

Unanimous Rules in the Laboratory*

Laurent Bouton	Aniol Llorente-Saguer	Frédéric Malherbe
Georgetown University,	Queen Mary	London Business School
Université libre de Bruxelles,	University of London	and CEPR
CEPR and NBER	and CEPR	

December 5, 2016

Abstract

We study the information aggregation properties of unanimous voting rules in the laboratory. In line with theoretical predictions, we find that majority rule with veto power dominates unanimity rule. We also find that the strategic voting model is a fairly good predictor of subject behavior. Finally, we exploit a framing effect to study how the presence of less sophisticated agents affects Veto's welfare properties.

JEL Classification: C92, D70

Keywords: Unanimity Rule, Veto Power, Constructive Abstention, Information Aggregation, Laboratory Experiments, Framing

*We thank participants of the ESA meetings in Tucson and Cologne, the Workshop on Political Economy and Voting Experiments in Granada, and seminar participants at the Berlin Colloquium in Behavioral Economics, Boston University, City University London, King's College London, London Behavioural and Experimental Group, Max Planck Institute, Paris School of Economics, University of Copenhagen, and University of Vienna. We extend particular thanks to David Austen-Smith, Jean-Pierre Benoit, Micael Castanheira, David Myatt, Nikos Nikiforakis, and Tom Palfrey. We would also like to thank Erica Gross, Lucas Meier, and Nico Meier for their excellent assistance with running the experiments. We gratefully acknowledge financial support from the London Business School (RAMD8871).

1 Introduction

In many sensitive situations, group decisions are required to be unanimous. Examples include a number of international organizations that would not exist without granting some sort of veto power to their members.¹ They also include partnerships and other unlimited liability companies, and criminal trials by jury in the US. The central question in this paper is: what voting system is best in such situations?

When agents have no uncertainty about their preferred alternative, all unanimous rules are equivalent –a proposal to reform the status quo is only accepted if it is Pareto improving (Wicksell 1967 [1896] and Buchanan and Tullock 1962). Unanimous rules are, however, not equivalent when agents are uncertain about the merits of a proposal and share common objectives. This is because voting then ought to aggregate the information dispersed among agents. The problem is that unanimous decision making is believed to aggregate information poorly (Feddersen and Pesendorfer 1998, Guarnaschelli, McKelvey, and Palfrey 2000).² This raises the question of whether a group necessarily sacrifices information aggregation when it grants veto power to its members.

In this paper, we compare the performance, in the laboratory, of two of the most widely used unanimous rules: *unanimity rule* and *majority rule with veto power* (henceforth Unanimity and Veto).³ Under Unanimity, agents must consent or dissent. The reform is then adopted if and only if no one dissents. Under Veto, agents can consent, dissent, or veto. The proposal is then accepted provided that no one vetoes and a (simple or qualified) majority consents. The main difference is that under Unanimity agents cannot convey negative information about the reform without blocking it altogether. The intense debate during the early years of the United Nations Security Council on the impossibility of dissenting without vetoing illustrates that this difference is far from innocuous (Sievers and Daws 2014).

And indeed, we find that, in contrast to Unanimity, Veto consistently aggregates infor-

¹See, e.g., Zamora (1980), Posner and Sykes (2014), and Maggi and Morelli (2006); and Bouton, Llorente-Saguer, and Malherbe (2016) for a discussion.

²Coughlan (2001), Duggan and Martinelli (2001), Persico (2004), and Bouton, Llorente-Saguer, and Malherbe (2016), however, highlight cases where unanimous decision making features good information aggregation properties.

³Among international organizations Unanimity is used, e.g., by the North Atlantic Treaty Organization (NATO), the European Council (for most sensitive topics, excluding Common Foreign and Security Policy), and the Southern Common Market (Mercosur). In contrast, Veto (or a close variation) is used, e.g., by the European Council (for the Common Foreign and Security Policy), the United Nations Security Council.

mation well in the laboratory. Hence, our findings provide empirical support to our previous theoretical result that Veto Pareto dominates Unanimity (Bouton, Llorente-Saguer, and Malherbe 2016). This provides a rationale for the use of Veto in practice and sheds light on the evolution of decision-making practices in the United Nations Security Council and the Council of the European Union. It also suggests that it would be beneficial for voting bodies that currently use Unanimity to adopt Veto instead.

Our experiment design follows the typical setup considered in the information aggregation voting literature. There are two possible states of the world (Red or Blue). Agents observe a binary private signal (red or blue) that is correlated with the realized state. They have a common objective: they are all rewarded if the group decision (Red or Blue) matches the state (decision Red represents the status quo). To make the group decision, they hold a simultaneous vote according to a pre-specified voting system: Unanimity or Veto.

Theoretically, the welfare performance of these voting rules depends on the information structure. To understand this idea, note that under both rules, any single agent can enforce the status quo. If the red signal is sufficiently informative relative to the blue signal, enforcing the status quo when observing a red signal is a weakly dominant strategy. In this case, where the red signal is *decisive*, information aggregation is relatively straightforward and both Veto and Unanimity are efficient. When the red signal is *not decisive*, however, information aggregation is a more subtle problem and Veto outperforms Unanimity because it offers the possibility of revealing a negative signal without pinning down the outcome.

We consider both cases in the laboratory. First, in the case where a red signal is not decisive, we use equally informative signals as in Feddersen and Pesendorfer (1998) and Guarnaschelli, McKelvey and Palfrey (2000). We find that groups using Veto make about a third the number of mistakes as those who use Unanimity. This difference is due to a dramatic reduction of type II errors. That is, using Veto makes it much less likely that agents will reject a good reform (or in the typical jury interpretation, acquit a guilty defendant). In the case where a red signal is decisive, we find that performances under Veto and Unanimity do not differ significantly. Our data therefore provides strong empirical support for the theoretical predictions.

We then analyze subject behavior in detail. This is important because, unless we can convince ourselves that the model is a sufficiently good predictor of subject behavior, we can hardly extrapolate our welfare results to variations in group size and information structure, for instance, let alone draw policy implications. Overall, despite some heterogeneity, we find that the model predicts aggregate behavior fairly well.

Finally, to inform the comparison between Unanimity and Veto, we run control treatments with two alternative voting rules: simple majority rule (henceforth Majority) and unanimity rule under the constructive abstention regime (henceforth Constructive Abstention). Beyond being a standard benchmark in the literature, the Majority treatments are useful to assess the consequences of equilibrium multiplicity under Veto. Constructive Abstention is strategically equivalent to Veto but changes actions' focality. We exploit this framing difference to study the sensitivity of Veto's welfare properties to the presence of less sophisticated agents.

Related Literature

Our paper is the first to compare the information aggregation properties of different unanimous voting rules in the laboratory.

Guarnaschelli, McKelvey, and Palfrey (2000) documents evidence of strategic voting under Unanimity in the Condorcet Jury setup. They show that, in line with theoretical predictions, subjects vote against their signal and that this improves information aggregation with respect to sincere voting. Goeree and Yariv (2011) also find evidence of strategic voting. In addition, they find that allowing for communication among agents before the vote substantially reduces the impact a voting rule has on the group decision, even when theory predicts that it should not be the case.⁴ Our paper contributes to this literature in at least two ways. First, and most importantly, we expand the set of unanimous rules beyond Unanimity by considering Veto and Constructive Abstention. Second, we consider information structures (i.e. the case where the red signal is decisive) for which unanimous rules are optimal mechanisms.⁵

⁴Other recent papers on information aggregation in committees include Battaglini, Morton and Palfrey (2010), Morton and Tyran (2011), Bhattacharya, Duffy and Kim (2014), Grosser and Seebauer (2013), Fehrler and Hughes (2014), Le Quement and Marcin (2015), Mattozzi and Nakaguma (2015), Herrera, Llorente-Saguer and McMurray (2016), Mengel and Rivas (2016). See Palfrey (2015) for an overview.

⁵Bhattacharya, Duffy and Kim (2014) consider biased information structures but not the case where the red signal is decisive.

Veto power has been studied and compared in private value environments. Relevant papers include Wilson and Herzberg (1987), Haney, Herzberg and Wilson (1992), Kagel, Sung, and Winter (2010), Battaglini, Nunnari, and Palfrey (2012), and Nunnari (2014). In this literature, veto rights constrain the set of implementable policies. In common-value environments such as ours, agents do not use veto for their purely private benefit. Instead, vetoing is a way to convey negative information about the proposal. When negative signals are precise enough, veto rights improve information aggregation.

Finally, our paper also contributes to the literature on framing that builds on Tversky and Kahneman (1981). Applications to political science have shown, for example, that the outcome of a vote can be affected by the ordering on the ballots (Miller and Krosnick, 1998), the use of core-value electoral platforms (Brewer, 2001) or the emphasis in the initiatives' titles (Bütler and Maréchal, 2007). We contribute to this strand of literature by focusing on *equivalency framing effects*, which studies how different logically equivalent frames affect choices (see Druckman, 2001). A typical strategy in these studies is to frame the same problem in terms of gains or losses. What we do is different in that we manipulate actions' focality. Moreover, we do this in a strategic environment, in which, to a certain extent, sophisticated agents are able to compensate the actions of the agents that are affected by the frame.

2 Theory

2.1 The Model

A group of $n \geq 3$ agents (with n odd)⁶ must vote over two possible alternatives, *Blue* and *Red*.

Information structure. There are two states of nature, $\omega \in \{\omega_B, \omega_R\}$, which materialize with equal probability. The actual state of nature is not observable, but each agent privately observes an imperfectly informative signal: either s_B or s_R (the *blue* or *red* signal, respectively). Conditional on the state of nature, the signals are independently drawn. The probability that an agent will observe signal s_B is higher in state ω_B than in state ω_R , and the converse is true for s_R . We denote the probability of receiving signal s

⁶That n is odd only simplifies the exposition.

in state ω by $\Pr(s|\omega)$.

Preferences. Agents have *common values*: they all prefer decision *Red* in state ω_R and *Blue* in state ω_B .⁷ We capture this with the following von Neumann-Morgenstern utility function: $u : \{\omega_B, \omega_R\} \times \{Blue, Red\} \rightarrow \mathbb{R}$, with $u(Red, \omega_R) = u(Blue, \omega_B) = 1$, and $u(Red, \omega_B) = u(Blue, \omega_R) = 0$.

Voting systems. The group makes a decision by taking a simultaneous vote. We mainly focus on two voting systems: Majority with Veto power (V) and Unanimity (U). A voting system $\Psi \in \{U, V\}$ is defined as a set of possible actions A_Ψ and an aggregation rule d_Ψ that maps agents' actions into a group decision: $d_\Psi : \{a \in A_\Psi\}^n \rightarrow \{Blue, Red\}$. We denote by X_a the total number of agents playing action a . Agents do not communicate before making their decision.⁸

Definition 1 *Voting system “Veto” is defined by: $V \equiv \{A_V, d_V\}$, where:*

$$\begin{aligned} A_V &= \{b, r, v\} \\ d_V &= \begin{cases} Blue & \text{if } X_v = 0 \text{ and } X_b > X_r \\ Red & \text{otherwise.} \end{cases} \end{aligned}$$

The group decision is *Blue* if and only if no one plays v and there is a majority that plays b . The decision is *Red* otherwise. Hence, we interpret b as a vote for *Blue*, r as a vote for *Red*, and v as a veto (“against *Blue*”). To highlight the differences and similarities between the voting systems, it is convenient to define Unanimity using the same aggregation rule as in the definition above (d_V) and to label the different actions in a similar fashion.

⁷ Assuming that all agents behave strategically, the potential presence of agents that prefer the status quo for private reasons neither affects the comparison of unanimous rules nor the behavior of common value agents (Bouton, Llorente-Saguer, and Malherbe 2016).

⁸ This assumption is not innocuous. In our setup, if communication were allowed, full information sharing would be possible under both Unanimity and Veto (i.e. there would exist an equilibrium in which agents truthfully reveal their information at the communication stage and the revealed information is used to reach an optimal decision at the voting stage). However, this appealing feature disappears when agents differ (sufficiently) in their disutility from wrong decisions (Coughlan 2001, Austen-Smith and Feddersen 2006). But, in that case, Veto still dominates Unanimity (Bouton, Llorente-Saguer and Malherbe 2016). Moreover, there are indeed plenty of situations in which communication is difficult and/or costly. See Persico (2004) and Section 6 in Bouton, Llorente-Saguer and Malherbe (2016) for a thorough discussion of hurdles to communication in related contexts.

Definition 2 *Voting system “Unanimity” is defined by: $U \equiv \{A_U, d_U\}$, where: $A_U = \{b, v\} \subset A_V$ and $d_U = d_V$, with X_r necessarily equal to 0.*

Under Unanimity, agents can play b (vote for *Blue*) or v (*veto Blue*). The group decision is *Blue* if and only if everyone plays b (votes for *Blue*).

Strategy and equilibrium concept. Formally, we define an agent’s strategy as a function $\sigma : \{s_B, s_R\} \rightarrow \Delta(A_\Psi)$. In particular, $\sigma_a(s)$ denotes the probability with which an agent who receives signal s plays a . Following Feddersen and Pesendorfer (1998), we focus on responsive symmetric Bayesian Nash equilibria.⁹

2.2 Equilibrium Analysis and Welfare Properties

In Bouton, Llorente-Saguer, and Malherbe (2016), we characterize the equilibrium under Veto and prove welfare results in a more general version of the model.¹⁰ A key aspect of the welfare analysis is that unanimous rules should be compared according to their information aggregation properties. To do so, we can rely here on the concept of *right decision*.

Definition 3 *The right decision maximizes agents’ expected utility given the realized signal profile. A voting system is efficient if it leads to the right decision always being taken.*

What the right decision is depends on the number of signals of each color, and on their relative precision. When the information structure is sufficiently biased, it can be the case that a single red signal is *decisive*.¹¹

Definition 4 *We say that a red signal is decisive if, unless all signals are blue, the right decision is Red (i.e. keeping the status quo).*

In the experiment, we consider a case where a red signal is decisive, and a case where it is not. In this section, we formally state the relevant theoretical results on which our

⁹In our context, a responsive profile is such that (i) at least some agents play action b with positive probability, and (ii) not all of them play b with probability 1. This ensures that, in equilibrium, some pivot probabilities are strictly positive and agents affect the outcome of the vote with positive probability.

¹⁰That is, we consider all admissible parameters and we allow for the presence of private value agents.

¹¹The relevant condition is $\frac{\Pr(s_R|\omega_R)}{\Pr(s_R|\omega_B)} > \left(\frac{\Pr(s_B|\omega_B)}{\Pr(s_B|\omega_R)}\right)^{n-1}$.

hypotheses are based, and we briefly discuss why we chose these cases. We do not repeat the proofs in this paper, but we give the details of where to find them in Appendix A2.

2.2.1 Baseline case: a red signal is not decisive

The baseline case is the most interesting. In the corresponding parameter region, we have chosen to focus on signals that are equally informative. This corresponds to the case originally studied by Feddersen and Pesendorfer (1998) and then taken to the laboratory by Guarnaschelli, McKelvey, and Palfrey (2000). Comparing our experimental results to theirs offers a simple robustness check in terms of subject behavior.

Lemma 1 *Assume $\Pr(s_B|\omega_B) = \Pr(s_R|\omega_R)$.*

- **Unanimity** *admits a unique equilibrium: $\sigma_b(s_B) = 1$, $\sigma_v(s_R) = 1 - \alpha^*$ and $\sigma_b(s_R) = \alpha^*$ where $\alpha^* \in (0, 1)$;*
- **Veto** *admits two equilibria, one in pure strategies: $\sigma_b(s_B) = \sigma_r(s_R) = 1$, and one in mixed strategies, which are equivalent to those in the equilibrium under Unanimity.*

As Feddersen and Pesendorfer (1998) have shown, it is not an equilibrium under Unanimity for agents to play b with a blue signal and v with a red signal. This is because agents are only pivotal when all other agents play b . Under this strategy, that would imply that they all have received a blue signal, which gives a strong incentive to disregard one's red signal and play b instead. This is why, in equilibrium, agents with a red signal mix (probability α^* is such that they are indifferent between playing b and v).

Under Veto, there are two equilibria. In the pure strategy equilibrium, agents simply play the color of their signal. Doing so, the group always makes the right decision. Hence, there is no reason to deviate. The second equilibrium mimics that under Unanimity. If no other agent ever plays r , playing r becomes strategically equivalent to playing b ; a single vote for *Red* (instead of *Blue*) cannot change the group decision.

Our welfare criterion is an agent's expected utility. Given that agents equally dislike both types of errors, it corresponds to the ex-ante probability to make the right decision.

Proposition 1 *Assume $\Pr(s_B|\omega_B) = \Pr(s_R|\omega_R)$;*

- if agents coordinate on the pure strategy equilibrium under *Veto*, they make the right decision with a strictly higher probability than under *Unanimity*;
- if agents coordinate on the mixed strategy equilibrium under *Veto*, the two systems are welfare equivalent.

Under *Veto*, the mixed strategy equilibrium is Pareto dominated by the pure strategy one, and it is also unstable.¹² This is why we see it as a less credible predictor of agents behavior. Beside comparing the realized welfare properties of *Veto* and *Unanimity*, running the experiment offers us an empirical test of such an equilibrium selection prediction. To this purpose, running a control treatment under simple majority rule offers a natural benchmark for such a test because, in its unique equilibrium, strategies are equivalent to the pure strategy equilibrium under *Veto*.¹³

That *Veto* strictly dominates *Unanimity* extends to any biased signal structures as long as the red signal is not decisive. Note however that the Pareto dominant equilibrium may then include mixed strategies.¹⁴

2.2.2 Extreme case: a red signal is decisive

When the red signal is decisive, *Veto* no longer strictly dominates *Unanimity*. Under both rules, it is a weakly dominant strategy for agents with red signals to play v . Taking this into account, agents with blue signals optimally choose to play b , and the group always makes the right decisions.

Lemma 2 *If a red signal is decisive, **Veto** and **Unanimity** admit a unique equilibrium where $\sigma_b(s_B) = \sigma_r(s_R) = 1$.*

Proposition 2 *If a red signal is decisive, **Veto** and **Unanimity** are welfare equivalent.*

Even though they are welfare equivalent (they are both efficient), *Veto* has a larger action set and, as we discuss below, the sincere action does not necessarily coincide with equilibrium strategies. This could undermine *Veto*'s information aggregation properties.

¹²If other agents with red signals play r with strictly positive probability, playing r or b are no longer strategic equivalent, and the equilibrium unravels.

¹³See Appendix A1 for a formal definition and the characterisation of equilibrium.

¹⁴This is similar to what happens under the simple majority rule.

2.2.3 Sincere voting and framing effects

To analyze departures from strategic behavior, it is worth discussing the notion of “sincere voting”, which is often used in the literature. In a related paper, Austen-Smith and Banks (1996) describe sincere voting as “*an individual’s optimal voting decision based solely on her own private information*”. That is, agents vote sincerely if they select “*the alternative yielding their highest expected payoff conditional on their own signal*”.

In our setup, contrarily to theirs, such a definition of sincere voting does not nail down the agents’ choice of action under Veto. If they receive a blue signal, playing b is the unambiguous sincere action. But, if they receive a red signal, there are two candidates for the sincere action as both playing r or v favors *Red*.

This ambiguity leaves the door open to different definitions. In the spirit of previous literature, we consider agents that are affected by the focality of an action’s label –i.e., the focal action for decision *Red* is r .¹⁵ Given the way we have defined Veto, this corresponds to a vote for *Red*. But it is also possible to frame the voting system in such a way that playing r amounts to vetoing *Blue*.

This frame is *unanimity rule under the constructive abstention regime* (Constructive Abstention), which we formalize as follows.¹⁶

Definition 5 *Voting system “Constructive Abstention” is defined by: $CA \equiv \{A_{CA}, d_{CA}\}$, where $A_{CA} = \{b, c, r\}$ and*

$$d_{CA} = \begin{cases} \text{Blue} & \text{if } X_r = 0 \text{ and } X_b > X_c \\ \text{Red} & \text{otherwise.} \end{cases}$$

Under Constructive Abstention, agents can play b , c , or r , and the group decision is *Blue* if and only if (i) no one plays r and (ii) more agents play b than c . This system is strategically equivalent to Veto, in the sense that it also has three possible actions and the aggregation rule is equivalent. They only differ in their labeling (c and r correspond respectively to r and v under Veto; b does not change). Hence, playing r corresponds to vetoing under Constructive Abstention, whereas it is a vote for *Red* under Veto.¹⁷ Since

¹⁵An alternative definition would be that sincere voters vote for the strongest action that favors their most preferred outcome (based on their own information). Here, this would correspond to playing v when receiving a red signal.

¹⁶We named this system after the notion of constructive abstention introduced by the Treaty of Amsterdam (1997).

¹⁷To avoid confusion, we will stick to the terminology “vote for Blue”, “vote for Red”, “veto”.

r is now the focal action corresponding to the red signal, we have that sincere agents do veto when they receive it (they still vote for *Blue* with a blue signal). Hence, while Veto and Constructive Abstention lead to identical outcomes in our strategic voting model, this will not necessarily be the case in the presence of sincere agents.

First, consider the baseline case. Since sincere actions and equilibrium strategies are the same under Veto, the presence of sincere agents is therefore inconsequential. This is not true, however, under Constructive Abstention because these agents use their veto, which is inefficient. In the extreme case, it is under Constructive Abstention that the presence of sincere agents does not affect the outcome, and it is under Veto that it does; because sincere agents vote for *Red* instead of vetoing when they get the red signal.

Ultimately, what we are interested in is the comparison with Unanimity. The presence of sincere agents affect negatively the performance of Unanimity in the baseline case (those who receive a red signal veto excessively), but not in the extreme case. This leads to the following empirical predictions: in the presence of sincere agents, (i) Veto outperforms both Unanimity and Constructive Abstention in the baseline case, (ii) Constructive Abstention and Unanimity outperform Veto in the extreme case.

This said, it is important to note that sincere agents' impact on welfare depends on the ability of strategic agents to adjust their strategies so that the group behavior resembles the optimal one (we refer to this mechanism as *compensation*). Unfortunately, in the baseline case under Constructive Abstention, sincere agents who receive a red signal use their veto power, which strategic agents' cannot overturn. By contrast, under Veto, sincere agents do not use their veto power when it is efficient to do so. This leaves room for correction by strategic agents. The presence of sincere agents should thus be less consequential under Veto in the extreme case, than under Constructive Abstention in the baseline case.

3 The Experiment

3.1 Design and Procedures

To test our theoretical predictions and potential framing effects, we ran controlled laboratory experiments. Experiments were conducted at the BonnEconLab at the University of Bonn between June and September 2012. We ran a total of 48 sessions, each comprised

of 18 subjects. No subject participated in more than one session. Students were recruited through the online recruitment system ORSEE (Greiner 2004), and the experiment was programmed and conducted using the software z-Tree (Fischbacher, 2007).

Subjects were introduced to a game with the same structure as the one presented in Section 2.1. Following the experimental literature on the Condorcet Jury Theorem initiated by Guarnaschelli, McKelvey, and Palfrey (2000), we did not refer to states of the world or signals but to jars and balls respectively. There were two jars, the *Blue jar* (representing state ω_B) and the *Red jar* (representing state ω_R). Each jar contained a total of 100 red and blue balls. The proportion of red and blue balls in each jar varied across treatments.

Each time the game was played, one of the jars was randomly selected with equal probability by the computer. The subjects were not told which jar had been selected, but they were privately shown a ball randomly and independently drawn from the selected jar. Hence, a blue ball corresponds to s_B and a red ball corresponds to s_R . After seeing their ball, the subjects had to vote. The possible votes and the aggregation rule varied across treatments.

If the group decision matched the color of the jar, the payoff for all members of the group was 100 talers. Otherwise, it was 10 talers.

We had two treatment variables, which led to a 2×4 design and eight different treatments. The first variable was the *voting rule*. We experimented using the four voting rules described above: Veto (V), Unanimity (U), Majority (M), and Constructive Abstention (CA). Their framing was the following. To vote, subjects had to click a button of their choice. In V treatments, subjects had to choose among *blue*, *red*, and *veto*. If a subject vetoed, the group decision was the *Red* jar. If nobody vetoed, the group decision was the jar whose color had received the most votes (in a language that can be common to our four voting systems, we refer to these actions as voting for *Blue*, voting for *Red* and vetoing, respectively). In U treatments, subjects had to choose between *blue* and *red*. The group decision (that is, the jar that was selected by aggregating the votes) was *Blue* if and only if all subjects chose *blue* (these actions correspond to voting for *Blue* and vetoing). In M treatments, subjects also had to choose between *blue* and *red*, but the group decision was the jar whose color had received the most votes (these actions correspond to voting

for *Blue* and voting for *Red*). In CA treatments, subjects had to choose between *blue*, *abstain*, and *red*. If a subject chose *red*, the group decision was the *Red* jar. If nobody chose *red*, the group decision was the *Blue* jar, as long as there were more votes for *Blue* than abstentions (these actions correspond to voting for *Blue*, voting for *Red* and vetoing, respectively).

The second variable that varied across treatments was the *information structure* –the likelihood of getting the right signal in either state. In Setting 1 (the baseline case), this likelihood was the same in both states: $\Pr(s_B|\omega_B) = \Pr(s_R|\omega_R) = 0.7$. In Setting 2 (the extreme case), signal precision differed. In fact, the red signal was decisive: $\Pr(s_B|\omega_B) = 0.99$ and $\Pr(s_R|\omega_R) = 0.3$.

We ran six sessions for each treatment. Each session consisted of 50 rounds played by the same 18 subjects. In each round, these subjects were randomly split into two groups of 9, and the game was played separately in each group. Table 1 provides an overview of the different treatments.

Treatment	Voting system	Setting	% blue balls in <i>Blue</i> jar	% red balls in <i>Red</i> jar
V1	Veto	1	70%	70%
U1	Unanimity	1	70%	70%
M1	Majority	1	70%	70%
CA1	Constr. Abs.	1	70%	70%
V2	Veto	2	99%	30%
U2	Unanimity	2	99%	30%
M2	Majority	2	99%	30%
CA2	Constr. Abs.	2	99%	30%

Table 1: Treatment overview.

All experimental sessions were organized along the same procedure: subjects received detailed written instructions, which an instructor read aloud (see Appendix A5). Before starting the experiment, students were asked to answer a questionnaire to confirm their full understanding of the experimental design. After the questionnaire, subjects began to play. At the end of each round, each subject received the following information: (i) the jar that was selected by the computer, (ii) the group decision, (iii) the number of votes for each alternative, and (iv) their payoff for that period.

To determine payment at the end of the experiment, the computer randomly selected

five periods; the total amount of talers earned in these periods was converted to euros with a conversion rate of 0.025. In total, subjects earned an average of 12.99€, including a show-up fee of 3€.

3.2 Equilibrium Predictions and Alternative

Equilibrium strategies. Table 2 summarizes the predictions of behavior drawn from Lemmas 1 and 2. In the baseline case, there are 2 equilibria under V and CA . In the table below, and henceforth, when we refer to the model predictions, we assume that agents coordinate on the Pareto dominant equilibrium (the one in pure strategies).

		Baseline case (1)			Extreme case (2)				
		%	for <i>Blue</i>	for <i>Red</i>	veto	%	for <i>Blue</i>	for <i>Red</i>	veto
Veto	blue ball		100	0	0		100	0	0
	red ball		0	100	0		0	0	100
Unanimity	blue ball		100	—	0		100	—	0
	red ball		77	—	23		0	—	100
Majority	blue ball		100	0	—		66	34	—
	red ball		0	100	—		0	100	—
Constr. Abs.	blue ball		100	0	0		100	0	0
	red ball		0	100	0		0	0	100

Table 2: Equilibrium Predictions. Predicted probability (in percentage) of playing each action for each signal (ball) received, voting rule, and setting.

Sincere voting. As explained above, the sincere voting hypothesis assumes that subjects vote for the action that corresponds to their signal color. Note that this implies identical predictions in the baseline and extreme cases. Table 3 summarizes the predictions.

4 Experiment: Results and Analysis

The main purpose of this section is to present our empirical analysis of the dominance of Veto over Unanimity. First, we compare average payoffs and information aggregation scores. Second, we delve into subject behavior to assess whether the results can be attributed to the differences in behavior predicted by the model and/or by the sincere voting hypothesis. We first consider the baseline case and then the extreme one. Finally,

		%	for <i>Blue</i>	for <i>Red</i>	veto
Veto	blue ball		100	0	0
	red ball		0	100	0
Unanimity	blue ball		100	—	0
	red ball		0	—	100
Majority	blue ball		100	0	—
	red ball		0	100	—
Constr. Abs.	blue ball		100	0	0
	red ball		0	0	100

Table 3: Predicted probability (in percentage) of playing each action for each signal (ball) received under sincere voting.

we exploit the framing effects (Veto versus Constructive Abstention) to push further the analysis of sincere voting.

All the non-parametric tests we refer to are two-sided and use averages at the matching group level as their unit of analysis. To allow for learning in the initial periods, we focus on the second half of the experiment. That is, we present in the main text our analysis of rounds 26 to 50.¹⁸ Our statements about statistical significance are at the 10% confidence level. Unless stated explicitly, they also holds for the whole 50 rounds.

4.1 Does Veto dominate Unanimity – and if so, why?

4.1.1 Baseline case

Payoffs. Table 4 displays realized average payoffs under the two rules and compares them to the model predictions and those under the sincere voting hypothesis. The results and predictions under Majority are also displayed as a point of comparison. To facilitate interpretation, we present payoffs in terms of the proportion (or probability) of mistakes.

	Experiment	Model predictions	Sincere voting
Veto	12.7	9.9	9.9
Unanimity	38.3	34.0	48.0
Majority	12.3	9.9	9.9

Table 4: Proportion (in percent) of mistakes in baseline case treatments.

We find that the proportion of mistakes is roughly 3 times larger under Unanimity than under Veto. This difference is statistically significant (Mann-Whitney, $z = 2.898$,

¹⁸See Appendix A3 for a comparison with the first half.




Figure 1: Information aggregation score in treatments V1, U1 and M1.

$p = 0.004$),¹⁹ which provides strong support for the hypothesis that Veto strictly dominates Unanimity in this case. In contrast, payoffs under Veto and Majority are very close (and not statistically different: Mann-Whitney, $z = 0.081$, $p = 0.935$).

Information aggregation. Ultimately, we are interested in whether the right decisions are taken. A simple way to assess this is to look at all the cases where the realized signal profile includes a given number of balls of each color (say, for instance, 6 blue and 3 red) and compute the proportion of times that the group made the right decision (Blue in this example) in these cases. There is, however, a caveat; there are very few observations for some signal profile realizations (for instance, 9 blue, 0 red), which can give us a noisy picture. To circumvent this issue, we simulated 10,000 group decisions for each possible number of blue balls received in a group *based on actual individual behavior* (see Appendix A4 for details). With these, we computed an information aggregation score, that corresponds to the proportion of decisions that were right.

Figure 1 plots these scores for Veto, Unanimity, and Majority for each possible number of blue balls in the draw (0 to 9). In addition, the figure includes the information aggregation score predicted by our two benchmarks: the model prediction and the sincere voting hypothesis.

First, consider Veto (see V1, on the left panel). Both benchmarks predict perfect scores. We find that essentially no mistakes are made when there are 3 blue balls or less in the draw. When there are 4 or 5 blue balls, approximately 20% of decisions are mistakes. The natural interpretation is that in such cases, it only takes the deviation

¹⁹See the Mann-Whitney tests for all pairwise comparisons in Appendix A3.

from equilibrium of a single individual for the group decision to be wrong. Finally, as the number of blue balls goes to 6 and more, the proportion of mistakes steadily goes down again.

Under Unanimity: (i) the model predicts poor information aggregation, and (ii) the sincere voting hypothesis predicts even worse information aggregation on average, but no mistake when there is a majority of red balls (i.e., there is no type I error –or in the jury interpretation, an innocent is not convicted). First, we find that when there is a majority of red balls, there are almost no mistakes. This is very close to the sincere voting outcome, but it stands in sharp contrast to the model prediction (about a 10 to 20% mistake rate depending on the ball draw). However, when there is a majority of blue balls, there are many more mistakes than what the model predicts (for example, only 4% of decisions are right with 5 blue balls, compared to more than 30% according to the model). In fact, the realized information aggregation score lies between the model prediction and sincere voting (except for 9 blue balls), closer to sincere voting.

In line with the theoretical prediction, we find that Veto does aggregate information better than Unanimity. When there are 5 blue balls or more, the difference is economically large, and statistically significant.²⁰ When there is a majority of red balls, both systems have close-to-perfect scores.²¹ This means that the gains from using Veto instead of Unanimity essentially materialize through a drastic reduction of errors of type II (i.e. false negatives such as not adopting a good reform, or acquitting a guilty defendant).

To assess the absolute performance of Veto, it is useful to compare it to Majority (M1, right panel), which is well known for its good information aggregation properties (at least in this case). We can see that Veto does almost as well as Majority. The only significant difference occurs with 6 blue balls and more, where Majority does slightly better.

Behavior. Can the dominance of Veto over Unanimity be attributed to the predicted differences in behavior? To answer this question, it is helpful to delve further in the comparison between Veto and Majority. Table 5 presents the average frequency at which agents who received a given signal played a given action.

²⁰We performed non-parametric tests separately for each number of balls based on the simulations.

²¹When there are 3 or 4 blue balls, Unanimity does slightly better on average, but these differences are not statistically significant. When there are 2 blue balls or less, Veto does very slightly better (the difference is statistically significant).

		%	for <i>Blue</i>	for <i>Red</i>	veto
Veto	blue ball	96.4	(100)	3.1 (0)	0.5 (0)
	red ball	3.2	(0)	94.5 (100)	2.3 (0)
Unanimity	blue ball	92.2	(100)	-	7.8 (0)
	red ball	52.9	(76.7)	-	47.1 (23.4)
Majority	blue ball	95.6	(100)	4.0 (0)	-
	red ball	5.3	(0)	94.8 (100)	-

Table 5: Aggregate behavior in the baseline case (treatments V1, U1, and M1). Each cell indicates the percentage of voting blue, red, or veto given the color of the received ball. Model predictions are indicated between brackets.

First, we note that behavior under Unanimity and Majority is overall consistent with previous studies (Guarnaschelli, McKelvey, and Palfrey 2000, and Goeree and Yariv 2011). Under both Veto and Majority, the model predicts that agents play the color of their ball –blue if they receive a blue signal and red if they receive a red signal. We find that subject behavior is fairly close to this (they play accordingly 96% and 94% of the time, respectively). These average frequencies are remarkably close to what we observe under Majority. They are not significantly different (Mann-Whitney, $z = 1.121$, $p = 0.262$ both for red and blue signals). However, while the proportion of deviations is almost identical under the two systems, the deviations themselves are different. Even though they do not generate significant differences in overall average payoff, it provides an explanation for why Veto gets a slightly lower information aggregation score than Majority when there is a majority of blue balls. A veto in such cases is indeed very likely to overturn the right decision that would have otherwise been made by the majority of non-vetoing players.

Figure 2 provides a useful representation of individual behavior. Start with Majority on the right panel. The horizontal axis gives the frequency at which agents with a blue signal voted for *Blue*. The vertical axis gives the frequency at which agents with a red signal voted for *Red*. Hence, the equilibrium strategy (represented by an orange triangle) and sincere voting behavior (represented by a green diamond) are on the top right corner. Each hollow circle in the graph corresponds to the number of subjects who played at those frequencies: the larger this number, the bigger the circle. We can see that the vast majority of subjects *always* play as predicted (84%). And indeed, average behavior (represented by the red circle) is very close to the top right corner.

Now, turn to Veto, on the left panel. Unlike in the case of Majority, there are now




Figure 2: Individual behavior in treatments V1 and M1. Each hollow circle in the graph corresponds to the observed frequency of play: its size represents the number of subjects who actually adopted that frequency of play. The red circle represents the average frequency of play observed, the orange triangle represents the symmetric equilibrium prediction and the green diamond represents the sincere voting prediction.

three possible actions. To facilitate comparison, we abstract from cases where subjects vetoed (which corresponds to less than 2% of total votes). That is, we report frequencies conditional on playing r or b .²² Overall, the picture is remarkably similar to that of Majority: a vast majority of subjects *always* vote as predicted (79% of total votes).

To sum up, both at the aggregate and the individual level, behavior under Veto is remarkably in line with that under Majority, and arguably pretty close to the model predictions. We interpret these as a reasonable validation of our equilibrium selection assumption, and as being consistent with the hypothesis that predicted behavior is driving the good information aggregation performance. Note, however, that equilibrium strategies correspond here to the sincere actions. We will return to this issue in Section 4.2.

Now consider what happens under Unanimity. As we can see in Table 5, subjects massively vote for *Blue* when they receive a blue signal (92%), which is the action predicted by the model (and by the sincere voting hypothesis). However, when they receive a red signal, they veto (i.e., they play r under Unanimity) 47% percent of the times, which is substantially higher than the model prediction (23%). This qualitative feature is in line with the previous findings of Guarnaschelli, McKelvey, and Palfrey (2000) and Goeree

²²91% of the subjects in treatment V1 never vetoed in the second half. Out of the 10 subjects that did veto at least once, 5 did so less than 10% of the time.




Figure 3: Individual behavior in treatment U1. Each hollow circle in the graph corresponds to the observed frequency of play: its size represents the number of subjects who actually adopted that frequency of play. The red circle represents the average.

and Yariv (2011). Overall, there is a larger proportion of votes that are not in line with the model prediction. Since these deviations lean toward a higher proportion of veto, this offers a natural interpretation for why Unanimity generates almost no errors of type I and many more errors of type II than predicted.

Figure 3 depicts individual behavior under Unanimity. Here subjects cannot vote for *Red*. Accordingly, the vertical axis gives the frequency at which agents with a red signal veto. We find that most agents with a blue signal vote for *Blue* a majority of the time. However, there are a number of subjects that sometimes veto when they receive a blue ball. This is not easy to rationalize, but given that the group decision is most often Red anyway, this is not necessarily costly (in the sense that the subject is unlikely to be pivotal).²³ We observe two opposite subject clusters for red signals. Some (31%) always veto in that case. Others (41%) always vote for *Blue*. A possible interpretation is that agents specialize instead of mixing. Another is that some agents vote sincerely (i.e., veto) and the others compensate. A conclusion one could draw from this latter interpretation is that there are too many subjects playing sincerely (and the others cannot fully compensate), which drives the redistribution of errors towards type II and an overall performance that

²³Based on the idea that agents are more likely to make mistakes if payoffs are not too different, Guarnaschelli, McKelvey, and Palfrey (2000) show that quantal response equilibrium can indeed account for some of the departures from the symmetric Bayesian Nash equilibrium.




Figure 4: Information aggregation score in treatments V2, U2 and M2.

is poorer than that predicted by the model.²⁴ Before delving further in the issue of sincere voting, we now turn to the extreme case.

4.1.2 Extreme case

Payoffs. In the extreme case, the model predicts identical outcomes under Veto and Unanimity. Table 6 displays realized and predicted average payoffs in this case. We find that average payoff is higher under Unanimity than under Veto (13% of mistakes compared to 10%), but this difference is not statistically significant (Mann-Whitney, $z = 1.046$, $p = 0.295$). The table also reports the results for Majority, which does much and significantly less well (Mann-Whitney, $z = 2.822$, $p = 0.005$, for the comparison with Veto).²⁵

	Experiment	Model predictions	Sincere voting
Veto	13.3	6.3	45.1
Unanimity	10.0	6.3	6.3
Majority	29.3	24.2	45.1

Table 6: Proportion of mistakes in the extreme case (treatments V1, U2 and M2).

Information aggregation scores. Figure 4 displays information aggregation scores. Here, the right decision is Blue if and only if all the balls are blue.

We find close to perfect information aggregation under both Veto and Unanimity. Most (of the overall few) mistakes happen when there are many blue balls. When there are 5,

²⁴Specialization has been observed in other experiments on information aggregation. See, e.g., Bouton, Castanheira, and Llorente-Saguer (2016).

²⁵This result adds the case where a red signal is decisive to the experimental literature that compares Unanimity and Majority.

6, 7, or 8 blue balls (but not 9), the slight difference in score in favor of Unanimity is statistically significant.

Interestingly, Veto does much better than what sincere voting predicts. Similarly, aggregation scores in Majority (M2 in the right panel) are very close to the model predictions and much higher than under the sincere voting hypothesis when they differ. Both observations are other big hints of strategic behavior. Let us examine this in further details.

		%	for <i>Blue</i>	for <i>Red</i>	veto
Veto	blue ball	86.6	(100)	12.4 (0)	1.0 (0)
	red ball	3.4	(0)	12.2 (0)	84.4 (100)
Unanimity	blue ball	99.5	(100)	-	0.5 (0)
	red ball	9.8	(0)	-	90.9 (100)
Majority	blue ball	66.9	(66)	33.1 (34)	-
	red ball	2.3	(0)	97.7 (100)	-

Table 7: Aggregate behavior in the extreme case (treatments V2, U2, and M2). Each cell indicates the percentage of voting blue, red, or veto given the color of the received ball. Model predictions are indicated between brackets.

Behavior. Table 7 displays the voting frequencies in the extreme case. First, under Veto, note that 84% of the subjects that received a red signal did not vote for *Red* (the sincere action) but decided to veto instead, which is in line with strategic behavior. Still, we observe more departures from the model prediction than in the baseline case. For both signals, a striking 12% of agents vote for *Red* (the model predicts 0%). Under Unanimity, the sincere actions correspond to equilibrium strategies and voting frequencies are overall close to the predictions. Here, the notable departure from the model prediction is that subjects who received a red signal play blue 10% of the time. For both systems, a tentative interpretation of the departures could be that some subjects are reluctant to nail down the group decision (in a sense, they could be *pivotal averse*). Under Unanimity, this could account for the asymmetry in deviations (most deviations are subjects with a red ball voting for *Blue*). Under Veto, such reluctance could help to explain the substantial proportion of agents voting for *Red* with a red ball. Of course, this latter behavior is also consistent with sincere voting. However, we also have a substantial proportion of subjects with blue balls that vote for *Red*. This cannot be accounted for by sincere voting, but




Figure 5: Individual behavior in treatments V2 and U2. Each hollow circle in the graph corresponds to the observed frequency of play: its size represents the number of subjects who actually adopted that frequency of play. The red circle represents the average frequency of play observed, the orange triangle represents the symmetric equilibrium prediction and the green diamond represents the sincere voting prediction.

could reflect strategic compensation for those with a red ball who do not veto. Finally, under Majority, 33% of agents with a blue signal voted for *Red*. This is strikingly close to the model prediction (34%).

Let us now compare individual behavior under Veto and Unanimity (see Figure 5). To facilitate such a comparison we abstract, for V2, from the action that is further away from the model prediction (i.e., playing b with a red signal and playing v with a blue signal) and is indeed played at a very low frequency. On the vertical axis, we display the frequency at which v is played conditional on r or v being played, and on the horizontal axis we display the frequency of b conditional on b or r .²⁶ Under Unanimity, model predictions and sincere voting coincide. We find that a very large fraction of agents always act accordingly (i.e., their behavior corresponds to the top right corner). Still, a non-negligible fraction of agents vote for *Blue* with a red ball, which is not consistent with equilibrium behavior or sincere voting. This is the main reason for the type II errors one observes when there are 7 or 8 blue balls (see Figure 4).

Under Veto, we find more heterogeneity in behavior than in the baseline case. The behavior of only a very few subjects is consistent with sincere voting (they are on the

²⁶The vertical axis is therefore different than that of V1 in Figure 2.




Figure 6: Individual behavior in treatment M2. Each hollow circle in the graph corresponds to the observed frequency of play: its size represents the number of subjects who actually adopted that frequency of play. The red circle represents the average

bottom right corner): only 4% of subjects consistently vote sincerely. Still, a substantial proportion of agents do not veto when they receive a red signal (even though this is a weakly dominant strategy). We also observe a number of subjects that always veto with a red ball but that mix with a blue ball. This behavior is consistent with compensating behavior.

Finally, under Majority, agents overwhelmingly vote red when they receive a red ball. With a blue ball, behavior seems consistent with at least two interpretations: some degree of specialization (instead of randomization according to the model symmetric equilibrium) and/or some sincere voting with countervailing compensation. On average, voting frequencies are almost spot on the equilibrium prediction.

To sum up, we find more departures from the model prediction than in the baseline case. Some are consistent with sincere voting (12% vote for *Red* with a red signal under Veto, for instance) and some are not (10 % vote for *Blue* with a red signal under Unanimity). However, the performances are barely affected in terms of average payoff or information aggregation. The key reason for this is that most departures consist of voting for *Blue* or *Red* (as opposed to veto). These deviations impair information aggregation but do not preclude it: since these actions do not nail down the group decision, it is still possible that the right decision will be made by the group. Hence, they can be interpreted

as noise that slightly affects average payoff. Now, an interesting question arises: what would happen if sincere voting implied exerting one’s veto power?

4.2 Sincere Voting and Framing Effects

As explained in Section 2.2.3, Constructive Abstention is strategically equivalent to Veto but its labeling makes vetoing the sincere action for agents with a red signal. In this section, we exploit this framing difference to explore how the presence of sincere agents affects outcomes and welfare.

4.2.1 Payoffs and information aggregation.

	Baseline case			Extreme case		
	Experiment	Model	Sincere	Experiment	Model	Sincere
Veto	12.7	9.9	9.9	13.3	6.3	45.1
Constr. Abs.	36.3	9.9	48.0	15.7	6.3	6.3
Unanimity	38.3	34.0	48.0	10.0	6.3	6.3

Table 8: Proportion of mistakes under Veto, Constructive Abstention, and Unanimity.

Table 8 displays the payoffs. First, we find that Veto strongly and significantly dominates Constructive Abstention in the baseline case (Mann-Whitney, $z = 2.892$, $p = 0.004$), which is consistent with the presence of sincere agents. Looking at information aggregation scores (see Figure 7) tells us that this comes from the cases with 5 blue balls or more, where the difference is economically strong and statistically significant.²⁷ This suggests that the impossibility for strategic voters to compensate a veto exerted by a sincere agent is indeed relevant. We also find that Constructive Abstention does not do significantly better than Unanimity (Mann-Whitney, $z = 0.326$, $p = 0.744$). In fact, Constructive Abstention’s information aggregation scores are very similar to those of Unanimity.²⁸ Rather than the impossibility for strategic agents to compensate, it may then be that the reason why Constructive Abstention does not do well is that agents coordinate on the Pareto dominated equilibrium.

Second, in the extreme case, Constructive Abstention does slightly less well than Veto

²⁷It is not statistically significant for 0, 3, or 4 blue balls.

²⁸Focusing on the differences that are statistically significant, Constructive Abstention does slightly better when there are 0, 1, 5, 6, 7, or 8 blue balls, and U does better when there are 4.




Figure 7: Information aggregation score in V, U and CA treatments.

(and Unanimity), but the differences are not statistically significant.²⁹ Note that information aggregation scores (Figure 7) in V2 and CA2 are also very close (the small differences are not statistically significant). This is interesting, because, in the presence of sincere agents, we would expect Constructive Abstention to do better than Veto. This raises a series of questions: is sincere voting under Veto in this case less prevalent than under Constructive abstention in the baseline case? Is it that strategic agents are able to compensate? Is it linked to strategic uncertainty or equilibrium uniqueness? Inspecting voting behavior will help us partially answer these questions.

4.2.2 Behavior and interpretation.

Baseline case Voting frequencies under V1 and CA1 are very different (see Table 9). The key points to note are the following. Among agents with a red ball, while 28% veto in CA1 (where it is the sincere action) only 2% do so in V1 (where it is not). In CA1, we also have that 10% of subjects with a red ball vote for *Blue* (this is the natural compensating action), and that 16% with a blue ball did not vote for *Blue*.³⁰ Finally, the difference in

²⁹Mann-Whitney, $z = 1.529$, $p = 0.126$ (Constructive Abstention versus Veto), and $z = 1.376$, $p = 0.169$ (Constructive Abstention versus Unanimity). Note, however that the difference with Unanimity is significant if one considers the full sample (Mann-Whitney, $z = 1.684$, $p = 0.092$). The difference mainly comes from an relative increase in payoffs under Constructive Abstention in the second half, which may reflect differences in learning across these treatments.

³⁰This is hard to rationalize, but given the observed voting frequencies, the unconditional probability that another agent vetoes is fairly high. Hence, the likelihood of being pivotal is very low, and thus the

voting behavior between CA1 and U1 is hardly consistent with subjects coordinating on the Pareto-dominated equilibrium. We can therefore rule out such an hypothesis.

	Ball	Baseline case (1)						Extreme case (2)					
		%	for <i>Blue</i>	for <i>Red</i>	Veto		%	for <i>Blue</i>	for <i>Red</i>	Veto			
V	blue		96.4 (100)	3.1 (0)	0.5 (0)			86.6 (100)	12.4 (0)	1.0 (0)			
	red		3.2 (0)	94.5 (100)	2.3 (0)			3.3 (0)	12.2 (0)	84.4 (100)			
CA	blue		84.5 (100)	10.6 (0)	5.0 (0)			86.7 (100)	11.4 (0)	2.0 (0)			
	red		10.2 (0)	62.2 (0)	27.6 (0)			3.8 (0)	6.1 (0)	90.1 (100)			
U	blue		92.2 (100)	-	7.8 (0)			99.5 (100)	-	0.5 (0)			
	red		52.9 (77)	-	47.1 (23)			9.8 (0)	-	90.9 (100)			

Table 9: Aggregate behavior under Veto, Constructive Abstention, and Unanimity (treatments V1, CA1, U1, V2, CA2, U2). Each cell indicates the percentage of voting blue, red, or veto given the color of the received ball.

Figure 8 displays individual behavior. First, consider CA1, located on the right panel. There is some clustering around the top right corner, which corresponds to the sincere actions. But this is far from overwhelming as this represents only 6% of the subjects (though this number grows to 10% if we look at agents that choose the sincere actions at least 80% of the times with both signals). Overall, we observe very dispersed behavior. Still, even in low proportion, sincere voting can help explain why CA1 fails to aggregate information better than U1. This is because a veto nails down the group decision. In other words, there is no way for strategic agents to compensate.

Extreme case We can compare this to what happens in the extreme case. First, in V2, 12% of subjects with a red ball vote for *Red* (the sincere action). Furthermore, under CA2, 6% of subjects with a red ball also vote for *Red* (not the sincere action). And, finally, inspecting individual behavior (see Figure 8) one can hardly argue in favor of a clustering in the corresponding low-right corner (there are 4 subjects that always choose the sincere action, though). In any case, the proportion of votes that can appear to be sincere in V2 is much lower than in CA1 (recall that 28% subjects with a red ball veto in that case).³¹ Assuming these do indeed reflect sincere voting, this raises the question of why it is more prevalent in CA1 than in V2.

difference in expected payoff of one's own action is low.

³¹The difference seems much smaller in the first half of the experiment. In particular, among agents with red balls, 25% that vote for Red in V2 and 31% veto in CA1. Learning seems therefore stronger under in V2.




Figure 8: Individual behavior in V and CA treatments. Each hollow circle on the graph corresponds to the observed frequency of play; its size represents the number of subjects who actually adopted that frequency of play. The red circle represents the average frequency of play observed, the orange triangle represents the symmetric equilibrium predictions and the green diamond represents the sincere voting prediction.

It has been established in other contexts that framing can affect choices between options. For instance, the insight that people are more likely to select a default option has revolutionized retirement savings in the US (many companies now offer their employee the option to opt out instead of having to opt in). In the Decision Theory literature, such bias has been related to the concept of *decision avoidance*, which is relevant to our findings if one interprets the focal action in our experiment as the default option.

Decision avoidance means that “[the default option] may be chosen in order to avoid a difficult decision” (Dean, Kibris, and Masatlioglu 2014; see also Tversky and Shafir 1992).³² For instance, Dean (2009) finds that subjects facing larger choice sets are more likely to select the default option. Applied to our context, this leads to the hypothesis that agents are more likely to choose the focal action (i.e., vote sincerely) if they face a more complex situation. Our findings are consistent with such an hypothesis, assuming that subjects find the extreme-case game less complex than the baseline one. This perceived complexity level could be due to the fact that the former presents a weakly dominant strategy and an obvious best response (even though computing the posterior involves non-trivial calculations), whereas the latter does not (here, no calculation is really needed if one understands the logic behind the Condorcet Jury Theorem). Strategic uncertainty, or complexity, may therefore also play a role for the difference in behavior we have observed across the two frames.

5 Conclusion

Our main finding is that Veto dominates Unanimity in the laboratory. We therefore provide empirical support for our previous theoretical results (Bouton, Llorente-Saguer, and Malherbe, 2016). Overall, we find that subject behavior is close to the model predictions, and that deviations are not too costly. This provides support for external validity and suggests that it would, indeed, be beneficial for voting bodies that currently use Unanimity to adopt Veto instead.

We have also studied the sensitivity of Veto’s welfare properties to the presence of less sophisticated agents. To do so, we have exploited the fact that Constructive Abstention

³²This result is often referred to as “status quo” bias (See for instance Kahneman, Knetsch, and Thaler 1991). Note that, in our context, status quo refers to something completely different. This is why we prefer to use the phrase “default option”.

is strategically equivalent to Veto, but that it changes actions focality. Our experimental results are consistent with the presence of sincere voters and confirm our empirical prediction that Veto's welfare performance are better when the action of vetoing is less focal. This finding is probably not relevant to experienced voters such as members of the Council of the European Union or the UN Security Council, who are unlikely to be affected by mere relabelling. But, for other committees including less experienced or sophisticated agents, our results provide a cautionary tale: when a voting rule is such that some votes have more weight than others, one must be careful in choosing which is the focal vote.

References

- [1] Austen-Smith, D. and J. Banks (1996). Information Aggregation, Rationality, and the Condorcet Jury Theorem, *American Political Science Review*, 90: 34–45.
- [2] Austen-Smith, D. and T. Feddersen (2006). Deliberation, Preference Uncertainty and Voting Rules. *American Political Science Review*, 100(2): 209–217.
- [3] Battaglini, M., Morton, R. and T. Palfrey (2008). Information Aggregation and Strategic Abstention in Large Laboratory Elections, *American Economic Review*, Papers & Proceedings, 98(2): 194–200.
- [4] Battaglini, M., Morton, R. and T. Palfrey (2010). The Swing Voter's Curse in the Laboratory, *Review of Economic Studies*, 77: 61–89.
- [5] Battaglini, M., S. Nunnari, and T. Palfrey (2012). Legislative Bargaining and the Dynamics of Public Investment, *American Political Science Review*, 106 (2): 407–429.
- [6] Bhattacharya, S., J. Duffy and S. T. Kim (2014). Compulsory versus Voluntary Voting: An Experimental Study, *Games and Economic Behavior*, 84, 111–131.
- [7] Bouton, L., Castanheira, M., and A. Llorente-Saguer (2016). Divided Majority and Information Aggregation: Theory and Experiment, *Journal of Public Economics*, 134: 114–128.

- [8] Bouton, L., A. Llorente-Saguer and F. Malherbe (2016). Get Rid of Unanimity: The Superiority of Majority with Veto Power, *Journal of Political Economy*, forthcoming.
- [9] Brewer, P. R. (2001). Value words and lizard brains: do citizens deliberate about appeals to their core values? *Political Psychology*, 22: 45–64.
- [10] Büttler, M., and M.A. Maréchal (2007). Framing effects in political decision making: evidence from a natural voting experiment. Working paper no. 1940, CESifo, Munich.
- [11] de Condorcet N.C. (1785). *Essai sur l'Application de l'Analyse à la Probabilité des Decisions Rendues à la Pluralité des Voix*, L'Imprimerie Royale, Paris.
- [12] Dean, M., O. Kibris, and Y. Masatlioglu (2014). Limited Attention and Status Quo Bias, mimeo.
- [13] Druckman, J. N. (2001). The implications of framing effects for citizen competence. *Political Behavior*, 23, 225–56.
- [14] Duggan, J., and C. Martinelli. (2001). A Bayesian model of voting in juries, *Games and Economic Behavior*, 37(2): 259-294.
- [15] Feddersen, T., and W. Pesendorfer. (1997). Voting Behavior and Information Aggregation in Elections with Private Information, *Econometrica*, 65: 1029-1058.
- [16] Feddersen, T., and W. Pesendorfer. (1998). Convicting the Innocent: the Inferiority of Unanimous Jury Verdicts, *American Political Science Review*, 92: 23-35.
- [17] Fehrler, S., and N. Hughes (2014). How transparency kills information aggregation: Theory and experiment, *mimeo*.
- [18] Goeree, J. and L. Yariv (2011). An Experimental Study of Collective Deliberation, *Econometrica*, 79(3): 893-921.
- [19] Grosser, J. and M. Seebauer (2013). The curse of uninformed voting: An experimental study, *mimeo*.
- [20] Guarnaschelli, S., McKelvey, R., and T. R. Palfrey (2000). An Experimental Study of Jury Decision Rules, *American Political Science Review*, 94(2): 407-423.

- [21] Haney, P. J., Herzberg, R., & Wilson, R. K. (1992). Advice and consent: unitary actors, advisory models, and experimental tests. *The Journal of Conflict Resolution*, 36(4), 603–633.
- [22] Herrera, H., A. Llorente-Saguer and J. McMurray (2016). The Marginal Voter’s Curse, CEPR Discussion Paper 11463.
- [23] Kagel, J., H. Sung and E. Winter (2010). Veto Power in Committees: An Experimental Study, *Experimental Economics*, 13(2):167–188.
- [24] Kahneman, D., Knetsch, J.L., and R. Thaler (1991). Anomalies: The endowment effect, loss aversion, and status quo bias. *The journal of economic perspectives*, pp.193-206.
- [25] Le Quement, M. T. and I. Marcin (2015). Communication and voting in heterogeneous committees: An experimental study, *mimeo*.
- [26] Maggi, G. and M. Morelli (2006). Self-Enforcing Voting in International Organizations, *American Economic Review*, 96(4): 1137-1158.
- [27] Mattozzi, A., and M. Nakaguma (2015). Public versus Secret Voting in Committees, *mimeo*.
- [28] McLennan, A. (1998). Consequences of the Condorcet Jury Theorem for Beneficial Information Aggregation by Rational Agents, *American Political Science Review*, 92: 413-418.
- [29] McMurray, J. (2013). Aggregating Information by Voting: The Wisdom of the Experts versus the Wisdom of the Masses, *Review of Economic Studies*, 80 (1): 277-312.
- [30] Mengel, F. and J. Rivas (2016). “Common value elections with private information and informative priors: theory and experiments.” Manuscript, University of Bath.
- [31] Miller, J. M., and J. Krosnick (1998). The Impact of Name Order on Election Outcomes. *Public Opinion Quarterly*, 62(3):291–330.
- [32] Morton, R. B., and J.-R. Tyran (2011). Let the experts decide? Asymmetric information, abstention, and coordination in standing committees, *Games and Economic Behavior*, 72(2): 485-509.

- [33] Nunnari, S. (2014). Dynamic Legislative Bargaining with Veto Power: Theory and Experiments, *mimeo*.
- [34] Palfrey, T. (2015). Experiments in Political Economy. In: Kagel, J., Roth, A. (Eds.), *Handbook of Experimental Economics* vol. 2. Princeton University Press.
- [35] Posner, E. and A. Sykes (2014). Voting Rules in International Organizations, *Public Law and Legal Theory Working Papers*, 458.
- [36] Sievers, L. and S. Daws (2014), *The Procedure of the UN Security Council*, New York: Oxford University Press.
- [37] Tversky, A., and D. Kahneman (1981). The framing of decisions and the psychology of choice. *Science*, 211: 453–458.
- [38] Tversky, A. and E. Shafir (1992). Choice under conflict: The Dynamics of Deferred Decision, *Psychological Science*, 3(6):358-361, 1992.
- [39] Wilson, R. K., & Herzberg, R. (1987). Negative decision powers and institutional equilibrium: experiments on blocking coalitions. *The Western Political Quarterly*, 40(4), 593–609.
- [40] Zamora, S. (1980). Voting in International Economic Organizations, *American Journal of International Law*, 74(3): 566-608.

Appendices

Appendix A1. Majority

To highlight the differences and similarities between the voting systems, it is convenient to define Majority using the same aggregation rule as in the definition above (d_V) and to label the different actions in a similar fashion.

Definition 6 *Voting system “Majority” is defined by: $M \equiv \{A_M, d_M\}$, where $A_M = \{b, r\} \subset A_V$ and $d_M = d_V$, with X_v necessarily equal to 0.*

Under Majority, agents can play b (vote for *Blue*) or play r (vote for *Red*). The group decision is *Blue* if and only if there are strictly more votes for *Blue* than for *Red*.

The next Lemma characterizes the equilibria under Majority for the regions of interest.

Lemma 3 *If $\Pr(s_B|\omega_B) = \Pr(s_R|\omega_R)$, **Majority** admits a unique equilibrium: $\sigma_b(s_B) = 1$, $\sigma_v(s_R) = 1$. If a red signal is decisive, **Majority** admits a unique equilibrium where $\sigma_b(s_B) = 1 - \beta^*$, $\sigma_r(s_B) = \beta^*$ and $\sigma_r(s_R) = 1$ where $\beta^* \in (0, 1)$.*

Appendix A2. Proofs

Proof of Lemma 1. Feddersen and Pesendorfer (1998) establishes the result for Unanimity (see p. 26). Bouton, Llorente-Saguer, and Malherbe (2016) establishes the result for Veto (Propositions 6 and 9). ■

Proof of Propositions 1 and 2. See Theorem 1 in Bouton, Llorente-Saguer, and Malherbe (2016). ■

Proof of Lemma 2. See Propositions 6 and 9 in Bouton, Llorente-Saguer, and Malherbe (2016). ■

Proof of Lemma 3. See Feddersen and Pesendorfer (1998) for the case where a red signal is not decisive. The case where a red signal is decisive is a straightforward extension. ■

Appendix A3. Additional Data and Tests

Behavior in First and Second Half of the Experiment

			First Half			Second Half		
			% Blue	% Red	% veto	% Blue	% Red	% veto
Baseline	Veto	Blue	95.2	4.5	0.4	96.4	3.1	0.5
		Red	4.7	91.3	4.0	3.2	94.5	2.3
	Unanimity	Blue	90.6	-	9.4	92.2	-	7.8
		Red	44.1	-	55.9	52.9	-	47.1
	Majority	Blue	95.6	4.4	-	95.6	4.4	-
		Red	4.1	95.9	-	5.2	94.8	-
	CA	Blue	82.6	12.4	5.1	84.5	10.6	5.0
		Red	9.8	59.2	31.0	10.2	62.2	27.6
Extreme	Veto	Blue	85.3	12.7	2.0	86.6	12.4	1.1
		Red	3.1	24.6	72.3	3.4	12.2	84.4
	Unanimity	Blue	98.9	-	1.1	99.5	-	0.5
		Red	13.2	-	86.8	9.8	-	90.2
	Majority	Blue	69.0	31.0	-	66.9	33.1	-
		Red	3.3	96.7	-	2.3	97.7	-
	CA	Blue	84.3	12.5	3.2	86.7	11.4	2.0
		Red	5.3	8.9	85.8	3.8	6.1	90.1

Table 10: Aggregate behavior in the first and second half (periods 1-25 and 26-50 respectively).

Percentage of Mistakes in First and Second Half of the Experiment

		V	U	CA	M
Baseline	First half	16.0	45.3	44.7	8.7
	Second half	12.7	38.3	36.3	12.3
Extreme	First half	19.7	13.3	20.3	30.3
	Second half	13.3	10.0	15.7	29.3

Table 11: Realized percentage of mistakes in each half of the experiment.

Non-parametric tests on Mistakes

	Setting 1			Setting 2		
	U	V	CA	U	V	CA
M	<	=	<	>	>	>
	$z = 2.898$ $p = 0.004$	$z = 0.081$ $p = 0.935$	$z = 2.892$ $p = 0.004$	$z = 2.898$ $p = 0.004$	$z = 2.822$ $p = 0.005$	$z = 2.415$ $p = 0.016$
U	—	>	=	—	=	=
	—	$z = 2.898$ $p = 0.004$	$z = 0.493$ $p = 0.622$	—	$z = 1.046$ $p = 0.295$	$z = 1.376$ $p = 0.168$
V	—	—	<	—	—	=
	—	—	$z = 2.892$ $p = 0.004$	—	—	$z = 0.243$ $p = 0.808$

Table 12: Mann-Whitney tests on the average realized information aggregation in the second half of the experiment. The sign $>$ ($<$) indicates that the amount of mistakes in the ‘row’ voting rule are strictly higher (lower) than the one in the ‘column’ voting rule. $=$ indicates that there are no significant differences. The only comparison which changes when considering all periods is U vs A in Setting 2: U does significantly better in that case.

Appendix A4. Methodology for the simulations

A simple way to do this would be to compute the proportion of right decisions that would make a group if all 9 subjects would adopt strategies that match voting frequencies. However, this would miss the point that heterogenous behavior can affect outcomes. This is why we base our measure of information aggregation on *individual voting frequencies*.

Here is how we do. We first compute individual voting frequencies for each subject based on the last 25 periods.³³ Then, for each possible realized signal profile (i.e. for each number of blue balls going from 0 to 9), we run 10,000 simulations where members of a matching group (i.e. subjects in a session) are divided into two random groups and randomly assigned the different signals. For each of these simulations we aggregate votes and compute the outcomes. Finally, we pool the results by treatment and compute the proportion of group decisions that coincide with that of the fully informed dictator.

³³In setting 2, three subjects never received a red ball in the second half of the experiment. For these subjects, instead of having the average in the last half of the experiment we use the average for the whole experiment.

Appendix A5. Instructions

Thank you for taking part in this experiment. Please read these instructions very carefully. It is important that you do not talk to other participants during the entire experiment. In case you do not understand some parts of the experiment, please read through these instructions again. If you have further questions after hearing the instructions, please give us a sign by raising your hand out of your cubicle. We will then approach you in order to answer your questions personally. Please do not ask anything aloud.

During this experiment you will earn money. How much you earn depends partly on your own decisions, partly on the decisions of other participants, and partly on chance. Your personal earnings will be paid to you in cash as soon as the experiment is over. Your payoffs during the experiment will be indicated in Talers. At the end of the instructions we are going to explain you how we are going to transform them into euros.

After the experiment, we will ask you to complete a short questionnaire, which we need for the statistical analysis of the experimental data. The data of the questionnaire, as well as all your decisions during the experiments will be anonymous.

The experiment you are participating in is a group decision making experiment. The experiment consists of 50 rounds. The rules are the same for all rounds and for all participants. At the beginning of each round you will be randomly assigned to a group of 9 participants (including yourself). You will not know the identity of the other participants. In each round you will only interact with the participants in your group. Your group will make a decision based on the vote of all group members. (important to say this here, because we say “before voting” later) The decision is simply a choice between two jars, the blue jar and the red jar. In what follows we will explain to you the procedure in each round.

The Jar. There are two jars: the blue jar and the red jar. The blue jar contains 7 blue balls and 3 red balls. The red jar contains 7 red balls and 3 blue balls. At the beginning of each round, one of the two jars will be randomly selected. We will call this the selected jar. Each jar is equally likely to be selected, i.e., each jar is selected with a 50% chance. You will not be told which jar has been chosen when making your decision.

The Sample Ball. Before voting, each of you receives a piece of information that may or may not help you decide which is the correct jar. After a jar is selected for your group, the computer will show each of the participants in your group (including yourself) the color of one ball randomly drawn from that jar. We will call this ball your sample ball. Since you are 9 in your group, the computer separately performs this random draw 9 times. Each ball will be equally likely to be drawn for every member of the group. That is, if the color of the selected jar for your group were red, then all members of your group would draw their sample balls from a jar containing 7 red and 3 blue balls. If the color of your group’s jar were blue, then all members of your group would draw their sample balls from a jar containing 3 red and 7 blue balls. Therefore, if the selected jar is blue, each member of your group has a 70% chance of receiving a blue ball. And if the selected jar is red, each member of your group has a 70% chance of receiving a red ball.

You will only see the color of your own sample ball. This will be the only information you will have when you vote.

Your Vote. Once you have seen the color of your sample ball, you can vote.

[Treatment M & U] You must vote for one of the two jars. That is, you must vote for *Blue* or vote for *Red*.

[Treatment CA] You must either vote for one of the two jars or abstain. That is, you must vote for *Blue*, vote for *Red* or Abstain.

[Treatment V] You must either vote for one of the two jars or veto the blue jar. That is, you must vote for *Blue*, vote for *Red* or Veto Blue.

You can vote for either option by clicking below the corresponding button. After making your decision, please press the ‘OK’ key.

Group Decision. The group decision will be set according to...

[Treatment M] ...majority. The group decision depends on the number of blue and red votes:

- If a majority of the group votes blue, the group decision is blue.
- Otherwise, if a majority of the group votes red, the group decision is red.

[Treatment U] ...unanimity. If you or anyone in your group votes red, the group decision is red. Otherwise, the group decision is blue. That is, the group decision is blue if and only if you and everybody in your group vote blue.

[Treatment CA] ...unanimity with possibility of abstention and majority quorum. If you or anyone in your group votes red, the group decision is red. In case there is no vote for *Red*, the group decision depends on the number of blue votes and abstentions:

- If less than a majority of the group abstains, the group decision is blue.
- Otherwise, if a majority of the group abstains, the group decision is red.

[Treatment V] ...majority rule with veto. If you or anyone in your group vetoes blue, the group decision is red. In case there is no one who vetoes blue the group decision depends on the number of blue and red votes:

- If a majority of the group votes blue, the group decision is blue.
- Otherwise, if a majority of the group votes red, the group decision is red.

Payoff in Each Round. If your group decision is equal to the correct jar, each member of your group earns 100 Talers. If your group decision is incorrect, each member of your group earns 10 Talers.

Information at the end of each Round. Once you and all the other participants have made your choices, the round will be over. At the end of each round, you will receive the following information about the round:

- Total number of votes for *Blue*
- Total number of votes for *Red*
- [Treatment CA] Total number of abstentions
- [Treatment V] Total number of vetoes on blue
- Group decision
- Selected jar
- Your payoff

Final Earnings. At the end of the experiment, the computer will randomly select 5 rounds and you will earn the payoffs you obtained in these rounds. Each of the 50 rounds has the same chance of being selected. The total number of talers accumulated in these 5 selected rounds will be transformed into euros by multiplying your earnings in talers by a conversion rate. For this experiment the conversion rate is 0.025, meaning that 100 talers equal 2.5 Euros. Additionally, you will earn a show-up fee of 3.00 Euros. Everyone will be paid in private and you are under no obligation to tell others how much you earned.