NBER WORKING PAPER SERIES

UNANIMOUS RULES IN THE LABORATORY

Laurent Bouton
Aniol Llorente-Saguer
Frédéric Malherbe

Working Paper 21943
http://www.nber.org/papers/w21943

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2016

We thank participants of the ESA 2012 meetings in Tucson and Cologne, and seminar participants at the Berlin Colloquium in Behavioral Economics, Boston University, City University 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, and Nikos Nikiforakis. 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). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Unanimous Rules in the Laboratory
Laurent Bouton, Aniol Llorente-Saguer, and Frédéric Malherbe
NBER Working Paper No. 21943
January 2016
JEL No. C92,D70

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 observed subject behavior. There are, however, cases where organizing the data seems to require a mix of strategic and sincere voting. This pattern of behavior would imply that the way majority rule with veto power is framed may significantly affect the outcome of the vote. Our data strongly supports such an hypothesis.

Laurent Bouton
Georgetown University
Department of Economics
37th & O Streets, NW
Washington,DC 20057
and Université libre de Bruxelles (ECARES)
and CEPR
and also NBER
boutonllj@gmail.com

Aniol Llorente-Saguer
School of Economics and Finance
Queen Mary, University of London
Mile End Road, London E1 4NS
United Kingdom
a.llorente-saguer@qmul.ac.uk

Frédéric Malherbe
London Business School
Regent's Park
NW1 4SA London
United Kingdom
and CEPR
fmalherbe@london.edu

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. 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. In contrast, 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 in-

[1] See e.g. Zamora (1980). Posner and Sykes (2014), and Maggi and Morelli (2006).

[2] Coughlan (2001), Persico (2004), and Bouton, Llorente-Saguer, and Malherbe (2015), however, highlight cases where unanimous decision making features good information aggregation properties.

[3] Among international organizations, for example, Unanimity is used 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 by the European Council (for the Common Foreign and Security Policy), the United Nations Security Council, the ALADI, the CARICOM, and others.

formation well. Hence, our findings provide empirical support to our previous theoretical result that Veto Pareto dominates Unanimity (Bouton, Llorente-Saguer, and Malherbe 2015). 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).[4] To make the group decision, they hold a simultaneous vote according to a pre-specified voting rule.

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. When the red signal is sufficiently informative (relative to the blue signal), both rules are efficient. This is because enforcing the status quo when observing a red signal is then a weakly dominant strategy. Information aggregation is therefore trivial. When the red signal is not very informative, information aggregation is a more subtle problem. As soon as this is the case, Veto outperforms Unanimity, because it offers the possibility of revealing a negative signal without pinning down the outcome.

Accordingly, we consider two treatments that span these two situations. In the first treatment (which corresponds to the latter case), the two signals are equally informative (as in Feddersen and Pesendorfer, 1998, and Guarnaschelli, McKelvey and Palfrey, 2000). Hence we expect Veto to perform as well as majority rule (henceforth Majority, which we include in the comparison given its good information properties) and to strictly dominate Unanimity. In the second treatment (*biased signals*), the red signal is much more informative than the blue one, and thus we expect Veto to perform as well as Unanimity and expect both to strictly dominate Majority.

As mentioned above, our main finding is that Veto dominates Unanimity in the labo-

[4] 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 2015).

ratory. When signals are equally informative, we find that groups using Veto do as well as those using Majority, and that they 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). When the signals are strongly biased, we find that performances under Veto and Unanimity do not differ significantly (and both rules largely dominate Majority). 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, we find that the model predicts subject behavior fairly well.

However, it is important to stress that deviations from model predictions (even by a very limited proportion of agents) can have a decisive impact on outcomes when unanimous rules are used. From a welfare point of view, it seems therefore crucial that veto power be used strategically and not merely because it is "focal". This implies that the way Veto is framed can significantly affect the outcome of the vote. We test such an hypothesis by considering a rule that differs from Veto only in the way actions are labeled. This rule is unanimity rule under the constructive abstention regime, i.e. abstention is allowed but a majority of yeses is required for the proposal to be accepted. Our findings strongly support this hypothesis. When the rule is such that some votes have more weight than others, one must be very careful in choosing which is the focal vote.

To the best of our knowledge, our paper is the first to compare the information aggregation properties of different unanimous voting rules in the laboratory. In addition, it contributes to the experimental literature on information aggregation in three ways. First, in the strand of papers that study unbiased information structures, it adds Veto to the comparison between Unanimity and Majority (see Guarnaschelli, McKelvey, and Palfrey 2000, and subsequent papers such as Goeree and Yariv 2011).[5] Second, to the strand

[5]Overall, the experimental literature on information aggregation in committees has been growing in recent years. See, for instance, Battaglini, Morton and Palfrey (2010), Morton and Tyran (2011), Bhattacharya, Duffy and Kim (2013), Grosser and Seebauer (2013), Fehrler and Hughes (2014), Herrera, Llorente-Saguer and McMurray (2015), Le Quement and Marcin (2015), and Mattozzi and Nakaguma

that studies biased information structures, it adds the case where Unanimity theoretically dominates Majority. Finally, it is the first that considers framing issues in a Condorcet Jury context.

Veto power has been studied and compared in private value environments. Papers in that literature 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 right constrains the set of implementable policies (by eliminating policies that are not favored by the voters with veto rights). In common-value environments such as ours, agents do not use veto for their purely private benefit. Instead, vetoing is a way to convey very strong negative information about the reform. The existence of such strongly negative signals can be very useful to information aggregation.[6]

2 Theory

2.1 The Model

A group of $n \geq 3$ agents (with n odd)[7] must vote over two possible alternatives, B (Blue) and R (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 in state ω by $\Pr(s|\omega)$.

Preferences. Agents have *common value*: they all prefer decision R in state ω_R and B in

(2015).

[6]Note that in our companion paper we consider both private and common value motives (Bouton, Llorente-Saguer, and Malherbe 2015).

[7]That n is odd only simplifies the exposition.

state ω_B. We capture this with the following von Neumann-Morgenstern utility function:

		State of the world	
		ω_R	ω_B
Group	R	1	0
Decision	B	0	1

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

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

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

The group decision is B if and only if no one plays v and there is a majority that plays b. The decision is R otherwise. Hence, we interpret b as a vote for B (we sometimes refer to it as a vote "for Blue"), r as a vote for R (or "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 and Majority using the same aggregation rule as in the definition above (d_V) and to label the different actions in a similar fashion. (See BLM for a discussion.)

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, voters 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).

Definition 3 *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, voters 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.

Strategy and equilibrium concept. Formally, we define an agent's strategy as a function $\sigma : \{s_B, s_R\} \to \triangle (A_\Psi)$. 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.[8]

Right Decision. If agents would observe all signals, they would agree on what the right (i.e., efficient) decision is.

Definition 4 *The right decision maximizes agents' expected utility given the realized signal profile.*

2.2 Equilibrium Analysis and Welfare Properties

In Bouton, Llorente-Saguer, and Malherbe (2015), we characterize the equilibrium under Veto in a more general version of the model.[9] A key aspect of the welfare analysis is that unanimous rules should be compared according to their information aggregation properties. In this paper, we focus on two representative cases: one in which Unanimity aggregates information well and one in which it aggregates it poorly. Veto, to the contrary, does well in both cases.

2.2.1 Unbiased signals

First, we consider a case where blue and red signals are equally informative (i.e. $\Pr(s_B|\omega_B) = \Pr(s_R|\omega_R)$). This facilitates comparisons with the literature, since this case corresponds to that originally studied by Feddersen and Pesendorfer (1998) and then taken to the laboratory by Guarnaschelli, McKelvey, and Palfrey (2000). In this case, the right decision is simply the one that corresponds to the majority of signals.

[8] In our context, a responsive profile is such that (i) at least some common value 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 –agents affect the outcome of the vote with positive probability.

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

Lemma 1 *If* $\Pr(s_B|\omega_B) = \Pr(s_R|\omega_R)$*:*

- *In the unique equilibrium under* ***Majority****, we have* $\sigma_b(s_B) = \sigma_r(s_R) = 1$*;*

- *In the unique equilibrium under* ***Unanimity****, we have* $\sigma_b(s_B) = 1$*,* $\sigma_v(s_R) = 1 - \alpha^*$ *and* $\sigma_b(s_R) = \alpha^*$ *where* $\alpha^* = \dfrac{\left(\frac{\Pr(s_B|\omega_B)}{\Pr(s_B|\omega_R)}-1\right)\left(\frac{\Pr(s_R|\omega_R)}{\Pr(s_R|\omega_B)}-\left(\frac{\Pr(s_B|\omega_B)}{\Pr(s_B|\omega_R)}\right)^{\frac{1}{n-1}}\right)}{\left(\frac{\Pr(s_R|\omega_R)}{\Pr(s_R|\omega_B)}-1\right)\left(\left(\frac{\Pr(s_B|\omega_B)}{\Pr(s_B|\omega_R)}\right)^{\frac{n}{n-1}}-1\right)} \in (0,1)$*;*

- *The two equilibria above are the only equilibria under* ***Veto.***

Proof. Feddersen and Pesendorfer (1998) establishes the result for Majority (see Appendix B) and Unanimity (see p. 26). Bouton, Llorente-Saguer, and Malherbe (2015) establishes the result for Veto (Proposition 2). ■

Under Majority, an agent is pivotal only if there is a tie without his vote. Assume that agents play b when they receive a blue signal and r when they receive a red signal. An agent is then pivotal when, without counting his own vote, there is a tie between red and blue signals. Since the signals are equally informative, and agents dislike both types of errors equally, the best response is to vote according to the color of his signal. Such a strategy profile aggregates information perfectly in the sense that it maximizes expected utility. The group thus always ends up making the right decision, and there is no incentive to deviate.

Under Unanimity, agents are only pivotal if all other agents play b. Assume that agents who observe the blue signal play b and agents with a red signal play v. Now, consider the strategic position of an agent with a red signal. Given the assumption, this agent can only be pivotal if all other agents have received a blue signal. Her posterior beliefs therefore combine her own red signal with the $n-1$ blue signals of other agents. That constitutes a vast majority of blue signals; the agent should conclude that it is much more likely that we are in the blue state (ω_B). The best response is therefore to play b, and such a strategy profile cannot be an equilibrium. In equilibrium, agents with a blue signal always play b, but agents with a red signal mix between b and v (probability α^* is such that she is indifferent between the two). Feddersen and Pesendorfer (1998), who first derived this result, have shown that such an equilibrium does not always lead to the right decision.

Under Veto, there are two equilibria. The first equilibrium mimics the equilibrium under Majority. The idea is simply that this equilibrium is still feasible (the actions are available) and always leads to the right decision. There is therefore no incentive for an agent to deviate (recall that they have a common objective). The second equilibrium mimics that under Unanimity. Again, this is feasible because all actions are available. The reason that there is no incentive to deviate is that a deviation by a single agent cannot improve the outcome. When no other agent ever plays b (which is the case in this equilibrium), playing b becomes strategically equivalent to playing r; a single vote for R (instead of B) cannot change the group decision. Thus the equilibrium under Unanimity must be an equilibrium under Veto.

2.2.2 Biased signals

Second, we look at a case where Unanimity is efficient, but Majority performs badly. This happens when the red signal is precise enough that a benevolent dictator would choose R unless all signals are blue. In this case, the right decision is B if and only if all signals are blue.

Lemma 2 *If*

$$\frac{\Pr(s_R|\omega_R)}{\Pr(s_R|\omega_B)} > \left(\frac{\Pr\left(s_B|\omega_B\right)}{\Pr\left(s_B|\omega_R\right)}\right)^{n-1},$$

- *in the unique equilibrium under* ***Majority****, we have* $\sigma_r^{M*}(s_R) = 1$, $\sigma_r^{M*}(s_B) = \eta^*$ *and* $\sigma_b^{M*}(s_B) = 1-\eta^*$, *with* $\eta^* = \frac{\left(\frac{\Pr(s_B|\omega_B)}{\Pr(s_B|\omega_R)}-1\right)\left(\frac{\Pr(s_R|\omega_R)}{\Pr(s_R|\omega_B)}-\left(\frac{\Pr(s_B|\omega_B)}{\Pr(s_B|\omega_R)}\right)^{\frac{n+1}{n-1}}\right)}{\left(\frac{\Pr(s_R|\omega_R)}{\Pr(s_R|\omega_B)}-1\right)\left(\left(\frac{\Pr(s_B|\omega_B)}{\Pr(s_B|\omega_R)}\right)^{\frac{2n}{n-1}}-1\right)} \in (0,1)$;
- *in the unique equilibrium under* ***Unanimity*** *and* ***Veto****, we have* $\sigma_b^{U*}(s_B) = \sigma_r^{U*}(s_R) = 1$.

Proof. See Bouton, Llorente-Saguer, and Malherbe (2015) for Veto (Proposition 2) and Unanimity (Supplementary Appendix). See Appendix A1 for Majority. ■

Here the red signal is so precise (relative to the blue one) that one red signal is sufficient information to conclude that decision R is better than decision B (even if all other signals are blue). Hence, under both Unanimity and Veto, agents who receive a red signal must play v. Taking this into account, an agent with a blue signal will optimally choose to vote blue; as a result, B will be chosen only when all agents have received the blue signal. In this equilibrium, the group always ends up making the decision that an informed benevolent dictator would have taken.

Under Majority, agents who receive a red signal cannot enforce option R (because playing v is not allowed). They will thus play r. Taking this into account, agents with blue signals prefer to downweigh them and play r with positive probability. Doing so, they "compensate" for the relative imprecision of the blue signal. In this equilibrium, the group does not always make the informed benevolent dictator's decision.

2.2.3 Welfare

Our welfare criterion is an agent's ex-ante utility. Given that agents equally dislike both types of errors, it corresponds to the ex-ante probability of choosing R in state ω_R and B in state ω_B.

Definition 5 *A voting system Ψ (strictly) dominates another voting system Ψ', if there exists an equilibrium under Ψ in which the probability of making the right decision is (strictly) higher than in all equilibria under Ψ'.*

The following proposition captures the dominance of Veto over Unanimity and Majority very well (see Bouton, Llorente-Saguer, and Malherbe (2015) for a thorough welfare analysis).

Proposition 1 *For all n, Veto dominates both Unanimity and Majority. It strictly dominates Unanimity when signals are unbiased, and it dominates Majority in the biased signal case.*

Proof. The proof of the first part is based on McLennan (1998). First, note that any strategy profile in Majority or Unanimity can be reproduced under Veto. This is because Veto's action set includes those of the other systems. Second, recall from McLennan (1998) that, in a pure common value environment, a strategy profile that maximizes ex-ante utility must be an equilibrium. Since the aggregation rules are the same, it follows that there cannot exist an equilibrium under Majority or Unanimity that yields a strictly higher utility than the one that produces the highest utility under Veto. The strict dominance of Veto over Majority in the biased case also follows, because the unique equilibrium under Majority is not an equilibrium under Veto. The strict dominance of Veto (and Majority) over Unanimity when signals are unbiased has been proven in Bouton, Llorente-Saguer, and Malherbe (2015) (and Feddersen and Pesendorfer 1998). ■

2.2.4 Sincere voting and framing

Sincere Voting. 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, voting b is the unambiguous sincere action, according to the definition. If they receive a red signal, however, there are two candidates for the sincere action: r or v, which both favor R.

This ambiguity leaves the door open to different definitions. One of them is what we call *label-sincere voting*, whereby label-sincere voters are affected by the focality of an action's label –i.e., the focal action for decision R is r.[10] Given the way we have defined Veto, this corresponds to a vote for R. But as we show below, we could have equally framed the voting system in such a way that playing r would amount to vetoing B. Framing issues are, therefore, potentially relevant.

Framing. We can formalize *Unanimity rule under the constructive abstention regime* (henceforth Constructive Abstention) as follows.

Definition 6 *Voting system "Constructive Abstention" (CA) is defined by:* $V \equiv \{A_{CA}, d_{CA}\}$, *where* $A_{CA} = \{b, c, r\}$ *and*

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

Even though other voting bodies have already used this system, we named it after the notion of constructive abstention introduced by the Treaty of Amsterdam (see Bouton,

[10] An 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.

Llorente-Saguer, and Malherbe 2015).

In the language of this voting system, playing b equals a vote for B, playing c equals an abstention, and playing r equals a vote for R. The group decision is B if and only if (i) no one plays r and (ii) more agents play b than c.

CA is in fact isomorphic to Veto. Action b is the same under both mechanisms, action c under CA corresponds to action r under V, and action r under CA corresponds to action v under V. In our wording, playing c is in fact a vote for R, and playing r is a veto (against Blue). To avoid confusion, we will stick to this terminology (*vote for Blue*, *vote for Red*, *veto Blue*).

Given this, V and CA lead to identical outcomes in the context of the strategic voting model studied above. However, if (some) agents depart from strategic voting, Veto and Constructive Abstention may lead to different behavior and outcomes. As implied above, under the label-sincere voting assumption, the focal action for R (which is still r) now corresponds to a veto. Since a veto nails down the group decision, framing may dramatically impact the system performance and is worth testing in the laboratory. Throughout the remainder of the paper, we refer to *label-sincere voting* simply as *sincere voting.*

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: Unanimity (U), Majority (M), Veto (V), and Constructive Abstention (CA). Their framing was the following. To vote, subjects had to click a button of their choice. In U treatments, voters 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 played *blue*. 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. In V treatments, on top of *blue* and *red*, subjects could choose to "*veto blue*". 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 CA treatments, subjects had to choose between *blue*, *red*, and *abstain*. If a subject played *red*, the group decision was the Red jar. If nobody played red, the group decision was the Blue jar, as long as there were more votes for Blue than abstentions.

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 unbiased case), this likelihood was the same in both states: $\Pr(s_B|\omega_B) = \Pr(s_R|\omega_R) = 0.7$. In Setting 2 (biased case), signal precision differed. The red signal was much more precise: $\Pr(s_B|\omega_B) = 0.99$ 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 Rule	Setting	% Blue Balls in Blue Jar	% Red Balls in Red 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 A4). 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€. Each experimental session lasted approximately 45 minutes.

3.2 Equilibrium Predictions and alternative

Equilibrium strategies. Table 2 summarizes the predictions of behavior drawn from Lemmas 1 and 2. In the unbiased signal 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.

	Ball	Unbiased signals (1) % Blue	% Red	% Veto	Biased signals (2) % Blue	% Red	% Veto
Veto	Blue	100	0	0	100	0	0
	Red	0	100	0	0	0	100
Unanimity	Blue	100	–	0	100	–	0
	Red	77	–	23	0	–	100
Majority	Blue	100	0	–	66	34	–
	Red	0	100	–	0	100	–
Constr. Abs.	Blue	100	0	0	100	0	0
	Red	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 with the name of the color of the signal. Note that this implies identical predictions in the unbiased and biased signal case. Table 3 summarizes the predictions.

	Ball	% Blue	% Red	% Veto
Veto	Blue	100	0	0
	Red	0	100	0
Unanimity	Blue	100	–	0
	Red	0	–	100
Majority	Blue	100	0	–
	Red	0	100	–
Constr. Abs.	Blue	100	0	0
	Red	0	0	100

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

4 Experiment: Results and Analysis

The main purpose of this section is to present an 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 alternative hypotheses such as sincere voting. Finally, we explore the framing effects by comparing the results under Constructive Abstention with those under Veto.

All non-parametric tests presented within the paper are two-sided and use averages at the matching group level as their unit of analysis. To allow for learning in the initial periods, this section will focus on the second half of the experiment – that is, we present and analyze the data from rounds 26 through 50. Unless stated explicitly, our statements about statistical significance are robust to considering the whole sample.

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

4.1.1 Unbiased signal structure

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 probability of making a mistake (and hence receiving the lower payoff).

	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 of Mistakes in unbiased treatments V1, U1, and M1.

We find that the proportion of mistakes is roughly 3 times larger under Unanimity than under Veto. This difference is statistically significant (Mann-Whitney, $z = 2.898$, $p = 0.004$),[11] 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 agents are making the right decision. 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

[11] See the Mann-Whitney tests for all pairwise comparisons in Appendix A2. From here on, we claim significance at the 10% confidence level.

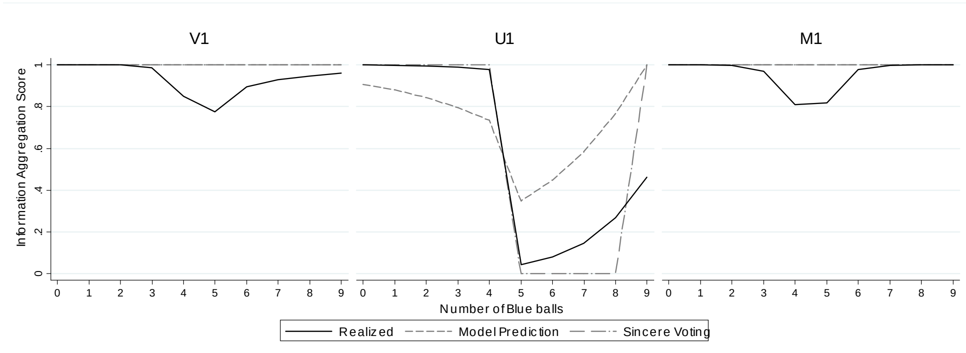

Figure 1: Information aggregation score in unbiased treatments V1, U1 and M1.

each possible number of blue balls received in a group *based on actual individual behavior* (see Appendix A3 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) and Majority (M1, right panel). In these cases, both benchmarks predict perfect scores. We find that groups make essentially no mistakes when there are 3 blue balls or less in the draw. In both cases, approximately 20% of decisions made are mistakes when there are 4 or 5 blue balls. The natural interpretation is that in such a case, if others play according to the equilibrium, it only takes a single individual deviation for the group decision to be wrong. When there are 6 blue balls or more, Veto scores slightly lower than Majority.[12]

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 is almost never a mistake. This is very close to the sincere voting outcome, but it stands in sharp contrast to the model prediction (about a 10 to

[12]We performed non-parametric tests separately for each number of balls based on the simulations. The tests show no significant differences in the information aggregation score between Majority and Veto when the number of blue balls is equal to 5 or lower. Majority does significantly better when the number of blue balls in the group is 6 or higher.

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. When there is a majority of red balls, both systems have close-to-perfect scores.[13] The implication of this finding is 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).

Behavior. Can the dominance of Veto over Unanimity be attributed to the predicted differences in behavior? To answer this question, it is helpful first to compare subjects' behavior under Veto to that under Majority. Table 5 presents the average frequency at which agents who received a given signal played a given action.

		% vote for Blue	% vote for Red	% 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 (33.4)
Majority	Blue ball	95.6 (100)	4.0 (0)	-
	Red ball	5.3 (0)	94.8 (100)	-

Table 5: Aggregate behavior under unbiased signals (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.

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

[13]When 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).

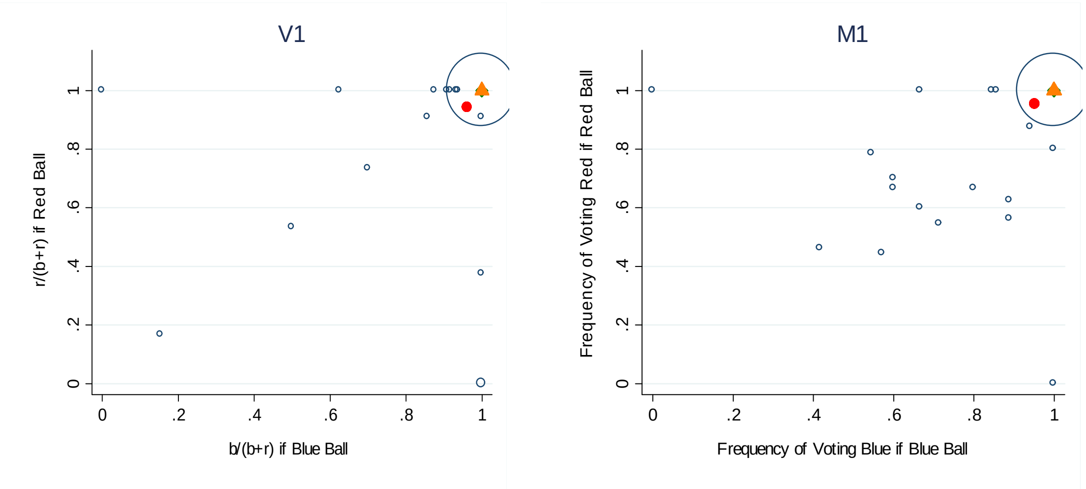

Figure 2: Individual behavior in unbiased treatments V1 and M1. Each hollow circle in the graph corresponds to the observed frequence of play: its size represents the number of subjects who actually adopted that frequence of play. The red circle represents the average frequency of play observed, the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

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 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 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 *red* or *blue*.[14] 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 reasonably 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 these are consistent with sincere voting too (more on this later).

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 *red* under Unanimity) 47% percent of the times, which is substantially higher than the model prediction (33%). This qualitative feature is in line with the previous findings of Guarnaschelli et al. (2000) and Goeree 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, these "mistakes" are not necessarily costly (in the sense that the subject was not pivotal).[15] 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 (we will further discuss asymmetric equilibria

[14] 91% 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.

[15] Based 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.

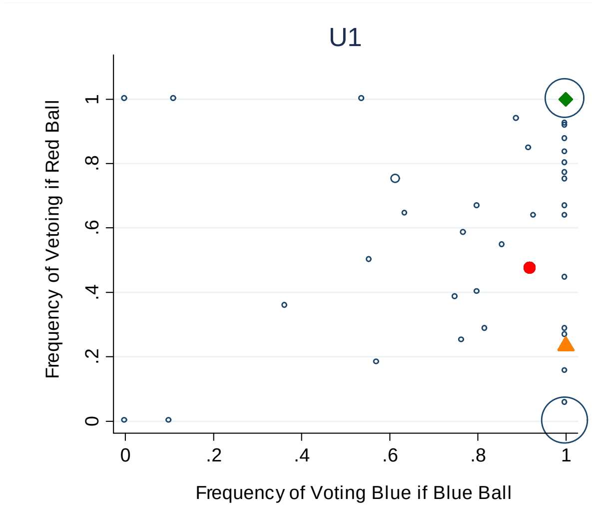

Figure 3: Individual behavior in treatment U1. Each hollow circle in the graph corresponds to the observed frequency of play: its size represents the number of subjects who actually adopted that frequency of play. The red circle represents the average.

later). Another is that some agents vote sincerely (i.e., veto) and the others "compensate". One possible conclusion we can 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 is poorer than that predicted by the model.[16]

4.1.2 Biased signal structure

Payoffs. In the biased case, equilibrium strategies are identical 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, $z = 1.046$, $p = 0.295$). Also as predicted, Veto clearly and significantly dominates Majority (Mann-Whitney, $z = 2.822$, $p = 0.005$).

Information aggregation scores. Figure 4 displays information aggregation scores for the biased case. Here, a benevolent dictator would only choose Blue if all balls were blue. The model predicts such a group decision under both Veto and Unanimity. Hence,

[16] Specialization has been observed in other experiments on information aggregation. See, e.g., Bouton, Castanheira, and Llorente-Saguer (2015a).

	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 under biased signals (treatments V1, U2 and M2).

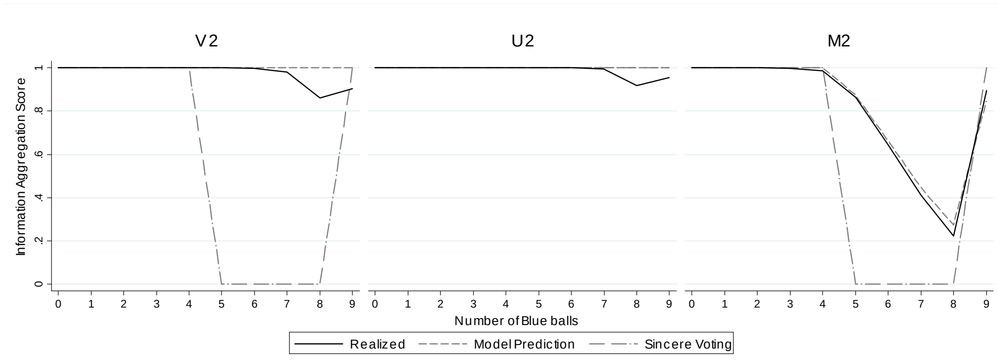

Figure 4: Information aggregation score in biased treatments V2, U2 and M2.

the model predicts perfect information aggregation. This, however, is not the case under Majority because agents cannot single-handedly enforce Red (by construction, they cannot veto). As a result, the model predicts poor information aggregation. The sincere voting predictions are that Unanimity aggregates information perfectly but both Veto and Majority score equally poorly.

We find that both Veto and Unanimity aggregate information well and that Majority scores poorly. In all cases, scores are arguably fairly close to the model prediction. Interestingly, Veto does much better than what sincere voting predicts. Under both Veto and Unanimity, most (of the overall few) mistakes happen when there are many blue balls. Even though Veto does not do as well as Unanimity in these cases, this does not affect average payoffs significantly (though statistically, Unanimity gets a significantly higher score when there are 5, 6, 7, or 8 blue balls). Under majority, most mistakes also happen when there are between 5 and 8 blue balls. Within this window, the score is also (economically and statistically) significantly lower than under Veto.

Behavior. Table 7 displays the voting frequencies in the biased signals case.

Under Majority, voting frequencies are remarkably close to the model prediction. For the other two, however, we observe a higher proportion of departures from the model

		% Vote Blue	% Vote Red	% 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 under biased signals (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.

prediction than in the unbiased case. Under Veto, the striking departure is that approximately 12% of votes are for Red (the model predicts 0%) for agents with blue or red signals. Under Unanimity, the notable departure is that subjects who received a red signal play blue 10% of the time. A tentative interpretation could be that some subjects are reluctant to nail down the group decision. Under Unanimity, this could account for the asymmetry in deviations (most deviations are subjects with a red ball voting for Blue). Under Veto, this 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. Interestingly, 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 could reflect strategic compensation for those with a red ball who play red.

Let us now compare individual behavior under Veto and Unanimity (see Figure 5). To facilitate such a comparisons 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.[17] 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).

[17] The vertical axis is therefore different than that of V1 in Figure 2.

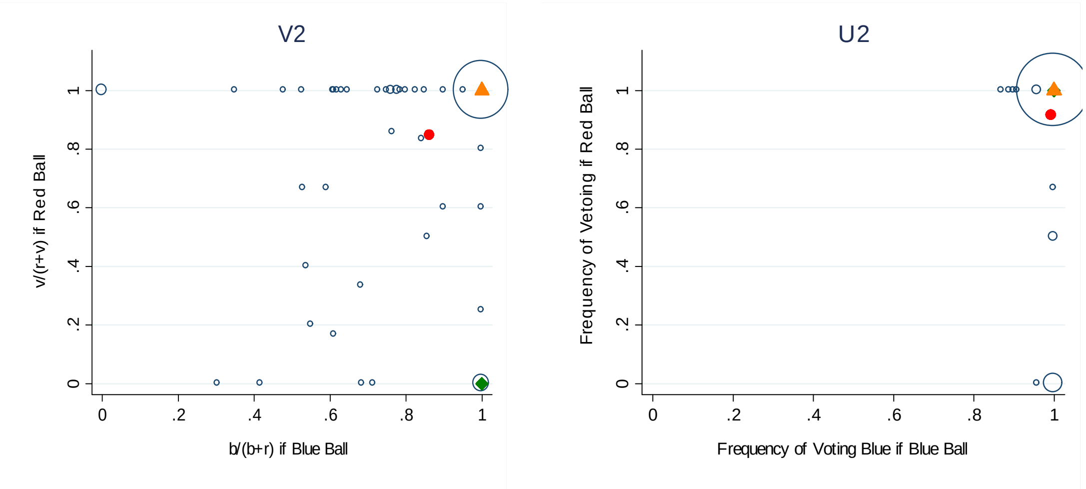

Figure 5: Individual behavior in biased treatments V2 and U2. Each hollow circle in the graph corresponds to the observed frequence of play: its size represents the number of subjects who actually adopted that frequence of play. The red circle represents the average frequency of play observed, the orange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

Under Veto, we find more heterogeneity in behavior than in the previous cases. The behavior of only a very few subjects is consistent with sincere voting (they are on the 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 the tentative interpretation of compensating behavior mentioned above.

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 unbiased 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 arguably barely affected in terms of average payoff or information aggregation. The key reason for this is that most departures consist of voting

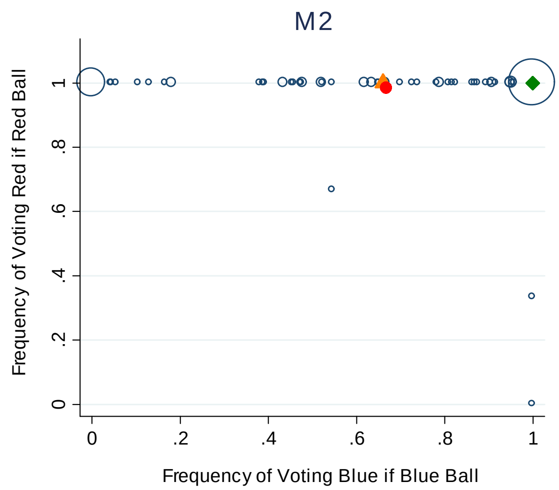

Figure 6: Individual behavior in biased treatment M2. Each hollow circle in the graph corresponds to the observed frequence of play: its size represents the number of subjects who actually adopted that frequence of play. The red circle represents the average

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 Framing effects

As explained above, Constructive Abstention is strategically equivalent to Veto. That is, both systems have three possible actions and, aside from their relabeling, they have identical aggregation rules. The model predictions are therefore identical under these two systems. Under sincere voting, however, labels play a role: whereas playing *red* corresponds to voting for Red under Veto, it corresponds to vetoing under Constructive Abstention. In this subsection, we compare payoffs, information aggregation, and behavior under these two systems.

	Unbiased signals			Biased signals		
	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

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

4.2.1 Payoffs.

The model prediction is that performance is identical under Veto and Constructive Abstention for both the unbiased and biased cases. However, if agents vote sincerely, Veto dominates in the unbiased case and Constructive Abstention dominates in the biased one. We find that Veto strongly and significantly dominates Constructive Abstention in the unbiased case (Mann-Whitney, $z = 2.892$, $p = 0.004$). In fact, Constructive Abstention does not even significantly over-perform against Unanimity (Mann-Whitney, $z = 0.493$, $p = 0.622$). In the biased signal case, average payoffs are slightly larger under Veto than under Constructive Abstention, but the difference is not statistically significant (Mann-Whitney, $z = 0.243$, $p = 0.808$).[18]

4.2.2 Information aggregation.

Figure 7 compares information aggregation scores for Constructive Abstention to those (already displayed earlier) under Veto and Unanimity.

When signals are unbiased (V1, U1, and CA1), Veto outperforms Constructive Abstention. The dominance is economically strong and statistically significant for 5 blue balls or more.[19]

We have already established that when signals are biased (V2, U2, and CA2), Veto performs very well. What we find here is that Constructive Abstention performs very well too. Given that this is the prediction according both to the model and the sincere voting hypothesis, this is not surprising. The direct comparison between Veto and Constructive Abstention does not deliver any significant differences. Economically, the only notable finding is that Veto does slightly better with 9 blue balls.

Overall, such an analysis strongly suggests that in contrast with Veto, Constructive

[18] In the biased case, we do not find a significant difference between CA and U (Mann-Whitney, $z = 1.376$, $p = 0.169$). However, if we consider the whole sample, we do find that U statistically outperforms CA (Mann-Whitney, $z = 1.684$, $p = 0.092$).

[19] It is not statistically significant for 0, 3, or 4 blue balls.

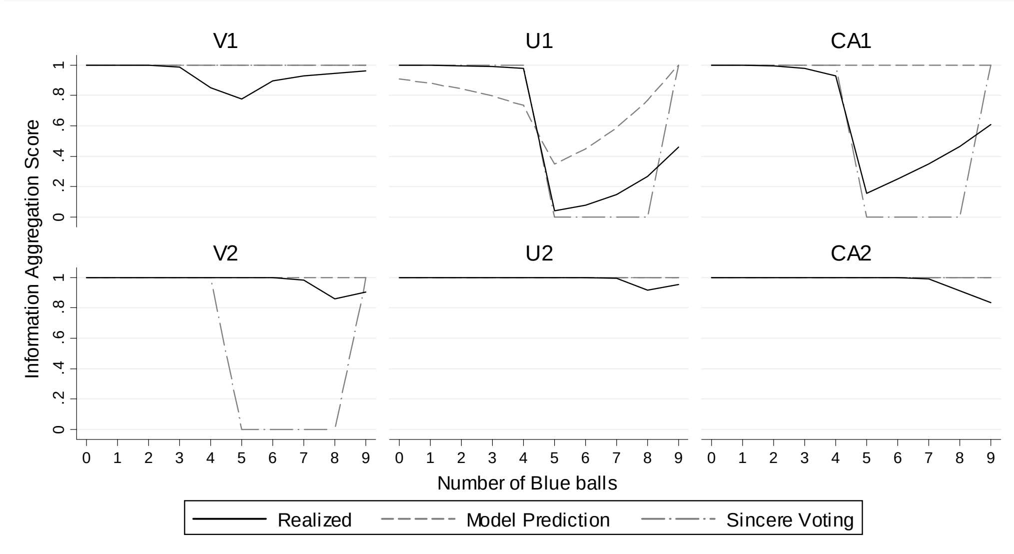

Figure 7: Information aggregation score in V, U and AC treatments.

Abstention inherits the information aggregation properties of Unanimity rather than Majority. We can see this intuitively through information aggregation scores for CA1 and CA2, which show similar patterns to those of U1 and U2 (see Figure 7) respectively; ; in contrast, as already established, those of V1 and V2 are very close to those of M1 and M2.[20] Since this is not consistent overall with either the model prediction or the sincere voting hypothesis, the next logical step toward explaining these results is to compare behavior.

4.2.3 Behavior.

Unbiased signals. Voting frequencies under V1 and CA1 are very different (see Table 9). In CA1, subjects that receive a red ball veto (by playing *red*) almost 28% of the times. Since vetoing pins down the group decision, doing so is potentially much more costly than other deviations from the model prediction (voting for Blue when receiving a red ball, or conversely). Interestingly, the frequency of veto with a blue ball is also relatively large (5%). These results are difficult to interpret. Assuming that other subjects play according to the model predictions, such a deviation is indeed extremely costly in ex-

[20]Formal tests provide some support to the claim that information aggregation is no different under CA than under U. In the biased case, they are significantly different only when there are 9 blue balls (U does better). In the unbiased case, the results are mixed. They are in favour of CA when there are 0, 1, 5, 6, 7, or 8 blue balls and in favour of U when there are 4. When there are 2, 3, or 9, the difference is not significant.

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

		Unbiased signals			Biased signals		
		% Blue	% Red	% Veto	% Blue	% Red	% Veto
Veto	Blue ball	**96.4**	3.1	0.5	**86.6**	12.4	1.0
	Red ball	3.2	**94.5**	2.3	3.3	12.2	**84.4**
Constr. Abs.	Blue ball	**84.5**	10.6	5.0	**86.7**	11.4	2.0
	Red ball	10.2	**62.2**	27.6	3.8	6.1	**90.1**
Model	Blue ball	100	0	0	100	0	0
prediction	Red ball	0	100	0	0	0	100

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

The sincere voting hypothesis helps organize the data in some way, then, but not completely. Figure 8 displays individual behavior. First, consider CA1, located on the right panel. There is some clustering around the top right corner, which is where sincere voters would appear. Only 6% of the subjects always played accordingly (though this number grows to 10% if we look at agents that vote accordingly at least 80% of the times with both signals). Overall, we observe very dispersed behavior.

Biased signals. Behavior under V2 and CA2 looks more similar (see Figure 8). We see two relevant discrepancies. First, the proportion of subjects with red balls that always veto (the action predicted by the model) is a bit lower in V2 than in CA2. Since playing *red* is the focal action under Veto, this could be result of sincere voting by some subjects. However, even if 4 subjects always vote sincerely, one can hardly argue in favor of a clustering in the corresponding low-right corner. Second, as illustrated in Table 9, Constructive Abstention exhibits a higher frequency of veto with blue balls (2%, versus 1% under Veto). These are small percentages, but the actions associated with them are potentially very costly. They account for the poorer information aggregation score of Constructive Abstention when there are 9 blue balls.

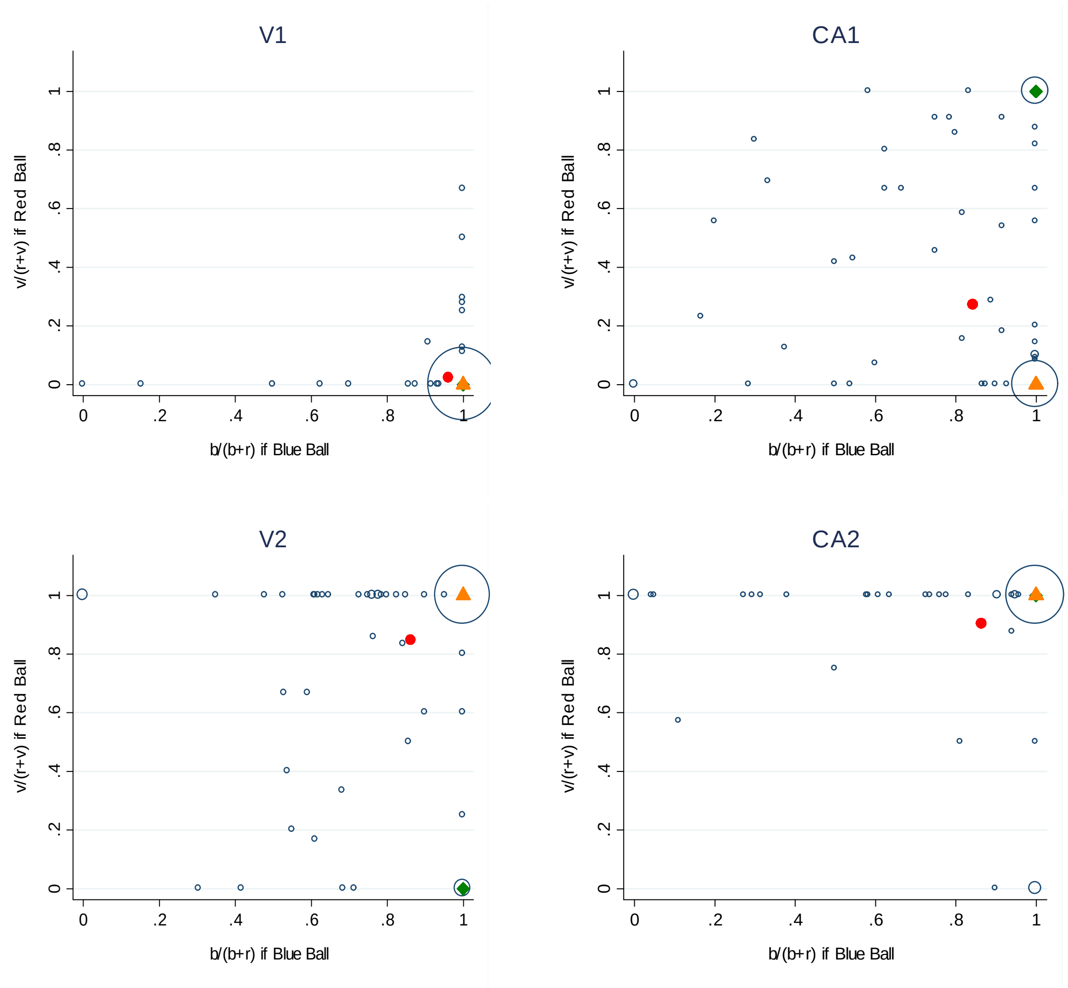

Figure 8: Individual behavior in V and CA treatments. Each hollow circle on the graph corresponds to the observed frequency of play; its size represents the number of subjects who actually adopted that frequency of play. The red circle represents the average frequency of play observed, the orange triangle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

Interpretation. Overall, we cannot explain these findings by pointing to a constant proportion of sincere voting across treatments. Given our notion of sincere voting, an interpretation of our findings is that the proportion of agents playing sincerely can vary based on the way the problem is framed.

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).[21] 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 biased-signal game less complex than the unbiased 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).

5 Conclusion

In this paper, we have explored empirically the performance of different unanimous rules. 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, 2015). Furthermore, we find that, when it is framed as majority rule with veto power, subject behavior falls in line with the model predictions and deviations are not too costly. This provides support for external validity and suggests that it would, indeed,

[21] This 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 perfer to use the phrase "default option".

be beneficial for voting bodies that currently use Unanimity to adopt Veto instead.

However, when Veto is framed as unanimity rule under the constructive abstention regime, we observe more deviations from equilibrium predictions and weaker results. This shows that the way in which a voting rule is presented can have a dramatic impact on the outcome of the vote. A tentative but potentially important takeaway is that when the rule is such that some votes have more weight than others, one must be very careful in choosing which is the focal vote. This is an interesting avenue for future research.

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 (2015a). Divided Majority and Information Aggregation: Theory and Experiment, *Journal of Public Economics*, forthcoming.

[8] Bouton, L., Castanheira, M., and A. Llorente-Saguer (2015b). Multicandidate Elections: Aggregate Uncertainty in the Laboratory, *Games and Economic Behavior*, forthcoming.

[9] Bouton, L., A. Llorente-Saguer and F. Malherbe (2015). Get Rid of Unanimity: The Superiority of Majority with Veto Power, *CEPR Discussion Paper 10408*.

[10] 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.

[11] Dean, Kibris, and Masatlioglu (2014). Limited Attention and Status Quo Bias, mimeo.

[12] Feddersen, T. and Pesendorfer, W. (1997). Voting Behavior and Information Aggregation in Elections with Private Information, *Econometrica*, 65: 1029-1058.

[13] Feddersen, T. and Pesendorfer, W. (1998). Convicting the Innocent: the Inferiority of Unanimous Jury Verdicts, *American Political Science Review*, 92: 23-35.

[14] Fehrler, S., and N. Hughes (2014). How transparency kills information aggregation: Theory and experiment, *mimeo*.

[15] Goeree, J. and L. Yariv (2011). An Experimental Study of Collective Deliberation, *Econometrica*, 79(3): 893-921.

[16] Grosser, J. and M. Seebauer (2013). The curse of uninformed voting: An experimental study, *mimeo*.

[17] Guarnaschelli, S., McKelvey, R., and T. Palfrey (2000). An Experimental Study of Jury Decision Rules, *American Political Science Review*, 94(2): 407-423.

[18] 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.

[19] Herrera, H., A. Llorente-Saguer and J. McMurray (2015). The Marginal Voter's Curse, *mimeo*.

[20] Kagel, J., H. Sung and E. Winter (2010). Veto Power in Committees: An Experimental Study, *Experimental Economics*, 13(2):167–188.

[21] 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.

[22] Le Quement, M. T. and I. Marcin (2015). Communication and voting in heterogeneous committees: An experimental study, *mimeo*.

[23] Maggi, G. and M. Morelli (2006). Self-Enforcing Voting in International Organizations, *American Economic Review*, 96(4): 1137-1158.

[24] Mattozzi, A., and M. Nakaguma (2015). Public versus Secret Voting in Committees, *mimeo*.

[25] McLennan, A. (1998). Consequences of the Condorcet Jury Theorem for Beneficial Information Aggregation by Rational Agents, *American Political Science Review*, 92: 413-418.

[26] 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.

[27] 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.

[28] Nunnari, S. (2014). Dynamic Legislative Bargaining with Veto Power: Theory and Experiments, *mimeo*.

[29] Posner, E. and A. Sykes (2014). Voting Rules in International Organizations, *Public Law and Legal Theory Working Papers*, 458.

[30] Romme, A. G. L. (2004). Unanimity Rule and Organizational Decision Making: A Simulation Model, *Organization Science*, 15(6): 704-718.

[31] Sievers, L. and S. Daws (2014), *The Procedure of the UN Security Council*, New York: Oxford University Press.

[32] Tversky, A. and E. Shafir (1992). Choice under conáict: The Dynamics of Deferred Decision, *Phychological Science*, 3(6):358-361, 1992.

[33] Wilson, R. K., & Herzberg, R. (1987). Negative decision powers and institutional equilibrium: experiments on blocking coalitions. *The Western Political Quarterly*, 40(4), 593–609.

[34] Zamora, S. (1980), Voting in International Economic Organizations, *American Journal of International Law*, 74(3): 566-608.

Appendices

Appendix A1. Proof of Lemma 2

Proof. Under Majority, given that $\Pr(s_B|\omega_B) \neq \Pr(s_R|\omega_R)$, we know from Austen-Smith and Banks (1996) that $\sigma_b(s_B) = 1 = \sigma_r(s_R)$ need not be an equilibrium. It is easy to check that indeed sincere voting is not an equilibrium since voters receiving a blue signal would want to deviate and vote red. The same reason explains why there is no equilibrium such that $\sigma_b(s_B) = 1$, and $\sigma_r(s_R) \in (0,1)$. Therefore, any responsive equilibrium must take the following form: $\sigma_b(s_B) \in (0,1)$, and $\sigma_r(s_R) = 1$.

To characterize the equilibrium, we need to define the expected payoff of the two possible actions for an agent with signal s. The probability of being pivotal (changing the outcome from R to B) in state ω when playing b is:

$$\pi^\omega = \frac{1}{2}\frac{(n-1)!}{\frac{n-1}{2}!\frac{n-1}{2}!}\left(\tau_b^\omega\right)^{\frac{n-1}{2}}\left(\tau_r^\omega\right)^{\frac{n-1}{2}},$$

where $\tau_r^\omega = \Pr(s_R|\omega) + \sigma_r(s_B)\Pr(s_B|\omega)$, and $\tau_b^\omega = \Pr(s_B|\omega)(1-\sigma_r(s_B))$. Given agents' preferences, the expected payoff of action b for an agent with signal s is then:

$$G(b|s) = \Pr(\omega_B|s)\,\pi^{\omega_B} - \Pr(\omega_R|s)\,\pi^{\omega_R}. \tag{1}$$

Given that an r-vote never changes the group decision, the expected payoff for action r is 0: $G(r|s) = 0$.

The following strategy profile

$$\sigma_r^{M^*}(s_R) = 1,\ \sigma_r^{M^*}(s_B) = \eta^* \in (0,1)\,,\text{ and } \sigma_b^{M^*}(s_B) = 1-\eta^*,$$

requires (i) $G(b|s_B) = 0$, and (ii) $G(b|s_R) \leq 0$. Since $G(b|s_R) \leq G(b|s_B)$, we have that (ii) is necessarily satisfied when (i) is satisfied. From (1), we have that $G(b|s_B) = 0$ if

$$\Pr(\omega_B|s_B)\,\pi^{\omega_B} - \Pr(\omega_R|s_B)\,\pi^{\omega_R} = 0$$

From $\tau_r^\omega = \Pr(s_R|\omega) + \eta^*\Pr(s_B|\omega)$, and $\tau_b^\omega = \Pr(s_B|\omega)(1-\eta^*)$, we have

$$\eta^* = \frac{\Pr(s_R|\omega_R) - \left(\frac{\Pr(s_B|\omega_B)}{\Pr(s_B|\omega_R)}\right)^{\frac{n+1}{n-1}}\Pr(s_R|\omega_B)}{\Pr(s_B|\omega_B)\left(\frac{\Pr(s_B|\omega_B)}{\Pr(s_B|\omega_R)}\right)^{\frac{n+1}{n-1}} - \Pr(s_B|\omega_R)},$$

or

$$\eta^* = \frac{\left(\frac{\Pr(s_B|\omega_B)}{\Pr(s_B|\omega_R)} - 1\right)\left(\frac{\Pr(s_R|\omega_R)}{\Pr(s_R|\omega_B)} - \left(\frac{\Pr(s_B|\omega_B)}{\Pr(s_B|\omega_R)}\right)^{\frac{n+1}{n-1}}\right)}{\left(\frac{\Pr(s_R|\omega_R)}{\Pr(s_R|\omega_B)} - 1\right)\left(\left(\frac{\Pr(s_B|\omega_B)}{\Pr(s_B|\omega_R)}\right)^{\frac{2n}{n-1}} - 1\right)}.$$

Finally, we need to prove that $\eta^* \in (0,1)$. Given that $\frac{\Pr(s_R|\omega_R)}{\Pr(s_R|\omega_B)} > 1$ and $\frac{\Pr(s_B|\omega_B)}{\Pr(s_B|\omega_R)} > 1$ are satisfied by assumption, $\eta^* \in (0,1)$ if and only if

$$\frac{\Pr(s_R|\omega_R)}{\Pr(s_R|\omega_B)} > \left(\frac{\Pr(s_B|\omega_B)}{\Pr(s_B|\omega_R)}\right)^{\frac{n+1}{n-1}}.$$

This is necessarily satisfied when $\frac{\Pr(s_R|\omega_R)}{\Pr(s_R|\omega_B)} > \left(\frac{\Pr(s_B|\omega_B)}{\Pr(s_B|\omega_R)}\right)^{n-1}$, which holds by assumption. ■

Appendix A2. Non-parametric tests on payoffs

	Setting 1			Setting 2		
	U	V	A	U	V	A
M	$>$ $z = 2.913$ $p = 0.004$	$=$ $z = 0.893$ $p = 0.371$	$>$ $z = 2.898$ $p = 0.004$	$<$ $z = 2.903$ $p = 0.004$	$<$ $z = 2.892$ $p = 0.004$	$<$ $z = 2.903$ $p = 0.004$
U	–	$=$ $z = 2.913$ $p = 0.004$	$=$ $z = 0.326$ $p = 0.744$	–	$=$ $z = 0.732$ $p = 0.464$	$>$ $z = 2.432$ $p = 0.015$
V	–	–	$>$ $z = 2.898$ $p = 0.004$	–	–	$=$ $z = 1.529$ $p = 0.126$

Table 10: Mann-Whitney tests on the average realized information aggregation in the second half of the experiment.

Appendix A3. 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 behaviour 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 votings frequencies for each subject based on the last 25 periods.[22] Then, for each possible realized signal profile (i.e. for

[22] In 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.

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.[23] Finally, we pool the results by treatment and compute the proportion of group decisions that coincide with that of the fully informed dictator.

Appendix A4. 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

[23] PERHAPS: explain that we keep matching group constants because they constitute an indenpent realization.

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 A] 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 A] ... 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 A] 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.