Sequential Voting with Externalities: Herding in Social Networks

Noga Alon* Moshe Babaioff† Ron Karidi‡ Ron Lavi§ Moshe Tennenholtz¶

February 7, 2012

Abstract

We study sequential voting with two alternatives, in a setting with utility externalities: as usual, each voter has a private preference over the candidates and likes her favorite candidate to win, but additionally, a voter values voting for the chosen winner (which is determined by the majority or super-majority of votes). This model aims to capture voting behavior (”likes”) in social networks which are publicly observed and sequential, and in which people care about their ”public image” as determined by their votes and the socially accepted outcome (the chosen winner). Unlike in voting with no externalities, voters act strategically although there are only two alternatives, as they rather vote against their preferred candidate if the other is to win. We present two rather surprising results that are derived from the strategic behavior of the voters. First, we show that in sequential voting in which a winner is declared when the gap in votes is at least some large value M, increasing M does *not* result in the aggregation of preferences of more voters in the decision, as voters start a herd on one candidate once a small lead in votes for that candidate develops. Furthermore, the threshold lead for such a herd to start is *independent* of M. Secondly, we show that there are cases in which sequential voting is strictly better than simultaneous voting, in the sense that it chooses the most preferred alternative with higher probability.

1 Introduction

The purpose of this paper is to analyze the process of sequential voting in social networks, which has become a popular method for online preference aggregation. In such networks, users have the option to vote either “like” or “dislike” on a certain specific content (e.g., picture, movie, blog post) posted in the network. As a user enters to watch content, she sees all previous votes (previous “likes” and “dislikes”), and then casts her own vote. The social network provides information on the identities of the different voters, so the votes are *not* anonymous. While the online setting is our main motivational example, there are other examples of sequential voting with observable votes. One is voting in various committees, where committee members sequentially announce their

*Tel-Aviv University and Microsoft Research. Email: nogaa@post.tau.ac.il
†Microsoft Research, Silicon Valley. Email: moshe@microsoft.com
‡Microsoft Israel Innovation Labs. Email: ronkar@microsoft.com
§Technion – Israel Institute of Technology and Microsoft Research. Email: ronlavi@ie.technion.ac.il
¶Technion – Israel Institute of Technology and Microsoft Research. Email: moshet@microsoft.com

opinion on the matter at hand. Another is the USA presidential primaries in which voting in different states happen at different times and is observable by later states and the candidates.

Such voting situations seem to exhibit utility externalities that should not be ignored. For example, in social networks, people may be unhappy to realize, in retrospect, that most of their friends voted differently than they did. In a committee voting for accepting a new faculty member, a committee member might get some dis-utility for voting against a candidate that ends up becoming a colleague. In presidential primaries, a state might get some dis-utility from voting against a candidate that ends up in a position of power.

In our model the utility of each voter is affected by two considerations: her own private preference over the two alternatives A and B, and an externality that makes her dislike an outcome in which she votes against the chosen winner. The utility does *not* directly depend on the identity of the candidate that is preferred by the majority of voters, it depends on the preferences of others only indirectly, through their votes. As far as we know this is the first paper that studies such a model theoretically, we discuss how our paper relates to prior literature at the end of this section.

Specifically, a voter that prefers A over B has a positive utility of α if she votes A and A is chosen, and a positive utility of $\gamma < \alpha$ if she votes B and B is chosen. Additionally, the voter has arbitrary positive utilities $\beta, \delta < \gamma$ for the cases where she votes A and B is chosen, and for the case where she votes B and A is chosen, respectively. Under the assumption of a negative externality for voting against the chosen winner, this utility structure captures a wide range of cases: from the case in which the voter cares only about conforming with the winner ($\gamma = \alpha > \beta = \delta$), to the case in which the voter puts equal importance on the two ingredients of voting for the winner and caring that her preferred alternative will win ($\alpha >> \gamma = \beta = \delta$).

In our model the utility of each voter is completely subjective. This is in contrast to all previous papers on sequential voting we are aware of, in which voters also want the outcome of the voting to be "objectively correct", in the sense that an unknown state of nature directly affects the utility of voters from the chosen alternative. Thus, in this prior work one alternative may be objectively better than the other, and the voting behavior mixes both strategic issues of conformity and informational issues of what is objectively the better alternative. In contrast, in our model, a voter's utility depends only on the voter's preference, her vote, and the chosen alternative, and there is no direct informational aspect. In our setting strategic considerations of voters are very important, as voting for one alternative or the other can influence future voters whose votes in turn affect the final outcome. In particular, truthful voting is not a dominant strategy.

We first study a setting with an infinite number of voters, aiming to capture the case of very large social networks. In this case we assume a symmetric prior probability, with probability $\frac{1}{2}$ a voter prefers A over B, independently of the other voters. (We also study a model with a finite population and non-symmetric prior, see additional details below.) Since the population is infinite, we assume that the voting stops when one alternative leads by at least M votes over the other alternative, where M is some arbitrarily large number. The alternative with such a big lead is accepted by the society as the "chosen alternative".

One might hope that as M grows the chosen alternative will depend on the preferences of a larger number of voters. Our first main result shows that this is *not* the case, and *herding* starts once some alternative gains a lead of magnitude that is independent of M. We show that from the first time one of the alternatives has a small lead (whose magnitude depends only on the utility parameters α, β, γ and *not* on the winning bar M) and onwards, all voters will vote for that temporary leader, regardless of their preferences, ensuring that this leader will be chosen. More formally, we show that this herd behavior is the *unique* symmetric subgame-perfect equilibrium of the voting game. Consequently, even if in the sequence of *true preferences* A is the first to have a lead of M, a small advantage for B in the beginning of the vote sequence will start a herd for B and make B the winner. Additionally, due to the strategic behavior of the voters, increasing M will *not* result in the preferences of larger number of voters influencing the outcome.

Intuitively, the herd is a consequence of a strategic forward-looking behavior, as follows. Consider a situation where the vote difference is $M-1$ for B, and suppose that the next voter prefers A. Then since the probability that B will be chosen (even if the current voter will reduce the difference to $M-2$ votes) is so high, the voter at this point will prefer to vote with the majority, unless her utility parameters are so extreme to neutralize the high probability. Thus, effectively, the threshold reduces to actually become $M-1$ and not M. This logic then continues. Interestingly, this avalanche does not stop at a point that depends on the winning threshold M, neither does it continue until a fixed difference of one or two votes. Instead, we show that the exact gap in which the herding starts depends only on the utility parameters (α, β and γ, but not δ).

Given this result a possible conclusion could be that simultaneous voting should be preferred, and we move to consider that option when the population is finite. (If the situation is sequential by nature, as is the case in online voting, one can simply hide the current vote history, to achieve the same simultaneity effect.) Quite surprisingly, if the prior is not completely symmetric and there is a small bias towards one of the alternatives (say, B), we show that this conclusion need not be correct. Towards this end, we study a setting with a finite number of voters (which is more natural in the context of simultaneous voting), in which the winner is the alternative with the majority of votes. We first reproduce a similar herding result for sequential voting, presenting a simple condition on the utility parameters α, β, γ that is necessary and sufficient for the unique subgame-perfect equilibrium to result with a herd in which all voters vote the same as the first voter, which votes according to her true preference. Once again, this condition does not depend on the number of voters, hence a "herd that follows the first voter" happens even if the number of voters becomes arbitrarily large, as long as the utility parameters satisfy a simple inequality.

Our second main result then states that, for any number of voters n, there exist utility parameters, and a non-symmetric prior probability with a slight advantage for B, such that, with these parameters, (1) sincere voting is *not* a Nash equilibrium of simultaneous voting, (2) the profile of strategies generating a "herd that follows the first voter" is the unique subgame-perfect equilibrium in sequential voting, and, most interestingly, (3) with these strategies the probability that the "correct" outcome (the outcome that the majority of voters truly prefer) is chosen is strictly higher than the probability it will be chosen in *any* Nash equilibrium of simultaneous voting. The intuition behind this result is that, with these parameters, non-informative voting in fact happens at larger extent in simultaneous voting. More specifically, in the parameters we choose, the most accurate Nash equilibrium of simultaneous voting is the one where all vote for B (each voter ignores her preference altogether and submit a non-informative vote of B), which is not that accurate as the prior probability gives only a slight advantage to B. Sincere voting of the first voter followed by the same vote by every later voter in sequential voting, is then enough to make the sequential voting method more accurate. While this result holds for any finite number of voters, we should emphasize that the result is meaningful when the number of voters is not too large, as the measure of the utility parameters for which the above conclusion is true in the domain of all utility parameters goes to zero as the number of voters goes to infinity.

1.1 Related Work

There is a large literature on herds and informational cascades in the context of sequential decision making, starting with Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992). Strategic considerations in such models are very different than in sequential voting, as these decision makers' utility is completely independent of the decisions made by players acting later. More specifically, these papers study a setting where players wish to make a "correct" action, that is determined by an unknown state of nature. Each player receives a signal about the state of nature, and can infer from his own signal, as well as from actions of previous players, a probability distribution on the possible states of nature. A well-established result shows that this informational uncertainty causes herding, when the aggregate information inferred from previous actions is of better quality than a player's own signal. This is clearly a very different setting than the one we study here. A good introduction to this literature is given in Easley and Kleinberg (2010).

More relevant to our setting is the literature on sequential voting, starting with Dekel and Piccione (2000). This literature combines the informational aspect with the standard assumption in voting models that voters have private preferences regarding the various alternatives. Importantly, in most of these models, the utility of a voter is assumed to be independent of her vote, given the chosen alternative. Thus, the kind of externalities we consider (a voter prefers her vote to be aligned with the chosen alternative) are not covered by this literature.

Two interesting papers on sequential voting do acknowledge the importance of conformity as a main component in voters' utilities. Callander (2007), in particular, discusses this issue in length. He relies on theories from social psychology (for example Aronson, Wilson and Akert (1997)) as well as many experimental work (for example Bartels (1988)) to justify and explain the motivation to vote for the winner. In his model, a player's utility depends both on whether she voted for the winner, and whether the winner is the "correct" one, given an unobserved state of nature. This model is different than our model in three important aspects. First, in his model voters do not have personal preferences, rather they have identical preferences for the correct candidate to be chosen and for the alignment with the winner. In contrast, in our model each voter likes her preferred candidate to win. Second, as discussed above, in our model utilities are completely subjective, which allows us to focus on this aspect which seems to be the most important aspect in social networks. Third, in his model the subjective component in the utility is binary, either the voter votes for the winner, or not. In our model, exactly because we focus on the subjective part, we study a much more refined non-binary utility, with four parameters that capture general utility structures.

There are also some technical differences between Callander (2007) and our paper. Mainly, for an infinite population the rule used to determine the winner is different. Callander's rule always requires infinite voting sequences to determine the chosen alternative, while our rule will likely choose a winner in finite time (our rule is very simple: when the gap in votes becomes large enough, voting ends and the leader is chosen). Some of his results are qualitatively different than ours. In his model, when the elections are "tight", simultaneous voting is better (in the information aggregation sense) than sequential voting. In our model, due to the refined utility structure, we show that even in "tight" elections there exist utility parameters for which sequential voting will be better. This happens as some parameters eliminate the informative voting equilibrium in simultaneous voting, while for the same parameters, in sequential voting, at least the first voter votes sincerely (and determines the outcome as all other follow), which in small populations is a significant difference.[1]

Ali and Kartik (2011) also recognize the importance of allowing the utility of a voter to depend on previous votes. However they only allow the utility to increase as more votes become aligned to the "correct" alternative (determined by state of nature), and not as they become aligned to the *chosen* alternative (as this is a subtle point, in Appendix A we present an elaborate example that demonstrates this difference). In fact, their model is conceptually closer to the classic literature on sequential voting than Callander's model.

Dasgupta, Randazzo, Sheehan and Williams (2008) conduct an experiment in the lab on a special case of the model we consider, with five voters, specific numbers $\alpha, \beta, \gamma, \delta$, and a symmetric prior. Yet the focus of our paper is inherently different than the focus of their experiments: we study infinite population with symmetric prior and finite population with a non-symmetric prior, while their experiments are for a small (and specific) finite population with symmetric priors. The

[1] In Callander (2007), the technical definition of "tight" is subtle, while here "tight" simply means that the prior is close to symmetric. However, if we think of the state of nature as the true majority, and a player's signal is her own preference, then if our setting is "tight" by our simple criterion then it is also "tight" by Callander's more complex criterion.

focus of their paper is mainly on comparing complete and incomplete information.

We are aware of two recent papers in the computer science literature that analyze sequential voting games in *complete information*. Desmedt and Elkind (2010) study a setting where voters are able to abstain, and they will not vote if their vote will not make a difference. They characterize the unique subgame perfect Nash equilibrium of sequential voting with two alternatives, and provide some counterintuitive examples when there are more alternatives. Xia and Conitzer (2010) suggest a dynamic programming algorithm to find the subgame perfect Nash equilibrium of sequential voting with any number of alternatives and for any anonymous voting rule. They use this algorithm in computer simulations, and observe that most voters prefer the winner of the equilibrium play over the alternative that would have won if all voters were to vote sincerely.

Paper Organization: The remainder of this paper is organized as follows. Section 2 outlines our model and notations. In Section 3 we discuss herding in an infinite population of voters, and in Section 4 we compare the accuracy of sequential and simultaneous voting in a finite population of voters. Section 5 concludes.

2 Model and Preliminaries

There are two alternatives A and B, and a set S of voters that are choosing one alternative as a winner. We study two variants of the problem, where $S = \mathcal{N}$ which means that the population is infinite and countable (in Section 3), and where $S = [n]$ for $n \geq 3$ that is odd, which means there is finite population of n voters (in Section 4).

Each voter in S has a private preference (type) about the alternative she likes, either A or B. We denote the type by $\sigma \in \{A > B, B > A\}$, where $A > B$ means that the voter prefers A over B (he "likes A"), and $B > A$ means that the voter prefers B over A (he "likes B"). The private preferences of the different voters are drawn independently and identically from some prior distribution, each voter prefers A with probability $p_A \in (0, 1)$ and prefers B with probability $p_B = 1 - p_A$. We represent the prior by the pair (p_A, p_B).

Unlike in traditional voting scenarios, in our model the voters not only care about the chosen alternative but also about the alignment of their votes with the chosen alternative (conforming with the social choice). Thus, in our model, the utility of a voter depends on the alternative she likes, her vote, and the chosen alternative. The utility of a voter that likes alternative $X \in \{A, B\}$, votes for alterative $V \in \{A, B\}$, and alternative $P \in \{A, B\}$ is picked, is $v(X, V, P)$. We study symmetric utilities, and denote:

- $v(A, A, A) = v(B, B, B) = \alpha$. In words, if a voter likes A, votes for A, and A wins, the utility of the voter is α. The same holds for the symmetric case of a voter that likes B, votes for B, and B wins.

- $v(A, B, B) = v(B, A, A) = \gamma$. In words, if a voter likes A, votes for B, and B wins, the utility of the voter is γ. The same holds for the symmetric case of a voter that likes B, votes for A, and A wins.
- Similarly, we denote $v(A, A, B) = v(B, B, A) = \beta$, and $v(A, B, A) = v(B, A, B) = \delta$.

We assume that $\alpha > \gamma > \max\{\beta, \delta\} > 0$. In other words, a voter obtains maximal utility if she votes for her preferred alternative, and that alternative wins. However, a voter prefers to vote for the chosen alternative rather than voting against the chosen alternative, even if this means that she votes for her less preferred alternative.

We consider two voting schemes. In *sequential voting* each of the voters observes the votes of all her predecessors before casting her publicly observed vote. In *simultaneous voting* each of the voters does not observe any other vote before casting her vote.

Simultaneous voting with finite population of n voters, for odd n, is modeled as a simultaneous move game in which each voter casts a vote and the alternative with majority of votes is selected. In this game a strategy for agent i is just a mapping from type to an alternative (vote). The strategy s_i of player i is a function $s_i : \{A > B, B > A\} \to \{A, B\}$ where $s_i(\sigma)$ is the vote of i when her type is $\sigma \in \{A > B, B > A\}$. We next move to formalize the sequential voting game.

2.1 Sequential Voting

We model sequential voting as an extensive-form game with perfect but incomplete information. Players in S are indexed by consecutive natural numbers starting at 1 (in Section 3 there are countably many players indexed by the natural numbers $\mathcal{N}$, and in Section 4 there are only n such players, indexed 1 to n). Players move sequentially, one at a time (roll-call voting): player 1 moves first, player 2 moves second, and so on. In her move, the action of player i is to vote for either alternative A or alternative B. A strategy s_i of player $i \in S$ is a function from her type and the history of votes of players $1, 2, \ldots, i-1$ to an alternative, her own vote. Formally, $s_i : \{A > B, B > A\} \times \{A, B\}^{i-1} \to \{A, B\}$ where $s_i(\sigma, h)$ is the vote of i when her type is $\sigma \in \{A > B, B > A\}$ and the history of the $i-1$ previous votes is h. Note that in this paper we only consider pure strategies.

When considering a finite population of voters (Section 4) we assume that n is odd and that the alternative with majority of votes is selected. With infinite population (Section 3) we only consider sequential voting and assume that at the first point in which one alternative leads by at least M votes over the other alternative, it is selected.

2.2 Solution Concepts

Given a voting game, and assuming that players are strategic, we aim to identify the equilibrium outcomes, and their properties. The most basic and standard equilibrium notion is Nash equilibrium: A set of strategies $\{s_i(\cdot)\}_{i\in S}$ is Nash equilibrium if for any player $i \in S$, $s_i(\cdot)$ is a best response to $s_{-i}(\cdot)$, that is, the player's expected utility given that the other players play $s_{-i}(\cdot)$ is maximized by playing $s_i(\cdot)$ (where the expectation is taken over the random types of other players for which there is uncertainty).

This solution concept fits the case of simultaneous voting and we indeed use it for that case. Yet, for sequential voting it suffers a significant drawback as it does not take into account the fact that the game is played sequentially, and players strategies can encode "non-credible" threats. Such threats allow some non-intuitive strategies to be in equilibrium. For example, the strategy in which a player always votes B, regardless of her type, is Nash equilibrium in the finite population case (and the infinite case with any $M > 1$). This is because, if all other players always vote for B, any fixed player will strictly prefer to vote B over A, as B will be chosen regardless of her vote, and, even if the player likes A, $\gamma > \beta$. Such a profile is a Nash equilibrium even if all players prefer A!

The standard solution is to add the requirement of subgame perfection, which is defined as follows. A subgame of the original game is defined by determining a player $i^* \in S$, and a tuple of votes $s_1, ..., s_{i^*}$ (i.e., determining a history that player $i^* + 1$ sees). In this subgame, only the players after i^* play, but the votes $s_1, ..., s_{i^*}$ affect the final outcome, as the subgame starts with a non-zero vote count determined by $s_1, ..., s_{i^*}$. Every choice of i^* and the initial votes determines a different subgame.[2]

Definition 1. *A set of strategies $\{s_i(\cdot)\}_{i\in S}$ (of the original game with all players) is a* subgame perfect equilibrium *if it is a Nash equilibrium in all possible subgames.*

The above example, where all players always vote B, is not a subgame perfect equilibrium. For example, in the finite case if the first $n-1$ players vote A, and the last player likes A, then it is not a best response for her to play B. Playing A will cause the game to end with A winning and the player voting for the winner getting the maximal possible utility of α, and not just $\delta < \alpha$. Thus, her "commitment" to *always* vote B is not credible. Preceding players realize this fact, which in turn can change their own strategies.

3 Equilibrium Analysis of Sequential Voting in an Infinite Population

In this section we study sequential voting in an infinite population, to capture the motivating story of voting in a social network. Since the population is infinite, to determine the end of the game we parameterize it by an integer $M > 0$, the "winning threshold". The game ends when one candidate leads by at least M votes over the other candidate. In this case, the candidate with

[2]Although we have incomplete information, all information sets in our game are essentially of size 1, as, given the actions of all voters, utilities are independent of voters' preferences. For this reason, the solution concept of subgame perfection is more appropriate than Bayesian perfect equilibrium or sequential equilibrium.

the larger number of votes is declared "the winner". Note that there exist strategies that ensure that the difference in vote count will never reach M, and in this case no winner is declared and the sequence of votes is infinite. We assume that players prefer the outcome in which they vote for their preferred alternative over the outcome where the voting never ends. Formally, this means that when the sequence of votes is infinite, the players have utility $u_0 < \alpha$. We note that if this inequality is reversed, one equilibrium is to keep the vote difference to be at most 1 by voting against the currently leading alternative, or for the preferred candidate when tie. This will ensure that the no alternative will ever be chosen.

In this section we only study the case where the two alternatives are a-priori symmetric, that is, the prior probability is ($\frac{1}{2}$,$\frac{1}{2}$). If one alternative, say B, has higher prior probability $p >\frac{1}{2}$ then clearly the probability that the majority of a finite population n will prefer B converges to one as the size of the population n goes to infinity, thus it seems that there is no much point in conducting any voting.

As discussed in Section 2.2 the solution concept used for sequential voting is of subgame perfect equilibrium. Designing a tuple strategies that for a subgame perfect equilibrium is not straightforward, our first result presents such a tuple that uses very simple and natural strategies.

Definition 2. *In the r-*threshold strategy*, a player that likes A votes A if the number of votes for B minus the number of votes for A is at most r, otherwise it votes for B. Similarly, a player that likes B votes B if the number of votes for A minus the number of votes for B is at most r, otherwise it votes for A.*

Let k^* be the largest integer such that $\frac{1}{k+1}\alpha + \frac{k}{k+1}\beta > \gamma$ (and, thus, $\frac{1}{k^*+2}\alpha + \frac{k^*+1}{k^*+2}\beta \leq \gamma$). Note that k^* does not depend on M, only on the utilities α, γ, and β.

Theorem 1. *The set of strategies in which every player plays the k^*-threshold strategy form a subgame perfect equilibrium.*

Proof. The proof relies on the following well-known fact about random walks, commonly termed the "Gambler's Ruin Problem" (see for example Spitzer (2001)). Consider an infinite sequence of i.i.d. random variables $x_1, x_2, ...$ where $Pr(x_i = -1) = Pr(x_i = 1) =\frac{1}{2}$. Let $S_n = \sum_{i=1}^{n} x_i$, and $T = \min\{n \,:\, S_n = a \text{ or } S_n = -b\}$ for two nonnegative integers a, b. Then, $Pr(S_T = a) = b/(a+b)$ and $Pr(S_T = -b) = a/(a+b)$. In other words, if an infinite random walk starts at 0 and ends when it either reaches a or $-b$, then the probability it ends at a is $b/(a+b)$ and the probability it ends at $-b$ is $a/(a+b)$.

Fix a player i, and suppose without loss of generality that i likes A. Fix any sequence of votes of players $1, ..., i-1$, and let x be the number of votes for B minus the number of votes for A in this sequence. Assuming that the game did not end, we have $-M < x < M$. We need to show that the the tuple of strategies in which every player plays the k^*-threshold strategy is a Nash equilibrium in this game for every player $i' \geq i$, that is, that player i' maximizes her expected utility by following the k^*-threshold strategy, assuming all players with index at least i also follow this strategy, and assuming any arbitrary actions of players $1, ..., i-1$. In fact it suffices to show this claim only for player i, as, then, taking a subgame in which i' is the first player will yield the claim for any $i' > i$. Thus, we will show that if all players with index at least $i+1$ follow the k^*-threshold strategy then player i maximizes her expected utility by following the k^*-threshold strategy as well.

Suppose first that $x \geq k^* + 1$. If $x \geq k^* + 2$ then regardless of i's vote, all subsequent players will vote B and therefore B will be the winner with probability 1. Since $\gamma > \beta$, i should vote B as the k^*-threshold strategy instructs her to do. If $x = k^* + 1$ and i votes B, once again B will be the winner with probability 1. If i votes A, on the other hand, B will lead by only k^* votes, which implies that with positive probability A will win. Since all subsequent players follow the k^*-threshold strategy, we can determine this probability by looking at the Gambler's ruin problem, interpreting $x_i = 1$ as a vote for B and $x_i = -1$ as a vote for A. The random walk will terminate with B as the winner after n steps if $S_n = 1$ since this will mean that at this point B leads by k^*+1 and from this point on all subsequent players vote B. Similarly, the random walk will terminate with A as the winner after n steps if $S_n = -(2k^* + 1)$ since at this point A leads by $k^* + 1$. Thus, we have the parameters $a = 1$ and $b = 2k^* + 1$. Therefore i's expected utility from voting A is $\frac{1}{2k^*+2}\alpha + \frac{2k^*+1}{2k^*+2}\beta$. Since $\frac{1}{2k^*+2} < \frac{1}{k^*+2}$ and $\alpha > \beta$,

$$\frac{1}{2k^* + 2}\alpha + \frac{2k^* + 1}{2k^* + 2}\beta < \frac{1}{k^* + 2}\alpha + \frac{k^* + 1}{k^* + 2}\beta \leq \gamma$$

where the last inequality follows from the definition of k^*. Thus, i should again vote B, in line with the k^*-threshold strategy.

Now suppose $x = k^*$. Similar to above, if i votes B, B wins with probability 1, and if i votes A we enter a random walk with stopping probabilities $a = 2$ and $b = (k^* - 1) + k^* + 1 = 2k^*$. Since $\frac{2}{2k^*+2}\alpha + \frac{2k^*}{2k^*+2}\beta = \frac{1}{k^*+1}\alpha + \frac{k^*}{k^*+1}\beta > \gamma$, i should vote A as the k^*-threshold strategy suggests. For $x < k^*$ the probability that A wins when i votes A only increases, and the probability that B wins if i votes B becomes smaller than 1, and therefore the same inequality again shows that the expected utility from voting A is larger than the expected utility from voting B. Thus, the k^*-threshold strategy maximizes the expected utility of player i, and the theorem follows. □

Non-symmetric subgame perfect equilibria. The tuple of strategies in which every player plays the k^*-threshold strategy is completely symmetric in the sense that all voters act the same, in a way that is completely symmetric between bidders and between types. In fact, in the next section we show that the k^*-threshold strategy is the *unique* symmetric strategy that forms an subgame perfect equilibrium.

However there are other, non-symmetric strategies that form other subgame perfect equilibria. Suppose, for example, that $\frac{1}{2}\alpha + \frac{1}{2}\beta < \gamma$. Then, $k^* = 1$, and in the k^*-threshold strategy, players start to vote for the leading candidate regardless of their type once the vote difference is at least 2.

For these utilities, however, the following strategy is an additional example of a subgame perfect equilibrium: A player that likes "B" always votes B. A player that likes A also almost always votes B, except for the case where the number of votes for A minus the number of votes for B is $M-1$, in which it votes A. It is clear that a player that likes B is always best responding, since, assuming all subsequent players follow the strategy, her vote for B will lead to the winning of B with probability 1, regardless of the initial history.

Now consider a player that likes A, and note that it is enough to examine the case where the vote difference is either $M-1$ or $M-2$, assuming all subsequent players follow the strategy. If it is $M-1$, the maximal utility is clearly obtained by voting A, as the strategy instructs. If it is $M-2$, a vote for B will lead to the winning of B with probability 1, for a resulting utility of γ, and a vote for A lead to the winning of A with probability $\frac{1}{2}$ (if the next player likes A) and to the winning of B with probability $\frac{1}{2}$ (if the next player likes B), for a resulting utility of $\frac{1}{2}\alpha + \frac{1}{2}\beta < \gamma$. Therefore voting for B is indeed a best response. We conclude that this strategy indeed forms a non-symmetric subgame perfect equilibrium.

3.1 Uniqueness under Symmetry

Since the entire setup is completely symmetric, it makes sense to focus on symmetric strategies. In our setting, symmetry has three implications: all players use the same strategy; this strategy depends only on the difference between the votes to A and to B (and not on the order of votes or any other property of the history), and the two types "likes A" and "likes B" are symmetric in their voting decisions. Formally,

Definition 3. *A tuple of strategies $\{s_i(\cdot)\}_{i=1}^{\infty}$ is symmetric if there exists a function $s : \{-M+1, -M+2, \ldots, 0, \ldots, M-2, M-1\} \times \{A > B, B > A\} \to \{A, B\}$ such that,*

1. *For any history h and any type $\sigma \in \{A > B, B > A\}$, the strategy s_i of voter i satisfies $s_i(h, \sigma) = s(h_A - h_B, \sigma)$, where h_A, h_B denote the number of votes to A, B in h, respectively.*
2. *For any integer $0 < x < 2M$, $s(M-x, A > B) = A$ if and only if $s(-M+x, B > A) = B$.*

For example, symmetry implies that if the vote difference is zero, a player that likes A votes A if and only if a player that likes B votes B. The k^*-threshold strategy is clearly symmetric. We next show it is the only symmetric strategy that that forms a subgame perfect equilibrium. Recall that k^* is the largest integer such that $\frac{1}{k+1}\alpha + \frac{k}{k+1}\beta > \gamma$. Thus, $\frac{1}{k^*+2}\alpha + \frac{k^*+1}{k^*+2}\beta \leq \gamma$. We prove the theorem for the case that the inequality is strict. If it is in fact an equality, additional technicalities of the indifference between k^* and k^*+1 make the picture more messy, without adding much essence.

Theorem 2. *If $\frac{1}{k^*+2}\alpha + \frac{k^*+1}{k^*+2}\beta < \gamma$, the k^*-threshold strategy is the unique symmetric strategy that forms a subgame perfect equilibrium.*

In the remainder of this section we prove this theorem. We start the proof with the following claim: if at some vote difference x, the strategy votes for a fixed outcome, say A (i.e., $s(x, B > A) = s(x, A > B) = A$), then for all larger vote differences $y > x$, the strategy will also always vote for A (i.e., $s(y, B > A) = s(y, A > B) = A$).

Claim 1. *If there exists $-M < x < M$ such that $s(x, B > A) = s(x, A > B) = A$ then for any $y > x$, $s(y, B > A) = s(y, A > B) = A$.*

Proof. Consider a game that starts with a vote difference of x for A. Since $s(x, B > A) = s(x, A > B) = A$, B can never win, and the only two possible outcomes are that A wins and that the voting process goes on infinitely. We first argue that $s(y, A > B) = A$ for any $y > x$ by induction on $y = M - 1, ..., x$. The base case of $s(M - 1, A > B) = A$ is clear since after such a vote A wins and α is the maximal possible utility for a player of type $A > B$. Thus, assume that $s(M - 1, A > B) = ... = s(y + 1, A > B) = A$. If $s(y, A > B) = B$, a profitable deviation for a player who faces a vote difference of y is to vote A, since this strictly increases the probability that A wins (and the player's expected utility is the probability that A wins times α plus the probability of no winner times u_0). Thus it must be that $s(M - 1, A > B) = ... = s(x, A > B) = A$. Thus, A wins with probability 1, regardless of $s(y, B > A)$ for any $y > x$. Therefore, when the vote difference is y, a player with $B > A$ will have utility γ if she votes for A, and β if she votes for B. Since $\gamma > beta$ it must be that $s(y, B > A) = A$ and the claim follows. □

This claim immediately implies the existence of a threshold T, $0 < T \leq M$, such that $s(x, A > B) \neq s(x, B > A)$ for every $-T \leq x \leq T$, $s(x, A > B) = s(x, B > A) = B$ for every $x < -T$, and $s(x, A > B) = s(x, B > A) = A$ for every $x > T$. To see this, note that if there exists $0 < x < M$ such that $s(x, A > B) = s(x, B > A) = A$, take the minimal such x, and then the claim implies that $T = x - 1$ (notice that by symmetry $s(-x, A > B) = s(-x, B > A) = B$). If no such x exists then we take $T = M$.

We next show that $T = k^*$. First, suppose towards a contradiction that $T > k^*$. Then, when the difference in votes is T in favor of B, the equilibrium strategy instructs one of the types (either a type that likes A or a type that likes B) to vote A. If the player votes B at this point, B will be chosen, and therefore if the player likes B she will maximize her utility by voting B. Hence the strategy instructs a player that likes A to vote A. In this case, and assuming that the other players follow the strategy, her expected utility is $\frac{2}{2T+2}\alpha + \frac{2T}{2T+2}\beta$ where the probabilities are again given by the gambler's ruin problem. Since $T > k^*$ we have $\frac{2}{2T+2}\alpha + \frac{2T}{2T+2}\beta = \frac{1}{T+1}\alpha + \frac{T}{T+1}\beta < \gamma$, hence deviating and playing B (by this causing B to win) improves the player's utility, a contradiction. Thus $T \leq k^*$.

Second, suppose towards a contradiction that $T < k^*$, and consider the decision of a player that likes A when the vote difference is $T+1$ for B. If she votes B as the strategy instructs her, B will be chosen for a resulting utility of γ. If she votes A, her expected utility will be $\frac{1}{T+1}\alpha + \frac{T}{T+1}\beta > \gamma$, as explained where the last inequality follows since $T+1 \leq k^*$. Thus deviating from the recommended strategy strictly increases the player's utility, a contradiction. We therefore conclude that $T = k^*$.

It remains to argue that when the vote difference is smaller than k^*, players are truthful. Consider a player that likes A, when the vote difference is k^* for B. Similar to above, if she votes B her utility is γ and if she votes A her utility is $\frac{1}{k^*+1}\alpha + \frac{k^*}{k^*+1}\beta > \gamma$, thus she votes truthfully. If the vote difference for B is smaller than k^*, the left hand side increases as the probability that A will win increases, and the right hand side decreases as the probability that B will win is now smaller than one, and $\delta < \gamma$. Thus, in this case as well the player is truthful. This concludes the proof of theorem 2.

4 Comparison of Sequential vs. Simultaneous Voting

As the herding phenomenon in sequential voting might result in picking a winner that does not match the true preferences of the voters, one might wonder whether it is better to move to simultaneous voting. (In a sequential setting, one can hide the vote count in order to make the process equivalent to simultaneous voting.) To study this question, we move to a model with finite number of players, n, as in the infinite case simultaneous voting does not make much sense. In this case, therefore, the winner is simply the candidate with the majority of votes. To avoid further assumptions on the utility function when there is a tie, we will assume that n is any arbitrary odd integer.

In a similar way to what we already observed regarding sequential voting, all voting for the same candidate (either A or B) is always a Nash equilibrium. The advantage of simultaneous voting is that sincere voting, in which each player votes for her preferred candidate, is also *sometimes* a Nash equilibrium. The (simple) necessary and sufficient condition for that always holds if the prior probability that a player prefers A or B is equal to $\frac{1}{2}$. However if the prior probability for preferring one of candidates is slightly higher than $\frac{1}{2}$, there are cases in which sincere voting is not be a Nash equilibrium. The analysis in this case is a quite surprising, as we show that if this happens, there are situations in which sequential voting predicts the "correct" outcome more accurately than simultaneous voting. Throughout, we denote by p the prior probability that a player prefers B, and assume without loss of generality that $p \geq \frac{1}{2}$.

4.1 Equilibria of Sequential Voting

We will show that a very simple condition on our parameters ensures that the profile in which every voter votes with current leader (and to her preferred candidate if tie) is the unique subgame perfect equilibrium in sequential voting. More specifically, in the "vote with current leader" strategy, each player votes for the leading candidate (so far), unless there is a tie in which case she votes for her true preference. If all players follow this strategy, the first vote completely determines the outcome as all other voters will herd, following the first voter. The following condition on the parameters of the model ensures that the profile of strategies in which every voter plays "vote with current leader" is a subgame perfect equilibrium.

$$\gamma \geq p\alpha + (1-p)\beta \tag{1}$$

This equation is generically satisfied in a non-negligible fraction of the domain of all parameters. For example, if the parameters of the utility are obtained using some joint distribution such that γ is uniformly distributed between α and β, and $p \in [0.5, 0.75]$, the probability that the equation will be satisfied is at least $1/4$.

Theorem 3. *The profile of strategies in which every voter plays "vote with current leader" is a subgame-perfect equilibrium if and only if Equation (1) holds. Moreover, if the equation holds with a strict inequality, this is the unique subgame-perfect equilibrium.*

Proof. Assume Equation (1) holds with a strict inequality (if it holds with an equality, a similar argument can be constructed by omitting the word "unique" throughout). Fix any player $1 \leq i \leq n$, and consider the subgame that starts with player i, with previous vote counts h_A, h_B for candidates A, B. We need to show that it is a Nash equilibrium for all voters to play "vote with current leader" in this subgame.

If $|h_A - h_B| \geq 2$, the leading candidate will continue to lead regardless of i's vote, therefore subsequent players will vote for this candidate as well, and thus it will win regardless of i's vote. Thus, i's unique best response is to vote for this candidate, as suggested by the "vote with current leader" strategy.

If $h_A = h_B$, and i votes for her preferred candidate (as suggested by "vote with current leader"), this candidate will win as subsequent players will vote for this candidate as well. This is clearly the unique best response of i in this case.

The remaining case is $|h_A - h_B| = 1$. If the leading candidate is the preferred candidate of i, once again the unique best response of i is to vote for this candidate, as by doing so this candidate is guaranteed to win. Otherwise, suppose that i prefers the other candidate, and consider two cases:

- Player i prefers A and the leading candidate is B: In this case, if i votes for B, B will be chosen and her utility will be γ. If i votes for A, the next voter $i+1$ will face a tie, and will therefore vote according to her own true preference. The expected utility for voter i if she votes A will therefore be $(1-p)\alpha + p\beta \leq p\alpha + (1-p)\beta < \gamma$ (where the first inequality follows since $\alpha > \beta$ and $p \geq 0.5 \geq 1-p$, and the second inequality is Equation (1)). Thus, her unique best response is to vote for B in this case.

- Player i prefers B and the leading candidate is A: If i votes for A then A will win and her utility will be γ. If she votes for B her expected utility will be $p\alpha + (1-p)\beta < \gamma$, implying again that i's unique best response is to vote for the leading candidate.

We conclude that player i will follow the "vote with current leader" strategy, if all subsequent players will do so as well. This implies that everyone playing the "vote with current leader" strategy is a subgame-perfect Nash equilibrium of this game. To see that it is the unique subgame-perfect Nash equilibrium, suppose towards a contradiction that there exists a different subgame-perfect Nash equilibrium, and let i be the player with the largest index whose strategy is different than the "vote with current leader" strategy. However, the above argument shows that, for any history h, the unique best response of i when players $i+1, ..., n$ are playing the "vote with current leader" strategy is to play the same strategy, a contradiction.

To conclude the proof, it only remains to show that if Equation (1) does not hold, everyone playing the "vote with current leader" strategy is not a subgame-perfect Nash equilibrium. To see this, consider any voter with an even index, for example player $n-1$. If the vote difference in the history that i sees is one for A, and i prefers B, "vote with current leader" instructs i to vote A, for a resulting utility γ. If player i deviates, and votes for her preferred candidate B, the next voter sees a tie, and since she and all subsequent voters follow the "vote with current leader" strategy, she will vote for her preferred candidate, and this candidate will be the winner. Thus, i's expected utility in this case is $p\alpha + (1-p)\beta > \gamma$ (as Equation (1)) does not hold). This implies that everyone playing the "vote with current leader" strategy is not a subgame-perfect Nash equilibrium, and the theorem follows. □

4.2 Equilibria of Simultaneous Voting

There are four possible pure strategies for a player in simultaneous voting: vote sincerely, vote exactly opposite of her type, or vote for a fixed alternative (independent of her type). Thus, there are four possible *symmetric* equilibria: "all vote for A", "all vote for B", all vote sincerely, and all vote opposite. There could potentially be additional asymmetric equilibria, but in this section we show that this cannot happen if $\beta \geq \delta$.

Lemma 1. *Fix an arbitrary pure Nash equilibrium of simultaneous voting. Then, if one player is sincere, no other player votes for a fixed alternative (independent of her preference).*

Proof. Assume towards a contradiction that voter 1 votes sincerely and voter 2 votes for some fixed alternative X, independent of her preference. We will show that one of the two voters is not best responding. Fix the strategies of all voters but the first two. Given these strategies and the prior, assume that the probability that X leads by exactly 1 vote is η_X, and assume that the probability that Y leads by exactly 1 vote is η_Y. Let μ_X be the probability that X leads by at least 3 votes and let μ_Y be the probability that Y leads by at least 3 votes (since n is odd, these are all the possible cases).

Given that voter 2 is always voting for X, for voter 1 that prefers Y to rather vote for Y and not to X it must hold that

$$\alpha \cdot \mu_Y + \beta \cdot \mu_X + \beta \cdot \eta_X + \alpha \cdot \eta_Y \geq \gamma \cdot \mu_X + \delta \cdot \mu_Y + \gamma \cdot \eta_X + \gamma \cdot \eta_Y$$

Given that voter 1 is always voting sincerely, for voter 2 that prefers Y to rather vote for X and not to Y it must hold that

$$\delta \cdot \mu_Y + \gamma \cdot \mu_X + \gamma \cdot \eta_X + \eta_Y(P_Y \cdot \delta + P_X \cdot \gamma) \geq \alpha \cdot \mu_Y + \beta \cdot \mu_X + \eta_X(P_Y \cdot \alpha + P_X \cdot \beta) + \alpha \cdot \eta_Y$$

Summing up the two equations we get that it must hold that

$$\beta \cdot \eta_X + \eta_Y(P_Y \cdot \delta + P_X \cdot \gamma) \geq \eta_X(P_Y \cdot \alpha + P_X \cdot \beta) + \gamma \cdot \eta_Y$$

We next argue that this leads to a contradictions.

First, as P_Y and P_X are positive, $P_Y + P_X = 1$, $\eta_X \geq 0$ and $\alpha > \beta \geq 0$, we conclude that $\beta \cdot \eta_X \leq \eta_X(P_Y \cdot \alpha + P_X \cdot \beta)$.

Second, $\eta_Y > 0$, as otherwise this implies that the majority among the other $n - 2$ voters is never Y (if $\mu_Y > 0$ then $\eta_Y > 0$ as each voter votes independently of the others), and in this case player 1 is not best responding – she will improve utility by always voting for X.

Now, since $\gamma > \delta \geq 0$, we conclude that $\eta_Y(P_Y \cdot \delta + P_X \cdot \gamma) < \gamma \cdot \eta_Y$, and we have a contradiction.

□

Lemma 2. *If $\beta \geq \delta$, in any arbitrary pure Nash equilibrium of simultaneous voting, no player always votes opposite to her preference.*

Proof. Assume by contradiction that player i votes opposite in some Nash equilibrium. Let μ_0 be the probability that there is a tie in the votes of the others. Note that if there is no tie, since n is odd, there is a gap of at least two votes between the two alternatives after the others have voted, and the voter is not pivotal. Let μ_A be the probability that the votes of the others already determined that A will be chosen, and Let μ_B be the probability that the votes of the others already determined that B will be chosen. For a voter that prefers B to rather vote for A and not to B it must hold that

$$\alpha\mu_0 + \beta\mu_A + \alpha\mu_B \leq \gamma\mu_0 + \delta\mu_B + \gamma\mu_A$$

For a voter that prefers A to rather vote for B and not to A it must hold that

$$\alpha\mu_0 + \alpha\mu_A + \beta\mu_B \leq \gamma\mu_0 + \delta\mu_A + \gamma\mu_B$$

Summing up the two equations we get that it must hold that

$$\alpha(2\mu_0 + \mu_A + \mu_B) + \beta(\mu_A + \mu_B) \leq \gamma(2\mu_0 + \mu_A + \mu_B) + \delta(\mu_A + \mu_B)$$

but this contradicts the fact that $\alpha > \gamma > 0$, $\beta \geq \delta > 0$, $2\mu_0 + \mu_A + \mu_B > 0$, and $\mu_A + \mu_B \geq 0$. $\square$

Corollary 1. *If $\beta \geq \delta$, there are only three possible Nash equilibria: all vote sincerely, "all vote for A", and "all vote for B".*

Without the assumption that $\beta \geq \delta$, this result is not necessarily true. For example, if $\alpha = \gamma = \delta > \beta$, all vote opposite is a Nash equilibrium as in this case any convex combination between γ and δ is larger than any convex combination between α and β.

4.3 Comparison of Accuracy of Equilibria Outcomes

We now ask in what situations sequential voting is better than simultaneous voting, in the sense that the probability that the "correct" outcome (the outcome preferred by the majority of voters) is chosen is higher in sequential voting, compared to the most accurate equilibrium of simultaneous voting. Of the three equilibria discussed above, "all vote for A" is less accurate than "all vote for B" since $p \geq \frac{1}{2}$, so in cases that "all vote for B" is less accurate than sequential voting, so is "all vote for A". We will show that there are parameters in which sincere voting is not a Nash equilibrium, and therefore the most accurate Nash equilibrium of simultaneous voting is "all vote for B". We will then show that, for any number of players, there exist parameters of the model in which the outcome of this equilibrium is less accurate than the outcome of "vote with current leader", which is in these parameters the unique subgame perfect equilibrium of sequential voting. This will give us the desired conclusion that there always exist parameters for which sequential voting is better than simultaneous voting.

For $m \leq n$ let $P(m, n)$ be the probability that at least m out of n voters prefer B (recall that each voter independently and identically prefers B with probability p). Formally,

$$P(m, n) = \sum_{i=m}^{n} \binom{n}{i} p^i (1-p)^{n-i} \tag{2}$$

Lemma 3. *Sincere voting is a Nash equilibrium of simultaneous voting if and only if*

$$\begin{aligned} \gamma \cdot P\left(\frac{n-1}{2}, n-1\right) + \delta\left(1 - P\left(\frac{n-1}{2}, n-1\right)\right) \leq \\ \alpha\left(1 - P\left(\frac{n+1}{2}, n-1\right)\right) + \beta \cdot P\left(\frac{n+1}{2}, n-1\right) \end{aligned} \tag{3}$$

Proof. First consider a voter that prefers A. If every other voter votes sincerely, then voting sincerely (for A) will give her utility of $\alpha \cdot \left(1 - P\left(\frac{n+1}{2}, n-1\right)\right) + \beta \cdot P\left(\frac{n+1}{2}, n-1\right)$, while voting for B will give her utility of $\gamma \cdot P\left(\frac{n-1}{2}, n-1\right) + \delta \cdot \left(1 - P\left(\frac{n-1}{2}, n-1\right)\right)$. Thus, sincere voting a best response in this case if and only if Eq. (3) holds. This suffices to show the necessity part of the claim.

To show sufficiency, consider next a voter that prefers B. In this case, if every other voter votes sincerely, then voting sincerely (for B) will give her utility of $\alpha \cdot P\left(\frac{n-1}{2}, n-1\right) + \beta \cdot \left(1 - P\left(\frac{n-1}{2}, n-1\right)\right)$, while voting for A will give her utility of $\gamma \cdot \left(1 - P\left(\frac{n+1}{2}, n-1\right)\right) + \delta \cdot P\left(\frac{n+1}{2}, n-1\right)$. The latter expression is smaller than the left hand side of Eq. 3 as $\gamma > \delta$ and $P\left(\frac{n-1}{2}, n-1\right)$ is at least as large as the probability that at least $\frac{n-1}{2}$ out of $n-1$ vote for A, which is equal to $\left(1 - P\left(\frac{n+1}{2}, n-1\right)\right)$. Similarly, the former expression is larger than the left hand side of Eq. 3 as $P\left(\frac{n-1}{2}, n-1\right) > \left(1 - P\left(\frac{n+1}{2}, n-1\right)\right)$ and $\alpha > \beta$. Thus, if Eq. 3 holds, it is a best response to vote for B in this case as well, and the claim follows. □

Let $\epsilon_{n,p} = \binom{n-1}{\frac{n-1}{2}} \left(p(1-p)\right)^{\frac{n-1}{2}}$. Note that $\epsilon_{n,p} > 0$ for every n, p, and that $P\left(\frac{n-1}{2}, n-1\right) = P\left(\frac{n+1}{2}, n-1\right) + \epsilon_{n,p}$. Suppose that $\alpha = \gamma > \beta = \delta$. Then, using Eq. (3), a sufficient condition that sincere voting is not a Nash equilibrium of simultaneous voting is that $\alpha \cdot P\left(\frac{n-1}{2}, n-1\right) + \beta\left(1 - P\left(\frac{n-1}{2}, n-1\right)\right) > \alpha\left(1 - P\left(\frac{n+1}{2}, n-1\right)\right) + \beta \cdot P\left(\frac{n+1}{2}, n-1\right)$. Rearranging, we have

$$\alpha\left(2P\left(\frac{n+1}{2}, n-1\right) - 1 + \epsilon_{n,p}\right) > \beta\left(2P\left(\frac{n+1}{2}, n-1\right) - 1 + \epsilon_{n,p}\right).$$

Thus, in this case, sincere voting is not a Nash equilibrium of simultaneous voting if and only if

$$2P\left(\frac{n+1}{2}, n-1\right) + \epsilon_{n,p} > 1. \tag{4}$$

Since for any n, p, and $\alpha > \beta$ Equation (3) is continuous in γ and δ, there exist γ and δ such that $\alpha > \gamma > \beta > \delta$, and for these parameters, sincere voting is not a Nash equilibrium of simultaneous voting. These parameters can be determined such that Eq. (1) will be simultaneously satisfied, since there, for any fixed p and $\alpha > \beta$, one should choose any γ close enough to α. Recall that the equation ensures that the unique subgame perfect equilibrium of sequential voting is that all voters play "vote with current leader". Therefore it only remains to compare the accuracy of the most accurate symmetric Nash equilibrium in simultaneous voting excluding sincere voting, which is "all vote for B", to the accuracy of the unique subgame perfect equilibrium of sequential voting, in which all voters play the "vote with current leader" strategy. The probability that "all vote for B" is accurate is simply the probability that B is the correct choice, which is $P\left(\frac{n+1}{2}, n\right)$. Computing the probability that the profile in which all play the "vote with current leader" strategy chooses the correct outcome is also not complicated:

Lemma 4. *In sequential voting, the probability that the correct outcome is chosen when all voters play the "vote with current leader" strategy is*

$$p \cdot P\left(\frac{n-1}{2}, (n-1)\right) + (1-p) \cdot \left(1 - P\left(\frac{n+1}{2}, (n-1)\right)\right) \tag{5}$$

Proof. If the first voter prefers B and votes for B, then B is chosen. This happens with probability

p. The probability that it is the correct choice equals the probability that at least $(n-1)/2$ out of the remaining $n-1$ voters prefer B, which equals to $P\left(\frac{n-1}{2},(n-1)\right)$. Similarly, if the first voter prefers A and votes for A, then A is picked. This happens with probability $1-p$. The probability that it is the correct choice equals one minus the probability that B is the correct choice, which is $P\left(\frac{n+1}{2},(n-1)\right)$. □

Merging these two observations, we have,

Lemma 5. *The profile of strategies in which all voters play "vote with current leader" is more accurate than all vote for B if and only if*

$$\frac{1}{2} > P\left(\frac{n+1}{2}, n-1\right) \tag{6}$$

Proof. Let $m = \frac{n-1}{2}$. By the previous observations it must hold that

$$p \cdot P(m, n-1) + (1-p) \cdot (1 - P(m+1, n-1)) > P(m+1, n) \tag{7}$$

Now, $P(m+1,n) = p \cdot P(m,n-1) + (1-p) \cdot P(m+1,n-1)$ follows by a simple probabilistic reasoning: obtaining at least $m+1$ B's out of n attempts can be achieved by either obtaining a B in the first attempt and then obtaining at least m B's in $n-1$ attempts, or by obtaining an A in the first attempt and then obtaining at least $m+1$ B's in $n-1$ attempts. Thus, rearranging Equation (7) we have $1 > 2 \cdot P(m+1, n-1)$, and the claim follows. □

Note that $P\left(\frac{n+1}{2}, n-1\right)$, the RHS of Equation (6), is increasing in p. Additionally, it is smaller than $\frac{1}{2}$ for $p=\frac{1}{2}$, and it is 1 for $p=1$, thus there exists $p_n \in (\frac{1}{2}, 1)$ such that the equation holds for any $p \in [\frac{1}{2}, p_n)$. In addition, p_n monotonically decreases to $\frac{1}{2}$ as n grows, since $p > \frac{1}{2}$. All this gives us:

Theorem 4. *For every odd number of players n there exist p and $\alpha > \gamma > \beta > \delta$ such that, for these parameters,*

- *"All vote B" is the most accurate Nash equilibrium of simultaneous voting.*
- *The profile of strategies in which all voters play "vote with current leader" is the unique subgame perfect equilibrium of sequential voting. The equilibrium outcome of this case is that all voters herd, vote the same as the first voter, which votes sincerely.*
- *The probability that the"vote with current leader" profile chooses the correct outcome is larger than the probability that "all vote B" chooses the correct outcome.*

However, the measure of the domain of all such parameters, relative to the domain of all valid parameters, goes to zero as n goes to infinity.

Proof. Fix any odd n, and any $\alpha > \beta$. Let $p = p_n$ be such that $1 = 2 \cdot P\left(\frac{n+1}{2}, n-1\right)$, and let $\alpha > \gamma > \beta > \delta$ be such that Eq. (1) and Eq. (4) are both satisfied with a strict inequality. (Such parameters exist as explained above in the paragraph that follows Eq. (4).) Thus, the first two requirements hold. In order to satisfy the third requirement as well, one needs to slightly decrease p, so that $1 > 2 \cdot P\left(\frac{n+1}{2}, n-1\right)$. Since everything here is continuous, this can be done while maintaining the previous two conditions. Since p_n goes to $\frac{1}{2}$ as n increases, the second part of the theorem follows. □

5 Concluding Remarks

We study a model which differs from classic models on sequential decision making and on sequential voting in two important aspects: First, in our model utilities are completely subjective and voters do not care about the "correct" alternative; Second, utilities have a significant externality, as a player dislikes to vote against the chosen winner. Similar to prior models, herding behavior also exists in our model. Specifically, with an infinite population and a required vote difference M for winning, players will start voting for the current leader in votes, once a small threshold gap k^* has been reached. We show that this threshold gap depends only on the details of the utility function, but does *not* depend at all on the required vote difference, M. Thus, due to the strategic behavior of the voters, increasing M does *not* result in aggregation of more votes in the decision.

Additionally, with finite population we show that there exist cases in which sequential voting yields more accurately the correct winner than simultaneous voting. This happens when the prior probability that a player will prefer B over A is only slightly larger than $\frac{1}{2}$ (which is exactly the case in which voting is necessary to determined the alternative with a realized majority). We show that for any number of players, there exist parameters of the utility function that cause sincere voting to be a non-equilibrium strategy for simultaneous voting, while in sequential voting, with these parameters, the "vote with current leader" profile is the unique subgame-perfect equilibrium. Furthermore, this profile chooses the alternative preferred by the majority of voters with probability higher than that of the best Nash equilibrium of simultaneous voting. This is notable especially when the population of voters is finite, and of relatively small size.

Since our motivating example is social networks, an interesting extension is to examine how the actual network structure affects the magnitude of the herding phenomenon. This is related to a series of works on diffusion in networks, for example as in Morris (2000). There, on the one hand, utility externalities are "built-in" in the model, as a node's utility is affected by the choices of her neighbors. But on the other hand, the dynamics are usually assumed to be myopic, and thus herding does not occur in these works. It would be challenging to integrate our insights here with the network structure.

References

Ali, S. and N. Kartik (2011). "Herding with collective preferences." *Economic Theory.* In press.

Aronson, E., T. D. Wilson and R. M. Akert (1997). *Social Psychology.* New York: Longman.

Banerjee, A. (1992). "A simple model of herd behavior." *The Quarterly Journal of Economics*, **107**(3), 797.

Bartels, L. M. (1988). *Presidential Primaries and the Dynamics of Public Choice.* Princeton, NJ: Princeton University Press.

Bikhchandani, S., D. Hirshleifer and I. Welch (1992). "A theory of fads, fashion, custom, and cultural change as informational cascades." *Journal of political Economy*, pp. 992–1026.

Callander, S. (2007). "Bandwagons and momentum in sequential voting." *Review of Economic Studies*, **74**(3), 653–684.

Dasgupta, S., K. Randazzo, R. Sheehan and K. Williams (2008). "Coordinated voting in sequential and simultaneous elections: some experimental evidence." *Experimental Economics*, **11**(4), 315–335.

Dekel, E. and M. Piccione (2000). "Sequential voting procedures in symmetric binary elections." *Journal of Political Economy*, **108**(1), 34–55.

Desmedt, Y. and E. Elkind (2010). "Equilibria of plurality voting with abstentions." In *Proceedings of the 11th ACM conference on Electronic Commerce (EC)*, pp. 347–356.

Easley, D. and J. Kleinberg (2010). *Networks, crowds, and markets: Reasoning about a highly connected world.* Cambridge Univ Pr.

Morris, S. (2000). "Contagion." *The Review of Economic Studies*, **67**(1), 57.

Spitzer, F. (2001). *Principles of random walk.* Springer Verlag.

Xia, L. and V. Conitzer (2010). "Stackelberg voting games: Computational aspects and paradoxes." In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pp. 921–926.

A Comparison to the model of Ali and Kartik (2011)

As discussed in the introduction, Ali and Kartik (2011) also recognize the importance of allowing the utility of a voter to depend on previous votes. However they only allow the utility to increase as more votes become aligned to the "correct" alternative (determined by state of nature), and not as they become aligned to the *chosen* alternative. This is a subtle point, and we wish to demonstrate the issue via an example.

Suppose a situation where all players like A. One can model the players' preferences as the state of nature, saying that in this case the correct alternative is A. Suppose player i indeed votes A, but the majority of votes is for B, assume that in this case i's resulting utility is X. Now, if some of the others change their vote to A and the majority becomes A, suppose i's utility increases to $Y > X$. The reason for this increase in utility is not clear, as initially i votes for the loser, and also the correct alternative loses. Then some players change their vote and the winner changes. So now i votes for the winner, and also the correct alternative wins. The reason that i's utility increases can be that she now votes for the winner ("first possible reason"), or that the winner is now the correct one ("second possible reason"). One cannot distinguish between these two reasons by looking at this scenario alone.

Next we introduce an additional scenario that enables us to distinguish the two reasons. Suppose everyone besides i like B (i's preference does not matter). Now the "correct" alternative is B. Suppose as before, that i votes A and the majority of votes is for B, which results in a utility of X' for i. As before, some of the other voters change their vote to A, the majority becomes A, and i's utility is now Y'. Is Y' smaller or larger than X'? The first possible reason ("he now votes for the winner") still holds in the new setting, so it must be that $Y' > X'$. The second possible reason ("the winner is now the correct one") does not hold in the new setting, so it is possible that $Y' \leq X'$. Ali and Kartik (2011) *require* $Y' \leq X'$. Thus, the reason for the first scenario cannot be the first possible reason. Our paper, on the other hand, focuses *only* on the first reason. Indeed, their result resembles a classical herding result, showing that a herd will start when the information regarding state of nature deduced from *previous* votes out-weights the strength of the signal the current voter gets regarding state of nature. Thus, herding in their model results from informational uncertainties, not from the utility externality.