Hypothetical Thinking and Information Extraction: Strategic Voting in the Laboratory *

Ignacio Esponda
(NYU Stern)

Emanuel Vespa
(NYU)

June 7, 2012

Abstract

We test for strategic behavior in common-value elections by using an experimental design that distinguishes the mistake of failing to account for the information content of others' strategies from other mistakes that arise naturally when making inferences. Depending on the treatment, between 50-80% of subjects fail to vote strategically. This mistake is robust to experience and hints about pivotality, and it is mainly driven by the failure to make inferences from hypothetical, rather than actual, events. Finally, the mistake also arises in more general settings where players have no private information.

*Research support from the Center for Experimental Social Science (CESS) at NYU is gratefully acknowledged. We thank Katherine Baldiga, Ryan Booth, Alessandra Casella, Stephen Coate, Pedro Dal Bó, Kfir Eliaz, Guillaume Fréchette, Philippe Jehiel, Navin Kartik, Alessandro Lizzeri, Rebecca Morton, Andy Schotter, Matan Tsur, Alistair Wilson, Jonathan Woon, Sevgi Yuksel, and several seminar participants for helpful comments, and Qingyuan Gao for research assistance. Esponda: 44 W Fourth Street KMC 7-76, New York, NY 10012, iesponda@stern.nyu.edu; Vespa: 19 W Fourth Street Off. 502, New York, NY 10012 , emanuel.vespa@nyu.edu.

1 Introduction

A large experimental literature on the winner's curse has documented that bidders often fail to bid optimally (e.g., Thaler (1988), Kagel and Levin (2002)). A common explanation is that bidders fail to anticipate that winning the auction conveys (bad) news about the value of the object, simply because other bidders with private information have been reluctant to bid as high.[1] A recent theoretical literature (e.g., Jehiel (2005), Eyster and Rabin (2005), Crawford and Iriberri (2007), Jehiel and Koessler (2008), Esponda (2008)) has formalized and generalized this mistake by postulating that players fail to extract information from other players' strategies.[2]

An election with common values (e.g., Feddersen and Pesendorfer, 1997) is yet another environment where optimality requires players to engage in information extraction, and this type of behavior is known as *strategic* or *pivotal voting*: A strategic voter votes for the alternative that is best conditional on both her private information and the information that she infers, given her beliefs about the equilibrium strategies of others, from the hypothetical event that her vote is pivotal.[3] Starting with Guarnaschelli et al. (2000), the experimental voting literature finds evidence consistent with strategic voting. The different conclusions reached by the auction and voting literatures seem puzzling given their common underlying structure.

While insightful on various dimensions, the common-value voting and auction experimental literatures have focused mostly on testing predicted behavior in these environments, but less on understanding the underlying cause of this behavior.[4] In particular, these literatures do not identify whether players are or are not sophisticated enough to engage in information extraction. The reason is that additional mistakes can occur at the information extraction stage, either because players form

[1]For example, suppose a bidder bids b and believes that other bidders follow the symmetric and increasing strategy β, mapping signals to bids. Then she should infer that, in the hypothetical event that she wins the auction (which is the only event in which her bid is relevant), other bidders are bidding below b and, consequently, each of their signals is below $\beta^{-1}(b)$.

[2]See Kagel and Levin (1986) and Holt and Sherman (1994) for early formal analyses in the context of auctions.

[3]Because 'pivotal' voting is a shortcut that voters may (but need not) take in order to vote optimally (see Section 3.1), we use the more general term 'strategic' to describe sophisticated subjects. Also, we focus on common value elections and ignore the experimental literature on elections with private values, where 'strategic' has a different meaning and refers either to optimal turnout or to voting for a less preferred candidate in an election with more than two candidates (e.g., Eckel and Holt (1989), Levine and Palfrey (2007)).

[4]An exception in the auction literature is Charness and Levin (2009), discussed in Section 2.

incorrect beliefs about the strategies of other players or because they make mistakes when using Bayes' rule to update their beliefs. This problem is exacerbated in voting games by the multiplicity of equilibria.

We conduct experiments in a voting environment with two alternatives and private information.[5] Our first main contribution is to isolate whether voters are strategic in the sense that they engage in information extraction from others' strategies. We cast a majority voting game as a decision problem. Each subject plays against two computers, and subjects know both the primitives of the game and the strategies of the computers, but not the actual votes of the computers. Computers' strategies, which are correlated with the identity of the best alternative, are chosen so that the inferences subjects need to make from them (e.g., Bayesian updating) are trivial. Consequently, the experiment provides a test of strategic behavior that is not contaminated by inferential mistakes.[6]

We find that a substantial fraction of subjects (79% in a treatment where subjects receive no feedback after each round) are non-strategic after playing for several rounds. The mistake of failing to extract information is fairly robust to experience, feedback, and even hints about pivotality (e.g., 50% of subjects are non-strategic when they do receive detailed feedback after each round of play). In addition, not all non-strategic behavior can be explained by the standard notion of sincere voting. This last finding (also obtained by Charness and Levin (2009) for the acquiring-a-company setup) suggests that other models of naivete should be considered when testing for non-strategic behavior.

Our second main contribution is to provide an explanation for what makes information extraction difficult in this environment. We turn the original simultaneous-voting experiment into a sequential-voting experiment: Each (new) subject now votes after having observed the realized votes of the computers. In this context, there is also information to be extracted from knowledge of others' actual votes. However, information must now be extracted from *actual* events, in contrast to the simultaneous-voting experiment, where information must be extracted from *hypothetical* events (e.g., the event where one's vote makes a difference). In the sequential environment,

[5]The voting environment allows us to better isolate, compared to a trading environment, whether people extract information from others' strategies (see Section 2).

[6]This simplification of the voting game is appropriate to test whether subjects apply the logic of information extraction, since that logic is the same whether the action from which information must be inferred is followed by a human player or any other entity.

we find that only 22% of the subjects are now unable to extract information from the computers' votes. We then ask the subjects who were able to extract information in the sequential experiment to participate in the simultaneous-voting experiment and, strikingly, we find that 79% of them vote non-strategically. Therefore, our data suggests that the mistake is mainly driven not from failure to make inferences from others' actions per se, but rather from failure to make inferences from hypothetical events.

Our final contribution is to show that the failure to extract information is both relevant and present in private values settings. We conduct a third experiment in which players vote simultaneously, but where the strategies of the computers are uninformative about the best alternative. Nevertheless, the rules of the election are modified so that information about other players' strategies can still be used to extract information about the *relevant* state of the world (which includes both the best alternative and the votes of the computers). We find that 65% of subjects are unable to extract information from the computers' strategies. This experiment illustrates that the mistake identified in this paper is relevant in more general settings and suggests the need for further work to formalize the more general mechanism and explore the underlying cognitive biases that make information extraction from hypothetical events such a challenging task.[7]

In Section 2, we review the related experimental literature. We present the experimental designs in Section 3 and the results in Section 4, and then briefly conclude in Section 5.

2 Related experimental literature

Common-value elections: Guarnaschelli et al. (2000) are the first to experimentally test models of strategic voting under asymmetric information. They find that voters in the 'jury model' (e.g., Feddersen and Pesendorfer, 1998) vote against their private information under unanimity rule but not under majority rule. Ali et al.

[7] Camerer and Lovallo (1999) relate the winner's curse to overconfidence and excess entry, Camerer et al. (2004) propose a model where people neglect the strategic complexity of the relevant environment and test it for several classes of related games, and Huck et al. (2011) show that subjects in a game best respond to the aggregate play of their opponent over various games, as in Jehiel's (2005) analogy-based expectation equilibrium. For another private information environment in which subjects fail to fully extract information from others' actions, see Vespa and Wilson (2012).

(2008) confirm the robustness of these findings and find relatively few differences in aggregate behavior between simultaneous and sequential voting. Battaglini et al. (2008, 2010) test the strategic abstention model of Feddersen and Pesendorfer (1996) and find that uninformed voters often abstain, thus delegating the decision to more informed voters, and, in contexts with partisan bias, uninformed independents vote to offset these biases.

All these findings, while noisy, tend to be consistent with models of strategic voting but do not necessarily imply that voters are able to extract information.[8] The reason is that the experiments test the joint hypotheses of strategic voting and equilibrium behavior. Once we drop the identifying assumption that players have correct beliefs about each others' strategies, then any vote is rationalizable and their findings are also consistent with ours.[9] In addition, these papers test strategic voting against one very specific alternative, sincere voting. Therefore, a voter voting against her signal is interpreted as being strategic. But non-strategic players may also have reasons to vote against their signals. For example, when unanimity is required to convict a defendant, a single vote to acquit lets the defendant go free. Some non-strategic subjects may prefer to avoid making the ultimate decision by voting to convict even if they have an 'innocent' signal.[10]

Common-value auctions: Kagel and Levin (1986), Kagel et al. (1987), and Levin et al. (1996) report overbidding in common-value auctions, implying that bidders fall prey to the winner's curse–see Kagel and Levin (2002) for a survey. More recently, Charness and Levin (2009) seek to understand what explains the observed overbidding. They ingeniously transform a trading game into a decision problem that retains the adverse selection problem but strips the environment of strategic uncertainty.[11] However, as they recognize, the nature of their game implies that subjects must engage in non-trivial computations of expected profits and conditional expectations. To mitigate this concern, Charness and Levin (2009) study a simpler version of their experiment, where subjects essentially choose one of two prices. In

[8]Eyster and Rabin (2005) use their model of cursedness to explain the errors made by subjects in the experiment of Guarnaschelli et al. (2000).

[9]In fact, the game has multiple equilibria and Guarnaschelli et al. (2000) find that election *outcomes* differ markedly from the symmetric Nash equilibrium prediction.

[10]Similarly, Morton and Tyran (2011) argue that delegation (to more informed voters) is likely to be a voting norm rather than the result of pivotal calculations, and Esponda and Pouzo (2012) present a model where non-strategic voters sometimes vote against their signals.

[11]Walker et al. (1987) are the first to use computerized competitors in an auction experiment.

this simpler version, one cannot distinguish between subjects making correct or naive inferences about value because both types choose the same low price.[12] Charness and Levin (2009) find that about half of the subjects "overbid" by choosing the high price, therefore providing a lower bound for the failure of information extraction and highlighting that overbidding cannot solely be explained by naive inferences. An advantage of our voting setup is that it can distinguish between naive and sophisticated inferences while requiring trivial computations.[13]

Finally, Ivanov et al. (2010) conduct an experiment where subjects play a two-player second-price auction, where the common value of the item is the maximum of the players' signals. Despite eliminating strategic uncertainty (by having players play against their own past strategies), they still find that subjects overbid. However, their design is not intended to identify whether subjects can extract information about the value of the item. The reason is that overbidding is also prevalent in private-value second-price auctions (e.g., Kagel et al., 1987), in which subjects know their value of the item.

3 Experimental designs

We first describe the three main voting problems faced by (different) subjects and then provide details about the experimental designs.

[12]In one treatment, they have 50 cards with a value of 20 and 50 cards with a value of 119. Subjects make an offer for a randomly chosen card of unknown value. If the offer is higher or equal than the card value, they receive 1.5 times the card value minus their offer; otherwise they receive nothing. Since an offer is only relevant in the hypothetical event in which the offer is accepted, a "strategic" subject realizes that an offer of 20 yields $1.5 \times 20 - 20 = 10$ with probability $1/2$ and nothing otherwise, which, assuming she is not too risk loving, she prefers to offering 119 and getting 59.5 with probability $1/2$ and losing 89 with probability $1/2$. On the other hand, a 'sincere' or 'cursed' (Eyster and Rabin (2005)) subject makes the same computation for an offer of 119 but incorrectly believes that an offer of 20 yields an average prize of $1.5 \times ((20+119)/2) - 20 = 49.5$ with probability $1/2$ and nothing otherwise. This experiment cannot distinguish between strategic and sincere subjects because an offer of 20 is optimal for both types. The distinction could be made with three effective prices, but it would require subjects to compute non-trivial conditional expectations.

[13]Our voting setup is also robust to subjects' risk preferences, although Charness and Levin (2009) resolve this issue by conducting an additional experiment with lotteries.

3.1 The main voting problems

At the center of all of our designs is the following setup. There is a 10-ball jar with p red balls and $10-p$ blue balls, where $p \in \{1, ..., 9\}$. One ball is randomly selected and becomes the *selected ball*. The subject must cast a vote for either Red or Blue without observing the color of the selected ball. In addition, two computers observe the color of the selected ball and are programmed to follow specific rules for casting a vote in favor of Red or Blue that are contingent on the color of the selected ball. The remaining description of the setup is specific to each experiment.

Experiment I. Simultaneous voting: If the color chosen by a simple majority matches the color of the selected ball, the subject's payoff is \$2 for that round; otherwise, the payoff is \$0. In each round, we vary the rule followed by the computers as well as the composition of the jar. Before casting her vote, the subject receives information about the rule being followed by the computers and the number of red and blue balls in the jar, but *not* about the actual votes of the computers. We restrict attention to the case where both computers follow the same rule (but randomizations are independent), and where the rules take the following form, where $q \in \{.1, .25, .5, .75, .9\}$:

$$\begin{array}{ll} \text{If selected ball is red:} & \text{vote Red} \\ \text{If selected ball is blue:} & \text{vote Blue with probability } q \text{ and Red with probability } 1-q \end{array} \tag{1}$$

Table 1 shows the state-space representation for this voting problem.[14] There are six states of the world, each with its corresponding probability and the payoff from voting Red or Blue.[15] A sophisticated subject who knows the composition of the jar (p) and the strategies of the computers (q) can extract information about the probability of each state, as shown in Table 1. In particular, the second and third states have zero probability because both computers vote Red when the selected ball is red. Thus, comparing the payoffs in the last two columns, the optimal strategy is to vote Blue.

[14] In applications, the state of the world is often considered to include only the color of the selected ball, but, as shown by Experiment III, this narrow reading of the state space obscures the generality of the mistake studied in this paper.

[15] For simplicity, states (2) and (5) aggregate the cases where either the first computer votes Red and the second Blue or vice versa.

	State		Probability	Subject's Payoff	
	Selected Ball	Computers' Votes		Vote R	Vote B
(1)	red	Red/Red	p	2	2
(2)	**red**	**Red/Blue**	$\mathbf{0}$	**2**	**0**
(3)	red	Blue/Blue	0	0	0
(4)	blue	Red/Red	$(1-p)\times(1-q)^2$	0	0
(5)	**blue**	**Red/Blue**	$(\mathbf{1}-\mathbf{p})\times\mathbf{2q}(\mathbf{1}-\mathbf{q})$	**0**	**2**
(6)	blue	Blue/Blue	$(1-p)\times q^2$	2	2

Table 1: State-space representation for Experiment I.

In other words, if the selected ball is red, then both computers will vote Red and the payoff will be \$2 irrespective of the subject's vote. But if the selected ball is blue, then the subject can influence the outcome, and it is optimal to vote Blue. Equivalently, the vote of a subject can only affect the outcome if the selected ball is blue (i.e., conditional on being pivotal, the selected ball must be blue), and, therefore, it is optimal to vote Blue.[16] In particular, a strategic subject must ignore her private information (i.e., the number of red balls in the jar), a result that is typical of voting games with private information (e.g., Feddersen and Pesendorfer, 1997).

Two features of the setup allow us to isolate whether voters are strategic from whether voters have incorrect beliefs about others' strategies or make computational mistakes. First, subjects have no uncertainty about the primitives of the game (i.e., the composition of the jar) or about the strategies of the computers. Second, the rules followed by the computers make it unnecessary for strategic subjects to engage in non-trivial Bayesian updating computations or, more generally, to compare the probabilities of each of the states (beyond knowing that some are positive and some are zero).

Experiment II. Sequential voting: The second experiment is identical to Experiment I, except that the subject now learns the realized votes of the computers before having to cast her vote. If the realized votes are $\{RR\}$ or $\{BB\}$, then the

[16]The last argument, known as *pivotal voting*, is the standard explanation in the literature, and follows from the fact that the difference in expected utility from voting Red versus Blue is given by the probability of being pivotal (which is positive) times the expected difference in payoff from voting Red versus Blue *conditional* on being pivotal. However, as shown by the initial argument, there are other, equivalent ways to extract information. Our experiment does not seek to distinguish which of these equivalent arguments are followed by a sophisticated subject, which is why we generically refer to this behavior as being *strategic* rather than pivotal.

	State			Subject's Payoff	
	Selected Ball	Computers' Votes	Probability	Vote R	Vote B
(1)	red	Red/Red	$p \times q^2$	2	2
(2)	red	Red/Blue	$p \times (1-q) \times (1+q)$	2	2
(3)	**red**	**Blue/Blue**	$\mathbf{0}$	**2**	**0**
(4)	**blue**	**Red/Red**	$\mathbf{(1-p) \times q^2}$	**0**	**2**
(5)	blue	Red/Blue	$(1-p) \times (1-q) \times (1+q)$	2	2
(6)	blue	Blue/Blue	0	2	2

Table 2: State-space representation for Experiment III.

subject cannot affect the outcome and is therefore indifferent between voting for Red or Blue. In these cases, we cannot infer much from the subject's behavior. If the realized votes are $\{RB\}$, however, then a sophisticated subject learns that either the second or fifth states in Table 1 have been realized, and, from knowledge of the computers' strategies, infers that the fifth state was realized and, therefore, votes Blue. In words, a sophisticated subject votes Blue because she infers that computers can vote differently only if the selected ball is blue. The difference with Experiment I is that the subject no longer needs to make inferences from all hypothetical states, but she can now simply focus on the event that computers vote differently.

Experiment III. No private information: The third experiment is intended to highlight that the logic applied above to find the optimal vote still holds in settings where the rules of the computers are not correlated with the color of the selected ball. Accordingly, the experiment coincides with Experiment I (in particular, voting is simultaneous), except that the payoffs and the rules of the computers are modified as follows. The payoff is now \$2 if there is at least one vote for the color of the selected ball, and \$0 otherwise. Computer 1 votes Red with probability q. Computer 2 observes the realized vote of Computer 1 and votes Red with probability q if Computer 1 voted Red and votes Red with probability 1 if Computer 1 voted Blue. Thus, computers' votes are correlated with each other, but they are not correlated with the color of the selected ball. Table 2 shows the state-space representation for this problem.

Despite the fact that computers' strategies are not correlated with the color of the selected ball, the strategies do contain information about the likelihood of each state and the reasoning is similar to that applied in Experiment I. In particular, it is never the case that both computers vote Blue, and, therefore, the third and sixth states have probability zero. But, then, a sophisticated subject should vote Blue. This example illustrates that the relevant state-space includes both the primitive uncertainty (i.e., the color of the selected ball) and the actions of the computers, and, therefore, extracting information from others' strategies is a very general phenomenon that can be relevant even if other players have no information about the primitive. More importantly, if the failure to extract information from hypothetical events is the underlying cause of mistakes, then one should find evidence for it both in Experiments I and III.

3.2 Details of the experimental designs

Experiment I. Simultaneous voting: The experiment consists of three parts. In Part 1, we vary the jar composition and the voting rule, (p, q), over the values indicated above for a total of $9 \times 5 = 45$ rounds.[17] A screenshot of the interface displaying the case $\{p = 6, q = 0.9\}$ is presented in Figure 1.[18] Subjects face each of these 45 rounds in a random order. In addition, for values of $q \in \{.1, .75\}$, we inverted the computers' voting rules: "If the selected ball is red, vote Red with probability q and Blue with probability $1 - q$; If the selected ball is blue, vote Blue." In such cases, the optimal strategy is to vote for Red.[19] For simplicity in the exposition, we adopt the convention that a setup with (p, q) always corresponds to the voting rule described above in (1), where p corresponds to the number of red balls in the jar.

In Part 2 of the experiment, each subject is given incentives to provide useful advice to another randomly chosen subject regarding how to vote in each of two different situations: round 46 is $p = 7, q = .9$ and round 47 is $p = 3, q = .1$. We use the written advice both to encourage further reflection about the problem and to confirm our classification of strategic behavior.

Finally, Part 3 of the experiment is divided intro three short stages: a, b, and c. At the beginning of the first two stages, we ask each subject a question that provides a hint of the notion of pivotality. Question 3a below is asked at the first stage and

[17]The parameters p and q capture the main aspects of a voting game with private information, where voters must make decisions based on both their information and the information inferred from the behavior of other voters. Variation in (p, q) allows us to identify to what extent voters respond to either source of information or both.

[18]We conducted our experiments using z-Tree; see Fischbacher (2007).

[19]We introduced such variation to distinguish a strategic subject from someone who always votes for the same color.

question 3b below is asked at the second stage (NP means that the number of balls in the jar, which is not needed to answer the questions, is not provided):

Question 3a: Case $(p, q) = (NP, .5)$. What is the probability that the selected ball is blue if one computer votes Red and the other computer votes Blue?

Question 3b: Case $(p, q) = (NP, .5)$. Suppose that the selected ball is red. Can your vote change the color chosen by the majority?

At the beginning of stage 3c, subjects read an explanation (see the Online Appendix) of why voting for Blue is optimal when computers follow the rule in (1). In each of the 3 stages of Part 3, the question or explanation is followed by four additional rounds of the voting problem, where $(p, q) \in \{(8, .5), (8, .75), (8, .25), (2, .5)\}$. Hence, subjects play the voting game for a total of 12 other rounds, in addition to the 47 rounds of Parts 1 and 2.[20]

We conducted two versions of Experiment I. In the "Feedback" treatment, after the conclusion of each round, each subject receives information about the color of the selected ball, her own vote, the votes of the two computers, the vote of the majority, and her payoff for that round. In the "No Feedback" treatment, subjects receive no feedback until the end of the experiment.

Experiment II. Sequential voting: The experiment consists of five parts, and subjects are not given any feedback until the end of the experiment. Part 1 coincides with Part 1 of the simultaneous treatment, except that voting is sequential and the subject observes the realized votes of the computers before casting her vote. As before, we vary (p, q) for a total of 45 rounds. In Section 4.2, we explain how Part 1 is necessary for comparison with Experiment I, despite the fact that subjects face few pivotal cases. In Part 2, each subject is given incentives to provide useful advice to another randomly chosen subject regarding how to vote in each of three different situations: In round 46, subjects are told that their vote is pivotal, and strategic and sincere behavior do not coincide ($p = 7, q = .9$); in round 47, subjects are told that their vote is pivotal, and strategic and sincere behavior coincide ($p = 3, q = .1$); and, in round 48, subjects are told that their vote is not pivotal ($p = 4, q = .25$ and both computers voted for Red).

[20] The 20 subjects in Session 1 did not face the fourth round in the previous list.

Part 3 is designed to collect more information from pivotal cases. In Part 3 (rounds 49 through 93), subjects face another set of 45 rounds of the voting game, as in Part 1. The difference is that subjects now have to provide a voting recommendation, rather than an actual vote, to a randomly matched partner who will later face, at the end of the experiment, the same 45 rounds. Subjects are asked to provide a partial voting recommendation for one of the following three scenarios: computers voted differently, both computers voted Blue, and both computers voted Red. For each of the 45 rounds, one of these three scenarios is randomly selected and subjects can only recommend to their partner how to vote in that particular scenario. For each round, we select an scenario in the following way: with probability .9, the pivotal scenario is chosen; with probability .1, the scenario is chosen according to the primitives of that round. We explained to subjects that their partners would later face each round and that the actual votes of the computers would be determined by the primitives of that round. In particular, partners only receive advice if the realized scenario is one for which the subject provided advice in Part 3.[21]

In Part 4, subjects are asked once again to submit incentivized written advice for the same three cases of Part 2, except that now, as in the simultaneous treatment, they must give advice to a subject who must cast her vote without knowing how the computers voted. This part allows us to test if experience with extracting information under sequential voting facilitates hypothetical thinking under simultaneous voting. Finally, in Part 5 subjects play the 45 rounds for which their partners provided advice during Part 3.

Experiment III. No private information: As explained earlier, the experiment coincides with Experiment I, except for the payoffs and the rules of the computers. In particular, there are 45 rounds in Part 1, advice in Part 2 (rounds 46 and 47), and another 12 rounds in Part 3. The hints provided in Part 3 are, of course, tailored to this experiment:

Question 3a: Case $(p, q) = (NP, .5)$. What is the probability that both computers vote Blue?

Question 3b: Case $(p, q) = (NP, .5)$. Suppose that the selected ball is red. Can

[21] We do not follow the strategy method (which asks for a recommendation for all contingent scenarios in a same round) because it may help subjects to think hypothetically, thus biasing our results from Part 4.

you get a payoff of \$0 if you vote Blue?

Finally, in stage 3c, subjects read an explanation of the optimal vote and play another four rounds. Subjects are only provided with feedback at the end of the experiment.

Subjects and payments: All sessions were run at NYU's Center for Experimental Social Science (CESS), where each of our 237 subjects only participated in one of the experiments. In each session, before the experiment began, subjects were asked incentivized questions to test their understanding of the instructions. On average, sessions had 20 subjects, and the total number of subjects was 60 for Experiment I without feedback, 58 for Experiment I with feedback, 58 for Experiment II, and 61 for Experiment III. For experiments I and III, payoffs were calculated by randomly selecting 7 out of 45 rounds from Part 1 and 3 out of 12 rounds from Part 3, and adding these payoffs to the payoffs obtained from answering the questions after the instructions and in Part 3.[22] For Experiment II, payoffs were calculated by randomly selecting 7 out of 45 rounds from Part 1 and 7 out of 45 rounds from Part 5 (adding both the subject and the partner payoffs).[23] On average, subjects received \$25 in Experiment I, \$34 in Experiment II, and \$28 in Experiment III. Sessions took between between 60 and 90 minutes in Experiments I and III and around 120 minutes in Experiment II. The instructions and additional details about how the experiment was conducted are provided in the Online Appendix.

4 Analyses and Results

Our main objective is to separate subjects into those that are strategic and those that are not. As shown below, the experimental designs make it easy to classify subjects because strategic subjects make essentially no mistakes (once they become strategic). In particular, we do not need to rely on a noisy statistical model to classify subjects within a treatment. Of course, we do account for sample randomness when testing whether strategic behavior differs across experiments.[24]

[22] Two subjects also received additional payoffs from Part 2, as explained in the Online Appendix.

[23] Some subjects received additional payoffs from Parts 2 and 4. We incentivized subjects in the same manner as in Part 2 of the simultaneous treatment.

[24] We report the p-value of one-sided z-tests between parentheses.

4.1 Experiment I: simultaneous voting

As explained in Section 3, in order to simplify the presentation we adopt the convention that all strategies take the form given by (1) above (in the experiment, the blue and red labels are interchanged). As argued above, a strategic subject should vote Blue.

Definition 1. A vote is optimal given $(p, q) \in \Omega$ if it is a vote for Blue.

We find two types of non-strategic behavior. Under sincere voting, a subject votes for the alternative that she considers best using only her information about the primitives and not taking into account the strategies of the computers.

Definition 2. A vote is sincere given $(p, q) \in \Omega$ if it is a vote for Red whenever $p > 5$, a vote for Blue whenever $p < 5$, and a vote for either Red or Blue whenever $p = 5$.

Also, because subjects experience changes in both the composition of the jar and the strategies of the computers, it is possible that some non-strategic subjects may systematically respond to both p and q. A reasonable hypothesis is that some subjects may try to "conform" to the votes of the computers by voting for the color that the computer votes with highest probability. Such conforming behavior is in line with prior evidence by Goeree and Yariv (2007), who find that, in a related setting, a significant number of subjects non-strategically conform to the choices of previous subjects.

The probability that a computer votes Red (p_R^c) is given by the probability that it votes Red and the ball is red plus the probability that it votes for Red and the ball is blue:

$$p_R^c(p, q) = 1 \times (p/10) + (1 - q) \times (1 - p/10). \tag{2}$$

Definition 3. A vote is conforming given $(p, q) \in \Omega$ if it is a vote for Red whenever $p_R^c(p, q) > .5$ and a vote for Blue whenever $p_R^c(p, q) < .5$.

Figure 2 summarizes whether a vote is Red or Blue for each type of behavior (strategic, sincere, and conforming) and for each of the 45 (p, q) pairs. In Region 1, sincere and confirming votes are for Red, while strategic votes are for Blue. In Region 2, conforming votes are for Red, while sincere and strategic votes are for Blue. In Region 3, all three types of votes are for Blue.[25]

[25] For simplicity, the figure arbitrarily shows sincere votes being always Blue if $p = .5$, but in the analysis we do follow Definition 2.

We find four representative types of behavior in our data, each of which is illustrated by one of the subjects in Figure 3. For each subject, the figure plots the vote for each of the first 45 (p, q) rounds, where p is plotted on the x-axis and q on the y-axis. A triangle indicates that the subject voted Blue, while a cross indicates that she voted Red. The number displayed to the right of the triangle or cross indicates the round number in which the subject encountered such (p, q) pair.

Subject 91 always votes Blue and is therefore strategic from the first round. Subject 46 votes Blue in all rounds after round 16; so we will say that subject 46 becomes strategic at $t = 17$. Subjects 35 and 59 are not strategic. All of subject 35's votes are sincere. On the other hand, subject 59 is a mix between sincere and conforming behavior (and, therefore, even farther from strategic behavior compared to a sincere voter). In particular, subject 59 behaves sincerely, except for low values of q and for values of p a bit below 5, where the probability that computers vote Red is above .5.

Restricting attention to Part 1 and characterizing subjects as being of a particular type if their votes agree with that type 90% of the rounds, it follows that 33% of the subjects are strategic or become strategic (e.g. subjects 91 and 46) and about 57% are non-strategic subjects that are either sincere, conforming, or a mix of these two types (e.g. subjects 35 and 59). The remainder 10% of subjects seem hard to categorize.[26] In what follows, we discuss the evidence for strategic and non-strategic behavior separately.

4.1.1 Strategic behavior

As illustrated by Subject 46 in Figure 3, we must account for the fact that subjects may become strategic only after several rounds of play. In addition, we must account for the fact that even strategic subjects may occasionally make a mistake and vote non-optimally.

Definition 4. A subject becomes z-strategic at round t in Part 1 if t is the first round in which her vote is optimal (i.e. Blue) in round t as well as in z% of the cases from round t until round 45.[27] A subject is strategic in Part 2 if her vote is optimal

[26] Only 5% cannot be categorized if we relax the definition to require consistency 80%, rather than 90%, of the rounds.

[27] Our definition of strategic behavior, which implicitly requires that a subject who becomes sophisticated in a round remains sophisticated in all later rounds, is justified by the fact that, once sophisticated, our subjects make essentially no mistakes: For $z = 95$, there are only 2 strategic subjects who do not satisfy the z-requirement at a later round.

in both rounds 46 and 47. A subject is strategic in Part $3k$, where $k \in \{a, b, c\}$, if her vote is optimal in all of the 4 rounds of that stage.

Figure 4 reports the main findings of the experiment. The left panel shows the cumulative percentage of subjects that become z-strategic for each of the 45 rounds in Part 1, for $z = 85$, $z = 95$ and 100. We only plot results until round $t = 40$, since that leaves 5 additional rounds that we can use to classify behavior. In round 1, only 10% of subjects play optimally every remaining period (100-strategic), while the percentage increases to 33% by round $t = 40$. Figure 4 also shows that more permissive definitions of strategic behavior that allow for a large fraction of mistakes ($z = 85$) also result in about 33% of subjects becoming strategic by round $t = 40$. This convergence by round $t = 40$ captures the fact that strategic subjects make almost no mistakes once they become strategic.[28]

The right panel of Figure 4 disaggregates the previous results by treatment and also provides information about the percentage of strategic subjects in Parts 2 and 3 of the experiment. In both treatments, the percentage of strategic subjects is close to 10% in the initial round. However, in the Feedback treatment, the fraction rapidly increases to reach 47% by round $t = 40$ compared to only 19% in the No Feedback treatment. There is a significant effect of feedback on the proportion of strategic subjects by round $t = 40$ (p-value < 0.001). In addition, while learning to be strategic flattens out in both treatments, learning stops much earlier in the No Feedback treatment.[29]

For both treatments, there is a small jump in strategic behavior in Part 2 (from 19% to 22% in the No Feedback treatment and from 47% to 50% in the Feedback treatment), which is explained by subjects who either become strategic in the last 5 rounds of Part 1 or when providing written advice.[30] In addition, very few subjects become strategic as a result of answering questions Q3a and Q3b in Parts 3a and 3b,

[28]The difference between $z = 100$ and $z = 85$ observed in the initials periods is explained by the fact that several subjects become strategic in later rounds.

[29]While the literature emphasizes the role of feedback in reducing strategic uncertainty (e.g., Armantier, 2004), our finding shows that feedback can help with cognitive tasks even in settings without uncertainty. Unfortunately, we do not have enough variation in the data to identify the events that trigger learning, so a full learning analysis is outside the scope of this experiment.

[30]Under the assumption that subjects who are not sophisticated by round 45 would not have become sophisticated if they had played the additional rounds in Part 2 *without* having to write down their advice, then the jump due to writing down the advice is not statistically significant (p-values 0.43 and 0.33 with and without feedback respectively). We make a similar assumption when testing the effect of feedback in Part 3.

respectively and the effect of these hints is not statistically significant (in fact, there is a decline in the Feedback treatment explained by a subject who became strategic in Part 2 but makes a mistake in Part 3a).[31] In contrast, the written explanation of pivotality in Part 3c has a larger effect, increasing the percentage of strategic behavior from 28% to 57% (p-value 0.00) in the No Feedback treatment and from 52% to 62% (p-value 0.06) in the Feedback treatment. Curiously, the percentage of strategic subjects increases to 75% if we only look at the first round of Part 3c, which is identical to the example that we use to explain the idea behind pivotal voting. But many subjects subsequently fail when the red and blue labels are switched.[32]

Finally, Table 3 provides further details by classifying subjects into strategic and not strategic at every part of the experiment and shows the transition from non-strategic to strategic behavior, aggregated over both treatments. There are four main findings. First, almost without exception, subjects classified as strategic in earlier parts continue to be classified as strategic in later parts. Second, subjects classified as strategic in Part 2 are more likely to answer questions correctly (35 of 42 subjects for Q3a and all 42 for Q3b) compared to non-strategic subjects (38 for Q3a and 62 for Q3b out of 76). Third, very few non-strategic subjects become strategic after answering Q3a (1 of 76) and Q3b (5 of 76). Fourth, non-strategic subjects are more likely to become strategic if they answer questions Q3a and Q3b correctly.

4.1.2 Non-strategic Behavior

In this section, we characterize behavior of non-strategic subjects. Our main finding is that most of the subjects that are classified as non-strategic can be classified as either sincere, conforming, or a mixture of these two types. First, we restrict attention to rounds in Regions 1 and 3 in Figure 2, where sincere and conforming votes coincide: Red in Region 1 and Blue in Region 3. Of the 79 subjects that are classified as non-strategic in Part 1 (at $t = 40$), 67 subjects vote in a manner consistent with sincere and conforming types in 90% of the rounds in Regions 1 and 3. Hence, the behavior of only 12 subjects cannot be explained in these regions.[33]

[31] The p-values are (i) for Q3a: 0.50 and 0.25; (ii) for Q3b: 0.35 and 0.1, with and without feedback, respectively.

[32] One possible reason why we obtain 75% and not 100% strategic behavior is that the explanation comes late into the experiment and, unlike the previous questions, subjects receive no direct monetary reward for carefully reading the explanation.

[33] When we decrease 90% to 80%, only 6 subjects cannot be explained.

Second, we evaluate to what extent these 67 subjects that behave as sincere and conforming types in Regions 1 and 3 can be classified as either sincere or conforming. We do so by focusing on Region 2, where a sincere vote is Blue but a conforming vote is Red. The top left panel of Figure 5 shows the frequency plot of the proportion of Blue votes for rounds (p, q) in Region 2 played in Part 1.[34] A proportion close to 1 indicates mostly sincere behavior, while a proportion close to 0 indicates mostly conforming behavior. While there is a spike in the frequency both near 0 and 1, there is also a spike around .5, suggesting that several subjects behave sometimes as conforming and sometimes as sincere. The bottom left panel of Figure 5 shows the same information, but restricted to the subset of subjects that remain non-strategic throughout the entire experiment. A similar picture emerges, except that now the spike around .5 disappears, suggesting that those subjects that are persistently non-strategic tend to behave more purely as either conforming or sincere subjects. The remaining panels disaggregate the findings for both treatments, showing that the spike in the middle of the distribution is mainly due to those subjects that do not receive feedback.

It is clear from our findings that, despite a substantial fraction of the subjects being sincere, there is also a non-negligible fraction of non-strategic subjects who do respond to the strategies of the computers. Therefore, while conforming behavior is likely to be specific to our setup, the general point is that one should look for alternatives beyond sincere behavior when testing for strategic behavior.

4.2 Experiment II: sequential voting

We adjust the definition of optimal voting to account for the fact that conclusions can only be drawn when a subject's vote is pivotal. That is, *if* computers voted for different colors, a vote is optimal given $(p, q) \in \Omega$ if it is a vote for Blue. Given the primitives, by design, subjects should face few pivotal cases in Part 1. In fact, the actual median number of cases in which the subject's vote is pivotal was 7. Moreover, the median number of pivotal cases in which optimal behavior differs from both sincere and conforming behavior (Rounds in Region 1 of Figure 2) was 1. While we do not use this part to classify subjects, Part 1 accomplishes two important objectives. First, we can compare Part 2 in the simultaneous and sequential treatments, since at that

[34] If $p = 5$, then a conforming vote is Red but a sincere vote can be either Red or Blue. Therefore, we only count this vote as evidence of sincere behavior if is is a Blue vote.

stage subjects have played for the same number of rounds. Second, Part 1 familiarizes subjects with the problem for which they will be giving advice in Part 3. Of particular importance is that, since the computers' votes are observed, Part 1 makes it clear to subjects that the probability of being pivotal is small, therefore allowing us to test whether the results for the simultaneous treatment are driven by the low stakes of the environment.

Figure 6 shows the findings for the sequential voting experiment. The first two rounds in the figure constitute Part 2. Applying Definition 4 to Part 2, we find that 78% of our subjects are classified as strategic, a number that is substantially higher compared to the simultaneous treatment without feedback (21%) (p-value 0.01).

In Part 3, subjects faced, on average, 41 pivotal cases. We modify Definition 4 in order to restrict attention to pivotal rounds.

Definition 5. A subject becomes z-strategic at round t in Part 3 if t is the first round in Part 3 in which her vote is both pivotal and optimal (i.e. Blue) in round t as well as in z% of the pivotal cases from round t until round 93.[35]

Figure 6 reports the findings using the above definition for three different values of z. The solid line tracks the cumulative percentage of subjects who become 100-strategic in Part 3. The starting value of 57% shows a drop with respect to the percentage of strategic subjects in Part 2 (72%). This drop is explained by a few mistakes made in Part 3 by a small group of subjects, as shown by the $z = 95$ dashed line.[36] In fact, the set of subjects classified as becoming sophisticated by the end of Part 3 for the $z = 100$ and $z = 95$ criteria differ only in two subjects who made a mistake in only one round.[37] Figure 6 also shows that most subjects have already become strategic before Part 3, and that the percentage of subjects that have become strategic by the end of Part 3 is around 80%.

The previous evidence shows that a large majority of our subjects are able to extract the right information from the votes of the computers and that the differences with respect to the simultaneous treatment (with and without feedback) are quite

[35]To avoid noise coming from the last rounds, we only consider subjects who satisfy the definition at most starting in Round 88, five rounds before Part 3 is over.

[36]The fact that more strategic subjects make mistakes in the sequential treatment may be explained by the lower incentives that exist as a result of subjects submitting only a recommendation to a partner.

[37]These subjects made a mistake in rounds 88 and 93.

large. In addition, these findings rule out the lack of incentives as a plausible explanation for the substantial lack of strategic behavior in the simultaneous treatments. In the sequential experiment, subjects gain experience (particularly in the first 45 rounds) regarding the low probability of being pivotal. Nevertheless, most subjects respond to the small incentives and behave strategically. The fact that subjects can behave strategically despite low stakes is also of independent interest for the voting literature, where a realistic feature of the environment is that people have a very small chance of affecting the outcome of an election.

An additional way to inspect the difference between the sequential and the simultaneous environments is to look at Part 4 of the sequential treatment, where subjects submit written advice for three simultaneous rounds after having faced 93 sequential rounds. We classify a subject as strategic in Part 4 if they recommend voting for Blue in all three rounds. As shown by Figure 6, only 21% of our subjects are classified as strategic in Part 4. Thus, even in the case where subjects have experience with the sequential version of the experiment, the proportion of strategic subjects falls by approximately 55 percentage points when voting becomes simultaneous.[38] Consequently, while most subjects correctly extract information when voting is sequential, only a small percentage (comparable to the percentage of strategic subjects in Experiment I without feedback) correctly extract information when voting is simultaneous.[39]

4.3 Experiment III: no private information

Because voting is simultaneous, we use the same definitions as in Experiment I. Overall, our findings for this experiment (summarized in Figure 7) are qualitatively similar to those for Experiment I. According to the left panel of Figure 7, only 8% of subjects are classified as 100−strategic at round $t = 1$. By round $t = 40$, the proportion is 36% (22 out of 61 subjects). Allowing for a larger fraction of mistakes ($z = 85$) results in about 44% of subjects becoming strategic by round $t = 40$.

The right panel of Figure 7 presents the data for all parts and to help comparisons also reproduces results for Experiment I. Although the proportion of strategic subjects

[38]The difference is statistically significant (p-value < 0.001). Not surprisingly, all of these subjects were also classified as strategic in Part 3.

[39]Although the results of Part 5 (where subjects receive advice before some rounds) are not reported in Figure 6, it is interesting to note that all subjects classified as becoming 100-strategic by the end of Part 3 would still be in the same category if we used Definition 5 for Part 5. In fact, the percentage of 100-strategic subjects increases to 90% towards the end of Part 5.

by round $t = 40$ is statistically higher compared to Experiment I without feedback (p-value 0.02), the main qualitative finding remains: there is a relatively large proportion of subjects (approximately 64%) who are classified as non-strategic.

In Part 2 of Experiment III, the proportion of strategic subjects jumps from 36% to 47% (p-value 0.11). Recall that, according to Definition 4, a subject is classified as strategic in Part 2 whenever her vote is optimal in cases 46 and 47, regardless of the quality of their advice. Indeed, when we inquire into the causes for the difference, the jump is largely due to a few subjects whose advice is optimal for the wrong reasons.

Hints and the explanation have a similar effect as in Experiment I. All but one of the subjects classified as strategic in Part 2 answer Q3a correctly and, while 69% of non-strategic subjects also succeed in giving a correct answer, the percentage of strategic subjects reaches only 38% and is not statistically different than the reported figure for round 40 of Part 1 (p-value 0.43). Similarly, almost half of our subjects who are classified as non-strategic in Part 3a answer Q3b correctly, but there is no significant effect from this hint (p-value 0.23). Finally, providing an explanation has a positive significant effect (p-value 0.05) on the percentage of strategic subjects. According to the four rounds faced after the explanation, 31 of our 61 subjects (51%) are classified as strategic, with a higher percentage (67%) that answer the first round correctly.[40]

4.4 Robustness checks

We begin by using the subjects' written advice to confirm our findings and then rule out two alternative explanations to our findings: misunderstanding of the instructions and failure to choose appropriately between lotteries.

4.4.1 Written advice

Previously, we used the choices during the advice rounds to conclude that it was rarely the case that encouraging further reflection about the problem led to more strategic behavior. We now also use the specific written advice to infer whether people provide

[40]Recall that the first case of the final four subjects face is identical to the one provided in the explanation. Although slightly lower, the percentage of overall success in the first case in Experiment III (67%) is not statistically different than that of Experiment I (74%).

the correct logic.[41] In all experiments, *all* subjects classified as strategic in previous sections (where only voting choices were used to classify subjects) provided a correct explanation of optimal behavior when giving written advice to a potential partner. Therefore, the written advice confirms the findings in previous sections.[42]

4.4.2 Understanding of the instructions

As with all experiments, a possible concern is that subjects did not understand the instructions. There are several reasons why a lack of understanding of the instructions is unlikely to explain our findings. First, we provided detailed instructions followed by several incentivized questions about the instructions as well as explanations of the right answers (see the Online Appendix for details).[43] Second, most subjects answer these questions correctly. Of the 118 subjects that participate in Experiment I, 98 subjects answer these questions correctly (either in the first or second trial).[44] When we take out one session where the questions were not incentivized, 88 subjects answer correctly and only 9 answer incorrectly. Similarly, of the 58 and 61 subjects that participate in Experiments II and III, 53 and 54 subjects answer correctly, respectively.

Third, in Experiment I, feedback seems to mostly affect people that provide correct answers to the questions in the instructions. Of the 58 subjects in the feedback treatment, 52 correctly answer the questions. Of these subjects, 27 become strategic. Of the 6 subjects who answer incorrectly, only 1 becomes strategic. This result suggests that the reason why more people become strategic in the Feedback treatment of Experiment I is not the lack of understanding of the instructions.

Finally, in Experiment II (sequential voting), subjects are asked to give advice for the case where both computers vote for the same color. According to the instructions, in this case the vote of the subject will be irrelevant in determining the outcome. We

[41] For details on different ways to use advice see Schotter (2003) and references therein. Our advice data was classified by us first and then by an independent research assistant. We verified that both accounts agree. The protocol that was given to the research assistant to process the data is available from the authors upon request.

[42] Also, through the written advice we discovered that some subjects were making "conforming" calculations consistent with Definition 3.

[43] In one session of the simultaneous No Feedback treatment, we did not provide incentives to answer the questions and obtained a higher proportion of incorrect answers.

[44] If the answer to a question is correct, the subject gets \$1 and learns that her answer is correct. If the answer is incorrect, the subject is provided with an explanation of the correct answer and is given a second chance to answer a similar question. Then, irrespective of how she answers this second time, the subject begins the experiment.

test whether subjects understand this aspect of the instructions by comparing the percentage of times that a subject votes the color voted by the computers when her vote is irrelevant.[45] While we find a bias in the direction of voting for the same color voted by the computers, we find that this percentage is actually (slightly) higher for subjects classified as strategic (94%) compared to non-strategic subjects (88%). For further reassurance, in the final session of Experiment II, we include a question after the instructions (see the Online Appendix) that more explicitly tests subjects' understanding of the notion of a simple majority (even if their vote is not in the majority) and the payoff implications. This question is answered correctly by all of our subjects, and we observe essentially no difference in the results between the session where this question is included and the sessions where it is not.

4.4.3 Choosing between lotteries

The mistake of failing to extract information from others' strategies can be viewed as one explanation behind the more general failure to reduce a decision problem to a choice between lotteries. For example, for Experiment I, it is easy to check that voting Red is equivalent to the compound lottery

$$\frac{p}{10} \cdot [1] \oplus (1 - \frac{p}{10}) \left[q^2 \cdot 1 \oplus (1 - q^2) \cdot 0\right] \tag{3}$$

and voting Blue is equivalent to

$$\frac{p}{10} \cdot [1] \oplus (1 - \frac{p}{10}) \left[\left(q^2 + 2q(1-q)\right) \cdot 1 \oplus (1 - q^2 - 2q(1-q)) \cdot 0\right]. \tag{4}$$

By inspection, the second lottery dominates the first. In deriving our conclusions, an implicit assumption is that people satisfy the dominance axiom when facing lotteries (3)-(4). As shown in the literature, this assumption is generally true when the choices are transparent, as in (3)-(4), but may not be true when the framing is more complicated (e.g., Tversky and Kahneman (1986), Beard and Beil (1994), Camerer (2003, Chapter 5)). Thus, the voting experiments can be more generally viewed as a particular framing under which the dominance axiom is not satisfied.[46]

[45] About half of our subjects explicitly write that their recommendation is irrelevant when computers vote for the same color, but, of course, we cannot draw conclusions from those subjects who do not explicitly mention the irrelevance.

[46] Another benefit of our experiment is that, in more general voting settings, the comparison between the relevant lotteries would be complicated by other factors, such as failure of the independence

5 Conclusion

The experiments in this paper provide direct evidence that a substantial amount of people fail to extract information from others' strategies in a voting environment. Depending on the amount of feedback received, between 50% and 80% of subjects behave non-strategically when voting is simultaneous. In addition, mistakes are mainly driven by difficulty extracting information from hypothetical events, since most subjects are indeed able to extract information when they have knowledge of others' actual votes. Finally, we show that the failure to extract information from others' strategies is a more general phenomenon that is also present in settings without private information. Overall, our findings suggest that information extraction from hypothetical events is a challenging task, and that, despite facing facing a simple voting problem and obtaining feedback, hints, an explanation of optimal behavior, and even experience with the sequential version of the experiment, a substantial amount of people still remain unable to apply the right logic.

References

Ali, S.N., J.K. Goeree, N. Kartik, and T.R. Palfrey, "Information aggregation in standing and ad hoc committees," *The American Economic Review*, 2008, *98* (2), 181–186.

Armantier, O., "Does observation influence learning?," *Games and Economic Behavior*, 2004, *46* (2), 221–239.

Battaglini, M., R.B. Morton, and T.R. Palfrey, "Information aggregation and strategic abstention in large laboratory elections," *The American Economic Review*, 2008, *98* (2), 194–200.

_ , _ , and _ , "The Swing Voter's Curse in the Laboratory," *Review of Economic Studies*, 2010, *77* (1), 61–89.

Beard, T.R. and R.O. Beil, "Do people rely on the self-interested maximization of others? An experimental test," *Management Science*, 1994, pp. 252–262.

axiom (see Eliaz et al. (2006)) or risk preferences.

Camerer, C., *Behavioral game theory: Experiments in strategic interaction*, Princeton University Press, 2003.

_ **and D. Lovallo**, "Overconfidence and excess entry: An experimental approach," *The American Economic Review*, 1999, *89* (1), 306–318.

Camerer, C.F., T.H. Ho, and J.K. Chong, "A cognitive hierarchy model of games," *The Quarterly Journal of Economics*, 2004, *119* (3), 861–898.

Charness, G. and D. Levin, "The origin of the winner's curse: a laboratory study," *American Economic Journal: Microeconomics*, 2009, *1* (1), 207–236.

Crawford, V.P. and N. Iriberri, "Level-k Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?," *Econometrica*, 2007, *75* (6), 1721–1770.

Eckel, C. and C. Holt, "Strategic voting in agenda-controlled committee experiments," *American Economic Review, Vol. 79, No. 4, p. 763, 1989*, 1989.

Eliaz, K., D. Ray, and R. Razin, "Choice shifts in groups: A decision-theoretic basis," *The American economic review*, 2006, *96* (4), 1321–1332.

Esponda, I., "Behavioral equilibrium in economies with adverse selection," *The American Economic Review*, 2008, *98* (4), 1269–1291.

_ **and D. Pouzo**, "Conditional retrospective voting in large elections," *working paper*, 2012.

Eyster, E. and M. Rabin, "Cursed equilibrium," *Econometrica*, 2005, *73* (5), 1623–1672.

Feddersen, T. and W. Pesendorfer, "Voting behavior and information aggregation in elections with private information," *Econometrica*, 1997, pp. 1029–1058.

_ **and** _ , "Convicting the innocent: The inferiority of unanimous jury verdicts under strategic voting," *American Political Science Review*, 1998, pp. 23–35.

Feddersen, T.J. and W. Pesendorfer, "The swing voter's curse," *The American economic review*, 1996, pp. 408–424.

Fischbacher, U., "z-Tree: Zurich toolbox for ready-made economic experiments," *Experimental Economics*, 2007, *10* (2), 171–178.

Goeree, J.K. and L. Yariv, "Conformity in the Lab," *working paper*, 2007.

Guarnaschelli, S., R.D. McKelvey, and T.R. Palfrey, "An experimental study of jury decision rules," *American Political Science Review*, 2000, pp. 407–423.

Holt, C.A. and R. Sherman, "The loser's curse," *The American Economic Review*, 1994, *84* (3), 642–652.

Huck, S., P. Jehiel, and T. Rutter, "Feedback spillover and analogy-based expectations: A multi-game experiment," *Games and Economic Behavior*, 2011, *71* (2), 351–365.

Ivanov, A., D. Levin, and M. Niederle, "Can relaxation of beliefs rationalize the winner's curse?: an experimental study," *Econometrica*, 2010, *78* (4), 1435–1452.

Jehiel, P., "Analogy-based expectation equilibrium," *Journal of Economic theory*, 2005, *123* (2), 81–104.

_ **and F. Koessler**, "Revisiting games of incomplete information with analogy-based expectations," *Games and Economic Behavior*, 2008, *62* (2), 533–557.

Kagel, J.H. and D. Levin, "The winner's curse and public information in common value auctions," *The American Economic Review*, 1986, pp. 894–920.

_ **and** _ , *Common value auctions and the winner's curse*, Princeton Univ Pr, 2002.

_ **, R.M. Harstad, and D. Levin**, "Information impact and allocation rules in auctions with affiliated private values: A laboratory study," *Econometrica*, 1987, pp. 1275–1304.

Levin, D., J.H. Kagel, and J.F. Richard, "Revenue effects and information processing in English common value auctions," *The American Economic Review*, 1996, pp. 442–460.

Levine, D.K. and T.R. Palfrey, "The paradox of voter participation? A laboratory study," *American Political Science Review*, 2007, *101* (1), 143.

Morton, R.B. and J.R. Tyran, "Let the experts decide? Asymmetric information, abstention, and coordination in standing committees," *Games and Economic Behavior*, 2011, *72* (2), 485–509.

Schotter, A., "Decision making with naive advice," *American Economic Review*, 2003, pp. 196–201.

Thaler, R.H., "Anomalies: The winner's curse," *The Journal of Economic Perspectives*, 1988, *2* (1), 191–202.

Tversky, A. and D. Kahneman, "Rational choice and the framing of decisions," *Journal of business*, 1986, pp. 251–278.

Vespa, E. and A.J. Wilson, "Communication with Multiple Senders and Multiple Dimensions: An Experiment," *working paper*, 2012.

Walker, J.M., V.L. Smith, and J.C. Cox, "Bidding behavior in first price sealed bid auctions:: Use of computerized Nash competitors," *Economics Letters*, 1987, *23* (3), 239–244.

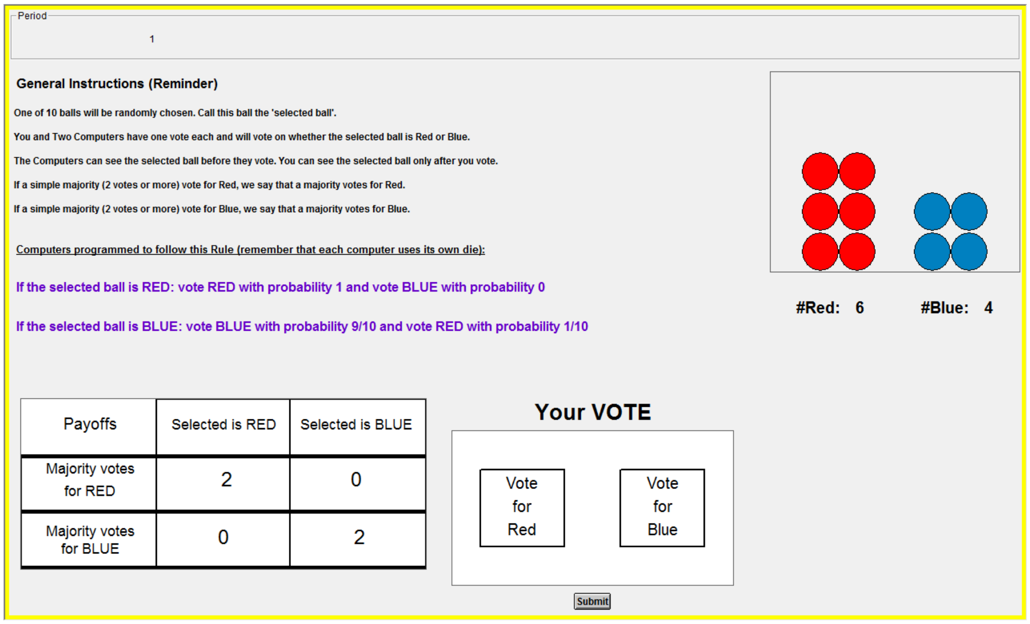

Figure 1: Screen shot for the experiment.

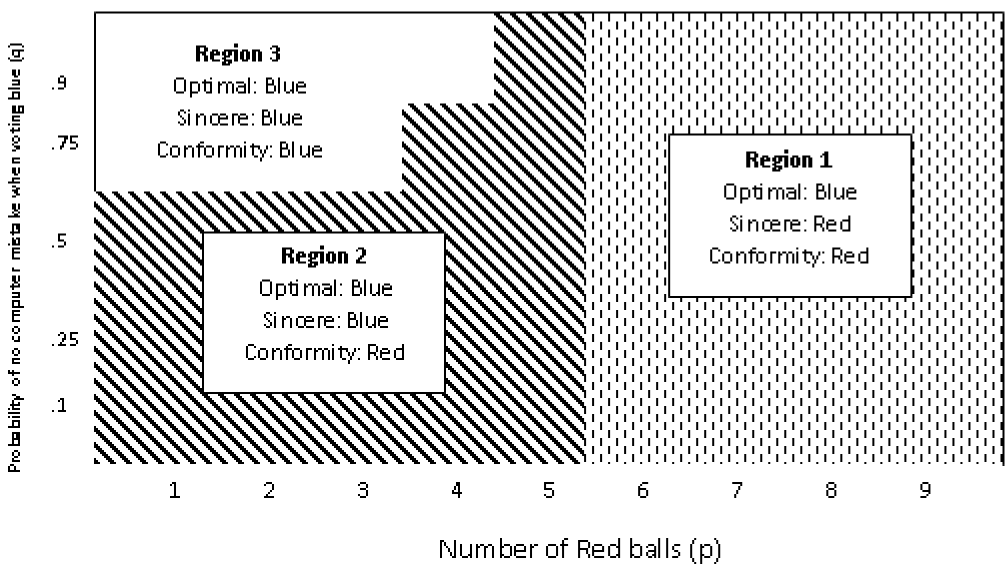

Figure 2: Regions of voting behavior by type.

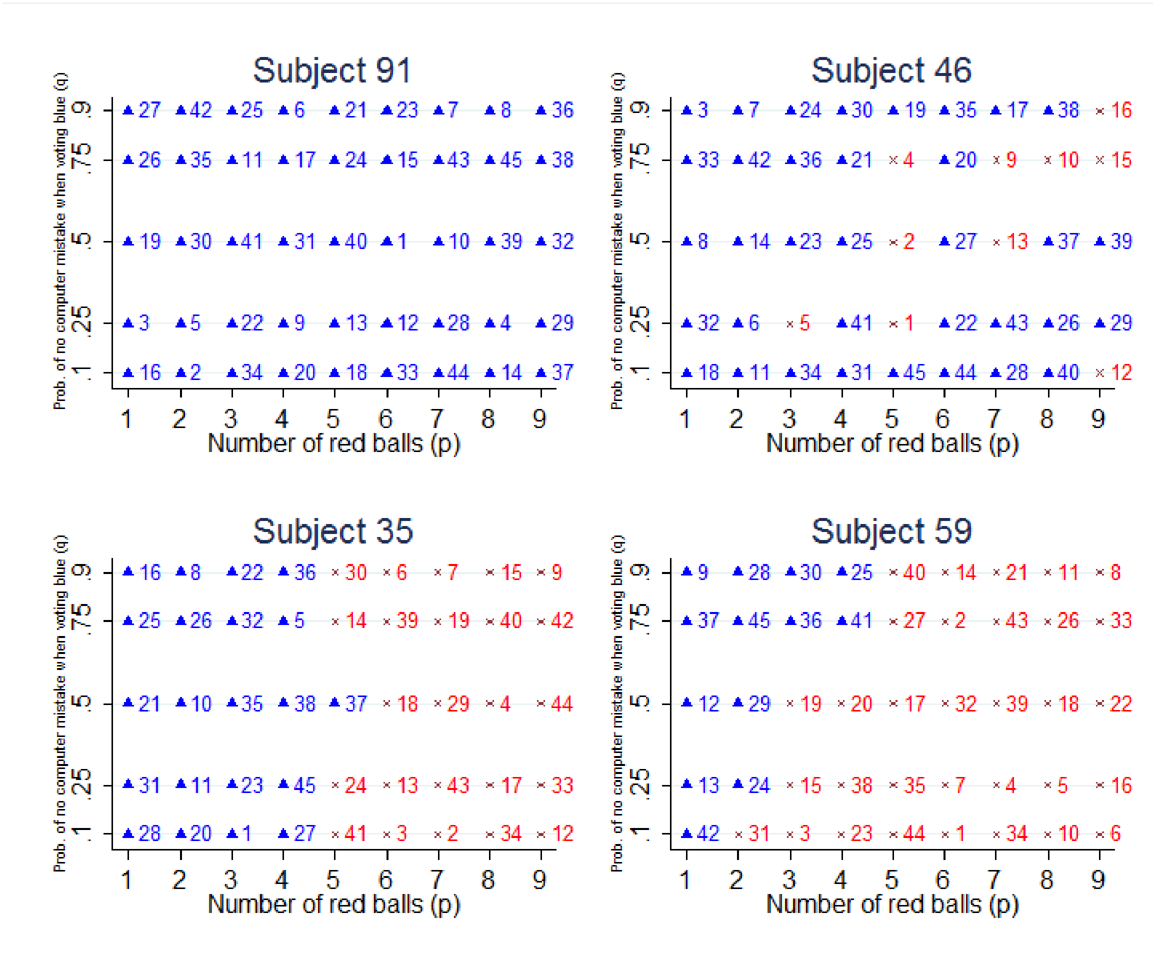

Figure 3: Representative types in the population.
A triangle represents a vote for Blue, a cross a vote for Red. The accompanying number denotes the round.

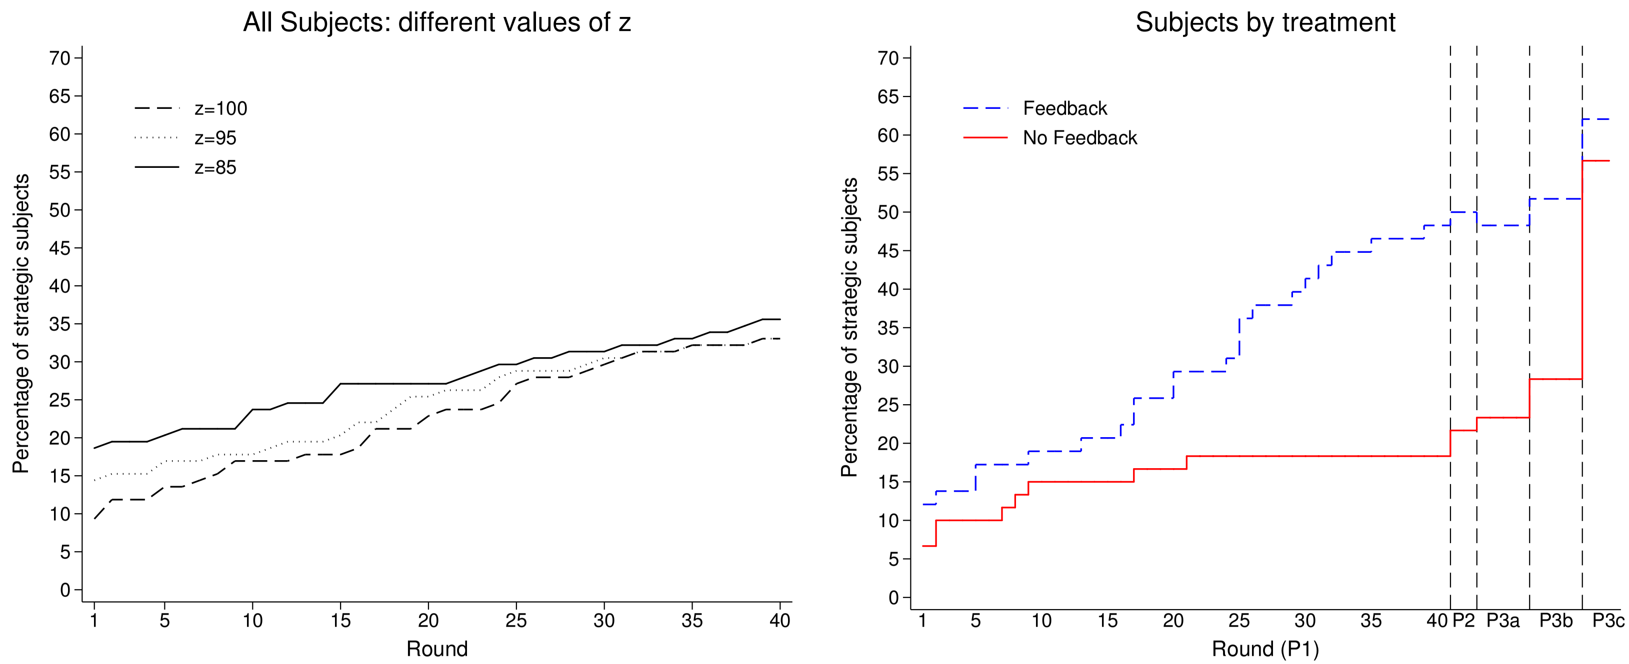

Figure 4: Strategic behavior in Experiment I.

			Strategic according to..							
			Part 2		Part 3a		Part 3b		Part 3c	
			Strat.	Not Strat.	Strat.	Not Strat.	Strat.	Not Strat.	Strat.	Not Strat.
Part 1: t=40	Strategic		39	0	39	0	39	0	39	0
	Not Strategic		3	76	3	76	8	71	31	48
Part 2: Advice	Strategic	Q3a corr.	-	-	35	0	35	0	34	1
		Q3a incorr.	-	-	6	1	7	0	7	0
	Not Strat.	Q3a corr.	-	-	1	37	4	34	22	16
		Q3a incorr.	-	-	0	38	1	37	7	31
Part 3a	Strategic	Q3b corr.	-	-	-	-	42	0	41	1
		Q3b incorr.	-	-	-	-	0	0	0	0
	Not Strat.	Q3b corr.	-	-	-	-	5	57	26	36
		Q3b incorr.	-	-	-	-	0	14	3	11
Part 3b	Strategic		-	-	-	-	-	-	46	1
	Not Strategic		-	-	-	-	-	-	24	47
Total			42	76	42	76	47	71	70	48

Table 3: Transitions from non-strategic to strategic behavior in Experiment I.

The table shows the transition from non-strategic to strategic behavior in all treatments of Experiment I. For example, row 1 and column 1 indicate how subjects classified as either strategic or not in Part 1 (the row) are then classified according to Part 2 (the column). The second and third rows disaggregate results depending on the answers to questions Q3a and Q3b asked at the beginning of Parts 3a and 3b, respectively. For example, the numbers in row 2 and column 2 indicate: (i) whether subjects correctly answered question Q3a depending on how they were classified in Part 2, and (ii) how subjects transition into being classified as strategic or not according to behavior in the 4 rounds of Part 3a.

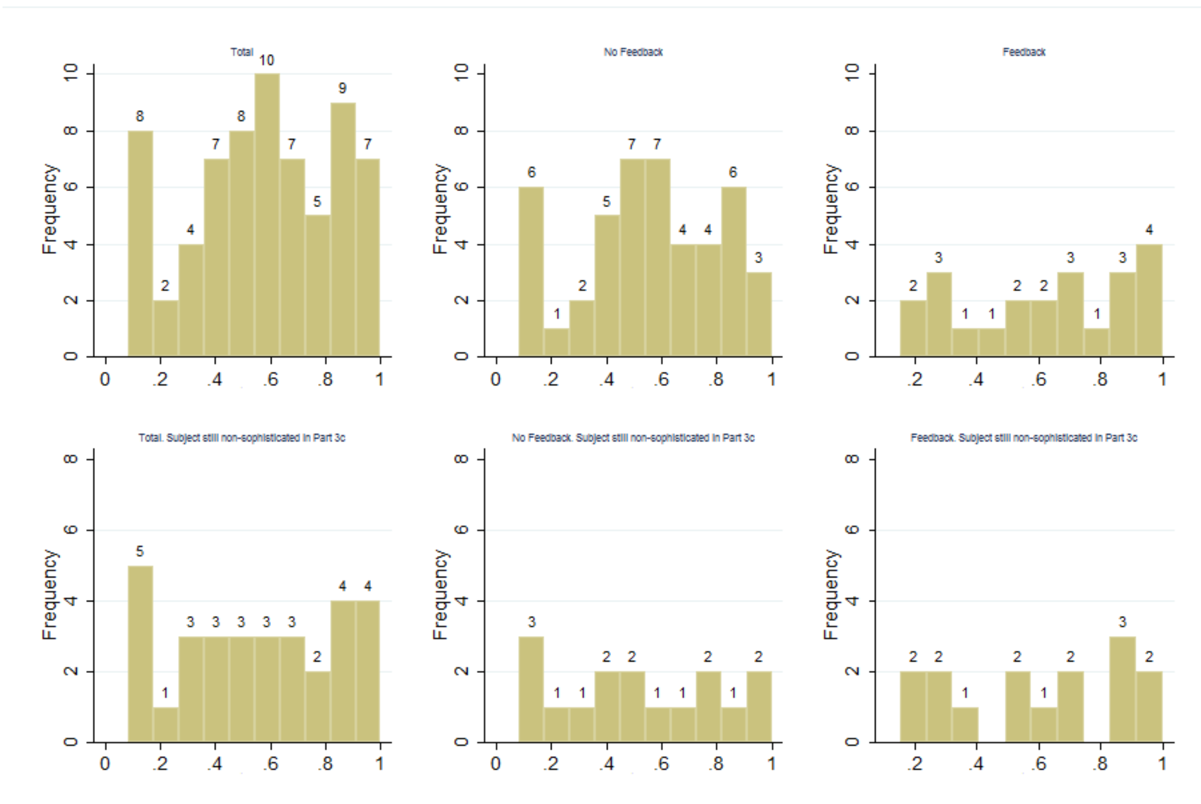

Figure 5: Non-strategic behavior: sincere vs. conformity.

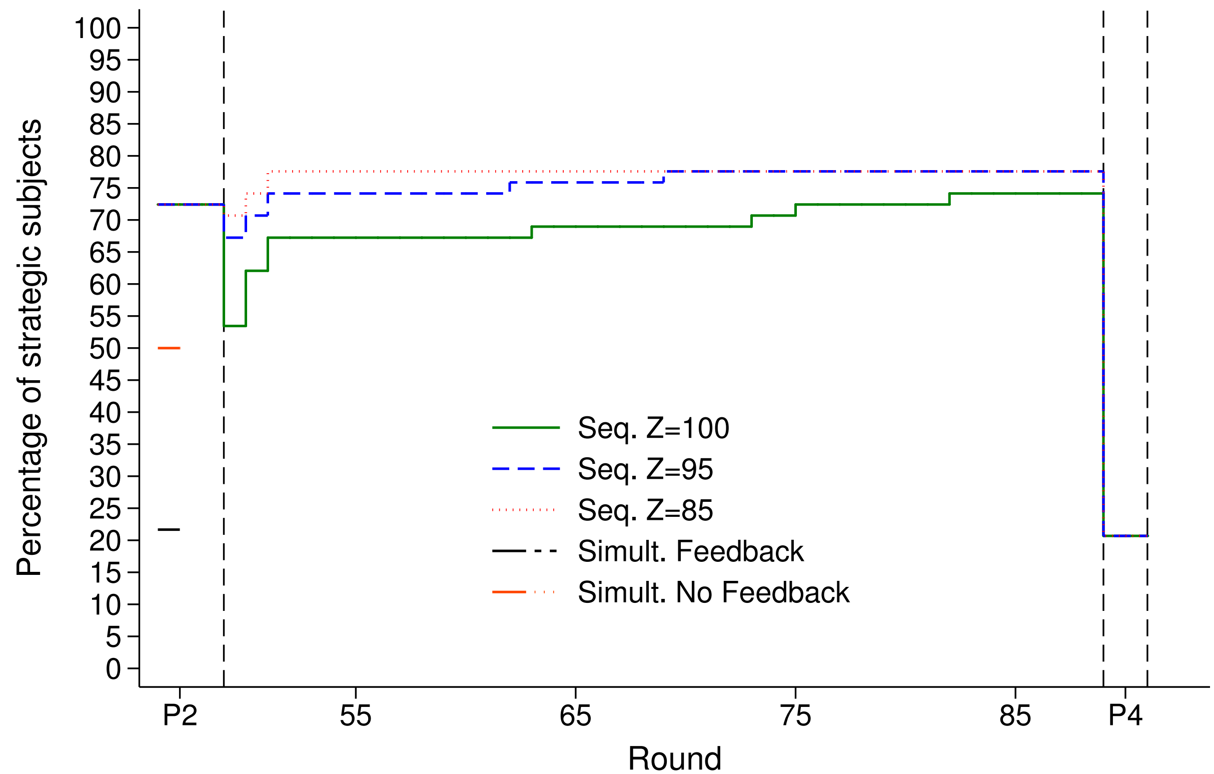

Figure 6: Strategic behavior in sequential experiment.

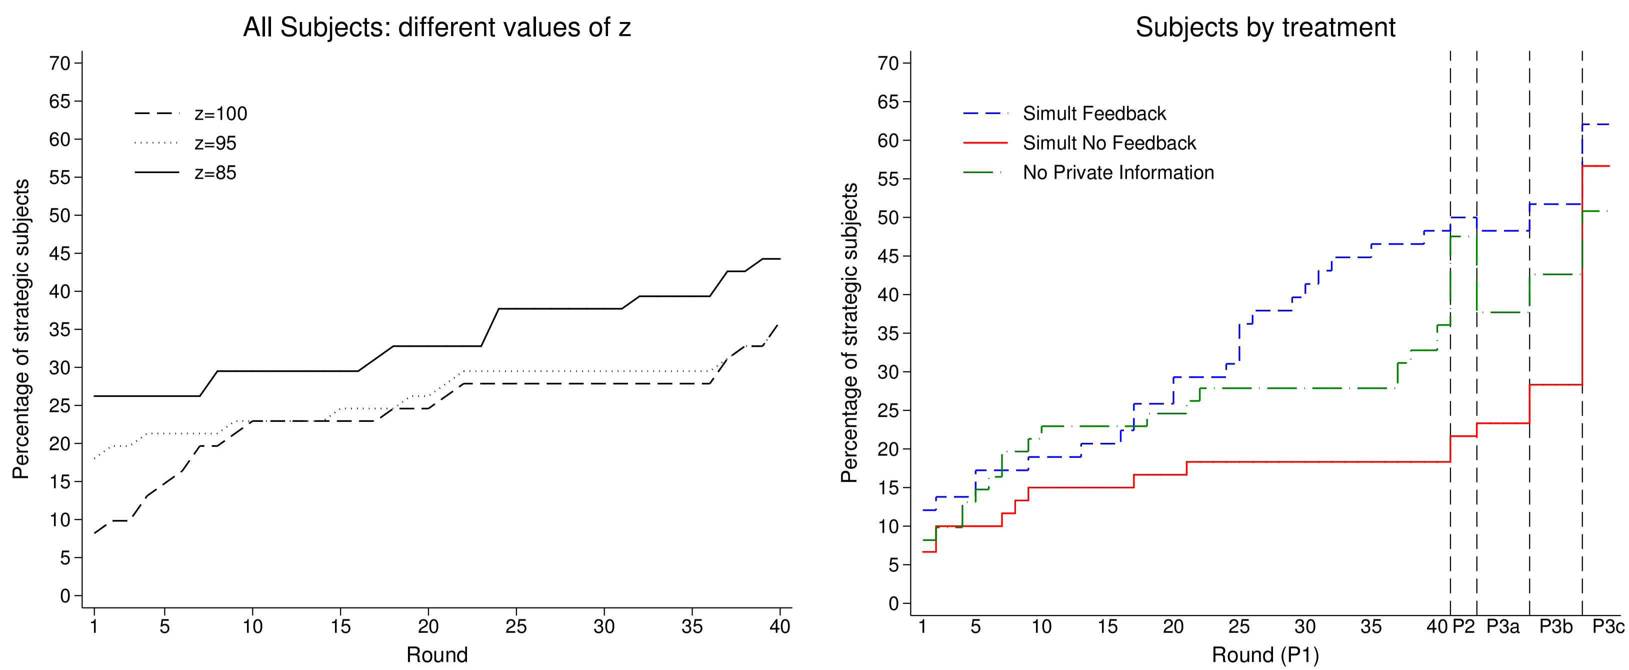

Figure 7: Strategic behavior: Experiment I vs III