Information Revelation in Constant-Sum Games: Elections and Beyond

Navin Kartik Francesco Squintani Katrin Tinn

October 2026

Motivation

Voters are poorly informed about policy; political candidates less so

$\rightarrow$ elections, in particular platforms can convey candidates' information

Two views of competition between office-seekers:

1. Promotes information revelation and efficiency

   "returns to an informed political entrepreneur from providing the information" (Wittman 1989)

2. Induces pandering to voters' prior beliefs

   "incentive to follow popular beliefs ... instead of their own information" (Heidhues & Lagerlöf 2003)

Our question: How well do elections aggregate *candidates'* private information?

(cf. large literature on aggregating voters' information)

This paper

Downsian elections

two office-motivated candidates with private information about policy consequences

Main result

- In any equilibrium, voter welfare is as if one candidate were always elected
- So elections cannot efficiently aggregate more than one candidate's information
  (and some equilibrium achieves that)

General theoretical backbone

- Two-player constant-sum Bayesian games with type-independent payoffs
  $\rightarrow$ under "completeness", every "identifiable" eqm is an ex-post eqm

Pandering?

- Despite office motivation, candidates may anti-pander (overreact to info)
- Some pandering may actually benefit voters

Elections with Informed Candidates

Model

Voter: utility $u(x, \theta)$ over policy $x \in X$ and unknown state $\theta \in \Theta$
(median, representative, or a single decision maker)

Two candidates, A and B; each i has a private signal $s_i \in S_i$ about θ

1. Candidates commit to platforms x_A, x_B
2. Voter observes platforms and elects a candidate
3. Winner implements his platform

Candidates are purely office motivated: payoff 1 if elected; 0 otherwise

Welfare: voter's ex-ante expected utility

Binary Example (1)

$\theta, x, s_i \in \{0, 1\}$; voter wants policy to match the state

Heidhues & Lagerlöf (2003)

Prior $\Pr(\theta = 1) = p > \frac{1}{2}$ $\quad \rightarrow$ policy 1 is "popular"

Conditionally iid signals with accuracy $q \in (p, 1)$

Welfare benchmarks

$$\underbrace{p}_{\text{no information}} < \underbrace{q}_{\text{one signal}} < \underbrace{q + q(1-q)(2p-1)}_{\text{both signals}}$$

Both candidates follow their signals. Platforms disagree: who wins?

- Signals cancel $\implies$ policy 1 wins
- Efficient use of both signals

Not an equilibrium: with $s_i = 0$, deviate to $x_i = 1$ $\quad \rightarrow$ pandering

Binary Example (2)

Equilibria:

	candidates	who wins	welfare
both pander	$x_i = 1$ always	either	p
one candidate	$x_A = s_A$; B uninformative	A always	q
symmetric mixed	$x_i = 1$ if $s_i = 1$ and w.p. φ if $s_i = 0$	each w.p. $\frac{1}{2}$, at every platform pair	$q - \varphi(q - p)$



None of these beats efficient use of one signal (Heidhues & Lagerlöf)

A general limit? Other policy spaces, info structures, equilibria?

The Limit to Information Aggregation

Suppose distr of (s_A, s_B) satisfies completeness (defined soon; holds in the binary example)

Proposition

In any equilibrium, there is a candidate i such that voter welfare is the same as if i were always elected.

Theorem

The voter's best equilibrium has welfare $\max\{v_A^*, v_B^*\}$.

v_i^*: welfare from efficient use of i's signal alone

- Cannot efficiently aggregate more than one candidate's information
- Voter commitment does not help

Why? Candidates play a constant-sum game in which signals do not enter payoffs

Constant-Sum Bayesian Games with Type-Independent Payoffs

Setting

Two players, A and B; private types $s_i \in S_i$, with $(s_A, s_B) \sim F$

Simultaneous actions $x_i \in X_i$

Payoffs $u_i(x_i, x_{-i})$, with $u_A + u_B = 0$ $\quad \rightarrow$ constant sum and type-independent

Elections: types are signals; u_i is i's win probability, for a given voter strategy

All Bayes-Nash equilibria have the same payoffs (U_A^*, U_B^*)

Definition

An equilibrium is ex post if $u_i(x_i, x_{-i}) = U_i^*$ for all on-path (x_A, x_B).

$\rightarrow$ no on-path action does better, even after seeing the opponent's action

For now: finite types and actions

Two Conditions

Completeness

For each i and every $g : S_{-i} \to \mathbb{R}$: $\quad \mathbb{E}[g(s_{-i}) \mid s_i] = 0 \ \ \forall s_i \implies g = 0.$

$\iff$ the joint distribution is a square matrix with full rank

- Generic given $|S_A| = |S_B| < \infty$, though fails under independence
  binary example: holds iff signals are informative ($q \neq \frac{1}{2}$)
- Also holds for continuous exponential families such as normal (under regularity conditions)

Identifiability

Strategy σ_i is identifiable if for every $g : X_i \to \mathbb{R}$:

$\mathbb{E}[g(x_i) \mid s_i] = 0 \ \ \forall s_i \implies g = 0$ on path.

- Any pure strategy is identifiable
- Not identifiable: mixing that does not vary with type

Main Theoretical Result

Identifiable equilibrium: some player's strategy is identifiable, e.g., pure

Theorem

Under completeness, any identifiable equilibrium is an ex-post equilibrium.

All on-path action profiles give the same payoffs,
though types differ in actions and beliefs

$\Longrightarrow$ third parties can make only limited *use* of the players' information

A Lemma

Recall: U_i^* is i's payoff, the same in every equilibrium

Lemma

In any equilibrium, for every type s_i and every on-path action x_i:

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

Every type gets U_i^* from *any* action played by *any* type
(not assuming completeness or identifiability)

Why?

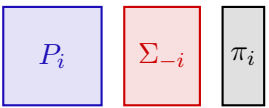
1. i's ex-ante action distribution ξ_i secures U_i^*; else opponent gains by ignoring his type
2. So each type's equilibrium payoff is exactly U_i^*
3. Against any type's beliefs, ξ_i yields $\geq U_i^*$
   and is an average of on-path actions, each yielding $\leq U_i^*$
   $\implies$ each on-path action yields U_i^*

(cf. Viossat (2006) on correlated equilibria)

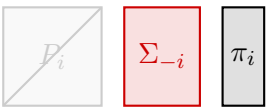
Proof of Theorem

Suppose $-i$'s strategy σ_{-i} is identifiable; fix an on-path x_i

$P_i := [F(s_{-i} \mid s_i)]$ $\quad$ $\Sigma_{-i} := [\sigma_{-i}(x_{-i} \mid s_{-i})]$ $\quad$ $\pi_i := [u_i(x_i, x_{-i}) - U_i^*]$



Lemma: every type s_i gets U_i^* from x_i



Completeness, applied to $\Sigma_{-i}\pi_i$
$\implies$ payoff is U_i^* against every opponent *type*



Identifiability, applied to π_i
$\implies$ payoff is U_i^* against every on-path *action*
$\rightarrow$ ex post

Uses completeness for i and identifiability for $-i$ only.

Neither Condition Can Be Dropped

Matching pennies: $X_A = X_B = \{0,1\}$, $u_A(x_A, x_B) = \mathbb{1}\{x_A = x_B\}$

Example

$S_A = S_B = \{0,1\}$, F uniform

Equilibrium: $s_i \mapsto s_i$ (identifiable)

Not ex post

Types independent: completeness fails

$$P_i = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$$

Example

Complete information (completeness holds)

Equilibrium: both mix uniformly

Not ex post

Strategies not identifiable

$$\Sigma_i = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} \end{pmatrix}$$

Nb: the Lemma holds in both

Infinite Types

Paper: general type and action spaces, as normal signals require
maintained: $F \ll F_A \otimes F_B$ (rules out perfectly correlated continuous signals)

Completeness: same condition, for bounded test functions g → bounded completeness

What is "full rank" with infinitely many types?

- Strong linear independence of the conditional distributions

Details

Same two-step proof structure: identical logic, just with technical details

Back to Elections

The Welfare Bound

To show: in any eqm, welfare is as if one candidate were always elected, so $\leq \max\{v_A^*, v_B^*\}$

Fix any equilibrium, possibly mixed; payoff = win probability

Completeness $\overset{\text{Lemma}}{\Longrightarrow}$ for on-path x_i and any opponent signal s_{-i},

▸ Completeness needed

$$\Pr(i \text{ wins} \mid x_i, s_{-i}) = U_i^*$$

$\Longrightarrow$ whether i wins is independent of i's platform and of the state

$\Longrightarrow$ welfare $= U_A^*\, v_A + U_B^*\, v_B \;\leq\; \max\{v_A, v_B\}$ (v_i: welfare if i always elected)

Voter optimality $\Longrightarrow$ welfare $= \max\{v_A, v_B\}$ ("as if"); $v_i \leq v_i^* \Longrightarrow$ bound

Identifiable equilibrium $\overset{\text{Theorem}}{\Longrightarrow}$ ex post: one candidate always wins, or voter always indifferent

(e.g., pure candidate strategies)

▸ Without ex-post

Pandering?

Normal-quadratic specification

$X = \Theta = S_i = \mathbb{R}; \quad u(x, \theta) = -(x - \theta)^2; \quad \theta \sim \mathcal{N}(0, 1); \quad s_i \mid \theta \overset{\text{iid}}{\sim} \mathcal{N}(\theta, 1/\beta)$

β: signal precision *relative* to the prior; completeness holds

Voter's preferred policy is the posterior mean. Note that

$$\mathbb{E}[\theta \mid s_i] = \frac{\beta}{1+\beta}\, s_i, \qquad \mathbb{E}[\theta \mid s_A, s_B] = \frac{2\beta}{1+2\beta}\left(\frac{s_A + s_B}{2}\right)$$

A strategy y_i is

- unbiased if $y_i(s_i) = \mathbb{E}[\theta \mid s_i]$
- pandering (underreaction) if $y_i(s_i)$ is between 0 and $\mathbb{E}[\theta \mid s_i]$
- anti-pandering (overreaction) if $y_i(s_i)$ lies beyond $\mathbb{E}[\theta \mid s_i]$

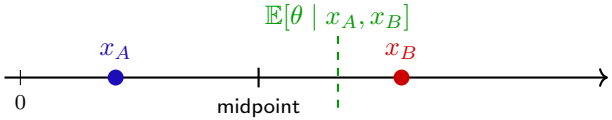
Unbiased Strategies?

Suppose both candidates are unbiased $\quad \to$ both signals revealed

Voter's posterior mean given platforms:

$$\mathbb{E}[\theta \mid x_A, x_B] = \underbrace{\frac{2(1+\beta)}{1+2\beta}}_{>1} \cdot \frac{x_A + x_B}{2}$$

$\to$ more extreme than the average platform $\quad (\because$ two signals $\implies$ less weight on prior)



$\implies$ voter elects the more extreme candidate

$\implies$ not an eqm, $\because$ profitable to overreact (but not to underreact/pander)

Binary example: could only distort toward the prior

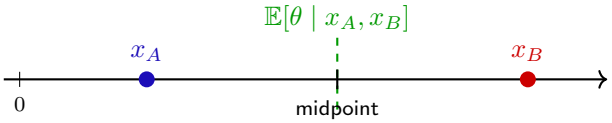
Anti-Pandering Equilibrium

Proposition (Normal-quadratic)

There is a symmetric, fully-revealing anti-pandering equilibrium: both candidates play

$$y_i(s_i) = \mathbb{E}[\theta \mid s_i, s_{-i} = s_i] = \frac{2\beta}{1+2\beta}\, s_i,$$

and each is elected with probability $1/2$ regardless of platforms.



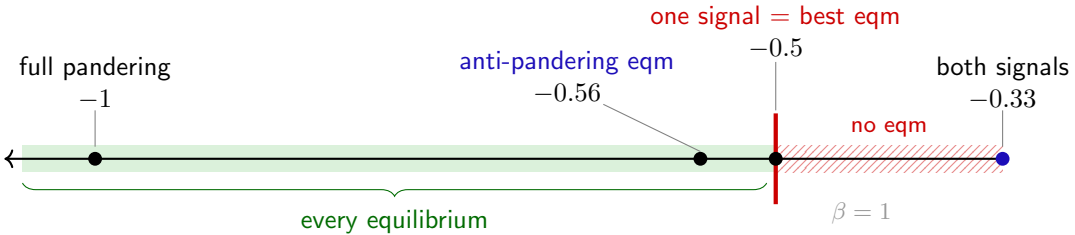
posterior = midpoint, for any platforms
$\implies$ voter indifferent

- Office motivation need not induce pandering; platforms diverge with prob 1
- Unique candidate strategies in any symmetric (where both win with pos prob), fully-revealing, continuous eqm

▸ Details

Equilibrium Welfare

$$\underbrace{-1}_{\text{full pandering}} < \underbrace{-\frac{1+4\beta}{(1+2\beta)^2}}_{\text{anti-pandering eqm}} < \underbrace{-\frac{1}{1+\beta}}_{\text{one signal}} < \underbrace{-\frac{1}{1+2\beta}}_{\text{both signals}}$$



- Anti-pandering reveals both signals, yet is worse than efficient use of one
- Best eqm: one candidate is unbiased and always wins
- Inefficiency of best eqm, as a share of the value of both signals: $\downarrow$ in β, from $\frac{1}{2}$ to 0 ▸ Plot

The Benefits of Pandering

Under unbiased strategies the more extreme candidate wins

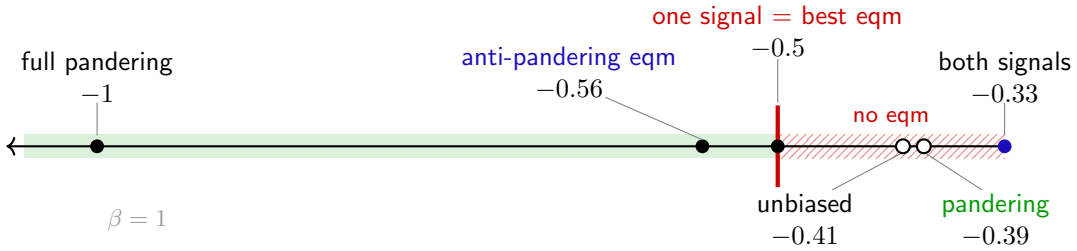
▸ Strategies

→ a "winner's curse": voter would, on average, prefer winner's policy to be more moderate

Proposition (Normal-quadratic)

If both candidates pander via $y_i(s_i) = \mathbb{E}[\theta \mid s_i, |s_{-i}| \leq |s_i|]$ and the voter best responds, voter welfare exceeds that in any equilibrium.

→ each proposes the best policy conditional on winning



Robustness

Welfare bound uses: two candidates, pure office motivation, voter sees only platforms

Other game forms

- "As if" result holds, e.g., with voter commitment or sequential platforms
- But the level can change: sequential platforms can even aggregate fully
- Informed voter: payoffs no longer type-independent (Laslier & Van der Straeten 2004)

Beyond office motivation

- Approximate bound when almost office motivated ▸ Mixed motives
- Substantial policy motivation can aggregate more (Ambrus, Baranovskyi, Kolb 2021)

Beyond normal-quadratic

- Anti-pandering equilibria in conjugate exponential families ▸ Details

Related Literature

Constant-sum games

- von Neumann (1928); Viossat (2006); Kattwinkel, Niemeyer, Preusser, Winter (2022)

Elections with privately informed candidates

- Heidhues & Lagerlöf (2003); Laslier & Van der Straeten (2004); Loertscher (2012); Millner, Ollivier, Simon (2020)
- Policy-motivated candidates: Schultz (1996); Martinelli (2001); Ambrus, Baranovskyi, Kolb (2021)

Pandering and anti-pandering with reputation concerns

- Canes-Wrone, Herron, Shotts (2001); Maskin & Tirole (2004); Prendergast & Stole (1996); Levy (2004)

Non-strategic extremity from Bayesian updating

- Glaeser & Sunstein (2009); Roux & Sobel (2015)

Conclusion

Conclusion

Two-player constant-sum Bayesian games with type-independent payoffs:
under completeness, every identifiable equilibrium is ex post

Takeaways

- Purely adversarial players' information is hard for third parties to *use*, even when revealed
- Downsian elections cannot efficiently aggregate both candidates' information;
  sharp bound: efficient use of one candidate's signal
- Office motivation need not induce pandering; yet some pandering would benefit voters
- Beyond elections: competition in dual spheres; arbitration; regulation

Thank you!

Without Completeness

$X = S_A = S_B = \{1, 2, 3, 4\}$; $\Theta = S_A \times S_B$, uniform prior; signals reveal own coordinate: $(s_A, s_B) = \theta$

$\Longrightarrow$ signals independent: completeness fails

Voter's utility is 1 if the "correct" policy is chosen, else 0

Correct policy

$s_A \,/\, s_B$	1	2	3	4
1	1	2	3	1
2	1	2	3	2
3	3	2	3	4
4	1	4	4	4

Elected candidate

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

Equilibrium: each candidate plays $s_i \mapsto s_i$; voter elects as in the right table

- Voter gets the correct policy in every state: first best, above $\max\{v_A^*, v_B^*\}$
- Each candidate wins with prob $1/2$ whatever his platform
  (posterior on opponent's signal is uniform)

(Back)

Completeness and Strong Linear Independence

Characterizes *completeness* (all integrable g), slightly stronger than bounded completeness

$\{F(\cdot \mid s_i)\}_{s_i \in S_i}$ is strongly linearly independent (SLI) if for every finite signed measure $\mu \ll F_i$,
$\int_{S_i} F(\cdot \mid s_i)\, \mu(\mathrm{d}s_i) = 0 \implies \mu = 0$

Proposition

Completeness for $i \iff$ SLI of $\{F(\cdot \mid s_{-i})\}_{s_{-i} \in S_{-i}}$.

- Finite types: SLI is linear independence
- Countable types: a linearly independent family can have a vanishing infinite mixture (Example 4)
- Convex independence is not enough, even with finite types (Example 5):
  $F = \frac{1}{40} \begin{pmatrix} 4 & 2 & 1 & 3 \\ 2 & 4 & 3 & 1 \\ 1 & 3 & 4 & 2 \\ 3 & 1 & 2 & 4 \end{pmatrix}$ has rank 3; an identifiable equilibrium is not ex post

(Back)

Mixing: Welfare Bound without Ex-Postness

The bound does not need ex-postness, only $\Pr(i \text{ wins} \mid x_i, s_{-i}) = U_i^*$

Example: $\Theta = X = \{1, 2\}$, $u(x, \theta) = \mathbb{1}\{x = \theta\}$, uniform prior; any complete signal structure

- Equilibrium: both candidates mix uniformly whatever their signals;
  voter elects A iff $x_A = x_B$
- Not identifiable, not ex post: win probability is 1 or 0 depending on the platform pair
- Yet given the opponent's signal, each wins with prob $1/2$ whichever policy he proposes
- Welfare is as if one candidate were always elected

(Back)

Anti-Pandering: Uniqueness

Normal-quadratic. Consider any equilibrium in which both candidates win with positive probability and use continuous, fully-revealing pure strategies

Then there are $c \in \mathbb{R}$ and i such that

$$y_i(s_i) = \frac{2\beta}{1+2\beta}\, s_i + c \qquad \text{and} \qquad y_{-i}(s_{-i}) = \frac{2\beta}{1+2\beta}\, s_{-i} - c$$

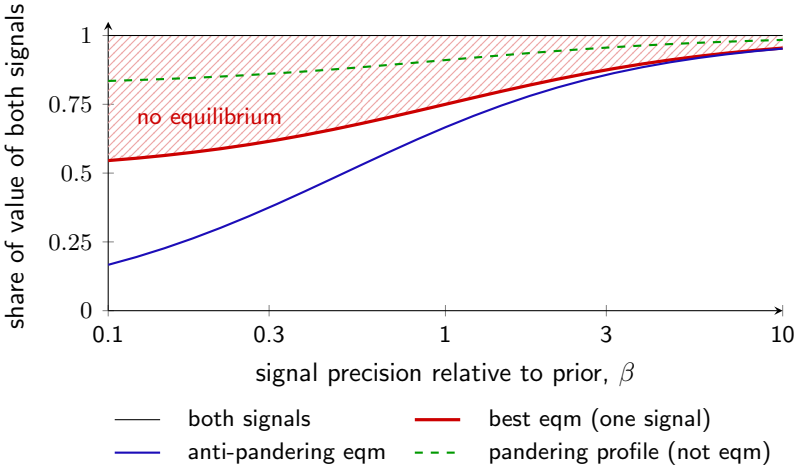
(symmetry $\implies c = 0$; welfare is c^2 lower than in the symmetric equilibrium)

Logic:

- Ex-post theorem $\implies$ voter indifferent after almost all on-path platform pairs
- $\implies \mathbb{E}[\theta \mid x_A, x_B] = (x_A + x_B)/2$ on path
- Continuity and full revelation then pin down the strategies

(Back)

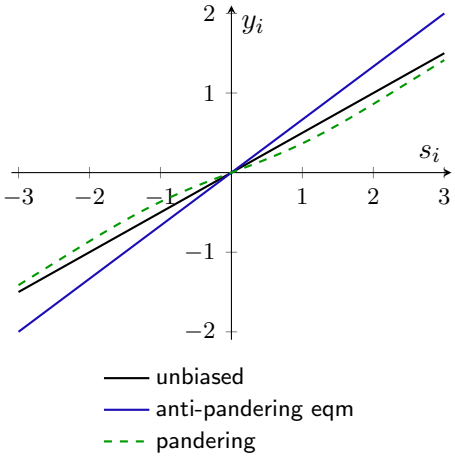
How Much Is Lost?



0 = no information, 1 = both signals used efficiently

(Back)

Unbiased, Anti-Pandering, and Pandering Strategies



$\beta = 1$

Proposition

The pandering profile maximizes voter welfare among all profiles in which the voter's best response elects i whenever $|s_i| > |s_{-i}|$.

Logic: optimal platform conditional on winning is $\mathbb{E}[\theta \mid s_i,\ i \text{ wins}]$

(Back)

Mixed Motives

Normal-quadratic, with candidate i's payoff

$$u_i = -\rho_i \left(x_W - \theta - b_i\right)^2 + (1 - \rho_i)\, \mathbb{1}\{W = i\}$$

W: winner; b_i: ideological bias; $\rho_i \in [0, 1]$: weight on policy

Proposition

There is a fully-revealing equilibrium in which one candidate i plays $y_i(s_i) = \frac{\beta}{1+\beta} s_i + b_i$, the other plays $y_{-i}(s_{-i}) = s_{-i} - \frac{1+\beta}{\beta} b_i$, and the voter always elects i.

Proposition

As both $(\rho_i, b_i) \to 0$, among equilibria in nondecreasing strategies with $|y_i(s_i)| \leq C(1 + |s_i|)$, the supremum of voter welfare converges to the baseline bound $-1/(1+\beta)$.

Finite spaces, continuous payoffs: upper hemicontinuity of equilibria suffices

(Back)

Anti-Pandering Beyond Normal-Normal

Asymmetric precisions, $s_i \mid \theta \sim \mathcal{N}(\theta, 1/\beta_i)$: fully-revealing equilibrium with

$$y_i(s_i) = \frac{2\beta_i}{1 + \beta_A + \beta_B}\, s_i$$

both candidates anti-pander when $|\beta_A - \beta_B| < 1$

Conjugate exponential families (quadratic loss)

posterior mean is linear in prior mean and signals

$\implies$ fully-revealing equilibrium with overreaction; voter randomizes uniformly on path

Beta-Bernoulli: $\theta \sim \mathrm{Beta}(a, b)$ with $a \neq b$; $s_i \in \{0, 1\}$ with $\Pr(s_i = 1 \mid \theta) = \theta$

- Unbiased strategies are not an equilibrium: deviation is to overreact
- Symmetric fully-revealing equilibrium: $y(1) = \dfrac{a+2}{a+b+2}$, $y(0) = \dfrac{a}{a+b+2}$

(Back)