%	101	68%	165	65%	173
Paid, smaller amount (correct)	45%	115	48%	113	29%	70	31%	83
Paid, the same amount	10%	25	8%	19	2%	6	3%	9
Paid, larger amount	1%	3	1%	3	1%	3	0%	1

Finally, as in Studies 1 and 2, we randomized the order in which we asked the first two objective knowledge questions about when costs began to exceed income and when depletion is projected to occur. Results in this study suggest a strong order effect on accuracy for both of these questions. Replicating our finding in Study 1, those who saw the question about the date when costs exceeded income first were less likely to answer correctly (47%) than those who saw that question second (64%; $b = 0.35$, $z = 5.41$, $p < .001$). In this study only, we also find a significant effect of order on accuracy for the depletion date question such that participants who

answered this question first were more likely to answer correctly (80%) than those who answered the question about income and costs first (73%; $b = 0.18$, $z = 2.38$, $p = .017$). We return to potential interpretations in the General Discussion.

Discussion

In this study, we tested an intervention that drew participants' attention directly to the state of the inflows into the trust funds, asking them to consider whether Social Security would continue to collect payroll taxes after depletion and how that tax revenue would be used. Notably, a very large majority – 90% of the sample – believed that Social Security would continue to receive these inflows, suggesting this specific aspect of Social Security policy is well-understood. Importantly, however, those asked to reflect on this *before* answering a question about what would happen to benefits as a result of depletion were both less likely to think benefits would go away completely and more likely to choose the correct answer (that benefits would still be paid out, but only partially). Thus, this intervention, designed specifically to prompt participants to reflect on the fact that the trust funds would continue to receive inflows (i.e., income from payroll taxes), was successful in further combatting the widely-held misperception that Social Security benefits (i.e., outflows) will dry up when the trust funds do.

This study identified a promising additional technique to reduce a common misunderstanding about the SSA trust fund. However, the lack of an effect from the stock-flow manipulation is worth considering further. Given the effectiveness of the intervention, it is possible the stock-flow difference was eliminated in the intervention-before condition. To examine whether the intervention-after conditions were consistent with Studies 1 and 2 (rather than being considered a failure to replicate), we conducted two sets of additional analyses. First, we pooled data from Study 1 with data from the stock and plain flows conditions from Study 2 to

estimate the magnitude of the stock-flow effects observed in prior studies. The point estimate of the stock-flow difference from this pooled analysis was 8.1%, 95% CI [3.9%, 12.2%]. We also calculated the simple effect of presentation within the intervention-after conditions of Study 3 (in other words, for the 510 participants whose survey followed the same exact path as those in Studies 1 and 2). While the simple effect was in the expected direction, with more participants in the stock condition choosing this incorrect answer option, it was not significant ($b = 0.05$, $z = 0.58$, $p = .56$). The point estimate for the stock-flow difference for this group was 2.4%, 95% CI [-6.1%, 9.9%]. Notably, the 95% confidence interval here includes the point estimate of the stock-flow difference from our pooled analysis with data from Studies 1 and 2. This indicates that if the true effect size were 8.1% (as we found in Studies 1 and 2), it would not be unusual to observe a point estimate of 2.4% in this smaller sample.

Second, we conducted a post hoc power analysis to estimate our statistical power to detect a stock-flow difference in this intervention-after group. Given a total sample size of 510, we calculated the statistical power if the true effect size were an 8% difference in the proportion of participants choosing the "benefits will go away completely" option. This analysis indicated that we had less than 50% power to detect an effect of this magnitude in the intervention-after condition of Study 3. While this study does not by itself provide evidence for stock-flow differences, these results are completely consistent with the stock-flow effects observed in prior studies and suggest that the most parsimonious explanation for the null result in this study may be Type II error due to less-precise estimates.

General Discussion

The Social Security trust funds hold and invest the surplus from payments that have been made into the system after benefits have been paid out. Because current projections suggest that

the trust funds will be depleted by 2035, it is critical that researchers and policymakers alike determine how best to communicate information about the situation to constituents and beneficiaries. Yet, media headlines often emphasize how the trust funds will run dry, dwindling down to a balance of $0, leading many everyday consumers to erroneously arrive at the conclusion that benefit payouts will dry up too. In this paper, we argue that one reason for these problematic interpretations is due to "stock-flow" reasoning errors: converting between stocks (the trust funds balance) and flows (the taxes paid in and benefits paid out over time) can be a difficult task to undertake.

In two pre-registered experiments, we found preliminary evidence of differences in inferences regarding the Social Security trust funds based on reasoning about stocks versus flows. Namely, results from Study 1 suggest that presenting the information as a stock leads to higher accuracy regarding when depletion of the trust funds will occur, but lower accuracy in terms of what will happen to benefits payouts as a result (i.e., relative to those in the flows condition, more respondents in the stock condition think that benefits will go away completely). These results hold practical importance as they suggest that framing trust funds information in terms of flows may help overcome misconceptions. Such framing, however, is not a panacea: more than half of those in the flows condition still incorrectly answer questions about benefits payouts.

Study 2 aimed to assess the replicability of Study 1. Using updated stimuli that more closely resembled information presented in the Trustees reports, we again found that framing information in terms of flows reduced the proportion of respondents who think benefits will go away completely after depletion. In Study 2, we also tested an enhanced presentation format, in which we explicitly showed participants a line that represented "payable benefits" (referred to as

enhanced flow condition), with the purpose of highlighting when depletion occurs as well as what happens to benefits after that point. Responses to our central outcome variables were no different in this condition than they were in the standard flows condition, indicating just how difficult it may be to fully de-bias consumer perceptions of the SSA trust funds even with more explicit information.

Finally, Study 3 provides supportive process evidence by testing an intervention aimed at further reducing the misconception that benefits stop completely when the trust funds become depleted. This intervention was simple: we asked two targeted questions that asked participants to explicitly consider whether the funds' inflows (income from payroll taxes) would continue, and if so, what that revenue would be used for. Despite its simplicity, the intervention resulted in a large reduction in the number of people who thought benefits would go away completely for those asked these questions before (vs. after) the questions about what happens to benefits. While there may be ways to improve understanding even further, the results of this study point to this kind of targeted reflection being a promising technique for harnessing citizen beliefs about the continuity of taxes to reinforce expectations for the continuity of benefits. It is also informative about process by indicating that participants are capable of the necessary stock-flow reasoning here. They may merely be unlikely to engage in it unless otherwise prompted (e.g., by the intervention questions). This may provide a promising avenue for improving qualitative stock-flow reasoning problems more generally, which have previously proven stubbornly resistant to accuracy interventions (e.g., Cronin et al., 2009).

Unexpectedly, in both Studies 1 and 3 (though not Study 2), we also found that participants were more likely to correctly report the year in which costs began to exceed revenues when they did so after reporting when the trust fund would be depleted. In fact,

participants who did not enter the correct date for when costs started exceeding income were likely to enter the depletion date. This may suggest that when encountering the question regarding flows after the question regarding stocks, respondents were more likely to distinguish between the two representations. In contrast, when encountering the question regarding flows first, they may not distinguish the two metrics and instead answer with respect to the stock instead.

We cautiously note here that this project was not without its limitations. Specifically, although we targeted a representative sample, our studies were conducted solely on Amazon's Mechanical Turk, leaving open questions of generalizability. It is worth noting that findings using online convenience samples like Amazon Mechanical Turk are strongly correlated with findings using (typically much-more expensive) probability samples (Coppock, Leeper, & Mullinix, 2018; Mullinix et al., 2015; Peyton, Huber, & Coppock, 2021; Snowberg & Yariv, 2021), especially when using screeners such as the CloudResearch approved pool we used (Litman, Robinson, & Abberbock, 2017; Peer et al., 2021). Second, our studies did not include incentive-compatible outcomes nor consequential choices and so it is theoretically possible that incentives could have enhanced accuracy (although unless they also would have reduced the difference between conditions, they would not pose a threat to our estimate of the effect of stock vs. flow presentation or the effect of our intervention; prior research suggests enhanced motivation is not sufficient; e.g., Cronin et al., 2009). Finally, even though we made an effort to generate stimuli that closely resembled material used by Social Security, our materials were still written to be more accessible and understandable than the status quo; it is possible that differences between conditions would be muted were we to use actual SSA materials.

Despite these limitations, the results are nonetheless initially promising and offer opportunities for future research. Specifically, further work is needed to better understand how to even further improve de-biasing attempts. We reason that it may be difficult to alter consumer perceptions about the trust funds precisely because the stock framing of the trust funds may be so strongly ingrained. Indeed, media headlines themselves may reinforce a focus on stocks rather than flows (cf. Jerit & Barabas, 2006). Further, because stock-flow reasoning is so difficult, future interventions may need to be more involved and/or occur over multiple timepoints (rather than the single-shot interventions we employed in these studies, e.g., the "enhanced flow" intervention in Study 2 or the reflection intervention in Study 3). Finally, future work may also wish to study not just perceptions of benefit payouts, but also preferences related to retirement benefit timing, as well as preferences regarding SSA policies regarding taxes and benefits.

Policy decisions about the Social Security trust funds affect members of the public on many levels. Many workers incorporate expected OASI benefits into their retirement planning, meaning changes to taxes and benefits calculations or amounts may have wide-ranging consequences for prospective retirees. It is important that the public fully understands the situation in order to make informed decisions – both for their own retirement planning and in forming policy preferences. This work contributes to a deeper understanding of how communication about this complex policy topic can influence public understanding and discourse.

References

Booth Sweeney, L., & Sterman, J. D. (2000). Bathtub dynamics: initial results of a systems thinking inventory. System Dynamics Review: The Journal of the System Dynamics Society, 16(4), 249– 286.

Brunstein, A., Gonzalez C., & Kanter, S. (2010). Effects of domain experience in the stock-flow failure. System Dynamics Review, 26(4), 347–354.

Coppock, A., Leeper, T. J., & Mullinix, K. J. (2018). Generalizability of heterogeneous treatment effect estimates across samples. *Proceedings of the National Academy of Sciences*, *115*(49), 12441-12446.

Cronin, M. A., Gonzalez, C., & Sterman, J. D. (2009). Why don't well-educated adults understand accumulation? A challenge to researchers, educators, and citizens. Organizational Behavior and Human Decision Processes, 108(1), 116–130.

Diamond, Peter A., and Peter R. Orszag (2003). *Reforming social security: A balanced plan*. Washington, DC: Brookings Institution.

Goldstein, D. G., Hershfield, H. E., & Benartzi, S. (2016). The illusion of wealth and its reversal. *Journal of Marketing Research*, *53*(5), 804-813.

Harman, J. L., Weinhardt, J. M., Beck, J. W., & Mai, I. (2021). Interpreting time-series COVID data: reasoning biases, risk perception, and support for public health measures. *Scientific Reports*, *11*(1), 1-11.

Hines Jr, James R., and Timothy Taylor (2005). "Shortfalls in the long run: Predictions about the Social Security trust fund." *Journal of Economic Perspectives* 19, no. 2, 3-9.

Holahan, William L., and Mark C. Schug (2000). "A simple exposition of the Social Security Trust Fund." *The Journal of Economic Education* 31, no. 4, 340-348.

Jerit, Jennifer, and Jason Barabas, (1996), "Bankrupt Rhetoric: How Misleading Information Affects Knowledge about Social Security," *Public Opinion Quarterly*, 70(3), 2006, 278–303.

Kiely, Eugene (2022). "Examining Rick Scott's Claim That Medicare, Social Security Will Soon Go 'Bankrupt'" Factcheck.org, March 29. https://www.factcheck.org/2022/03/examining-rick-scotts-claim-that-medicare-social-security-will-soon-go-bankrupt/

Kotlikoff, Laurence J., and Jeffrey D. Sachs (1998). "The personal security system: A framework for reforming social security." *Federal Reserve Bank of St. Louis Review,* Mar, 11-13.

Litman, L., Robinson, J., & Abberbock, T. (2017). TurkPrime.com: A versatile crowdsourcing data acquisition platform for the behavioral sciences. *Behavior Research Methods*, *49*(2), 433-442.

Luttmer, E. F., & Samwick, A. A. (2018). The welfare cost of perceived policy uncertainty: Evidence from social security. *American Economic Review*, 108(2), 275–307.

Mullinix, K. J., Leeper, T. J., Druckman, J. N., & Freese, J. (2015). The generalizability of survey experiments. *Journal of Experimental Political Science*, *2*(2), 109-138.

National Commission on Social Security Reform (1982). *Report of the National Commission on Social Security Reform*. Washington, DC: Social Security Administration.

Newell, B. R., Kary, A., Moore, C., & Gonzalez, C. (2016). Managing the budget: Stock-flow reasoning and the CO2 accumulation problem. Topics In Cognitive Science, 8, 138–159.

Padilla, L., Hosseinpour, H., Fygenson, R., Howell, J., Chunara, R., & Bertini, E. (2022). Impact of COVID-19 forecast visualizations on pandemic risk perceptions. *Scientific Reports*, *12*(1), 1-14.

Parker, Kim, Rich Morin, and Juliana Menasce Horowitz (2019). *Looking to the Future, Public Sees an America in Decline on Many Fronts*. Pew Research Center. https://www.pewresearch.org/social-trends/2019/03/21/retirement-social-security-and-long-term-care/

Pattison, David (2015). "Social Security trust fund cash flows and reserves." *Social Security Bulletin* 75, 1.

Paul, Trina (2022). “Will Social Security run out of money? Here’s what could happen to your benefits if Congress doesn’t act.” CNBC, July 8. https://www.cnbc.com/select/will-social-security-run-out-heres-what-you-need-to-know/

Peer, E., Rothschild, D., Gordon, A., Evernden, Z., & Damer, E. (2021). Data quality of platforms and panels for online behavioral research. *Behavior Research Methods*, 1-20.

Peyton, K., Huber, G. A., & Coppock, A. (2021). The generalizability of online experiments conducted during the COVID-19 pandemic. *Journal of Experimental Political Science*, forthcoming.

Quinby, Laura, and Gal Wettstein (2021). "Does Media Coverage of the Social Security Trust Fund Affect Claiming, Saving, and Benefit Expectations?." *SSRN*.

Reinholtz, N. Maglio, S. J. & Spiller, S. A. (in press). Stocks, flows, and risk response to pandemic data. Journal of Experimental Psychology: Applied.

Smetters, Kent (2004). "Is the Social Security trust fund a Store of Value?." *American Economic Review* 94, no. 2, 176-181.

Snowberg, E., & Yariv, L. (2021). Testing the waters: Behavior across participant pools. *American Economic Review*, *111*(2), 687-719.

Social Security and Medicare Boards of Trustees. 2022. *The 2022 Annual Report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Federal Disability Insurance Trust Funds.* Available at https://www.ssa.gov/OACT/TR/2022/tr2022.pdf.

Spiller, S. A., Reinholtz, N., & Maglio, S. J. (2020). Judgments based on stocks and flows: Different presentations of the same data can lead to opposing inferences. Management Science, 66(5), 2213–2231.

Sterman, J. D. (2008). Risk communication on climate: Mental models and mass balance. Science, 322, 532–533.

Sterman, J. D., & Booth Sweeney, L. (2007). Understanding public complacency about climate change: Adults' mental models of climate change violate conservation of matter. *Climatic Change*, 80(3), 213-238.

Villanova, D. (2022). Linear Biases and Pandemic Communications. *Medical Decision Making*, 42(6), 765-775.

Walker, E. A., Reno, V. P., & Bethell, T. N. (2014). *Americans Make Hard Choices on Social Security: A Survey with Trade-Off Analysis.* National Academy of Social Insurance. https://www.nasi.org/wp-content/uploads/2014/11/Americans_Make_Hard_Choices_on_Social_Security.pdf

Werschkul, Ben (2021). "Social Security expert details the 'worst case' scenario — and why millennials are most vulnerable." Yahoo! Finance, September 7. https://finance.yahoo.com/news/social-security-biggest-risk-millennials-most-vulnerable-153719368.html

Appendices

Appendix 1: Study 1 Materials

Trust funds descriptions [seen by everyone]

Introduction

On the next page, we'll be showing you some information about the Social Security Old-Age and Survivors Insurance (OASI) and Disability Insurance (DI) trust funds (also known as the OASDI trust funds), based on the 2022 Trustees Report from the Social Security Administration (SSA). Even if you have seen information about this topic before, please make sure to read carefully as the next part of the survey will involve answering some questions about what you read.

Description

The Social Security Administration uses accounts called "trust funds" to store income that's collected through Social Security taxes and then eventually uses that money to pay out benefits. The OASI Trust Fund pays retirement and survivors benefits and the DI Trust Fund pays disability benefits, but the two are often referred to together as the OASDI Trust Funds. At the end of the year in 2021, the OASDI trust funds held $2.85 trillion in total.

Social Security uses the OASDI trust funds to make benefits payments and pay administrative expenses. The trust funds receive income through two sources: tax revenue collected from workers and interest that comes from the investment of the money in US Government securities. In 2021, Social Security's total income from both of these sources was $56 billion lower than its total costs (benefits payments plus administrative expenses). This was the first time in many years that total income was lower than total costs. Social Security predicts that in future years, total income will continue to be lower than total costs. Because of this continued projected deficit, the trust funds balance is projected to reach $0 at some time in 2035.

Stock condition

The chart below provides some more information on the situation of the OASDI trust funds. The chart depicts the historical and projected OASDI trust fund balance for the period 1994 through 2034, in trillions of dollars (scaled to the current dollar). The solid line shows the historical trust fund balance, and the dotted line shows the projected balance.

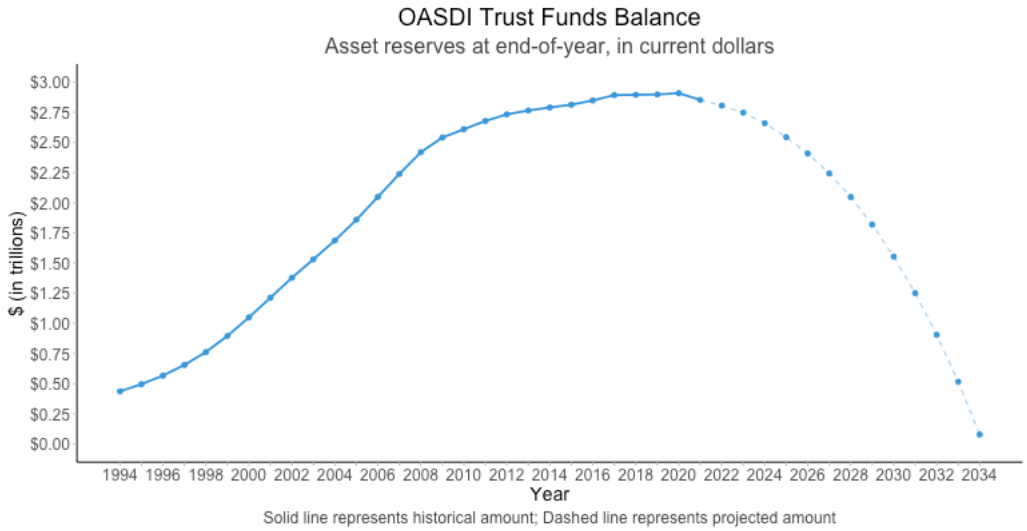

Flows condition

The chart below provides some more information on the situation of the OASDI trust funds. The chart depicts the historical and projected OASDI total income (from payroll taxes and interest) and expenditures for the period 1994 through 2034, in trillions of dollars (scaled to the current dollar). The solid lines show the historical trust funds income and expenditures, and the dotted lines show the projected income and expenditures.

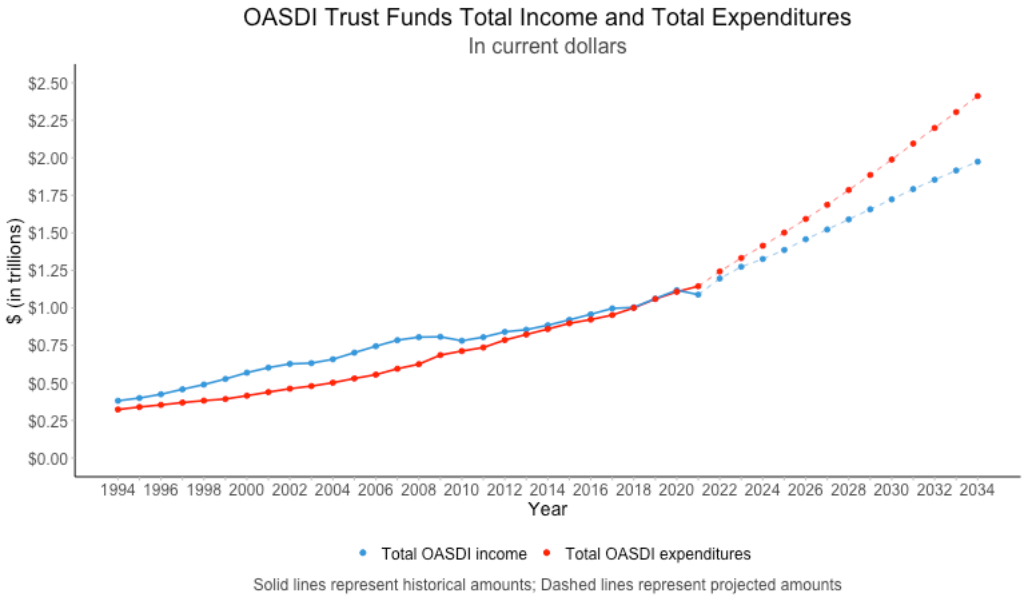

Key outcome measures
[order of Q1 and Q2 randomized]

Question Introduction

Now we'd like to ask you a few questions about your understanding of and reactions to the information you just saw.

For the questions in this section, please answer them with the assumption that no changes are made to the Social Security system regarding taxes or how benefits are calculated.

Q1. Date: costs exceed income
According to your understanding, in what year did or will the Social Security OASDI trust funds' total costs begin to exceed total income?

If you don't think total costs have ever or will ever exceed total income, you can select the "Never" option at the bottom of the dropdown.

▼ 1994 (1994) ... 2050 (2050); Never (-99)

Q2. Date: depletion
According to your understanding, in what year did or will the Social Security OASDI trust funds become depleted? In other words, in what year did or will the trust funds run out of money?

If you don't think the trust funds have ever or will ever be depleted, you can select the "Never" option at the bottom of the dropdown.

▼ 1994 (1994) ... 2050 (2050); Never (-99)

Q3. What happens to benefits (multiple-choice)
Assuming the government does not take any action to increase the amount of income that Social Security collects based on tax revenues – in your view, what is most likely to happen to Social Security benefits if the trust funds are depleted?

If you aren't sure, please select the option that reflects your best guess.

- Social Security will no longer be able to pay out benefits (1)
- Social Security benefits will still get paid out and will be a smaller amount (2)
- Social Security benefits will still get paid out and will be the same amount (3)
- Social Security benefits will still get paid out and will be a larger amount (4)

Q4. Benefits amount **[only asked of those who chose Social Security benefits will still get paid out and will be a smaller/larger amount above]**

You just indicated that Social Security benefits will most likely still get paid out and will be a different amount if the trust funds are depleted.

Again, assuming the government does not take any action to increase the amount of income that Social Security collects based on tax revenues – for someone whose benefits are currently projected to be $1,000 per month, what monthly amount might they expect after depletion of the trust funds?

$[text box].00 per month

Appendix 2: Study 2 Materials

Trust funds descriptions [seen by everyone]

Introduction

On the next page, we'll be showing you some information about the Social Security Old-Age and Survivors Insurance (OASI) and Disability Insurance (DI) trust funds (also known as the OASDI trust funds), based on the 2022 Trustees Report from the Social Security Administration (SSA). Even if you have seen information about this topic before, please make sure to read carefully as the next part of the survey will involve answering some questions about what you read.

Description

The Social Security Administration uses accounts called "trust funds" to store income that's collected through Social Security taxes and then eventually uses that money to pay out benefits. The OASI Trust Fund pays retirement and survivors benefits and the DI Trust Fund pays disability benefits, but the two are often referred to together as the OASDI Trust Funds. At the end of the year in 2021, the OASDI trust funds held $2.85 trillion in total.

Social Security uses the OASDI trust funds to make benefits payments and pay administrative expenses. The trust funds receive income through two sources: tax revenue collected from workers and interest that comes from the investment of the money in US Government securities. In 2010, Social Security's non-interest income (i.e., income from taxes) was $49 billion lower than its total costs (benefits payments plus administrative expenses). This was the first time in many years that non-interest income was lower than total costs. Non-interest income has continued to be lower than total costs since then, and Social Security predicts that this will be the case in future years. Because of this continued projected deficit, the trust funds balance is projected to reach $0 at some time in 2035.

Stock condition

The chart below provides some more information on the situation of the OASDI trust funds. The chart depicts the historical and projected OASDI trust fund ratio for the period 2000 through 2050. The "trust fund ratio" is the value of trust fund asset reserves at the start of a year expressed as a percentage of the projected costs for the ensuing year.

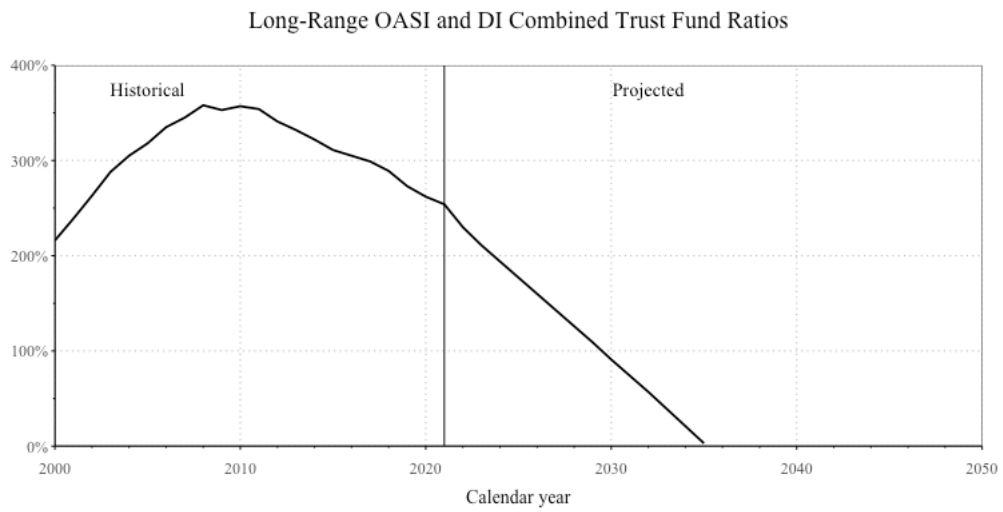

Plain flows condition

The chart below provides some more information on the situation of the OASDI trust funds. The chart depicts the historical and projected year-by-year relationship between OASDI income (excluding interest) and cost (including scheduled benefits) for the period 2000 through 2050. The figure shows all values as percentages of taxable payroll.

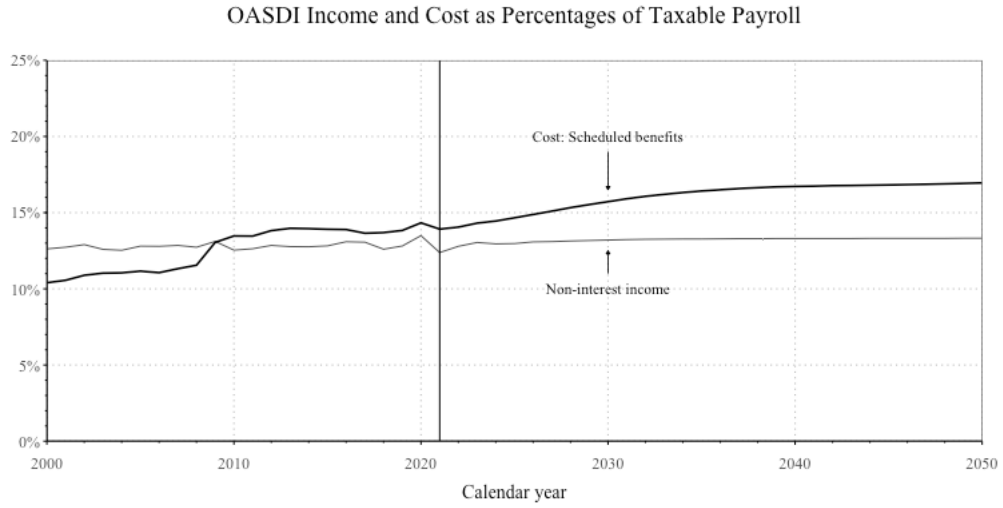

Enhanced flows condition

The chart below provides some more information on the situation of the OASDI trust funds. The chart depicts the historical and projected year-by-year relationship between OASDI income (excluding interest), cost (including scheduled benefits), and expenditures (including payable benefits) for the period 2000 through 2050. The figure shows all values as percentages of taxable payroll.

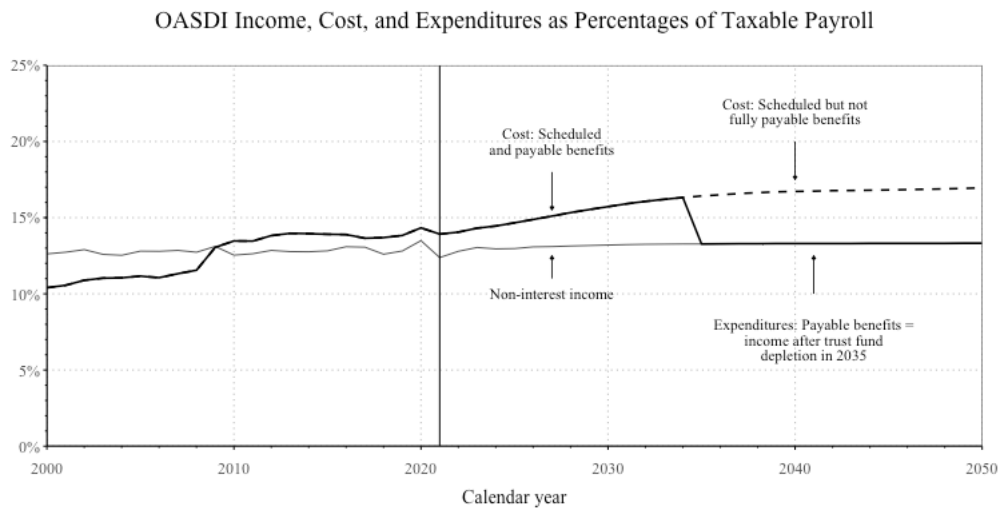

Key outcome measures
[order of Q1 and Q2 randomized]

Question Introduction
Now we'd like to ask you a few questions about your understanding of and reactions to the information you just saw.

For the questions in this section, please answer them with the assumption that no changes are made to the Social Security system regarding taxes or how benefits are calculated.

Q1. Date: costs exceed income
According to your understanding, in what year did or will the Social Security OASDI trust funds' total costs begin to exceed non-interest income?

If you don't think total costs have ever or will ever exceed non-interest income, you can select the "Never" option at the bottom of the dropdown.

▼ 2000 (2000) ... 2050 (2050); Never (-99)

Q2. Date: depletion
According to your understanding, in what year did or will the Social Security OASDI trust funds become depleted? In other words, in what year did or will the trust funds run out of money?

If you don't think the trust funds have ever or will ever be depleted, you can select the "Never" option at the bottom of the dropdown.

▼ 2000 (2000) ... 2050 (2050); Never (-99)

Q3. What happens to benefits (multiple-choice)

Assuming the government does not take any action to increase the amount of income that Social Security collects based on tax revenues – in your view, what is most likely to happen to Social Security benefits if the trust funds are depleted?

If you aren't sure, please select the option that reflects your best guess.

- Social Security will no longer be able to pay out benefits (1)
- Social Security benefits will still get paid out and will be a smaller amount (2)
- Social Security benefits will still get paid out and will be the same amount (3)
- Social Security benefits will still get paid out and will be a larger amount (4)

Q4. Benefits amount **[only asked of those who chose Social Security benefits will still get paid out and will be a smaller/larger amount above]**

You just indicated that Social Security benefits will most likely still get paid out and will be a different amount if the trust funds are depleted.

Again, assuming the government does not take any action to increase the amount of income that Social Security collects based on tax revenues – for someone whose benefits are currently projected to be $1,000 per month, what monthly amount might they expect after depletion of the trust funds?

$[text box].00 per month

Appendix 3: Study 3 Materials

Trust funds descriptions [seen by everyone]

Introduction

On the next page, we'll be showing you some information about the Social Security Old-Age and Survivors Insurance (OASI) and Disability Insurance (DI) trust funds (also known as the OASDI trust funds), based on the 2022 Trustees Report from the Social Security Administration (SSA). Even if you have seen information about this topic before, please make sure to read carefully as the next part of the survey will involve answering some questions about what you read.

Description

The Social Security Administration uses accounts called "trust funds" to store income that's collected through Social Security taxes and then eventually uses that money to pay out benefits. The OASI Trust Fund pays retirement and survivors benefits and the DI Trust Fund pays disability benefits, but the two are often referred to together as the OASDI Trust Funds. At the end of the year in 2021, the OASDI trust funds held $2.85 trillion in total.

Social Security uses the OASDI trust funds to make benefits payments and pay administrative expenses. The trust funds receive income through two sources: tax revenue collected from workers and interest that comes from the investment of the money in US Government securities. In 2021, Social Security's total income from both of these sources was $56 billion lower than its total costs (benefits payments plus administrative expenses). This was the first time in many years that total income was lower than total costs. Social Security predicts that in future years, total income will continue to be lower than total costs. Because of this continued projected deficit, the trust funds balance is projected to reach $0 at some time in 2035.

Stock condition

The chart below provides some more information on the situation of the OASDI trust funds. The chart depicts the historical and projected OASDI trust fund balance for the period 1994 through 2034, in trillions of dollars (scaled to the current dollar). The solid line shows the historical trust fund balance, and the dotted line shows the projected balance.

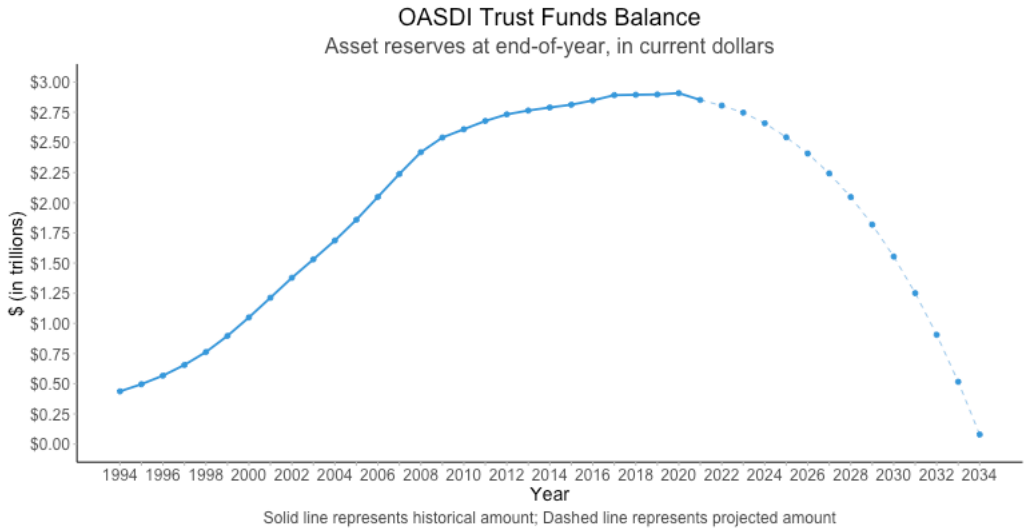

Flows condition

The chart below provides some more information on the situation of the OASDI trust funds. The chart depicts the historical and projected OASDI total income (from taxes and interest) and expenditures for the period 1994 through 2034, in trillions of dollars (scaled to the current dollar). The solid lines show the historical trust funds income and expenditures, and the dotted lines show the projected income and expenditures.

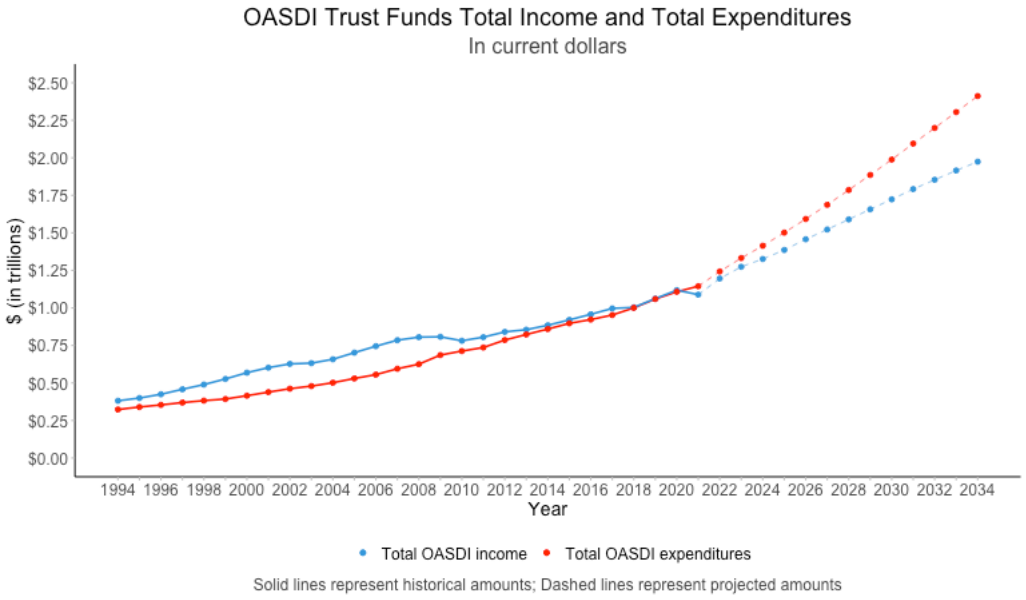

Key outcome measures
[order of Q1 and Q2 randomized]

Question Introduction

Now we'd like to ask you a few questions about your understanding of and reactions to the information you just saw.

For the questions in this section, please answer them with the assumption that no changes are made to the Social Security system regarding taxes or how benefits are calculated.

Q1. Date: costs exceed income
According to your understanding, in what year did or will the Social Security OASDI trust funds' total costs begin to exceed total income?

If you don't think total costs have ever or will ever exceed total income, you can select the "Never" option at the bottom of the dropdown.

▼ 1994 (1994) ... 2050 (2050); Never (-99)

Q2. Date: depletion
According to your understanding, in what year did or will the Social Security OASDI trust funds become depleted? In other words, in what year did or will the trust funds run out of money?

If you don't think the trust funds have ever or will ever be depleted, you can select the "Never" option at the bottom of the dropdown.

▼ 1994 (1994) ... 2050 (2050); Never (-99)

Q3. What happens to benefits (multiple-choice)
Assuming the government does not take any action to increase the amount of income that Social Security collects based on tax revenues – in your view, what is most likely to happen to Social Security benefits if the trust funds are depleted?

If you aren't sure, please select the option that reflects your best guess.

- Social Security will no longer be able to pay out benefits (1)
- Social Security benefits will still get paid out and will be a smaller amount (2)
- Social Security benefits will still get paid out and will be the same amount (3)
- Social Security benefits will still get paid out and will be a larger amount (4)

Q4. Benefits amount **[only asked of those who chose Social Security benefits will still get paid out and will be a smaller/larger amount above]**

You just indicated that Social Security benefits will most likely still get paid out and will be a different amount if the trust funds are depleted.

Again, assuming the government does not take any action to increase the amount of income that Social Security collects based on tax revenues – for someone whose benefits are currently projected to be $1,000 per month, what monthly amount might they expect after depletion of the trust funds?

$[text box].00 per month

Intervention Questions
[participants randomly assigned to answer these questions either before or after questions Q3 and Q4 above]

Q1. Income/Inflows Reflection
Most of the money that goes into the OASDI trust funds comes from a dedicated payroll tax. Employees and employers each pay 6.2% of wages up to the taxable maximum of $147,000 (in 2022).

Assuming that Congress doesn't change the Social Security system of taxation, do you expect Social Security to continue to collect payroll taxes from workers and employees if the trust funds are depleted?

- Yes, Social Security will continue to collect payroll taxes (1)
- No, Social Security will not continue to collect payroll taxes (2)

Q2. Outflows Reflection
If the trust funds are depleted and if Social Security continues to collect payroll taxes, how do you expect those payroll taxes to be used?

[text box]