Journal of Economic Behavior & Organization 84 (2012) 681–700

Contents lists available at SciVerse ScienceDirect

Journal of Economic Behavior & Organization

journal homepage: www.elsevier.com/locate/jebo

Pre-election polls as strategic coordination devices[☆]

Costel Andonie[a,*], Christoph Kuzmics[b,1]

[a] Department of Economics and Finance, City University of Hong Kong, Hong Kong
[b] IMW, Bielefeld University, Bielefeld, Germany

ARTICLE INFO

Article history:
Received 3 May 2012
Received in revised form 21 August 2012
Accepted 10 September 2012
Available online 19 September 2012

JEL classification:
C72
D83

Keywords:
Symmetry
Focal points
Voting
Duverger's law
Condorcet loser

ABSTRACT

In the unique *attainable* equilibrium of a voting model with one minority candidate and two similarly appealing majority candidates, majority voters are unable to coordinate their support and the minority candidate (Condorcet loser) is elected. Suppose a random sample of voters is asked about their preferences prior to the election. We show that there always exists an equilibrium of this two stage game in which all poll participants are truthful, resulting in a high likelihood of a majority candidate winning the election. This equilibrium is unique if the sample size of the poll is Poisson distributed or fixed and odd.

1. Introduction

Elections, in which voters act strategically, typically have multiple equilibria (e.g. Palfrey, 1989; Myerson and Weber, 1993; Cox, 1994; Myerson, 2002). Thus, elections involve a coordination problem. Elections are typically preceded by opinion polls. In this paper we are interested in understanding how a single opinion poll, in which participants act strategically, could serve the voters as a coordination or equilibrium selection device.[2] Furthermore, under the assumption that the outcome of the opinion poll is used by voters as an equilibrium selection device and that this is common knowledge among voters, we are interested in how strategic poll participants behave in the opinion poll.

☆ We thank the editor and the referees for their comments and suggestions. We also thank David Austen-Smith, Vincent Crawford, Eddie Dekel, Julio Gonzalez Diaz, Daniel Diermeier, Timothy Feddersen, Bard Harstad, Wojciech Olszewski, Alessandro Pavan, Tuan Hwee Sng, Joel Sobel, Bruno Strulovici, Asher Wolinsky and audiences at Australian National U, City U of Hong Kong, Northwestern MEDS, U of Bristol and Washington U in St. Louis for their comments.

* Corresponding author.
E-mail addresses: candonie@cityu.edu.hk (C. Andonie), christoph.kuzmics@uni-bielefeld.de (C. Kuzmics).

1 Present address: IMW, Bielefeld University, Bielefeld, Germany.

2 Of course, in real-world elections, there are other potential coordination devices that voters could use to coordinate their actions, for example public endorsements (e.g. Ekmekci, 2009), campaign spending levels (e.g. Myerson et al., 1998) etc. However, it will emerge in the analysis that an opinion poll is preferred by the majority voters at least to any exogenous one (that does not depend on voters' actions), such as a public coin toss. The reason is that a poll potentially contains information about the structure of the electorate, and thus, it allows voters to make better informed and thus more efficient ex-ante decisions as a group. While, in this paper, we abstract away from the potential information opinion polls could provide, we briefly discuss this issue in Section 3.1.

0167-2681/$ – see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.jebo.2012.09.014

Real-world opinion polls, in addition to possibly serving as a coordination device, also possibly provide information about the structure of the electorate. In this paper, as we are interested only in the equilibrium selection problem and to facilitate the analysis of the strategic implications of this aspect of opinion polls alone, we shall abstract away from this second (possibly important) role opinion polls might have. In fact we shall argue that opinion polls can play an important role in elections *even if they do not provide any information*.

In this paper we choose to work with the simplest election model that allows us to formally study the strategic implications of the assumed common knowledge fact that the opinion poll is used as a selection or coordination device. The formal model goes back to Myerson and Weber (1993) and Forsythe et al. (1993) and can be briefly described as follows.[3] There are three candidates: A, B, and C and three groups of voters: AB, BA, and C. The letters identifying the three groups of voters indicate their favorite candidates: A followed by B for AB group etc. AB and BA voters, as a group, form a majority of, say, 60%, while the C group voters represent 40% of the electorate. The model assumes that majority voters support either A or B, while C voters support only C. This implies that a majority candidate (either A or B) victory is possible only if majority voters (AB and BA voters) coordinate their support[4]; in the case of a vote split, for example half of the majority group supports A while the other half supports B, candidate C, which is a Condorcet loser, would win with 40%.

We assume that such an election is preceded by an opinion poll. A random sample of m individuals from the electorate is taken and each participant is asked to indicate their favorite candidate, that is to choose one letter (message) from the set $\{A, B, C\}$. The outcome of the poll is then publicly announced and observed before the election takes place.

As it will emerge from the analysis, what individuals know about the sample size m has important implications about the equilibria in the polling game.[5] As it will also emerge from the analysis, solving for all equilibria for an arbitrary distribution of the poll size is not easy. Our strategy in this paper to tackle this problem, is to first solve the game under the simplifying assumption of a Poisson distributed poll size.[6] This has the advantage of making the analysis very crisp and clear-cut, while at the same time already providing the key strategic considerations in the polling game.

We then study the case of a commonly known fixed poll size. It turns out that one there has to carefully distinguish between odd and even poll sizes. This distinction remains important even if the poll size is large (even if it tends to infinity). In addition to being of independent interest, the case of a commonly known fixed poll size provides insight into what might happen with an arbitrary poll size distribution. This is so, because, as it turns out, regardless of what the commonly known poll size is, there is always a truthful equilibrium. That is there is always an equilibrium, in which all polled individuals indicate truthfully their favorite candidate. Given that this is so for all poll sizes, we can conclude that there is a truthful equilibrium for any arbitrary distribution of poll sizes. In fact, as long as an individual believes all other poll participants to be truthful, and regardless of what she believes the poll size distribution is or what other poll participants' beliefs about the poll size distribution are, she has a unique best action, which is to be truthful.

The solution concept we employ in this paper is symmetric Nash equilibrium. More precisely, we restrict attention to attainable equilibria.[7] This is done because, if we believe that achieving coordination among the majority supporters is, in principle, difficult, we need to appeal to a solution concept (for the voting game) that reflects that. Indeed, in the absence of a polling stage the voting game alone, as we shall see, has a unique attainable equilibrium, in which majority supporters are not able to overcome the coordination problem.

The two-stage game – an opinion poll followed by an election – typically has multiple attainable equilibria. We shall focus on those that could be called focal, using the term in the spirit of (Schelling, 1960) as the way to play the game suggested by the structure of the game itself. To be more specific we postulate that there is a most efficient, and thus focal, way to behave in the voting stage as a function of poll results. We then analyze the behavior in the opinion poll under the assumption of such focal behavior in the voting stage.

The main results are as follows. We show that if the size of the opinion poll is Poisson distributed, and not known to the poll participants, then the two-stage game has a unique equilibrium, in which all polled individuals truthfully reveal their "type," i.e. voters who favor, say, candidate A announce this in the poll. The intuition for the result is as follows. Consider first a minority voter poll participant. Our model suggests that majority voters are able to coordinate whenever there is a majority candidate who is leading in the polls. Otherwise, by attainability, they are unable to focus attention on any particular candidate. Therefore, there are two events of interest to this polled minority voter. Consider first the event in which one of the majority candidates leads in the poll by only one message. By sending a misleading message, such as "I favor a majority candidate," she can create a tie, by which she triggers a non-coordinated outcome and therefore gains in

[3] Several historical elections are frequently used in the literature to motivate the formal model described here. These include the 1970 New York senatorial elections (e.g. Forsythe et al., 1993; Fey, 1997; Myatt and Fisher, 2002; Myatt, 2007), the 1912 US presidential elections (e.g. Myerson and Weber, 1993), and the 1997 UK general elections (e.g. Myatt and Fisher, 2002; Myatt, 2007).

[4] The game played by the AB and BA types is a multi-player version of the "battle of sexes." We may want to think of AB as the row player and BA as the column player. If they coordinate their actions on either A or B, they both get positive payoffs, though AB would like to coordinate on A, while BA would like to coordinate on B. If they do not coordinate, they both get a payoff of zero.

[5] It is not clear what actual polled individuals are told about the poll size. If told nothing, an individual may then perhaps believe that, as common poll sizes are 100 and 1000, the poll size has an equal chance of being one of these two numbers. Of course, many other beliefs could be reasonable.

[6] The Poisson distribution was introduced to voting by Myerson (1998, 2000).

[7] The voting game is not a completely symmetric game. Some care is needed to properly define the symmetries that are inherent in the game. Strategy profiles that respect those symmetries are then called attainable. The term attainable was coined by Crawford and Haller (1990).

this situation. By lying, however, she also runs the risk of breaking the tie between the majority candidates with the result of a coordinated outcome and a payoff loss. Therefore, her incentives to deviate from a truthful message come down to comparing the probability of a tie with the probability of a near-tie between the number of messages favoring each of the majority candidates. By the attainability restriction and the Poisson assumption, the number of messages favoring either of the two candidates must follow the same Poisson distribution. Moreover, their distribution is independent of each other. A simple statistical property (that is proven in the appendix) shows that the probability of having a tie between two i.i.d. (discrete) random variables is always larger than the probability of a near-tie. Therefore, a minority voter will never find it in her interest to lie in the poll. In other words, sending a truthful message is a dominant strategy.[8] A similar argument applies to the majority voters: sending a misleading message always yields a negative expected payoff. Therefore, for the Poisson case, the polling game has a unique attainable equilibrium in dominant strategies.

We then consider the case of a fixed poll size. We show that if the size of the opinion poll is fixed at some odd number that is commonly known, then the game has again a unique equilibrium, very much like the one in the Poisson case. If the size of the poll is fixed at some even number, then there are multiple equilibria: the truthful one, as well as two additional ones, in which the reported fraction of minority supporters is always strictly lower than the actual fraction. Thus, the poll results in these equilibria are not to be taken completely at face value. The complications for the case of a fixed poll size arise from the fact that, unlike the Poisson case, the number of messages favoring the majority candidates are no longer independent. Finally, we show that the polling game has a truthful equilibrium for any distribution of the poll size. This is an immediate consequence of the fact that the game has a truthful equilibrium for both the cases in which the poll size is odd and even.

The paper proceeds as follows. In Section 2 we introduce the (pure coordination) polling and voting model. As part of this section we provide the solution concept in Section 2.1 as well as the solution of the voting stage in Section 2.2. All equilibria are then presented in Section 3, with those for Poisson random poll sizes in Section 3.3, for fixed poll sizes in Section 3.4, and for arbitrarily distributed poll sizes in Section 3.5. Finally, we conclude the paper in Section 4. All proofs are relegated to the appendix, while some simple proofs are simply omitted.

1.1. Literature review

Historically, the literature studying coordination in plurality rule elections revolved around an empirically observed phenomenon, later stated as a theoretical conjecture, called "Duverger's law" (Duverger, 1954). In short, Duverger's law predicts that, in plurality rule elections, only two candidates will get non-negligible shares of votes. Duverger's explanation is that rational voters will direct their support toward only two candidates that are thought to have good chances of winning; while the other candidates will receive few or even no votes as voters will not "waste" their votes on almost sure losers.[9] The first one to provide a convincing formal proof of the law was Palfrey (1989). Considering a setup of a three-candidate election with rational voters, under quite weak assumptions, Palfrey (1989) showed that only two types of equilibria survive as the population size becomes arbitrarily large. The first one (called Duvergerian) has only two candidates receiving most of the votes, with the third candidate's share vanishing in the limit. The second type (called non-Duvergerian), however, has all three candidates receiving positive shares, with one candidate coming on top and no clear runner-up i.e. the other two candidates receive positive and approximately equal shares (see also Cox, 1994, 1997) Palfrey (1989) argues that the second case – the non-Duvergerian equilibrium – is an "exceptional" case and, thus, it is not expected to occur in real elections. Myerson and Weber (1993) showed that the non-Duvergerian equilibrium survives ϵ-perturbations in the structure of the electorate, and therefore cannot be dismissed as "exceptional". They argue: "Thus, Duverger's law cannot be derived exclusively from analyzes of voting equilibria. This result suggests that some of the interpretative discussion in Palfrey (1989) should be reconsidered. Any derivation of Duverger's law would seem to require some additional assumption of dynamic stability or persistence to eliminate equilibria of the type just illustrated" (Myerson and Weber, 1993, p. 106). Fey (1997), following on the suggestion of Myerson and Weber (1993), investigated the dynamic stability of both Duvergerian and non-Duvergerian equilibria. While he finds that the Duvergerian equilibrium is dynamically stable, the non-Duvergerian one is not. Elections, however, are not usually repeated. The US presidential elections, for instance, take place only once every 4 years, and then usually with one or more different candidates (or, at least older and wiser candidates). Of course, elections are typically preceded by (many) opinion polls, but opinion polls are not elections. They do not "count," that is, whatever is said in an opinion poll is cheap talk. Moreover, not everyone is asked to participate in the opinion polls. Thus, a dynamic analysis as in Fey (1997), justified by a sequence of pre-election polls, is not exactly appropriate. Here, we take this point seriously and provide a strategic analysis of behavior in both the opinion poll as well as the voting game.

Myatt and Fisher (2002) and Myatt (2007) showed that the predictions of Palfrey's model change dramatically if voters are uncertain about the real support of each candidate. In such a setup, and in which voters receive information by way of a private signal, they show that in equilibrium all candidates receive votes. The equilibrium fits neither the Duvergerian nor the non-Duvergerian type. As Myatt (2007) shows, the result relies crucially on the private information structure. Allowing voters to have access to a public signal with a good precision in relation to the private signal will yield the same equilibrium predictions as in Palfrey (1989). Furthermore, the information is exogenous. Even though voters receive their

[8] This is only true in the class of attainable strategy profiles. If all strategy profiles are considered then sending a truthful message is not dominant.

[9] See for example Riker (1982) for a history and detailed discussion of Duverger's Law.

Table 1
Experimental Results for elections with and without polls.

	Number of elections	% of time A won	% of time B won	% of time C won
Elections without polls	48	6.25	6.25	87.5
Elections with polls (overall)	48	45.83	20.83	33.33
A leads B in polls	24	87.5	0	12.5
Tie between A and B in polls	11	9.09	0	90.91
B leads A in polls	13	0	76.92	23.08

Note. Data source: Forsythe et al. (1993). The labels of voter types and candidates have been changed to fit the description of our setup.

private information from their counterparts, i.e. other voters, the model does not take into account the strategic incentives of senders (other voters) to manipulate the beliefs and thus influence the voting behavior of others. The assumption is that players truthfully reveal their information. In contrast, our model focuses exclusively on the public information side of the model. The key contribution of our paper is the endogenous information structure that properly takes into account the incentives of poll-participants to truthfully report their preferences. The model does not feature aggregate uncertainty.[10] Given a poll outcome, the voting behavior in our setup is trivial. This simplification allows us to focus on the strategic behavior of poll-participants, as they realize that their messages potentially influence the voting decisions and, therefore, the election outcome.

The setup we employ in this paper has been used by Forsythe et al. (1993) to run lab experiments in order to study the impact of polls, election histories, and ballot position on elections. The experimental results for the voting model, both with and without polls, are presented in Table 1.[11] We note the starkly different voting behavior in the two settings. In the game without polls, the outcome is, in a majority of cases (more precisely 87.5% of the time), non-Duvergerian, i.e. majority voters divide their votes between A and B. In the game with polls however, we see a completely different picture: the outcome is now, most of the time (more precisely 66.66%), Duvergerian, i.e. the majority coordinated their votes. The experimental evidence suggests that majority voters use the poll outcome as a tool to coordinate their actions, resulting in better outcomes for the majority group. It is this idea that we try to explore in this paper.

As the experimental evidence of Forsythe et al. (1993) suggests, polls can have a significant influence on an election's outcome. Goeree and Grosser (2007) and Taylor and Yildirim (2010) are two recent papers that touch on this idea. They consider a two-candidate voting model with costly voting to study the effect of the information provided by polls[12] on total welfare. They show that, on one hand, more precise information on electoral preferences increases turnout, and consequently increases the aggregate cost. On the other hand, more precise information makes the election more likely to be tied, therefore reducing the aggregate benefit. These two effects work in the same direction of reducing total welfare. While these two papers make the case for restricting the use of pre-election polls, as is currently done in some countries, they do not allow voters to be strategic in the polling stage, and simply assume they are all sincere. As Taylor and Yildirim (2010) put it, "Before firm policy conclusions can be drawn, however, future research must investigate the endogeneity of the source of information about political preferences. In particular, given the feedback we find between information and equilibrium voting behavior, it is important to understand the incentives for individuals to report their true preferences to pollsters and the incentives for pollsters and pundits to disclose publicly and fully any information they obtain." (Taylor and Yildirim, 2010, p. 363) One of the contributions of our work is providing a setup that properly takes into account voters' incentives to truthfully reveal their preferences in polls.[13]

2. The (pure coordination) model

The basic setup we use is a standard three-candidate voting model, similar to the models in e.g. Palfrey (1989), Myerson and Weber (1993), Fey (1997), and Myatt (2007). In what follows, we will use Palfrey (1989)'s notation. There are three candidates: A, B, and C and a (countably) infinite number of voters. Candidates A and B can be described as ideologically similar, while candidate C's political position is at the other end of the political spectrum. The population of voters is made up of three distinct groups denoted AB, BA, and C. With $\nu \in (0, 1)$ a known constant, their preferences can be described as follows:

1 AB voters: $u_{AB}(A) = 1$, $u_{AB}(B) = \nu$ and $u_{AB}(C) = 0$
2 BA voters: $u_{BA}(A) = \nu$, $u_{BA}(B) = 1$ and $u_{BA}(C) = 0$
3 C voters: $u_C(A) = 0$, $u_C(B) = 0$ and $u_C(C) = 1$

[10] This assumption however can be relaxed without changing the main implications of the model. We discuss this in Section 3.3.
[11] The game with polls is identical to the game without, except that, before the actual voting stage in the former, voters are asked to state their voting intentions (or preferences). The results of the poll are publicly announced, after which voting takes place.
[12] More precisely, they look at the role information (about the structure of the electorate) plays on elections outcomes, and thus the welfare. The information presumably comes from polls although they do not model that explicitly.
[13] Strategic behavior in opinion polls in different contexts has also been studied by Burke and Taylor (2008), Meirowitz (2005) and Morgan and Stocken (2008).

The share of AB and BA voters is q each, while the share of C voters is $1 - 2q$, where $q \in (1/4, 1/3)$ and known by everyone.[14] The winner is the candidate with a majority of votes, with ties broken (uniformly) randomly.

Note that the combined share of AB and BA voters is more than 50%. We say that the group of AB and BA voters are in a majority. Because $1/4 < q < 1/3$, a majority victory is possible only if AB and BA voters coordinate their support; if they do not coordinate then the minority candidate, C, wins.

Before the voting stage, a random sample of m voters ($m \geq 1$) is selected and asked who their favorite candidate is. A polled person must state one of three messages – A, B, or C – where the meaning of these messages will only become clear in equilibrium (as we shall see later). Sending any of these three messages is costless for the polled person, and is thus cheap talk. The polled persons' incentives therefore derive entirely from the effect the poll has on the eventual election outcome. We assume that the outcome of the poll is publicly announced (and observed by everyone) before the voting stage. Thus, our game has 3 stages:

1 Each player in the electorate is assigned a type. Types are drawn from a trinomial distribution with parameters $(q,q,1 - 2q)$.
2 (*Polling stage*) A random sample of m voters is selected and asked to reveal their preferences (their types) by sending a message from the set $\{A, B, C\}$.
3 (*Voting stage*) The poll results are publicly announced. The actual voting takes place and the candidate with the largest share of votes is elected.

2.1. Solution concept

The game in the polling stage is thus a normal form game $\Gamma = (I, S, u)$ with $I = \{1, \ldots, m\}$, $S_i = S = \{A, B, C\}$, and u derived from the outcome of the eventual voting stage. We shall discuss the voting stage in the next section.

For $i \in I$, let $\theta(i) \in \{AB, BA, C\}$ denote player i's type. Given the model's assumptions, the two types AB and BA are in an exactly symmetric position. One could therefore say that they are, in some sense, of the same type. This type could be characterized as the type who "prefers one of the two majority candidates over the other and considers the minority candidate her least preferred choice." Formally, any two players of this type are symmetric in the sense of Harsanyi and Selten (1988) (see also Alos-Ferrer and Kuzmics, forthcoming).

Given this, we shall restrict our attention to *attainable* Nash equilibria, where an attainable strategy profile is understood to be a strategy profile that satisfies the symmetry restrictions inherent in the game at hand.[15]

Before we provide the definition of an attainable strategy profile for the polling game, we first explain why we impose these restrictions on the possible equilibria we consider. Observe first this very important fact: the game we discuss here is an (almost) truly one-shot game. Every election is different. Suppose that in an election involving candidates Ann, Bob, and Charlie, majority voters somehow managed to coordinate their support on Bob. This will not help voters in another election that has the same general characteristics but with candidates Herbert, Edgar, and Josephine, unless they followed some peculiar convention regarding, say, the first initials of candidates or the length of their first name.[16]

Consider, for instance, a case in which we would like exactly one half of all C types to send message C in the poll and the other half to send message A. How would we achieve this? We could announce in a newspaper that this is what we would like to achieve. But how will any given polled C type now know whether she is supposed to send message C or to send A? Unless we talk to every polled C type separately, the best we can hope for is that all C types randomize by attaching equal probability of one-half to both messages.

The restrictions imposed by symmetries in the polling game are as follows. First, all voters of the same type are required to use the same strategy. Second, AB and BA types' strategies must be symmetric. Finally, all C types must attach the same probability to messages A and B. Formally, let x_{AB}, x_{BA}, x_C denote the (mixed) strategies $\in X \equiv \Delta(\{A, B, C\})$ that AB, BA, and C types use, respectively.[17] Then a strategy profile is attainable if $x_{AB}(C) = x_{BA}(C)$, $x_{AB}(A) = x_{BA}(B)$ (implying $x_{AB}(B) = x_{BA}(A)$), and $x_C(A) = x_C(B)$.

2.2. The voting stage

The voting stage game has been already extensively analyzed in the voting literature. In what follows, we shall briefly review some of the results from these papers. The analysis below (and notation) follows Palfrey (1989) and Fey (1997).

---

[14] Further below we discuss the case in which voters are uncertain about the structure of the electorate.

[15] The term *attainable strategy profile* was coined by Crawford and Haller (1990) in their study of repeated (symmetric) pure coordination games. Examples of the use of attainable equilibria, in some form or another, are given in Farrell (1987), Crawford and Haller (1990), Blume (2000), Blume and Gneezy (2000), Blume and Gneezy (2010), Blume and Franco (2007), and Blume et al. (2009). In the special case of voting in Poisson games, Myerson (1998, 2000) has been led to focus on symmetric strategies as well.

[16] This game is repeated, but not in the same way as the game about which side of the road to drive on. In that case the strategies Left and Right are always exactly the same. Over time a convention such as driving on the right (or on the left) could easily evolve. No such thing is possible in our elections.

[17] Note that we index the strategies by voter type and not identities. This is used for notational simplicity and it is justified by our restriction that all voters of the same type must use identical strategies.

Table 2
Equilibrium payoffs.

Type of equilibrium	u_{AB}	u_{BA}	u_C
No coordination	0	0	1
Majority coordinates on A	1	ν	0
Majority coordinates on B	ν	1	0

Let $\sigma^i : \{AB, BA, C\} \to \{A, B, C\}$ denote the voting strategy of a player i. As argued in the previous section, we will restrict attention to symmetric strategies i.e. strategies that do not depend on voters' identities, but on their type affiliation only. Therefore, we will drop the superscript i from σ. Given σ, we can define probabilities π_A, π_B and π_C that a randomly chosen voter casts a vote in favor of A, B or C respectively. These probabilities generate a discrete distribution over the set of all feasible outcomes, in which the probability of an outcome (N_A, N_B, N_C) is the trinomial probability[18]: $N!/(N_A!\,N_B!\,N_C!)(\pi_A)^{N_A}(\pi_B)^{N_B}(\pi_C)^{N_C}$. The probability distribution over the set of outcomes together with the payoff structure determine the incentives of each voter to vote A, B or C.

As it is commonly done in the literature, we assume that voters do not play weakly dominated strategies: C voters do not support A or B, while AB and BA voters do not vote for C. Therefore, in equilibrium, C voters will cast their support for C, while AB and BA voters will choose between A and B.

The following lemma, which is similar to the results in Palfrey (1989), Fey (1997) and Myerson and Weber (1993), spells out all (pure strategy) equilibria of the voting game.

Lemma 1. *The voting stage game has three (pure strategy) equilibria:*

1 $\sigma(AB) = \sigma(BA) = A$; $\sigma(C) = C$, *i.e. all majority voters coordinate on A.*
2 $\sigma(AB) = \sigma(BA) = B$; $\sigma(C) = C$, *i.e. all majority voters coordinate on B.*
3 $\sigma(AB) = A$; $\sigma(BA) = B$; $\sigma(C) = C$, *i.e. majority voters divide their votes between A and B.*

As in Palfrey (1989), we will focus on the case of arbitrarily large electorates. Then the equilibrium payoffs are as follows.

Lemma 2. *As the number of voters becomes arbitrarily large ($N \to +\infty$), the equilibrium payoffs of each type are as in Table 2.*

In the absence of a polling stage, the only *attainable* equilibrium is the one in which majority support is split between A and B. This equilibrium prediction is strongly supported by the experimental evidence of Forsythe et al. (1993). Running a series of 48 elections, they find that in 42 elections the majority voters were not able to coordinate their votes (see Table 1). Intuitively, the strong symmetry of the game makes it impossible for them to focus on one of the two majority contenders. Therefore, the only "reasonable" prediction would involve an equal split of their votes between A and B.

3. Equilibria

3.1. Focal equilibrium map

The opinion poll results in an outcome of the form (m_A, m_B, m_C), where m_j denotes the number of j messages, $j \in \{A, B, C\}$. We assume $m_A + m_B + m_C = m$. To economize on notation we denote the outcome of the poll also by m, i.e. $m = (m_A, m_B, m_C)$.

Attainability poses some restrictions even in the voting stage. First, suppose the outcome of the poll is such that $m_A = m_B$. Then the symmetry between AB and BA types is still unbroken. Thus, in this case, whatever weight AB types attach to A in their voting strategy must coincide with the weight BA types attach to B. Even if $m_A \neq m_B$, attainability imposes some restrictions. In fact, the probability weight AB types attach to, say, A in the case that $m_A = j$ and $m_B = k$ with $j \neq k$ must coincide with the weight BA types attach to B in the symmetric case in which $m_A = k$ and $m_B = j$.

The intuitive reasoning behind these restrictions can be explained as follows. Given that A and B are arbitrary labels, the candidates are identified only after and through the polls. If $m_A = m_B$ then there is still no way to distinguish between candidates A and B. If $m_A > m_B$, we can speak of the "leading" (majority) candidate. Voting decisions can depend on this characteristic of candidates, but cannot depend on the arbitrary labels.

Note that, as explained in the previous section, voting A or B is weakly dominated for a C type, while voting C is weakly dominated for AB and BA types. We shall restrict attention to voting equilibria in undominated strategies. Thus C types must vote for C. Voters of types AB and BA have a choice to make between voting for A or B only. This choice can, in principle, depend in an almost arbitrary way on the particular outcome of the poll. In order to discuss this more carefully consider the following definitions.

Let $\mathcal{M}$ be the space of possible *poll outcomes*. A typical element of $\mathcal{M}$ is, thus, (m_A, m_B, m_C), a vector indicating the number of polled voters who declared preference for the candidates A, B, and C, in that order. Let $\mathcal{E} = \{A, B, C\}$ denote the set of

[18] We denote the population size by N, and denote by N_A, N_B and N_C the number of votes candidates A, B and C get. As a convention, we use capital letters to denote the voting stage outcome and lower-case letters to denote the poll outcome.

(three) equilibria of the voting stage (see previous section): (equilibrium A) all majority voters vote for A, (equilibrium B) all majority voters vote for B, and (equilibrium C) majority voters vote for their preferred candidate. The equilibria are, thus, denoted by the candidate that will win the election in the respective equilibrium.

Denote by $e : \mathcal{M} \to \mathcal{E}$ a mapping from the set of outcomes of the poll to the set of equilibria in the voting stage. Call e an *equilibrium map*. In any attainable equilibrium of the polling game, and as intuitively explained above, we must have that the equilibrium map e is *attainable* if it is characterized by these properties:

1 $e(j, k, l) = C$ if $j = k$;
2 if $e(j, k, l) = A$ then $e(k, j, l) = B$;
3 if $e(j, k, l) = B$ then $e(k, j, l) = A$;
4 if $e(j, k, l) = C$ then $e(k, j, l) = C$.

Call an equilibrium map *e responsive* if there are two poll outcomes $m, m' \in \mathcal{M}$ such that $e(m) \neq e(m')$.[19]

Note that there is one and only one non-responsive attainable equilibrium map. This is the one in which, regardless of the poll outcome, the voting game equilibrium is C, the one in which every voter votes for their preferred candidate.[20] There are many responsive equilibria. Consider, for instance, the case in which only 3 voters are asked to participate in the poll. Then there are 10 possible different poll outcomes and, thus, $3^{10} = 59\,049$ possible equilibrium maps, of which $3^4 = 81$ are attainable, of which $3^4 - 1 = 80$ are responsive. We shall concentrate our analysis on one of all these responsive attainable equilibrium maps, that one could argue to be *focal* in the sense of Schelling (1960). As a focal selection among all possible equilibrium maps, we appeal to efficiency. In fact the most efficient equilibrium map can also be characterized as the most preferred equilibrium map of the majority voters given the C types' fixed behavior of voting for C. Thus, this focal equilibrium can be found through "group-think" or the idea that you should assume that everyone else (of your type) behaves like you do (and then behave optimally for yourself and, thus, for the group). If that is an equilibrium, then this is sound. This is reminiscent of Kant's categorical imperative and also of Feddersen and Sandroni (2006) model in which voters' preferences have an ethical component and depend on "how they believe they should behave given the behavior of other agents". The difference here is, that this ethical component is only an equilibrium selection device, that does not enter individuals' preferences.

Efficiency does not yield a unique equilibrium map. Any equilibrium map, in which equilibrium C is chosen only when it has to be chosen to satisfy attainability, is most efficient. For $m = 3$ there are still $2^4 = 16$ such most efficient equilibrium maps. Throughout this paper we shall focus on the following simple and intuitive equilibrium map, denoted by e^*[21]:

$$s_{AB} = \begin{cases} \text{vote A}, & \text{if } m_A = m_B \\ \text{vote A}, & \text{if } m_A > m_B \\ \text{vote B}, & \text{if } m_A < m_B \end{cases} \qquad s_{BA} = \begin{cases} \text{vote B}, & \text{if } m_A = m_B \\ \text{vote B}, & \text{if } m_B > m_A \\ \text{vote A}, & \text{if } m_B < m_A \end{cases}$$

In the remainder of this section, we shall provide two more justifications for using this simple and intuitive equilibrium map e^*.

Consider the following somewhat more realistic setup. There are two ex-ante equally likely states of the world: in one state the proportion of AB types is $q_{AB} = q + \epsilon$ and that of BA types is $q_{BA} = q - \epsilon$, in the other the proportion of AB types is $q_{AB} = q - \epsilon$ and that of BA types $q_{BA} = q + \epsilon$, where ϵ is positive and small. Thus in one state there are slightly more AB voters than BA voters, while in the other state the situation is reversed. This model still exhibits the same symmetries as the original model. Given, however, that the poll is now informative about which of the two states of the world is correct, majority voters can use the poll outcome to focus their attention on the *unique* equilibrium map that is most efficient. The proof of the following lemma can be found in the appendix.

Lemma 3. *Consider the ϵ-perturbed model and let $x = (x_{AB}, x_{BA}, x_C)$ be an attainable strategy profile employed by all polled individuals. If $x_{AB}(A) > x_{AB}(B)$ (and hence, by attainability, $x_{BA}(B) > x_{BA}(A)$), then the unique most efficient attainable equilibrium*

[19] While responsive equilibrium maps are very natural in our very symmetric setup, they are also feasible, possible, maybe even plausible in non-symmetric setups.

[20] A non-responsive attainable equilibrium in our context is somewhat similar to a "babbling" equilibrium in general cheap-talk games, as first studied by Crawford and Sobel (1982). There are, however, some differences. In a babbling equilibrium either the same message is sent in all states of the world or messages are sent randomly, such that ignoring them is indeed a best response. In a non-responsive attainable equilibrium in our context the strategies that guide which messages to send is completely unrestricted as playing the equilibrium C is an equilibrium no matter what happens in the message stage.

[21] For the case of Poisson distributed poll sizes our results still hold for the following more general equilibrium maps:

$$s_{AB} = \begin{cases} \text{vote A}, & \text{if } |m_A - m_B| \le k \\ \text{vote A}, & \text{if } m_A > m_B + k \\ \text{vote B}, & \text{if } m_A < m_B - k \end{cases} \qquad s_{BA} = \begin{cases} \text{vote B}, & \text{if } |m_A - m_B| \le k \\ \text{vote B}, & \text{if } m_B > m_A + k \\ \text{vote A}, & \text{if } m_B < m_A - k \end{cases}$$

where $k \geq 0$, i.e. majority voters coordinate only if either A or B lead in the poll by at least $(k + 1)$ messages.

Table 3
Votes cast by AB and BA type voters in the elections with polls.

Poll outcome	% of votes cast for A	% of votes cast for B	% of votes cast for C
A leads B	90.63	8.33	1.04
Tie between A and B	51.14	40.91	6.82
B leads A	16.35	81.73	0.00

Note. Data source: Forsythe et al. (1993). The labels of voter types and candidates have been changed to fit the description of our setup.

map (from the point of view of the majority voters) is e^* as given above. If $x_{AB}(A) < x_{AB}(B)$ (and hence $x_{BA}(B) < x_{BA}(A)$), then the unique most efficient attainable equilibrium map is e^{**}, derived from the above given e^* by reversing all inequality signs. Finally, if $x_{AB}(A) = x_{AB}(B)$ (and hence $x_{BA}(B) = x_{BA}(A) = x_{AB}(A) = x_{AB}(B)$), then both e^* and e^{**} are most efficient attainable equilibrium maps.

Note that the third case of the above lemma cannot arise in an attainable equilibrium. Suppose that majority voters mix equally between messages A and B, $x_{AB}(A) = x_{AB}(B)$ and $x_{BA}(B) = x_{BA}(A)$, and suppose they use the equilibrium map e^*.[22] In this case an AB type poll participant has incentives to deviate to a strategy that puts weight 1 on message A. This is because, by putting maximum weight on message A, the voter will increase the likelihood of breaking the tie in favor of candidate A and thus making it more likely majority voters coordinate on A in the voting stage, therefore strictly improving her expected payoff.

For the analysis of the polling game it is immaterial whether we use equilibrium map e^* or e^{**}. We choose to work with the more intuitive e^*. Yet, for any polling equilibrium we find, there is a corresponding symmetric equilibrium using equilibrium map e^{**}.

While Lemma 3, thus, shows that equilibrium map e^* is essentially the unique most efficient equilibrium map in the two-state model, the next lemma shows that, under the assumption that equilibrium map e^* is used the set of equilibria of the two-state model coincides with the set of equilibria of our original model, provided ϵ is small enough.

Lemma 4. *Consider the two-state model described above and our original setup proposed in Section 2. If ϵ is small enough, than the set of equilibria is the same in both models.*

The proof of this lemma is based on a continuity argument and follows immediately from the analysis of the equilibria of the limit model, which is carried out in the next sections.

Note that, the poll in the two-state setup, in addition to serving as a coordination device, potentially transmits useful information about the structure of the electorate, and thus it allows majority voters to make (ex-ante) more efficient decisions as a group. More specifically, whenever A or B leads in the poll, majority voters, by using the equilibrium map e^* obtain an (ex-ante) aggregate payoff higher than what they would have obtained had they chosen to coordinate on either A or B with, say, a fifty-fifty chance (as from using the outcome of a coin toss as a coordination device).[23] In this sense, they would prefer the poll as a coordination device to any exogenous coordinate device.

A final justification for focusing attention on the particular equilibrium map e^* is an empirical one. In the experiments of Forsythe et al. (1993) the behavior as summarized in Table 3 was found. Note that almost every majority supporter voted for the leading candidate after the poll identified a leading candidate.[24]

3.2. Pivot events

Let ρ denote the probability of a randomly drawn poll participant to send message A. Thus, given an attainable strategy profile $x \in X$, we have $\rho = qx_{AB}(A) + qx_{BA}(A) + (1-2q)x_C(A)$, or equivalently, using attainability, we have $\rho = q(1 - x_{AB}(C)) + (1-2q)x_C(A)$.

Before we identify and analyze equilibria of the polling game under two different assumptions about the number of polled voters in the next two subsections, we shall here discuss the pivotal events strategic polled individuals face. Note foremost, that any particular voter who is asked to participate in the poll, has $m-1$ "opponents" of uncertain types. To facilitate easy reading we reserved the letter $n = m-1$ for this number. Thus, whenever we speak of poll outcomes we shall use the letter m (and m_A, m_B, and m_C when needed), while whenever we speak of the behavior of a certain polled player's opponents we shall use the letter n (and n_A, n_B, and n_C when needed).

[22] Observe that this implies $x_{AB}(A) < 1$ and $x_{AB}(B) < 1$.

[23] Suppose, for example, that the outcome of the poll is such that $m_A > m_B$. Then the state in which AB group is slightly larger than the BA group is more likely and, since, majority voters would coordinate their support on candidate A, they obtain an expected aggregate payoff higher than what they would have obtained in case of using the outcome of a coin toss as a coordination device.

[24] While the voting stage in Forsythe et al. (1993) is the same in our model, in the sense that it is a second stage after a first polling stage, the opinion polling stage is subtly different in Forsythe et al. (1993) and our model. In Forsythe et al. (1993) there are exactly 4 AB types, 4 BA types, and 6 C types, and all of them participate in both the poll and the voting stage. In contrast, in our model the voters in the opinion poll are randomly chosen from the whole population and may well not have exactly the same number of AB and BA types participating in it. The attainable equilibria in Forsythe et al. (1993) are, thus, different from the ones in our model. In particular, there is certainly no truthful attainable equilibrium in Forsythe et al. (1993). It would be interesting to work out the attainable equilibria of the model in Forsythe et al. (1993) and use the data provided there to test them.

Consider first a polled C type. A C type simply wants to have a tie between the number of A and B messages. Suppose a C type is currently planning to send message C, but is contemplating sending message A.[25] Switching from the truthful message C to lying is detrimental to the C type's payoff if and only if without this C type's message there is a tie between the number of A and B messages. This event can be described as LOSS = $\{n_A = n_B\}$. The incurred payoff loss in this case is exactly 1. Switching from the truthful message C to lying is beneficial to the C type if and only if without this C type's message there is exactly one more B message than there are A messages. This event can be described as GAIN = $\{n_A = n_B - 1\}$. The incurred payoff gain in this case is also exactly 1. The probability of these two events depends on ρ (and on the number of polled people). In all other events the message of this C type has no bearing on the result. The incentive of a C type to lie (send message A) is, thus, given by $\Delta^C(\rho) = P(\text{GAIN}) - P(\text{LOSS})$, or, equivalently, by

$$\Delta^C(\rho) = P(\{n_A = n_B - 1\}) - P(\{n_A = n_B\}),$$

where the probability P of the two events depends on ρ. If $\Delta^C > 0$ a C type will lie and send message A or B,[26] while if $\Delta^C < 0$ a C type will be truthful and send message C.

Note that, without restricting attention to attainable strategies and for a given (realization of the) number of other individuals in the poll, clearly a C type finds that all messages A, B, and C are best-replies to some (non-attainable) strategy profile. Thus, none of these messages are (even weakly) dominated for a C type.

Now consider an AB type.[27] An AB type wants to avoid a tie. There are three events, in which the message sent by an AB type has an impact on the outcome of the election. One in which, without this AB type's message, there is a tie between the number of A and B messages, $\{n_A = n_B\}$, one in which there is exactly one more A message than there are B messages, $\{n_A = n_B + 1\}$, and one in which there is exactly one more B message than there are A messages, $\{n_A = n_B - 1\}$. Suppose an AB type is currently planning to send message A, but is contemplating sending message B. In the event $\{n_A = n_B + 1\}$ this creates a tie and induces a payoff loss of exactly 1. Sending message A induces candidate A to be elected, sending message B induces candidate C to be elected. In the event $\{n_A = n_B - 1\}$ this breaks the tie and induces a payoff gain of exactly v. Sending message A induces a tie and candidate C to be elected, sending message B induces candidate B to be elected. Note, however, that given any attainable strategy profile, both events $\{n_A = n_B + 1\}$ and $\{n_A = n_B - 1\}$ are equally likely. Thus, considering only these two events sending message A is at least as good as sending message B. In the third event, $\{n_A = n_B\}$, sending message A induces candidate A to be elected, while sending message B induces candidate B to be elected. An AB type, however, prefers candidate A over B. Thus, altogether, an AB type will never find it in her interest to send message B.

Sending message B for an AB type, while not weakly dominated if opponents are allowed to choose from all (non-attainable) strategies, is weakly dominated if we restrict attention to attainable strategy profiles of the opponents. These considerations eliminate message B as a possible equilibrium message for AB types. We now explore whether and when sending message C might be a viable option.

Thus, suppose an AB type is currently planning to send message A, but is contemplating sending message C. Switching from the truthful message A to lying and sending message C is detrimental to an AB type's payoff if and only if doing so creates a tie between the number of A and B messages. This happens in the event LOSS = $\{n_A = n_B\}$ with a resulting payoff loss of exactly 1. Switching from the truthful message A to lying and sending message C is beneficial to the AB type if and only if without this AB type's message there is exactly one more B message than there are A messages. This event can be described as GAIN = $\{n_A = n_B - 1\}$. The induced payoff gain is exactly v, as sending message A induces candidate C to be elected, while sending message C induces candidate B to be elected. Thus, the incentives of an AB type to lie are very similar to those of a C type and are given by $\Delta^{AB}(\rho) = vP(\text{GAIN}) - P(\text{LOSS})$, or, equivalently, by

$$\Delta^{AB}(\rho) = vP(\{n_A = n_B - 1\}) - P(\{n_A = n_B\}).$$

Again $\Delta^{AB} > 0$ induces an AB type to lie and send message C, while $\Delta^{AB} < 0$ induces an AB type to be truthful and send message A. BA types will behave analogously.

Note that if $v = 1$, the incentives to lie for AB types (as well as BA types) and C types coincide. That is, in this case, whenever a C type wants to lie also every AB type (and BA type) wants to lie and vice versa. Of course, this means that C types send message A or B, while AB- and BA types send message C, as lying is not the same for both types.

If $v < 1$, then, as $\Delta^{AB}(\rho) \leq \Delta^C(\rho)$ for all ρ, it is still true that whenever a C type is truthful so is every AB- and BA type. However, the opposite is not necessarily true. If a C type has an incentive to lie an AB type might well have an incentive to be truthful.

To finally analyze the equilibria in this model, essentially all we need to do is to compute the probability of the two events for all possible $\rho \in [0, 1/2]$. We will do this for two different assumptions about what individuals know about the poll size when they are asked to report on their preferences in the poll.

[25] For a C type, sending message B is, by symmetry, equivalent to sending message A.

[26] In an attainable strategy profile a C type will, in this case, have to randomize 1/2, 1/2 on A and B.

[27] The discussion for a BA type is, by symmetry, completely analogous.

3.3. The Poisson case

In this section we assume that the poll size is Poisson distributed $P(m)$, with $m > 0$ but finite. By the *decomposition property* of the Poisson distribution (see e.g. Myerson, 1998), the number of A (or B) messages is distributed as $m_A \sim P(\rho m)$ (or $m_B \sim P(\rho m)$), while the number of C messages is $m_C \sim P((1-2\rho)m)$. Interestingly, in the Poisson case, if we fix some particular voter participating in the poll, her "opponents" are also Poisson distributed with the same distribution. Thus, from this voter's point of view we also have $n_A \sim P(\rho m)$, $n_B \sim P(\rho m)$), and $n_C \sim P((1-2\rho)m)$. This is called the *environmental equivalence property* of the Poisson distribution (Myerson, 1998, Theorem 2).

Furthermore, the random variables n_A, n_B and n_C are mutually independent. Given this we obtain the result that the Poisson version of the game has a unique equilibrium.[28]

Proposition 1. *Suppose the poll size follows a Poisson distribution. Then the polling game has a unique attainable equilibrium (the truthful one) in which AB types send message A; BA types send message B and C types send message C.*

The proof of this proposition is given in the appendix. We here provide a brief sketch of the proof. We have already noted that the incentive of a C type to lie is governed by whether Δ^C is smaller than, equal to, or greater than zero, where $\Delta^C(\rho) = P(\mathrm{GAIN}) - P(\mathrm{LOSS})$, and where LOSS is the event of a tie between the number of A messages and B messages sent by opponents, while GAIN is the event of a near tie, in which there is one more B message than there are A messages sent by opponents. Thus, the sign of $\Delta^C(\rho)$ depends on which of the two events, a tie or near tie, is more likely. Note that, given attainability, a randomly chosen opponent is equally likely to send an A message or a B message. This probability is given by ρ, which depends on the attainable strategy profile used by the opponents. Intuitively, if the number of A messages and that of B messages is governed by the same independent distribution, as is the case if the poll size is Poisson distributed, then we should expect a tie more often than a near tie. This is indeed the case and at the heart of the proof. In fact, it turns out that this is true for any ρ. Given this, we, thus have that $\Delta^C(\rho) < 0$ for all $\rho \in [0, 1/2]$. Thus, C types have a strict incentive to send the truthful message C regardless of what everyone else is doing, provided they use an attainable strategy profile. Thus, when restricting attention to attainable strategies C types find message C strictly dominant. Given that $\Delta^C(\rho) < 0$ for all ρ and that $\Delta^{AB} \leq \Delta^C$ we, thus, must also have that $\Delta^{AB} < 0$. Therefore, AB-, and, by symmetry, also BA types, also find sending their truthful messages A and B, respectively, strictly dominant. Thus, there is one and only one equilibrium, which is such that all polled voters send their respective truthful messages.

The analysis can be easily extended to the case of a random q. Suppose that q is not known with certainty to the poll participants but it is known to be drawn from some distribution that puts positive probability mass only on states with $1/4 < q < 1/3$.[29] Using the results of the non-random setup, conditioned on each state, every poll participant has a dominant strategy to send a truthful message. This in turn implies that it is also dominant to send a truthful message if there is uncertainty about the actual q, as long as q stays within the bounds of our original setup i.e. $1/4 < q < 1/3$. Therefore, we proved[30]

Proposition 2. *Suppose the poll size follows a Poisson distribution. Suppose q is randomly drawn from a distribution that puts positive probability mass only on states with* $1/4 < q < 1/3$. *Then the polling game has a unique attainable equilibrium (the truthful one) in which AB types send message A; BA types send message B and C types send message C.*

This proposition provides one answer to the question whether C types can do anything in the polls to try to prevent coordination among majority voters, by perhaps sending misleading messages in the polls. In the Poisson case, Proposition 1 demonstrates that the best C types can do is to be truthful. Individual incentives of C types are such that sending misleading messages A or B is not in their interest. One could ask whether C types could do any better if they acted collectively, in a specific sense. The answer, in the Poisson case is again no. It turns out that in the Poisson case the incentives of an individual C type are very much aligned with that of the group of C types. To see this, suppose instead of C types individually and independently choosing their message in the poll, one representative C type recommends a strategy to all C types, which they then follow blindly. This representative C type shares the incentives of any other C type and assumes that AB and BA types are truthful. It can be shown that this representative C type would also recommend to every C type to be truthful. To summarize

Observation 1. *In the Poisson case, a representative C type, assuming AB and BA types to be truthful and that her recommendation be followed blindly, would recommend to all C types to be truthful.*

The proof follows from simple algebra and it is omitted. This statement, as we see in the next section, is not generally true when poll sizes are fixed.

---

[28] Recall that we assume that voting stage behavior is given by the equilibrium map e^* of Section 3.1.

[29] If $0 < q < 1/3$, then majority voters will condition their behavior only on those states with $1/4 < q < 1/3$. This is because for $q < 1/4$ both of the majority candidates are sure losers. So this case effectively reduces to $1/4 < q < 1/3$. The case of $1/4 < q < 1/2$, however, requires a more careful analysis and it is the subject of a separate paper that discusses such an extension.

[30] Recall that we assume that voting stage behavior is given by the equilibrium map e^* of Section 3.1.

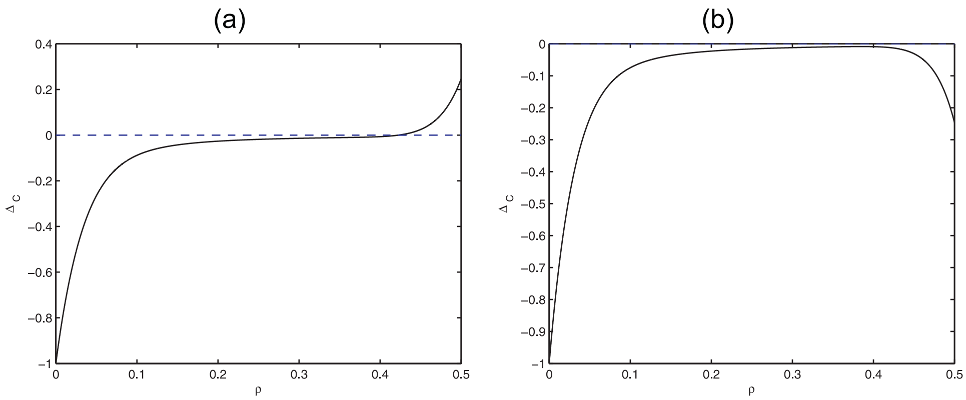

Fig. 1. Incentives $\Delta^C(\rho)$ for even and odd poll sizes (a) $\Delta^C(\rho)$ for an even poll size ($m = 10$). (b) $\Delta^C(\rho)$ for an odd poll size ($m = 11$).

3.4. The trinomial case

Suppose now that the poll size $m \geq 1$ is non-random and commonly known. Then the number of A, B and C messages, m_A, m_B, and m_C follow a trinomial distribution $T_m(\rho, \rho, 1-2\rho)$. From a particular voter's point of view, who is participating in the poll, her "opponents" messages are also from a trinomial, but with $n = m - 1$ instead of m. Thus, n_A, n_B, and n_C follow a trinomial distribution $T_n(\rho, \rho, 1-2\rho)$.

The incentives of a C type to lie are still given by $\Delta^C(\rho) = P(\text{GAIN}) - P(\text{LOSS})$. If $\Delta^C(\rho) > 0$ a C type finds it in her interest to send message A or B. If $\Delta^C(\rho) < 0$ she prefers to send message C. An AB type's incentives to lie are similarly given by $\Delta^{AB}(\rho) = vP(\text{GAIN}) - P(\text{LOSS})$. The two pivot events are given by LOSS = $\{n_A = n_B\}$ and GAIN = $\{n_A = n_B - 1\}$ as before.

In order to find all attainable equilibria of the polling game we need to, as in the Poisson case, determine whether or not a tie, event LOSS, is more likely than a near tie, event GAIN. If it was the case that the random variables n_A and n_B are independent from each other, we could use the same proof as in the Poisson case. This, however, is not true. There is a certain amount of negative correlation between n_A and n_B.[31] This complicates the proofs substantially and, indeed, changes the set of equilibria.

It turns out, it is important to differentiate between two cases. One in which the poll size is an even number, and one in which it is an odd number. This also complicates the writing of this paper, as when the poll size, m, is odd, the number of opponents of any one polled person, n, is even.

Proposition 3. *If the poll size m is odd there is a unique attainable equilibrium (the truthful one), in which C types send message C, AB types message A, and BA types message B.*

Proposition 4. *If the poll size m is even and $v < 1$, then:*

1 *there is no attainable equilibrium in which all types mix completely*,
2 *there is an attainable equilibrium with the property that AB and BA types are truthful and C types randomize (with equal mass on messages A and B)*,
3 *there is an attainable equilibrium with the property that C types lie (attach 1/2 probability on messages A and B each) and AB and BA types randomize*,
4 *there is no equilibrium in which all types lie (C types attaching 1/2 probability on messages A and B each; and AB and BA types sending message C)*,
5 *there is a truthful attainable equilibrium.*

The detailed proofs of these two propositions are given in the appendix. We shall here provide a rough sketch of these proofs. The incentive of a C type, as captured by function Δ^C, for an example of each of the two cases is given in Fig. 1(a) and (b).

Suppose n is odd or even and $\rho = 0$, i.e. everyone sends a C message. Thus, the probability of a tie $n_A = n_B$ is 1, as we have must have $n_A = n_B = 0$, and that of a near tie is 0. Thus, $\Delta^C(0) = -1$. This is bourn out in both Fig. 1(a) and (b) as in both cases $\Delta^C(\rho)$ intersects the y-axis at exactly -1.

Now consider the case of n even (and, thus, an odd sample size $m = n + 1$) and suppose $\rho = 1/2$, i.e. no-one sends message C, and messages A and B are sent with probability 1/2 each. Then a near-tie is must have probability 0. A tie, while not receiving probability 1, does have positive probability. Thus we also have that $\Delta^C(1/2) < 0$ as can be seen on the right side of Fig. 1(b). The difficulty in the proof is to show that not only is $\Delta^C(\rho)$ negative for the two extreme values of ρ of 0 and 1/2, but also *for*

[31] The correlation coefficient is $corr(n_A, n_B) = -\rho/(1-\rho)$.

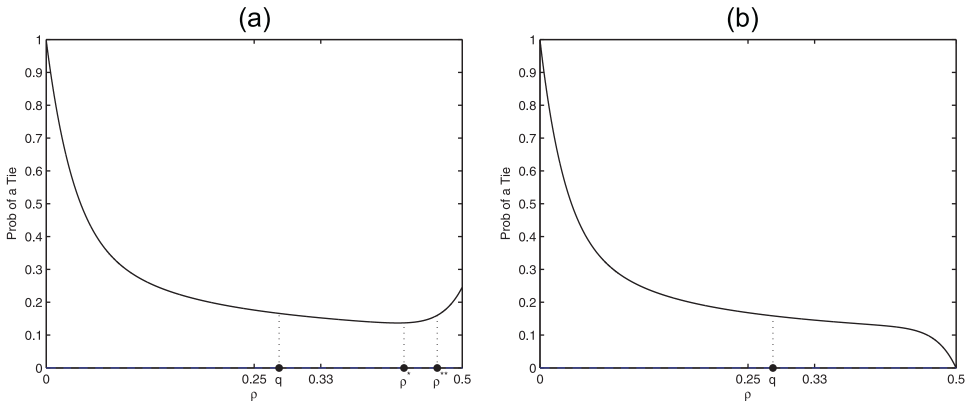

Fig. 2. The probability of a tie in the poll for even and odd poll sizes. (a) Probability of a tie in the poll ($m = 10$). (b) Probability of a tie in the poll ($m = 11$).

all values of ρ in-between. As can be seen in Fig. 1(b) the value of $\Delta^C(\rho)$ is closer to 0 for intermediate values of ρ than for extreme values. Nevertheless, it can be shown, as we do in the appendix, that $\Delta^C(\rho) < 0$ for all values of $\rho \in [0, 1/2]$ and, thus, sending message C is strictly dominant for a C type, just as it was in the Poisson case. As we also have, as in the Poisson case, that $\Delta^{AB}(\rho) \leq \Delta^C(\rho)$, it is also true that message A is strictly dominant for an AB type and message B is strictly dominant for a BA type. Thus, there is a unique attainable equilibrium, and this is truthful. Every type indicates their preferred candidate in the poll. This is Proposition 3.

Now consider the case of n odd (and, thus, an even sample size $m = n + 1$) and suppose $\rho = 1/2$, i.e. no-one sends message C, and messages A and B are sent with probability 1/2 each. As n is odd, the probability of a tie is 0. The probability of a near tie, while not 1, is positive. Thus, we must have that $\Delta^C(1/2) > 0$. This can be seen on the right side of Fig. 1(a).

This and the fact that Δ^C is continuous in ρ, implies that there must be a value ρ^* at which $\Delta^C(\rho^*) = 0$. It can be shown that this ρ^* is unique and must satisfy $\rho^* > 1/3$. This has a few implications. First, it implies that there exists a truthful equilibrium. This is because if individuals report truthfully in the poll their preferred candidate, the induced ρ is equal to q, which is less than or equal to 1/3. However, for any $\rho \leq 1/3$ we now have $\Delta^C(\rho) < 0$. Furthermore, as $\Delta^{AB}(\rho) \leq \Delta^C(\rho)$ for all $\rho \in [0, 1/2]$ we also have $\Delta^{AB}(\rho) < 0$. Thus, indeed no type has an incentive to lie. This is part 5 of Proposition 4.

Now suppose every type were to lie. This means AB types send message C, as do BA types, while C types randomize 1/2, 1/2 between messages A and B. Then ρ, the probability of a randomly chosen voter to send message A, is given by $(1 - 2q)1/2$ as the proportion of C types is $1 - 2q$ and only half of them send message A. The proportion of C types, however, is between 1/3 and 1/2. Thus, $(1 - 2q)1/2 \leq 1/4$, which in turn is less than ρ^*. Thus, if everyone were to lie, the resulting ρ is such that everyone has an incentive to be truthful. Thus, there is no lying equilibrium. This is part 4 of Proposition 4.

Note that, for $\nu < 1$ and $\rho > 0$, $\Delta^{AB}(\rho)$ is strictly less than $\Delta^C(\rho)$. This implies that $\Delta^{AB}(\rho)$ and $\Delta^C(\rho)$ cannot be equal to zero for the same value of ρ. This implies that there is no equilibrium in which both C and AB types randomize, proving part 1 of Proposition 4.

Furthermore the value ρ^* identified above characterizes an equilibrium, in which C types randomize between message C and an even mixture of messages A and B (as they are indifferent), and AB types, who, at ρ^*, must have an incentive to be truthful send message A, while BA types send message B. This is part 2 of Proposition 4.

Finally, again as $\Delta^{AB}(\rho) < \Delta^C(\rho)$, for $\nu < 1$ and $\rho > 0$, and by the fact that $\Delta^{AB}(1/2) = \nu\Delta^C(1/2) > 0$, there is another value of ρ, $\rho^{**} > \rho^*$, which again can be shown to be unique, such that $\Delta^{AB}(\rho^{**}) = 0$, while, necessarily, $\Delta^C(\rho^{**}) > 0$. This characterizes another equilibrium, in which AB and BA types randomize and C types lie by sending a 1/2, 1/2 mixture of messages A and B. This is part 3 of Proposition 4.

We now turn to the question as to what strategy a representative C type would recommend to all C types if they follow this recommendation blindly, and assuming that AB and BA types are truthful. In the case of an odd poll size, as in the Poisson case, the answer is again that the representative C type will recommend to be truthful.

Observation 2. *If the poll size m is odd, a representative C type, assuming AB and BA types to be truthful and that her recommendation be followed blindly, would recommend to all C types to be truthful.*

The formal proof of this observation is omitted but the intuition can be seen in Fig. 2(b). Fig. 2(b) shows that the probability of a tie, $\{m_A = m_B\}$, is decreasing in ρ, the probability of a randomly chosen polled person to send message A. C types would like to maximize the probability of a tie between A and B, or, equivalently, they want to have ρ as small as possible. Therefore, given the truthful behavior of majority voters, the best they can do is to be truthful as well, in which case $\rho = q$.

In the case of an even poll size, of the three equilibria, the truthful one is still perhaps the most immediate and perhaps most focal. However, one can ask the question which type prefers which equilibrium as well as, again, what behavior in the poll a representative C type would recommend to all C types, assuming they follow this recommendation blindly, and assuming AB and BA types are truthful. This amounts to comparing the probability of a tie in the poll between A and B messages across the equilibria, and generally, for all values of ρ, as is done in the following proposition.

Proposition 5. *Suppose poll size m is even and* $0 < \nu < 1$*, then:*

1 *the likelihood of coordination is largest in the "mixing" equilibrium (AB and BA's are truthful, and C's are mixing across A, B and C).*
2 *for large values of* ν*, the likelihood of coordination is larger in the lying equilibrium (C types are lying while AB and BA are mixing) than in the truthful equilibrium.*
3 *for low values of* ν*, the likelihood of coordination is smaller in the lying equilibrium (C types are lying while AB and BA are mixing) than in the truthful equilibrium.*

Parts of this proposition are reflected graphically in Fig. 2(a), in which the probability of a tie in the poll, $\{m_A = m_B\}$, is drawn as a function of ρ, the probability of a randomly selected polled person to send message A.

It is interesting to note, that in the case of majority voters sending truthful messages, the representative C type would recommend lying to all C types, in order to increase the chance of a tie.

Observation 3. *If the poll size m is even, a representative C type, assuming AB and BA types to be truthful and that her recommendation be followed blindly, would recommend to all C types to be mixing 50–50% between A and B.*

The proof is given in the appendix. Indeed, if AB and BA types were truthful themselves, then all C types would also find it in their interest to follow this recommendation, provided all other C types did so as well. However, in that case, AB and BA types would no longer find it in their interest to be truthful and would prefer to lie themselves. Thus, this strategy profile in which no-one sends message C is not an equilibrium.

Fixing the behavior of AB and BA types to be truthful and comparing the two equilibria that have this feature, we note that the one, in which also C types are truthful has a higher probability of a tie, and is, thus, between the two, preferred by all C types. For small values of ν, the cost of mis-coordination for majority voters, C types would prefer if they somehow could induce the lying equilibrium, in which AB and BA types randomize, while C types lie.

There is one feature that all equilibria share.

Corollary 1. *In any attainable equilibrium of any of the models in this section, the expected proportion of C-messages is less than or equal to the proportion* $(1 - 2q)$ *of C types. Thus C types are never over-represented in polls.*

In fact, we can be more specific. In any truthful equilibrium, the expected proportion of C messages is exactly equal, of course, to the proportion of C types. In the other two equilibria in the case of a fixed and even poll size, the expected proportion of C messages is strictly less than the proportion of C types. Moreover, as the (even) poll size tends to infinity the expected proportion of C messages tends to zero. Thus, in large polls with a fixed and even poll size there are two equilibria (in addition to the truthful one), in which hardly anyone professes a preference for the minority candidate, even though the actual proportion of voters with such a preference is between a third and a half.

3.5. Arbitrarily distributed poll size

Using the results we obtained in the previous section for fixed odd and even poll sizes, we immediately have the following corollary.

Corollary 2. *Suppose the poll size follows an arbitrary discrete distribution. Then the polling game has a truthful attainable equilibrium.*

The proof of this corollary follows from the fact that both the odd and even case feature a truthful equilibrium. Suppose that the size of the poll is uncertain and follows some discrete distribution $F(m)$. Conditional on the poll size being a fixed number and on every other poll participant sending a truthful message, it is a best response for a voter participating in the poll to also send a truthful message. Since this holds for any fixed poll size, averaging across the entire possible states of the distribution, it is still a best response for a voter participating in the poll to send a truthful message, if all other poll participants are sending truthful messages. Thus every poll participant sending a truthful message constitutes an equilibrium.

The corollary shows that even if a voter selected to participate in the poll is not aware of the actual poll size, if she believes that other voters truthfully report their preferences, she will also have incentives to truthfully report her preferences.

We also know that for some distributions of poll size, such as the Poisson, the truthful equilibrium is the only one. We know furthermore that there are distributions, such as one putting probability 1 on a fixed even poll size, for which there are multiple (three) equilibria. We do not yet know whether other equilibria could emerge for some particular poll size distributions.

4. Conclusion

The main purpose of our paper was to understand how a single opinion poll can serve as a coordination mechanism in elections. In order to understand this role of public opinion polls, we used a voting model in which there are two similarly appealing majority candidates and a minority candidate. We have first shown that, in this model, polls are *critical* to achieving coordination among majority supporters. Absent any coordination device, the only attainable equilibrium of the voting game features a non-coordinated outcome in which the minority candidate wins the election. By introducing a polling stage before the actual voting game, we have shown how majority voters can use the polling information to coordinate their actions and, thus, prevent the minority candidate from being elected into office.

Our analysis suggests that polls have a significant influence on election outcomes. They play the role of a selection device helping voters focus attention on one of the equilibria. Since, typically, polls are much smaller as compared to the total electorate, a straw vote in the poll is more likely to influence the election outcome than an actual vote. If one were to make an analogy to sequential elections (for example primaries) and think of the polling stage as a first voting round, then our results would suggest that earlier rounds have a disproportionately large impact on the final outcome.

Appendix A. Proof of Lemma 3

Given a poll outcome $m=(m_A, m_B, m_C)$, after some algebra, it can be shown that

$$P(q_{AB} = q + \epsilon, q_{BA} = q - \epsilon \mid m) > P(q_{AB} = q - \epsilon, q_{BA} = q + \epsilon \mid m)$$

if and only if

$$(m_A - m_B)[\log(\rho + \epsilon(x_{AB}(A) - x_{AB}(B))) - \log(\rho - \epsilon(x_{AB}(A) - x_{AB}(B)))] > 0,$$

where $\rho = q(x_{AB}(A) + x_{AB}(B)) + (1 - 2q)x_C(A)$. For the first part of the lemma, note that if the poll participants use an attainable strategy profile with the property that $x_{AB}(A) > x_{AB}(B)$, then an outcome of the poll with $m_A > m_B$ indicates that the state $(q_{AB} = q + \epsilon, q_{BA} = q - \epsilon)$ is more likely than $(q_{AB} = q - \epsilon, q_{BA} = q + \epsilon)$ and therefore the efficient equilibrium map is for the majority voters to coordinate on candidate A.[32] Similarly, if the poll outcome has $m_B > m_A$, then the state $(q_{AB} = q - \epsilon, q_{BA} = q + \epsilon)$ is more likely and therefore the efficient equilibrium map is for majority voters to coordinate on candidate B; finally if $m_A = m_B$, attainability requires the two groups of AB and BA voters to play the no-coordination equilibrium.

The second part of the lemma can be proved by a similar argument. Finally, for the last part, if $x_{AB}(A) = x_{AB}(B)$ then $P(q_{AB} = q + \epsilon, q_{BA} = q - \epsilon | m) = P(q_{AB} = q - \epsilon, q_{BA} = q + \epsilon | m)$ independent of the poll outcome. In this case, the poll does not transmit any information about the structure of the electorate. Therefore, no matter the outcome of the poll, the after-poll beliefs of majority voters about the two states are identical to their prior beliefs (i.e. both states are equally likely). Thus, in this case, the efficient equilibrium map is for the majority voters to coordinate on either A or B, whenever either A or B lead in the poll. □

Appendix B. Poll participants can be ignored in the voting stage

Given that the poll participants potentially have less information (after the poll results are made public) than those that did not participate in the poll, theoretically we should allow them to use different strategies in the voting stage. However, as we show below, a finite number of voters (i.e. the poll participants) cannot change the election outcome as the size of the electorate grows arbitrarily large, and therefore they can be ignored in the voting stage. The strategies used in our analysis can be thought of representing the voting strategies of the non-poll participants.

As shown in Lemma 1 the voting game alone has three equilibria. Consider, for example, the equilibrium in which all majority group voters go for A, while C voters vote for C. Let us denote by $(\mathbf{1}_{AB}, \mathbf{1}_{BA}, \mathbf{1}_C)$ the results of a trinomial draw where, for instance, $\mathbf{1}_{AB}$ indicates that an AB type was drawn. The number of AB and BA type voters after N draws is $N_{AB} = \sum_{n=1}^{N} \mathbf{1}_{AB}$ and $N_{BA} = \sum_{n=1}^{N} \mathbf{1}_{BA}$ respectively. Therefore, in the equilibrium considered, A gets a number of $N_{AB} + N_{BA}$ votes, B gets 0 votes while C gets $N_C = N - (N_{AB} + N_{BA})$ votes. Then, we have the following result.

Lemma 5. *Let k be some integer constant. Then*

$$\lim_{N\to+\infty} P(N_{AB} + N_{BA} > N_C + k) = 1$$
$$\lim_{N\to+\infty} P(N_{AB} + N_{BA} > k) = 1$$

Therefore a finite number of voters that would either switch to B or C cannot change the outcome of the election in the limit.

Proof. The first probability can be re-written as

$$\begin{aligned} P(N_{AB} + N_{BA} > N_C + k) &= P(N_{AB} + N_{BA} > N - (N_{AB} + N_{BA}) + k) \\ &= P\left(\frac{N_{AB} + N_{BA}}{N} > \frac{1}{2}\left(1 + \frac{k}{N}\right)\right) \end{aligned}$$

[32] This is because it is more likely that there are more AB voters in the electorate than BA voters, and therefore the equilibrium that involves coordination on A yields a higher expected aggregate payoff for majority voters than the equilibrium in which they coordinate on B.

By the law of large numbers $(N_{AB}+N_{BA})/N \to 2q$ and $1/2(1+k/N) \to 1/2$, and since $2q > 1/2$ the first limit immediately follows from a straightforward application of the definition of convergence in probability. For the second limit, we can use a similar argument to write

$$P(N_{AB}+N_{BA} > k) = P\left(\frac{N_{AB}+N_{BA}}{N} > \frac{k}{N}\right)$$

and since $(N_{AB}+N_{BA})/N \to 2q$ while $k/N \to 0$ we obtain $\lim_{N\to+\infty} P(N_{AB}+N_{BA} > k) = 1.$ □

Appendix C. Proof of Proposition 1

In order to prove Proposition 1 the following lemma is useful.

Lemma 6. *Consider a randomly chosen individual (AB- or C type) to participate in the poll, where the poll size is Poisson, $P(m)$, distributed. Fix the probability of a randomly chosen "opponent" to send message A (or equivalently B) of ρ (induced by a given strategy profile $x \in X$). Let* GAIN = $\{n_A = n_B - 1\}$, *and* LOSS = $\{n_A = n_B\}$. *Then the incentive of a C type to lie is given by*

$$\begin{aligned}\Delta^C &= P(\mathrm{GAIN}) - P(\mathrm{LOSS})\\ &= \sum_{k=0}^{\infty} e^{-\rho m}\frac{(\rho m)^k}{k!}e^{-\rho m}\frac{(\rho m)^{k+1}}{(k+1)!} - \sum_{k=0}^{\infty} e^{-\rho m}\frac{(\rho m)^k}{k!}e^{-\rho m}\frac{(\rho m)^k}{k!},\end{aligned}$$

The incentive of an AB type to lie is given by

$$\begin{aligned}\Delta^{AB} &= vP(\mathrm{GAIN}) - P(\mathrm{LOSS})\\ &= v\sum_{k=0}^{\infty} e^{-\rho m}\frac{(\rho m)^k}{k!}e^{-\rho m}\frac{(\rho m)^{k+1}}{(k+1)!} - \sum_{k=0}^{\infty} e^{-\rho m}\frac{(\rho m)^k}{k!}e^{-\rho m}\frac{(\rho m)^k}{k!}.\end{aligned}$$

Proof. The proof of this Lemma follows directly from the definition of Δ^C and Δ^{AB} and the probability mass function of the Poisson distribution. □

Lemma 6 allows us to prove Proposition 1 as follows. To simplify notation we define

$$p_k = e^{-\rho m}\frac{(\rho m)^k}{k!}, \quad k \geq 0.$$

Then the incentives to lie, by Lemma 6, can be written as

$$\begin{aligned}\Delta^C &= \sum_{k=0}^{\infty} p_k p_{k+1} - \sum_{k=0}^{\infty} p_k^2\\ \Delta^{AB} &= v\sum_{k=0}^{\infty} p_k p_{k+1} - \sum_{k=0}^{\infty} p_k^2\end{aligned}$$

Since $v < 1$, we have $\Delta^{AB} < \Delta^C$. So if we show that $\Delta^C < 0$ for all $\rho \in (0, 1/2)$, then the unique equilibrium must involve every polled voter sending a truthful message. To show $\Delta^C < 0$ we note that for each $l > 0$:

$$\sum_{k=0}^{l-1} p_k p_{k+1} - \sum_{k=0}^{l} p_k^2 < -\frac{1}{2}p_0^2$$

Letting $l \to \infty$, we obtain

$$\sum_{k=0}^{\infty} p_k p_{k+1} - \sum_{k=0}^{\infty} p_k^2 \leq -\frac{1}{2}p_0^2 < 0$$

Therefore, $\Delta^C < 0$ for all $\rho \in (0, 1/2)$ as required. □

Appendix D. Proof of Propositions 3 and 4

In order to prove Propositions 3 and 4 a few lemmata are needed.

Lemma 7. *Consider a randomly chosen individual (AB- or C type) to participate in the poll. Fix the probability of a randomly chosen "opponent" to send message A (or equivalently B) of* ρ *(induced by a given strategy profile* $x \in X$*). Let* GAIN = $\{n_A = n_B - 1\}$ *and* LOSS = $\{n_A = n_B\}$*. Then, the incentive of a C type to lie by sending message A is given by*

$$\begin{aligned}\Delta^C &= P(\text{GAIN}) - P(\text{LOSS})\\ &= \sum_{k=0}^{\lfloor \frac{n-1}{2} \rfloor} \frac{n!}{k!(k+1)!(n-2k-1)!}\rho^{2k+1}(1-2\rho)^{n-2k-1} - \sum_{k=0}^{\lfloor \frac{n}{2} \rfloor} \frac{n!}{k!^2(n-2k)!}\rho^{2k}(1-2\rho)^{n-2k},\end{aligned}$$

The incentive of an AB type to lie is given by

$$\begin{aligned}\Delta^{AB} &= vP(\text{GAIN}) - P(\text{LOSS})\\ &= v\sum_{k=0}^{\lfloor \frac{n-1}{2} \rfloor} \frac{n!}{k!(k+1)!(n-2k-1)!}\rho^{2k+1}(1-2\rho)^{n-2k-1} - \sum_{k=0}^{\lfloor \frac{n}{2} \rfloor} \frac{n!}{k!^2(n-2k)!}\rho^{2k}(1-2\rho)^{n-2k}.\end{aligned}$$

Proof. The proof of this Lemma follows directly from the definition of Δ^C and Δ^{AB} and the probability mass function of the Trinomial distribution. □

Because most of the proofs use an induction argument on n, we will use notation Δ_n^C and Δ_n^{AB} and, occasionally, $\Delta_n^C(\rho)$ and $\Delta_n^{AB}(\rho)$ to emphasize the dependence on n and ρ. To make the proofs easier to read we define

$$\begin{aligned}a_n &= \sum_{k=0}^{\lfloor \frac{n}{2} \rfloor} \frac{n!}{k!k!(n-2k)!}\rho^{2k}(1-2\rho)^{n-2k}\\ b_n &= \sum_{k=0}^{\lfloor \frac{n-1}{2} \rfloor} \frac{n!}{k!(k+1)!(n-2k-1)!}\rho^{2k+1}(1-2\rho)^{n-2k-1}.\end{aligned}$$

Thus, a_n is the probability of a tie (the event LOSS), while b_n is the probability of a near-tie (the event GAIN), for a given number of "other" poll participants. With this notation, we can write $\Delta_n^C = b_n - a_n$ and $\Delta_n^{AB} = vb_n - a_n$ for all $n \geq 0$. The following lemmata establish properties of a_n and b_n needed to analyze voters' incentives.

Lemma 8. *For all* $n \geq 0$ *we have the following recurrence equations*

$$\begin{aligned}a_{n+1} &= (1-2\rho)a_n + 2\rho b_n\\ b_{n+1} &= \frac{n+1}{n+2}[(1-2\rho)b_n + 2\rho a_n]\end{aligned}$$

with the initial starting points $a_0 = 1$ *and* $b_0 = 0$.

The proof of Lemma 8 follows from some straightforward algebra and is omitted. Note that, while the second recurrence relation does not, the first one has an intuitive explanation, which is as follows. The probability of having a tie in $(n+1)$ trials is the probability of having a tie in n trials and a C draw in the $(n+1)$st plus the probability of having a near-tie in n trials and having the appropriate message A or B in the $(n+1)$st trial such that a tie is created.

Lemma 9. *Let* $0 < \rho < 1/4$ *and* $b_n < a_n$ *(or* $1/4 < \rho < 1/2$ *and* $b_n > a_n$*). Then*

$$b_{n+1} < \frac{1}{2}a_n + \frac{1}{2}b_n < a_{n+1}.$$

Proof. Since $b_n < a_n$ and $0 < \rho < 1/4$, we have $(1-2\rho)b_n + 2\rho a_n < 1/2 a_n + 1/2 b_n < (1-2\rho)a_n + 2\rho b_n$. Using Lemma 8 yields the result. □

Lemma 10. *For each* $n \geq 0$*, we have the following recurrence equation:*

$$a_{n+2} - b_{n+2} = \alpha_n a_n - \beta_n b_n,$$

where $\alpha_n > 0$ *and* $\alpha_n - \beta_n = 2(1/(n+3) + 1/(n+2))(1/(2+1/(n+2)) - \rho)$.

The proof of Lemma 10 is based on straightforward algebra and is omitted.

Lemma 11. *If* $\rho \in (0, 1/3)$ *then* $a_n > b_n$ *for all* $n \geq 0$.

Proof. The proof is by induction. For $n = 0$ the statement is: $a_0 > b_0$ for $\rho \in (0, 1/3)$; while for $n = 1$ it reads: $a_1 > b_1$ for $\rho \in (0, 1/3)$. These can be immediately verified to be true.

Suppose it holds for some n and we want to prove it holds for $n + 2$: $a_{n+2} > b_{n+2}$ for $\rho \in (0, 1/3)$. Using Lemma 10, we can write $a_{n+2} - b_{n+2} = \alpha_n a_n - \beta_n b_n$, where $\alpha_n > 0$ and $\alpha_n - \beta_n = 2(1/(n+3) + 1/(n+2))(1/(2 + 1/(n+2)) - \rho)$.

Since $\rho \in (0, 1/3)$, we obtain $\alpha_n > \beta_n$ which combined with the induction hypothesis $a_n > b_n$ immediately implies $\alpha_n a_n > \beta_n b_n$ or $a_{n+2} > b_{n+2}$. □

Using the above results, we are now in a position to determine the sign of both Δ_n^C and Δ_n^{AB}. This is done in the next three lemmata.

Lemma 12.

1 *For* $\nu < 1$, $\Delta_n^C(\rho) > \Delta_n^{AB}(\rho)$ *for all* $\rho \in (0, 1/2)$.
2 *If n is odd,* $\Delta_n^C(\rho) < 0$ *for* $\rho \in [0, 1/3]$ *and* $\Delta_n^C(\rho) > 0$ *for* $\rho \geq 1/(2 + 1/n)$.
3 *If n is even,* $\Delta_n^C(\rho) < 0$, $\Delta_n^{AB}(\rho) < 0$ *for all* $\rho \in [0, 1/2]$.

D.1. Proof of Part 1

This follows immediately from Lemma 7. □

D.2. Proof of Part 2

Suppose n is odd. Since $\Delta_n^C = b_n - a_n$, using Lemma 11, we obtain $\Delta_n^C(\rho) < 0$ for all $\rho \leq 1/3$. For the second claim, we need to prove that $a_n < b_n$ for $\rho \geq 1/(2 + 1/n)$. The proof is by induction.

For $n = 1$ the statement is: $a_1 < b_1$ for $\rho > 1/3$. Suppose it holds for some n and we prove it holds for $n + 2$ i.e. $a_{n+2} < b_{n+2}$ for $\rho > 1/(2 + 1/(n+2))$. Using Lemma 10, we can write $a_{n+2} - b_{n+2} = \alpha_n a_n - \beta_n b_n$, where $\alpha_n > 0$ and $\alpha_n - \beta_n = 2(1/(n+3) + 1/(n+2))(1/(2 + 1/(n+2)) - \rho)$. Because $\rho > 1/(2 + 1/(n+2))$, we have $0 < \alpha_n < \beta_n$. On the other hand, $\rho > 1/(2 + 1/(n+2)) > 1/(2 + 1/n)$ implies $a_n < b_n$ by the induction hypothesis. Combining these two last inequalities, we obtain $a_{n+2} < b_{n+2}$ for all $\rho > 1/(2 + 1/(n+2))$. Since $\Delta_n^C = b_n - a_n$, we obtain $\Delta_n^C(\rho) > 0$ for all odd n and $\rho \geq 1/(2 + 1/n)$. □

D.3. Proof of Part 3

Suppose n is even. We need to prove that $a_n > b_n$ for all $\rho \in (0, 1/2)$. Again we do the proof by induction.

For $n = 0$ the statement is: $a_0 > b_0$ for all $\rho \in (0, 1/2)$. Suppose it holds for some even n and we want to prove it holds for $n + 2$, i.e. $a_{n+2} > b_{n+2}$ for all $\rho \in (0, 1/2)$. Using Lemma 10, we can write $a_{n+2} - b_{n+2} = \alpha_n a_n - \beta_n b_n$, where $\alpha_n > 0$ and $\alpha_n - \beta_n = 2(1/(n+3) + 1/(n+2))(1/2 + 1/(n+2)) - \rho)$.

If $\rho < 1/(2 + 1/(n+2))$ then $\alpha_n > \beta_n$. Combining this inequality with the induction hypothesis $a_n > b_n$ we obtain $a_{n+2} > b_{n+2}$ for $\rho < 1/(2 + 1/(n+2))$. Since n is even, then $n + 1$ is odd. By Part 2 we have $a_{n+1} < b_{n+1}$ for $\rho > 1/(2 + 1/(n+1))$. Because $1/(2 + 1/(n+1)) > 1/4$ we apply Lemma 9 to obtain $a_{n+2} > b_{n+2}$ for $\rho > 1/(2 + 1/(n+1))$. Because $\rho < 1/(2 + 1/(n+2))$ and $\rho > 1/(2 + 1/(n+1))$ cover all ρ's in $(0, 1/2)$ we are done. Since $\Delta_n^C = b_n - a_n$, we have that $\Delta_n^C(\rho) < 0$ for all even n and $\rho \leq 1/2$. □

Lemma 13. *Consider the case of n odd, with* $n \geq 1$. *Then* $\Delta_{n+2}^C(\rho) < 0$ *for* $\rho < \rho_{n+2}^*$ *and* $\Delta_{n+2}^C(\rho) > 0$ *for* $\rho > \rho_{n+2}^*$, *where* ρ_n^* *increasing in n with* $1/3 \leq \rho_n^* < 1/2$ *and* $\rho_1^* = 1/3$.

Proof. First we note that by Lemma 12 part 2, $\Delta_n^C = b_n - a_n > 0$ for all odd n and $\rho > 1/(2 + 1/n)$. Therefore, we must have $\rho_n^* \leq 1/(2 + 1/n$. The proof is by induction.

If $n = 1$, then $\Delta_1^C = b_1 - a_1 = 3\rho - 1$. Suppose the statement is true for some n, and we want to show it holds for $n + 2$. Using Lemma 10, we have $\Delta_{n+2}^C = b_{n+2} - a_{n+2} = \beta_n b_n - \alpha_n a_n$, where $\alpha_n > 0$ and $\alpha_n - \beta_n = 2(1/(n+3) + 1/(n+2))(1/(2 + 1/(n+2)) - \rho)$.

For any $\rho < \rho_n^*$, by the induction hypothesis we have $\Delta_n^C < 0$ or $b_n < a_n$. At the same time $\beta_n < \alpha_n$ (this is because $\rho < \rho_n^* \leq 1/(2 + 1/n) < 1/(2 + 1/(n+2))$). Combine these two inequalities to obtain $\beta_n b_n < \alpha_n a_n$ or $\Delta_{n+2}^C < 0$ for all $\rho < \rho_n^*$. Using Lemma 12 part 2, we have $\Delta_{n+2}^C > 0$ for $\rho > 1/(2 + 1/(n+2))$. If we show that Δ_{n+2}^C is increasing in ρ for $\rho > \rho_n^*$, then there must exist a threshold $\rho_{n+2}^* > \rho_n^*$ such that $\Delta_{n+2}^C(\rho) < 0$ for $\rho < \rho_{n+2}^*$ and $\Delta_{n+2}^C(\rho) > 0$ for $\rho > \rho_{n+2}^*$.

To show that Δ^C_{n+2} is increasing in ρ, we use Lemma 15 to compute its derivative as

$$\begin{aligned}
\frac{d\Delta^C_{n+2}}{d\rho} &= \frac{db_{n+2}}{d\rho} - \frac{da_{n+2}}{d\rho} \\
&= 4(n+2)(a_{n+1} - b_{n+1}) - \frac{1}{\rho} b_{n+2} \\
&= 4(n+2)(a_{n+1} - b_{n+1}) - \frac{1}{\rho}\frac{n+2}{n+3}[(1-2\rho)b_{n+1} + 2\rho a_{n+1}] \\
&= \left(4 - 2\frac{1}{n+3}\right) a_{n+1} - \left(4 + \frac{1}{\rho}\frac{1}{n+3}(1-2\rho)\right) b_{n+1} \\
&= \left(4 - 2\frac{1}{n+3}\right) [(1-2\rho)a_n + 2\rho b_n] - \left(4 + \frac{1}{\rho}\frac{1}{n+3}(1-2\rho)\right) \frac{n+1}{n+2}[(1-2\rho)b_n + 2\rho a_n] \\
&= \widehat{\beta}_n b_n - \widehat{\alpha}_n a_n,
\end{aligned}$$

where

$$\widehat{\beta}_n = \left(4 - 2\frac{1}{n+3}\right) 2\rho - \left(4 + \frac{1}{\rho}\frac{1}{n+3}(1-2\rho)\right) \frac{n+1}{n+2}(1-2\rho)$$

$$\widehat{\alpha}_n = \left(4 + \frac{1}{\rho}\frac{1}{n+3}(1-2\rho)\right) \frac{n+1}{n+2} 2\rho - \left(4 - 2\frac{1}{n+3}\right)(1-2\rho)$$

For all $\rho > \rho^*_n > 1/3$ we have $\widehat{\beta}_n > \widehat{\alpha}_n > 0$. At the same time we know $b_n > a_n$ (by the induction hypothesis), therefore $\widehat{\beta}_n b_n > \widehat{\alpha}_n a_n$. But this means $(d\Delta^C_{n+2})/(d\rho) > 0$ for all $\rho > \rho^*_n$. □

We note that the proof implies that Δ^C_n is increasing in ρ for $\rho > \rho^*_n$. Using this observation we can prove the following

Lemma 14. *Consider the case of n odd, with $n \geq 1$. Then $\Delta^{AB}_n(\rho) < 0$ for $\rho < \rho^{**}_n$ and $\Delta^{AB}_n(\rho) > 0$ for $\rho > \rho^{**}_n$, where $\rho^*_n < \rho^{**}_n <$ 1/2.*

Proof. Note that $\Delta^{AB}_n = \Delta^C_n - (1-v)b_n$. From the previous lemma we have $\Delta^C_n < 0$ for all $\rho < \rho^*_n$, which implies $\Delta^{AB}_n < 0$ for all $\rho < \rho^*_n$. Consider now the case of $\rho > \rho^*_n$. Using Lemma 15 we can compute the derivative of Δ^{AB}_n as

$$\begin{aligned}
\frac{d\Delta^{AB}_n}{d\rho} &= \frac{d\Delta^C_n}{d\rho} - (1-v)\frac{db_n}{d\rho} \\
&= 4n(a_{n-1} - b_{n-1}) - \frac{1}{\rho} b_n - (1-v)[2n(a_{n-1} - b_{n-1}) - \frac{1}{\rho} b_n]
\end{aligned}$$

For $v = 1$ we have $(d\Delta^{AB}_n)/(d\rho) = (d\Delta^C_n)/(d\rho) > 0$ for $\rho > \rho^*_n$, while for $v = 0$ we have

$$\begin{aligned}
\frac{d\Delta^{AB}_n}{d\rho} &= 4n(a_{n-1} - b_{n-1}) - \frac{1}{\rho} b_n - [2n(a_{n-1} - b_{n-1}) - \frac{1}{\rho} b_n] \\
&= 2n(a_{n-1} - b_{n-1}) > 0
\end{aligned}$$

Because of the linearity (in v), we obtain $(d\Delta^{AB}_n)/(d\rho) > 0$ for all $v \in (0, 1)$, meaning Δ^{AB}_n is increasing in ρ for $\rho > \rho^*_n$. Since $\Delta^{AB}_n(1/2) > 0$, we conclude that there exists a unique $\rho^{**}_n > \rho^*_n$ with the desired property. □

Previous lemmata finally allow us to prove Propositions 3 and 4.

D.4. Proof of Proposition 3

Follows directly from Lemma 12 part 3.

D.4.1. Proof of Proposition 4

Part 1 follows from Lemma 12 part 1. To see part 2, suppose C types randomize. From Lemma 13, there exists a unique $\rho^*_n > 1/3$ such that $\Delta^C_n = 0$. For this particular ρ, using Lemma 12 part 1, we have $\Delta^{AB}_n < 0$. Thus AB and BA types must be truthful and we have $x_{AB}(A) = x_{BA}(B) = 1$. Thus $\rho = q + (1-2q)x_C(A) \in [q, 1/2]$ for $x_C(A) \in [0, 1/2]$, where $x_C(A) = 1/2$ is the maximal probability a C type can attach to A in an attainable equilibrium. Since $\rho^*_n > 1/3$, there exists a unique $x_C(A) > 0$ that makes $\rho = \rho^*_n$. To see part 3 suppose AB (and BA) types randomize. By Lemma 14, there exists a unique ρ^{**}_n such that $\Delta^{AB}_n = 0$. For this particular ρ, using Lemma 12 part 1, we have $\Delta^C_n > 0$. Thus C types must lie i.e. play $x_C(A) = x_C(B) = 1/2$. Thus, we have $\rho = 1/2 - qx_{AB}(C) \in [1/2 - q, 1/2]$ for $x_{AB}(C) \in [0, 1]$. Since $\rho^{**}_n > 1/3$, there exists a unique $x_{AB}(C) > 0$ that makes $\rho = \rho^{**}_n$. To see part 4 note that in order to have an equilibrium in which all types lie we would have $\rho = 1/2 - q$, which is in the interval [1/6, 1/4] for $q \in [1/4, 1/3]$. Part 4 then follows from Lemma 12 part 2. Part 5, finally, follows from Lemma 12 part 2 and part 1. □

Appendix E. Proof of Proposition 5

The proof uses the notation of the previous section (Proof of Lemma 12). First we note the following results

Lemma 15. *For all $n \geq 1$ and $0 < \rho < 1/2$, we have*

$$\frac{da_n}{d\rho} = 2n(b_{n-1} - a_{n-1}) \quad \text{and} \quad \frac{db_n}{d\rho} = -2n(b_{n-1} - a_{n-1}) - \frac{1}{\rho}b_n.$$

Proof. The result can be proved by considering the odd and even cases separately, simply differentiating the sums, and closely following the steps of Lemma 8. □

Lemma 16. *For all $n \geq 0$, we have $a_{2n}(\rho = 0) = 1$, $a_{2n}(\rho = 1/4) = \binom{4n}{2n}(1/2)^{4n}$, and $a_{2n}(\rho = 1/2) = \binom{2n}{n}(1/2)^{2n}$, which implies $a_{2n}(\rho = 1/4) < a_{2n}(\rho = 1/2)$.*

Proof. We prove only the result for $\rho = 1/4$. Taking $\rho = 1/4$ in the recurrence equations of Lemma 8, we have $a_{n+1} = (a_n + b_n)/2$ and $b_{n+1} = (n+1)/(n+2)(a_n + b_n)/2$. Summing the two equations yields $2(a_{n+1} + b_{n+1})/2 = (2n+3)/(n+2)(a_n + b_n)/2$. If we denote by $u_n = (a_n + b_n)/2$, the equation becomes $2u_{n+1} = (2n+3)/(n+2)u_n$. Iterating from $n = 0 \ldots n-1$ and using $u_0 = 1/2$, we obtain $u_{n-1} = \binom{2n}{n}(1/2)^{2n}$. Since $a_{2n} = u_{2n-1}$ for all $n \geq 1$, the proof is finished. We note that u_{n-1} is decreasing in n, therefore $u_{2n-1} < u_{n-1}$, which is equivalent to $a_{2n}(\rho = 1/4) < a_{2n}(\rho = 1/2)$ □

Using the above two lemmata, we have the following

E.1. Proof of Proposition 5

Observe that $da_m/d\rho = 2m(b_{m-1} - a_{m-1}) = 2(n+1)\Delta_n^C$

By Lemma 13, we have $da_m/d\rho < 0$ for $\rho < \rho_n^*$ and $da_m/d\rho > 0$ for $\rho > \rho_n^*$, where ρ_n^* is given by $\Delta_n^C(\rho_n^*) = 0$. Therefore a_m is decreasing over $(0, \rho_n^*)$ and increasing over $(\rho_n^*, 1/2)$, thus it has a minimum at $\rho = \rho_n^*$ (see Fig. 2(a)) proving the first part of the proposition. Parts 2 and 3 follow from Lemma 16 and the fact that ρ_n^{**} is decreasing in ν with $\rho_n^{**}(\nu = 1) = \rho_n^*$ and $\rho_n^{**}(\nu = 0) = 1/2$. □

Appendix F. Proof of Observation 3

From previous section, we know that $a_m(\rho)$ is decreasing over $(0, \rho^*)$ and increasing over $(\rho^*, 1/2)$, where $1/3 < \rho^* < 1/2$. Moreover, from Lemma 16 we have $a_m(1/4) < a_m(1/2)$.

Since C's as a group would like to maximize the probability of a tie, subject to the constraint that majority voters are truthful, they will choose $\rho = 1/2$. As $\rho = q + (1-2q)x_C(A)$, this can be achieved by setting $x_C(A) = 1/2$. □

Appendix G. Supplementary data

A supplementary file containing omitted proofs can be found at http://dx.doi.org/10.1016/j.jebo.2012.09.014.

References

Alos-Ferrer, C., Kuzmics, C. Hidden symmetries and focal points. Journal of Economic Theory, forthcoming.
Blume, A., 2000. Coordination and learning with a partial language. Journal of Economic Theory 95, 1–36.
Blume, A., Franco, A., 2007. Decentralized learning from failure. Journal of Economic Theory 133, 504–523.
Blume, A., Gneezy, U., 2000. An experimental investigation of optimal learning in coordination games. Journal of Economic Theory 90, 161–172.
Blume, A., Gneezy, U., 2010. Cognitive forward induction and coordination without common knowledge: an experimental study. Games and Economic Behavior 68, 488–511.
Blume, A., Duffy, J., Franco, A., 2009. Decentralized organizational learning: an experimental investigation. American Economic Review 99, 1178–1205.
Burke, J., Taylor, C., 2008. What's in a poll? Incentives for truthful reporting in pre-election opinion surveys. Public Choice 137, 221–244.
Cox, G.W., 1994. Strategic voting equilibria under the single non-transferable vote. American Political Science Review 88, 608–621.
Cox, G.W., 1997. Making Votes Count. Cambridge University Press, Cambridge.
Crawford, V., Haller, H., 1990. Learning how to cooperate: optimal play in repeated coordination games. Econometrica 58, 571–596.
Crawford, V., Sobel, J., 1982. Strategic information transmission. Econometrica 50, 1431–1451.
Duverger, M., 1954. Political Parties: Their Organization and Activity in the Modern State. Wiley, New York.
Ekmekci, M., 2009. Manipulation through political endorsements. Journal of Economic Theory 144, 1227–1248.
Farrell, J., 1987. Cheap talk, coordination and entry. Rand Journal of Economics 18, 34–39.
Feddersen, T., Sandroni, A., 2006. A theory of participation in elections. The American Economic Review 96, 1271–1282.
Fey, M., 1997. Stability and coordination in Duverger's law: a formal model of pre-election polls and strategic voting. American Political Science Review 91, 135–147.
Forsythe, R., Myerson, R., Rietz, T.A., Weber, R.J., 1993. An experiment on coordination in multicandidate elections: the importance of polls and election histories. Social Choice and Welfare 10, 223–247.
Goeree, J., Grosser, J., 2007. Welfare reducing polls. Economic Theory 31, 51–68.
Harsanyi, J., Selten, R., 1988. General Theory of Equilibrium Selection in Games. MIT Press, Cambridge, MA.

Meirowitz, A., 2005. Polling games and information revelation in the Downsian framework. Games and Economic Behavior 51, 464–489.
Morgan, J., Stocken, P., 2008. Information aggregation in polls. American Economic Review 98, 864–896.
Myatt, D.P., Fisher, S.D., 2002. Tactical coordination in plurality electoral systems. Oxford Review of Economic Policy 18 (4), 504–522.
Myatt, D.P., 2007. On the theory of strategic voting. Review of Economic Studies 74, 255–281.
Myerson, R.B., 1998. Population uncertainty and Poisson games. International Journal of Game Theory 27, 375–392.
Myerson, R.B., 2000. Large Poisson games. Journal of Economic Theory 94, 7–45.
Myerson, R.B., 2002. Comparison of scoring rules in Poisson voting games. Journal of Economic Theory 103, 219–251.
Myerson, R.B., Weber, R.J., 1993. A theory of voting equilibria. American Political Science Review 87, 102–114.
Myerson, R., Rietz, T., Weber, R., 1998. Campaign finance levels as coordinating signals in three-way, experimental elections. Economics and Politics 10, 185–217.
Palfrey, T.R., 1989. A mathematical proof of Duverger's law. In: Ordeshook, P.C. (Ed.), Models of Strategic Choice in Politics. University of Michigan Press, pp. 69–92.
Riker, W.H., 1982. The two-party system and Duverger's law. American Political Science Review 76, 753–766.
Schelling, T.C., 1960. The Strategy of Conflict. Harvard University Press, Cambridge.
Taylor, C., Yildirim, H., 2010. Public information and electoral bias. Games and Economic Behavior 68, 353–375.