Social Choice to Solve Social Dilemmas

Lawrence R. De Geest[a,b] and John H. Miller[c,d]

[a]Economics, Suffolk University, 8 Ashburton Place, Boston, MA. 02108

[b]Lyft, 185 Berry Street Suite 400, San Francisco, CA. 94107

[c]Social and Decision Sciences, Carnegie Mellon University, Pittsburgh, PA. 15213

[d]Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501

Abstract

Social systems often rely on institutions to solve problems such as social dilemmas. Here, we explore whether agents in these systems can use a social-choice mechanism to self-organize an institutional solution to a social dilemma. We find that, notwithstanding the existence of productive institutions, such self-organization can fail. Systems composed of rational agents succumb to a proliferation of possible equilbiria, many of which lead to inferior outcomes. Systems composed of small groups of adaptive agents are able to evolve effective solutions, but as group size increases such solutions become far less likely. Given that social dilemmas lie at the heart of some of the most challenging social problems, such as global warming, these results suggest that alternative approaches to creating and maintaining useful institutional solutions may be needed.

1 Introduction

Social systems often confront conditions whereby an individual's self-interest is at odds with society's interest. Such social dilemmas arise in the context of many social challenges, including pollution, resource extraction, epidemic response, and the provision of public goods. Evidence from both small-scale societies and public goods experiments show that social dilemmas may be resolved through the emergence of ad hoc institutions that rely on peer enforcement to align individual with social incentives (Ostrom, 1990; Bowles and Gintis, 2011; Fehr and Schurtenberger, 2018; De Geest and Kingsley, 2021). There is no guarantee that such institutions will emerge and, if they do, that they will be socially productive (Abbink et al., 2017). Here, we explore the difficulty of developing socially productive institutions using prototypical social choice mechanisms (SCMs) akin to decentralized voting.

The question we explore is whether social dilemmas can be resolved using a penalty-based institution where the key parameters of that institution are derived by allowing individuals to "vote" using various prototypical social choice mechanisms (SCMs). SCMs aggregate a set of individual choices into a single social choice that is imposed on all of the individuals in the group. Arrow (1950) showed that there is no SCM that can support a small set of seemingly reasonable desiderata. Notwithstanding this result, SCMs are often used for social choices—the ones we analyze below are motivated, in part, by current practice (Fehr and Williams, 2018). For example, international climate change agreements are often derived using SCM-equivalent mechanisms that allow any party to veto a potential agreement (Finus et al., 2005; Lal Pandey, 2014).

Here, we explore the resolution of a social dilemma through the implementation of an institution that is determined by the choices, mediated by a SCM, of the individuals exposed to the dilemma. This approach is inspired by Ostrom (1990)'s work on "new institutionalism" and the evolution of self-organizing institutions to solve collective action problems in common pool resources. The work here provides an alternative framework by which to explore such phenomena.

We consider systems of agents from both a rational and adaptive perspective.

The rational analysis explores the Nash equilibria associated with this system. The results of this analysis demonstrate that the successful resolution of a social dilemma may be more difficult than intuition suggests. We also explore a dynamic system whereby agents adapt their behavior driven by indirect forces that impact the system on evolutionary time scales. We find that under some conditions this adaptive system can discover institutions that effectively solve the underlying social dilemma. The above results replicate some key results from the experimental literature on public goods with peer punishment, where subjects do not always converge on socially optimal institutions (Chaudhuri, 2011).

2 Public Goods, Institutions, and Social Choice Mechanisms

We focus our analysis on a simple linear public goods game that captures the fundamental notion of an N-person social dilemma (Fehr and Gächter, 2000). In this game, each of N agents must allocate an individual endowment between a private and public good (below, we normalize the total endowment each agent must allocate to 1.0). Each unit of the endowment held in the private good returns a value of 1 to the agent, while any endowments contributed to the public good get aggregated and α times this total is given to each agent, regardless of that agent's contribution to the public good. Thus, the payoff to each agent is

$$\pi_i = (e_i - c_i) + \alpha \left(c_i + \sum_{i \neq j} c_j \right)$$

where e_i is agent i's endowment and c_i is their contribution to the public good.

This scenario leads to a social dilemma when $1/N < \alpha < 1$, with the lower bound ensuring the potential of a social gain when everyone contributes to the public good (since $N\alpha > 1$) and the upper bound ensuring that a selfish individual has the incentive to never contribute to the public good (since $\alpha < 1$). Thus, if agents only care about their own payoffs, then each contributes nothing to the public good, implying that all of the agents are worse off than they would have been had they contributed everything to the public good.

One way to resolve the above social dilemma is to establish an institution that realigns individual incentives so that each agent wants to contribute to the public good. Such institutions range from altering the preferences of the agents to be prosocial ("it takes a village") to imposing a penalty for asocial behavior.

Below we explore an institution that imposes a penalty, P, on any agent that does not contribute at least a threshold amount, T, to the public good. Such a Threshold-Penalty institution (TPI) can, if appropriately designed, alter an agent's incentives enough so that fully contributing to the public good becomes incentive-compatible (Chaudhuri, 2011). Under a TPI, the payoff to agent i is

$$\pi_i = (e_i - c_i) + \alpha \left(c_i + \sum_{i \neq j} c_j \right) - \phi(c_i < T)P$$

where $\phi(c_i < T)$ is equal to 1 if $c_i < T$ and 0 otherwise.

Given that any contribution to the public good is costly to the agent (since $\alpha < 1$), in the presence of a TPI an agent considers only two possible behaviors: either contribute nothing or T to the public good. If an agent contributes nothing it may fail to meet the threshold and would be subject to the penalty. If the agent contributes the threshold it avoids the penalty, but loses $(1 - \alpha)T$ by paying the threshold versus adding T to its private good holdings. Thus, an agent's optimal contribution, c^*, is

$$c^* = \begin{cases} 0 & \text{if } (1 - \alpha)T > P, \\ T & \text{otherwise.} \end{cases}$$

To solve the social dilemma, we need $c^* = e_i$, which occurs when the TPI sets $T = e_i$ and $P \geq (1 - \alpha)T$.

2.1 Social Choice Mechanisms (SCM)

While the introduction of a suitable TPI seems like an easy solution[1] to the social dilemma, the question we explore below is the difficulty of implementing such an institution. If the social system is controlled by, say, a benevolent overlord or if there is an easy way to form binding contracts among the agents, then creating a useful TPI is straightforward.

A far more common means by which to create social policy is to allow the agents within the social system itself to self-organize using some form of a voting system. Could such a mechanism form a TPI that can resolve the social dilemma facing these agents?

To analyze this issue, we consider a system in which N agents "vote" on the TPI that will be imposed on the group by having each agent submit its preferred T and P choices to a SCM. The SCM then aggregates these two sets of submitted values into the TPI's final threshold and penalty values. Given this TPI, agents then decide on their public good contributions using the previously derived rule.

We explore four prototypical SCMs: MEAN, MEDIAN, MAX, and MIN. Each SCM outputs a single value based on a set of N inputs. MEAN (the mean of the inputs) and MEDIAN (the median of the inputs) were chosen as they both represent a notion of averaging the inputs, though they differ in terms of the impact of outliers on the average. MEDIAN also captures the importance of the median voter in certain types of elections (Black, 1948; Downs, 1957). MAX and MIN are SCMs that output either the maximum (MAX) or minimum (MIN) of the inputs, and thus each pushes the social choice toward an extreme. Both MIN and MAX can be very sensitive to their inputs, since any agent in the group has the ability to alter the social choice in the extreme direction.

[1] Below we ignore any difficulties associated with using a TPI, such as identifying and, if necessary, penalizing each agent's public contribution.

2.2 Nash equilibria

We first analyze the behavior of rational agents given by the Nash equilibria of the system. In the simple public goods game, a selfish agent treats the contributions to the public good of the other agents as exogenous, and therefore maximizes its own returns by holding on to its full endowment. In our modified public goods game, each agent first submits a desired threshold and penalty level to the SCM. The SCM then generates a TPI for the group, and each agent then contributes its payoff-maximizing amount to the public good. Insofar as the agent can influence the TPI, the contributions to the public good of the other agents are no longer exogenous, and the game becomes more strategically complex.

One new element of the strategy in this game is an agent's ability to understand its influence on the SCM. If the agent knows the values that the other agents will submit to the SCM—perhaps a difficult task depending on the size of the group and the flow of information—then the agent can calculate the exact influence of its submitted values on the TPI *ceteris paribus*. Even with more limited information, agents may have some insights into their influence on the TPI. For example, under MIN (MAX), an agent knows that if it submits a value lower (higher) than the current social value, it will be pivotal in the social choice. Under MEAN and MEDIAN, knowing the current social choice is sufficient to predict at least the potential direction of change.

Once an agent recognizes its ability to alter the TPI, it should take the (now *endogenous*) reaction of the other agents to that new TPI into account. In the game without a TPI, an agent always wants to contribute nothing to the public good, so blindly transferring such reasoning into the new game would imply a desire for TPIs that do not compel public good contributions. However, more sophisticated agents should recognize that, since all agents optimize their contributions given the announced TPI, having a TPI that forces everyone, including itself, to contribute to the public good could lead to a higher payoff to the agent (as well as the other agents).

In our analysis of the Nash equilibria, we assume that our rational agents have full knowledge about how their submissions to the SCM will alter the TPI *ceteris paribus* and that agents fully incorporate how that TPI will influence the public good contributions of the other agents. These assumptions should promote the ability of the rational system to implement TPIs that can resolve the social dilemma, thus providing a nice benchmark for this system.[2]

To simplify the rational analysis we restrict both the number of agents and the strategy space. We assume a system with five agents and constrain their choices to L possible, evenly spaced values between 0 and 1: $\{0.0, 1/(L-1), 2/(L-1), \ldots, 1.0\}$. For example, if $L = 3$ agents can make contributions, or submit threshold and penalty values to the SCM, on $\{0.0, 0.5, 1.0\}$. Since we constrain possible contributions to discrete increments, agents must choose the best contribution given the constrained set, which may differ from the previously derived c^*. Therefore, agents will contribute either zero or the lowest contribution possible consistent with meeting the threshold. Each agent has L^2 possible strategies that give all the various combinations of T and P that can be submitted to the SCM. This implies that with N agents, the payoff matrix will have L^{2N} possible elements. Thus, with $L = 3$ and $N = 5$, each agent must choose among 9 possible strategies resulting in $59,049$ possible combinations of strategies across the five agents. The combinatorics quickly escalate as either L or N increases.

Table 1 gives the best-response contribution for an agent facing the indicated TPI in a group of five agents when $\alpha = 0.5$ and contributions are constrained to $\{0.0, 0.5, 1.0\}$. As seen in the table, the best responses break into various blocks. When either T or P are zero, agents wish to contribute zero as either the threshold is non binding, or it binds but the penalty is zero. The horizontal boundaries between the blocks are driven by the discrete increments of possible contributions, for example, the best an agent can do to meet a threshold in the range of 0.1–0.5, is to contribute 0.5 (similarly, a contribution of 1.0 is required to meet thresholds between 0.6–1.0). The vertical boundaries to the left of each block are tied to α. When $\alpha = 0.5$, contributing 0.5 (1.0) to avoid violating a threshold is payoff improving as long as the penalty is above 0.25 (0.5). As α increases, the critical penalty values

[2] Obviously, assumptions that lessen this possibility could also be considered.

decrease and the boundaries move to the left. There may be TPIs where agents are indifferent between two contribution levels, for example, this occurs in Table 1 for thresholds between 0.6–1.0 and a penalty of 0.5.

Table 1: Best Response Contributions. Agents are in a group of five facing a TPI with the given threshold (row) and penalty (column) values, and are constrained to contribute either 0.0, 0.5, or 1.0. The interior of the table gives an agent's best response contribution (coded as $a = 0.0$, $b = 0.5$, and $c = 1.0$) when $\alpha = 0.5$.

	0.00	0.10	0.20	0.30	0.40	0.50	0.60	0.70	0.80	0.90	1.00
0.00	a–	a–	a–	a–	a–	a–	a–	a–	a–	a–	a–
0.10	a–	a–	a–	-b-	-b-	-b-	-b-	-b-	-b-	-b-	-b-
0.20	a–	a–	a–	-b-	-b-	-b-	-b-	-b-	-b-	-b-	-b-
0.30	a–	a–	a–	-b-	-b-	-b-	-b-	-b-	-b-	-b-	-b-
0.40	a–	a–	a–	-b-	-b-	-b-	-b-	-b-	-b-	-b-	-b-
0.50	a–	a–	a–	-b-	-b-	-b-	-b-	-b-	-b-	-b-	-b-
0.60	a–	a–	a–	a–	a–	a-c	–c	–c	–c	–c	–c
0.70	a–	a–	a–	a–	a–	a-c	–c	–c	–c	–c	–c
0.80	a–	a–	a–	a–	a–	a-c	–c	–c	–c	–c	–c
0.90	a–	a–	a–	a–	a–	a-c	–c	–c	–c	–c	–c
1.00	a–	a–	a–	a–	a–	a-c	–c	–c	–c	–c	–c

A Nash equilibrium in this system occurs when, for each of the five agents, unilaterally altering the values it sends to the SCM does not result in a TPI that leads to a higher payoff to that agent once every agent optimizes its contribution given that TPI. Table 2 provides some key properties of the Nash equilibria in this system across the four SCMs and various α values. Depending on the conditions, 40–89% of the cells in the payoff matrix represent Nash equilibria, with MEDIAN and MEAN near the lower end and MIN and MAX near the upper end of that range. All of MAX's equilibria solve the free rider problem with contributions of 1.0. Under MAX, any agent has the ability to drive the threshold and penalty values to 1.0 and, given that the resulting TPI will maximize every agent's public good contributions, doing so becomes a dominant strategy. MIN has 89% of its equilibria leading to contributions of 0.0. Under MIN, a single agent can only drive the threshold and penalty values down and, in essence, veto a binding TPI. Thus, MIN can easily get trapped in best-response wells characterized by TPIs with low threshold or penalty values that do not compel contributions. When $\alpha < 0.5$, only one of MIN's 52,489 possible equilibria supports full contributions (with a threshold and penalty of 1.0) and when $\alpha \geq 0.5$, 32 such equilibria emerge (with a threshold of 1.0 and penalties of either 1.0 or 0.5). Almost all of MEAN's equilibria result in contributions of 0.5 when $\alpha = 0.3$, and as α increases past 0.6 between 61–66% of the equilibria support full contributions. MEDIAN's equilibria result in a modal contribution of 0.5 across the range of α, with some of this mass shifting to support full contributions as α increases (constituting around 34% of the equilibria when $\alpha \geq 0.5$). Around 20% of MEDIAN's equilibria imply contributions of 0.0.

Table 2: Key properties of the Nash equilibria in a simple system with five agents and three strategy levels, under various αs and SCMs. The SCM used to derive the threshold and penalty levels given the submitted values is in the first column, the percent of the $59,049$ possible strategic combinations that result in a Nash equilibrium is given in the second column, the expected mean per-capita contribution assuming each possible equilibrium is equally likely is shown in the third column, and the final column gives the frequency distribution of the mean per-capita contributions (by percent of the total equilibria).

SCM	%Nash	Exp C	Frequency of Contributions
$\alpha = 0.3$			
MAX	75.4%	1.00	1.0 (100.0%)
MEAN	39.5%	0.54	0.0 (1.1%), 0.5 (90.3%), 1.0 (8.6%)
MEDIAN	45.9%	0.44	0.0 (21.8%), 0.5 (68.6%), 1.0 (9.6%)
MIN	89.1%	0.01	0.0 (98.1%), 0.5 (1.9%), 1.0 (0.0%)
$\alpha = 0.4$			
MAX	75.4%	1.00	1.0 (100.0%)
MEAN	41.2%	0.68	0.0 (1.0%), 0.5 (61.2%), 1.0 (37.8%)
MEDIAN	45.9%	0.44	0.0 (21.8%), 0.5 (68.6%), 1.0 (9.6%)
MIN	89.1%	0.01	0.0 (98.1%), 0.5 (1.9%), 1.0 (0.0%)
$\alpha = 0.5$			
MAX	86.5%	1.00	1.0 (100.0%)
MEAN	43.7%	0.77	0.0 (0.9%), 0.5 (44.3%), 1.0 (54.7%)
MEDIAN	48.5%	0.58	0.0 (18.2%), 0.5 (47.6%), 1.0 (34.2%)
MIN	88.9%	0.01	0.0 (98.4%), 0.5 (1.6%), 1.0 (0.1%)
$\alpha = 0.6$			
MAX	86.5%	1.00	1.0 (100.0%)
MEAN	51.4%	0.83	0.0 (0.1%), 0.5 (39.3%), 1.0 (60.7%)
MEDIAN	48.5%	0.58	0.0 (18.2%), 0.5 (47.6%), 1.0 (34.2%)
MIN	88.9%	0.01	0.0 (98.4%), 0.5 (1.6%), 1.0 (0.1%)
$\alpha = 0.7$			
MAX	86.5%	1.00	1.0 (100.0%)
MEAN	56.2%	0.82	0.0 (0.1%), 0.5 (35.7%), 1.0 (64.2%)
MEDIAN	48.5%	0.58	0.0 (18.2%), 0.5 (47.6%), 1.0 (34.2%)
MIN	88.9%	0.01	0.0 (98.4%), 0.5 (1.6%), 1.0 (0.1%)
$\alpha = 0.8$			
MAX	86.5%	1.00	1.0 (100.0%)
MEAN	59.0%	0.83	0.5 (34.7%), 1.0 (65.3%)
MEDIAN	48.5%	0.58	0.0 (18.2%), 0.5 (47.6%), 1.0 (34.2%)
MIN	88.9%	0.01	0.0 (98.4%), 0.5 (1.6%), 1.0 (0.1%)
$\alpha = 0.9$			
MAX	86.5%	1.00	1.0 (100.0%)
MEAN	59.8%	0.83	0.5 (34.2%), 1.0 (65.8%)
MEDIAN	48.5%	0.58	0.0 (18.2%), 0.5 (47.6%), 1.0 (34.2%)
MIN	88.9%	0.01	0.0 (98.4%), 0.5 (1.6%), 1.0 (0.1%)

Of central concern to this research is the ability of the social system to solve the underlying social dilemma. A key measure of this ability is given by the mean per-capita contribution of the agents. Table 3 gives the expected mean per-capita contribution across the various SCMs and αs, assuming that each Nash equilibrium is equally likely. Based on this measure, the SCMs can be ranked with MAX ¿ MEAN ¿ MEDIAN ¿ MIN, with mean per-capita contributions of 1.0, 0.54-0.83, 0.44–0.58, and 0.01 respectively (with the higher values in these ranges being associated with larger αs). Thus, assuming rational agents being driven to Nash equilibria, the outcome of this self-organizing institutional system is dependent on the SCM used to generate the TPI, with MAX completely solving the free riding problem, MEAN and MEDIAN producing partial solutions, and MIN succumbing to the social dilemma with complete free riding in the vast majority of outcomes.

Table 3: Expected mean per-capita contribution of Nash equilibria by SCM and α. The expected mean per-capita contribution is calculated over the different SCM's (rows) and αs (columns) possible Nash equilibria assuming each equilibrium is equally likely. The system is composed of five agents, constrained to contribution, threshold, and penalty values on $\{0.0, 0.5, 1.0\}$. The TPI is determined by the outcome of the SCM given each agent's submitted values for the threshold and penalty. Agents then optimize their contributions given the resulting TPI.

	0.30	0.40	0.50	0.60	0.70	0.80	0.90
MAX	1.000	1.000	1.000	1.000	1.000	1.000	1.000
MEAN	0.538	0.684	0.769	0.803	0.821	0.826	0.829
MEDIAN	0.439	0.439	0.580	0.580	0.580	0.580	0.580
MIN	0.009	0.009	0.009	0.009	0.009	0.009	0.009

The different performance of the SCMs in preventing free riding is tied to a few factors. The performance of MAX and MIN at the extremes of possible contributions is not surprising given that these two SCMs tend to ratchet the resulting social choice in the extreme direction that each promotes. Thus, MAX produces binding TPIs that require high contributions, while MIN typically produces non-binding TPIs that discourage contributions. MEAN and MEDIAN result in outcomes that tend to occupy the middle range of potential contributions. MEAN performs better than MEDIAN for any given α, with MEDIAN being confined to a narrower band of moderate outcomes between 0.44–0.58 than MEAN, which ranges ranges from 0.54–0.83. The different behavior that these two mechanisms engender in this system is somewhat surprising, as inputs are restricted to $\{0.0, 0.5, 1.0\}$, so the impact of possible outliers, a key difference between these two mechanisms, is limited. However, while the impact of outliers is limited, these latter two SCMS can generate very different outputs even when the inputs are restricted.

Table 4 shows the frequency of particular outputs generated by a SCM when facing the ensemble of all possible input patterns of size five over three levels. Three of the SCMs[3] are constrained to output values drawn from $\{0.0, 0.5, 1.0\}$, with MIN (MAX) having roughly 87% of its values at 0.0 (1.0) and 13% at 0.5, and MEDIAN having 58% of its outputs at 0.5, with the remainder being spread equally across the two extremes. MEAN has a more diffuse distribution across $\{0.0, 0.1, \ldots, 0.9, 1.0\}$ centered at 0.5.

Table 4: The frequency of outputs generated by MEAN, MEDIAN, MAX, and MIN, given all possible combinations of inputs (constrained to either 0.0, 0.5, or 1.0) across five agents.

Output	MAX	MEAN	MED	MIN
0.00	1	1	51	211
0.10	0	5	0	0
0.20	0	15	0	0
0.30	0	30	0	0
0.40	0	45	0	0
0.50	31	51	141	31
0.60	0	45	0	0
0.70	0	30	0	0
0.80	0	15	0	0
0.90	0	5	0	0
1.00	211	1	51	1

Based on the data in Table 4, the likelihood that a given contribution violates[4] the threshold set by a SCM is shown in Table 5. Of particular interest here is the

[3]This holds for MEDIAN as long as the number of agents is odd.

[4]Given that contributions are confined to $\{0.0, 0.5, 1.0\}$, when MEAN produces, say, a threshold of 0.1, agents that want to avoid the penalty must contribute at least 0.5.

difference between MEAN versus MEDIAN. MEAN results in a higher likelihood of violation for a given contribution level than MEDIAN does, which is consistent with the observation that MEAN results in TPIs that encourage higher contributions relative to MEDIAN. Also, the wider variety of possible penalty values produced by MEAN can also be important, as an agent's willingness to violate a given threshold can be sensitive to small differences in penalties.

Table 5: Frequency of a given contribution violating the threshold level set by a given SCM assuming a random set of five, three-level inputs.

Contribution	MAX	MEAN	MEDIAN	MIN
0.0	99.6%	99.6%	79.0%	13.2%
0.5	86.8%	39.5%	21.0%	0.4%

The analysis of Nash equilibria in this system suggests that solving the public goods social dilemma by using a decentralized SCM may be more difficult than anticipated, even when there exists a Pareto superior TPI. The system is rife with equilibria and, depending on the SCM, many of these equilibria result in Pareto inferior outcomes. The two SCMs that yield to extreme values, not surprisingly, reinforce such extremes and either encourage social maximization in the case of MAX or result in free-riding in the case of MIN. The two SCMs more akin to democratic voting, MEAN and MEDIAN, favor intermediate levels of contributions, and even though they are structurally similar, they produce relatively different outcomes as α increases.

Nash equilibrium entail a static analysis of the system, so having a more dynamic analysis of this system may be useful. A dynamic analysis may still result in the system settling on a Nash equilibrium, though it might favor particular equilibria versus the analysis above that assumed they were equally likely—equilibrium selection could potentially resolve the social dilemma. Moreover, the dynamics could exhibit stable points that are not Nash equilibrium.[5] A dynamic system also provides an oppor-

[5] One way to ameliorate a social dilemma is to develop reciprocity by repeatedly playing the game with the same opponents (Fehr and Gächter, 2000). In the dynamic model here, agents are randomly mixed together for each game neutralizing this avenue to cooperation.

tunity to lessen our assumptions about each agent's cognitive abilities and access to information. Finally, a dynamic model allows us to analyze larger-dimensional systems that may exhibit different patterns than lower-dimensional ones. Given these opportunities, we next consider a dynamic system driven by adaptive agents.

2.2.1 Adaptive Agents

To explore the behavior of a more dynamic system, we consider adaptive agents. To model adaptive agents we use a simple genetic algorithm (Holland, 1975). Genetic algorithms use analogs of natural evolution to drive the artificial evolution of a population of agents. Thus, each agent in a population is exposed to a problem and receives a payoff (analogous to fitness in natural systems). The agents are then evolved using reproduction by performance (selection), that is, agents with higher payoffs are more likely to be reproduced. Members of the new population are subject to some variation via "genetic operators," in this case, a simple mutation of the strategic parameters. Genetic algorithms tend to identify good solutions to the underlying problem being explored and are also relatively robust to systems with nonlinear, noisy, and complex search spaces.

The genetic algorithm used below has the following form. A population of 100 agents must each adapt two strategic parameters that determine the threshold and penalty choices that the agent will submitted to the SCM. As in the analysis of the Nash equilibria above, agents are constrained to L levels of threshold and penalty values (as well as contribution amounts) in discrete increments of $1/(L-1)$ on $[0.0, 1.0]$. At the start of the algorithm each agent's two parameters are randomized on the stated interval. The population of agents is then evolved for twenty-five generations.[6] At the start of each generation, agents are randomly drawn from the population (without replacement) and put into groups of size five. Each group then plays the public goods game defined above, with each agent submitting its threshold and penalty values to the SCM. Once the TPI is determined, each agent makes its best response contribution to the public good and receives its payoff.[7] Agents

[6] Adding more generations to the evolution did not significantly alter the results.

[7] We found that the evolving agents could easily adapt to contributing the optimal contribution

participate in twenty such games during a given generation and accumulate the payoffs from each of these twenty games.

At the end of each generation, a new population of 100 agents is created using selection and variation. For selection, two agents are randomly drawn (with replacement) from the population, and a copy of the one with the higher payoff is placed into the new population. This tournament selection process is performed 100 times, resulting in 100 agents reproduced from the old population biased by performance. To introduce some variation, each agent in the new population has a 25% chance of undergoing mutation. If an agent is mutated, one of its strategic parameters (either the threshold or penalty) was randomly chosen and altered by a uniformly chosen mutation amount drawn from $\{-1/(L-1), 0, 1/(L-1)\}$ (if the resulting parameter is pushed outside of $[0,1]$ it is reset to the nearest boundary). Once the new population is finalized, a new generation begins and the process above is iterated. Data is collected in the twenty-fifth generation of each experiment, and each experiment is repeated for 10,000 trials.

Table 6 gives the mean of the per-capita contribution levels observed in each group in the twenty-fifth generation of the algorithm across all 10,000 trials.[8] This table is the adaptive analog to the Nash equilibria results given in Table 3. The obvious difference between these two approaches is that the adaptive system results in MEAN, MEDIAN, and MIN leading to much higher contributions to the public good. Adaptation nearly solves the free rider problem across all levels of α under the MEAN, MEDIAN, and MAX SCMs, whereas before this only occurred under MAX, with MEAN and MEDIAN leading to intermediate contribution levels.[9]

---

when facing a fixed T and P.

[8]In the adaptive system, each agent evolved the threshold and penalty values that it submitted to the SCM. Across all levels of α, the mean of the MAX, MEAN, and MEDIAN values of the threshold parameter were between 0.83–0.95, with MEDIAN at the top of that range, followed by MEAN, and then MAX. The penalty parameters for these three SCMs bifurcated depending on α. For $\alpha \leq 0.5$ the range of penalty values was similar to those of the threshold values, though both MEAN and MEDIAN were near the upper end and MAX was near the lower end of the range. For $\alpha > 0.5$, the range of penalty values was between 0.62–0.84, with a similar pattern of the three SCMs as before. MIN had thresholds around 0.67 for $\alpha \leq 0.5$ and around 0.98 for higher αs, with penalties between 0.79–0.82 across all levels of α.

[9]In the Nash analysis, all of the equilibria for MAX had full contributions, while in the adaptive

Recall that under MIN, out of the 52,489 possible equilibria, only 1 (for $\alpha < 0.5$) or 32 (for $\alpha \geq 0.5$) support full contributions, with over 98% of the remaining equilibria leading to complete free riding, implying an expected contribution of about 0.01. However, under adaptation, MIN achieves contributions of around 40% when $\alpha \leq 0.5$ and around 82% when $\alpha > 0.5$.[10]

Table 6: Expected mean per-capita observed contribution of adaptive agents by SCM and α. The behavior of the agents is governed by the genetic algorithm described in the text. The per-capita public good contributions made across all of the groups in the twenty-fifth generation of the algorithm were collected and averaged across 10,000 trials. Each group was composed of five randomly chosen agents, constrained to contribution, threshold, and penalty values on $\{0.0, 0.5, 1.0\}$. The group's TPI is determined by the SCM given each agent's submitted values for the threshold and penalty, and agents optimize their contributions given the resulting TPI. Each agent participated in twenty such groups during a given generation.

	0.30	0.40	0.50	0.60	0.70	0.80	0.90
MAX	0.997	0.997	0.997	0.998	0.998	0.998	0.998
MEAN	0.989	0.993	0.995	0.996	0.996	0.997	0.997
MEDIAN	0.989	0.989	0.990	0.994	0.994	0.994	0.994
MIN	0.399	0.413	0.425	0.818	0.819	0.819	0.819

To understand better the behavior of the adaptive system, consider the system under the MIN SCM. Under MIN, almost all of the (many) Nash equilibria lead to contributions of zero. Notwithstanding this result, under adaptation the system tends toward moderate levels of contributions when $\alpha \leq 0.5$ and relatively high levels when $\alpha > 0.5$. Recall that the accumulated payoffs that an agent receives during its interactions across its various groups within a given generation drive selection in the adaptive system. Under MIN, the TPI parameters are determined by the lowest values submitted by the agents. Given this, consider two agents, one of which submits high values for both the threshold and penalty and the other of which submits low values. When an agent with low values is placed in a group, its values prevail in the resulting TPI and all of the agents in that group react by not contributing—implying

analysis the system does *slightly* worse, reflecting the impact of mutation on the agents.

[10]These bifurcated contribution levels are likely tied to the different number of possible full-contribution equilibria (1 versus 32) as well as an algorithmic choice that, when agents are indifferent between contribution levels, they choose the lower amount.

that each group member will receive a payoff of 1.0. Thus, agents that submit low values always receive a payoff of 1.0 in any game. When an agent with high values is placed in a group, one of two things can happen. If the agent is mixed with at least one low-value agent, the resulting TPI implies no contributions and payoffs of 1.0 to each agent. However, when the high-value agent is mixed with other high-value agents, the resulting TPI induces high contributions leading to payoffs of $N\alpha > 1.0$ to each agent. Thus, high-value agents will either get the same payoff as low-valued agents (when the group has low-value submissions) or a higher one (when the group has only high-valued submissions). Given this, selection will tend to favor high-value agents and push the system toward the higher-contributing equilibria. There are, of course, selective wells in MIN that can trap the system, for example, if all the agents have low submission values, then an agent with high values never ends up in an all-high-value group. Even in this case, genetic drift might allow more high-value agents to arise over time—since such agents always do at least as well as low-valued agents, if not better—potentially allowing the system to flip to a higher contribution level.

Given that evolution has the ability to overcome MIN's extreme behavior and produce low-levels of free riding, its ability to avoid free riding in the three less extreme SCMs is not surprising.

As noted, the behavior of some of the SCMs in the evolutionary system depends on whether $\alpha \leq 0.5$. For example, for $\alpha \leq 0.5$, MIN generates thresholds around 0.69 and penalties around 0.82, while for $\alpha > 0.5$ the thresholds jump to around 0.98 and penalties decline slightly to around 0.79. For MAX, while the thresholds remain around 0.83 throughout, penalties go from 0.83 for low to 0.62 for high αs. MEAN and MEDIAN also show declining penalties for higher αs. As α increases, for contributions of zero to remain optimal agents require either a higher threshold or a lower penalty, so the observed patterns are consistent with TPIs that promote contributions to the public good.[11]

[11]The existence of the 0.5 α boundary is tied to the discrete nature of the contributions. For all levels of α, a threshold of 0.5 is binding when penalties are between 0.35–0.5 thus, given the discrete nature of contributions, this requires penalties of either 0.5 or 1.0. For thresholds of 1.0 to be binding, $\alpha > 0.5$ requires penalties of 0.5 or 1.0, while for $\alpha < 0.5$ the required penalty is 1.0.

We also considered the impact of group size on the adaptive system. We considered groups of size 5, 10, 20, 25, and 50[12] with strategies confined to $\{0.0, 0.1, \ldots, 0.9, 1.0\}$. We re-scaled the αs so that the return from the public bank if everyone donated was constant across group size. For example, if $\alpha = 0.3$ in a group size of 5, it became 0.15 for a group size of 10, resulting in $\alpha \times GS = 1.5$ in either case. We found that group size had little impact under MAX, with groups of all sizes maintaining donations of around 1.0. The other three SCMs produced monotonically decreasing mean donations as group size increased above 5. For example, with a group size of 50 and $\alpha \times GS = 1.5$, the mean donations were 0.63, 0.70, and 0.04 for MEAN, MEDIAN, and MIN respectively. These contribution levels are more akin to the those in the five-agent system with equally-probably Nash equilibria than the five-agent adaptive system. With large groups, the two extreme SCMs, MIN and MAX, are likely to achieve their extreme values as it takes just one bad apple in the case of MIN or one good peach in the case of MAX, to drive the system to either complete or no free riding, respectively. For MEAN and MEDIAN, larger group sizes tend to minimize the impact of any single agent on the outcome, thus an agent's submissions to these SCMs have little relation to the resulting payoffs that the agent will receive, thus weakening the selective forces.[13]

When $\alpha = 0.5$ and the penalty is 0.5, agents are indifferent between contributing a threshold of 1.0 or contributing 0 and paying the penalty and in our implementation of the system agents choose to contribute 0 in such cases. These various factors can account for the bifurcations we observe around $\alpha = 0.5$.

[12] These are all divisors of the population size of 100.

[13] We explored other parametric changes as well. For example, evolving the system for 50 generations instead of 25, resulted in a slight increase in mean donations under MEAN and MEDIAN. Altering the population size from 100 to 200, made little difference. We also increasing the possible levels of strategic choices to eleven (implying choices on $\{0.0, 0.1, \ldots, 0.9, 1.0\}$) versus the three ($\{0.0, 0.5, 1.0\}$) used in the previously analyzed five-agent model. Under Nash, we sampled the 167.6 million possible payoff cells 10 million times and found that the mean per-capita contributions were similar to those in Table 3, with the only major difference being slightly higher contributions across MEAN, MEDIAN, and MAX, while MIN's contributions were confined to a narrower range across α. Under the adaptive system with eleven levels and five agents, MAX, MEAN, and MEDIAN behave very similar to what was observed in the three-level system shown in Table 6, while MIN results in higher contributions in the range of 0.72–0.95 as α increases. These minor changes are likely the result of finer grained strategic levels allowing more flexibility in TPI parametric choices.

2.2.2 Conclusions

The ability to solve social dilemmas, by whatever means, is key if we are to confront some of society's most pressing problems. While some social dilemmas have, on occasion, been solved by the introduction of various institutions, such solutions have typically involved relatively small groups and the creation of situation-specific institutions (Ostrom, 1990; Fehr and Schurtenberger, 2018). The implementation of a more transferable institution—such as the threshold-penalty institution generated by the decentralized "votes" of the agents in the social system we explored here—could, if such a system can discover and maintain a useful institution, potentially provide a more general means by which to reduce the negative impacts of social dilemmas.

To model such a system, we analyzed the behavior of four mechanisms to generate social choices, each of which took the votes of the individual agents in a group and aggregated them into the key parameters of the governing institution that is imposed on the group. These SCMs covered a broad swath of behavior, ranging from extreme mechanisms such as MAX and MIN designed to explore the limiting cases, to mechanism more closely associated with the usual notions of democratic voting systems such as MEDIAN and MEAN.

Our results indicated that the SCM used to aggregate votes and form the final institution mattered. Depending on the analytic view and the underlying parameters that characterize the social system, the ability of a given SCM to create an effective institution that resolves the social dilemma differed in predictable ways.

One branch of our analysis focused on the static behavior of the system as captured by rational agents pursuing Nash equilibria. Such equilibria abound in this system. Under the two extreme SCMs, the equilibria were concentrated on institutions that either solved the underlying social dilemma (MAX) or almost always resulted in complete free riding (MIN). MIN, in essence, gives each agent the ability to veto a TPI that forces contributions to the public good, similar to many international climate change agreement efforts (Lal Pandey, 2014). Under the more voting-like SCMs, MEAN and MEDIAN, the system favored a wider variety of equilibria that, on average, lead to partial contributions and hence partial solutions to the social dilemma. Notwithstanding the structural similarities between MEAN and MEDIAN, we observed systematic differences in the system's response to these two SCMs.

An alternative analytic branch considered a dynamic version of the system driven by agents adapting their votes based on prior payoffs. When group size was relatively small, these adaptive agents evolved institutions that resulted in a better resolution of the social dilemma than one might expect given the proliferation of equilibria under rational agents. Even in worlds governed by MIN, where the vast majority of rational equilibria result in complete free riding, the system exhibited moderate to relatively high levels of contributions to the public good (depending on the return to public contributions, α). The other three SCMs were able to avoid almost all free riding (modulo the noise from mutation), a result that is not too surprising given MIN's success.

The adaptive system was able to solve the free-rider problem because agents advocating strong institutions (those that require large contributions if high penalties are to be avoided) have an inherent payoff advantage over more laissez-faire agents. This advantage arises because whenever laissez-faire agents are in a group, the weak institution that emerges results in relatively low contributions and individual payoffs. While agents advocating strong institutions cannot avoid such groupings, on occasion these latter agents are grouped together, and the resulting strong institution leads to both high contributions and individual payoffs. Given this, advocates of strong institutions will always do at least as well as, but likely better than, laissez-faire agents. However, as group size increases an agent's influence on the final form of the institution wanes under all but the most extreme voting systems, and the system reverts to less productive outcomes.

Thus, the two branches of our analysis suggest that the ability of agents in a social system to self-organize productive institutions that can avoid a social dilemma may be more difficult than intuition might suggest. Even in the simplified system explored above, where a seemingly obvious institutional solution exists, acquiring that institution may be difficult, however compelling it may be once it is established. Rational agents can succumb to a proliferation of possible equilibria, many of which lead to inferior outcomes, and while adaptive agents may be able to evolve institutions that induce more promising outcomes, this ability diminishes as group size increases.

These results suggest that deriving productive institutional solutions to social dilemmas may not be easily accomplished using the typical decentralized mechanisms for creating social choice. Given that social dilemmas lie at the heart of some of the most critical problems facing social systems, finding an alternative means to enact such institutions may be needed.

References

Abbink, K., L. Gangadharan, T. Handfield, and J. Thrasher (2017). Peer punishment promotes enforcement of bad social norms. *Nature communications 8*(1), 1–8.

Arrow, K. J. (1950). A difficulty in the concept of social welfare. *Journal of Political Economy 58*(4), 328–346.

Black, D. (1948). On the rationale of group decision-making. *Journal of Political Economy 56*(1), 23–34.

Bowles, S. and H. Gintis (2011). A cooperative species. In *A Cooperative Species*. Princeton University Press.

Chaudhuri, A. (2011). Sustaining cooperation in laboratory public goods experiments: a selective survey of the literature. *Experimental Economics 14*(1), 47–83.

De Geest, L. R. and D. C. Kingsley (2021). Norm enforcement with incomplete information. *Journal of Economic Behavior & Organization 189*, 403–430.

Downs, A. (1957). *An economic theory of democracy.* Harper & Row: New York.

Fehr, E. and S. Gächter (2000). Cooperation and punishment in public goods experiments. *American Economic Review 90*(4), 980–994.

Fehr, E. and I. Schurtenberger (2018). Normative foundations of human cooperation. *Nature Human Behaviour 2*(7), 458–468.

Fehr, E. and T. Williams (2018). Social norms, endogenous sorting and the culture of cooperation. Working paper, CESifo.

Finus, M., J.-C. Altamirano-Cabrera, and E. C. Van Ierland (2005). The effect of membership rules and voting schemes on the success of international climate agreements. *Public Choice 125*(1-2), 95–127.

Holland, J. (1975). *Adaptation in natural and artificial systems.* The University of Michigan Press: Ann Arbor.

Lal Pandey, C. (2014). The limits of climate change agreements: From past to present. *International Journal of Climate Change Strategies and Management 6*(4), 376–390.

Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action.* Cambridge University Press.