

Perceptions of Program Abuse and Support for Social Insurance*

Scott E. Bokemper
Yale University

Albert H. Fang
Yale University

Gregory A. Huber
Yale University

July 4, 2018

Abstract

Do perceptions of abuse in public social insurance programs undercut program support? Answering this causal question is difficult because perceptions of program abuse can arise from multiple potential causes including prior opposition to the program. Examining the case of disability insurance, we circumvent these challenges using multiple laboratory experiments involving a novel simulated political economy to study the interplay between labor market shocks, program abuse, perceptions of abuse, and preferences for benefit levels. We find that negative labor market shocks that preclude injured workers from returning to work at their pre-injury wage upon recovery increases the probability of staying on disability instead of working at a lower-wage job despite being healthy. Further, when benefits are costly, learning about program abuse causes workers unaffected by labor market shocks to prefer lower benefit levels. Our results demonstrate an important channel by which shocks to market employment diminish support for government social insurance.

Keywords: political economy of social insurance; program abuse perceptions; support for social insurance; labor market shocks; free-riding; laboratory experiment

Abstract Word Count: 150

Article Word Count: 11,692

*Scott E. Bokemper is a Postdoctoral Associate at the Center for the Study of American Politics and the Institution for Social and Policy Studies at Yale University (sbokemp1@gmail.com; <http://www.sbokemper.com/>). Albert H. Fang is a Postdoctoral Associate at the Center for the Study of American Politics and the Institution for Social and Policy Studies at Yale University (albert.fang@yale.edu; <http://albertfang.com>). Gregory A. Huber is Forst Family Professor of Political Science at Yale University (gregory.huber@yale.edu; <https://huber.research.yale.edu>).

“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, 2015¹

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.² 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 levels of program-specific support?

Against the broader backdrop of austerity politics, these questions have grown in political 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 the individual risk associated with adverse events that prevent employment. Because program beneficiaries subsidize their own care and benefits are limited to defined social risks, these risk pooling programs are often popular, especially vis-à-vis more traditional redistributive and 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 (in the case of disability insurance, a long-term injury that prevents work). This abuse may in turn undercut support for social insurance by leading citizens to believe beneficiaries are not deserving of assistance.³

Recent changes in the market economy may have contributed to growing concerns about disability insurance program abuse. In the United States, disability insurance rolls have

¹Hiltzik (2015).

²See, for example, Joffe-Walt (2013*a,b,c*), McCoy (2017*b*), and Hiltzik (2013).

³In 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.

grown substantially over the last two decades, with a large uptick among working-age men and for certain “hard to diagnose” medical conditions. Scholars have hypothesized that concentrated economic declines in certain geographic areas and for certain occupations have created situations in which some formerly working individuals can no longer find similarly remunerative employment. If these individuals who could hold a job nonetheless obtain disability insurance benefits, it may depress support for the program as a whole.

Empirically, support for social insurance programs in the contemporary United States and elsewhere is far from universal, and survey evidence shows that individuals are concerned about the deservingness of disability beneficiaries in general (Fang and Huber 2018). Past research also shows that the usage of disability increases in areas when economic conditions deteriorate (Autor and Duggan 2003; Autor, Dorn and Hanson 2013; Black, Daniel and Sanders 2002; Charles, Li and Stephens Jr. 2017). Further, a handful of survey studies have examined citizens’ willingness to support individuals with health disorders.

However, no research to date has examined the interplay between economic conditions, program abuse, and public support for disability insurance. These causal relationships are difficult to uncover with the analysis of administrative data and surveys for several reasons. Perhaps the biggest challenge for observational analyses is that perceptions of disability insurance program abuse originate from many sources. For example, individuals may report perceptions of program abuse not because they have observed abuse or believe it occurs regularly, but because they are already opposed to the program. Similarly, we do not know about individuals’ exposure to actual disability insurance beneficiaries and whether those beneficiaries are abusing the program. Linking economic shocks to program abuse and attitudes is further complicated because it is often not clear how a particular economic event affects workers across skill levels in the economy or how to isolate the effect of shocks on disability insurance take-up from the direct effect of shocks on program support.

To address these limitations and to causally identify the effects of perceiving program abuse on support for disability insurance, we designed and implemented a series of laboratory

experiments where subjects participate in a novel disability insurance game. In this game, participants take the role of workers in a laboratory economy and their choices affect their monetary compensation. In this way, our game allows us to use incentives to create baseline preferences over outcomes, providing a degree of control not feasible outside the lab setting (see, e.g., Woon 2012; Duch, Przepiorka and Stevenson 2015; Hamman, Weber and Woon 2011; Smith 2006). Workers face a risk of suffering an injury that prevents them from working and automatically qualify for disability benefits if they become injured. Subsequently, workers may recover from an injury and, if they do so, choose to continue to collect benefits or re-enter the labor market. In a baseline experiment, we examine how a negative labor market shock affects individuals' usage of the disability insurance program. In particular, we test whether introducing a "bad" labor market causes workers who recover from an injury but who cannot find employment at their past wage level to continue to collect disability benefits instead of returning to work at a new lower wage. We find that healthy workers who are unable to return to a job at their prior wage remain on disability about 13% of the time.

In a second experiment, we build on this baseline experiment to understand whether abuse of the disability insurance program depresses support for program benefits. Participants vote regularly on benefit levels after receiving feedback about the health status of another randomly selected participant receiving disability benefits. In this experiment, we also vary whether higher benefit levels impose greater costs on workers active in the labor market (i.e., taxes). We find that when disability benefits are costly for other workers, observing program abuse (which increases when workers cannot find a job at their prior wage) reduces support for more generous disability benefits.

These experiments help to understand the nature of public support for disability insurance and also reveal how shocks to market employment can induce diminished support for disability insurance through increased use of disability as a substitute for reentering a less rewarding labor market. In contrast to survey evidence or analysis of aggregate trends in program use,

our experimental results allow us to isolate the effect of labor market conditions on worker behavior and downstream support for disability. Importantly, our incentivized experiment accounts for several important features of real labor markets and disability programs. These results therefore show why workers who themselves might benefit from disability programs nonetheless reduce the generosity of these programs in light of program abuse. Moreover, these results demonstrate how a single observation of abuse can cause citizens to reduce support for disability insurance, even when most beneficiaries are not taking advantage of the program.

Our work relates to a broader survey literature showing an association between perceptions of beneficiary deservingness and support for welfare benefits provision (Aarøe and Petersen 2014; Appelbaum 2001; Gilens 1999; Petersen 2012; Petersen et al. 2011; van Oorschot 2000). That work, however, cannot provide causal evidence that variation in program abuse diminishes support for redistribution. Instead, it often finds that individual-level characteristics, like ideology, explain program support, or that features of a particular beneficiary affect people’s perceptions of that beneficiary’s deservingness. Our work supplements that prior work by isolating the role of deservingness, manipulated via exogenous change in the labor market, in shaping program support.

Additionally, our experiments contribute to a broad literature in experimental economics, social psychology, and political science on how the possibility of free-riding may upset cooperation in public goods provision (Ledyard 1995; Sell and Wilson 1991). Prior laboratory experiments in this literature have examined how risk exposure affects support for social insurance, how variation in program design affects program abuse, and how observing free riding affects contributions in public goods games. It is unclear from existing work, however, whether the effect of free-riding extends beyond the scope of bounded public goods decisions to explain political support for social insurance programs. Nor has extant work combined the dynamic interplay of economic conditions, program abuse, and program support. This study addresses the need to assess the joint effects of economic shocks and learning about

program abuse on political support for social insurance in a unified experimental setting. Our work reveals that program abuse in social insurance can be induced in an incentivized political economy, and that observing actual abuse reduces program support when providing benefits is costly for others. We discuss the broader implications of this result and our work more generally in the conclusion.

1 Background: The Economy and Social Security Disability Insurance

In the United States, Social Security Disability Insurance (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.⁴ Once an individual qualifies for SSDI benefits, they receive cash payments until they recover, return to work, reach retirement age, or die (Bound and Burkhauser 1999).

Over the past 45 years, the number of SSDI beneficiaries has nearly quadrupled and the percentage of the working-age (20 to 66) 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 and Duggan 2003; Autor, Dorn and Hanson 2013; Black, Daniel and Sanders 2002; Charles, Li and Stephens Jr. 2017). Using merged administrative and economic data, Autor and Duggan (2003) find that counties that experi-

⁴<https://www.ssa.gov/disabilityfacts/materials/pdf/factsheet.pdf>

enced 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, Daniel and Sanders 2002). Broadly, this research suggests that disability benefits can serve as a substitute for labor when economic conditions deteriorate⁵, 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.)

The rapid growth of SSDI has led to greater political conflict over the program vis-à-vis many other social insurance programs. Many social insurance programs enjoy broad political support because they insure people against a number of common risks and have non-discretionary eligibility determinations (Page and Shapiro 1992; Cook, Barabas and Page 2002). The SSDI program, by contrast, is subject to growing elite criticism and calls for program reform.⁶ 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 and Huston 2018). While this move temporarily averted the program’s fiscal collapse, experts suggest the program will likely need to be reformed with stricter eligibility criteria and improved incentives to return to work to avoid continuing financial instability (Autor 2015).

2 Program Abuse, Deservingness, and Support for Social Insurance

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

⁵This work is also consistent with recent popular media coverage (see, e.g., McCoy 2017*a,b*).

⁶See, for example, Spiraling (2017).

“fraction of one percent,”⁷ 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 (2018) 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. More generally, social programs that citizens believe have greater rates of fraud and abuse receive less public support (Cook and Barrett 1992).

Beliefs about fraud and abuse in social insurance may be reinforced by selective media accounts of individuals defrauding the program. For example, in 2013 a Minnesota man was convicted of defrauding the SSDI program after pretending to have dementia in order to receive over \$140,000 in disability benefits.⁸ While these sorts of individuals are not typical disability beneficiaries, selective exposure to such accounts of program abuse may have outsized effects on attitudes about a program.

Program abuse is one of many potential factors shaping public perceptions of beneficiary deservingness. A rich observational and experimental literature has examined the antecedents of perceptions of beneficiary deservingness (Aarøe and Petersen 2014; Appelbaum 2001; Gilens 1999; Jensen and Petersen 2017; Petersen 2012; 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 and Petersen 2014; Petersen 2012). That is, citizens do not want to give benefits to someone who could work instead of continuing to collect benefits. Focusing more specifically on disability insurance, analyzing data from a survey experiment conducted on the 2016 CCES, Fang and Huber (2018) 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 this pattern arises because harder to diagnose

⁷<https://www.ssa.gov/disabilityfacts/materials/pdf/factsheet.pdf>. Accessed 25 Jun 2018.

⁸Browning (2013).

impairments are subject to greater abuse by otherwise healthy individuals.

We expect that citizens, in turn, may be less willing to support social insurance benefits for those who are perceived as being less in need of help. Deservingness perceptions matter because lack of deservingness implies free riding, which upsets the reciprocity necessary to sustain the provision of collective goods like public social insurance. While social insurance differs from public goods in several important ways (as discussed previously), a core challenge facing the provision of both public goods and social insurance is the free-rider problem: People have an incentive not contribute to fund the collective good and instead may benefit from the contributions of others (Olson 1965). Recent experiments on free riding in social psychology have shown that individuals are able to detect free-riders. Delton et al. (2012), for example, found that people are more likely to categorize someone as a free rider if they intentionally draw a benefit from a collective effort without contributing or if they purposefully withhold their contribution to a public good. Perceiving other public good beneficiaries as free-riders can disrupt the reciprocity required to sustain the provision of public goods because people are willing to give help to those who are likely to return help in the future, but not to those who would not (Axelrod 1984; Nowak 2006; Trivers 1971). That is, individuals do not want to pay a cost to provide a benefit to someone else when that person is unlikely to do so when the roles are reversed. Consequently, perceiving free-riding by others may lead individuals to want to punish free-riders either by decreasing their support for benefit levels or by abandoning the public good entirely, even if it imposes additional costs on others or on oneself.

3 Economic Experiments on Social Insurance Abuse

Past research has used incentivized experiments 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 (Barber, Beramendi and Wibbels 2013; Esarey,

Salmon and Barrilleaux 2012) or whether social insurance should be provided (Ahlquist, Hamman and Jones 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, Beramendi and Wibbels (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 to 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 program abuse affected program support.

Although past research 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 provided at the group level (Sell and Wilson 1991). However, other research has found that

the effect of individual information is only present when contributions are identifiable (i.e., are linked to participant-specific identifier), but not when they are anonymous (Croson and Marks 1998). This work shows that the potential for unobservable 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.

4 Disability Insurance Game

Building on this work, we designed a series of incentivized laboratory experiments to examine the interaction between economic conditions, the abuse of disability insurance, and public support. At the core of these experiments is a disability insurance game we designed to mirror key features of real world disability insurance programs like SSDI.

In the disability insurance game, we provide participants with group-level information about the number of other participants utilizing a laboratory disability insurance program. Further, we also show participants information about one randomly selected individual who is drawing disability benefits. Unlike the past work on information and public goods, this individual is someone who is not contributing to fund the disability program, but 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). Further, this recovered individual no longer needs the benefit they are taking, which is a core criterion for deservingness (van Oorschot 2000). Thus, our design manipulates the perception of deservingness that is directly attributable to the utilization of disability insurance when one does not actually need the benefits. 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. Moreover, the design of the game allows us to capture meaningful behavioral choices and preference expressions that map onto substantive quantities of interest. For clarity, we provide an overview of the baseline disability insurance game here, and describe relevant modifications to the design for each experiment when discussing each set of analyses.

Participants in the game take the role of workers in a laboratory economy.⁹ In this economy, workers have a low skill, medium skill, or high skill level that qualifies them to hold certain jobs.¹⁰ 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 4 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 4 tokens of effort and can look for a job. Importantly, a worker who has recovered from an injury can continue to collect disability benefits and also receive 4 additional tokens despite not exerting effort by working.

⁹We refer to participants interchangeably as “participants,” “players,” “subjects,” and “workers” throughout the paper, regardless of whether they are working in a given period in the game.

¹⁰See Online Appendix A for the full text of the instructions and Online Appendix B for screenshots of the disability insurance game interface.

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

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 does not have a job at their skill level decides whether to take a new job. 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.

5 Quantities of Interest and Main Hypotheses

We use the disability insurance game to test three hypotheses about the political economy of disability insurance. First, we attempt to corroborate the effect of deteriorating economic conditions on the take-up and abuse of disability insurance. Prior work shows workers may grow accustomed to earning a certain wage and will be less willing to re-enter the labor at a lower wage (Clark 1999; Clark and Oswald 1996). Consequently, workers may choose to stay on disability until they can return to work at their previous wage level. To test this claim, 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. Thus, in the first experiment, we test the following hypothesis:

- **H1:** *The bad labor market will cause medium skill workers to be more likely to continue claiming disability benefits compared to when they can re-enter the workforce at a medium skill job.*

Importantly, 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.¹¹

¹¹Disability 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

Finding support for this hypothesis would corroborate past observational research about how a change in the labor market affects the use of disability insurance.

In a follow-on experiment where participants also regularly vote on disability insurance benefit levels, we investigate the political consequences of healthy workers continuing to claim disability. Specifically, we examine whether workers who are unaffected by the labor market shock will be more likely to reduce their support for disability insurance, operationalized as voting to reduce benefits, when others workers react to the labor market shock by choosing to forego work. Further, we test whether this observation of free riding leads to a reduction in support when benefits are not costly to the unaffected workers.

One argument about the underpinning of the deservingness heuristic is that individuals are upset by those who take advantage of support when it is not needed because it imposes unnecessary costs on others. This free-riding can therefore upset the delicate balance by which individuals bear costs to help others who are unlucky because it makes those who bear contemporary costs feel that others are exploiting their generosity (Delton et al. 2012; Petersen 2012). However, it is possible that the observation of free riding provokes a negative response even when benefits are not costly to unaffected workers. Thus, we view any evidence that participants react negatively to perceived abuse of disability insurance in this context as indicative of a direct negative response to learning someone receives benefits when they do not need them, apart from the cost they impose on others. We therefore test two additional hypotheses:

- **H2:** *Unaffected workers will be more likely to vote to reduce benefits when they observe a healthy worker claiming disability relative to when they do not.*
- **H3:** *When free riding is observed, unaffected workers will be more likely to vote to reduce benefits when they are costly compared to when they are not costly.*

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.

6 Experiment 1: Labor Market Shocks and Program Abuse Behavior

6.1 Participants and Procedure

We conducted 3 sessions in an incentivized experimental laboratory to examine how changes in the labor market affect disability insurance claims. 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 subject pool for an experimental lab and only participated in a single session of the experiment. 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 1 token. The average earnings for participants not including the show up payment were \$18.69 (S.D. = \$5.99) for a task that took about 1 hour.

6.2 Design

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.

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 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 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.

6.3 Results

We begin our analysis by examining worker labor market behavior, focusing specifically on the decision players make about whether or not to return to work after recovering from an injury. Initially we examine whether our labor market manipulation induced health workers to continue collecting disability benefits.

Averaging across all players, workers are injured (not healthy) and therefore receiving benefits in 35% (763 out of 2160) of the periods of play. Additionally, healthy workers are on disability 4.5% of the time.¹² 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 2 tokens relative to remaining on disability (net of 12 tokens). Did this intervention cause medium skill workers to forgo

¹²In periods in which workers are healthy, in 7% of cases healthy workers eschew working and instead remain on disability.

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 never 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 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 support for hypothesis 1 and demonstrate that we can create a simulated labor market that induces behavioral patterns similar to those that may explain the rise of disability claims in the real economy. Workers who are made to become accustomed

Table 2: Experiencing a Bad Labor Market Increases the Probability Medium Skill Workers Claim Disability Despite Being Healthy (Experiment 1)

	(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)	0.149*** (0.015)	0.218*** (0.044)	0.022 (0.076)
Post Bad Labor Market (1=Yes)	-0.002 (0.015)	-0.011 (0.042)	-0.015 (0.069)
Medium Skill Worker x Bad Labor Market			0.286*** (0.092)
Medium Skill Worker x Post Bad Labor Market			0.003 (0.085)
Constant	0.014 (0.011)	0.034 (0.031)	0.033 (0.030)
Player Fixed Effects?	Y	Y	Y
Observations	1,397	251	251
Number of Participants	36	36	36
R-squared	0.095	0.144	0.194

***p<0.01, **p<0.05, *p<0.1. 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. (Since worker skill does not vary within participant, the main effect of medium skill worker drops out of the fixed-effects models.)

to jobs earning a particular wage react to the disappearance of those jobs by remaining on disability, despite the fact that doing so means earning less in the short run. This effect exists despite there being no indignity for accepting a lower skill job in a simulated economy, suggesting the effect we observe may be a lower bound.

This pattern sets the stage for our primary interest in 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 experimental conditions.

7 Experiment 2: Observing Program Abuse and Political Support for Social Insurance

In Experiment 2, we turn our focus to the political consequences of observing abuse in a disability insurance program. We test whether people will vote to reduce benefits when they observe a healthy worker collecting disability compared to when they observe an injured worker collecting disability. Further still, we examine whether we observe this effect when workers do not bear the cost of funding benefits.

7.1 Participants and Procedure

We conducted 6 sessions to examine how healthy workers staying on disability affects public support for disability insurance benefits. We recruited participants from the same laboratory subject pool to participate in an experiment (subjects could not participate if they participated in the first experiment). As with Experiment 1, 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 (S.D. = \$5.71) for a task that took about 1 hour.

7.2 Design

We added a voting component to the disability insurance game from Experiment 1. Workers voted on benefit levels every 5 periods (e.g., in periods 5, 10, ..., 60) and could vote to increase, decrease, or leave benefits the same for disabled workers. A vote to increase (decrease) benefits raised (diminished) 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 7 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.

We implemented two different voting conditions in which participants voted on benefit levels of disability insurance. In the “Free Lunch” condition ($N=36$; 3 sessions), benefit levels had no effect on taxes. 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. By contrast, in the “Budget Constraint” condition ($N=36$; 3 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 both conditions, we implemented the same within-subjects labor market manipulation and kept the same distribution of worker skill level from Experiment 1.

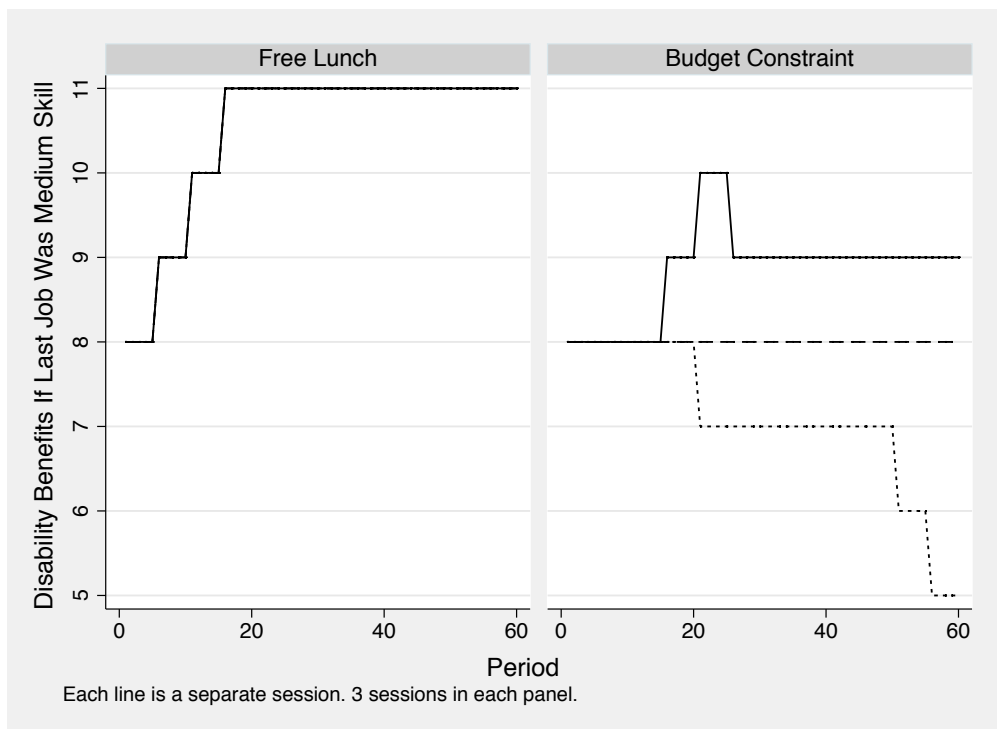
In the Free Lunch condition, there is no cost for any worker to enhancing the generosity of her own benefits because workers do not experience the cost of greater generosity. This is true both because workers do not have to pay for greater benefits for themselves, but also because we do not impose any cost of non-work by others on those who do work. That is, unlike in the real economy, the 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.

7.3 Results

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) condition and confirms our expectations

from hypotheses 3 about average benefit levels. 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 9 tokens for a medium skill job where they remain for the rest of the session. In one session (the middle line) benefits are never changed. Finally, in the third session (bottom line) benefits never increase and are eventually decreased to the minimum.

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 (Experiment 2)



At the individual level we see similar patterns. 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 to 38%

to increase.¹³ Regression analyses reported in the appendix that model either a trichotomous vote variable (-1=decrease, 0=no change, or 1=increase) or a binary indicator for voting to decrease benefits (1=yes, 0=no) as a function of treatment condition and period of play fixed effects yield similar results (see Appendix Table A1), as do analyses in which we partition by the current level of benefits (see Appendix Table A2).

This analysis exploits our between-session manipulation of the costliness of benefits. It does not, however, take 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 in Experiment 1, we find that this manipulation induces players who are newly healthy (i.e., recover from an injury in the period) to remain on disability. As column 1 of Table 3 shows, across all players, the bad labor market induces a 20 percentage point increase in the probability that newly recovered workers remain on disability compared to the pre-bad labor market rate of 3% (the effect of the post-bad labor market is statistically insignificant and predicts a 3 percentage point decrease from the pre-bad labor market rate). Column 2 adds benefit level fixed effects. While doing so could introduce post-treatment bias because benefit levels are endogenous, we find a qualitatively similar effect of 22 percentage points. Furthermore, this labor market effect is concentrated solely among medium skill workers. Low skill workers are no more likely to remain on benefits in the bad labor market condition and high skill workers never remain on disability (see Appendix Table A3). However, as column 3 of Table 3 shows, for medium skill workers, the bad labor market is associated with an approximately 35 percentage point increase in the chances of remaining on disability, whereas when the labor market recovers rates of staying on disability stays approximately the same as the pre-bad labor market. We also find a similarly sized increase of 38 percentage points when we condition on benefit levels (column

¹³If 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 to 38% and 38% respectively in the Budget Constraint condition.

4).¹⁴

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? Does observing greater rates of individuals on disability or specific instances in which healthy individuals remain on disability reduce support for those benefits? To assess hypothesis 2, we begin in Figure 2 by plotting 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. 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. 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.¹⁵

¹⁴In every session of the Free Lunch condition, benefits were at their maximum during the bad labor market, so we cannot estimate an effect of benefit levels on a medium skill worker's decision to claim benefits in the period they recovered from an injury. For the Budget Constraint condition, we estimate a bivariate regression predicting whether a medium skill worker stays on disability in the period they recovered from injury as a function of medium skill benefit levels. We find that a one token increase in disability benefits predicts a 13 percentage point increase in probability that a medium skill worker stays on disability ($t = 1.88, p < .10$). This effect increases slightly to 14 percentage points when controls are added for period and period squared ($t = 1.90, p < .10$). Because benefit levels are set endogenously, these estimates may suffer from post-treatment bias. See Appendix Table A4 for full regression estimates.

¹⁵When 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.

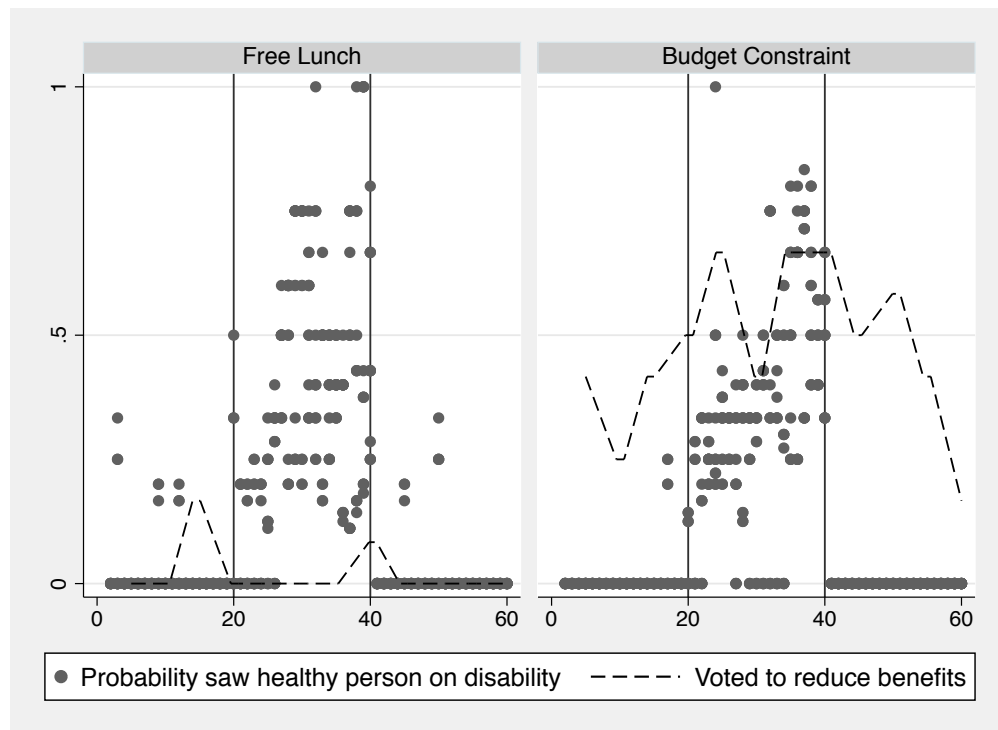
Table 3: Medium Skill Workers Remain on Disability in Period when Recovered from Injury during the Bad Labor Market (Experiment 2)

	(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)	0.197*** (0.048)	0.223*** (0.049)	-0.071 (0.048)	-0.046 (0.040)
Post Bad Labor Market (1=Yes)	-0.030 (0.028)	-0.014 (0.027)	-0.075 (0.064)	-0.057 (0.053)
Disability Benefits for Medium Skill Worker = 6		-0.016 (0.069)		-0.003 (0.075)
Disability Benefits for Medium Skill Worker = 7		-0.102* (0.054)		-0.129** (0.061)
Disability Benefits for Medium Skill Worker = 8		-0.017 (0.046)		-0.042 (0.050)
Disability Benefits for Medium Skill Worker = 9		-0.022 (0.065)		-0.036 (0.060)
Disability Benefits for Medium Skill Worker = 10		-0.037 (0.094)		-0.034 (0.090)
Disability Benefits for Medium Skill Worker = 11		-0.056 (0.098)		-0.074 (0.085)
Medium Skill Worker x Bad Labor Market			0.421*** (0.076)	0.424*** (0.077)
Medium Skill Worker x Post Bad Labor Market			0.069 (0.069)	0.064 (0.071)
Constant	0.030 (0.022)	0.058 (0.066)	0.034* (0.019)	0.081 (0.060)
Player Fixed Effects?	Y	Y	Y	Y
Observations	485	485	485	485
Number of Participants	72	72	72	72
R-squared	0.147	0.151	0.258	0.265

***p<0.01, **p<0.05, *p<0.1. 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. (Since worker skill does not vary within participant, the main effect of medium skill worker drops out of the fixed-effects models.)

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 very low. In the right panel where we display data from the Budget Constraint condition, the bad labor market also increases in the chances 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 Appendix Figure A1 and does not follow this pattern.) Formal statistical analyses in which we use OLS regression to predict each player's vote (1=vote to reduce, 0=otherwise) as a function of player fixed effects and other covariates

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 (versus Free Lunch) Condition (Experiment 2)



appear in Table 4 and confirms this inference.

Table 4: Non-Medium Skill Workers Vote to Reduce Benefit Levels During the Bad Labor Market (Experiment 2)

	(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	Budget Constraint Condition
Bad Labor Market (1=No Medium Skill Jobs)	0.094** (0.039)	0.087** (0.039)	0.089** (0.036)	0.107*** (0.040)	-0.021 (0.029)	0.185*** (0.066)
Post Bad Labor Market (1=Yes)	-0.010 (0.039)	-0.012 (0.038)	0.022 (0.037)	0.038 (0.040)	-0.042 (0.029)	0.076 (0.067)
Injured in Period (1=Yes)		-0.094*** (0.033)	-0.064** (0.032)	-0.061* (0.032)	0.024 (0.025)	-0.139** (0.058)
Disability Benefits for Medium Skill Worker				-0.024 (0.022)		
Constant	0.219*** (0.028)	0.258*** (0.031)	0.237*** (0.029)	0.444** (0.196)	0.033 (0.022)	0.452*** (0.054)
Observations	288	288	284	284	144	140
Number of Participants	24	24	24	24	12	12
R-squared	0.032	0.061	0.042	0.046	0.023	0.114

***p<0.01, **p<0.05, *p<0.1. The table reports estimates from OLS models with standard errors in parentheses.

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. 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 therefore introduce bias, per column 4). Partitioning by condition, in columns 5 and 6, however, shows that there is no average effect of the bad labor market in the Free Lunch condition, whereas the bad labor market increase the willingness to vote to reduce benefits by 19 points ($p < .10$) in the Budget Constraint condition. This provides further support for 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, we turn to our final analysis. This analysis focuses on the Budget Constraint condition and leverages the random assignment of worker feedback.¹⁶ 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 analysis appears 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.¹⁷

The .24 ($p < .01$) coefficient in column 1 means that observing a player who is healthy receiving disability increases the chances of voting to reduce disability by 24 points. Con-

¹⁶We 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 Appendix Table A5).

¹⁷There are no cases in which all players are working.

Table 5: Observing a Healthy Person on Disability Increases the Probability Low- and High-Skill Workers Vote to Reduce Disability Benefit Levels (Experiment 2)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	DV: Voted to Reduce Benefits (1=Yes, 0=No)							
	Low and High Skill Workers in Budget Constraint Condition		Low and High Skill Workers in Budget Constraint Condition, Excl. Cases where Benefit Levels Cannot Be Decreased				Matched Pairs (Unweighted) (Weighted)	
Observed a Healthy Worker Collecting Disability	0.244*** (0.087)	0.212** (0.084)	0.211*** (0.079)	0.212*** (0.080)	0.147* (0.088)		0.105 (0.091)	0.098 (0.084)
Injured in Period (1=Yes)		-0.202*** (0.059)	-0.146** (0.057)	-0.146** (0.058)	-0.133** (0.058)	-0.144** (0.057)	-0.093 (0.102)	-0.068 (0.133)
Disability Benefits for Medium Skill Worker								
Period					0.014* (0.008)			
Period Squared / 100					-0.018 (0.012)			
Number of Times Saw Healthy Worker on Disability (in Last 5 Periods) High Skill Worker						0.073*** (0.025)	0.108 (0.076)	0.171* (0.088)
Constant	0.437*** (0.032)	0.523*** (0.040)	0.510*** (0.038)	0.561 (0.461)	0.322*** (0.106)	0.497*** (0.039)	0.920*** (0.183)	0.858*** (0.101)
Player Fixed Effects?	Y	Y	Y	Y	Y	Y	N	N
Session x Period Fixed Effects?	N	N	N	N	N	N	Y	Y
Observations	144	144	140	140	140	140	89	89
Number of Participants	12	12	12	12	12	12		
R-squared	0.057	0.134	0.108	0.108	0.133	0.117	0.581	0.575

***p<0.01, **p<0.05, *p<0.1. Standard errors in parentheses. Sample restricted to low- and high-skill workers in the Budget Constraint (i.e., costly benefits) condition.

trolling 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. Above, we showed that the bad labor market increased the chances that healthy (medium-skill) workers would stay on disability. Controlling for Period and Period squared in column 5, which are jointly predictors of being on disability while healthy, therefore not surprisingly reduces the effect of observing a healthy worker (coefficient= $-.15$, SE=.09, $p<.10$). Finally, we also examine whether a different measure of treatment assignment, which is the number of times in this and the prior 4 rounds a worker saw feedback about a healthy worker being on disability, yields a similar result (this model fits the data slightly worse than the column 2 specification with the same covariates).

Taken together, these results confirm hypothesis 2 that learning of free riding by others

by observing healthy workers on disability increases the propensity to vote to reduce benefits. How large are these effects? In the costly benefits condition, 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.

In columns 7 and 8, 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.¹⁸ 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, etc.). To account for player attributes, we include an indicator for whether a player is high- (rather than low-) skilled. Per column 7, it appears that 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 8.

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 7 specification is similar in substantive magnitude to the result found in column 5—observing a healthy worker on

¹⁸This 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.

disability increases the chances of voting to reduce benefits by 36% (.11/.29).

8 Conclusion

Government social insurance programs have recently joined the long list of social services and redistributive policies targeted in conservative retrenchment efforts. Echoing language often seen in attacks on redistributive policies, calls for retrenchment have invoked concerns about undeserving beneficiaries, often by emphasizing that the programs are being abused by those who do not truly need them. Some of these claims seem motivated by the fact that program utilization has increased in economically depressed areas, which critics have used to argue that social insurance utilization may be a substitute form of welfare for those unable to find suitable rewarding jobs.

These claims raise two major questions at the core of contemporary debates about the political economy of social insurance. The first is whether perceptions of specific program abuse in fact depress public support for social insurance programs. That is, do perceptions that individuals are taking advantage of a specific program cause people to be less supportive of that program? This question is difficult to answer because we usually lack an exogenous source of variation in perceptions of program abuse, and so existing observational correlations between perceptions of abuse and program support struggle to rule out alternative explanations for that correlation (i.e., that conservatives both dislike welfare and believe it is abused because they are conservative). The second question is how economic conditions affect perceptions of program abuse. Focusing on disability insurance, the specific type of social insurance program that we examine in this paper, prior work has shown that deteriorating economic conditions are associated with an increase in the number of people claiming disability. That work has not directly assessed deservingness, nor how changes in program uptake are associated with perceptions of abuse.

Overall, no existing work has examined the interplay among economic conditions, percep-

tions of program abuse, and public support for social insurance programs. While a number of observational studies have examined parts of this larger dynamic environment, these are difficult questions to study observationally and their dynamic interactions are even more complicated. We address these challenges by designing and conducting a series of laboratory experiments in which subjects participate as workers in a simulated political economy where their economic and political choices affect their monetary compensation. Doing so allows us to isolate how changing economic conditions affect program utilization and how exogenously induced changes in perceptions of program abuse affect program support. Importantly, our experiment also moves beyond the scope of existing economic laboratory experiments which have examined how risk exposure affects support for social insurance, how variation in program design affects program abuse, and how observing free-riding affects contributions in public goods games. Prior work, however, has not examined the joint effects of economic shocks and learning about program abuse on political support for social insurance in a unified experimental setting.

In our baseline experiment, which does not incorporate any political dynamics, workers probabilistically experience an injury that automatically causes them to become unemployed and thus, beneficiaries of a laboratory disability insurance program. Workers who recover from an injury choose to return to work or to stay on disability despite being healthy. However, some workers may be affected by a negative labor market shock that precludes them from returning to work at a job with their pre-injury wage. This increases the relative attractiveness of staying on disability compared to returning to work. This intervention therefore proxies economic changes in which the returns to labor market participation decrease for some workers. We find that healthy, medium-skill workers who are unable to return to a job at their prior wage are 28 percentage points more likely to remain on disability in the period in which they recover from being injured than unaffected medium-skill workers (baseline rate of 4% for medium-skill workers).

Our second experiment adds a political process in which participants also vote on dis-

ability benefit levels at regular intervals. All workers benefit from a disability system that provides compensation when injured but requires workers to pay taxes when working. In addition to replicating the effect of the labor market shock on program abuse, we also assess the causal effect of perceived program abuse—identified by providing information about a randomly selected disability beneficiary’s health status—on subjects’ voting decisions on benefit levels. We find clear evidence that when disability benefits are costly for workers (i.e., when raising benefits also raises taxes), participants who observe program abuse reduce their support for more generous benefits. Among workers not affected by a labor market shock, the chances of voting to reduce benefits increases by 52% when they observe a healthy person on disability compared to someone who observes an injured person on disability.

These experiments therefore 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. This causal identification exploits the fact that we randomize which beneficiary a participant receives feedback about, allowing us to isolate concerns about abuse—here perturbed randomly—from general labor market and learning effects. Moreover, unlike survey experiments, a methodological contribution and advantage of our laboratory experimental design is that workers trade off among costly personal benefits (taxes, salaries, and disability benefits). Furthermore, these dynamics map onto theoretically motivated and policy relevant quantities of interest.

Empirically, our work helps to understand how shocks to the labor market can affect perceptions of program abuse. We find that shocks to the labor market cause increases in program abuse and have subsequent downstream effects on program support. While in our design this learning takes place directly (because individuals are told about the health status of a single program beneficiary in each period), a related question for subsequent work is whether individuals infer such abuse absent such direct feedback from the simple fact of

upticks in program usage. Other work suggests individuals do learn about program abuse and infer the deservingness of others, but it could also be the case that individuals infer that rising disability rolls in a poor economy indirectly reveal program abuse. Nonetheless, an important advantage of our experimental design is that it simultaneously incorporates both economic shocks that affect program abuse and perceptions of abuse that affect program support.

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.¹⁹ 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.

For example, we find that labor market shocks that reduce the wages participants earn cause them to forgo employment that is more remunerative than receiving disability benefits. This is despite the fact that in our experimental labor market there are no differences in prestige or non-material rewards to higher wage jobs. If those rewards were present (see, e.g., Kalleberg 2011; Weber 1905), then people would be even less willing to move to a

¹⁹One minor difference between our disability insurance program and actual disability insurance is that we allow workers to immediately claim disability when injured, whereas actual SSDI beneficiary status requires individuals to have a severe long-term disability interfering with work and to have worked long enough and recently enough, and is rarely granted until at least six months after disability onset (a person must have the disability at least 5 months since disability onset before applying for benefits, which if granted are paid one month after applying). Our experimental design is a simplification implemented for tractability in the lab setting.

less prestigious job, likely strengthening the (first stage) effect of labor market shocks on disability program abuse.²⁰ Alternatively, this effect may be overestimated because there is no stigma associated with claiming disability benefits in our laboratory economy (Moffitt 1983; Yaniv 1997). That is, participants may be significantly less likely to stay on disability if they suffered a social cost for continuing to claim benefits or for being discovered claiming benefits while they were healthy.

Similarly, we experimentally manipulate learning by randomly informing players about the health status of a single disability beneficiary. But outside of the lab setting individuals may learn from general patterns in their community or society at large. One key question, broached somewhat in prior survey work, is how individuals form inferences about larger program abuse. We also 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.

These objections aside, we believe this work also provides the foundation for subsequent experimental and observational research about how individuals form beliefs about program abuse and the interplay between economic conditions, these beliefs, and support for government social policy. First, our work focuses on how individual instances of program abuse affect perceptions of overall program abuse and support for those programs. But outside

²⁰We similarly expect that our estimates of the effect of labor market shocks on disability program abuse are conservative estimates and downward biased if the “work” that participants perform in our simulated laboratory economy does not require enough “real effort.”

of the experimental setting, learning takes place in a variety of different ways. 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. Similarly, it would be ideal to isolate real world circumstances, such as labor market shocks, that affect real levels of abuse in government programs.

Both approaches involve significant problems of measurement and causal inference, however. On the measurement side, it is often difficult to clearly quantify individuals' exposure to information (e.g., peer networks, etc.) 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.

References

- Aarøe, Lene and Michael Bang Petersen. 2014. “Crowding out culture: Scandinavians and Americans agree on social welfare in the face of deservingness cues.” *The Journal of Politics* 76(3):684–697.
- Ahlquist, John S., John R. Hamman and Bradley M. Jones. 2017. “Dependency Status and Demand for Social Insurance: Evidence from Experiments and Surveys.” *Political Science Research and Methods* 5(1):31–53.
- Appelbaum, Lauren D. 2001. “The influence of perceived deservingness on policy decisions regarding aid to the poor.” *Political Psychology* 22(3):419–442.
- Autor, David H. 2015. The Unsustainable Rise of the Disability Rolls in the United States: Causes, Consequences, and Policy Options. In *Social Policies in an Age of Austerity: A Comparative Analysis of the US and Korea*, ed. John Karl Schloz, Hyunpyo Moon and Sang-Hyop Lee. Northampton, MA: Edward Elgar Publishing pp. 107–136.
- Autor, David H., David Dorn and Gordon H. Hanson. 2013. “The China Syndrome: Local Labor Market Effects of Import Competition in the United States.” *American Economic Review* 103(6):2121–2168.
- Autor, David H. and Mark G. Duggan. 2003. “The Rise in the Disability Rolls and the Decline in Unemployment.” *The Quarterly Journal of Economics* 118(1):157–206.
- Axelrod, Robert. 1984. *The Evolution of Cooperation*. New York: Basic Books.
- Barber, Benjamin, Pablo Beramendi and Erik Wibbels. 2013. “The Behavioral Foundations of Social Politics: Evidence from Surveys and a Laboratory Democracy.” *Comparative Political Studies* 46(10):1155–1189.
- Black, Dan, Kermit Daniel and Seth Sanders. 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, Scott E. and Peter DeScioli. 2018. “The Problem of Dishonesty in Government

Safety Nets.” Working Paper.

<http://pdescioli.com/papers/bokemper.descioli.safety.nets.draft.2018.pdf>.

Bound, John and Richard V. Burkhauser. 1999. Economic analysis of transfer programs targeted on people with disabilities. In *Handbook of Labor Economics*, ed. Orley C. Ashenfelter and David Card. Vol. 3, Part C Amsterdam: Elsevier Science B.V. chapter 51, pp. 3417–3528.

Browning, Dan. 2013. “Alzheimer’s was Minnesota man’s path to fame, then jail.” *Minneapolis Star Tribune*. (Feb. 17)

<http://www.startribune.com/alzheimer-s-was-minnesota-man-s-path-to-fame-then-jail/191553521/>.

Charles, Kerwin Kofi, Yiming Li and Melvin Stephens Jr. 2017. “Disability Benefit Take-Up and Local Labor Market Conditions.” *Review of Economics and Statistics*
https://doi.org/10.1162/REST_a_00692.

Clark, Andrew E. 1999. “Are wages habit-forming? Evidence from micro data.” *Journal of Economic Behavior and Organization* 39(2):179–200.

Clark, Andrew E. and Andrew J. Oswald. 1996. “Satisfaction and Comparison Income.” *Journal of Public Economics* 61(3):359–381.

Congressional Budget Office. 2016. “Social Security Disability Insurance: Participation and Spending.” Washington, D.C. (June)

<https://www.cbo.gov/sites/default/files/114th-congress-2015-2016/reports/51443-ssdiparticipationspending-onecol-2.pdf>.

Cook, Fay Lomax and Edith J. Barrett. 1992. *Support for the American Welfare State*. New York: Columbia University Press.

Cook, Fay Lomax, Jason Barabas and Benjamin I. Page. 2002. “Invoking public opinion: Policy elites and social security.” *Public Opinion Quarterly* 66(2):235–264.

Croson, Rachel and Melanie Marks. 1998. “Identifiability of individual contributions in a threshold public goods experiment.” *Journal of Mathematical Psychology* 42(2-3):167–190.

Delton, Andrew W., Leda Cosmides, Marvin Guemo, Theresa E. Robertson and John Tooby.

2012. “The psychosemantics of free riding: Dissecting the architecture of a moral concept.” *Journal of Personality and Social Psychology* 102(6):1252–1270.
- Duch, Raymond, Wojtek Przepiorka and Randolph Stevenson. 2015. “Responsibility Attribution for Collective Decision Makers.” *American Journal of Political Science* 59(2):372–389.
- Esarey, Justin, Tim Salmon and Charles Barrilleaux. 2012. “Social Insurance and Income Redistribution in a Laboratory Experiment.” *Political Research Quarterly* 65(3):685–698.
- Fang, Albert H. and Gregory A. Huber. 2018. “Perceptions of Deservingness and the Politicization of Social Insurance: Evidence from Disability Insurance in the United States.” Working Paper. http://albertfang.com/research/ssdi_paper.pdf.
- Gilens, Martin. 1999. *Why Americans Hate Welfare*. Chicago: University of Chicago Press.
- Hamman, John R., Roberto A. Weber and Jonathan Woon. 2011. “An Experimental Investigation of Electoral Delegation and the Provision of Public Goods.” *American Journal of Political Science* 55(4):737–751.
- Hiltzik, Michael. 2013. “‘60 Minutes’ shameful attack on the disabled.” *Los Angeles Times* (October 7). <http://articles.latimes.com/2013/oct/07/business/la-fi-mh-disabled-20131007>.
- Hiltzik, Michael. 2015. “Rand Paul steps up the GOP attack on Social Security.” *Los Angeles Times* (Jan 19). <http://www.latimes.com/business/hiltzik/la-fi-mh-rand-pauls-supremely-cynical-attack-on-social-security-20150119-column.html>.
- Jensen, Carsten and Michael Bang Petersen. 2017. “The deservingness heuristic and the politics of health care.” *American Journal of Political Science* 61(1):68–83.
- Joffe-Walt, Chana. 2013a. “Americans on Disability Play an Increasingly Important Role in the Economy.” *National Public Radio* (Mar 29). <https://www.npr.org/2013/03/29/175722025/americans-on-disability-play-an-increasingly-important-role-in-the-economy>.
- Joffe-Walt, Chana. 2013b. “Expanded Definition of Disability Created Million Dollar Opportunity for Lawyers.” *National Public Radio* (Mar 26).

<https://www.npr.org/2013/03/26/175396983/expanded-definition-of-disability-created-million-dollar-opportunity-for-lawyers>.

Joffe-Walt, Chana. 2013c. “Moving People from Welfare to Disability Rolls is a Profitable, Full-Time Job.” *National Public Radio* (Mar 27).

<https://www.npr.org/2013/03/27/175502085/moving-people-from-welfare-to-disability-rolls-is-a-profitable-full-time-job>.

Kalleberg, Arne L. 2011. *Good Jobs, Bad Jobs: The Rise of Polarized and Precarious Employment Systems in the United States, 1970s to 2000s*. New York: Russell Sage Foundation.

Ledyard, John O. 1995. Public goods: A survey of experimental research. In *Handbook of Experimental Economics*, ed. John H. Kagel and Alvin E. Roth. Princeton University.

McCoy, Terrence. 2017a. “Disabled and disdained: How disability benefits divided this rural community between those who work and those who don’t.” *The Washington Post*. (July 21).

<http://www.washingtonpost.com/sf/local/2017/07/21/how-disability-benefits-divided-this-rural-community-between-those-who-work-and-those-who-dont/>.

McCoy, Terrence. 2017b. “Disabled, or just desperate?: Rural Americans turn to disability as jobs dry up.” *The Washington Post*. (Mar 30).

<http://www.washingtonpost.com/sf/local/2017/03/30/disabled-or-just-desperate/>.

Moffitt, Robert. 1983. “An economic model of welfare stigma.” *American Economic Review* 73(5):1023–1035.

Morton, William R. and Barry F. Huston. 2018. “Social Security: What Would Happen If the Trust Funds Ran Out?” Washington, D.C.: Congressional Research Service.

<https://fas.org/sgp/crs/misc/RL33514.pdf>.

Nowak, Martin A. 2006. “Five rules for the evolution of cooperation.” *Science* 314(5805):1560–1563.

Olson, Mancur. 1965. *The Logic of Collective Action*. Cambridge: Harvard University Press.

- Page, Benjamin I. and Robert Y. Shapiro. 1992. *The Rational Public: Fifty Years of Trends in Americans' Policy Preferences*. Chicago: University of Chicago Press.
- Petersen, Michael Bang. 2012. "Social welfare as small-scale help: evolutionary psychology and the deservingness heuristic." *American Journal of Political Science* 56(1):1–16.
- Petersen, Michael Bang, Rune Slothuus, Rune Stubager and Lise Togeby. 2011. "Deservingness versus values in public opinion on welfare: The automaticity of the deservingness heuristic." *European Journal of Political Research* 50(1):24–52.
- Sell, Jane and Rick K. Wilson. 1991. "Levels of information and contributions to public goods." *Social Forces* 70(1):107–124.
- Smith, Kevin B. 2006. "Representational Altruism: The Wary Cooperator as Authoritative Decision Maker." *American Journal of Political Science* 50(4):1013–1022.
- Spirling, Gene B. 2017. "The Fuzzy Claims Used to Justify Cutting Social Security Disability Insurance." *The Atlantic*. (May 23).
<https://www.theatlantic.com/business/archive/2017/05/ssdi/527802/>.
- Trivers, Robert L. 1971. "The Evolution of Reciprocal Altruism." *The Quarterly Review of Biology* 46(1):35–57.
- van Oorschot, Wim. 2000. "Who should get what, and why? On deservingness criteria and the conditionality of solidarity among the public." *Policy & Politics* 28(1):33–48.
- Weber, Max. 1905. *The Protestant Ethic and the Spirit of Capitalism*. New York: Scribner.
- Woon, Jonathan. 2012. "Democratic Accountability and Retrospective Voting: A Laboratory Experiment." *American Journal of Political Science* 56(4):913–930.
- Yaniv, Gideon. 1997. "Welfare fraud and welfare stigma." *Journal of Economic Psychology* 18(4):435–451.