WORLD DEVELOPMENT REPORT 2017

BACKGROUND PAPER

GOVERNANCE *and* THE LAW

The Governance Game

Sheheryar Banuri
University of East Anglia

David Bulman, Luis F. Lopez-Calva, Ezequiel Molina, Abla Safir, and Siddharth Sharma
The World Bank

Background Paper for the World Development Report 2017

The Governance Game

Sheheryar Banuri[1]

Ezequiel Molina

Luis F. Lopez-Calva

Siddharth Sharma

David Bulman

Abla Safir

January 2017

Abstract: In this paper, we use the lab to test a series of policy proposals designed to constrain rent-seeking behaviour in a policymaking context. The baseline governance game is conducted in the following way: subjects are randomly assigned to groups of four, with one subject randomly selected to be the "policymaker", while the other three are the "citizens". Citizens are informed that they can use their endowments to contribute to a group account. Any amount contributed to the group account are doubled. Once citizens have made their contribution decisions, the policymaker observes the contribution decisions of each citizen, and the total amount in the group account. The policymaker formulates a distribution "policy" to distribute the tokens among all four group members. The game is repeated for 20 rounds. With this basic framework, we implement and test the effect of three institutions designed to constrain policymaker rent-seeking behaviour: voting, policy commitment, and punishment. The results show that voting and enforced commitment are the most effective policy mechanisms to constrain rent-seeking, and improve citizen welfare. We find policymaker punishment regimes to be largely ineffective, both in reducing rent-seeking and improving welfare of citizens.

[1] Banuri: School of Economics and Centre for Behavioral and Experimental Social Science, University of East Anglia, Norwich, UK, NR4 7TJ (e-mail: s.banuri@uea.ac.uk). The author has no relevant or material financial interests that relate to the research described in this paper. The findings, interpretations, and conclusions expressed in this paper are entirely those of the author and do not necessarily represent the views of the World Bank, its Executive Directors, or the countries they represent.

1 Introduction

An extensive literature discusses policy mechanisms designed to constrain policymaker behaviour using laboratory experiments (for an overview, see Abbink and Serra, 2012; Banuri and Eckel, 2012). This is an important area of inquiry particularly since policymakers typically have access to large amounts of resources, which they can use to provide targeted benefits to favoured groups, or even to themselves (see Acemoglu and Robinson, 2001; Grossman and Helpman, 2001; Lohmann, 1995; Bennedsen and Feldmann, 2006; Jain 2001; among others). Relatively little is known about the mechanisms through which various policies are successful in constraining policymaker behaviour, particularly in the rent-seeking context. This paper reports an overview of the results of a large series of lab experiments designed to shed light on the mechanisms yielding success in constraining policymaker behaviour and improving citizen welfare. Using the World Development Report 2017 framework as the guide (World Bank, 2017), I test three common policy prescriptions designed to improve citizen welfare: voting, credibility of policy commitment, and sanctions (i.e. punishment) using a novel game designed to capture the governance problem.

Typically, constraining policymaker behaviour is critically important in improving citizen welfare outcomes. This is because policymakers have access to large levels of public resources, and while improving citizen welfare is the core objective for a majority of policymakers, this often comes in conflict with individual incentives to provide targeted benefits to favoured groups, or to retain private benefits when in public office. Therefore, policies designed to reduce the incentive to engage in rent-seeking behaviour, and improve the overall level of citizen welfare are needed.

Implementing policies are costly however, particularly when there are a number of implementation concerns that are not trivial. The laboratory can be used to test different forms of policies in a controlled setting so as to isolate the mechanism yielding benefits. Recognizing the fact that impact of governance policies on policymaker and citizen behaviour are difficult to identify, measure, and disentangle, I use a laboratory approach to test the impact of the aforementioned policies (under different implementation arrangements) on development outcomes. I use a novel game inspired from standard social dilemma games used in experimental economics: public goods (Ledyard, 1995; Gangadharan et al. 2015; Erkal et al. 2011) and trust games (Berg et al. 1995; Banuri et al. 2016), to construct the "governance" game. The game is designed to capture key elements of governance and collective action problems. This governance game contains modifications on existing games (specifically the public goods, and trust games) in the economics literature used to address social dilemmas. Next, I implement a series of treatments (10 in total) designed to capture key elements of policy mechanisms constraining policymaker behaviour: voting (prospective and retrospective); policy commitment; and punishment (unilateral or group-based, with and without communication).

The results show that voting is very effective in constraining policymaker behaviour due to the availability of incentives that allow policymakers to continue in their roles. Furthermore, I also find suggestive evidence for retrospective voting models yielding higher levels of citizen welfare. Similarly, communication combined with policy commitment is effective in constraining policymaker behaviour and improving welfare, even in the absence of voting mechanisms. Finally, punishment institutions are not very effective (unlike the public goods literature: e.g. Fehr and Gachter, 2000) in yielding improvements in welfare, regardless of the institutional arrangement of punishment.

In the next section I detail out the baseline governance game, which corresponds to a society with weak institutions. I then describe a series of treatments implementing different institutional arrangements for voting, policy commitment, and punishment. The I discuss an overview of the results from the experiment, and finally conclude with thoughts on the mechanism through which the effects operate.

2 Experimental design

As mentioned above, the baseline governance game is used to simulate an environment with weak governance institutions. The game is conducted in the following way: subjects are randomly assigned to groups of four, and retain membership in that group for the entirety of the experiment (fixed groups). Next, one subject from each group is randomly assigned to the "policymaker" role, while the other three are assigned to the "citizen" role (instructions are neutrally framed, however, with subjects referred to as "type X" or "type Y"). Subjects assigned to the role of policymaker maintain that role for the entirety of the experiment (i.e. essentially they are serving as a dictator). At the beginning of each round, each subject is endowed with 20 EMU (experimental monetary units) that will be used to contribute to a group account. All tokens contributed to the group account are doubled by the experimenter (which is common knowledge). Once citizens have made their contribution decisions, the policymaker observes the contributions of each citizen to the group account, and decides on a "distribution rule": that is, the policymaker distributes the tokens among all four group members. This is the key departure from the standard linear public goods game. The distribution rule (expressed in terms of % of group account assigned to each player) is then revealed to the citizens. Earnings are then calculated and reported to each group member individually, after which the round ends. Subject are informed that they will participate in this task for 20 rounds.

Hence, the game measures two key outcome variables used in the analysis. Citizen contributions and policymaker distribution rules. These are used to then construct citizen and policymaker earnings, which are then used to calculate welfare. In addition to this, I also measure (incentivized) beliefs of citizens regarding average contribution of other citizens, and beliefs about the percentage of group funds the policymaker will allocate to themselves. Once the governance game is complete, subjects take a short exit survey, are paid for all decisions and the experiment ends.

The experiment has an additional feature worth mentioning: subjects earn their endowments for use in this game by engaging in a coding task beforehand (Gangadharan et al., 2015; Charness et al. 2014; Erkal et al., 2011). This is implemented to inculcate ownership of the initial endowment. In addition to this, all subjects participate in an incentivized dictator game with a charity recipient at the start of the session, which is used to measure altruism. This data is then used in one of the treatments as described below.

Figure 1 displays the overall structure of the session. The governance game measures cooperation and generates outcomes relevant to our study: policymaker behaviour and citizen welfare. Furthermore, by using a real effort task, the game captures essential elements of the relationship between government and citizens through income taxation. Note that due to weak institutions managing the behaviour of the policymaker, the equilibrium of this game (through backward induction) is for the policymaker to keep all contributions made to the group account, and (anticipating this) for citizens to never contribute to the group account. Since this is a finitely repeated game, the equilibrium holds for each round in the game. Individual incentives yield the minimum level of social welfare, constituting a social dilemma.

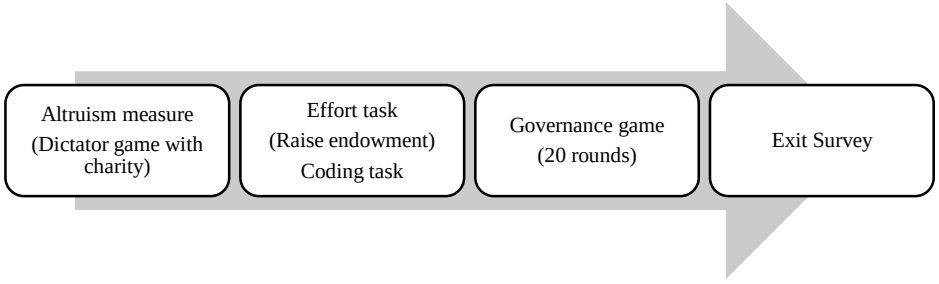

Figure 1: Experiment structure

One additional point to note is that the default contribution of the policymaker is to contribute their entire endowment to the group account. Since they are the ones with control over the group account, it is obvious that their dominant strategy is to contribute everything as contributions get doubled by the experimenter. Under their framework, I implement a number of treatments organized under our three policy mechanisms: voting, commitment, and punishment. Finally, the experiments were conducted

All sessions were conducted at the University of Valencia with a standard student subject pool. Each treatment reported comprised of 60 subjects. That is, each treatment contains 15 groups of four subjects (and hence, 15 independent observations per treatment). Sessions lasted 2.5 hours on average, resulting in an average payment of 24 euros per subject. All earnings were expressed in terms of EMU's (experimental monetary units), which were then exchanged for euros at the end of the experiment.

2.1 Voting treatments

I implement three treatments to isolate key aspects of voting mechanisms. The first is a simple change designed to mimic an environment with voting, but weak implementation. This treatment (called "Random replacement") is identical to the baseline treatment described above, but with the exception that the policymaker is randomly selected every four rounds. What this means is that every 4^{th} round, the program randomly selects a policymaker from the group. This includes the incumbent, so each group member has an equal chance of being selected as the policymaker for the next four rounds. Other than this, the treatment is identical to the baseline. This treatment corresponds to faux-democratic elections where citizens have no voice and policymakers cannot be held accountable, but are replaced at regular intervals.

As mentioned earlier, I implement two types of voting institutions. The first is called "Prospective voting." In this, subjects are asked to vote for a policymaker prior to the first round of the game, and then subsequently every fourth round after the first. Voting is based on two pieces of information about group members. The first is altruism: subjects are provided information on each group members contribution to charity in the earlier dictator game. It is important to note here that subjects were not aware that this data would be revealed to group members when they undertook the task. The second piece of information is ability. This is obtained from the practice rounds of the effort task undertaken earlier in the session. As with the dictator task, subjects are not aware that these data would be revealed to group members. Note that while altruism is a relevant variable for citizen decision-making, ability (in the effort task) is irrelevant (and largely treated as such by the subjects). Based on these two pieces of information, subjects must vote for a policymaker among their group members, and they cannot vote for themselves. To keep things simple, I also do not allow for abstentions. A simple majority rule applies, and in the event of a tie, the policymaker is randomly selected from the group members with the highest votes. In this treatment, voting is prospective in that subjects look to the future when casting their vote. Voting occurs every four rounds (as in the Random replacement treatment).

The second and final voting institution is called "Retrospective voting." This treatment is identical to the baseline, with a policymaker selected at random at the beginning of the game. However, every four rounds citizens can vote to retain or replace their policymaker, again through a simple majority voting mechanism. Citizens are provided with information on the performance of the policymaker, and are asked whether they would like to retain or replace the policymaker (no abstention allowed). Retaining the policymaker means that the policymaker continues in his/her role for an additional four rounds. Replacing the policymaker means that another policymaker is randomly selected from the remaining group members. This treatment corresponds to a strong form of voice and accountability where citizens have indirect control over policymaker behaviour through a voting mechanism.

2.2 Commitment treatments

Three treatments comprise my set of commitment treatments. The first is the baseline treatment (called "Baseline-no feedback", which is identical to the baseline described above, but with one minor modification: the distribution rule used by the policymaker is never revealed to the citizens. This treatment constitutes complete lack of information between the distribution policy and the citizens.

The second treatment manipulates information, and is called the "Non-credible commitment" treatment. Policymakers are asked to select a "distribution rule" which proposes a distribution (expressed in percentage of the group account) for all group members. This proposed distribution is then revealed to the citizens, who then choose their contributions. It is common information that the policymaker does not have to follow the rule, and the final rule that is implemented is not revealed to the citizens. In this manner, the only change to the "Baseline-no feedback" treatment is a proposed distribution rule that is not binding (i.e. cheap talk).

The third and final treatment is identical to the non-credible treatment outlined above, with the exception that the proposed distribution rule must be followed. Hence, policymakers propose a distribution rule that is automatically implemented within the round (called the "Credible commitment" treatment). These treatments manipulate the uncertainty associated with policymaker action, and the third treatment in particular is where policymakers would be expected to signal egalitarian distributions in the hope of attracting higher contributions.

2.3 Punishment treatments

The final policy mechanism that I test is of citizen punishment. This set constitutes four treatments. The first treatment is the "Baseline" which is the same as described in the voting treatments above. The second treatment is called "Standard punishment" where I implement a punishment regime (following the experimental economics literature), where each citizen can spend resources punishing the policymaker. At the end of each round, citizens have 5 EMU that they can spend on punishing the policymaker (interpreted as the effort cost for engaging in political action, with higher numbers indicating the intensity). For each EMU spent by the citizen, the policymakers' earnings are reduced by 3 EMU. Each citizen can unilaterally engage in punishing the policymaker.

The third treatment (called "Weakest link punishment") simulates a more realistic version of political action. As in standard punishment, citizens have a limited budget of 5 EMU to spend on punishing the policymaker. As earlier, subjects incur the full costs of punishment, but the lowest level of punishment is implemented for all citizens. This is akin to needing coordination and consensus in implementing punishment of policymakers. Hence, by adding the coordination problem punishment becomes more difficult to implement. The fourth and final treatment is called "Weakest link w/ communication", which is identical to the weakest link punishment treatment, but now citizens have 30 seconds to engage in free-form communication to agree on a punishment level. This treatment simulates a policy designed to improve voice and accountability by reducing costs due to mis-coordination.

3 Results

In this section I present an overview of the results from the experiments. I organize the results (as above) by our three policies designed to reduce policymaker rent-seeking: voting, policy commitment, and punishment. Unless otherwise noted, all reported estimates of treatment effects

are based on random-effects GLS regressions with group level clusters. The results demonstrate the effectiveness of each of the policy mechanisms in terms of (1) improvement in citizen welfare (measured through earnings), (2) reductions in policymaker rent-seeking, (3) citizen contributions to the public good, and (4) equity in public good resource distribution. Each of the following subsections will discuss treatment effects on the main variables of interest listed above.

3.1 Voting

Voting policies are designed to replace ineffective policymakers. As explained earlier, I test the effectiveness of two distinct voting mechanisms: prospective and retrospective voting (Kuklinski and West 1981; Lockerbie 1991), against two distinct baseline treatments: one where the policymaker retains their role throughout the sessions, and another where policymakers are replaced every four rounds, but through a random mechanism rather than through voting. In this manner, I isolate the effect of voting institutions on policymaker behaviour in the experiment.

Prospective voting entails subjects making decisions about the future, while retrospective voting entails subjects making decisions about past behaviour. The prospective treatment asks subjects to select their policymaker for the next four rounds based on two pieces of information: altruism (measured via donations to charity in an earlier elicitation), and effort (measured via an incentivized real effort task unrelated to the experiment). Retrospective voting entails subjects observing the actions of the policymaker over the past four rounds, and voting to retain or replace the incumbent. If incumbents are voted out, they are replaced by a randomly selected group member. Simple majority voting rules are used throughout, with a random draw determining the winner in case of ties. As mentioned earlier, all treatments are strongly balanced, such that 60 subjects (in 15 groups) participated in each treatment.

Figure 2 displays the core impact of voting on citizen and policymakers (in terms of earnings). The earnings reported in the figure are net of the endowment of 20 EMU for citizens, and hence is labelled as the surplus generated from participating in the game. For policymakers, the numbers reported are net of the total amount generated by policymaker contributions to the group account (that is, twice the endowment of the policymakers: 40 EMU). This is because policymakers contribute their entire endowment to the group account by default (which is 20 EMU). Hence the maximum benefit they generate for the group is 40 EMU, and any amount that they allocate to themselves above 40 EMU can be interpreted as rents. I use this interpretation to classify amounts as rents in the figure.

First, note that randomly selecting the policymaker every four rounds increases citizen surplus and reduces policymaker rents, but this effect is not statistically significant. Earnings for citizens (surplus) increase by 1.57 EMU (8% of the endowment: $p=0.44$), while earnings for policymakers (rents) decrease by 3.89 EMU (19% of the endowment: $p=0.36$). This suggests that simply replacing the policymaker at regular intervals is unlikely to have an effect on citizen welfare. Relative to the baseline, however, prospective voting improves average citizen surplus by 4.14 EMU (21% of the endowment: $p<0.05$) and reduces policymaker rents by 8.60 EMU (43% of the endowment: $p<0.05$), indicating an improvement over the baseline of a fixed policymaker. However, this increase in citizen surplus and decrease in policymaker rent-seeking is not significantly different from the random replacement treatment ($p=0.25$ and $p=0.15$), indicating that the voting mechanism (in isolation) is not yielding a significant improvement over simple random selection of policymakers.

I observe the largest impact in the retrospective voting treatment: citizen surplus increases by 5.52 EMU (28% of the endowment: $p<0.01$), while policymaker rents decrease by 12.30 ECU (62% of the endowment: $p<0.01$). Moreover, this improvement is significantly different from the random replacement treatment ($p<0.05$ for citizen surplus, and $p<0.01$ for policymaker rents). This

indicates that the voting mechanism has an independent effect on citizen welfare. While the result for retrospective voting is nominally better than prospective voting, they are not significantly different from each other.

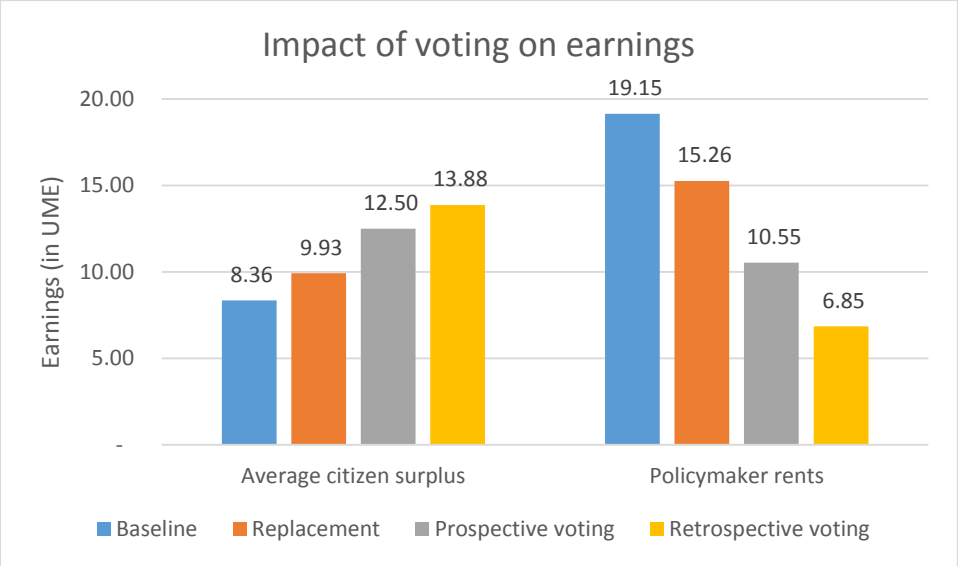

Figure 2: Impact of voting treatments on earnings

I next turn to how these impacts on citizen surplus and rent-seeking are generated. In the context of the experiment, there are two main channels for increases in citizen surplus and decreases in policymaker rents: citizen contributions, and policymaker allocations. That is, citizens can earn more (on average) if they increase their contributions (holding allocations constant), or policymakers retain a smaller percentage of the group fund. Figure 3(a) reports differences (by treatment) in the contributions of citizens, while figure 3(b) reports the average percentage allocated by the policymakers to the citizens.

Citizens contributions are nominally (but not significantly) higher (by 1 EMU on average: 5% of the endowment) in the voting treatments relative to the baseline and random replacement treatments ($p>0.20$ in all cases). Hence, I find little evidence in favour of voting increasing citizen contributions to public funds. Rather, the improvement in welfare observed earlier can be attributed to higher allocations by the policymakers to citizens. Prospective voting increases allocations to citizens by 2.82% over the baseline ($p<0.10$) and 2.13% over random replacement ($p=0.21$), whereas retrospective voting increases allocations to citizens by 4.07% ($p<0.01$) relative to the baseline, and by 3.38% ($p<0.05$) over random replacement. As in the welfare results above, retrospective voting is nominally higher in terms of policymaker allocations, but the difference is not significant ($p=0.35$). However, the results do provide evidence in favour of the increase in welfare being driven by increases in policymaker allocations to citizens.

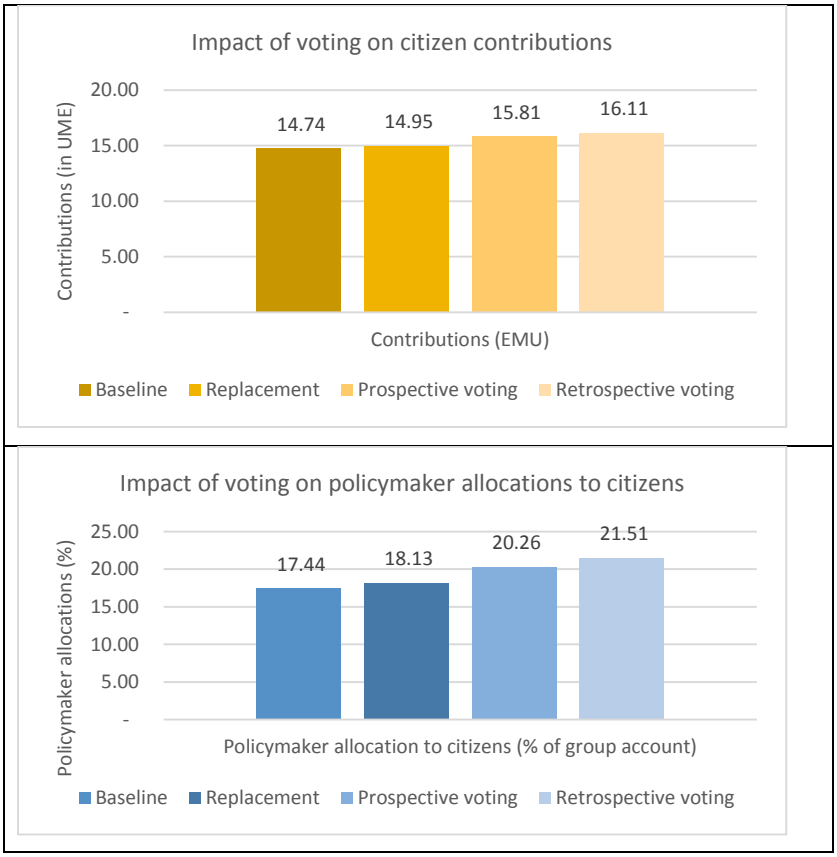

Figure 3 (a-b): Impact of voting treatments on citizen contributions (a) and policymaker allocations to citizens (b)

Figure 4 (a-b) provide direct evidence driving the differences between the voting treatments and random replacement. In the random replacement treatment, since policymakers know that there is only a 25% chance of being re-selected as policymaker, they increase the percentage allocated to themselves in the final round of their tenure. That is, right before policymakers get replaced, they start to increase rent-seeking behaviour. The voting mechanism serves as a disciplining device, in that policymakers do not increase allocations to self in the hopes of retaining their position. Note that in the final four rounds, when there is common knowledge that the game will end, rent-seeking behaviour is identical in all treatments.

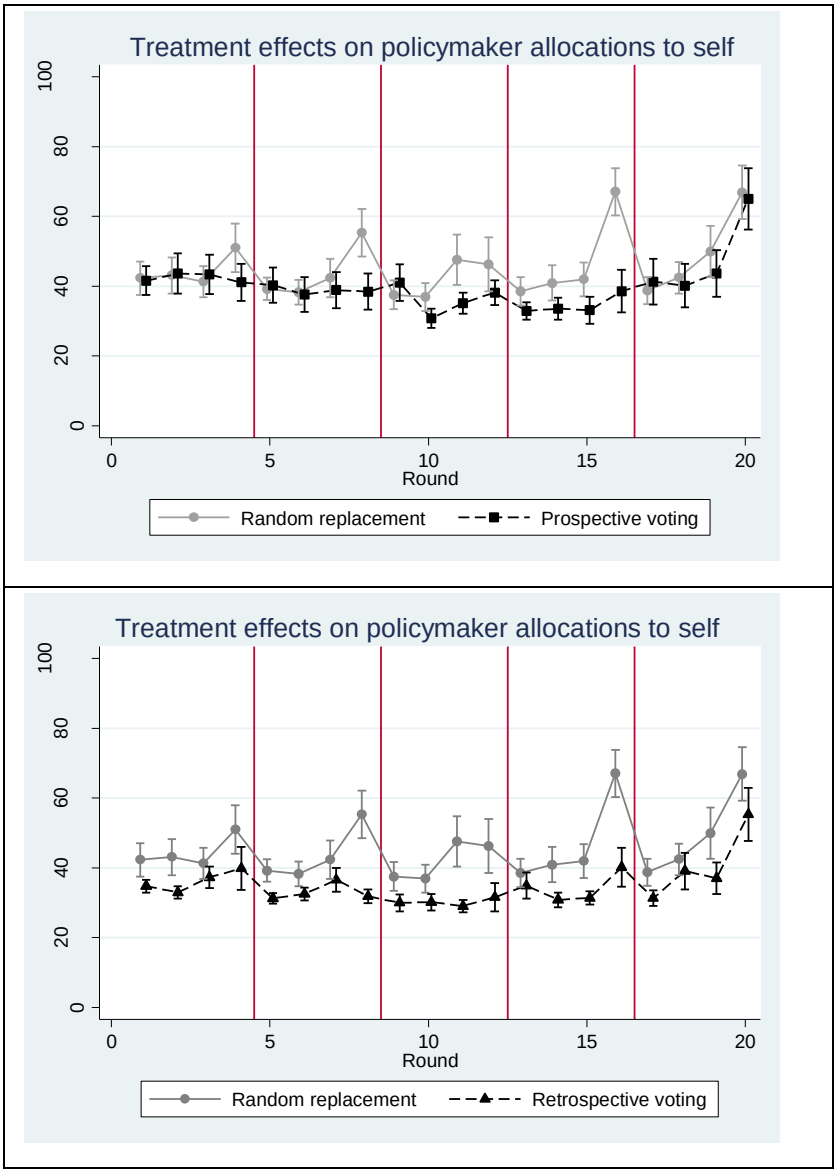

Figure 4 (a-b): Impact of voting treatments on policymaker allocations to self

Overall, the results show that the voting mechanism is effective, but results vary by type of voting institution (though not significantly). I find little evidence of differences between retrospective and prospective voting, but retrospective voting demonstrates some nominal improvements in welfare over prospective voting, in line with Lanoue (1994). What is clear, however, is that both voting institutions show significant improvements in citizen welfare, stemming from reductions in policymaker rent-seeking behaviour. Policymakers respond to the voting mechanism by reducing their appropriation levels, affecting earnings of citizens.

3.2 Commitment

Next I turn to policymaker commitment as a tool to discipline policymakers. These treatments focus on communicating the distribution rule to the citizens prior to citizens making their contribution decisions. Unlike the voting mechanisms described above, the policymaker is randomly selected at the beginning of the game, and retains his/her role throughout. I conduct three treatments: (1) a baseline (called "Baseline-no communication") where the policymaker does not have to provide any information prior to the contributions; (2) a non-credible commitment

treatment where the policymaker can propose a distribution (in terms of percentages of the group fund), but is free to change the distribution after observing contributions, and (3) a credible commitment treatment where the policymaker is bound by the distribution rule proposed prior to citizen contribution decisions. These treatments test whether strengthening policymaker commitment on promises reduce rent-seeking by policymakers, and increase contributions by citizens (through reducing uncertainty in benefits). One additional point to note is that the baseline used here is different from the one used in the voting and punishment policy experiment. This is because the baseline does not reveal the distribution, so the only feedback citizens get is the amount they received from the group account. The distribution rule is never revealed, which corresponds to an environment where citizens have no communication about the distribution policy. The non-credible commitment adds communication prior to contribution decisions, while the credible commitment adds communication and a commitment device to ensure that policymakers follow through on promises.

Figure 5 below displays the main findings: it compares differences in citizen surplus (stage earnings minus the endowment) and policymaker rents (stage earnings minus amount in group account attributable to policymaker contribution). Non-credible commitment increases citizen surplus by 2.62 EMU (13% of the endowment), but the increase is not statistically significant ($p=0.23$). Credible commitment, on the other hand, significantly improves citizen welfare: surplus increases by 4.04 EMU (20% of the endowment) and the effect is statistically significant ($p<0.05$). The difference between the credible and non-credible treatment effects is not significant ($p=0.52$) indicating that both communication and commitment need to be in place for increasing citizen welfare.

For policymakers, both non-credible and credible commitment treatments yield a reduction in rent-seeking behaviour. The non-credible commitment treatment reduces policymaker rents by 6.71 EMU, but the effect is marginally significant ($p<0.15$). However, credible commitment reduces policymaker rents to nearly zero ($p<0.01$). Moreover, the difference between the credible commitment and non-credible commitment is also significant ($p<0.01$). This provides evidence in favour of both communication in policy stances, and commitment in following through increasing citizen welfare.

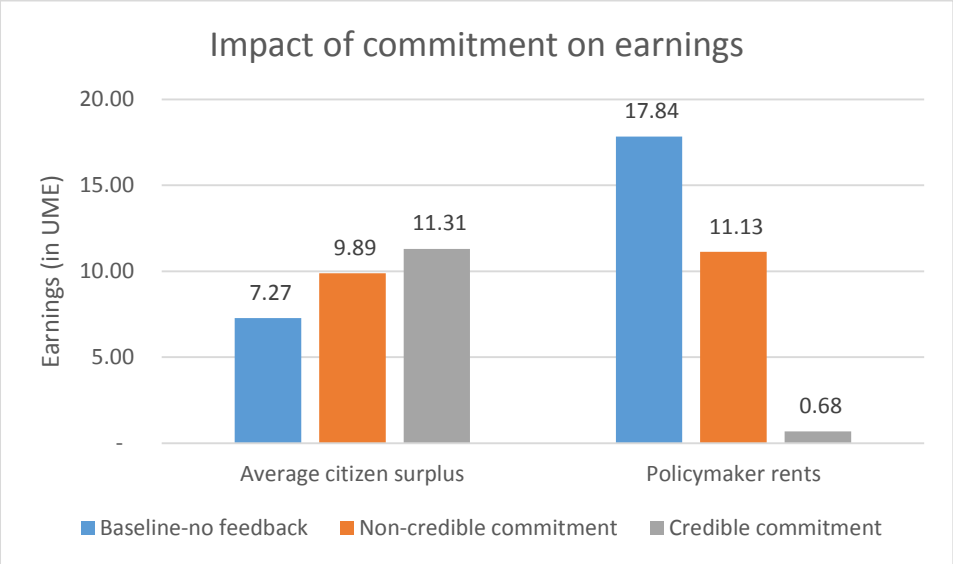

Figure 5: Impact of commitment policies on earnings

As stated in the previous subsection, improvements in citizen welfare can be attributed to either (a) an increase in citizen contributions, or (b) a reduction in policymaker rents. Figure 6 (a) displays

the effect of the treatments on citizen contributions, while 6 (b) reports the average percentage allocated by the policymakers to the citizens. I first note that the non-credible commitment treatment yields an increase in citizen contributions (of 0.38 EMU), while the credible comment treatment yields a reduction in citizen contributions (of 1.68 EMU), but neither of these estimated effects are significantly different from the baseline ($p=0.82$ and $p=0.20$ for the non- and credible commitment treatments respectively).

For distributions, I observe a different pattern. The percentage allocated to citizens by policymakers is not significantly different between the baseline and the non-credible commitment treatment (an increase of 0.92: $p=0.66$). However, I find that policymakers significantly increase their allocation to citizens in the commitment treatment (an increase of 3.24: $p<0.05$). The treatment effect is not significantly difference between the non-credible and credible treatments ($p=0.29$), however, indicating that both communication and commitment are needed to increase allocations to citizens, and thereby increase citizen welfare. There is evidence in favour of policymakers responding to the credible commitment treatment by choosing more equitable distribution rules in order to attract higher levels of contributions from the citizens. The (insignificantly) lower levels of contributions by citizens is consistent with the notion that citizens reduce contributions in favour of more equitable distribution rules.

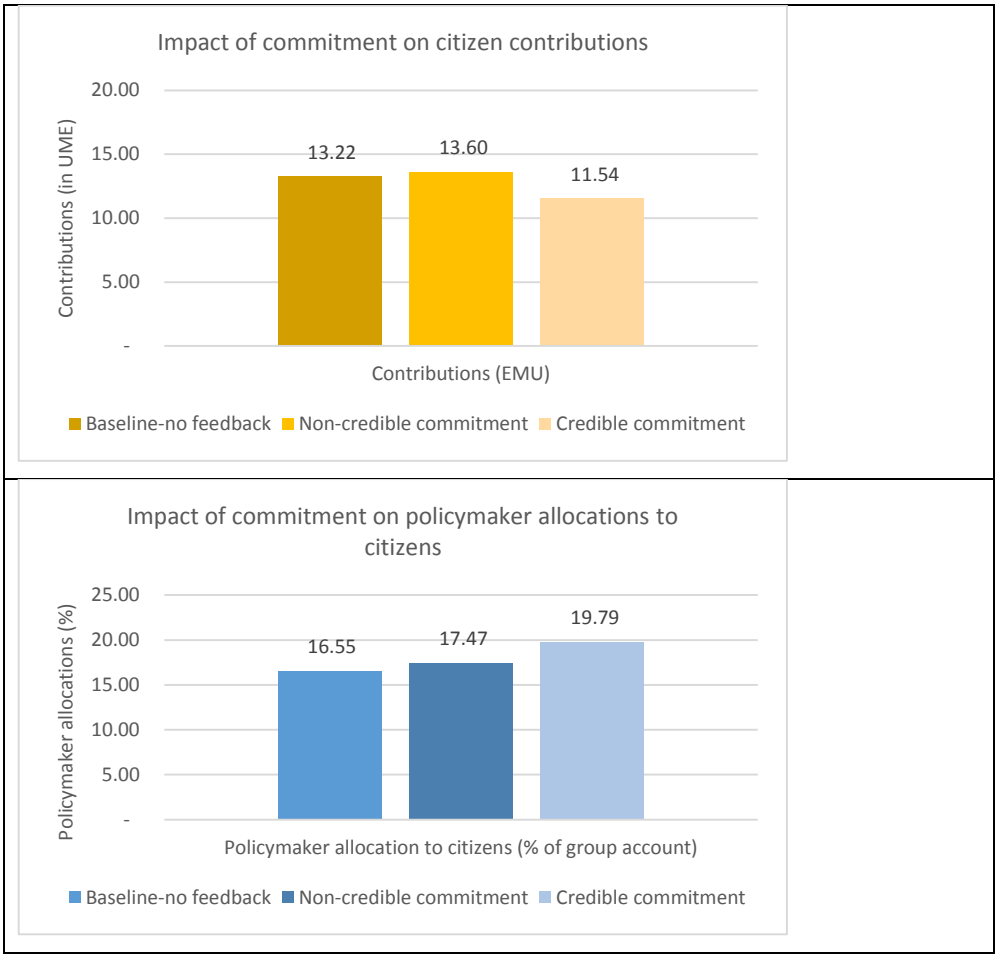

Figure 6 (a-b): Impact of commitment treatments on citizen contributions (a) and policymaker allocations to citizens (b)

Finally, I look at the difference between proposed and actual allocations to self by policymakers (in the non-credible commitment treat). Figure 7 displays the variation in policymakers' allocations to self over time. Figure 7(a) focuses on proposed allocations, while figure XX(b) focuses on

actual allocations (note that since proposed and actual are the same in the credible commitment treatment, these data don't change across the two graphs). In the non-credible commitment treatment, I find considerable divergence in proposed and actual allocations by policymakers over time.

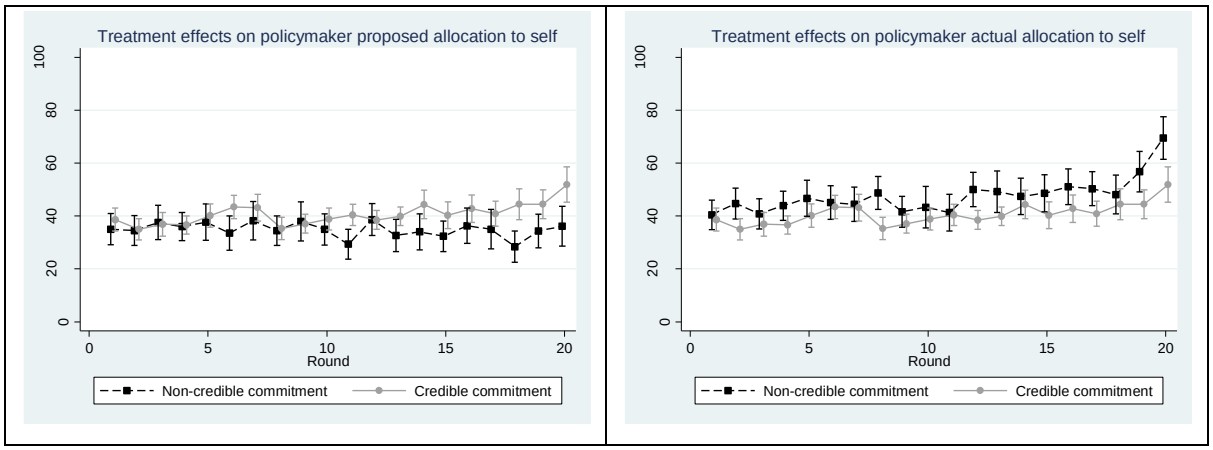

Figure 7 (a-b): Impact of commitment treatment on policymaker (a) proposed allocations to self (left) and (b) actual allocations to self (right)

In sum, the results show that when policymakers communicate their intentions, and they are committed to follow through, citizen welfare significantly improves. This is mainly due to a reduction in policymaker rent-seeking behaviour over time. Non-credible commitment institutions do not fare as well due to the lack of commitment, which policymakers exploit in the latter half of the experiment. Clearly, while communicating policies is important, it is observed that this cannot be without the preserved in the absence of commitment devices.

3.3 Punishment

The final policy mechanism I examine is that of punishment institutions. This is perhaps the most common policy mechanism (in conjunction with voting) to address policymaker underperformance. In the context of the experiment, I test three types of punishment regimes. The first is called "standard" punishment, which is the typical mechanism used to study punishment behaviour in public goods games (Fehr and Gachter, 2000). This punishment regime typically involves subjects spending some amount of money to reduce the payoff of any one group member. This corresponds to a well-functioning society where any individual can punish any other individual for under-contributing to the public good. I implement this regime by allowing each citizen to punish the policymaker unilaterally, with each EMU spent on punishment (corresponding to effort exerted in punishing) reducing the policymaker's payoff by 3 EMU (in line with previous experiments on punishment in trust and public goods games).

The standard punishment institution requires the assumption of well-functioning underlying institutions that allow any one member (or group) in a society to be able to punish policymakers. However, typically punishment requires coordination between a large group of agents in a society to be effective. In order to capture this (more realistic) regime, I use a second type of punishment regime called "weakest link" where each citizen chooses a punishment level unilaterally, but the lowest level of punishment is used to reduce payoffs for the policymaker. However, each agent incurs the full effort cost of punishment. This transforms punishment into a coordination game, where citizens try to coordinate on an appropriate punishment level, but there are clear incentives to choose lower levels of effort. The final treatment (weakest link with communication) alleviates

the coordination problem by allowing citizens to discuss and coordinate on an appropriate punishment level. Finally, the baseline treatment is identical to the baseline used in the voting treatments.

Figure 8 displays the core results of the treatments on citizen welfare and policymaker rents (net of punishment). I first note that the standard and weakest link punishment regimes reduce citizen surplus, while adding communication improves surplus. However, none of these treatment effects are significant ($p>0.20$ in all cases). Hence, punishment is not observed to improve citizen welfare over the course of the experiment.

Policymaker rents paint a different picture. Rents are significantly lower when punishment is available, and easy to implement. Standard punishment reduces policymaker rents by 17.40 EMU ($p<0.01$) relative to baseline. Weakest link punishment (where coordination is difficult) reduces rents by 1.17 EMU, but the difference is not significant ($p=0.81$). When communication is allowed (to facilitate coordination) rents are reduced by 17.90 EMU relative to baseline ($p<0.01$) and are not significantly different from standard punishment ($p=0.90$).

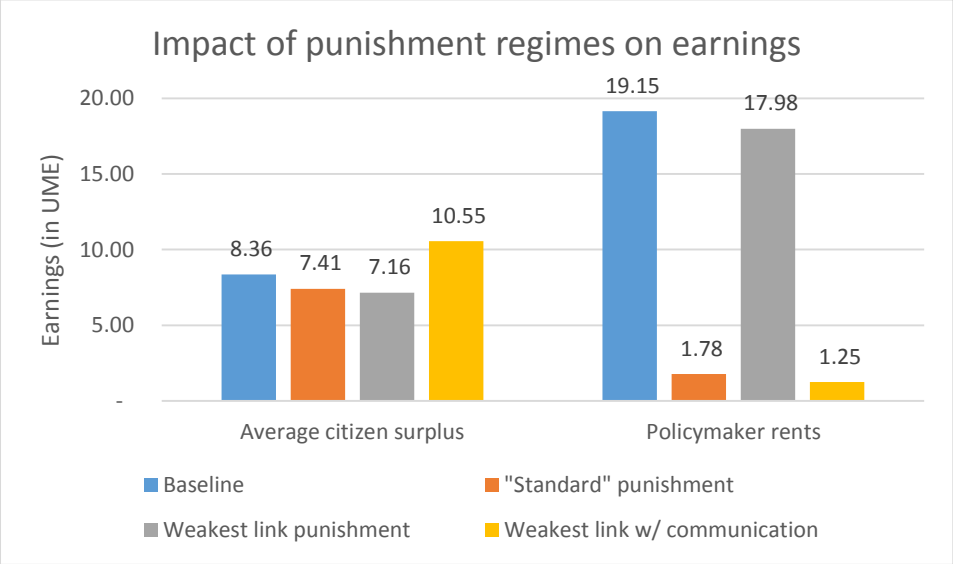

Figure 8: Impact of punishment regimes on earnings

Next, I turn to the treatment effects on citizen contributions to the group account. These results are displayed in figure 9 (a) below. Note that standard punishment has a marginally significant negative effect on contributions of 2.38 EMU relative to the baseline ($p<0.1$). Both the weakest link treatments lower contributions relative to baseline, but not significantly ($p>0.5$ in both cases). Hence, contributions remain largely unaffected by punishment, though the standard punishment does incur a decline, which is a puzzle. However, the differences between punishment regimes are not significant ($p>0.15$ in all cases).

Policymaker allocations to citizens (from group account funds – figure 9-b) are completely unaffected by punishment regimes. Standard punishment lowers allocation to citizens by 0.25%, but is not significant ($p=0.89$). Weakest link without coordination also lowers the allocation to citizens by 0.77%, but again is not significantly different from the baseline ($p=0.67$) or from standard punishment ($p=0.79$). Finally, weakest link punishment with communication does improve policymaker allocations to citizens (by 1.56%), however, this effect is also not significant ($p=0.43$).

Figure 9 (a-b): Impact of punishment regime on citizen contributions (a) and policymaker allocations to citizens (b)

One final point to note is the evolution of punishment across the regimes (presented in figure 10 below). The lowest expenditure by citizens is observed in the weakest link treatment (which is not a surprise given the difficulty in coordination and effort costs). However, standard and weakest link with communication work largely the same way, and punishment expenditure is nearly identical.

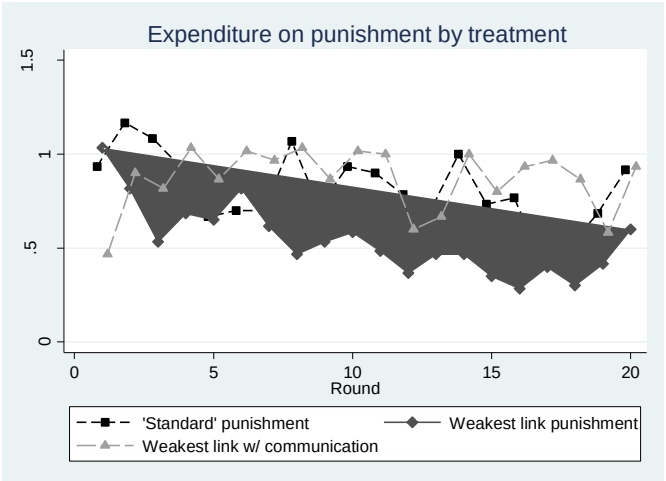

Figure 10: Citizen expenditure on punishment by treatment

Taken together, these suggest that subjects to engage in punishment when it is available and easy to implement, but any increase in welfare is offset by the effort costs of punishment. Furthermore, punishment is largely ineffective in improving policymaker allocation to citizens (i.e. changing policymaker rent-seeking behaviour). However, equity is improved mainly through the direct effects of punishment.

4 Conclusion

This paper presented an overview of the results obtained from a series of policy experiments on the governance context (the "Governance" game). Subjects were assigned to small groups of four, where one subject was asked to redistribute resources (generated via voluntary contributions) to group members. The experiment introduced three distinct policy tools to curb policymaker rent-seeking behaviour, and thereby improve governance and welfare. The policies tests were (i) voting, (ii) policy commitment, and (iii) punishment. The experiment also addresses a number of common implementation issues in each of these domains.

I found that voting institutions performed best in improving welfare and reducing policymaker rent-seeking behaviour. I provide evidence for the mechanism through which it operates, and offer direct evidence that the threat of removal from their position is a strong motivator for policymakers in constraining their behaviour. Gains in welfare were directly attributable to reductions in rent-seeking rather than through increasing citizen contributions, however.

I find very similar results for policy commitment. Asking policymakers to commit to policies prior to contributions, and enforcing these commitments also improved welfare. Policymakers were significantly more likely to pursue egalitarian distribution policies in order to attract and retain voluntary contributions. As before, improvements in citizen welfare were driven by greater redistribution rather than greater contributions. I also found that policy commitment was critical for welfare improvement. Simply communicating proposed policies without committing had very limited impact.

Finally, I find that (contrary to the literature in public goods and punishment) punishment institutions are largely ineffective in improving citizen welfare. This is because punishment institutions do not constrain policymakers from seeking rents. Hence punishment is effective in reducing inequality, but not in improving citizen welfare overall.

References:

Abbink, Klaus, and Danila Serra. "Chapter 4 Anticorruption Policies: Lessons from the Lab." New advances in experimental research on corruption. Emerald Group Publishing Limited, 2012. 77-115.

Acemoglu, Daron and James A Robinson. "Inefficient redistribution." American Political Science Association. Vol. 95. No. 03. Cambridge University Press, 2001.

Banuri, Sheheryar, Catherine C. Eckel, and Rick K. Wilson. "Deconstructing nepotism." (2012).

Banuri, Sheheryar, and Catherine Eckel. "Chapter 3 Experiments in Culture and Corruption: A Review." New advances in experimental research on corruption. Emerald Group Publishing Limited, 2012. 51-76.

Bennedsen, Morten, and Sven E. Feldmann. "Lobbying legislatures." Journal of political Economy 110.4 (2002): 919-946.

Berg, Joyce, John Dickhaut, and Kevin McCabe. "Trust, reciprocity, and social history." Games and economic behavior 10.1 (1995): 122-142.

Charness, G., Masclet, D., Villeval M.C. (2014). The dark side of competition for status. Management Science 60(1), 38–55.

Erkal, N., Gangadharan, L., Nikiforakis, N. (2011). Relative Earnings and Giving in a Real Effort Experiment. American Economic Review 101(7): 33303348.

Fehr, Ernst, and Simon Gächter. "Fairness and retaliation: The economics of reciprocity." The journal of economic perspectives 14.3 (2000): 159-181.

Gangadharan, Lata, Nikos Nikiforakis, and Marie Claire Villeval. "Equality concerns and the limits of self-governance in heterogeneous populations." Available at SSRN 2668518 (2015).

Grossman, Gene M., and Elhanan Helpman. Special interest politics. MIT press, 2001.

Jain, Arvind K. "Corruption: A review." Journal of economic surveys 15.1 (2001): 71-121.

Kuklinski, James H., and Darrell M. West. "Economic expectations and voting behavior in United States House and Senate elections." American Political Science Review 75.02 (1981): 436-447.

Lanoue, David J. "Retrospective and prospective voting in presidential-year elections." Political Research Quarterly 47.1 (1994): 193-205.

Ledyard, John (1995): "Public Goods: A Survey of Experimental Research", in: Handbook of Experimental Economics, ed. by John Kagel and Alvin Roth, Princeton University Press.

Lockerbie, Brad. "Prospective economic voting in US House elections, 1956-88." Legislative Studies Quarterly (1991): 239-261.

Lohmann, Susanne. "Information, access, and contributions: A signaling model of lobbying." Public Choice 85.3-4 (1995): 267-284.

World Bank. World development report 2017: Governance and the law. World Bank, 2017.

Appendix A

Table A.1: Voting regressions

Dependent variable(s):	Earnings		Contributions	Distribution (%)
Sample:	Citizens	Policymakers	Citizens	Citizens
	I	II	III	IV
Random replacement	1.569	-3.887	0.209	0.693
	(2.05)	(4.22)	(1.19)	(1.65)
Prospective voting	4.138**	-8.604**	1.067	2.820*
	(1.98)	(4.22)	(1.26)	(1.62)
Retrospective voting	5.518***	-12.300***	1.371	4.074***
	(1.54)	(3.69)	(1.08)	(1.27)
Constant	28.360***	59.150***	14.740***	17.440***
	(1.26)	(3.51)	(0.81)	(1.10)
Observations	3600	1200	3600	3600
Overall R2	0.043	0.054	0.008	0.025
P	0.002	0.000	0.553	0.004

Notes: Random effects GLS regressions. Table reports coefficients with standard errors in parentheses. * 10%, ** 5%, *** 1% significance level.

Table A.2: Commitment regressions

Dependent variable(s):	Earnings		Contributions	Distribution (%)
Sample:	Citizens	Policymakers	Citizens	Citizens
	I	II	III	IV
Non-credible commitment	2.618	-6.713	0.380	0.920
	(2.17)	(4.47)	(1.68)	(2.12)
Credible commitment	4.041**	-17.16***	-1.678	3.238**
	(1.75)	(4.05)	(1.32)	(1.58)
Constant	27.27***	57.84***	13.22***	16.55***
	(1.19)	(3.43)	(1.00)	(1.03)
Observations	2700	900	2700	2700
Overall R2	0.025	0.118	0.015	0.014
P	0.064	0.000	0.300	0.119

Notes: Random effects GLS regressions. Table reports coefficients with standard errors in parentheses. * 10%, ** 5%, *** 1% significance level.

Table A.2: Commitment regressions

Dependent variable(s):	Earnings (net of punishment)		Contributions	Distribution (%)
Sample:	Citizens	Policymakers	Citizens	Citizens
	I	II	III	IV
"Standard" punishment	-0.948	-17.37***	-2.382*	-0.247
	(1.99)	(4.49)	(1.24)	(1.73)
Weakest link punishment	-1.205	-1.168	-0.69	-0.766
	(2.09)	(4.77)	(1.14)	(1.77)
Weakest link punishment	2.192	-17.90***	-0.959	1.556
w/ communication	(2.45)	(4.55)	(1.66)	(1.95)
Constant	28.36***	59.15***	14.74***	17.44***
	(1.26)	(3.51)	(0.81)	(1.10)
Observations	3600	1200	3600	3600
Overall R2	0.015	0.130	0.014	0.007
P	0.590	0.000	0.287	0.737

Notes: Random effects GLS regressions. Table reports coefficients with standard errors in parentheses. * 10%, ** 5%, *** 1% significance level.