

# Perceptions of Program Abuse and Support for Social Insurance

American Politics Research  
1–17  
© The Author(s) 2020  
Article reuse guidelines:  
sagepub.com/journals-permissions  
DOI: 10.1177/1532673X20924784  
journals.sagepub.com/home/apr



Scott E. Bokemper<sup>1</sup>, Albert H. Fang<sup>2</sup>, and Gregory A. Huber<sup>1</sup>

## Abstract

Do perceptions of abuse in social insurance programs undercut program support? Answering this causal question is difficult because perceptions of program abuse can arise from multiple potential causes. Examining the case of disability insurance, we circumvent this challenges using laboratory experiments to study the interplay between program abuse and program support. Specifically, we test whether participants vote to reduce benefit levels when they observe program abuse, even if that abuse is not directly costly to them. We use a labor market shock to induce program abuse and show that the observation of a healthy worker receiving benefits causes workers who are unaffected by the shock to vote to lower benefits. This effect arises only when reducing benefit levels also reduces taxes. Our results demonstrate a causal link between program abuse and diminished support for social insurance, validating accounts that stress how violations of cooperative norms can undercut socially beneficial government programs.

## Keywords

political economy of social insurance, program abuse perceptions, support for social insurance, labor market shocks, free riding, laboratory experiment

“Over half the people on disability are either anxious or their back hurts. Join the club. Who doesn’t get up a little anxious for work every day and their back hurts?”

Senator Rand Paul (R-KY), January 14, 2015<sup>1</sup>

The claim that program abuse is rampant in the nation’s disability insurance program, as illustrated by the epigraph above, has become commonplace in contemporary American political discourse.<sup>2</sup> Do perceptions that disability benefits are being abused by people who could work decrease support for disability insurance? More generally, do perceptions of abuse in public social insurance programs undercut support for those programs?

These questions have grown in importance as government social insurance programs have become targets for retrenchment despite both their important function in a market economy and their design, in which benefits are available only to those who contribute to a program’s financing. Mandatory contributions allow individuals to collectively pool their resources to address individual risks. Because program beneficiaries subsidize their own care and benefits are limited to defined shocks, these risk pooling programs are often popular vis-à-vis more traditional redistributive welfare programs. At the same time, social insurance benefits raise the possibility of abuse if individuals claim benefits despite not needing them for the reasons specified by the program. This abuse may in

turn undercut support for social insurance by leading citizens to believe beneficiaries are not deserving of assistance.<sup>3</sup>

An extensive observational survey literature documents the correlation between beliefs about program abuse and the deservingness of program beneficiaries and support for both social insurance (e.g., government health care) and government benefits more generally (e.g., welfare) (Aarøe & Petersen, 2014; Petersen et al., 2011; van Oorschot, 2000). While this literature has documented important associations between beliefs about beneficiaries and support for government programs, existing work is unable to demonstrate a causal relationship between perceiving program abuse and program support because perceptions of program abuse may originate from many sources. Indeed, a key finding from prior research is that individual-level characteristics, like ideology, often explain program support. In addition, features of a particular beneficiary, like their race, often affect people’s perceptions of that beneficiary’s deservingness, but that effect could arise for multiple reasons, for example,

<sup>1</sup>Yale University, New Haven, CT, USA

<sup>2</sup>Independent Scholar, Seattle, WA, USA

## Corresponding Author:

Gregory A. Huber, Forst Family Professor of Political Science, Department of Political Science and Institution for Social and Policy Studies, Yale University, 77 Prospect Street, PO Box 208209, New Haven, CT 06520, USA.

Email: gregory.huber@yale.edu

beliefs about program abuse, affect or prejudice toward the potential beneficiary, and so on. Individuals may therefore report perceptions of abuse (or non-deserving beneficiaries) not because they have observed abuse (or undeserving behavior), but because they are already opposed to the program or have beliefs and attitudes about its beneficiaries. More generally, any analysis that relies on as-given expressed attitudes or beliefs about beneficiaries (e.g., whether they are abusing the program) to predict program support will produce biased estimates unless all factors that explain those beliefs and affect program support are accounted for.

The notion that deservingness is a key source of program support originates in two distinct potential causal pathways. The first is an account emphasizing self-interest—low deservingness indicates people are using a program despite not actually needing it or not having merited care (e.g., because they had not previously contributed), and that abuse reduces the chance that a person’s own subsequent meritorious claims will be met. For example, in Bokemper and DeScioli (2018), abuse of a “public fund” depletes the fund such that players with valid claims may not receive benefits. Program abuse might therefore undercut program support because as with any cooperative task, free-riding upsets an equilibrium of shared public goods production. The second causal pathway emphasizes instead an intrinsic motivation to enforce appropriate norms. In this account, individuals seek to punish those, or prevent abuse by, individuals who take advantage of others because such behavior violates core ideas about appropriate behavior (e.g., reciprocity), regardless of whether it affects one’s own self-interest in this particular case. For example, third party observers in public goods games may sanction other players who violate expectations about public goods production despite not having a stake in that particular public good (Fehr & Gächter, 2002).

In this article, we examine whether this second mechanism, an intrinsic concern for beneficiary deservingness, causes reduced program support. To do so, we construct an experimental setting that remove any direct material effect of program abuse on player welfare and instead ask whether observing program abuse reduces program support. This particular setting is a “hard case” for testing the intrinsic norms account because program abuse is costless, but the program supports players’ material self-interest. While prior experimental work has examined reactions to norm violations when players can sanction a specific individual, it is unknown whether reduced deservingness will lead individuals to reduce support for the overall program that is being abused.<sup>4</sup> To answer this question, we designed and implemented a series of laboratory experiments where subjects participate in a novel disability insurance game. The structure of the game first allows us to induce program abuse by some players and then randomly exposes other players to information about that abuse. We then assess whether observing abuse reduces support for disability insurance, even in a setting where all players in expectation benefit from the social

insurance program and abuse does not endanger access to benefits. Our baseline experiment demonstrates that by reducing the market wage some workers receive, we are able to induce program abuse, operationalized as claiming benefits when one could in fact participate in the labor market. Our core political experiments reveal that unaffected workers who are randomly informed about other players’ abusive behavior support lower benefits when funding those benefits requires higher taxes. This work therefore isolates the causal effect of perceiving program abuse on support for social insurance programs.

Our design necessarily abstracts from many of the details of real labor markets and social insurance programs. Nor do we incorporate features like beneficiary characteristics or social norms that may explain baseline differences in support for social insurance, although for more universal social insurance programs most individuals are eligible to receive benefits conditional on prior contributions. Our simplification, however, is an important advantage because it allows us the control necessary to isolate the causal effect of program abuse on program support. Furthermore, while our design is stylized, it retains the core logic that prior work has suggested is at play—individuals turning to social insurance to counteract diminished labor market returns (e.g., Autor & Duggan, 2003; Black et al., 2002) and observation of this abusive behavior undercutting subsequent support for social insurance. In addition, because our game uses monetary incentives, we can create baseline preferences over outcomes, thereby providing a degree of control not feasible outside the lab setting (see, e.g., Woon, 2012). Likewise, choices in this game are costly decisions, which reduces concerns that survey responses are pure “cheap talk” that do not correspond with actual behaviors.

While the experimental setting with costly choices provides researchers with a great deal of control, we remain sensitive to issues of external validity. For example, our subjects are undergraduates who participate in laboratory studies for modest amounts of money. How to extrapolate our results to other subjects and outside the laboratory setting is an issue we discuss in greater detail below. Nonetheless, given the importance of the question of how to sustain citizen support for social insurance and government spending more generally, as well as the key theoretical role that prior work has ascribed to ideas about deservingness and program abuse, we believe that isolating this causal relationship is of great theoretical and empirical importance, a key advantage of our approach. We return to issues of external validity in the conclusion.

## **Program Abuse, Deservingness, and Support for Social Insurance**

A core theoretical concept for understanding public support for government spending programs is the deservingness heuristic. Differences in perceptions of deservingness across recipients, as well as differences across programs (Cook &

Barrett, 1992), are correlated with support for government spending. Underpinning this empirical relationship is the idea that individuals are unwilling to engage in sharing or risk reduction when they fear they are being taken advantage of, and non-deservingness is a clear measure of such abuse.

While perceptions of deservingness are widely understood to explain support for government redistributive spending programs, many of which are at best marginally popular, the relationship between support for social insurance and perceptions of deservingness has not been fully developed. For example, welfare is not broadly popular, and welfare is a racialized policy domain in which individuals perceive minorities are over-represented among beneficiaries and have doubts about beneficiary deservingness that are correlated with beneficiary race (Gilens, 1999). By contrast, social insurance programs are often broadly popular because they insure people against a number of common risks and therefore have beneficiary pools that have both contributed to the program ahead of time and are more broadly representative of the general population. But past evidence about popularity of social insurance is confined to programs that have non-discretionary eligibility determinations (e.g., such as the work history and age requirement for receiving social security retirement benefits; Cook et al., 2002; Page & Shapiro, 1992). We know much less about the role of deservingness in affecting support for social insurance programs like Social Security Disability Insurance (SSDI) and unemployment insurance that combine important social insurance design features with discretionary decisions about a specific claim for benefits. For example, to be eligible to receive unemployment insurance, workers must have a demonstrated work history, but must also prove that they are unemployed “through no fault of their own,” a standard that requires judgment about the merits of a particular beneficiary’s request for benefits. SSDI similarly combines required pre-disability worker tax payments during a set employment period with a discretionary judgment about a claimant’s impairment status and their ability to work in the future.

This study therefore seeks to understand whether perceptions of deservingness are similarly useful in understanding support for social insurance programs when discretionary determinations are required to receive benefits. We focus on SSDI, whose structure we describe in greater detail below, because recent calls to reform SSDI are grounded in the belief that fraud and abuse are prevalent in the program. While the Social Security Administration claims that the amount of fraud is a “fraction of one percent,”<sup>5</sup> citizens generally believe that program abuse is far more prevalent. Analyzing a nationally representative sample from the 2016 Cooperative Congressional Election Study (CCES), Fang and Huber (2019) found that over half of the people they surveyed (62.8%) believe that at least 20% of SSDI recipients collect benefits because they do not want to work.

Our theorizing draws on prior work that has examined the role of deservingness in two distinct ways. First, a broad

range of studies in public opinion and political psychology have examined the correlates of perceptions of deservingness and the relationship between those perceptions and program support. Second, work in the behavioral tradition, often in the lab setting, has examined how recipient behavior in public goods games and the possibility of abuse affect cooperation and risk sharing. We review this work to focus our hypotheses and questions.

## Perceived Beneficiary Deservingness and Program Support

A rich observational and experimental literature has examined the antecedents of perceptions of beneficiary deservingness (e.g., Jensen & Petersen, 2017; Petersen et al., 2011; van Oorschot, 2000). For instance, experimental work on welfare has found that people are less supportive of providing benefits to someone who is not making an effort to find a job (Aarøe & Petersen, 2014; Petersen, 2012). Likewise, certain demographic markers, like a beneficiary’s race, are also associated with differences in perceptions of deservingness and support for providing welfare benefits (Gilens, 1999).

While race of beneficiary is an important explanation for support for many redistributive government programs, prior research suggests beneficiary race is not a good candidate to explain attitudes toward SSDI. Fang and Huber (2019) experimentally manipulated the race of a hypothetical SSDI beneficiary to examine whether respondents’ beliefs about whether that person deserved benefits were influenced by the person’s race. In contrast to prior work on welfare, they find that there is no difference between respondents’ ratings of a White beneficiary’s deservingness compared with a Black beneficiary. Furthermore, there were also no heterogeneous treatment effects by political party, that is, Democrats and Republicans were not asymmetrically affected by the race of a potential beneficiary, nor is the general beneficiary pool for SSDI perceived in racial terms.

Fang and Huber (2019) assessed beliefs about the percentage of SSDI beneficiaries who are African Americans compared with recipients of benefits from Temporary Assistance to Needy Families (TANF). Respondents believe that SSDI recipients are considerably less likely to be minorities compared with TANF recipients, and the modal respondent perceives African Americans to be underrepresented compared with their population proportion among those receiving SSDI (see Online Appendix Figure A4). Taken together, these results suggest SSDI is not a racialized policy domain. Racial attitudes are therefore unlikely candidates for explaining individual political attitudes and behavior regarding SSDI (and potentially social insurance programs more broadly). This, of course, makes the questions of why social insurance programs have become less popular over time, how individuals form judgments about beneficiaries, and what determines their level of support for these programs even more interesting—if race is not the driving factor, what

is? This article attempts to provide a more general theory for citizens' concerns about fraud and abuse in social insurance programs grounded in the psychology of free riding and cooperation that relates to existing work on how race of beneficiary is an important source of attitudes about beneficiary deservingness.

Most generally, perceiving that beneficiaries are undeserving undercuts the reciprocity needed to sustain the production of public goods because it calls into question the idea that one can rely on others when needed (e.g., Axelrod, 1984). This logic has been used to explain low public support for welfare, because those who are perceived as undeserving are both taking advantage of others by claiming benefits instead of working and not paying taxes that support government more generally. In this type of account, citizens care about undeserving beneficiaries because they do not want to bear the additional cost of financing benefits. This direct self-interest account is analogous to an individual in a public goods setting terminating their contributions because other people are free riding on their efforts. Thus, reducing support for social insurance when observing an undeserving beneficiary aligns with one's own self-interest as their material well-being improves by bearing less of the cost of the program.

Less clear, however, is how observing free riding in a social insurance context affects program support when withdrawing support does not change citizens' material payoffs, that is, if an individual does not personally bear a greater cost due to free riding, will they reduce their support when they observe an individual taking advantage of the programs? That is, concerns about deservingness could instead be driven by non-material concerns, like an opposition to people violating norms of acceptable behavior.

This norms account of deservingness in social insurance has not been fully explored. Because social insurance programs require beneficiaries to ex ante contribute to those programs prior to making claims, they undercut many of the concerns about voluntary contributions being insufficient to sustain public goods production (Olson, 1965). In addition, because these programs provide future benefits that diminish the downside risks associated with labor market and other shocks, concerns about deservingness may not be sufficient to reduce overall program support. Still, beneficiaries who claim social insurance benefits despite not really needing them are still taking advantage of the contributions of others to the risk pool, potentially reducing the willingness of actors to support the program in the first place (even if this free riding does not increase ex ante costs for the other individuals).

If deservingness helps in understanding support for social insurance programs, a key question is what markers of deservingness are most important for shaping those perceptions. Although inferences about deservingness may be made on the basis of recipient demographic characteristics, beliefs about deservingness may also be affected by observations of recipient behavior or knowledge of the basis on which recipients claim need. For example, focusing specifically on disability insurance, Fang and Huber (2019) analyze data from

a survey experiment conducted on the 2016 CCES and report that citizens judged a hypothetical recipient with a harder to diagnose impairment, a mood disorder, as less deserving relative to hypothetical recipients with conditions that are easier to diagnose, like a stroke-induced intellectual disability or chronic heart failure. They speculate that claims for harder to diagnose impairments are potentially subject to greater abuse by otherwise healthy individuals. Overall, however, studies of how deservingness affects support for social insurance programs remains scarce, and prior work has not directly manipulated observations of program abuse, despite the key role that abuse plays in undercutting beliefs about deservingness and therefore support for these public goods.

### Program Abuse, Cooperation, and Risk Sharing

A second relevant body of research employs behavioral experiments to study program abuse, cooperation, and risk sharing. Several disparate strands of research in this tradition are relevant to our theoretical expectations. First, scholars have used incentivized experiments specifically to study support for social insurance. Generally, participants earn income by completing a real effort task and subsequently vote on a tax rate to finance a social insurance program (e.g., Barber et al., 2013) or whether social insurance should be provided (Ahlquist et al., 2017). Generally, these experiments have found that participants who are exposed to a higher risk of losing their income are more supportive of social insurance. Barber et al. (2013), for example, found that participants conditioned their votes on a tax rate to fund a social program more on their risk of suffering income loss than they did on their earnings. While these studies shed light on the antecedents of support for social insurance programs, they do not examine how (potentially illegitimate) program beneficiary behavior affects support for social insurance.

Other work has examined how variation in program design affects program abuse. In two experiments, Bokemper and DeScioli (2018) examine how the rules of social insurance programs affect the rate of program abuse. These experiments use an unemployment game in which participants look for a job to pay taxes to fund a program that mimics real world unemployment insurance. Importantly, participants can make withdrawals from the program whether they are employed or not. In one experiment, employed participants were less likely to claim benefits when they had a limited number of requests compared with when requests were unregulated. In a second experiment, participants were more likely to be completely honest with their claims when they were subject to a probabilistic fine for claiming benefits when employed as opposed to a ban from receiving benefits in the next round. While these experiments allowed the researchers to study how program design affects abuse, these experiments did not address how directly observing program abuse affected program support.

Although work in this tradition has not examined how observations of free riding affect support for social insurance specifically, past experimental research on public goods provision has tested the effects of observing free riding on cooperation by investigating how providing information about past contributions affects one's future contributions. In one experiment, participants contributed more to a public good in later periods when they received feedback on the contributions of other individuals in each period relative to when feedback was not provided or was provided at the group level (Sell & Wilson, 1991). However, other research has found that the effect of individual information is present when contributions are identifiable (i.e., are linked to participant-specific identifier), but not when they are anonymous (Croson & Marks, 1998). This work shows that the potential for anonymous free riding in contribution behavior can undercut cooperation, but it is unclear if the potential for abuse in claiming collectively funded benefits generates similar opposition to those benefits.

### The Need for Additional Research

Overall, the balance of the evidence in prior research suggests that deservingness is conceptually central to understanding attitudes about government spending generally, but whether that logic explains support for social insurance programs is unknown. Moreover, what is missing from the literature on support for government spending is a direct test of how the observation of program abuse, a key marker of deservingness, affects program support.

In addition to these theoretical questions, we note that prior survey research on deservingness and program support suffers from three main limitations. First, these studies do not put subjects in a costly choice environment. Second, while these studies manipulate beneficiary characteristics—which proxy a bundle of perceptions—they do not directly manipulate perceptions of abuse. This leaves open the possibility that those markers alter both perceptions of deservingness and program support directly, without operating through the deservingness pathway.<sup>6</sup> Third, the overwhelming focus of survey-based research on deservingness is about redistribution; very little of this work is on social insurance or support for social insurance programs more generally, beyond whether a particular individual should receive benefits.<sup>7</sup>

Although incentivized behavioral experiments resolve the concern that survey responses reflect costless signals instead of costly choices, existing lab-based behavioral research on social insurance preferences has not examined how perceptions of abuse undercuts support for social insurance. More generally, that work has also not directly engaged the key institutional question of whether social insurance program design can counteract the negative effects of observing abuse on program support.

### Background: SSDI

In the United States, SSDI is a social insurance program that provides benefits to workers who are unable to engage in productive labor because of a medical condition. In 2016, approximately \$143 billion in benefits were paid to almost 9 million disabled workers (Congressional Budget Office, 2016). Individuals qualify for disability benefits by having a condition that prevents them from working and by having worked at a job that paid Social Security taxes for a sufficient amount of time.<sup>8</sup> Once an individual qualifies for SSDI benefits, they receive cash payments until they recover, return to work, reach retirement age, or die.

Over the past 45 years, the number of SSDI beneficiaries has nearly quadrupled and the percentage of the working-age (20–66 years) population claiming benefits has risen from 1.3% in 1970 to 4.5% in 2015 (Congressional Budget Office, 2016). While the rise of disability rolls has been attributed to multiple factors (see Autor, 2015), scholars have emphasized the role of economic conditions in the increased usage of SSDI. In particular, a growing body of literature in economics has studied the relationship between the health of the economy and the utilization of disability insurance (Autor & Duggan, 2003; Black et al., 2002). Using merged administrative and economic data, Autor & Duggan (2003) found that counties that experienced negative employment shocks saw a greater increase in unique SSDI applicants relative to counties that did not. Similarly, areas that have been most affected by the decline in the profitability of the coal industry have experienced larger increases in SSDI utilization (Black et al., 2002). Broadly, this research suggests that disability benefits can serve as a substitute for labor when economic conditions deteriorate, although it is unclear from this evidence whether those taking up disability in fact qualify for the benefits, that is, it could be that when the economy is performing well, people who legally could qualify for disability benefits continue to work, nor is it clear whether these dynamics have downstream political effects on program support.

The rapid growth of SSDI has led to growing political conflict over the program and calls for reform.<sup>9</sup> Highlighting the program's vulnerability, in 2015 Congress was forced to allocate more funding to the Disability Insurance Trust Fund to stave off impending insolvency (Morton & Huston, 2018).

### Disability Insurance Game

We designed a series of incentivized laboratory experiments to provide a direct test of the causal effect of observing program abuse (as a marker of deservingness) on support for social insurance. The experiments simulate a political economy in which individuals are workers who pay taxes into a social insurance scheme and collect disability insurance if they experience a labor market shock. Importantly, this game allows disability benefit recipients to choose to abuse the program, and we randomly manipulate whether actual

**Table 1.** Wage, Tax, and Benefit Levels by Job Skill Level, Baseline Game.

Job skill level	Effort	Wage	Tax	Working period net (wage-effort-tax)	Initial disability benefit amount
Low skill	4 tokens	24 tokens	6 tokens	14 tokens	6 tokens
Medium skill	4 tokens	32 tokens	8 tokens	20 tokens	8 tokens
High skill	4 tokens	48 tokens	12 tokens	32 tokens	12 tokens

program abuse behavior is observed by other players. In extensions of the baseline game, which we describe below, players vote on disability benefits (and taxes). In this section we provide an overview of the baseline disability insurance game and note key features of the game that are central to our research design. Although extrapolation from this stylized context is not without risks, the key advantage of control over mechanisms allows for theoretically informative causal inference that is otherwise highly difficult.

### Baseline Game

Participants in the game take the role of workers in a laboratory economy.<sup>10</sup> All participants are fully informed about their payoffs before making any decision that affected those payoffs. Unless otherwise noted, participants are fully informed about all features of the experimental design in the initial instructions. In this economy, workers are randomly assigned a static skill level—low skill, medium skill, or high skill—that qualifies them to hold certain jobs.<sup>11</sup> Workers can hold a job at or below their skill level. For example, a medium skill worker can hold a medium skill or low skill job, but not a high-skill job. Workers start the game by holding a job that matches their skill level.

Regardless of skill level, workers spend four tokens in effort to work at their job in each period. When working in a period, each worker faces a 1-in-5 chance of suffering an injury. If a worker is injured, they lose their ability to spend effort and lose their job. Consequently, workers collect disability benefits when they are injured. The amount of disability benefits a worker receives is determined by the skill level of their most recent job. A worker can continue to collect benefits for the duration of the game or until they return to work. Table 1 shows the wage, tax, and benefit schedule for a single period of the baseline disability insurance game.

In each period that a worker is injured, they have a 1-in-3 chance of recovering. When a worker recovers from an injury, they receive four tokens of effort and can look for a job. Importantly, a worker who has recovered from an injury can choose to continue to collect disability benefits and also receive four additional tokens despite not exerting effort by working. This is our proxy for program abuse—taking disability benefits when one could otherwise be working. These parameters were chosen so that social insurance is net welfare improving and, if risk neutral, it was always desirable to return to work rather than stay on disability.<sup>12</sup>

A period of the game progresses as follows:

1. Workers learn if they are healthy or injured.
2. Workers earn tokens based on their health and whether they are holding a job or are collecting disability benefits:
  - (a) A healthy worker who has a job at their skill level spends effort to work.
  - (b) A healthy worker who has a job below their skill level decides whether to take a new job if it is available. After this decision, they spend effort to work.
  - (c) A healthy worker who is on disability can either take a job for which they are qualified and spend effort to work or continue to collect disability.
  - (d) An injured worker collects disability.
3. Workers see how many workers held jobs and how many workers collected disability for that particular period. Workers also see whether another randomly chosen individual who is collecting disability is healthy.

This feedback described in Step 3 is randomized at the player level, so that each worker could receive information about a different worker on disability in any given period in which multiple workers were on disability (i.e., one player could learn of a healthy worker on disability, while the other learned about an injured player on disability). Both the total number of workers on disability and the potential observation of a healthy person on disability are proxies for the potential abuse of the disability system, although the second is much more informative because it provides the actual health status of a worker on disability. Our rationale for individual-level randomization is that, in addition to allowing clearer causal identification, it also proxies the idiosyncratic social processes by which individuals learn about the merits of others in their community who are on disability. (Alternatively, one can think of individuals as proxying different communities that learn about abuse at different times.)

### Features of the Design

We draw attention to three distinctive features of the disability insurance game. First, decision-making in our game is costly. Unlike prior experimental or observational survey

research, our game places participants in a setting in which decisions about benefits (and in some versions of the game, taxes to fund benefits) can affect their own well-being. Moreover, decision-making in our game takes place in a social insurance setting where all individuals can reasonably expect, in light of the set chance of being injured, to utilize the program at some point during the game.

Second, our design manipulates the perception of deservingness that is directly attributable to program abuse—the utilization of disability insurance when one could otherwise be working. In the game we provide participants with group-level information about the number of other participants utilizing a laboratory disability insurance program. Furthermore, we also show participants information about one randomly selected individual who is drawing disability benefits. Unlike past work on information and public goods, while this individual is someone who is not contributing to fund the disability program, this does not necessarily mean that they are free riding either. Specifically, we show a participant whether someone else who is claiming disability benefits is healthy or injured. If the person is injured, they are legitimately partaking in a social insurance program. By contrast, if this person is healthy, they can be categorized as a free rider because they are intentionally taking a benefit from a public good without contributing (Delton et al., 2012). Furthermore, this recovered individual no longer needs the benefit they are taking, which is a core criterion for deservingness (van Oorschot, 2000). By isolating and experimentally manipulating observed instances of abuse, we can gain causal leverage on the relationship between information about deservingness and support for disability insurance.

Third, we exogenously manipulate levels of program abuse by simulating labor market shocks that prior observational work has suggested leads to the use of disability insurance in lieu of engaging in the labor market. We impose a labor market shock that removes new medium skill jobs from the labor market. This leaves recovered medium skill workers with a decision to either work at a low skill job, which pays less than they had previously earned, or to continue to claim disability despite being recovered.<sup>13</sup> We expect this targeted labor market shock will cause more workers to stay on disability than if they could return to work at their previous wage level, allowing us to induce greater program abuse through this intervention and test its downstream effects. Importantly, this labor market shock is only targeted at medium skill workers such that workers at other skill levels are unaffected, analogous to how industry- or geographically targeted shocks affect some subset of the economy. By contrast, macroeconomic shocks that affect all workers may affect both program abuse and the perceived need for social insurance, thereby conflating abuse with program need. In addition, this abuse does not directly endanger any other player's benefits, because actual benefits are unaffected by utilization or abuse.

## Main Hypotheses

We use the disability insurance game to test the following three main hypotheses about the political economy of disability insurance:

**Hypothesis 1: Inferred abuse.** Situations that lead to inferences that there is greater program abuse will generically reduce program support irrespective of whether one observes abuse.

**Hypothesis 2: Observed abuse.** Observing program abuse will reduce program support.

**Hypothesis 3: Costliness enhancement.** The negative effect of inferring or observing abuse on program support is magnified when increases in program benefits are costly.

For all three hypotheses, we focus specifically on the behavior of workers who are unaffected by an exogenous shock (i.e., low- and high-skill workers unaffected by the labor market shock targeting medium skill workers) respond, because, as we articulate above, we cannot disentangle multiple effects for those directly affected by the shock.

## Baseline Validation Experiment: Labor Market Shocks Induce Program Abuse

We first conducted three sessions in an incentivized experimental laboratory to validate that targeted labor market shocks induce greater program abuse among affected workers. This setup is a common baseline to which political features are added in our subsequent core experiment. It is designed to create abuse whose political implications we study and is not meant to be a complete model of labor market decision-making. We emphasize the common features here, as they are used below.

### Participants and Procedure

Each session had 12 participants who interacted anonymously on networked computers. The experiment was implemented using HTML and Javascript in a web-based application. Participants were recruited from an existing university subject pool for an experimental lab and only participated in a single session of the experiment.<sup>14</sup> No deception was used in the session or at the facility more generally. In each session, we obtained informed consent and participants then read instructions and completed a comprehension quiz. After all participants in a session were finished with the instructions, they played the disability insurance game. At the conclusion of the game, participants received a show-up payment of \$12 and payment for additional tokens earned during the experiment. Tokens were converted to money at a rate of 2 cents for one token. The average earnings for

participants not including the show-up payment were \$18.69 ( $SD = \$5.99$ ) for a task that took about 1 hr.

There are two key manipulations in this design. First, we randomly assigned each participant's skill level such that each session had two low skill, eight medium skill, and two high-skill workers. Each participant's skill level remained the same throughout the duration of the session.

Second, we implemented a 60 period disability insurance game with a within-subjects labor market manipulation. In Periods 1–20, participants could return to work at low, medium, or high-skill jobs (depending on their skill level). In Periods 21–40, we removed new medium skill jobs from the labor market. This means that medium skill workers who recovered from an injury in these periods could only take a low skill job or continue to collect disability. Workers who were actively working a medium skill job could continue to work at it until they suffered an injury. Starting in Period 41, new medium skill jobs were once again available and remained so until the end of the experiment in Period 60. Participants were informed ahead of time that the labor market would change, but were not informed when that change would occur, what skill level jobs would be affected, or how long the change would last. However, when the labor market change occurs, workers at all skill levels saw that medium skill jobs were not available when they made the decision whether to re-enter the labor force.

### **Results: Targeted Labor Market Shocks Induce Targeted Program Abuse**

To validate our design, we examine whether labor market shocks alter the decisions players make about whether or not to return to work after recovering from an injury.

Averaging across all players, workers are injured (not healthy) and therefore receiving benefits in 35% (763 out of 2,160) of the periods of play. In addition, healthy workers are on disability 4.5% of the time.<sup>15</sup> Focusing on the 251 periods in which a worker first became healthy (i.e., was injured last period but is no longer injured), in 9.6% of cases the worker does not immediately return to work. This happens 7% of the time for low skilled workers, 13% of the time for medium skilled workers, and never for high-skilled workers.

Our experimental manipulation in Periods 21–40 removed medium skilled jobs, putting medium skill workers in a situation where returning to work would yield a smaller net payoff (14 tokens) than they were previously accustomed to earning (20 tokens), even though returning to work still increased the payoffs for that period by two tokens relative to remaining on disability (net of 12 tokens). Did this intervention cause medium skill workers to forgo work? In the aggregate, the chance of staying on disability when becoming healthy is 4% prior to the elimination of medium skill jobs, rises to 24% in the periods (21–40) when medium skill jobs are unavailable, and then falls again to 3% when all jobs are again available.

This effect is concentrated solely among medium skill workers. For medium skill workers the elimination of

medium skill jobs substantially increases the chances of staying on disability in the period they become healthy. This figure rises from 4% in Periods 1–20 to 32% in Periods 21–40 and then falls to only 3% when medium skill jobs are once again available. By contrast, high-skill workers never remain on disability, and for low skill workers the chances of staying on disability are relatively constant across these three settings (8%, 8%, and 5%, respectively).

Formal statistical analysis confirming this result appears in Table 2. In Column 1, we predict the chance that a healthy worker is on disability (1 = yes, 0 = no) for all cases in which a worker is healthy using indicators for the bad labor market (no medium skill jobs available, Periods 21–40) and the post-bad labor market condition (the excluded category is Periods 1–20 when all jobs are available). Our modeling approach is ordinary least squares (OLS) regression with player fixed effects, which account for all static individual-level player factors, including a worker's skill. The bad labor market increases by 15 points ( $p < .01$ ) the probability a healthy worker is on disability. In Column 2, we restrict our attention to periods in which a worker recovered from being injured and estimate an even larger effect of 22 points ( $p < .01$ ). Finally, in Column 3, we examine heterogeneity by worker skill. Consistent with the simple cross tabulations presented earlier, the bad labor market increases the chances a newly healthy medium skill worker stays on disability by  $28.6 + 2.2 \approx 31$  points ( $p < .01$ ). By contrast, the effect is a statistically insignificant 2 points for low- and high-skilled workers (those who are not directly affected by the bad labor market).

These results provide clear evidence that we can experimentally induce program abuse. In addition, this is a targeted intervention that affects only the behavior of medium skill workers. This result allows us to extend our game to investigate our primary interest of understanding how the increased incidence of healthy people forgoing work and instead receiving disability affects support for disability insurance, a question we can investigate in our subsequent experiment.

### **Core Experimental Design: Program Abuse and Support for Social Insurance**

Our core design incorporates a political process into the apolitical baseline validation experiment described above. We test whether a bad labor market, which we show above induces program abuse, causes unaffected workers to reduce their support for disability insurance and whether those effects are caused by the specific observation of program abuse. In addition, we examine whether these effects are contingent on unaffected workers bearing the cost of program abuse.

### **Participants and Procedure**

We conducted six sessions to examine how healthy workers staying on disability affects public support for disability insurance benefits. We recruited participants from the same

**Table 2.** Experiencing a Bad Labor Market Increases the Probability Medium Skill Workers Claim Disability Despite Being Healthy (Baseline Validation Experiment).

	(1)	(2)	(3)
	DV: Healthy but receiving disability in period (1 = yes, 0 = no)		
	Conditional on worker being healthy	Conditional on worker becoming healthy this period	Conditional on worker becoming healthy this period
Bad labor market (1 = no medium skill jobs)	.149*** (.015)	.218*** (.044)	.022 (.076)
Post bad labor market (1 = yes)	-.002 (.015)	-.011 (.042)	-.015 (.069)
Medium skill worker $\times$ bad labor market			.286*** (.092)
Medium skill worker $\times$ post bad labor market			.003 (.085)
Constant	.014 (.011)	.034 (.031)	.033 (.030)
Player fixed effects?	Yes	Yes	Yes
Observations	1,397	251	251
Number of participants	36	36	36
R <sup>2</sup>	.095	.144	.194

Note. The table reports estimates from fixed-effects models with standard errors in parentheses. The interaction terms in the models estimate simple effects, not interaction effects. (As worker skill does not vary within participant, the main effect of medium skill worker drops out of the fixed-effects models.)

\* $p < .1$ . \*\* $p < .05$ . \*\*\* $p < .01$ .

laboratory subject pool to participate in an experiment (subjects could not participate in multiple experiments). As with the baseline validation experiment, each session had 12 participants and participants only participated in one session of the experiment. After we obtained informed consent, participants read instructions, completed comprehension questions, and played the disability insurance game. Each participant received a \$12 show-up payment and received an additional 2 cents for each token they earned. The average earnings for participants from the experiment not including the show-up payment were \$19.54 ( $SD = \$5.71$ ) for a task that took about 1 hr. We implemented the same within-subjects labor market manipulation and kept the same distribution of worker skill level.

We add two key additions to our earlier game. First, we added an additional political process to the game where workers voted on benefit levels every five periods (e.g., in Periods 5, 10, . . . , 60) and could vote to increase, decrease, or leave benefits the same for disabled workers. Voting is a natural political mechanism for aggregating preferences. Voting directly on benefits rather than, for example, over candidates allows us to isolate our focus on the politically relevant choice dimension: benefit levels. A vote to increase (decrease) benefits raised (lowered) benefits for all disabled workers by one token, and deviations from initial benefits could not exceed three tokens. This was done to ensure that working was always more profitable than collecting disability at a given skill level. All decisions were by majority vote, which required a decision to increase or decrease benefits to receive at least seven votes (out of 12 participants) to change benefits. Voting took place immediately after workers received feedback about the overall

distribution of players working and on disability in that period and whether an individual worker collecting disability was healthy or injured.

Second, we randomized whether or not benefits were costly to provide at the session level. Sessions were randomly assigned to either a “Free Lunch” condition or to a “Budget Constraint” condition. In the Free Lunch condition ( $N = 36$ ; three sessions), benefit levels had no effect on taxes, and therefore there is no cost for any worker to enhancing the generosity of her own benefits because workers do not experience the cost of greater generosity. This is analogous to a political environment in which the additional costs of greater disability payments are borne by a third party or are simply unaccounted for by workers. However, this also means that any reduction in benefit levels does not similarly reduce taxes on workers, which means that reducing benefits only imposes costs (in the form of lost potential benefits). As such, this makes this condition a particularly conservative test case for whether observing abuse causes reduced support.

By contrast, in the Budget Constraint condition ( $N=36$ ; three sessions), increasing (decreasing) benefits also raised (diminished) the taxes employed workers pay by half a token each period. This is similar to a social insurance environment with a real Budget Constraint in which worker benefits are funded by worker taxes or in a way that makes the cost of changes in benefit levels more apparent to workers. In this environment, compared with the Free Lunch condition, it was therefore in expectation less costly to reduce benefits because reduced benefits were offset by reduced taxes. Note, however, that we did not make *program abuse* directly costly in any condition, that is, unlike in the real economy, the

**Table 3.** Medium Skill Workers Remain on Disability in Period when Recovered from Injury during the Bad Labor Market (Core Experiment).

	(1)	(2)	(3)	(4)
DV: Stay on disability in period				
Recovered from injury (1 = yes, 0 = no)				
Bad labor market (1 = no medium skill jobs)	.197*** (.048)	.223*** (.049)	-.071 (.048)	-.046 (.040)
Post bad labor market (1 = yes)	-.030 (.028)	-.014 (.027)	-.075 (.064)	-.057 (.053)
Disability benefits for medium skill worker = 6		-.016 (.069)		-.003 (.075)
Disability benefits for medium skill worker = 7		-.102* (.054)		-.129** (.061)
Disability benefits for medium skill worker = 8		-.017 (.046)		-.042 (.050)
Disability benefits for medium skill worker = 9		-.022 (.065)		-.036 (.060)
Disability benefits for medium skill worker = 10		-.037 (.094)		-.034 (.090)
Disability benefits for medium skill worker = 11		-.056 (.098)		-.074 (.085)
Medium skill worker × bad labor market			.421*** (.076)	.424*** (.077)
Medium skill worker × post bad labor market			.069 (.069)	.064 (.071)
Constant	.030 (.022)	.058 (.066)	.034* (.019)	.081 (.060)
Player fixed effects?	Yes	Yes	Yes	Yes
Observations	485	485	485	485
Number of participants	72	72	72	72
R <sup>2</sup>	.147	.151	.258	.265

Note. The table reports estimates from fixed-effects models with robust standard errors in parentheses. The interaction terms in the models estimate simple effects, not interaction effects. (As worker skill does not vary within participant, the main effect of medium skill worker drops out of the fixed-effects models.)

\* $p < .1$ . \*\* $p < .05$ . \*\*\* $p < .01$ .

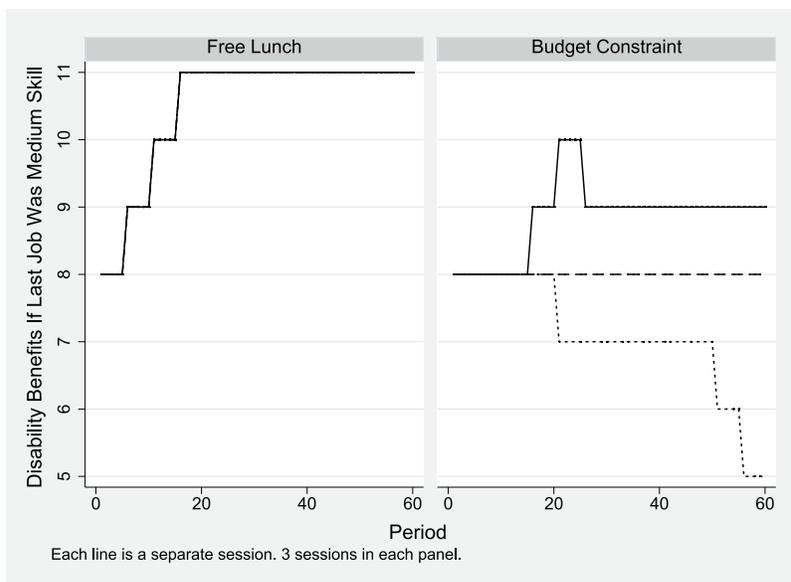
decision by a healthy worker to forgo working and instead remain on disability (which reduces tax revenue) does not affect any other worker's payoff. This means that there is no direct cost to any player of another player abusing benefit. Our design therefore "knocks out" a potential causal mechanism of direct self-interest whereby abuse by others reduces the availability of future benefits. This design therefore allows us to isolate the deservingness mechanism discussed above.

### Results: Observing Program Abuse Undercuts Program Support

First, confirming the results of the validation exercise, our manipulation of the bad labor market induces medium skill workers affected by that shock to abuse benefits. This analysis takes advantage of our within-session manipulation, which is the presence of the bad labor market in Periods 21–40. As a reminder, in these periods, there are no medium skill jobs available. As Table 3 shows (see Columns 3 and 4), and confirming the results of the baseline validation experiment, the bad labor market induces newly healthy workers to remain on disability and this effect is concentrated solely among medium skill workers (see Online Appendix Table A3, page A-27). Note that the results are unaffected by controlling for benefits levels, which may could induce post-treatment bias in this setting. See Online Appendix Table A4 (page A-28) for full regression estimates.

If the change in the labor market induces changes in the instances of individuals staying on disability, what effect does this have on support for generous disability benefits? Is this effect conditioned by whether providing benefits is costly? Does observing greater rates of individuals on disability or specific instances in which healthy individuals remain on disability reduce support for those benefits? We begin by examining support for benefits over the course of the game in both the Free Lunch and Budget Constraint conditions to assess the first order question of whether there are different dynamics in these two settings. This analysis reveals strong evidence that the potentiality or actuality of abuse only causes reductions of generous benefits when those benefits are costly to other players in the form of higher taxes. (We set aside for the moment the fact that there is within-session variation in abuse induced by the labor market manipulation.)

Figure 1 plots realized benefit levels, by period, for each of the three sessions in the Free Lunch (left panel) and Budget Constraint (right panel) conditions. In the Free Lunch condition, in all three sessions individuals vote to increase benefits whenever possible (in Periods 5, 10, and 15) and the benefits are maintained at their maximum for the remainder of the game. In the right panel, by contrast, we see a different pattern in the Budget Constraint condition. In one session (the top line), benefits are increased twice and then decreased back to nine tokens for a medium skill job, where they remain for the rest of the session. In one session (the middle line)



**Figure 1.** Realized benefit levels associated with a medium skill prior job, by session and period, in the “free lunch” (left panel) and “budget constraint” (right panel) conditions (core experiment).

benefits are never changed. Finally, in the third session (bottom line) benefits never increase and are eventually decreased to the minimum.

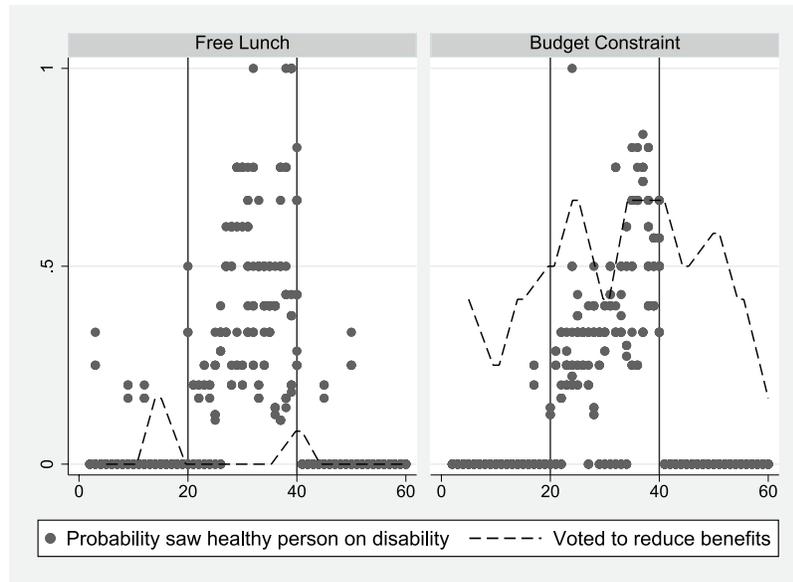
At the individual level we see similar patterns in voting. In the Free Lunch condition in cases in which benefits can be increased or decreased (that is, benefits are not already at the minimum or maximum), 3% of votes are to decrease benefits while 76% are to increase. By contrast, in the Budget Constraint condition, 39% of votes are to decrease compared with 38% to increase.<sup>16</sup> Regression analyses reported in the Online Appendix that model either a trichotomous vote variable ( $-1 = \text{decrease}$ ,  $0 = \text{no change}$ ,  $1 = \text{increase}$ ) or a binary indicator for voting to decrease benefits ( $1 = \text{yes}$ ,  $0 = \text{no}$ ) as a function of treatment condition and period of play fixed effects yield similar results (see Online Appendix Table A1, page A-25), as do analyses in which we partition by the current level of benefits (see Online Appendix Table A2, page A-26). In short, there is clear evidence that costliness of benefits is essential to induce a willingness to reduce (and not raise) benefit levels.

We next turn to examining how the prevalence of program abuse and the observation of that abuse explains support for benefit expansion and contraction. In Figure 2, we plot for each treatment average rates of voting to reduce benefits by low- and high-skilled workers (the dashed line; votes take place only every fifth round) and the chances that at the individual level these players were told a healthy (rather than injured) worker was on disability (the dots). As a reminder, we randomized at the player level the feedback they received about a worker on disability, so in periods and sessions where both healthy and non-healthy workers are on disability,

individual players could get different information.<sup>17</sup> We focus on low- and high-skill workers because medium skill workers have a different incentive during the bad labor market (Periods 21–40, between the vertical lines) to increase benefits because they cannot earn wages at their prior wage level.<sup>18</sup>

Focusing first on the Free Lunch condition (left panel), we see that the bad labor market increases substantially the chances that players are informed about a healthy worker on disability. Rates of voting to reduce benefits, however, are low and constant across the experiment. In the Budget Constraint condition (right panel), the bad labor market also increases the chance of seeing a healthy worker on disability. However, unlike in the Free Lunch condition, the bad labor market is also associated with an increased chance of voting to reduce benefits, and those rates decline again once the labor market recovers. (Parallel analysis for medium skill workers appears in Online Appendix Figure A1 [see page A-29] and does not follow this pattern.) Formal statistical analyses in which we use OLS regression to predict each player’s vote ( $1 = \text{vote to reduce}$ ,  $0 = \text{otherwise}$ ) as a function of player fixed effects and other covariates appear in Table 4 and confirms this inference.

Per Column 1, pooling across both conditions, non-medium skill workers are about 9 percentage points more likely ( $p < .05$ , two-tailed test) to vote to reduce benefit levels during the bad labor market, which confirms Hypothesis 1. This result holds controlling for whether a worker is injured (which is randomly assigned, Column 2), excluding cases where benefits cannot be further reduced (Column 3), and controlling for realized benefit levels (which is post-treatment, and may



**Figure 2.** Probabilities of observing another healthy worker on disability and average rates of voting to reduce benefits by low- and high-skill workers, by labor market condition and budget constraint (vs. free lunch) condition (core experiment).

**Table 4.** Non-Medium Skill Workers Vote to Reduce Benefit Levels During the Bad Labor Market (Core Experiment).

	(1)	(2)	(3)	(4)	(5)	(6)
	DV: Voted to reduce benefits (1 = yes, 0 = no)					
	Excluding medium skill workers	Excluding medium skill workers	Excluding cases where benefit levels cannot be decreased	Excluding cases where benefit levels cannot be decreased	Free lunch condition excluding cases where benefit levels cannot be decreased	Budget constraint condition excluding cases where benefit levels cannot be decreased
Bad labor market (1 = no medium skill jobs)	.094** (.039)	.087** (.039)	.089** (.036)	.107*** (.040)	-.021 (.029)	.185*** (.066)
Post bad labor market (1 = yes)	-.010 (.039)	-.012 (.038)	.022 (.037)	.038 (.040)	-.042 (.029)	.076 (.067)
Injured in period (1 = yes)		-.094*** (.033)	-.064** (.032)	-.061* (.032)	.024 (.025)	-.139** (.058)
Disability benefits for medium skill worker				-.024 (.022)		
Constant	.219*** (.028)	.258*** (.031)	.237*** (.029)	.444** (.196)	.033 (.022)	.452*** (.054)
Player fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes
Observations	288	288	284	284	144	140
Number of participants	24	24	24	24	12	12
R <sup>2</sup>	.032	.061	.042	.046	.023	.114

Note. The table reports estimates from OLS models with standard errors in parentheses. \* $p < .1$ . \*\* $p < .05$ . \*\*\* $p < .01$ .

therefore introduce bias, per Column 4). Partitioning by condition, in Columns 5 and 6, shows that there is no average effect of the bad labor market in the Free Lunch condition, whereas the bad labor market increases the willingness to vote to reduce benefits by 19 points ( $p < .01$ ) in the Budget Constraint condition. This confirms the aggregate graphical presentation in Figures 1 and 2, and Hypothesis 3.

In light of the finding that voting to reduce benefits takes place almost exclusively in the costly benefits condition and with greater regularity during the bad labor market, the

remainder of our analysis focuses on the interplay between the bad labor market, program abuse, and support for generous benefits only in the Budget Constraint (costly benefits) condition, with relevant parallel analysis from the Free Lunch condition discussed in footnotes. These analyses leverage the random assignment of worker feedback.<sup>19</sup> In particular, we examine for low- and high-skill workers whether observing a healthy (rather than injured) worker on disability can explain votes to reduce benefits. If free-riding concerns do in fact motivate the greater skepticism of

benefits in this condition, we can use this individual-level randomization to isolate the specific effect of learning about program abuse and test Hypothesis 2. Our analyses are summarized in Table 5. We begin in Column 1 with a simple OLS regression where we predict voting to reduce benefits (1 = yes) as a function of player fixed effects (which absorb all static player-level factors, including worker skill) and whether a worker observed a healthy (rather than injured) player on disability.<sup>20</sup>

The .24 ( $p < .01$ ) coefficient in Column 1 means that observing a player who is healthy receiving disability (rather than observing an injured player) increases the chances of voting to reduce disability by 24 points. Controlling for whether a player is personally injured reduces this effect slightly (Column 2), but excluding cases in which benefits cannot be further reduced has no effect (Column 3). In Column 4, we control for the (post-treatment) benefit levels for medium skill workers and this also does not materially affect our estimate.

Next, to rule out the possibility that time trends explain this effect, in Column 5 we estimate a model controlling for period and period squared. Above, we showed that the bad labor market increased the chances that healthy (medium skill) workers would stay on disability. Therefore not surprisingly controlling for time in this way reduces the effect of observing a healthy worker (coefficient =  $-.15$ ,  $SE = .09$ ,  $p < .10$ ), but it still remains marginally significant. Finally, we also examine a different measure of treatment assignment, which is the number of times in this and the prior four rounds a worker saw feedback about a healthy worker being on disability. This specification yields a similar result, but this model fits the data slightly worse than the Column 2 specification with the same covariates.

Taken together, these results confirm the hypothesis that learning of free riding by others by observing healthy workers on disability increases the propensity to vote to reduce benefits (Hypothesis 2). How large are these effects? When benefits are costly, 47% of votes by low- and high-skill workers are to reduce benefits. The .24 estimate from Column 1 means that observing a healthy worker on disability increases the predicted rate of voting to reduce benefits by a very large 52% ( $.24/.47$ ). Even the smallest estimate (from Column 5) still implies an increase of 31% in the rate of opposing benefit increases.

While we have only a single instrument that induces program abuse (the disappearance of medium skill jobs during the bad labor market), we can nonetheless undertake some preliminary analysis to partition how much of the effect of the bad labor market is attributable to the specific observation of program abuse, the causal pathway highlighted in Hypothesis 2, versus the inferences people may draw on the basis of the bad labor market causing an increase in the overall prevalence of people on disability. Empirically, the bad labor market induces a 1.2 unit increase ( $1.2/4.3 \approx 28\%$ ,  $p < .01$ ) in the number of people claiming disability from

a baseline of 4.3. But a change in the number of people on disability does not definitively reveal abuse because it does not allow a player to discern if this is due to chance or instead to someone taking advantage of the program. To distinguish the effect of observing program abuse (randomized at the player level) from the effect of inferring program abuse (where inferences are based on the size of disability rolls at any given point in time), we add to the model specification from Column 3 the number of people on disability in the round in which voting takes place. In this specification (Column 7), we find that it is the observation of abuse (coefficient =  $.20$ ,  $SE = .08$ ,  $p < .05$ ), rather than indirect signals about abuse (coefficient =  $.01$ ,  $SE = .02$ , *n.s.*), that predicts voting to reduce benefits. This test provides the strongest evidence in support of Hypothesis 2.

Finally, in Columns 8 and 9, we assess the robustness of this result to an even more conservative model specification. In particular, we examine specific session-period-level cases (e.g., Period 7 from our second session, accounted for with fixed effects) in which there are at least two workers who could receive different feedback about the health of workers receiving disability.<sup>21</sup> This matched pair design accounts for all factors up to that point that are shared across players (e.g., realized benefit levels, period of play). To account for player attributes, we include an indicator for whether a player is high- (rather than low-) skilled. Per Column 8, even in this highly restrictive definition of our sample (which discards a great deal of data), observing a healthy worker on disability increase the willingness to vote to reduce benefits, a finding that holds when we control for the probability of treatment (with inverse probability weights to account for varying probabilities of learning a worker receiving disability is healthy) in Column 9.

The estimates in these last two columns are statistically imprecise, unfortunately, but given the amount of data that is lost in constructing this comparison set, it is notable that we continue to see the same pattern as in the earlier specifications. Moreover, the baseline rates of voting to reduce benefits are much lower among cases included in this sample—in only 29% of cases do the player-period-level observations included in this analysis vote to reduce benefits. This means that the .11 coefficient from the Column 8 specification is similar in substantive magnitude to the result found in Column 5—observing a healthy worker on disability increases the chances of voting to reduce benefits by 36% ( $.11/.29$ ).

## Conclusion

Our experiments provide the first causal evidence about the effect of perceptions of program abuse on support for social insurance programs. We show that learning a particular beneficiary is taking benefits despite being able to work reduces overall program support, despite the fact that all workers benefit from social insurance when injured. Specifically,



when disability benefits are linked to taxes, participants who observe program abuse reduce their support for more generous benefits beyond the effects that are attributable to the number of people on disability. Among workers not affected by a labor market shock, the chances of voting to reduce benefits increases by 15 to 24 percentage points when they observe a healthy person on disability compared with someone who observes an injured person on disability.

In contrast to prior survey experimental work, a key advantage of our incentivized experimental design is that we gain tight control over participants' material incentives and the information they have about others' behaviors. At the same time, our experiments are not a perfect replication of the actual labor market environment or the external political environment surrounding support for disability insurance. In light of the difficulty of making strong causal inference in a more traditional observational setting, we view this tradeoff as desirable. Nonetheless, there are certain areas in which abstraction may mitigate or amplify the effects we detect. We consider several potential concerns about external validity here.

An important feature of our experimental design is that we randomly assign observations of abuse that affect program support. We acknowledge that the process of learning about program abuse may be different outside of the experimental setting. For example individuals may learn together from common media accounts or specific communities may experience salient instances of publicized abuse that change beliefs. While that specific learning process is different, the content of that information is very similar to what we provide in our setting, which is more amenable to clear causal inference. Similarly, factors like beneficiary race may shape interpretations of whether a particular case of program use is perceived as abusive or garners coverage in the first place. While prior survey experimental evidence finds little evidence that race shapes attitudes about particular beneficiaries, patterns of media coverage might be another potential cause of racialization over time.

Furthermore, in none of our designs is actual program abuse a threat to an individual's own access to benefits. If abuse did directly reduce the chances one could access benefits, it would likely magnify the effect of abuse on support for program retrenchment (Bokemper & DeScioli, 2018). By contrast, in our Budget Constraint condition, raising benefits also raises taxes, and abuse therefore represents simply others avoiding (higher) taxes and claiming (greater) benefits. Our choice to exclude such a direct self-interest mechanism was motivated by a desire to focus on how the psychology of deservingness affected program support, but such deservingness concerns would be intertwined with self-interest in a setting where cheating threatened benefits directly.

Although participants regularly voted to reduce benefits in the Budget Constraint condition, only 6% of votes in the Free Lunch condition were to reduce benefits. This result suggests that people are not insensitive to the costs of taking

an action to "punish" underserving behavior when they observe it. Participants in the Budget Constraint were in a cost-neutral environment in which a reduction of benefits was coupled with a reduction in taxes. In contrast, participants in the Free Lunch condition only imposed a cost on themselves (i.e., lost benefits when they were injured) by voting to reduce benefits. Even though free riding did not pose a direct cost on participants in either condition, it does appear that the cost of foregoing benefits in the Free Lunch condition deterred people from reducing benefits in response to observing norm violations. This result is somewhat inconsistent with prior work on punishment in which individuals are willing to bear a cost to punish free riding even though they will not receive a benefit from doing so (Fehr & Gächter, 2002). Understanding sensitivity to these costs in the willingness to punish is an important area for subsequent research.

One key question, broached somewhat in prior survey work, is how individuals form beliefs about larger program abuse. Our focus on the personal observation of program abuse is likely only part of how people form beliefs about program abuse. We do not consider the role of prior beliefs and orientations (e.g., partisanship) in explaining either perceptions of deservingness or program support. Our view is that while such differences are likely powerful sources of individuals' beliefs about disability insurance and government social policy more generally, understanding how individuals learn and react to stimuli in a dynamic environment is nonetheless very important. Tellingly, a great deal of political rhetoric surrounding debates about social policy retrenchment raises the possibility of program abuse, suggesting our environment captures the factors affecting individuals' willingness to support these programs after accounting for baseline individual-level differences in opinions.

It would be ideal to find observational settings in which exogenous sources of shifts in perceptions of abuse could be identified and their downstream consequences for program support measured. But this approach involves significant problems of measurement and causal inference. On the measurement side, it is often difficult to clearly quantify individuals' exposure to information (e.g., peer networks) or which claims of disability are actually valid. Isolating the causal role of those shocks is also difficult because we as analysts have a difficult time distinguishing the effect of changes in program abuse from factors that are correlated both with levels of program abuse and levels of program support. For example, macroeconomic shocks that drive certain workers to turn to disability insurance as a substitute for employment might also cause unaffected workers to become less willing to support social insurance, independent of any abuse. (The direction of bias could also move in the opposite direction, if unaffected workers in an area that experiences a labor market shock become more supportive of disability because they increase their assessment of the probability that

they might also need to turn to disability insurance.) This need to validate the exclusion restriction is an unfortunate and ubiquitous problem in observational studies.

Finally, our approach also provides a blueprint for studying the dynamics of public support and the political economy of other social insurance programs and distributive politics more generally. In a stylized and incentivized laboratory environment one can bring together features of real world economic behavior with the dynamics of politics. While we focus here on program abuse and support for disability insurance, there are numerous questions about the interplay between politics and economics that can be examined within this framework and our approach provides a general model for doing so.

### Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

### Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

### ORCID iDs

Albert H. Fang  <https://orcid.org/0000-0001-6265-5967>

Gregory A. Huber  <https://orcid.org/0000-0001-6804-8148>

### Supplemental Material

Supplemental material for this article is available online.

### Notes

1. Hiltzik (2015).
2. See, for example, McCoy (2017).
3. In this respect, the key element of the program is the discretionary determination of merit vis-à-vis programs like Social Security where more objective factors like age define eligibility.
4. We note more generally that almost all experimental work on deservingness has focused on whether specific beneficiaries warrant support, without considering the more politically consequential question of whether concerns about deservingness reduce overall program support.
5. See <https://www.ssa.gov/disabilityfacts/materials/pdf/factsheet.pdf>. Accessed 25 Jun 2018.
6. For example, it could be that individuals who dislike African Americans both rate them as less deserving and want to reduce support for welfare when a particular beneficiary is non-White, but that this effect arises because of antipathy not via reduced deservingness.
7. A key exception examines how a beneficiary's medical impairment affect perceived deservingness (e.g., Jensen & Petersen, 2017), but that research does not directly study attitudes about social insurance program size and funding.
8. <https://www.ssa.gov/disabilityfacts/materials/pdf/factsheet.pdf>
9. See, for example, Spirling (2017).
10. We refer to participants interchangeably as “participants,” “players,” “subjects,” and “workers” throughout the article, regardless of whether they are working in a given period in the game.
11. See Online Appendix A (page A-3) for the full text of the instructions and Online Appendix B (page A-16) for screenshots of the disability insurance game interface.
12. Because there is no real effort in this task, the relevant comparison is between expected tokens earned if one stays on disability versus expected tokens if one returns to work.
13. Medium skill workers earn more money in the short term by working a low skill job than they do continuing to collect disability benefits. If a worker cares only about their earnings, they should re-enter the labor market at the first available opportunity. Disability benefits are calculated from the wage of the last job, so re-entering the labor market at a low skill job exposes a worker to the risk of suffering an injury and consequently receiving lower disability benefits. However, a worker would have to believe that the duration of the bad labor market is longer than 18 periods for the payoff of staying on medium skill level disability to be greater than the payoff for returning to work at a low skill job, even with the risk of suffering a new injury.
14. As is common with university subjects pools, our participants are generally younger ( $M_{age} = 21.4$ ), more likely to be female (62%) and more left leaning (63% self-identified as liberal or very liberal; 59% self-identified as Democrats) than the general population. Because this profile means the sample is more likely to be generically supportive of government programs, we believe this sample would be less responsive to the effects of observing program abuse on program support than a more ideologically conservative sample, which suggests that any treatment effects we observe are biased downward (toward zero).
15. In periods in which workers are healthy, in 7% of cases healthy workers eschew working and instead remain on disability.
16. If we do not condition on benefit levels, 6% of cases in the Free Lunch condition are to reduce benefits and 20% are to increase, compared with 38% and 38%, respectively, in the Budget Constraint condition.
17. Note that this analysis therefore exploits our within-session manipulation across individuals of observed abuse. While one might be concerned about statistical power in across-session comparisons (because there are three session of each condition), this concern is largely obviated in analysis that compares across individuals within session because the treatment (observation of abuse) is randomized across players and periods within session.
18. When benefits were not at their maximum or minimum, 51% of votes by medium skill workers were to increase benefits, 27% of votes were to decrease benefits, and 22% of votes were to keep benefits the same.
19. We conducted a parallel analysis for the Free Lunch condition and we do not find a positive effect of seeing a healthy worker on disability on voting to decrease disability benefits (see Online Appendix Table A5, page A-28).
20. There are no cases in which all players are working.
21. This restriction eliminates cases where there is no variation in the health status of workers (other than the worker herself) about whom a participant could receive feedback, and holds constant the total number of people on disability, an indirect signal about abuse which we analyzed in Table 5, Column 7.

## References

- Aarøe, L., & Petersen, M. B. (2014). Crowding out culture. *The Journal of Politics*, 76(3), 684–697.
- Ahlquist, J. S., Hamman, J. R., & Jones, B. M. (2017). Dependency status and demand for social insurance. *Political Science Research and Methods*, 5(1), 31–53.
- Autor, D. H. (2015). The unsustainable rise of the disability rolls in the united states. In J. K. Scholz, H. Moon, & S.-H. Lee (Eds.), *Social policies in an age of austerity* (pp. 107–136). Edward Elgar.
- Autor, D. H., & Duggan, M. G. (2003). The rise in the disability rolls and the decline in unemployment. *The Quarterly Journal of Economics*, 118(1), 157–206.
- Axelrod, R. (1984). *The evolution of cooperation*. Basic Books.
- Barber, B., Beramendi, P., & Wibbels, E. (2013). The behavioral foundations of social politics. *Comparative Political Studies*, 46(10), 1155–1189.
- Black, D., Daniel, K., & Sanders, S. (2002). The impact of economic conditions on participation in disability programs: Evidence from the coal boom and bust. *American Economic Review*, 92(1), 27–50.
- Bokemper, S. E., & DeScioli, P. (2018). *The problem of dishonesty in government safety nets* [Working paper]. <http://www.sbokemper.com/papers/bokemperdescioli.pdf>
- Congressional Budget Office. (2016, June 15). *Social security disability insurance: Participation and spending*. <https://www.cbo.gov/publication/51443>
- Cook, F. L., Barabas, J., & Page, B. I. (2002). Invoking public opinion: Policy elites and social security. *Public Opinion Quarterly*, 66(2), 235–264.
- Cook, F. L., & Barrett, E. J. (1992). *Support for the American welfare state*. Columbia University Press.
- Croson, R., & Marks, M. (1998). Identifiability of individual contributions in a threshold public goods experiment. *Journal of Mathematical Psychology*, 42(2–3), 167–190.
- Delton, A. W., Cosmides, L., Guemo, M., Robertson, T. E., & Tooby, J. (2012). The psychosemantics of free riding. *Journal of Personality and Social Psychology*, 102(6), 1252–1270.
- Fang, A. H., & Huber, G. A. (2019). Perceptions of deservingness and the politicization of social insurance: Evidence from disability insurance in the United States. *American Politics Research*. Advance online publication. <https://doi.org/10.1177/1532673X19863439>
- Fehr, E., & Gächter, S. (2002). Altruistic punishment in humans. *Nature*, 415(6868), Article 137.
- Gilens, M. (1999). *Why Americans hate welfare*. University of Chicago Press.
- Hiltzik, M. (2015, January 19). Rand Paul steps up the GOP attack on social security. *Los Angeles Times*. <https://www.latimes.com/business/hiltzik/la-fi-mh-rand-pauls-supremely-cynical-attack-on-social-security-20150119-column.html>
- Jensen, C., & Petersen, M. B. (2017). The deservingness heuristic and the politics of health care. *American Journal of Political Science*, 61(1), 68–83.
- McCoy, T. (2017, March 30). *Disabled, or just desperate? Rural Americans turn to disability as jobs dry up*. The Washington Post. <https://www.washingtonpost.com/sf/local/2017/03/30/disabled-or-just-desperate/>
- Morton, W. R., & Huston, B. F. (2018). *Social security: What would happen if the trust funds ran out?* Congressional Research Service.
- Olson, M. (1965). *The logic of collective action*. Harvard University Press.
- Page, B. I., & Shapiro, R. Y. (1992). *The rational public: Fifty years of trends in Americans' policy preferences*. University of Chicago Press.
- Petersen, M. B. (2012). Social welfare as small-scale help: Evolutionary psychology and the deservingness heuristic. *American Journal of Political Science*, 56(1), 1–16.
- Petersen, M. B., Slothuus, R., Stubager, R., & Togeby, L. (2011). Deservingness versus values in public opinion on welfare. *European Journal of Political Research*, 50(1), 24–52.
- Sell, J., & Wilson, R. K. (1991). Levels of information and contributions to public goods. *Social Forces*, 70(1), 107–124.
- Spirling, G. B. (2017, May 23). The fuzzy claims used to justify cutting social security disability insurance. *The Atlantic*. <https://www.theatlantic.com/business/archive/2017/05/ssdi/527802/>
- van Oorschot, W. (2000). Who should get what, and why? On deservingness criteria and the conditionality of solidarity among the public. *Policy & Politics*, 28(1), 33–48.
- Woon, J. (2012). Democratic accountability and retrospective voting: A laboratory experiment. *American Journal of Political Science*, 56(4), 913–935.

## Author Biographies

**Scott E. Bokemper** is a postdoctoral associate at the Center for the Study of American Politics at Yale University and holds a PhD from Stony Brook University. His work focuses on the interaction between the psychology of cooperation and institutions.

**Albert H. Fang** is an independent scholar who conducts research in behavioral and psychological science, with a focus on application in politics and markets. He was previously a postdoctoral associate at the Center for the Study of American Politics at Yale University and holds a PhD from Columbia University.

**Gregory A. Huber** is a Forst Family professor and chair of Political Science at Yale University, where he is also the associate director of the Center for the Study of American Politics and the Director of the ISPS Behavioral Research Lab. His work focuses on behavioral and institutional questions in the field of American Politics.