

# Biased Sampling Equilibrium

Kyle Woodward\*

February 2020

## Abstract

In games with a large number of players, it may be difficult for an agent to correctly assess her opponents’ strategies. I propose an equilibrium concept, *biased sampling equilibrium* (BSE), in which agents respond to a finite biased sample of opponent play. Statistical bias in the sampling process leads to a novel equilibrium property: when the sampling process is homophilic, equilibrium strategies may not be isotone in type, and may even be significantly antitone. I apply BSE to a model of strategic voting, and find that extreme partisans may abstain from voting more frequently than moderate agents. Nonmonotonicity of turnout in ideological strength qualitatively matches behavior in recent British elections.

*To be quite frank, I did not believe it would happen.*

– R.S., regretful Leave supporter<sup>†</sup>

## 1 Introduction

In games with a large number of players, it may be difficult for agents to form correct beliefs regarding opponent behavior. When it is not possible to obtain perfect information regarding opponent strategies, realized actions will be consistent with the limited information the agent has. When information is limited, the manner in which it is acquired becomes crucial to determining agent beliefs, and is therefore an important piece of any equilibrium concept.

This paper contributes to the literature on misperception in games by explicitly modeling a source of bias in the belief formation process. I define a *biased sampling equilibrium* (BSE) so that agents are responding to finite and biased samples of opponent play. Agents observe a relatively small number of samples of opponent play, and these samples are taken from a distribution that does not match the underlying distribution of types. Agents respond as if their biased samples are a precise estimate of population play. When the sampling process is homophilic, so that agents draw

---

\*University of North Carolina at Chapel Hill; [kyle.woodward@unc.edu](mailto:kyle.woodward@unc.edu). I would like to thank a seminar audience at UNC for helpful feedback.

<sup>†</sup>“‘I thought I’d put in a protest vote’: the people who regret voting to leave.” *The Guardian*. 25 November 2017.

samples too frequently from those with similar types, equilibrium strategies can be nonmonotone in type even when, holding fixed beliefs, best responses are monotone in type.

BSE builds on the *sampling equilibrium* defined in Osborne and Rubinstein (2003). In a sampling equilibrium, agents respond to finite samples of population behavior. Because the sampling process is identical for all agents, all agents respond to identical beliefs. When best responses are monotone in type, this implies that equilibria are monotone in type (my Observation 1). BSE extends the notion of sampling equilibrium to allow for heterogeneous sampling procedures. When two types have different sampling processes they will in general have different beliefs, and responses will depend not only type but also on the beliefs induced by that type’s sample.

One key question in classical equilibrium analysis is whether equilibria are monotone (c.f. Athey (2001), McAdams (2003), Reny (2011), Woodward (2019)). Monotone Bayesian Nash equilibria exist in many games of interest, but are far from universal. Monotonicity typically results from single-crossing and quasisupermodularity of utility, but also implicitly requires that agents cannot draw too much adverse inference from their own type realization (c.f. McAdams (2007)). In BSE, adverse inference is not due to learning about opponent types, but follows from a restriction on inference itself. It is perhaps not surprising that BSE may be nonmonotone. However, BSE may be not only nonmonotone but also (nearly) globally antitone: with few exceptions, strategies are monotonically decreasing in type.

Equilibrium nonmonotonicity may result when agents’ samples are insufficiently reflective of population play, and beliefs regarding opponent play differ significantly between types. These beliefs will differ when sampling processes differ significantly between types. An individual’s social network is a natural pool from which to sample population play, in which case beliefs will depend on who an agent knows. Social networks are not complete nor even particularly diverse, and humans tend to exhibit homophilic preferences and associate with individuals similar to themselves (McPherson et al., 2001). Homophilic preferences extend to traits which are not directly observable, such as political preferences (Knoke (1990), Halberstam and Knight (2016)). While BSE places no assumptions on the direction or scale of bias in the sampling process, in application I assume the sampling process is homophilic. This induces a regular source of bias in action selection: agents sample others with similar preferences, potentially irreflective of the population at large. In a voting context, left-extremists confuse moderates for right-extremists, and right-extremists confuse moderates for left-extremists.

In the context of strategic voting, homophilic sampling implies that individuals with relatively partisan preferences are relatively unlikely to vote. An individual on one end of the left-right spectrum is likely to query the voting behavior of similarly extreme individuals. These individuals are unlikely to vote for the opposition candidate, and are therefore unlikely to suggest that the election is competitive. More moderate individuals are more likely sample from both sides of the political divide. They are therefore more likely to believe that the election is competitive, and that casting a ballot is worth the cost. Data from the British Election Study suggests recent British elections exhibit similar nonmonotone turnout rates, and extreme partisans have turned out at

relatively low rates.<sup>1</sup>

While motivated by application to voting models, this paper is most clearly related to the literatures on misperception and bounded rationality. As noted above, biased sampling equilibrium generalizes the sampling equilibrium concept of Osborne and Rubinstein (2003). Salant and Cherry (2019) also extend the sampling equilibrium concept, but focus on the method of drawing inference from a finite sample. I hold fixed the inference concept, and vary the sampling procedure. My framework is static, and hence my approach is distinct from the studies of learning from dynamic sampling, such as Oyama et al. (2015) and Häfner (2018). The realization of the biased sampling process defines a directed network between agents, hence there are natural ties to peer confirming equilibrium (Lipnowski and Sadler, 2019). A key distinction is that my sampling network is random and directed, while in peer confirming equilibrium it is fixed and undirected. When the information-sharing network is unknown, biased sampling equilibrium may serve as a reduced-form approximation to peer confirming equilibrium. This is a useful direction for future study.

## 2 Model

The *game* is  $G = (\Theta, F, A, R)$ . The set of types  $\Theta$  is a finite set with order  $\prec_\Theta$ , and the proportion of players with type  $\theta \in \Theta$  is  $F(\theta)$ .<sup>2</sup> Types  $\theta$  and  $\theta'$  are *adjacent* if there is no  $\theta''$  with  $\theta \prec_\Theta \theta'' \prec_\Theta \theta'$  or  $\theta' \prec_\Theta \theta'' \prec_\Theta \theta$ . Agents choose from a finite set of actions  $A$  with order  $\prec_A$ . A strategy profile is  $\sigma = (\sigma_\theta)_{\theta \in \Theta}$ , where  $\sigma_\theta \in \Delta A$ , and  $\sigma_{\theta a}$  gives the probability that an agent with type  $\theta$  plays action  $a$ .

The game is augmented by a *sampling process*  $S = (k, H)$ . The sampling process specifies  $k$ , the number of samples drawn, and  $H : \Theta \times \Theta \rightarrow [0, 1]$ , the distribution by which they are drawn. The sampling distribution is *unbiased* if  $H(\cdot; \theta) = F(\cdot)$ , and is *homophilic* if  $H(\cdot; \theta)/F(\cdot)$  is single-peaked at  $\theta$  and nonconstant. The sampling distribution is *nontrivial* if for any type  $\theta$  and any adjacent type  $\theta'$ ,  $H(\theta'; \theta) > 0$ . Given a sampling process  $S$ , the support of realized samples is contained in  $K = \{\alpha \in \mathbb{N}^{|A|} : \sum_a \alpha_a = k\}$ .

The agent's *reaction function* is  $R : K \times \Theta \rightarrow A$ , which gives his response to a distribution over opponent play, conditional on his own type.<sup>3</sup> Given a strategy profile  $\sigma$ , a sampling process  $S$ , and a type  $\theta$ , the distribution over realized samples  $\alpha \in K$  is a multinomial distribution over  $|A|$  outcomes, where the probability of outcome  $a$  is

$$\pi(a; S, \theta) = \sum_{\theta' \in \Theta} H(\theta'; \theta) \sigma_{\theta' a}.$$

---

<sup>1</sup>The literature on voter turnout typically finds that turnout increases with ideological strength (c.f. Geys (2006)). As illustrated in Figure 2, recent British elections have bucked this trend. See Larcinese (2009) for a model of voter turnout which supports monotone turnout rates in older (1997) British elections.

<sup>2</sup>The ordering on  $\Theta$  is inessential to the definition and existence of equilibrium, but necessary to empower some discussion of monotonicity.

<sup>3</sup>I make no assumption on how the response function  $R$  relates to type and sampled play. In many cases it is natural to assume that the agent has an underlying utility specification, and that  $R$  is a best response to some inference drawn from the sampled distribution  $s(\sigma; S, \theta)$ .

The realized sample distribution is used to define a biased sampling equilibrium.

**Definition 1** (Biased sampling equilibrium). *A strategy profile  $\sigma$  is a biased sampling equilibrium (BSE) if for all types  $\theta \in \Theta$  and all actions  $a \in A$ ,  $\Pr(R(\alpha; \theta) = a | \theta) = \sigma_{\theta a}$ .*

In a BSE, the probability a type selects an action is consistent with the probability that this type witnesses a distribution over opponent play that induces the selection of this action. The sets of available types  $\Theta$  and actions  $A$  are both finite. Then because only a finite number of samples is taken, the set of potential distributions over samples is also finite.

**Theorem 1** (Equilibrium existence). *A biased sampling equilibrium exists.*

*Proof.* This follows from finiteness and direct application of Kakutani's fixed point theorem.  $\square$

## 2.1 Equilibrium antitonicity

When the sampling distribution is unbiased,  $H(\cdot; \theta) = F(\cdot)$ , BSE is equivalent to the sampling equilibrium of Osborne and Rubinstein (2003). My analysis of potential nonmonotonicity in BSE proceeds from the following observation.

**Observation 1** (Monotone sampling equilibrium). *If the sampling process is unbiased and the reaction function is monotone, so that  $\theta < \theta'$  implies  $R(\cdot; \theta) \preceq_{FOSD} R(\cdot; \theta')$ , all sampling equilibria are monotone. Because each type  $\theta$  derives identical beliefs over opponent play, the reaction function  $R$  is sufficient to guarantee monotonicity.*

A consequence of Observation 1 is that nonmonotonicity in a BSE is necessarily due to bias in the sampling process, and not to finite sampling error alone. While BSE may be nonmonotone, under fairly general circumstances it cannot be globally antitone.

**Lemma 1** (No antitone biased sampling equilibrium). *Suppose that the sampling distribution is increasing in type, and the response function is increasing and nonconstant in type (fixing distribution), and decreasing in distribution. Then there is no antitone biased sampling equilibrium.*

*Proof.* Suppose otherwise. Then the distribution over sampled play is decreasing in type, and a high type's sampled distribution suggests a (first order) higher response than a low type's sampled distribution. Because  $R$  is decreasing in distribution, the high type would have a higher response if he observed the low type's sample. Since  $R$  is nonconstant in type, it follows that either the high type's or low type's strategy does not satisfy the definition of BSE.  $\square$

The assumption that the reaction function  $R$  is nonconstant in type is crucial to proving that there is no antitone BSE. When  $R$  is derived from a utility maximization problem,  $R$  will be nonconstant in type when two types have different best responses to the same distribution over opponent play. Even when the sampling process is biased, resulting beliefs may not depend too strongly on the agent's type, and nonconstant responses may locally dominate any effect of shifting beliefs. To evaluate equilibrium antitonicity on a reasonable footing, I identify types at which  $R$  changes in order to eliminate them from consideration.

**Definition 2** (Inflection point). *Let  $\sigma$  be a strategy profile. A type  $\theta \in \Theta$  is an inflection point if there is an adjacent type  $\theta'$  and actions  $a$  and  $a'$  such that*

$$\Pr(R(\alpha; \theta) = a | \theta) \neq \Pr(R(\alpha; \theta') = a | \theta) \quad \text{and} \quad \Pr(R(\alpha; \theta') = a' | \theta) \neq \Pr(R(\alpha; \theta) = a' | \theta).$$

Holding fixed a sampled distribution of equilibrium play, a type is an inflection point if a slightly higher or lower type has a different response. Because the action space  $A$  is finite, the distribution of realized samples  $\alpha$  is discrete. When the response function  $R$  is derived from some best response to beliefs generated by sampled play, it is frequently locally constant in type: playing a particular action indicates that it generates greater utility than other actions, and utility is similar when types are similar. Then holding fixed the sample distribution, the response function frequently changes only at a handful of types  $\theta$ .

I show below that even when response functions are isotone, BSE may not be isotone. This nonisotonicity is limited, in the sense that equilibria may be monotonically decreasing at many types, but will not typically be monotonically decreasing at inflection points. By definition, responses tend to be isotone near inflection points. This isotonicity is discrete, a jump from one action to another, and it is difficult for sampling to be sufficiently limited to overpower this shift in incentives. I therefore slightly restrict the notion of antitonicity to hold only at types that are not inflection points.

**Definition 3** (Antitonicity). *A strategy profile  $\sigma$  is locally antitone at  $\theta \in \Theta$  if, for any  $\underline{\theta}, \bar{\theta}$  adjacent to  $\theta$  with  $\underline{\theta} < \theta < \bar{\theta}$ ,  $\sigma_{\underline{\theta}} \succ_{\text{FOSD}} \sigma_{\theta} \succ_{\text{FOSD}} \sigma_{\bar{\theta}}$ . A strategy profile  $\sigma$  is nearly antitone if it is locally antitone at all types  $\theta \in \Theta$  that are not inflection points.*

A precise characterization of when nearly antitone equilibria exist depends in a complicated way on the number of samples drawn, as well as the bias inherent in the sampling process. With sufficiently many relatively unbiased samples, biased sampling information is close to full information, and equilibrium will tend to be as monotone as the response function. With relatively few, moderately biased samples, local selection effects can dominate. When agents sample only opponents with identical types, equilibrium monotonicity may be regained.

### Illustration: a congestion game

There are  $t + 1$  types  $\theta \in \{0, 1/t, \dots, 1 - 1/t, 1\}$ , each with equal proportion in the population. Agents play a congestion game with actions  $a \in \{E, N\}$ . Playing  $N$  gives utility 0 and playing  $E$  gives utility  $\theta - \mu(E)$ , where  $\mu(E)$  is the fraction of the population action  $E$ . Agents with lower types are, ceteris paribus, less willing to play  $E$ , so I define  $N \prec_A E$ .

Agents draw  $k$  samples, and respond as if their sample perfectly reflects the distribution of play in the population. Given type  $\theta$ , an agent will choose action  $E$  if less than  $\theta$  of her sample has

chosen  $E$ , and will choose  $N$  otherwise. That is,

$$R(\alpha; \theta) = \begin{cases} E & \text{if } \frac{\alpha_E}{k} \leq \theta, \\ N & \text{otherwise.} \end{cases}$$

Given type  $\theta$ , let  $\bar{\alpha}_E(\theta)$  denote the maximum number of action  $E$  in a sample that will result in type  $\theta$  playing  $E$ . By construction,  $\bar{\alpha}_E(\theta) = \lfloor k\theta \rfloor$ . Provided the type space is sufficiently rich ( $t$  is sufficiently large)  $\bar{\alpha}_E$  is constant on local intervals. The probability that type  $\theta$  plays  $E$  is  $\Pr(\alpha_E \leq \bar{\alpha}_E(\theta))$ . Then because the sampling distribution is identical for all types, this probability is constant on local intervals when  $\bar{\alpha}_E$  is constant. Then in sampling equilibrium the probability of choosing action  $E$  is a step function which is monotone in type  $\theta$ .

**Lemma 2** (Inflection points in sampling equilibrium). *In a sampling equilibrium,  $\theta$  is an inflection point if  $\min\{k\theta - \lfloor k\theta \rfloor, \lceil k\theta \rceil - k\theta\} < k\varepsilon$ .*

*Proof.* Let  $\theta, \theta' \in \Theta$ , and assume  $\theta < \theta'$ . To be an inflection point,  $\theta$  and  $\theta'$  must be adjacent, therefore assume  $\theta + 1/t = \theta'$ . Fixing an realized sample  $\alpha$ , the response function  $R$  will differ between these two types if and only if there is  $\alpha_E$  such that  $k\theta < \alpha_E \leq k\theta' = k\theta + k\varepsilon$ . Then  $0 < \alpha_E - k\theta \leq k\varepsilon$ ; equivalently, there is an integer  $n$  such that  $-k\varepsilon \leq k\theta - n < 0$ . The left-hand inequality is easiest to satisfy at  $n = \lceil k\theta \rceil$ , thus  $\theta$  is an inflection point if  $\lceil k\theta \rceil - k\theta \leq k\varepsilon$ . A symmetric argument applies if we analyze  $\theta'$  instead, giving the desired result.  $\square$

In the congestion game, inflection points are those types which, in order to respond with entry, are willing to sample one less (or one more) entry than their adjacent neighbors. They are the types on either side of the discontinuities in the left panel of Figure 1. The same basic logic applies to BSE, but an equivalence is not guaranteed: the definition of inflection point depends on the strategy profile analyzed. Holding fixed a distribution over realized samples, the set of inflection points does not change between sampling equilibrium and BSE. However, the distribution of realized samples is not fixed, thus it cannot be ruled out that the set of inflection points is distinct between the two equilibrium concepts.

With a biased sampling procedure the sampling distribution is no longer constant in type. When the type space is sufficiently rich ( $t$  is sufficiently large compared to  $k$ ) this will result in nonmonotone entry probabilities: agents with low types sample other agents with low types, who are generally unwilling to enter. Although agents with higher types are *ex ante* more willing to enter, they tend to draw from stronger distributions over realized samples, depressing their entry probability.

**Lemma 3** (No isotone biased sampling equilibrium). *Suppose that the sampling distribution is homophilic, increasing in type, and nontrivial. If  $t > k$ , there is no isotone biased sampling equilibrium in the congestion game.*

*Proof.* Suppose otherwise. Since  $t > k$ , the pigeonhole principle implies that there are two adjacent types  $\theta$  and  $\theta'$  such that  $\bar{\alpha}_E(\theta) = \bar{\alpha}_E(\theta')$ ,  $\theta < \theta'$ , and  $\bar{\alpha}_E(\theta') \neq \bar{\alpha}_E(\theta' + 1/t)$ . Since  $H$  is homophilic,

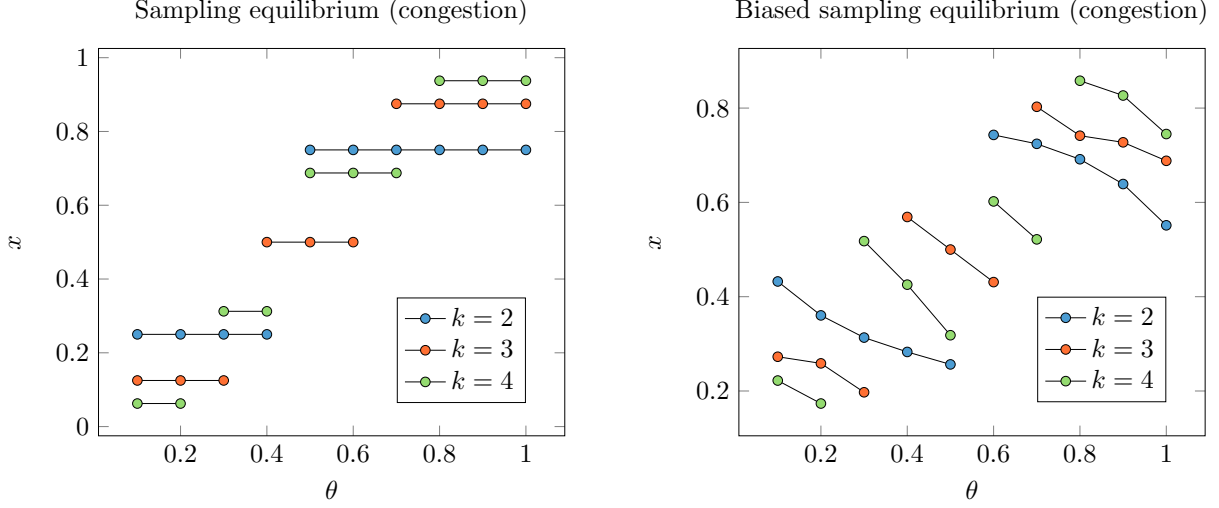


Figure 1: In a standard sampling equilibrium, strategies are isotone in type (Observation 1). In a biased sampling equilibrium, strategies are not isotone when the type space is rich (Lemma 3), and may be nearly antitone.

increasing in type, and samples adjacent types,  $\theta'$  is an inflection point. Then is the case that  $\Pr(\alpha_E \leq \bar{\alpha}_E(\theta)|\theta) > \Pr(\alpha_E \leq \bar{\alpha}_E(\theta')|\theta')$ . Then type  $\theta'$  plays  $E$  with lower probability than type  $\theta < \theta'$ , contradicting isotonicity.  $\square$

The probability distribution over agent responses to realized samples depends on the probability that an agent with type  $\theta$  samples no more than  $\bar{\alpha}_E(\theta)$  entries. Let  $\pi_E(\theta)$  be the probability that a single sample of entry is drawn,

$$\pi_E(\theta) = \sum_{\theta' \in \Theta} H(\theta'; \theta) \sigma_{\theta'E}, \text{ and } \Pr(R(\alpha; \theta) = E|\theta) = \sum_{n=0}^{\bar{\alpha}_E(\theta)} \binom{k}{n} \pi_E(\theta)^n (1 - \pi_E(\theta))^{k-n}.$$

The strategy profile  $\sigma$  is a BSE if

$$\sigma_{\theta E} = \sum_{n=0}^{\bar{\alpha}_E(\theta)} \binom{k}{n} \left[ \sum_{\theta' \in \Theta} H(\theta'; \theta) \sigma_{\theta'E} \right]^n \left( 1 - \sum_{\theta' \in \Theta} H(\theta'; \theta) \sigma_{\theta'E} \right)^{k-n}. \quad (1)$$

The solution to (1) is shown in Figure 1. BSE may be nearly antitone, and away from inflection points entry probabilities may be strictly decreasing in type. Concordant with the proof of Lemma 3, inflection points in BSE are clearly visible as points at which equilibrium strategies are increasing. Similarly, as  $k$  increases the number of inflection points also increases (in line with Lemma 2). As the number of inflection points increases, equilibrium strategies appear increasingly isotone.

### 3 A voting game

I now consider the implications of biased sampling on strategic voting. Reaction functions are derived from a utility model in which voting is costly, and players have differential strength of preference for one of two candidates. In a standard sampling equilibrium, partisans vote for their preferred candidates while moderate voters refrain from voting. Under biased sampling, willingness to vote is nonmonotone.<sup>4</sup> I show that nonmonotone participation is robust to the introduction of extreme candidates who appeal directly to agents who tend to abstain in a two-candidate system.

#### 3.1 Two candidates

Agents' types are drawn independently and uniformly from  $\{-1, -1 + 1/t, \dots, 1 - 1/t, 1\}$ ,  $t \in \mathbb{N}$ . There are two candidates for office,  $p \in \{-P, +P\}$ , where  $P \in (0, 1)$ . If candidate  $p$  is elected, an agent with type  $\theta$ 's utility is  $-|\theta - p|$ . Agents choose an action  $a \in A = \{-P, \emptyset, +P\}$ , and let  $-P \prec_A \emptyset \prec_A +P$ . Candidates are elected by simple majority, with random tiebreaking. Agents who choose to vote,  $a \neq \emptyset$ , incur positive cost  $c$ ,  $0 < c < P$ ;<sup>5</sup> agents who abstain from voting,  $a = \emptyset$ , incur no cost. If an agent with type  $\theta$  takes action  $a$  and candidate  $p$  is elected, his ex post utility is

$$u(a; \theta, p) = -|\theta - p| - \mathbb{1}[a \neq \emptyset]c.$$

I consider a biased sampling procedure in which agents sample uniformly from types that are similar to their own. Fix  $\delta > 0$ , and let  $N(\theta; \delta) = \#\{\theta'' : |\theta'' - \theta| < \delta\}$  be the number of types that are within  $\delta$  of type  $\theta$ . Define the sampling distribution  $H$  by

$$H(\theta'; \theta, \delta) = \begin{cases} \frac{1}{N(\theta; \delta)} & \text{if } |\theta' - \theta| < \delta, \\ 0 & \text{if } |\theta' - \theta| \geq \delta. \end{cases}$$

Note that if  $\delta > 2$ , then  $H(\cdot; \theta) = F(\cdot)$ , and the sample is unbiased.

The reaction function  $R$  is defined by utility maximization, subject to beliefs induced by the realized sample  $\alpha$ . After drawing a sample  $\alpha$ , the agent believes that the proportion of his sample voting for a particular candidate is equal to the proportion of the full population voting for that candidate. Because voting is costly, the agent will vote only if he is pivotal. Since he believes his sample is an accurate representation of population play, he will vote only if his sample consists of uniform abstention, or an equal number of individuals voting for each candidate.

For computational simplicity I focus on the two-sample case,  $k = 2$ . With two samples, the agent's realized sample may be written as  $\alpha = (\tilde{a}_1, \tilde{a}_2)$ . When the agent believes he is pivotal, he believes that if he abstains from voting the election will be randomly decided between his favored

<sup>4</sup>BSE strategies are nonmonotone in the order defined on the action space, which is distinct from monotonicity in turnout. It is straightforward to show that any strategy that is isotone in the defined order also has turnout monotonically increasing in ideological strength, and nonmonotonicity of turnout implies nonmonotonicity in the order defined on the action space.

<sup>5</sup>The assumption that  $c < P$  ensures that agents with types  $|\theta| > P$  will choose to vote in a sampling equilibrium. If  $c > P$  there is no voting in any sampling equilibrium, biased or not.



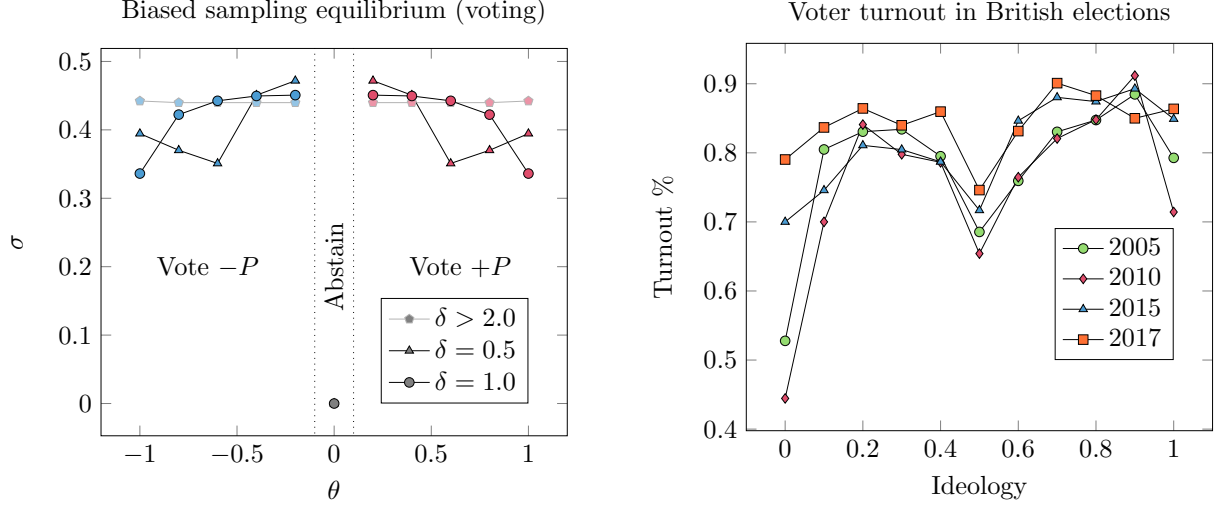


Figure 2: In a BSE with a large sampling window ( $\delta = 1.0$ ), the probability a type votes is inversely related to strength of preference. When the sampling window is small ( $\delta = 0.5$ ), the relationship is nonmonotone. With a large sampling window, biased sampling equilibrium qualitatively matches the gull-wing shape of turnout in British elections. *Simulation parameters:  $t = 5$ ,  $c = 0.125$ ; values below 0.001 have been censored.*

candidate  $p^*(\theta)$  and his disfavored candidate  $p^\times(\theta)$ . If he is pivotal he can sway the election to his favored candidate, and he will vote if

$$-|\theta - p^*(\theta)| - c \geq \frac{1}{2} ( -|\theta - p^*(\theta)| - |\theta - p^\times(\theta)| ).$$

I show in Appendix A.1 that this implies that the agent will vote if he is pivotal and  $|\theta| \geq c$ . Then

$$R((\tilde{a}_1, \tilde{a}_2); \theta) = \begin{cases} p^*(\theta) & \text{if } \tilde{a}_1 = \tilde{a}_2 = \emptyset, \\ p^*(\theta) & \text{if } \{\tilde{a}_1, \tilde{a}_2\} = \{-P, +P\}, \\ \emptyset & \text{otherwise.} \end{cases}$$

Note that it is never optimal for an agent to vote for his disfavored candidate,  $\sigma_{\theta p^\times(\theta)} = 0$  in any BSE.<sup>6</sup> Then to determine the probability with which an agent votes for a particular candidate, it is sufficient to know his type  $\theta$  and the probability  $\sigma_{\theta\emptyset}$  that he abstains from voting.

Define  $\pi_a(\theta) = \sum_{\theta' \in \Theta} H(\theta'; \theta) \sigma_{\theta' a}$ . Then in a BSE,

$$\sigma_{\theta p^*(\theta)} = \begin{cases} \pi_{\emptyset}(\theta)^2 + 2\pi_{p^*(\theta)}(\theta) \pi_{p^\times(\theta)}(\theta) & \text{if } |\theta| \geq c, \\ 0 & \text{otherwise;} \end{cases} \quad \sigma_{\theta p^\times(\theta)} = 0; \quad \text{and } \sigma_{\theta\emptyset} = 1 - \sigma_{\theta p^*(\theta)}. \quad (2)$$

Finding a BSE amounts to finding a fixed point of the system implied by (2) and the probabilities  $\pi_a(\cdot)$ . Figure 2 illustrates equilibrium under two sampling distributions.

<sup>6</sup>This is particular to the two candidate case, and is discussed in the four candidate analysis in Section 3.2.

It is natural to suspect that agents with stronger ideological preferences are more likely to vote than agents with weaker ideological preferences. However, Figure 2 shows that this need not be the case — neither in practice, nor in BSE. The intuition for nonmonotone turnout in BSE is especially straightforward in the case of  $\delta = 1$ . Consider the sample drawn by type  $\theta = 1$ . This sample contains only votes for  $+P$  and abstentions. The sample drawn by the adjacent type  $\theta' = 1 - 1/t$  also contains only votes for  $+P$  and abstentions, but these are observed from a different sample window. The sample drawn by type  $\theta'$  includes (potentially) type  $\tilde{\theta} = 0$ , which abstains with probability 1. Since neither  $\theta$  nor  $\theta'$  ever samples an individual voting for candidate  $-P$ , they are pivotal only when they observe two abstentions. Since type  $\theta'$  observes abstentions with higher probability, he is more likely to vote in spite of being less partisan.

When the sampling window is relatively large,  $\delta = 1.0$ , Figure 2 shows that BSE is not only nonmonotone, but nearly antitone. Starting from left-partisan preferences, abstention decreases and voting for  $-P$  increases. Strategies are monotone at the inflection points  $\theta \in \{-1/t, 0, 1/t\}$ , where abstention first increases, then decreases in favor of voting for  $+P$ . To the right of the inflection points, abstention increases and voting for  $+P$  decreases.

### 3.2 Four candidates

To assess the robustness of nonmonotone voter turnout, I consider a model with four candidates,  $p \in \{-X, -P, +P, +X\}$ ,  $0 < P < X$ , and extend  $\prec_A$  so that  $-X \prec_A -P$  and  $+P \prec_A +X$ . All other assumptions are retained from the two candidate analysis. Relatively extreme candidates appear attractive to relatively partisan individuals, and the possibility of electing a closely-aligned candidate might return a monotone relationship between partisanship and turnout, and the introduction of extreme candidates therefore serves as a robustness check on turnout nonmonotonicity.

Agents' responses are determined by believed pivotality, where the agent believes he is pivotal if and only if he samples two abstentions, or two votes for distinct candidates. Unlike the two candidate case, the four candidate case admits the possibility of strategic voting and not only a strategic decision of whether or not to vote. With two candidates, if an agent is pivotal he can swing the election to his favored candidate. With four candidates, it is possible that neither of the two sampled votes is for the agent's favored candidate. Because the agent believes his sample represents population play, he cannot vote for his favored candidate in hopes of generating a three-way tie, and he can affect the election only by voting for one of the two candidates voted for in his sample.

Of the two candidates appearing in his sample, the agent should vote for the one he prefers, or not vote at all. Whether or not the agent votes is no longer a question of pivotality alone, but also a question of the benefit of determining the outcome of the election, which depends in turn on the candidates appearing in his sample. As a clear example, if  $\theta = (P + X)/2$ , the agent is indifferent between the feasible candidates if  $\{\tilde{a}_1, \tilde{a}_2\} = \{+P, +X\}$ , but strictly prefers candidate

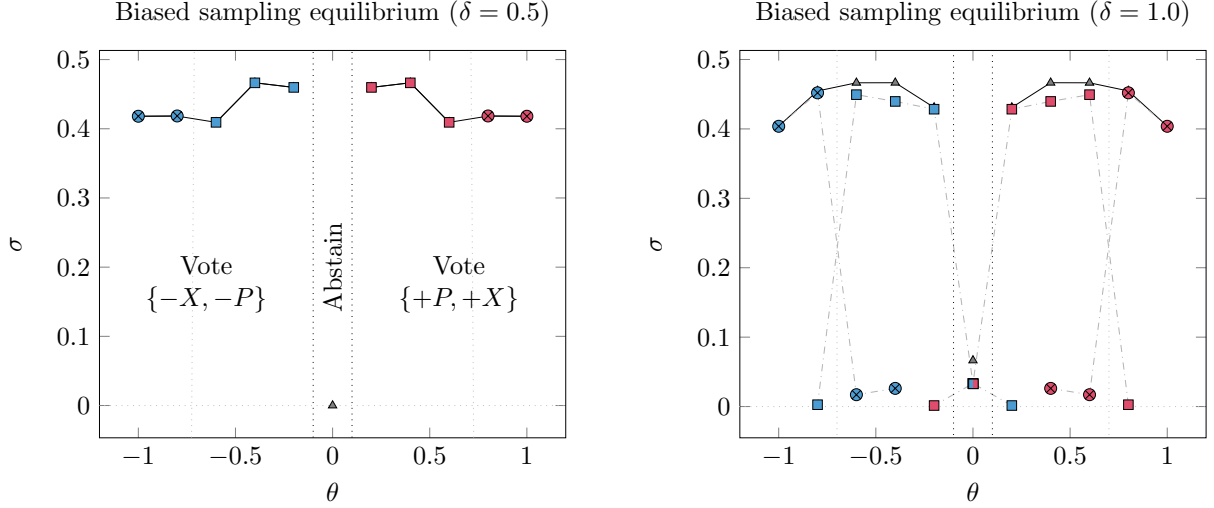


Figure 3: With four candidates and a relatively small sample window ( $\delta = 0.5$ ), agents only ever sample others who are voting for their most-preferred candidate, or one adjacent, or who abstain from voting. Pivotal voting is truthful, and turnout is equal to the proportion of agents voting for their most-preferred candidate. With a larger sample window ( $\delta = 1.0$ ), agents occasionally sample others who are voting for a strongly dispreferred candidate, and will occasionally vote for a less-favored candidate to win. Turnout is greater than the proportion of agents voting for their most preferred candidate. *Simulation parameters:  $t = 5$ ,  $P = 0.5$ ,  $X = 1.0$ ,  $c = 0.125$ ; values below 0.001 have been censored.*

$+X$  if  $\{\tilde{a}_1, \tilde{a}_2\} = \{-X, +X\}$ . The response function can then be represented as

$$R((\tilde{a}_1, \tilde{a}_2); \theta) = \begin{cases} p^*(\theta) & \text{if } \tilde{a}_1 = \tilde{a}_2 = \emptyset \text{ and } |\theta| \geq \underline{\tau}(c), \\ p^*(\theta; \tilde{a}_1, \tilde{a}_2) & \text{if } \emptyset \neq \tilde{a}_1 \neq \tilde{a}_2 \neq \emptyset \text{ and } \theta \in I(\tilde{a}_1, \tilde{a}_2; c), \\ \emptyset & \text{otherwise.} \end{cases}$$

Here,  $\underline{\tau}(c)$  is the minimum type necessary to participate in the election when no one else participates, and  $I(\tilde{a}_1, \tilde{a}_2; c)$  is the interval of types  $\theta$  that both prefer candidate  $p^*(\theta; \tilde{a}_1, \tilde{a}_2)$ , and find the distinction between  $\tilde{a}_1$  and  $\tilde{a}_2$  to be worth the cost of participating in the election.<sup>7</sup>

BSE in the four candidate voting model is depicted in Figure 3, in the cases of uniform sampling with windows  $\delta = 0.5$  and  $\delta = 1.0$ . The results are qualitatively similar to the two candidate case, but share some notable distinctions which illustrate the features of BSE. First, although turnout is relatively low for partisans, it is higher when there are four candidates than when there are two. When there are four candidates, even relatively partisan agents may believe they are pivotal not because they sample only abstention, but because they sample other agents who are voting for a more moderate candidate. The introduction of these new pivotal events increases turnout for most types.

Second, not only does strategic voting occur in equilibrium, but equilibrium turnout for extreme

<sup>7</sup>Detailed calculations for the response function are given in Appendix A.2.

candidates can increase as voters become less partisan. With a homophilic sampling process, a less partisan voter is more likely to sample from agents on the other side of the political spectrum. These voters occasionally believe that they are pivotal between an extreme candidate on their own side and a moderate candidate on the other side. Provided their views are not too moderate, they prefer the extreme candidate to the moderate candidate. Since less partisan voters are more likely to believe they are pivotal in this way, they may show more support for the extreme candidate than more partisan adjacent types. More moderate voters make the opposite calculation, that they would rather support a moderate on the opposite side than an extreme candidate on their own, and there is occasional voting across the political divide.

Third, the introduction of two extreme candidates depresses turnout for relatively moderate voters. Because turnout in all cases is relatively low, the highest probability event in which the agent is pivotal is when he samples two abstentions. For more partisan voters the introduction of new events indicating pivotality increases turnout, reducing the probability that a moderate agent samples two abstentions. The reduced probability of this event outweighs any effect of new events indicating pivotality, and turnout decreases for moderate voters.

## 4 Conclusion

I have defined the concept of *biased sampling equilibrium*, in which agents' strategies are consistent with a biased finite sample of opponent play. Unlike standard equilibrium concepts, where a isotone relationship between type and strategy is a focal target of analysis, games may admit no isotone biased sampling equilibrium. This nonmonotonicity can be quite strong, and equilibrium strategies can be strictly antitone away from well-defined inflection points. Applied to strategic voting, homophilic sampling may provide a novel explanation for low turnout of strong partisans. When agents are likely to obtain information from others like them, strong partisans are overly certain that their preferred candidate is going to win and that their vote is not pivotal. This induces a gull wing relationship between ideology and voter turnout, qualitatively matching data from recent British elections. Biased sampling equilibrium is likely to be of use in other situations where observed behaviors are not isotone with respect to a natural order on types.

## References

- S. Athey. Single crossing properties and the existence of pure strategy equilibria in games of incomplete information. *Econometrica*, 69(4):861–889, 2001.
- B. Geys. ‘Rational’ theories of voter turnout: a review. *Political Studies Review*, 4(1):16–35, 2006.
- S. Häfner. Stable biased sampling. *Games and Economic Behavior*, 107:109–122, 2018.
- Y. Halberstam and B. Knight. Homophily, group size, and the diffusion of political information in social networks: Evidence from twitter. *Journal of public economics*, 143:73–88, 2016.

- D. Knoke. Networks of political action: Toward theory construction. *Social Forces*, 68(4):1041–1063, 1990.
- V. Larcinese. Information acquisition, ideology and turnout: theory and evidence from Britain. *Journal of Theoretical politics*, 21(2):237–276, 2009.
- E. Lipnowski and E. Sadler. Peer-confirming equilibrium. *Econometrica*, 87(2):567–591, 2019.
- D. McAdams. Isotone equilibrium in games of incomplete information. *Econometrica*, 71(4):1191–1214, 2003.
- D. McAdams. On the failure of monotonicity in uniform-price auctions. *Journal of Economic Theory*, 137(1):729–732, 2007.
- M. McPherson, L. Smith-Lovin, and J. M. Cook. Birds of a feather: Homophily in social networks. *Annual review of sociology*, 27(1):415–444, 2001.
- M. J. Osborne and A. Rubinstein. Sampling equilibrium, with an application to strategic voting. *Games and Economic Behavior*, 45(2):434–441, 2003.
- D. Oyama, W. H. Sandholm, and O. Tercieux. Sampling best response dynamics and deterministic equilibrium selection. *Theoretical Economics*, 10(1):243–281, 2015.
- P. J. Reny. On the existence of monotone pure-strategy equilibria in Bayesian games. *Econometrica*, 79(2):499–553, 2011.
- Y. Salant and J. Cherry. Statistical inference in games. 2019.
- K. Woodward. Equilibrium with monotone actions. 2019.

## A Response functions for voting model

### A.1 Two candidate case

Suppose that the agent’s type is  $\theta$ , his favored candidate is  $p^*(\theta) = p^*$ , and his disfavored candidate is  $p^\times(\theta) = p^\times$ . Denote the agent’s sample draws by  $(\tilde{a}_1, \tilde{a}_2)$ . Consider the following cases:

- $\tilde{a}_1 = \tilde{a}_2 = \emptyset$ . Then the voter believes no one is voting, and he is pivotal. He should vote whenever

$$-|\theta - p^*| - c \geq \frac{1}{2}(-|\theta - p^*| - |\theta - p^\times|) \iff -|\theta - p^*| - 2c \geq -|\theta - p^\times|.$$

- $\tilde{a}_1 \neq \tilde{a}_2 = \emptyset$ . Then the voter believes half the population is voting for candidate  $\tilde{a}_1$ , and half is abstaining. He is not pivotal, and will not vote.

- $\{\tilde{a}_1, \tilde{a}_2\} = \{-P, P\}$ . Then the voter believes half the population is voting for each candidate, and he is pivotal. As in the case of total abstention, he should vote whenever

$$-|\theta - p^\star| - 2c \geq -|\theta - p^\times|.$$

- $\tilde{a}_1 = \tilde{a}_2 \neq \emptyset$ . Then the voter believes that the entire population is voting for the same candidate. He is not pivotal, and will not vote.

From these, it follows that the voter will vote (for his preferred candidate  $p^\star$ ) if and only if  $\tilde{a}_1 = \tilde{a}_2 = \emptyset$  or  $\{\tilde{a}_1, \tilde{a}_2\} = \{-P, P\}$  and  $-|\theta - p^\star| - 2c \geq -|\theta - p^\times|$ . This defines the reaction function  $R$  given in section 3.1.

For simulation, it is sufficient to identify a threshold  $\theta^\star$  that determines when a voter will abstain, regardless of her sampled actions. By symmetry, it is sufficient to assume that  $\theta > 0$ ,  $p^\star = P$ , and  $p^\times = -P$ . Note that if a voter of type  $\theta = P$  always prefers to abstain, then all types prefer to abstain: lower  $\theta > 0$  are further from  $P$  and closer to  $-P$ , and are less likely to value voting; higher  $\theta > 0$  are further from  $-P$ , but identically further from  $P$ . Then it is sufficient to assume that the threshold  $\theta^\star < P$ , otherwise no threshold exists and voting is never optimal.

$$-|\theta - P| - 2c \geq -|\theta - (-P)| \iff \theta - P - 2c \geq -\theta - P \iff \theta \geq c.$$

Then a voter will vote if her observed sample is as described above, and (by symmetry)  $|\theta| \geq c$ .

## A.2 Four candidate case

With four candidates and  $k = 2$  samples, there are 25 possible sample draws. The set of relevant samples is constrained by the observation (from the two candidate case) that a voter will vote only if his samples disagree with one another,  $\tilde{a}_1 \neq \tilde{a}_2, \emptyset$  and  $\tilde{a}_2 \neq \tilde{a}_1, \emptyset$ , or his samples indicate that the full population is abstaining,  $\tilde{a}_1 = \tilde{a}_2 = \emptyset$ .

As in the analysis of the two candidate case, by symmetry we may assume that  $\theta > 0$ . The voter's preferred candidate is  $+P$  if  $\theta \in (0, (P + X)/2]$ , and  $+X$  if  $\theta > (P + X)/2$ . Unlike in the two candidate case, it is now important to consider the exact value of  $P$ .

- $\tilde{a}_1 = \tilde{a}_2 = \emptyset$ . Then the voter believes the entire population is abstaining, and he is pivotal. His expected utility from not voting is

$$\begin{aligned} \theta < P : \quad & \frac{1}{4}(-|\theta - P| - |\theta + P| - |\theta - X| - |\theta + X|) = -\frac{1}{2}(P + X); \\ P \leq \theta < X : \quad & \frac{1}{4}(-|\theta - P| - |\theta + P| - |\theta - X| - |\theta + X|) = -\frac{1}{2}(\theta + X); \\ X \leq \theta : \quad & \frac{1}{4}(-|\theta - P| - |\theta + P| - |\theta - X| - |\theta + X|) = -\theta. \end{aligned}$$

Because the rest of the voting population is abstaining, if an agent votes he should vote for

his most-preferred candidate. Then

$$\begin{aligned}
\theta < P : \quad & -|\theta - P| - c \geq -\frac{1}{2}(P + X) \iff \theta \geq c + \frac{1}{2}(P - X); \\
P \leq \theta < \frac{1}{2}(P + X) : \quad & -|\theta - P| - c \geq -\frac{1}{2}(\theta + X) \iff 2(P - c) + X \geq \theta; \\
\frac{1}{2}(P + X) \leq \theta < X : \quad & -|\theta - X| - c \geq -\frac{1}{2}(\theta + X) \iff \theta \geq \frac{2}{3}c + \frac{1}{3}X; \\
X \leq \theta : \quad & -|\theta - X| - c \geq -\theta \iff X \geq c.
\end{aligned}$$

For the first condition to be (somewhere) consistent, it must be that  $c + (P - X)/2 < P$ , or  $c < (P + X)/2$ . For the second condition to be (somewhere) consistent, it must be that  $2(P - c) + X > (P + X)/2$ , or  $c < (3P + X)/4$ . For the third condition to be (somewhere) consistent, it must be that  $(2c + X)/3 \leq (P + X)/2$ , or  $c \leq (3P + X)/4$ . For the fourth condition to be (somewhere) consistent, it must be that  $c \leq X$ . Since  $X > P$ , the tightest condition is  $c < (P + X)/2$ .

Therefore, if  $c < (P + X)/2$  and the voter samples two abstentions, he will vote for his preferred candidate (by symmetry) whenever  $|\theta| \geq c + (P - X)/2$ .

- $\tilde{a}_1 \neq \tilde{a}_2$ , and  $\tilde{a}_1, \tilde{a}_2 \neq \emptyset$ . Then the voter is pivotal between candidates  $\tilde{a}_1$  and  $\tilde{a}_2$ ; if he votes, he will vote for the one which is closer to  $\theta$ . Denote this candidate by  $p^\star$ , and the other candidate by  $p^\times$ . The agent will vote if

$$-|\theta - p^\star| - c \geq \frac{1}{2}(-|\theta - p^\star| - |\theta - p^\times|) \iff -|\theta - p^\star| - 2c \geq -|\theta - p^\times|.$$

If  $\tilde{a}_1, \tilde{a}_2 = \pm P$ , the two candidate analysis applies, and the agent will vote if  $\theta \geq c$ . The same is true if  $\tilde{a}_1, \tilde{a}_2 = \pm X$ .

There are 12 ways that  $\tilde{a}_1$  can differ from  $\tilde{a}_2$ , neither being an abstention. 4 are handled above, leaving 8 further cases to analyze. It is more straightforward to complete the analysis in terms of which candidate an agent prefers, conditional on a drawn sample. Noting that an agent with type  $\theta > 0$  will never prefer candidate  $-X$ , the number of cases may be further reduced.

---

Now suppose that  $p^\star = P$  and  $p^\times = -X$  (recall that we are maintaining the assumption that  $\theta > 0$ ). The agent votes if

$$-|\theta - P| - 2c \geq -(\theta + X).$$

If  $\theta > P$ , this holds if  $P - 2c \geq -X$ , or if  $c < (P + X)/2$ . Since  $c < (P + X)/2$  by assumption, all  $\theta > P$  vote for  $P$  after observing this split sample. If  $\theta < P$ , the above inequality holds if  $\theta \geq c + (P - X)/2$ . Since  $c < (P + X)/2$  by assumption, the right-hand side is strictly below  $P$ . Then all  $\theta \in (2c + (P - X), P]$  vote for  $P$  after observing this split sample. It follows that all types  $\theta \geq c + (P - X)/2$  vote for  $P$  after observing this split sample.

---

Now suppose that  $p^* = -P$  and  $p^\times = X$ . The agent votes if

$$-|\theta + P| - 2c \geq -|\theta - X|.$$

Since  $p^* = -P$ , it must be that  $\theta < X$ , and the above becomes  $-\theta - P - 2c \geq \theta - X$ . In turn, this is  $\theta \leq (X - P)/2 - c$ .

---

Now suppose that  $p^* = X$  and  $p^\times = -P$ . The agent votes if

$$-|\theta - X| - 2c \geq -(\theta + P).$$

If  $\theta > X$ , this holds if  $c < (X + P)/2$ , which holds by assumption, and the agent votes after observing this split sample. If  $\theta < X$ , the above inequality holds if  $\theta \geq c + (X - P)/2$ . Note that  $p^* = X$  and  $p^\times = -P$  if and only if  $\theta \geq (X - P)/2$ , so any agent with type  $\theta > c + (X - P)/2$  is such that  $p^* = X$  and  $p^\times = -P$ .

---

Under the assumption that  $\theta > 0$ , it cannot be the case that  $p^* = -X$  and  $p^\times = P$ .

---

Now suppose that  $p^* = P$  and  $p^\times = X$ . The agent votes if

$$-|\theta - P| - 2c \geq -|\theta - X|.$$

Since  $p^\times = X$ , it must be that  $\theta < (P + X)/2 < X$ . Then  $-|\theta - X| = \theta - X$ . If  $\theta < P$ , the above inequality is  $\theta - P - 2c \geq \theta - X$ , which holds if and only if  $c \leq (X - P)/2$ . If  $\theta > P$ , the above inequality is  $P - \theta - 2c \geq \theta - X$ , which holds if and only if  $\theta \leq (X + P)/2 - c$ . Then the agent votes for  $p^*$  if  $\theta < P$  and  $c \leq (X - P)/2$ , or if  $\theta \in [P, (X + P)/2 - c]$ .

---

Now suppose that  $p^* = X$  and  $p^\times = P$ . The agent votes if

$$-|\theta - X| - 2c \geq -|\theta - P|.$$

Since  $p^* = X$ , it must be that  $\theta > P$ , so  $-|\theta - P| = P - \theta$ . If  $\theta < X$ , the above inequality is  $\theta - X - 2c \geq P - \theta$ , or  $\theta \geq (X + P)/2 + c$ . If  $\theta \geq X$ , the above inequality is  $X - \theta - 2c \geq P - \theta$ , or  $c \leq (X - P)/2$ . Then the agent votes for  $p^*$  if  $\theta \geq X$  and  $c \leq (X - P)/2$ , or if  $\theta \in [(X + P)/2 + c, X]$ .

---



Now suppose that  $p^* = -P$  and  $p^\times = -X$ . The agent votes if

$$-|\theta + P| - 2c \geq -|\theta + X|.$$

Since  $\theta \geq 0$ , this is  $-P - \theta - 2c \geq -\theta - X$ , or  $c \leq (X - P)/2$ .