Int J Game Theory (2007) 35:379–394
DOI 10.1007/s00182-006-0050-5

ORIGINAL PAPER

The welfare effects of public opinion polls

Esteban F. Klor · Eyal Winter

Revised: 15 May 2006 / Published online: 1 November 2006

Abstract This paper presents an experimental study of the effects of polls on voters' welfare. The analysis shows that polls have a different effect on closely divided and lopsided divided electorates. The data show that in closely divided electorates (and only for these electorates) the provision of information on the voters' distribution of preferences significantly raises the participation of subjects supporting the slightly larger team relative to the smaller team. This causes a substantial increase on the frequency of electoral victories of the larger team. As a consequence, we observe a steep decrease in the welfare of the members of the smaller team because they vote more often and yet they loose the elections more frequently. Polls are detrimental to aggregate welfare in closely divided electorates because the decrease in the payoffs of the minority is stronger than the increase in the payoffs of the majority. In lopsided divided electorates polls don't have a significant different effect on the voters' turnout conditional on their team size. We do observe an increase on the frequency of electoral victories of the larger team after the provision of information, but this is in part due to smaller teams' members voting less frequently and saving the participation costs. As a consequence, while polls have a negative effect on the relative payoffs of the minority for these electorates as well, they have a positive effect on total welfare.

E. F. Klor (✉)
Department of Economics, Centre for Economic Policy Research (CEPR),
The Hebrew University of Jerusalem, Jerusalem, Israel
e-mail: eklor@mscc.huji.ac.il
URL: http://www.economics.huji.ac.il/facultye/klor/klor.htm

E. Winter
Department of Economics and The Center for the Study of Rationality,
The Hebrew University of Jerusalem, Jerusalem, Israel
e-mail: mseyal@mscc.huji.ac.il

Keywords Voter turnout · Public Opinion polls · Experimental economics

JEL Classification C72 · C92 · D72 · H41

1 Introduction

A large number of democratic countries ban the publication of opinion polls for a fixed period of time before elections.[1] The main objective of the restriction, according to the arguments put forward by its supporters, is to guarantee a fair electoral process. That is, supporters of the ban on the publication of polls contend that polls unfairly influence people and can thus be manipulated in favor of a given candidate. This clearly creates an uneven playing field that undermines the democratic process. According to this line of reasoning, imposing restrictions on the publication of polls isn't qualitatively different from existing regulations on political advertising and campaign finance. On the other side of the debate, individuals that oppose these restrictions claim that there isn't clear evidence of any systematic effect of polls on the electorate. Moreover, they regard any restriction on the publication of polls as an unjustifiable and unconstitutional abridgement of freedom of speech, expression, and the press.

Clearly, some of the issues raised in this lively debate are not quantifiable and, therefore, don't directly allow for a cost-benefit analysis. Some of the arguments, however, are quantifiable and easier to approximate via a formal theoretical model that delivers testable predictions. Do polls affect the voters' preferences and behavior? Is that influence strong enough to significantly affect electoral outcomes? And finally, is the publication of public opinion polls desirable from a purely instrumental point of view?

This paper presents an experimental analysis of the welfare effects of the provision of information on the electorate's distribution of preferences. The experiment, based on the seminal theoretical contribution of Palfrey and Rosenthal (1983), compares the subjects' participation decisions in an election when they know the exact distribution of preferences of the electorate to their decisions when they only know their own preferences. Our experiment is especially designed to identify the effects of polls on the voters' behavior. A careful analysis of those effects allows us to assess the resulting changes in aggregate welfare.

The current paper is closely related to recent theoretical contributions by Goeree and Großer (2006) and Taylor and Yildirim (2005) on the desirability of public opinion polls. These two papers conclude that polls are unambiguously detrimental to aggregate welfare. Intuitively, in a private-values model with costly-voting aggregate welfare is maximized when only one individual that supports the alternative supported by a majority of the population votes. This outcome minimizes voting costs and maximizes aggregate payoffs. According

[1] According to the Foundation for Information/ESOMAR a ban on the publication of opinion polls exists in 30 out of 66 countries surveyed in their study published in 2003. Nowadays, a lively debate is being conducted in several countries like Canada, France, Ireland, The Philippines and Russia.

to the equilibrium strategies, however, the provision of information on the electorate's preferences doesn't induce a decrease in turnout. In particular, polls cause individuals supporting the alternative supported by a majority of the population to participate less frequently because they free ride on the voting of other individuals supporting the same alternative. At the same time, polls stimulate the participation of individuals supporting the alternative supported by a minority of the population to offset the advantage of the other alternative. As a consequence, there isn't a significant decrease in turnout and its associated voting costs, but there is a significant increase in the likelihood that the electorate chooses the "wrong" alternative; that is, the alternative that doesn't maximize aggregate payoffs.

The experimental results show stark differences between closely divided electorates and lopsided divided electorates. In closely divided electorates (and only for these electorates) the provision of information significantly raises the participation of subjects supporting the slightly larger team relative to the smaller team—we refer to this behavior as the bandwagon effect of polls.[2] This behavior contradicts the qualitative predictions of the unique quasi-symmetric Nash equilibrium of the baseline theoretical model underlying the experiment. Moreover, the observed bandwagon effect of information provision is qualitatively opposed to the equilibrium behavior predicted by the models of information revelation through public opinion polls mentioned above.

In closely divided electorates the provision of information has two opposite effects on aggregate welfare. On the one hand, polls cause an increase in the participation of subjects from both teams. On the other, due to the bandwagon effect the frequency of victories of the slightly larger team is significantly higher than without the provision of information. The data show that the provision of information has a negative effect on aggregate payoffs in electorates that are closely divided. Moreover, polls cause a significant redistribution of payoffs across teams of different sizes. In particular, we observe an important decrease in the relative payoff of the minority versus the majority. Intuitively, even though the provision of information causes an increase in the participation of voters in the minority, their frequency of electoral victories decreases because of the high participation rate of voters' in the majority. This brings the minority's share of the overall payoffs to levels significantly lower than the minority's share of the electorate.

When preferences are lopsided in favor of one alternative the revelation of information causes some free riding in the behavior of subjects supporting the larger team, but subjects supporting the smaller team decrease their frequency of electoral participation with similar proportions. We do observe an increase in the frequencies of electoral victories of the larger teams after the provision of information, but this is in part due to smaller teams voting less frequently and saving the participation costs. These two effects have a positive impact on

[2] This behavior is extensively analyzed experimentally, empirically and theoretically by Klor and Winter (2006). While not explicitly studied, the bandwagon effect also appears in Cason and Mui (2005).

aggregate welfare. When decomposed across teams, however, we do observe that polls cause a decrease in the share of the payoffs received by the smaller teams.

Summarizing, our results show that the impact of polls on aggregate welfare depends on the distribution of preferences of the electorate. Polls have a negative impact on aggregate welfare when the electorate is closely divided and a positive impact on aggregate welfare for the rest of the electorates.[3] Decomposing the impact of polls on the payoffs of subjects that belong to the larger team relative to the payoffs of subjects that belong to the smaller team reveals that polls have a significant negative effect on the minorities' share of the payoffs for every distribution of preferences. It is noteworthy that this important disfranchisement of the minorities has been overlooked in the heated debate on the desirability of public opinion polls.[4]

The paper proceeds as follows. In the next section we describe the theoretical framework behind the experiment and show its predictions for the chosen parameter values. Section 3 presents a detailed description of the experiment especially designed to test the theoretical model. The experimental results appear in Sect. 4. The last section concludes.

2 Theoretical model and predictions

The theoretical framework we consider is the same as the one used in Klor and Winter (2006). The framework is based on the seminal contribution of Palfrey and Rosenthal (1983). Accordingly, $n \geq 3$ risk neutral individuals have to decide between two alternatives via simple plurality rule; that is, the alternative with the greater number of votes is chosen. In the event of a tie each alternative is selected with equal probability. The chosen alternative applies to all the individuals.

Each individual has preferences over the two alternatives. Let B denote the utility difference to an individual between the event that her favored alternative is elected and the event that the other alternative wins the election. Each individual has to decide whether to vote or abstain.[5] Let us denote by s_i the strategy of individual i (say $s_i = 1$ when individual i votes and $s_i = 0$ otherwise). All individuals make their strategy choices simultaneously. There is a positive

[3] In an independent study, Großer et al. (2005) examine the welfare implications of endogenous voter participation using a different experimental design that includes floating voters. They find that polls don't have a significant effect on aggregate electoral welfare. Their test, however, doesn't differentiate between closely divided electorates and the rest.

[4] Respect for the minorities is an important guiding principle when evaluating electoral systems. Casella et al. (2005), for example, show that a system of storable votes is highly desirable in particular because it substantially helps minorities.

[5] In the present framework voting against one's preferred alternative is strictly dominated by not voting. Therefore, we rule out this possibility and, whenever we say that an individual votes, we imply that she is voting in support of her preferred alternative.

cost $C > 0$ associated with the act of voting. B and C are common knowledge and identical to all the individuals. We assume that $B > 2C$.

In this setup, a rational individual votes if and only if

$$B \times P(1, s_{j \neq i}) - C \geq B \times P(0, s_{j \neq i}),$$

where P denotes the probability that individual i's preferred alternative is chosen and $s_{j \neq i}$ is a profile that describes the strategy of all the individuals excluding individual i.

Clearly, a rational individual participates in the election only if, given the other individuals' strategies, her participation affects the probability that her preferred alternative is chosen. In other words, an individual may turn out to vote only when she is pivotal.

We analyze the game above under two different frameworks regarding the individuals' information about the distribution of preferences. The first scenario focuses on a symmetric private value model of voting. Accordingly, each voter knows the alternative that she favors and that the probability that any other individual prefers any given alternative is the same for both alternatives. The individuals' probability distributions are stochastically independent. This game has a unique Bayesian Nash equilibrium (BNE). In this equilibrium all the individuals vote for their preferred alternative with the same probability; this probability is strictly between zero and one. The symmetric equilibrium strategies are a direct consequence of the symmetric common prior over the individuals' distribution of preferences.

In the second scenario the number of voters favoring each alternative is commonly known. This is exactly the framework analyzed by Palfrey and Rosenthal (1983). This complete information game has multiple Nash equilibria. The solution concept that generates unique predictions for the game is that of totally quasi-symmetric mixed strategy Nash equilibrium (QSNE). According to this equilibrium concept all the individuals supporting the same alternative use the same strategy. Moreover, this strategy involves voting with a probability strictly between zero and one. Note that individuals supporting different alternatives are not necessarily mixing with the same probability.

For the purposes of our experimental study we focus on electorates of seven individuals and set $B = 10$ and $C = 4$. The unique BNE for these parameter values is for each subject to vote with probability 0.0807. We choose an odd number of participants in each electorate to rule out equilibrium in pure strategies (except for the case where all the participants share the same preferences). Symmetric equilibria when the distribution of preferences is commonly known do not exist for this configuration of the parameters either. In fact, with seven players and two alternatives there exists a unique totally mixed quasi-symmetric Nash equilibrium for each possible distribution of preferences.

Table 1 presents the equilibrium voting probabilities for the individuals as a function of the size of their team. The table differentiates between the two different frameworks regarding the individuals' information about the distribution

Table 1 Predictions from the unique BNE and the unique QSNE

Distribution of preferences	3 vs. 4	2 vs. 5	1 vs. 6	0 vs. 7
Bayesian Nash equilibrium				
Probability that a supporter of the large team votes			0.0807	
Probability that a supporter of the small team votes			0.0807	
Probability that the large team wins the election	0.53263	0.59745	0.66102	0.72256
Quasi-symmetric Nash equilibrium				
Probability that a supporter of the large team votes	0.0873	0.070805	0.13988	0.036508
Probability that a supporter of the small team votes	0.1229	0.24001	0.86012	
Probability that the large team wins the election	0.49067	0.44286	0.45349	0.614604

of preferences. The table also shows the resulting probability of victories of the different teams.

As already noted, all the individuals participate in the election with the same probability according to the unique symmetric BNE of the symmetric private value model of voting. On the contrary, according to the unique QSNE of the game when the distribution of preferences is commonly known, individuals in the minority vote with higher probability than individuals in the majority. Moreover, the difference in the probability of voting between individuals in the two teams increases as the electorate becomes more lopsided divided. This non-intuitive result is a direct consequence of the mixed strategies equilibrium's requirement that, in deciding how to randomize between voting or abstaining, individuals in one team select their choice probability so as to make individuals in the other team indifferent between voting and abstaining. This implies that individuals have to expect that with high enough probability the number of votes in support for each team would be equal, or differ by only one vote. To satisfy that requirement individuals supporting the large team should vote with lower probability than individuals supporting the small team.

The second striking fact presented in Table 1 is the significant difference between the two frameworks analyzed regarding the probability that the large team wins the election. When the distribution of preferences is not known the probability that the large team wins the election is always higher than half, and it increases monotonically with the difference in the number of supporters of the two teams. On the contrary, when the distribution of preferences is commonly known, the large team it's not always more likely to win the election. This is a direct consequence of the differences in the probability of voting across teams. In fact, due to free riding of individuals in the large team and an underdog effect in the behavior of individuals in the small team, the small team is more likely to win the election for every distribution of preferences but the most lopsided one of 7 versus 0. This fact, already mentioned by Palfrey and Rosenthal (1983), is the main cause behind the inefficient effect of polls theoretically described in Goeree and Großer (2006) and Taylor and Yildirim (2005). That is, polls tend to increase turnout and the likelihood that the minority wins the elections.

3 Experimental design

This section briefly describes our experimental design. A fuller account appears in Klor and Winter (2006). The experiment was run at the RatioLab—The Center for Rationality and Interactive Decision Theory at The Hebrew University of Jerusalem. The 84 subjects in this experiment were recruited from the pool of undergraduate and graduate students from The Hebrew University and had no previous experience in experiments related to voters' participation.

In each session 21 subjects participated as voters. The experiments were conducted via computers. Before the experiment started an experimental administrator read the instructions aloud. We also asked several hypothetical questions at the end of the instructions to check subjects' comprehension of the procedure (the instructions and the questionnaire are available from the authors upon request). The experiment began after all subjects had solved all questions successfully. The experiment lasted for about ninety minutes. Each subject received 80 tokens as a participation fee and subsequent earnings according to the payoffs specified in the experiment. Average earnings were equal to 244 tokens. We converted each token to NIS 0.25 and paid the subjects in cash in private at the end of the session.[6] Throughout the experiment we ensured anonymity and effectively isolated each subject in a cubicle to minimize any interpersonal influence that could stimulate uniformity of behavior. Communication among subjects was not allowed throughout the session.

Each experimental session entailed 20 independents rounds. In each round we randomly divided 21 subjects into three electorates of seven participants each. At the beginning of each round an equal probability rule randomly assigned each subject to one of two teams: Green or Blue. A subject earns ten tokens if the team she prefers is selected by majority voting in an election. Voting entails a cost of four tokens.

The sequence of events is as follows. In the first stage of each round each subject knows only her preferred color. She decides whether to vote or abstain. After all the participants make their decisions we proceed to the second stage of the round. In this stage subjects are told the electorate's distribution of preferences. Note that subjects don't receive any information on the subjects' participation decisions in the round's first stage. Subjects have to decide again whether or not to vote. After all the subjects choose an action, they learn the selected teams of the first and second elections, their corresponding payoffs for the round, and their cumulative payoffs – no information is provided on the number of subjects that voted for a given team. Ties are always broken by an equal probability rule.

At the end of each round subjects are randomly rematched between electorates and each subject's preferred color is again randomly decided. We employ a random and anonymous reassignment procedure to avoid repeated game effects and to minimize subject's learning.

[6] That is, subjects on average earned NIS 61 for roughly 90 min of their time. The hourly minimum wage in Israel is slightly below NIS 20. The current exchange rate is NIS 4.55 per US dollar.

According to our experimental design every subject casts two votes in each round. This particular setup allows us to use each electorate as its own control for the tests on the welfare effects of polls, thus avoiding the additional noise caused by the pairing of groups. Since we can keep the subjects in the laboratory for at most 90 min, and we use relative large groups of seven individuals each, we are limited in the number of rounds repetitions that we manage to run in each experimental session. This limits the number of available observations. Consequently, it is extremely important for us that all the observations are as reliable as possible. We believe that this is better achieved by having two votes in each round, rather than by comparing teams that vote without information on the distribution of preferences to teams that vote knowing the exact distribution of preferences.

4 Results

This section presents the main results of the paper. To clarify the exposition we divide this section into two subsections. The first subsection, which draws heavily on Klor and Winter (2006), presents the main effects that the provision of information has on subjects' participation decisions. The analysis differentiates between closely divided and lopsided divided electorates. The second subsection analyses the impact that the provision of information on the subjects' distribution of preferences has on the subjects' payoffs.

4.1 The effect of information on subjects' participation

Figure 1 depicts the average turnout rate before the provision of information and the average turnout rate after information is revealed, as a function of the size of the teams. The figure presents the average turnout separately for each different distribution of subjects' preferences.

The figure clearly shows that for closely divided electorates revealing information on subjects' preferences causes an important increase on the participation of all the subjects: Subjects belonging to teams of size three and four vote more often after learning the distribution of preferences. The effect, however, is stronger for subjects that belong to the slightly larger team. The turnout rate of subjects that belong to a team with four supporters is more than twenty percent higher than the turnout rate of subjects that belong to a team of three supporters. Evaluated at the averages, an individual that belongs to the larger team is more than 10 percent more likely to vote than otherwise.[7]

As noted above, according to the theoretical predictions presented in Table 1, members of the small team should vote with a higher probability than the members of the large team in order to offset the electoral advantage of the large

[7] This is the result reached by Klor and Winter (2006) using a regression analysis. The estimation controls also for the subjects' participation decision in the first stage of each round, a round trend, and a subject specific random effect.

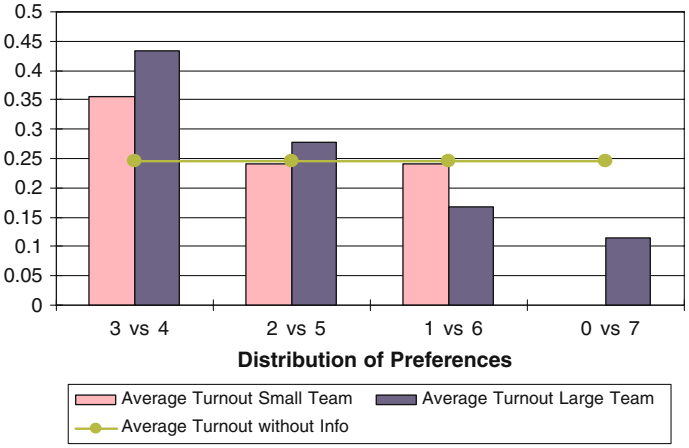

Fig. 1 Average turnout rate by distribution of preferences and team size

team. Moreover, the provision of information should induce a decrease in the turnout rate of the large team because of free riding of its members. Therefore, the observed bandwagon effect of polls contradicts the quantitative and qualitative predictions of the theoretical model underlying the experiment.

We don't observe a bandwagon effect in electorates with a more lopsided division of preferences. Revealing the distribution of preferences in electorates that aren't closely divided doesn't affect the turnout rate of subjects supporting the small team but lowers the turnout of subjects supporting the large team. For example, we see a decrease in the participation of subjects after learning that they belong to a team of size seven. A similar phenomenon occurs for subjects that belong to a team of six subjects. Note that these subjects turn out in a frequency lower than the frequency of a subject that is the sole supporter of an alternative. An analogous situation occurs when the subjects' preferences are divided between teams of five and two members. None of these effects, however, are statistically significant at standard confidence values.

Based on a regression analysis, Klor and Winter (2006) show that, in lopsided divided electorates, the best predictor for a subject participation decision after the provision of information is the subject's turnout decision before the provision of information. That is, there are subjects that like to participate in elections, and these are therefore the subjects that turn out before the provision of information and after learning that the electorates are not closely divided. In particular, in lopsided elections subjects that voted in the first stage of a round without knowing the distribution of preferences are more than 25% more likely to vote in the second stage, regardless of whether they belong to the large or small team.

Summarizing, in closely divided electorates subjects in the majority vote with a significantly higher frequency than subjects in the minority. In the rest of the electorates we don't observe significant differences on the voting behavior of subjects in the majority relative to subjects in the minority. As already

pointed out in the introduction, these results contradict the theoretical predictions of traditional models of participation based on Nash equilibria. They cannot be accounted either by the alternative quantal response equilibrium concept (Goeree and Holt 2005).

The anomaly observed in the subjects' behavior in the laboratory may be consistent with several interesting alternatives to the traditional rational choice approach.[8] Some of these alternatives propose different behavioral assumptions on the voters in order to explain the bandwagon effect observed in sequential voting contests. One of the most prominent approaches is based on the observation that voters show a desire to vote for the winner.[9] Basically, this approach adds an extra component to the voters' utility function to capture their desire to support the winning candidate (Callander 2004). Usually, the related studies focus on the presidential primaries in the United States, where the candidates for the presidential elections are determined through a sequence of elections in single states that span over a period of several months (Bartels 1988). Alternatively, the model can also be applied to case studies where voting takes place over several different time zones. In these cases, the voting behavior of individuals living in areas that are still voting might be affected by the electoral results of areas that closed their voting booths earlier.

While the approach based on the desire to support the winning candidate manages to explain the bandwagon effect of polls, it is unable to account for the fact that, in simultaneous elections, the bandwagon effect is only present in closely divided electorates. We, therefore, propose below two conjectures that may explain the existence of the bandwagon effect only in closely divided electorates. Our first conjecture is that the voters' utility function shows greater sensitivity to losses than to gains when the individual belongs to the majority.[10] In particular, an individual in the majority may "suffer" more in utility terms from an electoral defeat than an individual in the minority. As a consequence, individuals in the majority may be more willing to vote than individuals in the minority, especially when an electoral defeat is more likely. This behavior would result in the emergence of a bandwagon effect only in closely divided electorates.

Our second conjecture places emphasis on the fact that voters may miscalculate the probability of casting a pivotal vote for the different distributions of preferences. Riker and Ordeshook (1968) claim that this may be the main reason for the positive effect of closeness on the turnout rate. Informal interviews with the subjects conducted after the experiment suggest that this factor

[8] See Feddersen (2004) for a recent literature review on models of turnout.

[9] This observation was established over fifty years ago by Berelson et al. (1954) in their analysis of presidential primaries. It was corroborated more recently by Herron (1998) using micro-level data of voters' desire to vote in the 1992 presidential elections.

[10] This conjecture in fact suggests that the individuals' utility function is reference dependent, as proposed by Kahneman and Tversky (1979). See Quattrone and Tversky (1988) for a discussion of the applicability of prospect theory to political science. We thank an anonymous referee for pointing out the connection between our experimental results and loss aversion.

has indeed some explanatory power for the bandwagon effect of polls in closely divided electorates.

4.2 Welfare analysis

The subjects' anomalous turnout behavior has direct implications on the desirability of public opinion polls. Theoretical models studying the welfare effects of polls conclude that polls unambiguously reduce aggregate welfare due to two detrimental effects: polls cause an increase in turnout (with the associated cost of voting), and they increase the likelihood that the small team wins the election.

In the laboratory we found that polls indeed cause a significant increase in turnout in closely divided electorates. Furthermore, we observe an increase in the frequency of electoral victories of the large team as a consequence of the bandwagon effect of polls. That is, contrary to the theoretical predictions the two effects of polls don't work in the same direction in the laboratory. Hence, without proper calculations we can only conclude that the welfare effect of polls in electorates that are closely divided is ambiguous.

In lopsided divided electorates we expect polls to have a small positive effect on total welfare. Intuitively, turnout slightly decreases as a consequence of the revelation of information, but the decrease is not significantly different for large and small teams. Thus, the probability of the larger team winning the elections should not be significantly affected by the publication of the poll.

Figure 2 depicts the frequency of minority victories for every distribution of preferences. The figure includes the frequencies predicted by the unique BNE and QSNE.[11]

Before we analyze the figure above, it is important to keep in mind that the strategy profile that maximizes the sum of the subjects' payoffs is for exactly one subject of the larger team to vote. Note, thus, that the unique quasi symmetric Nash equilibrium is clearly inefficient, since it predicts that the smaller team wins with frequencies higher than the larger team for every distribution of preferences but the most skewed one.

The experimental data obtained in elections when the distribution of preferences isn't known is consistent with the main characteristics of the BNE. Namely, the frequencies of victories of the small team are always below fifty percent, and they decrease monotonically as the difference in the support between the two teams increases.

After the provision of information, the data show a pattern that contradicts the QSNE predictions in two respects. First, for every distribution of preferences the large team wins with higher frequency than the small team. Moreover, the frequency of victories of the small team decreases together with decreases

[11] All the statistical tests performed in this section use the electorate as the unit of observation. In fact, for every electorate we have a pair of observations since we compare the electoral outcomes before and after the provision of information on the distribution of preferences.

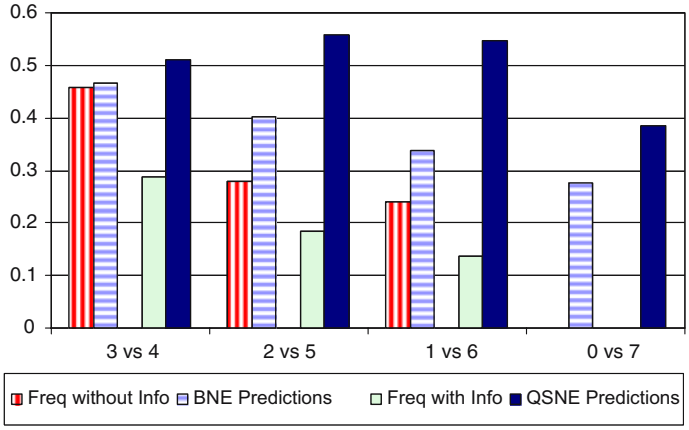

Fig. 2 Frequency of electoral victories of the small team by distribution of preferences

of the small team's relative size of the electorate. Second, the provision of information causes a decrease in the frequency of victories of the small team for every distribution of preferences. For closely divided electorates, revealing the electorate's distribution of preferences lowers the frequency of victories of the minority by almost 20 percentage points.[12] Given that members of the small team increase their participation after the provision of information, the significant change in the frequency of victories of the minority is a direct consequence of the bandwagon effect of polls. We observe an important decrease in the frequencies of victories of the small team for the other distributions of preferences as well.[13] This effect, however, is a consequence of the slight decrease in participation of all the subjects, both in the large and small teams.

The effects of information revelation on the frequencies of victories of the minorities have direct implications on the subjects' total payoffs. Table 2 depicts the effects of polls on aggregate welfare for every distribution of preferences. The table presents the average total payoffs divided by the maximum available payoffs by distribution of preferences.[14]

The table shows that the provision of information reduces the level of efficiency in closely divided electorates. Total average payoffs before the provision of information for electorates that are closely divided are significantly

[12] Without information on the distribution of preferences a team of three subjects wins 45.9% of the elections. With information, a team of three subjects wins only 28.6% of the elections. The difference between the two is highly statistically significant with $z = 2.75$ ($p < 0.01$) according to a two-sided sign test using the normal approximation to the binomial distribution.

[13] This change is statistically significant at the 5% level for a distribution of preferences of six subjects versus one. The difference isn't statistically significant at conventional levels for a distribution of preferences of two subjects against five. Both conclusions are reached using a two-sided sign test.

[14] The strategy profile that maximizes the sum of subjects' payoffs delivers payoffs of 36 tokens for electorates divided in teams of three and four member; payoffs of 46 tokens for electorates divided in teams of two and five members, and so forth for the rest of the electorates.

Table 2 Aggregate payoffs divided by efficient payoffs by distribution of preferences (standard deviations appear in parentheses)

Distribution of preferences	3 vs. 4	2 vs. 5	1 vs. 6	0 vs. 7
Total average payoffs without information	0.7832 (0.2009)	0.7480 (0.2727)	0.7574 (0.3868)	1
Total predicted payoffs by the BNE	0.9186	0.7753	0.7284	0.7321
Total average payoffs with information	0.7208 (0.1783)	0.8035 (0.2545)	0.8596 (0.2917)	1
Total predicted payoffs by the unique QSNE	0.8899	0.5517	0.4630	0.6364

higher than after the provision of information.[15] This decrease is caused by the increase in turnout of individuals in both teams. In the rest of the electorates the provision of information increases efficiency, mainly because we observe an important increase in the frequency of electoral victories of the larger team.[16] Whereas the actual level of efficiency before the provision of information is relatively close to the predictions of the BNE, we observe important differences between the actual level of efficiency after the provision of information and the predictions of the QSNE. The actual efficiency level after the provision of information is higher than the equilibrium predictions due to the significantly higher frequencies of victories of the majority relative to the QSNE predictions documented above in Figure 2.

This table presents, in fact, a reproduction of Table 3 in Großer et al. (2005) using a different experimental design. As in Großer et al. (2005), we observe that the level of efficiency monotonically decreases as the size of the small team increases. Moreover, our results show that revealing information on the distribution of preferences has an important detrimental welfare effect when the electorates are closely divided. We don't observe a significant effect in the rest of the electorates.

Table 2, however, masks important and unexpected differences between teams for a given distribution of preferences. To disentangle the effects of information on the payoffs of the different teams for a given distribution of preferences, we present in Fig. 3 the average share of total realized payoffs obtained by subjects in the small team.[17] The figure also presents the share of total payoffs predicted by the Bayesian Nash equilibrium and the quasi-symmetric Nash equilibrium.

This figure shows the most startling effect that emerges from the welfare analysis of this paper: The effect of information on the voters' payoffs is not homogenous across teams of different sizes, but it significantly lowers the payoffs of the

[15] The difference in efficiency is statistically significant at the 1% level with $z = -2.8$. This conclusion is based on a Wilcoxon matched-pairs signed-ranks test.

[16] The difference is statistically significant at the 10% level with $z = 1.65$ for distributions of two versus five. The difference is not statistically significant for distributions of one versus six and seven versus zero. All the results are based on a Wilcoxon matched-pairs signed-ranks test.

[17] That is, this figure disregards the efficiency effects of polls.

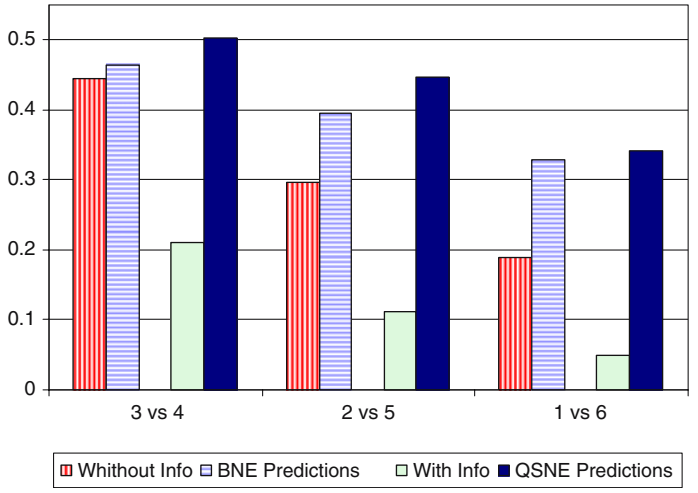

Fig. 3 Average share of total realized payoffs of the minority by distribution of preferences

small teams relative to the payoffs of the large teams. This effect is not expected from a comparison of the predictions of the BNE to those of the QSNE. On the contrary, according to the equilibrium predictions the provision of information should increase the relative payoffs of the small teams.

Before the provision of information the distribution of payoffs between the teams is almost equal to each team's share of the population for every distribution of preferences.[18] As a consequence of the provision of information on the electorate's distribution of preferences the share of realized payoffs obtained by the small team decreases significantly. We observe a decrease of 23 percentage points on the average share of the payoffs for a team of size three; a decrease of 18 percentage points on the average share of the payoffs for a team of size two; and a decrease of 14 percentage points on the average share of the payoffs for a team of size one. Whereas the decreases in the share of the payoffs obtained by the small team are only marginally statistically significant in lopsided divided electorates ($t = 1.16$ for a distribution of two against five and $t = 1.74$ for a distribution of one against six), the difference is highly statistically significant in closely divided electorates ($t = 3.46$).

As a direct consequence of the decrease in the payoffs of the small team we observe a very significant difference between the small teams' share of the realized payoffs in the laboratory and the one predicted by the unique equilibrium for every distribution of preferences. The difference between the observed and the predicted minority's share of the realized payoffs oscillates around 30 percentage points for each distribution of preferences. Hence, we can

[18] A team of size three constitutes 42.8% of the population and obtains 44.4% of the realized payoffs. A similar picture emerges for the other teams: A team of size two constitutes 28.6% of the population and obtains 29.6 of realized payoffs, whereas a team of size one constitutes 14.2% of the population and obtains 18.9% of the realized payoffs.

unambiguously conclude that public opinion polls significantly hurt minorities in the laboratory, contrary to the theoretical predictions.

5 Conclusions

This paper studied the welfare effects of the provision of information on the voters' distribution of preferences. The main findings are that closeness induces an important increase in turnout, and that the increase in turnout is significantly larger for the alternative with a slight majority according to the poll. As a consequence, polls increase voting costs and the frequency of victories of the large team. Overall, this has a negative effect on total welfare, and it especially lowers the minority's share of the total realized payoffs.

In lopsided divided electorates, polls have a positive welfare effect as they cause a decrease in overall participation and an increase in the majority's frequency of electoral victories. As was the case in electorates that are closely divided, also in lopsided divided electorates polls redistribute part of the electoral payoffs from the minority to the majority.[19]

Summing up, our analysis makes several contributions to the debate on the desirability or not of imposing restrictions on the publication of public opinion polls before elections. First, we show that polls do affect the voters' behavior but only in closely divided electorates. Second, we observe that polls do have a significant effect in overall welfare: They decrease welfare in closely divided electorates and they increase aggregate welfare in the rest of the electorates. Therefore, it is impossible ex-ante to provide a clear cut prediction on the desirability of polls based only on total aggregate welfare. If, on the contrary, the guiding normative principle takes equity considerations into account, we can conclude that polls are undesirable. They consistently disenfranchise the minority and significantly reduce its share of the payoffs to levels far below the minority's share of the electorate.

Acknowledgement We thank an anonymous referee for detailed comments that helped us improve the quality of the paper. The paper has also benefited from the comments and suggestions of seminar participants at ECARE and The Center for the Study of Rationality at The Hebrew University of Jerusalem. All errors and mistakes in the paper remain our own.

References

Bartels LM (1988) Presidential primaries and the dynamics of public choice. Princeton University Press, Princeton

Battaglini M, Morton R, Palfrey T (2005) Efficiency, equity and timing in voting mechanisms. Centre Econ Policy Res DP 5291

Berelson BR, Lazarsfeld PF, McPhee WN (1954) Voting. The University of Chicago Press, Chicago

[19] These results are consistent with the ones presented in Battaglini et al. (2005). Testing the effects of information revelation in a sequential voting game where voters have identical preferences, Battaglini et al. (2005) found that the revelation of information produces more efficient outcomes and increases inequities between informed and uninformed voters.

Callander S (2004) Bandwagons and momentum in sequential voting. Northwestern University, mimeo
Casella A, Palfrey T, Riezman R (2005) Minorities and storable votes. Centre Econ Policy Res DP 5278
Cason TN, Mui V-L (2005) Uncertainty and resistance to reform in laboratory participation games. Eur J Politi Econ 21(3):708–737
Feddersen TJ (2004) Rational choice theory and the paradox of not voting. J Econ Perspect 18(1):99–112
Goeree JK, Holt CA (2005) An explanation of anomalous behavior in models of political participation. Am Polit Sci Rev 99(2):201–213
Goeree JK, Großer J (2006) Welfare reducing polls. Econ Theory (forthcoming)
Großer J, Kugler T, Schram A (2005) Preference uncertainty, voter participation and electoral efficiency: an experimental study. Working Paper No. 2, University of Cologne
Herron MC (1998) Voting, abstention, and individual expectations in the 1992 Presidential election. Northwestern University, mimeo
Kahneman D, Tversky A (1979) Prospect theory: an analysis of decision under risk. Econometrica 47(2):263–291
Klor EF, Winter E (2006) On public opinion polls and voters' turnout. Cent Econ Policy Res DP 5669
Palfrey T, Rosenthal H (1983) A strategic calculus of voting. Public Choice 41(1):7–53
Riker W, Ordeshook P (1968) A theory of the calculus of voting. Am Polit Sci Rev 62(1):25–42
Quattrone GA, Tversky A (1988) Contrasting rational and psychological analyses of political choice. Am Polit Sci Rev 82(3):719–736
Taylor CR, Yildirim H (2005) Public information and electoral bias. Duke University, mimeo