Self-Interest, Inequality, and Entitlement in Majoritarian Decision-Making[*]

Daniel Diermeier[1] and Sean Gailmard[2]

[1]*Department of Managerial Economics and Decision Sciences, Kellogg School of Management, and Northwestern Institute on Complex Systems (NICO), Northwestern University; E-mail d-diermeier@kellogg.northwestern.edu*
[2]*Department of Political Science, Northwestern University; E-mail gailmard@northwestern.edu*

ABSTRACT

We experimentally test competing theories of three-player majoritarian bargaining models with fixed, known disagreement values. Subjects are randomly assigned to three roles: a proposer and two types of voters. Each role is randomly assigned a disagreement value, i.e. a given amount of money he/she will receive if the proposal is rejected. These values are known to all players before any decision is made. Proposers then make a take-it-or-leave-it offer on how to split a fixed, known amount of money among the players. If a majority of players accepts the proposal, the players' payoffs are determined by the proposal; if the proposal is rejected, each player receives his or her reservation value. We assess the ability of three behavioral hypotheses – self-interest, egalitarianism, and inequality-aversion – to account for our results. Our primary design variable is the proposer's reservation value, which allows us to obtain different implications from each hypothesis. We find that each hypothesis is inconsistent with our data in important respects. However, subjects strongly respond to changes in reservation values as if they were interpreted as a basic form of entitlement.

Since its publication in 1978, the Romer–Rosenthal model has long been recognized as a powerful tool to study political decision-making. Its key insight is that agents with proposal power may be able to bias the policy outcome away from the pivotal voter's ideal point.

The Romer–Rosenthal approach can easily be generalized to a large class of collective decision problems. The most widely used is the Baron–Ferejohn model of legislative bargaining (Baron and Ferejohn 1989a). The Baron–Ferejohn model is a n-player multi-period bargaining game under a given voting and amendment rule. In all variants of the model, a proposer is selected according to a known recognition rule. He then proposes a policy or an allocation of benefits ("money") to a group of voters.[1] Under the closed amendment rule, the proposal is then either accepted or rejected according to a given voting rule.[2] If the proposal is accepted, the game ends and all actors receive payoffs as specified by the accepted proposal. Otherwise, another proposer is selected, and so forth until a proposal is accepted or, in game forms with finitely many periods, the last period is reached.

Under a closed amendment rule, the Baron–Ferejohn model predicts that the proposer will propose a minimal winning coalition composed of the voters with the lowest "continuation values" which equal the player's equilibrium expected payoffs in case the proposal is rejected and bargaining continues. All other players will receive a payoff of zero. Proposals are always accepted in the first round. Note that the proposing party will always choose as a coalition partner the party with the lowest continuation value. The division of spoils will, in general, be highly unequal, especially if the players' discount factors are low.

Eventually, the usefulness of proposer-pivot models depends on how well it explains behavior in actual multi-person bargaining environments. Recent experimental work on related two-person bargaining games has shown that the models' predictions may fail to be supported in the laboratory. For example, a number of experimental studies have examined ultimatum games (Güth, Schmittberger, and Schwarze 1982) in which one player proposes a division of a fixed amount of money; the other player must either accept or reject, with rejection implying a zero payoff for both. In experiments on ultimatum games, proposers should take (almost) all of the money, yet the divisions are far more equal than predicted. Moreover, if proposers offer less than a threshold amount, the other player frequently rejects the offer (even if a significant amount of money is offered and anonymity is ensured) and thus receives a payoff of zero.[3] Experiments on sequential bilateral bargaining games with alternating offers result in similar outcomes (e.g. Roth 1995).

---

[1]   We denote proposers by male, voters by female pronouns.

[2]   Under a closed rule, the proposal is take-it-or-leave-it and may not be amended before voting, whereas under the open rule, the proposal may be replaced by a sequence of amendments.

[3]   These findings are quite robust to large financial stakes, of anywhere from $100 to a month's wages in experiments conducted in developing countries. The threshold amount for rejections may vary slightly by subject pool, but in many experiments is about 40% of the total pie. See Camerer (2003) for a detailed review of the literature.

Many explanations have been proposed for this finding. One of the most fruitful ones suggests that players are not trying to maximize their individual monetary payoff, but are influenced by moral motivations such as the desire to follow norms of fairness, even under experimental conditions that guarantee anonymity between players. Forsythe, Horowitz, Savin, and Sefton (1994) investigated this hypothesis by comparing ultimatum and dictator games. The dictator game differs from the ultimatum game in that the proposing player proposes a division between the two players, yet the other player (the "passive voter") cannot reject the proposal. In the original dictator game experiments, proposers gave significant shares to passive voters. However, subsequent experiments have shown that variations in the experiment may sharply reduce or even eliminate giving to passive actors (e.g. Hoffmann, Mc Cabe, Shachat, and Smith 1994). This holds, for example, in the so-called "dummy-player game" (Güth and Van Damme 1998, Kagel and Wolf 2001) where one of the recipients has veto power over the allocation (as in the ultimatum game), while the other player is passive (as in the dictator game). Moreover, once competition between responders is introduced behavior closely approximates self-interested behavior (Fischbacher, Fong, and Fehr 2003).[4]

To summarize, the proposer-pivot model shares important similarities with several other types of games with sequential decision-making. Coalition members resemble bargaining partners as in the ultimatum games, while non-coalition members appear to play the same role as passive players in the dictator or dummy-player game. But observed subject behavior in the experimental investigation of these related games seems to point in opposite directions. Therefore, it is important to explore the proposer-pivot model directly. Unfortunately, there has been little experimental research devoted to this task.

In the first experimental investigation of the Baron–Ferejohn model,[5] McKelvey (1991) studies a three-voter, three-alternative stochastic bargaining game under the closed rule. He concludes that the predicted solution to his game at best modestly explains the data. Proposers usually offer too much relative to the symmetric equilibrium and their proposals are accepted with too high a probability.

Diermeier and Morton (2005) use a finitely repeated version of the Baron–Ferejohn model to obtain sharper predictions than the original model investigated in McKelvey (1991). Similar to McKelvey, Diermeier and Morton find little support for the predictions of the Baron–Ferejohn model. First, in 30–40% of the cases, proposers allocate money to all players, not just to the members of the minimal winning coalition. Further, proposers do not seem to select the "cheapest" coalition partner (i.e., the one with the lowest continuation value). Rather, players frequently include all players in the coalition. Even if proposers select minimal winning coalitions, they appear to select their coalition

---

[4] In these so-called "market games" two or more responders submit the smallest offer they would accept before being matched with a given proposer (Fischbacher, Fong, and Fehr 2003, Grosskopf 2003, Güth, Marchand, and Rulliere 1997).

[5] Eavey and Miller (1984) performed related experiments on the explanatory power of the core in bargaining games. Their main finding was that the core performed quite poorly in accounting for actual bargaining outcomes.

partners randomly. Second, proposers do not seem to exploit their proposal power as predicted by the Baron–Ferejohn model. They consistently offer too much to the respective coalition members. Third, a significant percentage of first-period proposals above the continuation value are rejected. Diermeier and Morton's data, however, do reveal some consistent behavioral patterns. Proposers appear to first select a subset of players with which to form a coalition and then split the money equally among its members.[6]

In a recent paper Frechette, Kagel, and Lehrer (2003) compare open- versus closed-rule versions of the Baron–Ferejohn model and find some qualitative support for the model. In particular, as predicted by the model, delays are longer and distributions are more egalitarian under the open rule. However, some less obvious (but critical) aspects of the Baron–Ferejohn model are not well supported in the data. For example, in their design, proposers should propose minimal winning coalitions in both the open- and closed-rule case. However, only 4% of proposals correspond to this prediction. Even more troubling, under the open rule subjects accept proposals that offer them *less* than their continuation values!

Frechette et al. suggest that proposers and voters may rely on a "fair" reference point of $1/n$ share of the benefits when making decisions. Offers below that share are consistently rejected while shares above $1/n$ are usually accepted. This focal point interpretation may also explain why subjects may accept an amount less than their continuation value, for this happened mostly in cases where the respective continuation values were significantly higher than the fair reference point.

A common problem with all existing experimental investigations of the Baron–Ferejohn model is that determining a specific games continuation values is cognitively quite demanding.[7] This is a key difference between the easy-to-understand bargaining tasks in the tradition of the ultimatum game and previous experiments that investigate the Baron–Ferejohn model, and it may explain why many of the experimental results do not match theoretical predictions.

Our goal in this paper is to design an experiment that captures the key features of the proposer-pivot model, yet is as easy to understand as the ultimatum game. This will allow us to separate cognitive from motivational issues, such as fairness concerns. In our design, each player is directly assigned an ex ante known *disagreement value*, i.e. a given amount of money the player will receive if the proposal is rejected. So, there is no need for the subjects to calculate continuation values.[8] These disagreement values are, in essence, a "reduced form" representation of outside options induced by future

---

[6] A particularly striking finding are so-called "pittance coalitions" (about 25% of all coalitions). These are allocations where two players each receive $22 (out of a total payoff of $45) while the third player receives $1. Since Diermeier and Morton restrict players to full-dollar offers, players cannot split the $45 exactly into equal halves, but have to either propose either $22 or $23 to the coalition partner. It thus appears that proposers prefer to "waste" one dollar on the strategically irrelevant third player rather than allocating unequal payoffs among the proto-coalition members.

[7] Frechette, Kagel, and Lehrer (2003) do try to account for this in a second experiment designed to facilitate learning. Even in this case play does not converge to solution predicted in the Baron–Ferejohn model but to the $1/n$ reference allocation.

[8] Giving players hints or information about how to calculate continuation values appears to have little effect in game experiments (Camerer 2003).

bargaining interactions. By varying the disagreement values, particularly the *proposer's* disagreement value, as our treatment variables we can then test competing theories of sequential bargaining behavior.[9] The common feature in all tested theories is that predicted behavior (by proposers and voters) either *does not* depend on the proposer's disagreement value at all or can depend on it only in specific, testable ways.

We begin by considering theories of self-interested, egalitarian, and inequality-averse play as explanations for behavior in our data. Self-interested play is the benchmark case in which a utility function is based only on the player's own payoff. A voter will accept any offer that is at least as high as her reservation value. In a subgame-perfect equilibrium proposers select the $(N - 1)/2$ cheapest of the non-proposers and pay those players their disagreement values. Under egalitarian play, the proposer will allocate money equally among all group members (group egalitarianism) or among the coalition members (coalition egalitarianism), while voters will reject non-egalitarian allocations *even if their own payoff is substantial*. Models of inequality-averse play (pioneered by Bolton and Ockenfels 2000a, 2000b, Fehr and Schmidt 1999) retain subgame-perfect equilibrium as the solution concept, but change the utility function to include fairness or equity concerns. All else constant, players still prefer more money to less, but also prefer to have neither too little nor too much relative to others – this is the sense in which they are inequality-averse.

Each of these three benchmark cases has been supported by some previous experimental evidence. As emphasized by Fischbacher, Fong, and Fehr (2003), play can look markedly more self-interested and strategic with two or more responders than with only one. But as demonstrated by Frechette, Kagel, and Lehrer (2003), some results may be best explained by models of fair or at least inequality-averse players. Finally, evidence of egalitarian behavior has been observed or suggested in Diermeier and Morton (2005) as well as Frechette, Kagel, and Lehrer (2003). An additional contender to explain behavior in surplus division problems is efficiency (Charness and Rabin 2002): e.g., players may prefer one split over another because it gives higher aggregate payoffs. This approach has not been applied to majoritarian bargaining games, but our design does allow us to examine efficiency considerations as explanations of subject behavior.

PROCEDURES, DESIGN, AND HYPOTHESES

The experiments involved groups of three subjects that had to split 1,250 points among themselves by majority rule. One subject in each group, designated the proposer, proposed a split of the 1,250 points for all three group members. Then all three group members saw the entire proposal and voted on it. If a majority voted in favor, it passed and all subjects were paid according to the proposal. If a majority voted against, it failed and all subjects were paid according to pre-specified and known disagreement values.

The major innovation in our design is to use the proposer's disagreement value as our treatment variable, which allows for a clean assessment of several versions of the

---

[9]   There is some (largely unrelated) experimental research on bargaining in which the disagreement value is taken as a treatment variable (e.g. Binmore, Shaked, and Sutton 1989 and Schmitt 2004).

behavioral hypotheses above. With self-interested players using backward induction, equilibrium behavior will be driven entirely by the disagreement value of the pivotal voter (as long as the proposer's expected payoff is not strictly lower than his reservation value). In particular, results should *stay constant if the proposer's disagreement value is varied*. Also, manipulating disagreement values allows us the change the inequality players face if the offer is rejected. In conjunction with the observed offer this generates clear implications for voter behavior under a hypothesis of fair or inequality-averse players.

The three treatments, defined in terms of the disagreement value vector (proposer, voter 1,voter 2), were Treatment A = (60, 60, 40), Treatment B = (1,125, 60, 40), and Treatment C = (65, 60, 30).[10] To remind the reader we often refer to treatments in the text below by the proposer's disagreement value, DVp. Eight sessions of the majority rule experiment were run in a PC classroom in the Northwestern University Library. In four sessions the treatment order was ABBC, and in four it was BABC.[11] Participants were Northwestern University undergraduates. The sessions had 12 participants each, except session 1, which had 15 participants, so that a total of 99 subjects participated in 1,320 group-rounds[12] across all eight sessions. The experiment was computerized using the z–Tree software developed at the University of Zurich (Fischbacher 1999).

There were 40 total rounds in each session, and the disagreement values and roles in the groups (proposer or voter) changed every 10 rounds. Group membership, on the other hand, was randomly re-drawn after every round. Each subject participated in every treatment and occupied each role for at least 10 rounds of the session (but did not occupy every role in every treatment). Payments in the experiment were the sum total of points from all 40 rounds, multiplied by 0.001. Thus, in each round the pie to split was $1.25. Sessions lasted about 90 minutes, and subjects earned about $19 on average, plus a $5 participation fee. Subjects kept track of their results on personal history sheets. All of these procedures were clearly explained to subjects before the experiment began.

This design allows us to test the following hypotheses.

Hypothesis 1 (Self-interest)
 (a) Every proposer will be more likely to select a coalition that includes himself/herself and the one other team member with lowest disagreement value.
 (b) The payoff to the cheapest coalition partner is likely to be close to that player's disagreement value. The third player will receive a payoff of zero.

---

[10] We also introduced slight variation (with negligible impact on the subgame-perfect equilibrium) in responder disagreement values to prevent voters from relying on the simple heuristic that since voter disagreement values were the same, voter behavior should be the same. Disagreement values for the voters were permuted in different sessions.

[11] Generally, neither proposer nor voter behavior within a given treatment was affected by the sequencing. The one exception is that in the first run (but not the second) of the $DVp = 60$ treatment in sessions 3 and 5, proposers exhibited less self-interested behavior compared to other runs of the same treatment. The same phenomenon did not occur in sessions 1 or 7, when $DVp = 60$ was also run first, and did not appear in any other run of $DVp = 60$. We also found no important changes in proposer or voter behavior within a treatment as subject gained experience. In each treatment, late-round behavior looked like early-round behavior.

[12] Here 1,320 = (7 sessions with 12 subjects)(40 rounds/session)(4 groups/round) + (1 session with 15 subjects)(40 rounds)(5 groups/round).

(c) Proposals above an agent's reservation value are always accepted; below the reservation value they are always rejected.
(d) Variations in the proposer's reservation value should have no consequence for which allocations are proposed; i.e. there should be no significant variation across different treatments.

If this hypothesis is confirmed, we have strong evidence that in a cognitively simple environment, the self-interested, strategic behavior known from other multilateral experiments such as "market" games (e.g., Fischbacher, Fong, and Fehr 2003) is also found in the proposer-pivot model. In other words, the limits placed by institutional context on a voter's ability to sacrifice money for fairness (especially in our $DVp = 1,125$ treatment) would in this case create behavior that appears, to an outside observer, self-interested and strategic in accordance with conventional assumptions in rational choice modeling. The "anomalous" behavior in McKelvey (1991), Frechette et al. (2003), and Diermeier and Morton (2005) would then appear as a likely consequence of cognitively demanding game forms.

If Hypothesis 1 is disconfirmed, other-regarding behavior can take various shapes. One such example is egalitarianism.

Hypothesis 2 (Egalitarianism)
(a) Proposer will allocate money equally among all three group members (strong version – "group egalitarianism") or all members of the coalition (weak version – "coalition egalitarianism").
(b) Responders will reject non-egalitarian allocations even if their own proposed payoff is substantial.
(c) Variations in the proposer's reservation value should have no consequence for which allocations are proposed. In particular, proposers should allocate *less* to themselves than their reservation values in treatment B (1,125, 60, 40).

Bolton and Ockenfels (2000a), as well as Fehr and Schmidt (1999), recently proposed influential theories of other-regarding motivation. In these theories, players prefer more money to less, all else constant (to wit, group members' *relative* shares). But for a given monetary payoff, they prefer to have neither too little nor too much relative to others (they are "inequality-averse"). Unlike self-interested players, a inequality-averse player's utility depends on the entire payoff vector, not just her own payoff component. Different players may have different levels of inequality aversion, but (as in the self-interest theory) a given player uses the same function to evaluate any vector of payoffs. This leads to the following hypothesis.

Hypothesis 3 (Inequality aversion)
(a) A vector of payoffs should be preferred by a voter when her own monetary payoff is higher, all else constant, but should appear worse when inequality in that vector increases, all else constant.
(b) Voters' behavior depends only on the payoff vector, not whether the vector is generated by acceptance or rejection of the proposal.
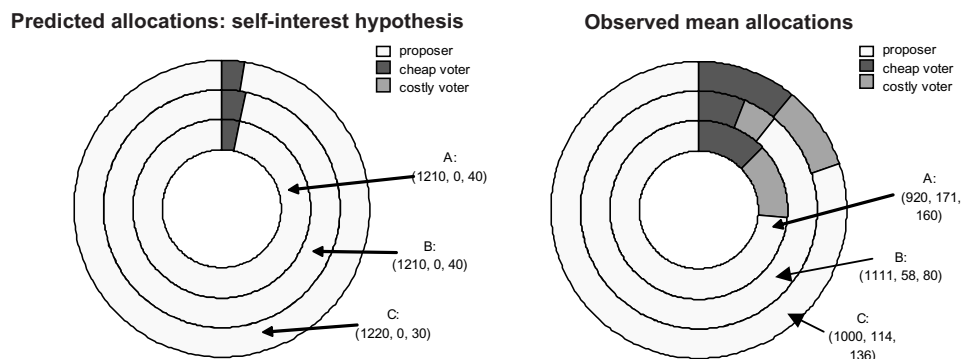
Predicted allocations: self-interest hypothesis

Observed mean allocations

Figure 1. Self-interest predictions and observed mean allocations, by treatment.[13]

RESULTS

According to our experimental results, *none* of the three hypotheses is confirmed by the data. One of the most striking findings is that there is a large effect of changes in the proposer's rejection value. That is, proposers allocate a significantly higher share to themselves in the case of $DVp = 1,125$ *and* responders are more likely to approve such proposals.

Cognitive Heuristics and Self-Interested Players

There are several ways to test the explanatory power of the self-interest theory in our design. First consider proposer behavior. Figure 1 shows the predicted offers by treatment next to the mean observed offer from that treatment.[14] According to the theory, proposer behavior should be very similar across treatments, differing only by 10 fewer points offered to the cheap voter in the treatment B ($DVp = 65$).

Based on the mean observed offers, proposer behavior does exhibit notable self-interest in all treatments. Indeed, in the $DVp = 1,125$ treatment, proposers offer themselves about 90% of the pie on average. Interestingly, this share is even higher than proposer's

---

[13] Disagreement values for treatments A, B, and C are (60, 60, 40), (1,125, 60, 40), and (65, 60, 30), respectively. In each triple the first number corresponds to the proposer's DV or allocation, the second number to the costly voter's, and the third number to the cheap voter's. Observed mean allocations in each treatment do not necessarily sum to 1,250 because the components are the arithmetic means of all observed offers in the treatment, computed without regard to the constraint of 1,250 total points.

[14] Since we did not find any noticeable effects of either treatment order or experience of subjects (with the one exception noted in a previous footnote) we aggregate all data from a given treatment, regardless of when it occurred in a session and regardless of how many rounds the subjects had played in it.

keep in the Dictator game, where a typical finding is that proposers allocate roughly 80% to themselves (cf. Camerer 2003).

Nevertheless, the figure reflects that actual allocations deviate from the predicted allocations in critical ways. While 40% of all proposals allocate positive amounts only to the proposer and the cheap voter (i.e., constitute "correct" minimal winning coalitions), 20% allocate positive amounts only to the proposer and the costly voter (i.e., constitute "incorrect" minimal winning coalitions), and 40% allocate positive amounts to all three players. In addition, the costly voter receives a nontrivial proposal on average.

Not only do the observed allocations depart noticeably from the theoretical ones for each treatment, but the observed allocations in the strategically identical $DVp = 1,125$ and $DVp = 60$ treatments are markedly different from each other. Subjects *modify their behavior* in response to changes in the proposer's disagreement value, changes that should be strategically irrelevant if the original self-interest model was correct. These findings are confirmed by Kolmogorov–Smirnov and Wilcoxon tests for difference in distributions. Tests comparing the $DVp = 60$ and 65 treatments to the $DVp = 1,125$ treatment reject the null hypothesis of no difference in the distribution of offers to each player at the $p = 0.003$ level or smaller.[15]

One may suspect that because $DVp = 1,125$ is close to the maximum payoff of 1,250, proposers may simply choose to propose 1,250 even if they expect to be rejected. This conjecture, however, is not supported by the data. The following figure shows proposers' offers to themselves, by treatment.

About 8% of proposers in $DVp = 1,125$ offer themselves the entire pie (2% do in the other treatments). About 27% more offer themselves 1,200 or more (but less than 1,250), and 48% more offer themselves 1,125 or more (1,125 is exactly 90% of the pie, and 1,200 is 96%). In each treatment the modal offer to the proposer is between 85% and 95% of the pie.

Moreover, in the infrequent cases when proposers do offer themselves the whole pie (54 occurrences overall), they do so almost exclusively when $DVp = 1,125$ (53 times out of 54), and their proposals are almost always rejected by the group (52 times out of 54). In the round following a rejection of a 1,250 offer when $DVp = 1,125$, proposers lower allocations to themselves about 34% of the time. On average the reduced allocation in the next round is 1,162 points. Finally, three subjects who nearly always offered themselves 1,250 under $DVp = 1,125$ account for 27 of these 53 whole-pie proposals. With these subjects removed (leaving 15 other subjects), the remaining proposers lower their own allocations about 70% of the time after 1,250 is rejected.

A problem with focusing solely on proposer data is that (under the null hypothesis that the theory is correct) proposers would best respond to expected voter decision rules.

---

15. While nonparametric, the Kolmogorov–Smirnov and Wilcoxon tests do require independence of sample observations. Strictly speaking our observations are not independent, a common problem when aggregating repeated observations in experiments. However, computed test statistics are so large this seems unlikely to matter. Moreover, the same results are confirmed below when (unconditional) correlation in an individual subject's observations is considered. Given the lack of notable trends in behavior over the course of the sessions as noted above, failure of independence because of experience, etc. seems to be unimportant.
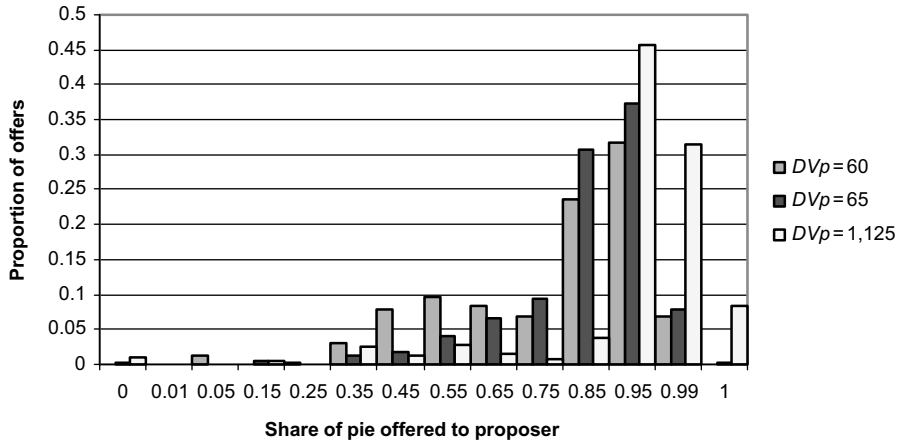
Figure 2. Proposer offers to self, by treatment.

Hence, departures of proposer behavior from theoretical predictions can only be viewed as casting doubt on the *joint* hypothesis of self-interested agents *and* subgame perfection. If voters depart from own-income maximization, a self-interested, backward-inducting proposer might do so as well, but for strictly self-interested reasons. It is therefore critical to investigate voter behavior directly. Here the predictions of the model refer to conditional probabilities; for example, acceptance rates conditional on receiving an offer above the reservation value.
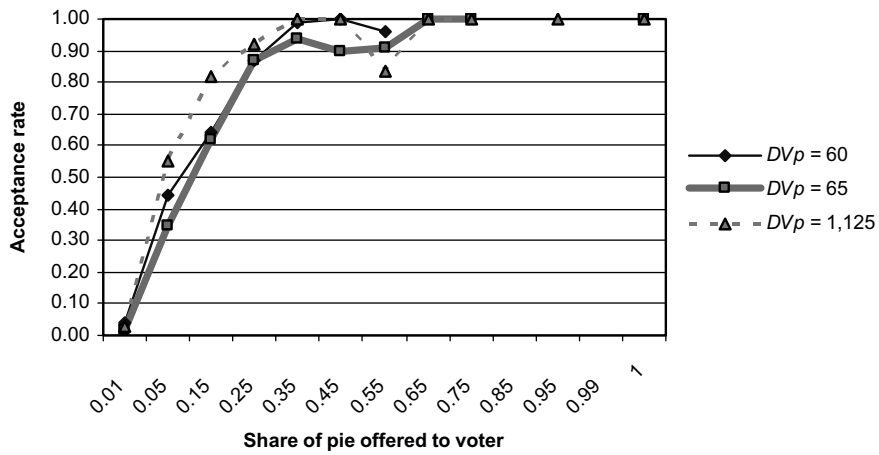

**Figure 3**. Conditional acceptance rates, by treatment.

Figure 4. Conditional acceptance rate for low–moderate offers, by treatment.

The change to $DVp = 1,125$ has a significant effect on voters as well, conditional on the type of offer received. Figure 3 provides evidence of this. Voters generally have a higher conditional probability of acceptance in the $DVp = 1,125$ treatment,[16] and in this sense are more accepting of unequal offers.[17] This is especially evident for offers in the range of 1% to 25% of the total pie, a range that contains 56% of all proposals. In Figure 4 we focus on this range. The figure pools the cases with a low disagreement value for the proposer ($DVp = 60/65$) for comparison to the high DVp case ($DVp = 1,125$). Consider, for example, the 5–15% range (62.5–187.5 points). This is a large segment of the data; it contains 29% of all offers in $DVp = 60$, 34% in $DVp = 65$, and 23% in $DVp = 1,125$. In this range acceptance rates are over 30% higher in $DVp = 1,125$ than in the other two treatments (a difference of 19 percentage points; p-value in tests for differences[18] across treatments is less than 0.001). Thus, proposers not only demand more when their reservation value is higher, voters are also willing to give it to them.

A reasonable modification of the proposer-pivot mode is to allow for some error or randomness in voting. We can capture such a variation as a random utility model. In this case the self-interest model yields an explicit functional form that makes it suitable for

---

[16] The sharp drop in the 45–55% range for $DVp = 1,125$ is probably due to small sample size. There are only 12 offers in this bin (less than 1% of all offers made), and only 24 offers above 35% across all sessions in this treatment. Two of these were rejected and they happened to fall in the 45–55% range; the other 10 offers in that range were accepted.

[17] Interestingly, there is a higher proportion of "no" votes, and rejections by the group in the $DVp = 1,125$ treatment than in the other treatments. The reason is that, although voters are conditionally more accepting, proposers are also demanding more.

[18] This test assumes that observations are independent; correlation among them would increase standard errors. Since the result is highly significant, and we uncover similar findings when subject-specific effects are taken into account, this caveat seems unlikely to matter substantively.

Table 1. Means and fixed effect logit results, self-interest model

Variable name	Mean value	Parameter estimate[19]	Standard error	*p*-Value
Vote "For" (dependent variable)	0.480	–	0.001	–
Own offer	106.54	0.031	0.002	0.000
Own DV	48.78	−0.041	0.007	0.000
$DV_p = 65$	0.251	−0.321	0.207	0.121
$DV_p = 1{,}125$	0.498	1.018	0.178	0.000
No. observations: 2,631[20]	No. subjects: 99	observations per subject: (min, mean, max): 10, 26.6, 30		

Log likelihood: −720.07.

a more detailed econometric analysis of vote choices. If the voting model is correct but voter utility is observed with error, the probability that voter i (conditioning on being pivotal) votes in favor of an offer giving him x_i is just $\Pr[\gamma_i + x_i + \varepsilon_{offer} > DV_i + \varepsilon_{dv}]$, or $F_{\Delta\varepsilon}(\gamma_i + x_i - DV_i)$, where γ_i is a voter-specific fixed effect to account for any fixed (but unobserved) disposition to vote in favor of the proposed offer. To test this hypothesis we estimate a fixed-effects logit model of the probability of observing i's actual voting pattern, conditional on the total number of "yes" votes by i. The variable *Own Offer* captures the points offered to the voter (out of 1,250), while the variable *Own DV* captures the voter's own disagreement value.

To account for the effect of the proposer's disagreement value, we include dummy treatment variables (with $DVp = 60$ as the baseline case). As discussed above, they should have no effect if the self-interest theory is correct. The results are found in Table 1. The parameter estimates for *own offer* and *own DV* do have the predicted signs and are highly significant. However, again we see the strongly significant effect of the $DVp = 1{,}125$ treatment compared to the baseline $DVp = 60$. Relative to $DVp = 60$, $DVp = 65$ reduces the probability of favorable votes, but only at a *p*-value of 0.12. Denoting the parameter estimate for variable k by β_k and estimating the model with the theoretical restrictions $\beta_{own\ offer} = -\beta_{own\ DV}$ and $\beta_{65} = \beta_{1,125} = 0$ produces a log likelihood of −759.68. The χ^2 (with three degrees of freedom) statistic from the likelihood ratio

---

[19] The marginal effects are difficult to obtain, since they depend on the value of the fixed effects, which are not directly estimated in conditional logit because the estimators would be inconsistent (Wooldridge 2002). Assuming a specific value of say $\gamma_i = 0$ would allow marginal effects to be pinned down, but it is not possible to assert that this is the marginal effect given the mean fixed effect, since the distribution of them is unknown.

[20] A total of 2,640 votes actually took place. Nine observations were dropped because the voter made the same decision in each round.

test with a null hypothesis of non-constraining restrictions[21] is 79.22 – a value that has a probability of essentially zero if the null is true.[22]

To summarize, while proposers and voters respond to monetary incentives and there is much evidence of self-interested behavior, the pure self-interest theory in the proposer-pivot model cannot account for the observed behavior. The self-interest hypothesis not only implies certain behavior, but it also predicts which behavior should *not* occur: i.e. which aspects of the choice environment should be ignored by decision-makers. In our experiments, behavior is strongly affected by strategically irrelevant features of the decision context, in particular the proposer's reservation values. In the case of a high proposer reservation value voters are more accepting of unequal offers, and proposers offer more unequal distributions. Note that these findings deviate not only from the self-interest theory, but from *any* theory in which voting behavior depends only on the voter's own disagreement value and proposed share. This includes models with noisy best response behavior such as Quantal Response Equilibrium (McKelvey and Palfrey 1995, 1998).[23]

Egalitarianism and Efficiency

Both the group and coalition versions of the egalitarian model perform poorly enough in explaining our data that a more detailed econometric analysis seems unnecessary. In particular, the dependence of proposer and voter behavior on the proposers' disagreement value is as damaging to the egalitarian hypothesis as the self-interest hypothesis. In addition, while coalition formation does not follow the purely self-interest model, over 80% of all proposals offer some voter 100 points (8%) of the pie or less, and over 60% offer one voter nothing, thus resulting in a minimal winning coalition. Within minimal winning coalitions, proposers offer themselves almost eight times as much as the coalition member. Voters receiving nonzero offers when the other voter receives a zero offer support the proposal over 71% of the time. In short, neither version of the egalitarian model has much explanatory power in our data. Behavior is too opportunistic and sensitive to context.

The change in voter behavior also provides some information about the explanatory power of efficiency considerations. Charness and Rabin(2002) have proposed that aggregate payoffs may have an influence on proposal and acceptance behavior in surplus division problems. Engelmann and Strobel (2004) found efficiency considerations to

---

[21] The test statistic is twice the difference in the log likelihoods for the model without the constraints and the model with the theoretical constraints imposed. If the voting model is correct, imposing the constraints should not matter much, and the log likelihoods should be similar, producing a small test statistic.

[22] However, using only $\beta_{own\ offer} = -\beta_{own\ DV}$ and leaving the parameters of the treatment dummies unrestricted produces a χ^2 (1 d.f.) of only 1.8 with a *p*-value of 0.180.

[23] Of course, if the cost of errors changes across treatments, it could change behavior across treatments. It may be that the voter's errors, for example, are less costly for *proposers* in $DVp = 1,125$ than $DVp = 60$ or 65, pointing to a potential extensive form quantal response effect. But this cannot help to explain *voter* behavior across treatments, since at the voting stage there is no relevant asymmetric information and no relevant beliefs about the future behavior of others in the group.

be important in explaining surplus division. In our treatments, the efficiency loss from rejection under $DVp = 60$ or 65 is far greater than it is under $DVp = 1,125$. Therefore, if efficiency or aggregate payoff concerns had an independent effect on behavior, we would expect voters to be *more* accepting, all else constant, in $DVp = 60/65$ than in $DVp = 1,125$. Instead they are significantly less accepting in these treatments, conditional on the offer received.

In sum, neither the SPNE with self-interested play, nor an egalitarian model, nor efficiency considerations can account for the substantial impact of the changes in the proposer's reference points.

Inequality Aversion

A promising alternative explanation of behavior in pie–splitting situations is that players are strategic in the sense of subgame perfection, but are inequality averse. This issue is potentially important in our experiment because rejection of the proposal in the (1,125, 60, 40) treatment creates much more inequality in payoffs than rejection in the (60, 60, 40) treatment. This would suggest that a voter might be more likely to accept an offer close to her disagreement value in the former treatment than the latter. Inequality aversion also has a variety of other implications that we can test in our data.

Two of the most prominent and empirically successful models of inequality aversion are due to Bolton and Ockenfels (2000a, 2000b) and Fehr and Schmidt (1999). These models differ in the specific way inequality enters individual utility functions. Fehr and Schmidt propose that, for a fixed number of players n, the utility to individual i from the monetary payoff vector $x = (x_1, x_2, \ldots x_n,)$ is

$$u_i(x) = x_i - \alpha_i \Sigma_{j \neq i} \max(x_j - x_i, 0) - \beta_i \Sigma_{j \neq i} \max(x_i - x_j, 0)$$

with $0 < \beta_i < 1$ and $\alpha_i > \beta_i$.[24] In other words, utility is the sum of (a) i's own monetary payoff, (b) bilateral "disadvantageous" inequality (in which j receives more than i), and (c) bilateral "advantageous" inequality (in which i receives more than j). Each type of inequality is negatively weighted to reflect that it reduces utility. The difference in weights captures the intuition that players dislike disadvantageous inequality more than advantageous inequality.

Fehr and Schmidt's model, like the self-interest model, readily produces a functional form suitable for econometric testing assuming a random utility model. We use conditional fixed effects logit to estimate a model of the probability of observing i's vote pattern, conditional on the explanatory variables given by the theory, the total number of "yes" votes by voter i, and a fixed effect γ_i for voter i to capture unobserved heterogeneity in voter dispositions to vote in favor.[25] This model, unlike the self-interest one, does not

---

[24] In the original Fehr–Schmidt theory, the α and β terms are scaled by $n - 1$. Since these scaling terms are the same for both terms and all voters in our design we omit them for simplicity. We thank an anonymous referee for pointing this out.

[25] Note that this formulation constraints the parameters α and β to be the same across all players. This is consistent with the "calibration" approach to explaining previous experimental findings,

Table 2. Means and fixed effect logit results, Fehr–Schmidt inequality aversion

Variable name	Mean value	Parameter estimate	Standard error	*p*–Value
Vote "For" (dependent variable)	0.480	–	0.001	–
Own offer	106.54	0.040	0.003	0.000
Own DV	48.78	−0.005	0.021	0.803
Disadvantageous inequality in the offer	983.52	0.001	0.001	0.132
Disadvantageous inequality in the DV pie	554.37	0.001	0.0002	0.000
Advantageous inequality in the offer	61.02	−0.009	0.002	0.000
Advantageous inequality in the DV pie	11.29	−0.036	0.021	0.085

No. observations: 2,631 No. subjects: 99 Observations. per subject (min, mean, max): 10, 26.6, 30

Log likelihood: −710.17.

include treatment dummies because they are perfectly collinear with the disagreement pie inequality measures. The results are captured in Table 2. Several important features stand out. First, voters are not more sensitive to disadvantageous inequality than advantageous inequality. Second, two of the parameter estimates have incorrect signs (though their *p*-values indicate at best borderline significance) – the sign on *Disadvantageous Inequality in the Offer* should be negative, while the sign on *Advantageous Inequality in the DV pie* should be positive. This suggests that subjects do not treat inequalities in the disagreement values in the same way as inequalities in the proposal. We can directly test for this conjecture by estimating a restricted model in which voters only respond to the difference in inequality between offer and disagreement value. That is, in the restricted model (*Own Offer–Own DV*), (*Disadvantageous Inequality in DV pie–Disadvantageous Inequality in the Offer*), and (*Advantageous Inequality in DV pie–Advantageous Inequality in the Offer*) are the explanatory variables. The corresponding likelihood ratio test yields a χ^2 (three degrees of freedom) statistic of 20.46, with a *p*-value of 0.0001. Hence, in contrast to the predictions of the Fehr–Schmidt model, voters distinguish between

used in Fehr and Schmidt (1999). We use this approach because with so few observations on each individual, we cannot rely on asymptotic justifications of estimators for subject-specific structural parameters.

inequalities in disagreement values and inequalities in proposals: they seem to be more tolerant of inequality in disagreement values than in proposed allocations.

One explanation for this finding based on other experiments is that responders care about how inequality is generated, in addition to its magnitude (Blount 1995, Hibbing and Alford 2004). For example, Blount (1995) finds that unequal offers generated by computer are more acceptable than equally unequal offers generated by other subjects. In addition, Hibbing and Alford (2004) show that, when allocations are generated by chance, responders are more accepting than when those same allocations are generated by human choice. In our experiments, inequality in the *offer* is the decision of a participant in the experiment, but the inequality in the *disagreement values* is assigned exogenously, which explains why voters are more tolerant of it. These findings imply that the expected payoff for proposers depends importantly on the decision context. It follows that, by manipulating a decision *process* to create an inherent strategic advantage, proposers can improve their overall payoffs.[26]

The Equity–Reciprocity–Competition (ERC) model of Bolton and Ockenfels (2000a, 2000b) also includes inequality aversion as part of agents' utility functions, but with a different functional form. Agents compare their own share of the total pie with an even three-way split of the pie: the further their share is from the even-split "social reference point" (whether above or below), the worse off they are – all else constant. Formally, in our setting the utility of a monetary payoff vector x is[27]

$$u_I(x) = v_i(x_i, \sigma_i) \equiv v_i(x_i, x_i / \Sigma_j x_j),$$

with the key assumptions being that (i) fixing σ_i, v_i is increasing and strictly concave in x_i, and (ii) fixing x_i, v_i is strictly concave with a maximum at $\sigma_i = 1/n$.

The ERC model does not specify an explicit functional form. Here, we assume that $u(x) = (x_i)^{.5} + (\sigma_i - 1/3)^2$ for some outcome vector x. Results with other functional forms (other power coefficients, logs, etc.) yield similar results. Results of the (unconstrained) conditional fixed effects logit estimation are listed in Table 3. Again, we find some significant deviations from the model's predictions. Four of six coefficients have the wrong sign (but none is significantly different from zero by conventional benchmarks): the one for (*Own DV*)$^{0.5}$ should be negative, that for *Own DV/Sum of DV's* should be positive, that for ((*Own offer*)/(*Sum of offers*))^2 should be positive, and that for ((*Own DV*)/(*Sum of DV's*))$^{0.5}$ should be negative. As in the case of the Fehr–Schmidt model, we wish to know whether voters evaluate the proposed offer the same way they evaluate the disagreement value pie. A likelihood ratio test (χ^2 test statistic, 3 d.f., is 20.84; *p*-value > 0.0001) suggests that the hypothesis that the theoretical restrictions are not constraining does not hold.

---

[26] Other procedural aspects of experimental design have been shown to benefit proposers, by making responders more accepting of low offers (Güth and Tietz 1986; Hoffman, McCabe, and Smith 1996; Hoffman and Spitzer 1985; List and Cherry 1999; Schmitt 2004). For example, if proposers "earn" the right to be proposers by doing well in a trivia contest, rather than being randomly chosen, responders are more accepting of low offers.

[27] This simplified version assumes that $\Sigma_i x_i > 0$, which is satisfied in our context.

Table 3. Means and fixed effect logit results, ERC inequality aversion (parametric)

Variable name	Mean value	Parameter estimate	Standard error	*p*-Value
Vote "For" (dependent variable)	0.480	–	0.001	–
(Own offer)$^{0.5}$	7.687	0.615	0.064	0.000
(Own DV)$^{0.5}$	6.931	0.052	0.160	0.746
(Own offer)/ (Sum of offers)	0.090	−6.197	6.860	0.366
(Own DV)/ (Sum of DVs)	0.172	−5.762	4.009	0.151
((Own offer)/ (Sum of offers))^2	0.025	−5.436	10.737	0.613
((Own DV)/ (Sum of DVs))^2	0.050	2.102	10.520	0.842
No. observations: 2,631	No. subjects: 99	Observations per subject (min, mean, max): 10, 26.6, 30		

Log likelihood: −578.88.

Thus, key components of the Fehr–Schmidt and Bolton–Ockenfels (ERC) models are inconsistent with our data. Voters do not value the proposed offer or the disagreement pie in the ways predicted by the theories. More fundamentally, they do not evaluate the characteristics of the disagreement pie in the same way they evaluate the same characteristic in the proposed offer: they are more concerned with inequality in the offers than in the disagreement pie. This issue generalizes beyond existing inequality-aversion models. *Any theory* in which utility is solely a function of the vector of payoffs along the path of play will require voters to evaluate proposals the same way they evaluate the disagreement value pie, a prediction not supported by our data. In other words, our data suggest that voters use a different function to evaluate endogenously vs. exogenously generated inequality. Conventional specifications of rational choice (including these models of inequality aversion) require that a voter apply the same function to any possible vector of payoffs.

ACCOUNTING FOR THE FINDINGS: THE IMPORTANCE OF ENTITLEMENTS

We can summarize the key findings of our experiment as follows:

1. Agents do not behave according to the proposer-pivot model with purely self-interested agents. Rather, their behavior is significantly influenced by strategically

      irrelevant variables that according to the theory should not matter – most importantly the proposer's reservation value.

2. However, agents' behavior does respond to monetary incentives. The probability a voter supports a proposal increases in the amount offered to her. Both proposed and accepted allocations are frequently highly unequal, suggesting some recognition of proposal power. In over 40% of all proposals, proposers offer zero to the voter with the higher disagreement value, and proposers offer larger shares to the cheaper voter about 60% of the time.

3. Inequality aversion also plays a role in behavior, but alone cannot account for voters' evaluations of offers. In particular, voters appear to be less concerned with inequality in the disagreement value compared to inequality in the proposed allocation.

4. When the proposer's disagreement value is high (1,125 out of 1,250 points), proposers are demand more and voters are more accepting of low offers.

Taken together these findings suggest a theory that in addition to self-interest and inequality aversion incorporates a basic form of *entitlement*. Subjects seem to consider an agent's reservation value as determining an entitlement and are more willing to accept less generous offers if a proposer's reservation value is high.

    This approach would explain the dramatic percentage increase in voters' acceptance rates (more than 30%) if the proposer's reservation value changes from low to high. According to the same approach proposers make more unequal offers that are nevertheless accepted. This change in proposer behavior may reflect a self-interested calculation as proposers recognize an opportunity to exploit this willingness on the side of the voters, or it may be due to the fact that proposers feel entitled to a larger share when their disagreement values are high.

    Considerations of entitlement suggest that voters are willing to punish proposers who act "too greedily." This seems to be at least partially the case: voters tolerate a greater share for the proposer up to a point, but beyond it they are less likely to accept offers with larger proposer allocations. Table 4 displays fixed-effect logit results from a model of vote choice. Specifically, we examine the effect of proposer surplus – the excess of the offer to the proposer over the proposer's disagreement value – on the decision to support the proposal, controlling for the voter's own surplus. Voters are less likely to support a proposal as the proposer's surplus increases. Furthermore, the negative effect on the probability of approval is stronger when the proposer's surplus is larger. Similar results hold when we examine only observations in which the voter's offer exceeds that of the other voter, and the voter's own offer exceeds the voter's disagreement value, to control for the voter's own entitlement considerations. These results suggest that voters support offers that are consistent with a proposer's entitlement, but are more likely to reject proposals as the proposer's share increases above that benchmark.

    So far, we have focused on entitlements as a factor in the relationship between proposer and voter. However, the entitlement considerations also factor into the relationship between the two non-proposing voters. Consider a proposal where voter A receives a higher share than voter B, but both voters receive more than their reservation value. If the entitlement hypothesis is correct, voter B should be more likely to accept such an

Table 4. Means and fixed effect logit results, vote choice and proposer surplus

Variable name	Mean value	Parameter estimate	Standard error	*p*-Value
Vote "For" (dependent variable)	0.480	–	0.001	–
Offer − DV	57.76	0.031	0.001	<0.000
Proposer surplus	437.00	−0.000541	0.00026	0.042
(Proposer surplus)^2	454,248.90	−5.12× e^{-7}	2.58× e^{-7}	0.047
No. observations: 2,631	No. subjects: 99	Observations per subject (min, mean, max): 10, 26.6, 30		

Log likelihood: −727.61.

offer if B's reservation value is lower than A's, compared to the case where B's reservation value is higher. That is, according to B, A would be "entitled" to a higher share because of his higher reservation value. On the other hand, if B's reservation value was lower, *the same offer* would presumably violate B's sense of entitlement.

Indeed, voters do respond to violations of such a perceived entitlement. When offered less than the other voter, voters are significantly (*p*-value < 0.001) more likely to accept a proposal if the other voter's reservation value is higher – supporting the view that reservation values serve as entitlements. An entitlements-based explanation would also suggest that, conditional on receiving a greater proposed allocation than the other voter, a voter is more likely to accept if her disagreement value is higher as well: an advantageous (and unequal) offer may be easier to accept if it corresponds to entitlement. This difference in conditional acceptance rates, while present when a voter's proposal does not exceed her disagreement value by a very wide margin (which we define as 250 points), is not as strong as when the voter's offer is less than the other voter's. In other words, if offers are not too high, voter responses to proposals that advantage them relative to the other voter are consistent with an entitlement-based approach. These concerns, however, are not as important as when proposals disadvantage them relative to the other voter. In this sense, voters appear to protect their own perceived entitlement first.. But they do less to protect the entitlement of other agents (i.e., the other voter) from the proposer's violations.

Table 5 shows that these results are robust when we control for the voter's Offer–disagreement value margin. In particular, as indicated by the *p*-value of 0.023 on the final parameter estimate, voters are significantly less likely to vote in favor of a proposal if their disagreement value is greater than the other voter's but their offer is lower.

It is important to note that both entitlements and other-regarding considerations can be swamped by sufficiently high individual payoffs (cf. Bolton and Ockenfels 2000a): voters

Table 5. Means and fixed effects logit results. voter–voter comparison, $DVp = 60$ or 65

Variable name	Mean value	Parameter estimate	Standard error	*p*-Value
Vote "For" (dependent variable)	0.527	–	–	–
Offer − DV	96.98	0.023	0.002	0.000
=1 if Other Offer > Own offer	0.414	−0.734	0.311	0.018
=1 if Own DV > Others's DV	0.500	−0.044	0.411	0.915
=1 if Own DV greater *and* Own offer lower	0.207	−1.11	0.491	0.023
No. observations: 1,280	No. subjects: 89	Observations per subject (min, mean, max): 10.14.4; 20		

appear to have a "price" for their sense of entitlement. In $DVp = 60$ or 65, with Own offer – Own DV > 250, almost 97% of all offers (212 out of 219) are accepted.[28] In short, when a voter's offer is high enough (roughly 1/5 of the total pie), other aspects of the proposal have no effect on the acceptance decision.[29]

ENTITLEMENTS IN BILATERAL BARGAINING

In order to investigate whether our conception of entitlements factor into other circumstances we conducted a set of bilateral bargaining experiments (for more detail see Diermeier and Gailmard 2005). Specifically, we ran three sessions of bilateral proposer games, including the standard ultimatum game, with subjects drawn as before from the Northwestern undergraduate student body. These sessions had 12, 10, and 10 participants, respectively, and took place under the same conditions described above: each session had 40 rounds with three treatments per session, each subject participated in each treatment, and subjects interacted anonymously at computer terminals in a PC classroom at Northwestern. The pie to split was again 1,250 points, and we again varied disagreement values as the treatment variable. In particular we used the following pairs of disagreement values (proposer, responder): Treatment A = (0, 0), which corresponds to the ultimatum game; Treatment B = (120, 65); and Treatment C = (1, 150, 60). The

---

[28] The conditional acceptance rate is somewhat higher (about 97.5%) if $DVp = 1,125$ is included, but voters are generally more accepting in that treatment in any case.

[29] This is reminiscent of the self-serving evaluations of equity uncovered in Knez and Camerer (1995) and Fehr and Schmidt (1999).

Figure 5. Conditional acceptance rates, by treatment – bilateral bargaining.

ultimatum game condition was included to make our results comparable to the existing literature and exclude treatment effects. In two sessions the treatment order was ABCA, and in one it was CABC.

If the entitlement approach is correct, we should expect different acceptance depending on the proposer's disagreement value. Responders should be more likely to accept the same offer if the proposer's disagreement value is high. Indeed, this is the case as Figure 5 indicates. First, observed acceptance rates in the ultimatum case are similar to the existing literature. More importantly, as in the case of majoritarian decision-making, acceptance rates increase substantially as we change the proposer's disagreement value from 120 to 1,150. Thus, we find the same entitlement effect as in the proposer-pivot case.

The effect of the change in treatment on proposer behavior is also similar to the majoritarian decision-making case: the strategically irrelevant increase in the proposer's disagreement value shifts proposed allocations strikingly in favor of the proposer. In the $DVp = 0$ and $DVp = 120$ treatments, the mean allocation to the proposer is about 73% of the pie; in the $DVp = 1,150$ case it is about 93% of the pie. In short, the results from the bilateral experiments clearly corroborate the importance of our conception of entitlements.

CONCLUSION

We propose an experimental test of proposer-pivot models using a simple three-player game that instantiates both the Romer–Rosenthal (1978) model and the Baron–Ferejohn (1989a) model. In the game both proposers and voters are assigned a known disagreement value, i.e. a fixed amount of money that they will receive if the proposed allocation is

not passed by majority vote. We focus on three different assumptions about the extent to which subjects exhibit fairness concerns in their behavior: self-interested, egalitarian, and inequality-averse behavior. The key design variable in our experiments is the variation in the proposer's reservation value. In the experiment the proposer's reservation value varies between about 5–10% of the pie (low-value condition) and almost 90% (high-value condition).

We find that none of the proposed explanations is able to account for key features of the data. Proposers are much more opportunistic than an egalitarian model would imply, to the point of taking more than 90% of the pie over 90% of the time in some treatments. At the same time, proposers do not exploit their proposal power enough to confirm the self-interest hypothesis, particularly in the low-value treatments, even if these treatments occur after many rounds of the high-value treatment. Moreover, both proposer and voter behavior is sensitive to changes in the proposer's reservation value, which is fundamentally inconsistent with *any* explanations based solely on the pivotal voter's reservation value. Finally, effects of treatments on conditional acceptance rates are the opposite of effects predicted by models with inequality aversion.

A closer analysis of the data suggests that subjects interpret the reservation values as a basic form of entitlement. That is, everything else equal, they are more accepting of unequal offers if the proposer's reservation value is higher. Similarly, voters are more tolerant of offers that allocate more money to the other non-proposing voter if that agent's reservation value is higher. In turn, in the case of higher proposer reservation values, proposers demand significantly more of the total payoff for themselves, even if this results in somewhat higher overall rejection rates. The effects of both other-orientation and a sense of entitlement are strongest when the offers to the respective voter are low. At a given point, roughly at 1/5 of the total payoff, subjects accept offers with high probability irrespective of distributional or entitlement concerns.

The importance of entitlements is confirmed by experiments on bilateral interactions. Again, responders are more willing to accept the same offer if the proposer has a higher reservation value. Otherwise subjects behave exactly as in known experiments such as the ultimatum game.

Our analysis suggests that the ethical aspects of decision-making are an important ingredient of any descriptive theory of majoritarian decision-making. In addition to the known fairness and inequality concerns, voters also significantly respond to a basic form of entitlement. The findings suggest that real subjects exhibit a complex interaction between self-interest and different forms of ethically motivated behavior, an interaction we are just beginning to understand. While this paper has demonstrated the inadequacy of standard categories of explanation and the importance of entitlements, much more work in theory and observation is needed to integrate self-interest, inequality, and entitlements into a unified and coherent theory.

REFERENCES

Baron, David P., and John A. Ferejohn. 1989a. "Bargaining in Legislatures." *American Political Science Review* 89: 1181–1206.

Baron, David P., and John A. Ferejohn. 1989b. "The Power to Propose." In *Models of Strategic Choice in Politics*, ed. Peter C. Ordeshook. Ann Arbor: University of Michigan Press, 343–66.

Binmore, Kenneth, Avner Shaked, and John Sutton. 1989. "An Outside Option Experiment." *Quarterly Journal of Economics* 104: 753–70.

Blount, Sally. 1995. "When Social Outcomes Aren't Fair: The Effect of Causal Attributions on Preferences." *Organizational Behavior and Human Decision Processes* 63: 131–44.

Bolton, G., and A. Ockenfels. 2000a. "ERC: A Theory of Equity, Reciprocity, and Competition." *American Economic Review* 90 (March): 166–93.

Bolton, G., and A. Ockenfels. 2000b. "Strategy and Equity: An ERC analysis of the Güth–van Damme Game." *Journal of Mathematical Psychology* (42)2: 215–26.

Camerer, Colin. 2003. *Behavioral Game Theory*. Princeton, NJ: Princeton University Press.

Charness, Gary, and Matthew Rabin. 2002. "Understanding Social Preferences with Simple Tests." *Quarterly Journal of Economics* 117: 817–69.

Diermeier, Daniel, and Sean Gailmard. 2005. "Entitlements in Bilateral Bargaining." Mimeo, Center for Business, Government, and Society, Northwestern University.

Diermeier, Daniel, and Rebecca Morton. 2005. "Experiments in Majoritarian Bargaining" 2005. In *Social Choice and Strategic Decisions: Essays in Honor of Jeffrey S*, Banks, eds. D. Austen-Smith and J. Duggan. Berlin: Springer, 201–36.

Eavey, Cheryl, and Gary Miller (1984). "Fairness in Majority Rule Games with a Core." *American Journal of Political Science* 28: 570–86.

Engelmann, Dirk, and Martin Strobel. 2004. "Inequality Aversion, Efficiency, and Maximin Preferences in Simple Distribution Experiments." *American Economic Review* 90: 857–69.

Fehr, E., and K. Schmidt. 1999. "A Theory of Fairness, Competition, and Cooperation." *Quarterly Journal of Economics* 114(3): 769–816.

Fischbacher, Urs. 1999. "z-Tree: Zurich Toolbox for Readymade Economics Experiments." Mimeo, University of Zurich.

Fischbacher, Urs, Christina Fong, and Ernst Fehr. 2003. "Fairness and the Power of Competition." Mimeo. University of Zurich and Carnegie Mellon University.

Forsythe, Robert, Joel L. Horowitz, N. E. Savin, and Martin Sefton. 1994. "Fairness in Simple Bargaining Experiments." *Games and Economic Behavior* 6: 347–69.

Frechette, G., J. H. Kagel, and S. Lehrer. 2003. "Bargaining in Legislatures: An Experimental Investigation of Open versus Closed Amendment Rules." *American Political Science Review* 97: 221–232.

Grosskopf, Brit. 2003. "Reinforcement and Directional Learning in the Ultimatum Game with Responder Competition." *Experimental Economics* 6: 141–158.

Güth, W., and R. Tietz. 1986. "Auctioning Ultimatum Bargaining Positions: How to Act Rational if Decisions are Unacceptable." In *Current Issues in West German Decision Research*, ed. R. W. Scholz. Frankfurt: Verlag Peter Lang, 173–85.

Güth, W., and E. Van Damme. 1998. "Information, Strategic Behavior and Fairness in Ultimatum Bargaining: An Experimental Study." *Journal of Mathematical Psychology* 42(2/3): 227–47.

Güth, W., N. Marchand, and J. L. Rulliere. 1997. "On the Reliability of Reciprocal Fairness – An Experimental Study." Mimeo, Humboldt University of Berlin.

Güth, W., R. Schmittberger, and B. Schwarze. 1982. "An Experimental Analysis of Ultimatum Bargaining." *Journal of Economic Behavior and Organization* 75: 367–88.

Hibbing, John, and John Alford. 2004. "Accepting Authoritative Decisions: Humans as Wary Cooperators." *American Journal of Political Science* 48: 62–76.

Hoffman, Elizabeth, and Matthew Spitzer. 1985. "Entitlement, Rights, and Fairness: An Experimental Investigation of Subjects' Concepts of Distributive Justice." *Journal of Legal Studies* 14: 259–97.

Hoffman, Elizabeth, Kevin McCabe, Jason Shachat, and Vernon Smith. 1994. "Preferences, Property Rights and Anonymity in Bargaining Games." *Games and Economic Behavior* 7: 346–80.

Hoffman, Elizabeth, Kevin McCabe, and Vernon Smith. 1996. "On Expectations and Monetary Stakes in Ultimatum Games." *International Journal of Game Theory* 25: 289–301.

Knez, Marc, and Colin Camerer. 1995. "Outside Options and Social Comparison in Three-Player Ultimatum Experiments." *Games and Economic Behavior* 10: 65–94.

Kagel, John and Katherine Wolfe (2001). "Tests of Fairness Models Based on Equity Considerations in a 3 Person Ultimatum Game." *Experimental Economics* 4: 203–219.

List, John, and Todd Cherry. 1999. "Learning to Accept in Ultimatum Games: Evidence from an Experimental Design that Generates Low Offers." *Experimental Economics* 3: 11–31.

McKelvey, Richard D. 1991. "An Experimental Test of a Stochastic Game Model of Committee Bargaining." In *Laboratory Research in Political Economy*, ed. Thomas R. Palfrey. Ann Arbor: University of Michigan Press, 139–68.

McKelvey, Richard D., and Thomas R. Palfrey. 1995. "Quantal Response Equilibrium for Normal Form Games." *Games and Economic Behavior* 10: 6–38.

McKelvey, Richard, and Thomas R. Palfrey. 1998. "Quantal Response Equilibrium in Extensive Form Games." *Experimental Economics* 1: 9–41.

Romer, Thomas, and Howard Rosenthal. 1978. "Political Resource Allocation, Controlled Agendas, and the Status Quo." *Public Choice* 33: 27–43.

Roth, Alvin E. (1995). "Bargaining Experiments." In *Handbook of Experimental Economics*, Alvin E. Roth and John Kagel, eds. Princeton: Princeton University Press.

Schmitt, Pamela. 2004. "On Perceptions of Fairness: The Role of Valuations, Outside Options, and Information in Ultimatum Bargaining Games." *Experimental Economics* 7: 49–73.

Wooldridge, Jeffrey. 2002. *Econometric Analysis of Cross Section and Panel Data*. Boston: MIT Press.