

DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

CALIFORNIA INSTITUTE OF TECHNOLOGY


PASADENA, CALIFORNIA 91125

AN EXPERIMENTAL STUDY OF JURY DECISION RULES

Serena Guarnaschelli

Richard D. McKelvey

Thomas R. Palfrey



SOCIAL SCIENCE WORKING PAPER 1034

March 2000

An Experimental Study of Jury Decision Rules¹

Serena Guarnaschelli
California Institute of Technology

Richard D. McKelvey
California Institute of Technology

Thomas R. Palfrey
California Institute of Technology

October 1997 (current version March 31, 2000)

¹Support of the National Science Foundation (Grant #SBR-9617854) is gratefully acknowledged. We thank Tara Butterfield for research assistance, John Patty for help in running the experiments, and Tim Reed and Charles Smith for writing the computer program for the experiments. We also thank Tim Feddersen, Susanne Lohmann, Krishna Ladha, the audiences at several academic conferences and seminars, and three referees for their comments.

Abstract

We present experimental results on individual decisions in juries. We consider the effect of three treatment variables: the size of the jury (three or six), the number of votes needed for conviction (majority or unanimity), and jury deliberation. We find evidence of strategic voting under the unanimity rule, where the form of strategic behavior involves a bias to vote guilty to compensate for the unanimity requirement. A large fraction of jurors vote to convict even when their private information indicates the defendant is more likely to be innocent than guilty. This is roughly consistent with the game theoretic predictions of Feddersen and Pesendorfer (FP) [1998]. While individual behavior is explained well by the game theoretic model, at the level of the jury decision, there are numerous discrepancies. In particular, contrary to the FP prediction, we find that in our experiments juries convict fewer innocent defendants under unanimity rule than under majority rule. We are able to simultaneously account for the individual and group data by using Quantal Response Equilibrium to model the error.

1 Introduction

Recent research in political science has addressed from a theoretical point of view the question of how individuals behave in juries. The classic result in this area was the Condorcet Jury Theorem, which dealt with a model of a jury in which all jurors have identical preferences (they want to convict a guilty defendant, and acquit an innocent defendant), but they differ in their probability of making a correct decision. In this setting, the Condorcet Jury Theorem states that if the jury decision is made by majority rule, then the probability that the jury makes a correct decision is higher than that of any individual, and the probability of making a correct decision goes to one as the size of the jury becomes very large.

Recently, Austen-Smith and Banks [1996] challenged the foundations of the Condorcet Jury Theorem by questioning whether any game theoretic basis could be given for the type of behavior assumed by the Condorcet Jury Theorem. They demonstrated that if individuals start from a common prior about guilt of the defendant, and then obtain private information, it is generally not a Nash equilibrium to vote sincerely (i.e., based only on one's private information). Subsequent literature by Wit [1996] and McLennan [1996] reestablished that the conclusions of the Condorcet Jury Theorem still hold if individuals vote strategically according to a symmetric mixed strategy equilibrium of the game.

In a recent paper, Feddersen and Pesendorfer (FP) [1998] analyzed the Nash equilibrium of the Condorcet Jury Theorem, with particular attention to the effect of the decision rule. They compared the unanimous decision rule with decision rules requiring only a majority or super majority to convict. They concluded that unanimous rule results in probabilities of convicting an innocent defendant that are higher than those for majority rule, and which do not go to zero as the number of jurors goes to infinity.

The model of Feddersen and Pesendorfer assumes that the jury decision is a simultaneous move game, in which all jurors vote without any communication beforehand. As Coughlan [1997] shows, if it is possible to have a "straw poll" prior to the vote, then

there exist equilibria in which voters reveal their information in the straw poll, and then vote optimally based on the pooled information in the actual vote. This type of behavior would eliminate the unattractive aspects of unanimity, as then decisions under majority rule should be identical to those under unanimity.

Despite the active theoretical literature on juries, there has been relatively little experimental work to investigate the implications of these theories. The one exception is the paper by Ladha, Miller and Oppenheimer [1996], who run experiments in which a jury makes a sequence of decisions under majority rule. They find evidence of strategic behavior and also evidence that asymmetric Nash equilibria are sometimes played.

In this paper, we run experiments to test the Feddersen-Pesendorfer and Coughlan predictions. In particular, we focus on the questions of whether jurors really vote strategically, and whether unanimity rule leads to more convictions of innocent defendants than majority rule. We are thus concerned with the effects of three treatment variables: size of majority needed for conviction, effect of a straw poll, and the size of the jury. In each of these treatments we consider the individual and group behavior.

Our basic findings are that there is evidence of strategic behavior that is roughly consistent with the Feddersen-Pesendorfer game theoretic model. That is, a large percentage (between 30% and 50%) of jurors vote to convict even when their private information indicates the defendant is more likely to be innocent than guilty. The percentage increases with the size of the jury, as is predicted. While individual behavior is explained well by the game theoretic model, at the level of the jury decision, there are numerous discrepancies. In particular, contrary to the FP prediction, we find that in our experiments juries convict fewer innocent defendants under unanimity rule than under majority rule. We are able to simultaneously account for the individual and group data by using Quantal Response Equilibrium to model the error.

2 The Condorcet Jury Model

The Condorcet jury model is meant to apply to a general class of group decision problems where the members of group have a common interest, but hold different beliefs about the true state of the world. By a common interest, what is meant is that *if* the state of the world were common knowledge, then all group members would agree about which decision to make. The differences in beliefs create an information aggregation problem which creates potential obstacles for the group to reach a consensus and to make the “correct” decision. This class of decision problems has potential application to many real world settings, including juries in criminal and civil trials, corporate strategy decisions by boards of directors or partners, hiring and tenure decisions by faculty committees, examinations judged subjectively by committees, and so forth. The first of these, trial juries was the subject of a recent paper in this *Review* by Feddersen and Pesendorfer (1997), and it is the main motivation for this paper.

The reader should keep in mind several things while reading this paper. First, the model we study is meant to apply to a broad class of settings, and therefore cannot capture all the interesting institutional details of one single setting, like a trial jury. Second, it is an approximation, which leaves out some contextual components that may affect behavior in specific applications. Third, it is meant to be simple, so that we can isolate and study certain phenomena of theoretical interest. Fourth, it is flexible, in the sense that one can analyze the model using a wide range of different assumptions about the degree of sophistication or rationality of the members of the group. In this sense, while it is unquestionably a formal theoretical model, one can explore the implications of bounded rationality as well as the implications of rational choice within the confines of the same model. Thus it provides a very nice framework for conducting experiments to compare rational choice and bounded rationality.

2.1 Model Structure and Notation

We consider a game with a set $N = \{1, 2, \dots, n\}$ of n players (jurors), and let $1 \leq k \leq n$ represent the number of votes needed for conviction. The game begins by nature choosing a state of the world in $\Omega = \{G, I\}$ ¹, with probability s and $1-s$, respectively. The players do not observe the state that is selected, but each obtains some private information about the state. If the true state is G , then each juror observes an independent Bernoulli random variable which is g with probability p , and i with probability $1-p$. If the true state is I , then each juror observes an independent Bernoulli random variable which is i with probability p and g with probability $1-p$. After observing their private information, jurors then vote for one of two actions in $X = \{C, A\}$.² If k or more jurors vote for C the group decision is C , and otherwise the decision is A . The utility $u : X \times \Omega \mapsto \mathfrak{R}$ of each player is defined by $u(A, I) = u(C, G) = 0$, $U(C, I) = -q$, and $U(A, G) = q - 1$, where $0 < q < 1$.

In all of our examples, we will assume that $s = .5$, $q = .5$, and $p = .7$. We will be concerned with two different voting rules, majority rule, in which case k is the least integer greater than $n/2$, and unanimity, in which case $k = n$.

In our analysis, we will distinguish between two kinds of behavior, which we call *naive* and *strategic*. By naive, we mean that voters ignore the group strategy aspect of the decision problem and simply vote as if they were the only juror. In the above setting, that means if they receive a guilty signal they vote to convict, and if they receive an innocent signal they vote to acquit. This is the kind of behavior that was assumed by Condorcet.

The second kind of behavior we consider is *strategic*. Contemporary game theorists, including some political scientists (e.g. Austen-Smith and Banks) argue that we should assume the individuals behave strategically rather than naively. They also prove formally that naive and strategic behavior can have dramatically different logical implications in

¹Read “G” = Guilty, “I” = Innocent

²Read “C” = Convict, “A” = Acquit

the Condorcet jury model. In particular, strategic behavior by jurors is modeled using game theory, which predicts that under certain specific circumstances it is optimal for voters to vote *against* their signal. Our predictions about strategic behavior in these experiments are given by the choice probabilities at Nash equilibria.

We also consider statistical versions of both of these types of behavior. Since we observe all kinds of choices in our experiment, it is necessary to introduce an error component to individual choices. For the case of naive behavior, we do this by assuming that naive subjects vote with their signal with some fixed probability, γ , and make an error with probability $1 - \gamma$. The probability of correct choice is assumed to be independent of the signal received and the same across all treatments. This becomes a free parameter of the naive behavior model which allows us to fit the data to the model by standard estimation methods. For the case of strategic behavior, we incorporate the error structure into the equilibrium concept, by using Quantal Response Equilibrium (QRE). The QRE model (which is explained in more detail below) assumes that players may deviate with some probability from best responses and the probability of deviation depends on the expected payoff difference between the best response and the deviation. We use a Logit parameterization of QRE, which includes a free parameter, λ , which determines payoff responsiveness. Higher values of λ in the strategic model correspond approximately to higher values of γ in the naive model.

2.2 Strategic Behavior: Nash Equilibrium

To characterize equilibria, a strategy for a voter is a function $\sigma : \{g, i\} \mapsto [0, 1]$ taking signals into probability of voting for conviction. There are trivial equilibria to the above game in which voters ignore their information. However, of special interest are symmetric “informative” equilibria to the above game. Symmetric equilibria require that jurors with the same signal adopt the same (mixed) strategy. Informative equilibria are those in which the jurors do not ignore their information.

As shown by Feddersen and Pesendorfer [1998], for the case of unanimity, the unique

symmetric informative equilibrium requires that $\sigma(g) = 1$, and that the probability of the defendant being guilty conditional on player i receiving an innocent signal, and all other players voting guilty must be equal to q . I. e.,

$$q = \frac{(1-p)g_G^{n-1}}{(1-p)g_G^{n-1} + pg_I^{n-1}}$$

where

$$g_G = p\sigma(g) + (1-p)\sigma(i) \tag{1}$$

$$g_I = (1-p)\sigma(g) + p\sigma(i) \tag{2}$$

are the probabilities that an individual votes to convict when the defendant is guilty or innocent, respectively. Using $\sigma(g) = 1$, this implies that

$$\sigma(i) = \frac{D_n p - (1-p)}{p - D_n(1-p)}$$

where

$$D_n = \left(\frac{(1-q)(1-p)}{qp} \right)^{\frac{1}{n-1}}$$

Also, the probability that the jury votes incorrectly to convict an innocent defendant ($Pr[C|I] = (g_I)^n$), and to acquit a guilty defendant ($Pr[A|G] = 1 - (g_G)^n$), are determined from the above formula. Table 1 (b) gives the values of $\sigma(i)$, $Pr[A|G]$, and $Pr[C|I]$, respectively for unanimity rule for certain values of n .

For majority rule, the formulae are slightly more complicated, but the symmetric equilibrium is simpler, namely $\sigma(g) = 1.0$ and $\sigma(i) = 0.0$, regardless of the size of n .³ We can then compute the corresponding probabilities of convicting an innocent defendant

³This is true for both the case of odd and even n , although it is somewhat more difficult to prove for even n since the rule is no longer symmetric ($n/2$ votes to acquit suffices to acquit, but more than $\frac{n}{2}$ votes to convict are necessary to convict).

($Pr[C|I] = \sum_{k > \frac{n}{2}} \binom{n}{k} (1-p)^k p^{n-k}$) and of acquitting a guilty defendant ($Pr[A|G] = \sum_{k \leq \frac{n}{2}} \binom{n}{k} p^k (1-p)^{n-k}$). These numbers are reported in Table 1 (a) for the parameter values used in our experiment.

In the case of a straw poll, the game expands so that the jurors now have two votes. A strategy is now a specification of how to vote in the straw poll, as a function of the juror's signal, and then a specification of how to vote in the final vote as a function of the signal and the observed outcome of the straw poll for the final vote. Let $\sigma_0 : \{g, i\} \mapsto [0, 1]$ be the probability that the voter votes for conviction on the straw poll. Define $N_0 = N \cup 0$ to be the possible outcomes of the straw poll. Then let $\sigma_1 : \{g, i\} \times N_0$. For the case of a straw poll, Coughlan (1997) shows that for the parameters used here, there is a fully informative equilibrium in which all jurors reveal their signals on the straw vote (i. e., $\sigma_0(i) = 0$, and $\sigma_0(g) = 1$), and then vote based on the majority outcome of the straw vote in the final vote (i. e., $\sigma_1(s, M(\sigma_0)) = M(\sigma_0)$, for all signals $s \in \{i, g\}$, where $M(\sigma_0)$ is zero or one or some appropriate mixing probability according to whether the majority outcome of the straw poll was to convict or acquit, or a tie).

2.3 Strategic Behavior: Quantal Response Equilibrium

The above solutions all assume no error. McKelvey and Palfrey (1995, 1998) propose a general way to incorporate decision error, Quantal Response Equilibrium (QRE), which is a statistical version of Nash equilibrium. The basic idea is that it is unreasonable to expect individuals to always behave perfectly in accordance with rationality, and always choose best responses to the other players. Instead, they choose better responses more often than worse responses. Thus, rather than a deterministic model, QRE specifies a probability distribution, $\sigma^*(\cdot)$ over actions. The probabilities are ordered by the expected payoffs of the actions, $EU(\cdot)$ according to some specific function, called a quantal

response function, which is just the statistical version of a best response function.⁴ Thus, actions with higher expected payoff will be played more frequently and actions with lower expected payoffs will be played less frequently. So, for any individual i and any pair of actions available to i , say a and b , $\sigma_i^*(a) > \sigma_i^*(b)$ if and only if $EU_i(a) > EU_i(b)$. The “E” of QRE stands for equilibrium, in the sense that the expected payoffs, $EU(\cdot)$, are themselves derived from the equilibrium probabilities, $\sigma^*(\cdot)$. One can think of an iterative process in which a given profile of choice probabilities for all the players results determines a profile of expected payoffs for every action, which in turn (via the quantal response function) generate a new profile of choice probabilities. A QRE is just a fixed point of this iterative process. Thus QRE retains the rational expectations flavor of Nash equilibrium, but relaxes the assumption that players choose optimal responses.

There are several ways to justify a formal model that has the above properties. The idea that individuals choose stochastically rather than deterministically has been proposed for a long time, for example to motivate reinforcement learning and discrete choice econometrics. Alternatively, one can “rationalize” stochastic choice if players have stochastic utility functions. Harsanyi (1973) proposed a model in this vein where the game payoff matrix is viewed as just an approximation of the utilities of the player over outcomes in the game, and each player’s actual utility vary about these means according to some statistical distribution. Thus, in a QRE, for every action an individual might choose, there is a privately observed payoff disturbance for that action, and one then looks at the Bayesian equilibrium to the corresponding game of private information. This is equivalent to “smoothing out” the best response curves of the players and then looking at a fixed point of these smoothed out response functions, which is exactly what QRE does.

Here we focus on a particularly tractable form of QRE, called *Logit QRE*, where the quantal response functions are Logit curves. That is, for any pair of actions a and b we

⁴In this paper, we use a Logit specification for the quantal response function. This is explained below.

let $\ln[\sigma(a)/\sigma(b)] = \lambda \cdot [EU(a)/EU(b)]$, where λ is a response parameter.

In our game players have different information, and in some cases make a sequence of decisions. Therefore, we turn to the extensive form of the game and represent profiles of action probabilities as behavior strategies, and apply QRE to the “agent” form of this game.⁵ Formally, let $p = (p_1, \dots, p_n)$ be a completely mixed profile of behavior strategies, where $p_i = \{p_{ijk}\}$ and p_{ijk} is the probability that player i , with signal $j \in \{i, g\}$ votes for $k \in \{c, a\}$. Let $\bar{u}_{ijk}(p)$ denote the expected utility to player i from taking action k with signal j , given p . Then p^* is a Logit equilibrium if and only if, for all i, j, k ,

$$p_{ijk}^* = \frac{e^{-\lambda \bar{u}_{ijk}(p^*)}}{\sum_l e^{-\lambda \bar{u}_{ijl}(p^*)}}$$

where again $\lambda > 0$ is a free parameter determining the slope of players’ logit response curves. As we vary λ from 0 to ∞ , we can map out a family of QRE’s which correspond to different levels of rationality (or, more precisely: “payoff responsiveness”). When $\lambda = 0$, response curves are completely flat, so all strategies are used with equal probability (pure error, or zero rationality). When λ approaches ∞ , logit response curves converge to standard best response curves, so players use only optimal strategies (no error, or perfect rationality). This family of QRE’s has several interesting properties that are described in McKelvey and Palfrey (1995, 1998). For example, if we consider a convergent sequence of logit equilibria for a sequence of λ values converging to ∞ , the limit point must be a Nash equilibrium of the underlying game. In this sense, Nash equilibrium is just a very special *boundary* case of QRE, which corresponds to perfect rationality.

We also consider an alternate model of errors combined with strategic behavior, called the Noisy Nash Model (NNM), which also looks at statistical variation around the Nash equilibrium, but differs from QRE in two ways. First, it does not incorporate the rational expectations assumption of QRE. The NNM model assumes that individuals follow Nash behavior with some fixed probability γ (to be estimated), and choose randomly⁶ with

⁵See McKelvey and Palfrey (1998) for details.

⁶That is, they vote to convict or to acquit with equal probability.

probability $1 - \gamma$. Like QRE, in the limiting case when γ approaches 1, the prediction approaches Nash equilibrium. However, for intermediate values of the error term, it will often differ from QRE. The difference is twofold. First, NNM assigns the same probability of deviating from Nash equilibrium ($1 - \gamma$) regardless of the expected utility loss from such a deviation. Second, NNM is not an equilibrium model. Recall that QRE is defined as a fixed point of in terms of choice probabilities and Logit responses, That is, each player’s errors (deviations from Nash play) affect the expected payoffs of all the other players, and hence will indirectly affect all other players’ Logit responses. In contrast, under NNM, there is no such “feedback”, so that one player’s deviations from Nash play has no indirect effect on any other player’s deviations from Nash play. See McKelvey and Palfrey (1998), McKelvey, Palfrey, and Weber (forthcoming), and Fey, McKelvey, and Palfrey (1996) for further discussion of the differences between QRE and NNM. Similar to QRE, we can map out a family of NNM predictions by varying the free parameter γ from 0 to 1.⁷ When $\gamma = 0$, all strategies are used with equal probability (pure error, or zero rationality). When γ approaches 1, predictions of NNM converge to Nash equilibrium.

The symmetric portion of the Logit AQRE correspondence for jury sizes $n = 3$ and $n = 6$ and both the majority and unanimity voting rules are displayed as the thick solid curves in Figures 1 and 2.⁸ Each graph is on the unit square of mixed behavior strategies of a representative player. The horizontal and vertical axes of each of the four graphs correspond to the probability of voting to convict, given innocent and guilty signals, respectively. At the center of each unit square is the “pure error” Logit equilibrium that corresponds to $\lambda = 0$. As λ increases, the equilibrium curves converge to the symmetric Nash equilibrium, which is on the upper boundary of the unit square (the upper left corner, in the case of majority rule).

In a similar fashion, one constructs the correspondences defined by the NNM model

⁷As with QRE, we limit attention to the NNM corresponding to the symmetric mixed strategy Nash equilibrium.

⁸Later in the paper we present and discuss the asymmetric components of this correspondence. The large dots in the figures are explained in the data analysis section.

and the Naive (non-strategic) model, by varying γ between 0 and 1. Referring to Figure 1, the NNM correspondence is simply the line segment (not drawn) connecting the center of the probability square (pure error) to the symmetric Nash equilibrium. The Naive model correspondence is simply the line segment connecting the center to the upper left corner of the strategy space. This vertex corresponds to “honest” voting.

In the appendix we are able to fully characterize, and compute the symmetric quantal response equilibrium correspondence for the 3-person unanimous jury game with a straw vote. The characterization of the majority rule QRE correspondence is similar, and is not included in the appendix. Unfortunately, our efforts to compute the symmetric quantal response equilibrium correspondence for the majority rule juries with a straw vote were unsuccessful. Also, we found that the 6-person jury game with a straw vote is too complex to compute the QRE correspondence, using our numerical methods.

3 Experimental Design

We conducted a total of four experiments, using as subjects undergraduate and graduate students at the California Institute of Technology. Each experiment used twelve subjects (plus one subject that was used as a monitor). The experiment was divided into four sessions. Between sessions, two treatment variables were varied, the decision rule (Majority or Unanimity) and the existence of a straw poll (Yes or No). The treatment variables were varied according to the design in Table 2, which gives the particulars of each experiment. In each session the subjects participated in a sequence of fifteen “matches”⁹. In each match the subjects were randomly matched in groups of size n , where n was one of the treatment variables, and a jury game similar to that described in the previous section was conducted. In that table, for each session, the values of each treatment variables and the number of sessions is given. For example “U/N (15)” means 15 matches with

⁹The last three sessions of Experiment CJ1 were truncated to ten matches each, due to one particularly slow subject.

Unanimity decision rule and No straw poll. All matches in the same experiment were run with the same number of subjects. Subjects were paid in cash at the end of the experiment. Subjects were paid a “show up fee” of \$5.00, plus whatever they earned in the experiment.

In the experiments that were conducted, the subjects were not told that the experiment was intended to represent a jury decision. The states of the world were called the Red Jar and Blue Jar instead of Guilty and Innocent, and instead of choosing to Convict or Acquit, the subjects were instructed to guess whether the true jar was Red or Blue. The complete instructions are given in the appendix.

Briefly, each match proceeded as follows: The subjects were told that there were two jars, a Red Jar and a Blue Jar. The Red jar contains seven red balls and 3 blue balls, and the Blue Jar contains seven blue balls and 3 red balls. One of the jars was selected for each group. The jar which was selected was determined by the roll of a die by the monitor (a subject chosen at random from the group of subjects at the beginning of the experiment).¹⁰ The subjects were not told which jar was selected, but were each allowed to choose one ball at random from the jar that was selected.¹¹ After choosing a ball, they then voted, for either the Red Jar or the Blue Jar. Two decision rules were investigated; majority rule and unanimity. The decision rule, which had been explained to them prior to the session, was used to determine the group decision, and their payoffs were determined based on whether the group decision was correct or incorrect. They received fifty cents if the group decision was correct, and five cents if it was incorrect.

¹⁰The die was rolled once for each group in each match, so in each match, different groups could have different states.

¹¹This was accomplished by placing the balls in a random order on their computer screen, with the colors hidden. Subjects then used the mouse to select one of the balls and reveal its color. To convince them that this procedure was conducted honestly, prior to the experiment, we generated the order of the samples for each match, each group, each possible state, and each subject. The samples in the experiment were generated according to this list. Subjects recorded which ball they selected in each match, and were free to peruse the list after the match to verify that there were the correct number of balls of each color, and that the ball they selected was of the correct color.

4 Results

4.1 Jury behavior without deliberation

Table 3 (a) shows the realizations of the voter strategies for the case when there is no straw poll. Recall that in this instance the prediction is that under majority rule, voters should vote the same direction as their signal (convicting with a guilty signal, and acquitting with an innocent signal). Under unanimity, the prediction is that with a guilty signal one should vote to convict, and with an innocent signal, one should mix, voting with a probability of either .314 or .651 to convict, depending on whether the group size is 3 or 6.

The data provide some support these predictions, for both the majority rule treatment and the unanimity treatment. Under majority rule, the subjects vote the same direction as their signals over 94% of the time, with the only exception being the case of innocent signals in a jury of size 6, in which case the subjects err 21% of the time. This seems like a surprising result, but can be explained by the fact that the Nash equilibrium in the 6-person majority rule experiments is a *weak* Nash equilibrium. In the symmetric equilibrium of that game, those voters receiving innocent signals are indifferent between voting to acquit and voting to convict. In the QRE, this indifference leads to a prediction that voters with innocent signals will vote against their signal more frequently than voters with guilty signals, for every positive value of λ .

For the case of unanimity, voters frequently vote to convict when they get an innocent signal. Jurors with a guilty signal still tend to vote guilty – strongly so for a jury of size 3, and less strongly with a jury of size 6. However, when they get an innocent signal, the subjects vote to convict at a rate of 36% for the 3 person jury, and 48% for the 6 person jury. For the case of the three person jury, this is very close to the Nash predicted value of .314. For the case of a six person jury, the rate is significantly below the predicted value of .651.

Since the Nash equilibrium of the game requires that in some cases a pure strategy

is adopted, any observations in which subjects do not follow that strategy are enough to statistically reject the Nash equilibrium as a model of behavior. Thus, any game theoretic model that is to explain the data must incorporate a model of where error comes from. In section 2, we proposed 3 alternative models of behavior that incorporated error: *Naive (non-strategic)*; *Logit QRE*; and *NNM*.

Table 4 gives the results of estimating the free parameters in these three models. Note that because of the symmetry of the game, all three models make identical predictions for the three person majority rule case. Also, the aggregate choice frequencies from all the non-deliberation data are superimposed (large dots) in Figures 1 and 2. The first (and perhaps most important) thing to observe is that the Naive model does very poorly. In other words, voters are behaving as if they understand the strategic subtleties of the decision problem. For the NNM model we estimate γ to be in the 0.90 to 0.93 range for three person. Since $(1 - \gamma)$ is the probability of choosing randomly, this corresponds to an error rate of less than five percent. For six person juries, we estimate γ to be in the 0.75 to 0.78 range, corresponding to an error rate slightly greater than ten percent. For the case of six person juries, the QRE fits significantly better than the NNM under both majority and unanimity rules. In the case of three person juries under unanimity, the fits of QRE and NNM are almost identical, with no significant difference between the two. The Naive model is rejected in favor of the QRE for all treatments where the two models make distinct predictions.

4.2 Jury Behavior with deliberation

As is evident from Table 3(b), in juries with a straw poll, the final vote can no longer be predicted as well using the equilibrium of the Feddersen-Pesendorfer model. On average individuals getting a guilty signal vote to acquit about 15% of the time, independent of the treatments. Those getting an innocent signal vote to convict (against their signal) between 16 – 37%, and this percentage depends on the treatment. This is a higher rate of voting against their signals than for those who obtain a guilty signal, but we do not

get the same differences between majority rule and unanimity as we had in the case of no straw poll.

Table 5 presents the analysis of the straw poll sessions based on the Coughlan equilibrium. Recall that this equilibrium predicts that subjects will reveal their signal in the straw poll, and then vote based on the majority outcome of the straw poll in the final vote. Here we see that the subjects for the most part do use the straw poll to reveal their signal. Over 90% of the subjects in every cell (except one cell which is 89.7%) reveal their signal in the straw poll. Possibly of some interest is the additional finding that false revelation of innocent signals occurs with greater frequency than false revelation of guilty signals in all four treatments. Overall, false revelation of innocent signals is about twice as frequent as false revelation of guilty signals.

In the final vote, when the outcome of the straw poll does not end up with a tie, in all treatments voters vote with the public signal 84% of the time or above. One might expect the numbers to be higher here. In an equilibrium of the Coughlan type, individuals should ignore their own signals, and only pay attention to the public signal. Table 6 gives the result of a probit analysis of the final vote against the individual's private signal and the publicly available information. For the public information for an individual (the variable "PubInfo"), we use the number of other individuals who voted to convict in the straw poll plus the information of the individual (1 for a guilty signal, 0 for an innocent signal). If voters are following the Coughlan equilibrium, then they should base their vote on the public information, and ignore their private information. The probit analysis indicates that this is not the case. The individual's private information has a significant effect on the final vote for all combinations of treatment variables.

There are several possible explanations for why the straw poll does not work perfectly. The simplest explanation we have is based on the results of the treatment without a straw poll. Even in the simplest, 3-person majority rule juries, some voters fail to vote sincerely. This source of error causes a small effect if there is no straw poll, and the QRE model showed in the last section how these small effects could be accounted for using an

equilibrium model with errors. However, with a straw poll, the relatively small effects of errors and the initial voting stage become compounded in the second stage since the noisy behavior in the straw poll means that voters are not sure how to interpret the results of the straw poll. If individuals believe that there is some likelihood that others will not perfectly follow the first stage equilibrium, then they *should* give some weight to their own private information on the second stage. This in turn has a snowball effect on straw poll behavior because they know that other voters will be uncertain how to interpret their (straw) vote. Since QRE is an equilibrium model, it can capture this effect of compounding errors.

Appendix 2 presents a QRE analysis of the straw vote game for the 3-person unanimity treatment, including a table that presents the maximum likelihood estimates of λ , and γ . Our computational algorithm was unable to compute the QRE correspondence for majority rule game or the 6-person games. The main findings is summarized as follows, with a more detailed account contained in the appendix.

There are several important features of the QRE correspondence. First, for higher values of λ , there are multiple symmetric equilibria, corresponding to the various Nash equilibria. Second, the graph is not well behaved, as it contains two points of bifurcation. Third, one of these corresponds to the informative equilibrium studied by Coughlan [1998], and it is the component of the QRE correspondence that most closely matches the data. Fourth, this component has the feature that voters are conditioning their final vote on their own signal as well as the vote outcome in the final round. In particular, for any fixed number of votes to convict in the straw vote, the probability of voting to convict in the final stage is higher if one observed a guilty signal than if one observed an innocent signal. Fifth, the probability of voting to convict in the final vote is monotonic in the number of straw votes to convict.

These features of the QRE are consistent with the data, and consistent with the simple probit analysis of Table 6. However, while the main qualitative predictions are QRE are found in the data, the quantitative fit is less successful. In fact, the maximum

likelihood fit of the NNM is better than the fit of the QRE model. Further details and discussion are in Appendix 2.

4.3 Jury Accuracy

Table 7 presents a summary of the accuracy of the final decision of the jury as a function of the experimental treatment variables. We compare the actual data with the Nash equilibrium predictions of error rates, given in Table 1. First, in the experimental data, under unanimity the probability of voting to convict an innocent defendant goes down from .190 to .029 in our data as the jury becomes larger (this difference is significant at the .05 level using a difference of proportion test, with a t value of $t = 2.223$). The Nash theory predicts the opposite (from .14 to .19). Furthermore, the error rate for six person unanimous juries is lower than the error rate for majority juries (.03 vs .30), in the innocent state. This difference (significant at the .01 level, with $t = 2.924$) is also counter to the Nash theory. As for majority juries, error rates decline with larger juries in the guilty state, and increase with larger juries in the innocent state. This is exactly the opposite of what is predicted to happen in the Nash equilibrium.

We view the contradictions with the aggregate theoretical predictions as surprising, especially given that the individual choice frequencies are not very different from the theoretical choice frequencies. We interpret this to mean that the accuracy of jury decisions is not a robust phenomenon. That is, small changes in individual choice behavior can result in large changes in the probability of an erroneous judgment. This is especially true for unanimous juries, where a small amount of juror decision error can produce a much larger number of acquittals than is predicted by the Nash equilibrium.

As evidence in support of the above claim, we have computed the expected jury accuracy under the QRE model. These are reported in Table 7 (a). To compute these values, we used the maximum likelihood QRE estimates of $\sigma(i)$ and $\sigma(g)$ from Table 4, and substituted in to equations (1) and (2) to get values of g_G and g_I . Then the probabilities $P(C|I)$ and $P(A|G)$ are computed by the corresponding binomial formulas

used in Table 1. We see that the jury accuracy implications of the QRE estimates match the data better than the Nash predictions. All of the above discrepancies except one are explained by the QRE predictions. In particular, the QRE predicts that under unanimity, the probability of convicting an innocent defendant should go down (from .19 to .07) as the size of the jury increases. This is consistent with the data.

Turning now to the experiments with a straw poll, we find that the presence of a straw vote increased the accuracy of judgments in the guilty state, but there was essentially no effect in the innocent state. Thus, it appears that the probability of convicting an innocent defendant does not improve with deliberation, but the probability of acquitting in the guilty state declines substantially. Comparing to the theoretical predictions, in this case we can compute the Nash predictions. With a straw vote, the fully informing Nash equilibrium predicts that both majority and unanimous juries should have identical accuracies to each other, and they should also be identical to the case of majority rule without a straw poll. This is true since the equilibrium strategy is to reveal your signal in the straw poll, and then for all voters to vote on the final ballot based on whether the number of reported guilty signals exceeds $\frac{n}{2}$.

4.4 Jury Accuracy in the Logit Equilibrium

As pointed out earlier, the game-theoretic analysis of Feddersen and Pesendorfer (1997) implies that large unanimous juries will convict innocent defendants with fairly high probability, while the probability of such errors in large majority juries will vanish. Even in juries of size six, the probability of convicting an innocent defendant is predicted to be higher with unanimity rule than with majority rule. This does not happen in our experiment. We find in both the three and six person juries that errors of this kind are more prevalent with majority juries than with unanimous juries.

We show here that the Logit equilibrium predicts that we should be observing this. Specifically, we can show that for any Logit parameter, $\lambda < \infty$, the probability of conviction goes to zero in large unanimous juries. In contrast, for any $\lambda > 0$, the probability

of a majority-rule jury error in the symmetric informative Logit equilibrium goes to 0 as jury size increases, regardless of the state of the world. We first show that the probability a voter votes innocent in a Logit equilibrium is bounded below by an expression that is independent of n .

Theorem 1 *Fix $\lambda < \infty$. For every $\delta > 0$, there exists $N(\delta, \lambda)$ such that for all $n > N(\delta, \lambda)$, the probability of acquittal in any Logit AQRE is greater than $1 - \delta$ regardless of whether the defendant is innocent or guilty.*

Proof. Writing p_{iga}^* and p_{igc}^* for the probability voter i votes to acquit and convict, respectively with a guilty signal, it follows that in any logit AQRE,

$$p_{iga}^* = \frac{1}{1 + e^{\lambda[\bar{u}_{igc}(p^*) - \bar{u}_{iga}(p^*)]}} > \frac{1}{1 + e^\lambda}$$

since $\bar{u}_{igc}(p^*) - \bar{u}_{iga}(p^*) < 1$. A similar argument establishes that $p_{iia}^* > \frac{1}{1+e^\lambda}$.

The theorem now follows immediately, because the lower bound on both p_{iga} and p_{iia} are independent of n . ■

Theorem 2 *Fix $\lambda < \infty$, and consider the Logit AQRE of the unanimous jury game with a straw poll. For every $\delta > 0$, there exists $N(\delta, \lambda)$ such that for all $n > N(\delta, \lambda)$, the probability of acquittal is greater than $1 - \delta$ regardless of whether the defendant is innocent or guilty.*

Proof. The argument from the previous theorem applies to the final stage of the straw poll in exactly the same way as it applies to the case with no straw poll. Therefore, the lower bound on p_{iga} and p_{iia} identified in the previous theorem is the same, and is independent of n . The result follows immediately. ■

Therefore, we see that (Logit) equilibrium behavior is entirely consistent with traditional jurisprudential theory that argues for unanimous juries as a protection against the conviction of innocent defendants. The reason is that the Nash equilibrium conviction/acquittal probabilities are not robust to decision errors. In order for the probability

of convicting an innocent to increase, the probability of voting to acquit must go to zero extremely fast in the number of jurors (on the order of $1/n$) since the probability of conviction is equal to $(1 - p_{ia}^*)^n$. If the probability of voting to acquit goes to 0 any slower than this, then defendants will always be acquitted by large juries. That is, unanimous juries will become completely uninformative.

In our unanimous jury data, we find a couple of interesting facts that match up with these results. First, both with and without straw polls, innocent defendants are wrongly convicted *less* frequently in our large juries than our small juries. Second, both with and without straw polls, the probability of acquitting a guilty defendant *increases* in n . Both of these are the opposite of what the Nash theory predicts, but are consistent with the Logit equilibrium, which predicts that unanimous juries will become more heavily biased toward acquittal as jury size increases. With majority rule this acquittal bias does not occur, but at the nontrivial cost of roughly 50% higher wrongful conviction rates of innocent defendants.

4.5 Individual Behavior, Heterogeneity, and Asymmetric Equilibria

The theoretical work underlying this experiment focused entirely on symmetric informative equilibria. As noted in passing earlier, there are many equilibria in these voting games, all of which satisfy standard refinement criteria such as perfection, properness and stability. In this section we investigate whether the data could be interpreted as evidence either for uninformative equilibria or for asymmetric equilibria.

With regard to uninformative equilibria, there is clear evidence against this. The unique uninformative equilibrium for the unanimity game has all voters always voting to acquit, independent of their actual signal. Table 7 clearly shows that voting is in fact very informative in the unanimity games with and without straw polls. Thus we conclude that the data does not provide evidence of this kind of behavior.

The question of asymmetric equilibria is more subtle, and more problematic as well.

A necessary condition for the presence of asymmetric equilibria is that there be some evidence of heterogeneity in the observed decision rules of different jurors, within the same treatment. We find strong evidence for heterogeneity in our data. As a simple way to categorize voter behavior, we classify voters into one of three strategy-types, based on how they vote when they receive an innocent signal.¹² Strategy-type-1 jurors always vote sincerely; that is, $\sigma(i) = 0$. We call these honest voters “Sincere.” Strategy-type-2 voters mix when they receive an innocent signal; that is, $\sigma(i) > 0$. We call these voters “Mixers.” Strategy-type-3 jurors always vote to convict, independent of their signal; that is, $\sigma(i) = 1$. We call these voters “Rednecks.”¹³ To implement this classification, we simply use actual frequency of voter choices.

Table 8 presents the breakdown of strategy types by session. Overall, 25% of the voters are always sincere, 56% mix, and 19% always vote to convict.

The numbers in parentheses indicate the number of guilty votes and the total number of innocent signals observed by all voters classified in that cell. For example, in session 1, there were four voters who always voted sincerely when they received an innocent signal in the unanimity/no-deliberation treatment. Those voters received a total of 31 innocent signals and voted to acquit in every instance. In that same session, seven subjects mixed – that is, they neither always voted to acquit nor always voted to convict when they received an innocent signal. That set of voters received a total of 57 innocent signals and voted to convict 24 times.

To illustrate the range of possible asymmetric equilibria in jury games, consider the simplest case of three juror unanimity juries with no deliberation. In this case there is an equilibrium with two Rednecks and one sincere voter. To see this is an equilibrium, first look at the Rednecks. Either of them is pivotal if and only if the sincere voter votes to convict, which happens only if the sincere voter received a guilty signal. Thus the

¹²Since nearly all jurors always vote to convict with a guilty signal, we do not break down the individual strategy types based on that.

¹³There are no voters who always vote to acquit.

best response is to vote to convict with a guilty signal, and to voter either to convict or to acquit (indifferent) with an innocent signal. For the sincere voter, since the other two voters are always voting to convict, conditioning on being pivotal is uninformative, so sincere behavior is strictly optimal for either signal. Using similar reasoning, one can easily show that for *any* probability $p \in [0, 1]$, there is an equilibrium with one sincere voter, one redneck, and one Mixer voting to convict with a innocent signal with probability p (and always voting to convict with a guilty signal). An even wider range of asymmetric equilibria exist with 6 person juries. Ladha[1998] has shown that the pure strategy asymmetric Nash equilibria maximize the accuracy of information aggregation in jury games.

While we have not fully characterized the asymmetric equilibria in this paper, the point is clear: There are lots of equilibria, and so the empirical restrictions of equilibrium behavior are limited. But we point out that in all of these equilibria for three person juries, the aggregate probability of a vote to convict, given an innocent signal, is above $1/3$.

The issue of asymmetric equilibria is further complicated if one takes into account the possibility that subjects may make errors. In Figure 3, we have plotted the full QRE correspondence for the three person unanimity game under the assumption that iteratively dominated strategies are not played (i. e., a player receiving a guilty signal always votes guilty). Even under this assumption, it is evident from the figure that the full QRE correspondence is quite complicated. First, it is evident that the QRE correspondence has several places where it bifurcates. This is not a feature of the QRE correspondence, and happens in this game because the game is symmetric, and hence not a generic game. In generic games, it is shown in McKelvey and Palfrey [1995, 1998] that the principal branch of the QRE correspondence selects a unique equilibrium (i.e, there is no bifurcation) . For non generic games, such as the ones for these jury games, such a selection is no longer possible. In this case, the forks, or bifurcations in the QRE correspondence are places where the equilibrium correspondence branches to connect to

the asymmetric equilibria. Thus, if we drop the focus on symmetric equilibria, the QRE gives us no guidance as to which asymmetric equilibrium should be selected. Further, if subjects make errors, then in addition to the limiting points of the correspondence, one predicts that other points on the correspondence could occur. If there is a failure to coordinate on any equilibrium, then matters are further complicated.

Figure 3 deals only with the three person jury. For the case of a six person jury, things become even more complicated. Figure 4 plots the symmetric part of the QRE correspondence for the six person jury using unanimity rule. This graph is the same as that in Figure 2 (b), but viewed from a different projection into the coordinate axes. In this view, we see that even the principal branch of the symmetric part of the QRE correspondence is not monotonic in λ . Thus, if one follows the QRE principal branch of the QRE starting at $\lambda = 0$, one reaches a point (at about $\lambda = 15$) where the curve bends backward for a while. In this region, the symmetric branch of the QRE is multiple valued. For the six person case, we have not been able to compute the full QRE correspondence that would include the asymmetric QRE. But since there is no unique selection of a symmetric QRE for some values of λ , this may even further increase the probability of a lack of coordination, and increase the tendency towards asymmetric behavior.

5 Learning

The jury game we implemented in the laboratory is very complicated. We ask subjects to make choices based on limited partial information in an uncertain environment with asymmetric information. They are not guided by any natural clues that would make it easy to connect the group decision problem they are solving to real-life situations that they have experienced. This is intentional, since we do not wish to bias their decisions or distort their induced preferences in ways that are difficult to predict or measure. Furthermore, in order to make the most of this information, they need to anticipate how other subjects in the room will make decisions based on their (different) partial

information.

Given the complexity of the task, we expect to find trends in the data. For example, the concept of conditioning your decision on the event that your vote is pivotal is a very subtle idea, yet is the main idea underlying the equilibrium of these jury games. Thus it is conceivable that in early rounds subjects are less aware of this equilibrium effect and experience is required before it affects their decision making. The difficulty of the task was one of the factors motivating the repeated trial feature in our design.

For the majority rule experiments, the aggregate data are a close match to the equilibrium prediction that voters will be sincere. This is not surprising since the simple intuitive rule of voting sincerely corresponds to the optimal decision rule when one conditions on being pivotal. In the 6-person juries with majority rule, we find significantly more deviations from sincere voting, but this can be explained by the fact that for 6-person juries, voting sincerely is only a weak equilibrium. That is, with an innocent signal, a juror is indifferent between voting guilty and voting innocent.

To test whether the observed phenomenon of strategic voting is partly the result of learning, we focus on the decisions by voters with innocent signals in the unanimity treatments with no straw poll. This is the only case in our experiment where voters are predicted to vote against their signal in the Nash equilibrium. For these voters, we compare the frequency of voting to convict for the first five rounds of data and the last five rounds of data. What we find is that for all four sessions, the frequency of voting to convict in the last five rounds is higher than in the first five rounds and it is closer to the Nash equilibrium prediction.

As expected, the time trend data from other voter types in other treatments does not show much of a pattern. Breaking it down by session, treatment and signal (a total of 28 cases), we find that in 61% of the cases the decision frequencies move closer to the Nash equilibrium or remain constant¹⁴ in the later periods, while in 54% of the cases

¹⁴In all but one of the cases where there is no change between the first and last periods, all the early data are exactly at the Nash equilibrium.

the decision frequencies are further from the Nash equilibrium in the later periods.

6 Conclusions

We find that in these experimental juries, there is a definite tendency to vote strategically in the case of a unanimous rule and no straw vote. With 3-person juries, voters with innocent signals vote more than one-third of the time for conviction. This is very close to the Nash equilibrium prediction. With 6-person juries, this frequency increases significantly, also as predicted by the theory, although the increase is not as great as Nash equilibrium predicts.

While individual behavior is explained well by the game theoretic model, at the level of the jury decision, there are numerous discrepancies. In particular, contrary to the FP prediction, we find that in our experiments juries convict fewer innocent defendants under unanimity rule than under majority rule, and larger juries convict fewer innocent defendants than smaller juries under unanimity. We are able to simultaneously account for the individual and group data by using Quantal Response Equilibrium to model the error.

With deliberation, which we implement experimentally with non-binding straw polls, jurors tend to reveal their signals in the straw poll, but not to the full extent predicted by the theory. Most voters then use the information from the straw vote to make their final vote. Thus, introduction of straw polls largely (but not completely) eliminates the strategic voting in the case of unanimous juries. Surprisingly, straw polls lead to *more* strategic voting with majority juries than with unanimous juries. The fact that there is a residual of strategic revelation even with straw polls results in some departure from the “sincere” equilibrium path of the final stage. As expected, introduction of straw polls increases the frequency with which juries make correct decisions in the case of a guilty defendant.

REFERENCES

- Austen-Smith, D., and J. S. Banks, "Information aggregation, rationality, and the Condorcet jury theorem," *American Political Science Review*, 90 (1996): 34-45
- Coughlan, P., "In defense of unanimous jury verdicts: Communication, mistrials, and sincerity," mimeo, Caltech, 1997.
- Feddersen, T., and W. Pesendorfer, "Convicting the innocent: The inferiority of unanimous jury verdicts under strategic voting," *American Political Science Review*, 92(1998): 23-36.
- Fey, M., R. McKelvey, and T. Palfrey, "An experimental study of the constant sum centipede game," *International Journal of Game Theory*, 25 (1996): 269-87.
- Ladha, K., "Collective hypothesis testing by Bayesian actors under majority-rule voting," *mimeo* University of Chicago, (April 1998)
- Ladha, K., G. Miller, and J. Oppenheimer, "Information aggregation by majority rule: Theory and experiments," Draft 2s, *mimeo* Washington University, (June, 1996)
- McKelvey, R., and T. Palfrey "Quantal response equilibria for normal form games," *Games and Economic Behavior*, 10 (1995): 6-38.
- McKelvey, R., and T. Palfrey "Quantal response equilibria for extensive form games," *Experimental Economics*, 1 (1998): 9-41.
- McKelvey, R., T. Palfrey, and R. Weber, "The effects of payoff magnitude and heterogeneity on behavior in 2x2 games with unique mixed strategy equilibria" *Journal of Economic Behavior and Organization*, (forthcoming 2000).
- McLennan, A., "Consequences of the Condorcet jury theorem for beneficial information aggregation by rational agents," *American Political Science Review*, 92(1998): 413-18.

Wit, J., "Rational choice and the Condorcet jury theorem," *Games and Economic Behavior*, 22(1998): 364-76.

Appendix 1: Experiment instructions

DECISION MAKING EXPERIMENT

(Condorcet Jury)

This is an experiment in group decision making, and you will be paid for your participation in cash, at the end of the experiment. Different subjects may earn different amounts. What you earn depends partly on your decisions, partly on the decisions of others, and partly on chance.

The entire experiment will take place through computer terminals, and all interaction between you will take place through the computers. It is important that you not talk or in any way try to communicate with other subjects during the experiments.

We will start with a brief instruction period. During the instruction period, you will be given a complete description of the experiment and will be shown how to use the computers. If you have any questions during the instruction period, raise your hand and your question will be answered so everyone can hear. If any difficulties arise after the experiment has begun, raise your hand, and an experimenter will come and assist you.

In this experiment, one subject will act as a monitor. The monitor will be paid a fixed amount (\$20.00) for the experiment. The monitor will assist in running the experiment by generating random numbers for use in the experiment, and serving as someone who can check that the experiment is run correctly. We will select the monitor by having you each select an envelope. Open your envelope, and read the card inside to determine if you are the monitor.

[PASS OUT ENVELOPES]

Will the monitor please go to the master computer, and we will now pass out record sheets to the subjects in the experiments.

[PASS OUT RECORD SHEETS]

Please put your name at the top of the record sheet. The experiment you are participating in is broken down into 4 separate sessions. Each session will consist of 15 matches. At the end of the last match of the last session, you will be paid the total amount you have accumulated during the course of the 4 sessions. Everyone will be paid in private and you are under no obligation to tell others how much you earned. Your earnings are denominated in FRANCS. Your DOLLAR earnings are determined by multiplying your earning in FRANCS by a conversion rate. In this experiment, the conversion rate is .005, meaning that 100 FRANCS is worth 50 cents.

We will now begin the computer instructions for the first session. We will teach you about the experiment and how to use the computer by going through a short practice match. During the instruction session, do not hit any keys until you are told to do so, and when you are told to enter information, type exactly what you are told to type. You are not paid for the practice match.

[START JURY PROGRAM]

When the computer prompts you for your name, type your full name. Then hit the ENTER key, and confirm by typing Y.

FIRST SCREEN

You now see the first screen of the experiment. At the top center of the screen, you see your subject number. Please record that on your record sheet now. This session consist of 15 matches. In each match, you will be matched with another group of subjects. Your group number and the number of subjects you are matched with is on the top right side of your screen. In this session, each of you is matched with two other subjects, in groups of size 3.

There are two jars, which we call the Red Jar and the Blue Jar. The Red Jar contains 7 Red balls and 3 Blue balls. The Blue Jar contains 7 Blue balls, and 3 Red balls. At the beginning of the match, we will randomly choose one of the two Jars for your group. This will be done by the monitor, who will roll a dice to determine the color. If an odd number is rolled, the jar will be Blue, if an even number is rolled, the jar will be Red. Different groups may have different colors. You will not be told which Jar has been chosen. You may now press a key to advance to the next screen.

SECOND SCREEN

To help you determine which Jar has been selected for your group, each member of your group will be allowed to select one ball, at random, from the Jar. This is done on the computer by randomly ordering the balls in the Jar for each subject, but leaving the colors hidden. You can then use the mouse to click on one of ten balls. When you select a ball, its color will be revealed to you, and displayed on your screen, and you must record this information on your record sheet. This is called your "Sample". You are not told the colors of the sample balls drawn by the other members of your group. Note that the balls are numbered. This is just so that you can remember and record which ball you selected.

Each subject in your group selects a ball on their own, and only sees their own sample ball. The balls are ordered differently for each subject. However, all the members of your group have the same number of Red and Blue balls. That is, if the true color of the jar for your group were Red, then all members of your group are drawing their sample balls from a list containing 7 Red and 3 Blue balls. If the true color of your group's jar is Blue, then all members of your group would be drawing their sample balls from a list containing 3 Red and 7 Blue balls. Please select a ball now, and record its color and number on your record sheet. When you have finished recording, press a key to continue.

THIRD SCREEN

After everyone has drawn their sample, you will be asked to vote for either a Red Jar or a Blue Jar.

[READ EITHER A OR B]

A Unanimity instructions: If (and only if) everyone in your group votes Red, your group decision is Red. Otherwise, your group decision is Blue. That is, if anyone in your group votes Blue, the group decision is Blue.

B Majority rule instructions: If a majority (two or more) of your group votes Red, then the group decision is Red. Otherwise, your group decision is Blue. That is,

your group's decision will be whichever color receives more votes.

The table on this screen tells you your payoffs. If your group's decision is the same as the true color, then the decision is correct, and each member of your group earns 100 francs. If your group decision is different from the true color of the Jar, then the group decision is incorrect, and each member of the group earns 10 francs. You may vote by selecting, with your mouse, either the Red or the Blue Jar, and then confirming. In the actual experiment, you may vote however you like. In the practice match, please vote Red if you are on the left side of the room, and Blue if you are on the right side of the room. Please vote now.

FOURTH SCREEN

After everyone has voted, the vote is tallied, and you and the others in your group will be told what the result of the vote was, and what the true color of the Jar for your group was. Please hit a key to proceed

FIFTH SCREEN

After all decisions have been made, you are given a summary of the decisions up to now. Please record this information on your record sheet. This marks the end of the match. After you have recorded all of the information, please hit a key to continue.

SIXTH SCREEN – DISPLAY OVERHEAD

We will now illustrate in more detail how the sample is drawn. Your computer screen now shows the balls you drew your sample from with all of the colors revealed. The ball that you drew is outlined with a white box.

Prior to the experiment, all of the lists of sample balls for each subject, and each match, and each possible true Jar color have been generated by a random number generator. We have a printout of this information that will be available to the monitor. You now see on the overhead projector screen a copy of the first page of this data, which is for the practice match. Find your subject number and the true color of your groups jar. Compare the display on your computer screen to that on the overhead projector. The color of the balls should be identical. Please look at the overhead projector screen. Note also that the ordering of the balls for all subjects is different, even if they are drawing from the same Jar.

During the actual experiment, you will not be shown this screen. However, your samples will be done in exactly the same way. We have printed out the similar sheet for all of the matches to follow. If you record the number of the ball that you drew each match, then at the end of the experiment, if you wish, you may ask the monitor to check your sample against this sheet to verify that the ball you selected was the correct color.

After a match is over, and everyone has recorded the information about their match, you will be randomly re-matched into another group consisting of different subjects. A new jar will be randomly drawn for your group, you will get a new sample from a jar of the group color, you will vote, and a group decision will be determined as described above. After we have finished 15 matches in this manner, the first session of the experiment will be over, and we will give you instructions for the next session.

Are there any questions before we begin with the actual experiment

[EXPERIMENTER TAKE QUESTIONS]

We will now begin with the actual experiment. If there are any problems from this point on, raise your hand and an experimenter will come and assist you.

[START EXPERIMENT]

RULES FOR EXPERIMENT SESSION #2

This experiment session will also last for 15 matches. The rules are the same as in the first session, with one exception.

[READ A OR B]

A Recall that in the previous session, your group decision was Red if and only if all three members of your group voted Red. Otherwise your group decision was Blue. In this session there will be a slight change in the voting rule. Now, if a majority (two or more) of your group votes Red, then the group decision is Red. Otherwise, your group decision is Blue. That is, your group's decision will be whichever color receives more votes.

B Recall that in the previous session, your group decision was Red if two or more members of your group voted Red. Otherwise your group decision was Blue. In this session there will be a slight change in the voting rule. Now, if (and only if) everyone in your group votes Red, your group decision is Red. Otherwise, your group decision is Blue. That is, if anyone in your group votes Blue, the group decision is Blue.

Are there any questions?

RULES FOR EXPERIMENT SESSION #3

This experiment session will also last for 15 matches. The rules are the same as experiment session #2, with one exception. In this experiment, each group will conduct one preliminary vote before making the final vote. In the preliminary vote, you may vote for either the Red Jar or the Blue Jar. You do not receive any earnings for the preliminary vote. The number of Blue and Red votes in the preliminary vote will be revealed to you and the other members of your group. After you have been told this preliminary vote outcome, your group will conduct the final vote, in the same manner as in experiment session #2. We will conduct a practice session to illustrate how the preliminary vote is conducted.

Are there any questions?

RULES FOR EXPERIMENT SESSION #4

This experiment session will also last for 15 matches. The rules are the same as experiment session #1, with one exception. As in the previous experiment, each group will conduct one preliminary vote before making the final vote. You do not receive any earnings for the preliminary vote. The outcome of this preliminary vote will be revealed

to you and the other members of your group, including the exact number of red votes and blue votes. You are not paid anything for this outcome. You must record this information on your information and record sheet. Then your group will conduct the final vote, in the same manner as experiment session #1. That is,

[READ A OR B]

A Unanimity instructions: If (and only if) everyone in your group votes Red, your group decision is Red. Otherwise, your group decision is Blue. That is, if anyone in your group votes Blue, the group decision is Blue.

B Majority rule instructions: If a majority (two or more) of your group votes Red, then the group decision is Red. Otherwise, your group decision is Blue. That is, your group's decision will be whichever color receives more votes.

Are there any questions?

7 Appendix 2: Solution for the case: $N = 3$, unanimity, with straw poll

7.1 Notation and strategies

Denote by:

- $s = \{G, I\}$ the signal (either guilty or innocent);
- p = the probability of receiving the correct signal;
- $v_t = \{C_t, A_t\}$ the vote (either convict or acquit) in the straw poll ($t = 0$) and in the decisive final vote ($t = 1$);
- $i = \{0, 1, 2\}$ the number of other convicting votes in the straw poll;
- $j = \{0, 1, 2\}$ the number of other matching signals.

Payoffs are defined as follows:

- 0 for making the right decision (that is, convicting a guilty defendant and acquitting an innocent defendant);
- q for acquitting a guilty defendant (where $q \in [0, 1]$);
- $q - 1$ for convicting an innocent.

Each agent's strategy consists of a pair of probabilities σ_0, σ_1 , specifying the probability to vote to convict in the straw poll (as a function of one own's signal) and the probability to vote to convict in the final vote (as a function of one own's signal, one own's vote in the straw poll, others' vote in the straw poll):

$$\sigma_0(s) = \frac{e^{\lambda EU_{SP}(C_0, s)}}{e^{\lambda EU_{SP}(C_0, s)} + e^{\lambda EU_{SP}(A_0, s)}} \quad (3)$$

$$\sigma_1(s, v_0, i) = \frac{e^{\lambda EU_{FV}(C_1, s, v_0, i)}}{e^{\lambda EU_{FV}(C_1, s, v_0, i)} + e^{\lambda EU_{FV}(A_1, s, v_0, i)}} \quad (4)$$

where EU_{SP} (EU_{FV}) is the expected utility for casting a given vote in the straw poll (final vote, respectively).

7.2 Expected utility for the straw poll

$$EU_{SP}(v_0, s) = \begin{cases} \sum_{j=0}^2 \sum_{i=0}^2 \pi(j) \mu(i, j) EU_{FIN}(G, v_0, i) & \text{if } s = G \\ \sum_{j=0}^2 \sum_{i=0}^2 \pi(j) \mu(i, 2-j) EU_{FIN}(I, v_0, i) & \text{if } s = I \end{cases}$$

where $\pi(j)$ is the (conditional) probability of other j matching signals:

$$\pi(j) = \begin{cases} p^3 + (1-p)^3 & \text{if } j = 2 \\ 2[p^2(1-p) + p(1-p)^2] & \text{if } j = 1 \\ p(1-p)^2 + p^2(1-p) & \text{if } j = 0 \end{cases}$$

that is, in a more compact form,

$$\pi(j) = (2 - |j - 1|)[p^{j+1}(1-p)^{2-j} + p^{2-j}(1-p)^{j+1}];$$

$\mu(i, j)$ is the (conditional) probability that i other jurors voted to convict in the straw poll when the number of their matching signals is j :

$$\mu(i, j) = \begin{cases} [\sigma_0(G)]^j [\sigma_0(I)]^{2-j} & \text{if } i = 2 \\ j\sigma_0(G) + (2-j)\sigma_0(I) - 2[\sigma_0(G)]^j [\sigma_0(I)]^{2-j} & \text{if } i = 1 \\ (1 - \sigma_0(G))^j [1 - \sigma_0(I)]^{2-j} & \text{if } i = 0 \end{cases}$$

and EU_{FIN} is the weighted combination of the expected utilities from voting either to convict or acquit in the final vote, where weights are given by the strategy for the final vote:

$$EU_{FIN}(s, v_0, i) = \sigma_1(s, v_0, i) EU_{FV}(C_1, s, v_0, i) + [1 - \sigma_1(s, v_0, i)] EU_{FV}(A_1, s, v_0, i).$$

7.3 Expected utility for the final vote

In order to make the analysis clearer, consider first the case of a guilty signal (so for now, $s = G$). By the end of the straw poll, every juror can observe all the votes (so that he knows j), but still does not know what the true state is (either guilty or innocent) and what signals the other jurors received (where the number of other guilty signal can be either 0, 1 or 2). This leads to $2 \cdot 3 = 6$ information sets. The conditional probability of being in any of these sets is given by:

$$\hat{\beta}(i, k) = \frac{\beta(i, k)}{\sum_{l=1}^6 \beta(i, l)}$$

where

$$\beta(i, k) = \begin{cases} (2 - |k - 2|)p^{4-k}(1-p)^{k-1}\mu(i, 3-k) & \text{if } k = 1, 2, 3 \\ (2 - |k - 2|)p^{k-1}(1-p)^{4-k}\mu(i, 3-k) & \text{if } k = 4, 5, 6. \end{cases}$$

For an intuition about these coefficients, notice that we ordered the information sets so that the defendant is guilty in the first three of them ($k = 1, 2, 3$); and innocent otherwise. So $\beta(\hat{i}, k)$ are the conditional probabilities for the signal being correct (as in this case we are taking $s = G$) for $k = 1, 2, 3$; and incorrect otherwise.

For the case of $s = I$ we can define similar coefficients:

$$\hat{\epsilon}(i, k) = \frac{\epsilon(i, k)}{\sum_{l=1}^6 \epsilon(i, l)}$$

where

$$\epsilon(i, k) = \begin{cases} (2 - |k - 2|)p^{4-k}(1-p)^{k-1}\mu(i, k-1) & \text{if } k = 1, 2, 3 \\ (2 - |k - 2|)p^{k-1}(1-p)^{4-k}\mu(i, k-1) & \text{if } k = 4, 5, 6. \end{cases}$$

Notice first that the unanimity rule in the final vote implies that, whenever a juror votes to acquit, then independently of what the other jurors voted, the payoffs are determined (0 if the true state is innocent, $q - 1$ otherwise). So:

$$EU_{FV}(A_1, s, v_0, i) = \begin{cases} (q - 1) \sum_{k=1}^3 \hat{\beta}(i, k) & \text{if } s = G \\ (q - 1) \sum_{k=1}^3 \hat{\epsilon}(i, k + 3) & \text{if } s = I. \end{cases}$$

The formulas for EU_{FV} when the final vote is to convict are more complicated. So consider first the case in which the juror received a guilty signal and voted to convict in the straw poll:

$$EU_{FV}(C_1, G, C_0, i) = (q - 1) \sum_{k=1}^3 \hat{\beta}(i, k)[1 - \theta(i, k)] - q \sum_{k=1}^3 \hat{\beta}(i, k + 3)[\theta(i, k)]$$

To understand this formula, notice that the final payoff is 0 only if the true state is guilty *and* everyone is voting to convict. The payoff is $q - 1$ (as shown in the first addend) if the true state is guilty (so that the sum is over $k = 1, 2, 3$) but someone voted to acquit (with probability as expressed in $1 - \theta$). The second addend represents the other case, when the true state is innocent but everyone voted to convict. The θ s are defined as follows:

$$\theta(i, k) = \begin{cases} \sigma_1(G, A_0, 1)^{3-k} \sigma_1(I, A_0, 1)^{k-1} & \text{if } i = 0 \\ \frac{1}{2} [\sigma_1(G, C_0, 1)^x \sigma_1(I, C_0, 1)^{1-x} \sigma_1(G, A_0, 2)^{1-w} \sigma_1(I, A_0, 2)^w + \\ \sigma_1(G, C_0, 1)^{1-w} \sigma_1(I, C_0, 1)^w \sigma_1(G, A_0, 2)^x \sigma_1(I, A_0, 2)^{1-x}] & \text{if } i = 1 \\ \sigma_1(G, C_0, 2)^{3-k} \sigma_1(I, C_0, 2)^{k-1} & \text{if } i = 2 \end{cases}$$

where $x = 1.5 - |k - 1.5|$ and $w = 1.5 - |k - 2.5|$.

The case for a guilty signal and a straw poll vote to acquit ($EU_{FV}(C_1, G, A_0, i)$) is similar, but for the fact that everyone else observes one vote to convict less than in the previous case. So the formulas are as above, but for subtracting an observed convicting

vote in the θ 's (as an example, $\sigma_1(G, C_0, 1)$ becomes $\sigma_1(G, C_0, 0)$).

The case for an innocent signal is similar, but for the fact that it involves the $\hat{\epsilon}$ coefficients and the exponents in the formulas are switched. More precisely, consider the case of a straw poll vote to convict:

$$EU_{FV}(C_1, I, C_0, i) = (q - 1) \sum_{k=1}^3 \hat{\epsilon}(i, k + 3)[1 - \eta(i, k)] - q \sum_{k=1}^3 \hat{\epsilon}(i, k)[\eta(i, k)]$$

where

$$\eta(i, k) = \begin{cases} \sigma_1(G, A_0, 1)^{k-1} \sigma_1(I, A_0, 1)^{3-k} & \text{if } i = 0 \\ \frac{1}{2} [\sigma_1(G, C_0, 1)^w \sigma_1(I, C_0, 1)^{1-w} \sigma_1(G, A_0, 2)^{1-x} \sigma_1(I, A_0, 2)^x + \\ \sigma_1(G, C_0, 1)^{1-x} \sigma_1(I, C_0, 1)^x \sigma_1(G, A_0, 2)^w \sigma_1(I, A_0, 2)^{1-w}] & \text{if } i = 1 \\ \sigma_1(G, C_0, 2)^{k-1} \sigma_1(I, C_0, 2)^{3-k} & \text{if } i = 2 \end{cases}$$

Using the above analysis, we constructed a numerical algorithm to compute the Logit QRE correspondence.

Figures X and XX show the Logit equilibrium correspondence for the 3-person Unanimity and Majority Rule jury treatments with a straw poll, respectively. The curves labeled I and G graph the Logit equilibrium probability of voting to convict in the first (straw) voting stage, depending on whether an innocent or guilty signal was received, respectively. The other curves show the Logit equilibrium probabilities of voting to convict in the second (final) stage as a function of the signal received (i or g), one's own vote in the straw poll (C or A) and the number of conviction votes in the straw poll by the *other* jurors (0, 1, or 2).

Several points about these figures are notable. First, the equilibrium correspondence involves two bifurcations. For small values of λ , there is a unique equilibrium. At a critical point of approximately $\lambda = .xx.xx$ we pick up several new equilibria, which branch off simultaneously. The original part of the correspondence continues smoothly and eventually converges to the mixing probabilities of the babbling equilibria. That is, player's vote uninformatively in the straw vote and in the final vote they play the symmetric equilibrium of the game as if not straw vote had taken place. The new equilibria that branch off correspond to informative equilibria, with one of these branches converging to the full communication equilibrium, in which voters vote informatively in the straw poll, and, in the final vote, vote according to the majority outcome in the straw poll.

We focus our attention on the branch converging to the full communication equilibrium. There are two reasons for doing this. First, it seems like the most sensible equilibrium in the game with a straw poll, and is the one that Coughlan (1998) emphasizes. Second, when we fit QRE to the data, we find that this is the equilibrium closest to the data.

The key feature to observe in the QRE graph is the fact that voters will condition on their signal as well as the outcome of the straw poll, in contrast to the full communication

equilibrium. Moreover, this conditioning follows a very logical pattern. For any combination of how the other jurors vote in the straw poll, the QRE predicts that a voter is more likely to vote to acquit if they received an innocent signal than if they had received a guilty signal. For example, consider the QRE at our estimated value of $\lambda = 38.42$ (see table). If both other jurors voted to acquit in the straw poll, a voter with an innocent signal is predicted to vote to acquit approximately 75% predicted to vote to acquit slightly less than half the time. Similar comparisons hold for the cases where one other voter voted to acquit in the straw poll and where neither other voter voted to acquit. These qualitative comparisons are all borne out in the data.

While the QRE picks up the qualitative features of conditional voting in the second stage, it does not provide a very good fit quantitatively. In fact, the NNM estimates provide a better fit to the data than the QRE, even though NNM does not make the right qualitative predictions. The estimated choice probabilities from the QRE model is superimposed on the figures.

The reason why the quantitative fit of QRE is not as tight as the NNM fit is not entirely clear. QRE predictions of choice probabilities vary substantially across information sets. For some of these information sets, the QRE choice probabilities converge to 0 or 1 very quickly (especially in the second stage), while others converge very slowly (especially in the first stage). However, the data in the first stage has *less* noise than the data in the second stage! Therefore, we can't get a good fit both the first stage and to the second stage with a single value of λ .

There is a second possible explanation for this. In the QRE, the frequency of voting informatively should be much higher with a guilty signal than with an innocent signal. This reflects the fact that voters who receive innocent signals are less likely to be pivotal in the second stage, so that their vote in the first stage is less likely to make a difference. Consequently, the QRE choice probabilities conditional on an innocent signal converge to the informative Nash equilibrium more slowly than the QRE choice probabilities conditional on a guilty signal. The data do indeed support this qualitative prediction, but the difference is not statistically significant (.067 for innocent signals and .046 for guilty signals).

Given this post hoc analysis, there would be several fixes that to improve the QRE fit, such as estimating a separate value for λ in the first and second stages, for for different signals, but these fixes are ad hoc and we have not pursued them. ¹⁵

¹⁵A second, more interesting alternative is to estimate a version of QRE that incorporates a second "tremble" term, similar to NNM, but which (unlike NNM) is accounted for in the computation of equilibrium. This would effectively bound the QRE a fixed distance from 0) and 1. We have not pursued this due to issues of computational feasibility.

Table 1: Solution Values Without a Straw Poll When $p = .7$

n	Majority				Unanimity			
	$\sigma(g)$	$\sigma(i)$	$Pr[C I]$	$Pr[A G]$	$\sigma(g)$	$\sigma(i)$	$Pr[C I]$	$Pr[A G]$
3	1.0000	0.0000	.2160	.2160	1.0000	0.3143	.1404	.4989
6	1.0000	0.0000	.0705	.2557	1.0000	0.6511	.1863	.4849
12	1.0000	0.0000	.0386	.1178	1.0000	0.8244	.2073	.4776
24	1.0000	0.0000	.0115	.0314	1.0000	0.9120	.2173	.4740

Table 2. First and Second Treatments by Experimental Session

Experiment Number	Number of Subjects	Group Size	Session			
			1	2	3	4
CJ1	12	3	U/N (15)	M/N (10)	M/Y (10)	U/Y (10)
CJ2	12	3	M/N (15)	U/N (15)	U/Y (15)	M/Y (15)
CJ3	12	6	U/N (15)	M/N (15)	M/Y (15)	U/Y (15)
CJ4	12	6	M/N (15)	U/N (15)	U/Y (15)	M/Y (15)

Note: U = Unanimity rule, M = Majority Rule, N = No Straw Poll, Y = Straw Poll

Table 3: Proportion Voting to Convict by Signal
and Voting Rule, Controlling for Straw Poll

Rule	n	Without Straw Poll		With Straw Poll	
		Blue (<i>Innocent</i>)	Red (<i>Guilty</i>)	Blue (<i>Innocent</i>)	Red (<i>Guilty</i>)
Majority	3	.057 (157)	.972 (143)	.205 (146)	.838 (154)
	6	.209 (172)	.979 (188)	.248 (202)	.835 (158)
Unanimity	3	.360 (186)	.954 (174)	.161 (149)	.834 (151)
	6	.478 (186)	.897 (174)	.372 (183)	.864 (177)

Note: Number of observations in parentheses

Table 4: QRE, NNM, and Naive Model Estimates of 3- and 6-Person Groups

		Three-Person					Six-Person				
		f_i	Nash	QRE	NNM	$Naive$	f_i	Nash	QRE	NNM	$Naive$
Majority	$\sigma(g)$.972	1.000	0.964	0.964	0.964	.979	1.000	0.960	0.889	0.889
	$\sigma(i)$.057	0.000	0.036	0.036	0.036	.209	0.000	0.158	0.111	0.111
	$\lambda \gamma$			19.016	0.928	0.928			23.204	0.778	0.778
	$\lambda_{lo} \gamma_{lo}$			14.536	0.856	0.856			19.016	0.708	0.708
	$\lambda_{hi} \gamma_{hi}$			21.642	0.953	0.953			29.757	0.837	0.837
	$-\mathcal{L}^*$			53.738	53.738	53.738			110.201	125.580	125.580
Unanimity	$\sigma(g)$.954	1.000	0.958	0.952	0.792	.897	1.000	0.887	0.878	0.704
	$\sigma(i)$.360	0.314	0.407	0.332	0.209	.478	0.651	0.535	0.614	0.297
	$\lambda \gamma $			9.197	0.903	0.583			12.041	0.755	0.407
	$\lambda_{lo} \gamma_{lo} $			7.689	0.827	0.496			11.921	0.643	0.311
	$\lambda_{hi} \gamma_{hi} $			10.893	0.954	0.662			12.530	0.845	0.498
	$-\mathcal{L}^*$			154.897	154.350	184.227			188.016	194.057	219.409

Note: See Table 3 for number of observations.

Table 5: Voting with a Straw Poll:
 Proportion Voting to Convict by Prior Information
 (Straw Vote by Signal, and Final Vote by Straw Poll Result)

		Straw Vote		Final Vote		
Rule	n	Blue (<i>Innocent</i>)	Red (<i>Guilty</i>)	Blue (<i>Innocent</i>)	Tie	Red (<i>Guilty</i>)
Majority	3	.103 (146)	.968 (154)	.111 (144)	—	.917 (156)
	6	.089 (202)	.994 (158)	.135 (156)	0.606 (66)	.877 (138)
Unanimity	3	.067 (149)	.954 (151)	.160 (156)	—	.868 (144)
	6	.049 (183)	.989 (177)	.167 (120)	.657 (102)	.971 (138)

Note: Number of observations in parentheses

Table 6: Probit Analysis of Final Vote in Groups with Straw Poll,
by Voting Rule and Group Size

Variable	Decision Rule			
	Majority		Unanimous	
	$n = 3$	$n = 6$	$n = 3$	$n = 6$
Constant	-1.891 (-9.468)	-2.213 (-10.080)	-1.799 (-9.363)	-2.039 (-9.285)
Public information (Straw Poll)	1.016 (7.355)	0.618 (8.540)	0.810 (6.407)	0.723 (9.296)
Private sSignal	0.874 (4.068)	1.154 (6.487)	1.136 (5.276)	0.751 (4.037)
Number of observations	300	360	300	360
Log likelihood	-109.82	-138.46	-110.26	-130.13
% correctly predicted	89.0	82.2	83.7	84.17

Note: t statistics in parentheses.

Table 7: Jury Decision vs True State: Proportion Incorrect Decisions
by True State, Group Size, and Voting Rule

Rule	n	Without Straw Poll						With Straw Poll			
		True State			True State			True State		True State	
		<i>Innocent</i>		QRE	<i>Guilty</i>		QRE	<i>Innocent</i>		<i>Guilty</i>	
Data	Nash	Data	Nash		Data	Nash		Data	Nash		
Major	3	.175 (40)	.216	.234	.300 (60)	.216	.234	.226 (53)	.216	.191 (47)	.216
	6	.296 (27)	.071	.177	.212 (33)	.256	.221	.176 (34)	.071	.115 (26)	.256
Unan	3	.190 (63)	.140	.187	.526 (57)	.499	.502	.140 (50)	.216	.360 (50)	.216
	6	.029 (34)	.186	.069	.731 (26)	.485	.772	.107 (28)	.071	.437 (32)	.256

Note: Number of observations in parentheses.

Table 8: Breakdown of Juror Types in Unanimous Juries without Deliberation
(Innocent Signals Only)

Juror Type	Session 1	Session 2	Session 3	Session 4	Total
1: Sincere	4 (0/31)	4 (0/35)	3 (0/27)	0 (0/0)	12 (93/93)
2: Mixers	7(24/57)	5 (20/40)	6 (19/41)	10 (32/75)	27 (95/213)
3: Punitive	1 (5/5)	3 (18/18)	3 (23/23)	2 (15/15)	9 (61/61)

Table A-1: Model Estimates for 3-Person Juries with Straw vote

	n	f_i	QRE	NNM	Nash	
Unanimity	$\sigma(g)$	151	.954	.882	.904	1.000
	$\sigma(i)$	149	.067	.293	.096	0.000
	$\sigma(gC0)$	31	.419	.574	.096	0.000
	$\sigma(gC1)$	47	1.000	0.995	0.904	1.000
	$\sigma(gC2)$	66	.955	1.000	0.904	1.000
	$\sigma(gA0)$	2	1.000	0.501	0.096	0.000
	$\sigma(gA1)$	2	.500	0.895	0.904	1.000
	$\sigma(gA2)$	3	0.000	1.000	0.904	1.000
	$\sigma(iC0)$	4	.250	0.093	0.096	0.000
	$\sigma(iC1)$	3	0.000	0.099	0.096	0.000
	$\sigma(iC2)$	3	0.000	0.751	0.904	1.000
	$\sigma(iA0)$	49	.061	0.259	0.096	0.000
	$\sigma(iA1)$	68	.074	0.378	0.096	0.000
	$\sigma(iA2)$	22	.682	0.978	0.904	1.000
				38.42	0.807	
			37.28	0.756		
			40.00	.850		
			279.63	190.62		





Figure 1: Symmetric Part of QRE Correspondence under Majority Rule





Figure 2: Symmetric Part of QRE Correspondence under Unanimity

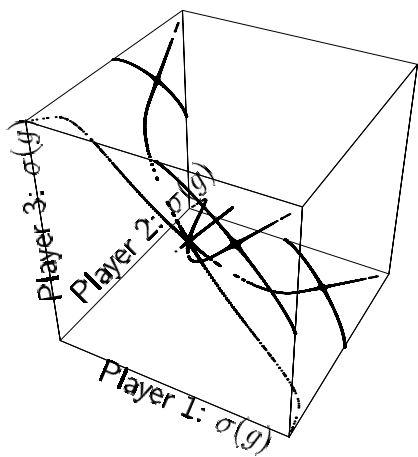



Figure 3: QRE Correspondence for 3 Person Unanimity Jury





Figure 4: Symmetric part of QRE Correspondence for 6 Person Unanimity Jury