The Problem of Dishonesty in Government Safety Nets

Scott E. Bokemper
Center for the Study of American Politics
Institution for Social and Policy Studies
Yale University

Peter DeScioli
Department of Political Science
Stony Brook University

Abstract: Many governments provide benefits to unemployed citizens but someone might claim unemployment even when they actually have a job. Here we use economic experiments to examine how the rules of social insurance institutions affect people's honesty. In an unemployment game, participants interact in a difficult job market where they can claim unemployment to request help. In Experiment 1, participants were more honest about unemployment when they repeatedly interacted with the same partner compared to when they switched partners. In Experiments 2 and 3, participants could request benefits from a public fund that was financed by taxes on employed workers. When benefits were unregulated, participants frequently (67%) made false unemployment claims. However, when we applied common regulations from real-world unemployment programs, including limits, bans, and fines, participants were dishonest at considerably lower rates (33%-40%). We discuss implications for the efficient relief of hardship and the experimental study of social insurance institutions.

Keywords: Social Insurance, Dishonesty, Cooperation, Experimental Economics

Introduction

The concept of social insurance is simple: People help others who suffer hardships such as unemployment, illness, or damage from natural disasters, and they receive help when they are the ones in trouble, so everyone gains some protection against hardship. All societies have at least some basic forms of social insurance such as helping family or friends, and many societies have developed highly elaborate institutions for mitigating hardship. In hunter-gatherer societies, foragers who find food frequently share it with those who do not (Gurven 2004). In agricultural societies from past eras, some governments stockpiled food in granaries that they would distribute to citizens during famines (Shiue 2005). In modern societies, people hit by hard times look for help from family, friends, charities, and, increasingly, the government (e.g., Larrimore, Dodini, and Thomas 2016). For example, in the United States, over half of the federal budget is spent to relieve hardships, including illness, injury, poverty, and unemployment (DeSilver 2017). In many developed countries, the government has become in large part a social insurance institution.

Here we use economic experiments to investigate a key issue at the core of social insurance: the problem of dishonesty in claims of hardship. The basic problem is easy to appreciate for anyone who has been asked for help. Does this person really need help? Do they really need help as much as they say, or do they only seek to gain from another's generosity? The dilemma is a difficult one, and errors on either side are costly. To turn down a genuine claim is to abandon someone in need, add further injury to the victim, and shirk any responsibilities to help. But to accept a false claim is to squander resources, reward deception, and invite many more opportunists to the door.

We use methods from experimental economics to study this problem in a laboratory microcosm. In experimental economics, researchers design incentivized games to create small-scale economic systems for controlled experiments, similar to studying a miniature plane in a wind tunnel (Kagel and Roth 2015; Morton and Williams 2010; Ostrom 1998; Smith 1982). Experimenters use these methods to study people's *behavior*, such as cooperation or investment, and also to study economic and political *institutions*—how different rules affect the outcomes—such as comparing different types of markets, auctions, regulations, committees, or elections. Experimenters have studied a wide array of economic and political institutions; however, the experimental study of social insurance institutions is less developed than other major areas such as market design, voting mechanisms, or managing common pool resources.

We designed a new economic game—the unemployment game—to study social insurance institutions. Participants are workers who look for a job and use their wages to pay the bills and improve their health. At the end of the game, the worker's health determines how much real money they earn. In some rounds, the worker is unemployed and if they cannot pay their bills, then they suffer the hardships of poverty which do substantial damage to their health. To prevent hardship, workers can ask for help from peers (Experiment 1) or a public fund that is financed by taxes on other workers (Experiments 2-3). Importantly, the worker's employment is not known to others, so they could ask for help even if they actually found a job and are at no risk of hardship.

We use the unemployment game to trace the problem of dishonesty through a broad arc of social insurance institutions found across many societies. In Experiment 1, we begin with the most basic form of social insurance by looking at how peers help each other through unemployment. We test a key prediction from reciprocity theory that participants will be more honest about their employment when they repeatedly interact with the same partner compared to when they interact with a different person each period. This mirrors the difference between helping in small close-knit groups versus large anonymous communities. In Experiment 2, we examine a form of government social insurance in which a public fund collects taxes from the employed workers and distributes money to workers who claim unemployment. We test whether participants are more honest when there is a limit on the number of times they can claim unemployment compared to when the unemployment benefits are unregulated. In Experiment 3, we examine a public fund in which participants know that there is a chance (1/3) that each claim of unemployment will be checked to see if they are really unemployed. We test whether penalties deter dishonesty and if deterrence is greater when the penalty is a fine compared to a ban from claiming benefits in the next period.

Reciprocity and covert cheating

In the most basic form of social insurance, family and friends help each other through hard times. This kind of helping is observed across cultures including in small-scale foraging societies (Gurven 2004) and among people with little access to formal social insurance institutions (Morduch 1999). When people rely on peers for support in hard times, the recipient's honesty is typically enforced by a kind of reciprocity (Andreoni and Miller 1993; Axelrod 1984; Trivers 1971). If a helper discovers that the recipient made false claims about their needs, then they generally refuse to help in the future.

Importantly, the stability of reciprocity depends on repeated interactions and the ability to detect cheating. Repeated interactions are important because it is the potential withdrawal of help in the future that creates an incentive for people to help each other and to be honest about their needs. Coate and Ravallion (1993) applied the logic of reciprocity to social insurance in a game theoretic model in which two households with uncertain income can help each other through hard times. The model shows that reciprocity can sustain social insurance in repeated interactions. Similarly, in an economic experiment in which participants received uncertain income each period, they were more willing to give money to an unlucky partner when there was a greater chance that they would interact with the same partner again in the future (Charness and Genicot 2009). Reciprocity also depends on the ability to detect cheating, and research from psychology finds that people closely monitor cheating and they withhold future cooperation in response (Cosmides 1989; Delton et al. 2012).

However, in many real-world situations, an individual does not know for sure if someone cheated. For instance, someone who claims hardship might be lying to fraudulently gain benefits. This uncertainty makes reciprocity more difficult to sustain. For example, a game theory model showed that reciprocal strategies perform worse in noisy environments where other players' choices are uncertain (Bendor 1993). Similarly, in economic experiments, participants who were strictly reciprocal in a noisy environment earned low payoffs after reciprocating mistaken defections (Fudenberg, Rand, and Dreber 2012). In situations of uncertainty, a reciprocal player can only probabilistically guess whether their partner is cheating. These uncertainties frequently apply to social insurance and hardship, including claims of unemployment.

Mass society and government social insurance

Reciprocity is also more challenging to sustain in larger groups in which it is difficult to monitor individuals' reputations (Boyd and Richerson, 1988). As societies grew in size from small-scale communities to immense, anonymous societies, this created new challenges for maintaining reciprocity and hence the social insurance institutions that depend on it. When people are anonymous in large groups, they are less constrained by reputation and the repercussions for future interactions. For a stranger roaming through a crowd, each interaction approximates a one-shot game in which cheating has no long-term costs. This means that a self-interested stranger is expected to exaggerate their needs, and in response, helpers will attribute little credibility to these claims, thus compromising the basis for social insurance.

In response to these distinctive challenges, people have devised a variety of institutions to provide social insurance in mass societies. Among them is government social insurance, in which centralized government programs collect funds though mandatory taxes and then deliver benefits to citizens who suffer hardships. This includes unemployment insurance to protect against job loss, which is especially important in modern economies characterized by an extreme and ever-shifting division of labor (Moene and Wallerstein 2001; Rehm, Hacker, and Schlesinger 2012). Government unemployment programs face a difficult problem of dishonesty because they need to evaluate claims of unemployment for strangers in a noisy environment. Consider, for instance, a government that gave benefits to anyone who claimed unemployment, no questions asked. Most likely, a swarm of unscrupulous opportunists would devour the funds and there would be nothing left for those who are actually unemployed.

To prevent exploitation and inefficiency, policymakers create regulations that promote honest claims and efficient targeting, including rules for who is eligible to receive benefits, how often they can receive them, and what penalties apply for breaking the rules. Equipped with a battery of regulations, many unemployment programs appear to do fairly well at preventing fraud. For instance, one report estimated that in 2016 the U.S. government provided about $600 million in unemployment benefits to citizens who misrepresented their earnings, which represents only 1.9% of program spending (United States Department of Labor 2016). Importantly, however, false claims are difficult to measure precisely because people actively conceal them. This also makes it difficult to assess the effectiveness of specific regulations for curbing dishonest claims.

The problem of dishonesty is critical for the efficiency of social insurance institutions because false claims divert funds from other uses, including meeting other citizens' genuine needs. This is particularly important because social insurance programs are a perennial source of political conflict whose funding is frequently attacked and constrained by political opponents. When funds are politically limited, their efficient use is even more crucial for the effective relief of hardship. Moreover, the problem of dishonesty is also important for public support of social insurance programs. Surveys and experiments find that public opinion about social programs is shaped by whether the recipients deserve to receive benefits (Aarøe and Petersen 2012; Gilens

1999; Petersen 2012; van Oorschot 2006). Hence, the prevalence of dishonesty is likely to influence the public's support for social spending.

Economic experiments and social insurance institutions

There is of course a massive literature on social insurance spanning multiple fields and methods, including observational analyses, economic models, opinion surveys, and so on. In comparison, the *experimental* study of social insurance institutions is in a nascent state consisting of a relatively small set of experiments.

In previous experiments, researchers examined when people vote to expand social insurance programs. Generally, they found that participants are more supportive of social insurance when they have a higher probability of losing their income (Ahlquist, Hamman, and Jones 2017; Barber, Beramendi, and Wibbels 2013; Esarey, Salmon, & Barrilleaux 2012). For instance, one study examined people's relative support for programs that aim to promote equality per se versus programs that specifically protect against the risk of low income (Barber, Beramendi, and Wibbels 2013). In a series of rounds, participants earned income and voted on a tax rate for their group. In the beginning, the tax revenue was allocated to increase equality, then the program changed to augment low income, and finally it changed to a hybrid institution that mixed both functions. Under the program focused on equality, participants who earned greater average income voted for a lower tax rate; in the program focused on supplementing low income, participants who were more vulnerable to low income preferred a higher tax rate; when the institution served both functions, participants' tax rate preferences were driven by their exposure to risk, but not their relative income.

Another literature on tax evasion examines how specific regulations affect whether individuals honestly report their income. In a typical experiment, participants receive income and decide how much to report to the tax collection agency (reviewed in Alm 2012; Kirchler et al. 2010). Participants are audited with some probability and they pay a penalty for underreporting their income. Broadly, this research finds that tax compliance depends on several factors including the tax rate, the frequency of audits, and the severity of penalties (e.g., Alm, Jackson, and McKee 1992; Friedland, Maital, and Rutenberg 1978; Spicer & Thomas 1982).

Finally, experimenters have studied a kind of social insurance among peers, rather than implemented by collective institutions. In the solidarity game, three players each receive either a fixed income (2/3 chance) or nothing (1/3 chance), and those who receive income can decide to share with the others (Büchner, Coricelli, and Greiner 2007; Ockenfels and Weimann 1999; Selten and Ockenfels 1998; Trhal and Radermacher 2009). Specifically, before finding out if they received income, each player decides how much money to share if they receive it. Generally, these experiments find that participants share substantial amounts with the unlucky group members. For instance, one study found that participants gave 15-20% of their income to unlucky group members (Büchner, Coricelli, and Greiner 2007).

In the present experiments, we designed an unemployment game that expands in several respects on previous research. First, we introduce the concept of hardship, which is a high cost to health that occurs when a player's income is insufficient to cover their expenses, representing the damage of poverty such as hunger, illness, or displacement. Hence, unlike the solidarity game, giving money in the unemployment game is not only a zero-sum transfer of money, because money that is sent to the unemployed prevents the costs of hardship, hence improving aggregate economic efficiency. Second, to examine dishonest claims, a player can claim unemployment whether or not they have a job, and others cannot know for sure whether a claim is honest. This kind of noisy cooperation has been studied in a few previous experiments (e.g. Fudenberg, Rand, and Dreber 2012), but not for the particular case of social insurance. Third, we introduce a public fund that collects taxes and provides unemployment benefits, and we manipulate the regulations and penalties for preventing false claims. While there is a large literature on punishment in cooperative games (reviewed in Balliet et al. 2011) and on penalties for tax evasion (Alm 2012; Kirchler et al. 2010), we extend this research on enforcement to the important case of social insurance institutions.

The Unemployment Game

Consider a simple model of social insurance with two players: a worker and a helper. The worker looks for a job and they either find one with some probability or they are left unemployed. Then the worker chooses whether to *Claim* unemployment to ask for money to pay the bills or remain *Silent*. If the worker asks for money, then the helper decides whether to *Give* money or *Deny* the request. The helper does not know whether the worker is really unemployed, but they do know the chances of finding a job.

Figure 1 shows this situation in a payoff matrix. The payoffs represent a player's general health which can be increased by income and decreased by bills and hardship. The payoffs are denominated in general health to reflect the fact that hardship can damage not only someone's finances but also their physical and psychological wellbeing. The worker's chance of finding a job is p(employment) = .5. The helper does not know whether the worker found a job and hence which payoffs apply, but they do know the chance of finding a job. (In game theory terms, it is a Bayesian game in which the helper has incomplete information about the worker's employment). The payoffs reflect parameters for wages, bills, and the costs of hardship. The helper and an employed worker receive a wage of w = 20, while an unemployed worker receives 0. (In this example, the helper is always employed with wages they could share; later we allow both players to act as worker and helper.) The helper and worker both pay bills of b = 6, deducted from their wages. If the worker does not have enough money to pay the bills, then they pay the cost of hardship, h = 12, which reflects debt from unpaid bills and the additional costs of poverty due to deprivation of food, shelter, and medicine.

If the worker asks for help and the helper chooses to *Give*, then the helper gives the worker enough money to pay the bills, $+b$ for the worker and $–b$ for the helper, which prevents the worker from suffering hardship. (The helper's choice to *Give* transfers money only if the worker chooses to *Claim*, not if they are *Silent*.) The payoffs in Figure 1 sum these parameters accordingly. For instance, the helper earns the wage minus the bills (20 – 6 = 14) when they do not give money; when they do give money (in response to *Claim*), they earn less (20 – 6 – 6 = 8).

In this model, so far, a self-interested helper is expected to always deny requests since giving is costly. (*Deny* is a weakly dominant strategy.) Moreover, this is economically inefficient when the worker is unemployed, since the aggregate payoffs are greater when the helper gives money to an unemployed worker (aggregate of 8 + 0 = 8 versus 14 + -12 = 2) because giving prevents the additional cost of hardship.

Figure 1. A Simplified Unemployment Game

		Employed (p = .5)			Unemployed (1 - p)	
		Worker			**Worker**	
		Claim	*Silent*		*Claim*	*Silent*
Helper	*Give*	8, 20	14, 14	*Give*	8, 0	14, -12
	Deny	14, 14	14, 14	*Deny*	14, -12	14, -12

Note: This is a simplified version with fixed roles for helper and worker. We extend the model in the text and experiment to a repeated version in which both players act as worker and helper.

Now consider a repeated version of the game in which both players can act as worker and helper each period. Both players search for a job (p = .5), decide whether to claim unemployment, and, if employed, decide whether to help their partner. In this case, a helper who is willing to give money to an unemployed worker could potentially benefit in the future if their partner returns the favor. Hence, when repeated, the unemployment game becomes a variant of the well-known repeated prisoner's dilemma (Axelrod, 1984). Hence, it has multiple equilibria including a selfish equilibrium in which players always ask for help while denying help to others, and a cooperative equilibrium in which players honestly signal their unemployment and give money to a partner who is unemployed, yielding greater payoffs for both players in the long run.

However, cooperation could be more difficult to sustain in this variant of a social dilemma. False claims pose a special difficulty because a helper cannot observe the worker's employment status, so they can only estimate probabilistically whether the partner is lying based on the chances of employment. On the other side, this means it is more tempting for a worker to make false claims to gain additional money because this cheating is concealed.

In the present experiments, we use the unemployment game to pose participants with these difficult dilemmas at the core of social insurance. We programmed custom software to create an interactive, online unemployment game for experiments (see demo of the software here: http://unemploymentgame.droppages.com/). We observe how participants claim unemployment and provide unemployment benefits in variations of the game with different social insurance institutions.

Experiment 1

In Experiment 1, we start with the simplest form of social insurance in which peers can help each other through hard times. Motivated by reciprocity theory (e.g., Axelrod 1984), we test whether participants are more honest when they interact with the same partner repeatedly compared to interacting with a new partner each round. As discussed above, previous experiments found that repeated interaction promotes cooperation. However, players in the unemployment game cannot know for sure whether their partner is being honest about unemployment, creating an additional obstacle for cooperation. Importantly, the difference between repeated and one-shot interactions mirrors a key difference in real-world social insurance systems found across societies: Some forms of social insurance rely on personal relationships and reputations, whereas other forms such as government social insurance are implemented by impersonal institutions in anonymous communities where it can be difficult to monitor individuals' reputations.

Based on reciprocity theory, the reputation hypothesis predicts that participants will be more honest in repeated interactions than in one-shot interactions. If people rely on the same neighbor for help, then the neighbor can detect probabilistically if they are making deceptive claims. Specifically, all players know that the chance of being unemployed is 50%, so a neighbor who asks for help substantially more than half of the time is likely to be exaggerating their needs. In this instance, there is a possible incentive for workers to be honest because excessive requests could lead their neighbor to suspect cheating. On the other hand, an employed worker could be tempted to ask for help anyway to make some extra money, since their partner cannot detect any particular instance of lying, which could potentially make honesty difficult to sustain.

Methods

We recruited undergraduates (N = 142) to participate in a laboratory experiment (59% female; mean age = 20 years, S.D. = 2 years. Participants earned $5 for attendance and they earned additional money from the experimental game (mean = $16.76, S.D. = $4.08). Participants were seated at private computer terminals to interact in groups on a computer network in a laboratory designed for economic experiments. Participants read the instructions, answered comprehension questions, and then played the unemployment game. Participants' final health in the game determined their additional earnings (20 cents per health point; see Supporting Information S4 for instructions).

We programmed an online version of the unemployment game using HTML and Javascript, and participants played the game on an internet browser. In the game, participants are workers who look for jobs to pay the bills and buy goods to boost their health. A worker begins with 20 health points. Each month, a worker has a 50% chance of finding a job. An employed worker receives 20 in wages which they can use to pay the bills, which cost 6 per month, and to buy goods that add to their health. If a worker does not find a job, then they cannot pay the bills so they suffer the additional damage of hardship, which deducts 12 from their health.

To avoid this damage, a worker can claim unemployment to ask their neighbor for help. Each worker is matched with a neighbor, who is another participant playing a worker in the game. If the worker requests help and their neighbor found a job, then the neighbor can choose to send 6 to cover the worker's bills, which prevents the cost of hardship. Both neighbors can claim unemployment (whether they found a job or not), and they make these decisions simultaneously before both players' choices are revealed; similarly, players decide simultaneously whether to send money if their neighbor claimed unemployment.

After both workers decide whether to request and send help, they pay their bills, spend any remaining money on goods, and then advance to the next month. (To focus on social insurance, players cannot insure themselves by saving money across months.) This sequence repeats for a total of 24 months (participants were unaware of the exact number of months to prevent end game effects).

Figure 2. Interface of the Unemployment Game with Peer Safety Net

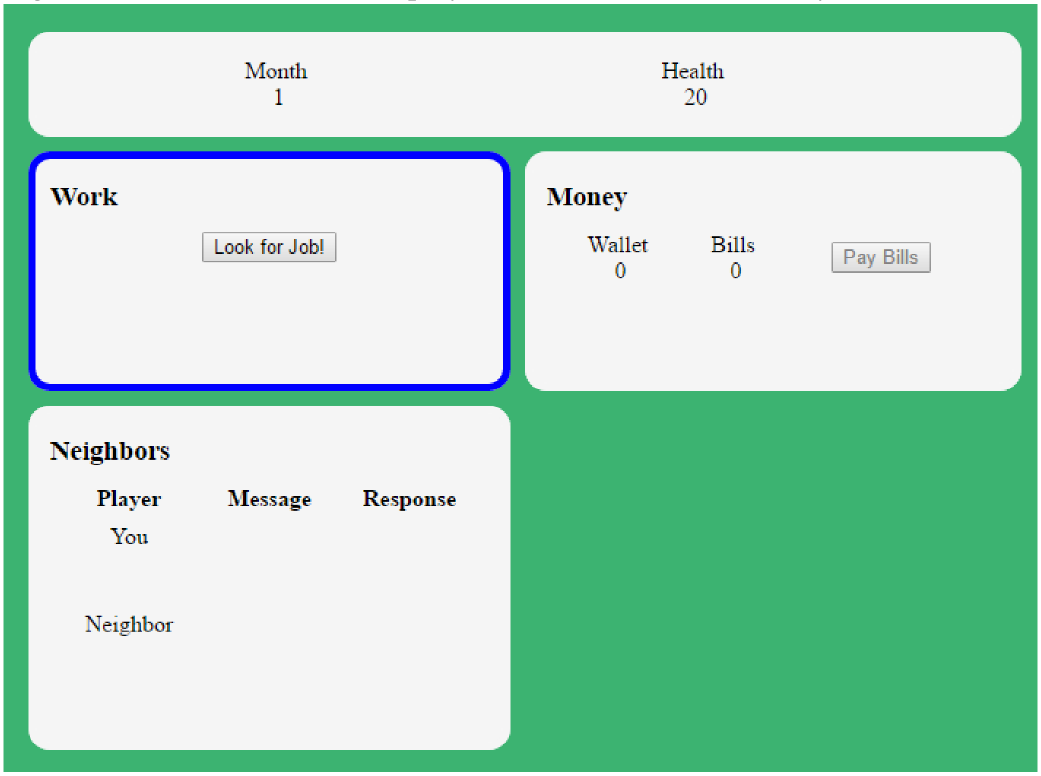

Participants were assigned to the partner condition (*n* = 76) or stranger condition (*n* = 66). In the partner condition, participants were paired with the same neighbor for all 24 months. In the stranger condition, participants were paired randomly with a different neighbor each month (drawn from a set of at least 7 other participants).

The main dependent variable is dishonesty: the percentage of time that a participant claimed unemployment when they actually had a job. Hence, this dependent measure could potentially vary from 0%, if a player never claimed unemployment when they had a job, to 100%, if they always claimed unemployment when they had a job. We also examine how often an employed player helped someone who requested it. This dependent measure could vary from 0%, if an employed player never helped someone who requested it, to 100%, if they always helped when asked.

Results

Figure 3 shows participants' dishonest claims, helping, and final health for the stranger and partner conditions. In the stranger condition, participants frequently made false claims of unemployment (mean = 76% dishonest); however, in the partner condition, participants were substantially less dishonest (mean = 45% dishonest), $t(140) = 7.55$, $p < .001$. Participants were also more likely to help a requester in the partner condition than the stranger condition, $t(140) = 7.72$, $p < .001$. Further, participants earned greater total health in the partner condition than the stranger condition, $t(140) = 3.06$, $p < .01$. We also examined whether participants' dishonesty and helping changed as the rounds progressed in the game; a regression analysis showed no change in dishonesty and a small decline in helping in both conditions (see Supporting Information S1). Overall, these results show that participants created more effective social safety nets, despite uncertainty about each other's honesty, when they could establish reputations over time in repeated interactions.

Figure 3. Mean (S.E.) Dishonesty, Helping, and Health by Condition

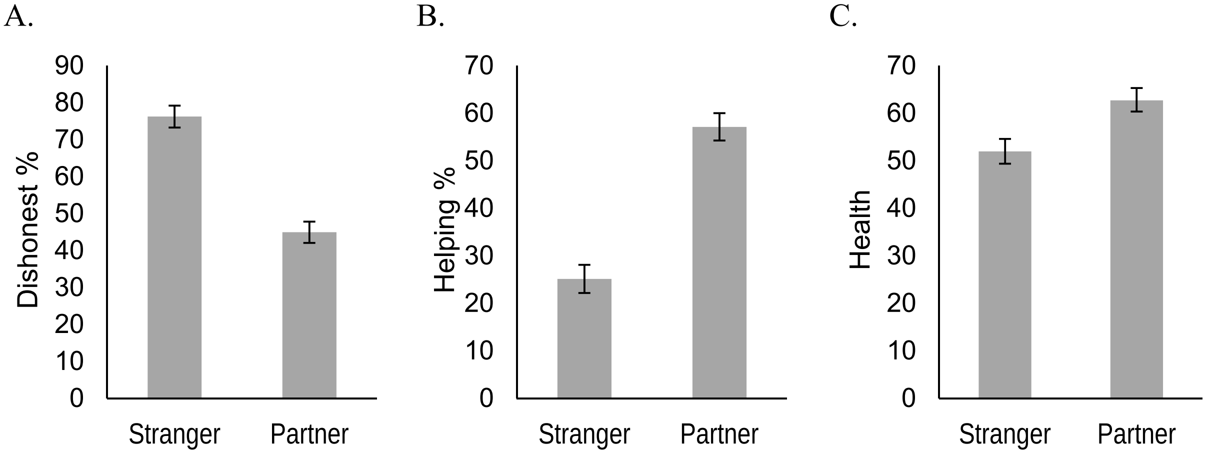

Did liars and helpers make more money?

We examine whether participants who lied more often made more money. Figure 4 (panel A) plots the best-fit lines for participants' final health by how often they were dishonest. The plot suggests that lying paid off in the stranger condition but not the partner condition. We confirm this pattern in a regression analysis (Table 1, model 1). In the stranger condition, participants who were more dishonest earned more health. In the partner condition, a participant's dishonesty had no significant effect on health, $t = 0.53$, $p = .60$ (test of the combined coefficients for dishonesty and the partner X dishonesty interaction). We also look at whether helping paid off. The plot (Figure 3, panel B) and regression (Table 1, model 2) indicate that players who helped more often made less money in the stranger condition but not in the partner condition, which showed no significant effect of helping on health, $t = 0.29$, $p = .77$ (again, testing the combined coefficients for dishonesty and the partner X dishonesty interaction).

Figure 4. Final Health by Dishonest % and Helping %

A. B.

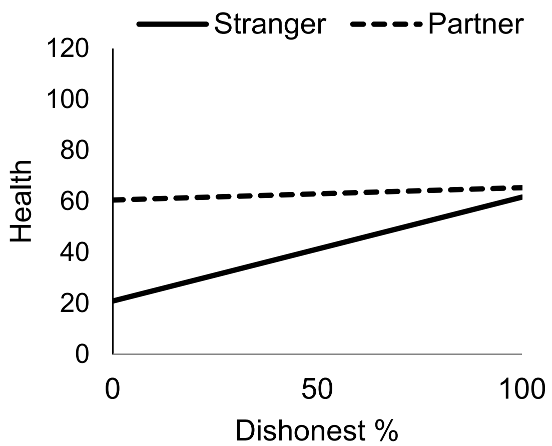

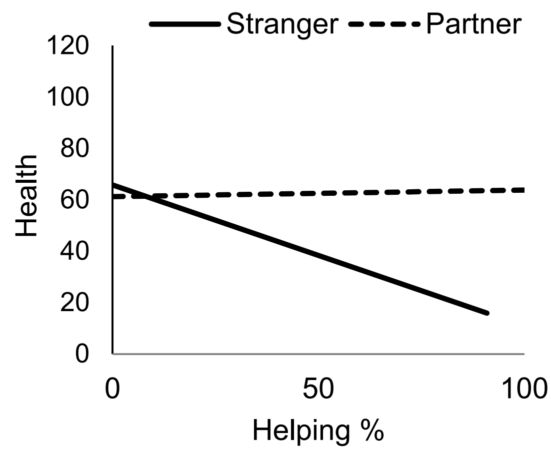

Table 1. Final Health Predicted by Dishonest % and Helping %

	(1)			(2)		
	Coeff.	Std. Err.	P value	Coeff.	Std. Err.	P value
Partner Condition	39.66	(9.38)	0.000	-4.57	(6.37)	0.475
Dishonest %	0.41	(0.10)	0.000			
Partner X Dishonest %	-0.36	(0.14)	0.009			
Helping %				-0.55	(0.10)	0.000
Partner X Helping %				0.57	(0.13)	0.000
Constant	20.89	(8.12)	0.011	65.82	(3.37)	0.000

Note: OLS regression. n = 142. R^2 = 0.162 for model 1 and 0.240 for model 2.

Did participants who received more help also give more help?

To examine reciprocity, we test whether participants who received help from their neighbor were more likely to give help in return. Table 2 reports a logistic regression of a participant's decision to help (coded 0 or 1) when they had a job, received a request, and had requested help in the previous round, which their neighbor provided or declined (coded 0 or 1). The model includes the partner condition, whether the participant received help in the previous round (lagged received help), the interaction, and a random effect for each participant to account for repeated observations. Consistent with reciprocity, the results show that participants who received help in the previous round were more likely to give help. Specifically, the model indicates that in the partner condition, the participant's probability of giving help increased from 52% to 70%, when their own previous request for help was denied or fulfilled, respectively; in the stranger condition, the participant's probability of giving help increased from 21% to 40% when their previous request was denied or fulfilled, respectively.

Table 2. Giving Help Predicted by Received Help

	Coeff.	Std. Err.	P value
Partner Condition	1.75	(0.28)	0.000
Lagged Received Help	1.12	(0.32)	0.001
Partner X Lagged Received Help	-0.18	(0.42)	0.665
Constant	-1.65	(0.20)	0.000

Note: Logistic regression with random effect for participant (*n* = 142 participants). *n* = 1,038 helping decisions.

Economic efficiency

To examine economic efficiency, we compare participants' earnings to what would happen if all players always cooperated (honest and helpful) or if they always defected (dishonest and never help). With full cooperation, players would earn an average of 80 total health in both conditions (see Supporting Information S3 for efficiency calculations). Participants earned 22% less than this cooperative benchmark in the partner condition, and they earned 35% less in the stranger condition. With full defection, players would earn an average of 44 total health. Participants earned 43% more than the defection benchmark in the partner condition, and they earned 18% more in the stranger condition.

Discussion

Overall, participants were more honest and provided more help when they were paired with the same partner every round, compared to a different stranger each month. Participants also made more money and achieved greater economic efficiency in the partner condition than the stranger condition. Overall, these results show that participants were able to create peer safety nets despite their uncertainty about each other's needs, but only when they could monitor each other's reputation in repeated interactions. In contrast, when participants changed neighbors each moth, they lied at high rates and also provided little help. This shows how peer helping breaks down in anonymous interactions, which points to the difficult challenge of sustaining social insurance in mass societies in which reputations are difficult to monitor.

Experiment 2

Government institutions offer a potential solution for sustaining social insurance in anonymous societies. Governments can collect taxes to distribute benefits to people who have fallen on hard times. This ensures that more help is available for people who need it, but does not by itself address the problem of dishonesty.

A common policy tool for social insurance programs is limiting the number of times that an individual can receive benefits in a fixed period of time. By limiting access, beneficiaries are responsible for managing their own claims. If individuals can efficiently manage limited benefits, then they will tend to avoid making false claims in order to preserve their limited benefits for when they are genuinely unemployed, promoting honest claims. Alternatively, people might not adequately plan for the future and claim benefits when they are employed, discounting the future.

In Experiment 2, we examine participants' dishonesty when unemployment benefits are provided by a centralized public fund financed by taxes on the employed. We also test the resource management hypothesis that participants will be more honest about unemployment when benefits are limited, compared to unregulated benefits.

Methods

We recruited undergraduates (N = 128) to participate in a laboratory experiment (58% female; mean age of 19.6 years (S.D. = 1.42 years). Participants earned $5 for participation and they earned additional money from the experimental game (mean = $16.95, S.D. = $5.20).

Participants played an unemployment game similar to Experiment 1 except instead of peer helping, unemployment benefits are provided by a public fund (see demo: http://publicfund.droppages.com/). The public fund is financed by taxes on employed workers. When a worker finds a job, they earn 20 and then pay -6 in taxes to the public fund. As before, if a worker cannot pay their bills, they suffer a hardship cost of -18.

Before paying bills, workers can claim unemployment to ask for help from the public fund. Importantly, workers can claim unemployment whether they found a job or not. If there are funds available, any worker who claims benefits receives them. If there are not enough funds, the public fund randomly determines who receives the remaining benefits. At the end of the round, any remaining money in the public fund carries over to the next round. (At the end of the game, any remaining money is split evenly.) Unlike Experiment 1, workers knew that the game lasts for 24 months, because they needed this information to manage benefits in the limits condition (described below). As in Experiment 1, we examine dishonesty and health. We also examine the frequency of hardship to assess the efficiency of the social insurance institution.

Figure 5. Interface of the Unemployment Game with Public Fund

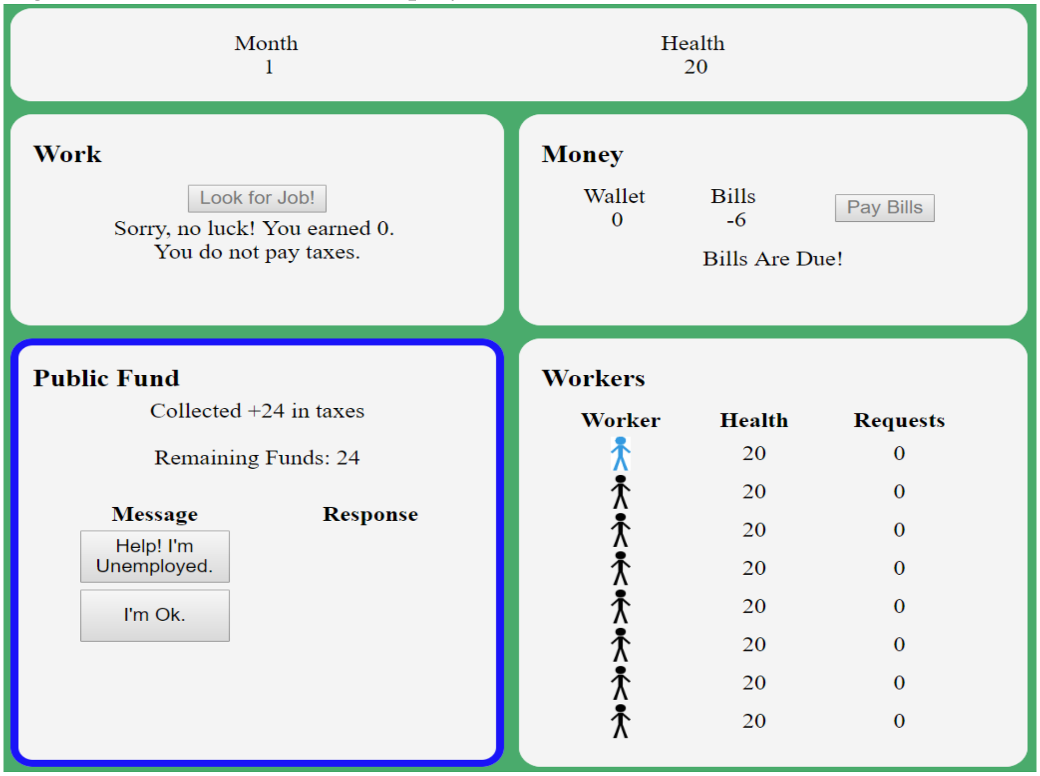

Participants were assigned to an unregulated public fund (n = 64) or a public fund that provided limited benefits (n = 64), a policy tool commonly used in unemployment insurance programs. In the unregulated condition, workers could claim benefits every month as long as funds were available. In the limited condition, workers could receive benefits only 12 times in 24 months. (A claim that was denied due to insufficient funds did not count against the limit.) We examine participants' overall honesty and whether participants are more honest when benefits are limited.

Results

Figure 6 shows participants' dishonest claims, hardship, and final health for the unregulated and limited conditions. Participants were less dishonest when benefits were limited (mean = 38% dishonest) compared to when they were unregulated, (mean = 67% dishonest), $t(126) = 7.67$, $p < .001$. Participants also suffered hardship less often in the limited condition than the unregulated condition, $t(126) = 4.25$, $p < .001$, and they earned more health (and money) in the limited condition than the unregulated condition, $t(126) = 3.03$, $p < .01$. These results show that when unemployment benefits were regulated by limits, participants were less dishonest about their employment, they suffered less often from hardship, and they achieved better economic outcomes.

Figure 6. Mean (S.E.) Dishonesty, Hardship, and Health by Condition

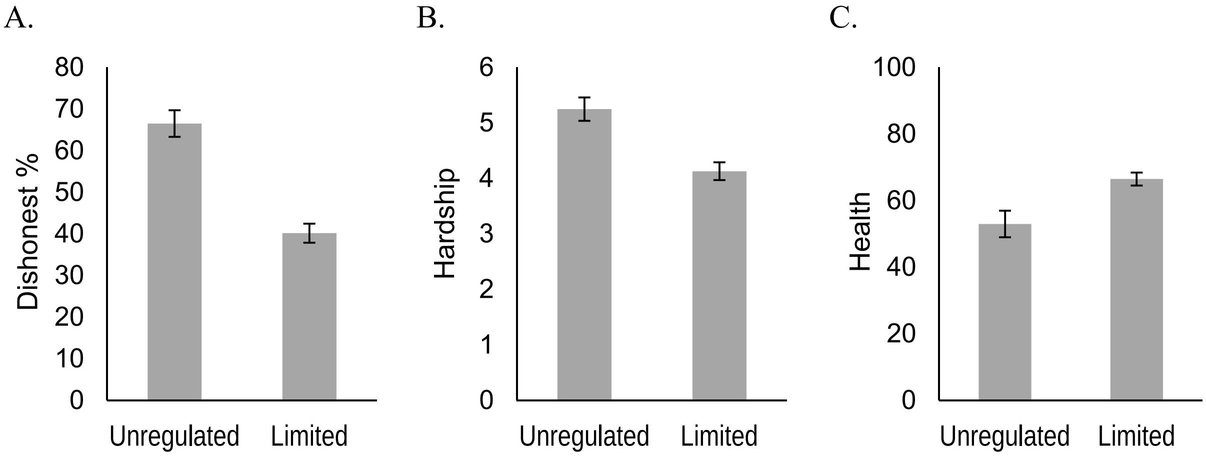

Did dishonesty change over time?

Figure 7 shows the percentage of dishonest claims by round aggregated across all participants. The aggregate dishonesty in each round is the percentage who requested benefits out of the total number who had a job and could request benefits (i.e., excluding those who reached their limit in the limited condition). Looking at the trends in the figure, when benefits were unregulated, participants became steadily more dishonest over time up to about 90% in the final rounds, presumably as they learned about the incentives and others' dishonesty. When benefits were limited, participants' dishonesty remained steady at relatively low levels (~40%) up until the last two rounds, when it suddenly rose all the way to 95%. To look closer, we conducted a logistic regression of a participant's decision to be dishonest (coded 0 or 1, when they had a job) with predictors for the limited condition, round, the interaction, and random effects for participant to account for individual-level variation (Table 3, model 1). The results show that dishonesty increased over time, but the significant interaction indicates that dishonesty increased less when benefits were limited (predicted increase of 14.9% in the probability of lying over all rounds) compared to unregulated (predicted increase of 41.5% in the probability of lying over all rounds).

Finally, there was a sudden increase in dishonesty in the final rounds of the limited condition. This endgame effect makes sense strategically since players knew that these were the last rounds. (They needed to know the number of rounds in this experiment in order to manage their limited unemployment benefits.) Hence, a strategic self-interested player might as well use up their remaining claims whether they found a job or not. To examine this, we reanalyzed the time trends while excluding the last two rounds when endgame effects are expected (Table 3, model 2). As before, in the unregulated condition, dishonesty increased over time (predicted increase of 37.2% in the probability of lying over all rounds), but now in the limited condition, there was no significant change over time in dishonesty (predicted increase of 0.4% over all rounds). Hence, dishonesty was stable over time in the limited condition until the endgame effect in the final rounds. Moreover, we suggest that this large endgame effect points to participants' alert opportunism; even after 22 rounds of being relatively honest, participants quickly became nearly entirely dishonest in the final rounds when restraint no longer served their self-interest.

Figure 7. Dishonest % by Round

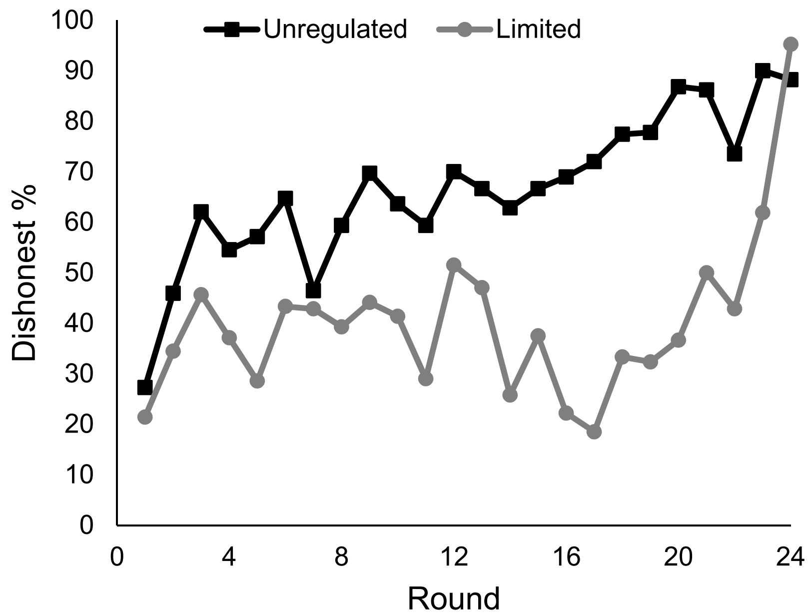

Table 3. Dishonesty Predicted by Limited and Round

	(1) All Rounds			(2) Rounds 1-22		
	Coeff.	Std. Err.	P value	Coeff.	Std. Err.	P value
Limited Condition	-0.49	(0.30)	0.107	-0.28	(0.31)	0.369
Round	0.11	(0.01)	0.000	0.10	(0.02)	0.000
Limited X Round	-0.07	(0.02)	0.000	-0.10	(0.02)	0.000
Constant	-0.38	(0.22)	0.084	-0.33	(0.22)	0.143

Note: Logistic regression with random effect for participant (n = 128 participants). Model 1, n =1,499 opportunities to lie. Model 2, n = 1,393 opportunities to lie.

Did liars make more money?

Next, we analyze whether people who were more dishonest gained more total health. Figure 8 shows the best-fit lines for participants' final health by dishonesty. The plot indicates that participants who were more dishonest made more money in the unregulated condition, but they made less money in the limited condition. We confirmed this with a regression analysis (Table 4). There is a significant interaction, indicating that the effect of dishonesty differed by condition. Specifically, in the unregulated condition, the effect of dishonesty does not significantly differ from zero; in the limited condition, dishonesty significantly reduced the player's health, $t = 3.86$, $p < .001$ (combined coefficients for dishonesty and the dishonesty X limited interaction).

Figure 8. Health by Dishonest %

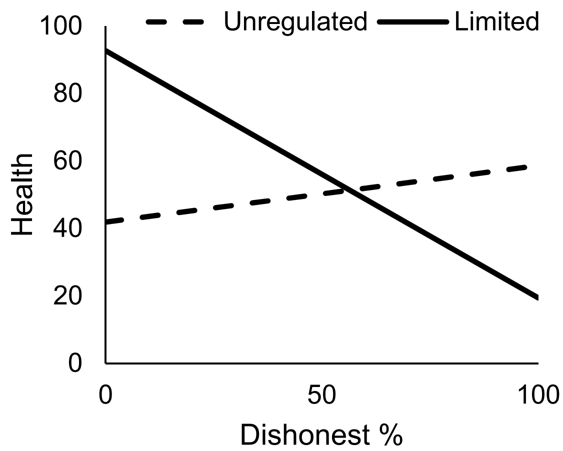

Table 4. Health Predicted by Dishonest % and Limited Condition

	Coeff.	Std. Err.	P value
Limited Condition	51.80	(11.39)	0.000
Dishonest %	0.17	(0.12)	0.157
Limited X Dishonest %	-0.89	(0.22)	0.000
Constant	41.89	(8.35)	0.000

Note: OLS regression. $n = 128$. $R^2 = .177$.

Economic efficiency

To examine economic efficiency, we compare participants' earnings to what would happen if all players were always honest or if they were always dishonest. With complete honesty, players would earn on average 102 health in both conditions. Participants earned 35% less than this honesty benchmark in the limited condition, and they earned 48% less in the unregulated condition. With complete dishonesty, players would earn on average 32 health in the limited condition and 35 health in the unregulated condition. Participants earned 108% more than this dishonesty benchmark in the limited condition, and they earned 52% more in the unregulated condition.

Discussion

In sum, we find that participants were more honest when unemployment benefits were limited compared to when they were unregulated. As a result, the social insurance institution performed better with limits, allowing participants to suffer less often from hardship and to earn more health and money. Thus, we find support for the resource management hypothesis that people can efficiently manage a limited supply of benefits, even when they could be tempted to claim benefits when employed.

Experiment 3

In Experiment 3, we examine another common policy tool: penalties. Many social insurance programs penalize people who claim benefits they are not eligible to receive. Two common penalties are fines and bans from receiving benefits in the future. Also, governments generally cannot perfectly check all claims so some false claims go undetected. We examine whether probabilistic fines and bans deter dishonesty, and whether one kind of penalty is more effective.

In general, penalties can deter misbehavior because they decrease the expected payoffs. However, penalties can also backfire (Bowles, 2008), such as when they lead to resentment, defiance, and retaliation instead of compliance (Nikiforakis, 2008). Hence, we examine to what extent fines and bans are effective deterrents of dishonesty, as well as whether they can also increase cheating.

Moreover, we test whether these factors differ for fines versus bans, potentially making one type of penalty more effective. Fines take money away from an offender which violates their sense of ownership, whereas bans prevent future benefits which the offender does not own. People especially value and defend what they own (DeScioli & Wilson, 2011; Kahneman, Knetsch, & Thaler, 1991), so we expect that a fine is perceived as more harsh than a ban with equivalent economic impact. As Adam Smith (1759/2010) observed: “To be deprived of that which we are possessed of, is a greater evil than to be disappointed of what we have only the expectation.” Hence, the additional harshness of a fine could potentially make it a more powerful deterrent; or, oppositely, a fine might provoke greater defiance than a ban.

Methods

We recruited undergraduates for a laboratory experiment (n = 128; 55% female; mean (S.D.) age of 19.8 (1.9) years. Participants earned $5 for showing up, and they earned additional money from the experimental game (mean = $20.98, S.D. = $5.87). The procedures and measures were the same as Experiment 2, except we modified the employment game to include fines or bans.

Participants played an eight-player unemployment game similar to Experiment 2, except a worker’s claims of unemployment are checked with a 1/3 probability. (When there are no funds available in the public fund, the claim is turned down without being checked.) If the claim was false, then the worker suffers the penalty.

Participants were assigned to the fine condition or ban condition. In the fine condition, the penalty for false claim is -15 health, which consists of -6 to pay back the benefits and an additional fine of -9. In the ban condition, the penalty prevents the worker from receiving benefits and bans them from claiming benefits in the next month. Importantly, we designed the amounts of the fine and ban so that the expected costs are roughly the same (about -5), holding constant the monetary incentives associated with the different penalties.

We examine to what extent these penalties deter false claims. We also assess whether the penalties can backfire by observing participants’ dishonesty immediately after a penalty. Finally, we examine whether fines, since they transgress ownership, are more potent deterrents than bans.

Results

Figure 9 shows participants' dishonest claims, hardship, and health in the fine and ban conditions. Participants were less dishonest when the penalty was a fine compared to a ban, although this difference was only marginally significant, $t(126) = 1.58$, $p = .059$. (This analysis excludes rounds when a participant was banned since they could not claim benefits, genuinely or falsely.) We also examine how many participants never lied at all. We find that more participants never lied in the fine condition (23%) than the ban condition (6%), $\chi^2(1) = 7.48$, $p < .01$. Additionally, participants suffered hardship less often with fines than bans, $t(126) = 8.27$, $p < .001$, and they earned more health with fines than bans, $t(126) = 6.86$, $p < .002$.

Unlike Experiment 2, we did not find a significant relationship between honesty and health in either condition (see Supporting Information S2); this is because the costs of the penalties meant that dishonesty was only slightly more profitable than honesty on average. We also analyzed whether dishonesty changed over rounds; a regression showed no change over time except in the final round of the ban condition when dishonesty spiked because there was no subsequent round for the ban to affect (see Supporting Information S2). This strategic endgame effect is similar to what we observed in Experiment 2, again showing participants' alert opportunism.

Overall, these results indicate that the fine was slightly more effective than the ban at deterring dishonesty, preventing hardship, and promoting economic health.

Figure 9. Mean (S.E.) Dishonesty, Hardship, and Health by the Penalty Condition

A. B. C.

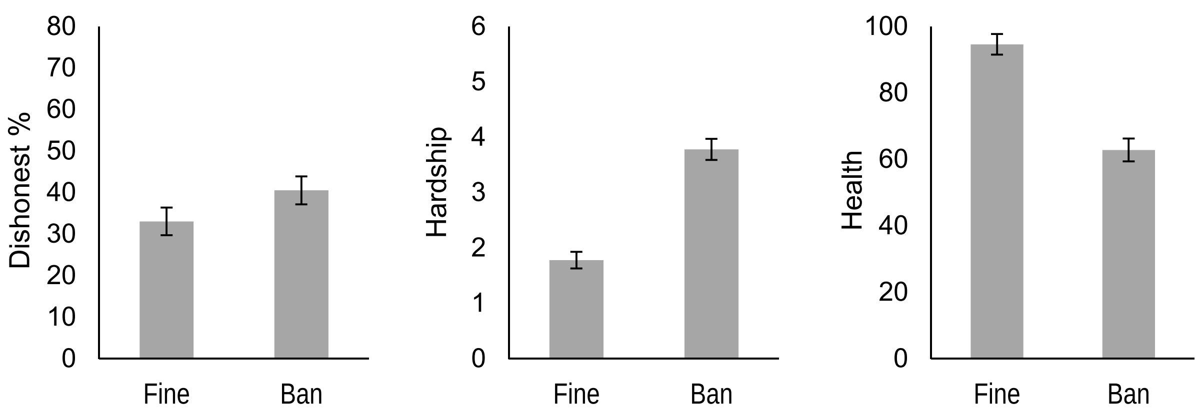

We further compare these cases to the institutions from Experiment 2 with unregulated and limited benefits. Compared to unregulated benefits, participants lied considerably less when there were fines, $t(126) = 7.23$, $p < .001$, and bans, $t(126) = 5.60$, $p < .001$, showing that both penalties reduced dishonesty. Compared to limits, fines and bans were similarly effective: rates of dishonesty did not differ from the limited condition for either penalty, ps > .10.

How did punishment affect dishonesty?

We next test how a penalty affected participants' dishonesty in the next round. Table 5 shows a logistic regression of a participant's dishonesty with predictors for the type of penalty, whether they were penalized in the previous round (lagged penalty), the interaction, and a random effect for participant. The results show that participants were *more* likely to lie after they were penalized in the previous round, and this effect did not differ by the type of penalty. Specifically, the model estimates that a participant's probability of lying increased after a fine from 32% to 64%, and it increased after a ban from 36% to 72%. This supports the hypothesis that people can react to penalties with defiance, rather than exclusively being deterred by their costs. Even so, these penalties did deter dishonesty overall and in the long-run (as shown by reduced dishonesty relative to the unregulated public fund) but the *immediate* effect of a penalty was an increase in dishonesty in the next round.

We also conducted additional analyses to address potential alternative explanations for increased dishonesty. One alternative is that participants thought that being audited in the current round meant they were less likely to be audited in the next round (despite the instructions stating that these chances were independent). A second is that participants were motivated to recover the health they lost as a result of the penalty. Briefly, our additional analyses did not support either of these alternatives (see Supporting Information S2).

Table 5. Dishonesty Predicted by Previous Penalties

	Coeff.	Std. Err.	P value
Ban Condition	0.24	(0.25)	0.333
Lagged Penalty	1.70	(0.47)	0.000
Ban X Lagged Penalty	0.22	(0.69)	0.747
Constant	-0.98	(0.18)	0.000

Note: Logistic regression with random effect for participant (n = 128 participants). n = 1,481 opportunities to lie.

Economic efficiency

We compare participants' earnings to what would happen if all players were always honest or if they were always dishonest. With complete honesty, players would earn an average of 102 health. Participants earned 8% less than this honesty benchmark in the fine condition, and they earned 38% less in the ban condition. With complete dishonesty, players would earn an average of 84 health in the fine condition and 5 health in the ban condition. Participants earned 12% more than this dishonesty benchmark in the fine condition, and they earned 1102% more in the ban condition. We point out that when players are dishonest, the ban leads to much less efficiency than the fine because the ban exposes a player to hardship in the next round without benefiting other players in the group, unlike the fine which funnels money back to the public fund.

Discussion

Overall, fines and bans were effective deterrents, restraining dishonesty at moderate levels (~30-40%). Fines appeared to be slightly more effective: Participants were less dishonest under fines than bans with only marginal significance, but more participants were completely honest with fines (23%) than bans (6%). Consequently, participants in the fine condition achieved greater earnings and suffered less hardship. These results support the hypothesis that fines are a more effective deterrent than (payoff-equivalent) bans because a fine directly takes the violator's money, transgressing their sense of ownership. We also found some evidence that penalties backfired in the short-term. Participants were more likely to lie immediately after being fined or banned, supporting the idea that participants resented being punished.

General Discussion

In three experiments, we examined the problem of dishonesty at the core of social insurance institutions. In Experiment 1, we found that participants were generally honest and helpful when they interacted repeatedly with the same partner. Despite the potential for dishonesty, participants were able to create and sustain a peer safety net to protect them from hardship when they did not find a job. However, when they interacted with a different individual each round, participants were dishonest and provided little help when asked—the peer safety net fell apart. These results point to a critical difference between small-scale communities versus large-scale anonymous societies: Social insurance is difficult to sustain without the binding force of reputation.

In Experiment 2, we introduced a public fund and a mandatory income tax, analogous to government social insurance. We found that when the public fund was unregulated, participants falsely claimed unemployment at high rates (67% of the time). But when the public fund was regulated with limited benefits, participants were substantially less dishonest (40%). This finding provides experimental evidence about the functioning of this common policy tool. Consistent with the common rationale for setting limits, participants were able to effectively manage their own finite supply of unemployment benefits, leading to greater honesty, health, and economic efficiency.

In Experiment 3, we found that audits with penalties also helped to deter dishonest claims of unemployment. The fine reduced dishonesty a little more than the ban, and generally, the fine yielded greater economic efficiency by channeling the fines back to the public fund where they could help other workers. However, fines and bans also triggered a backfire effect in which participants were more likely to lie immediately after being penalized.

More generally, we suggest that the unemployment game can provide a flexible framework for the experimental study of social insurance institutions. One key element of the game is the concept of hardship, which is an extra cost incurred when a player cannot pay their expenses, representing deprivations of basic necessities such as food, shelter, and medicine. Players can mutually gain by helping each other to avoid these additional costs, which captures the basic function of social insurance. Moreover, by preventing hardship, social insurance can promote economic efficiency. Traditional theories of redistribution, without a cost of hardship, assume that transfers of wealth from the rich to the poor necessarily reduce economic efficiency (Mankiw 2013; Okun 1975).

However, to promote efficiency, the rules of social insurance institutions need to discourage the employed from falsely claiming benefits. Policymakers face the challenge of determining the right form and amount of punishment for false claims. To maximize social welfare, the punishment should be set just high enough to be effective, but not so high that it inflicts more damage with no further benefits (Bentham, 1780/2000). While the current experiments examined the form of punishment, future research can examine the amounts. It is possible that lower penalties might be similarly effective if they reinforce the social expectation that someone should not be dishonest about their needs.

While some policies may promote social welfare better than others, another challenge for policymakers is gaining public support for a policy. Citizens have different beliefs about the role of government and who deserves benefits (e.g. Petersen 2012), which can make it difficult to develop policies that appeal to a majority. The unemployment game can be adapted to study public support by allowing players to vote for policy proposals about how the public fund collects taxes and distributes benefits.

Importantly, we used carefully controlled laboratory experiments to examine when individuals will misrepresent their needs. However, the social insurance programs that people rely in on their daily lives are immensely more complicated than the public fund institution in these experiments. Thus, the experiments presented here are necessarily simplifications of social insurance institutions to study the interaction between program design and program abuse. Inherently, our experiments do not map directly onto specific programs in the real world, but they do provide insight into common policy tools that reduce fraud.

In closing, we note that economies around the world are rapidly changing and so are the needs of workers. New technologies reshape the economy and make many jobs obsolete. Government institutions struggle to keep up with the quickly shifting needs of citizens. The present experiments illustrate how difficult problems such as evaluating workers' claims of need and preventing hardship can be studied in controlled laboratory experiments. The methods of experimental economics might eventually help design the social safety nets of the future.

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 Survey.” *Political Science Research and Methods* 5 (1): 31-53.

Alm, James. 2012. “Measuring, Explaining, and Controlling Tax Evasion: Lessons from Theory, Experiments, and Field Studies.” *International Tax and Public Finance* 19 (1): 54-77.

Alm, James, Betty R. Jackson, and Michael McKee. 1992. “Estimating the Determinants of Taxpayer Compliance with Experimental Data.” *National Tax Journal* 45 (1): 107-114.

Andreoni, James, and John H. Miller. 1993. “Rational Cooperation in the Finitely Repeated Prisoner’s Dilemma: Experimental Evidence.” *The Economic Journal* 103 (418): 570-585.

Balliet, Daniel, Laetitia B. Mulder, Paul A.M. Van Lange. 2011. “Reward, Punishment, and Cooperation: A Meta-Analysis.” *Psychological Bulletin* 137 (4): 594-615.

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.

Bendor, Jonathan. 1993. "Uncertainty and the Evolution of Cooperation." *Journal of Conflict Resolution* 37 (4): 709-734.

Bentham, Jeremy. 2000. *Principles of Morals and Legislation*. Kitchner, Ontario, Canada: Batoche Books. (Original work published 1780).

Bowles, Samuel. 2008. "Policies Designed for Self-Interested Citizens May Undermine "the Moral Sentiments": Evidence from Economic Experiments." *Science*, 320 (5883): 1605-1609.

Boyd, Robert, and Peter J. Richerson. 1988. "The Evolution of Reciprocity in Sizable Groups." *Journal of Theoretical Biology* 132 (3): 337-356.

Büchner, Susanne, Giorgio Coricelli, and Ben Greiner. 2007. "Self-Centered and Other-Regarding Behavior in the Solidarity Game." *Journal of Economic Behavior and Organization* 62 (2): 293-303.

Charness, Gary and Garance Genicot. 2009. "Informal Risk Sharing in an Infinite-Horizon Experiment." *The Economic Journal* 119 (537): 796-825.

Coate, Stephen and Martin Ravallion. 1993. "Reciprocity Without Commitment: Characterization and Performance of Informal Insurance Arrangements." *Journal of Development Economics* 40 (1): 1-24.

Cosmides, Leda. 1989. "The Logic of Social Exchange: Has Natural Selection Shaped How Humans Reason? Studies with the Wason Selection Task." *Cognition* 31(3): 187-276.

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.

DeScioli, Peter and Bart J. Wilson. 2011. "The Territorial Foundations of Human Property." *Evolution and Human Behavior* 32(5): 297-304.

DeSilver, Drew. 2017. "What Does the Federal Government Spend Your Tax Dollars on? Social Insurance Programs, Mostly." Pew Research Center, Washington D.C. Retrieved from: www.pewresearch.org/fact-tank/2017/04/04/what-does-the-federal-government-spend-your-tax-dollars-on-social-insurance-programs-mostly/.

Esarey, Justin, Timothy C. Salmon, and Charles Barrilleaux. 2012. "Social Insurance and Redistribution in a Laboratory Experiment." *Political Research Quarterly* 65 (3): 685-698.

Friedland, Nehemiah, Shlomo Maital, and Aryeh Rutenberg. 1978. A Simulation Study of Income Tax Evasion. *Journal of Public Economics* 10 (1): 107-116.

Fudenberg, Drew, David G. Rand, and Anna Dreber. 2012. "Slow to Anger and Fast to Forgive: Cooperation in an Uncertain World." *American Economic Review* 102 (2): 720-749.

Gerber, Alan S. and Donald P. Green. 2012. *Field Experiments: Design, Analysis, and Interpretation*. New York, NY: W.W. Norton.

Gilens, Martin. 1999. *Why Americans Hate welfare: Race, Media, and the Politics of Antipoverty Policy*. Chicago, IL: University of Chicago Press.

Gurven, Michael. 2004. "To Give and to Give Not: The Behavioral Ecology of Human Food Transfers." *Behavioral and Brain Sciences* 27 (4): 543-559.

Kagel, John H. and Alvin E. Roth. 2016. *The Handbook of Experimental Economics, Volume 2*. Princeton, NJ: Princeton University Press.

Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler. 1991. "Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias." *Journal of Economic Perspectives* 5(1): 193-206.

Kirchler, Erich, Stephan Muehlbacher, Barbara Kastlunger and Ingrid Wahl. 2010. "Why Pay Taxes? A Review of Tax Compliance Decisions." *Developing Alternative Frameworks for Explaining Tax Compliance* 15-31.

Larrimore, Jeff, Sam Dodini, and Logan Thomas. 2016. *Report on the Economic Well-Being of U.S. Households in 2015.* Washington: Board of Governors of the Federal Reserve System. Retrieved from: https://www.federalreserve.gov/2015-report-economic-well-being-us-households-201605.pdf.

Mankiw, N. Gregory. 2013. "Defending the One Percent." *Journal of Economic Perspectives* 27 (3): 21-34.

Moene, Karle Ove and Michael Wallerstein. 2001. "Inequality, Social Insurance, and Redistribution." *American Political Science Review* 95 (4): 859-874.

Morudch, Jonathan. 1999. "Can Informal Insurance Patch the Safety Net?" *The World Bank Observer* 14 (2): 187-207.

Morton, Rebecca B. and Kenneth C. Williams. 2010. *Experimental Political Science and the Study of Causality: From Nature to the Lab.* Cambridge: UKL Cambridge University Press.

Nikiforakis, Nikos. 2008. "Punishment and Counter-punishment in Public Goods Games: Can We Really Govern Ourselves?." *Journal of Public Economics* 92 (1-2): 91-112.

Ockenfels, Axel and Joachim Weimann. 1999. "Types and Patterns: An Experimental East-West-German Comparison of Cooperation and Solidarity." *Journal of Public Economics* 71 (2): 275-287.

Okun, Arthur M. 1975. *Equality and Efficiency: The Big Tradeoff*. Washington, D.C.: Brookings Institution Press.

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

Ostrom, Elinor. 1998. "A Behavioral Approach to the Rational Choice Theory of Collective Action: Presidential Address, American Political Science Association, 1997." *American Political Science Review* 92 (1): 1-22.

Palfrey, Thomas R. 2009. "Laboratory Experiments in Political Economy." *Annual review of Political Science* 12: 379-388.

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.

Rehm, Phillipp, Jacob S. Hacker, and Mark Schlesinger. 2012. "Insecure Alliances: Risk, Inequality, and Support for the Welfare State." *American Political Science Review* 106 (2): 386-406.

Selten, Reinhard, and Axel Ockenfels. 1998. "An Experimental Solidarity Game." *Journal of Economic Behavior and Organization* 34 (4): 517-539.

Shiue, Carol H. 2005. “The Political Economy of Famine Relief in China 1740-1820.” *Journal of Interdisciplinary History* 36 (1): 33-55.

Smith, Adam. 2010. *The Theory of Moral Sentiments*. London, UK: Penguin. (Original work published 1759.).

Smith, Vernon L. 1982. “Microeconomic Systems as an Experimental Science.” *American Economic Review* 72 (5): 923-955.

Spicer, Michael W. and J. E. Thomas. 1982. “Audit Probability and Tax Evasion Decision: An Experimental Approach.” *Journal of Economic Psychology* 2 (3): 241-245.

Trhal, Nadja and Ralf Radermacher. 2009. “Bad Luck vs. Self-Inflicted Neediness- An Experimental Investigation of Gift Giving in a Solidarity Game.” *Journal of Economic Psychology* 30 (4): 517-526.

Trivers, Robert L. 1971. “The Evolution of Reciprocal Altruism.” *The Quarterly Review of Biology* 46 (1): 35-57.

United States Department of Labor. 2016. *Benefit Accuracy Measurement Annual Report.* Retrieved from: https://oui.doleta.gov/unemploy/bam/2016/IPIA_2016_Benefit_Accuracy_Measurement_Annual_Report.pdf.

van Oorschot, Wim. 2006. “Making the Difference in Social Europe: Deservingness Perceptions Among Citizens of European Welfare States. *European Journal of Social Policy* 16 (1): 23-42.