

Public opinion polls, voter turnout, and welfare:

An experimental study

Jens Großer Florida State University

Arthur Schram University of Amsterdam

Abstract

We experimentally study the impact of public opinion poll releases on voter turnout and welfare in a participation game. We find higher overall turnout rates when polls inform the electorate about the levels of support for the candidates, than when polls are prohibited. Distinguishing between allied and floating voters, our data show that this increase in turnout is entirely due to floating voters. When polls indicate equal levels of support for the candidates, turnout is high and welfare is low (compared to the situation without polls). In contrast, when polls reveal more unequal levels of support, turnout is lower with than without this information, while the effect of polls on welfare is non-negative. Finally, many of our results are well predicted by quantal response (logit) equilibrium.

Keywords: Public opinion polls, voter turnout, welfare, participation games, laboratory experiments

Jens Großer is an Assistant Professor at the Departments of Political Science and Economics, Florida State University and Experimental Social Science at Florida State (xs/fs), 531 Bellamy Hall, Tallahassee, FL 32306-2230 (jgrosser@fsu.edu). Arthur Schram is a Professor of Economics at the University of Amsterdam and the Center for Research in Experimental Economics and political Decision making (CREED), Roetersstraat 11, 1018 WB Amsterdam, the Netherlands (schram@uva.nl).

Financial support from the EU-TMR Research Network ENDEAR (FMRX-CT98-0238) is gratefully acknowledged. Tamar Kugler provided very useful input at early stages of this project. We would like to thank Jason Barabas, Ronald Bosman, Jordi Brandts, Thorsten Giertz, Jon Rogers, Joep Sonnemans, and Ramzi Suleinam for helpful comments. The research project this study is built on has benefited from comments by seminar participants at the University of Amsterdam in 1999; University of Cologne; Public Choice Society and ESA Meetings, San Antonio; European Public Choice Society Meeting, Paris in 2001; and ENDEAR Workshops, University of Amsterdam in 2001 and 2003, Tinbergen Institute Amsterdam in 2004; GEW Tagung, Cologne in 2005; 1st World Meeting of the Public Choice Society, Amsterdam; International ESA Meeting, Rome in 2007; University of Limerick 2008; and the University of Catania 2009.

In many countries, public opinion polls provide the electorate with information about voter preferences in upcoming elections. In contrast, about equally many countries prohibit the release of such information in a given period prior to Election Day.¹ Neither policy can claim a clear foundation in scientific research: “After at least 60 years of research, a rich literature has developed concerning the question ‘do polls influence behavior?’. Yet no conclusive or unambiguous answer to the question can be given, whether related to vote choice, turnout, or opinions on issues” (Irwin and Van Holsteyn 2000: 22). Consequently, policy makers have no basis to adequately evaluate the effects of public opinion poll releases.

This paper uses game theory and laboratory experiments in an attempt to overcome this gap in our knowledge. We study the effects of public opinion polls that provide the electorate with information about the aggregate levels of support for two candidates in majoritarian elections with costly turnout. Voters may use this information when deciding whether to vote or abstain. For example, if a poll indicates that a vast majority of the electorate supports either of the two candidates, some voters may assume that the outcome of the election is obvious with or without their vote and choose to abstain, as would be predicted by rational choice theory (e.g., Downs 1957). Other voters who support the strong candidate may decide to jump on the bandwagon and vote where they would otherwise have abstained. Both examples of responses to a poll release show how knowledge about public opinion may influence the decision to vote or abstain.

Of course, all kinds of information may affect the turnout decision. This has been studied both theoretically and empirically (see Diermeier and van Mieghem 2008; Gerber and Green 2000; Goeree and Großer 2007; Großer and Schram 2006; Lassen 2005; and the references

¹ For 2002, a survey of 66 countries shows that 36 have no embargo on poll releases and 30 ban publication in a period ranging from a day to a month before elections (“The Freedom to Publish Opinion Poll Results - Report on a Worldwide Update”, Foundation for Information/ESOMAR/WAPOR 2003). Examples of countries with bans are Canada, Mexico, and Switzerland.

therein). In these studies, being informed typically increases the probability of turning out. As mentioned above, we focus on one particular type of information: the level of support for each candidate, which directly determines the ‘level of disagreement’ within the electorate. Disagreement is high if both candidates have similar levels of support, and it is low if the vast majority of voters prefer the same candidate (Feddersen and Sandroni 2006).^{2,3}

Because polls provide information about the level of disagreement, voters may be able to make better judgments about the turnout decisions of other voters, than they would without polls. Most importantly for our study, polls give voters an indication about the expected closeness of

² To avoid unnecessary noise in the experiment, our polls provide perfect information about the levels of support. In practice, so-called ‘trial-heat polls’ become increasingly more informative about these levels in the course of a campaign (Erikson and Wlezien 1996). Though public opinion polls typically contain more information, the levels of support for the various candidates and the resulting expectations about the closeness of the election are undoubtedly crucial information in polls and heavily discussed in the media.

³ It is important to note that our paper focuses on the effects of poll releases *after* all voters have decided on which of the candidates to support. That is, voter preferences have already formed at the time of polling. Preference formation typically takes place (and may be affected by polls) at earlier stages in the campaign (e.g., Huckfeldt and Sprague 1995; Lau and Redlawsk 2006; Mutz 1998, 2006; Zaller 1992). Moreover, polls may also influence voters’ *candidate choices* by creating bandwagon or underdog effects (Simon 1954, Fleitas 1971; Ceci and Kahn 1982, Irwin and Van Holsteyn 2000). For game theoretical and experimental evidence on how polls affect candidate choice, see McKelvey and Ordeshook (1985) and Forsythe et al. 1993). In our experiment, voters cannot respond to a poll by changing to another candidate. This allows us to isolate the effect of information on the turnout decision without having to correct for changes in preferences.

elections (Duffy and Tavits 2008), if one takes into account the probabilities of voting for the two candidates (Palfrey and Rosenthal 1983; 1985). A voter may use this information when deciding whether or not to vote. In the rational turnout model of Downs (1957), for example, a person will vote if her expected benefits exceed her costs. Because expected benefits change if the perception of being pivotal changes—for example, as a consequence of information about the level of disagreement—polls may affect turnout in such models (e.g., Brown and Zech 1973; Gärtner 1976).

The positive relationship between closeness and voter turnout has been extensively tested and finds widespread support in empirical studies of voting (for comprehensive surveys of mainly observational studies, see Blais 2000 and Matsusaka and Palda 1993; and for a field experiment see Gerber and Green 2000). Blais (2000: 60) summarizes the findings as follows: “(...) The verdict is crystal clear (...): closeness has been found to increase turnout in 27 out of 32 different studies that have tested the relationship, in many different settings and with diverse methodologies. There is strong reason to believe that, as predicted by rational choice theory, more people vote when the election is close.”

There are two important elements that distinguish our study from observational studies. First, these studies often use actual votes cast as an *ex post* measure of closeness. In contrast, laboratory control allows us to use the ‘true’ level of disagreement as an *ex ante* measure of closeness, which is difficult to estimate in the field.⁴ Second, by controlling the level of disagreement we can design and compare elections that are identical in all aspects but one: whether or not polls provide the electorate with information about voter preferences. This allows us to directly measure how turnout depends on the content of information provided by public opinion polls. For example, polls may affect behavior differently if a high level of disagreement is revealed than if

⁴ Matsusaka and Palda (1993: 861) argue that: “The ideal [closeness] measure would be survey predictions from opinion polls taken the day before the election; (...)”

this level is low. Studies based on observational data typically do not allow for such direct measurement, because either a prediction has been published for any given election or it has not.⁵ In addition to the direct effect of poll releases on turnout, we also study their consequences for social welfare. For this purpose, we define social welfare as the aggregate net benefits of all voters, that is, aggregate benefits minus aggregate voting costs. This depends on the election outcome, which affects voters differently depending on the candidate they support and on their individual turnout decisions (Palfrey and Rosenthal 1983, 1985; Ledyard 1984). In our experiment we control the costs, benefits, and levels of support in elections and, therefore, can accurately measure social welfare for any configuration of voting decisions. This also allows us to measure welfare disaggregated for minorities and majorities.

A change in turnout—e.g., through information in polls—has two potential effects on social welfare. First, it can directly affect aggregate voting costs. For example, all else held constant, higher turnout decreases welfare because it increases these costs. Second, a change in turnout can indirectly affect welfare by altering the chances of the majority-preferred candidate (we will assume that aggregate benefits are larger if this candidate wins;⁶ cf. Campbell 1999 and Großer and Giertz 2006 for alternative assumptions). Whether this indirect effect on social welfare is positive or negative depends on *how* polls change the decisions to vote. For example, if responses to polls are independent of the voters' preferences, then higher turnout increases the majority's chances (which increases expected aggregate benefits). However, such independence is not ob-

⁵ In the words of Simon (1954: 246): “(...) we must carefully distinguish between: (a) what the outcome would have been *in the absence of a published prediction*; and (b) what the outcome actually was *after a prediction has been published*” (italics in the original).

⁶ This assumption holds, for example, if each voter's benefits from her preferred candidate winning are the same (and the same holds for losing). Since the aggregation of individual well-being is beyond the scope of our paper, we only consider this case.

vious. Polls may also change the relative turnout between two supporter groups in favor of the minority, decreasing the majority's chances due to stronger free riding incentives in this group (Goeree and Großer 2007; Taylor and Yildirim 2010). Whether the release of public opinion polls ultimately increases or decreases social welfare is an important theoretical and empirical question. Our model and controlled laboratory experiment allow us to answer this question by systematically separating and examining the various direct and indirect effects of polls on turnout and welfare.

In analyzing the consequences of public opinion polls, it is important to distinguish between *allied* and *floating* voters. To a large extent, uncertainty about the level of disagreement is caused by floating voters, who typically decide whom to support on a case-by-case basis shortly before or on Election Day.⁷ The importance of these voters for election outcomes has long been recognized (Converse 1962; Daudt 1961; Zaller 2004; and the references therein). In fact, they are a main reason why pollsters conduct public opinion polls in the first place. In contrast to floating voters, allied voters have stable preferences across elections. Importantly, allied and floating voters may respond differently to information in polls (e.g., the long-term attachment of allied voters to their group may override this short-term, election-specific information). Moreover, the interaction between the two different voter types may influence their behavior. Consequently, we systematically compare the effects of poll releases for electorates where allied and floating voters coexist to electorates with only floating voters.⁸

⁷ For example, in 2004 the pollster Populus categorized 35% of the UK electorate as floating voters (Times Online, 07.09.2004, "Boost for Kennedy as Blair and Howard slip").

⁸ Note that electorates with only allied voters would make polls meaningless in our setup, because the level of disagreement is known with and without polls providing this information.

Participation Games

The participation game of Palfrey and Rosenthal (1983, 1985; henceforth PR83 and PR85, respectively) provides a suitable framework for our purpose.⁹ This game allows us to hold constant the electorate size, voting costs, and benefits from election outcomes. We then systematically vary, one at a time, (i) the level of disagreement within the electorate and (ii) whether or not polls inform subjects about this level before elections. Moreover, we create voter alliances with candidates by (iii) keeping the preferred candidate of allied voters constant across elections. To study the effects of poll releases on turnout and welfare we combine the participation games with complete information (PR83) and incomplete information (PR85). In the former case, voters know others' preferences over the two candidates, with incomplete information they do not. The situation we model is where the preferences of some voters are private information unless pollsters publish them in the run-up to elections. A more formal description of the model is available on request from the authors.

In the participation game, each voter from an electorate of size E prefers (supports) one of two exogenous candidates, A and B . We denote the number of voters preferring each candidate (i.e., the levels of support) by N_A and N_B , respectively, where $N_A + N_B = E$. Voters simultaneously and independently decide on whether to turn out to vote at a cost, $c > 0$, or abstain (at no

⁹ To the best of our knowledge, we are the first to theoretically and experimentally investigate uncertainty about the level of disagreement and its resolution through poll releases in the participation game framework. For other theoretical studies that use this game to investigate the effect of polls see Goeree and Großer (2007) and Taylor and Yildirim (2010). For applications that investigate other types of information, see Diermeier and van Mieghem (2008) and Großer and Schram (2006). In a related experimental study, Levine and Palfrey (2007) systematically test the game's predictions with cost uncertainty for two different levels of disagreement under varying electorate sizes. They do not, however, study the effects of polls on turnout and welfare.

cost). The elections are decided by simple majority rule (with random tie-breaking). Each voter supporting the winning candidate receives the same benefit, which is independent of her turnout decision. This benefit is larger than the benefit received by each voter supporting the defeated candidate. Assuming that the stakes are the same for all voters, aggregate benefits are larger if the majority-preferred candidate wins. Whereas in PR83 the levels of support are common knowledge, the electorate only knows a common probability distribution of these levels in PR85.¹⁰ The equilibrium turnout probabilities of these games can be derived and, in turn, used to compute the candidates' winning probabilities and expected welfare. Specific predictions for our experimental parameters are given in the following sections.

To accommodate the possibility of allied voters, we modify the participation game such that

- (a) each candidate $i = A, B$ has a minimal level of support, $\underline{N}_i \geq 1$, implying a maximal level of support of $\bar{N}_i = E - \underline{N}_{-i}$, $i \neq -i$, and
- (b) there is a discrete probability distribution over all possible electoral compositions (N_i, N_{-i}) from the set $\{(\underline{N}_i, E - \underline{N}_i), (\underline{N}_i + 1, E - \underline{N}_i - 1), \dots, (E - \underline{N}_i, \underline{N}_i)\}$, with $prob(\cdot) > 0$ for each element in the set.

To allow for our distinction between allied and floating voters, a repeated setting is needed where the allied voters constitute the minimal levels of support and stay together for all elections

¹⁰ In PR85 there is also incomplete information about the voting costs of others in the electorate, which allows for Bayesian Nash equilibria in cutoff strategies. Because our focus is on the resolution of uncertainty about the level of disagreement, we avoided confronting subjects in the laboratory with additional sources of uncertainty and used constant costs and benefits across voters. It should be noted, however, that one interpretation of the quantal response equilibrium we will use to derive theoretical predictions is that it captures subjective perceptions of the costs and benefits of voting, creating uncertainty about these perceptions in the electorate.

without changing their preferences. In contrast, the preferences of floating voters are randomly drawn anew before each election. Note that this distinction does not depend on whether or not pollsters publish information.

Our paper contributes to a small literature studying the participation game experimentally (e.g., Bornstein 1992; Cason and Mui 2005; Duffy and Tavits 2008; Großer and Schram 2006; Levine and Palfrey 2007; Schram and Sonnemans 1996a, 1996b). In all these experiments, relatively high turnout is observed, albeit lower than in most general elections around the world. Except for a study with cost uncertainty (Levine and Palfrey 2007) the standard (Bayesian) Nash equilibrium concept finds little empirical support in the experimental literature on participation games. However, Cason and Mui (2005), Goeree and Holt (2005), and Levine and Palfrey (2007) show that quantal response equilibrium (QRE) can account for the data in many cases.¹¹ For example, contrary to Bayesian Nash equilibrium, QRE can predict substantial turnout in large elections (Levine and Palfrey 2007). This may be up to the order of 50%. Therefore, our theoretical analysis will focus on QRE.¹²

Experimental design

Procedures and treatments

The computerized¹³ experiment was run at the CREED laboratory of the University of Amsterdam. 288 undergraduate students participated in 12 sessions of 24 subjects. Each session lasted

¹¹ QRE is a solution concept for games developed by McKelvey and Palfrey (1995). This will be explained in the next section.

¹² A description of the (Bayesian) Nash equilibria for our games is available online at <http://myweb.fsu.edu/jgrosser/research.htm>. These equilibria are indeed poor predictors of the behavior observed in our experiment. This includes equilibria in pure strategies, which do not yield testable predictions about our treatment effects.

¹³ The experimental software was programmed using RatImage (Abbink and Sadrieh 1995).

about 2 hours (see the online appendix for the read-aloud instructions). Earnings in the experiment were expressed in tokens. At the end of a session, token earnings were transferred to cash at a rate of 4 tokens to one Dutch Guilder. Subjects earned an average of 56.01 Dutch Guilders ($\approx \text{€}25.42$).

In each session, the 24 subjects were randomly divided into two electorates of $E = 12$ voters. Each electorate consisted of two supporter groups $i = A, B$. There was no interaction of any kind between subjects in different electorates, and this was known to all of them. Given that we do not know the structure of the correlations across observations, we treat the electorate as the independent unit of observation. Hence, each session provides us with two independent observations. We employed a full 2×2 between subject treatment design with three sessions (six electorates) per cell. Our first treatment variable manipulated the *information* about the realized level of support for each candidate (i.e., N_A and N_B). This information was either given at the beginning of each round ('informed'; i.e., a poll is released) or not at all ('uninformed'; i.e., no poll is released). Our second treatment variable manipulated *voter alliances*. In one treatment ('floating') there were only floating voters, while in the other ('mixed') there were 3 allied voters for each candidate (i.e., $N_A = N_B = 3$) plus 6 floating voters. Each subject was either an allied or floating voter throughout the experiment and knew her type right from the start. Note that floating voters were always reallocated within the same electorate.

Uncertainty about the level of disagreement

Information and voter alliance were both varied between subjects. To create uncertainty, we varied the level of disagreement within subjects. In any given round, each candidate was supported by a minimum of $\underline{N}_i = 3$ voters and a maximum of 9. Any integer number of supporters $N_A, N_B \in \{3, 4, \dots, 9\}$ was possible, where $N_A + N_B = E$. This means that the level of disagreement was highest when each of the two supporter groups consisted of 6 voters and lowest when a mi-

nority of 3 voters faced a majority of 9. We will represent the level of disagreement by the size of the minority (note that this level indeed increases with this size).

A randomized allocation procedure was used to determine for each subject in the electorate which candidate she supported. In this way, the levels of support for the two candidates were determined. Specifically, we used the following two steps:

- Step 1: 3 subjects were allocated to each candidate. Each subject in the electorate had an equal chance of being chosen to support either candidate.
- Step 2: The remaining 6 subjects were independently and randomly allocated, with equal probability for each candidate.

This procedure was known to all subjects. The way it was applied is different for our ‘floating’ and ‘mixed’ treatments. In ‘floating’ sessions, both steps were performed at the beginning of each round and, importantly, subjects did not know at which step they were allocated to a candidate. In ‘mixed’ sessions, step 1 determined the 6 subjects to take the role of allied voters and their allocation to the candidates. This step was performed only once, at the beginning of the first round, while step 2 reallocated the 6 floating voters at the beginning of each round. Notice that step 2 produces a binomial distribution of support levels with $p = 0.5$, where electoral composition (6,6) occurs with probability .3125, (5,7) and (7,5) each with .2344, (4,8) and (8,4) each with .0938, and (3,9) and (9,3) each with .0156.

Each session consisted of 100 decision rounds.¹⁴ The electoral composition was varied in a random, but predetermined manner across rounds. 33 rounds used the composition (6,6), 23 used (5,7), 22 used (7,5), 9 used (4,8), 9 used (8,4), 2 used (3,9), and 2 used (9,3). Whether subjects knew the actual levels of support for each candidate when making their decisions, depends on the information treatment. In the ‘informed’ sessions, the actual levels of support for the ‘own’ and ‘other’ candidate were announced at the beginning of each round, while in the ‘uninformed’ ses-

¹⁴ Due to a computer crash, one session had to be stopped after 94 rounds.

sions these were never released.¹⁵ Hence, subjects in ‘uninformed’ basically faced the same decision problem in each round.¹⁶

Payoff parameters

In each round, each subject supporting the winning candidate received 4 tokens and each supporter of the defeated candidate received 1 token. As the cost of turnout was 1 token (independent of a subject’s type), negative payoffs were avoided. Table 1 summarizes treatments and parameters.

Equilibrium predictions

For the parameters of our experiment, we can numerically derive (Bayesian) Nash equilibria and quantal response equilibria (QRE; McKelvey and Palfrey 1995). QRE is a generalization of Nash equilibrium that allows for errors in decision making (i.e., boundedly rational behavior). For example, erroneous perceptions of costs and benefits due to cognitive limitations or mistakes made while choosing an alternative could both cause the type of error that is assumed in QRE. Because decisions are noisy, players do not find ‘best’ responses to others’ decisions, as assumed in the standard Nash equilibrium. Instead, they make *stochastic* ‘better’ responses. These are systematic, however, in the sense that decisions that yield higher expected payoffs are made more frequently than less lucrative decisions. The degree to which errors enter the decisions is repre-

¹⁵ Groups were labeled ‘own’ and ‘other’ to avoid floating voters associating with either group.

¹⁶ However, ‘uninformed’ allied and floating voters have different information about the level of disagreement. An allied voter knows that there are at least 3 (at most 9) voters for a candidate. A floating voter, on the other hand, knows that there are at least 4 voters supporting her own candidate and at most 8 supporting the other. As a consequence, in the ‘mixed’ treatment allied voters have an expected level of support for the own candidate of 6, whereas for floating voters Bayesian updating yields an expected own level of support of 6.5. When there are only floating voters, this expectation is 6.25.

sented by a non-negative noise parameter μ . In our specification, if $\mu = 0$, decision makers make no errors and QRE reduces to the standard Nash equilibrium. At the other extreme, as $\mu \rightarrow \infty$ decision making becomes completely random. In the participation game this would mean that each voter turns out with probability 0.5.

QRE has become a common tool in analyzing games and laboratory experiments (for examples related to participation games and other binary choice games see Cason and Mui 2005; Goeree and Holt 2005; Levine and Palfrey 2007). To analyze our game, we follow the tradition in this literature and use the logit specification of QRE (a.k.a. the ‘logit equilibrium’; see McKelvey and Palfrey 1995)¹⁷ and consider only quasi-symmetric equilibria, where all voters facing the same decision problem vote with equal probability.¹⁸

Figure 1 shows the equilibrium turnout probabilities per treatment as a function of the noise level μ . For the informed treatments (IF and IM), these probabilities are aggregated as weighted averages of the probabilities for the distinct levels of support used in our experiment. Because in these treatments pre-election information is exactly the same for allied and floating voters, we show one aggregated turnout probability. For the uninformed treatments (UF and UM) slight differences in the expected levels of disagreement may occur between the two voter types (see footnote 16). Hence, we show separate turnout probabilities for allied and floating voters.

In the Bayesian Nash equilibrium ($\mu = 0$) for the uninformed treatments, average expected turnout is substantially higher when there are allied and floating voters (53% in UM) than when

¹⁷ The logit specification we use assumes that errors follow a double exponential distribution (see Goeree and Holt 2005).

¹⁸ Voters only observe aggregate turnout. This imperfect monitoring after elections may give rise to ‘private sequential equilibria’ (Mailath, Matthews, and Sekiguchi 2002). For our setup, however, it is infeasible to compute such repeated game equilibria.

all voters are floating (10% in UF). The latter turnout is close to the Nash equilibrium for the informed treatments (11%). The high turnout in UM is entirely due to the very high turnout probability of allied voters (94%). In contrast, the 12% voting probability for floating voters in UM is similar to that in the other treatments.

When noise is introduced, the equilibrium predictions across treatments quickly converge. From a noise level of approximately $\mu = 0.3$ upwards, equilibrium turnout probabilities are virtually identical across voter types and treatments (recall that in the informed treatments these are aggregated probabilities; disaggregated turnout probabilities differ across the distinct levels of support). These logit equilibria are increasing in μ and give qualitatively different predictions than (Bayesian) Nash equilibrium. In particular, except for allied voters in UM higher turnout is predicted than in the (Bayesian) Nash equilibrium. A noise level of 0.3 lies below the levels reported for previous participation game experiments. Goeree and Holt (2005) estimate $\mu = 0.8$ in early rounds and $\mu = 0.4$ in later rounds for the Schram and Sonnemans (1996a) data based on game parameters similar to ours. Hence, we will use the turnout probabilities for $\mu \geq 0.3$ for our logit predictions.

In the following, we focus on these logit equilibria for $\mu \geq 0.3$ to derive qualitative *theoretical results* (TR) relevant to our experiment. These equilibria specify turnout probabilities (disaggregated per level of support in informed treatments). We use them to compute equilibrium predictions for winning probabilities and expected welfare. We do so in aggregate; per level of support; and per level of disagreement. Moreover, by comparing the equilibria for uninformed and informed electorates, we can establish the predicted effects of poll releases. Confronting our TR with our experimental results will allow us to evaluate the comparative statics predicted by logit equilibrium. In presenting our TR we do not distinguish between voter alliance treatments, because they provide the same qualitative predictions.

Theoretical results for our logit equilibria ($\mu \geq 0.3$):

TR1 (aggregate turnout): *Polls do not affect aggregate expected turnout.*

TR2 (turnout per level of support): *Polls increase turnout probabilities for intermediate levels of support and decreases it for relatively small and large levels.*

Argument. The online appendix shows that when voters know the levels of support, turnout probabilities are highest for intermediate levels. When voters are uninformed, turnout probabilities are the same for all levels of support and lie in between the distinct levels predicted for informed electorates. Hence, disclosure of the levels of support in polls will increase (decrease) turnout if these levels are (not) at an intermediate level.

TR3 (turnout per level of disagreement): *Polls increase expected turnout for high levels of disagreement and decrease it for low levels.*

Argument. This follows directly from TR2 because high levels of disagreement imply intermediate levels of support.

TR4 (winning probabilities): *Polls somewhat diminish the advantage of the majority-preferred candidate for all levels of disagreement.*

Argument. The majority-preferred candidate always has better chances in the election than the minority opponent. The probability of winning strongly increases in the level of support in uninformed electorates. It also increases, but to a lesser extent, in the level of support in informed electorates. In other words, having more supporters increases the probability of winning, but the chances are highest when it is not known that there are many supporters.

TR5 (welfare): *Polls have little effect on expected social welfare, but minority welfare slightly increases and majority welfare slightly decreases when polls are released.*

Argument. This is a direct consequence of TR2 and TR4: for example, if polls reveal a low level of disagreement, turnout in both the minority and majority is suppressed (which in itself has a

positive effect on social welfare). If they reveal a high level of disagreement, turnout increases in both supporter groups (with a negative effect on social welfare). In both cases the changes in turnout probabilities increase the relative probability of the minority-preferred candidate winning, which negatively affects social welfare. For revealed low levels of disagreement, the positive cost effect is larger than this negative effect and the net effect on social welfare is small, but positive. When polls reveal high levels of disagreement, both effects are negative. Hence, polls sometimes increase and sometimes decrease social welfare. In aggregate, the effect of polls on expected social welfare is negligible. Supporters of the majority candidate suffer welfare losses however, whereas supporters of the minority candidate slightly benefit.

Experimental results

The presentation and analysis of our experimental results is organized as follows. We start with the observed effects of polls on turnout rates in aggregate, per level of support, and per level of disagreement. Then, we investigate turnout rates for each voter type. Finally, we examine the influence of polls on the chances of minorities and majorities and on welfare. Many of our statistical tests are based on nonparametric statistics. For the reasons mentioned above, these tests are conducted at the electorate level (qualitative conclusions are based on one-tailed tests). In addition, random effects probit estimations are used to analyze turnout behavior at the individual level. To provide a benchmark for our data, we sometimes present the specific equilibrium predictions for the cases concerned.¹⁹ Laboratory findings are summarized as *experimental results* (ER).

Polls and turnout

Aggregate turnout: Figure 2 shows turnout rates averaged over blocks of 20 rounds each.

¹⁹ These are the predictions underlying figure 1 and TR1-TR5. The online appendix gives a table providing a full overview of our equilibria per treatment, using $\mu = 0.4$ and 0.8 for the logit equilibria and including $\mu = 0$ for the (Bayesian) Nash equilibria.

We observe higher rates when electorates are informed about the levels of support than when they are not. This holds for all blocks of rounds: the turnout rate is always higher in IF than in UF and always higher in IM than in UM. Though all treatments start at similar levels, a difference of approximately 10%-points exists between informed and uninformed electorates from the second block onward for both comparisons.

ER1: *Polls increase aggregate turnout levels by 22-28%.*

Support. Wilcoxon-Mann-Whitney tests reject the null hypothesis of no difference in average aggregate turnout in favor of higher rates for informed electorates at the 5% significance level for the IF-UF comparison and at the 10% level for IM-UM. The increase in turnout is approximately 28% when all voters float and 22% when there are allied voters.

Note that ER1 rejects our TR1 that polls have no effect on aggregate turnout. This does not mean that the logit equilibria for $\mu = 0.4$ and 0.8 are far off the mark, however. For all treatments, the predicted turnout levels lie between 30% and 31% ($\mu = 0.4$) or between 38% and 39% ($\mu = 0.8$), close to the turnout rates in our data. Hence, logit equilibrium can explain aggregate turnout reasonably well, but it fails to predict the boost in turnout caused by poll releases. ER1 may be considered to be one of our main results. We will discuss possible explanations in our subsection on voter alliances and turnout, after analyzing the determinants of individual turnout.

Turnout per levels of support and disagreement: Figure 3 shows predicted and observed turnout rates per level of support (left panel) and level of disagreement (right panel). For the uninformed treatments (UF and UM), there is no reason to expect turnout to vary across conditions, because subjects do not know the actual levels of support. This is confirmed in the left panel. For the informed treatments (IF and IM), figure 3 (right panel) shows that observed turnout increases in the level of disagreement. The left panel reveals that this holds for both the minority and majority (with one exception, observed turnout decreases when moving away from the support level

of 6 in either direction). This panel also shows that observed turnout is always lower in the minority than in the opposing majority (compare turnout for 3 supporters to that for 9 supporters, for example). Finally, observe the very high turnout rates of 49% in IF and 58% in IM for the highest disagreement level of 6; the former can only be justified by the logit equilibrium for unreasonably high noise levels and the latter cannot be explained by any of our equilibria (see figure 1; we will return to this observation in the conclusions).

A first, general result on the observed effects of polls on turnout per level of support can be derived from figure 3. This is that polls increase turnout for intermediate levels of support and decrease turnout for small and large levels of support. For a more detailed analysis of these effects at the individual level, we use random effects probit estimations. Specifically, we estimate a panel model explaining the individual decision to turn out to vote or abstain for each of our four treatments separately.²⁰ As independent variables, we consider the various levels of support a voter encounters (which can only affect behavior in our informed treatments); a time trend; the voter type; the previous turnout decision; and whether or not the previous decision was pivotal *ex post*. Note that if subjects use quasi-symmetric equilibrium strategies, their turnout decisions should not be affected by past events. We distinguish between allied and floating voters in our

²⁰ Following a suggestion by an anonymous referee, we also estimated panel models using pooled data for the pairwise comparisons UF-IF, UM-IM, IF-IM, and UF-UM, allowing us to test the treatment effects directly using interaction terms (e.g., whether voters decide differently in elections with polls than in those without polls). The results confirm our analysis from comparing separated data. Because we investigate treatment effects mainly using nonparametric statistics (due to the relatively low number of independent observations in laboratory studies), we do not report the results of the extended panel models here. These are available online at <http://myweb.fsu.edu/jgrosser/research.htm>.

mixed treatments. Then, for the treatments with only floating voters (denoted by superscript F) the panel model is given by²¹

$$D_{i,t}^F = \beta_0^F + \beta_1^F \frac{t}{100} + \beta_2^F D_{i,t-1} + \beta_3^F \Delta LS_{i,t}^< + \beta_4^F \Delta LS_{i,t}^> + \beta_5^F PIV_{i,t-1}^0 + \beta_6^F PIV_{i,t-1}^1 + \varepsilon_{i,t} + \gamma_i \quad (1)$$

and for the treatments with a mix of allied and floating voters (superscript M) by

$$D_{i,t}^M = \beta_0^M + \beta_1^M \frac{t}{100} + \beta_2^M D_{i,t-1}^{fl} + \beta_3^M D_{i,t-1}^{al} + \beta_4^M AL_i + \beta_5^M \Delta LS_{i,t}^{<,fl} + \beta_6^M \Delta LS_{i,t}^{>,fl} + \beta_7^M \Delta LS_{i,t}^{<,al} + \beta_8^M \Delta LS_{i,t}^{>,al} + \beta_9^M PIV_{i,t-1}^{0,fl} + \beta_{10}^M PIV_{i,t-1}^{1,fl} + \beta_{11}^M PIV_{i,t-1}^{0,al} + \beta_{12}^M PIV_{i,t-1}^{1,al} + \varepsilon_{i,t} + \gamma_i, \quad (2)$$

where i denotes the voter, and t denotes the round. $D_{i,t}^F$ ($D_{i,t}^M$) is a dummy variable equal to 1 if i voted in t , and 0 otherwise. AL in equation (2) is a dummy variable equal to 1 if i is an allied voter, and 0 otherwise. Superscripts fl and al indicate whether the independent variable concerns an allied or floating voter (we omit superscripts fl in equation (1) where there are only floating voters). $\Delta LS_{i,t}^<$ and $\Delta LS_{i,t}^>$ measure the absolute difference in the levels of support between the two groups, i.e. $\Delta LS_{i,t}^< = \Delta LS_{i,t}^> = |N_A - N_B| \in \{0,2,4,6\}$, where superscripts ‘<’ (‘>’) indicate that i is in the minority (majority), with $\Delta LS_{i,t}^< = 0$ ($\Delta LS_{i,t}^> = 0$) if i is in the majority (minority).²²

$PIV_{i,t-1}^0$ is a dummy variable equal to 1 if i was pivotal *ex post* in the previous round and she abstained (denoted by superscript ‘0’), and 0 otherwise. Similarly, $PIV_{i,t-1}^1$ is a dummy variable equal to 1 if i was pivotal *ex post* in the previous round and she voted (denoted by superscript

²¹ To be more precise, the right hand sides of equations (1) and (2) specify the linear (random) utility underlying the probit estimations.

²² Note that $\Delta LS_{i,t}^<$ and $\Delta LS_{i,t}^>$ are inversely related to the level of disagreement as measured by the size of the minority (*SoM*). Specifically, $SoM = 6 - (\Delta LS_{i,t}^<)/2 = 6 - (\Delta LS_{i,t}^>)/2$.

‘1’), and 0 otherwise. $\varepsilon_{i,t}$ and γ_i are error terms, where the latter is a random effect used to correct for the panel structure in our data. Table 2 presents the maximum likelihood estimates of the coefficients of our models.

Table 2 helps us distinguish between various effects of poll releases on voter turnout. First, consider electorates with only floating voters. As already shown in figure 3, turnout in IF increases in the level of disagreement (measured by $-\Delta LS_{i,t}^{<,fl}$ and $-\Delta LS_{i,t}^{>,fl}$). The coefficients -0.20 and -0.13 are both negative and highly significant, and indicate that the minority responds more strongly to differences in support than the majority (see the left panel in figure 3). As expected, the levels of disagreement do not affect turnout in UF. Interestingly, poll releases suppress other influences on the decision to vote: previous turnout ($D_{i,t-1}^{fl} = 1$) and having previously been pivotal ($PIV_{i,t-1}^{0,fl} = 1$ or $PIV_{i,t-1}^{1,fl} = 1$) increase the probability of voting in UF, but no such effects are observed in IF. It seems that subjects use information from previous elections to determine their choice, if there is no current information available. Very similar results hold when there are allied voters in the electorate (see the next subsection for a comparison between both voter types). The main difference is that now having been pivotal in the previous round significantly increases turnout of informed voters in three out of four cases, i.e., information from previous rounds is still used even if poll information is available. Finally, we find a downward trend across rounds in all treatments except IM.

The result that information about the level of disagreement in polls (partially) redirects the attention away from retrospective information is interesting. Subjects seem to use clues about whether or not their decision will matter. If no polls are available, information about having been pivotal in previous rounds gives them a more general clue about the importance of their vote in upcoming elections. As a result, all of the independent variables on ex post pivotalness have a significant effect on turning out. If polls are available, the use of retrospective information de-

depends on the treatment. When there are only floating voters, this clue has little or no added value to the information in polls, because the composition of supporter groups changes across elections; our estimated coefficients are indeed small and insignificant. However, when there are also allied voters, ex post pivotalness is, to some extent, informative about the balance of turnout between the two groups of allied supporters. If, for example, the allied voters of one candidate generally vote less often, this decreases the probability of being pivotal in any election. Ex post pivotalness is a signal that this situation may also be relevant for upcoming elections and may therefore affect the decision to vote, in particular that of floating voters. Indeed, the estimated coefficients for the independent variables representing lagged pivotalness of floating voters are positive and significant in our mixed treatment, even when polls are released. In addition, allied voters may react to (e.g., reciprocate) perceived turnout behavior of co-allies. Although these decisions cannot be observed directly (because only aggregate election outcomes are made public), some information about relative turnout levels for the allied groups can be deduced from previous rounds. This may explain why one of the variables of lagged pivotalness of allied voters is positive and significant in IM.

Finally, the effect of the lagged dependent variable also follows intuition. When polls are prohibited, a voter basically faces the same problem in every election. Assume that she starts with a certain belief about being pivotal and updates this belief using the observed pivotalness in elections. After a few elections, the weight of new observations becomes low and it becomes more likely that her turnout decision in two subsequent elections will be the same. In some sense, voting becomes ‘habitual’ and based on predispositional considerations. As a consequence, the coefficients of the lagged dependent variable are positive and highly significant in the treatments without polls. In contrast, polls provide information that is specific to the upcoming elections, which can be used to make decisions that better reflect their current interests.²³ Then, correlation

²³ We would like to thank an anonymous referee for pointing out this aspect of polls.

of decisions across elections is much less likely. Indeed, we observe no significant coefficients for the lagged dependent variable in elections with poll releases.

The following result summarizes our findings.²⁴

ER2: *Polls cause turnout to increase as the level of disagreement increases, but this effect is stronger for the minority than for the majority. Turnout probabilities are higher in the majority than in the opposing minority. Without polls, both having voted and having been pivotal in the previous election increases turnout.*

We can compare this result to TR2 and TR3 and the equilibrium turnout probabilities shown in figure 3. The observed inverse U-shape in the relationship between the level of support and turnout for informed treatments supports the logit equilibrium predictions in TR2, but our finding that turnout rates are not higher in the minority than in the opposing majority rejects the predictions shown in figure 3. A similar result is reported in Großer and Giertz (2006) and Klor and Winter (2008) who also consider cases where voting costs are constant. In contrast, Levine and Palfrey's (2007) experiment with uncertain voting costs supports the predicted higher turnout probabilities for the minority.²⁵ As for the level of disagreement, ER2 supports the predictions that turnout increases in this level and that polls increase turnout for high levels and decrease it for low levels (TR3).

²⁴ ER2 is also supported by nonparametric tests using the electorate as the unit of observation.

²⁵ Though this suggests that the cost setup may be important, one explanation for the discrepancy may be the following: subjects in Levine and Palfrey's study faced the same level of disagreement in fifty consecutive rounds. This stable learning environment is not present in our environment (nor in the other two mentioned), because our purpose is to study frequent changes in preferences across elections. Consequently, our experimental design does not provide the best suitable environment to study learning towards equilibrium. In the conclusions, we will discuss an alternative explanation.

To conclude, this subsection has shown that poll releases have strong effects on voter turnout. Polls often redirect voters' attention away from past events towards the current level of disagreement. This information yields important regularities in observed voter turnout; most strikingly, turnout increases in the level of disagreement. Hence, as in the field (Blais 2000; Matsusaka and Palda 1993), closeness matters for informed electorates in the laboratory.²⁶

Voter alliances and turnout

Aggregate turnout: Figure 2 not only shows an increase in aggregate turnout rates due to polls, but also that the presence of allied voters boosts voting. Turnout differences between electorates with and without allied voters start out relatively small, but they increase to an average of 6.7%-points and 6.0%-points in the last three blocks of 20 rounds for the UM-UF and IM-IF comparisons, respectively. However, only the latter difference is statistically significant.

ER3: *When there are polls, electorates with allied voters have higher aggregate turnout levels in later rounds.*

Support. Across all rounds, Wilcoxon-Mann-Whitney tests cannot reject the null hypothesis of no difference in turnout rates for the UM-UF and IM-IF comparisons at the 10%-level. Considering blocks 3-5 only, rates are significantly higher in IM than in IF (5% level) but the difference UM-UF is not significant.

Turnout per voter type: Figure 4 gives observed turnout rates per voter type. It shows that polls increase turnout of floating voters by 29% (when there are only floating voters) and by 55% (when there are allied and floating voters), but leave turnout by allied voters unaffected. One way to summarize the data underlying figure 4 is that turnout is 40% or somewhat higher for allied voters (independent of polls) and for informed floating voters. Uninformed floating voters, however, vote at a lower rate of roughly 30%.

²⁶ As expected, the rate of ex post pivotalness in our experiment is increasing in the level of disagreement.

ER4: *The increase in aggregate turnout levels through polls is entirely due to floating voters.*

Support. A Wilcoxon-Mann-Whitney test cannot reject the null hypothesis of no difference in average aggregate turnout between allied voters in the IM-UM comparison, but rejects it between floating voters in favor of higher rates in IM than UM at the 5% level. Recall from ER1 that when there are no allied voters, floating voters turn out significantly more in IF than in UF.

To understand ER4, we start with a rational choice approach based on the equilibrium predictions described in the previous section. Note that with polls, all voters have the same information before deciding whether or not to vote and there are no differences in equilibrium turnout probabilities. Without polls, however, floating voters are slightly better informed about the true levels of support for each candidate than allied voters. Recall from footnote 16 that this is because each floating voter knows which candidate she herself supports, in addition to the commonly known preferences of the allied voters. Specifically, a perfectly rational allied voter uses Bayesian updating to determine beliefs about the levels of support and expects equal support levels of 6 for both candidates. On the other hand, a rational floating voter expects a larger level of support for the candidate she prefers (6.25 vs. 5.75 in UF and 6.5 vs. 5.5 in UM). This yields (slightly) higher free-riding incentives—that is, incentives to abstain—for floating voters than for allied voters. The question is then how much this affects predicted turnout probabilities. This depends on the equilibrium concept used. As shown in figure 1, in the Bayesian Nash equilibrium ($\mu = 0$) without polls, floating voters vote at a much lower rate than allied voters (12.3% vs. 93.4% in UM, respectively). In the logit equilibrium ($\mu \geq 0.3$), however, differences in turnout between both voter types disappear and the small differences in expectations about the levels of support can no longer explain turnout differences.²⁷ Note that the observed turnout difference

²⁷ For much larger electorates than the ones considered here the differences would also disappear in the Bayesian Nash equilibrium.

between both types (14.0%-points) is nowhere near the 81.1%-points predicted by the Bayesian Nash equilibrium. In combination with our earlier observation that the logit equilibrium describes our data much better than the Bayesian Nash equilibrium, this leads us to conclude that our equilibria alone cannot explain why floating voters respond more strongly to polls than allied voters do.

Our favorite alternative explanation is that floating voters are subject to a ‘false consensus effect’, which describes the tendency of people to believe that others have tastes that are similar to their own (e.g., Mullen et al. 1985; Ross, Greene, and House 1977).²⁸ In the context of elections, voters tend to overestimate the level of support for the own candidate relative to the levels for other candidates (e.g., Brown 1982; Morwitz and Pluzinsky 1996).²⁹ In our uninformed treatments, this would mean that floating voters tend to believe that relatively many of the other floating voters support the same candidate. A high estimate of equal minded voters, in the absence of polls, increases free-riding incentives and may lead to lower turnout by floating voters than when polls provide accurate information about the levels of support. A similar false consensus effect would not be observed for allied voters, because even without polls they are perfectly informed about the fact that exactly three allied voters support each candidate.

²⁸ As pointed out by Dawes (1990), the consensus effect need not necessarily be ‘false’ but could be the outcome of using one’s own taste when Bayesian updating beliefs about others. For a theoretical application to elections, see Goeree and Großer (2007).

²⁹ For example, Brown (1982) conducted a survey among 179 psychology students in the run-up to the 1980 U.S. presidential elections. While Anderson supporters estimated 30.4% of their classmates to support Anderson, the estimates for this candidate by Carter and Reagan supporters were only 24.4% and 22.0%, respectively. Similarly, Carter (Reagan) supporters estimated that 38.3% (44.5%) favor their own candidate versus 34.1% and 33.6% estimated by Anderson and Reagan supporters (35.4% and 37.9% by Anderson and Carter supporters).

A next question is whether the turnout decision of floating voters is affected by the presence of allied voters. Figure 4 shows that without polls, their turnout rates are relatively low and very similar (31.3% in UF and 30.4% in UM). When polls are released, floating voters turn out somewhat more when there are allied voters than in their absence (47% vs. 40%), but this difference is statistically insignificant at the electorate level. All in all, the effect of allied voters on the turnout decisions of floating voters appears to be limited.³⁰

Finally, we can directly compare the two voter types when they coexist within the same electorate. Doing so, we find that allied voters turn out 44% more than floating voters when there are no polls (UM) and 8% less when there are (IM). Though these differences are not statistically significant at the electorate level (Wilcoxon signed ranks tests, 10% level), they are at the individual level, after correcting for other factors. This follows from the estimated coefficients of the dummy variable AL in table 2, which is 0.70 in UM and -0.23 in IM (in both cases significantly different from zero at the 1%-level). Whereas the lower turnout by floating voters in UM can be explained by the false consensus effect described above, the lower turnout by allied voters in IM still needs to be explained. Because both voter types respond very similarly to the information in polls (in the last column of table 2 the coefficients of $-\Delta LS_{i,t}^{<,fl}$ and $-\Delta LS_{i,t}^{<,al}$ are very similar, as are those of $-\Delta LS_{i,t}^{>,fl}$ to $-\Delta LS_{i,t}^{>,al}$), polls themselves do not seem to be the source of this difference. Our conjecture is that allied voters turnout less in IM than floating voters do because—as discussed above—they respond to perceived turnout decisions of their co-allies in earlier rounds.

³⁰ As discussed above, however, our random effects probit estimations do reveal some significant differences between floating voters in IM and IF (see table 2), specifically in how they respond to having been pivotal ex post in previous elections.

Winning rates

Logit equilibrium for $\mu \geq 0.3$ predicts that turnout probabilities will be higher in the minority than in the opposing majority when polls are released, but that this difference is insufficient to give the minority a higher probability of winning the election. In contrast, observed turnout rates are not higher in the minority (see figure 3), from which it follows directly that the majority wins more often. Figure 5 compares observed majority winning rates with and without polls. It also gives theoretical predictions, including the social welfare maximizing winning probabilities.³¹ For comparison, we add the case where both candidates have an equal level of support of 6. For this case, we define the winning rate of the ‘majority’ as 50%.

Figure 5 shows that majorities win more than 50% of the time and (with one exception) the chance that they win is increasing in the level of support. For the uninformed treatments, this is a direct consequence of equal average turnout rates across the levels of support (since subjects cannot respond to what they don’t know). Moreover, only for support levels of 7 do the winning rates depend significantly on whether or not polls are released (see both panels in figure 5). When informed, a majority of 7 voters wins significantly more often than when there are no polls (Wilcoxon-Mann-Whitney tests, 5% level for IF and 10% level for IM). These observations give:

ER5: *Majorities win the elections more often than the opposing minorities. Polls further increase this difference in close elections.*

Support. This follows from the discussion above.

Comparing this result to TR4 and to the theoretical predictions depicted in figure 5, we observe that the logit equilibrium accurately predicts the comparative statics.

³¹ In case of unequal levels of support, social welfare is maximized when one majority-voter votes and everybody else abstains. When support levels are equal social welfare is maximized when everybody abstains.

Welfare effects

We consider the effects of polls on welfare (aggregate net benefits) at both the level of support ('group welfare') and for the whole electorate ('social welfare'). The latter may depend on the level of disagreement. Using equilibrium turnout probabilities to compute expected payoffs, we can determine group and social welfare in equilibrium. From these, we can predict welfare effects, calculated as welfare with polls minus welfare without polls, and compare these predictions to our data.

Group welfare effects: Figure 6 shows predicted and observed group welfare effects across levels of support. It reveals a stark contrast between theory and our data. The logit equilibria predict that polls benefit the minority and harm the majority (i.e., the lines start above zero and end up below zero). What we observe is quite different: for majorities, welfare is most often increased by poll releases. Whereas 5 of the 6 majorities shown in figure 6 are better off with polls, this holds for only 2 of the 6 minorities.

Social welfare effects: Table 3 shows the effects polls have on social welfare, distinguishing between the cases with only floating voters (columns 2 to 6) and with allied and floating voters (columns 7 to 11). Rows present the logit predictions for $\mu = 0.4$ and 0.8, and our data.

The predicted effect of polls on social welfare is always negative (i.e., in equilibrium polls are harmful for the electorate). The magnitude of this effect is monotonically decreasing in the level of disagreement for $\mu = 0.4$ (i.e., polls are more harmful in lopsided elections), but smaller and more or less independent of this level for $\mu = 0.8$. In stark contrast, our data only show a negative social welfare effect of polls for the highest level of disagreement. Because this effect counts relatively heavily in the weighted average, it yields a small negative effect on the weighted average of social welfare across all elections (-0.08 when there are only floating voters and -0.02 when there are allied and floating voters).

We summarize our findings in the following result:

ER6: *Polls mostly increase majority welfare and decrease or barely affect minority welfare.*

Without allied voters, polls decrease social welfare when the level of disagreement is highest.

Support. Wilcoxon-Mann-Whitney tests reject the null hypothesis of no difference between group welfares in informed and uninformed electorates in favor of lower welfare with polls for minorities of 5 for both the IF-UF and IM-UM comparisons (1% and 10% significance level) and for support level 6 for the IF-UF comparison (10% level). Equal group welfare is rejected in favor of higher welfare with polls for majorities of 7 (5% level for IF-UF; 10% level for IM-UM) and for majorities of 8 for IM-UM (5% level). For all other minorities and majorities, the null hypothesis cannot be rejected (10% level). Moreover, a Wilcoxon-Mann-Whitney test rejects the null hypothesis of no difference between aggregate social welfares in informed and uninformed electorates in favor of lower welfare with polls for a level of disagreement of 6 for IF-UF (10% significance level), but the same test cannot reject the null for all other levels and comparisons (10% level).

The logit equilibrium predicts the observed effects of polls on expected social welfare well (see TR5). However, it does not capture the observed increase in majority welfare caused by polls.

Though the opposite is predicted by logit equilibrium, our data suggest that the majority-preferred candidate should favor poll releases and the minority-preferred candidate should oppose them.

Conclusions

Our aim in this paper has been to shed light on the effects of public opinion polls on voter turnout and welfare in two-candidate majoritarian elections. In a novel setup, we have theoretically and experimentally compared participation games with and without poll releases in the run-up to elections. Polls inform the electorate about the level of support for each candidate. A major advantage of this simple setup is that it can accommodate a variety of more specific applications in which researchers and political engineers may be interested (e.g., various electoral systems).

Even without extensions, however, our study allows us to draw many conclusions about the consequences of poll releases.

Our laboratory subjects strongly react to information in polls: knowing the level of support for each candidate increases aggregate voter turnout, but distinct elections are affected differently. The turnout boost is most striking in *ex ante* dead heats. For these elections, polls yield severe welfare losses because turnout is well above the socially optimal level. In contrast, polls decrease turnout in elections where a landslide victory is expected. In cases with unequal support levels, they generally cause majority welfare to increase (because they win more often) and either cause minority welfare to decrease or leave it unaffected. When distinguishing between allied and floating voters, we find that the increase in aggregate turnout through polls is caused entirely by floating voters. As a plausible explanation for this result we conjectured that there may be a false consensus effect. Based on their own taste, uninformed floating voters may overestimate the number of other ‘floaters’ supporting the same candidate. This would then increase the incentive to free-ride, causing a drop in turnout by this voter type.

We used quantal response (logit) equilibria to derive our theoretical predictions. These are generally supported by our laboratory data. In particular, most turnout levels and patterns across levels of support and disagreement are explained very well. However, in addition to our finding that the increase in aggregate turnout through polls is only due to floating voters, the following two observations remain at odds with logit equilibrium: (i) majority voters turn out at higher rates than the opposing minority voters in informed electorates (and as a consequence, polls generally increase majority but not minority welfare); (ii) turnout in informed electorates is 50% or higher when both candidates have the same levels of support. There are various possible reasons why logit equilibrium fails here. We favor an explanation based on ‘group think’, where some voters base their turnout decision on (group) rule-utilitarian goals rather than individualistic goals (Coate and Conlin 2004; Feddersen and Sandroni 2006). For example, in Coate and

Conlin's study, a voter is assumed to maximize the welfare of the supporter group to which she belongs, as opposed to maximizing her individual payoff. Using this as a point of departure, one can derive Nash equilibrium turnout probabilities based on group welfare maximization for the informed electorates in our study. These predictions match quite well with the comparative statistics we observe across support and disagreement levels and can improve on the two observations left unexplained by logit predictions. They correctly predict the higher turnout rates in the majority and an average turnout of 50% when both candidates have the same level of support. A further improvement of the fit between predictions and data may be obtained by combining the logit and group think models. To study this, however, an alternative experimental design would be much more suitable and we therefore leave this for future research.

The positive relationship between closeness and voter turnout that we observe in our experiment when public opinion polls are released is an important indication of the external validity of our setup. This increases our confidence that our results are also meaningful outside of the laboratory. Nevertheless, there are limitations to the generalizability of our results. For example, we have ignored the possibility of strategic poll responses that may bias information and yield less representative polls than in our setup. Moreover, we have considered only majority rule. As pointed out by Irwin and van Holsteyn (2000), results may be different for a situation with proportional representation (a framework studied in an experimental participation game by Schram and Sonnemans 1996a). Because this electoral system typically yields multiparty environments, public opinion polls may affect welfare not only through voter turnout but also through strategic voting (i.e., polls may act as a coordination device for coalition formations; e.g., Forsyth et al. 1993). Future research may also relax other assumptions in our study (e.g., by using larger electorates with a larger variety of levels of disagreement). Such extensions would increase the contribution of our setup and add even more to our understanding of the consequences of poll re-

leases. The theoretical and experimental results presented in this paper are intended to be a first important step in this direction.

Appendix (See our online appendix at <http://myweb.fsu.edu/jgrosser/research.htm>)

References

- Abbink, Klaus, and Abdolkarim Sadrieh. 1995. "RaImage: Research Assistance Toolbox for Computer-Aided Human Behavior Experiments." University of Bonn. Discussion Paper B-325.
- Blais, André. 2000. *To Vote or Not to Vote? The Merits and Limits of Rational Choice Theory*. Pittsburgh: University of Pittsburgh Press.
- Bornstein, Gary. 1992. "The Free-Rider Problem in Intergroup Conflicts Over Step-Level and Continuous Public Goods." *Journal of Personality and Social Psychology* 62 (4): 597-606.
- Brown, Clifford E. 1982. "A False Consensus Bias in 1980 Presidential Preferences." *Journal of Social Psychology* 118 (1): 137-38.
- Brown, Kenneth M., and Charles E. Zech. 1973. "Welfare Effects of Announcing Election Forecasts." *Public Choice* 19 (1): 117-23.
- Campbell, Colin M. 1999. "Large Electorates and Decisive Minorities." *Journal of Political Economy* 107 (6): 1199-217.
- Cason, Timothy N., and Vai-Lam Mui. 2005. "Uncertainty and Resistance to Reform in Laboratory Participation Games." *European Journal of Political Economy* 21: 708-37.
- Ceci, Stephen J., and Edward L. Kain. 1982. "Jumping on the Bandwagon With the Underdog: The Impact of Attitude Polls on Polling Behavior." *Public Opinion Quarterly* 46 (2): 228-42.
- Coate, Stephen and Michael Conlin. 2004. "A Group Rule-Utilitarian Approach to Voter Turnout: Theory and Evidence." *American Economic Review* 94 (5): 1476-1504.
- Converse, Philip E. 1962. "Information Flow and the Stability of Partisan Attitudes." *Public Opinion Quarterly* 26 (4): 578-99.
- Daudt, Hand. 1961. *Floating Voters and the Floating Vote. A Critical Analysis of American and English Elections Studies*. Leiden: H.E. Stenfert Kroese.
- Dawes, Robyn M. 1990. "The Potential Nonfalsity of the False Consensus Bias." In R.M. Hogarth (ed.). *Insights In Decision Making: A Tribute to Hillel J. Einhorn*. Chicago: University of Chicago Press: 179-99.
- Diermeier, Daniel, and Jan A. Van Mieghem. 2008. "Coordination and Turnout in Large Elections." *Mathematical and Computer Modelling* 48 (9-10): 1478-96.
- Downs, Anthony. 1957. *An Economic Theory of Democracy*. New York: Harper and Row.
- Duffy, John, and Margit Tavits. 2008. "Beliefs and Voting Decisions: A Test of the Pivotal Voter Model." *American Journal of Political Science* 52 (3): 603-18.
- Erikson, Robert S., and Christopher Wlezien. 1996. "Of Time and Presidential Election Forecasts." *PS: Political Science and Politics* 29 (1): 37-9.
- Feddersen, Timothy J., and Alvaro Sandroni. 2006. "A Theory of Participation in Elections." *American Economic Review* 96 (4): 1271-82.

- Fleitas, Daniel W. 1971. "Bandwagon and Underdog Effects in Minimal-Information Elections." *American Political Science Review* 65 (2): 434-38.
- Forsythe, Robert, Roger B. Myerson, Thomas A. Rietz, and Robert J. Weber. 1993. "An Experiment on Coordination in Multi-Candidate Elections: The Importance of Polls and Elections Histories." *Social Choice and Welfare* 10 (3): 223-47.
- Gärtner, Manfred. 1976. "Endogenous Bandwagon and Underdog Effects in a Rational Choice Model." *Public Choice* 25 (1): 83-9.
- Goeree, Jacob K., and Jens Großer. 2007. "Welfare Reducing Polls." *Economic Theory* 31 (1): 51-68.
- Goeree, Jacob K., and Charles A. Holt. 2005. "An Explanation of Anomalous Behavior in Models of Political Participation." *American Political Science Review* 99 (2): 201-13.
- Großer, Jens, and Thorsten Giertz. 2006. "Candidates, Voters, and Endogenous Group Formation: An Experimental Study." Princeton University. Working Paper.
- Großer, Jens, and Arthur Schram. 2006. "Neighborhood Information Exchange and Voter Participation: An Experimental Study." *American Political Science Review* 100 (2): 235-48.
- Huckfeldt, R. Robert, and John Sprague. 1995. *Citizens, Politics and Social Communication: Information and Influence in an Election Campaign*. Cambridge, Massachusetts: Cambridge University Press.
- Irwin, Galen A., and Joop J.M. Van Holsteyn. 2000. "Bandwagons, Underdogs, the Titanic, and the Red Cross: The Influence of Public Opinion Polls on Voters." Presented at the 18th World Congress of the International Political Science Association, Quebec, Canada.
- Klor, Esteban F., and Eyal Winter. 2008. "On Public Opinion Polls and Voters' Turnout." Hebrew University of Jerusalem. Working Paper.
- Lang, Kurt, and Gladys Engel Lang. 1984. "The Impact of Polls on Public Opinion." *Annals of the American Academy* 472 (March): 129-42.
- Lassen, David Dryer. 2005. "The Effect of Information on Voter Turnout: Evidence from a Natural Field Experiment." *American Journal of Political Science* 49 (1): 103-18.
- Lau, Richard R., and David P. Redlawsk. *How Voters Decide: Information Processing During Election Campaigns*. Cambridge: Cambridge University Press.
- Ledyard, John O. 1984. "The Pure Theory of Large Two-Candidate Elections." *Public Choice* 44 (1): 7-41.
- Levine, David K., and Thomas R. Palfrey. 2007. "The Paradox of Voter Participation? A Laboratory Study." *American Political Science Review* 101 (1): 143-58.
- Malaith, George J., Steven A. Matthews, and Tadashi Sekiguchi. 2002. "Private Strategies in Finitely Repeated Games with Imperfect Public Monitoring." *Contributions to Theoretical Economics* 2 (1): Article 2.

- Matsusaka, John G., and Filip Palda. 1993. „The Downsian Voter Meets the Ecological Fallacy.” *Public Choice* 77 (4): 855-78.
- McKelvey, Richard D., and Peter C. Ordeshook. 1985. “Elections with Limited Information: A Fulfilled Expectations Model Using Contemporaneous Poll and Endorsement Data as Information Sources.” *Journal of Economic Theory* 36 (1): 55-85.
- McKelvey, Richard D., and Thomas R. Palfrey. 1995. “Quantal Response Equilibria for Normal Form Games.” *Games and Economic Behavior* 10 (1): 6-38.
- Morwitz, Vicki G., and Carol Pluzinski. 1996. “Do Polls Reflect Opinions or Do Opinions Reflect Polls?” *Journal of Consumer Research* 23 (1): 53-67.
- Mullen, Brian, Jennifer L. Atkins, Debbie S. Champion, Cecelia Edwards, Dana Hardy, John E. Story and Mary Vanderklok. 1985. “The False Consensus Effect: A Meta-Analysis of 115 Hypothesis tests.” *Journal of Experimental Psychology* 21 (3): 262-83.
- Mutz, Diana C. 1998. *Impersonal Influence: How Perceptions of Mass Collectives Affect Political Attitudes*. Cambridge: Cambridge University Press.
- Mutz, Diana C. 2006. *Hearing the Other Side: Deliberative Versus Participatory Democracy*. Cambridge: Cambridge University Press.
- Palfrey, Thomas R., and Howard Rosenthal. 1983. “A Strategic Calculus of Voting.” *Public Choice* 41 (1): 7-53.
- Palfrey, Thomas R., and Howard Rosenthal. 1985. “Voter Participation and Strategic Uncertainty.” *American Political Science Review* 79 (1): 62-78.
- Ross, Lee, David Greene, and Pamela House. 1977 “The False Consensus Effect: An Egocentric Bias in Social Perception and Attribution Processes.” *Journal of Experimental Psychology* 13 (3): 279-301.
- Schram, Arthur, and Joep Sonnemans. 1996a. “Voter Turnout as a Participation Game: An Experimental Investigation.” *International Journal of Game Theory* 25 (3): 385-406.
- Schram, Arthur, and Joep Sonnemans. 1996b. “Why People Vote: Experimental Evidence.” *Journal of Economic Psychology* 17 (4): 417-42.
- Simon, Herbert A. 1954. “Bandwagon and Underdog Effects and the Possibility of Elections Predictions.” *Public Opinion Quarterly* 18 (3): 245-53.
- Taylor, Curtis, and Huseyin Yildirim. 2010. “Public Information and Electoral Bias.” *Games and Economic Behavior* 68 (1): 353-75.
- Zaller, John. 1992. *The Nature and Origins of Mass Opinions*. Cambridge, Massachusetts: Cambridge University Press.
- Zaller, John. 2004. “Floating Voters in U.S. Presidential Elections, 1948-2000.” In W.E. Saris and P.M. Sniderman (eds.). *Studies in Public Opinion: Attitudes, Nonattitudes, Measurement Error, and Change*. Princeton, NJ: Princeton University Press: 166-212.

Table 1: Summary of treatments and parameters

Treatment	Acronym	# Floating voters	# Allied voters	Poll release
Uninformed Floating	UF	12	0	No
Uninformed Mixed	UM	6	6	No
Informed Floating	IF	12	0	Yes
Informed Mixed	IM	6	6	Yes

Notes: All treatments had 100 rounds and electorates of 12 voters, with a minimum (maximum) of 3 (9) in each group. A victory (defeat) paid 4 (1) to each voter in the group and the individual costs of voting were equal to 1. We have observations from 6 independent electorates per treatment.

Table 2: Random effects probit estimations of voter turnout

Constant and independent variables	Coefficients			
	UF	IF	UM	IM
Constant	-0.93 (16.18)***	0.07 (1.00)	-1.03 (14.77)***	0.29 (4.80)***
$t / 100$	-0.40 (6.23)***	-0.20 (3.44)***	-0.18 (2.76)***	-0.05 (0.82)
$D_{i,t-1}^{fl}$	0.48 (9.20)***	0.00 (0.04)	0.87 (11.47)***	-0.02 (0.30)
$D_{i,t-1}^{al}$	-	-	0.39 (5.27)***	-0.00 (0.03)
AL_i	-	-	0.70 (7.14)***	-0.23 (3.00)***
$\Delta LS_{i,t}^{<,fl}$	0.00 (0.32)	-0.20 (13.78)***	-0.02 (0.82)	-0.26 (9.63)***
$\Delta LS_{i,t}^{>,fl}$	-0.00 (0.07)	-0.13 (11.45)***	0.01 (0.52)	-0.19 (12.56)***
$\Delta LS_{i,t}^{<,al}$	-	-	0.02 (0.88)	-0.24 (13.26)***
$\Delta LS_{i,t}^{>,al}$	-	-	0.02 (1.17)	-0.14 (8.17)***
$PIV_{i,t-1}^{0,fl}$	0.27 (5.87)***	0.01 (0.20)	0.14 (2.10)**	0.17 (2.35)**
$PIV_{i,t-1}^{1,fl}$	0.19 (3.20)***	0.00 (0.06)	0.15 (1.76)*	0.17 (2.41)**
$PIV_{i,t-1}^{0,al}$	-	-	0.12 (1.71)*	-0.04 (0.57)
$PIV_{i,t-1}^{1,al}$	-	-	0.37 (4.89)***	0.15 (2.16)**


Notes: The dependent variable is the voters' binary choice between voting (= 1) and abstaining (= 0). The independent variables in column 1 are defined in the main text. Absolute z -values are given in parentheses. * (**; ***) indicates significance at the 10% (5%; 1%) level. Results on the random effects estimates are available on request.

Table 3: Effects of polls on social welfare

Level of disagreement	Only floating voters					Allied and floating voters				
	3	4	5	6	Weighted average	3	4	5	6	Weighted average
logit $\mu = .4$	-0.35	-0.19	-0.12	-0.05	-0.12	-0.33	-0.18	-0.11	-0.05	-0.11
logit $\mu = .8$	-0.05	-0.03	-0.05	-0.05	-0.05	-0.05	-0.03	-0.05	-0.05	-0.04
Observed	0.72	0.00	0.01	-0.36	-0.08	0.14	0.45	0.07	-0.45	-0.02


Notes: Entries in the cell represent the predicted or observed effects of poll releases on social welfare for the situation depicted by the column. The average effect is weighted by the relative frequency with which the levels of disagreement occur in our experiment.

Figure 1: Equilibrium turnout per treatment




Notes: The lines give the logit equilibria for varying μ from 0 to 2 and the discrete cases 10, 100, and 1000 (which show up as a ‘jump’ upwards, with turnout probabilities close to 0.5). Where distinct lines can no longer be visually separated, one line represents all.

Figure 2: Aggregate turnout rates



Notes: UF=uninformed electorates with only floating voters;
IF=informed electorates with only floating voters; UM=uninformed electorates with allied and floating voters; IM=informed electorates with allied and floating voters.

Figure 3: Turnout rates



Notes: The bars in both figures show turnout rates for informed electorates per level of support (left panel) and per level of disagreement (right panel). In the left panel, the lines with markers give turnout levels for uninformed electorates. The remaining lines show for informed electorates our logit predictions for noise levels $\mu = 0.4$ and 0.8 . The level of disagreement is measured by the size of the minority.

Figure 4: Turnout rates for allied and floating voters

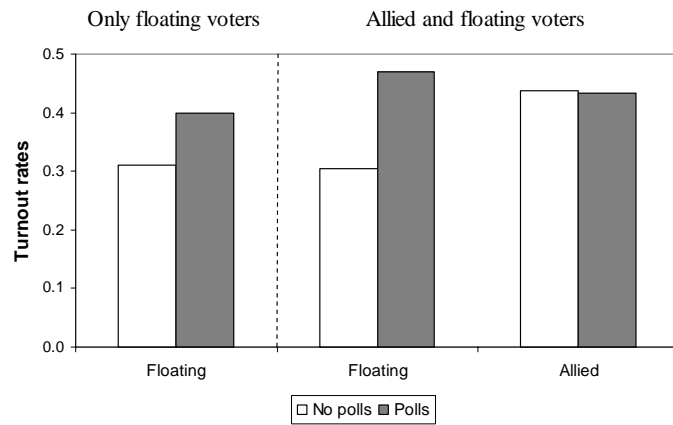




Figure 5: Winning rates



Notes: Bars show the observed frequency of majority wins for the level of support shown. Lines show theoretical predictions based on logit equilibrium and social welfare maximization. For presentational reasons, the logit predictions for $\mu = 0.4$ are not shown. These lie parallel to but below the logit predictions for $\mu = 0.8$.

Figure 6: Effects of polls on group welfare



Notes: The figure shows the effects of polls on group welfare, separately for only floating voters (left panel) and for allied and floating voters (right panel). Effects are measured as group welfare with polls – group welfare without polls. Bars show observed values of this difference across support levels and lines show the logit predictions for μ equal to 0.4 and 0.8.