Responsibility Attribution for Collective Decision Makers*

Raymond Duch
Nuffield College Oxford

Wojtek Przepiorka
University of Oxford

Randolph Stevenson
Rice University

December 7, 2013

*We would like to thank Yuval Heller for his very helpful comments. The experiments reported in this paper were conducted at Nuffield CESS (Nuffield College Centre for Experimental Social Sciences). We thank Inaki Sagarzazu and Akitaka Matsuo for their excellent research assistance. This work was partly supported by the Swiss National Science Foundation [grant number 100017_124877].

Abstract

We argue that individuals have general responsibility attribution heuristics that apply to collective decisions made, for example, by families, teams within firms, boards in international organisations or coalition governments. We conduct laboratory and online survey experiments designed to tease out the heuristics subjects use in their responsibility attribution for collective decision makers. The lab experiments comprise a collective dictator game with weighted voting power of the decision makers and a punishment possibility for the recipients. Our results show that recipients punish unfair allocations and mainly target the decision maker with proposal power and with the largest vote share. We find rather weak evidence that decision makers with veto power are targeted or that recipients engage in punishment proportional to weighted voting power. The survey experiment tests whether subjects indeed believe that the decision maker with proposal power has the most influence on the collective decision outcome. The results confirm this conjecture. We discuss the implications of our findings for theories of vote choice in contexts of multi-party governing coalitions.

Word count: 8,732 (excluding title page, abstract, footnotes and references)

1 Introduction

The act of attributing responsibility concerns a broad range of behaviour and occurs with considerable frequency. Given the range of social behaviour in which responsibility attribution occurs it is not surprising that it has attracted considerable attention from social psychologists and cognitive scientists (Gerstenberg, Speekenbrink and Cheung, 2011; Ross and Nisbett, 1991; Weiner, 1995). But understanding how individuals make these responsibility attribution decisions has also been a concern of other disciplines such as economics (Charness, 2000; Barro, 1973; Besley, 2006) and political science (Ferejohn, 1986; Iyengar, 1991; Kiewiet, 1983). This scholarship for the most part has focused on single, or unified, actors as the target of responsiblity attribution – a firm, a political party, a presidential candidate, or an individual.

There is though a subset of responsibility attribution decisions that have received much less scholarly attention – specifically situations in which outcomes are the result of collective decisions reached by a vote amongst the decision makers. These situations can be formally defined features of decision making, which is the case for boards of directors, remuneration committees, or coalition cabinets, for example. But they can also be more informal features of collective decision-making processes in which, for example, the voting rules are not formally defined but rather are informally recognised according to the role decision makers occupy in a decision making entity.

The challenge we frequently face in such situations is to determine how to assign responsibility to individual members of a decision making body. Which family member had the most influence on deciding where to go for our summer vacation? Who amongst the members of the company's remuneration committee should get most of the blame for our miserly salary increase? In response to the government's poor handling of the economy, how should voters apportion blame to the individual political parties making up the coalition government? This essay is concerned with explaining how individuals

distribute responsibility amongst members of a collective decision making entity, who differ in the influence they can have on the outcome of the collective decision.

Since individuals frequently make responsibility attribution decisions in general, and decisions about distributing responsibility amongst members of a collective decision making entity in particular (Messick, Moore and Bazerman, 1997), our conjecture is that they have developed heuristics that facilitate their decision making. By heuristics we mean strategies that 'guide information search and modify problem representations to facilitate solutions' (Goldstein and Gigerenzer, 2002).

1.1 Punishment in a Collective Dictator Game

We chose the collective dictator game with weighted voting power of the dictators (i.e., decision makers) and a punishment possibility for the recipients as our experimental vehicle. How recipients in our game distribute punishment amongst the individual decision makers is the critical behaviour we study in order to understand responsibility attribution. To our knowledge, no other experimental studies have directly addressed this issue. In a standard dictator game, the decision maker (DM) is an individual, but there is a body of experimental research in which dictator game decisions are made by a group rather than an individual (Engel, 2010). We exploit the proven finding that DMs often keep a larger share for themselves and, given the opportunity, recipients punish DMs for their inequitable allocation (Guth, Schmittberger and Schwarze, 1982; Fehr and Gachter, 2000; Fehr and Fishbacher, 2004; Henrich et al., 2006; Dawes et al., 2007). However, our primary goal is to understand how the average recipient attributes responsibility for an inequitable outcome to individual members of the decision making body, and we use recipients' punishment decisions to measure their attributions.

Our expectations regarding the overall punishment of these collective decisions are informed by a well-developed theoretical literature on other-regarding preferences (Bolton

and Ockenfels, 2000; Falk, Fehr and Fischbacher, 2006; Fehr and Schmidt, 1999, 2006; Rabin, 1993). Most formal models of other-regarding preferences imply that an unequal distribution of monetary payoffs causes disutility in inequity-averse individuals and conceives punishment as a means these individuals have to alleviate the perceived inequity. However, these models provide less theoretical guidance as to how individuals would distribute their punishment amongst the DMs in a decision making body whose collective decision led to an inequitable outcome.[1] With our collective dictator game with punishment we hope to contribute to the previous literature an experimental paradigm that allows studying individuals' punishment behavior and their responsibility attribution heuristics.

The features of our game are designed to ensure that recipients punish with considerable frequency, i.e., make lots of responsibility attribution decisions. In order to minimize confounding strategic calculations on the part of subjects, the game mimics in every way possible a one-shot interaction. For example, our design eliminates any reputational incentives that recipients might exploit in order to discipline DMs. We also implement costless punishment for recipients on the grounds that this would maximise the expenditure of punishment points. Finally, we structured the payoffs, including deductions due to punishment, to encourage DM allocations that would maximize punishment and thus responsibility attributions by recipients.

[1]There is though the literature in social psychology concerning the heuristics that people employ to ascribe causality to events that happen in a social context (e.g. Weiner, 1995). The findings speak to the set of heuristics we expect shape responsibility attribution in collective decision making contexts. This literature finds, for instance, that people tend to overestimate the causal effect of dispositional (i.e. personal, internal) factors and underestimate contextual (i.e. impersonal, external) factors (Gilbert, Pelham and Krull, 1988; Jones, 1990; Ross and Nisbett, 1991). Other findings suggest that individuals are rather insensitive to the impact of decision rules on group decision making (Messick, Moore and Bazerman, 1997; Allison and Messick, 1985).

1.2 Laboratory and online survey experiments

We report results from a set of experiments designed to investigate responsibility attribution heuristics in collective decision making contexts in which a proposal is adopted by majority voting. Our aim is to identify which of four possible heuristics individuals employ for distributing responsibility to individual members of a collective decision making body, if they can attribute responsibility (1) to the DM with the most voting power, (2) proportional to the DM's voting power, (3) to the DM with veto power, or (4) to the DM with proposal power. Accordingly, in our experiments, we vary the distribution of DMs voting wights from relatively equal to unequal where a single party has a majority vote weight (i.e. veto power). We also vary the voting weight of the proposing DM (i.e. agenda setter) and whether or not recipients know the proposing DM's identity. The experimental approach we adopt allows us to disentangle the effect of, for instance, agenda setting power versus veto power in a very controlled setting, which is difficult to achieve based on observational data. Only having confirmed the nature of this heuristic reasoning in the lab will we then be prepared to address external validity (Morton and Williams, 2009, 253).

We conduct two laboratory and one online survey experiment. In our first experiment, we find that the average amount recipients spend on punishment does not vary substantially across information conditions, but recipients punish unfair allocations and thereby mainly target the DM with agenda power. We find rather weak evidence that DMs with veto power are targeted or that recipients engage in vote-proportional punishment. Our second experiment further explores the motives behind this agenda setter effect and also varies the outcome of an unfair proposal receiving a no-majority vote. We find that the proposing DM is punished most for an unfair allocation even in conditions where a no-majority vote would result in an equitable distribution of the original amount rather than in a reversion to zero. In the online survey experiment we test whether recipients

indeed believe that the DM with agenda power has the most influence on the collective decision outcome. Our results confirm this conjecture.

We end this paper with a discussion of the implications of our results for theories of voting behaviour in multi-party governing coalition contexts. Our discussion builds on an existing body of work that demonstrates that significant numbers of voters in coalitional contexts exercise a "coalition-directed vote" (Kedar, 2005; Duch, May and Armstrong, 2010). Coalition-directed voting implies that voters assess the impact of a vote for a party on the coalitions that form and the policies they enact. This requires voters to map observed characteristics of a governing coalition into shares of administrative responsibility. Our results suggest that the party in the governing coalition with proposal power will be seen by the voter as having a disproportionate amount of administrative responsibility.

2 Experiment 1: Responsibility Attribution

Our first experiment was conducted with 96 student and non-student participants. Each of the four experimental sessions comprised 20 rounds. In each round, subjects played the collective dictator game with punishment, which had the following characteristics:

1. Five DMs are chosen at random and are given an endowment which equals the number of participants in the experimental session in £($n = 25$ in sessions 1 and 4, and $n = 23$ in sessions 2 and 3). The recipients receive no endowment. Moreover, one of four voting weight distributions is randomly chosen and each DM is randomly assigned a voting weight from this distribution.

2. One of the DMs is chosen at random to propose an allocation a of the money between the DMs and the recipients in units of £1 ($a \in 0, 1, 2, ...n$). If the proposal is accepted, a and $n-a$ is split equally amongst the DMs and recipients, respectively.

Thus, each DM's payoff at this stage would be $\pi_i = \frac{a}{5}$, and each recipient's payoff $\pi_j = \frac{n-a}{n-5}$. For instance, if the proposing DM proposes to allocate £15 to the DMs and £10 to 20 recipients, each DM's payoff would be $\pi_i = £3$ and each recipient's payoff would be $\pi_j = £0.5$.

3. The five DMs cast all their votes for the proposed allocation or all against it. Abstention is not allowed. If the allocation receives at least 51 votes then it passes. Otherwise the proposing DM has to propose a different allocation, which is also voted on. No other communication is allowed. If three consecutive proposals fail to obtain 51 votes, the group is disbanded and no one is paid anything (i.e. $\pi_i = \pi_j = £0$).

4. In case of a successful proposal, each recipient is given 30 "deduction points" and decides independently and at no cost how much to deduct from each DM's payoff. A DM is deducted the average of the deduction points assigned to him or her by all recipients times £0.1. Let p_{ij} be the deduction points assigned by recipient j to DM i. As a result, the final payoff for DM i is $\pi_i = \frac{a}{5} - \frac{0.1}{n-5}\sum_{j=1}^{n-5} p_{ij}$. Continuing from the example above, if 10 recipients assign 15 deduction points to a DM, five assign ten points and five zero points, then the DM's payoff is reduced by £1 [£0.1*(10*15+5*10+5*0)/20 = £1], and his or her final payoff is $\pi_i = £2$.

In each session, there were four voting weight distribution treatments which we randomised across rounds: (53%, 29%, 10%, 6%, 2%); (48%, 19%, 14%, 11%, 8%); (38%, 21%, 17%, 13%, 11%); (23%, 21%, 20%, 19%, 17%). Except for the first session (see Section 2.2.4) the voting weight distribution was common knowledge in each round. In case of a successful proposal, recipients were also informed of the amounts allocated to the DMs and to the recipients and that this allocation had received a majority of votes; they were not informed of the voting decisions of individual DMs nor of the exact number of votes constituting the majority. Moreover, to minimize learning effects, DMs were

given no feedback on the amount of their deduction – they only learned their total payoff at the end of the session. If a proposed allocation was unsuccessful, recipients were only informed that the allocation failed to receive a majority of votes.[2]

2.1 Behavioural predictions

Our experimental design is informed by the extensive theoretical and empirical work on decision-making in dictator and ultimatum games (e.g. Camerer, 2003).[3] Our interest here is in how recipients in the collective dictator game attribute responsibility to the individual DMs responsible for the collective offers. To our knowledge, no experimental evidence has directly addressed this issue. Unlike most other-regarding games, in our rendition recipients are making two decisions: They decide on how much to punish and also on how to distribute that punishment amongst the individual DMs. In order to gain insight into the later decision, which is the focus of this essay, we require high levels of punishment resulting from their first decision. We thus set the parameters of our game to make it likely that DMs' collective offers will be relatively selfish and recipients will punish quite aggressively.

An implicit assumption here is that the recipients' first decision, the magnitude of their punishment, is not conditioned on the fact that the offer is the result of a collective, rather than individual, decision. Similarly, we conjecture that overall punishment is not significantly affected by features of the collective decision-making process. However, we do expect overall punishment to be negatively correlated with the generosity of the collective

[2]Section 2.1 in the online appendix contains a detailed description of the procedure and design as well as the instructions given to subjects in Experiment 1.

[3]In most of this literature, the initial offers are made by an individual while in our experiments they are made by a collective decision making group. The few experiments on collective decision making in other-regarding games typically compare the initial offers of individual decision makers with those of a group. A major result of these experiments is that groups are more rational and selfish than single individuals (Engel, 2010). Groups are found to be more selfish in a dictator game than individuals and the most selfish group member has the strongest influence on group decisions (Luhan, Kocher and Sutter, 2009). Groups are also found to offer less, and accept lower offers than individuals in ultimatum games (Bornstein and Yaniv, 1998).

offer. Our expectations are based on a game theoretic analysis of a simplified version of the game that subjects played in our experiments. In our analysis, we assume that only one DM i interacts with only one recipient j. We now briefly summarise the results of this analysis. The full analysis, as well as a discussion of how the analysis extends on the multi-party game implemented in our experiment, can be found in Section 1 of the online appendix.

If we assume exclusively self-regarding actors, our dictator game with punishment has multiple Nash equilibria. In most of these equilibria, the DM keeps all of the allocation and, since punishment is costless, the recipient is indifferent as to whether and how much to punish the DM; expending punishment points does not affect the recipient's payoff. However, if we assume that actors are minimally inequity averse, then we can identify a unique subgame perfect Nash equilibrium in our game.[4] In this equilibrium, the recipient is made "better off" as a result of reducing the DM's payoff in order to bring it closer to an equitable outcome.

One possible way to formalize inequity aversion is the Fehr-Schmidt model (Fehr and Schmidt, 1999, 2006), in which a DM's utility depends on their payoff and the payoff of the recipient in the following way: $u_i = \pi_i - \alpha_i \max(\pi_j - \pi_i; 0) - \beta_i \max(\pi_i - \pi_j; 0)$. DM i's payoff is denoted by π_i, α_i is the so called 'envy' parameter, and β_i the so called 'guilt' parameter. It is assumed that $\alpha_i \geq 0$, $\alpha_i \geq \beta_i \geq 0$ and $\beta_i < 1$. That is to say, no DM i likes having less than recipient j, no DM i likes having more than recipient j but dislikes having less more than he or she dislikes having more, and DM i's dislike for having more never outweighs his or her utility from having π_i, respectively. The recipient's utility function can be written accordingly by interchanging subscripts.[5]

[4]Subgame perfection is an equilibrium refinement concept in game theory that excludes Nash equilibria in which players' responses to out-of-equilibrium behavior are not utility maximizing.

[5]There are other models of other-regarding preferences which could be used instead of the Fehr-Schmidt model, and which, with all likelihood, would yield qualitatively similar results (e.g. Rabin (1993); Bolton and Ockenfels (2000); Falk, Fehr and Fischbacher (2006)). We use the Fehr-Schmidt model here because it is very tractable in this setup and allows us to give clear intuitions to readers.

Assuming weak inequity aversion, by setting $\alpha_{i,j}$ and $\beta_{i,j}$ to 0.01 (or any arbitrarily small value), results in a game with a unique subgame perfect Nash equilibrium. In this equilibrium, 1) the recipient punishes a non-egalitarian allocation with the amount necessary to bring the DM's payoff as close to their own payoff as possible, and 2) the DM maximizes their utility by allocating all for themselves (in which case the recipient expends all their punishment points). This minor adjustment to preferences generates equilibrium behaviour on the part of DMs and recipients that we require for the experiment; specifically the maximum allocation of punishment points by recipients.

The equilibrium behaviour of the recipient does not change with variations in α_j and β_j. A DM's behaviour, however, is sensitive to variations in β_i, the 'guilt' parameter in their utility function. The DM will choose an egalitarian allocation ($a_i = 5$) for $\beta_i > 0.5$ and will allocate all to DMs ($a_i = 25$) for a $\beta_i < 0.5$. The DM will be indifferent between the two actions if $\beta_i = 0.5$. Under the plausible assumption that subjects are heterogenous in their other-regarding preferences, we are likely to observe both types of behavior in our experiment.[6]

Utility functions that incorporate inequity-aversion, such as the Fehr-Schmidt model, help explain the overall punishment behaviour of subjects in our collective dictator experiment. But inequity-averse recipients in these models are entirely indifferent as to whom they punish in the group of DMs as long as their punishment reduces inequity. This implies that recipients' punishment will be unaffected by the amount and type of information they are given about the decision making process and the role assumed by individual DMs.

In contrast, we contend that if information about the collective decision making process is communicated, then recipients will use this information to distribute responsibility to individual DMs. In a recent paper, Bartling and Fischbacher (2012) suggest that be-

[6]In fact, we see considerable out-of-equilibrium behaviour from our subjects. About 25 percent of DMs split the allocation equally with recipients and only 5 percent keep all of the allocation. The remainder falls between these two equilibria (see Section 2.2.3 in the online appendix).

sides inequity aversion and reciprocity, responsibility attribution could be another possible motive for punishment. In addition to responding to the inequity or unkindness of the collective DM allocation, recipients can distribute their punishment points over the individual decision makers in the collective decision making group. Bartling and Fischbacher (2012, 74) suggest a simple measure that attributes "most responsibility to the player whose action had the largest impact on the probability that unfair allocation results."

Our expectation is that individuals have acquired decision making short-cuts (i.e. heuristics) that facilitate this "distribution of responsibility" task. They do not, for example, calculate the Banzhaf (1965) voting power index for each DM. Rather, individuals use selective, and simplified, information about the decision making process in order to decide how to attribute responsibility across the individual members of a collective decision making body. In the next section we explore four features of the decision making process individuals likely employ for attributing responsibility: agenda setting power; the largest DM; the voting weights of DMs, and veto power.

2.2 Results

2.2.1 Punishment Levels

We first examine whether this experimental set-up produces behaviour that punishes unfair allocation decisions similar to previous experiments with an individual DM. In each round, 18 or 20 recipients evaluated each of the five DMs with respect to the collective allocation decision. This gives a total of 7,600 evaluation decisions (76 recipients $\times$ 5 DMs $\times$ 20 rounds). Figure 1 shows the average deduction points used on punishment (i.e., in total per recipient and round) for the amounts the DMs collectively chose to keep for themselves (across all experimental conditions). The lowest observed amount that DMs keep for themselves is £5 (everyone receives £1) and the maximum amount is £25 (nothing for the recipients). The graph shows that there is a strong positive cor-

relation between the amount kept by DMs and punishment of DMs by recipients. This result is consistent with findings from (single) dictator game experiments with punishment (Bering, 2008; Henrich et al., 2006; Fehr and Fishbacher, 2004) and confirms our conjecture that recipients respond to an overall collective offer in a fashion similar to what we observe when they receive an offer from a single DM.

Figure 1: Average total deduction points used for amount kept by DMs

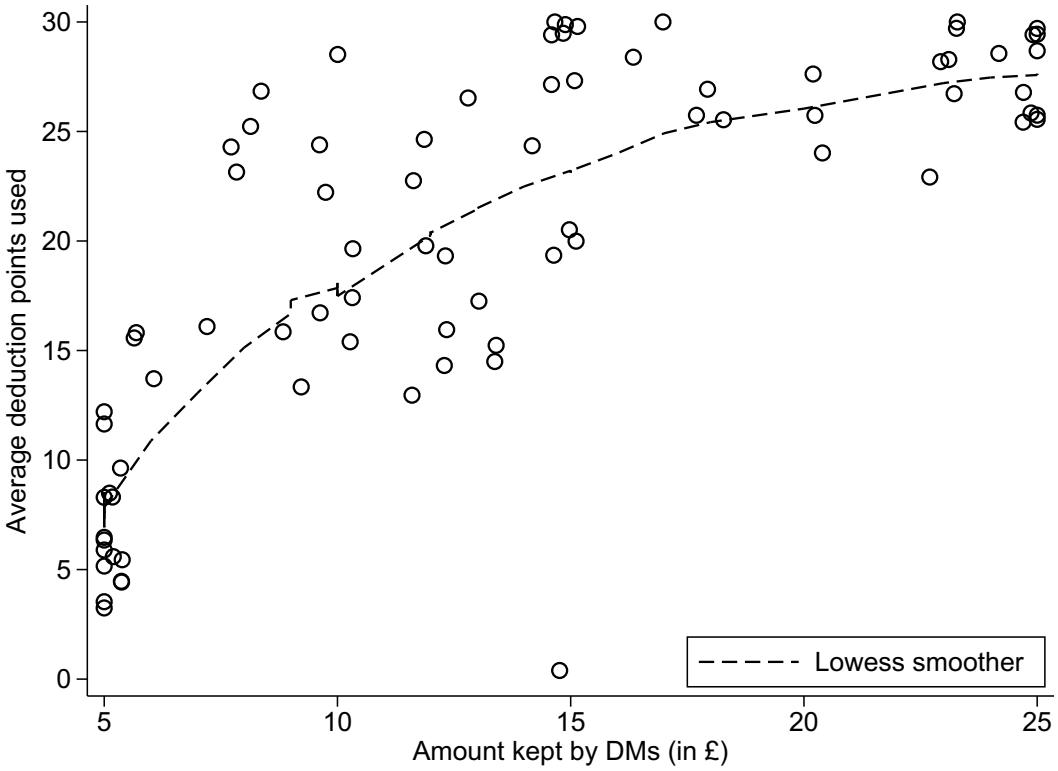

2.2.2 Distribution of punishment

We now turn to the second type of recipients' behavior – the way recipients distribute punishment amongst the individual DMs – to identify the responsibility attribution heuristics they use. Recipients in our experiment appear to use six different heuristics:

- "No punishment" – Recipients used no deduction points.

- "Punish all DMs equally" – Recipients used deduction points and gave equal shares to all DMs. In the vast majority of cases, they used all 30 of their available points (i.e., six points to each DM).

- "Punish only proposer" – Recipients only used deduction points on the proposer. In 85% of such cases they used all (30) of their points on the proposer. In 12% of the cases, the proposer was also the DM with the largest voting weight; in 5% of the cases, the proposer also had a majority of the vote weight.

- "Punish proposer and DM with most weight" – Recipients split their deduction points between the proposer and the DM with the largest vote weight (which was not the proposer in these cases). On no other DM were points used.

- "Punish equally, but more to proposer:" – Recipients used most points on the proposer, more than on other DMs, on each of which they used the same amount.

- "Punish proportional to vote weight, except for proposer" – Recipients did not fall into any other category and used their deduction points in a way that resulted in a positive relationship with vote weight. However, they used disproportionately more points on the proposer.

Clearly, the punishment heuristics our subjects used heavily emphasised the proposer. Recipients paid much less attention to the relative vote weights of the DMs. Indeed, in only 24% of the cases overall did subjects use a punishment heuristic that depended at all on the vote weights of the DMs and in about half of these the only vote weight information that was used was the identity of the largest DM.

Figure 2 consolidates the heuristics above into those with generous or "fair" allocations (1-9 to the DMs), "mixed" allocations (10-14 to the DMs), and selfish or "unfair" allocations (15 or more to the DMs). Figure 2a makes it clear that as allocations become selfish, the majority of subjects appear to blame the proposer; even when this proposer

is not the largest DM. Figure 2b summarises the information recipients employ in their punishment decisions. Clearly, as allocations become less fair (i.e. mixed), recipients use more of the available information overall. However, as they become unfair, they focus less on making distinctions about the sizes of all DMs – only about 20% employ this punishment heuristic in the unfair allocation cases. In most unfair cases (80%), recipients focus exclusively on punishing the proposer without considering the sizes of DMs. Finally, we have cases in which some information about DM size is used (proportional allocations and special punishment of the largest parties). But in all these cases recipients also accounted for the proposer.

Figure 2: Punishment Heuristics and Fairness

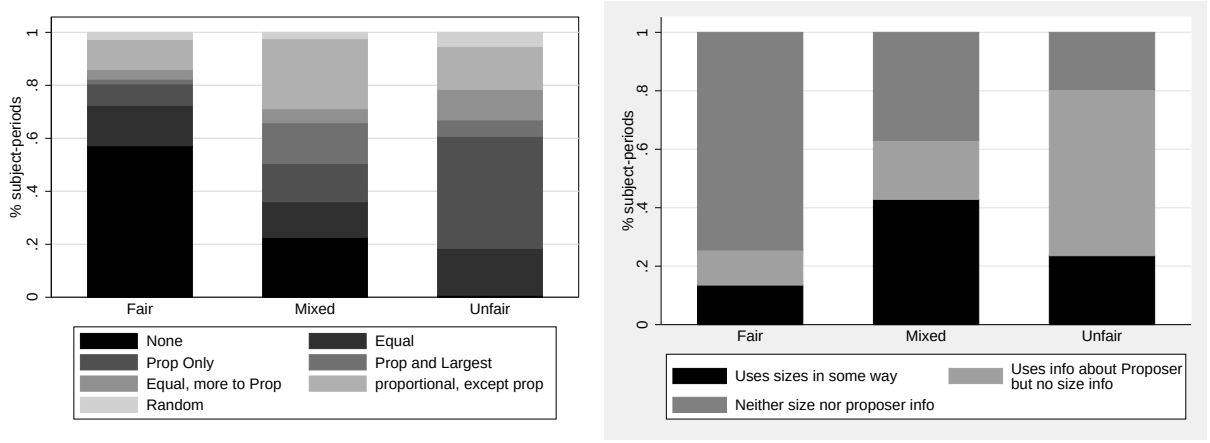

(a) Heuristics by Fairness

(b) Information Content of Heuristics

2.2.3 Multivariate Model

Figure 2 quite compellingly suggests that responsibility attribution is focused on DMs with proposal power rather than on the voting weights of DMs. Accordingly, we estimate a series of statistical models in which we model how much recipients punished DMs with different characteristics, while controlling for the total amount the DMs kept for

themselves.[7] We include a full set of indicator variables for DMs' voting weights and a dummy variable for whether the particular DM was the proposer. Note that the indicator for the largest DM heuristic will be the indicator for the largest voting weight in each distribution (i.e., 23, 38, 48, or 52).

Each row in our data records the number of deduction points (out of 30) that a recipient allocated to a DM in an experimental period. Since there are five DMs in each period, we have five allocations of deduction points from each recipient in each period. In order to identify heuristics about responsibility attribution we focus on the share of the total deduction points that the recipient allocates to each DM. Such shares will sum to one for each recipient who allocated at least one deduction point to one of the five DMs. Hence, our dependent variable is a five-part composition that gives the share of the total deduction points allocated to each DM by a recipient. We follow Atchison's (1986) recommendations in estimating this as compositional data (see Section 2.2.1 in the online appendix for estimation details).

Figure 3 gives a substantive picture of the results for a case in which the DMs' collective choice was to keep £20 for themselves and a recipient's total punishment was 30 points.[8] The circles with the light grey confidence bands indicate the estimated share of punishment allocated to proposing DMs with the corresponding voting weight on the x-axis. The circles with the dark confidence bands are estimates of the average punishment allocated to non-proposing DMs (these are averaged over different proposer treatments for this distribution). For example, in the upper-left distribution (2, 6, 10, 29 53), when

[7]We have also run models without this control, which would be an appropriate specification if the DM's allocation is a "collider" variable in the causal link between our treatments an punishment decisions. The results, however, do not change appreciably for this specification (see Section 2.2.5 in the online appendix).

[8]Substantive results calculated at other allocations give the same message, though consistent with the descriptive results reviewed in the last section, the impact of the proposer and largest parties relative to more equal allocations declines as DMs are more generous. In Section 2.2.6 in the online appendix, we provide similar graphs that give the estimates separately for cases in which each other DM was the proposer (e.g., separate estimates for the predicted share of punishment for a DM with vote weight 6 in distribution 1 when each of the other DMs - with weights 2, 10, 29, and 53 respectively was the proposer).

the DM with a weight of 2 is the proposer, recipients allocate about 20 percent of their deduction points to this DM. This compares to the circle at essentially zero, which is the average estimated share of deduction points allocated to non-proposing DMs with a weight of 2. Note, moreover, that the estimated deduction points for a DM proposer with vote weight 53 is very similar to that of the average estimated deduction points for non-proposers with this weight – both approach 100%.

There are three results that stand out in Figure 3. First, a proposing DM gets punished more than non-proposing DMs with the same vote weight. Although the simulated 95% confidence bands sometimes overlap, in all cases are the point estimates for proposers and non-proposers with the same vote weight statistically different from one another. Thus, consistent with our analysis of these data in the previous section, subjects use proposal power as an important guide to attributing responsibility to DMs. Secondly, the plurality DM receives more punishment, and punishment increases as the size of the plurality DM increases in size. Finally, the results very clearly show that, other than the plurality effect, voting weight is not used by subjects as a general cue for attributing responsibility. There is no clear pattern between voting weight and punishments that applies across the distributions, or even within individual distributions.

Two more subtle results can be discerned in Figure 3. The first is that there appears to be an interaction effect between proposal status and DM size. Specifically, we see that when both the plurality signal and proposal power signal point to the same DM, subjects almost universally punish only this DM (with a predicted share of punishment almost at 100% even in the most egalitarian distribution of vote weights). Finally, our estimates do not suggest that subjects gave any special weight to the fact that in the first distribution the plurality DM could veto any proposal. To see this more clearly, Figure 4 gives the same estimates as in Figure 3 but only for the plurality DM in each distribution. The lower fitted line captures estimated punishment shares that are solely accounted for by

Figure 3: The Impact of Voting Weights and Proposal Powers on Punishment

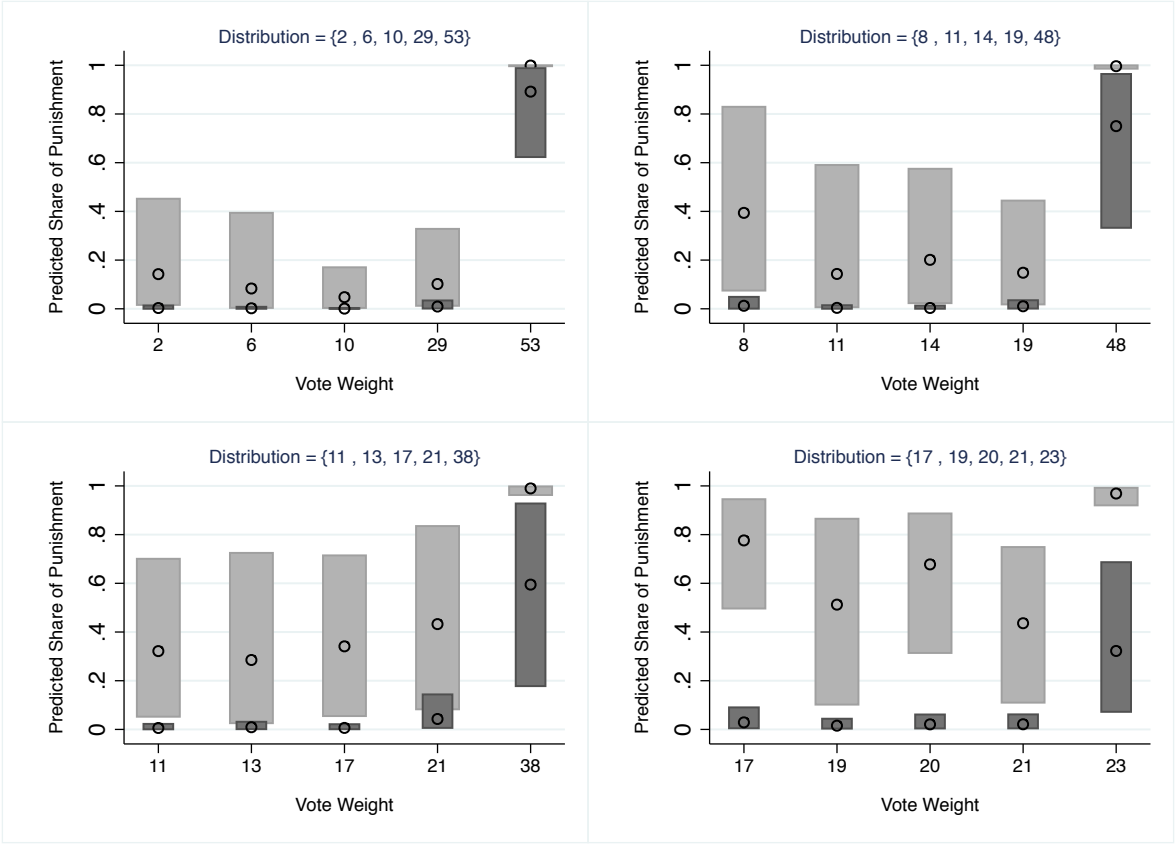

Dark gray bars are estimates for non-proposers and light gray bars are for proposers. The y-axis is the share of a recipient's total deduction points that was allocated to a DM with the indicated characteristics. The hollow circles are point estimates of predicted punishment shares for a typical recipient (who allocated all 30 of her possible deduction points in a situation in which the DMs kept £20 of the initial endowment). These predicted effects are derived from estimated compositional model described in Section 2.2.1 and 2.2.4 of the online appendix).

the plurality size effect.

Examining the estimates in this way, one can immediately see that recipients, while they did punish larger plurality DMs more than smaller plurality DMs, did not give an additional increment of punishment to (plurality) DMs having veto power. If this were the case then the estimated punishment for the veto DM (53 voting weight) would have been significantly above the plurality effect fitted line. Had recipients been employing a veto power heuristic, one might have expected them to only punish this DM, even when she did not have proposal power. However, given the point estimate and confidence band

Figure 4: The Impact of Voting Weights and Proposal Powers on Punishment

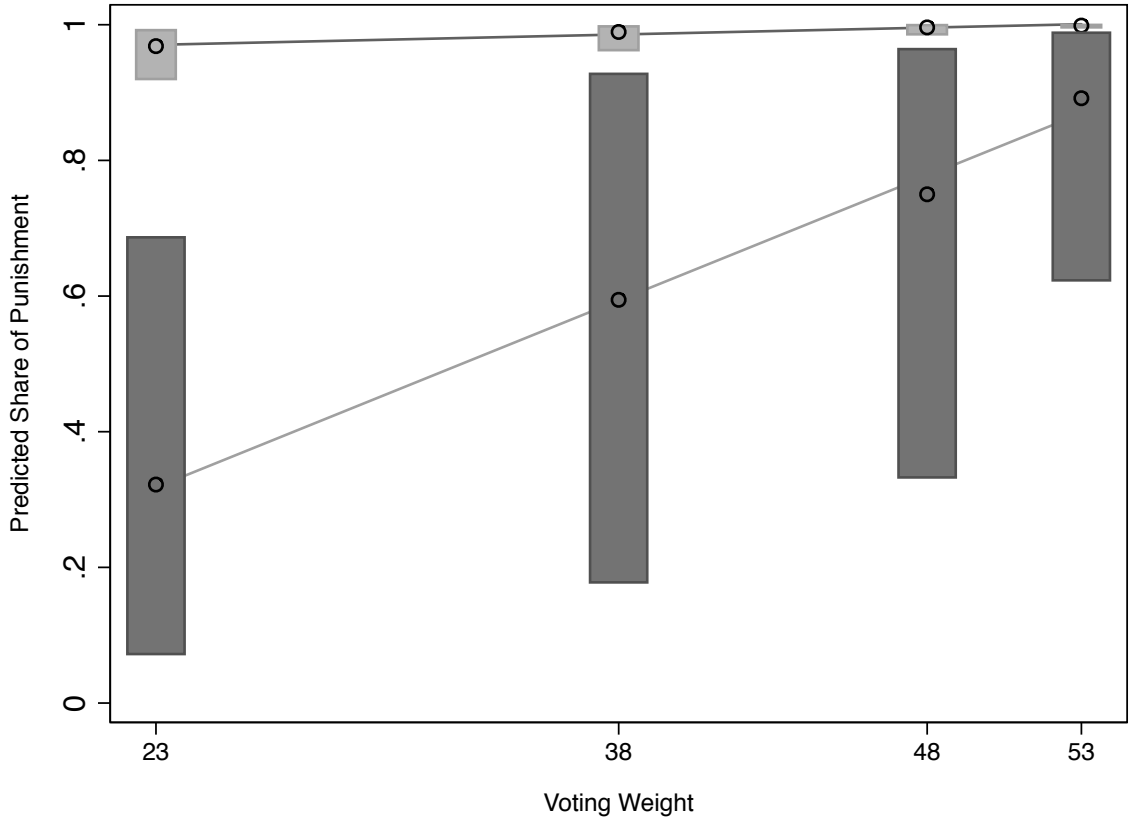

Dark gray bars are estimates for non-proposers and light gray bars are for proposers. See note to Figure 3 for more details.

this is clearly not the case (deduction point shares are about 85% when the DM does not have proposal power, which is significantly lower than the essentially 100% when the DM has proposal power). Thus, unlike the case of proposal power, subjects seem less able to recognise the special significance of veto power.

2.2.4 Information Treatments

Figure 1, presented earlier, suggests that recipients respond to a collective dictator offer in much the same way they would respond to one from a single dictator. We also conjectured that recipients' overall punishment is not conditioned on information they receive regarding the decision-making process (e.g. DM characteristics). Our view is that

recipients respond to the DMs' decision with an overall punishment and then use what information they can glean about the process to apportion responsibility. An alternative perspective is that overall punishment is the sum of recipients' punishment decisions for each individual DM and may therefore depend on the possibility to attribute responsibility to individual DMs. This perspective implies that information about the decision-making process determines the overall level of punishment – limited information should reduce overall punishment. We can test this conjecture with our experimental data.

In addition to the full-information treatment that we referred to in the last two sections, we implemented two additional information treatments in our experiment. In the no-information treatment, subjects neither knew the DMs' voting weights nor which one of the DMs was the proposer; in the semi-information treatment, subjects knew the DMs' voting weights but not which of the DMs was the proposer (see Section 2.1 in the online appendix for details). In order to test whether overall punishment varied depending on the amount of information provided to recipients, we regress the total number of deduction points used by a recipient in one round on indicator variables for each information treatment.[9]

There is at best weak evidence that levels of punishment are conditioned by the information context. The coefficient estimates are 26.7, 27.1 and 29.2 for the no-information, semi-information and full-information treatments, respectively. Hence, while the order of the effects suggest that punishment is higher in environments with more information the pairwise differences between coefficients are insignificant (at the 5% level for two-sided test).

This result provides an interesting insight into responsibility attribution for policy decision made by multi-party governing coalitions. There are claims that multi-party governing coalitions, i.e., the information context, make it more difficult for voters to

[9]Note that in the no-information treatment, DMs kept either £20 or £25 – accordingly the regression we run compares punishment levels across information treatments for amounts kept by DMs of £20 or more.

assess responsibility for policy decisions and hence, there is an attenuation of overall responsibility attribution (Powell and Whitten, 1993). Our information treatment results suggest that there is a tendency for overall punishment to decrease as the amount of information about the decision-making process decreases, but in our case, this decrease is substantially insignificant.

2.3 Discussion

Our first experiment was designed to 1) determine how recipients' overall punishment is related to an offer from a collective decision-making body; and 2) understand the heuristics employed to distribute responsibility amongst individual DMs. In the case of the first point, we find that the responses of our recipients to a collective dictator game offer is similar to the responses to an offer from a single dictator game. Moreover, varying information about the decision-making process, in particular about the DMs in a collective decision making game, has very little effect on overall punishment.

With respect to the second point, we assess the importance of four responsibility attribution heuristics: the relative weights of the DMs; the plurality DM; the DM with veto power over a decision; and the agenda setting DM. Rather surprisingly, with the exception of punishment for the plurality DM, this experiment provides little support for the notion that individuals apportion responsibility for collective decisions according to the voting weights of the DMs. The results are also surprising in that the subjects exhibited a clear tendency to hold the agenda setter accountable. Less surprising, but very prevalent, is the tendency for subjects to punish the plurality DM; and, as the plurality DM's vote gets larger, this punishment increases in a very linear fashion. Indeed, given this linear trend, it is not clear that there is a "bonus" for being above a majority threshold (i.e., a veto player) that comes *in addition* to being the largest DM. Overall, it seems that subjects are looking for clear cues like proposal power and plurality weights. When

they find these cues they use them while ignoring other information. When they do not find any cues – or find limited cues – there is a very modest decline in overall punishment levels and, as one would expect, punishment is randomly allocated to DMs.

One is tempted to conclude from this that subjects in our experiment believe that proposal power and plurality status confer significant influence on collective decisions while vote weight and veto power are less reliable indicators of policy influence. But this is not necessarily the case. There are two alternative explanations that we believe deserve careful attention. First, given the setup of our game, it is possible that higher-level reasoning will lead recipients to punish mainly the proposer; and second, recipients' focus on the proposer may be due to reasons other than our hypothesised one – for example the proposer could be a focal point for the recipients expression of anger rather than a deliberately chosen target of responsibility attribution. Our next two experiments address these issues.

3 Experiment 2: Reasoning about Strategic Proposal Power

In contrast to our conclusion that subjects are using a proposal heuristic it is possible that they are instead undertaking a more sophisticated calculation, that results in the same behaviour (i.e., focusing punishment on the proposer). Specifically, it could be that recipients understand that the proposing DMs in our game is a policy dictator. That is, the proposer can propose her ideal allocation because in the third round of voting this allocation will be preferred to the reversion outcome, which is zero for each DM. Hence, punishing the proposer may not reflect the operation of a heuristic, but a more sophisticated level of reasoning in the game. We test this possibility with an additional lab experiment. In this experiment, we implement a treatment in which the reversion

point of the DM bargaining game is no longer zero, but rather an equitable allocation (Safe condition). Thus, in this treatment, the proposer cannot leverage the "reversion to zero" outcome in order to ensure her preferred allocation wins a majority vote. In contrast, for non-treated rounds, we continue to have a defeated proposal result in zero payoffs (Loss condition). If the recipients' punishment behaviour reflects reasoning about the leverage that the reversion to zero outcome gives the proposer, then proposers should receive fewer deduction points in the Safe treatment than in the Loss treatment.

3.1 Procedure and design

Twenty-four subjects participated in our second experiment, in which we varied three factors across 20 rounds. As in the previous experiment, we varied the vote distribution randomly assigned to the five DMs, as well as the identity of the proposing DM, in each round. Unlike in our first experiment, the proposer faced a binary decision. The proposer could choose between Allocation I, where both DMs and recipients receive £1 each, and Allocation II, where DMs receive £3 each and recipients receive £0.50 each. If the proposer chose Allocation I, Allocation I was implemented (Outcome 1). If Allocation II was proposed it was voted on, and if it received at least 51 votes, then it passed (Outcome 2). The third factor varied the outcome if Allocation II did not receive a majority vote (Outcome 3). In some rounds, if Allocation II received less than 51 votes, Allocation I was implemented (Safe condition); in other rounds, no one was paid anything (Loss condition). It was announced in each round what the Outcome 3 could be. In this case we used the "Strategy Method" to maximise the number of observations. Specifically, in each round the 19 recipients were presented with the three possible outcomes on their screens and were asked to enter the amount of deduction points they wanted to use on each of the five DMs for each of the three outcomes. For each of the outcomes, they had 30 deduction points to divide across DMs. This resulted in $19 \times 3 \times 5 = 280$ punishment

decisions per round.

3.2 Results

We only present a brief summary of the results of this experiment; primarily focusing on the plausibility of the conjecture that recipients expect proposers to exploit their leverage from controlling the proposal in the last round of the game. The results summarised in Figure 5 suggest that this is not the reasoning that motivates recipients' punishment of the proposing DM (PDM). The most informative comparisons in this regard are between punishment targeted at a proposer in the Loss condition versus the Safe condition. As is clear in every case, the punishments are essentially the same across the Safe and Loss conditions, regardless of the other parameters of the game.

Figure 5: Average deduction points assigned to proposing and non-proposing DM by treatments

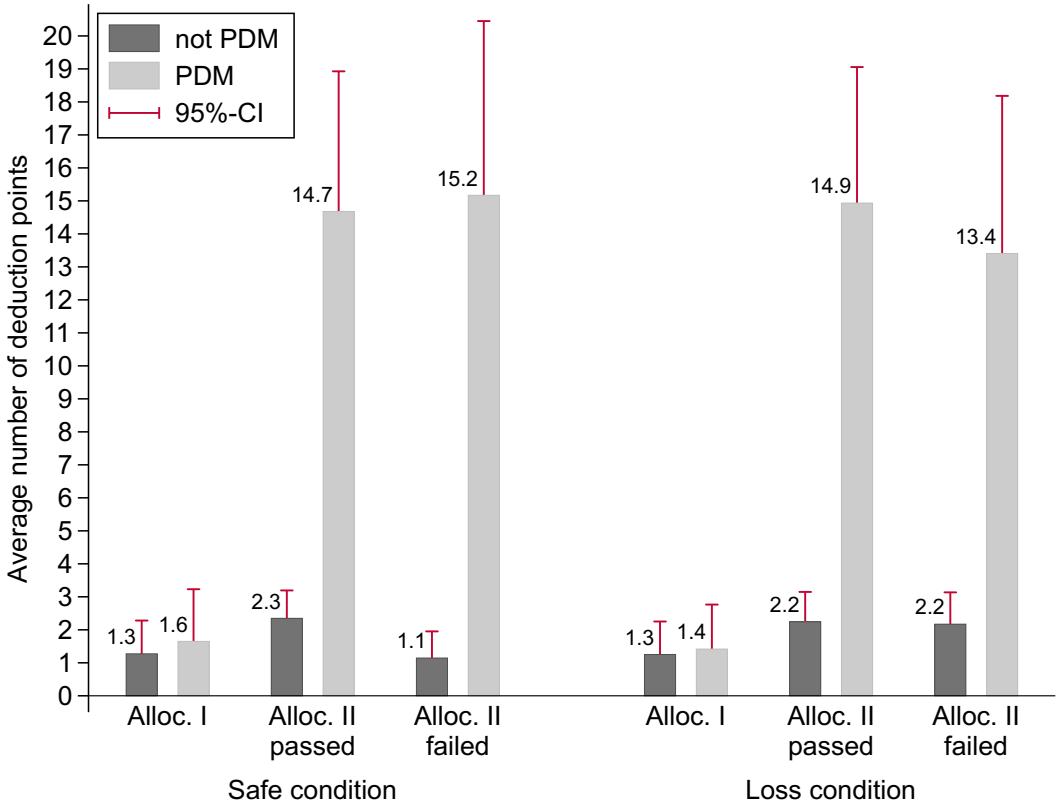

Thus, these results support our earlier conjecture that individuals employ a proposal power heuristic when attributing responsibility for collective outcomes, rather than working through the strategic logic of the particular decision making situation they face. Of course, it could be (and is likely in our opinion) that one of the reasons the proposer heuristic is used here is that in many similar situations it "works" because in those situations proposal power really is valuable. But it would seem, certainly based on the experiment we just described, that individuals employ this proposal power heuristic irrespective of whether or not proposers are especially advantaged in the current situation they face.

4 Experiment 3: From Punishment to Decision-Making Influence

Our interpretation of the results given above is that recipients use a proposer heuristic because they think the proposer has the greatest influence on the decision. of an expression of anger than a rational assessment of responsibility. Emotional psychological models of voting behaviour certainly suggest that the expression of anger in vote choice decisions are associated with an ability to focus blame on a specific target (Brader, Groenendyk and Valentino, 2010; Ekman, 2003; Huddy, Feldman and Cassese, 2007). Perhaps the subjects who are angrier about unfair allocations seek a focus for that anger rather than trying to assess who actually influenced the policy outcome.[10]

Our third experiment aims to determine whether the proposer-centric punishment identified in the previous experiments results because individuals think that the proposing

[10]One can speculate that these two different "mental models" – proposers actually affect policy outcomes versus proposers represent an emotional focal point for blame – can have quite different implications for voter behaviour. In the former case, it will be difficult for proposers to escape responsibility for policy outcomes and this will likely affect their policy proposals. The emotional explanation likely allows for more room to manoeuvre on the part of decision makers in the sense that a proposer can attempt to place public attention on other focal points.

DM has more influence over the collective decision. The experiment is designed to recover expectations regarding decision making influence. It does not invoke punishment or reward on the part of subjects. This is an online survey experiment based on a sample of 1004 UK respondents.[11]

4.1 Procedure and design

Respondents were asked to guess the outcome of a collective decision that had been decided by the weighted vote of five DMs prior to the survey. Depending on how closely their guess matched the collective decision, they could earn from 0 to 30 SSI points (more accurate guesses paid more).

The outcome of the collective decision was described to the respondents with a picture that illustrated the information (these screen shots are included in Section 3.1 in the online appendix). Respondents were told that five DMs had been given a total of £30 and the DMs were asked to decide, based on a majority weighted voting rule, how much should be given to two different charities – an animal shelter and a soup kitchen. The decision making situations differed in three ways. First, in the picture respondents were shown a one-dimensional space that located each DM's preferred donation amounts for the animal shelter (the balance going to the soup kitchen): £4, £10, £16, £21, £28. Second, the voting weight associated with each of the five DMs was also included in the picture. Three different voting weight distributions were used: [.02, .06, .10, .29, .53]; [.11, .13, .17, .21, .38]; and [.17, .19, .20 .21, .23]. The voting weights in each distribution added to 1 and they determined how much each DM's vote counted when the donation proposal was voted on. Finally, the picture also indicated which DM was chosen to propose the allocation of the £30.

[11]The Internet panel sample was provided by Survey Sampling International (SSI). Their panelists are compensated with SSI points which are exchanged for money. A complete description of the survey and sample characteristics is available from the authors.

Each respondent was asked to guess the outcome for three collective decisions taken under the three different distributions of voting weights assigned to the five DMs. Within each distribution, the particular allocation of weights over the five DMs (i.e., over the five ideal points) that a respondent saw were randomised (there were 120 possible assignments of weights to positions). In addition, the identity of the proposer that each respondent saw (for each decision) was randomised (so there are 25 possible assignments of a weight and position combination to a proposer in each decision situation). Given the 1,004 respondents, this results in about 40 respondents per proposer/position combination.[12]

For each respondent, for each of the three questions, we calculated the spatial distance between the respondent's guess about the collective decision and the ideal points of each of the five DMs (below, we refer to this variable as distance). This gives us five data points for each respondent for each of the three questions. Taken together, these data points contain information about the respondent's beliefs about the relative influence of different DMs (with different seat weights, positions, and agenda powers) had on the collective decision.[13]

The goal of the empirical analysis reported below is to explore which characteristics (or combinations of characteristics) of DMs condition respondent's beliefs about the outcomes of collective decisions. Building on the results from the lab experiments reported above,

[12]While there are technically 120 $\times$ 5 cells in our design matrix for each decision (i.e., assignments of weights to positions $\times$ assignment of proposer) and so a small number of respondents in each cell, this is not consequential for our study since we are only concerned with two aspects of vote weights: (1) whether there is parametric relationship between weights and how close respondents think the policy will be to a given DM, and (2) whether there is a "largest DM" effect. Since we can examine the first of these by fitting parametric models across cells of this design matrix, the information from a relatively small number of respondents in each cell is easily aggregated to bear on the main questions of interest here.

[13]We do not assume that a respondent who believes the policy outcome will be near the ideal point of a given DM necessarily thinks that the party is influential. It could be, for example, that the respondent believes the policy will end up being in the middle of the policy space, because two relatively influential DMs with opposing and extreme preferences compromise on the policy, leaving it close to a more centrally located but non-influential DM. In the empirical analysis we attempt to isolate this sort of incidental influence from beliefs about real influence by focusing on how the addition or removal of DM characteristics (like agenda powers) change respondent beliefs while holding constant the ideological positions of the parties.

we focus on two kinds of characteristics: agenda setting powers – both positive (proposal) and negative (veto) – and the distribution of voting weights (including which DM has the largest weight as well as other more extensive uses of size cues, e.g., proportional influence). The results reported in the following are based on the estimation of three separate regressions (one for each question) of our distance variable on measures of the agenda powers, vote weights, and policy preferences of DMs (see Section 3.3 in the online appendix for details).

4.2 Results

Figure 6 summarises the multivariate results for the survey experiment. The graph presents the distance between respondents' predicted policy outcome and the policy position of the DM with the indicated characteristics. There are three graphs corresponding to each of the three distribution treatments. Each graph presents the distances between the predicted policy outcomes and the DM policy positions for DMs that were proposers (with darker confidence bars) and for those who were non-proposers (with lighter confidence bars). A line has been fitted between these two sets of distances excluding the DM with the most weight.

First note that none of the slopes of the lines graphed are statistically significant from zero. They are not even close. So working from the top line in the first graph (proposers for Distribution 1) to the bottom line in the third graph (non-proposers Distribution 3) the p-values for each line are: .37, .29, .30, .38, .56, and .87. We have no evidence that the distance between the expected policy and the DM's position becomes smaller as voting weight becomes larger.

Note moreover that the distance between the respondent's predicted policy outcome and the policy position of the proposing DM is always smaller than it is for the non-proposing DMs. And in most cases the confidence intervals on the predictions do not

overlap. Hence it is not surprising that in the three graphs in Figure 6, the fitted line for the proposer is consistently lower than it is for the non-proposer. Our principal finding from the lab experiment is confirmed: responsibility attribution is strongly associated with agenda power.

Being the largest DM also sometimes matters, and even more if one is also the proposer. Figure 6 shows the interactions that result from the non-linear model. The fitted-line indicates how much closer the expected policy is to the largest DM than it would be if one projected only a size effect assuming proportional influence. In two cases the point estimate for the largest DM is below the line. Hence the largest vote weight heuristic, for which we do have evidence, is quite distinct from a voting weight heuristic, for which there is limited evidence. However, in the most equal case (Distribution 3), the point estimate for the largest DM is not different from the point estimates for the other DMs. In this distribution case there is no largest DM effect.

Finally, being the veto player does not appear to matter. The point estimates for the majority DM (for both proposing and non-proposing DMs) is not really different from those for the large but non-majority DM (who has .38 vote weight). Thus, there is no real majority or negative veto effect apparent when one accounts for being the largest.

4.3 Discussion

The lab experiments were designed to recover responsibility attribution heuristics subjects employ to punish individual DMs who take a collective decision that affects them. The online experiment, with a representative sample of the UK population, was designed to test whether the attribution heuristics recovered in the lab experiment result from an assessment of the relative importance of individual DMs for the collective decision. Accordingly, in this online experiment, subjects were asked to anticipate collective decision outcomes under a number of different treatments that varied the characteristics of the

Figure 6: Predicted Distances between Respondent's Guess and DM Position

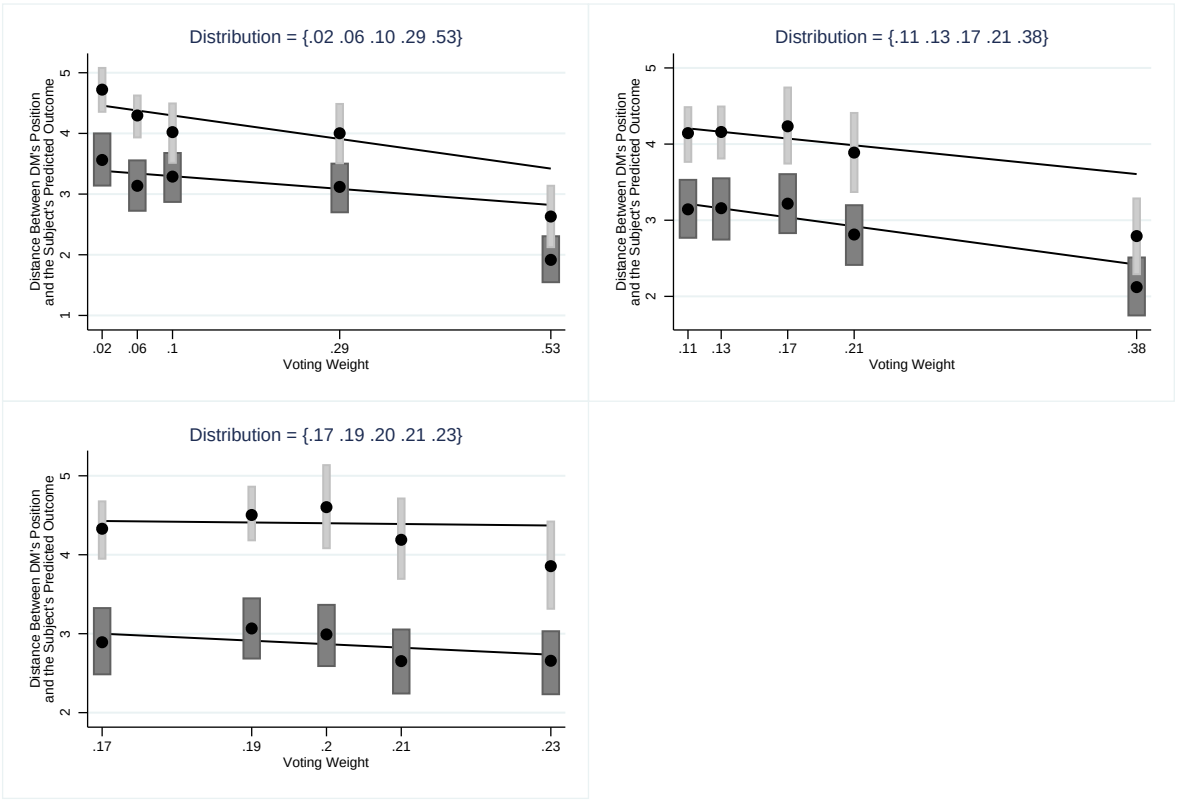

Lighter confidence bars are for non-proposers; Darker bars are for for proposers.

Fitted lines are calulated using only cases other than the DM with the most voting weight. None of the slopes are statistically significant.

DMs. The results confirm the two negative results of the lab experiment that voting weights and negative agenda power are of limited importance for responsibility attribution. The results also confirm the two positive lab experiment findings: respondents are inclined to attribute responsibility for collective decisions to the DM with agenda power and in some cases to the DM with the largest percentage of votes.

5 Responsibility Attribution and the Coalition Directed Vote

Our conclusions regarding responsibility attribution for collective decision makers are of political importance because voters in contexts with multiparty governing coalitions

face a challenge similar to those of the subjects in our experiments. In order to exercise a coalition directed vote they need to attribute responsibility to the parties in the governing coalition or anticipate the administrative responsibility of parties that are likely to form the governing coalition after the election (Duch, May and Armstrong, 2010; Duch and Stevenson, 2008). Our experiment captures the important elements of this voting decision.

First, key features of the decision making context in these lab experiments resemble those characterising coalition cabinet decision making. Most importantly, each party's portfolio allocations in the cabinet typically reflect their relative seat strength in the legislature. We can think of the number of ministerial portfolios as the party's voting weight in cabinet decision making. It is widely assumed that some form of portfolio-weighted majority voting is employed for reaching decisions in coalition cabinets. And cabinets have a procedure for bringing forward proposals for a formal or informal vote – our contention is that the public identifies a party as the proposer and that this is typically the Prime Minister's party. Accordingly, the DMs, or "parties" in our experiment, are randomly assigned the voting weights, one DM is randomly assigned to be the proposer, and the group of DMs take their decisions by majority vote.

Second, coalition directed voters, like the recipients in our experiment, need to figure out how to attribute responsibility to individual DMs (parties) in a collective decision making entity (the cabinet). Whereas our recipients had punishment points to distribute, citizens can support or oppose different political parties and can adjust those levels based on observed policy outcomes. Importantly, we think of this level of support or opposition as an unobserved sentiment that individuals hold that might find expression in any number of different activities and behaviours (e.g., voting, political contributions, time devoted to campaigns or party activities, persuading friends and neighbors, or even simply making relevant statements in everyday conversation).

But the behaviour of the recipients in our experimental setup does not entirely correspond to the prevailing characterisations of how voters attribute responsibility for multiparty coalition policy outcomes. It is fair to say that most efforts to characterise how voters attribute responsibility to the individual parties in a governing coalition have favoured a weighted voting model. This notion that influence over coalition outcomes should reflect the proportion of seats allocated to coalition parties is in line with Gamson's "Law of Proportionality" (Gamson, 1961, 382) and is widely accepted as the basis for bargaining amongst coalition partners for portfolio allocations (Bueno de Mesquita, 1979; Browne and Frendreis, 1980), also because there is evidence to this effect.

Anderson (1995, 210) finds 'that more responsibility in the government results in the economic variables having stronger effects on party support.' Duch and Stevenson (2008) argue that the voters condition their economic vote on the distribution of seat shares within cabinet government. Kedar (2009) similarly contends that voters will equate seat share with administrative responsibility for cabinet policy decisions. It is entirely possible, then, that individuals' attributions of responsibility are increasing in the sizes of the parties. However, our experimental results are at odds with this literature. They quite clearly suggest that voters are not likely to attribute responsibility for collective decisions according to the relative voting weights of the individual DMs.

While the experimental results suggest that voters do not calibrate their responsibility attribution according to relative voting weights, they clearly seem to attribute a disproportionate responsibility to the largest party in the coalition cabinet. This conclusion is not inconsistent with other empirical findings. The largest party is the target of disproportionate voter responsibility attribution (Anderson, 1995) although this effect is typically characterised as an artefact of the relative voting weight effect discussed above. Our experimental findings suggest that the continuous effect might be overstated in the literature. The experimental treatments allowed us to carefully distinguish between the

largest DM effect and the (continuous) effect of DM voting weight, and the evidence is in favour of the former. In our view, empirical models of voter responsibility attribution for parties in governing coalitions should give more attention to distinguishing these two effects. Our intuition is that these efforts would favour the largest party heuristic for responsibility attribution.

The most intriguing positive result from the experiments is that subjects punished disproportionately, irrespective of their voting weight, the decision maker with proposal power. We do not pretend to account for the genesis of the proposal power heuristic in this essay; this is clearly a subject for further research. Nevertheless, some candidates suggest themselves. There is strong evidence from experimental psychology that individuals bias to favour omissions over commissions that cause harm (Cushman, Young and Hauser, 2006; Spranca, Minsk and Baron, 1991). These biases might account for the prevalence of a positive as opposed to negative agenda setting heuristic: People morally accept "passive" actions (i.e., the veto player who simply does not use his veto power) much easier than "active" actions (the proposer's actions). In general, the most "active" action taken by the DMs in our experiment can be associated with the DM with proposal power – the responses by the other DMs are reactions and hence in some sense more "passive." Hence, the proposal power heuristic may be founded on the notion that proposing is the most active of the actions taken by DMs in these collective decision making situations. Further experimental research will be necessary in order to better understanding which of these (or other) perspectives best explain the prevalence of the proposer power heuristic.

With respect to the coalition-directed vote, the implication of this result is that voters will attribute responsibility to the party with proposal power in the governing coalition. In one respect this is not surprising since there is a substantial body of literature suggesting that vote outcomes in collective decision making institutions is strongly shaped by the preferences of the agenda setter (Weingast and Marshall, 1988; Cox and Magar, 1999).

With respect specifically to coalition governance there is evidence that Ministers, because they have agenda power related to their portfolios, strongly influence policy outcomes within coalition governments (Laver and Shepsle, 1996). What is novel though is that our experiments demonstrate that voters likely attribute responsibility to the party in a coalition with positive agenda power.

In light of our findings, the question whether voters are informed about positive agenda power clearly deserves further exploration. Which cabinet party, if any, is analogous to the proposer in our experiment? An obvious candidate here is the chief executive, or Prime Ministerial, party in the coalition. A review of the literature on coalition decision making suggests that the Prime Ministerial party has disproportionate influence on coalition policy outcomes and is quite likely to be viewed by the public as exercising positive agenda power. Moreover, there is some evidence to suggest that a high proportion of voters in Western Democracies are better able to identify the name of the Prime Ministerial party.[14] All of this is simply to point out that there is some limited evidence indicating that in the minds of the voters positive agenda power is associated with Prime Ministerial party.

The challenge is to determine whether this positive agenda power result, identified in the controlled experimental setting, helps us explain responsibility attribution in the general voting population? Two empirical issues are of particular interest: Firstly, are voters who recognise the importance of proposal power more likely to hold policy makers accountable? Secondly, are those who recognise the importance of proposal power more likely to attribute responsibility to the individual party in the coalition with proposal power. In a follow-up study based on the online survey experiment data, the authors [] model the economic vote and establish that it exhibits both of these characteristics.[15]

[14]Evidence that voters can easily distinguish Prime Ministerial parties from their partners comes from the Dutch survey conducted by Fortunato, Lin and Stevenson (2013). 75 percent of respondents correctly identified the PM party (out of a choice of 10 parties) and almost 60 percent identified the junior cabinet partner.

[15]That paper uses the data from the online experiment described in Section 4. The authors use a random coefficient model to estimate an individual-specific effect of proposal power on responsibility

Only those individuals who understood the value of proposal power in the experiment exercised an economic vote (i.e., rewarded or punished the parties based on their evaluation of the economy). Moreover, those who recognised the importance of proposal power focused all of their economic vote on the Conservative Party; the Prime Ministerial party in the current Conservative-Lib-Dem governing coalition with clear proposal power for economic policy. This finding suggests that the attribution behaviour we identify in the lab maps very nicely onto actual voting behaviour in the general population and hence strengthens the external validity of our results.[16]

6 Conclusions

This essay reports the results of three experiments that identify the precise heuristics that individuals employ in allocating responsibility for collective decisions arrived at by majority vote. We accomplish this with a novel collective dictator games in which recipients can punish individual decision makers forming the collective dictator. Recipients respond to the collective offer in a fashion very similar to how recipients respond in conventional dictator games, in which the offer comes from a single person. As the collective offer gets worse, punishment rises. Our information treatment results suggest that at best there is a weak but substantially insignificant decrease in overall punishment when the amount of information about the decision-making process decreases. For the most part, these information treatments suggest that the responsibility attribution reflex is one of first responding to the collective decision with an overall punishment (irrespective

attributions (which was possible because each individual made three attributions of policy responsibility).

[16]Fortunato, Lin and Stevenson (2013) report evidence from a 2012 Dutch survey that while perceived party size is an important predictor of voters' attributions of responsibility to a party, there is a distinct and large separate effect for being the largest. Moreover, the overwhelmingly largest predictor of responsibility attributions is perceived prime ministerial status. Indeed, the impact of perceived status as a junior partner actually has a negative impact on perceived policy-making responsibility. Further, as we saw in our experimental results, the impact of being perceived as the largest party is secondary to the perception that the party holds the prime ministry.

of the process by which the decision was taken) and then deciding on how to distribute this punishment amongst the individual DMs (which is sensitive to the decision making process).

However, our experiments were primarily designed to shed light on how individuals distribute responsibility amongst individual DMs in a collective decision making body. We find that the voting weight heuristic is employed with limited frequency, and individuals do not seem to favour the heuristic that assigns responsibility to the DM with negative agenda power. On balance, these two heuristics play a minor role in the attribution of individual responsibility for collective decision making. On the other hand, individuals clearly favour agenda power and the largest vote weight as heuristics for attributing responsibility for members of a collective decision making body. Quite different responsibility attribution experiments produce this same result. Two differently designed laboratory experiments and a third online survey experiment demonstrate that if individuals have the opportunity to hold individual DMs responsible for a group decision, they primarily attribute responsibility to the proposer and the DM with the largest vote weight.

The online experiment also demonstrates that individuals hold proposers responsible because they believe they actually have a disproportionate impact on the outcome (rather than simply constituting a focal point for an individual's anger or enthusiasm). The second laboratory experiment was also designed to determine whether strategic reasoning on the part of recipients explained the decision to focus punishment on the proposing DM. The results indicate this is clearly not the case; again reinforcing our claim that individuals have internalised an agenda setting heuristics when deciding which DM to punish for a collective decision governed by a majority voting rule.

In sum, the experimental results reported in this essay provide a unique understanding of how individuals hold decision makers responsible for the decisions they make collec-

tively. Examples of such collective decision making bodies include boards of directors of firms, legislatures, families, and international organisations. We contend that multi-party coalition government is one of these general collective decision making entities. Accordingly, our experimental results suggest that the party with agenda setting power in the cabinet along with the party with the largest voting weight will bear much of the responsibility for coalition decisions. Veto players will not be held accountable and voters will not apportion responsibility according to the voting weights of parties in the coalition.

References

Aitchison, J. 1986. *The Statistical Analysis of Compositional Data (Monographs on Statistics and Applied Probability)*. Chapman and Hall.

Allison, S.T. and D.M. Messick. 1985. "The Group Attribution Error." *Journal of Experimental Social Psychology* 21:563–579.

Anderson, Christopher J. 1995. *Blaming the Government: Citizens and the Economy in Five European Democracies*. Sharpe.

Banzhaf, J.F. 1965. "Weighted Voting Doesn't Work: A Mathematical Analysis." *Rutgers Law Review* 19:317–343.

Barro, Robert. 1973. "The Control of Politicians: An Economic Model." *Public Choice* 14:19–42.

Bartling, Bjorn and Urs Fischbacher. 2012. "Shifting the Blame: On Delegation and Responsibility." *Review of Economic Studies* 79(1):67–87.

Bering, Jesse M. 2008. "The Effects of Perceived Anonymity on Altruistic Punishment." *Evolutionary Psychology* 6(1):487–501.

Besley, Timothy. 2006. *Principled Agents*. Oxford University Press.

Bolton, Gary E. and Axel Ockenfels. 2000. "A Theory of Equity, Reciprocity and Competition." *American Economic Review* 90:166–193.

Bornstein, Gary and Ilan Yaniv. 1998. "Individual and Group Behavior in the Ultimatum Game: Are Groups More 'Rational' Players?" *Experimental Economics* 1:101–108.

Brader, T., E.W. Groenendyk and N. Valentino. 2010. "Fight or Flight? When Political Threats Arouse Public Anger and Fear." manuscript University of Michigan.

Browne, Eric C. and John P. Frendreis. 1980. "Allocating Coalition Payoffs by Conventional Norm: An Assessment of the Evidence from Cabinet Coalition Situations." *American Journal of Political Science* 24(4):753–68.

Bueno de Mesquita, Bruce. 1979. "Coalition Payoffs and Electoral Performance in European Democracies." *Comparative Political Studies* 12(1):61–81.

Camerer, Collin F. 2003. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press.

Charness, Gary. 2000. "Responsibility and Effort in an Experimental Labor Market." *Journal of Economic Behavior and Organization* 42(3):375–84.

Cox, Gary and Eric Magar. 1999. "How Much is Majority Status in the U.S. Congress Worth?" *American Political Science Review* 93:299–309.

Cushman, Fiery, Liane Young and Marc Hauser. 2006. "The Role of Conscious Reasoning and Intuition in Moral Judgement: Testing Three Principles of Harm." *Psychological Science* 17:76–105.

Dawes, Christopher T., James H. Fowler, J.H. Johnson, R. McElreath and O Smirnov. 2007. "Egalitarian Motives in Humans." *Nature* 446:794–796.

Duch, Raymond M., Jeff May and David Armstrong. 2010. "Coalition-Directed Voting in Multi-Party Democracies." *American Political Science Review* 104(4):698–719.

Duch, Raymond M. and Randy Stevenson. 2008. *The Economic Vote: How Political and Economic Institutions Condition Election Results*. Cambridge: Cambridge University Press.

Ekman, Paul. 2003. *Emotions Revealed: Recognizing Emotions and Feelings to Improve Communications and Emotional Life*. Henry Holt and Company.

Engel, Christoph. 2010. "The behaviour of corporate actors: How much can we learn from the experimental literature?" *Journal of Institutional Economics* 6(4):445475.

Falk, Armin, Ernst Fehr and Urs Fischbacher. 2006. "A Theory of Reciprocity." *Games and Economic Behavior* 54(1):293–315.

Fehr, E. and K.M. Schmidt. 1999. "A Theory of Fairness, Competition and Cooperation." *Quarterly Journal of Economics* 114:817–68.

Fehr, E. and K.M. Schmidt. 2006. *Handbook on the Economics of Giving, Reciprocity and Altruism*. North Holland chapter The Economics of Fairness, Reciprocity, and Altruism – Experimental Evidence and New Theories.

Fehr, E. and S. Gachter. 2000. "Cooperation and Punishment in Public Goods Experiments." *American Economic Review* 90:980–994.

Fehr, Ernst and Urs Fishbacher. 2004. "Third-party Punishment and Social Norms." *Evolution and Human Behavior* 25:63–87.

Ferejohn, John. 1986. "Incumbent Performance and Electoral Control." *Public Choice* 50(3):5–25.

Fortunato, David, Nick Lin and Randy Stevenson. 2013. "What do voters know about cabinet Formation." Rice University Working Paper.

Gamson, William. 1961. "A Theory of Coalition Formation." *American Sociological Review* 26:373–82.

Gerstenberg, Tobias, Maarten Speekenbrink and Catherine Lok Sze Cheung. 2011. "Rational Order Effects in Responsibility Attributions.".

Gilbert, Daniel T., Brett W. Pelham and Douglas S. Krull. 1988. "On cognitive busyness: When person perceivers meet persons perceived." *Journal of Personality and Social Psychology* 54(5):733–740.

Goldstein, Daniel G. and Gerd Gigerenzer. 2002. "Models of Ecological Rationality: The Recognition Heuristic." *Psychological Review* 109(1):75–90.

Guth, W., R. Schmittberger and B. Schwarze. 1982. "An Experimental Analysis of Ultimatum Bargaining." *Journal of Economic Behavior and Organization* 3:367–388.

Henrich, J., R. McElreath, A. Barr, J. Ensminger, C. Barrett and A. Bolyanatz. 2006. "Costly punishment across human societies." *Science* 312:1767–1770.

Huddy, L., S. Feldman and E. Cassese. 2007. On the Distinct Political Effects of Anxiety and Anger. In *The Affect Effect: Dynamics of Emotion in Political Thinking and Behaviour*, ed. W.R. Neuman, G.E. Marcus, A.N. Crigler and M. MacKuen. The University of Chicago Press.

Iyengar, Shanto. 1991. *Is Anyone Responsible? How Television Frames Political Issues.* Chicago: University of Chicago Press.

Jones, Edward E. 1990. *Interpersonal Perception.* New York: Macmillan.

Kedar, Orit. 2005. "When Moderate Voters Prefer Exterme Parties: Policy Balancing in Parliamentary Elections." *American Political Science Review* 99(2):185–199.

Kedar, Orit. 2009. *Voting for Policy, Not Parties: How Voters Compensate for Power Sharing.* Cambridge University Press.

Kiewiet, D. Roderick. 1983. *Macroeconomics and Micropolitics: The Electoral Effect of Economic Issues.* University of Chicago Press.

Laver, Michael and Kenneth A. Shepsle. 1996. *Making and Breaking Governments: Cabinets and Legislatures in Parliamentary Democracies*. Cambridge: Cambridge University Press.

Luhan, Wolfgang J., Martin G. Kocher and Matthias Sutter. 2009. "Group polarization in the team dictator game reconsidered." *Experimental Economics* 12(1):26–41.

Messick, D. M., D. A. Moore and M. H. Bazerman. 1997. "Ultimatum bargaining with a group: Underestimating the importance of the decision rule." *Organizational Behavior and Human Decision Processes* 69(2):87–101.

Morton, Rebecca and Kenneth Williams. 2009. *From Nature to the Lab: Experimental Political Science and the Study of Causality*. Cambridge University Press.

Powell, G. Bingham and Guy Whitten. 1993. "A Cross-National Analysis of Economic Voting: Taking Account of the Political Context." *American Journal of Political Science* 37:391–414.

Rabin, Matthew. 1993. "Incorporating Fairness into Game Theory and Economics." *American Economic Review* 83(5):1281–1302.

Ross, Lee and Richard E. Nisbett. 1991. *The person and the situation: Perspectives of social psychology*. New York: McGraw-Hill.

Spranca, Mark, Elisa Minsk and Jonathan Baron. 1991. "Omission and Commission in Judgment and Choice." *Journal of Experimental Social Psychology* 27:76–105.

Weiner, Bernard. 1995. *Judgments of Responsibility: A Foundation for a Theory of Social Conduct*. New York: Guilford.

Weingast, Barry R. and William J. Marshall. 1988. "The Industrial Organization of Congress: or, Why Legislatures, Like Firms, Are Not Organized as Markets." *Journal of Political Economy* 96:132–163.