Squid Voting Game

Rational Indecisiveness in Sequential Voting*

Yuichiro Kamada† and Yosuke Yasuda‡

October 4, 2022

Abstract

We consider a model of common-value sequential voting in which voters are differentiated in their information. We ask whether the intuition as in the simultaneous-voting case—voters with no information would vote so as not to influence the outcome—would be valid to imply long voting in our sequential setting. We find that any voting outcome, including short voting, can arise in equilibrium, and hence the intuition from the simultaneous voting does not apply. We discuss conditions under which long voting results. We also show that a voter may vote against her information in equilibrium, and that may improve welfare.

*We thank Yasushi Asako, Yukio Koriyama, Dmitriy Kvasov, Constantine Sorokin and seminar participants at Boston College, Stony Brook University, UC Berkeley, University of Glasgow and Waseda University for their helpful comments. All remaining errors are ours. Yasuda gratefully acknowledges research support from Grant-in-Aid for Scientific Research (KAKENHI), 18KK0342. He would also like to thank the ISEG, the University of Lisbon, for providing an excellent research environment as a visiting researcher, which enables him to proceed with much academic research, including this project. Masanori Kobayashi provided excellent research assistance.

†University of California, Berkeley, Haas School of Business, Berkeley, 2220 Piedmont Avenue, CA 94720. E-mail: `y.cam.24@gmail.com`

‡Osaka University, Department of Economics, 1-7 Machikaneyama, Toyonaka, Osaka 560-0043 Japan. E-mail: `yosuke.yasuda@gmail.com`

1 Introduction

In episode 2 of the blockbuster Netflix show *Squid Game*, voting was held to determine if the 201 "players" in the game would continue or stop the game. The votes were cast sequentially, in the order of the players' jersey numbers. All players watched each other's vote, and the current vote counts were displayed on a screen in front of them. The voting was close and took a long time, and the 201st voter, a fragile elderly man, took his turn.[1]

Sequential voting is used on various occasions. Besides the *Squid Game* show, it is used in the roll call vote in United Nations, the US Senate, and the Japanese Diet, among others.[2] In those settings, voters would not be perfectly informed of the favorability of options (each Senate or Diet member may face uncertainty about the future evolution of the economy, and it is still episode 2 of the show in *Squid Game*). How would voters vote if they are uncertain about the favorability of the options?

Voter behavior with uncertainty has been discussed in static frameworks, notably by the seminal work by Feddersen and Pesendorfer (1996). They analyzed a common-value framework and showed that, in equilibrium, a voter with imprecise information about the state of the world (so she is uncertain which option is better) would prefer to abstain, thereby essentially delegating the decision to other voters who might be more informed. If the same incentive as in Feddersen and Pesendorfer (1996) was at work in the sequential setting with no abstention option (as in all the above examples), then we would expect a *long voting*: the pivotal voter, if she was uninformed, would vote against the current majority, thereby having the voting continued and delegating the decision to the subsequent voters. In other words, the pivotal voter would be *rationally indecisive*.

We show that this intuition is not entirely correct. To understand the incentives in the sequential setting, we analyze a stylized model of sequential voting, abstracting away from various aspects that may exist in different contexts of real-world voting. In our model, there are two states of the world with equal probability, α and β. Conditional on each state, each voter can (independently) be perfectly informed about the state with a small probability, or else they are completely uninformed. Each voter votes for either A or B until a majority has cast votes for one of the options, which determines the outcome. Voters receive payoff 1 if the outcome "matches" the state and 0 otherwise. Our first main result says "anything goes": For any sequence of votes that can feasibly realize in the game, there is an equilibrium

[1] Which option did he vote? Watch *Squid Game*.

[2] In the Japanese Diet, the roll call vote, known as *Kimei-tohyo*, is done by each Diet member coming up to the front stage and giving to *Sanji* a wooden plate that is either white (representing the "yes" vote) or blue (representing the "no" vote). The *Sanji* would then put the plates on a meter according to their color, and this meter is visible to everyone in the Diet. The votes can easily be counted at any moment because there are lines on the meter per every five plates.

that achieves such a sequence if no one is informed. That is, not only long voting but also any length of voting can be an outcome of our voting game.

Why does anything go? In our voting model, Feddersen and Pesendorfer (1996)'s intuition is indeed at work: If a pivotal voter is uninformed, she would like to delegate the decision to the subsequent voters. We call this effect the *delegation effect*. This, however, is only one side of the story. There is another side: *A voter's optimal vote depends on what she has learned from the preceding voters' choices and the effect of her vote on the subsequent voters' choices.* This learning makes it appealing for the pivotal voter to end the game (because that means she votes for the current majority), and we call this effect the *posterior effect*. If this pivotal voter is supposed to vote to end the game conditional on being uninformed, the posterior effect implies that it is costly for her not to end the game. This is because not doing so means that she essentially informs the subsequent voters that she knew the state (even though she does not!) and this state is the one that she thinks is less likely. We obtain the "anything goes" result due to the balance of these two effects.

To show the "anything goes" result, we construct equilibria that use "consistent strategies": the voter votes for A if she is informed of the state α and votes for B if she is informed of the state β. We further show that any consistent equilibria induce the same payoff to voters.

Furthermore, we show the existence of inconsistent equilibria. The welfare in inconsistent equilibria can take a wide range: we provide the tight upper and lower bounds of the equilibrium payoffs achievable by inconsistent equilibria. In particular, the upper bound is the maximum payoff achievable by any social choice function. The welfare from consistent equilibria positions in between these bounds.

Given that the "anything goes" does not give us any prediction about what happens in our voting game, we seek a way to enrich the model to obtain some prediction. Specifically, we find that the following two modifications to the model would imply a "long voting" as a *unique* equilibrium outcome. First, we assume the last voter trembles slightly (as in *Squid Game*). Second, we assume that each voter experiences a small cost by ending the game. We show that each of these two modifications tips over the balance of the posterior effect and delegation effect to make the pivotal voter rationally indecisive, and the voting goes to the last voter under any equilibria if no voter is uninformed.

Our model is stylized to capture the essence of the implications of uncertainty in sequential voting. As such, for each of the real-world voting applications, there are ways in which the model may not perfectly fit it. For example, in *Squid Game*, although our model may perfectly capture the timing of moves and observation structure, it could be questionable if our payoff function represents the episode's situation well.[3] In United Nations, different countries may have different information and preferences. In US Senate and Japanese Diet, different politicians may belong to different parties and hence may have different preferences. These voters, however, may still have common uncertainty about the state of the world, and the voting may have a common-value component. The fact that our model is stylized enables us to capture the key incentives faced by the voter with uncertainty in a common-value sequential environment by abstracting away from various details.

There are many ways to enrich the model. In Section 6, we provide various discussions and generalizations of the model. For example, we discuss how robust our "anything goes" result is to various types of asymmetries in the model and to the case when voters can only observe vote counts, not the entire history of voting.

1.1 Literature Review

We ask how the differentiated precision of signals across voters affects the behavior in a sequential voting setup. This question is asked in the simultaneous voting setting by Feddersen and Pesendorfer (1996) that has an otherwise similar setup as ours. As we discussed, the fact that the voting is sequential changes the picture quite a bit.

In what follows, we review the papers on sequential voting. First, Dekel and Piccione (2000)'s seminal work considers a model of sequential voting with an information structure and payoffs that are more general than ours. Their focus, however, is on history-independent strategies that are equilibria under simultaneous voting and whether they also constitute equilibria under a sequential setting. We, in contrast, consider history-dependent strategies. They also have an analysis on a common-value case, which is the same environment as ours. Under such a restriction, however, they restrict attention to the case of two signals, and thus the question of how the differentiated precision of signals across voters affects the behavior—which is our main point—cannot be answered in their model.[4]

One can think of our model as that of herding. The main difference is that in the canonical model of herding (e.g., Banerjee (1992)), each player's payoff is independent of the actions by other players. In our model, however, each voter's payoff depends on other voters'

[3]One (admittedly farfetched) interpretation is the following: In the show, all players were in serious financial debt. If the players had continued the game, they would have received a large monetary prize if they had won (otherwise, they would have been killed). If they had stopped the game, they would not have received money and gone back to a life of financial ruin. No one was supposed to know what sort of games they would play to determine the winner upon continuation. In a sense, every player's situation was similar. Of course, players might have different beliefs from each other about the likelihood of their own winning, which would have introduced some private-value component to the model.

[4]See also their extension to the three-option case with a choice of the timing of moves (Dekel and Piccione, 2014).

behavior because it affects the voting outcome. Ali and Kartik (2012) consider a general model of interdependent preferences that includes elections as a special case and show that herding occurs under certain restrictions. Their model allows for only two signals, and thus does not answer the question about the differentiated precision of signals. Also, no signal is perfectly informative about the state in their model. In our model, however, there is a signal that perfectly reveals the state. Hence, an action by a voter may substantially change the belief held by the subsequent voters, even if those voters initially had a very skewed belief.

Callander (2007)'s sequential election model can also be thought of as a herding model with interdependent preferences. He assumes the players with "conformity" incentives (the desire to vote for the majority) and shows that herding can occur in equilibrium. We do not assume conformity, and introducing such an incentive would change the result.[5] Another difference is that, in his model, there are infinitely many voters, and thus no voter would be able to unilaterally end the game. In contrast, our model has finitely many voters, and hence someone becomes pivotal. We show that the pivotal voter faces the following trade-off: ending the game with the outcome that has a higher posterior versus not ending the game to delegate the decision to the subsequent voters. Such a trade-off does not exist in Callander (2007)'s model.

2 Model

There are voters $1, 2, \ldots, 2n+1$, where n is a non-negative integer. The voters sequentially vote on either A or B. We assume the majority voting rule: A results as soon as $n+1$ voters vote for A, and B results as soon as $n+1$ voters vote for B. Formally, let the set of *sequences* be

$$\mathcal{S} := \bigcup_{k=n+1}^{2n+1} \{(X_1, \ldots, X_k) \in \{A, B\}^k \,||\{l \in \{1, 2, \ldots, k\} | X_l = X_k\}| = n+1\}.$$

That is, a sequence is a path of actions that ends the game. For any $S = (X_1, \ldots, X_k) \in \mathcal{S}$, call X_k the *outcome* of S.

The state of the world is either α or β. The prior probability for each state is $1/2$. When the state is α, each voter, when she is called upon to vote, independently receives signal a with probability $\varepsilon > 0$; otherwise (i.e., with probability $1 - \varepsilon$), she receives no signal, or formally we say she receives signal $\emptyset$. Similarly, when the state is β, each voter, when voting,

[5]Indeed, Appendix Appendix A.6.1 considers a version of an incentive to end the game and shows that long voting occurs in the unique equilibrium.

independently receives signal b with probability ε and signal $\emptyset$ with probability $1-\varepsilon$. Given this simple information structure, signals a and b perfectly reveal the state of the world.

All voters receive the payoff of 1 if the outcome of the realized sequence coincides with the state, i.e., the outcome is A when the state is α or it is B when the state is β; otherwise, all voters receive the payoff of 0. Note that the model is characterized by parameters (n, ε).

For voter k, a *history* is an enumeration of actions by $k-1$ voters before k, and her own signal. Thus, the set of voter k's histories is

$$H_k = \{A, B\}^{k-1} \times \{a, b, \emptyset\}.$$

Voter k's mixed strategy is a mapping $\sigma_k : H_k \to \Delta(\{A, B\})$. Let Σ_k be the set of all strategies of voter k and let $\Sigma = \times_k^{2n+1} \Sigma_k$.

In this paper we consider the concept of sequential equilibrium. We sometimes refer to it as "equilibrium" for short.

We say that a strategy σ_k is *consistent* if $\sigma_k(\cdot, a)(A) = 1$ and $\sigma_k(\cdot, b)(B) = 1$. Although consistency may be a natural property, it will turn out that there are equilibria that are not consistent and, in fact, they may have a greater expected payoff than any equilibria that are consistent.

Given any pure strategy profile $\sigma \in \Sigma$, there is a unique sequence S that realizes if no voter receives a signal. We call such S the *default sequence of σ*.

3 Consistent Strategies

Our first set of results pertains to equilibria that use consistent strategies. Section 3.1 shows that any sequence can be a default sequence of a consistent equilibria, and Section 3.2 shows that any such equilibria have the same welfare.

3.1 Anything Goes

Our first result is rather striking. It says: anything goes!

Theorem 1. *For any $n \geq 0$, $\varepsilon \in (0, 1]$ and any sequence $S \in \mathcal{S}$, there is a pure consistent strategy equilibrium σ in the model with (n, ε) such that S is the default sequence of σ.*

That is, any path of actions can be a default sequence of an equilibrium. This in particular implies that paths in which the voting ends early (e.g., $n+1$ straight As) can also realize in equilibrium. This conclusion contrasts with the intuition of Feddersen and Pesendorfer (1996) that a voter who does not have information would like to leave the decision to other voters.

The proof constructs an equilibrium for each S. A rough intuition for the construction and why it works is as follows. If voter k without information is supposed to play A in the default sequence S but plays B, since such a deviation can only be caused when k has signal b, the subsequent voters with no information will rationally update their belief to assign probability 1 to state β. Moreover, even in the case when some of the subsequent voters receive signal a and play A, we let the voters without information to have a belief that assigns probability 1 to β and play B.[6] This will make it hard for the subsequent voters to use the information they receive. Hence, voter k has an incentive to follow the default sequence when she does not receive a signal. We denote the strategy profile constructed in this way by σ^S.

When we presented this result in seminars, one common reaction was that the result would be trivial because the voter with no information assigns equal probability to each state so must be indifferent, and this should be why "anything can happen." This argument misses the point that, in consistent equilibria, each voter's action changes the belief of the subsequent voters. Thus, each voter's best response depends on the action sequence by the preceding voters (which affects her posterior belief) as well as on the effect of her action on the subsequent voters' beliefs. Indeed, the voters do not necessarily have posterior beliefs that assign equal probability to each state, and thus they are not necessarily indifferent. In fact, on the path of play of any pure consistent equilibrium, a voter is indifferent if and only if she is *deciding*: Under σ^S, say that voter k is *deciding* if when the history up to k has been consistent with S, if no voter in $\{1, \ldots, 2n+1\}$ receives a signal, the outcome varies by k's vote, and one of k's vote immediately ends the game. For example, when $n = 2$, if $S = (A, B, A, B, A)$, then only $k = 5$ is deciding. Note that $k = 4$ is not deciding because the outcome is A irrespective of his action (given that no one receives a signal and everyone else follows σ^S). If $S = (A, A, B, B, B)$, then $k = 3, 4, 5$ are all deciding.

Now, consider voter $k < 2n + 1$ who is deciding and suppose, without loss of generality, that the game ends if she plays A.[7] Since k is deciding, if she plays B and no subsequent voter receives a signal, then all the subsequent voters will play B and thus the last voter $2n+1$ gets to play. The proof of Theorem 1 shows that voter k is indifferent between playing A and B under this situation. In what follows, we provide three explanations for why this is true.

The first explanation is due to the comparison of the benefit of deciding on the currently superior outcome and postponing the decision to benefit from the subsequent voters' infor-

[6] Such a belief can be shown to satisfy consistency of Kreps and Wilson (1982).

[7] If $k = 2n + 1$, then k is indifferent because the posterior assigns equal probability to each state and k's vote determines the outcome.

mation: Voter k assigns a high posterior on the state being α because A has been played n times while B has been played less than n times. This makes playing A attractive for voter k. We call this effect the *posterior effect*. On the other hand, passing the decision to the subsequent voters makes it possible to use the information obtained by those voters. This makes playing B attractive. This intuition is analogous to the one in Feddersen and Pesendorfer (1996) that would make the voter "rationally indecisive," and we call this effect the *delegation effect*.

Why do these two effects cancel each other? If k plays A, then the game ends with the outcome A with the probability that is the posterior belief on α, which we denote by $P(\alpha)$. Note that $P(\alpha) > 1/2 > 1 - P(\alpha)$ as we discussed. Hence, the benefit from playing A due to the higher posterior on α can be expressed as $P(\alpha) - (1 - P(\alpha))$. If instead she plays B, then voter k gets payoff 1 if and only if the game ends with outcome A when the state is α and at least one of the subsequent voters receives a signal, or when the state is β. Thus, the benefit from playing B due to the information that a subsequent voter may receive can be expressed as $P(\alpha) \cdot (1 - (1 - \varepsilon)^{2n+1-k})$, which is equal to $P(\alpha) - P(\alpha)(1 - \varepsilon)^{2n+1-k}$. Now, since $P(\alpha) = \frac{(1-\varepsilon)^{k-1-n}}{(1-\varepsilon)^{k-1-n}+(1-\varepsilon)^n}$ and $1 - P(\alpha) = \frac{(1-\varepsilon)^n}{(1-\varepsilon)^{k-1-n}+(1-\varepsilon)^n}$, we have:

$$1 - P(\alpha) = P(\alpha)(1 - \varepsilon)^{2n+1-k}.$$

This is why the two effects cancel each other.

The second explanation is based on a state-by-state comparison: If the state is α, with probability $(1 - \varepsilon)^{2n+1-k}$, playing B decreases the payoff by 1 (relative to playing A). If the state is β, with probability 1, playing B increases the payoff by 1 (again, relative to playing A). The ratio of the posteriors on the two states is $(1 - \varepsilon)^{k-1-n}$ to $(1 - \varepsilon)^n$, that is, 1 to $(1 - \varepsilon)^{2n+1-k}$.

The third explanation is due to the consideration of the event in which voter k becomes pivotal: If k chooses A, then that ends the game. If she chooses B, then if at least one of the subsequent voters plays A then the outcome is A. If all the subsequent voters play B, then that means that (i) n other voters played A when they were supposed to play A given no signal, and (ii) n other voters played B when they were supposed to play B given no signal. Hence, conditional on this event, voter k is indifferent between A and B. Hence, her expected payoff from playing B would not change even if we modified the game and hypothetically assigned outcome A to this event. With this modification, the outcome from k's playing B is always A. Thus, voter k must be indifferent between playing A and playing B.

Let us formally state the result about when a voter is indifferent.

Remark 1. Fix (n, ε). For any $S \in \mathcal{S}$, consider σ^S and voter k who observed a history that does not contradict S. When voter k does not receive a signal, she is indifferent between voting for A and B if and only if k is deciding.

One may wonder our equilibrium construction may hinge "too much" on the observation of the particular default sequence. However, the "anything goes" result continues to hold even if we allow voters to observe only the vote counts, where we modify our equilibrium construction. We present this result in Section 6.4.

Finally, we note that the result depends on the symmetry of the model. However, we still believe the "anything goes" result is useful for two reasons. First, in Section 6.1, we examine the case with asymmetry in the prior, the signal probabilities, and the voting rule. We show that if the asymmetry is introduced in the way that "favors" a particular outcome (A or B), then any sequence that induces the favored outcome is a default sequence of some pure consistent equilibrium. Thus, "anything goes" is still "half true" with a restriction on the outcome. Second, as in our explanation after its statement, our "anything goes" theorem gives us a way to understand the incentives faced by the voters in the most simplified setting. The two effects we identified—the posterior effect and the delegation effect— will help us understand why a particular asymmetry would imply the half "anything goes" result that we present in Section 6.1.

3.2 Welfare Invariance

Take an arbitrary pure consistent equilibrium σ. Suppose without loss of generality that σ induces the outcome of A with probability 1 if no voter receives a signal. If the state is α, then no one receives signal b and thus those who are supposed to play A on the default sequence of σ would never switch their actions to B. Consequently, the outcome becomes A with probability 1. On the other hand, if the state is β, there are $n+1$ chances to overturn the outcome. This occurs if some voter who is supposed to play A on the default sequence of σ receives signal b and thus plays B. Note that if such an event occurs, then by the Bayes rule, all the subsequent voters must have a belief that assigns probability 1 to β, and thus the outcome will be B. Therefore, the outcome becomes B with probability $1-(1-\varepsilon)^{n+1}$ for any such σ. Overall, the ex ante payoff under σ is

$$\frac{1}{2}\cdot 1+\frac{1}{2}(1-(1-\varepsilon)^{n+1}) = 1-\frac{(1-\varepsilon)^{n+1}}{2}.$$

Thus, we have shown the following.

Theorem 2. *For any (n, ε), pure consistent strategy equilibrium in the model with (n, ε) has the ex ante payoff of* $1 - \frac{(1-\varepsilon)^{n+1}}{2}$.

This welfare invariance can be suitably extended to the settings with asymmetry (Corollary 2 in Section 6.1): all the pure consistent equilibria have the same payoff.

4 Inconsistent Strategies

So far, we have restricted our analysis to equilibria that use consistent strategies. Perhaps surprisingly, there exist equilibria that use inconsistent strategies. This section is devoted to an analysis of those equilibria. Specifically, we consider two types of inconsistent strategies. In Section 4.1, we consider a strategy that lets the voter vote for B given signal a and for A given signal b. In Section 4.2, we consider a strategy that lets the voter vote for a particular action regardless of the signal. In both sections, the main issue is the welfare implications of the equilibria that use inconsistent strategies.

4.1 Voting against a Signal Improves Welfare

The following example shows that there exists an equilibrium that uses inconsistent strategies in which a voter votes against her signal, and the equilibrium achieves a strictly larger ex ante payoff than the one under any consistent equilibrium.

Example 1. Assume there are 3 voters. The following strategy profile constitutes an equilibrium, which we denote by $\bar{\sigma}$.

- Voter 1 chooses B if she receives a; otherwise, she plays A.
- Voter 2 uses a consistent strategy that plays the opposite action to voter 1's if he receives no signal.
- Voter 3 uses a consistent strategy that plays the same action as voter 2's if she receives no signal.

Note that the default sequence of this strategy profile $\bar{\sigma}$ is (A, B, B). Here we explain that voter 1 has a strict incentive to vote against her signal under $\bar{\sigma}$.[8] Consider voter 1. Suppose she receives signal a. This happens only if the state is α, and hence, the subsequent voters receive either a or ϕ. Therefore, if voter 1 follows the above strategy by choosing B, her expected payoff becomes 1 since voters 2 and 3 play A with probability 1. If she instead

[8] The proof of Theorem 3 fully checks the voters' incentives to show that this strategy profile is indeed an equilibrium.

plays A, then her payoff is $1-(1-\varepsilon)^2$. Hence, she strictly prefers to play B. Suppose voter 1 receives signal b. If she plays A, her payoff is 1. If she plays B, her payoff is $1-(1-\varepsilon)^2$. Hence, she strictly prefers to play A.

We compute the ex ante expected payoff from $\bar{\sigma}$. First, if the state is β, then the outcome is B with probability 1. Second, if the state is α, then outcome A realizes if at least one voter out of the three voters receives signal a. Thus, the payoff is

$$\frac{1}{2}\cdot 1+\frac{1}{2}[1-(1-\varepsilon)^3].$$

Note that the payoff from any pure consistent equilibrium is

$$\frac{1}{2}\cdot 1+\frac{1}{2}[1-(1-\varepsilon)^2].$$

Hence, the payoff is strictly greater under $\bar{\sigma}$ than in any pure consistent equilibrium. The intuition is that, under $\bar{\sigma}$, there are three chances to overturn the outcome when the state is α. □

The equilibrium construction in the above example can be generalized. The next theorem formalizes this point. Let $\bar{R}_n := 1-\frac{(1-\varepsilon)^{2n+1}}{2}$. Note that when $n=1$, $\bar{R}_n$ is the equilibrium payoff we derived in the above example.

Theorem 3. *For any n, there is an equilibrium σ that achieves the payoff $\bar{R}_n$.*

The proof in the Appendix shows that a strategy profile generalizing $\bar{\sigma}$ in Example 1 is an equilibrium.

We now show that $\bar{R}_n$ is the highest feasible payoff that can be achieved. To formalize this statement, given n, consider a social choice function $f:\{a,b,\emptyset\}^{2n+1}\to\Delta(\{A,B\})$ that returns a probability distribution over outcomes for each realization of signals of all the $2n+1$ voters. Note that for any strategy profile σ in our original game, there is a social choice function f that achieves the same distribution over outcomes conditional on any realization of the state and the signals. Let F_n be the space of all social choice functions. Notice that a social choice function $f\in F_n$ determines the ex ante expected payoff to each voter. Let this payoff be $R_n(f)$.

Proposition 1. *For all $f\in F_n, R_n(f)\le\bar{R}_n$.*

Proof. Fix any $f\in F_n$. Let $f(\emptyset,\dots,\emptyset)(A)=q$. Then, conditional on the signal realization $(\emptyset,\dots,\emptyset)$, the voters' expected payoff is $\frac{1}{2}q+\frac{1}{2}(1-q)=\frac{1}{2}$. The probability that the signal profile $(\emptyset,\dots,\emptyset)$ realizes is $(1-\varepsilon)^{2n+1}$. Hence, the ex ante expected payoff is at most $(1-(1-\varepsilon)^{2n+1})\cdot 1+(1-\varepsilon)^{2n+1}\cdot\frac{1}{2}=1-(1-\varepsilon)^{2n+1}\cdot\frac{1}{2}=\bar{R}_n$. □

We note that the payoff $\bar{R}_n$ can be realized by a social choice function that, whenever there is at least one signal, the outcome accords with that signal (and returns an arbitrary outcome when there is no signal). The intuition is that the social choice function asks every voter if they have received a signal and returns the outcome according to the voter's response. What Theorem 3 shows is that we can construct an equilibrium corresponding to such a social choice function. In contrast, Theorem 2 shows that the consistent strategy equilibria correspond to asking only $n+1$ voters about their signals.

4.2 Ignoring a Signal Leads to Inefficiency

The following example shows that there exists an equilibrium in which some voters ignore their signals, which induces a payoff lower than that of pure consistent equilibria.

Example 2. Consider the case with 5 voters and the following strategy profile, which we denote by $\underline{\sigma}$.

- Voters $k = 1, 2$: Voter 1 plays A irrespective of the signal. Voter 2 plays A irrespective of the signal if voter 1 has played A. Otherwise, he plays a consistent strategy in which he plays A if no signal is received.
- Voters $k = 3, 4$: If all the actions observed so far are 2 times of A followed by $k-3$ times of B, then play B irrespective of the signal. Otherwise, play a consistent strategy in which she plays A if no signal is received.
- Voter $k = 5$: Play a consistent strategy in which she plays A if no signal is received.

The default sequence of this strategy profile $\underline{\sigma}$ is (A, A, B, B, A). Note that $\underline{\sigma}$ has the ex ante payoff of $\frac{1}{2} + \frac{1}{2}\varepsilon = \frac{1+\varepsilon}{2}$. We can show that when $\varepsilon > 0$ is sufficiently small, $\underline{\sigma}$ constitutes an equilibrium with the belief that, after any deviation, any voter who has not received a signal assigns probability 1 to α. Here, let us explain why voters have incentives to ignore the signals on the path of play.[9]

First, voters 1 and 2 play A even if they receive signal b. The reason is that if they play B, then the subsequent voters believe that the state is α. For the outcome to be B under such a situation, there need to be at least two subsequent voters who receive signal b, which happens with probability $O(\varepsilon^2)$.[10] In contrast, if they follow $\underline{\sigma}$ and play A, then the outcome

[9] The proof of Theorem 4 fully checks the voters' incentives to show that this strategy profile is indeed an equilibrium.

[10] We note that the argument depends on $\varepsilon > 0$ being small. To see this, consider voter 2 who observes voter 1's A and receives signal b. If he plays A, then the outcome will be B if and only if voter 5 receives signal b, which happens with probability ε. Hence, his payoff in this case is ε. If he plays B instead, then the

will be B if and only if voter 5 receives signal b, which happens with probability ε. Thus, it is indeed optimal to ignore the signal and play A.

Second, voters 3 and 4 play B even if they receive signal a. The reason is that even if they do so, the outcome will be A with probability 1 according to $\underline{\sigma}$ given that the state is α, which means that playing B ensures that the payoff is 1. Thus, ignoring the signal and playing B is indeed optimal. □

The equilibrium construction in the above example can be generalized. The next theorem formalizes this point. Recall that the equilibrium constructed in the above example achieves the payoff of $\frac{1+\varepsilon}{2}$.

Theorem 4. *For any $n \geq 2$, there is $\bar{\varepsilon} > 0$ such that for all $\varepsilon \in (0, \bar{\varepsilon})$, there is an equilibrium that achieves the payoff of $\frac{1+\varepsilon}{2}$.*

The proof in the Appendix shows that a strategy profile generalizing $\underline{\sigma}$ in Example 2 is an equilibrium.

We now show that the payoff identified in Theorem 4 is the worst possible payoff. The meaning of "worst" needs an explanation because we do not consider all social choice functions. For example, imagine social choice functions that always return the outcome that is against the signals obtained by the voters. Such social choice functions induce a payoff that is obviously strictly below the payoff that any equilibrium of our sequential voting game can achieve. For this reason, we seek to obtain the lowest possible pure strategy equilibrium payoff that can be achieved in our sequential voting game.

Proposition 2. *Given any (n, ε), the expected payoff from any pure strategy equilibrium is at least $\frac{1+\varepsilon}{2}$.*

The proof is involved. We show that the last voter on the default sequence uses a consistent strategy[11], and use this property to bound the payoffs of the previous voters inductively. We note that the payoff $\frac{1+\varepsilon}{2}$ corresponds to the payoff from the social choice function that asks only one voter about her signal and returns the outcome that accords with her signal if she receives one. This payoff may appear quite low, but it is strictly greater than an equilibrium payoff in the simultaneous-move version of the game: If $n > 1$ (i.e., there are 3 or more voters), the simultaneous-move voting has an equilibrium in which every voter votes for A regardless of the signal, and this equilibrium has the payoff of $\frac{1}{2}$.

---

outcome will be B if and only if two or three out of three subsequent voters receive signal b. The probability of this event is $3\varepsilon^2(1-\varepsilon)+\varepsilon^3$, and hence his payoff in this case is $3\varepsilon^2(1-\varepsilon)+\varepsilon^3$. The latter payoff is greater than the former if and only if $\varepsilon < 3\varepsilon^2(1-\varepsilon)+\varepsilon^3$, or $(2\varepsilon-1)(\varepsilon-1) < 0$, which holds if and only if $\varepsilon > \frac{1}{2}$.

[11] This means that we need to rule out the strategy of, for example, playing A to end the game given no signal and playing B to continue the game given signal a or b.

Let us summarize this section. There are various kinds of inconsistent strategy equilibria, and their welfare ranges from the best payoff achievable under any social choice function, which corresponds to asking all the $2n+1$ voters about their signals, to the one that corresponds to asking just a single voter about her signal. The payoff under the consistent equilibria corresponds to asking $n+1$ voters about their signals and thus positions between those two bounds.

5 Long Voting

A simple application of the Feddersen and Pesendorfer (1996)'s intuition might lead one to conjecture long voting, but the "anything goes" result shows that long voting is only one of many possible predictions. In this section, we consider two modification of our model that leads to the unique prediction of long voting.

5.1 When the Last Voter Is "Crazy"

Consider a modified game in which voter $2n+1$ is constrained to choose a mixed action (he trembles, as in the last voter in *Squid Game*). Formally, for any $\xi \in (0, \frac{1}{2}]$, a *ξ-tremble game* is a game that is the same as the one in the main section, except that voter $2n+1$, upon moving, chooses a probability distribution over $\{A, B\}$ such that each action receives probability at least ξ if he does not receive a signal. We say that a strategy profile in a ξ-tremble game is *pseudo-pure* if every voter from 1 to $2n$ uses a pure strategy.

Theorem 5. *In any ξ-tremble game, there is no consistent pseudo-pure strategy equilibrium such that the game ends before reaching voter $2n+1$ when no voter receives a signal.*

The intuition for the proof is as follows. Consider voter k who is deciding (see the discussion after Theorem 1 for the definition of "deciding"), and suppose that this voter is not the last voter in the game (i.e., $k < 2n+1$). Suppose without loss of generality that the game ends if she plays A. We know that the posterior effect and the delegation effect make actions A and B attractive, respectively, and consequently, k is indifferent between playing A and B. This is based on the fact that all the subsequent voters receiving no signal will play B after k plays B. However, if there is a chance that, after k's playing B, a subsequent voter receiving no signal plays A, then voter k is no longer indifferent between the two actions. This is because, since state α is more likely than β, the posterior effect makes playing B more attractive as well, which makes voter k "rationally indecisive". This is why it is impossible to sustain k's playing A as an equilibrium action.

5.2 Small Incentive Not to End

Consider a model in which we change the payoff function of each voter so that if voter k votes for X when the history has n times of X, then k incurs a positive cost $c > 0$. Call this game the *pivot-aversion game.*

Theorem 6. *In any pivot-aversion game, there is no pure consistent strategy equilibrium such that the game ends before reaching voter* $2n + 1$ *when no voter receives a signal.*

The intuition is simple. In the original voting game, as long as voters play a pure consistent equilibrium, any voter who is deciding is indifferent (see Remark 1). In the pivot-aversion game, such a voter incurs an additional cost from ending the game. Hence, the indifference is broken in favor of continuing the voting, and the voter becomes "rationally indecisive" whenever she has an option of not ending the game. This is why the game lasts until it reaches the last voter.

Although pivot aversion seems natural, one may argue that it is an ad hoc assumption to obtain the long voting result. In Appendix A.6.1, we consider a version of incentives in which voters want to end the game and show a long voting result.

6 Discussions

This section provides several discussions. Section 6.1 introduces asymmetries to our base model and examines how our "anything goes" result and the welfare invariance are affected. In Section 6.2, we compare our consistent equilibria with an equilibrium in the static version of the voting game. Section 6.3 discusses a model with abstention, and Section 6.4 considers a restriction to strategies that depend only on the vote count to date.

6.1 General Asymmetric Voting

In this section, we explore how the result in Theorem 1 is affected by an introduction of asymmetry in the environment. For this purpose, we modify the model as follows:

1. The prior probability on α is $p \in (0, 1)$.

2. There are N voters. K votes are necessary to achieve outcome A and L votes are necessary to achieve outcome B, where $K + L - 1 = N$.

3. The probability of receiving signal a at state α is $\varepsilon_a \in (0, 1)$ and that of receiving signal b at state β is $\varepsilon_b \in (0, 1)$.

We call this model a *general voting model with* $(p, K, L, \varepsilon_a, \varepsilon_b)$. Our model in Section 2 is a general voting model with $(\frac{1}{2}, n+1, n+1, \varepsilon, \varepsilon)$.

We extend the definition of sequences to this general environment by letting the set of *sequences* be

$$\mathcal{S} := \left(\bigcup_{k=K}^{N} \{(X_1, \ldots, X_k) \in \{A, B\}^k \| \{l \in \{1, 2, \ldots, k\} | X_l = A\}| = K, X_k = A\} \right) \cup$$

$$\left(\bigcup_{k=L}^{N} \{(X_1, \ldots, X_k) \in \{A, B\}^k \| \{l \in \{1, 2, \ldots, k\} | X_l = B\}| = L, X_k = B\} \right).$$

As before, for any $S = (X_1, \ldots, X_k) \in \mathcal{S}$, call X_k the *outcome* of S.

Proposition 3. *Consider the general voting model with* $(p, K, L, \varepsilon_a, \varepsilon_b)$ *and fix a sequence* $S \in \mathcal{S}$*. Then, the following holds.*

1. *If* $\frac{1-p}{p} < \frac{(1-\varepsilon_a)^{L-1}}{(1-\varepsilon_b)^{K-1}}$*, then there is a pure consistent strategy equilibrium* σ *such that* S *is the default sequence of* σ *if and only if the outcome of* S *is* A*.*

2. *If* $\frac{1-p}{p} = \frac{(1-\varepsilon_a)^{L-1}}{(1-\varepsilon_b)^{K-1}}$*, then there is a pure consistent strategy equilibrium* σ *such that* S *is the default sequence of* σ*.*

3. *If* $\frac{1-p}{p} > \frac{(1-\varepsilon_a)^{L-1}}{(1-\varepsilon_b)^{K-1}}$*, then there is a pure consistent strategy equilibrium* σ *such that* S *is the default sequence of* σ *if and only if the outcome of* S *is* B*.*

That is, introducing asymmetry breaks the "anything goes" result in a particular direction. The following corollary provides some special cases that are helpful in understanding this general result.

Corollary 1. *Consider the general voting model with* $(p, K, L, \varepsilon_a, \varepsilon_b)$ *and suppose that one of the following holds:*

1. $p < \frac{1}{2}$*,* $K = L$*, and* $\varepsilon_a = \varepsilon_b$*.*

2. $p = \frac{1}{2}$*,* $K < L$*, and* $\varepsilon_a = \varepsilon_b$*.*

3. $p = \frac{1}{2}$*,* $K = L$*, and* $\varepsilon_a > \varepsilon_b$*.*

Then, there is a pure consistent strategy equilibrium σ *such that* S *is the default sequence of* σ *if and only if the outcome of* S *is* B*.*

On the one hand, the "if" direction of the proof of the proposition (that is, to show that any sequence with outcome B can be a default sequence) closely follows that of Theorem 1.

On the other hand, the "only if" direction of the proof of the proposition (that is, to show that no sequence with outcome A can be a default sequence) closely follows that of Theorem 5 because the proof is about breaking the indifference due to the balance between the posterior effect and the delegation effect, and the voter being "rationally indecisive". To understand this point, consider a pure consistent strategy profile in which the default sequence is the K times of A. When $p = \frac{1}{2}$, $K = L(= n+1)$ and $\varepsilon_a = \varepsilon_b$, voter $n+1$ was indifferent between playing A and B due to the balance of the two effects.

Now, consider the three settings in Corollary 1 and suppose for simplicity that the first $K-1$ voters have voted for A and it is voter K's turn. This voter would be indifferent between voting for A and B in the perfectly symmetric model. In part 1 ($p < \frac{1}{2}$), the posterior effect is undermined, and hence there is a strict incentive to vote for B. In part 2 ($K < L$), the posterior effect is given by the $K-1$ times of A, while the delegation effect comes from $L-1$ times of B. Hence, the delegation effect is stronger, and voter K has a strict incentive to vote for B. In part 3 ($\varepsilon_a > \varepsilon_b$), since signal a is more likely than signal b, after $K-1(= L-1)$ times of straight As, the posterior effect is not large (as B has been relatively unlikely to be chosen under state β), and the delegation effect is strong (as A will relatively likely be chosen under state α).

To summarize, if the asymmetry is introduced in a way that "favors" a particular outcome in light of the balance between the posterior effect and the delegation effect, then no sequence that induces the unfavored outcome can be a default sequence of any pure consistent equilibrium, while any sequence that induces the favored outcome is a default sequence of some pure consistent equilibrium. In this sense, the result shows a limitation as well as the robustness of the "anything goes" result.

The welfare equivalence result (Theorem 2) can be extended as follows:

Corollary 2. *Consider the general voting model with $(p, K, L, \varepsilon_a, \varepsilon_b)$. The voters' payoff from any pure consistent strategy equilibrium is*

$$\begin{cases} 1-(1-p)(1-\varepsilon_b)^K & \text{if} \quad \frac{1-p}{p} < \frac{(1-\varepsilon_a)^{L-1}}{(1-\varepsilon_b)^{K-1}} \\ 1-p(1-\varepsilon_a)^L & \text{if} \quad \frac{1-p}{p} > \frac{(1-\varepsilon_a)^{L-1}}{(1-\varepsilon_b)^{K-1}} \end{cases}.$$

Examining the case with $p = \frac{1}{2}$ and $\varepsilon_a = \varepsilon_b$ provides a further interesting insight.

Corollary 3. *Consider the general voting model with N voters such that $p = \frac{1}{2}$ and $\varepsilon_a = \varepsilon_b$. The voters' payoff under pure consistent strategy equilibria is strictly increasing in*

$\max\{K, L\}$. *Therefore, it is maximized when* $K = N$ *or* $L = N$*, and minimized when* $|K - L| \leq 1$.

This corollary shows that, in terms of the equilibrium payoffs, the majority voting rule performs the worst while the unanimous rule is the best.

6.2 Comparison with Static Voting

Call the consistent strategy in which the voter assigns probability $\frac{1}{2}$ to each action regardless of the history of actions the *random voting strategy*.

Proposition 4. *For any* n*, the strategy profile in which every voter uses the random voting strategy is an equilibrium.*

The intuition is simple. Given the random strategy profile, the only case in which a given voter becomes pivotal is when n other voters vote for A and the remaining n voters vote for B. Under such an event, the posterior of the voter, if she does not receive a signal, assigns probability $\frac{1}{2}$ to each state. Hence, the voter is indifferent between A and B, and thus randomizing is indeed optimal. It is straightforward that the voter receiving a signal finds it optimal to vote for the option that is consistent with her signal.

Now, we compare the welfare from pure consistent equilibria with that of the random voting strategy. Given ε, let the payoff from a pure consistent strategy profile be $C(\varepsilon)$ and the payoff from the random voting strategy profile in the static model be $RV(\varepsilon)$.

Proposition 5. *For any* $n \geq 1$*, there exists* $\bar{\varepsilon} > 0$ *such that for all* $\varepsilon \in (0, \bar{\varepsilon})$*,* $RV(\varepsilon) < C(\varepsilon)$*.*

That is, the expected payoff from a pure consistent equilibrium is greater than that of the random strategy profile. Although the proof involves some tedious calculus, the intuition is simple. For the sake of argument, imagine that n is large. When $\varepsilon > 0$ is very small, one voter receiving a signal is much more likely than more voters receiving a signal. Hence, consider the case in which exactly one voter receives a signal, say a. In such a case, the event in which this voter would become pivotal under the random voting strategy is small because n is large, and we would need exactly n voters voting for A and exactly n voters voting for B. However, under a pure consistent strategy profile, if this voter happens to be on the default sequence that specifies her to play B, then she can overturn the outcome because her action, voting for A, affects the subsequent voters' actions. This substantially improves efficiency. In summary, the fact that the voting is sequential makes the probability of a voter becoming pivotal high, and this is the reason why the expected payoff from a pure consistent equilibrium is greater than that of random strategy profile.

6.3 Abstention Model

The main section considered a model without abstention. Such a model would approximate the situation as in *Squid Game* as well as roll call vote in United Nations, US Senate and Japanese Diet. In other applications, however, abstention is allowed. This section discusses a model with abstention.

Formally, there are N voters. The action space is now $\{A, B, \Phi\}$, where Φ stands for abstention. The voting ends with outcome A at voter k if, after k's vote, the number of votes on A minus the number of votes on B is strictly greater than $N - k$. Similarly, the voting ends with outcome B at voter k if, after her vote, the number of votes on B minus the number of votes on A is strictly greater than $N - k$. If the number of votes for A and that for B are equal after N voters vote, then outcomes A and B realize with equal probability.

We say that voter k's strategy is *consistent in the abstention model* if k chooses the "consistent" action if she receives a signal, i.e., $\sigma_k(\cdot, a)(A) = 1$ and $\sigma_k(\cdot, b)(B) = 1$.

In Appendix A.10, we show that there is an equilibrium in which all voters use a strategy that is consistent in the abstention model under any history, and abstain on the path of play if no signal is received. The Appendix also provides an example of an equilibrium in which there is a voter who does not use a strategy that is consistent in the abstention model: in particular, she votes for Φ if her signal is a and A if her signal is ϕ (and votes for B if the signal is b). In both examples, the ex ante payoff turns out to be the highest payoff achievable by any social choice function.

Finally, we note that, in the standard voting models in the literature, introduction of abstention would strictly improve the welfare. In our model, however, the highest payoff achievable by any social choice function can be obtained *without* the introduction of abstention as we have seen in Section 4.

6.4 Vote Count Model

One may wonder if our "anything goes" result hinges "too much" on the observation of the particular default sequence. However, the same result can be proved even if we allow voters to observe only the vote counts.

To state this formally, change the base model so that voter k's strategy is a mapping from the numbers of votes for A and k's signals to distributions over her actions: $\sigma_k : \{0, \ldots, k-1\} \times \{a, b, \emptyset\} \to \Delta(\{A, B\})$. Given this definition of strategy, we can define sequential equilibria. Call this model the *vote count model*.

Theorem 7. *In the vote count model, for any $n \geq 0$, $\varepsilon \in (0, 1]$ and any sequence $S \in \mathcal{S}$, there is a pure consistent strategy equilibrium σ in the model with (n, ε) such that S is the* *default sequence of σ.*

As for Theorem 1, the proof is constructive. The equilibrium construction, however, is different because the voters cannot condition their votes on who deviated first from the default sequence when the observed vote counts contradict the default sequence. For this reason, we consider an alternative construction. There is still a default sequence. If a voter with no signal sees the vote counts that are implied by the default sequence, then she votes according to the default sequence. Otherwise, the voter votes for the option that has received more votes than what the default sequence implies.[12]

7 Conclusion

This paper considered a common-value sequential voting in which voters are differentiated in the precision of information about the favorability of the outcomes. We identified two effects, the posterior effect and the delegation effect, and the balance of these two effects resulted in the "anything goes" result. The welfare is invariant as long as we consider the class of pure consistent equilibria where consistency means that the voter's vote reflects her information, if any. There are inconsistent equilibria as well, and the welfare from the inconsistent equilibria varies quite a bit. Two modifications of the model are discussed, where the modified models imply long voting because the pivotal (deciding) voter becomes "rationally indecisive" due to the modification.

We end the paper by re-stressing that our model is stylized. Although we discussed some extensions (such as asymmetries, abstention or the vote count model), we understand there are various features of real sequential voting that we do not capture in this paper. We view such discrepancies as a wonderful possibility for future work. In such future work, the two effects we found—the posterior effect and the delegation effect—would be essential building blocks in understanding the incentives of the voters, as they were in this paper.

References

Ali, S. N. and N. Kartik (2012): "Herding with collective preferences," *Economic Theory*, 51, 601–626.

Banerjee, A. V. (1992): "A simple model of herd behavior," *The Quarterly Journal of Economics*, 107, 797–817.

[12]We note that this strategy profile is an equilibrium in the model in which the entire history of past votes is observable as well.

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

Dekel, E. and M. Piccione (2000): "Sequential Voting Procedures in Symmetric Binary Elections," *Journal of Political Economy*, 108, 34–55.

——— (2014): "The strategic dis/advantage of voting early," *American Economic Journal: Microeconomics*, 6, 162–79.

Feddersen, T. J. and W. Pesendorfer (1996): "The Swing Voter's Curse," *The American Economic Review*, 86, 408–424.

Kreps, D. M. and R. Wilson (1982): "Sequential equilibria," *Econometrica: Journal of the Econometric Society*, 863–894.

A APPENDIX

A.1 Proof of Theorem 1 and Remark 1

Proof. Fix $k \in \{n+1, n+2, \ldots, 2n+1\}$ and take a sequence of actions $S^* = (X_1, \ldots, X_k) \in \{A, B\}^k$ such that X_k appears exactly $n+1$ times. Consider the consistent strategy profile σ satisfying the following conditions:

1. $\sigma_i(X_1, \ldots, X_{i-1}, \emptyset)(X_i) = 1$ for each $i = 1, \ldots, k$.
2. For any $i = 1, \ldots, 2n+1$ and $i' = 1, \ldots, \min\{k, i-1\}$, take any sequence $S = (X_1, \ldots, X_{i'-1}, X'_{i'}, \ldots, X'_{i-1}) \in \{A, B\}^{i-1}$ such that $X'_{i'} \neq X_{i'}$. That is, i' is the first voter who does not follow the sequence $(X_1, \ldots, X_k)$. Then, all the subsequent voters choose the same action as $X'_{i'}$, i.e.,

$$\sigma_i(S, \emptyset)(X'_i) = 1.$$

Now we check incentives. First, voters who have received a signal take a best response given condition 2 above.

Second, under any histories described in condition 2, by letting i have a belief that the state is α if $X'_{i'} = A$ and β if $X'_{i'} = B$, it is straightforward to see that i is taking a best response.

Third, under any histories described in condition 1, without loss of generality, let $X_i = A$. Suppose that there has been m times of A's before i. Voter i's posterior on α is then

$$P(\alpha) = \frac{(1-\varepsilon)^{i-1-m}}{(1-\varepsilon)^{i-1-m} + (1-\varepsilon)^m},$$

and her posterior on β is

$$P(\beta) = 1 - P(\alpha) = \frac{(1-\varepsilon)^m}{(1-\varepsilon)^{i-1-m} + (1-\varepsilon)^m}.$$

Suppose that i plays B. If the state is β, then outcome B realizes with probability 1 given that the voters follow σ. If the state is α, the outcome becomes A if and only if at least $n+1-m$ subsequent voters receive signal a. Thus, her payoff is

$$1 - P(\alpha) + P(\alpha) \cdot Y,$$

where $Y := \text{Prob}(n+1-m \text{ or more signals out of } 2n+1-i)$.

Now suppose that she plays A. We consider the following two (exhaustive) cases.

1. Suppose that $X_k = A$.

   Then i's payoff when she plays A is

$$P(\alpha) + (1 - P(\alpha))(1 - (1 - \varepsilon)^{n-m}).$$

   Thus, the payoff from playing A is no less than the payoff from playing B if and only if:

$$P(\alpha) + (1 - P(\alpha))(1 - (1 - \varepsilon)^{n-m}) \geq 1 - P(\alpha) + P(\alpha)Y \iff$$
$$P(\alpha)[1 + (1 - \varepsilon)^{n-m} - Y] \geq (1 - \varepsilon)^{n-m} \iff P(\alpha) \geq \frac{(1 - \varepsilon)^{n-m}}{(1 - \varepsilon)^{n-m} + 1 - Y}.$$

   Now, notice that

$$\begin{aligned} P(\alpha) &= \frac{(1 - \varepsilon)^{i-1-m}}{(1 - \varepsilon)^{i-1-m} + (1 - \varepsilon)^{m}} \\ &= \frac{(1 - \varepsilon)^{n-m}}{(1 - \varepsilon)^{n-m} + (1 - \varepsilon)^{n+1+m-i}}. \end{aligned}$$

   Hence, the payoff from playing A is no less than the payoff from playing B if and only if:

$$(1 - \varepsilon)^{n+1+m-i} \leq 1 - Y,$$

   or

$$(1 - \varepsilon)^{n+1+m-i} \leq \text{Prob}(n - m \text{ or less signals out of } 2n + 1 - i).$$

   Now, let $Z_1 := (n + 1) + (m - i)$ and $Z_2 := (n + 1) + (n - i)$. Notice that $Z_1 \leq Z_2$ because $m \geq n$. The payoff from playing A is no less than the payoff from playing B if and only if:

$$(1 - \varepsilon)^{Z_1} \leq \text{Prob}(Z_2 - Z_1 \text{ or less signals out of } Z_2). \tag{1}$$

   Note that

$$\begin{aligned} &\text{Prob}(Z_2 - Z_1 \text{ or less signals out of } Z_2) \\ &\geq \text{Prob}(\text{The last } Z_1 \text{ voters do not receive a signal}) \qquad (2) \\ &= (1 - \varepsilon)^{Z_1}. \end{aligned}$$

   Thus, eq. (1) indeed holds, and therefore, the payoff from playing A is no less than the payoff from playing B.

2. Suppose that $X_k = B$.

Suppose the state is α, which happens with probability $P(\alpha)$. Since there have been $m+1$ times of A right after i plays A and there are $n+1$ times of B in the given sequence, there will be $(n+1)-(i-(m+1))$ times of B at which the voter, if she receives signal a, can change the outcome to A. The probability of at least one voter out of $(n+1)-(i-(m+1))$ voters receiving signal a is $1-(1-\varepsilon)^{(n+1)-(i-(m+1))}$.

Suppose the state is β, which happens with probability $1-P(\alpha)$. In this case, i's playing A ensures that there is probability 1 that B will be chosen as the outcome.

Thus, overall, the payoff from i's voting for A is

$$1-P(\alpha)+P(\alpha)(1-(1-\varepsilon)^{(n+1)-(i-(m+1))}).$$

Thus, the payoff from playing A is no less than the payoff from B if and only if:

$$1-P(\alpha)+P(\alpha)(1-(1-\varepsilon)^{(n+1)-(i-(m+1))}) \geq 1-P(\alpha)+P(\alpha)Y \iff$$

$$(1-\varepsilon)^{n+2+m-i} \leq 1-Y,$$

or

$$(1-\varepsilon)^{n+2+m-i} \leq \text{Prob}(n-m \text{ or less “}a\text{” signals out of } 2n+1-i).$$

Now, let $Z_1 := n+2+m-i$ and $Z_2 := 2n+2-i$. Notice that $Z_1 \leq Z_2$ because $Z_2 - Z_1 = 2n+2-i-(n+2+m-i) = n-m \geq 0$. The payoff from playing A is no less than the payoff from B if and only if:

$$(1-\varepsilon)^{Z_1} \leq \text{Prob}(Z_2 - Z_1 \text{ or less “}a\text{” signals out of } Z_2 - 1). \tag{3}$$

Note that

$$\begin{aligned}
&\text{Prob}(Z_2 - Z_1 \text{ or less “}a\text{” signals out of } Z_2 - 1)\\
&\geq \text{Prob}(Z_2 - Z_1 \text{ or less “}a\text{” signals out of } Z_2)\\
&\geq \text{Prob}(\text{The last } Z_1 \text{ voters do not receive an “}a\text{” signal})\\
&= (1-\varepsilon)^{Z_1}.
\end{aligned}$$

Thus, eq. (3) indeed holds, and therefore, the payoff from playing A is no less than the payoff from B.

Overall, playing A is a best response for voter i. □

Proof of Remark 1. The voters on the default sequence is indifferent if and only if eq. (2) holds with equality in Case 1 (i.e., $X_k = A$). This happens when either $Z_1 = Z_2$ or $Z_1 = 0$. Note that $Z_1 = Z_2$ is equivalent to $m = n$, and $Z_1 = 0$ is equivalent to $i - 1 = n + m$, so i has seen n times of B. Since eq. (2) is in the case when $X_i = X_k = A$, these are exactly the cases when voter i is deciding. □

A.2 Proof of Theorem 3

Proof. For any number of voters $2n + 1$, the following strategy profile $\bar{\sigma}$ is an equilibrium and achieves the ex ante payoff $\bar{R}$.

- For any $k = 1, \dots n$, voter k chooses B if she receives a; otherwise she plays A.
- For any $k = n + 1 \dots 2n + 1$, voter k uses a consistent strategy. If he does not receive a signal and observes that voters $1 \dots n$ play A and voters $n + 1 \dots k - 1$ play B, then he plays B. Otherwise, k plays A.

Note that this is a pure strategy profile and is not consistent. Its default sequence is $(A, \dots, A, B, \dots, B)$, where A continues n times and then B continues $n + 1$ times. To see that each voter chooses a best response, observe first that it is immediate that any voter who receives a signal is taking a best response. Also, observe that once there is a deviation from the default sequence, then voters take a best response given a belief that the first deviator from the default sequence has received a signal and followed the equilibrium strategy. So suppose that the voters 1 through $k - 1$ have followed the default sequence, and suppose that voter k receives no signal. The posterior on α is $P(\alpha) = \frac{(1-\varepsilon)^{k-1}}{(1-\varepsilon)^{k-1}+1}$. If k follows the specified strategy, then her payoff is

$$P(\alpha)(1 - (1 - \varepsilon)^{2n+1-k}) + (1 - P(\alpha)) \cdot 1.$$

If k deviates, then her payoff is $P(\alpha)$. Since $P(\alpha) < \frac{1}{2} < 1 - P(\alpha)$, the former is larger than the latter. Hence, k is taking a best response.
Finally, since the payoff is 0 under $\bar{\sigma}$ if and only if the state is α and no one receives a signal, the expected payoff is $1 - \frac{(1-\varepsilon)^{2n+1}}{2}$. □

A.3 Proof of Proposition 2

Proof. Fix a pure strategy equilibrium σ. Take the default sequence of σ. Suppose without loss of generality that the last voter on this default sequence, whom we denote voter k^*, plays A.

First, we show the following lemma.

Lemma 1. *Suppose that the votes by voters $1, \ldots, k^* - 1$ have followed the default sequence of σ. Then, voter k^* votes for B if and only if she receives signal b.*

Proof of Lemma 1. Consider the history in which voters $1, \ldots, k^* - 1$ have followed the default sequence of σ. Given the assumption that k^* plays A given no signal, it suffices to check the cases when she receives a signal.

1. First, we show that k^* plays B if she receives signal b. We show this using induction on k. To do this, fix $k' \in \{k^*, \ldots, 2n+1\}$. Suppose, as an induction hypothesis, that for every $k > k'$, voter k plays B if he receives signal b.[13] Suppose that voter k' receives signal b. Given this signal, her posterior probability on β is 1. If k' votes for A, then her expected payoff is 0. If k' votes for B, then there is a strictly positive probability that all the subsequent voters receive signal b and, in that case, the game ends with outcome B by the induction hypothesis. Hence, her expected payoff is strictly positive. Thus, playing B is a unique best response for k', and hence, voter k' plays B if she receives signal b. This shows that voter k^* plays B if she receives signal b.

2. Next, we show that k^* plays A if she receives signal a. To see this, suppose to the contrary that k^* plays B if she receives signal a. The outcome will be A with probability 1 after k^* plays B because otherwise, playing B would give k^* a payoff strictly less than 1, while she would get the payoff of 1 if she played A, making her choice B suboptimal. Now, suppose that k^* did not receive a signal. Let $P(\alpha)$ be the posterior of k^* at such an information set. Then, the expected payoff of voter k^* is $P(\alpha)$ if she plays A. If she instead plays B, then the outcome will be A with probability 1 if the state is α as we have concluded. The argument in item 1 above implies that if the state is β and all the subsequent voters receive signal b, which happens with probability $\varepsilon^{2n+1-k^*} > 0$, then the outcome is B. Hence, the expected payoff of voter k^* is at least $P(\alpha) + (1 - P(\alpha))\varepsilon^{2n+1-k'}$ if she plays B, and this is strictly greater than $P(\alpha)$ because $1 - P(\alpha) > 0$ and $\varepsilon > 0$. Hence, playing A is suboptimal, which contradicts the assumption that k^* votes for A.

□

Given Lemma 1, once k^* plays B, the subsequent voters assign posterior probability 1 to state β, so the outcome will be B with probability 1.

Lemma 2. *Under σ, if voter 1 receives signal a, then she expects that the outcome will be A with probability 1.*

[13] Note that this hypothesis is vacuously true when $k' = 2n + 1$.

Proof of Lemma 2. We use induction. Fix $k \leq k^*$. Suppose as an induction hypothesis that for every $k' \in \{k+1, \dots, k^*\}$, if the actions by voters $1, \dots, k'-1$ have followed the default sequence and k' receives signal a, then the outcome will be A with probability 1. Then, suppose that the actions by voters $1, \dots, k-1$ have followed the default sequence and k receives signal a. If k plays the action specified in the default sequence, then either (i) no subsequent voters receive a signal, or (ii) there is at least one subsequent voter who receives signal a. In case (i), by the definition of the default sequence, the outcome will be A with probability 1. In case (ii), the first voter who receives signal a expects that the outcome will be A, which follows from the induction hypothesis. Hence, the outcome will be A in either case. Therefore, k's expected payoff is 1 if k plays the action specified in the default sequence. Since k under σ must be doing at least as good as playing any action, this implies that, under σ, k expects the payoff of 1, and hence she expects that the outcome will be A with probability 1. This completes the induction argument. Therefore, we have shown that, if voter 1 receives signal a, then she expects that the outcome will be A with probability 1. □

Lemma 3. *Under σ, if any voter receives signal b when the votes so far followed the default sequence, then she expects that the outcome will be B with probability at least ε.*

Proof of Lemma 3. We use induction. Fix $k \leq k^* - 1$. Suppose as an induction hypothesis that for every $k' \in \{k+1, \dots, k^*-1\}$, if the actions by voters $1, \dots, k'-1$ have followed the default sequence and k' receives signal b, then the outcome will be B with probability at least ε. Then, suppose that the actions by voters $1, \dots, k-1$ have followed the default sequence and k receives signal b. If k plays the action specified in the default sequence, then either (i) no subsequent voters in $\{k+1, \dots, k^*-1\}$ receive a signal, or (ii) there is at least one subsequent voter in $\{k+1, \dots, k^*-1\}$ who receives signal b. In case (i), by Lemma 1, voter k^* will play B if she receives signal b, and thus there is probability ε that the outcome will be B as we concluded after Lemma 1. In case (ii), the first voter who receives signal b expects that the outcome will be B with a probability of at least ε, which follows from the induction hypothesis. Hence, the outcome will be B with at least ε probability in either case. Therefore, k's expected payoff under σ must be at least ε if she plays the action specified in the default sequence. Since k under σ must be doing at least as good as playing any action, this implies that, under σ, k expects a payoff of at least ε, and hence she expects that the outcome will be B with probability at least ε. This completes the induction argument. Therefore, we have shown that, if any voter receives signal b, then she expects that the outcome will be B with probability at least ε. □

Finally, consider voter 1. If the state is α, then either (i) no voters in $\{1, \dots, k^*\}$ receive a signal, or (ii) there is at least one voter in $\{1, \ldots, k^*\}$ who receives signal a. In case (i), the outcome will be A with probability 1 by the choice of the default sequence. In case (ii), we have shown that the outcome will be A with probability 1. Hence, voter 1's expected payoff is 1 conditional on this event.

If the state is β, then either (i) no voters in $\{1, \ldots, k^* - 1\}$ receive a signal, or (ii) there is at least one voter in $\{1, \ldots, k^* - 1\}$ who receives signal b. In case (i), voter k^* will have seen the history that has followed the default sequence. If she receives no signal, which happens with probability $1 - \varepsilon$, then she plays A by assumption, and thus the outcome will be A with probability 1. If she receives signal b, which happens with probability ε, then the outcome will be B with probability 1 as we concluded after Lemma 1. Overall, voter 1 expects the probability that the outcome will be B to be ε conditional on (i). In case (ii), take the first voter who has received signal b. We have shown that this voter assigns a probability of at least ε that the outcome will be B. Thus, conditional on the event that (ii) realizes, voter 1 expects that the probability that the outcome will be B is at least ε. Therefore, if the state is β, the outcome will be B with at least ε probability.

Overall, since the prior on the states α and β are both $\frac{1}{2}$, voter 1's expected payoff is at least

$$\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot \varepsilon = \frac{1 + \varepsilon}{2}.$$

□

A.4 Proof of Theorem 4

Proof. The next example shows that there is an inconsistent equilibrium that achieves a strictly lower expected payoff than that of consistent equilibria.

Consider the following strategy profile, which we denote by $\underline{\sigma}$.

- Voter $k = 1, \ldots, n$: If all the actions observed so far are A, then play A irrespective of the signal. Otherwise, play a consistent strategy in which A is played if no signal is received.

- Voter $k = n + 1, \ldots, 2n$: If all the actions observed so far are n times of A followed by $k - 1 - n$ times of B, then play B irrespective of the signal. Otherwise, play a consistent strategy in which A is played if no signal is received.

- Voters $k = 2n + 1$: Play a consistent strategy in which A is played if no signal is received.

Note that this strategy profile has the ex ante payoff of $\frac{1}{2} + \frac{1}{2}\varepsilon = \frac{1+\varepsilon}{2}$. We show that, if $n \geq 2$, $\underline{\sigma}$ constitutes an equilibrium with a belief that, after any deviation, any voter who has not received a signal assigns probability 1 to α.

Consider voter k. Fix any history of actions by the previous voters. First, suppose that k receives signal a. Then, since the outcome will be A if k follows $\underline{\sigma}_k$, playing $\underline{\sigma}_k$ induces the expected payoff of 1, which is the highest possible payoff in this game. Hence, $\underline{\sigma}_k$ is a best response.

Second, suppose that k receives signal b.

- Suppose that the voters so far have followed $\underline{\sigma}$.
  - If $1 \leq k \leq n$, then if she follows $\hat{\sigma}_k$ then her payoff is ε. If instead she plays B then at least n voters from the set of subsequent voters have to receive signal b in order for k to expect the payoff of 1, and otherwise she receives the payoff of 0. Hence, her payoff is $O(\varepsilon^n)$. Since $n \geq 2$, this implies that there exists $\bar{\varepsilon} > 0$ such that for all $\varepsilon < \bar{\varepsilon}$, following $\underline{\sigma}_k$ is a best response.
  - If $n + 1 \leq k \leq 2n$, then playing A ends the game with outcome A, so it induces the expected payoff of 0. Hence, playing B is a best response.
  - $If k = 2n + 1$, then playing B induces the payoff of 1 with probability 1, so it is a best response to play B.
- Suppose that there is a voter who has deviated from $\underline{\sigma}$. Then, the play by the subsequent voters will not be affected by the action taken by voter k. Hence, it is a best response for voter k to provide additional vote for B. Therefore, playing B is a best response.

Third, suppose that voter k did not receive a signal.

- Suppose that the voters so far have followed $\underline{\sigma}$. Then, the posterior belief on α is $\frac{1}{2}$.
  - Suppose $k \leq n$. If k follows $\underline{\sigma}_k$, then her payoff is $\frac{1}{2} + \frac{1}{2}\varepsilon$. If she does not follow $\underline{\sigma}_k$, then her payoff is $\frac{1}{2} + \frac{1}{2}O(\varepsilon^n)$. Since $n \geq 2$, following $\underline{\sigma}_k$ is a best response.
  - Suppose $n + 1 \leq k \leq 2n$. If k follows $\hat{\sigma}_k$, then her payoff is $\frac{1}{2} + \frac{1}{2}\varepsilon$. If she does not follow $\underline{\sigma}_k$, then her payoff is $\frac{1}{2}$. Hence, following $\underline{\sigma}_k$ is a best response.
  - Suppose $k = 2n + 1$. Then, k is indifferent between the two actions. So following $\hat{\sigma}_k$ is a best response.
- Suppose that there is a voter who has deviated from $\underline{\sigma}$ before k's move. Then, the play by the subsequent voters will not be affected by the action taken by voter k. Hence it is a best response for k to provide additional vote for A because k's belie assigns probability 1 to α.

□

A.5 Proof of Theorem 5

Proof. Fix a pseudo-pure consistent strategy profile. Suppose that if no voter gets a signal, on the path of play, the game ends before reaching voter $2n+1$. Let k be the voter who ends the game in this case. Suppose without loss of generality that k plays A.

Lemma 4. *For voter k, playing A is suboptimal.*

Proof. Voter k has seen n times of action A and $k-n-1$ times of action B. The posterior on α is

$$P(\alpha) = \frac{(1-\varepsilon)^{k-n-1}}{(1-\varepsilon)^{k-n-1} + (1-\varepsilon)^n} = \frac{1}{1+(1-\varepsilon)^{2n+1-k}}. \tag{4}$$

If voter k plays A, then the game ends and her expected payoff becomes $P(\alpha)$.

Let k' ($\geq k+1$) be the voter who ends the game if k plays B.

Case 1: $k' = 2n+1$ If voter k plays B, then all the subsequent voters before $2n+1$ also choose B (when they receive no signal). Let q be the probability that voter $2n+1$ assigns to A (when he does not receive a signal). By the craziness assumption, $q \geq \xi > 0$. Then, k's payoff if she plays B is

$$P(\alpha)\left[1-(1-q)(1-\varepsilon)^{2n+1-k}\right] + (1-P(\alpha))\left[\varepsilon + (1-\varepsilon)(1-q)\right].$$

Let this be $f(q)$. Note that

$$f(0) = P(\alpha)[1-(1-\varepsilon)^{2n+1-k}] + (1-P(\alpha)) = 1 - P(\alpha)(1-\varepsilon)^{2n+1-k} = P(\alpha).$$

Then, for $q > 0$,

$$\begin{aligned}
\frac{f(q)-P(\alpha)}{q} &= \frac{f(q)-f(0)}{q} \\
&= P(\alpha)(1-\varepsilon)^{2n+1-k} - (1-P(\alpha))(1-\varepsilon) \\
&= \frac{(1-\varepsilon)^{(2n+1-k)}}{1+(1-\varepsilon)^{2n+1-k}} - \frac{(1-\varepsilon)^{(2n+1-k)}}{1+(1-\varepsilon)^{2n+1-k}}(1-\varepsilon) \\
&= \frac{\varepsilon(1-\varepsilon)^{(2n+1-k)}}{1+(1-\varepsilon)^{2n+1-k}} > 0.
\end{aligned}$$

So, $f(q) > P(\alpha)$ for any $q > 0$. This means that k would be better off playing B than A.

Case 2: $k' \leq 2n$ If voter k plays B, then all the subsequent voters before k' also choose B and voter k' chooses A (when they receive no signal). In this case, if the state is α, then A realizes with probability 1 since voter k' always chooses A. We show that, if the state is β, then B realizes with strictly positive probability. Note that, if the subsequent voters from k'

to $2n+1$ all receive signal b, which occurs with probability $\varepsilon^{2n+2-k'} > 0$, then they all play B by consistency. Consequently, B realizes. This implies that B realizes with a probability of at least $\varepsilon^{2n+2-k'}$, irrespective of the voters' strategies. Hence, the voter's payoff is at least

$$P(\alpha) \cdot 1 + (1 - P(\alpha)) \cdot \varepsilon^{2n+2-k'},$$

which is strictly greater than $P(\alpha)$. Therefore, k would be better off playing B than A. □

Lemma 4 implies that the pseudo-pure consistent strategy profile that we fixed is not an equilibrium. This completes the proof. □

A.6 Proof of Theorem 6

Proof. Fix a pure consistent equilibrium in a pivot-aversion game and suppose, towards contradiction, that its default sequence has length $l \leq 2n$. Without loss of generality, let A be the action that l takes when she has observed actions specified in the default sequence and does not receive a signal. Then, her posterior belief about the state α is

$$P(\alpha) = \frac{(1-\varepsilon)^{l-1-n}}{(1-\varepsilon)^n + (1-\varepsilon)^{l-1-n}}.$$

By choosing A, her payoff is $P(\alpha) - c$. If chooses B, then all subsequent voters have a belief that assigns probability 1 to state β unless they receive signal a. Since voting for A ends the voting with outcome A, this implies that all subsequent voters vote for A if and only if they receive signal a. This implies that l's expected payoff from choosing B is

$$P(\alpha)(1 - (1-\varepsilon)^{2n+1-l}) + (1 - P(\alpha))$$

(note that we do not subtract c because there have been $l-1-n$ times of B, and $l-1-n \leq 2n-1-n < n$), which is equal to

$$\frac{(1-\varepsilon)^{l-1-n}(1-(1-\varepsilon)^{2n+1-l}) + (1-\varepsilon)^n}{(1-\varepsilon)^n + (1-\varepsilon)^{l-1-n}} = \frac{(1-\varepsilon)^{l-1-n}}{(1-\varepsilon)^n + (1-\varepsilon)^{l-1-n}} = P(\alpha).$$

Since $c > 0$, choosing B is strictly better than choosing A, which implies that voter l is not taking a best response. Contradiction. □

A.6.1 Further Discussion for Section 5.2

In Section 5.2, we considered incentives not to end the game. Here, we provide a model with incentives to end the game, and show a long voting result. We note that the objective of this section is not to declare that long voting is a ubiquitous result. Rather, we would like to point out that the long voting result is not an artifact that only results from the incentive not to end.

We consider the following model. Suppose that, when voter k does not receive a signal, independently with probability $\eta > 0$, she is a *decisive type* that receives additional payoff $D > 1$ from voting for X when the history has n times of X. Suppose D does not depend on the action (thus, the last voter $2n+1$ would simply ignore this additional D payoff). Call this game the *decisive-type game.*

Theorem 8. *In any decisive-type game, there is no pure consistent strategy equilibrium such that the game ends before reaching voter* $2n$ *when no voter receives a signal.*

Proof. The proof closely follows that of Theorem 5. Fix a pure consistent strategy profile. Suppose that if no voter gets a signal and no voter is a decisive type, on the path of play, k is the voter who ends the game and $k < 2n$. Suppose without loss of generality that k plays A.

Lemma 5. *For voter* k*, playing* A *is suboptimal.*

Proof. Voter k has seen n times of action A and $k-n-1$ times of action B. Let l be the number of voters before k who has seen n times of action A. The posterior on α is

$$P(\alpha) = \frac{(1-\varepsilon)^{k-n-1}}{(1-\varepsilon)^{k-n-1} + (1-\varepsilon)^{n}} = \frac{1}{1+(1-\varepsilon)^{2n+1-k}}.$$

If voter k plays A, then the game ends and her expected payoff becomes $P(\alpha)$.

Let k' be the voter who ends the game if k plays B.

Case 1: $k' = 2n+1$ If voter k plays B, then all the subsequent voters before $2n+1$ also choose B (when they receive no signal and are not a decisive type). Let q be the probability that voter $2n+1$ assigns to A (when he does not receive a signal). Then, k's payoff if she plays B is

$$P(\alpha)\left[1-(1-\eta)^{2n-k}(1-\varepsilon)^{2n+1-k}(1-q)\right] + (1-P(\alpha))\left[(1-(1-\varepsilon)\eta)^{2n-k}(\varepsilon+(1-\varepsilon)(1-q))\right].$$

Let this be $f(q)$. Note that

$$\begin{aligned} f(0) &= P(\alpha)\left[1-(1-\eta)^{2n-k}(1-\varepsilon)^{2n+1-k}\right]+(1-P(\alpha))(1-(1-\varepsilon)\eta)^{2n-k} \\ &> P(\alpha)\left[1-(1-\eta)^{2n-k}(1-\varepsilon)^{2n+1-k}\right]+(1-P(\alpha))(1-\eta)^{2n-k} \\ &= \frac{(1-\varepsilon)^{k-n-1}-(1-\eta)^{2n-k}(1-\varepsilon)^{n}+(1-\varepsilon)^{n}(1-\eta)^{2n-k}}{(1-\varepsilon)^{k-n-1}+(1-\varepsilon)^{n}} \\ &= \frac{(1-\varepsilon)^{k-n-1}}{(1-\varepsilon)^{k-n-1}+(1-\varepsilon)^{n}} = P(\alpha). \end{aligned}$$

Also,

$$f(1) = P(\alpha)+(1-P(\alpha))(1-\eta)^{2n-k}\varepsilon > P(\alpha).$$

So $f(q) > P(\alpha)$ for both $q = 0, 1$. This means that k would be better off playing B than A.
Case 2: $k' \leq 2n$ If voter k plays B, then all the subsequent voters before k' also choose B and voter k' chooses A (when they receive no signal). In this case, if the state is α, then A realizes with probability 1 since voter k' always chooses A. We show that, if the state is β, then B realizes with strictly positive probability. Note that, if no voter after k until $2n$ is a decisive type and the subsequent voters from k' to $2n+1$ all receive signal b, which occurs with probability $(1-\eta)^{2n-k}\varepsilon^{2n+2-k'} > 0$, then they all play B by consistency. Consequently, B realizes. This implies that B realizes with a probability of at least $(1-\eta)^{2n-k}\varepsilon^{2n+2-k'}$, irrespective of the voters' strategies. Hence, the voter's payoff is at least

$$P(\alpha)\cdot 1+(1-P(\alpha))\cdot(1-\eta)^{2n-k}\varepsilon^{2n+2-k'},$$

which is strictly greater than $P(\alpha)$. Therefore, k would be better off playing B than A. □

The lemma implies that the pure consistent strategy profile that we fixed is not an equilibrium. This completes the proof. □

A.7 Proof of Proposition 3

Fix the general voting model with $(p, K, L, \varepsilon_a, \varepsilon_b)$ and a sequence $S \in \mathcal{S}$ that ends with A. To prove the three parts of the proposition, we show the following two claims:

(A) If $\frac{1-p}{p} \leq \frac{(1-\varepsilon_a)^{L-1}}{(1-\varepsilon_b)^{K-1}}$, then there is a pure consistent strategy equilibrium σ such that S is the default sequence of σ.

(B) If $\frac{1-p}{p} > \frac{(1-\varepsilon_a)^{L-1}}{(1-\varepsilon_b)^{K-1}}$, then there is no pure consistent strategy equilibrium σ such that S is the default sequence of σ.

By symmetry, these are enough to show the three parts of the proposition.

Proof for claim (A). Fix $k \in \{\min\{K, L\} + 1, \min\{K, L\} + 2, \dots, N\}$ and take a sequence of actions $S^* = (X_1, \dots, X_k) \in \{A, B\}^k$ such that X_k appears exactly K times if $X_k = A$ and L times if $X_k = B$. Consider the consistent strategy profile σ satisfying the following conditions:

1. $\sigma_i(X_1, \dots, X_{i-1}, \emptyset)(X_i) = 1$ for each $i = 1, \dots, k$.

2. For any $i = 1, \dots, N$ and $i' = 1, \dots, \min\{k, i-1\}$, take any sequence $S = (X_1, \dots, X_{i'-1}, X'_{i'}, \dots, X'_{i-1}) \in \{A, B\}^{i-1}$ such that $X'_{i'} \neq X_{i'}$. That is, i' is the first voter who does not follow the sequence $(X_1, \dots, X_k)$. Then, all the subsequent voters choose the same action as $X'_{i'}$, i.e.,
$$\sigma_i(S, \emptyset)(X'_{i'}) = 1.$$

Now we check incentives. First, voters who have received a signal take a best response given condition 2 above.

Second, under any histories described in condition 2, by letting i have a belief that the state is α if $X'_{i'} = A$ and β if $X'_{i'} = B$, it is straightforward to see that i is taking a best response.

Third, under any histories described in condition 1, let $X_i = A$. Suppose that there have been m times of A's before i. Voter i's posterior on α is then

$$P(\alpha) = \frac{p(1-\varepsilon_a)^{i-1-m}}{p(1-\varepsilon_a)^{i-1-m} + (1-p)(1-\varepsilon_b)^m},$$

and her posterior on β is

$$P(\beta) = 1 - P(\alpha) = \frac{(1-p)(1-\varepsilon_b)^m}{p(1-\varepsilon_a)^{i-1-m} + (1-p)(1-\varepsilon_b)^m}.$$

Suppose that i plays B. If the state is β, then outcome B realizes with probability 1 given that the voters follow σ. If the state is α, the outcome becomes A if and only if at least $K - m$ subsequent voters receive signal a. Thus, her payoff is

$$1 - P(\alpha) + P(\alpha) \cdot Y,$$

where $Y := \text{Prob}(K - m \text{ or more "}a\text{" signals out of } N - i)$.

Now suppose that she plays A.

1. Suppose that $X_k = A$. Then i's payoff when she plays A is

$$P(\alpha) + (1 - P(\alpha))(1 - (1-\varepsilon_b)^{K-1-m}).$$

Thus, the payoff from playing A is no less than the payoff from B if and only if:

$$P(\alpha) + (1 - P(\alpha))(1 - (1 - \varepsilon_b)^{K-1-m}) \geq 1 - P(\alpha) + P(\alpha)Y \iff$$

$$P(\alpha)[1 + (1 - \varepsilon_b)^{K-1-m} - Y] \geq (1 - \varepsilon_b)^{K-1-m} \iff P(\alpha) \geq \frac{(1 - \varepsilon_b)^{K-1-m}}{(1 - \varepsilon_b)^{K-1-m} + 1 - Y}.$$

Now, notice that

$$\begin{aligned} P(\alpha) &= \frac{p(1 - \varepsilon_a)^{i-1-m}}{p(1 - \varepsilon_a)^{i-1-m} + (1 - p)(1 - \varepsilon_b)^m} \\ &= \frac{(1 - \varepsilon_b)^{K-1-m}}{(1 - \varepsilon_b)^{K-1-m} + \frac{1-p}{p}\frac{(1-\varepsilon_b)^{K-1}}{(1-\varepsilon_a)^{i-1-m}}}. \end{aligned}$$

Hence, the payoff from playing A is no less than the payoff from B if and only if:

$$\frac{1-p}{p}\frac{(1 - \varepsilon_b)^{K-1}}{(1 - \varepsilon_a)^{i-1-m}} \leq 1 - Y.$$

Now, recall that we have $\frac{1-p}{p} \leq \frac{(1-\varepsilon_a)^{L-1}}{(1-\varepsilon_b)^{K-1}}$ by assumption. Thus, it suffices to show that

$$\frac{(1 - \varepsilon_a)^{L-1}}{(1 - \varepsilon_a)^{i-1-m}} \leq 1 - Y,$$

or

$$(1 - \varepsilon_a)^{L+m-i} \leq \text{Prob}(K - 1 - m \text{ or less “}a\text{” signals out of } N - i).$$

Now, let $Z_1 := L + m - i$ and $Z_2 := N - i$. Notice that $Z_1 \leq Z_2$ because $Z_2 - Z_1 = N - L - m = K - 1 - m \geq 0$. The payoff from playing A is no less than the payoff from B if and only if:

$$(1 - \varepsilon_a)^{Z_1} \leq \text{Prob}(Z_2 - Z_1 \text{ or less “}a\text{” signals out of } N - i). \tag{5}$$

Note that

$$\text{Prob}(Z_2 - Z_1 \text{ or less “}a\text{” signals out of } Z_2)$$

$$\geq \text{Prob}(\text{The last } Z_1 \text{ voters do not receive an “}a\text{” signal}) = (1 - \varepsilon_a)^{Z_1}.$$

Thus, eq. (5) indeed holds, and therefore, the payoff from playing A is no less than the payoff from B.

2. Suppose that $X_k = B$.

Suppose the state is α, which happens with probability $P(\alpha)$. Since there have been $m+1$ times of A right after i plays A and there are L times of B in the given sequence, there will be $L-(i-(m+1))$ times of B at which the voter, if she receives signal a, can change the outcome to A. The probability of at least one voter out of $L-(i-(m+1))$ voters receiving signal a is $1-(1-\varepsilon_a)^{L-(i-(m+1))}$.

Suppose the state is β, which happens with probability $1-P(\alpha)$. In this case, i's playing A ensures that there is probability 1 that B will be chosen as the outcome.

Thus, overall, the payoff from i's voting for A is

$$1-P(\alpha)+P(\alpha)(1-(1-\varepsilon_a)^{L-(i-(m+1))}).$$

Thus, the payoff from playing A is no less than the payoff from B if and only if:

$$1-P(\alpha)+P(\alpha)(1-(1-\varepsilon_a)^{L-(i-(m+1))}) \geq 1-P(\alpha)+P(\alpha)Y \iff$$

$$(1-\varepsilon_a)^{L-(i-(m+1))} \leq 1-Y,$$

or

$$(1-\varepsilon_a)^{L-(i-(m+1))} \leq \text{Prob}(K-1-m \text{ or less "}a\text{" signals out of } N-i).$$

Now, let $Z_1 := L-i+(m+1)$ and $Z_2 := N+1-i$. Notice that $Z_1 \leq Z_2$ because $Z_2 - Z_1 = N+1-i-(L-i+(m+1)) = K-(m+1) \geq 0$. The payoff from playing A is no less than the payoff from B if and only if:

$$(1-\varepsilon_a)^{Z_1} \leq \text{Prob}(Z_2 - Z_1 \text{ or less "}a\text{" signals out of } Z_2+1). \tag{6}$$

Note that

$$\begin{aligned}
&\text{Prob}(Z_2 - Z_1 \text{ or less "}a\text{" signals out of } Z_2 - 1)\\
&\geq \text{Prob}(Z_2 - Z_1 \text{ or less "}a\text{" signals out of } Z_2)\\
&\geq \text{Prob}(\text{The last } Z_1 \text{ voters do not receive an "}a\text{" signal})\\
&= (1-\varepsilon_a)^{Z_1}.
\end{aligned}$$

Thus, eq. (6) indeed holds, and therefore, the payoff from playing A is no less than the payoff from B.

Overall, playing A is a best response for voter i. □

Proof for claim (B). The proof closely follows that of Theorem 5. Fix a pure consistent strategy profile. Suppose that if no voter gets a signal, on the path of play, k is the voter who ends the game. Suppose for contradiction that k plays A.

Lemma 6. *For voter k, playing A is suboptimal.*

Proof. Voter k has seen $K-1$ times of action A and $k-K$ times of action B. The posterior on α is

$$P(\alpha) = \frac{p(1-\varepsilon_a)^{k-K}}{p(1-\varepsilon_a)^{k-K} + (1-p)(1-\varepsilon_b)^{K-1}} = \frac{1}{1+\frac{1-p}{p}\frac{(1-\varepsilon_b)^{K-1}}{(1-\varepsilon_a)^{k-K}}}$$
$$= \frac{1}{1+\frac{1-p}{p}\left(\frac{1-\varepsilon_b}{1-\varepsilon_a}\right)^{K-1}(1-\varepsilon_a)^{2K-1-k}}.$$

If voter k plays A, then the game ends and her expected payoff becomes $P(\alpha)$.

Let k' be the voter who ends the game if k plays B.

Case 1: $k' = K+L-1 (= N)$ If voter k plays B, then all the subsequent voters before $K+L-1$ also choose B (when they receive no signal). Let q be the probability that voter $K+L-1$ assigns to A (when he does not receive a signal). Then, k's payoff if she plays B is

$$P(\alpha)\left[1-(1-q)(1-\varepsilon_a)^{K+L-1-k}\right] + (1-P(\alpha))\left[\varepsilon_b + (1-\varepsilon_b)(1-q)\right].$$

Let this be $f(q)$. Note that

$$f(0) = P(\alpha)[1-(1-\varepsilon_a)^{K+L-1-k}] + (1-P(\alpha)) = 1 - P(\alpha)(1-\varepsilon_a)^{K+L-1-k}$$
$$= 1 - \frac{(1-\varepsilon_a)^{K+L-1-k}}{1+\frac{1-p}{p}\left(\frac{1-\varepsilon_b}{1-\varepsilon_a}\right)^{K-1}(1-\varepsilon_a)^{2K-1-k}}$$
$$= \frac{1+(1-\varepsilon_a)^{K-1-k}\left[\frac{1-p}{p}\left(\frac{1-\varepsilon_b}{1-\varepsilon_a}\right)^{K-1}(1-\varepsilon_a)^{K} - (1-\varepsilon_a)^{L}\right](1-\varepsilon_a)^{2n+1-k}}{1+\left(\frac{1-\varepsilon_b}{1-\varepsilon_a}\right)^{n}(1-\varepsilon_a)^{2n+1-k}}$$
$$> \frac{1}{1+\left(\frac{1-\varepsilon_b}{1-\varepsilon_a}\right)^{n}(1-\varepsilon_a)^{2n+1-k}} = P(\alpha),$$

where the last inequality is due to:

$$\frac{1-p}{p}\left(\frac{1-\varepsilon_b}{1-\varepsilon_a}\right)^{K-1}(1-\varepsilon_a)^{K} - (1-\varepsilon_a)^{L} > 0 \iff \frac{1-p}{p} > \frac{(1-\varepsilon_a)^{L-1}}{(1-\varepsilon_b)^{K-1}}.$$

Also,

$$f(1) = P(\alpha) + (1 - P(\alpha))\varepsilon_b > P(\alpha).$$

So $f(q) > P(\alpha)$ for both $q = 0, 1$. This means that k would be better off playing B than A.
Case 2: $k' \leq K + L (= N - 1)$ If voter k plays B, then all the subsequent voters before k' also choose B and voter k' chooses A (when they receive no signal). In this case, if the state is α, then A realizes with probability 1 since voter k' always chooses A. We show that, if the state is β, then B realizes with strictly positive probability. Note that, if the subsequent voters from k' to $2n+1$ all receive signal b, which occurs with probability $\varepsilon_b^{K+L-k'} > 0$, then they all play B by consistency. Consequently, B realizes. This implies that B realizes with a probability of at least $\varepsilon_b^{K+L-k'}$, irrespective of the voters' strategies. Hence, the voter's payoff is at least

$$P(\alpha) \cdot 1 + (1 - P(\alpha)) \cdot \varepsilon_b^{K+L-k'},$$

which is strictly greater than $P(\alpha)$. Therefore, k would be better off playing B than A. □

The lemma implies that the pure consistent strategy profile that we fixed is not an equilibrium. This completes the proof. □

A.8 Proof of Proposition 4

Proof. Fix the strategy profile in which every voter chooses the random voting strategy, and consider voter k who has not received a signal and is about to choose between A or B. The only case in which k's vote affects k's payoff is when exactly n other voters choose A and n other voters choose B. When this happens, k's posterior on each state is

$$\frac{\frac{1}{2}(\varepsilon + \frac{1}{2}(1-\varepsilon))^n(\frac{1}{2}(1-\varepsilon))^n}{\frac{1}{2}(\varepsilon + \frac{1}{2}(1-\varepsilon))^n(\frac{1}{2}(1-\varepsilon))^n + \frac{1}{2}(\varepsilon + \frac{1}{2}(1-\varepsilon))^n(\frac{1}{2}(1-\varepsilon))^n} = \frac{1}{2}.$$

Hence, k is indifferent between choosing A and choosing B.

If voter k receives a signal, then it is obviously a best response to assign probability 1 to the corresponding action given that her action does not affect the actions by the subsequent voters. □

A.9 Proof of Proposition 5

Proof. First, we know that the conclusion of the statement is true for $n = 1$. Hence, below we consider the case of $n \geq 2$. We have

$$C(\varepsilon) = \frac{1}{2} \cdot 1 + \frac{1}{2}(1 - (1-\varepsilon)^{n+1}).$$

On the other hand, we have

$$
\begin{aligned}
RV(\varepsilon) = {} & \text{Prob(no voter receives a signal)} \cdot \text{Prob}(n+1 \text{ or more voters who receive no signal vote for } X), \\
& + \text{Prob(1 voter receives a signal)} \cdot \text{Prob}(n \text{ or more voters who receive no signal vote for } X) \\
& + \text{Prob(2 voters receive a signal)} \cdot \text{Prob}(n-1 \text{ or more voters who receive no signal vote for } X) \\
& + \cdots \\
& + \text{Prob}(n-1 \text{ voters receive a signal}) \cdot \text{Prob(2 or more voters who receive no signal vote for } X) \\
& + \text{Prob}(n \text{ voters receive a signal}) \cdot \text{Prob(1 or more voters who receive no signal vote for } X) \\
& + \text{Prob}(n+1 \text{ or more voters receive a signal}) \cdot 1
\end{aligned}
$$

where X is the action corresponding to the realized state. We can show that $RV(\varepsilon)$ has the following upper bound $\overline{RV}(\varepsilon)$:

$$
\begin{aligned}
RV(\varepsilon) \leq \overline{RV}(\varepsilon) := {} & \text{Prob(no voters receive a signal)} \cdot \frac{1}{2} \\
& + \text{Prob(1 voter receives a signal)} \cdot \frac{1+g(n)}{2} \\
& + \text{Prob(2 or more voters receive a signal)} \cdot 1,
\end{aligned}
$$

where $g(n)$ is the probability that exactly n voters out of $2n$ voters, conditional on not receiving a signal, vote for A, and we have

$$
g(n) = \frac{(2n)!}{n!n!}\left(\frac{1}{2}\right)^n\left(\frac{1}{2}\right)^n = \frac{(2n)!}{n!n!}\left(\frac{1}{2}\right)^{2n}.
$$

Note also that

$$
\text{Prob(1 voter receives a signal)} = (2n+1)\varepsilon(1-\varepsilon)^{2n}.
$$

Hence,

$$
\overline{RV}(\varepsilon) = (1-\varepsilon)^{2n+1}\frac{1}{2} + (2n+1)\varepsilon(1-\varepsilon)^{2n}\frac{1+\left[\frac{(2n)!}{n!n!}\left(\frac{1}{2}\right)^{2n}\right]}{2} + O(\varepsilon^2).
$$

Thus,

$$
C(\varepsilon) - \overline{RV}(\varepsilon)
$$

$$
= \frac{1}{2}(1-(1-\varepsilon)^{n+1}) - \left[-(1-(1-\varepsilon)^{2n+1})\frac{1}{2} + (2n+1)\varepsilon(1-\varepsilon)^{2n}\frac{1+\left[\frac{(2n)!}{n!n!}\left(\frac{1}{2}\right)^{2n}\right]}{2} + O(\varepsilon^2)\right]
$$

$$= 1 - \frac{1}{2}\left[(1-\varepsilon)^{n+1} + (1-\varepsilon)^{2n+1} + (2n+1)\varepsilon(1-\varepsilon)^{2n}\left(1 + \left[\frac{(2n)!}{n!n!}\left(\frac{1}{2}\right)^{2n}\right]\right)\right] - O(\varepsilon^2)$$

Now, let

$$h(\varepsilon) := (1-\varepsilon)^{n+1} + (1-\varepsilon)^{2n+1} + (2n+1)\varepsilon(1-\varepsilon)^{2n}\left(1 + \left[\frac{(2n)!}{n!n!}\left(\frac{1}{2}\right)^{2n}\right]\right).$$

We have

$$\begin{aligned} h'(\varepsilon) = &- (n+1)(1-\varepsilon)^n - (2n+1)(1-\varepsilon)^{2n} \\ &+ (2n+1)\left(1 + \left[\frac{(2n)!}{n!n!}\left(\frac{1}{2}\right)^{2n}\right]\right)\left((1-\varepsilon)^{2n} - 2n\varepsilon(1-\varepsilon)^{2n-1}\right). \end{aligned}$$

Hence,

$$\begin{aligned} h'(0) &= -(n+1) - (2n+1) + (2n+1)\left(1 + \left[\frac{(2n)!}{n!n!}\left(\frac{1}{2}\right)^{2n}\right]\right) \\ &= -(n+1) + (2n+1)\left(\frac{(2n)!}{n!n!}\left(\frac{1}{2}\right)^{2n}\right). \end{aligned}$$

Now, note that $\frac{(2n)!}{n!n!}\left(\frac{1}{2}\right)^{2n} \leq \frac{1}{2}$. To see this, first note that when $n = 1$, this inequality holds. Now, suppose this inequality holds when $n = k$. Then, we have

$$\begin{aligned} \frac{(2(k+1))!}{(k+1)!(k+1)!}\left(\frac{1}{2}\right)^{2(k+1)} &= \frac{(2k)!}{k!k!}\left(\frac{1}{2}\right)^{2k} \times \frac{(2k+1)(2k+2)}{(k+1)^2}\frac{1}{2^2} \\ &\leq \frac{1}{2} \times \frac{2k+1}{2k+2}\frac{2k+2}{2k+2} \leq \frac{1}{2}. \end{aligned}$$

Hence, we have

$$h'(0) \leq -(n+1) + (2n+1)\frac{1}{2} = -\frac{1}{2}.$$

Note also that $h(0) = 2$. Thus, there is $\bar{\varepsilon} > 0$ such that for all $\varepsilon < \bar{\varepsilon}$, $h(\varepsilon) < 2 - \frac{\varepsilon}{3}$. This implies that for all $\varepsilon < \bar{\varepsilon}$,

$$C(\varepsilon) - RV(\varepsilon) \geq C(\varepsilon) - \overline{RV}(\varepsilon) > 1 - \frac{1}{2}(2 - \frac{\varepsilon}{3}) - O(\varepsilon^2) = \frac{\varepsilon}{6} - O(\varepsilon^2).$$

Hence, there is $\hat{\varepsilon} > 0$ such that for all $\varepsilon < \hat{\varepsilon}$, $C(\varepsilon) - RV(\varepsilon) > 0$, and thus $C(\varepsilon) > RV(\varepsilon)$. This completes the proof. □

A.10 Examples for the Abstention Model

Example 3. Consider the strategy profile in which all voters use a strategy that is consistent in the abstention model, vote for Φ when receiving signal ϕ and all past actions have been Φ, vote for A when receiving signal ϕ and the first past action that was not Φ was A, and vote for B when receiving signal ϕ and the first past action that was not Φ was B. If voter k receives signal ϕ and the first past action that was not Φ was A, then k's belief assigns probability 1 to state α. Symmetrically, if voter k receives signal ϕ and the first past action that was not Φ was B, then k's belief assigns probability 1 to state β. We show that this is an equilibrium.

First, suppose that voter k receives signal a. Consider the case in which all past actions have been Φ or the first past action that was not Φ was A. In this case, if k chooses A, then her payoff is 1 with probability 1. Since 1 is the highest feasible payoff, choosing A is indeed a best response. Consider the case in which the first past action that was not Φ was B. In this case, there is a positive probability that k's action affects the outcome of the game, and this probability is independent of k's action. Moreover, when k's action affects the outcome of the game, it is her best response to vote for A. Thus, it is a best response for the voters receiving signal a to play A regardless of the history.

A symmetric argument shows that it is a best response for the voters receiving signal b to play B regardless of the history.

Suppose that voter k receives signal ϕ. Consider the case in which all the past actions are Φ. If $k = N$, then she is indifferent among all actions, so playing Φ is a best response. So suppose $k < N$. If k plays Φ, then her payoff is

$$(1-\varepsilon)^{N-k}\frac{1}{2} + (1-(1-\varepsilon)^{N-k}) \cdot 1 = 1 - \frac{1}{2}(1-\varepsilon)^{N-k}$$

because if no one receives the signal then the payoff is $\frac{1}{2}$, while if there is at least one voter receiving a signal then the payoff is 1. If she plays A, then her payoff is at most

$$1 - \frac{1}{2}\left[(1-\varepsilon)^{N-k} + \frac{1}{2}(N-k)\varepsilon(1-\varepsilon)^{N-k-1}\right]$$

because the cases in which her payoff is not 1 include the situations in which the state is β and there is only zero or one voter who receives a signal after k. It is easy to see by inspection that the payoff from playing Φ is bigger than the upper bound of the payoff from playing A. In the same way, the payoff from playing Φ is bigger than the payoff from playing B. Hence, playing Φ is a best response.

Finally, suppose that voter k receives signal ϕ and consider the case in which there is a past action that was not Φ. Suppose without loss that the first such action was A. In this case, k assigns probability 1 to state α. By playing A, she can guarantee the payoff of 1, which is the highest payoff, so this is a best response. □

Example 4. Suppose $N = 4$. Consider the following strategy profile.

- Voter 1 votes for Φ if her signal is a, A if her signal is ϕ, and B if her signal is b.
- Voter 2 plays a strategy that is consistent in the abstention model. If he receives signal ϕ, then he plays
  - Φ if 1 played A;
  - A if 1 played Φ;
  - B if 1 played B.
- Voter 3 plays a strategy that is consistent in the abstention model. If she receives signal ϕ, then she plays
  - Φ if the action sequence was (A, Φ);
  - A if the action sequence was $(\Phi, \cdot)$ or (A, A);
  - B if the action sequence was $(B, \cdot)$ or (A, B).
- Voter 4 plays a strategy that is consistent in the abstention model. If he receives signal ϕ, then he plays
  - Φ if the action sequence was (A, Φ, Φ);
  - A if the action sequence was $(\Phi, \cdot, \cdot)$ or (A, A, B);
  - B if the action sequence was $(B, \cdot, \cdot)$ or $(A, B, \cdot)$ or $(A, \cdot, B)$.

Voters 3 and 4 have information sets that can be reached with probability zero when they do not receive a signal. In such a case, they assign probability 1 to state α under the history in which the above specification says they play A, and they assign probability 1 to state β under the history in which the above specification says they play B. We show that this is an equilibrium when $\varepsilon > 0$ is sufficiently small.

Let us check the incentives. The incentives of voters 2, 3, and 4 are straightforward so we check the incentive of voter 1. If she receives signal a or b, it is again straightforward that following the given strategy is a best response. So consider the case in which she receives signal ϕ.

In this case, if she plays A, then her payoff is 1 if the state is α while, if the state is β, her payoff is at least $\frac{1}{2} \cdot 3\varepsilon(1-\varepsilon)^2$ because if at least one of the subsequent voters receives signal b then there is probability $1/2$ that the payoff is 1.

If she plays Φ, then her payoff is 1 if the state is α. If the state is β, the payoff is $O(\varepsilon^2)$ because the only cases in which the payoff is 1 is when at least two subsequent voters receive signal b.

If she plays B, then her payoff is 1 if the state is β. If the state is α, the payoff is $O(\varepsilon^2)$ because the only cases in which the payoff is positive is when at least two subsequent voters receive signal a.

Overall, when voter 1 receives signal ϕ, it is a best response to play A when $\varepsilon > 0$ is sufficiently small. $\square$

In both examples, the ex ante payoff is

$$(1-\varepsilon)^N \frac{1}{2} + (1-(1-\varepsilon)^N) \cdot 1.$$

Note that this is the highest payoff achievable by any social choice function.

A.11 Proof of Theorem 7

Proof. Fix a sequence $S = (X_1, \ldots, X_k)$ such that X_k appears exactly $n+1$ times. Consider the consistent strategy profile σ satisfying the following conditions:

1. $\sigma_i(m, \emptyset)(X_i) = 1$ if $m = |\{X_l = A | 1 \leq l \leq i-1\}|$.

2. $\sigma_i(m, \emptyset)(A) = 1$ if $m > |\{X_l = A | 1 \leq l \leq i-1\}|$.

3. $\sigma_i(m, \emptyset)(B) = 1$ if $m < |\{X_l = A | 1 \leq l \leq i-1\}|$.

Now we check incentives. First, voters who have received a signal take a best response given Cases 2 and 3 above.

Second, under any histories described in Case 2 above, by Bayes rule, i has a belief that the state is α. Hence, it is straightforward to see that i is taking a best response. A symmetric argument holds for Case 3 as well.

Third, under any histories described in Case 1 above, without loss of generality, let $X_k = A$. Suppose that, in the default sequence, there have been s_A times of A and s_B times of B, and there will be s^A times of A and s^B times of B. Voter i's posterior on state α is then

$$P(\alpha) = \frac{(1-\varepsilon)^{s_B}}{(1-\varepsilon)^{s_A} + (1-\varepsilon)^{s_B}}. \tag{7}$$

Case 1: Suppose that $X_i = A$. Then, if i plays A, her payoff is

$$P(\alpha) \cdot 1 + (1 - P(\alpha))(1 - (1 - \varepsilon)^{s^A}) = 1 - (1 - P(\alpha))(1 - \varepsilon)^{s^A}.$$

If, instead, i plays B, then her payoff is at most:

$$P(\alpha)(1 - (1 - \varepsilon)^{s^B}) + (1 - P(\alpha)) \cdot 1 = 1 - P(\alpha)(1 - \varepsilon)^{s^B}$$

because, for the outcome to be A, there must be at least one time in the future when a voter who is supposed to play B in the default sequence receives signal a (and thus votes for A).

The former payoff is no less than the latter upper bound if and only if

$$(1 - P(\alpha))(1 - \varepsilon)^{s^A} \leq P(\alpha)(1 - \varepsilon)^{s^B}$$

which, by (7), is equivalent to:

$$(1 - \varepsilon)^{s_A}(1 - \varepsilon)^{s^A} \leq (1 - \varepsilon)^{s_B}(1 - \varepsilon)^{s^B},$$

or $s_A + s^A \geq s_B + s^B$.

Now, notice that, since S is the default sequence that ends with A and i plays A in S, we have that $s_A + 1 + s^A = n + 1$ and $s_B + s^B \leq n$. Hence, we indeed have $s_A + s^A \geq s_B + s^B$, completing the proof for this case.

Case 2: Suppose that $X_i = B$. Then, if i plays B, her payoff is

$$P(\alpha) \cdot 1 + (1 - P(\alpha))(1 - (1 - \varepsilon)^{s^A})$$

as in Case 1. If, instead, i plays A, then her payoff is at most:

$$P(\alpha) \cdot 1 + (1 - P(\alpha))(1 - (1 - \varepsilon)^{s^A})$$

because, for the outcome to be B, there must be at least two times in the future when a voter who is supposed to play A in the default sequence receives signal b (and thus votes for B). This event is less likely than the event in which there is at least one time in the future when a voter who is supposed to play A in the default sequence receives signal b (and thus votes for B). This latter event happens with probability at most $1 - (1 - \varepsilon)^{s^A}$. Now, since the former payoff is equal to the latter upper bound, it is indeed voter i's best response to play B. $\square$