



University of Cyprus
Department of Economics

Working Paper 05-2018

Beyond Outcomes: Experimental Evidence on the Value of Agreement

Philippos Louis, Matias Núñez, and Dimitrios Xefteris

Department of Economics, University of Cyprus, P.O. Box 20537, 1678 Nicosia, Cyprus
Tel.: +357-22893700, Fax: +357-22895028, Web site: <http://www.ucy.ac.cy/econ/en>

Beyond Outcomes: Experimental Evidence on the Value of Agreement

Philippos Louis*

University of Cyprus

Matias Núñez†

Université Paris-Dauphine

Dimitrios Xefteris‡

University of Cyprus

May 2018

Abstract

Does an individual assign a higher value to a group decision she has explicitly agreed with? Or does she only care about the intrinsic features of the outcome? Since it is difficult to address this question in natural settings, we employ a laboratory experiment where, after the group collectively decides on an issue, each individual may propose a revision to the group decision. We find that outcomes generated by congruent mechanisms –i.e. procedures that incentivize subjects to agree with the winning alternative– are revised to a far lesser extent compared to outcomes generated by outcome-wise identical mechanisms that encourage disagreement.

JEL classification codes: D71, D72.

Keywords: implementation; mechanism design; consensus; agreement; congruence; experiment; endorsements.

1 Introduction

In many instances of collective decision making, participants' preferences are aggregated by the means of procedures that push them towards agreement. Such *congruent* decision rules are used, for example, in several international organizations. The council of the E.U. decides by consensus on important issues like granting E.U. membership, while in the U.N. security council, consent of all permanent members is

*Department of Economics, University of Cyprus, P.O. Box 20537, Nicosia 1678, Cyprus. email: louis.philippos@ucy.ac.cy

†Université Paris-Dauphine, PSL Research University, CNRS, Lamsade, 75016 Paris, France. email: matias.nunez@dauphine.fr

‡Department of Economics, University of Cyprus, P.O. Box 20537, Nicosia 1678, Cyprus. email: xefteris.dimitrios@ucy.ac.cy

necessary for the adoption of a resolution. In other cases decision rules are *incongruent*: They encourage participants to exaggerate their disagreement, not only with each other, but also with respect to the implemented alternative. For example, for the setting of the London Interbank Offered Rate (LIBOR), banks submit proposals and the decision is made via an averaging process. Galton (1907) first noticed that such rules incentivize participants to propose extreme alternatives in order to bring the implemented alternative closer to their ideal one. Hence, the outcome is often very different from what any of the individual participants proposed.

From a standard rational choice perspective it makes no difference whether a decision mechanism is congruent or not. What matters is the outcome it produces and its welfare properties. While this assumption was key to shaping a formal approach to collective decision making with important findings over the years, its empirical relevance has not been unquestioned.¹ A long line of research, especially in organizational studies, looks at variables beyond the outcome itself and supports the view that decision rules can be important in determining the acceptance of a decision among the concerned individuals.² Indeed, a collective choice –in a variety of contexts– is typically not final and, hence, the degree to which it is deemed acceptable by the concerned parties largely determines the quality of its implementation and its persistence when opportunities for revision arise. Individuals may also enjoy procedural utility: They may value a group choice not only for *what* it is, but also for *how* it is taken (Sen, 1997).³ In contrast to the classical approach, these views suggest that the design of good institutions for group decisions should take into account their congruence dimension. Of course, whether the agreement fostered by a decision rule truly affects how group members value an outcome ex-post or not is still an open empirical question –one we attempt to answer in this paper.

A major problem with establishing the importance of agreement in natural settings, above and beyond the importance of the outcome itself, is the endogeneity of outcomes: Different rules typically produce different outcomes.⁴ Individuals may find one outcome more acceptable than the other not because of the way it was chosen, but simply because it is “better” in some dimensions. Furthermore, in real life the contexts in which different rules apply are very divergent and the rules themselves differ in a number of aspects apart from the agreement incentives that they provide (e.g., formal versus informal

¹See Maskin (2008) for an excellent introduction to mechanism design and implementation theory under this classic approach.

²See Mason and Mitroff (1981), Schweiger, Sandberg, and Ragan (1986), Priem, Harrison, and Muir (1995), and Hartnett (2011).

³See Sen (1997), Frey, Benz, and Stutzer (2004) and Frey and Stutzer (2005).

⁴For example, all decision rules in Schweiger et al. (1986) and similar studies typically lead to different resolutions.

deliberation, time constraints, and revision procedures). Therefore, we need to turn to more controlled environments to effectively isolate the value of agreement in group decisions.

Overcoming the endogeneity issue described above, even in the lab, is not trivial. To achieve this we use classic (Moulin, 1980) and more recent (Yamamura and Kawasaki, 2013; Núñez and Xefteris, 2017) findings regarding mechanism design in the single-peaked context. In particular, we compare subjects' behavior under two decision mechanisms: the Median Approval and the Simple Mean. These two rules are *outcome-wise identical*: For any given preference profile, they result in the same unique Nash equilibrium outcome. At the same time, they largely differ in the incentives for agreement. Considering that the outcome space is the unit interval, the Median Approval rule allows each individual to approve any subset of alternatives (i.e., give one vote to as many alternatives as she wants) and the outcome coincides with the median of the distribution of the votes cast by all voters. Since players have single-peaked preferences (i.e., each one is characterized by an ideal policy and prefers that the outcome is as close as possible to her ideal policy), incentives are such that in equilibrium the implemented policy is included in all individuals' sets of approved alternatives: Everyone agrees with the outcome. According to the Simple Mean rule, each individual reports a number, and the outcome coincides with the mean of the reports. The incentives lead individuals to vote for extreme alternatives. In equilibrium, this leads to an exaggerated disagreement between individual votes and the implemented policy. Thus, the Median Approval is a congruent mechanism and the Simple Mean an incongruent one. However, they both apply to the same class of problems, and, more importantly, they produce identical outcomes.

After subjects make a collective decision under one of these two mechanisms (first stage), they move to a random dictatorship phase (second stage), where each of them is allowed to propose any revision of the original decision, and each of the proposals gets implemented equiprobably. In this manner, we can elicit the post-decision individual respect towards the original outcome. Evidently, the nature of this test urges individuals to propose a revision of the outcome to their liking. But if differences in the magnitudes of these revisions are observed between the Median Approval and the Simple Mean treatment, they should be attributed to the mechanisms themselves and they indicate heterogeneous levels of post-decision outcome persistence. That is, since both mechanisms deliver similar outcomes, a potential difference in post-decision outcome persistence should be due to factors that are not related to the outcome itself.

In the first stage we do not find important differences between the outcomes of the Median Approval

and the Simple Mean treatment: Conditional on the group’s preference profile, the two mechanisms implement very similar outcomes, just as theory predicts. We do find, though, that the support for the winning alternative is much larger under the Median Approval treatment than under the Simple Mean treatment. For instance, in about 80% of the cases a subject endorsed the implemented alternative in the Median Approval treatment while in the Simple Mean treatment individual votes were on average one-third of the total measure of the alternatives’ space away from the outcome. In the second stage, while the payoff maximizing strategy is simply to propose one’s ideal policy –like any dictator game– we observe that under both the Median Approval and the Simple Mean treatment the proposal of a subject is, essentially, a convex combination of her ideal policy and the original group decision. That is, we find that the outcomes of a collective decision process are respected by individuals to a certain degree and are, hence, somewhat persistent. Importantly, *the weight assigned to the original decision is about 30% in the Median Approval treatment and about 10% in the Simple Mean treatment*, with this large difference being statistically significant at any conventional level. Since our design ensures that the two mechanisms deliver similar outcomes, the difference in the persistence of the outcomes in the two treatments must be attributed to the features of the different mechanisms used.

We construct a measure of agreement of a voter’s strategy with the outcome and we find that it has important across- and within-treatment explanatory power. In particular, we find that the distance between the implemented outcome and the closest vote of an individual (i.e., the voted alternative that is closer to the outcome than any other voted alternative) has large across- and within-treatment variation and strongly relates to the degree of outcome persistence in the second stage of the experiment. Hence, we provide evidence that not only supports the idea that congruent mechanisms deliver more persistent outcomes, but additionally, that agreement itself –even in the context of the same mechanism– has a part in explaining the persistence of group decisions.⁵ We extend this analysis to the group-level by defining a compatible measure of agreement of all voters’ strategies with the outcome, and we find that, while this measure also explains part of the treatment effect, it is much weaker compared to its individual-level counterpart. That is, we find that when an individual declares an outcome acceptable, the individual is less prone to propose significant revisions. At the same time, the strategies of the rest of the players do not play a significant role as far as her revision proposal is concerned.

⁵Our results are robust to alternative measures of agreement between a voter’s strategy and the implemented outcome –such as the average distance between one’s voted alternatives and the implemented outcome– but somewhat smaller in magnitude.

This last finding is arguably of independent interest as it points towards a potential novel route through which procedures affect individuals. It has been suggested (Sen, 1997; Frey et al., 2004) that procedures affect individuals' utility: a) directly, through their attitudes about specific procedures (e.g., whether they are fair, just, democratic, etc.) and b) indirectly, through the way they are treated by others within a procedure. In our study we detect a factor that acts orthogonally to these suggested factors: The incentives provided by the procedure shape individuals' within-procedure behavior, and the within-procedure behavior of an individual is found to have a significant effect on her attitude towards the procedure's outcome.

In what follows we discuss the relevant literature (section 2), provide a discussion regarding the two decision rules (section 3), detail our experimental design (section 4), present our results (section 5), elaborate on the treatment effect (section 6), and conclude (section 7).

2 Relevant literature

There has been a growing interest in the use of lab experiments to measure the effects of different collective decision processes on the effectiveness or acceptance of decision outcomes. Walker, Gardner, Herr, and Ostrom (2000) show that voting can increase efficiency through coordination in a common pool resource game. Dal Bó, Foster, and Putterman (2010) show how the effects of a policy on cooperation are stronger when it is chosen democratically. A similar effect is found for the performance of sanctioning institutions in public good games that are voted on instead of imposed exogenously (Tyran and Feld, 2006; Sutter, Haigner, and Kocher, 2010; Markussen, Putterman, and Tyran, 2013; Kamei, Putterman, and Tyran, 2015; Kamei, 2016) or chosen by an elected –versus imposed– leader (Grossman and Baldassarri, 2012). Markussen, Reuben, and Tyran (2014) also find that a scheme of intragroup competition is more effective in enhancing cooperation when it is chosen democratically. Beyond social dilemmas, Mellizo, Carpenter, and Matthews (2014) find that democratic processes can lead to higher effort in the workplace when compensation schemes are chosen by voting. As we measure the effect of the group choice process directly on outcome persistence, our results can help explain this positive effect of democratic institutions.

With the exception of Walker et al. (2000), these papers compare exogenous to endogenous choice. By contrast, since we focus on agreement, the processes we compare are all endogenous. Furthermore, our design allows us to experimentally control for differences in the outcomes produced by different

mechanisms, giving a clean identification of the effect of a given mechanism on the persistence of the outcome. Previous work typically controls for such effects econometrically.

Scholars in both management and psychology have long been interested in the effect of different decision processes on group decisions (see for instance Mason and Mitroff, 1981; Schweiger et al., 1986; Schweiger, Sandberg, and Rechner, 1989; Priem et al., 1995; and Hartnett, 2011). Another long stream of literature in organizational studies examines the role of conflict in teams and groups. Two extensive meta-studies (De Dreu and Weingart, 2003; De Wit, Greer, and Jehn, 2012) find that conflict is in general negatively correlated with group performance, as measured by different metrics. While there is some support for the idea that conflict can be beneficial in specific contexts (Jehn, 1994; Jehn, Chadwick, and Thatcher, 1997; Jehn and Mannix, 2001), this is not universally true for the relationship between conflict and group satisfaction. Our work complements this literature by looking at conflict that is created –or mitigated– by the decision processes used. We do not measure the effect of disagreement on the outcome; in fact we use a design that minimizes any possibility for such an effect. This allows us to obtain estimates of the causal effects of disagreement on the persistence of the outcome. To our knowledge, such an incentivized elicitation of satisfaction with the outcome has not been applied in this literature.

Our results can be interpreted as evidence for procedural utility: Processes matter above and beyond their explicitly associated outcomes. The idea has its origins in social psychology (Thibaut and Walker, 1975; Lind and Tyler, 1988), but it has also been advanced by economists (Sen, 1997; Frey et al., 2004; Frey and Stutzer, 2005). Research in this area focuses mainly on moral characterizations of processes, such as whether participants are treated equally or fairly, and the effect the moral characterization of a process may have on outcomes and their acceptance. In our case, we look at processes that, while resulting in divergent levels of agreement, can arguably not be ranked in terms of how fair they are. Hence, procedural utility in our case is independent of procedural justice. Our finding that an individual’s agreement with an outcome, as expressed through her vote, makes her less likely to change it when given the opportunity seems to be parallel with what Corazzini, Kube, Sebastian, Maréchal, André, and Nicolo (2014) find. Their paper looks at how the process of electing leaders may incentivize them to keep their promises. This of course relates to the literature on lying aversion (Gneezy, 2005; Fischbacher and Föllmi-Heusi, 2013). In our experiment subjects are not aware when voting that they will have an opportunity to revise the outcome. Hence, their outcome-revision behavior seems more

connected to a desire for self-consistency or self-concept maintenance (Mazar, Amir, and Ariely, 2008).

To properly choose the decision context and the employed mechanisms, one has to turn to the theoretical literature. We have chosen to focus on the single-peaked domain since a) it is quite intuitive and easy to explain in the lab, and b) it is the only one, to our knowledge, for which outcome-wise identical congruent and incongruent mechanisms exist. Renault and Trannoy (2005) and Yamamura and Kawasaki (2013) analyze the properties of the Simple Mean mechanism and show that the unique Nash equilibrium outcome under the average voting rule must be equivalent to the median of $(t_1, t_2, \dots, t_n, \frac{1}{n}, \frac{2}{n}, \dots, \frac{n-1}{n})$, where (t_1, t_2, \dots, t_n) is the vector of the n players' ideal policies. They also prove that in equilibrium most players select an extreme announcement (hence, the Simple Mean mechanism is incongruent). Núñez and Xefteris (2017) prove how to implement the same outcome using the Median Approval mechanism that leads players to endorse the implemented alternative (hence, the Median Approval mechanism is congruent). Finally, Gershkov, Moldovanu, and Shi (2017) show how to make the same decision through sequential quota procedures. A study of the comparative effects of simultaneous versus sequential mechanisms is beyond the scope of this analysis, but it presents itself as an interesting avenue of research for the future.

3 The two mechanisms

As described in the introduction, we focus on two mechanisms: the *Simple Mean* mechanism and the *Median Approval* mechanism. This section reviews their definitions and the equilibrium prediction for the situation tested experimentally. For a formal derivation of these results, we refer the reader to Renault and Trannoy (2005) and Yamamura and Kawasaki (2013) for the Simple Mean mechanism and to Núñez and Xefteris (2017) for the Median Approval mechanism.

We consider a committee with three individuals; a decision $x \in [0, 100]$ needs to be made. Individual preferences over outcomes are summarized by the utility function $u_i(x) = 100 - |x - t_i|$. Individual utility is maximized at $x = t_i$, so that t_i denotes the individual's most preferred decision. The larger the difference between the decision x and t_i , the smaller the individual utility. Our arguments do not depend on the precise shape of the utility functions and extend as long as individual preferences have a unique preferred decision (i.e., single-peaked preferences).

In order to ease the comparison, we focus on the case with $t_1 < t_2 < t_3$ with $t_1 \leq \frac{200}{3}$ and $t_3 \geq \frac{100}{3}$. In this case, both mechanisms under consideration admit a unique equilibrium outcome and a simple

Equilibrium outcome




Figure 1: Equilibrium outcome as a function of t_2 with $t_1 \leq \frac{200}{3}$ and $t_3 \geq \frac{100}{3}$.

derivation of the equilibrium strategies. This equilibrium outcome, as depicted by Figure 1, is the median of the peaks t_1, t_2, t_3 jointly with $\frac{100}{3}$ and $\frac{200}{3}$.

Simple Mean mechanism: Each player $i \in N$ simultaneously submits a value $s_i \in [0, 100]$. For each vector of announcements $s = (s_1, s_2, s_3)$, the outcome $\theta_{\mathbf{SM}}(s)$ equals:

$$\theta_{\mathbf{SM}}(s) = \frac{s_1 + s_2 + s_3}{3}.$$

Equilibrium Behavior: In equilibrium, each player casts a strategy that minimizes the distance between the outcome $\theta_{\mathbf{SM}}(s)$ and her own peak. Player 1, the player with the lowest peak, always announces 0 since she anticipates that the outcome is higher than her peak and wants to shift the outcome as much as possible to the left. Similarly, Player 3, the player with the highest peak, always announces 100 since she wants to shift the outcome as much as possible to the right.

The strategy of the median player, Player 2, depends on the value of her type t_2 . If the median type is low ($t_2 \leq \frac{100}{3}$), then Player 2 announces 0 and, by symmetry, if the median type is high (i.e., if $t_2 \geq \frac{200}{3}$), then Player 2 announces 100. Finally, if the median peak is centered ($\frac{100}{3} \leq t_2 \leq \frac{200}{3}$), the median player plays a strategy that allows to obtain t_2 as an outcome: This strategy equals $3t_2 - 100$.

In any equilibrium, the outcome is equal to $f(t_1, t_2, t_3) = \text{median}(t_1, \frac{100}{3}, t_2, \frac{200}{3}, t_3)$. The equilibrium is unique since slightly altering one's announcement affects the final outcome independent of the announcement of the rest of the players (see Proposition 3 in Yamamura and Kawasaki, 2013 for a precise statement of the conditions that lead to a unique equilibrium).

Median Approval mechanism: Each player $i \in N$ simultaneously submits an interval $b_i = [b_i^-, b_i^+]$ with $b_i^- \leq b_i^+$. Player $i \in N$ casts one vote for each alternative included in her chosen interval. Let $\mu(b_i) = b_i^+ - b_i^-$ denote the measure of b_i and, for each set of intervals $b = (b_1, b_2, b_3)$, $\mu(b) = \mu(b_1) + \mu(b_2) + \mu(b_3)$ the measure of b . For each $x \in [0, 100]$ and each set b , $s_x(b) = \#\{i \in N \mid x \in b_i\}$ denotes the score of x at b . Note that if $\mu(b) = 0$, each announcement is a singleton. If $\mu(b) > 0$, the distribution of votes $\phi : \mathcal{B}^n \times [0, 100]$ is denoted by $\phi(b, z) = \frac{1}{\mu(b)} \int_0^z s_x(b) dx$.

For each vector of announcements $b = (b_1, b_2, b_3)$, the outcome $\theta_{\mathbf{MA}}(b)$ equals:

$$\theta_{\mathbf{MA}}(b) = \begin{cases} \text{median}(b_1, b_2, b_3), & \text{if } \mu(b) = 0 \\ \min\{z^* \in [0, 100] \mid \phi(b, z^*) = \frac{1}{2}\}, & \text{otherwise} \end{cases}$$

Figure 2 depicts the computation of the median of the announced intervals. After plotting the intervals (Figure 2b), we plot the vote distribution (Figure 2c) –that is, the number of votes that each alternative gets by the players. The median of the intervals coincides with the point that divides the area below the vote distribution into two equal parts.

Equilibrium Behavior. In a similar fashion to the Simple Mean mechanism, each player chooses a strategy that minimizes the distance between the outcome and her own peak. The unique equilibrium outcome is also equal to $f(t_1, t_2, t_3) = \text{median}(t_1, \frac{100}{3}, t_2, \frac{200}{3}, t_3)$.

Player 1 announces an interval b_1 that ranges from 0 to $f(t_1, t_2, t_3)$. Namely, she votes for the outcome $f(t_1, t_2, t_3)$ and for any alternative located to its left. By symmetry, Player 3 approves the interval b_3 that goes from $f(t_1, t_2, t_3)$ until 100, voting for that outcome and all the alternatives located to its right.

The median player, Player 2, plays a strategy that depends on the value of t_2 . When $t_2 < \frac{100}{3}$ ($t_2 > \frac{100}{3}$) then she votes for $b_2 = [0, \frac{100}{3}]$ ($b_2 = [\frac{200}{3}, 0]$), and the outcome is equal to $\theta_{\mathbf{MA}} = \frac{100}{3}$ ($\theta_{\mathbf{MA}} = \frac{200}{3}$). When $\frac{100}{3} \leq t_2 \leq 50$ ($50 \leq t_2 \leq \frac{200}{3}$), she can vote any alternative from 0 to $4t_2 - 100$ (from $4t_2 - 200$ to 100) –that is, also t_2 – inducing the implementation of her ideal policy.

To better understand the equilibrium behavior, consider an example with $t_1 < t_2 = 40 < t_3$ and the strategy profile $b_1 = [0, 40]$, $b_2 = [0, 60]$ and $b_3 = [40, 100]$. These strategies lead to the implementation

Announcements	b_i^-	b_i^+
Individual 1	0	40
Individual 2	30	50
Individual 3	30	90

(a) Individuals report intervals.

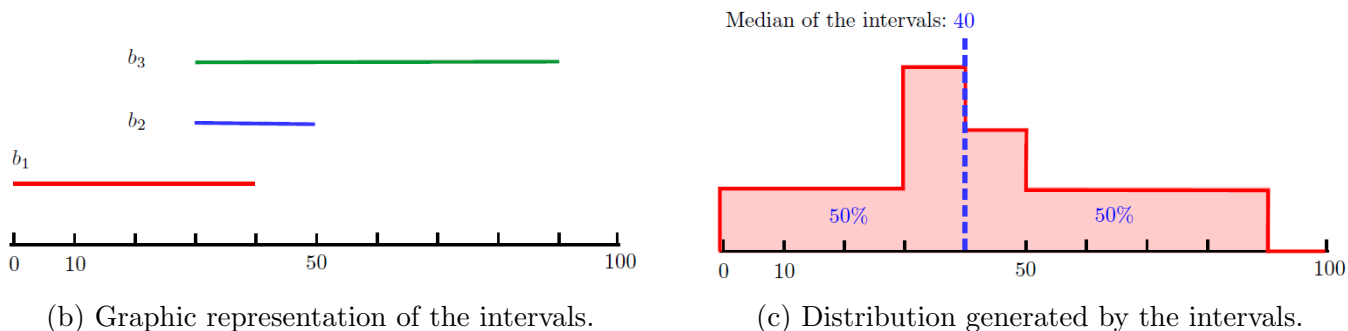


Figure 2: Computing the median of the intervals.

of alternative 40 –which is voted by all players– since in total $40+60+60=160$ units of votes are cast, and half of them are given to alternatives to the left (right) of 40. To see why this is an equilibrium consider deviations of the first player. If, for example, she expands her interval to $b'_1 = [0, 46]$, then the total votes cast will be $46+60+60=166$, and the implemented alternative will thus move from 40 to 41. Since $t_1 < t_2 = 40$, this is not a profitable change for player 1. If she shrinks her interval to $b''_1 = [0, 20]$, then the total votes cast will be $20+60+60=140$, and the implemented alternative will thus move from 40 to 45. Since $t_1 < t_2 = 40$ again this is not a profitable deviation. Of course these deviations are just indicative of what may happen: Players have a variety of different options to choose from. These few cases though are sufficient to show that both by voting for alternatives to the right of the implemented outcome and by not voting for alternatives to the left of the implemented alternative, a player can shift the outcome to the right. Hence, if a player's ideal policy is to the left, her only best response is to vote for the outcome and all alternatives to its left.

4 Experimental Design

The experimental design is geared towards answering our main question on the influence of agreement on the persistence of outcomes. For this, we use a between-subject design with two treatments, each of which has two parts. In the first part, subjects make collective decisions using a different decision rule in

each treatment. The decision rules are such that theory predicts the same outcome with different levels of agreement. In the second part subjects make individual proposals, one of which is chosen randomly for each group, to replace the outcome of the first part. Given no difference between treatments in the outcomes of the first part, a treatment effect in the second part can be interpreted as an effect of agreement on the stability of the outcome. We now explain the details of the experiment and our design choices.

The experiment took place at the University of Cyprus Lab of Experimental Economics (UCY LExEcon). A total of 90 subjects, all students of the University of Cyprus, participated in six equally sized sessions, with three sessions per treatment.⁶ Recruitment was done using ORSEE (Greiner, 2015). The experiment was computerized, and the software was programmed and run using zTree (Fischbacher, 2007). An outline of the design is presented in Table 1.

TABLE 1: THE TWO EXPERIMENTAL TREATMENTS

Treatment	Part A (20 periods)	Part B (20 periods)	N	# of Sessions	Subjects per session	Group size
MA	Interval voting with median as outcome	Random dictator game	45	3	15	3
SM	Single vote with mean as outcome	Random dictator game	45	3	15	3

Timing For both treatments, subjects received written instructions for part A after entering the lab.⁷ These were also read aloud to establish common knowledge. After part A finished, instructions for part B were distributed and read aloud. At the end of part B, subjects were informed about their profits and paid privately before leaving the lab.

⁶Two pilot sessions were completed before the main experiment to finalize the design, fine-tune some of the parameters, and receive feedback on the instructions. Data from these pilot sessions are not included in any of our analysis.

⁷A translation of the instructions, originally in Greek, can be found in the Appendix.

Collective choice - Part A In each round of part A, subjects are placed in groups of three. Each group needs to choose collectively an integer between 1 and 100 as the group's *destination*. Each group member has an individual *starting point*, that is, a different integer between 1 and 100. The payoff in each period is then 100 points minus the distance between the *destination* and the subject's *starting point*. *Starting points* are common knowledge and are different for every subject in each period.⁸ Groups are reshuffled in each period, and subjects do not know the identity of the other group members.

Treatment - Median Approval (MA) In treatment *MA*, to choose the *destination*, each group member chooses an interval of integers between 1 and 100, and each location in the interval receives a single vote. The collective choice is the maximum median of the distribution of votes. Subjects choose the interval by moving specific bars in their screen that mark the lower and higher limits of the interval of votes.

Treatment - Simple Mean (SM) In treatment *SM* subjects can vote for a single location by choosing an integer between 1 and 100. The collective choice is the mean of all three votes. Voting takes place by moving a bar to the specific location that the subjects wishes to vote for.

Voting, information and time limit Voting in both treatments lasts for $60 + x$ seconds,⁹ where x is a number between 1 and 10, chosen randomly in each round and not known to the subjects. During this time, each subject is informed about her and others' starting points and can enter her votes (interval or single vote, depending on the treatment). She can also observe the votes entered by other group members in real time. At any given point in time the software calculates the *destination* and the group members' payoffs. These are shown on the screen as a clock counts down from one minute. At 10 seconds, a text starts blinking indicating that time is almost up, after which it turns red and indicates that voting may finish at any moment. The *destination* for the period is determined by the votes entered when the $60 + x$ seconds finish. After that, a screen appears informing subjects about the results of the voting: each subject's votes, the final *destination*, and subjects' payoffs.

⁸We chose starting points in a way that maximized power for the experiment in terms of detecting a treatment effect in the second part. To that end we ran simulations using many different sets of starting points and hypothesizing a treatment effect of magnitude and variance similar to what we found in the pilot session. We then chose the set of starting points where the effect was stronger in a linear regression similar to the one corresponding to the first column of Table 3. More details on the exact process are available upon request.

⁹For the first two periods this is extended to $90 + x$ to allow subjects to get familiarized with the voting environment.

Random Dictator Game - Part B Part B is identical for both treatments. In each round of part B, subjects are placed in the same groups as in the corresponding round of part A. Again, the group needs to choose a collective *destination* from the same *starting points* as in the respective round of part A. Unlike part A, the choice is now made by a random dictator: Each group member proposes a location by choosing an integer between 1 and 100. One of the three proposals is chosen randomly as the group's destination for the round, and payoffs are determined in the same manner as in part A. Before making a proposal, subjects are reminded of all *starting points*, all votes, and the chosen *destination* in the corresponding round of part A. They make a proposal by clicking on a location on the screen and then on a 'submit' button. They can revise their proposal as many times as they wish before clicking 'submit'. For each location they click, the software calculates and shows all players' payoffs if that proposal is selected. They cannot see the proposals of the other group members, and they are not informed about the others' proposals and the final outcome until the end of all rounds.

Payments After part B is completed a screen informs subjects about the outcomes of all rounds in both parts. One round from each part is chosen randomly, and payoffs in that round are used to determine the subjects' payment for the experiment. Subjects receive €1 for every 15 points earned in the selected round of each part, plus an additional €3 as a participation fee. Subjects earned €13.21 on average across all sessions.

5 Results

5.1 Part A

5.1.1 Structure of the data

The only parameters that differed between subjects and rounds were the subjects' starting points, which determine their payoffs. Nevertheless, the exact same set of parameters was used across all six sessions. That is, for any combination of starting points used for a group in a specific round of a session, there was another group in all other sessions with the same starting points in the same round. Furthermore, the exact same sequence of parameters was assigned to subjects in all sessions.

The collective choice of a group is a function of the group members' votes (i.e., single votes or vote intervals). These may depend on the specific combination of group members' starting points. Hence, the outcome of the collective choice can be seen as a draw from a distribution that depends on these

starting points and the treatment. Since we used the same parameters in both treatments, we have three draws for each treatment with the same starting points. The statistical tests we use to compare results in both treatments take advantage of this structure by averaging out the three observations per treatment and using *paired* non-parametric tests.

5.1.2 Voting process

We use a voting process with real-time feedback to allow for fast within-round learning.¹⁰ A random ending point is used to discourage extreme “snipping” behavior, which was observed in a pilot session with fixed ending points.¹¹ One difference between the voting mechanisms that is worth noting concerns the degree to which a single voter can affect the outcome, given the others’ votes. In *MA* the outcome can, theoretically, move up to 98 points by a single individual’s change in votes, but only for very particular choices of the rest of the players.¹² In most scenarios, a single voter’s power over the outcome is quite limited: When two voters approve of many alternatives, the median of the induced vote distribution is moderately responsive to a change in the strategy of the third voter. On the other hand, in *SM*, it is always possible for a single voter to move the outcome to any point within a range of 33 points. Moreover, in *SM*, it is also practically easier to move the outcome since an individual’s vote can change to any other with one direct move, while in *MA* a change to a different strategy involves a two-step process: A subject needs to change one end of the interval before making a change to the other.

In the two panels of Figure 3, we show the provisional outcomes across time for 10 randomly chosen groups in each treatment. The patterns are fairly typical for all groups. In *MA*, movements are more gradual. Substantial movements happen mostly in the first 30 seconds. In *SM*, there is more volatility throughout the round. There are often substantial moves of the provisional outcome between the 40th and 60th second. After that movements are rare and small in magnitude. The difference in volatility reflects the preceding discussion. The increase in volatility towards the end of the round in *SM* could reflect some residual “snipping” attempts.

¹⁰Moreover, it has been shown that feedback exchange among players prior to the group decision point helps diminish outcome-related institutional differences (see, for instance, Goeree and Yariv, 2011 and Gerardi and Yariv, 2007), which is desirable in our case.

¹¹With a fixed end point, many subjects would significantly change their votes in the last seconds of voting in an effort to achieve a more favorable outcome. The term “snipping” has been used in online auctions to describe bidders that only submit a bid in the last moment to avoid driving up the price through a bidding war. See, for example, Ockenfels and Roth (2006).

¹²This happens in the extreme scenario where all voters cast a single vote on location 1. Then one of them can switch and vote the interval [99,100], moving the outcome from 1 to 99.

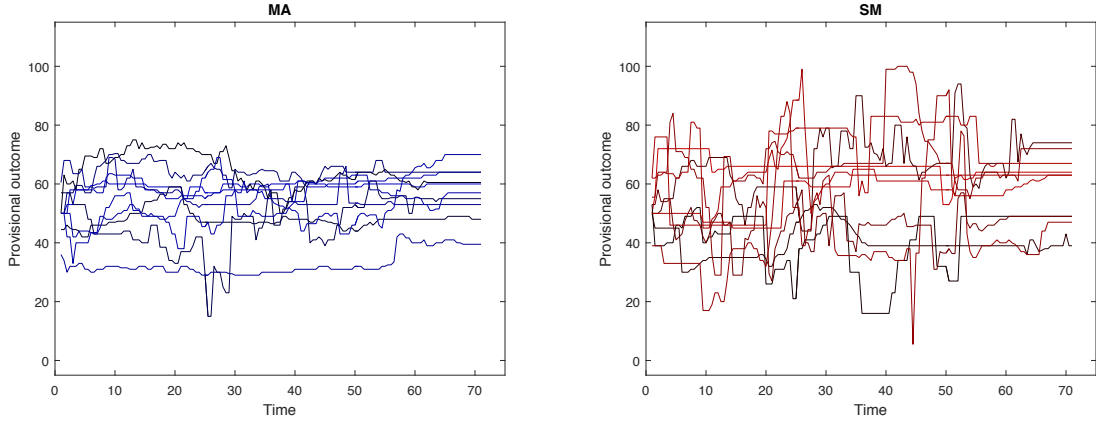


Figure 3: Volatility of collective choices. The left panel shows 10 randomly chosen groups and how the provisional outcome changes across time during the voting process in the MA treatment. The right panel shows the same for 10 groups with the same starting points in the *SM* treatment.

5.1.3 Agreement

Our premise for the choice of experimental design is that the *MA* mechanism is more congruent than the *SM* mechanism. We now quantify this claim and show how this difference is reflected in the data. To this end, we construct two measures of agreement with the outcome: *individual agreement* and *group agreement*. The first one refers to the agreement of the outcome with an individual's votes, while the second one refers to the agreement of the outcome with the votes of all group members. These will also be of interest later when we try to understand the underlying forces that give rise to the treatment effects we find in part B.

Given the differences of the two mechanisms, we need to select a measure of agreement that is applicable both to interval and single votes and that can provide a smooth estimate of outcome endorsement by the individual or the group. We do that by taking the distance between the subject's vote that is closest to the outcome and the group's outcome and subtracting it from 100. Formally, let V_i be the set of all locations that subject i votes for (in *SM* this set is a singleton). Then,

$$\text{Individual agreement of } i = 100 - \min\{|x - \text{outcome}| \text{ s.t. } x \in V_i\}$$

This gives an individual measure of agreement with a range from 0 (complete disagreement with the outcome) to 100 (the outcome is a location the individual voted for). For *group agreement*, we simply take the average of *individual agreement* across all group members. Figure 4 shows the empirical cdf's for these measures for each treatment. It is clear that treatment *MA* displays substantially higher levels

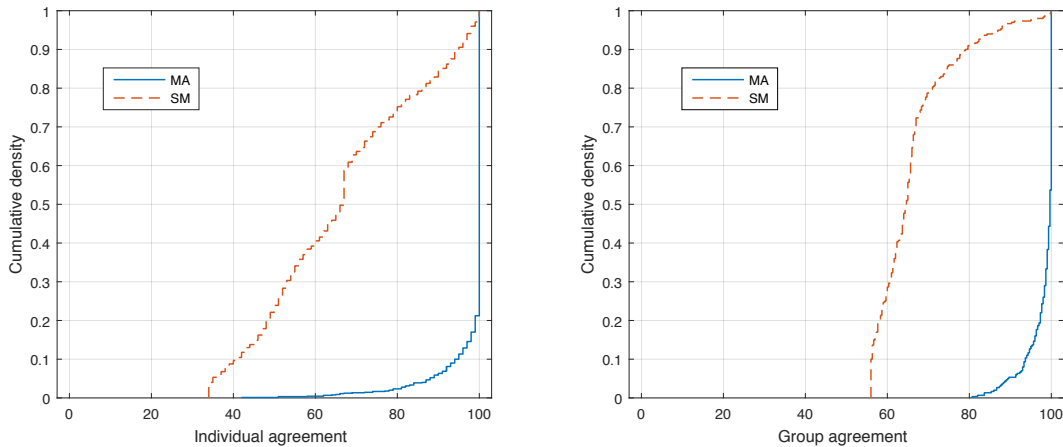


Figure 4: Measures of agreement. The left panel shows the empirical cdf for *individual agreement* in each treatment. The right panel shows the empirical cdf for *group agreement* in each treatment.

of both individual and group agreement.¹³ In particular, about 80% of individual strategies in *MA* exhibit the highest degree of outcome endorsement, while in nearly half of the cases, the outcome is unanimously approved by all group members. Finally, there is a stochastic dominance relationship in the agreement levels between treatments –both at the individual and the group level– which indicates that subjects vote for alternatives closer to the outcome in *MA* compared to *SM*, even if one focuses on cases where the outcome is not fully endorsed.

5.1.4 Pareto efficiency

Another concern one might have is whether these voting protocols allow subjects to reach “reasonable” outcomes. We discuss the mechanisms’ outcomes in more detail below, but as a first approach we look at the frequency of non-Pareto outcomes. Pareto outcomes lie between the lowest and highest starting points in a group. As a result, a non-Pareto outcome can only occur if at least two group members deviate significantly from best-responding.¹⁴ Overall, we observe non-Pareto outcomes in less than 5% of collective choices in the experiment: 8 out of 300 outcomes in *MA* and 20 out of 300 outcomes in *SM*. This small difference can be attributed to the more volatile nature of the *SM* mechanism and the large majority of non-Pareto outcomes (20 out of 28) is mainly observed in the initial ten rounds. We do not exclude these observations from our analysis, but our main results would only be strengthened if we did.

¹³Alternative measures can be constructed, using, for instance, the average instead of minimum distance of votes from the outcome. Our results do not change qualitatively. Still, measures based on the minimum distance are stronger predictors of proposals in part B both across and within treatments (see section 5.2).

¹⁴In some cases, non-Pareto outcomes can only occur if all three group members are far from best-responding.

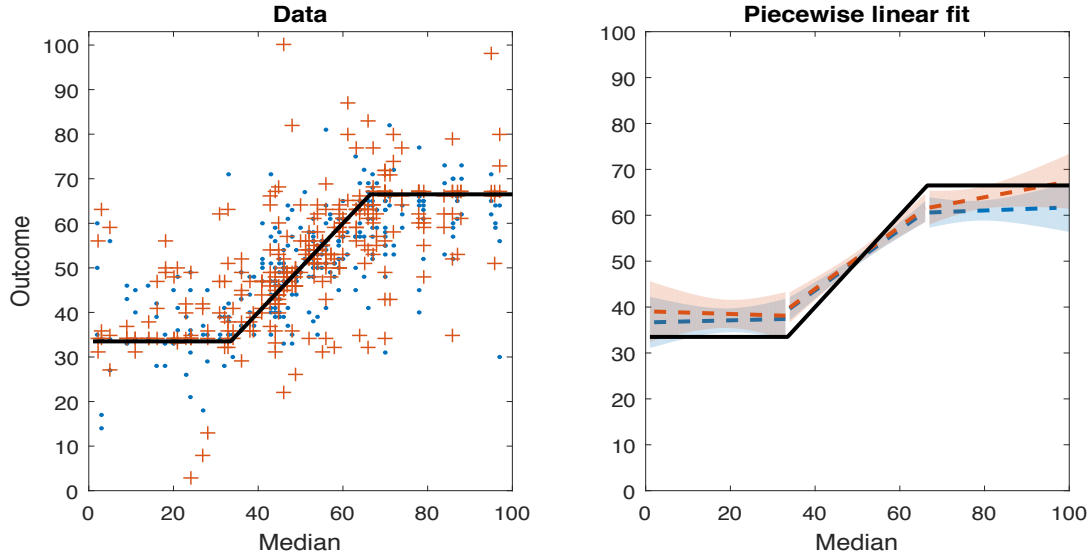


Figure 5: Collective choices for all groups in all rounds. Median refers to the group members’ median starting point. The solid line in both panels corresponds to the Nash equilibrium outcome. For our choice of parameters this equilibrium depends entirely on the position of the median. In the left panel, data from the *MA* and *SM* treatments are indicated by dots and crosses respectively. In the right panel, dashed lines correspond to linear fits to the data with respect to the three Nash equilibrium regions, for the *MA* and *SM* treatments in blue and red, respectively. The corresponding shaded areas indicate 95% confidence intervals.

Despite the differences between the collective choice institutions, and even the differences in the paths that groups take to reach decisions, the outcomes in both treatments look remarkably similar. We discuss this next.

5.1.5 Comparison of part A outcomes

The left panel of Figure 5 shows the collective choices of all groups, in all rounds, for both treatments. As can be seen from the graph the Nash equilibrium does a relatively good job of predicting the outcomes in both treatments. Starting points in all groups are chosen so that the equilibrium outcome is constant and equal to 33 ($\approx \frac{100}{3}$) when the median starting point in a group is below 33. For values higher than that, but lower than 67 ($\approx \frac{200}{3}$), the equilibrium outcome coincides with the median starting point. The equilibrium is again constant and equal to 67 when the median exceeds that value. As we observe, the collective choice tends to be very close to the median when it lies between 33 and 67. When the median is below or above this interval the dependency disappears, and the collective choice hovers around 33 and 67.

The above is supported by the piecewise linear fits shown in the right panel of Figure 5. From this

graph we also note that the outcomes in both treatments tend to be closer to the center of the range compared to the predicted Nash outcome. Most importantly, though, for our research question, we do not observe any systematic differences in the outcomes across the two treatments.

TABLE 2: SUMMARY STATISTICS OF COLLECTIVE CHOICES IN PART A FOR BOTH TREATMENTS.

Treatment	average absolute deviation from Nash ^a	average deviation from Nash to center ^b	average efficiency (% of max) ^c	average inequality (Gini coeff.) ^d
MA (N = 300)	6.47 (std 6.36)	3.56 (std 8.34)	95.3 (std 4.95)	8.80% (std 4.35)
SM (N = 300)	6.72 (std 8.32)	3.21 (std 10.20)	94.9 (std 5.91)	9.58% (std 4.87)
statistical difference (p-value) ^e	0.677	0.736	0.418	0.008

There are 300 observations in each treatment. Each observation k refers to a group making a collective choice in a specific round.

^a The absolute deviation from Nash is calculated as: $|outcome_k - Nash_k|$.

^b The deviation from Nash to center is calculated as:
 $(outcome_k - Nash_k) \times sign(50.5 - Nash_k)$.

^c Efficiency is calculated as: $\sum_i pay_{ik}(outcome_k) / \sum_i pay_{ik}(med_k)$, where $pay_{ik}(x) = 100 - |start_{ik} - x|$ and $med_k = median(start_{ik})$.

^d The Gini coefficient is calculated as:

$$\sum_i \sum_j |pay_{ik}(outcome_k) - pay_{jk}(outcome_k)| / 6 \sum_i pay_{ik}(outcome_k).$$

^e Wilcoxon signed rank test.

We further explore this issue by comparing the outcomes across treatments in three dimensions: their location, their efficiency, and their degree of inequality. Any of these dimensions could affect how an individual evaluates the collective choice. Table 2 summarizes the outcomes in each treatment across these dimensions.

Location As one can see from Figure 5, outcomes in the two treatments seem to lie close to each other in a statistical sense. This is further supported by comparing the distribution of absolute deviations from Nash and the distribution of deviations from Nash to the center (see Table 2, columns 2 and 3). In both cases the means are very close and according to a Wilcoxon signed-rank test the distributions

are not significantly different.¹⁵ Outcomes appear to be slightly more “noisy” in the *SM* treatment; nevertheless, this difference is not significant.¹⁶ We conclude that the outcomes in the two treatments do not differ substantially in terms of location.

Efficiency In our setup, maximum efficiency is achieved when the collective choice coincides with the median starting point. Since collective choices are close to the Nash equilibrium outcome, which in turn coincides with the median for a broad range of observations, it is not surprising that high levels of efficiency are achieved in both treatments. Outcomes are slightly more efficient in the *MA* treatment, but not significantly so.

Inequality We use the Gini index to measure income inequality associated with the outcome of a collective choice. In our setup, this measure is minimized when the collective choice coincides with the midpoint between the two more extreme group members. This point is also the maxmin choice: It maximizes the lowest payoff achieved by any group member. Outcomes in the *MA* treatment are slightly less unequal than in *SM*. The difference is small in magnitude but statistically significant. Still, this seems to be the consequence of some extreme observations, and the difference becomes non-significant if we do not include outcomes from the first five rounds or exclude the few non-Pareto outcomes from the sample. We conclude that there are some small differences in outcomes between treatments in terms of inequality. As we show in the next section these variations are not able to account for the large treatment effect we find in part B.

5.2 Part B

5.2.1 Structure of the data

In each round in part B, subjects are put in the same group as in the corresponding round in part A. They are shown the group’s votes and outcomes in part A and are assigned the same starting points. They are asked to propose a new destination for part B. Proposals are not restricted in any way and can be any point between 1 and 100. For each group, one of the proposals is selected randomly to determine payoffs. Hence, for each individual, we have 20 proposals from different groups, which gives 900 observations per treatment.

¹⁵Throughout the text we refer to differences as being statistically significant at the 1% level. We also report the p-value for the corresponding test.

¹⁶Wilcoxon signed-rank test: $VAR_{MA} = VAR_{SM}$, $p - val = .959$.

To summarize the data, we compute the deviation of the proposal from the corresponding part A result for each subject and each round. This is normalized to 0 when these coincide and 1 when the proposal coincides with the subject’s corresponding starting point, which is also the payoff maximizing choice (formally, our deviation measure is given by $\frac{\text{proposal} - \text{part A result}}{\text{starting point} - \text{part A result}}$). These deviations are presented for each treatment in Figure 6. Values larger than 1 indicate deviations away from the part A result to the direction of one’s starting point, which lie even farther away from the part A result compared to the subject’s starting point; and values smaller than zero indicate deviations away from the part A result to the opposite direction of the subject’s starting point. Naturally, the majority of proposals (83.8% in total) take values between 0 and 1.

The large jumps in the empirical cumulative density functions in the right panel of Figure 6 at 1 indicate that a large fraction of proposals in both treatments (52.7% in total) coincide with the individuals’ starting points. Nevertheless, a substantial number of proposals do not. In particular, 31.1% proposals have a deviation smaller than 1 but greater than or equal to 0. This suggests that the outcome in part A affects subjects’ proposals in part B. In fact, while small, there is a noticeable jump in the cumulative distribution of proposal deviations at 0 in both treatments. It is also worth noting that the above observations cannot be attributed to a subset of subjects consistently proposing their own starting points. Heterogeneity in behavior is, of course, present. Still, only 1 out of 90 subjects proposed his/her starting point in all rounds of part B, while all subjects proposed their own starting points at least once.

5.2.2 Treatment effect on outcome stability

We observe significant differences in the distribution of proposals across treatments.¹⁷ In particular, 366 (40.67%) proposals equal the individuals’ starting points in treatment *MA*, while this number goes up to 583 (64.78%) in treatment *SM*. This difference across treatments suggests a strong treatment effect. In other words, it appears that subjects are more willing to sacrifice some of their own payoff and propose something closer to the decision taken in part A, if the latter is the outcome of the *MA* mechanism. We find further support for the existence of this treatment effect by fitting linear regressions that explain proposals.

In the first column of Table 3, we report results from a linear regression that explains subjects’ proposals as a convex combination of their starting points and part A results, and the interaction of

¹⁷Two sample Kolmogorov-Smirnov test: $MA \neq SM - p < 0.0001$.

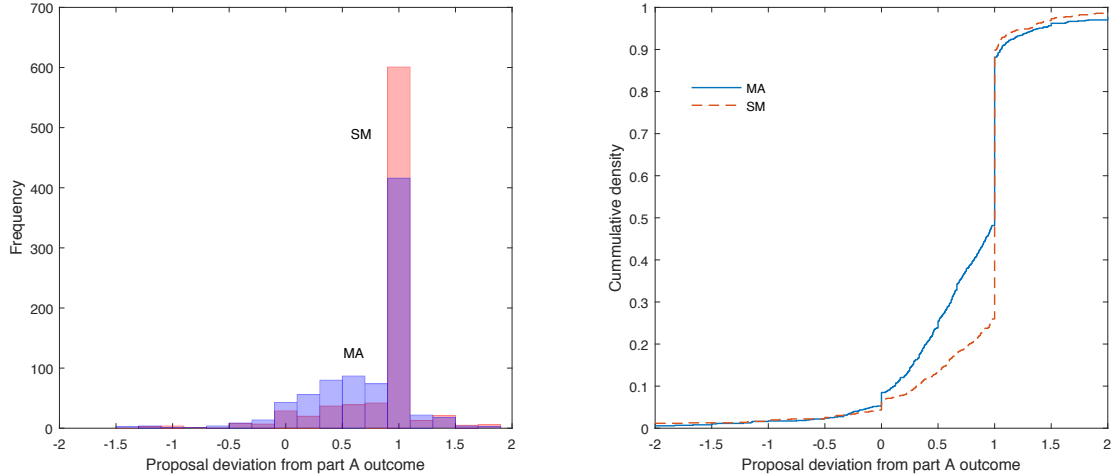


Figure 6: Subjects’ part B proposals. The graph represents proposals as deviations from the corresponding part A outcome. These are normalized to 0 when the proposal coincides with the part A outcome and 1 when the proposal coincides with the subject’s corresponding starting point. The left panel shows the histogram of the deviations for each treatment. In the right panel, we show the corresponding empirical cumulative density.

these variables with the treatment dummy (which takes value 0 in *MA* and value 1 in *SM*). The estimated values verify what we see in Figure 6, and we find all coefficients to be highly significant. Subjects put a substantial weight on the result of part A in choosing their proposal in treatment *MA*. In treatment *SM* this effect is reduced by two-thirds.¹⁸

Subsequently, we introduce two new explanatory variables. The first is the *efficiency maximizer*, which corresponds to the point that, if chosen, maximizes the group’s sum of payoffs. This coincides with the median starting point. If a subject cares about efficiency, she would be expected to put some weight on this point in her proposal. The second variable is the *inequality minimizer*, which corresponds to the point that, if chosen, minimizes within-group inequality as measured by the Gini coefficient. This happens at the mid-point between the two more extreme starting points in each group. It is also the *maxmin* of each group –the point where the lowest payoff is maximized. Overall, a positive weight on its coefficient should capture subjects’ concerns for inequality. Given that we find some differences in part A outcomes with respect to inequality, we want to examine whether they can explain the treatment effects we find in part B.

In column (2), we report the regression results. The coefficient for the *efficiency maximizer* is essentially zero, leading us to conclude that subjects are not concerned about efficiency. The coefficient

¹⁸All results reported here and in subsequent regressions are robust to clustering errors at the session or subject level.

of the *inequality minimizer* is positive and slightly significant. Nevertheless, all other coefficients remain highly significant and at essentially the same magnitude as in column (1). We conclude that there is some degree of inequality aversion among the subjects, but this can in no way explain the large treatment effect that we find.

TABLE 3: REGRESSION RESULTS

Dependent variable: *proposals*

	OLS		2SLS	Reduced
	(1)	(2)	2nd stage	form
	(1)	(2)	(3)	(4)
Constant	.965 (1.069)	-.649 (1.267)	.125 (1.501)	2.901* (1.450)
Starting point	.729** (.013)	.722** (.013)	.720** (.014)	.726** (.049)
Part A result	.261** (.024)	.235** (.030)	.288** (.034)	
Nash				.223** (.057)
Efficiency maximizer		-.0002 (.018)		
Inequality minimizer		.063* (.025)		
Starting point × Treatment	.167** (.018)	.164** (.018)	.179** (.019)	.173** (.058)
Part A result × Treatment	-.170** (.020)	-.167** (.020)	-.186** (.022)	
Nash × Treatment				-.174** (.060)
Observations	1800	1800	1800	1800

Notes: *efficiency maximizer* is the point that maximizes the sum of payoffs for the group. *Inequality minimizer* is the point that minimizes inequality as measured by the Gini coefficient. Treatment is a dummy variable that takes value 1 for treatment *SM*.
 ** : $p - val < .01$
 * : $p - val < .05$

Next, we instrument the outcome of part A by the means of the Nash equilibrium prediction. Formally, since the part A result appears in the basic specification of column (1) both alone and in an interaction with the treatment we need to utilize two instruments: the Nash equilibrium prediction and its interaction with the treatment dummy. We present the second stage of the 2SLS estimation process in column (3) which gives us largely the same results as our benchmark specification.¹⁹ The additional insight provided by this exercise is that the large differences in the persistence of the group outcome across treatments cannot be attributed to any outcome-related differences. Indeed, column (3) shows that even if one focuses only on the part of the group-decision of part A that is explained by the Nash equilibrium prediction, one finds similar results to our benchmark specification.

Finally, in column (4), we present the reduced form of this two-stage approach: We report results from a linear regression that explains subjects' proposals as a convex combination of their starting points and the Nash equilibrium prediction, and the interaction of these variables with the treatment dummy. These results are particularly important, given that the part A outcome similarity between the *MA* and *SM* treatments is established in a stochastic manner. That is, despite the demonstrated affinity of outcome distributions –even when one controls for the exact preference profile– this coincidence is not deterministic: Two groups with identical preferences hardly ever arrive at exactly the same outcome, both across– and within–treatments. The reduced form results of column (4) establish that the desired outcome –the Nash equilibrium one– is a stronger predictor of the subjects' proposals in *MA* than in *SM*, hence reassuring us that the smaller weight that subjects assign to their starting point in *MA* compared to *SM* are not driven by outcome-related differences. It should be noted that these last findings carry an independent interpretation and broader implications as far as implementation of welfare optima is concerned. When a mechanism designer wants to implement a certain welfare optimum –in our case, this is the median of the set that contains the subjects' starting points plus points 33 and 67– and expects that after the voting procedure individuals will try to revise the outcome to their liking, then the mechanism designer should opt for a congruent mechanism, as it will enhance the probability that the post-revisions outcome will be as close as possible to the desired policy alternative.

¹⁹As is standard in the literature (see Angrist and Pischke, 2008) when we have more than one endogenous regressors –and, hence, we utilize two (or more) exogenous instruments– the relevant statistic for the first stage regression becomes the Cragg-Donald Minimum Eigenvalue Vector statistic. In our case this statistic takes a value of about 700, which is well above the critical threshold of 7.

6 Explaining the treatment effect

We have identified a strong treatment effect concerning outcome stability: Subjects are less inclined to overturn the group’s outcome in favor of their preferred choice in treatment *MA*, which used a congruent mechanism. A natural question to ask is which features of the collective choice mechanism drive this effect? As discussed, based on our experimental design and econometric analysis, we can exclude explanations that relate to the qualities of the outcomes produced by the two mechanisms. We therefore look closer at the effect of subjects’ agreement with the outcome, as expressed through their votes. To that effect, we use the measures introduced in section 5.1.3.

In the first column of Table 4, we report results from a regression that explains the deviation of subjects’ proposals from the part A outcomes as a function of the distance between the part A results and their starting points and the interaction of this variable with the treatment dummy.²⁰ If subjects were behaving in a simple utility-maximizing manner, there would be no treatment effect (in other words, the coefficient of the interaction would be insignificant, and the dependent variable would be fully explained by our first independent variable). In line with the results presented in the previous section, we find a strong treatment effect: The coefficient of the interaction is positive and significant, indicating that proposals are farther away from the part A result and closer to the subjects’ starting points in the *SM* treatment compared to the more moderate changes in the outcome observed in the *MA* treatment.

Next, we introduce our measure of *individual agreement* in to the regression (column 2). Recall that this measures a subject’s agreement with the group outcome, as expressed through her vote. As we observe, the treatment effect vanishes. Indeed, *individual agreement* seems to fully pick up the differences observed across the two treatments. As seen in Figure 4, *individual agreement* differs significantly across treatments. One worry here might be that this variable is correlated with some unobserved factor that differs across treatments and the effect captured in this regression is unrelated to *individual agreement*. To address this, we run regressions restricting the sample to each treatment separately (columns 3 and 4) and obtain very similar results.²¹

²⁰Due to the nature of our dependent variable, the standard choice here is the Poisson regression, but similar results can be obtained if one applies a negative binomial estimator instead. For all regressions of Table 4 tests of goodness of fit of the Poisson model were performed, indicating that our selection was the appropriate one.

²¹Results are robust to constructing the measures of agreement using *average* instead of *minimum* distance of vote to outcome (see section 5.1.3). The only difference is that, while with minimum distance the treatment effects become unambiguously insignificant (see, column 2), using average distance the interaction ceases to be significant at our chosen 1% level, but remains significant at the 5% level.

If we use *group agreement* (column 5) we observe that, while these measures have a small part in explaining the degree of deviations of the proposals from the original group’s choice both across- and within-treatments, the direct treatment effects remain significant.²² In fact, one can break down *group agreement* into two components: *individual agreement* and the agreement of other group members with the outcome. When these variables are introduced separately in to the regression, both the treatment effect and *other’s agreement* are not significant, while *individual agreement* remains highly significant.

In summary, we find that individuals who voted for alternatives close to the implemented outcomes exhibit a higher degree of commitment to policies near the group’s decision than similar individuals who voted for alternatives far from the outcome. The degree to which other group members agree with the implemented outcome does not seem to be relevant. These effects can explain both across- and within-treatment variation. In other words, it seems that when presented with the opportunity to revise a collective choice, individuals exhibit a tendency to remaining consistent with previously expressed preferences, irrespective of the mechanism used. Since congruent mechanisms incentivize individuals to vote for the outcome that is eventually implemented, these outcomes are less likely to be revised.

7 Concluding remarks

Our experimental approach aims to test whether procedures matter beyond consequences. After the first part of our experiment –in which a collective choice is reached through two mechanisms– the second part of the experiment tests how much the voters respect the collective decision reached in the first part of the protocol. The behavior we find in the second part of our experiment is at odds with self-regarding rational choice. Perhaps more surprising though, is the fact that *other-regarding* preferences have very little part in explaining the treatment effects we find. Frey et al. (2004), for instance, propose that procedural utility is derived from the allocative and redistributive properties of a mechanism, or from how one is treated in interaction with others. In the neutral context used in the lab, subjects remained *self-regarding*, even if not rational. The differences across treatments we observe are more in line with some preference for consistency, or aversion to cognitive dissonance (we refer the reader to Kamenica, 2012 for a discussion of these concepts). Of course, we do not preclude a role for other-regarding preferences in other settings, especially outside of the lab environment. But our findings suggest that these may not be the unique source of procedural utility. Further research in this direction is warranted.

²²The within-treatment results are not presented here to save space and are available by the authors upon request.

TABLE 4: POISSON REGRESSION RESULTS

Dependent variable: *proposal deviation from part A result*

	Full Sample		Treatment MA	Treatment SM	Full Sample	
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	1.958** (0.011)	2.279** (0.030)	2.474** (0.099)	2.63** (0.032)	2.273** (0.039)	2.278** (0.039)
Deviation from starting point	0.034** (.0003)	0.035** (.0004)	0.034** (.0005)	0.035** (.0004)	0.035** (.0004)	0.035** (.0004)
Starting point × Treatment	.004** (.0002)	.0003 (.0004)			.0009* (.0004)	.0003 (.0004)
Individual agreement		−.004** (.0003)	−.0056** (.001)	−.0037** (.0004)		−.0038** (.0004)
Group agreement					−.004** (.0005)	
Other’s agreement						.00002 (.0005)
Observations	1800	1800	900	900	1800	1800

Notes: The variable *Other’s agreement* is calculated as *Group agreement*, but ignoring one’s own agreement. Treatment is a dummy variable that takes value 1 for treatment *SM*.

**: $p - val < .01$

*: $p - val < .05$

In fact, whether our results are evidence of procedural utility –the mechanism affects preferences directly and hence welfare– or of a behavioral bias in decision making –akin to framing or priming effects– is up to interpretation. Either way, they suggest an important role for agreement between votes and outcomes in the persistence of collective decisions. A designer of choice rules that desires to enhance the persistence of social choice should take this in to account and select mechanisms that promote such agreement over ones that do not.

References

Angrist, Joshua D and Jörn-Steffen Pischke (2008). *Mostly harmless econometrics: An empiricist’s companion*. Princeton university press.

- Corazzini, Luca, Kube, Sebastian, Maréchal, Michel André, and Antonio Nicolo (2014). “Elections and deceptions: an experimental study on the behavioral effects of democracy”. In: *American Journal of Political Science* 58.3, pp. 579–592.
- Dal Bó, Pedro, Andrew Foster, and Louis Putterman (2010). “Institutions and behavior: Experimental evidence on the effects of democracy”. In: *American Economic Review* 100.5, pp. 2205–29.
- De Dreu, Carsten KW and Laurie R Weingart (2003). “Task versus relationship conflict, team performance, and team member satisfaction: a meta-analysis.” In: *Journal of applied Psychology* 88.4, p. 741.
- De Wit, Frank RC, Lindred L Greer, and Karen A Jehn (2012). “The paradox of intragroup conflict: a meta-analysis.” In: *Journal of Applied Psychology* 97.2, p. 360.
- Fischbacher, Urs and Franziska Föllmi-Heusi (2013). “Lies in disguise? an experimental study on cheating”. In: *Journal of the European Economic Association* 11.3, pp. 525–547.
- Frey, Bruno S and Alois Stutzer (2005). “Beyond outcomes: measuring procedural utility”. In: *Oxford Economic Papers* 57.1, pp. 90–111.
- Frey, Bruno S, Matthias Benz, and Alois Stutzer (2004). “Introducing procedural utility: Not only what, but also how matters”. In: *Journal of Institutional and Theoretical Economics JITE* 160.3, pp. 377–401.
- Galton, Francis (1907). “One vote, one value”. In: *Nature* 75.1948, p. 414.
- Gerardi, Dino and Leeat Yariv (2007). “Deliberative voting”. In: *Journal of Economic Theory* 134.1, pp. 317–338.
- Gershkov, Alex, Benny Moldovanu, and Xianwen Shi (2017). “Optimal voting rules”. In: *The Review of Economic Studies* 84.2, pp. 688–717.
- Gneezy, Uri (2005). “Deception: The role of consequences”. In: *American Economic Review* 95.1, pp. 384–394.
- Goeree, Jacob K and Leeat Yariv (2011). “An experimental study of collective deliberation”. In: *Econometrica* 79.3, pp. 893–921.
- Grossman, Guy and Delia Baldassarri (2012). “The Impact of Elections on Cooperation: Evidence from a Lab-in-the-Field Experiment in Uganda”. In: *American Journal of Political Science* 56.4, pp. 964–985. ISSN: 1540-5907.
- Hartnett, Tim (2011). *Consensus-oriented decision-Making: the CODM model for facilitating groups to widespread agreement*. new society publishers.
- Jehn, Karen A (1994). “Enhancing effectiveness: An investigation of advantages and disadvantages of value-based intragroup conflict”. In: *International journal of conflict management* 5.3, pp. 223–238.
- Jehn, Karen A and Elizabeth A Mannix (2001). “The dynamic nature of conflict: A longitudinal study of intragroup conflict and group performance”. In: *Academy of management journal* 44.2, pp. 238–251.
- Jehn, Karen A, Clint Chadwick, and Sherry MB Thatcher (1997). “To agree or not to agree: The effects of value congruence, individual demographic dissimilarity, and conflict on workgroup outcomes”. In: *International journal of conflict management* 8.4, pp. 287–305.

- Kamei, Kenju (2016). “Democracy and resilient pro-social behavioral change: an experimental study”. In: *Social Choice and Welfare* 47.2, pp. 359–378.
- Kamei, Kenju, Louis Putterman, and Jean-Robert Tyran (2015). “State or nature? Endogenous formal versus informal sanctions in the voluntary provision of public goods”. In: *Experimental Economics* 18.1, pp. 38–65.
- Kamenica, Emir (2012). “Behavioral Economics and Psychology of Incentives”. In: *Annual Review of Economics* 4.1, pp. 427–452.
- Lind, E Allan and Tom R Tyler (1988). *The social psychology of procedural justice*. Springer Science & Business Media.
- Markussen, Thomas, Louis Putterman, and Jean-Robert Tyran (2013). “Self-organization for collective action: An experimental study of voting on sanction regimes”. In: *Review of Economic Studies* 81.1, pp. 301–324.
- Markussen, Thomas, Ernesto Reuben, and Jean-Robert Tyran (2014). “Competition, Cooperation and Collective Choice”. In: *The Economic Journal* 124.574, F163–F195. ISSN: 1468-0297.
- Maskin, Eric S (2008). “Mechanism design: How to implement social goals”. In: *American Economic Review* 98.3, pp. 567–76.
- Mason, Richard O. and Ian I. Mitroff (1981). *Challenging strategic planning assumptions: Theory, cases, and techniques*. John Wiley & Sons Inc.
- Mazar, Nina, On Amir, and Dan Ariely (2008). “The dishonesty of honest people: A theory of self-concept maintenance”. In: *Journal of marketing research* 45.6, pp. 633–644.
- Mellizo, Philip, Jeffrey Carpenter, and Peter Hans Matthews (2014). “Workplace democracy in the lab”. In: *Industrial Relations Journal* 45.4, pp. 313–328. ISSN: 1468-2338.
- Moulin, Herve (1980). “On Strategy-proofness and Single Peakedness”. In: *Public Choice* 35, pp. 437–455.
- Núñez, Matias and Dimitrios Xeferis (2017). “Implementation via approval mechanisms”. In: *Journal of Economic Theory* 170, pp. 169–181.
- Ockenfels, Axel and Alvin E Roth (2006). “Late and multiple bidding in second price Internet auctions: Theory and evidence concerning different rules for ending an auction”. In: *Games and Economic behavior* 55.2, pp. 297–320.
- Priem, Richard L., David A. Harrison, and Nan Kanoff Muir (1995). “Structured conflict and consensus outcomes in group decision making”. In: *Journal of management* 21.4, pp. 691–710.
- Renault, Régis and Alain Trannoy (2005). “Protecting minorities through the average voting rule”. In: *Journal of Public Economic Theory* 7.2, pp. 169–199.
- Schweiger, David M., William R. Sandberg, and James W. Ragan (1986). “Group approaches for improving strategic decision making: A comparative analysis of dialectical inquiry, devil’s advocacy, and consensus”. In: *Academy of management Journal* 29.1, pp. 51–71.

- Schweiger, David M, William R Sandberg, and Paula L Rechner (1989). “Experiential effects of dialectical inquiry, devil’s advocacy and consensus approaches to strategic decision making”. In: *Academy of Management Journal* 32.4, pp. 745–772.
- Sen, Amartya (1997). “Maximization and the Act of Choice”. In: *Econometrica: Journal of the Econometric Society*, pp. 745–779.
- Sutter, Matthias, Stefan Haigner, and Martin G Kocher (2010). “Choosing the carrot or the stick? Endogenous institutional choice in social dilemma situations”. In: *The Review of Economic Studies* 77.4, pp. 1540–1566.
- Thibaut, John W and Laurens Walker (1975). *Procedural justice: A psychological analysis*. L. Erlbaum Associates.
- Tyran, Jean-Robert and Lars P Feld (2006). “Achieving compliance when legal sanctions are non-deterrent”. In: *The Scandinavian Journal of Economics* 108.1, pp. 135–156.
- Walker, James M, Roy Gardner, Andrew Herr, and Elinor Ostrom (2000). “Collective choice in the commons: Experimental results on proposed allocation rules and votes”. In: *The Economic Journal* 110.460, pp. 212–234.
- Yamamura, H. and R. Kawasaki (2013). “Generalized Average Rules as stable Nash mechanisms to implement generalized median rules”. In: *Social Choice and Welfare* 40, pp. 815–832.

A Instructions

The experiment was run in Greek. We present here a translation of the instructions done by the authors. Original instructions in Greek are available upon request.

A.1 Treatment *SM*

Thank you for participating in this session. Please remain quiet. The experimental session will be run using a computer and all answers will be given through it. Please do not talk to each other and keep quiet during the session. Please note that the use of mobile phones and other electronic devices is not permitted. Please read the instructions carefully, and if you have any questions, raise your hand. The answer that will be given will be announced to everyone.

General Instructions

During the experiment, you can win points. The points will be converted into euros. **1 euro = 15 points**. Each participant will receive a payment. The exact amount you will receive depends on the

decisions you will make during the experiment, the decisions of other participants and also on luck. In addition, you will receive the amount of **€3 as a show-up fee**. Following the completion of the experimental session, a fee will be paid privately in cash to each one of you. The experiment consists of two parts. The instructions below are for part A of the experiment. Following the completion of part A, the instructions of part B will be given. Your final earnings will be:

€3 show-up fee + earnings in part A + earnings in part B

Part A

Aim

Part A of the experiment consists of 20 periods. In each period, you will be in a three-member group with two other participants. The aim of the group is to choose a common **destination** from 100 consecutive locations (that is, an integer from 1 to 100) which will be the final decision of the group in the end of each period. The composition of the groups will change in every single period, and you will not be able to know the identity of the members of the group. The way the destination is chosen from the group will be explained below. First, we will explain the way in which the payoffs of each player are determined.

Starting points and payoffs

In each period, a specific destination (that is, an integer from 1 to 100) will be chosen as an **individual starting point**. The payoffs in each period depend on the distance between the final destination that will be chosen and your individual starting point: **The farther the destination is from the starting point of each player, the smaller his/her payoffs will be**. Specifically, each player's payoff (in points) will be calculated as follows:

$$Profits = 100 - |destination - starting point|$$

Example:

The group chose the location 49 as a destination.

Player 1's starting point is 20. Player 1's payoff is **71 points**. (The distance between the starting point of player 1 and the final (common) destination is $49 - 20 = 29$. So, the payoff is $100 - 29 = 71$).

Player 2's starting point is 50. Player 2's payoff is **99 points**. (The distance between the starting point of player 2 and the final (common) destination is $50 - 49 = 1$. So, the payoff is $100 - 1 = 99$)

Player 3's starting point is 95. Player 3's payoff is **54 points**. (The distance between the starting point of player 1 and the final (common) destination is $95 - 49 = 46$. So, the payoff is $100 - 46 = 54$)

The calculations above will be conducted automatically from the computer, and on the screen you will see your starting point, your group members' starting points, and their payoffs, depending on the chosen destination.




Figure 7

Attention!

- The starting point of each player will be different (unique).
- In each period, the starting points will change.
- All players' starting points will be shown in the screen with arrows.
- Each player's payoff will be indicated by a bar. The greater the payoff, the taller the bar.

Selection

The selection of the destination will be done as follows: Each group member can vote for exactly one location. The final destination will be the average of all locations voted for by the group members. (In case of a non-integer average, the final destination will be calculated by rounding to the nearest integer).

Example 1:

Player 1 votes 5.

Player 2 votes 80.

Player 3 votes 95.

The final destination will be the location 60 because the average of the chosen locations is: $\frac{5 + 80 + 95}{3} = 60$




Figure 8

Example 2:

Player 1 votes 30.

Player 2 votes 80.

Player 3 votes 95.

The final destination will be the location 68 because the average of the chosen locations is: $\frac{30 + 80 + 95}{3} \approx 68,33$

Voting Procedure

Every single period, the voting procedure will last $60+x$ seconds, where x is a random number from 1 to 10. In other words, following the completion of the voting, the procedure will stop randomly in one of the next 10 seconds.




Figure 9

You can specify the location you are voting by clicking on the white frame you will see on your screen.




Figure 10

At the same time, you will see what other group members are voting for and how the common destination is shaped. You can change your vote as many times as you want until the voting procedure is over.

The destination of each period **will be determined after the completion of the voting procedure**. Hence, make sure you have made your choice before the end of the 60-second period. In the first two periods, the duration will be $90+x$ seconds, so as to allow you plenty of time to get used to the procedure. The remaining time will be shown at the bottom of your screen.

At the end of the experiment, **one period from part A will be selected randomly and your payment will be based on your earnings in this period**. Hence, we encourage you to pay attention to all your decisions in all periods, since each of them can determine your final payment.

Before we start, there will be two trial periods to make sure everything is understood. These two trial periods cannot be chosen, and your decisions in those periods will not affect your payment.

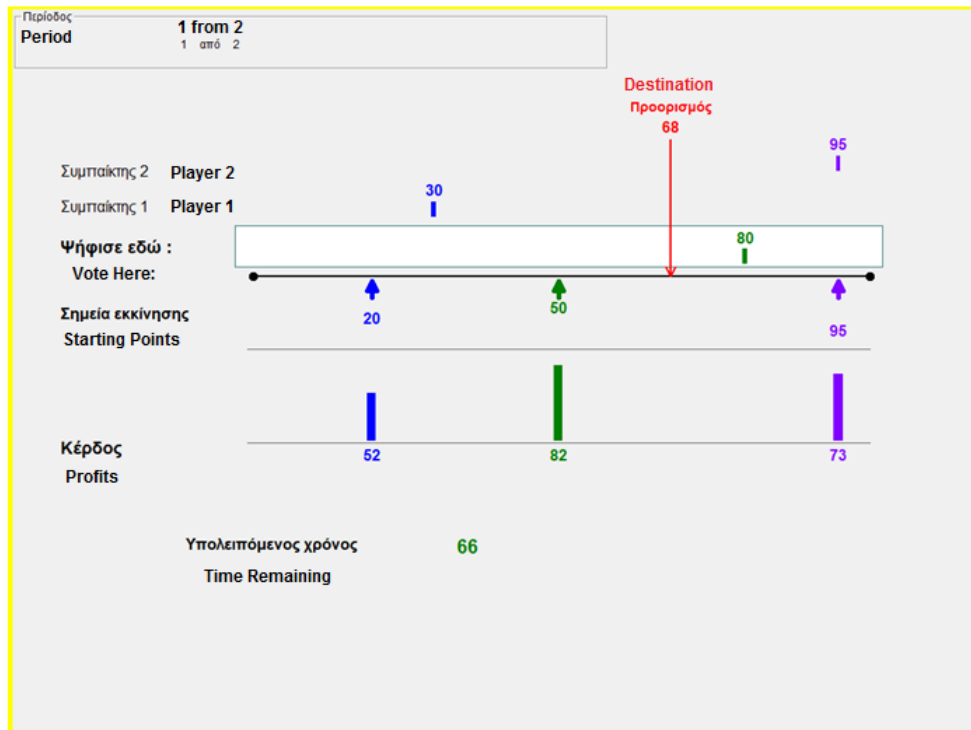


Figure 11

Part B

This part again consists of 20 periods.

In each period you will be placed in the same group you were in the corresponding period of part A, with the same individual starting points.

The aim of the group is to make a collective decision regarding your common destination.

At the top of your screen, you will see a figure in which the starting points, the votes of your teammates and the final chosen destination of the corresponding period of part A will be displayed.

This time the final destination will be determined according to a new procedure. Click in the white frame at the middle of the screen to propose a destination. Your proposal will appear as well as each group member's payoff if your proposal is selected. **You can change your proposal as many times as you like until you press the red button 'Submit'**. When you press the button, your submission will be confirmed, and you will move to the next period. In this part, the proposals of your teammates

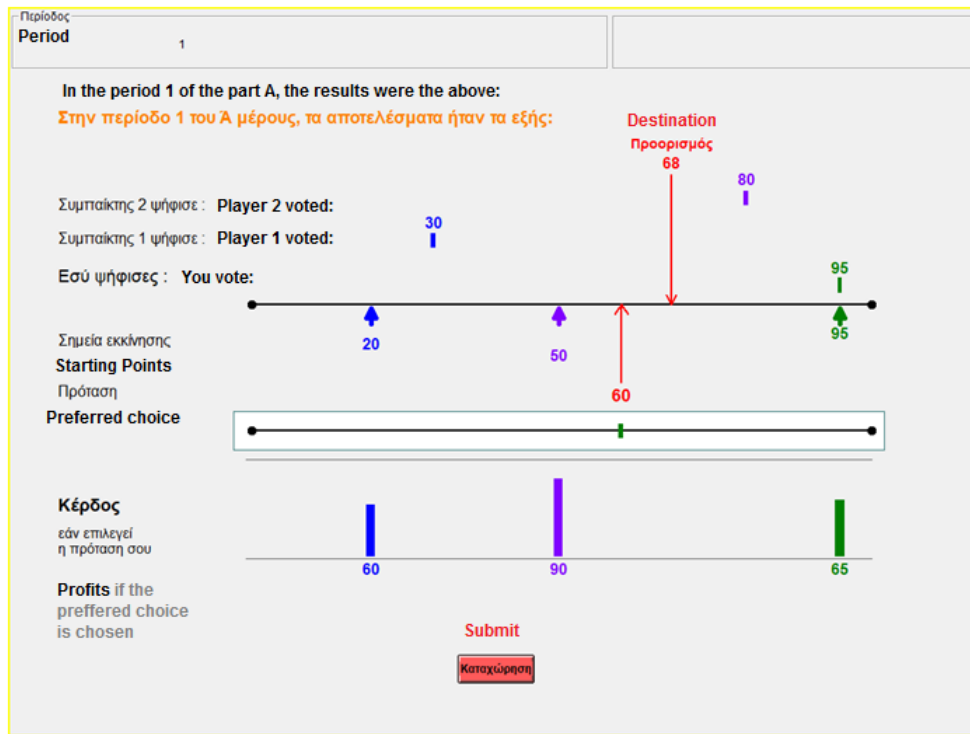


Figure 12

will be unknown to you until the end of the experiment.

One of the three proposals made by the members of each group will be chosen randomly and become the new common destination for this period, according to which group members' payoffs will be determined.

The profits will be calculated in the same way as in part A by taking into consideration the new destination.

At the end of the experiment, one period from part B will be selected randomly and your payment will be based on your earnings in this period. Hence, we encourage you to pay attention to all your decisions in all periods, since each of them can determine your final payment.

Following the completion of the experiment, you will see on your screen the chosen periods for the calculation of your profit, as well as your final profit.

B Treatment *MA*

Thank you for participating in this session. Please remain quiet. The experimental session will be run using a computer, and all answers will be given through it. Please do not talk to each other and keep quiet during the session. Please note that the use of mobile phones and other electronic devices is not permitted. Please read the instructions carefully and if you have any questions, raise your hand. The answer that will be given will be announced to everyone.

General Instructions

During the experiment, you can win points. The points will be converted into euros. **1 euro = 15 points**. Each participant will receive a payment. The exact amount you will receive depends on the decisions you will make during the experiment, the decisions of other participants, and also on luck. In addition, you will receive the amount of **€3 as a show-up fee**. Following the completion of the experimental session, a fee will be paid privately in cash to each one of you. The experiment consists of two parts. The instructions below are for part A of the experiment. Following the completion of part A, the instructions of part B will be given. Your final earnings will be:

€3 show-up fee + earnings in part A + earnings in part B

Part A

Aim

Part A of the experiment consists of 20 periods. In each period, you will be in a three-member group with two other participants. The aim of the group is to choose a common **destination** from 100 consecutive locations (that is, an integer from 1 to 100) which will be the final decision of the group in the end of each period. The composition of the groups will change in every single period, and you will not be able to know the identity of the members of the group. The way the destination is chosen from the group will be explained below. First, we will explain the way in which the payoffs of each player are determined.

Starting points and payoffs

In each period, a specific destination (that is, an integer from 1 to 100) will be chosen as an **individual starting point**. The payoffs in each period depend on the distance between the final destination that will be chosen and your individual starting point: **The farther the destination is from the starting point of each player, the smaller his/her payoffs will be**. Specifically, each player's payoff (in points) will be calculated as follows:

$$\text{Profits} = 100 - |\text{destination} - \text{starting point}|$$

Example:

The group chose the location 49 as a destination.

Player 1's starting point is 20. Player 1's payoff is **71 points**. (The distance between the starting point of player 1 and the final (common) destination is $49 - 20 = 29$. So, the payoff is $100 - 29 = 71$).

Player 2's starting point is 50. Player 2's payoff is **99 points**. (The distance between the starting point of player 2 and the final (common) destination is $50 - 49 = 1$. So, the payoff is $100 - 1 = 99$)

Player 3's starting point is 95. Player 3's payoff is **54 points**. (The distance between the starting point of player 1 and the final (common) destination is $95 - 49 = 46$. So, the payoff is $100 - 46 = 54$)

The calculations above will be conducted automatically from the computer, and on the screen you will see your starting point, your group members' starting points, and their payoffs, depending on the chosen destination.




Figure 13

Attention!

- The starting point of each player will be different (unique).

- In each period, the starting points will change.
- All players' starting points will be shown in the screen with arrows.
- Each player's payoff will be indicated by a bar. The greater the payoff, the taller the bar.

Selection

The selection of the destination will be done as follows: Each group member can vote up to 100 locations. These locations should be consecutive (e.g., someone can vote from 23 to 56). The destination will be the largest median of the distribution of the votes. That is, destination X will be selected if at least half of the votes have been given to locations to the left of $X + 1$ and at least half of the votes have been given to locations to the right of $X-1$. If there are more than two locations with this feature, the largest one will be selected.

Example 1:

Player 1 votes from 1 to 20 (20 votes).

Player 2 votes from 71 to 75 (5 votes).

Player 3 votes from 91 to 100 (10 votes).

The sum of the votes is 35.

Every single location from 1 to 20, from 71 to 75, and from 91 to 100 has been voted for once. The rest of the locations have not been voted for.

The destination will be location 18 (because the votes that have been given to locations to the left of $18 + 1$ are $18 > 35/2 = 17.5$, and the votes that have been given to locations to the right of $18-1$ are $18 > 35/2 = 17.5$).

Example 2:

Player 1 votes from 1 to 20 (20 votes).

Player 2 votes from 71 to 75 (5 votes).

Player 3 votes from 71 to 100 (30 votes).




Figure 14

The sum of the votes is 55.

Every single location from 71 to 75 has been voted for twice (by player 2 and player 3).

Every single location from 1 to 20 and from 76 to 100 has been voted for once.

The rest of the locations have not been voted for.

The destination will be location 74 (because the votes that have been given to locations to the left of $74 + 1$ are $28 > 55/2 = 27.5$, and the votes that have been given to locations to the right of $74 - 1$ are $28 > 55/2 = 27.5$).




Figure 15

Voting Procedure

Every single period, the voting procedure will last $60+x$ seconds, where x is a random number from 1 to 10. In other words, following the completion of the voting, the procedure will stop randomly in one of the next 10 seconds.

You can specify the locations you are voting for by clicking and by moving the green bars that you will see on your screen.

At the same time, you will see what your teammates are voting for and how the common destination is shaped. You can change your vote as many times as you want until the voting procedure is over.

The destination of each period **will be determined after the completion of the voting procedure**. Hence, make sure you have made your choice before the end of the 60-second period. In the first two periods, the duration will be 90+x seconds, so as to allow you plenty of time to get used to the procedure. The remaining time will be shown at the bottom of your screen.

At the end of the experiment, **one period from part A will be selected randomly, and your payment will be based on your earnings in this period**. Hence, we encourage you to pay attention to all your decisions in all periods, since each of them can determine your final payment.

Before we start, there will be two trial periods to make sure everything is understood. These two trial periods cannot be chosen, and your decisions in those periods will not affect your payment.

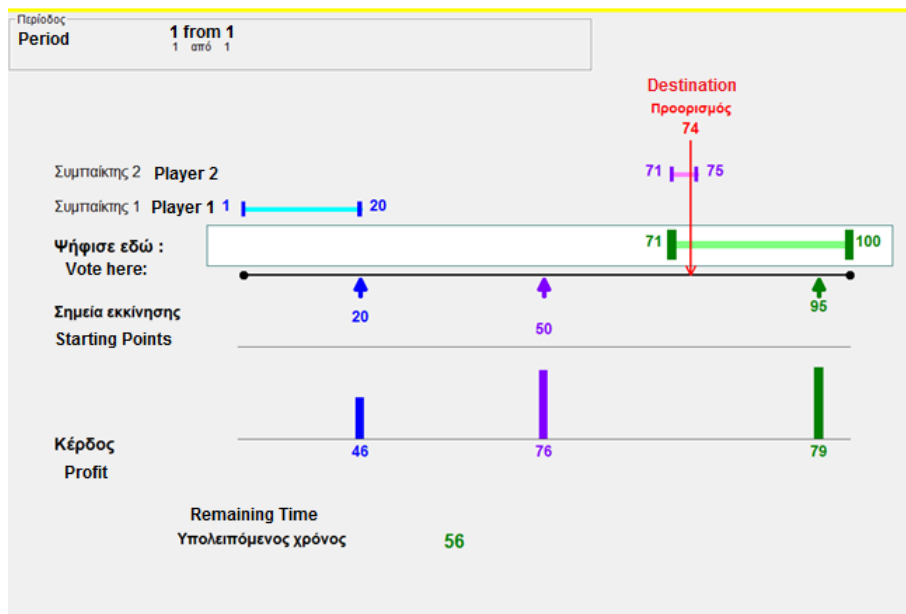


Figure 16

Part B

This part, again, consists of 20 periods.

In each period you will be placed in the same group you were in the corresponding period of part A, with the same individual starting points.

The aim of the group is to make a collective decision regarding your common destination.

At the top of your screen, you will see a figure in which the starting points, the votes of your teammates, and the final chosen destination of the corresponding period of part A will be displayed.

This time the final destination will be determined according to a new procedure. Click on the white frame at the middle of the screen to propose a destination. Your proposal will appear, as well as each group member's payoff if your proposal is selected. **You can change your proposal as many times as you like until you press the red button 'Submit'**. When you press the button, your submission will be confirmed and you will move to the next period. In this part, the proposals of your teammates will be unknown to you until the end of the experiment.

One of the three proposals made by the members of each group will be chosen randomly and become the new common destination for this period, according to which group members' payoffs will be determined.

The profits will be calculated in the same way as in part A by taking into consideration the new destination.

At the end of the experiment, one period from part B will be selected randomly, and your payment will be based on your earnings in this period. Hence, we encourage you to pay attention to all your decisions in all periods, since each of them can determine your final payment.

Following the completion of the experiment, you will see on your screen the chosen periods for the

calculation of your profit, as well as your final profit.

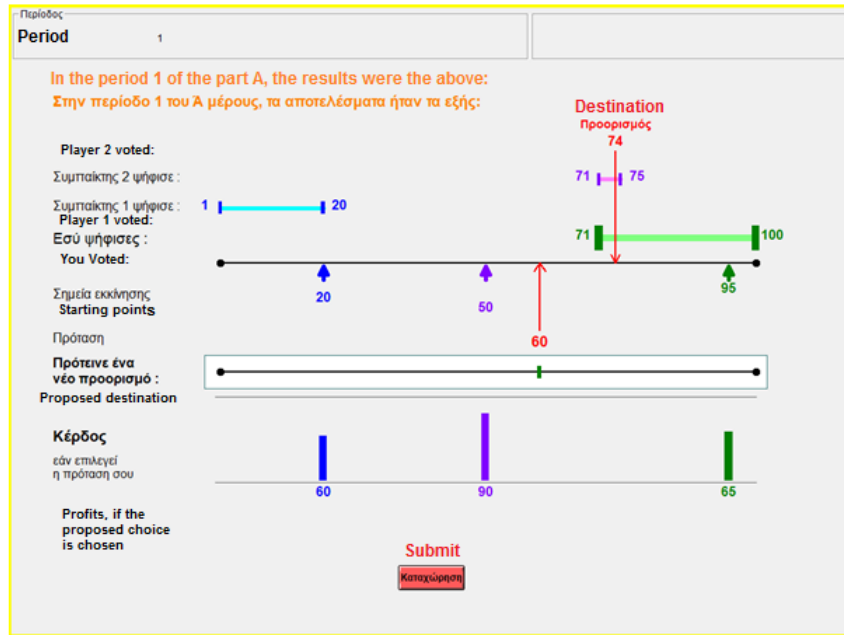


Figure 17