

Why Votes Have Value:
**Instrumental Voting with Overconfidence
and Overestimation of Others' Errors**

Ingolf Dittmann^a

Dorothea Kübler^b

Ernst Maug^c

Lydia Mechtenberg^d

November 14, 2013

Abstract

We perform an experiment in which subjects bid for participating in a vote. The setting precludes conflicts of interests or direct benefits from voting. The theoretical value of participating in the vote is therefore zero if subjects have only instrumental reasons to vote and form correct beliefs. Yet, we find that experimental subjects are willing to pay for the vote and that they do so for instrumental reasons. The observed voting premium in the main treatment is high and can only be accounted for if some subjects either overestimate their pivotality or do not pay attention to pivotality at all. A model of instrumental voting, which assumes that individuals are overconfident and that they overestimate the errors of others, is consistent with results from treatments that make the issue of pivotality salient to experimental subjects.

JEL Classification: C92, D72, G32

Keywords: Voting, experimental economics, overconfidence, pivotality, cognitive biases

^a Erasmus University Rotterdam, P.O. Box 1738, 3000 DR, Rotterdam, The Netherlands. E-mail: dittmann@ese.eur.nl. Tel.: +31 10 408 1283.

^b (Corresponding author) WZB Berlin, Reichpietschufer 50, 10785 Berlin, and Technical University of Berlin, Faculty of Economics and Management, Straße des 17. Juni 135, 10623 Berlin. E-mail: d.kuebler@ww.tu-berlin.de, Tel.: +49 30 25491 440.

^c University of Mannheim, Department of Business Administration, 68131 Mannheim, Germany. E-mail: maug@bwl.uni-mannheim.de. Tel: +49 621 181 1952.

^d University of Hamburg, Von-Melle-Park 5, 20148 Hamburg, Germany. Email: lydia.mechtenberg@wiso.uni-hamburg.de, Tel.: +49-30-25491-446.

1 Introduction

A longstanding literature investigates when and why people value participation in a vote, especially when they seem to have only a negligible impact on the outcome¹. Moreover, it is still open to debate whether purely instrumental reasons can explain why people often incur costs in order to vote. Alternatively, non-instrumental reasons such as expressive motives or a civic duty may be necessary to account for why people vote. To address these questions, we experimentally elicit the willingness to pay for a vote in a setting where conflicts of interest are absent, information is symmetric, and direct benefits from voting are negligible. Hence, we measure the willingness to pay for a vote when it should be zero. We thereby contribute to a lively debate in the literature about whether the instrumental approach to voting, as adopted in the pivotal voter model, is sufficient to explain the high turnout rates often observed in reality.² Indeed, both in political elections (see Matsusaka and Palda (1993)) and votes on worker representation (see Farber (2010)), turnout is so high that it becomes difficult to rationalize why individuals are willing to incur costs in order to vote.³ For shareholder votes, the value of the voting right is reflected in stock market data and has been shown to be substantial (Adams and Ferreira (2008)).

The literature addresses the challenges to the pivotal voter model in different ways. One strand proposes explanations that are based on alternative models of preferences, in particular expressive motives and other-regarding preferences.⁴ Another strand of the literature postulates

¹ This debate goes back to Downs (1957) and Tullock (1967). See also Friedman (1995) and Green and Shapiro (1994), pp. 47-48.

² As Ferejohn and Fiorina (1974) put it: “The question is (...) whether the principal explanation of the voting act is found in its investment aspect or its consumption aspect. If it is in the former, then one type of rational choice model is appropriate. If in the latter, another theory - the theory of consumer choice - is relevant.” A prominent example of instrumental explanations of voting is the pivotal-voter model of Palfrey and Rosenthal (1983), which relies on conflicts of interest in a political setting. The notion of non-instrumental voting goes back at least to Fiorina (1976) who credits Butler and Stokes (1969) with the distinction between instrumental and non-instrumental explanations and refers to the “older tradition” in political science (i.e., the tradition preceding Downs, 1957) as supporting a non-instrumental view of voting. The “D-term” in Riker and Ordeshook (1968) is generally thought to capture (unspecified) non-instrumental reasons to vote.

³ Farber (2010) finds turnout rates between 80% and 95% for votes among employees in U.S. private sector firms on whether or not they want to be represented by a union. Similar rates (between 75% and 80%) are observed for elections of worker representatives in German private sector firms (see Böckler Impuls 16/2006, http://www.boeckler.de/32014_84303.html).

⁴ For an early formalization of the expressive voting theory, see Brennan and Hamlin (1998). Tyran (2004) extends this theory and provides an experimental test. Feddersen (2009) provides experimental evidence for expressive ethical motivations of voters that increase with electorate size. Höchtl, Sausgruber and Tyran (2012) experimentally show that fairness preferences motivate rich voters to vote for redistribution. Morton and Tyran (2012) provide evidence from internet lab experiments that voters exhibit instrumental ethical motivations.

that individuals overestimate their probability of being pivotal.⁵ In this paper, we show that individuals value participating in the vote even though their motives are purely instrumental and their impact on the outcome is small. They pay a significant premium in order to vote if and only if the *outcome* of the vote affects their payoffs, even though their *participation* in the vote has only a negligible impact. To explain this seemingly paradoxical result, we suggest that individuals are overconfident in that they underestimate their own probability of making mistakes, while they overestimate the error rates of others. In addition, a significant minority of the experimental subjects seems to be inattentive to the relevance of their pivotality and therefore willing to overpay for the vote.

Our experiment is framed as shareholder voting and designed in a way that leaves little room for direct benefits from voting. Shareholder voting provides a natural context for implementing a common interest setting and for eliciting individuals' willingness to pay for participating in a vote. Our experiment considers a firm with two classes of shares that are auctioned off at the beginning. One class of shares votes and the other class does not vote. After observing a public signal about the quality of the current manager, those experimental subjects who bought the voting shares decide on the replacement of the manager. The quality of the manager in charge after the vote determines the size of a dividend that is paid out to all shareholders. Non-voting shares receive exactly the same dividends as voting shares and information is symmetric.

We develop several benchmark solutions of the game and assume instrumental voting throughout. We begin by assuming that subjects make no errors when they vote and when they predict their own and other subjects' behavior. Subsequent solutions successively relax this assumption. Hence, the starting point is that subjects make no errors and therefore all make the same choices in equilibrium. As a result, they are pivotal with probability zero, attach no value to voting, and the premium for voting shares is zero. In a second step we apply the Quantal Response Equilibrium model (QRE) of McKelvey and Palfrey (1995, 1998), in which individuals make random mistakes and follow strategies that are best responses in such a context. In this model, the probability of being pivotal becomes strictly positive. However, as long as all subjects have the same error rates and these are the same for all individuals, the voting premium subjects

⁵ Already Riker and Ordeshook (1968) argue that voters may overestimate their pivot probability. There is still little experimental evidence on this claim. One notable exception is the contribution of Duffy and Tavits (2008).

should pay in equilibrium is still zero, because they have the same preferences. Positive probabilities of being pivotal by themselves are therefore insufficient to generate a willingness to pay for the vote.

The first two benchmark solutions we consider predict that there should be no voting premium in equilibrium. However, our experimental subjects do attach a significant value to participating in the vote. We therefore analyze the possibility that individuals may have non-instrumental motivations to vote and conduct two control treatments in which voting does not affect payoffs. We find no significant voting premium in these treatments and conclude that subjects value their participation in the vote for purely instrumental reasons. We therefore need to reconcile instrumental voting with individuals' willingness to pay for a vote in a setting in which standard assumptions predict that a voting premium should not exist.

Our point of departure is the QRE model. We introduce overconfidence by assuming that individuals believe that only the other participants in the game make mistakes. In this model we obtain a voting premium, but it is orders of magnitude smaller than the voting premium in our experimental data. The reason is that a single shareholder is pivotal only with a relatively small probability, so even an overconfident shareholder should not be willing to pay much for participating in the vote. Overconfidence by itself is therefore insufficient to explain our results.

Next, we also assume that subjects overestimate the probability that other subjects make mistakes in the voting stage of the game. The level-k model of Stahl and Wilson (1995) and Nagel (1995) nicely captures both, overconfidence and overestimating the errors of others. As a consequence, individuals can legitimately assume higher probabilities of being pivotal in this model than in the QRE model with overconfidence. Accordingly, we find that this model is able to explain much higher voting premiums than the QRE model with overconfident players. It can even explain the entire premium in some of our experimental treatments.

The results for the level-k model are promising, because overconfidence together with an overestimation of the errors of other players can explain a significant portion of the voting premium. However, this explanation relies on individuals making counterfactual assumptions about the errors of other players and about their own errors. In all previous treatments individuals may hold incorrect beliefs, because they receive only limited feedback at the voting stage. We therefore repeat the experiment with a setting that provides more feedback about voting behavior. Voting premiums decline dramatically, by about 60% relative to the baseline treatment. However,

subjects' willingness to pay for the vote can still not be fully rationalized by beliefs that are consistent with the observed play of the game.

Finally, we repeat the experiment again, but now elicit subjects' beliefs about their pivotality directly. Four findings emerge from this treatment. First, bidding behavior in the game changes significantly relative to the baseline treatment, and it does so in almost exactly the same way as in the treatment with additional feedback. We therefore observe that providing subjects with additional information relevant to their pivotality and prompting them to reflect on their pivotality have the same behavioral impact. Second, the premium individuals are willing to pay for participating in the vote is correlated and mostly consistent with their beliefs about being pivotal. Only about one quarter of the subjects overbid relative to their beliefs. Third, on average individuals overestimate their probability for being pivotal. Fourth and last, beliefs about being pivotal vary markedly across subjects: The majority believes not to be pivotal at all, while about 30% hold beliefs that are higher than what can be rationalized by the level-k model.

We contribute by showing that non-instrumental motives are not necessary to explain why individuals value a vote. We suggest that individuals simultaneously underestimate their own error rates and overestimate the error rates of others. Both mistakes are consistent with the level-k model and they are complementary because none of them would generate any willingness to pay for a vote by itself. We also show that a significant minority of individuals reduce their willingness to pay for a vote if they are prompted to reflect on their pivotality. Interestingly, this reduction in the willingness to pay obtains independently of whether this prompt is associated with additional information or not.

A number of experimental studies investigate behavioral aspects of voting in the political sphere. The pivotal voter model of Palfrey and Rosenthal (1983) has been used by Schram and Sonnemans (1996a, 1996b) who test its comparative static predictions.⁶ Duffy and Tavits (2008) elicit subjects' beliefs and find that subjects overestimate the probability of being pivotal. Several other experimental papers vary the probability of being pivotal to test the theory of expressive voting.⁷ However, these studies differ from our paper in important respects: First, none of them

⁶ In contrast to these papers and a number of earlier contributions, Palfrey and Rosenthal (1985) and Levine and Palfrey (2007) introduce heterogeneous costs, which lead to a unique equilibrium of the pivotal voter model.

⁷ This theory holds that voters tend to vote more expressively the less they believe their vote to be decisive (see Brennan and Hamlin (1998) and Brennan and Lomasky (1993)). Feddersen, Gailmard, and Sandroni (2009) and

measures to what extent individuals value participation in a vote.⁸ Also, none of them excludes non-instrumental motives. Third, in all studies voters differ in their preferences over the outcome of the vote, unlike our setup. With the exception of Duffy and Tavits (2008), no other study elicits beliefs about pivotality. We are the first to show how pivotality relates to the value individuals attach to voting.

Another set of papers analyzes voter turnout when there is asymmetric information (Feddersen and Pesendorfer (1996, 1998, 1999) and Guarnaschelli, McKelvey, and Palfrey (2000)). Battaglini, Morton, and Palfrey (2006) provide a test of the swing-voter model by Feddersen and Pesendorfer (1996) and find that less informed voters delegate voting to better informed voters. In contrast to these studies, participants in our experiment have symmetric information, which simplifies the strategic considerations.

2 The model

2.1 Setup of the model

General setup. The game has N investors with an initial endowment of cash who can bid for shares in a company. There are two classes of shares: A-shares, which give shareholders a vote in the company, and otherwise identical B-shares that do not. A-shares and B-shares both pay out the same dividends and the number of A-shares equals the number of B-shares, which is $M < N$. At the beginning investors bid for shares in the company, and each investor can gain at most one share of each class. There are two periods $t = 1, 2$, and the firm pays a dividend D_t in each period. Dividends depend on the quality of the manager and on a state of nature. This basic setup is the same for both periods. After observing the dividend paid in the first period, shareholders vote on whether the manager will run the firm again in the second period or whether she will be replaced by a new manager.

Technology and dividends. Managers are drawn from a pool, and the number of good managers and bad managers in the pool is the same. If the manager is good, then the dividend in

Shayo and Harel (2010) support this hypothesis whereas Tyran (2004) as well as Kamenica and Egan (2010) contradict it.

⁸ We know of only one field experiment on the value of voting. Güth and Weck-Hannemann (1997) try to elicit the value of the voting right in the context of a political election. Schram's (1997) comment highlights the limitations of this approach. Apart from this paper, none of the studies on voting behavior attempts to gauge individuals' valuation of a vote.

period t is high ($D_t = H$) with probability $p \geq 0.5$. In the complementary state, the dividend is low ($D_t = L < H$) with probability $1 - p$. If the manager is bad, then $D_t = L$ with probability p and $D_t = H$ with probability $1 - p$. Investors know the probability p , but not the quality of the manager. Since firms draw a good manager or a bad manager from the pool with equal likelihood, the posterior probability of having a good manager is therefore p if $D_t = H$ and $1 - p$ if $D_t = L$.

Voting. Shareholders observe the dividend in the first period and then vote on the replacement of the manager by majority vote. Only owners of A-shares vote. They have one vote per share and cannot abstain from voting. For simplicity, we assume that M is odd, so the manager will be replaced whenever at least $(M + 1) / 2$ votes are cast for the replacement of the manager. Then a new manager is drawn from the same pool for the second period, so the new manager is again good or bad with equal probability. If fewer than $(M + 1) / 2$ votes are cast for the replacement of the manager, then the old manager stays in charge.

Initial allocation of shares. At the beginning, each investor can simultaneously submit a bid for an A-share and another bid for a B-share. No investor is allowed to bid for multiple shares of the same class.⁹ It is impossible not to bid, but it is possible to bid zero. The auctioneer collects the N bids for each class of shares. The shares for each class are then allocated to the investors who submitted the M highest bids. Investors pay a price equal to the $M+1^{st}$ bid submitted for this class of shares. If several investors submit identical bids so that the M^{th} bid is not unique, then the auctioneer allocates the shares by lot among these investors.

Sequence of events. We obtain the following extensive form of the game:

1. Nature draws the quality of the manager who runs the firm. Investors bid for shares in the firm.
2. The auctioneer allocates the shares of the firm to investors and sets prices. Investors pay the prices for the shares they receive.
3. The dividend for the first period is realized, becomes observable to all investors, and is paid to all shareholders.

⁹ We impose this restriction for several reasons. At the bidding stage, this assumption greatly simplifies the potential bids of the investors, and (together with a sufficiently high budget) makes sure that investors always have enough resources to bid for all the shares they want. At the voting stage, this restriction makes sure that there are always M voters, so that it is not possible to buy a majority of votes.

4. A-shareholders vote on the replacement of the manager.
5. The dividend for the second period is realized, becomes observable to all investors, and is paid to all shareholders.

The structure of the game is as simple as possible. We need a minimum of two periods, so that individuals can make inferences from the first dividend regarding the quality of the manager, and then make appropriate decisions contingent on these inferences. We use two classes of otherwise identical shares to empirically measure individuals' valuations of a vote. In particular, our setup is robust to risk-aversion, because risk-aversion affects the prices of the two types of shares, but has a negligible effect on the price difference, which is the quantity we are most interested in.¹⁰

We frame the game as a shareholder vote to rule out many reasons why individuals might value voting in practice. Our aim is not to model the microstructure of trading in dual-class shares. In real-world markets, individuals with superior abilities to price shares may act as arbitrageurs and take advantage of individuals who overvalue voting shares, and this may lead to different pricing results. Moreover, real voting shares give a voting *right* to their owners, which can, but need not be exercised.

In our experiment, there are no information asymmetries, no conflicts of interest, and no private benefits from voting. Hence, the incentives of all subjects are fully aligned and there is no scope for interest groups. Our setup also rules out the “voter’s illusion” (see Quattrone and Tversky (1984)), where the voter believes that if he incurs the costs of voting, others will do so as well. In our setting, the number of voters (or the “turnout”) is always constant and the question is not *how many* individuals vote, but *who* votes. Finally, our setup precludes notions of civic duty or social status that may otherwise affect voting behavior.

2.2 Nash equilibrium: Voting without errors

In this section we derive theoretical predictions for the voting premium implied by the instrumental voting hypothesis in a context where individuals have correct beliefs about the

¹⁰ For constant absolute risk aversion, there is no wealth effect and the price difference is independent of risk-aversion. However for more general models of risk-aversion, different prices lead to different wealth and therefore to a wealth effect. We argue that this wealth effect is negligible, because we consider the difference between two very similar payouts.

voting behavior of others and respond optimally and without errors. Thus, we calculate the Nash equilibrium. Throughout our theoretical analysis, we assume that individuals are risk-neutral.

Voting. In the first period, and also after a replacement of the manager in the second period, the quality of the manager is good or bad with equal probability. If the manager is good, the expected dividend is $pH + (1-p)L$; if she is bad, it is $(1-p)H + pL$. Denote by $\delta \in \{R, K\}$ the decision to replace (R) or to keep (K) the manager.

Table 1 shows the expected dividends conditional on the first period dividend D_1 and on the decision to replace δ . It follows from Table 1 that if $p > 1/2$, then A-shareholders wish to replace the manager if $D_1 = L$ and they wish to keep the manager if $D_1 = H$. If $p = 1/2$, they are indifferent between replacing and keeping the manager. In principle there are many asymmetric equilibria of this game where no shareholder is pivotal for the decision.¹¹ However, in the unique symmetric equilibrium in weakly undominated strategies, all A-shareholders vote for replacing the manager if $D_1 = L$ and for keeping her if $D_1 = H$ as long as $p > 1/2$.¹² We focus on this equilibrium here and ignore the asymmetric equilibria of the voting game. We also ignore non-responsive symmetric equilibria where all shareholders always vote for the same alternative. These equilibria are not in weakly undominated strategies.

Bidding and valuation. Using the expressions in Table 1 and the equilibrium replacement decision, we obtain the value of the shares in the Nash equilibrium as the sum of the expected dividends from the two periods:

$$V_{NE} = H + L + \left(\frac{1}{4} - p(1-p) \right) (H - L). \quad (1)$$

Whenever $p > 1/2$, we have $V_{NE} > H + L$. In contrast, if the manager could not be replaced by voting, then $V_{NE} = H + L$. The second term in equation (1) therefore reflects the benefits from voting to shareholders, which are zero if $p = 1/2$. The intrinsic value V_{NE} is the same for A-shares and for B-shares. The auction is a standard multi-unit Vickrey auction, so it is a weakly dominant

¹¹ For example, in one equilibrium $(M + 1)/2 + 1$ A-shareholders always vote to replace the manager and $(M - 1)/2 - 1$ A-shareholders always vote to keep her.

¹² In order to see this, note that in any completely mixed strategy profile there is a positive probability for each A-shareholder to be pivotal. But then the expected second period dividend is strictly higher when the manager is replaced compared to when she is kept (given that $D_1 = L$ and $p > 1/2$) and voting to replace her is then a strictly dominant strategy for each shareholder.

strategy for each risk-neutral investor to bid the intrinsic value V_{NE} for an A-share and also for a B-share. If all investors do this, then the unique price for the A-shares is $P_A = V_{NE}$, and, similarly for the B-shares $P_B = V_{NE}$. There is therefore no price difference between the two classes of shares. Since all investors bid the same amount, the auctioneer allocates the shares by lot.

3 Experimental setup

All treatments have $N = 8$ investors. There are $M = 5$ A-shares and 5 B-shares, and the dividends are $H = 20$ and $L = 0$. We analyze six different treatments. Each treatment was repeated with ten different groups, which were assigned to two subtreatments: In five groups per treatment, we only indicated the outcome of the vote (keep/retain the manager) after each round whereas in the other five groups, we also informed subjects about the exact distribution (in favor/against the manager) of the five votes. The division into these subtreatments was applied to all six treatments.

BASE: The baseline treatment has $p = 80\%$. With this parameter value, the difference between the good manager (generating an expected dividend of $E(D) = 16$) and the bad manager (with $E(D) = 4$) is substantial. We attach the translation of the experimental instructions for the baseline treatment at the end of this document.

NU: In the no-uncertainty treatment $p = 100\%$. Then $E(D) = 20$ with a high quality manager and $E(D) = 0$ with a low quality manager, so that the stakes have increased relative to the base treatment. In addition, the inference problem in this treatment is much simpler than in BASE, which should reduce the number of shareholders who erroneously vote against their interest. This treatment therefore serves as a robustness check on the baseline treatment.

NI: In the non-influential manager treatment $p = 50\%$. Then $E(D) = 10$ independently of the quality of the manager, and voting is inconsequential for the value of the firm. The inference problem is trivial. While A-shareholders have an impact on whether the manager is kept or replaced, they cannot influence the distribution of future dividends with their vote. We use this treatment to check that non-instrumental reasons do not motivate bidding behavior.

NP: In the non-pivotal treatment $p = 80\%$, just as in the baseline treatment. However, in addition to the 5 A-shares auctioned off at the beginning, there is a blockholder who owns 6 A-shares and who always votes in favor of the incumbent manager. As a result, A-shareholders are never pivotal and can never replace the manager, no matter how they vote. This treatment is similar to

the non-influential manager (NI) treatment with respect to the instrumental value of the vote, which is zero. However, in the non-pivotal treatment there are good and bad managers and the vote is meaningful in principle, hence, this treatment provides a different check that non-instrumental reasons do not motivate bidding behavior.

The non-influential manager treatment and the non-pivotal treatment also allow us to address a potential criticism of the experiment, which holds that subjects pay for participating in the vote for the mere entertainment value of voting: They might value voting simply to avoid looking at a blank screen while the other participants in the experiment vote. If this is true, then participants should also be willing to pay a premium in treatment NI and treatment NP, where their vote has no impact on the outcome. We show below that there is no reason for this concern in our experiment.

FEED: In the feedback treatment $p = 80\%$ as in the baseline treatment. FEED deviates from BASE only by providing subjects with additional feedback about the previous rounds in an aggregated and easy-to-read manner. Beginning with the sixth round of play, subjects see three pie charts at the end of each round. The first chart indicates the proportion of votes in favor of the manager after a low dividend, the second chart shows the proportion of votes against the manager after a high dividend, and the third chart displays the proportion of close voting outcomes, i.e., outcomes of three votes against two. This information helps subjects to get a good estimate of the voting behavior of others.

BE: In the belief-elicitation treatment $p=80\%$ as in BASE, and deviates from BASE only by asking subjects to report their beliefs about their pivot probability in each round. These beliefs are elicited after the subjects placed their bids for the A-shares and B-shares, but before they learn the outcome of the auction. In the instructions, we explain the concept of pivotality as being decisive for the outcome of the vote. Moreover, subjects are informed that they become decisive if and only if exactly two of the other four A-shareholders vote for keeping the manager while the other two vote for replacing him. Each subject has to indicate the probability that his vote will be decisive for the outcome if he wins an A-share.

Subjects are paid for their statements about their beliefs according to the quadratic scoring rule.¹³ A subject who has indicated an estimate p of his probability of being pivotal receives a payoff $4(1-(p-d)^2)$, where the dummy d is one when he is pivotal and zero otherwise. In the experiment, subjects enter p in percentage points as a number between 0 and 100. This treatment allows us to measure by how much subjects overestimate their pivot probability, how they are influenced by having to think about pivotality, and to which extent they respond optimally to their stated beliefs.

Table 2 provides an overview of the six different treatments and gives the parameter values for each treatment. The table also shows the equilibrium value of the shares from equation (1). Shares are most valuable in treatment NU with $V_{NE} = 25$ and least valuable in treatments NI and NP with $V_{NE} = 20$. The table also shows the benefit ΔU from making the correct decision relative to making an incorrect decision about replacing the manager (see Table 1):

$$\Delta U = 2 \left(\frac{1}{4} - p(1-p) \right) (H - L) \geq 0 \quad (2)$$

We ran independent sessions with ten different groups per treatment, and each session lasted for about 90 minutes. There were eight subjects per group, and they stayed together for the entire experiment. No subject participated in more than one session. The participants played the same game (treatment) for 30 rounds. Subjects' average earnings were 28 Euros, including a show-up fee of three Euros. The currency in the experiment was "points," which were exchanged into real money at the end of the experiment at the exchange rate of 1 point = 3 cents. The exchange rate was known to all participants. At the beginning of the experiment, each participant was given a budget of 150 points. In order to prevent investors from running out of money after a series of zero dividends and to avoid frustration from receiving a payoff of zero in several rounds, each individual received an additional 16 points at the end of each round.

All experiments were carried out in the computer lab at Technical University Berlin between January 2006 and October 2012. We used the software tool kit z-Tree to program the experiments (Fischbacher (2007)). Participants were students from different universities in Berlin. After the experiment, students were asked to fill in a questionnaire, and the vast majority

¹³ The payment rule also takes into account subjects who do not win an A-share and who therefore do not participate in the vote. They are assigned the role of the subject with the lowest bid among those who actually obtain an A-share.

of them stated that they did not have any experience with investing in the stock market. Therefore, we are confident that the majority of students were not aware of the price difference between voting and non-voting shares in listed companies.

4 Do subjects pay for a vote?

In this section, we present our experimental results for the baseline treatment and test whether the Nash equilibrium derived in Section 2.2 provides a good description of our findings. That equilibrium implies that subjects bid identical amounts for voting and non-voting shares. In what follows, we present our analysis for the data pooled across the two subtreatments with, respectively, without detailed feedback about the voting outcome, since we did not find any significant differences between the two subtreatments in any of the six treatments.

Prices. Table 3 shows statistics for the prices of voting and non-voting shares (P_A and P_B) as well as for the voting premium $P_A - P_B$ and the relative voting premium $(P_A - P_B) / V_{NE}$. We scale prices because the intrinsic values of the shares vary across treatments. The prices P_A and P_B are very volatile and therefore unsuitable for scaling, so we scale by the theoretical value of the shares in equation (1), which is 21.80 for the baseline treatment. Within each group of eight investors, the subjects observe other subjects' behavior indirectly via the outcome of the voting game and the prices of the two types of shares. Therefore, the observations of prices and bids within a session cannot be treated as independent observations. We use average prices for each group and present the results on these 10 independent observations. The same remark applies to subsequent tables. Panel A shows the results for the last 15 rounds and our discussion focuses on these results to avoid the influence of learning that takes place over the first ten rounds (see below). For completeness, we also report the same results for all 300 observations (10 groups with 30 rounds each) in Panel B, where we treat every round as if it were an independent observation.

The results in Table 3 strongly reject the hypothesis that the voting premium is zero. Individuals do value participating in the vote. The average (median) price for an A-share is 18.52 (18.27), and the average (median) price of a B-share is 15.61 (17.37). Hence, the voting premium is 2.91 on average, with a median of 2.63. If we look at the entire sample of all 300 rounds, then the voting premium is only slightly different with a mean of 2.82 and a median of 2.00. These

voting premiums are different from zero at all conventional significance levels, and, at about 13.4% of the expected value of the dividends, they are also economically significant.

Bidding behavior. Table 4 shows descriptive statistics on average bids for 10 independent observations across the last 15 rounds in Panel A. Panel B displays the results for 2,400 individual bids (10 groups with 8 subjects and 30 rounds). Our analysis of bids arrives at the same qualitative conclusions as the analysis of prices, even though bids are much more volatile than prices, with standard deviations that are six to eight times higher.

Bids vary between 0 and 300, the lower and upper bound implemented in the experiment. However, extremely high bids are very rare and only 14 of the 1,200 bids we analyze for the baseline treatment equal 300. Given that the maximum payout per share is 40 if both dividends turn out to be high, any bid above 40 appears irrational at first glance. However, the five winners in the auction do not pay their own bids but the sixth highest bid, so bidding more than the intrinsic value of the share, or more than 40, is merely a weakly dominated strategy.¹⁴ We therefore only interpret the median of the bids in order to give less weight to these outliers.

Median bids are higher by construction than average or median prices, because prices equal the sixth highest bid, i.e. the 25th percentile of all bids. Table 4, Panel B reveals that 53.8% of all 2,400 pairs of bids feature a positive voting premium and only in 10.8% of all cases do subjects bid less for an A-share than for a B-share. The remaining bids (35.4%) have the same price for both shares. Again, voting premiums are positive and statistically significant at all conventional significance levels and we conclude that our experimental subjects value participating in the vote.

Learning. We first check to what extent our experimental subjects learn to play the game during the experiment. Figure 1 shows the prices of voting and non-voting shares over time, and Figure 2 displays the frequency of incorrect votes, i.e. the frequency that A-shareholders vote for replacing the manager after a high dividend or for keeping the manager after a low dividend. Both figures refer only to the BASE treatment.¹⁵ Both figures indicate that subjects display some learning in the course of the experiment: prices increase over time and voting errors decrease over time. Figure 1 shows that learning with respect to prices takes place only during the first 10

¹⁴ The highest price ever paid in the baseline treatment in any of the 30 rounds of the game is 30, the highest price ever paid in any treatment is 35.

¹⁵ Similar observations apply to all other treatments, but are not reported here to conserve space.

periods; after period 10, average prices do not vary by much anymore. Figure 2 shows that there is little learning with respect to voting decisions. After observing a high dividend, the individual error rate falls from an average of about 18% during the first 10 periods to about 12% during the last 10 periods. After a low dividend there is hardly any change at all.¹⁶ Note that errors in actual voting outcomes in groups of five are much lower at around 3%, since three individuals must make errors in order to change the voting outcome.

Underbidding. Average prices are substantially below the risk-neutral value of 21.80 in the baseline treatment. The discount is 15% for A-shares and 28% for B-shares, and the differences to the equilibrium value V_{NE} are statistically highly significant for both shares. We hypothesize that these differences can be attributed to risk aversion. We regard each period as a separate gamble with three potential outcomes: both dividends are high (HH=40), both dividends are low (LL=0), and one is high and one is low (HL=LH=20). In Panel B of Table 1 we display formulae for the likelihood of each outcome. We then infer the coefficient of absolute risk aversion ρ from the following condition:

$$U(\text{Price}_A) = pr(40)U(40) + pr(20)U(20) + pr(0)U(0), \quad (3)$$

where $U(x) = -\exp(-\rho x)$ and $pr(x)$ is the probability for each outcome, which is given in Panel B of Table 2 for each treatment. A coefficient of absolute risk aversion of 0.028 solves this equation for the baseline treatment. In many experimental papers, constant relative risk aversion is assumed. One possibility in our context is to use the average profits across individuals in each round of the experiment, which is 26.76. With this calibration we obtain a coefficient of relative risk aversion of 0.75 (0.028 x 26.76) which is comparable in size to levels found in the experimental literature (see Goeree, Holt, and Palfrey 2002; Holt and Laury 2002).¹⁷

5 Can non-instrumental voting account for the premium?

We conclude from the discussion of the baseline treatment that subjects value participating in the vote, despite symmetric information, perfectly aligned interests, and a very low rate of inefficient voting outcomes. We hypothesize that our results are best explained within the framework of

¹⁶ The difference for individual errors τ between periods 1 to 15 and periods 16 to 30 is statistically highly significant for votes after a high dividend (χ^2 -test with p-value 0.8%). All other differences between the first 15 periods and the last 15 periods in the figures are not statistically significant.

¹⁷ For treatment NU we would obtain a value of 0.028 for the coefficient of absolute risk aversion and 0.81 for the coefficient of relative risk aversion, using the same calculations.

instrumental voting, because our setting does not provide much scope for expressive voting. Still, we have to consider non-instrumental voting as a potential explanation for our results. Our experimental subjects may attach a direct value to voting, and we consider several possible non-instrumental explanations in this section.

First, we consider *expressive voting*, which posits that voters enjoy voting because it gives them the opportunity to “make their voice heard.” The outcome and the ultimate effect of one’s own vote on the outcome are secondary in this case. Brennan and Hamlin (1998) epitomize the expressive motivation to vote as “voter participation just *is* the act of consumption that brings the voter to the poll” (p.156).¹⁸

With expressive voting we should observe that investors value voting independently of how much influence it has on their payoffs, but conditional on there being a meaningful issue on which opinions can be built and expressed. In particular, investors should be willing to pay a premium for voting shares in the non-pivotal treatment NP, but not in the non-influential manager treatment NI. In both treatments, the vote of the A-shareholders does not affect the expected second period dividend. However, in treatment NI voting has no influence because the quality of the manager does not affect dividend payments, so that there is no meaningful issue on which opinions could be expressed. By contrast, in treatment NP the vote has no influence because the manager is never replaced, but since management quality affects payoffs, the issue of the vote is meaningful.

The second potential non-instrumental explanation for our results is that subjects value the entertainment value of voting, because they prefer to engage in some activity rather than watching a blank computer screen while other subjects vote. Subjects may also value voting shares because they are labeled ‘A,’ which they may consider a label of superiority.¹⁹ If one or both of these explanations are true, subjects should also pay a premium in treatments NP and NI.

Table 5 shows the results for the non-influential manager treatment (NI) and Table 6 shows those for the non-pivotal treatment (NP). To test for differences between prices and premiums across treatments here as well as in the remainder of the paper, we report the p-values

¹⁸ Here, we refer only to the general definition of expressive voting by Brennan and Hamlin (1998) and Brennan and Lomasky (1993). Their so-called low-cost theory of expressive voting is not applicable to our setup. According to this theory, voters tend to vote more expressively the more they identify with the issue that is voted on and the less they think their vote matters,

¹⁹ All our experimental subjects were undergraduate students at a German university and were therefore not conditioned on US-style letter grades. This explanation is therefore less plausible in our context.

of three different tests: (1) a two-sample t-test; (2) pooled panel regressions, in which standard errors are clustered at the group level to account for within-group correlations;²⁰ (3) Mann-Whitney U tests (MWU tests). The MWU test provides information about the relative location of two distributions and should not be interpreted as a test on the difference of means or medians. From both tables we can see that the voting premium becomes very small and statistically indistinguishable from zero. The voting premium is as low in treatment NP as it is in treatment NI, and it is significantly smaller in either treatment than in the baseline treatment.

Another non-instrumental explanation for a voting premium may be that individuals enjoy voting for the winner (see Bartels (1988), Callander (2007), Fiorina (1974), and Schuessler (2000)). If this is the case, then we should see that individuals do vote for the winner and that they value votes more if they have a better chance of voting for the winner. Treatment NP should arguably be very attractive for individuals who want to vote for the winner, because the winner is known for sure. However, we do not observe a voting premium in this treatment and Table 4 shows that individuals do not vote for the winner. Therefore, we also do not find any support for this variant of non-instrumental voting. Therefore, the results from Table 5 and Table 6 do not lend any support to the non-instrumental voting hypotheses discussed above.

We conclude that individuals are not willing to pay for participating in a vote when the vote of all participants does not have any material consequences. The voting premium vanishes whenever the instrumental value of the vote is zero, either because the issue decided by the vote has no impact on the payoffs, or because the vote can never become decisive. This observation lends strong support to instrumental voting.

6 Voting with errors and overconfidence

The central result of the last two sections is that neither the Nash equilibrium with instrumental voting nor non-instrumental voting can explain the premium individuals pay on voting shares in our experiment. From now on, we focus on explanations based on instrumental voting.

One candidate explanation for the voting premium is that some individuals have higher error rates than others in the voting game. Anticipating this, the more competent voters, i.e., those with lower error rates, might want to exclude the less competent voters from the voting game by

²⁰ The p-values are computed by pooling the observations of two treatments and then regressing the variable of interest (price, premium, etc.) on a constant and a dummy variable that equals one for one of the treatments and zero otherwise. The p-value we report is for the test that the coefficient on the dummy variable equals zero.

bidding aggressively for the voting shares. To investigate this possibility, we take a closer look at actual voting behavior.

6.1 Voting in the baseline treatment

Let τ_{D_i} denote the probability that an A-shareholder deviates from the equilibrium prediction when voting on the dismissal of the manager so τ_L is the probability that a shareholder votes against dismissal after observing a low dividend, and τ_H is the probability that a shareholder votes for dismissal after observing a high dividend. Similarly, let e_L denote the probability that the A-shareholders vote to keep the manager after a low dividend and let e_H denote the probability that they fire the manager after a high dividend. Assuming that individual shareholders' errors are independent of each other, incorrect decisions from the vote are given by the binomial probability

$$e_{D_i} = 1 - B\left(\frac{M-1}{2}, M, \tau_{D_i}\right), \quad (4)$$

where $B\left(\frac{M-1}{2}, M, \tau_{D_i}\right)$ is the probability that at most $(M-1)/2$ out of M shareholders, i.e., less than half, make a mistake and the individual probability of making a mistake is τ_{D_i} .

Table 7 reports these errors for all treatments. The table shows that in the baseline treatment, A-shareholders vote for keeping the manager after observing a low dividend 18.3% of the time, averaged over the last 15 rounds. The frequency e_L that after a low dividend more than two A-shareholders vote against replacing the manager is 7.1%. After a high dividend, 11.8% of A-shareholders vote incorrectly, that is, for replacing the manager, which results in an actual dismissal in 3.0% of all cases.²¹

²¹ Our error rates of 18.3% and 11.8% are somewhat higher than those found by Guarnaschelli, McKelvey, and Palfrey (2000) who observe error rates in the range of 5% to 12%. In their experiment, every voter has a piece of privileged information and it is therefore important for the final outcome that everyone votes correctly. In our setting, there is no information asymmetry and the outcome is not affected by occasional errors. Our experimental setup therefore invites more inattentiveness than the setup in Guarnaschelli, McKelvey, and Palfrey (2000). In treatments NI and NP, voting errors are higher than in BASE. We conjecture subjects understand their votes are of no consequence to their payoffs in these treatments, and some of them therefore vote randomly. However both in NI and NP, subjects keep a manager more often after a low dividend than replace him after a high dividend. This tendency is much stronger in NI than in NP. This asymmetry in errors might reflect a moral preference for keeping the manager, even though the manager is computerized. In NP, it might also reflect a bandwagon effect, i.e., "voting for the winner." Subjects learn nothing about the quality of the manager from their first dividends in NI, whereas they do in NP. This difference may account for the larger errors in NI.

Further analyses, which are not reported in the tables, reveal that errors in voting are not due to specific individuals who persistently make such errors. We rather observe that most individuals experiment once in a while or occasionally become inattentive when they vote. As three correct votes are sufficient to reach the equilibrium outcome, these errors remain inconsequential in most cases. We also checked whether the voting premium increases after an erroneous decision in the previous period. We do not find any significant relation here, possibly because erroneous decisions are rare. Likewise, we find no relation between the number of wrong votes and the subsequent voting premium in the subtreatment with full information about the voting outcome. We also searched for a relationship between voting errors and the premium in bids, but did not find anything, presumably because of the high variation in bids. In sum, we do not find any statistical evidence for the supposition that actual error rates drive the bidding behavior that leads to the voting premium.

6.2 *The no-uncertainty treatment*

We complement the statistical reasoning in the previous section by an experimental test using the no-uncertainty treatment. In this treatment, a good manager always generates a high dividend and a bad manager always generates a low dividend. Therefore, the inference problem at the voting stage is greatly reduced, and we expect fewer errors at the voting stage. If the anticipation of voting errors drives the voting premium, then these errors should vanish or at least decline in NU.

Table 7 shows that there are indeed fewer errors in treatment NU than in BASE. After a low dividend, only 11.0% of the subjects vote for keeping the manager in the no-uncertainty treatment while 18.3% do so in the baseline treatment. After a high dividend, error rates are also lower in NU than in BASE.

Table 8 shows descriptive statistics for prices (Panel A) and bids (Panel B) averaged across the last 15 rounds for the no-uncertainty treatment. Given that the theoretical value V_{NE} of the shares is higher in the no-uncertainty treatment than in the baseline treatment (see Table 2), it is not surprising that prices and bids are also generally higher. The mean (median) relative price premium of voting shares in NU is 16.0% (12.5%) in Table 8, Panel A, which is higher than in the baseline treatment, where it is only 13.4% (12.1%, see Table 3, Panel A). Mean and median bids are slightly lower in NU than in BASE (compare Table 8, Panel B with Table 4, Panel A).

However, the differences between prices and voting premiums in BASE compared to NU are never significant.

Overall, we find that error rates in the voting game in treatment NU are slightly lower than in BASE. However, the voting premium is equal or even higher in NU than in BASE. We conclude that actual error rates cannot explain the voting premium.

6.3 The voting premium in Quantal Response Equilibrium (QRE)

Our next step extends the theoretical analysis in Section 3 and incorporates the possibility that experimental subjects make mistakes at the voting stage. McKelvey and Palfrey (1995, 1998) develop the concept of Quantal Response Equilibrium (QRE), which incorporates this aspect and which has been successfully applied to experimental data.²² A QRE is based on the assumption that players' strategies are best responses to the mistakes they anticipate other players to make. In the first step, we maintain the assumption that individuals correctly assess their own mistakes as well as those of other players.

We first derive the value of one share when individuals make errors. Using the probabilities e_H and e_L that the vote results in a wrong decision (see equation (4)), the manager is kept after a high dividend with probability $(1 - e_H)/2$ and after a low dividend with probability $e_L/2$. With probability $(1 - e_L)/2$ the manager is replaced after a low dividend, and with probability $e_H/2$ the manager is replaced after a high dividend. For all cases the payoffs are given in Table 1. The value of one share in a QRE is therefore:

$$V_{QRE} = H + L + \left(\frac{1}{4} - p(1-p) \right) (H - L) (1 - e_L - e_H). \quad (5)$$

For the equilibrium without errors, we have $e_L = e_H = 0$, and the value of the shares in QRE becomes equal to V_{NE} as in (1). If we plug the numbers from Table 7 into equations (4) and (5), we obtain $V_{QRE} = 21.62$, compared to $V_{NE} = 21.80$ for the baseline treatment. Similarly, in treatment NU $V_{QRE} = 24.93$, is the value of the shares with errors compared to $V_{NE} = 25.00$. Hence, the reduction in the theoretical values of the shares caused by the observed voting errors is small for both treatments BASE and NU.

Quantal response equilibrium predicts that subjects make fewer errors when the utility difference between a correct and an incorrect vote is larger. As a consequence, subjects who own

²² See, for example, Guarnaschelli, McKelvey, and Palfrey (2000) for an application of QRE to a voting game.

an A-share and a B-share should make fewer errors than subjects who merely own an A-share. We find the opposite: A-shareholders who do not own a B-share make errors with probability 12.4% while the probability is 16.4% for A-shareholders who also own a B-share, but the difference is not statistically significant. The probability of making a mistake is therefore symmetric across players and common knowledge by assumption. It follows that equation (5) applies to both classes of shares and $P_A = P_B = V_{QRE}$, so there is no price difference between voting A-shares and non-voting B-shares. This argument holds for any model with symmetric errors that are common knowledge. QRE alone can therefore not explain the voting premium in BASE and NU.

6.4 Can overconfidence explain the voting premium?

We depart from QRE by going one step further and assume that investors are overconfident in the following sense. At the voting stage they make errors as described above, and the probability of making a mistake is the same for all investors. However, the symmetry of errors is no longer common knowledge. Instead, each investor believes that only the other shareholders will make random errors at the voting stage, whereas he himself votes correctly without fail. We assume this extreme form of overconfidence as it will provide us with an upper bound on the voting premium. Overconfident investors value being pivotal at the voting stage, because, according to their beliefs it helps them to avoid the errors other investors would make.

Our notion of overconfidence involves that subjects have correct beliefs about the error rates of others and only underestimate their own error rates. Behavior at the voting stage is therefore not affected by overconfidence, and the actual error probabilities are still given by τ_L and τ_H . Only bidding strategies are affected by overconfidence. Furthermore, no false beliefs about *strategies* are involved here: Each individual correctly anticipates all bidding and voting strategies in equilibrium. Hence, the model we analyse in this subsection is still an equilibrium model. It shares the assumption that agents might be naïve about the behavior of their own “future selves” with a number of equilibrium models in the behavioral economics literature.²³

With overconfident investors, the valuation formulae for expected dividends and therefore the intrinsic valuation of a share in equation (5) do not change. However, overconfident investors

²³ In all models with agents who exhibit a behavioral bias of some kind on a later stage of the game, one can model (partial) naïveté of the agent about her own bias on previous stages. See, e.g., Bénabou (2013), p.436.

anticipate that they can improve on this valuation if they own a voting share and are pivotal in the voting game, i.e., if the remaining four votes lead to a tie of 2:2. We denote the probability of any A-shareholder to be pivotal conditional on the dividend payment D_I by π_{D_I} . Table 7 presents these probabilities for the individual treatments. An overconfident investor believes that whenever he is pivotal, he can change the probability of replacing the manager from $1 - \tau_L$ to 1 if $D_I = L$ and from τ_H to 0 if $D_I = H$. Hence, an overconfident investor overvalues each share conditional on being able to vote by ω :

$$\omega = \frac{1}{2}(\pi_H \tau_H + \pi_L \tau_L) \Delta U, \quad (6)$$

where ΔU is given in equation (2). Thus, ω denotes the increase in intrinsic value of a share from participating in the vote from the point of view of an overconfident investor. This will generally not be equivalent to the voting premium in equilibrium, because an investor's valuation of the B-shares depends on his ability to obtain an A-share in the auction.²⁴ An overconfident investor also values the non-voting share above its intrinsic value (5) conditional on obtaining an A-share, because, whenever he owns a voting share, this also increases his valuation of the non-voting share. The interdependence makes the analysis of the equilibria of the auction game somewhat tedious, and a complete characterization of these equilibria is beyond the scope of this paper.²⁵ In the following proposition, which we prove in the Appendix, we only summarize the theoretical results that are relevant for our empirical analysis in providing us with an upper bound of the theoretical voting premium in the QRE model with overconfidence.

Proposition 1: *Assume that each investor believes that the value of a share of the company increases by ω if he owns one A-share. Then: (i) There exists no symmetric equilibrium in pure strategies. (ii) There exists a continuum of asymmetric equilibria in weakly undominated pure strategies where the voting premium in prices and bids satisfies $0 \leq P_A - P_B \leq 2\omega$ and $0 \leq B_A - B_B \leq 2\omega$, respectively.*

²⁴ For interdependent valuations, auction theory recommends a combinatorial auction, where participants can bid on each possible bundle of the two shares. In such a setup truthful bidding remains the optimal strategy even in the presence of complementarities or substitution effects (Varian and MacKie-Mason (1995); for a survey see de Vries and Vohra (2003)). We decided against such a design because of its complexity and because a combinatorial auction could have an unwanted framing effect. When making a bid on the bundle of an A-share and a B-share, subjects could have been led to believe that A-shares and B-shares are complementary, i.e. the framing could have induced incorrect beliefs.

²⁵ A complete analysis of all pure-strategy equilibria of the auction game with overconfident players is available from the authors upon request.

Intuitively, Proposition 1 can be understood as follows. Investors are willing to pay $2(V_{QRE} + \omega)$ for both shares together. Prices below V_{QRE} cannot be sustained in equilibrium for any B-share, and, similarly, prices below $V_{QRE} + \omega$ cannot be sustained for any A-share. The reason is that those investors who bid low and receive no shares in equilibrium would otherwise be better off by bidding high for one type of share. However, the additional value ω of the B-share can be allocated in any possible way between the A-share and the B-share. The highest premium of 2ω results when the price of the A-share reflects the whole increase in value of both shares while the price of the B-share is just equal to V_{QRE} . Conversely, both shares may be priced at $V_{QRE} + \omega$ in equilibrium, so that the voting premium is zero. Any intermediate case is also possible, which gives rise to a continuum of equilibria.

Panel A of Table 9 displays the upper bound 2ω from Proposition 1. We obtain a maximum premium of $2\omega = 0.14$ for the BASE treatment which is less than 5% of the average observed premium of 2.91. The observed premium is significantly higher than this upper bound at the 1% level. The disparity between predicted and observed values is even larger for treatment NU. We conclude that overconfidence in the sense of underestimating one's own error rate alone cannot explain a meaningful portion of the observed voting premium.

7 Overestimating the errors of others: Level-k reasoning

The previous section shows that assuming that individuals make errors at the voting stage and that they are overconfident about their own error rates is insufficient to explain the size of the observed voting premium in our experiment. A potential reason is that this model is still too restrictive by assuming that subjects have correct beliefs about the errors of others. We therefore introduce the possibility that subjects hold false beliefs and overestimate the likelihood that other subjects make mistakes. We refer to the equilibrium of this model as equilibrium with overestimation of the errors of others. Note that we retain the assumptions that individuals make mistakes and that they are overconfident.

We capture the notion that subjects hold incorrect beliefs about the strategies of others by employing the level-k model. According to the formulation of Stahl and Wilson (1995), there are a number of possible player types. Level-0 types randomize uniformly over all possible actions. Level-1 types believe that all other players are level 0 and choose a best-response based on this

belief. All other types are of some higher level k and play a best response based on the belief that the other players are level $k-1$ or lower and that all lower-level types occur with positive probability.

It is not useful to apply the level- k model to the bidding stage of our game because level-0 players would be assumed to randomize over all prices between zero and 300 (upper bound of feasible bids) and the model could then generate a large range of predictions, depending on the assumed proportion of level-0 players and their randomization strategies. We therefore apply the level- k model to the voting stage only and assume that bids can be interpreted as best responses to voting behavior. This allows us to derive upper bounds on the voting premium based on Proposition 1 as before.

We proceed by deriving an upper bound for the voting premium in BASE. We assume that level-0 types randomize at the voting stage and higher-level players play best responses to their beliefs about the voting behavior of lower-level players. Level-1 types therefore always vote correctly to keep the manager after a high dividend and to replace her after a low dividend. As level-1 players believe that all other players randomize, they believe that others' error rates are $\tau_H = \tau_L = 0.5$ and that they themselves can prevent an error in 50% of all cases in which they are pivotal. The probability of being pivotal implied by their beliefs about the errors of other players is $\pi_H = \pi_L = 3/8$.²⁶ Accordingly, level-1 players believe that each share of the firm is worth more by $\omega = 3/8 \times 1/2 \times \Delta U$ when they own a voting share. The reasoning of level-2 types is similar to that of level-1 types and they also assume that they will always vote correctly. However, the value of a voting share is lower for level-2 types than for level-1 types. To see this, note that level-2 types believe that a certain proportion of players is level 1 and therefore votes correctly. The perceived probability of mistakes by other players is lower for level-2 types than it is for level-1 types. The same argument applies to all players with levels higher than 2. We therefore obtain the highest premium in the level- k model under the assumption that all players are level 1. From equation (6) and Proposition 1, the maximum voting premium is then $2\omega = 2 \times (3/8 \times 1/2 \times \Delta U) = 1.35$ for BASE, about ten times larger than the upper bound implied by

²⁶ This theoretical probability of being pivotal conditional on D_1 , π_{D_1} , is based on the assumption that shareholders' errors are independent and for $\tau = 0.5$ is calculated as

$$\pi_{D_1} = \binom{M-1}{(M-1)/2} \tau_{D_1}^{(M-1)/2} (1-\tau_{D_1})^{(M-1)/2} = \binom{4}{2} \tau_{D_1}^2 (1-\tau_{D_1})^2 = \frac{3}{8}.$$

overconfidence alone. For NU we have $2\omega = 3.75$ from the same calculation, about 30 times larger than the upper bound with overconfidence only (see Panel A of Table 9).

Panel B of Table 9 reports tests for this upper bound of the level-k model for treatments BASE and NU. The upper bound of 1.35 in treatment BASE is about half of the observed median premium of 2.63 and statistically significantly smaller: the Wilcoxon signed-rank test rejects the hypothesis that the upper bound holds with a p-value of 4.6%. However, for NU the upper bound of 3.75 implied by the level-k model is statistically indistinguishable from the observed premium, which has a mean of 3.99 and a median of 3.13. We therefore conclude that the level-k model can provide a satisfactory explanation for the evidence in treatment NU and a partial explanation for the evidence in the baseline treatment.

8 Feedback and beliefs

The question remains why individuals are willing to pay an even higher voting premium in BASE than is predicted by the level-k model. From the analysis in Section 7 we can rationalize the observed voting premium of 2.91 in BASE only if $2\omega \geq 2.91$, which would require a pivot probability of at least 81%.²⁷ The level-k model implies that subjects do not assume pivot probabilities above 37.5%. We therefore hypothesize that at least some experimental subjects in BASE overestimate their pivot probability by more than what the level-k model implies for level-1 players. Put differently, some subjects may hold implicit beliefs about others' error rates that are more extreme than could be explained by the level-k model. Such extreme beliefs may result if individuals receive little feedback about the game. We therefore analyze the belief formation of individuals by investigating the impact of additional feedback and by eliciting the beliefs about pivotality.

8.1 The impact of feedback

In treatment FEED we provide individuals with more easy-to-read feedback about the voting behavior of other players and the share of close votes in past play.²⁸ Subjects could directly estimate the average pivot probabilities and error rates of past rounds from this information. Table 7 shows the voting behavior in the feedback treatment, Table 9 provides test results for the

²⁷ If we set $\tau_H = \tau_L$ and $\pi_H = \pi_L$ in equation (6), then together with Proposition 1 we need that the maximum premium $2\omega = 2\tau\pi\Delta U$ exceed 2.91, which is equivalent to $\pi \geq 0.81$.

²⁸ We are grateful to an anonymous associate editor for suggesting a treatment along these lines.

voting premium and Table 10 summarizes the prices and bids in the feedback treatment. Voting behavior itself is not affected by feedback: comparing subjects' error rates τ_L and τ_H in FEED with those in BASE shows that voting behavior does not differ between these treatments. Also, the probabilities of being pivotal are very similar in treatments BASE and FEED. If participants would use the feedback to update their beliefs about others' choices correctly and still remain overconfident about their own error rates, then the voting premium should decline to the levels calculated in Panel A of Table 9.

Our results provide partial support for this hypothesized effect of feedback on beliefs. On the one hand, the average (median) voting premium in the feedback treatment is only 1.13 (1.10), which is economically and statistically significantly lower than the premium in BASE, although still statistically significantly different from zero at the 1% level. This decline in the voting premium suggests that the feedback helps subjects to update their beliefs about their pivotality and about others' error rates. The premium is now comfortably below the upper bound implied by the level-k model. Hence, overconfidence together with overestimation of the errors of others can potentially explain the results in FEED. On the other hand, Panel A of Table 9 shows that Quantal Response Equilibrium with overconfidence, but without overestimation of the errors of others would only result in a premium of 0.13 given the voting behavior in FEED and would therefore be insufficient to explain the voting premium.²⁹ Apparently, the feedback does not fully eliminate subjects' overestimation of others' error rates or their own pivot probability.

8.2 Prices and bidding behavior in the belief elicitation treatment

Eliciting subjects' beliefs allows us to further investigate the relationship between the belief to be pivotal in the vote and the voting premium. Treatment BE differs from BASE only in that we elicit the beliefs of our experimental subjects about their own pivot probability. Thus, they receive the same feedback as in BASE. The results for the prices, bids, and voting premiums are summarized in Table 11; summaries of voting behavior can be found in Table 7 and test results for the premium in Table 9.

²⁹ Note that in FEED, we provide our subjects with distributions of some or all past votes and information about the frequency of close elections, but with no *personalized* information about error rates, i.e., subjects are not reminded what their own vote was. Hence, a systematic underestimation of one's own error rate is sustainable even in FEED. Moreover, the feedback did not involve any judgments as to which votes were correct or incorrect. Given the probabilistic nature of all treatments (except NU) in which the dividend was an informative signal about the manager's quality, there is room for false beliefs that are not corrected, since we did not disclose the true quality of the manager.

We first observe that belief elicitation does not seem to influence voting behavior: Subjects' error rates, τ_L and τ_H , and pivot probabilities, π_L and π_H , are similar in treatment BE to those in BASE and FEED. However, belief elicitation significantly reduces the voting premium: It is much lower than in BASE, and the difference is weakly significant.³⁰ In fact, the voting premium in BE is almost identical to the one in FEED.

Figure 3 provides further evidence on individuals' bidding behavior in BASE, FEED, and BE by showing density plots of the bid premium, $\frac{Bid_A - Bid_B}{Theoretical\ Value}$. First, belief elicitation affects bidding behavior in a way that is remarkably similar to that in FEED. In fact, the graphs for the distributions of bid premiums are almost indistinguishable. Second, treatments BE and FEED differ from the baseline treatment mainly by a shift of the frequency distribution from large bid premiums towards the medium range. Panel C of Table 9 provides additional information on bidding behavior by grouping bids as large (premium exceeds 15% of the theoretical value), intermediate (premium between -15% and +15%), and small (premium below -15%). In all three treatments, about 6% of the bid premiums are small. However, in BASE, 36% of the bid premiums are large and 58% are intermediate, whereas in FEED and BE, only 14% of bid premiums are large, whereas about 80% are in the intermediate range. Hence, the difference between BASE on the one hand and FEED and BE on the other hand can be entirely accounted for by the fact that 22% of large bids in BASE are replaced by intermediate bids in FEED and BE. This shift in the distribution of bids has a significant impact on prices.

We therefore observe that prompts to reflect on pivotality (BE) and the provision of additional information relevant to correctly assessing pivotality (FEED) generate very similar outcomes that differ from BASE by changing the behavior of a little more than one fifth of the subjects. This suggests that an increase in the *salience* of pivotality suffices to prevent extreme voting premiums. As a result, the voting premium in BE as well as in FEED is consistent with the level-k model (see Panel B of Table 9).

8.3 Beliefs in treatment BE

After comparing the voting premiums in BE, BASE, and FEED, we now turn to the analysis of the stated beliefs in BE and ask to what extent they can be accounted for by the level-k model.

³⁰ A number of experiments have shown that belief elicitation can affect the behavior of experimental subjects. See Rutström and Wilcox (2009) and the references cited therein.

Table 12 provides the results from belief elicitation and Figure 4 contains the distribution of beliefs. Observe that the variation of beliefs is large. The distribution is bimodal, with one mode at a pivot probability of zero and another at 50%. Most individuals are remarkably consistent in their beliefs about their pivotality: Panel A of Table 12 shows that 57 (71.3%) of the 80 individuals in treatment BE always state the same belief during the last 15 rounds. Therefore, we first calculate the average reported belief π_r for each individual and then group individuals into five intervals depending on their average reported beliefs.

Of the 57 individuals who always report the same beliefs, 38 (47.5%) report $\pi_r = 0$ and a further 16 (20%) report $\pi_r = 0.5$. We cannot exclude the possibility that subjects who report $\pi_r = 0.5$ may simply regard the middle of the admissible range as a useful response without having worked through the implications of the game. However, their bidding behavior, which we analyze further below, suggests otherwise. A total of 59, or almost three quarters of the individuals, have average beliefs that are below the 37.5% cut-off that can be rationalized by the level-k model. The average belief of all subjects is 17.8% and exceeds the average frequency of being pivotal, which is 11.8%, but this difference is statistically not significant, and we cannot reject the hypothesis that average beliefs are consistent with observed frequencies.³¹ Hence, on average and for the large majority of individuals, beliefs are within the upper limit of the level-k model and therefore consistent with the notion that individuals overestimate the errors of others. This complements the analysis in the previous section showing that the level-k model is consistent with the observed voting premiums in BE.

8.4 Beliefs and bidding behavior

As a final step of the analysis, we investigate the relationship between beliefs, bids, and prices. Specifically, we ask if players use best responses based on their reported pivot probabilities when bidding for A- and B-shares in treatment BE. In the right part of Panel A of Table 12 we use the level-k model to calculate the *implied bid premium*. We define the implied bid premium as the largest voting premium that overconfident individuals should be willing to pay, given their reported pivot probabilities and overestimation of others' error rates according to the level-k model. We build on our analysis in Section 6.4 and set $\tau_H = \tau_L = \tau = 0.5$ and $\pi_H = \pi_L = \pi$ in

³¹ P-values are 0.178 for the t-test and 0.169 for the Wilcoxon signed-rank test.

equation (6), assuming that π corresponds to the reported pivot probability. Then, the upper limit on the bid premium implied by Proposition 1 and equation (6) becomes

$$\text{Premium} \leq 2\omega = 2\tau\pi\Delta U = \pi\Delta U = \text{Implied Premium} , \quad (7)$$

with $\Delta U = 3.6$. This implied bid premium is then averaged for each individual and the average implied bid premium across individuals is reported alongside the actual premium in Panel A of Table 12. We also report the percentage of individuals who overbid, i.e., who bid a premium for A-shares that is strictly larger than the upper limit from (7).

We find that the average premium of a participant increases with his reported pivot probability, with the exception of those reporting pivot probabilities above 50%. Moreover, the percentage of subjects who overbid relative to the implied premium is smaller for those who report a pivot probability of zero than for those who report positive pivot probabilities. In fact, of the 720 reports with a pivot probability of zero, 71% are associated with a bid premium of zero; hence most individuals seem to recognize the implications of not being pivotal.

In Panel B of Table 12 we investigate the relationship between actual premiums and implied premiums by running panel regressions of the actual bid premium against the implied bid premium. If the model with overconfidence and reported pivot probabilities has explanatory power, then we expect a consistent relationship between the actual and the implied premium with a regression coefficient of one. Indeed, the F-test can never reject the hypothesis that the slope coefficient of the regression is one, i.e., that the actual bid premium increases one for one with the implied premium.³² Overall, we conclude that there is a consistent relationship between the voting premium and the belief of being pivotal at the individual level.

Summing up, the premiums observed in FEED and BE can be accounted for by level-k reasoning despite the fact that some individuals state too high pivot probabilities and bid too high premiums even in these two treatments. We conjecture that the observed higher bid premium in BASE can be accounted for by a higher number of such individuals, since more subjects seem to hold extreme pivot beliefs or misperceptions of pivotality when the issue of pivotality is not salient.

³² There is significant within-subject variation in bid premiums, even though there is little within-subject variation in reported beliefs about pivotality. Hence subjects change their bidding behavior for reasons other than their beliefs about pivotality. Accordingly, there is a consistent relationship between implied and actual premiums, but the explanatory power of all three regressions is low and comes mostly from the variation between subjects.

9 Conclusion

In this paper we perform an experiment in which subjects pay for participating in a vote in a common interest setting. There is no asymmetric information, and no features that could trigger non-instrumental motivations for voting like a perceived civic duty or expressive motives. We find that individuals are willing to pay a significant premium for participating in the vote, even though standard game theory suggests that the value of participating in the vote is zero.

If we change a treatment variable such that there is no instrumental reason for voting at all, the voting premium vanishes. We conclude that individuals pay for participating in the vote for instrumental reasons. We then calibrate our experiments to a model that assumes that subjects are overconfident and believe that they can contribute to better voting decisions by voting themselves rather than allowing other participants in the experiment to vote. However, a version of the QRE model adapted to include this aspect explains only 4% to 5% of the voting premium we observe in our experiments.

We can explain most of the experimental evidence if we combine the following assumptions: subjects (1) are overconfident in the sense that they believe not to make mistakes that other subjects make, (2) they overestimate the actual probability that other subjects make mistakes and (3) they overestimate their own probability of being pivotal. Biases (1) and (2) are both captured by the level-k model. In one of the treatments we elicit subjects' beliefs about being pivotal and show that only subjects with excessively large beliefs of being pivotal are willing to pay for participating in the vote. We obtain large premiums for voting shares if these individuals become the marginal bidders in the auction. Moreover, if individuals have to reflect on their pivotality, their average willingness to pay for participating in the vote declines by more than 60%, although it remains large and significant.

Overestimation of one's pivot probability can explain the existence of a voting premium under the assumption of overconfidence. Moreover, it can explain the size of the premium under the additional assumption that subjects also overestimate the error rates of other subjects. The level-k model is sufficient to account for the size of the voting premium in treatments that make the issue of pivotality salient, either by easy-to-read feedback or by belief elicitation. Our contribution is to show that a sizeable voting premium can occur even if subjects have only instrumental reasons to vote.

Acknowledgements

We are grateful to Dana Novotny, Jana Stöver, and André Schmelzer for excellent research assistance as well as to Nina Bonge, Dietmar Fehr, Wei Min Wang, and Mark Henninger for their help with conducting the experiments. The paper has benefited from discussions with seminar audiences in Berlin, Copenhagen, Frankfurt/Main, Helsinki, Konstanz, London, Mannheim, Munich, Princeton, Tübingen, Warwick, and Vienna. We also thank Roland Bénabou, Thomas König, Beni Lauterbach, and Wendelin Schnedler for their comments on an earlier draft of this paper.

Funding

This project was funded by the Rudolf von Bennigsen-Foerder foundation and by Deutsche Forschungsgemeinschaft through CRC 649 “Economic risk” at Humboldt University, Berlin, and through CRC 504 “Rationality concepts, decision behavior, and economic modeling” at the University of Mannheim.

References

- Adams, Renée and Daniel Ferreira (2008): One Share-One Vote: The Empirical Evidence, *Review of Finance* 12, 51-91.
- Bartels, L.M. (1988): *Presidential Primaries and the Dynamics of Public Choice*. Princeton, (Princeton University Press).
- Battaglini, Marco, Rebecca Morton and Thomas R. Palfrey (2006): The Swing Voter's Curse in the Laboratory, *Review of Economic Studies* 77, 61-89.
- Bénabou, Roland (2013): Groupthink: Collective Delusions in Organizations and Markets. *Review of Economic Studies* 80, 429–462.
- Brennan, Geoffrey and Alan Hamlin (1998): Expressive voting and electoral equilibrium, *Public Choice* 95, 149-175.
- Brennan, Geoffrey and L. Lomasky (1993): *Democracy and Decision*, Cambridge (Cambridge University Press)
- Butler, David and Donald Stokes (1969): *Political Change in Britain*, New York (St. Martins).
- Callander, Steven (2007): Bandwagons and Momentum in Sequential Voting, *Review of Economic Studies* 74, 653-684.
- Downs, Anthony (1957): *An Economic Theory of Democracy*, New York (Harper)
- Duffy, John and Margit Tavits (2008): Beliefs and Voting Decisions: A Test of the Pivotal Voter Model, *American Journal of Political Science* 52, 603-618.
- Farber, Henry S. (2010): Rational choice and voter turnout: Evidence from union representation elections, Working paper, Princeton University.
- Feddersen, Timothy, Sean Gailmard and Alvaro Sandroni (2009): Moral bias in large elections: Theory and experimental evidence, *American Political Science Review* 103, 175-192.
- Feddersen, Timothy J. and Wolfgang Pesendorfer (1996): The Swing Voter's Curse, *American Economic Review* 86, 404-424.
- (1998): Convicting the Innocent: The Inferiority of Unanimous Jury Verdicts under Strategic Voting, *American Political Science Review* 92, 23-36.
- (1999): Abstention in Elections with Asymmetric Information and Diverse Preferences, *American Political Science Review* 93, 381-398.
- Fiorina, Morris P. (1974): *Representatives, Roll Calls, and Constituencies*, Lexington (Lexington Books)
- Fischbacher, Urs (2007): Z-Tree: Zurich Toolbox for Ready-Made Economic Experiments, *Experimental Economics* 10, 171-178.
- Friedman, Jeffrey, ed. (1995): *The Rational Choice Controversy*. New Haven, CT: Yale University Press.
- Green, Donald P., and Ian Shapiro (1994): *Pathologies of Rational Choice Theory*. New Haven, CT: Yale University Press.

- Goeree, Jacob, Holt, Charles A. and Thomas R. Palfrey (2002): Quantal Response Equilibrium and Overbidding in Private-Value Auctions, *Journal of Economic Theory*, 104(1), May 2002, 247-272
- Guarnaschelli, Serena, Richard D. McKelvey and Thomas R. Palfrey (2000): An Experimental Study of Jury Decision Rules, *American Political Science Review* 94, 407-423.
- Güth, Werner and Hannelore Weck-Hannemann (1997): Do People Care About Democracy? An Experiment Exploring the Value of Voting Rights, *Public Choice* 91, 27-47.
- Höchtel, Wolfgang, Sausgruber, Rupert and Jean-Robert Tyran (2012): Inequality aversion and voting on redistribution. *European Economic Review* 56, 1406-1421.
- Holt, Charles A., and Susan K. Laury (2002): Risk Aversion and Incentive Effects, *American Economic Review* 92(5), 1644-1655.
- Kamenica, Emir and Louise Egan (2010): Voters, Dictators, and Peons: Expressive Voting and Pivotality, Working paper, University of Chicago.
- Levine, David K. and Thomas R. Palfrey (2007): The Paradox of Voter Participation? A Laboratory Study, *American Political Science Review* 101, 143-158.
- Matusaka, John G. and Filip Palda (1993): The Downsian Model Meets the Ecological Fallacy, *Public Choice* 77, 855-878.
- McKelvey, Richard D. and Thomas R. Palfrey (1995): Quantal Response Equilibria for Normal-Form Games, *Games and Economic Behavior* 10, 6-38.
- (1998): Quantal Response Equilibria for Extensive Form Games, *Experimental Economics* 1, 9-41.
- Nagel, Rosemarie (1995): Unraveling in Guessing Games: an Experimental Study, *American Economic Review* 85, 1313-1326.
- Palfrey, Thomas R. and Howard Rosenthal (1983): A Strategic Calculus of Voting, *Public Choice* 41, 7-53.
- (1985): Voter Participation and Strategic Uncertainty, *American Political Science Review* 79, 62-78.
- Quattrone, George A. and Amos Tversky (1984): Causal versus Diagnostic Contingencies: On Self-Deception and on the Voter's Illusion, *Journal of Personality and Social Psychology* 46, 237-248.
- Ruström, E. Elisabeth, and Nathaniel T. Wilcox (2009): Stated beliefs versus inferred beliefs: A methodological inquiry and experimental test, *Games and Economic Behavior* 67, 616-632.
- Schram, Arthur (1997): Do People Care About Democracy? Comment, *Public Choice* 91, 49-51.
- Schram, Arthur and Joep Sonnemans (1996a): Voter Turnout as a Participation Game: An Experimental Investigation, *International Journal of Game Theory* 25, 385-406.
- (1996b): Why People Vote: Experimental Evidence, *Journal of Economic Psychology* 17, 417-442
- Schuessler, Alexander A. (2000): Expressive Voting, *Rationality and Society* 12, 87-119.

- Shayo, Moses and Alon Harel (2010): Non-Consequentialist voting, Working paper, Hebrew University of Jerusalem.
- Stahl, Dale O. and Paul W. Wilson (1995): On Players' Models of Other Players: Theory and Experimental Evidence, *Games and Economic Behavior* 10, 218-254.
- Tullock, Gordon (1967): Toward a Mathematics of Politics. Ann Arbor, Michigan: University of Michigan Press, chapter 7.
- Tyran, Jean-Robert (2004): Voting When Money and Morals Conflict: An Experimental Test of Expressive Voting, *Journal of Public Economics* 88, 1645-1664.
- Varian, Hal and J. K. MacKie-Mason (1995): Generalized Vickrey auctions, Technical report, University of Michigan.
- de Vries, Sven and Rakesh Vohra (2003): Combinatorial Auctions: A Survey, *INFORMS Journal on Computing* 15, 284-309.

Appendix: Proof of Proposition 1

Denote by Π^i the profits of investor i and by α^i, β^i the probability that the i^{th} investor obtains an A-share, respectively, a B-share in the auction. We denote bids by B_A^i and B_B^i and equilibrium prices by P_A and P_B . Then equilibrium profits for each investor can be written as:

$$\Pi^i = \alpha^i (V + \omega - P_A) + \beta^i (V + \alpha^i \omega - P_B) . \quad (8)$$

Part 1: Non-existence of symmetric pure strategy equilibria.

We prove this part by contradiction. Hence, assume that a symmetric pure strategy equilibrium exists where each investor bids the equilibrium prices P_A and P_B . Then each investor is rationed in equilibrium, and $\alpha = \beta = M/N$. Also, it must be the case that $\partial \Pi / \partial \alpha = \partial \Pi / \partial \beta = 0$. If $\partial \Pi / \partial \alpha < 0$ ($\partial \Pi / \partial \alpha > 0$), then the investor has an incentive to reduce (increase) his bid for the A-share, and analogously for the B-share. We therefore obtain the following prices for the conjectured equilibrium from (8):

$$\begin{aligned} \frac{\partial \Pi}{\partial \alpha} = 0 &\Rightarrow P_A = V + \left(1 + \frac{M}{N}\right) \omega , \\ \frac{\partial \Pi}{\partial \beta} = 0 &\Rightarrow P_B = V + \frac{M}{N} \omega . \end{aligned} \quad (9)$$

For any other set of prices, each investor would have an incentive to deviate from the conjectured equilibrium strategy. However, with these prices, we obtain $\Pi = -(M/N)^2 \omega < 0$. Hence, all investors would make negative profits. This cannot be the case, because each investor has the option not to bid in the auction and receive zero profits.

Part 2: Existence of asymmetric pure strategy equilibria.

Denote by $\varepsilon > 0$ the smallest increment by which investors can increase their bids. We now construct an asymmetric pure strategy equilibrium as follows.

$$\begin{aligned} B_A^i &= V + (1 + \gamma) \omega + \varepsilon , \quad i = 1 \dots M, & B_A^i &= V + (1 + \gamma) \omega , \quad i = M + 1 \dots N, \\ B_B^i &= V + (1 - \gamma) \omega + \varepsilon , \quad i = 1 \dots M, & B_B^i &= V + (1 - \gamma) \omega , \quad i = M + 1 \dots N, \end{aligned} \quad (10)$$

where $0 \leq \gamma \leq 1$. The parameter γ describes how the valuation premium for the B-share is allocated between the price for B-shares and the price for A-shares. Then $\alpha^i = \beta^i = 1$ for the M highest bidders, $i = 1 \dots M$, who win all the shares in equilibrium. Also, $\alpha^i = \beta^i = 0$ for the $M-N$

other bidders, $i=M+1 \dots N$, who never win a share in any auction. The profits of the winners in the auction are then:

$$\Pi = V + \omega - (V + (1 + \gamma)\omega) + V + \omega - (V + (1 - \gamma)\omega) = 0 . \quad (11)$$

The losers in the auctions make zero profits, too. The profits of a winner who would deviate by reducing his bid below the stipulated equilibrium price in one of the auctions (and without changing the strategy in the other auction) would be:

$$\begin{aligned} \Pi^{Winner}(\text{bid} < B_B^{M+1}) &= V + \omega - (V + (1 + \gamma)\omega) = -\gamma\omega \leq 0 \\ \Pi^{Winner}(\text{bid} < B_A^{M+1}) &= V - (V + (1 - \gamma)\omega) = -(1 - \gamma)\omega \leq 0. \end{aligned} \quad (12)$$

Bidding lower in both auctions results in profits of zero. If winners increase their bids above those in (10), this has no consequence unless one of the losing bidders also increases his price, in which case they would overpay. Hence, bidding higher is a weakly dominated strategy.

The losers of the auction could bid higher in one or both auctions. In order to win, they would have to increase their bids at least to B_A^M and B_B^M (then they would be rationed) or by an increment ε higher to win with probability one. The payoffs from bidding B_A^M and B_B^M are (note that ε , α , and β are all strictly positive):

$$\begin{aligned} \Pi(\text{bid } B_A^M \text{ for A}) &= \alpha^i (V + \omega - (V + (1 + \gamma)\omega + \varepsilon)) \\ &= -\alpha^i (\gamma\omega + \varepsilon) < 0 \\ \Pi(\text{bid } B_B^M \text{ for B}) &= \beta^i (V - (V + (1 - \gamma)\omega + \varepsilon)) = -\beta^i ((1 - \gamma)\omega + \varepsilon) < 0 \\ \Pi(\text{bid } B_A^M \text{ and } B_B^M) &= \alpha^i (V + \omega - (V + (1 + \gamma)\omega + \varepsilon)) + \beta^i (V + \alpha^i \omega - (V + (1 - \gamma)\omega + \varepsilon)) \\ &= -\alpha^i (\gamma\omega + \varepsilon) - \beta^i ((1 - \gamma - \alpha^i)\omega + \varepsilon) < 0. \end{aligned} \quad (13)$$

Bidding above B_A^M and B_B^M increases the probability of winning to one and generates even lower payoffs. Hence, the losers of the auction have no incentive to deviate by bidding higher for either the A-share or the B-share or both. Thus, the voting premium is

$$P_A - P_B = B_A^{M+1} - B_B^{M+1} = 2\gamma\omega . \quad (14)$$

Since γ can be any number in the unit interval, it holds that $0 \leq P_A - P_B \leq 2\omega$.

Figure 1: Prices over time in the baseline treatment

This figure shows, for each period of the BASE treatment, the prices P_A and P_B of voting A-shares and non-voting B-shares (averaged across the 10 groups). In addition, the plot shows a cubic spline that smoothes the observations across periods. The solid line represents the smoothed price of a voting A-share, and the broken line the smoothed price of a nonvoting B-share.

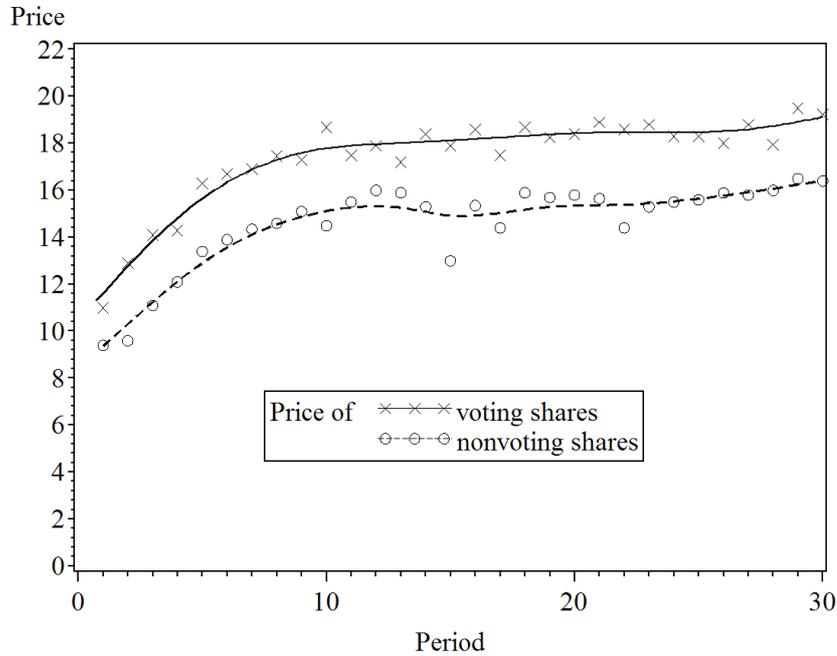


Figure 2: Incorrect votes over time in the baseline treatment

This figure shows, for each period of the BASE treatment, the frequency τ of wrong votes conditional on the observed first dividend D_1 . In addition, the plot shows a cubic spline that smoothes the observations across periods. The solid line represents the smoothed error rate after a low dividend, and the broken line represents the smoothed error rate after a high dividend.

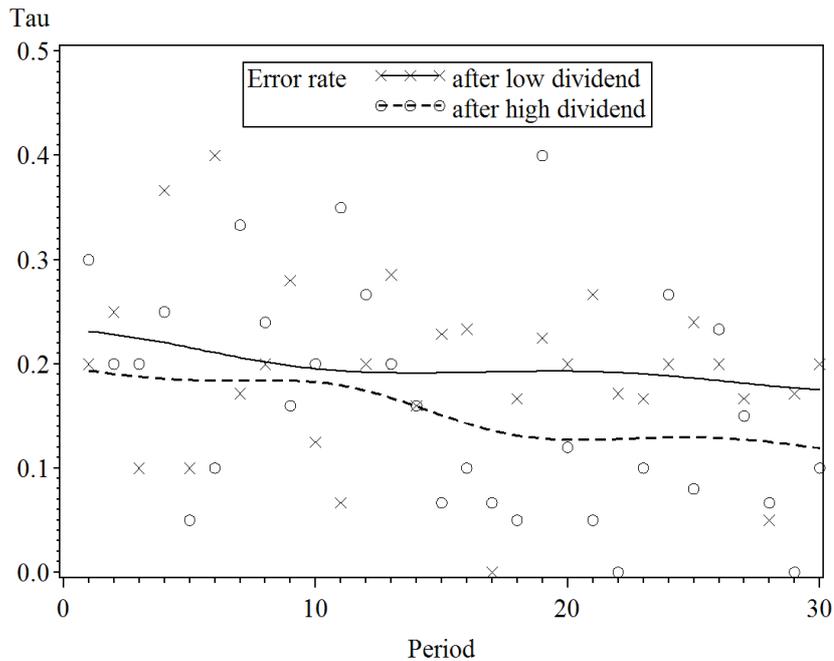


Figure 3: Density graphs of bid premiums in treatments BASE, FEED, and BE

The figure graphs the kernel densities of the bid premiums in treatments BASE, FEED, and BE. For each treatment, only bids for the last 15 rounds in each group were included in the graph. The kernel densities are estimated with a Gaussian kernel with bandwidth 0.07.

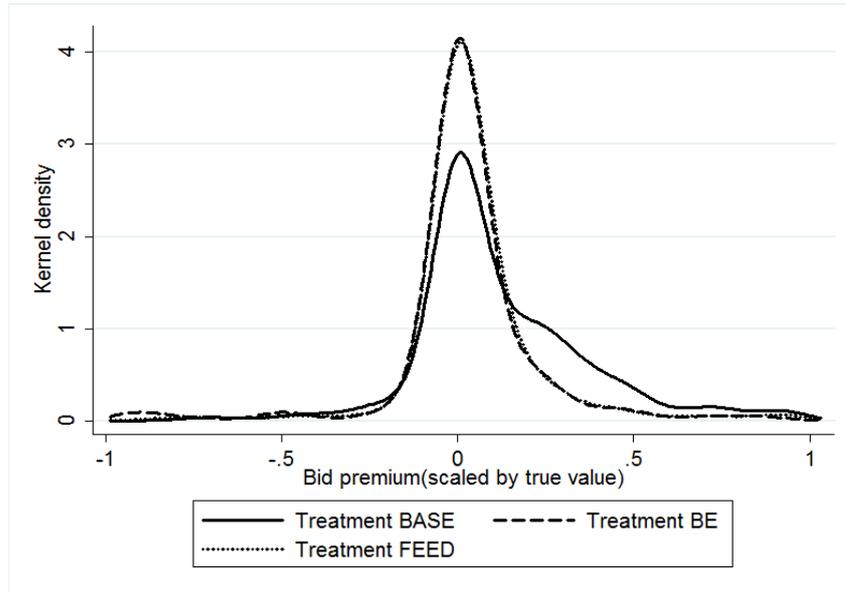


Figure 4: Distribution of beliefs stated in treatment BE

Subjects were asked in treatment BE to indicate their beliefs about the probability that they themselves are pivotal in the upcoming vote. This figure shows the histogram of the beliefs during the last 15 rounds.

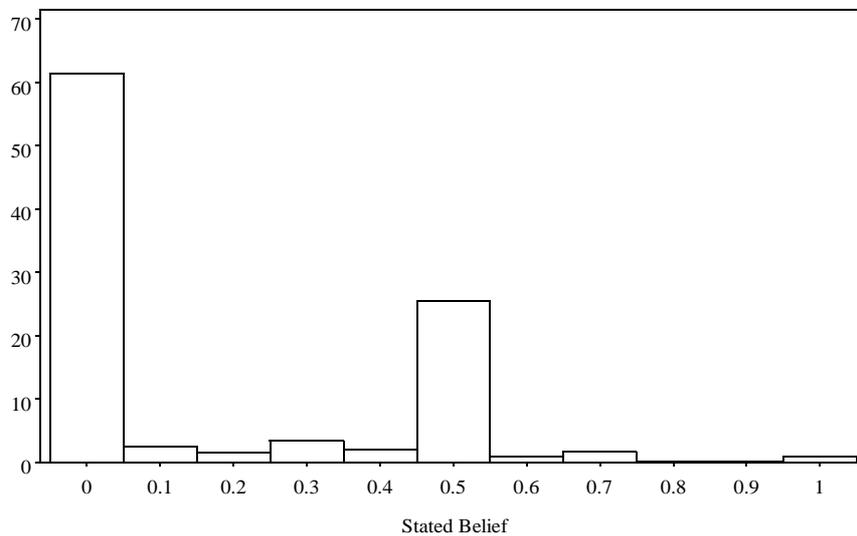


Table 1: Expected Dividends

Panel A shows the expected dividends in the second period, $E(D_2|D_1, \delta)$, conditional on the first-period dividend D_1 and the replacement decision δ where $\delta = K$ indicates that the old manager is kept, whereas $\delta = R$ indicates that the old manager is replaced. After observing a high dividend, , the posterior probability of the manager to be good is p , after a low dividend it is $1 - p$. The expected dividend with a good manager is $D^+ = pH + (1 - p)L$, with a bad manager it is $D^- = (1 - p)H + pL$. Panel B shows the probabilities of aggregate outcomes for each round of the play. “HH” is the outcome that both dividends are high, “LL” is the outcome that both dividends are low, and “HL” is the outcome that one dividend is high and one is low.

Panel A: Conditional expectations of second-period dividends

Dividend D_1	Replacement δ	Expected dividend $E(D_2 D_1, \delta)$
H	R	$\frac{H+L}{2}$
H	K	$pD^+ + (1-p)D^- = \frac{H+L}{2} + 2\left(\frac{1}{4} - p(1-p)\right)(H-L)$
L	R	$\frac{H+L}{2}$
L	K	$(1-p)D^+ + pD^- = \frac{H+L}{2} - 2\left(\frac{1}{4} - p(1-p)\right)(H-L)$

Panel B: Probabilities of aggregate outcomes

Outcome	Probability
HH	$\frac{1}{2}(p^2 + (1-p)^2)$
HL	$p(1-p) + \frac{1}{4}$
LL	$\frac{1}{4}$

Table 2: Overview of the different treatments

This table shows the main parameters and theoretical values for each treatment. Panel A shows p , the likelihood that the dividend is high (low) when the manager is good (bad), the number of shares held by a blockholder who always votes in favor of the manager, whether beliefs are elicited, whether additional feedback on voting behavior was provided, the number of groups, the number of rounds per group, and the number of subjects per group. Panel B shows the benefits from a correct decision ΔU from (5), the theoretical value of the shares, V_{NE} , from equation (1), and the probabilities of outcomes calculated from Panel B in Table 1.

Panel A: Parameters of the treatments

Treatment	Link p between skill and dividend	Shares held by block- holder	Belief Elici- tation	Feed- back	Group	Rounds per group	Subjects per group
Base treatment (BASE)	80%	0	no	no	10	30	8
No uncertainty (NU)	100%	0	no	no	10	30	8
Noninfluential manager (NI)	50%	0	no	no	10	30	8
Non pivotal (NP)	80%	6	no	no	10	30	8
Feedback (FEED)	80%	0	no	yes	10	30	8
Belief Elicitation (BE)	80%	0	yes	no	10	30	8

Panel B: Theoretical outcomes for the treatments

Treatment	ΔU	Theore- tical value of shares	Probability of outcome		
			HH	HL	LL
Base treatment (BASE)	3.6	21.8	34%	41%	25%
No uncertainty (NU)	10.0	25.0	50%	25%	25%
Noninfluential manager (NI)	0.0	20.0	25%	50%	25%
Non pivotal (NP)	3.6	20.0	34%	41%	25%
Feedback (FEED)	3.6	21.8	34%	41%	25%
Belief Elicitation (BE)	3.6	21.8	34%	41%	25%

Table 3: Prices in the baseline treatment (BASE)

This table contains statistics on four variables for the baseline treatment: the price of voting A-shares, the price of non-voting B-shares, the difference between these two prices (the premium), and the premium scaled by the equilibrium value of the shares V_{NE} (the relative premium). The table shows mean, median, standard deviation, and the p-values of two tests, the t-test for zero mean and the Wilcoxon signed-rank test. For the prices of voting and non-voting shares, the table also shows the mean difference between the price and the equilibrium value V_{NE} , and the p-value of the t-test for this difference to be zero. In Panel A, we first calculate one value of the respective variable for each group by averaging the variable across the last 15 rounds of the group. In the second step, we calculate the statistics across the 10 groups. Panel B shows the results for the pooled sample of all groups and rounds.

Panel A: Average of last 15 rounds

Variable	Obs.	Mean	Median	Standard deviation	T-test p-value	Wilcoxon p-value	Difference to Nash	
							Mean	T-test p-value
Voting A-share	10	18.52	18.27	4.65	0.000	0.005	3.28	0.053
Non-voting B-share	10	15.61	17.37	5.41	0.000	0.005	6.19	0.006
Premium	10	2.91	2.63	2.50	0.005	0.007		
Relative premium	10	13.4%	12.1%	11.5%	0.005	0.007		

Panel B: All rounds

Variable	Obs.	Mean	Median	Standard deviation	T-test p-value	Wilcoxon p-value	Difference to Nash	
							Mean	T-test p-value
Voting A-share	300	17.41	18.00	4.63	0.000	0.000	4.39	0.000
Non-voting B-share	300	14.60	15.76	5.30	0.000	0.000	7.20	0.000
Premium	300	2.82	2.00	3.31	0.000	0.000		
Relative premium	300	12.9%	9.2%	15.2%	0.000	0.000		

Table 4: Bids in the baseline treatment (BASE)

This table contains statistics on four variables for the baseline treatment: the bid for voting A-shares, the bid for non-voting B-shares, the difference between these two bids (the premium), and the premium scaled by the equilibrium value of the shares V_{NE} from equation (1) (the relative premium). The table shows mean, median, standard deviation, and the p-values of two tests, the t-test for zero mean and the Wilcoxon signed-rank test. For the premium, the table also shows the frequencies that the premium is negative or, respectively, positive. In Panel A, we first calculate one value of the respective variable for each subject by averaging the variable across the last 15 rounds. In the second step, we calculate the statistics across the 10 groups. Panel B shows the results for the pooled sample of all subjects and rounds.

Panel A: Average of last 15 rounds

Variable	Obs.	Mean	Median	Standard deviation	T-test p-value	Wilcoxon p-value	Frequency	
							negative	positive
Voting A-share	10	29.41	26.38	9.85	0.000	0.005		
Non-voting B-share	10	23.95	22.63	8.52	0.000	0.005		
Premium	10	5.6	3.89	4.22	0.003	0.005	0.0%	100.0%
Relative premium	10	25.0%	17.9%	19.4%	0.003	0.005		

Panel B: Average of all rounds

Variable	Obs.	Mean	Median	Standard deviation	T-test p-value	Wilcoxon p-value	Frequency	
							negative	positive
Voting A-share	2400	27.76	20.00	36.43	0.000	0.000		
Non-voting B-share	2400	22.65	18.00	32.18	0.000	0.000		
Premium	2400	5.11	1.00	17.39	0.000	0.000	10.8%	53.8%
Relative premium	2400	23.4%	4.6%	79.8%	0.000	0.000		

Table 5: Prices in the non-influential manager treatment (NI)

This table describes prices in the non-influential manager treatment (NI). It shows statistics on the price of voting A-shares, the price of non-voting B-shares, the difference between these two prices (the premium), and the premium scaled by the equilibrium value of the shares V_{NE} (the relative premium). The table shows mean, median, standard deviation, and the p-values of two tests, the t-test for zero mean and the Wilcoxon signed-rank test. In addition, it shows the p-values of the two-sample t-test, of the Mann-Whitney U test, and of a clustered regression comparing the prices and premiums in NI and BASE. For the t-test and the Mann-Whitney U test, we first calculate one value of the respective variable for each group by averaging the variable across the last 15 rounds of the group. In the second step, we calculate the statistics across the 10 groups. For the regression, we cluster at the group level and regress the price or relative premium on a treatment dummy that is one if the treatment is BASE and 0 otherwise.

Average prices of the last 15 rounds

Variable	Obs.	Mean	Median	Standard deviation	T-test p-value	Wilcoxon p-value	P-value two-sample comparison with BASE		
							T-test	Mann-Whitney U	Clustered regressions
Voting A-share	10	16.38	17.03	2.63	0.000	0.005	0.225	0.151	0.209
Non-voting B-share	10	15.98	16.47	2.72	0.000	0.005	0.851	0.545	0.846
Premium	10	0.40	0.28	0.67	0.089	0.138	0.012	0.015	0.005
Relative premium	10	2.0%	1.4%	3.3%	0.089	0.139	0.013	0.019	0.006

Table 6: Prices in the non-pivotal treatment (NP)

This table describes prices in the non-pivotal treatment (NP). It shows statistics on the price of voting A-shares, the price of non-voting B-shares, the difference between these two prices (the premium), and the premium scaled by the equilibrium value of the shares V_{NE} (the relative premium). The table shows mean, median, standard deviation, and the p-values of two tests, the t-test for zero mean and the Wilcoxon signed-rank test. In addition, it shows the p-values of the two-sample t-test, of the Mann-Whitney U test, and of a clustered regression comparing the prices and premiums in NP and BASE. For the t-test and the Mann-Whitney U test, we first calculate one value of the respective variable for each group by averaging the variable across the last 15 rounds of the group. In the second step, we calculate the statistics across the 10 groups. For the regression, we cluster at the group level and regress the price or relative premium on a treatment dummy that is one if the treatment is BASE and 0 otherwise.

Average prices of the last 15 rounds

Variable	Obs.	Mean	Median	Standard deviation	T-test p-value	Wilcoxon p-value	P-value two-sample comparison with BASE		
							T-test	Mann-Whitney U	Clustered regressions
Voting A-share	10	13.14	12.97	4.08	0.000	0.005	0.013	0.026	0.011
Non-voting B-share	10	12.85	12.73	4.02	0.000	0.005	0.212	0.112	0.199
Premium	10	0.29	0.33	0.93	0.346	0.201	0.010	0.028	0.005
Relative premium	10	1.5%	1.7%	4.7%	0.346	0.169	0.011	0.041	0.006

Table 7: Voting behavior

This table describes the voting outcomes for all five treatments based on data from the last 15 rounds of the experiment. The table shows for each treatment (1) the frequency τ_L that A-shareholders vote for keeping the manager after a low dividend, (2) the frequency e_L that a manager is actually kept after a low dividend, (3) the probability π_L of an individual to be pivotal after a low dividend (4) the frequency τ_H that A-shareholders vote for replacing the manager after a high dividend, (5) the frequency e_H that a manager is actually replaced after a high dividend, and (6) the probability π_H of an individual to be pivotal after a high dividend.

Treatment	τ_L	e_L	π_L	τ_H	e_H	π_H
Baseline treatment (BASE)	18.3%	7.1%	16.4%	11.8%	3.0%	7.3%
No uncertainty treatment (NU)	11.0%	0.0%	3.9%	10.6%	1.4%	5.8%
Non-influential manager treatment (NI)	56.3%	64.5%	43.4%	21.9%	8.1%	18.7%
Non-pivotal treatment (NP)	31.2%	15.1%	29.3%	23.8%	7.8%	23.4%
Feedback treatment (FEED)	20.5%	2.5%	16.5%	11.4%	1.4%	2.6%
Belief elicitation treatment (BE)	18.8%	1.2%	12.6%	15.1%	5.8%	9.6%

Table 8: Prices and bids in the no-uncertainty treatment (NU)

This table describes prices (Panel A) and bids (Panel B) for the no-uncertainty treatment (NU). Panel A shows statistics on the price of voting A-shares, the price of non-voting B-shares, the difference between these two prices (the premium), and the premium scaled by the equilibrium value of the shares V_{NE} (the relative premium). The table shows mean, median, standard deviation, and the p-values of two tests, the t-test for zero mean and the Wilcoxon signed-rank test. In addition, it shows the p-values of the two-sample t-test, of the Mann-Whitney U test, and of a clustered regression comparing the prices and premiums in NU and BASE. For the t-test and the Mann-Whitney U test, we first calculate one value of the respective variable for each group by averaging the variable across the last 15 rounds of the group. In the second step, we calculate the statistics across the 10 groups. For the regression, we cluster at the group level and regress the price or relative premium on a treatment dummy that is one if the treatment is BASE and 0 otherwise. Panel B displays the same statistics for bids.

Panel A: Average prices of the last 15 rounds

Variable	Obs.	Mean	Median	Standard deviation	T-test p-value	Wilcoxon p-value	P-value two-sample comparison with BASE		
							T-test	Mann-Whitney U	Clustered regressions
Voting A-share	10	20.90	20.77	4.37	0.000	0.005	0.254	0.273	0.242
Non-voting B-share	10	16.91	17.37	2.96	0.000	0.005	0.518	0.762	0.504
Premium	10	3.99	3.13	3.27	0.004	0.005	0.417	0.385	0.404
Relative premium	10	16.0%	12.5%	13.1%	0.004	0.005	0.640	0.650	0.631

Panel B: Average bids of the last 15 rounds

Variable	Obs.	Mean	Median	Standard deviation	T-test p-value	Wilcoxon p-value	P-value two-sample comparison with BASE		
							T-test	Mann-Whitney U	Clustered regressions
Voting A-share	10	25.89	24.70	5.78	0.000	0.005	0.346	0.496	0.651
Non-voting B-share	10	22.16	20.96	4.03	0.000	0.005	0.560	0.940	0.520
Premium	10	3.72	3.10	5.13	0.048	0.047	0.420	0.326	0.232
Relative premium	10	14.9%	12.4%	20.5%	0.048	0.047	0.420	0.326	0.199

Table 9: Tests of instrumental voting models with behavioral biases

This table tests predictions on the upper bound of the voting premium against the observed voting premium for treatments BASE, NU, FEED, and BE. Panel A generates predictions from the Quantal Response Equilibrium model (QRE) with overconfidence. Panel B generates predictions from the level-k model. Empirical probabilities of being pivotal are the observed frequencies of being pivotal in the vote. The upper bound 2ω follows from Proposition 1 and equation (6). In Panel A individuals' error probabilities τ_L and τ_H are equal to the observed frequencies while in Panel B τ_L and τ_H are set equal to 50%. Panel A and B also display the p-values for the t-test and the one-sided Wilcoxon signed-rank test that the observed voting premium does not exceed the upper bound. Panel C groups the voting premiums for treatments BASE, FEED, and BE into three intervals as small ($\leq -15\%$), medium (between -15% and $+15\%$ of the theoretical value), and large ($\geq +15\%$). Percentages are in terms of the theoretical value, which is 21.8 in all treatments.

Panel A: QRE model with overconfidence

Treatment	Model	Pivotality	Upper bound on premium	Actual premium		T-test p-value	Wilcoxon p-value
				Mean	Median		
BASE	QRE	empirical	0.14	2.91	2.63	0.003	0.007
NU	QRE	empirical	0.10	3.99	3.13	0.002	0.003
FEED	QRE	empirical	0.13	1.13	1.10	0.002	0.003
BE	QRE	empirical	0.14	1.11	1.00	0.022	0.037

Panel B: Voting premiums with overconfidence and overestimation of the errors of others

Treatment	Model	Pivotality	Upper bound on premium	Actual premium		T-test p-value	Wilcoxon p-value
				Mean	Median		
BASE	level-k	37.5%	1.35	2.91	2.63	0.040	0.046
BASE	-	50.0%	1.80	2.91	2.63	0.097	0.121
NU	level-k	37.5%	3.75	3.99	3.13	0.410	0.439
NU	-	50.0%	5.00	3.99	3.13	0.822	0.884
FEED	level-k	37.5%	1.35	1.13	1.10	0.785	0.963
FEED	-	50.0%	1.80	1.13	1.10	0.984	0.970
BE	level-k	37.5%	1.35	1.11	1.00	0.715	0.713
BE	-	50.0%	1.80	1.11	1.00	0.936	0.954

Panel C: Frequencies of bid premiums across treatments

Treatment	Prem. $\leq -15\%$	Prem. $\geq +15\%$	$-15\% < \text{Prem.} < 15\%$
BASE	5.67%	57.92%	36.42%
FEED	5.92%	79.75%	14.33%
BE	5.50%	80.75%	13.75%

Table 10: Prices in the Feedback treatment

This table describes prices in the feedback treatment (FEED). It shows statistics on the price of voting A-shares, the price of non-voting B-shares, the difference between these two prices (the premium), and the premium scaled by the equilibrium value of the shares V_{NE} (the relative premium). The table shows mean, median, standard deviation, and the p-values of two tests, the t-test for zero mean and the Wilcoxon signed-rank test. In addition, it shows the p-values of the two-sample t-test, of the Mann-Whitney U test, and of a clustered regression comparing the prices and premiums in FEED and BASE as well as in FEED and BE. For the t-test and the Mann-Whitney U test, we first calculate one value of the respective variable for each group by averaging the variable across the last 15 rounds of the group. In the second step, we calculate the statistics across the 10 groups. For the regression, we cluster at the group level and regress the price or relative premium on a treatment dummy.

Average prices of the last 15 rounds

Variable	Obs.	Mean	Median	Standard deviation	T-test p-value	Wilcoxon p-value	P-value two-sample comparison with BASE			P-value two-sample comparison with BE		
							T-test	Mann-Whitney U	Clustered regressions	T-test	Mann-Whitney U	Clustered regressions
Voting A-share	10	17.27	16.27	3.98	0.000	0.005	0.524	0.326	0.513	0.952	0.880	0.950
Non-voting B-share	10	16.14	15.23	4.15	0.000	0.005	0.811	0.650	0.806	0.959	0.880	0.958
Premium	10	1.13	1.10	0.76	0.002	0.005	0.056	0.162	0.042	0.963	0.762	0.962
Relative premium	10	5.2%	5.0%	3.9%	0.002	0.005	0.056	0.151	0.042	0.963	0.850	0.962

Table 11: Prices in the Belief Elicitation treatment

This table describes prices in the belief-elicitation treatment (BE). It shows statistics on the price of voting A-shares, the price of non-voting B-shares, the difference between these two prices (the premium), and the premium scaled by the equilibrium value of the shares V_{NE} (the relative premium). The table shows mean, median, standard deviation, and the p-values of two tests, the t-test for zero mean and the Wilcoxon signed-rank test. In addition, it shows the p-values of the two-sample t-test, of the Mann-Whitney U test, and of a clustered regression comparing the prices and premiums in BE and BASE. For the t-test and the Mann-Whitney U test, we first calculate one value of the respective variable for each group by averaging the variable across the last 15 rounds of the group. In the second step, we calculate the statistics across the 10 groups. For the regression, we cluster at the group level and regress the price or relative premium on a treatment dummy that is one if the treatment is BASE and 0 otherwise. Panel B displays the same statistics for bids.

Average prices of the last 15 rounds

Variable	Obs.	Mean	Median	Standard deviation	T-test p-value	Wilcoxon p-value	P-value two-sample comparison with BASE		
							T-test	Mann-Whitney U	Clustered regressions
Voting A-share	10	17.14	16.97	5.17	0.000	0.005	0.537	0.364	0.526
Non-voting B-share	10	16.03	15.80	4.68	0.000	0.005	0.854	0.762	0.851
Premium	10	1.11	1.00	1.31	0.026	0.037	0.064	0.104	0.052
Relative premium	10	5.1%	4.6%	6.0%	0.026	0.037	0.064	0.104	0.052

Table 12: Beliefs and bids in the Belief Elicitation treatment

In this table we analyze individuals' beliefs and the relationship between their beliefs and their bidding behavior in the belief elicitation treatment. Panel A calculates the arithmetic mean π_R of the beliefs for each individual and groups individuals into five intervals according to their average reported beliefs. For each interval, the table shows the number of subjects in that interval, the fraction of subjects in that interval as a percentage of the overall number of subjects in the treatment, the number of subjects who report the same beliefs in the last 15 periods of the experiment ("No variation..."), and the average beliefs across all subjects. The actual bid premium is the difference between the bid prices for A shares and B shares. The implied bid premium is calculated as $\pi_R * 3.6$ from equation (7). „Perc. overbid" is the percentage of bids for which the actual premium strictly exceeds the implied premium. Panel B reports regressions of the actual premium on the implied premium. Model (1) uses subject fixed effects, model (2) is a between estimation, and model (3) is an OLS regression with robust standard errors that are clustered at the group level. T-statistics are in parentheses below the coefficients. The F-test is for the hypothesis that the coefficient on the implied premium equals one.

Panel A: Relation between beliefs and bids

Reported beliefs (π_R)				Bid premium			
Range	No. of subjects	Perc. of subjects	No variation	Mean π_R	Actual	Implied	Perc. Overbid
			across periods				
$\pi_R = 0$	38	47.50%	38	0.00%	0.03	0.00	18.95%
$0 < \pi_R \leq 0.375$	21	26.25%	1	14.76%	1.24	0.59	34.92%
$0.375 < \pi_R < 0.5$	2	2.50%	1	42.50%	1.53	1.67	23.33%
$\pi_R = 0.5$	17	21.25%	16	50.00%	3.06	1.80	32.55%
$0.5 < \pi_R \leq 1$	2	2.50%	1	85.00%	1.00	2.52	0.03%
Overall	80	100.00%	57	17.82%	1.03	0.64	25.75%

Panel B: Relation between implied premiums and actual bid premiums

Specification	(1)	(2)	(3)
	Subject fixed effects	Between subjects	OLS, clustered & robust standard errors
Implied premium	1.019 (2.18)	1.335 (2.20)	1.283 (2.34)
Constant	0.373 (1.09)	0.170 (0.27)	0.203 (0.33)
R-squared	0.004	0.058	0.025
Observations	1,200	1,200	1,200
F-test Impl. Prem.=1	0.000	0.310	0.270
p-value	0.967	0.582	0.618

Instructions for the baseline treatment (not for publication)

This appendix contains a translation of the German instructions given to our experimental subjects. To conserve space, we only present the instructions for the baseline treatment. For the other three treatments we changed as little as possible. For instance for the treatment NI, we merely replaced the probability 0.8 of the baseline treatment with the probability 0.5.

Welcome! This is an experiment in behavioral economics. You and the other participants in the experiment will participate in a situation where you have to make a decision. In this situation, you can earn money that will be paid out to you in cash at the end of the experiment. How much you will earn, depends on the decisions that you and the other participants in the experiment make.

These instructions describe the situation in which you have to make a decision. The instructions are identical for all participants in the experiment. It is important that you read the instructions carefully so that you understand the decision-making problem well. If something is unclear to you while reading, or if you have other questions, please let us know by raising your hand. We will then answer your questions individually.

Please do not, under any circumstances, ask your question(s) aloud. You are not permitted to give information of any kind to the other participants. You are also not permitted to speak to other participants at any time throughout the experiment. Whenever you have a question, please raise your hand and we will come to you and answer it. If you break these rules, we may have to terminate the experiment.

Once everyone has read the instructions and there are no further questions, we will conduct a short quiz where each of you will complete two short tasks on your own. We will walk around, look over your answers, and solve any remaining comprehension problems. The only purpose of the quiz is to ensure that you thoroughly understand the crucial details of the decision-making problem.

As a matter of course, your anonymity and the anonymity of the other participants will be guaranteed throughout the entire experiment. You will neither learn about the identity of the other participants, nor will they learn about your identity.

1. General

The experiment consists of several rounds in which decisions must be made and information must be processed. With your decisions, you can earn points. These points represent your income and will be converted into euros at the end of the experiment and paid out in cash. The exact sequence of the stages of the experiment, the different decisions, and the payment modalities will be explained in detail in the following.

2. The allocation of roles

You will find yourself in a situation where you have to make decisions. This situation will be repeated for 30 rounds. This situation can be compared to a simple stock market. In this stock market, there is one single company issuing shares. You can buy shares of this company and thus take on the role of a shareholder. The amount of dividends paid out on the shares depends on how good or bad the manager of the company is. Chance plays an important role here: it determines whether the manager of the company is good or bad. The computer takes on the roles of company, manager, and it performs the chance moves.

3. Two possible stock market conditions

At the beginning of *each round*, the computer draws the quality of the manager randomly. You can visualize this as a drawing of one ball out of a hat that contains five balls labeled with the words “good manager” and another five labeled with the words “bad manager.” Neither you nor any of the other participants know the quality of the manager. You only know the following: there is a 50% probability of a good manager, and a 50% probability of a bad manager. The quality of the manager influences the probability that the company will make a profit. We first consider the situation that arises in the case of a good manager.

4. The stock market condition in the (unobservable) case of a good manager

If the manager is good (which neither you nor any of the others know), there is an 80% probability that the company will make a profit. You can visualize the situation with a good manager as if the computer picked a ball randomly out of a hat containing eight balls labeled with the words “profit”, and two balls labeled with the words “no profit.” Whether or not the company makes a profit is important for you only because it affects the dividend payments of the shares.

If the company makes a profit, the dividends on shares amount to 20 points per share. If the company does not make a profit, however, the dividends amount to 0 points per share.

Now we will look at the situation in the (also unobservable) case of a bad manager.

5. The stock market condition in the (also unobservable) case of a bad manager

If the manager is bad, there is a 20% probability that the company will make a profit. You can visualize the situation with a bad manager as if the computer picked a ball randomly out of a hat containing two balls labeled with the words “profit”, and eight balls labeled with the words “no profit.” Whether or not the company makes a profit is important for you only because it affects the dividend payments of the shares.

If the company makes a profit, the dividends on shares amount to 20 points per share, as before.

If the company does not make a profit, however, the dividends amount to 0 points per share.

Obviously, a dividend of 20 points per share is more likely if the manager is good. Conversely, a dividend of 0 points per share is more likely if the manager is bad.

Please keep in mind that the quality of the manager is unknown to all participants and that both qualities are equally likely at the beginning of each round.

6. The auctioning of shares

There are two types of shares: type-A shares and type-B shares. The next section will explain the difference between them. The dividends are the same for both types.

In every round, five type-A shares and five type-B shares are sold.

Apart from you, there are seven other persons who can buy shares of the company: namely, the other participants in the experiment. So in total, there are eight persons in the market who can bid for the shares. All these persons are subject to the same rules. These rules will be explained to you in the following.

You can buy a maximum of one share of each type per round. In total, therefore, you can buy a maximum of two shares per round: one type-A share and one type-B share. You can try to buy the shares of your choice by taking part in two auctions that are conducted simultaneously. In one auction, type-A shares are sold, and in the other auction, type-B shares are sold.

You will be asked to make a bid for each share type. In each of the two auctions, the following rule applies: the five participants in the experiment who submit the five highest bids for the type of shares offered in that auction will each receive one share of that type. The price they will each pay for their share will be the amount of the sixth-highest bid in that auction. (If two or more bids in one auction are identical, the computer will randomly determine their order.) Those participants who made the sixth-to-eighth-highest bids will not receive any share of the type offered and also will not pay anything in this auction.

The most sensible thing you can do in both of these auctions is to submit a bid that corresponds to your true valuation of the share in question, in other words, the maximum amount you would be willing to pay for it.

After all participants have placed their bids in the two simultaneous auctions, everyone will be informed about the price of type-A and type-B shares, and about the dividends per share (which is the same for both types of shares). You will always be given this information, no matter whether you have purchased shares or not. Furthermore, everyone will be told which share, if any, he or she has acquired. You will not receive any information on which shares, if any, others have acquired.

The winners of the auction will then receive the five type-A shares and the five type-B shares. If you have purchased at least one share, your interim earnings will amount to, in points:

The *dividends* on the shares you purchased, multiplied by the *number* of shares you purchased, minus the *price(s)* you had to pay.

You can also have negative earnings.

If you did not acquire any share, your interim earnings amount to zero.

You (and only you) will be told your interim earnings.

In the next section, the difference between A shares and B shares will be explained.

7. The share types

When you purchase a type-A share, you acquire the right to participate in a vote. You do not acquire such a right with type-B shares, however.

All five participants who have purchased a type-A share participate in a vote whose outcome determines whether or not the company manager will be retained. This vote takes place after all share purchases have been made and after you have been told your interim earnings from that round. At this point in time, you therefore know whether the shares paid out dividends of 0 or 20 points.

If the majority of the participants who own a type-A share are in favor of keeping the same manager, she will be retained. Retaining the same manager means that the computer does not again randomly generate the quality of the manager. If retained, the manager has the same quality as before the vote. If the majority of participants who have a type-A share vote to fire the manager, she will be fired and a new manager will be hired. Firing means that the computer will

again randomly generate the quality of the new manager. Then there is again a 50% probability that the new manager is good and a 50% probability that the new manager is bad.

8. The second dividend payout on shares

After the vote on the manager, the company will again become active on the market, either with the old or with a new manager depending on the outcome of the vote.

Depending on the quality of the manager and the corresponding probability of the company's success, the computer again determines the dividends on the shares. The dividends again amount to 20 points in the case that the company is successful and 0 points in the case that the company is unsuccessful. As before, the probability of earning 20 points is higher, when the manager is good (namely, 80% for a good manager in comparison to 20% for a bad manager).

You will then be told your final earnings from that round. In the next round, the situation will be repeated with a new manager. So all rounds have the same structure.

9. Your finances

At the beginning of the experiment, you receive an initial budget of 150 points. You will also be paid a fixed amount of 16 points each round, in addition to your profits (or losses) in that round. Your final payout consists of the sum of the payouts from all rounds. Furthermore, you receive 3 euro if you arrived 10 minutes before the beginning of the experiment. The exchange rate used to convert points into cents is:

1 point = 3 cents.

10. The sequence of events of one round at a glance

1. The manager's ability is determined randomly but not disclosed. The auctions for type-A shares and type-B shares take place.
2. All participants are informed about the prices and dividends of the shares. Furthermore, each of you individually receives information about which share(s) you have acquired.
3. The dividends of 0 or 20 points per share are announced. You are told your interim earnings from that round.
4. The vote on retaining or firing the manager takes place among all holders of type-A shares. Depending on the results of the vote, either the same manager is retained or a new manager is hired.
5. The result of the vote is announced to all participants.

6. The second dividend payout of 0 or 20 points per share is announced. You receive the total payout in this round, which is made up as follows: if you did not purchase any shares, you only receive the fixed amount of 16 points in this round. If you purchased one or two shares, your profits from this round consist of the dividends of the shares that you own, minus the price(s) you had to pay for these share(s), plus the fixed payment of 16 points.

This procedure is repeated 30 times.

After the end of the 30th round, you will be asked to fill in a questionnaire containing questions on your behavior during the experiment. After everyone has completed the questionnaire, you will be told your total earnings. You will not learn anything about the earnings of the other participants.

Please remain seated and wait quietly until we come to you and pay out your total earnings in cash. Please do not talk to the other participants. After your total earnings have been paid out, please leave the room quietly.

If there was anything you did not understand, please let us know by raising your hand. We will answer your questions on an individual basis.

Thank you for participating!