Information Revelation in Constant-Sum Games: Elections and Beyond*

Navin Kartik† Francesco Squintani‡ Katrin Tinn§

December 19, 2025

(First version: March 2012)

Abstract

We study two-player constant-sum Bayesian games with type-independent payoffs. Under a "completeness" statistical condition, any "identifiable" equilibrium is an ex-post equilibrium. We apply this result to a Downsian election in which office-motivated candidates possess private information about policy consequences. The ex-post property implies a sharp bound on information aggregation: equilibrium voter welfare is at best equal to the efficient use of a single candidate's information. In canonical specifications, politicians may "anti-pander" (overreact to their information), whereas some degree of pandering would be socially beneficial. We discuss other applications of the ex-post result.

---

*Earlier versions of this paper were circulated under the title "Information Revelation and Pandering in Elections". For helpful comments and discussion, we thank Nageeb Ali, Scott Ashworth, Cole Wittbrodt, Kfir Eliaz, Andrea Galeotti, Christian Hellwig, Matias Iaryczower, Alessandro Lizzeri, Tianhao Liu, Massimo Morelli, David Rahman, Refine.ink, Andres Santos, Jesse Shapiro, Balázs Szentes, Olivier Tercieux, as well as several seminar and conference audiences. We are especially grateful to Afonso Gonçalves da Silva and Johannes Hörner. Kartik acknowledges, with gratitude, financial support from the Sloan Foundation, the National Science Foundation (Grant SES-115593), and the hospitality of and funding from the University of Chicago Booth School of Business during a portion of this research.

†Yale University, Department of Economics. Email: nkartik@gmail.com. The author was affiliated with Columbia University during most of the work on this paper.

‡University of Warwick, Department of Economics. Email: F.Squintani@warwick.ac.uk.

§McGill University, Desautels Faculty of Management. Email: katrin.tinn@mcgill.ca

1. Introduction

For representative democracy to be effective, voters must select representatives whose policies enhance their welfare. A challenge is that citizens are often poorly informed on policy issues, as posited by Downs (1957) in his "rational ignorance" hypothesis and since supported by numerous studies starting with Campbell, Converse, Miller and Stokes (1960). Political candidates, by contrast, devote substantial resources and have broad access to policy experts and think tanks. Politicians can convey their information to the electorate through their electoral campaigns, and in particular, through their policy positions. Indeed, there is evidence that voters learn and/or refine their views during elections.[1] But when office-seeking politicians choose their positions strategically, how effectively do elections aggregate their information?

One prevalent view is that elections function well because even office-seeking politicians are impelled to choose policies that promote voters' interests. Indeed, Wittman (1989, p. 1400) influentially argued that political competition benefits the electorate because "there are returns to an informed political entrepreneur from providing the information to the voters, winning office, and gaining the [...] rewards of holding office." Concurrently, however, there are also concerns—raised both in popular circles and in academic work that we discuss subsequently—that competitive pressures drive politicians to pander to voters' opinions rather than provide valuable information. After all, the argument goes, it is hard to win an election by campaigning on policies with recondite merits; a politician is better off simply promising to do whatever voters believe is best from the outset. Pandering is viewed as inefficient because it would lead to policies that are excessively distorted toward the voters' less-informed opinions.

Our paper (re-)assesses the efficiency of elections when office-seeking politicians possess private information about policy consequences. Section 3 lays out an extension of the canonical Downsian model of elections (Downs, 1957; Hotelling, 1929). Our framework is quite general, but we maintain the Downsian assumption of two candidates making policy commitments to

[1] Le Pennec and Pons (2023) provide large-scale cross-country evidence that a substantial share of voters decide late in campaigns, consistent with meaningful voter learning. Earlier work includes experiments on deliberative polling (Fishkin, 1997), studies on the effects of information on voters' opinions (Zaller, 1992; Althaus, 1998; Gilens, 2001), work on framing in polls (Schuman and Presser, 1981), and experiments on priming (Iyengar and Kinder, 1987).

maximize their probability of winning the election. The key twist is that each politician has (imperfect) private information about policy consequences. In other words, they each have information about which policy would be best for a representative or median voter—hereafter, "the voter."

We show that Downsian elections are fundamentally limited in their ability to aggregate candidates' private information. Under reasonable conditions, Proposition 1 deduces that in any equilibrium, holding fixed the politicians' equilibrium strategies, voter welfare—ex-ante expected utility—equals the welfare from (hypothetically) electing the same candidate regardless of platforms. That implies a tight upper bound: equilibrium voter welfare is no higher than what can be obtained based on one politician's information alone (Theorem 2).

The key to establishing those conclusions is a general property of a class of constant-sum Bayesian games, studied in Section 2. Specifically, consider any two-player constant-sum Bayesian game with type-independent payoffs.[2] Our main theoretical result, Theorem 1, says that under a general completeness condition on the distribution of types, any equilibrium in which at least one player uses an "identifiable" strategy—for example, a pure strategy—must be an ex-post equilibrium. That is, even after observing the opponent's action, a player must be indifferent among all his own on-path actions. This ex-post property substantially limits the scope for how much private information a player's actions can reveal in purely competitive settings, such as a Downsian election.

To understand politicians' strategic incentives in more detail, we turn in Subsection 3.2 to a canonical one-dimensional normal-quadratic specification of the Downsian election. We assume there that the best policy for the voter—the "state" of the world—is drawn from a normal distribution; each candidate's private signal is the true state plus noise that is also normally distributed; and the voter's payoff is a quadratic loss function of the distance between the chosen policy and the state.

For that specification, we explain why it is not an equilibrium for each politician to propose a policy that is best for the voter based on his own information, i.e., to use an "unbiased

[2] In the electoral context, the two players are the politicians; office motivation (each candidate is maximizing their probability of winning) means that no matter the voter's strategy, the politicians are engaged in a constant-sum game. Their private information, while certainly relevant to the voter, does not directly affect the politicians' payoffs.

strategy" (in which case the election would aggregate more than one politician's information). We show that, perhaps contrary to intuition, politicians would have an incentive to deviate by "anti-pandering"—*overreacting* to their private information—as the rational voter would elect the more extreme politician under unbiased strategies. The voter would do so because each politician's estimate of the state based on his own signal places more weight on the prior than the voter's estimate after learning both politicians' signals.[3]

Building on the above logic, we identify in Proposition 2 a symmetric equilibrium that features anti-pandering by both politicians. In this equilibrium, politicians choose different platforms with probability one, yet—regardless of their platforms—are elected with equal probability. Although there are other equilibria, the anti-pandering equilibrium shows starkly that office motivation need not induce pandering (or underreaction to private information). Further, in terms of welfare, we show in Proposition 3 that a suitable degree of disequilibrium pandering would actually benefit the voter, contrary to perceptions that pandering is always harmful.

Although our main economic application concerns Downsian elections—or any equivalent setting in which two agents compete for their proposals to be selected by a decision-maker—the abstract ex-postness result of Theorem 1 has broader relevance. In Section 4 we develop another application, in which two firms compete for both private market share and a government action (e.g., procurement). The firms have private information about a fundamental that matters for the government's optimal allocation, for example the relative social value of their products. Although the government might benefit from learning about the fundamental from the firms' choices, we show that the ex-post property implied by our key statistical condition often precludes such benefit. We also discuss how the framework in Section 4 has broader applicability.

Related Literature. Our Theorem 1 and its ingredient Lemma 1 relate to work on the equilibrium properties of two-player constant-sum games (hereafter just "constant-sum games"), which dates back to Von Neumann (1928). Our results go beyond equilibrium payoff uniqueness or interchangeability (Nash, 1951), by establishing, under some conditions, an ex-post

[3] Glaeser and Sunstein (2009) and Roux and Sobel (2015) also identify this implication of Bayesian updating in a non-strategic group decision-making context.

property for a class of constant-sum games: Bayesian games with type-independent payoffs. Ex-postness is not a general equilibrium property of constant-sum games; simply consider "matching pennies". While the class of games Lemma 1 or Theorem 1 apply to is (very) restricted, we demonstrate relevant economic applications. There are two papers we are aware of with closely-related results.[4] First, Viossat (2006, Proposition 3.8) derives certain properties of correlated equilibria of complete-information constant-sum games, which, as we explain after Lemma 1, is connected to our lemma. He does not have an analog to Theorem 1.

Second, Kattwinkel, Niemeyer, Preusser and Winter (2022) study mechanisms without transfers when two agents, each with a finite set of types, have type-independent and opposing preferences over a binary allocation. Their Proposition 1 characterizes incentive compatibility of direct mechanisms, and their Proposition 3 (part 2) shows that under a full-rank condition, only constant mechanisms are incentive compatible. For finite type sets, these results are related to our Lemma 1 and Theorem 1, as elaborated in Section 2. We study more general games rather than just direct mechanisms; moreover, a treatment of infinite type sets is valuable for applications, including our main electoral one, whose leading specification (Subsection 3.2) has normally distributed types.

With regard to our main application, there is a small prior literature on electoral competition when candidates have policy relevant private information.[5] Heidhues and Lagerlof (2003) illustrate why candidates may have an incentive to pander to the electorate's prior belief; their setting is one with binary policies, binary states, and binary signals. We find that in our richer setting, the opposite may be true for a broad class of information structures. Plainly, with binary policies, one cannot see the logic of why and how candidates may wish to overreact to private information. Loertscher (2012) maintains the binary signal and state structure, but introduces a continuum policy space. His results are more nuanced, but at least when signals

[4] As part of their study of robust implementation, Pei and Strulovici (2025, Theorem 4, part 1) observe that if two agents' payoffs are state-independent (but not necessarily constant sum), then any equilibrium of the game with no information about the state can also be supported as an equilibrium when agents observe the state. The reason is that the state is merely a correlating device. However, this observation only addresses some of the equilibria when agents observe the state (which is sufficient for their purposes), and moreover, these equilibria need not be ex-post equilibria. Our Theorem 1's conclusion of ex-postness of all (identifiable) equilibria owes to its constant-sum and statistical-completeness assumptions.

[5] There are also models in which candidates have private information that is not policy relevant for voters, e.g., about the location of the median voter (Ottaviani and Sorensen, 2006; Bernhardt, Duggan and Squintani, 2007, 2009).

are sufficiently precise, the conclusions are similar to those of Heidhues and Lagerlof (2003).[6]

Laslier and Van de Straeten (2004) show that if voters in the Heidhues and Lagerlof (2003) model are endowed with sufficiently precise private information about the policy-relevant state, then there are equilibria in which candidates fully reveal their private information; see also Klumpp (2014) and Gratton (2014). By contrast, we are interested in settings in which there is little information voters have that candidates do not.

The anti-pandering equilibrium of our normal-quadratic model specification provides a new perspective on the classic issue of policy divergence. Unlike some other prevalent explanations (e.g., ideologically-motivated candidates with uncertainty about voter preferences, as in Wittman (1983) and Calvert (1985)), anti-pandering features office-motivated politicians diverging in order to maximize support from a risk-averse voter whose ideology is known.[7]

Building on earlier versions of the current paper, Millner, Ollivier and Simon (2020) introduce confirmation bias for voters in a continuum-policy ternary-state model. They find that confirmation bias can reduce equilibrium anti-pandering.

Schultz (1996) studies a model in which two candidates are perfectly informed about the policy-relevant state but are policy motivated. He finds that when the candidates' ideological preferences are sufficiently extreme, platforms cannot reveal the true state; however, because of the perfect information assumption, full revelation can be sustained when ideological preferences are not too extreme. Martinelli (2001) and Martinelli and Matsui (2002) derive further results with ideologically motivated candidates who are perfectly informed about a policy-relevant variable.

Ambrus, Baranovskyi and Kolb (2021) study a model related to our normal-quadratic specification, but with candidates who are policy motivated. We explain in Subsection 3.3 that our welfare result continues to hold, approximately, when the extent of policy motivation

[6] In the Supplementary Appendix, we show how overreaction or anti-pandering arises in a binary-signal model specification when the policies and the state lie in the unit interval. That specification permits a closer comparison with Heidhues and Lagerlof (2003) and Loertscher (2012).

[7] Other explanations for divergence include those based on increasing turnout (Glaeser, Ponzetto and Shapiro, 2005), campaign contributions (Campante, 2011), valence asymmetries (Groseclose, 2001; Aragones and Palfrey, 2002), signaling character, competence, or related mechanisms (Callander and Wilkie, 2007; Kartik and McAfee, 2007; Honryo, 2018), or more than two candidates (Palfrey, 1984).

is small. Ambrus et al. (2021) show that when policy motivation looms large and candidates' ideologies are sufficiently similar to the voter's, equilibria can aggregate more information. Our papers are complementary.

There are various other settings in economics and political science in which distortions arise because agents wish to influence their principals' beliefs. In particular, electoral models often feature a single politician seeking to build a reputation for either competence (e.g., Canes-Wrone, Herron and Shotts, 2001) or aligned preferences (e.g., Maskin and Tirole, 2004). While most such papers highlight the possibility of pandering—or even "over-pandering" as in Acemoglu, Egorov and Sonin (2013) and Kartik and Van Weelden (2019)—anti-pandering arises in Prendergast and Stole (1996), Levy (2004), and Bils (2023).

Finally, we note that our applications illustrate that ex-post equilibrium or incentive constraints leave limited scope for information revelation among two purely competitive players. Although in a very different setting, that is reminiscent of negative results like Jehiel, Meyer-ter Vehn, Moldovanu and Zame (2006), who show that generically only constant mechanisms are ex-post implementable under interdependent values and multidimensional signals.

2. Two-Player Constant-Sum Bayesian Games

Our applications are underpinned by a general result on two-player constant-sum Bayesian games with type-independent payoffs. This section develops that result.

Setting. There are two players, A and B. Each player $i \in \{A, B\}$ has a private type $s_i \in S_i$, where S_i is a nonempty standard Borel space. The type profile (s_A, s_B) is drawn from a common-prior probability measure F on $S_A \times S_B$, whose marginals F_A and F_B have supports S_A and S_B, respectively. Writing $-i$ for the player different from i as usual, we will denote by $F(\cdot \mid s_i)$ the regular conditional distribution of s_{-i} given s_i.[8]

After learning their types, players choose actions simultaneously; player i's action is denoted $x_i \in X_i$, where each X_i is a nonempty standard Borel space.[9] Player i's (von-Neumann–

[8] This exists and is unique almost everywhere (a.e., hereafter) because $S_A \times S_B$ is a standard Borel space (Durrett, 1995, pp. 229–230).

[9] As usual, each x_i can also be interpreted as player i's action plan in a sequential-move game.

Morgenstern) payoff is $u_i(x_i, x_{-i})$, with $u_A(\cdot) + u_B(\cdot) = 0$. So the game is constant sum and types do not directly affect payoffs. (Nevertheless, the players' types may affect the payoffs of third parties, as in our subsequent applications.) Assume payoffs are uniformly bounded: $|u_i(\cdot)| \leq K$, for some constant K. We denote each player i's space of mixed actions—randomizations over actions—by $\Delta(X_i)$, with generic element ξ_i, and extend payoffs to $(\Delta(X_A), \Delta(X_B))$ by linearity as usual, writing $u_i(\xi_i, \xi_{-i})$. A mixed strategy for player i is $\sigma_i : S_i \to \Delta(X_i)$. We study Bayes-Nash equilibria.

Identifiability conditions. Theorem 1 below requires the following statistical condition on the distribution of signals. As is common, we use notation like $F(g(s_i) = 0)$ as shorthand for $F(\{s_i : g(s_i) = 0\})$.

Condition 1. *For any $i \in \{A, B\}$ and any bounded measurable function $g : S_{-i} \to \mathbb{R}$, it holds that*

$$\mathbb{E}_{s_{-i}}[g(s_{-i}) \mid s_i] = 0 \text{ for } F_i\text{-a.e. } s_i \in S_i \implies F(g(s_{-i}) = 0 \mid s_i) = 1 \text{ for } F_i\text{-a.e. } s_i \in S_i. \quad (1)$$

In words, (1) says that if $g(s_{-i})$ has mean zero conditional on i's signal, then it must in fact equal zero almost surely (a.s., hereafter) conditional on i's signal. Since s_i is itself a random variable, conditional expectations and distributions are defined only up to F_i-null sets, and we therefore formulate (1) in an s_i-a.s. sense. With that caveat in mind, Condition 1 is equivalent to requiring that the family of conditional distributions $\{F(\cdot \mid s_i)\}_{s_i \in S_i}$ is *boundedly complete.*

Bounded completeness is a recognized concept in statistics (e.g., Lehmann, 1986, p. 144), which in our context captures a notion of richness in how variation in s_i affects player i's beliefs about his opponent's type s_{-i}. Specifically, the variation in s_i must identify or distinguish any nontrival bounded "feature" of the opponent's type. For finite type spaces, Condition 1 is equivalent to the matrix of joint probabilities of types s_A and s_B having full row and column rank, since bounded completeness rules out nontrivial linear relations among each player's conditional distributions. Crémer and McLean (1985, 1988) have prominently used a linear independence notion of richness in the context of full surplus extraction in mechanism design with finite type spaces. Appendix B discusses the connection of completeness with linear independence in infinite type spaces.

Plainly, if either $|S_A| > 1$ or $|S_B| > 1$, then Condition 1 is violated if the types are independent. But we are interested in settings in which each player's type is informative about the other's, stemming from both being informative about some underlying "state of the world". In those contexts, we view Condition 1 as a reasonable requirement. In particular, it follows from a well-known fact about complete families that if each type space $S_i \subset \mathbb{R}^n$ has a nonempty interior, Condition 1 holds when the distribution of $s_{-i} \mid s_i$ for each $i \in \{A, B\}$ is in a regular exponential family of distributions with continuous cumulant function.[10] This canonical class includes a variety of widely-used discrete and continuous distributions with bounded and unbounded supports, such as normal, exponential, gamma, beta, chi-squared, binomial, Dirichlet, and Poisson.

Theorem 1 also requires the following analogy to Condition 1 on players' strategies.

Definition 1. Strategy σ_i is *identifiable* if for any bounded and measurable $g : X_i \to \mathbb{R}$, it holds that

$$\mathbb{E}_{x_i}\left[g\left(x_i\right) \mid s_i\right] = 0 \text{ for } F_i\text{-a.e. } s_i \in S_i \implies \Pr\left(g\left(x_i\right) = 0 \mid s_i\right) = 1 \text{ for } F_i\text{-a.e. } s_i \in S_i,$$

where the left-hand-side expectation and right-hand-side probability are computed using σ_i.

Importantly for our applications, any pure strategy $\sigma_i : S_i \to X_i$ is identifiable because in that case $\mathbb{E}_{x_i}[g(x_i) \mid s_i] = g(\sigma_i(s_i))$. An example of a non-identifiable strategy is any non-pure strategy that does not vary with the player's signal.

The result. Since the game has type-independent payoffs, standard arguments for constant-sum games imply that all equilibria yield the same payoff vector (U_A^*, U_B^*). We say that an

[10] This follows from the well-known property that regular exponential families are boundedly complete when the natural parameter ranges over a set with nonempty interior (e.g., Lehmann, 1986, Theorem 1, p. 142). In particular, when each player i's type space $S_i \subset \mathbb{R}^n$ has nonempty interior and the conditional distribution of $s_{-i} \mid s_i$ admits a density of the form

$$f(s_{-i} \mid s_i) = \exp\bigl(s_i \cdot T(s_{-i}) - \psi(s_i)\bigr)\, h(s_{-i}), \tag{2}$$

and the cumulant function ψ is continuous—as is the case for familiar exponential-family distributions—Condition 1 is satisfied. For, in such families, the map $s_i \mapsto \mathbb{E}\left[g\left(s_{-i}\right) \mid s_i\right]$ is continuous for a bounded function g, and so $\mathbb{E}[g(s_{-i}) \mid s_i]$ vanishing on F_i-a.e. s_i implies it vanishes throughout the interior of S_i. hence $g(s_{-i}) = 0$ $F(\cdot \mid s_i)$-a.s. for every interior s_i, which yields (1) because in a regular exponential family the interior s_i have full F_i-measure.

equilibrium $\sigma^* := (\sigma_A^*, \sigma_B^*)$ is an *ex-post equilibrium* if $u_i(x_i, x_{-i}) = U_i^*$ for each $i \in \{A, B\}$ and for (σ^*, F)-a.e. pair (x_A, x_B). In other words, in an ex-post equilibrium, it holds that no type of either player would have an incentive to change their action even after learning the action played by the opponent.[11] We also say that an equilibrium σ^* is an *identifiable equilibrium* if either σ_A^* or σ_B^* is identifiable. In particular, any equilibrium in which at least one player is playing a pure strategy is identifiable.

Theorem 1. *If Condition 1 holds, then any identifiable equilibrium σ^* is an ex-post equilibrium.*

The theorem says that, subject to Condition 1, in any identifiable equilibrium the players are indifferent over all the action profiles that are played in equilibrium—even though different types of a player may be playing different (distributions over) actions and hold different beliefs about the opponent's type.

The theorem's proof requires the following lemma, which is of independent interest as it does not rely on either Condition 1 or identifiability of strategies. The lemma says that in any equilibrium, all types of a player obtain the same interim expected payoff (which is independent of the equilibrium), and any type would obtain that interim payoff regardless of which action it plays among all the actions taken by some type of that player.

Lemma 1. *Let (U_A^*, U_B^*) denote the payoffs in every equilibrium, and let σ^* be some equilibrium. Then for each $i \in \{A, B\}$, for (σ_i^*, F_i)-a.e. x_i and F_i-a.e. s_i, it holds that*

$$U_i^* = \mathbb{E}_{x_{-i}}[u_i(x_i, x_{-i}) \mid s_i],$$

where the expectation is taken with respect to the measure induced by σ_{-i}^* and F.

Proof. Since the game is constant sum with type-independent payoffs, and there is some equilibrium σ^*, it holds that each player i can secure U_i^*.[12] This implies that for each i,

[11] To be precise, this is up to a probability-zero caveat; to lighten the exposition, we frequently omit such caveats hereafter outside of formal statements.

[12] A mixed action $\xi_i \in \Delta(X_i)$ *secures* player i the payoff $U_i \in \mathbb{R}$ if for all $\xi_{-i} \in \Delta(X_{-i})$, it holds that $u_i(\xi_i, \xi_{-i}) \geq U_i$. The statement follows from standard logic for constant-sum games, which we detail for completeness. For either player i, consider the mixed action ξ_i^* defined as the ex-ante measure over X_i

conditional on F_i-a.e. types s_i, player i's interim equilibrium payoff must in fact equal U_i^*. We now make two observations: first, for F_i-a.e. types s_i, the conditional distribution of the opponent's actions induced by σ_{-i}^* and F must secure the opponent U_{-i}^*; for if not, there would be some action that yields s_i a payoff strictly larger than U_i^*. Second, (σ_i^*, F_i)-a.e. actions x_i must be a best response to any mixed action ξ_{-i} that secures U_{-i}^*; for if not, player $-i$ can obtain a payoff strictly larger than U_{-i}^* by playing the constant strategy $s_{-i} \mapsto \xi_{-i}$. Hence, for (σ_i^*, F_i)-a.e. x_i and F_i-a.e. s_i, it follows that $\mathbb{E}_{x_{-i}}[u_i(x_i, x_{-i}) \mid s_i] = U_i^*$; the expectation cannot be larger by the first observation earlier, nor smaller by the second. □

One way to appreciate the content of Lemma 1 is via its connection to *correlated* equilibrium of *complete-information* games. Consider a complete-information two-player constant-sum game G with action spaces X_A and X_B and payoff functions u_A and u_B as above. Any (objective) correlated equilibrium of this game is a Bayes-Nash equilibrium of our Bayesian game with a suitably-defined information structure; conversely, any Bayes-Nash equilibrium of our game is a correlated equilibrium of G. It follows from Lemma 1 that if $\rho \in \Delta(X_1 \times X_2)$ is a correlated equilibrium of G with payoffs (π_1, π_2), then for any $i \in \{A, B\}$ and ρ-a.e. x_i, x_i', it holds that $\mathbb{E}_{x_{-i}}[u_i(x_i', x_{-i}) \mid x_i] = \pi_i$, and hence x_i' is a best response to $\rho(\cdot \mid x_i)$. For finite games, this fact has been noted by Viossat (2006, Proposition 3.8).[13]

Lemma 1 also relates to Kattwinkel et al. (2022, Proposition 1). Our result is stronger for two reasons. First, we allow for infinite type sets. Second, our result applies to arbitrary action spaces and equilibrium strategies, not just direct mechanisms and truthful equilibria. When types may mix over their actions, Lemma 1 establishes that each type is indifferent among all the actions in any type's equilibrium mixture—not merely indifferent among all types' distributions. This is crucial for the proof of Theorem 1, which we now turn to.

Proof of Theorem 1. Let σ^* be an equilibrium in which player $-i$'s strategy is identifiable. Below, all expectations are with respect to the measure induced by (σ^*, F). Fix x_i in a

induced by the equilibrium strategy σ_i^*. If ξ_i^* does not secure the payoff U_i^*, then there is some mixed action ξ_{-i} such that $u_i(\xi_i^*, \xi_{-i}) < U_i^*$, or equivalently by the constant-sum property, $u_{-i}(\xi_{-i}, \xi_i^*) > U_{-i}^*$. But then playing the constant strategy $\sigma_{-i}(s_{-i}) = \xi_{-i}$ would be a profitable deviation for player $-i$ against σ_i^*.

[13] Moreover, because of the "conversely" point noted earlier in the paragraph, if we restricted to finite types and actions, Viossat's (2006) result could in turn be used to prove Lemma 1. Indeed, that indirect approach was used in earlier versions of our paper (Kartik, Squintani and Tinn, 2015, Appendix A) and by Kattwinkel et al. (2022, Lemma 1 and Proposition 1).

(σ_i^*, F_i)-full-measure subset of X_i. For F_i-a.e. s_i, we have

$$
\begin{aligned}
U_i^* &= \mathbb{E}_{x_{-i}}[u_i(x_i, x_{-i}) \mid s_i] \quad \text{by Lemma 1} \\
&= \mathbb{E}_{s_{-i}}\big[\mathbb{E}_{x_{-i}}[u_i(x_i, x_{-i}) \mid s_i, s_{-i}] \mid s_i\big] \quad \text{by the law of iterated expectation} \\
&= \mathbb{E}_{s_{-i}}\big[\mathbb{E}_{x_{-i}}[u_i(x_i, x_{-i}) \mid s_{-i}] \mid s_i\big] \quad \text{because } x_{-i} \text{ is independent of } s_i\text{, conditional on } s_{-i}.
\end{aligned}
$$

Now, applying Condition 1 (just for one player, i) with $g(s_{-i}) = \mathbb{E}_{x_{-i}}[u_i(x_i, x_{-i}) \mid s_{-i}] - U_i^*$, we get for F_i-a.e. s_i that

$$
\mathbb{E}_{x_{-i}}[u_i(x_i, x_{-i}) \mid s_{-i}] = U_i^* \quad \text{for } F(\cdot \mid s_i)\text{-a.s. } s_{-i}.
$$

Integrating over s_i with respect to F_i and using the law of total probability yields the above equality for F_{-i}-a.e. s_{-i}.

It then follows from the identifiability of σ_{-i}^*, applied with $g(x_{-i}) = u_i(x_i, x_{-i}) - U_i^*$, that for F_{-i}-a.e. s_{-i} we have

$$
u_i(x_i, x_{-i}) = U_i^* \quad \text{for } \sigma_{-i}^*(\cdot \mid s_{-i})\text{-a.s. } x_{-i}.
$$

Integrating over s_{-i} with respect to F_{-i} yields the above equality for $\big(\sigma_{-i}^*, F_{-i}\big)$-a.e. (x_{-i}, s_{-i}).

Since x_i was arbitrary in a (σ_i^*, F_i)-full-measure set, the result follows. □

The following two examples show that neither Condition 1 nor the qualification of identifiability can be dropped from Theorem 1. Both examples are based on a "matching pennies" payoff structure.

Example 1. Consider $X_A = X_B = S_A = S_B = \{0, 1\}$, $u_A(x_A, x_B) = \mathbb{1}\{x_A = x_B\}$, and F the uniform distribution. Condition 1 fails because s_A and s_B are independent. There is an identifiable equilibrium in which each player i plays the pure strategy $s_i \mapsto s_i$. This is, however, not an ex-post equilibrium. □

Example 2. Now consider a complete-information variant of the previous example. There are singleton type sets, $|S_A| = |S_B| = 1$; trivially, Condition 1 holds. Actions and payoffs are as in Example 1. The equilibrium in which both players uniformly randomize over their two actions is not an ex-post equilibrium; these mixed strategies are not identifiable, and hence the equilibrium is not identifiable. □

Note that Lemma 1 applies to both examples. In particular, in the equilibrium of Example 1, neither type of a player is playing an action that secures the equilibrium payoff; nevertheless, each type has the same interim payoff and is indifferent between its equilibrium action and the equilibrium action of the other type.

We close this section by fleshing out one implication of Theorem 1 that is useful for applications. Let Ω be some set of outcomes and $w : X_A \times X_B \to \Omega$ an outcome function. Assume that preferences depend on only the outcome, i.e., there is some $\tilde{u}_i : \Omega \to \mathbb{R}$ such that $u_i(x_A, x_B) = \tilde{u}_i(w(x_A, x_B))$ for all (x_A, x_B). Say that there are *strict preferences over outcomes* if each $\tilde{u}_i$ is injective, and say that an equilibrium σ^* has a *single outcome* ω^* if $w(x_A, x_B) = \omega^*$ for (σ^*, F)-a.e. (x_A, x_B). If there are strict preferences over outcomes, then plainly an ex-post equilibrium must have a single outcome. Hence, the following result follows directly from Theorem 1.

Corollary 1. *Assume strict preferences over outcomes. If Condition 1 holds, then any identifiable equilibrium has a single outcome.*

In fact, since the corollary holds for an arbitrary outcome space and function, it is equivalent to Theorem 1. For, if we define outcomes as the utility-equivalence classes of action profiles and the outcome function as mapping each action profile into its equivalence class, then we have strict preferences over outcomes by construction, and a single outcome corresponds exactly to ex-postness.

As detailed in in Appendix C, Corollary 1 implies Kattwinkel et al.'s (2022) Proposition 3.2. In a direct-mechanism setting with binary allocations that two agents with a finite number of types have opposed preferences over, those authors show that incentive compatibility requires a constant allocation probability, so long as the type distribution has full rank.

In the remainder of the paper, we use Theorem 1/Corollary 1 to study information revelation and aggregation in some economic settings.

3. Information Aggregation and Pandering in Elections

This section studies a model of a Downsian election with informed candidates.

Model. We consider an electorate that is represented in reduced-form by a single voter. The voter's preferences depend upon the implemented policy $x \in X$ and an unknown state of the world $\theta \in \Theta$, where both X and Θ are standard Borel spaces (e.g., subsets of Euclidean spaces). The voter can be interpreted as either some representative of a group, or a median voter if one exists—as is assured if, for example, $X \subset \mathbb{R}$ and voters have single-crossing expectational-differences preferences (Kartik, Lee and Rappoport, 2023)—or even just a single decision-maker.

The state is drawn from a probability measure F_θ. The voter's preferences are represented by a von-Neumann–Morgenstern utility function $u : X \times \Theta \to \mathbb{R}$. We assume there is a utility-maximizing policy in each state and an expected-utility-maximizing policy under the prior. A leading example that we will return to is the (one-dimensional) quadratic loss function: $X, \Theta \subset \mathbb{R}$ and $u(x, \theta) = -(x - \theta)^2$, with F_θ having finite expectation.

There are two electoral candidates, A and B. Given the state θ, each candidate $i \in \{A, B\}$ privately observes a signal $s_i \in S_i$, where S_i is a closed subset of $\mathbb{R}^n$, with $n \geq 1$. The joint conditional cumulative distribution of $(s_A, s_B) \in S_A \times S_B$ is denoted by $F_{s_A,s_B|\theta}$, and the conditional marginal for each candidate i by $F_{s_i|\theta}$. The measure F_θ and distributions $F_{s_A,s_B|\theta}$ induce a joint cumulative distribution F_{s_A,s_B} of signal profiles unconditional on the state, with marginals F_{s_A} and F_{s_B}; we assume that for each candidate i, the support of F_{s_i} is S_i. We assume that for either candidate $i \in \{A, B\}$ and any signal $s_i \in S_i$, there is a voter-optimal policy: $\max_{x \in X} \mathbb{E}_\theta[u(x, \theta) \mid s_i]$ exists.

Our results will require that the joint cumulative distribution F_{s_A,s_B} satisfy Condition 1. This is a reasonable requirement when both signals s_A and s_B are informative about the state; in particular, following the discussion after Condition 1, typical exponential families of signal distributions satisfy the requirement. A leading example that we will return to is the (one-dimensional) *normal-normal* structure: $\Theta = S_A = S_B = \mathbb{R}$, $\theta \sim \mathcal{N}(0, 1/\alpha)$, i.e., the state is normally distributed with mean 0 and precision $\alpha \in \mathbb{R}_{>0}$, and conditional on the state θ, each candidate i's signal is drawn independently from the normal distribution $\mathcal{N}(\theta, 1/\beta_i)$ with precision parameter $\beta_i \in \mathbb{R}_{>0}$.[14] The case of $\beta_A \neq \beta_B$ captures candidates having access to information of different quality. In general, F_{s_A,s_B} can satisfy Condition 1 even with signals being positively (or negatively) correlated conditional on the state.

After privately observing their signals, the candidates simultaneously choose their platforms x_A and x_B from the policy space X, with the objective of maximizing their respective probabilities of winning the election.[15] Upon observing the platforms (x_A, x_B), the voter updates her belief about the state θ and then elects the candidate whose platform provides the highest expected utility. The elected candidate $i \in \{A, B\}$ implements his platform x_i. Platforms are thus policy commitments in the Downsian tradition. Candidates are expected utility maximizers, with the elected candidate obtaining a utility of 1 and the other candidate 0. Hence, candidates are purely office motivated. All aspects of the model except the candidates' private signals are common knowledge.

Strategies, Equilibria, and Welfare. A pure strategy for a candidate i is a measurable function $y_i : S_i \to X$, with $y_i(s_i)$ the platform chosen by i when his signal is s_i. A strategy for the voter is a measurable function $w_A : X^2 \to [0, 1]$, where $w_A(x_A, x_B)$ represents the probability with which candidate A is elected when the platforms are x_A and x_B. Candidate B is elected with the complementary probability $w_B(x_B, x_A) := 1 - w_A(x_A, x_B)$.

We study (weak) perfect Bayesian equilibria (y_A, y_B, w_A) of the electoral game in which candidates play pure strategies—hereafter, simply *equilibria*.[16] Hence, as noted in Section 2,

[14] In this case, it is routine to verify that conditional on signal s_i, the distribution of signal s_{-i} is normal with mean $\frac{\beta_i}{\alpha+\beta_i} s_i$ and variance $\sigma^2 := \frac{1}{\alpha+\beta_i} + \frac{1}{\beta_{-i}}$. Hence, Equation 2 holds with

$$T(s_{-i}) = \frac{1}{\sigma^2} \cdot \frac{\beta_i}{\alpha + \beta_i} s_{-i}, \quad \psi(s_i) = \frac{1}{2\sigma^2} \left(\frac{\beta_i}{\alpha + \beta_i} s_i \right)^2, \quad \text{and} \quad h(s_{-i}) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2\sigma^2} s_{-i}^2}.$$

[15] We assume that both candidates can choose from the same set of platforms for notational simplicity. Our analysis in this section would hold equally well if each candidate i can only choose platforms from some subset $X_i \subset X$. One could use $X_A \neq X_B$ to capture asymmetries between the candidates, e.g., if there is an incumbent and a challenger, and the incumbent's history precludes him from choosing certain policies.

[16] Our leading specifications—such as the normal-normal structure—have continuous signals with atomless distributions, in which case it is salient to focus on equilibria with pure candidate strategies. Theorem 2 assures that such equilibria exist regardless of the model specification. Notwithstanding, we discuss equilibria in which candidates may mix in Subsection 3.3.

the players' strategies are identifiable in these equilibria. The voter elects candidate i if x_i is strictly preferred to x_{-i}. We allow the voter to randomize arbitrarily when indifferent. For the median voter interpretation, one may want to insist on uniform randomization when indifferent; our results would be unaffected by this requirement, modulo one caveat noted in fn. 20.

The notion of welfare we adopt is the voter's ex-ante expected utility, which we denote $v(y_A, y_B, w_A)$ as a function of the strategies.

Policy Commitment. In the Downsian tradition, our model assumes that candidates make commitments to the policies they would implement if elected.[17] In reality, while commitment may be imperfect, some degree of it is plausible and valuable to candidates; in their meta-study of earlier research, Pétry and Collette (2009) conclude that around 67% of campaign promises have historically been kept. The theoretical literature has proposed multiple rationales for commitment, most prominently that of re-election concerns (Alesina, 1988). Alternatively, if there is uncertainty about a candidate's quality of information and candidates have reputation concerns (perhaps because of re-election motives), then "flip flopping" or "vacillating" may be associated with poor quality information, resulting in stickiness akin to commitment (e.g., Majumdar and Mukand, 2004).

Alternative Model Interpretations. Although we focus on the elections interpretation, the model applies equally well to other political or organizational settings involving a principal or decision-maker (DM) and two agents (cf. Ambrus et al., 2021). The privately-informed agents make proposals and the DM selects one of the two proposals; each agent simply wants her own proposal to be selected, while the DM wants a proposal that best matches the state.

3.1. The Limit to Information Aggregation

Our welfare conclusions for our Downsian model stem from the following result.

Proposition 1. *Assume the signal distribution F_{s_A,s_B} satisfies Condition 1, and consider any equilibrium with candidates' strategies (y_A, y_B). There is a candidate $i \in \{A, B\}$ such that the*

[17] See Osborne and Slivinski (1996), Besley and Coate (1997) and subsequent work for non-Downsian "citizen-candidate" models.

voter's welfare in this equilibrium is the same as if candidate i were elected no matter which policies are proposed using (y_A, y_B).

To elaborate, consider any equilibrium (y_A, y_B, w_A). Denote by $v_i(y_A, y_B)$ the voter's welfare from electing candidate i no matter which policies are proposed. Since the voter always has the option of electing one candidate regardless of the platforms, it holds that $v(y_A, y_B, w_A) \geq \max\{v_A(y_A, y_B), v_B(y_A, y_B)\}$. Proposition 1 says that under Condition 1, the bound is tight:

$$v(y_A, y_B, w_A) = \max\{v_A(y_A, y_B), v_B(y_A, y_B)\}.$$

Put another way, insofar as equilibrium welfare is concerned, the voter may as well be ignoring one of the candidates and always electing the other. Crucially, this is an "as if": both candidates may in fact win with positive probability in equilibrium, as detailed in Subsection 3.2.

Proposition 1 is a straightforward application of Corollary 1. Given an arbitrary voter strategy w_A (recall $w_B \equiv 1 - w_A$), the two candidates are engaged in a constant-sum Bayesian game in which each candidate i's payoff is the probability $w_i(x_i, x_{-i})$ with which he wins the election. These payoffs are type independent because the candidates are office motivated and the voter can only infer their types from the platforms. As candidates have strict preferences over the winning probability outcome, Corollary 1 implies that under Condition 1, in any equilibrium the probability of a candidate winning is independent of which on-path platforms are proposed. Hence, there are only two possibilities on the equilibrium path. Either (i) one candidate wins with probability one, or (ii) both candidates win with a constant interior probability, regardless of their platforms. In the latter case, the voter is always indifferent between the candidates. It follows that in either case, the voter's ex-ante expected utility can be evaluated as if she always elects the same candidate.

Proposition 1 implies a sharp upper bound on the voter's welfare across all equilibria. To make that precise, let v_i^* denote the voter's welfare if candidate i were always elected with his platform chosen—based on his information alone—to maximize voter welfare. Proposition 1 implies that in any equilibrium, welfare is at most

$$\max\{v_A^*, v_B^*\}. \tag{3}$$

In other words, even when both candidates have socially valuable information, the voter's equilibrium welfare is, at best, determined by the efficient use of only one candidate's signal.

Theorem 2 below states that point and also observes that the upper bound (3) can be achieved. Say that candidate i is *better* (than his opponent) if $v_i^* \geq v_{-i}^*$. That is, if each candidate would choose the voter-optimal policy based on their information alone, playing $y_i^*(s_i) := \arg\max_{x \in X} \mathbb{E}_\theta [u(x,\theta) \mid s_i]$,[18] then the voter would prefer to ex-ante delegate policy-making to i rather than the opponent. (If $v_A^* = v_B^*$, then without loss we stipulate that A is the better candidate.)

Theorem 2. *If the signal distribution F_{s_A,s_B} satisfies Condition 1, then a voter-welfare maximizing equilibrium has welfare $\max\{v_A^*, v_B^*\}$. There is one such equilibrium in which the better candidate $i \in \{A, B\}$ is elected with probability one and plays y_i^*.*

There may be multiple equilibria that achieve the proposition's welfare bound, but a simple construction is as follows. The better candidate i plays y_i^*, and the opponent uninformatively chooses the prior-optimal policy, i.e., he plays $y_{-i}(s_{-i}) = \arg\max_{x \in X} \mathbb{E}_\theta[u(x,\theta)]$. It is then optimal for the voter to always elect i on path. We can stipulate that the voter also elects i if $-i$ chooses any other platform (and i chooses any of his on-path platforms) because she believes that the deviation by $-i$ is uninformative about s_{-i}.[19,20] Note that this construction does not require Condition 1; rather, the condition guarantees, by Proposition 1, that this equilibrium is welfare maximizing. The following example shows that Condition 1 cannot be dispensed with in Theorem 2.

Example 3. Let $X = S_A = S_B = \{1,2,3,4\}$, and $\Theta = S_A \times S_B$ with a uniform prior. In each state θ, the signal profile is deterministically $(s_A, s_B) = \theta$. Hence, the unconditional joint signal distribution is uniform on $S_A \times S_B$, violating Condition 1. The voter's utility is

[18] If there are multiple maximizers at any s_i, we can choose an arbitrary one.

[19] Any sequentially rational behavior by the voter after an observable deviation by i supports the equilibrium, as i has no incentive to deviate.

[20] Let x_i be an on-path platform of candidate i. Our construction entails the voter electing candidate i even if both candidates choose x_i. If one insists that the voter must randomize uniformly between the candidates when indifferent, then Theorem 2 is still valid with essentially the same construction so long as every on-path platform of candidate i has zero ex-ante probability. This is the case with a continuous policy space when there is a unique and distinct optimal policy after each signal of the better candidate i and the marginal distribution F_{s_i} is atomless. An example is the normal-normal information structure with quadratic-loss voter preferences.

1 if the "correct" action is taken in a state and 0 otherwise, with the correct action in each state—or equivalently, after each signal profile (s_A, s_B)—marked in blue in the table below. The table also shows, in magenta, a strategy for the voter, i.e., which candidate the voter elects following each platform pair (x_A, x_B).

$s_A\,/\,s_B$ $x_A\,/\,x_B$	1	2	3	4
1	1, A	2, B	3, B	1, A
2	1, B	2, A	3, B	2, A
3	3, A	2, B	3, A	4, B
4	1, B	4, A	4, A	4, B

This voter strategy and each candidate playing the strategy $s_i \mapsto s_i$ constitute an equilibrium: the voter is playing optimally because she obtains her preferred policy in every state; candidates are playing optimally because, no matter their signal, their posterior is uniform over the opponent's signal and they thus expect to win with probability 1/2 no matter their platform. This equilibrium achieves the voter's first-best welfare, which is larger than the welfare level $\max\{v_A^*, v_B^*\}$ because neither candidate's signal individually reveals the state. Evidently, the conclusions of Proposition 1 and Theorem 2 do not hold. □

3.2. The Normal-Quadratic Specification

To substantiate Proposition 1 and Theorem 2, we now elaborate on a leading specification of our Downsian model. The analysis of this subsection yields insights into how politicians' strategic incentives play out in equilibria—in particular, on whether office motivation necessarily leads to pandering, and whether pandering is detrimental to welfare.

Consider a one-dimensional *normal-quadratic* specification: $X, \Theta \subset \mathbb{R}$, the voter's utility is $u(x,\theta) = -(x-\theta)^2$, the state is $\theta \sim \mathcal{N}(0, 1/\alpha)$, and, conditional on the state θ, each candidate $i \in \{A, B\}$ receives an independent signal $s_i \sim \mathcal{N}(\theta, 1/\beta)$, with parameters $\alpha, \beta \in \mathbb{R}_{>0}$. Note that the unconditional joint signal distribution satisfies Condition 1. Subsection A.4 discusses how some of this section's themes generalize to broader informational structures.

Quadratic-loss utility implies that that voter's preferred policy given any information $\mathcal{I}$ is

$\mathbb{E}[\theta \mid \mathcal{I}]$. Hence, by standard properties of normal information,

$$y_i^*(s_i) = \mathbb{E}\left[\theta \mid s_i\right] = \frac{\beta}{\alpha+\beta} s_i, \tag{4}$$

and we refer to y_i^* as the *unbiased* strategy because it is the best estimate of state given s_i. We say that a strategy y_i displays *pandering* (or underreaction) if $s_i > 0 \implies y_i(s_i) \in [0, \mathbb{E}[\theta \mid s_i])$, $s_i < 0 \implies y_i(s_i) \in (\mathbb{E}[\theta \mid s_i], 0]$, and $y_i(0) = 0$. In other words, a candidate panders if for $s_i \neq 0$ his platform is distorted from his unbiased estimate toward the voter's prior expectation $\mathbb{E}[\theta] = 0$ of the best policy. Analogously, we say that y_i displays *anti-pandering* (or overreaction) if $s_i > 0 \implies y_i(s_i) > \mathbb{E}[\theta \mid s_i]$ and $s_i < 0 \implies y_i(s_i) < \mathbb{E}[\theta \mid s_i]$. We also say that a platform x is *more extreme* than platform x' if the former is further from the prior mean of 0, i.e., if $|x| > |x'|$. A strategy y_i is *informative* if it is not constant, and it is *fully revealing* if it is bijective. An equilibrium is *symmetric* if both candidates use the same strategy and both win with positive probability.

Unbiased Strategies. There are trivial equilibria in which candidates disregard their information, e.g., the "full pandering equilibrium" in which they each play $y_i(\cdot) = \mathbb{E}[\theta] = 0$, and the voter elects each candidate with some constant probability no matter the platforms. To tackle informative equilibria, a natural starting point is the profile of unbiased strategies. From Equation 4, we see that the voter would then infer from a platform x_i that $s_i = \frac{\alpha+\beta}{\beta} x_i$. As the expected value of θ conditional on both signals is

$$\mathbb{E}[\theta \mid s_A, s_B] = \frac{2\beta}{\alpha+2\beta}\left(\frac{s_A+s_B}{2}\right), \tag{5}$$

the voter's posterior expectation of the state given the platforms x_A and x_B is

$$\frac{2(\alpha+\beta)}{\alpha+2\beta}\left(\frac{x_A+x_B}{2}\right).$$

So the voter's preferred policy, which is the posterior expectation, has the same sign as the average of the two platforms but is more extreme (so long as the average is non-zero). The voter thus elects the more extreme candidate, and consequently, each candidate would benefit by deviating to a more extreme platform, i.e., by anti-pandering or overreacting to his information. Proposition 5 in Appendix A.1 provides a formal statement.

Equilibrium Anti-Pandering. Building on the above intuition, the next result identifies an anti-pandering equilibrium.

Proposition 2. *In the normal-quadratic specification, there is an anti-pandering equilibrium, which is symmetric and fully-revealing: both candidates play*

$$y_i(s_i) = \mathbb{E}[\theta \mid s_i, s_{-i} = s_i] = \frac{2\beta}{\alpha + 2\beta} s_i, \tag{6}$$

and each candidate is elected with probability $1/2$ *regardless of their platforms. The voter's welfare in this equilibrium is*

$$-\frac{\alpha + 4\beta}{(\alpha + 2\beta)^2}. \tag{7}$$

Moreover, any symmetric equilibrium in which both candidates use fully-revealing and continuous pure strategies has both candidates playing (6) *and voter welfare* (7).

In the equilibrium of Proposition 2, candidates can be viewed as choosing the unbiased platform based on a signal with twice the actual accuracy. Alternatively, each candidate's platform is the Bayesian estimate of the state assuming his opponent has received the same signal. That is despite each candidate i knowing that, in expectation, his opponent's signal is in fact more moderate than his own, as that expectation is just i's unbiased estimate of the state, $\frac{\beta}{\alpha+\beta} s_i$. When the voter conjectures that both candidates play the strategy (6), she is indifferent between the candidates no matter their platforms. For, whenever a candidate i increases his platform by any $\delta > 0$, formula (5) implies that the voter's posterior expectation increases by $\frac{2\beta}{\alpha+2\beta}\left(\frac{\alpha+2\beta}{2\beta}\frac{\delta}{2}\right) = \delta/2$.

An important implication of Proposition 2 is that office motivation does not necessarily lead to pandering. Moreover, the voter's welfare (7) in the anti-pandering equilibrium is higher than in the trivial full-pandering equilibrium, as the welfare in the latter is simply $-1/\alpha$. Consistent with Proposition 1, both the anti-pandering equilibrium and the trivial full-pandering equilibrium have the ex-post property for the candidates, and the voter's welfare in these equilibria is the same as if she always elected either candidate. Moreover, consistent with the construction we described for Theorem 2, there is yet another equilibrium: (either) candidate i plays the unbiased strategy (4), the other candidate $-i$ plays $y_{-i}(\cdot) = \mathbb{E}[\theta] = 0$, and the voter always elects candidate i.[21] The voter's welfare in this equilibrium can be straightforwardly computed as $-1/(\alpha+\beta)$, which is even higher than the anti-pandering equilibrium's welfare (7). Indeed, although it is infeasible to characterize all equilibria of the normal-quadratic specification, Theorem 2 tells us that that $-1/(\alpha+\beta)$ is the maximum equilibrium welfare.

The (Disequilibrium) Benefits of Pandering. Interestingly, in this normal-quadratic specification, an appropriate degree of non-equilibrium pandering would actually benefit the voter. To get some intuition for why, consider again the benchmark where both politicians play the unbiased strategy $y_i^*(\theta) = \mathbb{E}[\theta \mid s_i]$. As explained above, the voter would then select the politician with the most extreme platform. This implies a "winner's curse": the electoral winner, say i, would have received the most extreme signal, and so voter welfare would be improved if i were elected with a slightly more moderate platform. Such moderation can be achieved by underreacting to private information, i.e., by pandering—although that run counter to office motivation.

To formalize the point, consider the following strategy:

$$y_i(s_i) = \mathbb{E}\left[\theta \mid s_i, |s_{-i}| \leq |s_i|\right], \tag{8}$$

which features pandering because conditioning on the opponent having a more moderate signal makes a candidate underreact to his own signal. Lemma 2 in Appendix A.3 verifies that, and also shows that the voter's best response to both candidates playing (8) is to elect the candidate with the more extreme platform (which, recall, is also her best response to both candidates using unbiased strategies).

Proposition 3. *In the normal-quadratic specification, consider the strategy profile where both candidates pander by playing strategy* (8)*, and the voter best responds. This profile yields higher voter welfare than any equilibrium, as well as the non-equilibrium profile in which the candidates play the unbiased strategies* (4) *and the voter best responds.*

[21] In this normal-quadratic specification, the same outcome—i.e., that i plays (4) is always elected—can also be supported in a fully-revealing equilibrium with $y_{-i}(s_{-i}) = s_{-i}$. Candidate $-i$ is overreacting to his information here to such an extent that the voter never finds it optimal to elect $-i$ despite correctly inferring his information.

In fact, we prove in Appendix A.3 that the (non-equilibrium) strategy profile of Proposition 3 maximizes voter welfare within a broad class of profiles. Our takeaway is that an appropriate degree of pandering would benefit the voter.

3.3. Discussion

We now return to our general Downsian model with informed candidates and discuss the robustness of our welfare conclusions.

Candidates Mixing. Proposition 1, and hence the welfare bound of Theorem 2, also apply to equilibria in which candidates mix, so long as at least one candidate plays an identifiable strategy. This is because the ex-postness conclusion of Theorem 1 applies to such equilibria. We do not know whether equilibria in which both candidates play non-identifiable strategies—if they exist in a given specification—can overturn the conclusion of Proposition 1/Theorem 2.[22] In particular, even non-ex-post equilibria can still satisfy the welfare conclusions. To illustrate, consider a variant of Example 1: $\Theta = X = \{1, 2\}$, $u(x, \theta) = \mathbb{1}\{x = \theta\}$, a uniform prior on Θ, and any signal structure F_{s_A,s_B} that satisfies Condition 1. There is an equilibrium in which, regardless of their signals, both candidates mix uniformly over both policies, and the voter (being indifferent between both policies) plays $w_A(x_A, x_B) = \mathbb{1}\{x_A = x_B\}$. Neither candidate's strategy is identifiable and Theorem 1 does not apply; yet, the equilibrium trivially still satisfies the conclusion of Proposition 1.

Other Game Forms. Proposition 1 also applies much more generally than to the canonical Downsian game form we have considered. For concreteness, we only mention two variations:

1. The elected candidate does not necessarily implement their platform x_i, but instead some exogenous—possibly stochastic—function of x_i. For example, there could be a status quo policy x^0 (e.g., the ex-ante optimal policy), and the elected candidate i implements their platform x_i with some probability (which could depend on x_i) and x^0 otherwise. Alternatively, the candidate may always moderate after the election and implement $q_i \cdot x_i$

[22] In some specifications, we can deduce that they do not. For example, consider a binary-policy binary-signal setting (e.g., Heidhues and Lagerlof, 2003). Here the only non-identifiable strategies are uninformative and so an equilibrium in which neither candidate uses an identifiable strategy is clearly no better for voter welfare than efficiently aggregating one candidate's signal.

for some parameter $q_i > 0$.

2. Instead of choosing platforms simultaneously, candidate A chooses his platform x_A first, and candidate B, having observed x_A, then chooses x_B. The asymmetry in timing might reflect that one candidate is an incumbent and the other a challenger.

The reason Proposition 1 holds for these variations is that the candidates are still purely office motivated and the voter's decision cannot depend directly on their signals; hence, given any voter strategy, the candidates still face a constant-sum Bayesian game with type-independent payoffs, and Theorem 1 applies.

On the other hand, the welfare level obtained in Theorem 2 may not apply to these variations. In model variation #1 above, the upper bound on equilibrium welfare would have to be adjusted for the stochastic policy implementation; if the status quo is implemented with high probability, then evidently equilibrium voter welfare cannot be much higher than the expected utility from the status quo. More importantly, model variation #2 above generally allows for equilibria that achieve welfare higher than $\max\{v_A^*, v_B^*\}$. For, the level v_i^* obtains from efficiently using only candidate i's signal. Under natural specifications, there can even be full information aggregation in model variation #2: candidate A reveals his signal s_A via his platform; candidate B then proposes the best policy for the voter given both s_A and s_B; and the voter always selects candidate B. Certainly this raises questions about equilibrium refinements.[23] We do not pursue that issue; instead, we observe that the welfare bound of Theorem 2 is focal because it applies to the canonical Downsian game form.

Voter Commitment. For some interpretations of the model—such as decision-making in an organization, as mentioned earlier—it is plausible that the voter (decision maker) can commit ex ante to how she will select among the candidates' (agents') platforms (proposals). Since the constant-sum property between candidates holds for an arbitrary voter strategy, Proposition 1 and Theorem 2 also apply to this case. In other words, commitment cannot increase the maximum voter welfare.

[23] Indeed, there could be another equilibrium in which candidate A—without loss, the better candidate—proposes the best policy given his signal s_A (which reveals his signal); candidate B then proposes the *worst* policy for the voter given both candidates' signals; and voter always selects candidate A. This equilibrium's welfare is the same as that of Theorem 2.

Beyond Office Motivation. Since candidates' office motivation is a key assumption for our results, we conclude this subsection by discussing the robustness of Theorem 2's welfare conclusion to small departures from that assumption.

Consider a variant of our Downsian model in which the payoff of each candidate $i \in \{A, B\}$ is given by $u_i(x_A, x_B, \theta, W; \gamma_i)$, where the new notation $W \in \{A, B\}$ denotes the election's winner and γ_i is a commonly-known payoff parameter. Pure office-motivation corresponds to the utility $\mathbb{1}\{W = i\}$, but in general $u_i(\cdot)$ allows for a variety of *mixed motivations*, including policy motivation (a candidate cares about the winner's policy, in relation to the state) and platform motivation (he cares about his own platform, in relation to the state).

An election with mixed motivations is not generally a constant-sum game for the candidates; consequently, for arbitrary mixed motivations, voter welfare may be significantly different from the bound in Theorem 2. However, consider a family of mixed-motivations games in which each candidate i's payoff is parameterized by $\gamma_i \in \mathbb{R}^m$ such that $u_i(x_A, x_B, \theta, W; \vec{0}) = \mathbb{1}\{W = i\}$. That is, when $\gamma_i = \vec{0} \equiv (0, \ldots, 0)$, candidate i is purely office motivated. Under appropriate technical conditions, the Theorem of the Maximum assures that the equilibrium correspondence is upper hemicontinuous in the parameter (γ_A, γ_B), and hence the upper bound on voter welfare when $(\gamma_A, \gamma_B) \approx (\vec{0}, \vec{0})$, i.e., when candidates are almost office-motivated, is approximately that of Theorem 2.[24] Simple sufficient technical conditions are that all the spaces S_A, S_B, Θ, and X are finite and that each $u_i(\cdot)$ is continuous in $(x_A, x_B, \theta, \gamma_i)$.

We note that our leading one-dimensional normal-quadratic specification from Subsection 3.2 does not satisfy the aforementioned technical conditions; in particular, the policy space $X = \mathbb{R}$ is not compact. The Supplementary Appendix analyzes an extension of the normal-quadratic model with mixed motivations of the form

$$u_i(x, \theta, W; b_i, \rho_i) = -\rho_i(x_W - \theta - b_i)^2 + (1 - \rho_i)\mathbb{1}\{W = i\}. \tag{9}$$

So each candidate i has quadratic-loss policy utility with an ideological bias $b_i \in \mathbb{R}$ and places

[24] More precisely, we would be assured upper hemicontinuity of the set of Bayes-Nash equilibria. Although our solution concept is weak Perfect Bayesian equilibrium (in which candidates use pure strategies), Proposition 1 holds for Bayes-Nash equilibria too because its backbone, Theorem 1, guarantees the ex-post property for an arbitrary voter strategy. Note also that we implicitly restrict attention to equilibria of the perturbed games in which candidates use pure strategies, to ensure that this property is preserved in any limit.

weight $\rho_i \in [0,1]$ on policy utility. The Supplementary Appendix establishes that even though the equilibrium correspondence is not upper hemicontinuous, the upper bound on voter welfare when each $b_i \approx 0$ and $\rho_i \approx 0$ is still close to that of efficiently using only one candidate's signal. Moreover, there is an equilibrium that approximately achieves that welfare. So, the welfare conclusions of Theorem 2 still approximately hold.

4. Competition in Dual Spheres

To illustrate the implications of Theorem 1 beyond elections, we now develop an application involving competition in dual spheres. We frame it as two firms competing to both maximize their shares of a private market and to secure a government allocation. So, as in Section 3, we have two agents (politicians previously, now firms) competing for the favor of a principal (an electorate previously, now a government). Now, however, the agents' actions directly affect their own payoffs but not the principal's. At the end of this section we discuss how the framework is more broadly applicable beyond the framing with which we introduce it.

Consider two firms, A and B. Each firm $i \in \{A, B\}$ chooses a product or technology $x_i \in X_i$. There is an unknown state of the world $\theta \in \Theta \subset \mathbb{R}^n$, representing a variable that matters for a government's action (elaborated below); for instance, θ could reflect the relative social value of the firms. Each firm i observes a private signal $s_i \in S_i \subset \mathbb{R}^n$ about the state, drawn from some joint distribution conditional on the state $F_{s_A,s_B|\theta}$. Denote the unconditional joint distribution of signals by F_{s_A,s_B}.[25]

The firms' choices have consequences in two domains. First, they determine market shares in a private market, captured by a function $m : X_A \times X_B \to [0,1]$, where $m(x_A, x_B)$ is firm A's share. This private-market competition is zero sum: it contributes a payoff $m(x_A, x_B)$ to firm A and $1-m(x_A, x_B)$ to firm B. Note that the function m does not depend on the state θ; the state represents social value rather than appeal in the private market. This is reasonable when θ represents externalities, long-run reliability, or other attributes that the market does not price.

Second, a government observes some statistic $t \in \mathcal{T}$ of the firms' choices, generated by

[25] We suppress the technical conditions on Θ, each S_i, and the distributions, which follow those in the previous sections.

the map $\tau(x_A, x_B)$, and then chooses an action or allocation $a \in \mathcal{A}$. The government's payoff is given by $u_G(x_A, x_B, a, \theta)$. For instance, a may represent the share of public procurement allocated to firm A, and the payoff u_G may represent how well this allocation matches the state, with higher values of θ leading the government to prefer larger shares for firm A. That could be captured by $a \in [0, 1]$, $\theta \in \mathbb{R}$, and $u_G = -(a - \theta)^2$.

The firms care about both their private market share and the government action (e.g., public procurement share). For some bounded function $v : \mathcal{A} \to \mathbb{R}$, firm A's overall payoff is

$$u_A(x_A, x_B, a) := m(x_A, x_B) + v(a),$$

and firm B's, after normalization, is

$$u_B(x_A, x_B, a) := -m(x_A, x_B) - v(a).$$

We refer to this setting as one of *competition in dual spheres*, because the firms are competing for both market share and the government action.

Observe that any government strategy $\alpha : \mathcal{T} \to \Delta(\mathcal{A})$ induces a constant-sum Bayesian game with type-independent payoffs between the firms in which A's payoff is $m(x_A, x_B) + v_\alpha(x_A, x_B)$, where we define $v_\alpha : (x_A, x_B) \mapsto \mathbb{E}\left[v(\alpha(\tau(x_A, x_B)))\right]$, with the expectation over the government's randomization.[26] Denoting firm i's strategy by $\sigma_i : S_i \to \Delta(X_i)$, the following result follows immediately from Corollary 1.

Proposition 4. *Consider competition in dual spheres, with F_{s_A, s_B} satisfying Condition 1. In any Bayes-Nash equilibrium $(\sigma_A^*, \sigma_B^*, \alpha^*)$ in which firms play identifiable strategies, there exists a constant c such that (a.s.) on path, $m(x_A, x_B) + v_{\alpha^*}(x_A, x_B) = c$.*

Although stated for (Bayes-Nash) equilibria of the game, the result evidently holds for any government strategy to which the firms mutually best respond.

Proposition 4 establishes that—under Condition 1 and identifiability, qualifiers we omit in the subsequent paragraphs for brevity—competitive pressure in the private market ties the government's hands. Whatever the government's objective, its equilibrium actions must

[26] We assume this expectation is well-defined; we suppress such technical details in the rest of this section.

perfectly offset the market outcome in the sense that $v_{\alpha^*}(x_A, x_B) = c - m(x_A, x_B)$. In other words, the firms' dual competition forces a rigid relationship between the private market and the government action.

What does Proposition 4 imply about information revelation and government welfare? Intuitively, it places strong constraints. For brevity, we only discuss two illustrative specifications. In both cases, we assume the government only observes the market outcome: $\tau(\cdot) = m(\cdot)$. First, suppose the firms don't actually care about the government action: $v(\cdot) = 0$. Then it is immediate that in any equilibrium the market outcome must be constant on path. This implies that the government learns nothing, so it cannot tailor its action to the fundamental θ even if firms are arbitrarily well informed; moreover, even aside from its own action, any desire for the government to have market outcomes correlated with θ is defeated in equilibrium.

Second, suppose the government's action is $a \in [0, 1]$ with linear firm valuation $v(a) = a$. Then Proposition 4 says that in equilibrium $m + \mathbb{E}[\alpha(m)]$ is constant. With $\theta \in \mathbb{R}$ and government utility $u_G(a, \theta) = -(a - \theta)^2$, it is $m + \mathbb{E}[\theta \mid m]$ that must be constant. It will typically not be possible for the posterior mean to decrease and precisely offset any increase of the market share when the latter is informative. So, again, typically the equilibrium market outcome has to be uncorrelated with the fundamental and the government will not learn any of the firms' information.[27]

Remark 1. The framework of this section is in some respects quite general and relevant to various economic contexts, as we now illustrate:

1. It subsumes the Downsian election model from Section 3. Concretely, the Downsian one-dimensional quadratic specification obtains when $X_A = X_B = \Theta = \mathbb{R}$, $\tau(x_A, x_B) = (x_A, x_B)$, $\mathcal{A} = \{A, B\}$, $u_G(x_A, x_B, a, \theta) = -(x_a - \theta)^2$, and $m(\cdot) = 0$ and $v(a) = \mathbb{1}\{a = A\}$. This corresponds to the principal (voter) selecting a winning agent (candidate) to mini-

[27] To see the role of Condition 1 in the conclusions for both cases, consider a variation of Example 1. Let the distribution on $S_A \times S_B = \{0, 1\} \times \{0, 1\}$ be uniform (violating Condition 1) and take $\theta = \mathbb{1}\{s_A = s_B\}$. Let $X_A = X_B = \{0, 1\}$, and $m(x_A, x_B) = |x_A - x_B|$. Suppose each firm i plays the pure strategy $s_i \mapsto s_i$. From each firm's perspective, taking the other firm's strategy as given, either of its own actions induces a uniform lottery over market outcomes 0 and 1. But the realized outcome fully reveals θ; specifically, abusing notation, we have $m \in \{0, 1\}$ and $\mathbb{E}[\theta \mid m] = 1 - m$. So, we have an equilibrium (with identifiable firms' strategies) either when (i) $v(\cdot) = 0$ or (ii) $v(a) = a$ and the government optimally chooses $a = \mathbb{E}[\theta \mid m]$.

mize the distance between the winner's action and the state, while the agents just want to be selected.

2. Consider a modification of the above specification to $u_G(a, \theta) = -(\mathbb{1}\{a = A\} - \theta)^2$. Then the agents' actions are cheap-talk messages, and the principal wants to select agent A when the state is high. This is now a model of arbitration with a binary allocation, similar to Kattwinkel et al. (2022). Proposition 4 implies that under its conditions, the principal's (probabilistic) selection in any equilibrium—or, even in a stochastic mechanism that the principal commits to, as the result holds for any principal strategy—must be independent of the agents' signals, as those authors also show for the case of finite signals.

   But Proposition 4 can be applied to richer arbitration problems as well. Consider $\mathcal{A} = [0, 1]$, where a represents the arbitrator's ruling of a transfer from agent B to A, and any strictly increasing affine $v(a)$ and any $u_G(a, \theta)$. Assume the arbitrator has a unique prior-optimal ruling, $a^* := \arg\max_{a \in \mathcal{A}} \mathbb{E}_\theta [u_G(a, \theta)]$. Then Proposition 4 implies that under its conditions, the *expected* transfer in any equilibrium (or mechanism) is a^*, independent of the agents' signals. While that does not generally necessitate the realized transfer to be constant, it would if $u_G(\cdot, \theta)$ were strictly concave for each θ, for example.

3. The principal can be passive while caring about the agents' interaction and the state. For instance, agent B may be a regulator (or security force) chasing a non-compliant firm A (or interdicting a smuggler). Each agent i chooses a location $x_i \in X_i \subset \mathbb{R}^k$ and they have opposing preferences over their location gap or compliance: $m(x_A, x_B) = \|x_A - x_B\|$ for some norm $\|\cdot\|$, and $v(\cdot) = 0$. The principal/society takes no action ($|\mathcal{A}| = 1$) but has utility $u_G(x_A, x_B, \theta) = -\theta \|x_A - x_B\|$. So society would prefer tighter compliance when the stakes $\theta \in \Theta \subset \mathbb{R}_{\geq 0}$ are higher. Although the agents are informed about these stakes, Proposition 4 implies that under its conditions, in any equilibrium the realized compliance does not vary with signal profiles or stakes.[28]

[28] Even though the agents do not care about the stakes, absent Condition 1 there are examples in which society gets its first-best state-contingent outcome; cf. fn. 27.

5. Conclusion

We have developed a general result about equilibrium behavior in two-player constant-sum Bayesian games with type-independent payoffs. Under a completeness statistical condition on the distribution of types, any identifiable (Bayes-Nash) equilibrium must be ex post: each player is indifferent among all the actions played by any of his types, even after observing the opponent's action. This ex-postness property limits the extent to which adversarial players' information can be revealed to and used by third parties.

Motivated by the debate on whether political competition promotes information aggregation and informed choices by electorates, we have applied the above result to Downsian electoral competition between two office-motivated candidates who have private information about policy consequences. We find a sharp bound on the (median or representative) voter's welfare. Welfare in any equilibrium–under our statistical condition and pure/identifiable politician strategies—is effectively determined by just one candidate's platform strategy. Consequently, Downsian elections cannot efficiently aggregate more than one candidate's information, despite the availability of two informational sources. Moreover, the upper bound of efficiently aggregating the "better" candidate's information can be achieved in an equilibrium.

To substantiate the electoral welfare bound and to better understand politicians' strategic incentives, we have studied in more detail a normal-quadratic specification of our Downsian model. In that specification, there is a fully-revealing equilibrium in which candidates' anti-pander or overreact to their information. Furthermore, we find that an appropriate degree of (disequilibrium) pandering by candidates would actually benefit voters. These findings run counter to conventional views that candidates' pandering is an inevitable consequence of candidate office-motivation, and is necessarily harmful to voters.

As in most formal models of spatial electoral competition, we have restricted attention to two candidates and assumed that their information is exogenously given. Relaxing both these assumptions are interesting topics for future research. We note here that since a voter-optimal equilibrium of our model involves always electing the "better"—roughly, more informed—candidate, there can be strong incentives for candidates to observably acquire information.

While our primary application is to electoral competition, the logic of Theorem 1 applies more broadly. We have illustrated this with a model of competition in dual spheres, in which firms compete for both private market share and government actions. Even though the firms may be well informed about a fundamental that the government cares about, their zero-sum rivalry forces a rigid relationship between market outcomes and government actions—for instance, precluding state-contingent government procurement. This application underscores that the limits to information revelation and aggregation identified in this paper are not specific to elections, but are a general feature of institutions constrained by purely adversarial incentives.

A. Proofs and Other Material for Section 3

We omit proofs for Proposition 1 and Theorem 2 (and also Proposition 4 in Section 4), as they were explained in the main text.

A.1. Unbiased Strategies

Let us substantiate the discussion in Subsection 3.2 by showing that candidates cannot play unbiased strategies in an equilibrium of the normal-quadratic specification.

Proposition 5. *In the normal-quadratic specification, the profile of unbiased strategies cannot be supported in an equilibrium. In particular, candidates would deviate by overreacting to their information, whereas underreacting would be worse than playing the unbiased strategy.*

Proof. Assume both candidates use the unbiased strategy $y_i(s_i) = \frac{\beta}{\alpha+\beta}s_i$. Since this strategy is fully revealing, the voter correctly infers s_A, s_B for all signal realizations. The voter's expected utility from a platform x follows a standard mean-variance decomposition:

$$\begin{aligned}\mathbb{E}[u(x,\theta) \mid s_A, s_B] &= -\mathbb{E}\left[(x-\theta)^2 \mid s_A, s_B\right] \\ &= -\left[x^2 + \mathbb{E}\left[\theta^2 \mid s_A, s_B\right] - 2x\mathbb{E}\left[\theta \mid s_A, s_B\right]\right] \\ &= -\left[x^2 + \left(\mathbb{E}\left[\theta \mid s_A, s_B\right]\right)^2 - 2x\mathbb{E}\left[\theta \mid s_A, s_B\right]\right] - \mathbb{E}\left[\theta^2 \mid s_A, s_B\right] + \left(\mathbb{E}\left[\theta \mid s_A, s_B\right]\right)^2 \\ &= -\left[x - \mathbb{E}(\theta \mid s_A, s_B)\right]^2 - \operatorname{Var}\left(\theta \mid s_A, s_B\right). \end{aligned} \tag{10}$$

So the voter elects candidate i whenever x_i is closer to $\mathbb{E}[\theta \mid s_A, s_B]$ than is x_{-i}.

We now show that for any $i = A, B$ and s_i, candidate i can profitably deviate. By (10), if i plays as if he has received signal $\hat{s}_i$ (no matter his true signal), then i wins against any realization s_{-i} such that

$$(y_{-i}(s_{-i}) - \mathbb{E}[\theta \mid \hat{s}_i, s_{-i}])^2 > (y_i(\hat{s}_i) - \mathbb{E}[\theta \mid \hat{s}_i, s_{-i}])^2 .$$

Substituting from (4) and (5), this is equivalent to

$$\left(\frac{\beta}{\alpha+\beta}s_{-i} - \frac{\beta}{\alpha+2\beta}(\hat{s}_i + s_{-i})\right)^2 > \left(\frac{\beta}{\alpha+\beta}\hat{s}_i - \frac{\beta}{\alpha+2\beta}(\hat{s}_i + s_{-i})\right)^2,$$

or after algebraic simplification, $(\hat{s}_i)^2 > (s_{-i})^2$. Hence, i wins when he mimics a more extreme (i.e., larger in magnitude) signal than $-i$'s true signal. Since for any true signal s_i the conditional distribution of $-i$'s signal is normal with mean $\mathbb{E}[\theta \mid s_i] = \frac{\beta}{\alpha+\beta}s_i$, it follows that no matter his true signal, candidate i strictly increases his win probability by overreacting and strictly decreases it by underreacting. □

A.2. Anti-Pandering

Proof of Proposition 2. For the proposition's first statement, it suffices to verify that the voter is indifferent between the two candidates for any pair of platforms, assuming that both candidates play the strategy (6). Since the candidates' strategies are fully revealing, the voter correctly infers the candidates' signals from the platform pair. Furthermore, since the candidates' strategies each have range $\mathbb{R}$, there are no off-path platform pairs. Therefore, it suffices to show that for any s_i and s_{-i}, we have

$$-\mathbb{E}[(y_i(s_i) - \theta)^2 \mid s_i, s_{-i}] = -\mathbb{E}[(y_{-i}(s_{-i}) - \theta)^2 \mid s_i, s_{-i}],$$

or equivalently that $(y_i(s_i) - \mathbb{E}[\theta \mid s_i, s_{-i}])^2 = (y_{-i}(s_{-i}) - \mathbb{E}[\theta \mid s_i, s_{-i}])^2$.[29] Using (5) and (6), this latter equality can be rewritten as

$$\left(\frac{2\beta}{\alpha+2\beta}s_i - \frac{2\beta}{\alpha+2\beta}\left(\frac{s_i+s_{-i}}{2}\right)\right)^2 = \left(\frac{2\beta}{\alpha+2\beta}s_{-i} - \frac{2\beta}{\alpha+2\beta}\left(\frac{s_i+s_{-i}}{2}\right)\right)^2,$$

[29] That this latter equality is equivalent to the former follows from a standard mean-variance decomposition under quadratic loss utility as in the proof of Proposition 5.

which holds for any s_i, s_{-i}.

Next, we derive the equilibrium welfare (7) as follows:

$$\begin{aligned}\mathbb{E}[-(y_i-\theta)^2] &= \mathbb{E}\left[-\left(\frac{2\beta}{\alpha+2\beta}s_i-\theta\right)^2\right] \\ &= -\operatorname{Var}\left(\frac{2\beta}{\alpha+2\beta}s_i-\theta\right) \\ &= -\left(\frac{2\beta}{\alpha+2\beta}\right)^2\operatorname{Var}(s_i)-\operatorname{Var}(\theta)+2\left(\frac{2\beta}{\alpha+2\beta}\right)\operatorname{Cov}(s_i,\theta) \\ &= -\left(\frac{2\beta}{\alpha+2\beta}\right)^2\left(\frac{1}{\alpha}+\frac{1}{\beta}\right)-\frac{1}{\alpha}+2\left(\frac{2\beta}{\alpha+2\beta}\right)\frac{1}{\alpha} \\ &= -\frac{\alpha+4\beta}{(\alpha+2\beta)^2}.\end{aligned}$$

Finally, for the proposition's last statement, we prove something stronger that does not assume symmetry: in any equilibrium in which both candidates win with positive probability and use continuous fully-revealing pure strategies, there is $c\in\mathbb{R}$ and $i\in\{A,B\}$ such that

$$y_i(s_i)=\frac{2\beta}{\alpha+2\beta}s_i+c \quad \text{and} \quad y_{-i}(s_{-i})=\frac{2\beta}{\alpha+2\beta}s_{-i}-c.$$[30]

(Imposing symmetry, as in the proposition, implies $c=0$, which yields Equation 6.) To prove that, fix any equilibrium in which each candidate i uses a continuous and fully revealing strategy $\bar{y}_i$ and both win with positive probability. Denote the interior of the range of $\bar{y}_i$ by $\bar{X}_i$, noting that $\bar{X}_i$ is an open interval. Also denote $\bar{s}_i(x_i) := (\bar{y}_i)^{-1}(x_i)$. Theorem 1 and voter optimality imply that the voter is indifferent between both candidates after almost all on-path platform pairs. This implies that for almost all $x'_A\in\bar{X}_A$ and $x'_B\in\bar{X}_B$—hereafter we drop the "almost all" qualifier for brevity, understanding that some subsequent statements are up to measure zero sets, returning to the issue at the very end of the proof—we must have $\mathbb{E}[\theta\mid x'_A,x'_B]=\frac{x'_A+x'_B}{2}$, which implies $\frac{\beta}{\alpha+2\beta}(\bar{s}_A(x'_A)+\bar{s}_B(x'_B))=\frac{x'_A+x'_B}{2}$, or equivalently

$$\bar{s}_B(x'_B)=\frac{\alpha+2\beta}{2\beta}\left(x'_A+x'_B\right)-\bar{s}_A(x'_A). \tag{11}$$

[30] Using a very similar analysis to that in the first paragraph of this proof, it is readily verified that these strategies constitute an equilibrium, with the voter indifferent after any pair of platforms, and hence welfare is given by (7).

For small $\varepsilon > 0$ and $x_A \in \bar{X}_A$ and $x_B \in \bar{X}_B$, the same logic also holds for platforms $x_A + \varepsilon$ and $x_B - \varepsilon$, yielding

$$\bar{s}_B(x_B - \varepsilon) = \frac{\alpha + 2\beta}{2\beta}(x_A + x_B) - \bar{s}_A(x_A + \varepsilon). \tag{12}$$

Substituting $x'_B = x_B - \varepsilon$ and $x'_A = x_A$ into (11) and then equating that with (12) yields

$$\frac{\alpha + 2\beta}{2\beta}(x_A + x_B - \varepsilon) - \bar{s}_A(x_A) = \frac{\alpha + 2\beta}{2\beta}(x_A + x_B) - \bar{s}_A(x_A + \varepsilon),$$

or equivalently,

$$\bar{s}_A(x_A + \varepsilon) = \frac{\alpha + 2\beta}{2\beta}\varepsilon + \bar{s}_A(x_A). \tag{13}$$

The equality in (13) can only hold for all $x_A \in \bar{X}_A$ and small $\varepsilon > 0$ if there is a constant $c_A \in \mathbb{R}$ such that $\bar{s}_A(x_A) = \frac{\alpha+2\beta}{2\beta}x_A + c_A$ for all $x_A \in \bar{X}$, from which it follows that $\bar{y}_A(s_A) = \frac{2\beta}{\alpha+2\beta}s_A + c_A$ for all s_A. A symmetric argument establishes that $\bar{y}_B(s_B) = \frac{2\beta}{\alpha+2\beta}s_B + c_B$ for all s_B. But then (11) implies $c_B = -c_A$. Finally, note that continuity pins down the strategies even at measure zero sets of signals. □

A.3. The Benefits of Pandering

Proposition 6 below provides a result that is stronger than Proposition 3. Before that, we record the following.

Lemma 2. *Assume the normal-quadratic specification. The strategy (8) exhibits pandering. Moreover, if both candidates play (8), then the voter's best response is to elect the candidate with the more extreme platform.*

Proof. We first prove that the strategy (8) exhibits pandering. Using (5) and iterated expectations, and dropping the subscript on y_i for the reminder of this proof (since the strategy is common to both candidates), we can rewrite (8) as

$$\begin{aligned} y(s_i) &= \mathbb{E}\big[\mathbb{E}[\theta \mid s_i, s_{-i}] \mid s_i, |s_{-i}| \leq |s_i|\big] \\ &= \frac{\beta}{\alpha + 2\beta}\Big(s_i + \mathbb{E}\big[s_{-i} \mid s_i, |s_{-i}| \leq |s_i|\big]\Big). \end{aligned} \tag{14}$$

Plainly $y(0) = 0$. We will argue that if $s_i > 0$ then $y(s_i) \in \big(0, \mathbb{E}[\theta \mid s_i]\big)$. By a symmetric argument for $s_i < 0$, it follows that $y(\cdot)$ exhibits pandering.

Accordingly, fix any $s_i > 0$. It is straightforward from (14) that $y(s_i) > 0$. Next, algebraic manipulation of (14) and the equalities $\mathbb{E}[\theta \mid s_i] = \mathbb{E}[s_{-i} \mid s_i] = \frac{\beta}{\alpha+\beta} s_i$ shows that $y(s_i) < \mathbb{E}[\theta \mid s_i]$ is equivalent to

$$\mathbb{E}\big[s_{-i} \mid s_i, |s_{-i}| \leq |s_i|\big] < \mathbb{E}[s_{-i} \mid s_i].$$

This inequality holds because $s_{-i} \mid s_i$ is normally distributed with a mean $\frac{\beta}{\alpha+\beta} s_i > 0$, and a truncation to the interval $[-s_i, s_i]$ which is symmetric around 0 (hence centered below the mean) pulls the truncated mean towards 0.

Now we turn to the voter's best response. Define the function $l : \mathbb{R} \to \mathbb{R}$ by

$$l(s_i) := s_i - \mathbb{E}\big[s_{-i} \mid s_i,\ |s_{-i}| \leq |s_i|\big].$$

Using (5) again and the formulae above, some algebra yields

$$\mathbb{E}[\theta \mid s_A, s_B] - \frac{y_A(s_A) + y_B(s_B)}{2} = \frac{\beta}{2(\alpha + 2\beta)} \Big(l(s_A) + l(s_B) \Big). \tag{15}$$

Since strategy (8) is fully revealing, and under quadratic loss it is optimal for the voter to elect A over B if and only if

$$(y_A(s_A) - y_B(s_B)) \Big(\mathbb{E}[\theta \mid s_A, s_B] - \frac{y_A(s_A) + y_B(s_B)}{2} \Big) \geq 0,$$

Equation 15 implies that it is optimal for the voter to elect A if and only if

$$(y(s_A) - y(s_B)) \big(l(s_A) + l(s_B) \big) \geq 0. \tag{16}$$

Note that both $y\left(s_i\right)$ and $l\left(s_i\right)$ are odd and strictly increasing in s_i. Oddness follows from symmetry of the joint signal distribution around zero, while monotonicity follows because $s_{-i} \mid s_i$ is normally distributed with mean $\mathbb{E}\left[s_{-i} \mid s_i\right] = \frac{\beta}{\alpha+\beta} s_i$ and symmetric truncation on $\big[-|s_i|, |s_i|\big]$ preserves monotone dependence on s_i (by log-concavity of the normal density). Hence, when the realized s_A and s_B have the same sign, the term $l(s_A) + l(s_B)$ shares that sign, so whether inequality (16) holds is determined by the sign of $y\left(s_A\right) - y\left(s_B\right)$. When instead the realized signals have opposite signs, $y(s_A) - y(s_B)$ has the same sign as s_A (since y is sign-preserving), so whether the inequality holds is determined by the sign of $l(s_A) + l(s_B) = l(s_A) - l(-s_B)$. By oddness and monotonicity of both y and l, it follows that in both cases, inequality (16) holds if and only if $|s_A| \geq |s_B|$, which—because $|y(s)|$ is strictly increasing in $|s|$—is equivalent to $|y(s_A)| \geq |y(s_B)|$. That is, the voter elects the candidate with the more extreme platform. □

Lemma 2 implies that if both candidates pander using strategy (8) and the voter best responds, then the candidate with the more extreme signal wins. With that in mind, we now state the following result.

Proposition 6. *Assume the normal-quadratic specification. Consider the symmetric strategy profile in which each candidate i panders by playing (8) and the voter best responds. This profile maximizes voter welfare among all strategy profiles in which the voter's best response would lead to candidate i winning whenever $|s_i| > |s_{-i}|$.*

The intuition for Proposition 6 is as follows. The welfare-maximizing platform given any information $\mathcal{I}$ is $\mathbb{E}[\theta \mid \mathcal{I}]$. When the voter is selecting the candidate with the most extreme signal, the relevant information that candidate i has when he conditions on winning is his own signal, s_i, and that $|s_i| > |s_{-i}|$. Since the voter would optimally elect the candidate with the most extreme signal if both candidates used unbiased strategies, an implication of Proposition 6 is that both candidates playing the pandering strategy (8) provides higher voter welfare than both candidates playing unbiased strategies (and the voter best responding in each case), and hence also over any equilibrium—which is the statement of Proposition 3.

Proof of Proposition 6. By the law of iterated expectations, the voter's ex-ante utility can be expressed as

$$v(y_A, y_B, w_A) = -\mathbb{E}[(x-\theta)^2] = -\mathbb{E}[\mathbb{E}[(x-\theta)^2 \mid s_A, s_B]] = -\mathbb{E}\left[\left(x - \frac{\beta(s_A+s_B)}{\alpha+2\beta}\right)^2\right] - \frac{1}{\alpha+2\beta}$$

$$\begin{aligned} &= -\Pr(A \text{ wins})\, \mathbb{E}\left[\left(x_A - \frac{\beta(s_A+s_B)}{\alpha+2\beta}\right)^2 \,\middle|\, A \text{ wins}\right] \\ &\quad - \Pr(B \text{ wins})\, \mathbb{E}\left[\left(x_B - \frac{\beta(s_A+s_B)}{\alpha+2\beta}\right)^2 \,\middle|\, B \text{ wins}\right] - \frac{1}{\alpha+2\beta}. \end{aligned} \tag{17}$$

It is convenient to define $h_i(s_i) := \mathbb{E}\left[s_{-i} \mid s_i, i \text{ wins}\right]$. Using iterated expectations again and a mean-variance decomposition as in the proof of Proposition 5, it also holds that for any i,

$$\begin{aligned}
&\mathbb{E}\left[\left(x_i - \frac{\beta\left(s_A + s_B\right)}{\alpha + 2\beta}\right)^2 \,\middle|\, i \text{ wins}\right] \\
&\quad = \mathbb{E}\left[\mathbb{E}\left[\left(x_i - \frac{\beta\left(s_A + s_B\right)}{\alpha + 2\beta}\right)^2 \,\middle|\, s_i, i \text{ wins}\right] \,\middle|\, i \text{ wins}\right] \\
&\quad = \mathbb{E}\left[\left(x_i - \frac{\beta\left(s_i + \mathbb{E}\left[s_{-i} \mid s_i, i \text{ wins}\right]\right)}{\alpha + 2\beta}\right)^2 + \left(\frac{\beta}{\alpha + 2\beta}\right)^2 \operatorname{Var}\left[s_{-i} \mid s_i, i \text{ wins}\right] \,\middle|\, i \text{ wins}\right] \\
&\quad = \mathbb{E}\left[\left(x_i - \frac{\beta\left(s_i + h(s_i)\right)}{\alpha + 2\beta}\right)^2 \,\middle|\, i \text{ wins}\right] + \left(\frac{\beta}{\alpha + 2\beta}\right)^2 \mathbb{E}\left[\operatorname{Var}\left[s_{-i} \mid s_i, i \text{ wins}\right] \,\middle|\, i \text{ wins}\right].
\end{aligned} \tag{18}$$

Equations (17) and (18) imply

$$v(y_A, y_B, w_A) = -\left(\frac{\beta}{\alpha + 2\beta}\right)^2 L_V - L_E - \frac{1}{\alpha + 2\beta}, \tag{19}$$

where

$$L_V := \sum_{i=A,B} \Pr\left(i \text{ wins}\right) \mathbb{E}\left[\operatorname{Var}\left[s_{-i} \mid s_i, i \text{ wins}\right] \,\middle|\, i \text{ wins}\right], \tag{20}$$

$$L_E := \sum_{i=A,B} \Pr\left(i \text{ wins}\right) \mathbb{E}\left[\left(x_i(s_i) - \frac{\beta\left(s_i + h(s_i)\right)}{\alpha + 2\beta}\right)^2 \,\middle|\, i \text{ wins}\right]. \tag{21}$$

Our problem is to maximize (19) subject to i winning when $|s_i|| > |s_{-i}|$. Since (20) does not depend on platforms while (21) is bounded below by 0, a solution must satisfy for each i:

$$y_i(s_i) = \frac{\beta\left(s_i + h(s_i)\right)}{\alpha + 2\beta} = \mathbb{E}[\theta \mid s_i, i \text{ wins}].$$

Since the constraint is that i wins when $|s_i| > |s_{-i}|$, it follows immediately that the solution is for each candidate to use the strategy (8). $\square$

We remark that although we do not have a proof, we conjecture that Proposition 6 holds without the qualification that a candidate must win when he has the more extreme signal.[31]

A.4. Anti-Pandering Beyond the Normal-Normal Structure

The existence of an anti-pandering equilibrium like the one characterized in Proposition 2 holds beyond our normal-normal informational structure. The simplest extension is to an asymmetric normal-normal specification in which, conditional on the state θ, each candidate i receives an independent signal $s_i \sim \mathcal{N}(\theta, 1/\beta_i)$, with different precisions β_i. In this case, the anti-pandering equilibrium strategy takes the form $y_i(s_i) = \frac{2\beta_i}{\alpha+\beta_A+\beta_B} s_i$.

More generally, maintaining quadratic loss voter preferences, a fully-revealing anti-pandering equilibrium exists when the distributions of the state θ and signals s_i are conjugate and belong to an exponential family. The Supplementary Appendix explicitly derives such an equilibrium in a Beta-prior–Bernoulli-signals specification and shows that it has characteristics analogous to that of Subsection 3.2. The key general property of an exponential family is that the posterior expectation $\mathbb{E}[\theta \mid s_0, s_1, \ldots, s_n]$ of the state θ given a prior mean parameter, say s_0, and any number of signal realizations, $s_1, \ldots, s_n$, is linear in $s_0, s_1, \ldots$ and s_n, (Jewel, 1974). In our Downsian framework, suppose the two candidates' signals s_A and s_B are identically distributed conditional on the state θ. (Identical distributions are not necessary, but make the points below more transparent.) Then, there are constants w_0 and w_1 such that

$$\mathbb{E}[\theta \mid s_i] = \frac{w_0 s_0 + w_1 s_i}{w_0 + w_1} \quad \text{and} \quad \mathbb{E}[\theta \mid s_A, s_B] = \frac{w_0 s_0 + 2w_1 \left((s_A + s_B)/2 \right)}{w_0 + 2w_1}.$$

As a result, the following generalization of the existence result of Proposition 2 can be veri-

[31] For a suggestive heuristic, consider any symmetric strategy profile in which both candidates play the same strategy y that is symmetric around 0. For the unbiased strategy, we have the derivative $y'(\cdot) = \frac{\beta}{\beta+\alpha}$; for the overreaction strategy identified in Proposition 2, we have $y'(\cdot) = \frac{2\beta}{\alpha+2\beta}$. Presuming differentiability, one can verify that whenever $y'(\cdot) \in [0, \frac{2\beta}{\alpha+2\beta}]$, it would be optimal for the voter to elect the candidate with the most extreme platform and hence the most extreme signal. Thus, roughly speaking, the requirement that a candidate wins when he has the most extreme signal is satisfied as long as neither candidate overreacts by more than he would when conditioning on the opponent having received the same signal as he did. It appears unlikely that such a degree of overreaction could improve voter welfare.

fied:[32] there is an equilibrium with overreaction in which each candidate i plays

$$y_i(s_i) = \frac{2w_1}{w_0 + 2w_1} s_i + \frac{w_0}{w_0 + 2w_1} s_0,$$

and the voter randomizes uniformly after any pair of on-path platforms.[33]

B. Completeness and Strong Linear Independence

Using "signal" as a synonym for "type", recall that we noted in the main text after introducing Condition 1 that it is equivalent to a full row and column rank condition for finite signal spaces. This appendix clarifies that relationship more generally, with its main result being Proposition 7 below. The setting and notation for this appendix follow Section 2.

B.1. Strong Linear Independence

Let $\mathcal{B}(S_i)$ denote the Borel σ-algebra on S_i and let $\mathcal{M}(S_i)$ denote the set of finite signed measures on S_i. Define the linear operator $K_i : \mathcal{M}(S_i) \to \mathcal{M}(S_{-i})$ by setting, for any $\mu \in \mathcal{M}(S_i)$ and $B \in \mathcal{B}(S_{-i})$,

$$(K_i\mu)(B) \;:=\; \int_{S_i} F(B \mid s_i)\, \mu(\mathrm{d}s_i).$$

For any $\mu \in \mathcal{M}(S_i)$, the map $s_i \mapsto F(B \mid s_i)$ is bounded and measurable, so $(K_i\mu)(B)$ is a well-defined Lebesgue integral for each $B \in \mathcal{B}(S_{-i})$. Moreover, $K_i\mu$ is a finite signed measure because μ is finite and $F(\cdot \mid s_i)$ is a probability measure for each s_i.

Let $\mathcal{M}_{ac}(S_i) \subset \mathcal{M}(S_i)$ denote the subset of finite signed measures that are dominated by (i.e., absolutely continuous with respect to) F_i. Writing, as usual, that a measure $\mu = 0$ means $\mu(\cdot) = 0$, we define:

---

[32] As the prior density need no longer be symmetric around the mean (unlike with a normal prior) and signals may be bounded (unlike with normally distributed signals), the definitions of anti-pandering or overreaction have to be broadened from earlier. We now say that a strategy y_i has overreaction if for all s_i, $|y_i(s_i) - \mathbb{E}[\theta]| \geq |\mathbb{E}[\theta \mid s_i] - \mathbb{E}[\theta]|$ with strict inequality for some s_i. The focus on posterior expectations of the state is justified when the voter has a quadratic loss function. See the discussion in Roux and Sobel (2015) to get a sense of how asymmetric loss functions would affect the conclusions.

[33] While there may now be off-path platforms (unlike with normal distributions), as in the Beta-Bernoulli example in the Supplementary Appendix, the equilibrium can be supported with reasonable off-path beliefs.

Definition 2. The family $\{F(\cdot \mid s_i)\}_{s_i \in S_i}$ satisfies *Strong Linear Independence* (SLI) for player i if for every $\mu \in \mathcal{M}_{ac}(S_i)$, we have

$$K_i \mu = 0 \implies \mu = 0.$$

For short, we will say that there is *SLI for i* when the family $\{F(\cdot \mid s_i)\}_{s_i \in S_i}$ satisfies SLI for player i. SLI requires that no nonzero finite signed measure that is dominated by F_i creates a zero mixture of the conditional distributions $\{F(\cdot \mid s_i)\}_{s_i \in S_i}$. When S_i is countable, any measure is dominated by F_i (which has support S_i), so SLI is stronger than textbook linear independence because SLI does not restrict to finite mixtures.[34] SLI is equivalent to linear independence for finite S_i.

B.2. Completeness

As usual, write $L^1(F_i)$ for the measurable and F_i-integrable real-valued functions on S_i.

Definition 3. Player i's signals have a *complete* family of conditional distributions $\{F(\cdot \mid s_i)\}_{s_i \in S_i}$ if for every $g \in L^1(F_{-i})$ it holds that

$$\mathbb{E}[g(s_{-i}) \mid s_i] = 0 \text{ for } F_i\text{-a.e. } s_i \implies F\left(g(s_{-i}) = 0 \mid s_i\right) = 1 \text{ for } F_i\text{-a.e. } s_i.$$

For short, we will say that there is *completeness for i* when player i's signals have a complete family of conditional distributions. Completeness (when required for both players) is closely related to Condition 1 and captures the same idea; it is, however, slightly stronger in general because it allows for unbounded test functions g. Unbounded test functions are irrelevant when S_{-i} is finite, but we will see that they are essential for the connection of completeness to SLI when S_{-i} is infinite. When Definition 3 is restricted to bounded test functions g, we say that there is *bounded completeness* for i.

[34] Recall that even when S_i is infinite, $\{F(\cdot \mid s_i)\}_{s_i \in S_i}$ is *linearly independent* if every nonzero *finite* linear combination is nontrival, i.e., for any finite set $\{s_i^1, \ldots, s_i^K\} \subset S_i$ and any $c : \{1, \ldots, K\} \to \mathbb{R}$, it holds that

$$\sum_{k=1}^{K} c(k) F(\cdot \mid s_i^k) = 0 \implies c(\cdot) = 0.$$

B.3. Equivalence

Completeness for i concerns how informative s_i is about the *other player* $-i$'s signal. So, even with finite signal sets, completeness for i corresponds to linear independence of $\{F(\cdot \mid s_{-i})\}_{s_{-i} \in S_{-i}}$, not of i's own conditional distributions. Indeed, with finite signal sets for both players (albeit of different cardinality), we can have completeness for i but a failure of linear independence of $\{F(\cdot \mid s_i)\}_{s_i \in S_i}$, and vice-versa.[35]

With that in mind, we have the following equivalence.

Proposition 7. *For each player $i \in \{A, B\}$, it holds that*

$$\textit{Completeness for } i \iff \textit{Strong Linear Independence for } -i.$$

Proof. Fix an arbitrary player $i \in \{A, B\}$ and write j for $-i$.

(Completeness for $i \implies$ SLI for j). Assume completeness for i. Let $\mu \in \mathcal{M}_{ac}(S_j)$ such that $K_j\mu = 0$. We must show that $\mu = 0$.

Let $g(s_j) := \mathrm{d}\mu/\mathrm{d}F_j(s_j)$ be the Radon–Nikodym derivative of μ with respect to F_j, which exists because μ is dominated by F_j. Fix any measurable $A \subset S_i$. We have

$$\begin{aligned}
(K_j\mu)(A) &= \int_{S_j} F(A \mid s_j)\, g(s_j)\, F_j(\mathrm{d}s_j) \\
&= \mathbb{E}[\, \mathbb{1}_A(s_i)\, g(s_j)\,] \quad \text{using } F(A \mid s_j) = \mathbb{E}[\mathbb{1}_A(s_i) \mid s_j] \text{ and the law of iterated expectations} \\
&= \int_A \mathbb{E}[\, g(s_j) \mid s_i\,]\, F_i(\mathrm{d}s_i).
\end{aligned}$$

Since $K_j\mu = 0$, the left-hand side above is zero for all measurable $A \subset S_i$, and hence

$$\mathbb{E}[\, g(s_j) \mid s_i\,] = 0 \text{ for } F_i\text{-a.e. } s_i.$$

By completeness for i, it follows that

$$F(g(s_j) = 0 \mid s_i) = 1 \text{ for } F_i\text{-a.e. } s_i.$$

[35] In particular, linear independence fails for i when S_i has duplicate signals, while there is completeness for i when $|S_{-i}| = 1$. Conversely, if $|S_{-i}| > |S_i| = 1$, there is linear independence for i but not completeness for i.

Integrating over s_i with respect to F_i and using the law of total probability yields $g(s_j) = 0$ F_j-a.s., and therefore $\mu = 0$.

(SLI for $j \implies$ Completeness for i). Assume SLI for j. Let $g \in L^1(F_j)$ satisfy

$$\mathbb{E}[\, g(s_j) \mid s_i \,] = 0 \text{ for } F_i\text{-a.e. } s_i. \tag{22}$$

Define $\mu \in \mathcal{M}_{ac}(S_j)$ by

$$\mu(B) := \int_B g(s_j)\, F_j(\mathrm{d}s_j) \quad \text{ for any } B \in \mathcal{B}(S_j). \tag{23}$$

Fix any measurable $A \subset S_i$. We have

$$\begin{aligned}(K_j\mu)(A) &= \int_{S_j} F(A \mid s_j)\, g(s_j)\, F_j(\mathrm{d}s_j) \\ &= \mathbb{E}[\, \mathbb{1}_A(s_i)\, \mathbb{E}[\, g(s_j) \mid s_i \,]\,] \quad \text{using } F(A \mid s_j) = \mathbb{E}[\mathbb{1}_A(s_i) \mid s_j] \text{ and the law of iterated expectations} \\ &= 0 \qquad \text{by (22).}\end{aligned}$$

Thus $K_j\mu = 0$. SLI implies $\mu = 0$, which by (23) implies $g(s_j) = 0$ F_j-a.s. The law of total probability now implies $F(g(s_j) = 0 \mid s_i) = 1$ for F_i-a.e. s_i, which is completeness for i. □

Proposition 7 formalizes the sense in which SLI is the appropriate infinite-dimensional analog of full rank (of the other player's conditional distributions) for completeness. With infinite (even countable) signal spaces, finite linear combinations are no longer sufficient, as seen in Example 4 below. SLI instead requires injectivity of the linear operator mapping (absolutely continuous) finite signed measures on one player's signal space to mixtures of the corresponding conditional distributions.

The following example with countable signals illustrates how the relevant directions of Proposition 7 rely on completeness rather that bounded completeness and on SLI rather than linear independence.

Example 4. Consider $S_A = \{0, 1, \ldots\}$ and $S_B = \{1, 2, \ldots\}$. While we could use any signal distribution F that has a full-support marginal F_A, for concreteness take $F_A(0) = 1/2$ and

$F_A(s_A) = 3^{-s_A}$ for $s_A > 0$, and let the conditional distributions be

$$F(s_B \mid s_A) = \begin{cases} 2^{-s_B} & \text{if } s_A = 0 \\ \mathbb{1}\{s_B = s_A\} & \text{if } s_A > 0. \end{cases}$$

Bounded completeness holds for player B: for any signal s_B, the conditional distribution $F(\cdot \mid s_B)$ is supported on $\{0, s_B\}$; thus, if a bounded function $g(s_A)$ has $\mathbb{E}[g(s_A) \mid s_B] = 0$ for all s_B, then $g(0) = 0$ by considering large enough s_B, and consequently $g(\cdot) = 0$.

The family $\{F(\cdot \mid s_A)\}_{s_A \in S_A}$ does not satisfy SLI because for any $s_B \in S_B$, we have

$$F(s_B \mid 0) + \sum_{s_A=1}^{\infty} \left(-2^{-s_A}\right) F(s_B \mid s_A) = 2^{-s_B} - 2^{-s_B} = 0.$$

The family $\{F(\cdot \mid s_A)\}_{s_A \in S_A}$ is, however, linearly independent: if $\sum_{s_A=0}^{K} c(s_A) F(\cdot \mid s_A) = 0$ for any integer K, then $c(0) = 0$ by considering $s_B > K$, and consequently $c(\cdot) = 0$.

These observations show that bounded completeness for B does not imply SLI for A, and linear independence for A does not imply completeness for B (since SLI failing for A implies that completeness fails for B, by Proposition 7).[36] □

B.4. Insufficiency of Convex Independence

In their work on informational richness in mechanism design, Crémer and McLean (1988) discussed the role of both linear independence (for dominant-strategy mechanisms) and also convex independence (for Bayesian-incentive-compatible mechanisms) with finite types; McAfee and Reny (1992) extended the latter analysis to an infinite setting.

Even with finite signals, convex independence—being weaker than linear independence—is insufficient for our purposes. Formally, say that *convex independence* holds for player $i \in \{A, B\}$ if for all $s_i \in S_i$,

$$F(\cdot \mid s_i) \notin \operatorname{co}\left(\{F(\cdot \mid t_i) : t_i \in S_i \setminus \{s_i\}\}\right),$$

[36] The failure of completeness can also be directly verified using the test function g given by $g(0) = 1$ and $g(s_A) = -(1/2)(3/2)^{s_A}$ for $s_A > 0$.

where co$(\cdot)$ denotes the convex hull.

Consider the following joint distribution when each player has 4 signals:

$$F = \frac{1}{40}\begin{pmatrix} 4 & 2 & 1 & 3 \\ 2 & 4 & 3 & 1 \\ 1 & 3 & 4 & 2 \\ 3 & 1 & 2 & 4 \end{pmatrix}.$$

This is a symmetric distribution in which the marginal distribution of a player's signal is uniform. Hence, the matrix of conditional distributions for either player is just a rescaling of the above matrix (multiplying it by 4). For either player i, the family of conditional distributions $\{F(\cdot \mid s_i)\}_{s_i \in S_i}$ is convexly independent—each conditional distribution $F(\cdot \mid s_i)$ assigns highest probability to a distinct opponent signal—yet completeness (or equivalently here, bounded completeness) fails because the conditional matrix has rank 3 (as the sum of the first and third rows of the F matrix equals the sum of the second and fourth rows).

The conclusion of Theorem 1 fails under this joint distribution:

Example 5. Consider the above signal structure, labeling signals as $S_i = \{1, 2, 3, 4\}$ in the natural way. Let $X_A = X_B = \{0, 1\}$ and $u_A(x_A, x_B) = -\mathbb{1}\{x_A = x_B\}$. There is an identifiable equilibrium in which each player takes action 1 when he receives signal 1 or 3, and takes action 0 otherwise; the signal structure implies that each type of each player then faces a uniform distribution over the opponent's actions. This is, however, not an ex-post equilibrium. □

C. On Corollary 1

Here we make precise the connection between Corollary 1 and Kattwinkel et al. (2022, Proposition 3, part 2). They consider direct mechanisms, so $X_i = S_i$, and an outcome space $\Omega = [0, 1]$, interpreted as an allocation probability. Their primitive is preferences over outcomes given by $\tilde{u}_A(\omega) = \omega$ and $\tilde{u}_B(\omega) = -\omega$. Given any mechanism (or outcome function) w, the induced preferences are $u_i(x_A, x_B) := \tilde{u}_i(w(x_A, x_B))$, and hence there are strict preferences over outcomes. Kattwinkel et al. ask which mechanisms are incentive compatible in the sense that truthful reporting (i.e., the strategy $s_i \mapsto s_i$) forms an equilibrium. Since such strategies are identifiable (being pure), Corollary 1 implies that under Condition 1 only constant mechanisms are incentive compatible. This subsumes Kattwinkel et al. (2022, Proposition 3, part 2), which assumes finite type sets, in which case Condition 1 reduces to their full-rank condition.[37]

In the other direction, Kattwinkel et al.'s (2022) result can be combined with a revelation-principle argument to derive Corollary 1 for finite type sets and pure-strategy equilibria. However, to deal with infinite type sets or (identifiable) mixed-strategy equilibria, we believe an argument like ours is needed.

References

Acemoglu, Daron, Georgy Egorov, and Konstantin Sonin, "A Political Theory of Populism," *Quarterly Journal of Economics*, 2013, *128* (2), 771–805.

Alesina, Alberto, "Credibility and Policy Convergence in a Two-Party System with Rational Voters," *American Economic Review*, September 1988, *78* (4), 796–805.

Althaus, Scott L., "Information Effects in Collective Preferences," *American Political Science Review*, 1998, *92* (3), 545–558.

Ambrus, Attila, Volodymyr Baranovskyi, and Aaron Kolb, "A Delegation-Based Theory of Expertise," *American Economic Journal: Microeconomics*, 2021, *13* (4), 373–419.

Aragones, Enriqueta and Thomas Palfrey, "Mixed Equilibrium in a Downsian Model with a Favored Candidate," *Journal of Economic Theory*, 2002, *103*, 131–161.

Bernhardt, Dan, John Duggan, and Francesco Squintani, "Electoral competition with privately-informed candidates," *Games and Economic Behavior*, January 2007, *58* (1), 1–29.

_ , _ , and _ , "Private Polling in Elections and Voter Welfare," *Journal of Economic Theory*, September 2009, *144* (5), 2021–2056.

Besley, Tim and Stephen Coate, "An Economic Model of Representative Democracy," *Quarterly Journal of Economics*, 1997, *112*, 85–114.

[37] To be more precise: Condition 1 is equivalent to full rank when $|S_A| = |S_B| < \infty$. Since Condition 1 is stated as applying to both players, it cannot hold with finite type sets when $|S_A| \neq |S_B|$, because it requires full row and column rank. However, as seen in the proof of Theorem 1, the theorem—and hence also Corollary 1—only requires that Condition 1 hold for one player i when player $-i$'s equilibrium strategy is identifiable. So when both players' strategies are identifiable, as in the present discussion, it is sufficient that Condition 1 hold for either player. With finite type sets, that reduces to usual full rank (i.e., either full row or column rank).

Bils, Peter, "Overreacting and Posturing: How Accountability and Ideology Shape Executive Policies," *Quarterly Journal of Political Science*, 2023, *18* (2), 153–182.

Callander, Steven and Simon Wilkie, "Lies, Damned Lies, and Political Campaigns," *Games and Economic Behavior*, 2007, *60* (2), 262–286.

Calvert, Randall L., "Robustness of the Multidimensional Voting Model: Candidate Motivations, Uncertainty, and Convergence," *American Journal of Political Science*, 1985, *29*, 69–95.

Campante, Filipe R., "Redistribution in a model of voting and campaign contributions," *Journal of Public Economics*, 2011, *95* (7), 646–656.

Campbell, Angus, Philip E. Converse, Warren E. Miller, and Donald E. Stokes, *The American Voter*, Chicago: University of Chicago Press, 1960.

Canes-Wrone, Brandice, Michael Herron, and Kenneth W. Shotts, "Leadership and Pandering: A Theory of Executive Policymaking," *American Journal of Political Science*, July 2001, *45* (3), 532–550.

Che, Yeon-Koo, Wouter Dessein, and Navin Kartik, "Pandering to Persuade," *American Economic Review*, February 2013, *103* (1), 47–79.

Crémer, Jacques and Richard P. McLean, "Optimal Selling Strategies Under Uncertainty for a Discriminating Monopolist When Demands Are Interdependent," *Econometrica*, 1985, *53* (2), 345–361.

_ **and** _ , "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, 1988, *56* (6), 1247–1257.

Downs, Anthony, *An Economic Theory of Democracy*, New York: Harper and Row, 1957.

Durrett, Richard, *Probability: Theory and Examples*, 2nd ed., Duxbury Press, 1995.

Fishkin, James S., *The Voice of the People: Public Opinion and Democracy*, New Haven: Yale University Press, 1997.

Gilens, Martin, "Political Ignorance and Collective Policy Preferences," *American Political Science Review*, 2001, *95* (2), 379–396.

Glaeser, Edward L. and Cass R. Sunstein, "Extremism and Social Learning," *Journal of Legal Analysis*, Winter 2009, *1* (1), 263–324.

___, **Giacomo A. M. Ponzetto, and Jesse M. Shapiro**, "Strategic Extremism: Why Republicans and Democrats Divide on Religious Values," *Quarterly Journal of Economics*, November 2005, *120* (4), 1283–1330.

Gratton, Gabriele, "Pandering and Electoral Competition," *Games and Economic Behavior*, 2014, *84*, 163–179.

Groseclose, Timothy, "A Model of Candidate Location when One Candidate Has a Valence Advantage," *American Journal of Political Science*, 2001, *45*, 862–886.

Heidhues, Paul and Johan Lagerlof, "Hiding Information in Electoral Competition," *Games and Economic Behavior*, January 2003, *42* (1), 48–74.

Honryo, Takakazu, "Risky Shifts as Multi-sender Signaling," *Journal of Economic Theory*, 2018, *174*, 273–87.

Hotelling, Harold, "Stability in Competition," *Economic Journal*, March 1929, *39* (153), 41–57.

Iyengar, Shanto and Donald Kinder, *News That Matters: Television and American Opinion*, Chicago: University of Chicago Press, 1987.

Jehiel, Philippe, Moritz Meyer ter Vehn, Benny Moldovanu, and William R. Zame, "The Limits of Ex Post Implementation," *Econometrica*, 2006, *74* (3), 585–610.

Jewel, William S., "Credible Means are Exact Bayesian for Exponential Families," *ASTIN Bulletin*, 1974, *8*, 77–90.

Kartik, Navin and R. Preston McAfee, "Signaling Character in Electoral Competition," *American Economic Review*, June 2007, *97* (3), 852–870.

___ **and Richard Van Weelden**, "Informative Cheap Talk in Elections," *Review of Economic Studies*, 2019, *86*, 755–784.

___, **Francesco Squintani, and Katrin Tinn**, "Information Revelation and (Anti-)Pandering in Elections," April 2015. unpublished.

___, **SangMok Lee, and Daniel Rappoport**, "Single-Crossing Differences in Convex Environments," *Review of Economic Studies*, 10 2023, *91* (5), 2981–3012.

Kattwinkel, Deniz, Axel Niemeyer, Justus Preusser, and Alexander Winter, "Mechanisms without Transfers for Fully Biased Agents," 2022. arXiv:2205.10910.

Klumpp, Tilman, "Populism, Partisanship, and the Funding of Political Campaigns," August 2014. unpublished.

Laslier, Jean-François and Karine Van de Straeten, "Electoral Competition under Imperfect Information," *Economic Theory*, August 2004, *24* (2), 419–446.

Lehmann, Erich L., *Testing Statistical Hypotheses*, 2nd ed., Wiley, 1986.

Levy, Gilat, "Anti-herding and Strategic Consultation," *European Economic Review*, June 2004, *48* (3), 503–525.

Loertscher, Simon, "Location Choice and Information Transmission," 2012. mimeo, University of Melbourne.

Majumdar, Sumon and Sharun W. Mukand, "Policy Gambles," *American Economic Review*, September 2004, *94* (4), 1207–1222.

Martinelli, Cesar, "Elections with Privately Informed Parties and Voters," *Public Choice*, July 2001, *108* (1–2), 147–67.

_ **and Akihiko Matsui**, "Policy Reversals and Electoral Competition with Privately Informed Parties," *Journal of Public Economic Theory*, 2002, *4* (1), 39–61.

Maskin, Eric and Jean Tirole, "The Politician and the Judge: Accountability in Government," *American Economic Review*, September 2004, *94* (4), 1034–1054.

McAfee, R. Preston and Philip J. Reny, "Correlated Information and Mechanism Design," *Econometrica*, 1992, *60* (2), 395–421.

Millner, Antony, Hélène Ollivier, and Leo Simon, "Confirmation Bias and Signaling in Downsian Elections," *Journal of Public Economics*, 2020, *185*, 104–175.

Nash, John, "Non-Cooperative Games," *Annals of Mathematics*, 1951, *54* (2), 286–295.

Osborne, Martin J. and Al Slivinski, "A Model of Political Competition with Citizen-Candidates," *Quarterly Journal of Economics*, 1996, *111*, 65–96.

Ottaviani, Marco and Peter Norman Sorensen, "Reputational Cheap Talk," *RAND Journal of Economics*, Spring 2006, *37* (1), 155–175.

Palfrey, Thomas, "Spatial Equilibrium with Entry," *Review of Economic Studies*, 1984, *51*, 139–156.

Pei, Harry and Bruno Strulovici, "Robust Implementation with Costly Information," *Review of Economic Studies*, January 2025, *92* (1), 476–505.

Pennec, Caroline Le and Vincent Pons, "How do Campaigns Shape Vote Choice? Multi-country Evidence from 62 Elections and 56 TV Debates," *Quarterly Journal of Economics*, 03 2023, *138* (2), 703–767.

Pétry, François and Benoît Collette, "Measuring how political parties keep their promises: a positive perspective from political science," in Louis M. Imbeau, ed., *Do They Walk Like They Talk? Speech and Action in Policy Processes*, Springer, 2009, chapter 5, pp. 65–80.

Prendergast, Canice and Lars Stole, "Impetuous Youngsters and Jaded Old-Timers: Acquiring a Reputation for Learning," *Journal of Political Economy*, December 1996, *104* (6), 1105–34.

Roux, Nicholas and Joel Sobel, "Group Polarization in a Model of Information Aggregation," *American Economic Journal: Microeconomics*, November 2015, *7* (4), 202—32.

Schultz, Christian, "Polarization and Inefficient Policies," *Review of Economic Studies*, April 1996, *63* (2), 331–44.

Schuman, Howard and Stanley Presser, *Questions and Answers in Survey Questions*, San Diego: Sage Publications, 1981.

Viossat, Yannick, "The Geometry of Nash Equilibria and Correlated Equilibria and a Generalization of Zero-Sum Games," SSE/EFI Working Paper Series in Economics and Finance 641, Stockholm School of Economics August 2006.

Von Neumann, John, "On the Theory of Games of Strategy," *Mathematische Annalen*, 1928, *100*, 295–320.

Wittman, Donald, "Candidate Motivation: A Synthesis of Alternatives," *American Political Science Review*, 1983, *77*, 142–157.

_ , "Why Democracies Produce Efficient Results," *Journal of Political Economy*, 1989, *97*, 1395–1424.

Zaller, John R., *The Nature and Origins of Mass Opinion*, New York: Cambridge University Press, 1992.

Supplementary Appendix

D. Mixed Motives

This section substantiates the discussion in Subsection 3.3 of the paper by formally generalizing our main welfare conclusions to a normal-quadratic setting in which candidates are largely but not entirely office motivated. We will establish that when the parameters b_i and ρ_i defined in Equation 9 are sufficiently close to zero for each $i = A, B$, (i) there is an equilibrium that achieves welfare arbitrarily close to the level obtained by efficiently aggregating the signal of only one candidate (Proposition 8 below), and (ii) that welfare is an approximate bound on voter welfare in any equilibrium (Proposition 9 below).

In the context of a normal-quadratic mixed-motivation game, with candidates' payoffs as defined in Equation 9, we say that candidate's i strategy is *unbiased* if

$$y_i(s_i) = \frac{\beta}{\alpha+\beta}s_i + b_i. \tag{24}$$

Note that this refers to candidate i choosing a policy that maximizes *his* preference over policy given his signal, as opposed to the voter's.

Proposition 8. *In the normal-quadratic mixed-motivations game, there is a fully revealing equilibrium in which one candidate i plays the unbiased strategy (24), the other candidate $-i$ plays*

$$y_{-i}(s_{-i}) = s_{-i} - \frac{\alpha+\beta}{\beta}b_i, \tag{25}$$

and the voter elects candidate i no matter the pair of platforms.

Proof. Given the strategies (24) and (25), it follows that

$$\mathbb{E}[\theta \mid x_i, x_{-i}] = \frac{\beta(x_i - b_i)\frac{\alpha+\beta}{\beta} + \beta\left(x_{-i} + \frac{\alpha+\beta}{\beta}b_i\right)}{\alpha+2\beta} = \frac{\alpha x_i + \beta(x_i + x_{-i})}{\alpha+2\beta}.$$

Straightforward algebra then verifies that for any x_i and x_{-i},

$$(x_i - \mathbb{E}[\theta \mid x_i, x_{-i}])^2 < (x_{-i} - \mathbb{E}[\theta \mid x_i, x_{-i}])^2 \iff \beta < \alpha+\beta.$$

Hence it is optimal for the voter to always elect candidate i; clearly the candidates are playing optimally given this strategy for the voter. $\square$

As the equilibrium constructed in Proposition 8 is invariant to ρ_A and ρ_B, it has a number of interesting implications. First, the equilibrium exists when candidates are purely policy-motivated. Second, for $\rho_A = \rho_B = b_A = b_B = 0$, this equilibrium reduces to one that verifies the first statement of Theorem 2. Moreover, by taking $b_A = b_B = 0$ and $\rho_A = \rho_B = 1$, we see that there is also an equilibrium in which one candidate plays the unbiased strategy and always wins when both candidates are benevolent. Hence, the equilibrium of Proposition 8 continuously spans all three polar cases of candidate motivation.

Consider a normal-quadratic game with mixed-motivated candidates parameterized by $(\boldsymbol{\rho}, \boldsymbol{b})$, where $\boldsymbol{\rho} \equiv (\rho_A, \rho_B)$ and $\boldsymbol{b} \equiv (b_A, b_B)$. Let $\mathcal{E}(\boldsymbol{\rho}, \boldsymbol{b})$ denote the set of equilibria in which candidates play pure strategies, for consistency with our baseline model. Given any equilibrium $\sigma \equiv (y_A, y_B, w_A)$, let $v(\sigma)$ be the voter's welfare in this equilibrium. Note that the voter's welfare depends only on the strategies used and not directly on the candidates' motivations. Let $v^*(\boldsymbol{\rho}, \boldsymbol{b}) := \sup\{v(\sigma) : \sigma \in \mathcal{E}(\boldsymbol{\rho}, \boldsymbol{b})\}$ be the supremum of equilibrium voter welfare given candidate motivations. Plainly, $v^*(\mathbf{0}, \mathbf{0})$ is the welfare bound identified by Theorem 2.

Proposition 9. *In the normal-quadratic mixed-motivations game, as* $(\boldsymbol{\rho}, \boldsymbol{b}) \to (\mathbf{0}, \mathbf{0})$*, it holds that* $v^*(\boldsymbol{\rho}, \boldsymbol{b}) \to v^*(\mathbf{0}, \mathbf{0})$.

This result holds despite the equilibrium correspondence not being upper hemicontinuous. Indeed, observe that given any candidates' motivations with $b_A > 0$, there is an equilibrium in which both candidates use the constant strategy $y_i(s_i) = 1/b_A$; this is supported by suitable off-path beliefs such that any candidate whose platform differs from b_A loses for sure. The limit of these candidates' strategies, $\lim_{b_A \to 0} 1/b_A$, is not a valid strategy.

We require two lemmas to prove Proposition 9. Let

$$\mathcal{E}^*(\boldsymbol{\rho}, \boldsymbol{b}) := \{\sigma \in \mathcal{E}(\boldsymbol{\rho}, \boldsymbol{b}) : v(\sigma) = v^*(\boldsymbol{\rho}, \boldsymbol{b})\}$$

be the set of welfare-maximizing equilibria.[38] Given a strategy profile $\sigma \equiv (y_A, y_B, w_A)$ and an $\varepsilon > 0$, let $W^\sigma_\varepsilon(s_A, s_B)$ denote the set of candidates who win with probability at least ε when the signal realizations are s_A, s_B. The first lemma below says that in welfare-maximizing equilibria, a candidate cannot win with non-vanishing probability on a non-negligible set of bounded signal realizations while proposing arbitrarily extreme policies. To state it formally, denote by $\text{proj}_{\hat{S}_i}(E)$ the projection of $E \subset \mathbb{R}^2$ onto $\hat{S}_i \subset \mathbb{R}$.

[38] In what follows, we will proceed as if $\mathcal{E}^*(\boldsymbol{\rho}, \boldsymbol{b})$ is non-empty for all $(\boldsymbol{\rho}, \boldsymbol{b})$. If this is not the case, one can proceed almost identically, just by defining for any $\varepsilon > 0$, $\mathcal{E}^*_\varepsilon(\boldsymbol{\rho}, \boldsymbol{b}) := \{\sigma \in \mathcal{E}^*(\boldsymbol{\rho}, \boldsymbol{b}) : v(\sigma) \geq v^*(\boldsymbol{\rho}, \boldsymbol{b}) - \varepsilon\}$, and then applying the subsequent arguments for a sequence of $\varepsilon \to 0$.

Lemma 3. *Let* $\hat{S}_A \times \hat{S}_B$ *be a bounded set of signals of positive measure, and* $\varepsilon, \eta > 0$. *There exists* $k > 0$ *such that for any* $(\boldsymbol{\rho}, \boldsymbol{b})$, *if* $\sigma \equiv (y_A, y_B, w_A) \in \mathcal{E}^*(\boldsymbol{\rho}, \boldsymbol{b})$ *and there exists a measurable set* $E \subseteq \hat{S}_A \times \hat{S}_B$ *with measure at least* η *such that* $i \in W_\varepsilon^\sigma(s_A, s_B)$ *for almost all* $(s_A, s_B) \in E$, *then* $|y_i(s_i)| < k$ *for almost all* $s_i \in \mathrm{proj}_{\hat{S}_i}(E)$.

Proof. Take any $\sigma \in \mathcal{E}^*(\boldsymbol{\rho}, \boldsymbol{b})$. We have $v(\sigma) \geq -\operatorname{Var}(\theta) = -1/\alpha$, because $-\operatorname{Var}(\theta)$ is the welfare in a trivial equilibrium in which both candidates uninformatively choose policy 0. Since $E \subseteq \hat{S}_A \times \hat{S}_B$ and $\hat{S}_A \times \hat{S}_B$ is bounded, the voter's expected utility $\mathbb{E}[u(x, \theta) \mid s_A, s_B] \to -\infty$ as $|x| \to \infty$, uniformly over $(s_A, s_B) \in E$. Since the voter's utility conditional on any signal profile (in particular, those outside E) is bounded above by zero, if the lemma's conclusion were false then σ would have arbitrarily low welfare, a contradiction. □

The next lemma builds on the previous one to show that in welfare-maximizing equilibria, a candidate's platform cannot diverge on any set of signals of positive measure while still winning with non-vanishing probability.

Lemma 4. *In any sequence of welfare-maximizing equilibria* $\sigma^{\boldsymbol{\rho},\boldsymbol{b}} \equiv (y_A^{\boldsymbol{\rho},\boldsymbol{b}}, y_B^{\boldsymbol{\rho},\boldsymbol{b}}, w_A^{\boldsymbol{\rho},\boldsymbol{b}}) \in \mathcal{E}^*(\boldsymbol{\rho}, \boldsymbol{b})$, *as* $(\boldsymbol{\rho}, \boldsymbol{b}) \to (\mathbf{0}, \mathbf{0})$ *either:*

(1) for some i, $\Pr(i \text{ wins in } \sigma^{\boldsymbol{\rho},\boldsymbol{b}}) \to 0$ *as* $(\boldsymbol{\rho}, \boldsymbol{b}) \to 0$; *or*

(2) for any i *and almost all* s_i, $y_i^{\boldsymbol{\rho},\boldsymbol{b}}(s_i)$ *is bounded.*

Proof. Suppose the lemma is false. Then, without loss, there is a number $\delta > 0$ and a (sub)sequence of $(\boldsymbol{\rho}, \boldsymbol{b}) \to (\mathbf{0}, \mathbf{0})$ with equilibria $\sigma^{\boldsymbol{\rho},\boldsymbol{b}} \in \mathcal{E}^*(\boldsymbol{\rho}, \boldsymbol{b})$ such that: (i) for all $(\boldsymbol{\rho}, \boldsymbol{b})$ and $i \in \{A, B\}$, it holds that $\Pr(i \text{ wins in } \sigma^{\boldsymbol{\rho},\boldsymbol{b}}) > \delta$; and (ii) there exists a bounded set $S'_A \subset \mathbb{R}$ of signals with positive measure such that for every $\bar{s}_A \in S'_A$, either $y_A^{\boldsymbol{\rho},\boldsymbol{b}}(\bar{s}_A) \to +\infty$ or $y_A^{\boldsymbol{\rho},\boldsymbol{b}}(\bar{s}_A) \to -\infty$.

Fix any $\kappa > 0$ and $\varepsilon > 0$. By Lemma 3 (applied contrapositively), for $(\boldsymbol{\rho}, \boldsymbol{b})$ small enough the set of pairs $(s_A, s_B) \in S'_A \times [-\kappa, \kappa]$ such that $A \in W_\varepsilon^{\sigma^{\boldsymbol{\rho},\boldsymbol{b}}}(s_A, s_B)$ has arbitrarily small measure. Hence, by Fubini, there exists a subset $\tilde{S}_A \subseteq S'_A$ of positive measure such that for every $\bar{s}_A \in \tilde{S}_A$, the set of $s_B \in [-\kappa, \kappa]$ for which $A \in W_\varepsilon^{\sigma^{\boldsymbol{\rho},\boldsymbol{b}}}(\bar{s}_A, s_B)$ has arbitrarily small measure. Fix any such $\bar{s}_A$.

Since the distribution of $s_B \mid \bar{s}_A$ does not change with $(\boldsymbol{\rho}, \boldsymbol{b})$, and since the κ above can be taken sufficiently large that $\Pr(|s_B| > \kappa \mid \bar{s}_A)$ is small enough, it follows that

$$\text{for any } \xi > 0, \text{ if } (\boldsymbol{\rho}, \boldsymbol{b}) \text{ is small enough then } U_A(\bar{s}_A; \sigma^{\boldsymbol{\rho},\boldsymbol{b}}, \boldsymbol{\rho}, \boldsymbol{b}) < \xi. \tag{26}$$

However, by point (i) above, and because tail-signals have vanishing prior probability, we can choose a bounded set $\hat{S}_A \subset \mathbb{R}$ such that $\Pr(s_A \notin \hat{S}_A) \leq \delta/2$, and hence for any $(\boldsymbol{\rho}, \boldsymbol{b})$,

$$\Pr(A \text{ wins in } \sigma^{\boldsymbol{\rho},\boldsymbol{b}} \mid s_A \in \hat{S}_A) > \delta/2.$$

Consider candidate A with signal $\overline{s}_A$ deviating to a mixed action as follows: he draws $\hat{s}_A \in \hat{S}_A$ with the distribution of the prior F_{s_A} truncated on $\hat{S}_A$, and then follows the equilibrium prescription $y_A^{\boldsymbol{\rho},\boldsymbol{b}}(\hat{s}_A)$. Choose $k > 0$ large enough that $\Pr(|s_B| > k \mid s_A) \leq \delta/4$ uniformly over $s_A \in \hat{S}_A \cup \{\overline{s}_A\}$; this is possible because $s_B \mid s_A$ is normally-distributed with variance independent of s_A and mean linear in s_A, and $\hat{S}_A \cup \{\overline{s}_A\}$ is bounded. On $\{|s_B| \leq k\}$, the likelihood ratios of the conditional densities $f_{s_B|s_A}(\cdot \mid s_A)$ are uniformly bounded for $s_A \in \hat{S}_A \cup \{\overline{s}_A\}$. For each $s_A \in \hat{S}_A$, the event that A wins when he plays $y_A^{\rho,b}(s_A)$ depends only on s_B, and on $\{|s_B| \leq k\}$ its probability under $\overline{s}_A$ is bounded below by a positive multiple of its probability under s_A. Averaging over $s_A \in \hat{S}_A$, it follows that the probability with which A wins at signal $\overline{s}_A$ under the aforementioned deviation is bounded away from zero, uniformly in $(\boldsymbol{\rho}, \boldsymbol{b})$. Moreover, since $\hat{S}_A$ is bounded, Lemma 3 implies that the deviation's platforms are bounded uniformly in $(\boldsymbol{\rho}, \boldsymbol{b})$, and so as $(\boldsymbol{\rho}, \boldsymbol{b}) \to (\mathbf{0}, \mathbf{0})$, the policy-utility contribution to the deviation's expected payoff vanishes. Hence, the deviation yields candidate A with signal $\overline{s}_A$ a strictly positive expected payoff bounded away from zero, which given (26) would be a profitable deviation for small enough $(\boldsymbol{\rho}, \boldsymbol{b})$. □

Proof of Proposition 9. Let $\sigma_{\text{UB}}^{\boldsymbol{\rho},\boldsymbol{b}}$ be the equilibrium identified in Proposition 8 where, without loss, we take A to be the candidate who wins with probability one. Let $\sigma^{\boldsymbol{\rho},\boldsymbol{b}} \equiv (y_A^{\boldsymbol{\rho},\boldsymbol{b}}, y_B^{\boldsymbol{\rho},\boldsymbol{b}}, w_A^{\boldsymbol{\rho},\boldsymbol{b}}) \in \mathcal{E}^*(\boldsymbol{\rho}, \boldsymbol{b})$ be a sequence of welfare-maximizing equilibria as $(\boldsymbol{\rho}, \boldsymbol{b}) \to (\mathbf{0}, \mathbf{0})$. Applying Lemma 4 to this sequence, there are two cases:

(a) If Case 1 of Lemma 4 holds, then it is straightforward to verify that $v(\sigma^{\boldsymbol{\rho},\boldsymbol{b}}) \to v^*(\mathbf{0}, \mathbf{0})$. Intuitively, for $(\boldsymbol{\rho}, \boldsymbol{b}) \approx (\mathbf{0}, \mathbf{0})$, if i is winning with ex-ante probability approximately zero, then the voter's welfare cannot be much higher than if $-i$ wins with ex-ante probability one using the unbiased strategy, and Proposition 8 ensures that in a welfare-maximizing equilibrium it is not much lower either.

(b) If Case 2 of Lemma 4 holds, pick any subsequence of $\sigma^{\boldsymbol{\rho},\boldsymbol{b}}$ that converges pointwise almost everywhere and denote the limit by $\sigma^{\mathbf{0},\mathbf{0}}$.[39] Since payoffs are continuous, standard ar-

[39] More precisely, letting $\sigma^{\mathbf{0},\mathbf{0}} \equiv (y_A, y_B, w_A)$, we require that (i) $y_i^{\boldsymbol{\rho},\boldsymbol{b}}(s_i) \to y_i(s_i)$ for each i and almost all s_i and (ii) $w_A^{\boldsymbol{\rho},\boldsymbol{b}}(x_A, x_B) \to w_A(x_A, x_B)$ for each $(x_A, x_B) \in \mathbb{R}^2$. Case 2 of Lemma 4 assures that at least one subsequence converges in this sense. How each y_i is defined on zero-measure sets of signals is irrelevant. Note also that because the ex-ante probability of $\{s_i : s_i \notin [-k, k]\}$ can be made arbitrarily small by choosing $k > 0$ arbitrarily large, it follows that $v(\sigma^{\boldsymbol{\rho},\boldsymbol{b}}) \to v(\sigma^{\mathbf{0},\mathbf{0}})$.

guments imply that $\sigma^{\mathbf{0},\mathbf{0}}$ is an equilibrium of the limit pure-office-motivation game (intuitively, if the voter or a candidate with any signal has a profitable deviation, there would also have been a profitable deviation from $\sigma^{\boldsymbol{\rho},\boldsymbol{b}}$ for small enough $(\boldsymbol{\rho},\boldsymbol{b}) > (\mathbf{0},\mathbf{0})$). This implies that

$$\lim_{(\boldsymbol{\rho},\boldsymbol{b})\to(\mathbf{0},\mathbf{0})} v(\sigma^{\boldsymbol{\rho},\boldsymbol{b}}) = v(\sigma^{\mathbf{0},\mathbf{0}}) \leq v^*(\mathbf{0},\mathbf{0}).$$

Finally, the inequality above holds with equality because for all $(\boldsymbol{\rho},\boldsymbol{b})$, we have $v(\sigma^{\boldsymbol{\rho},\boldsymbol{b}}) \geq v(\sigma^{\boldsymbol{\rho},\boldsymbol{b}}_{\mathrm{UB}})$ as $\sigma^{\boldsymbol{\rho},\boldsymbol{b}}$ is welfare maximizing, and $v(\sigma^{\boldsymbol{\rho},\boldsymbol{b}}_{\mathrm{UB}}) \to v^*(\mathbf{0},\mathbf{0})$. □

E. A Beta-Bernoulli Specificaton

Here we repeat the analysis of Subsection 3.2 for the case in which the state follows a Beta distribution and each candidate gets a binary signal drawn from a Bernoulli distribution; the feasible set of policies is $[0,1]$ (or any superset thereof). This statistical structure is a member of the exponential family with conjugate priors. Aside from illustrating how the incentives to overreact exist even when the state distribution may not be unimodal and may be skewed, signals are discrete, etc., it also provides a closer comparison with the setting of Heidhues and Lagerlof (2003) and Loertscher (2012) than does our leading normal-normal specification.

Assume the prior distribution of θ is Beta(α,β), which is the Beta distribution with parameters $\alpha,\beta > 0$ whose density is given by $f(\theta) = \frac{\theta^{\alpha-1}(1-\theta)^{\beta-1}}{B(\alpha,\beta)}$, where $B(\cdot,\cdot)$ is the Beta function.[40] Thus θ has support $[0,1]$ and $\mathbb{E}[\theta] = \frac{\alpha}{\alpha+\beta}$. For reasons explained at the end of the section, we assume $\alpha \neq \beta$. (This rules out a uniform prior, which corresponds to $\alpha=\beta=1$.) Each candidate $i \in \{A,B\}$ observes a private signal $s_i \in \{0,1\}$; conditional on θ, signals are drawn independently from the same Bernoulli distribution with $\Pr(s_i = 1 \mid \theta) = \theta$. The policy space is any subset of $\mathbb{R}$ containing $[0,1]$.

It is well-known that the posterior distribution of the state given signal 1 is now Beta$(\alpha+1,\beta)$ (i.e., has density $f(\theta \mid s_i = 1) = \frac{\theta^{\alpha}(1-\theta)^{\beta-1}}{B(\alpha+1,\beta)}$); similarly the posterior given signal 0 is Beta$(\alpha,\beta+1)$. It is also straightforward to check that the posterior distribution of the state given two signals is as follows: if both $s_i = s_{-i} = 1$, it is Beta$(\alpha+2,\beta)$; if $s_i = 0$ and $s_{-i} = 1$, it is Beta$(\alpha+1,\beta+1)$; and if $s_i = s_{-i} = 0$, it is Beta$(\alpha,\beta+2)$.

It follows that

$$\mathbb{E}[\theta \mid s_i] = \frac{\alpha+s_i}{\alpha+\beta+1} \text{ and } \mathbb{E}[\theta \mid s_i, s_{-i}] = \frac{\alpha+s_i+s_{-i}}{\alpha+\beta+2}.$$

[40] If α and β are positive integers then $B(\alpha,\beta) = \frac{(\alpha-1)!(\beta-1)!}{(\alpha+\beta-1)!}$.

The above formulae imply that for any realization (s_A, s_B),

$$\text{sign}\left(\mathbb{E}[\theta \mid s_A, s_B] - \mathbb{E}[\theta]\right) = \text{sign}\left(\frac{\mathbb{E}[\theta \mid s_A] + \mathbb{E}[\theta \mid s_B]}{2} - \mathbb{E}[\theta]\right),$$

$$|\mathbb{E}[\theta \mid s_A, s_B] - \mathbb{E}[\theta]| > \left|\frac{\mathbb{E}[\theta \mid s_A] + \mathbb{E}[\theta \mid s_B]}{2} - \mathbb{E}[\theta]\right|. \tag{27}$$

Hence, both the posterior mean given two signals and the average of the individual posterior means shift in the same direction from the prior mean, but the former does so by more.

Consequently, if candidates were to play unbiased strategies and the voter best responds, then whenever $s_A \neq s_B$ there is one candidate who wins with probability one: the candidate i with $s_i = 1$ (resp., $s_i = 0$) when $\beta > \alpha$ (resp., $\beta < \alpha$). Of course, when $s_A = s_B$, both candidates would choose the same platform and win with equal probability. It is worth highlighting that when $s_A \neq s_B$, it is the candidate with the ex-ante *less* likely signal who wins, because ex-ante $\Pr(s_i = 1) = \mathbb{E}[\theta] = \alpha/(\alpha + \beta)$. This implies that unbiased strategies cannot form an equilibrium, but not because candidates would deviate when drawing the ex-ante less likely signal; rather, they would deviate when drawing the ex-ante *more* likely signal to the platform corresponding to the ex-ante less likely signal.[41] Notice that this profitable deviation given signal s_i is to an (on-path) platform x_i such that $|x_i - \mathbb{E}[\theta]| > |\mathbb{E}[\theta \mid s_i] - \mathbb{E}[\theta]|$; hence, it is a profitable deviation through overreaction rather than pandering.

Finally, we observe that there is a symmetric fully revealing equilibrium with overreaction in which both candidates play

$$y(1) = \frac{\alpha + 2}{\alpha + \beta + 2} \quad \text{and} \quad y(0) = \frac{\alpha}{\alpha + \beta + 2}.$$

This strategy displays overreaction because

$$y(0) < \mathbb{E}[\theta \mid s_i = 0] < \mathbb{E}[\theta] < \mathbb{E}[\theta \mid s_i = 1] < y(1).$$

It is readily verified that when both candidates use this strategy, $\mathbb{E}[\theta \mid s_A, s_B] = \frac{y(s_A)+y(s_B)}{2}$ for all (s_A, s_B), and hence each candidate would win with probability $1/2$ for all on-path platform pairs; a variety of off-path beliefs can be used to support the equilibrium.

Note that this overreaction equilibrium would exist even when $\alpha = \beta$. However, were $\alpha = \beta$, unbiased strategies would also constitute an equilibrium: for, given unbiased strategies,

[41] See Che, Dessein and Kartik (2013) for an analog where options that are "unconditionally better-looking" need not be "conditionally better-looking" .

both sides of (27) would be equal to each other (in fact, equal to zero) when $s_A \neq s_B$, and hence the voter could elect both candidates with equal probability no matter their platforms.