Available online at www.sciencedirect.com
ScienceDirect

Games and Economic Behavior 58 (2007) 1–29

GAMES and Economic Behavior

www.elsevier.com/locate/geb

Electoral competition with privately-informed candidates

Dan Bernhardt [a], John Duggan [b,*], Francesco Squintani [c]

[a] Department of Economics, University of Illinois
[b] Department of Political Science and Department of Economics, Wallis Institute of Political Economy, Harkness Hall, University of Rochester, Rochester, NY 14627
[c] Department of Economics, University College London

Received 25 August 2004

Available online 18 April 2006

Abstract

This paper formulates and analyzes a general model of elections in which candidates receive private signals about voters' preferences prior to committing to political platforms. We fully characterize the unique pure-strategy equilibrium: After receiving her signal, each candidate locates at the median of the distribution of the median voter's location, conditional on the other candidate receiving the *same* signal. Sufficient conditions for the existence of pure strategy equilibrium are provided. Though the electoral game exhibits discontinuous payoffs for the candidates, we prove that mixed strategy equilibria exist generally, that equilibrium expected payoffs are continuous in the parameters of the model, and that mixed strategy equilibria are upper hemicontinuous. This allows us to study the robustness of the median voter theorem to private information: Pure strategy equilibria may fail to exist in models "close" to the Downsian model, but mixed strategy equilibria must, and they will be "close" to the Downsian equilibrium.

JEL classification: C72; D72; D78

1. Introduction

The most familiar and widely-used model of elections in political science and political economy is the classical Downsian model (Hotelling, 1929; Downs, 1957; Black, 1958). The central

* Corresponding author.
E-mail address: dugg@troi.cc.rochester.edu (J. Duggan).

0899-8256/$ – see front matter
doi:10.1016/j.geb.2006.03.004

result is the median voter theorem: When two office-motivated candidates are completely informed about the median voter's ideal point, the unique electoral equilibrium is that both candidates locate there. The political science literature on "probabilistic voting" relaxes the assumption of complete information and assumes only that candidates are symmetrically informed, sharing a common prior distribution about the location of the median voter. It is a "folk theorem" that, in this environment, the unique Nash equilibrium is for both candidates to locate at the *median of the prior distribution* of the median voter's ideal point. In reality, however, candidates' platforms often diverge from their estimates of the median voter's preferred policy, as is widely documented in the empirical literature.[1]

Symmetric information is clearly very strong. It is evident that political candidates do not precisely know voters' policy preferences when selecting platforms, and that asymmetries in candidates' information about voters may arise from many sources—differences in the candidates' personal experiences, or different backgrounds of their political advisors, or different results from private polling conducted before the election. There is considerable supporting evidence. Eisinger (2003) finds that since the Roosevelt administration, private polls have been an integral part of the White House modus operandi. Medvic (2001) finds that 46 percent of all spending on US Congressional campaigns in 1990 and 1992 was devoted to the hiring of political consultants, primarily political pollsters. In addition, the major parties provide polling services to their candidates. Of course, private polling information is jealously guarded by candidates and parties. Indeed, Nixon had polls routinely conducted, but did not disclose results even to the Republican National Committee; and F.D. Roosevelt described private polling as his "secret weapon" (Eisinger, 2003).

In this paper, we develop a general model of elections in which candidates receive private polling signals. Each candidate receives a signal drawn from an arbitrary finite set of possible signals about the location of the median voter's ideal policy; each candidate updates about the location of the median voter and about her opponent's platform before choosing a platform from the real line; and the candidate whose platform is closest to the median voter wins. The setting we consider is quite general. In particular, we allow for any possible correlation in polling signal structure and any number of signal realizations. We allow for any family of conditional distributions of the random median policy such that the conditional distributions are continuous with connected supports. While we give results for a baseline model in which candidates have identical polling technologies, our most general results allow candidates to have different polling technologies, as might be expected when an incumbent runs against a challenger. Within this framework, we derive the existence and continuity properties of electoral equilibria, and we determine the ways in which the classical median voter theorem is and is not robust to the introduction of small amounts of asymmetric information.

The introduction of private polling to the model generates subtle informational incentives for candidates, and logic of the median voter theorem does not extend to the general private-information environment in the expected way. In particular, a candidate does not target the median voter conditional on his own signal. We prove that in the symmetric model, there is

[1] See, for example, the National Election Survey data estimating presidential candidates' platforms from 1964 to 1972 (Page, 1978, Chapters 3 and 4) and for 1984 and 1988 races (Merrill and Grofman, 1999, pp. 55–56). Budge et al. (2001) compare estimates of the US and British median voters based on survey data (such as the NES and British Election Survey) with estimates of candidates' platforms derived from speech and writing context analyses. They find clear evidence of divergence from the median policy, and no evidence of extremization. Poole and Rosenthal (1997) obtain similar findings using roll call voting to estimate Congress-persons' platforms (pp. 62–63).

at most one pure strategy equilibrium: After receiving a signal, a candidate updates the prior distribution of the median voter, *conditioning on both candidates receiving that same signal,* and locates at the median of *that* posterior distribution.[2] In the probabilistic voting model, where candidates have symmetric information, conditioning on one candidate receiving a signal is the same as conditioning on both receiving it, so we obtain the traditional probabilistic version of the median voter theorem as a special case. With private information, however, our result shows that strategic competition leads candidates to take positions that are more extreme than their own estimates of the median voter's ideal policy: Asymmetric information obviously leads to policy divergence, and the strategic effect magnifies the policy divergence already inherent in private information.

We give sufficient conditions for existence of the pure strategy equilibrium, the key being that conditional on a candidate receiving a signal, the probability that the opponent receives a signal weakly to the "left" should exceed the probability that the opponent receives a signal strictly to the "right," and vice versa. This limits the incentive for a candidate to move away from the equilibrium platform after any signal, and together with other background conditions, it ensures the existence of the pure strategy equilibrium. This key condition is actually necessary for existence in some environments. It becomes quite restrictive, however, when the number of possible signals is large, and we conclude that the pure strategy equilibrium often fails to exist in elections with fine polling information. In fact, we show that if we add arbitrarily small amounts of asymmetric information to the Downsian model, then the pure strategy equilibrium may cease to exist, highlighting the issue of robustness of the median voter theorem with respect to even small amounts of private information.

These considerations lead us to analyze mixed strategy equilibria. Despite discontinuities inherent in candidate payoffs, we prove that mixed strategy equilibria exist. We show that the (unique) mixed strategy equilibrium payoffs of our model vary continuously in its parameters, and we use this result to prove upper hemicontinuity of equilibrium mixed strategies. Imposing only our minimal functional form restrictions, we obtain characterization results for mixed-strategy equilibria. We prove that the supports of mixed strategy equilibria lie in the interval defined by the smallest and largest conditional medians, and we deduce the corollary that the equilibrium of the traditional probabilistic voting model is unique within the class of all *mixed* strategy equilibria. Furthermore, we show that the only possible atoms of equilibrium mixed strategies are at conditional medians. As a consequence, if there is a positive probability that the candidates converge on the same policy platform in equilibrium, then that platform must belong to the finite set of conditional medians.

Finally, we return to the issue of robustness of the median voter theorem. Our continuity results apply to the traditional probabilistic voting model and immediately yield robustness of the probabilistic version of the median voter theorem: When candidate beliefs about the median voter's location are "close" to some common distribution, mixed strategy equilibria must be "close" to the median of that distribution. Furthermore, even though the Downsian model is marked by fundamental discontinuities, the robustness result extends. Thus, in the Downsian model, the median voter theorem is fragile in terms of pure strategies, but robust in terms of mixed strategies: Mixed strategy equilibria exist and must be close to the median when small

[2] This result is reminiscent of the findings of Milgrom (1981), who shows that in a common-value second-price auction, the equilibrium bid of a type θ corresponds to the expected value of the good conditional on both types being equal to θ. Here, since candidates maximize the probability of winning, the relevant statistic is the median.

amounts of asymmetric information are added to the model.[3] Lastly, we give examples showing the robustness result for the Downsian model relies critically on complete information: It does not extend to general models with discontinuous conditional distributions.

2. Related literature

Bernhardt et al. (2005) (henceforth BDS) explicitly calculate the (essentially unique) mixed strategy equilibrium in a tractable version of our model of privately informed candidates. This allows them to derive comparative statics and investigate voter welfare. Ledyard (1989) first raised the issue of privately-informed candidates and considered several examples exploring the effects of the order of candidate position-taking, public polls, and repeated elections. Other papers have independently considered aspects of elections with privately-informed candidates. Chan (2001) studies a three-signal model that differs structurally from ours in that a valence term, common to all voters but unobserved by the candidates, is attached to each candidate. He only shows existence of a pure strategy equilibrium when signals are almost uninformative, circumventing non-convexities in candidate payoffs. Ottaviani and Sorensen (2003) consider a model of financial analysts who receive private signals of a firm's earnings and simultaneously announce forecasts, with rewards depending on the accuracy of their predictions. The case of two analysts can be interpreted as a model of electoral competition with privately-informed candidates similar to ours, except that the authors assume signals are continuously distributed. They prove existence of mixed strategy equilibrium when the policy space is finite, permitting the application of standard purification arguments to obtain a pure strategy equilibrium. For infinite policy spaces, they give a characterization of equilibrium strategies based on a first order analysis, but they do not prove existence, which our results indicate is a significant issue.

Presumably, given the difficulty of observing the structure of information in an election, the ideal analysis would model signals as generated by an arbitrary probability measure, conditional on the true location of the median voter. We choose to model signals as discretely distributed, providing a natural model and giving us a handle on issues of equilibrium existence and characterization. This also captures the simple special cases, e.g., two or three possible signals, that are likely to be considered in applications of the model. Moreover, since we allow for *arbitrary* discrete distributions of signals, we gain insight into the continuous signal model when viewed as a limit of discrete models. For example, because pure strategy equilibria typically fail to exist when the number of signals is large, we see that pure strategy equilibria in the continuous signal model, if they exist and if they are robust to discretization of the signal space, must be viewed as the limit of mixed, rather than pure, strategy equilibria.

More distantly related, in Heidhues and Lagerlöf (2003) there are two policy alternatives, all voters prefer the implemented policy to match the unknown state of the world, and candidates receive private signals correlated with this unknown state. They find that candidates bias their platforms toward the prior, and that lower signal correlation leads to larger bias and lower welfare. Martinelli (2001) studies a situation in which parties are better informed than voters about the optimal policies for voters, but voters have private information of their own. If the voters' information is biased, equilibrium results in less than full convergence even if parties know with certainty the optimal policy for the median voter. Martinelli and Matsui (2002) extend this model

[3] Banks and Duggan (2005) derive a related finding in the probabilistic voting model with expected plurality-maximizing candidates. There, the unique pure strategy equilibrium may cease to exist if a small amount of randomness is introduced to voter behavior, but mixed strategy equilibria exist and vary continuously in this respect.

to the case of two policy-motivated parties. Surprisingly, in all separating equilibria, when the left-wing party attains power, the policies it implements are to the right of the policies implemented by the right-wing party when it attains power.

There is also a large literature on the convergence of policy platforms in elections, as seen in the Downsian and probabilistic voting models. Policy convergence survives if we introduce policy preferences for the candidates (Calvert, 1985; Duggan and Fey, 2005).[4] Further variations of the basic model can generate policy divergence. For example, adding probabilistic voting and policy-motivation simultaneously to the basic model does so, as in any pure strategy equilibrium a wedge appears between the candidates' platforms (Wittman, 1983; Calvert, 1985). Also, if one candidate has a "valence advantage," so that all voters would vote for that candidate even if the opponent offers a slightly better policy, then equilibria in pure strategies do not exist, and mixed strategy equilibria obviously lead candidates to adopt distinct platforms (Groseclose, 2001 and Aragones and Palfrey, 2002).[5]

3. The electoral framework

Two political candidates, A and B, simultaneously choose policy platforms, x and y, on the real line, $\Re$. Voter utilities are symmetric. We assume that a median, denoted μ, is uniquely defined (e.g., either the number of voters is odd, or there is a continuum of voters with ideal policies distributed according to a density with convex support). Hence, policy z is majority-preferred to policy w if and only if z is preferred by the median voter, and we can summarize all relevant information about voters' preferences by the median policy μ. Candidate A wins the election if $|x-\mu|<|y-\mu|$ and loses if the inequality is reversed; if $|x-\mu|=|y-\mu|$, then the election is decided by a fair coin toss, so that each candidate wins with probability one half.

Candidates do not observe μ, but receive private signals, s and t, about its realization. These signals may reflect the personal experiences of the candidates, or they may be generated by private polls conducted by the candidates' campaign organizations or political parties. Let S denote the finite set of possible signals for A, and let T denote the finite set of possible signals for B. Candidates have a common prior distribution on $\Re \times S \times T$, where the distribution of μ conditional on signals s and t is $F_{s,t}$, and the marginal probability of signal pair (s,t) is $P(s,t)$. The marginal probability of signal s is $P(s)$, and that of signal t is $P(t)$, where we assume that $P(s)>0$ and $P(t)>0$. Conditional probabilities $P(\cdot|s)$ and $P(\cdot|t)$ are defined using Bayes rule. The model is fully general with respect to the correlation between candidates' signals, allowing for conditionally-independent signals and perfectly-correlated signals as special cases.

We impose minimal regularity conditions on the conditional distributions of the median. We assume that each $F_{s,t}$ is continuous, and that for all $a,b,c \in \Re$ with $a<b<c$, if $0<F_{s,t}(a)$ and $F_{s,t}(c)<1$, then $F_{s,t}(a)<F_{s,t}(b)<F_{s,t}(c)$. Thus, $F_{s,t}$ admits a density, denoted $f_{s,t}$, with convex support. Let $m_{s,t}$ be the uniquely-defined median of $F_{s,t}$. Given a subset $T' \subseteq T$, we define $F_{s,T'}$ by

$$F_{s,T'}(z)=\sum_{t\in T'}\frac{P(t|s)}{P(T'|s)}F_{s,t}(z),$$

[4] Banks and Duggan (2005) show that convergence obtains quite generally under expected plurality maximization as well.

[5] A different modeling approach is the "citizen-candidate model" (Osborne and Slivinski, 1996; Besley and Coate, 1997), where candidates' platforms are determined exogenously by their policy preferences.

when $P(T'|s) > 0$, and we let $m_{s,T'}$ denote the unique median of $F_{s,T'}$. We define $F_{S',t}$ and $m_{S',t}$ analogously. We write F_s for $F_{s,T}$, the distribution of μ conditional only on s, and m_s for the associated median; F_t and m_t are defined analogously.

Thus, the probability that candidate A wins when A uses platform x and receives signal s and B uses platform y and receives signal t, is

$$F_{s,t}\left(\frac{x+y}{2}\right) \quad \text{if } x<y, \qquad 1-F_{s,t}\left(\frac{x+y}{2}\right) \quad \text{if } y<x, \quad \text{and} \quad \frac{1}{2} \quad \text{if } x=y.$$

The probability that B wins has an analogous form and, of course, is equal one minus the probability that A wins. Assuming office-motivated candidates, this defines a Bayesian game, in which pure strategies for the candidates are vectors $X=(x_s)$ and $Y=(y_t)$, and given pure strategies X and Y, candidate A's interim expected payoff conditional on signal s, denoted $\Pi_A(X,Y|s)$, is

$$\sum_{t\in T:\ x_s<y_t} P(t|s)F_{s,t}\left(\frac{x_s+y_t}{2}\right) + \sum_{t\in T:\ y_t<x_s} P(t|s)\left(1-F_{s,t}\left(\frac{x_s+y_t}{2}\right)\right) + \frac{1}{2}\sum_{t\in T:\ x_s=y_t} P(t|s).$$

Candidate B's interim expected payoff is defined analogously.

A *pure strategy Bayesian equilibrium* is a strategy pair (X,Y) such that $\Pi_A(X,Y|s) \geqslant \Pi_A(X',Y|s)$ for all signals $s\in S$ and all strategies X'; and $\Pi_B(X,Y|t) \geqslant \Pi_B(X,Y'|t)$ for all signals $t\in T$ and all strategies Y'. That is, candidates' campaign platforms are chosen optimally given all information available to them.

Candidate A's ex ante expected payoff, the expected payoff before receiving a signal, is then just the ex ante probability of winning:

$$\Pi_A(X,Y) = \sum_{s\in S} P(s)\Pi_A(X,Y|s).$$

Candidate B's ex ante expected payoff, $\Pi_B(X,Y)$, is defined analogously. Clearly, $\Pi_A(X,Y) + \Pi_B(X,Y) = 1$ for all strategies X and Y. Thus, (X,Y) is a pure strategy Bayesian equilibrium if and only if $\Pi_A(X,Y') \geqslant \Pi_A(X,Y) \geqslant \Pi_A(X',Y)$ for all X' and all Y'. That is, the pure strategy Bayesian equilibria are equilibria of a two-player, constant-sum game. The constant sum property implies that equilibria are "interchangeable": If (X,Y) and (X',Y') are equilibria, then so are (X,Y') and (X',Y). All of these concepts extend to mixed strategies, defined in Section 5, where candidates use their information in a non-deterministic way.

At times, we impose more structure. Conditions (C1)–(C4) define our *Canonical Model* of polling, in which candidates employ identical polling technologies and signals exhibit a natural ordering structure. We first impose symmetry conditions (C1) and (C2).

(C1) $S=T$.

Thus, the same set I, with elements i, j, etc., can be used to index these sets. We then write $P(i,j)$ for $P(s_i,t_j)$, $F_{i,j}$ for F_{s_i,t_j}, and so on.

(C2) For all signals $i,j\in I$, $P(i,j)=P(j,i)$ and $F_{i,j}=F_{j,i}$.

Condition (C2) implies that signal i of candidate A can be identified with signal i of candidate B in the sense that they are equally informative. If candidates have equal access to resources and polling technologies, then conditions (C1) and (C2) are natural assumptions. In that case, we may be interested in equilibria in which candidates use information similarly: A *symmetric* pure strategy Bayesian equilibrium is an equilibrium (X,Y) in which $x_i=y_i$ for all $i\in I$.

Under (C1)–(C2), candidates' ex ante payoffs are symmetric in the sense that $\Pi_A(X,Y) = \Pi_B(Y,X)$ for all X and Y. Symmetry implies that if (X,Y) is an equilibrium, then so is (Y,X). With interchangeability, it follows that if (X,Y) is an equilibrium, then (X,X) and (Y,Y) are symmetric equilibria.

The next condition says that when one candidate receives any signal, the other candidate also receives that signal with positive probability.

(C3) For all signals $i \in I$, $P(i,i) > 0$.

The last condition defining the Canonical Model imposes a natural linear ordering structure on signals. It captures the idea that "higher" signals are correlated with higher values of μ.

(C4) There exists a linear ordering $\leqslant$ of I, with asymmetric part $<$,[6] such that for all signals $i, j \in I$,

$$i < j \quad \text{if and only if} \quad \text{for all } K \subseteq I \text{ with } P(K|i)P(K|j) > 0, \text{ we have } m_{i,K} < m_{j,K}.$$

Note that condition (C4) implies that $i < j$ if and only if the conditional medians satisfy $m_{j,i} < m_{j,j}$, or equivalently (in the Canonical Model), $m_{i,i} < m_{j,i}$.

A special case of the Canonical Model is the traditional *Probabilistic Voting Model*, in which information between the candidates is complete, so that $P(i|i) = 1$ for all $i \in I$, and the conditional distribution $F_{i,i}$ is common knowledge after signal pair (i,i) is realized. Since the realizations of the candidates' signals are common knowledge in this model, the electoral game following any signal pair (i,i), which we refer to as the ith "component game," can be analyzed independently as the game with strategy set $\Re$ for each candidate and payoffs

$$\Pi_A^i(x,y) = \begin{cases} F_{i,i}(\frac{x+y}{2}) & \text{if } x < y, \\ 1 - F_{i,i}(\frac{x+y}{2}) & \text{if } y < x, \\ \frac{1}{2} & \text{if } x = y. \end{cases}$$

From the probabilistic version of the median voter theorem, the unique pure strategy Bayesian equilibrium has each candidate locating at m_i after signal i.

The *Downsian Model* of elections specializes the Probabilistic Voting Model, but for one feature that places it outside the class of models that we consider. As in the Probabilistic Voting Model, candidates have complete information about each other's signals, i.e., $P(i|i) = 1$ for all $i \in I$. The exceptional feature is that conditional on a candidate's signal, the location of μ is known with certainty. Formally, $F_{i,i}$ is the point mass on $m_{i,i} = m_i$, violating our maintained assumption that $F_{i,i}$ is continuous. A candidate then wins with probability one if her platform is closer to μ than the other candidate's, generating expected payoffs

$$\Pi_A(X,Y|i) = \begin{cases} 1 & \text{if } x_i < y_i \text{ and } \frac{x_i+y_i}{2} > m_i, \\ 1 & \text{if } y_i < x_i \text{ and } \frac{x_i+y_i}{2} < m_i, \\ \frac{1}{2} & \text{if } x_i = y_i \text{ or } \frac{x_i+y_i}{2} = m_i, \\ 0 & \text{otherwise.} \end{cases}$$

[6] We say $\leqslant$ is a *linear order* if it is complete, transitive, and anti-symmetric, i.e., $i \leqslant j$ and $j \leqslant i$ implies $i = j$. As the asymmetric part, $i < j$ holds if and only if $i \leqslant j$ but not conversely.

By the median voter theorem, the unique pure strategy equilibrium has the candidates locating at m_i after receiving signal i.

4. Pure strategy equilibrium

This section provides a full characterization of the pure strategy equilibria of the Canonical Model: If a pure strategy Bayesian equilibrium exists, then it is unique; and after receiving a signal, a candidate locates at the median of the distribution of μ conditional on *both* candidates receiving that signal. Given a natural restriction on conditional medians, a corollary is that candidates take policy positions that are extreme relative to their expectations of μ given their own information. Specifically, candidates who receive high signals tend to overshoot μ, while those who receive low signals tend to undershoot. To derive the characterization, we first consider symmetric equilibria in Lemma 1. The proof is provided in Appendix A, as are all proofs of formal results that are omitted from the main text.

Lemma 1. *In the Canonical Model, if (X, Y) is a symmetric pure strategy Bayesian equilibrium, then $x_i = y_i = m_{i,i}$ for all $i \in I$.*

The intuition is simple. Suppose that the candidates choose the same policy following two signal realizations, i and j, and suppose for simplicity that these are the only realizations for which they choose this policy; the proof uses (C4) to rule out other cases. Then, conditional on those realizations, each candidate expects to win the election with probability one half. If the candidates are not located at the median policy conditional on signals i and j, then either candidate could gain by moving slightly toward that conditional median: If A deviates in this way, then A's expected payoff given other signal realizations for B varies continuously with A's location, but A's payoff given realization j would jump discontinuously above one half. Therefore, a slight deviation would raise A's payoff, which is impossible in equilibrium.

Theorem 1 *(Necessity). In the Canonical Model, if (X, Y) is a pure strategy Bayesian equilibrium, then $x_i = y_i = m_{i,i}$ for all $i \in I$.*

Proof. First, consider a symmetric equilibrium (X, Y), where $x_i = y_i$ for all $i \in I$. By Lemma 1, $x_i = y_i = m_{i,i}$ for all $i \in I$. Now suppose there is an asymmetric equilibrium (X, Y), where $x_i \neq m_{i,i}$ for some $i \in I$, and define the strategy $Y' = X$ for candidate B. Then, by symmetry and interchangeability, (X, Y') is a symmetric Bayesian equilibrium with $x_i \neq m_{i,i}$, contradicting Lemma 1. □

It is natural to suppose that lower signals indicate lower values of μ and higher signals indicate higher ones. Then Theorem 1 implies that private polling causes candidates to "extremize" their locations.

Corollary 1. *In the Canonical Model, suppose there exists a signal $c \in I$ such that $i < c$ implies $m_{i,i} < m_i$ and $c < i$ implies $m_i < m_{i,i}$. If (X, Y) is a pure strategy Bayesian equilibrium, then $x_i < m_i$ for $i < c$ and $m_i < x_i$ for $i > c$.*

To highlight this result, suppose that c corresponds to an uninformative signal in the sense that the conditional median $m_{c,c}$ is equal to the unconditional median. Then it follows that in a pure strategy equilibrium, private polling magnifies platform divergence: Candidates bias their locations in the direction of their private signals past the median given their own signal and away from the unconditional median.

Theorem 1 gives a necessary, not a sufficient, condition for the existence of a pure strategy equilibrium. While we soon give sufficient conditions for existence, the next example shows that pure strategy equilibria do not always exist. In fact, the example shows more: Pure strategy equilibria may not exist in Canonical Models arbitrarily close to the Downsian Model. Thus, the example demonstrates a type of fragility of the Downsian median voter theorem. We return to this issue in our analysis of mixed strategies. There we show that even if pure strategy equilibria fail to exist in models close to the Downsian Model, mixed strategy equilibria do exist and are necessarily "close" to the Downsian equilibrium.

Example 1 *(Fragility of Pure Strategy Equilibrium in the Downsian Model).* Consider the Downsian Model in which $I = \{-1, 1\}$ and $P(-1,-1) = P(1,1) = \frac{1}{2}$, with conditional distributions $F_{-1,-1}$ and $F_{1,1}$ with point mass on $m_{-1,-1}$ and $m_{1,1} = m_{-1,-1} + 1$, respectively. Let $F_{-1,1}$ and $F_{1,-1}$ be degenerate on $m_{-1,1} = m_{1,-1} = (m_{-1,-1} + m_{1,1})/2$. By the median voter theorem, the unique equilibrium is (X, Y) defined by $x_{-1} = y_{-1} = m_{-1,-1}$ and $x_1 = y_1 = m_{1,1}$. Now define the sequences $\{F^n_{i,j} \mid i, j = -1, 1\}$ of conditional distributions as follows. For each $n \geqslant 2$, let $P^n(-1,-1) = P^n(1,1) = \frac{1}{2} - \frac{1}{n}$ and $P^n(-1,1) = P^n(1,-1) = \frac{1}{n}$. Let $F^n_{-1,-1}$ be the uniform distribution on $\Re$ with density n centered at $m_{-1,-1}$; let $F^n_{1,1}$ be the uniform distribution with density n centered at $m_{1,1}$; and let $F^n_{-1,1} = F^n_{1,-1}$ be the uniform distribution with density n^2 centered at $m_{-1,1} = m_{1,-1}$. Note that the conditional medians of $F^n_{-1,-1}$ and $F^n_{1,1}$ are fixed at $m_{-1,-1}$ and $m_{1,1}$, respectively, for all n. Furthermore, the upper bound of the support of $F^n_{-1,1} = F^n_{1,-1}$ is $m_{-1,1} + \frac{1}{2n^2}$. By Theorem 1, the only possible pure strategy Bayesian equilibrium in the n^{th} game is (X, Y) defined above. But we claim that (X, Y) is not an equilibrium, because A can deviate profitably to strategy $\hat{X}^n$ defined by $\hat{x}^n_{-1} = m_{-1,-1} + \frac{1}{n^2}$ and $\hat{x}^n_1 = x^n_1$. To see this, note that the difference $\Pi_A(\hat{X}^n, Y \mid -1) - \Pi_A(X, Y \mid -1)$ equals

$$P^n(-1\mid -1)\left[1 - F^n_{-1,-1}\left(\frac{m_{-1,-1} + \hat{x}^n_{-1}}{2}\right)\right] + P^n(1\mid -1)(1) - \frac{1}{2},$$

where we use the fact that $m_{-1,1} + \frac{1}{2n^2} = \frac{1}{2}(m_{-1,-1} + \frac{1}{n^2} + m_{1,1})$, which in turn implies $F^n_{-1,1}(\frac{\hat{x}^n_{-1} + m_{1,1}}{2}) = 1$. After substituting, this difference equals

$$\left(1 - \frac{2}{n}\right)\left(\frac{1}{2} - \frac{1}{2n}\right) + \frac{2}{n} - \frac{1}{2} = \frac{1}{2n} + \frac{1}{n^2} > 0,$$

establishing the claim. Since $F^k_{i,j} \to F_{i,j}$ weakly and $P^k(i,j) \to P(i,j)$ for $i, j = 1, -1$, the sequence of perturbed models can be chosen arbitrarily close to the Downsian Model. Thus, introducing even arbitrarily small amounts of private information to the Downsian Model can lead to the non-existence of pure strategy equilibrium. □

We now provide sufficient conditions for the pure strategy equilibrium to exist. To simplify our arguments, we provide separate conditions on the priors over signal pairs and on the distribution of μ conditional on signal realizations. Condition (C5) is a regularity condition on the conditional distributions that reinforces the symmetry present in the Canonical Model. Condition (C6) imposes a stochastic dominance-like restriction on the conditional distributions.

(C5) For all signals $i, j \in I$ with $P(i, j) > 0$, we have $m_{i,j} = (m_{i,i} + m_{j,j})/2$.
(C6) For all signals $i, j, k \in I$ with $i < j < k$, and for all $z \in [m_{i,i}, m_{k,k}]$, we have

$$F_{i,j}\left(\frac{m_{i,i} + z}{2}\right) \geqslant F_{j,k}\left(\frac{m_{k,k} + z}{2}\right). \tag{1}$$

Because $m_{i,i} \leqslant m_{j,j} \leqslant m_{k,k}$, the range over which inequality (1) holds is necessarily non-empty. To interpret inequality (1), note that it holds over $[m_{i,i}, m_{k,k}]$ if and only if

$$F_{i,j}(z) \geqslant F_{j,k}\left(\frac{m_{k,k} - m_{i,i}}{2} + z\right)$$

holds over $[\frac{m_{i,i}}{2}, \frac{m_{i,i}+m_{k,k}}{2}]$. That is, the distribution conditional on signals j and k when shifted to the left by $(m_{k,k} - m_{i,i})/2 \geqslant 0$ must dominate the distribution conditional on signals i and j. Condition (C6) is stronger than stochastic dominance in that $F_{j,k}$ is shifted to the left, but it is weaker in that the inequality must hold only over a given range.

Lastly, we impose a restriction on priors over signals, formalizing the idea that conditional on a candidate's signal, the probability the other candidate receives the same signal is high enough. Indeed, the condition is weaker than that, because it only restricts "net" probabilities.

(C7) For all signals $i \in I$,

$$\sum_{j\in I:\ j\leqslant i} P(j|i) \geqslant \sum_{j\in I:\ j>i} P(j|i) \quad \text{and} \quad \sum_{j\in I:\ j<i} P(j|i) \leqslant \sum_{j\in I:\ j\geqslant i} P(j|i).$$

Clearly, (C7) is most restrictive for the "extremal" signals, for which $P(i|i) \geqslant 1/2$ is implied by the condition, and its intuitive restrictiveness depends on the number of possible signals. With just two signals, for example, it is satisfied whenever signals are not negatively correlated.

Theorem 2 *(Sufficiency). In the Canonical Model, conditions* (C5)–(C7) *are sufficient for the existence of the unique pure strategy Bayesian equilibrium. In that equilibrium candidates locate at $m_{i,i}$ following signal $i \in I$.*

If we strengthen (C6) by imposing equality in (1), as BDS (2005) do for a special case of our model, then inspection of the proof of Theorem 2 reveals that under (C5), condition (C7) actually becomes necessary for the existence of the unique pure strategy equilibrium. Since condition (C7) is difficult to sustain when the number of possible signals is large, this highlights the importance of mixed strategies in the electoral model, to which we now turn.

5. Mixed strategy equilibrium

We now consider mixed strategy equilibria in the electoral game. We let candidate A randomize over campaign platforms following signal s according to a distribution G_s. A mixed strategy for A is a vector $G = (G_s)$ of such distributions, and a mixed strategy for B is a vector $H = (H_t)$. We let $G_s(z)^-$ and $H_t(z)^-$ be the left-hand limits of these distributions, e.g., $G_s(z)^- = \lim_{w\uparrow z} G_s(w)$. Accordingly, G_s has an atom at x if and only if $G_s(x) - G_s(x)^- > 0$.

To extend our definition of interim expected payoffs, we denote the probability that A wins using platform x following signal s when B uses platform y following signal t as

$$\pi_A(x, y|s, t) = \begin{cases} F_{s,t}(\frac{x+y}{2}) & \text{if } x < y, \\ 1 - F_{s,t}(\frac{x+y}{2}) & \text{if } y < x, \\ \frac{1}{2} & \text{if } x = y, \end{cases}$$

so that $\pi_B(\cdot|s,t) = 1 - \pi_A(\cdot|s,t)$. Then, given mixed strategies (G, H), candidate A's interim expected payoff conditional on signal s is

$$\Pi_A(G, H|s) = \sum_{t \in T} P(t|s) \int \pi_A(x, y|s, t) G_s(\mathrm{d}x) H_t(\mathrm{d}y).$$

Candidate B's interim payoff $\Pi_B(G, H|t)$ is defined analogously. Abusing notation slightly, let $\Pi_A(X, H|s)$ be A's expected payoff from the degenerate mixed strategy with $G_s(x_s) - G_s(x_s)^- = 1$ for all $s \in S$, and let $\Pi_B(G, Y|t)$ be the analogous expected payoff for B.

A *mixed strategy Bayesian equilibrium* is a strategy pair (G, H) such that $\Pi_A(G, H|s) \geqslant \Pi_A(G', H|s)$ for all signals $s \in S$ and all strategies G', and $\Pi_B(G, H|t) \geqslant \Pi_B(G, H'|t)$ for all signals $t \in T$ and all strategies H'. As with pure strategies, ex ante expected payoffs correspond to the ex ante probability of winning,

$$\Pi_A(G, H) = \sum_{s \in S} P(s) \Pi_A(G, H|s) \quad \text{and} \quad \Pi_B(G, H) = \sum_{t \in T} P(t) \Pi_B(G, H|t).$$

Thus, mixed strategy Bayesian equilibria of the electoral game are equilibria of a two-player, constant-sum game. In the Canonical Model, the game is symmetric, and we define a *symmetric* mixed strategy Bayesian equilibrium as an equilibrium strategy pair (G, H) with $G = H$.

The strategy X can be a discontinuity point of $\Pi_A(\cdot, H|s)$ only if H_t puts positive probability on policies y such that $\pi_A(\cdot, y|s, t)$ is discontinuous at x_s for some $t \in T$. By continuity of the conditional distributions, there is only one such policy, namely $y = x_s$. Because each H_t has at most a countable number of atoms, candidate A's expected payoff function is continuous on all but perhaps a countable set of pure strategies. Further, in equilibrium, if X is a continuity point of $\Pi_A(\cdot, H|s)$ in the support of G_s, then the expected payoff from X conditional on signal s must be $\Pi_A(G, H|s)$. Candidate A must therefore be indifferent over all such policies.

The next theorem provides a general existence result for mixed strategy equilibria in which candidates use mixed strategies with supports bounded as follows. Let $\overline{m} = \max\{m_{s,t}\colon s \in S, t \in T\}$ and $\underline{m} = \min\{m_{s,t}\colon s \in S, t \in T\}$. The interval defined by these "extreme" conditional medians is $M = [\underline{m}, \overline{m}]$. We say (G, H) *has support in* M if the candidates put probability one on M following all signal realizations: For all $s \in S$, $G_s(\overline{m}) - G_s(\underline{m})^- = 1$; and for all $t \in T$, $H_t(\overline{m}) - H_t(\underline{m})^- = 1$.

Theorem 3. *There exists a mixed strategy Bayesian equilibrium with support in M. Under* (C1) *and* (C2)*, there exists a symmetric mixed strategy Bayesian equilibrium with support in M.*

Our existence proof uses a result due to Dasgupta and Maskin (1986),[7] and it relies on the assumption of continuous conditional distributions: If discontinuities are permitted, then their

[7] An alternative is to define a corresponding game with endogenous sharing rule, and then to apply results of Simon and Zame (1990) and Jackson et al. (2002) to generate a selection of win probabilities as a function of the candidates' policy locations for which an equilibrium exists. However, our model assumes "symmetric" discontinuities: when the

weak lower semicontinuity condition can be violated. A weaker sufficient condition for existence in symmetric games that could be used is Reny's (1999) diagonal better reply security. However, Example A.1 in Appendix A shows that Reny's condition can be violated even if only very limited discontinuities are present. As a result, prospects for a more general result using known sufficient conditions for existence in discontinuous games seem poor.

We now study the continuity properties of the mixed strategy equilibrium correspondence as we vary the parameters of the model; specifically, the candidates' marginal prior on $S \times T$ and the conditional distributions of μ. We index specifications of the model by γ, where the marginal probability of (s,t) in game γ is $P^\gamma(s,t)$, and the distribution of μ conditional on s and t in γ is $F^\gamma_{s,t}$, with median $m^\gamma_{s,t}$. To consider continuity properties, we assume that γ lies in a metric space Γ, and that the indexing is continuous: For each $s \in S$ and $t \in T$, if $\gamma_n \to \gamma$, then $P^{\gamma_n}(s,t) \to P^\gamma(s,t)$ and $F^{\gamma_n}_{s,t} \to F^\gamma_{s,t}$ weakly. Denote the interval defined by the extreme conditional medians in game γ by $M(\gamma)$, and note that by the assumption of continuous indexing, the correspondence $M\colon \Gamma \rightrightarrows \Re$ so-defined is continuous.

Theorem 3 establishes the existence of a mixed strategy equilibrium for all $\gamma \in \Gamma$. Since the electoral game is constant-sum, a candidate's ex ante expected payoff in game γ is the same in all mixed strategy equilibria. Denote these payoffs, or "values," by $v_A(\gamma)$ and $v_B(\gamma)$. Furthermore, each candidate has an "optimal" mixed strategy that guarantees the candidate's value, no matter which strategy the opponent uses. If (C1) and (C2) hold for game γ, then the game is symmetric, so $v_A(\gamma) = v_B(\gamma) = 1/2$ trivially. It is more challenging to prove that $v_A(\gamma)$ and $v_B(\gamma)$ vary continuously in the parameters of the game while allowing for asymmetries, as when candidates have different polling technologies. We establish this continuity result next.

Theorem 4. *The mapping $v_A : \Gamma \to \Re$ is continuous.*

Letting $\mathcal{B}$ denote the Borel probability measures over X with the topology of weak convergence, define the mixed strategy Bayesian equilibrium correspondence $E\colon \Gamma \rightrightarrows \mathcal{B}^{S\cup T}$ so that $E(\gamma)$ consists of all equilibrium mixed strategy pairs (G,H). We have shown that this correspondence has non-empty values. The next result establishes an important continuity property of the equilibrium correspondence.

Theorem 5. *The correspondence $E : \Gamma \rightrightarrows \mathcal{B}^{S\cup T}$ has closed graph.*

Despite the discontinuities present in the electoral game, Theorem 5 delivers a desirable robustness property for mixed strategy equilibria: If we perturb the game slightly, then mixed strategy equilibria cannot be far away from the mixed strategy equilibria for the original specification of the game. The maintained assumption that conditional distributions are continuous in each specification $\gamma \in \Gamma$ is critical for our proof. When we provide an analogous robustness characterization of the unique pure strategy equilibrium for the Downsian Model in Section 6, we must use a distinct line of argument.

To this point, we have established existence of mixed strategy Bayesian equilibria with support in M, but there remains the possibility of equilibria that put positive probability outside that interval. We now show that in the Canonical Model, even without (C4), all mixed strategy Bayesian equilibria must have support in M.

---

candidates are equidistant from a conditional median, the probability of winning is one half for each candidate. The sharing rule approach delivers a selection satisfying this property only in the symmetric version of our model.

Theorem 6. *Under* (C1)–(C3)*, if* (G, H) *is a mixed strategy Bayesian equilibrium, then it has support in* M.

The symmetry assumed in Theorem 6 is *essential* for the bounds derived on equilibrium strategies: Without symmetry, it is possible that one candidate, say B, essentially competes with just one type of the other candidate, A, allowing the other types of candidate A to win with probability one, even if they choose platforms outside the interval M.

Example 2 *(Symmetry Needed for Equilibrium Bounds).* Let $I = \{-1, 1\}$, let each $F_{i,j}$ be a uniform distribution with density 2, with priors on $I \times I$ and conditional medians as follows:

	$i=-1$	$i=1$
$j=1$	$P(-1,1)=\epsilon$ $m_{-1,1}=-1$	$P(1,1)=\epsilon^2$ $m_{1,1}=1$
$j=-1$	$P(-1,-1)=1-\epsilon-\epsilon^2-\epsilon^3$ $m_{-1,-1}=0$	$P(1,-1)=\epsilon^3$ $m_{1,-1}=-1$

The conditional distribution following signal pairs $(-1, 1)$ or $(1, -1)$ has support $[-1.25, -0.75]$; the conditional distribution following $(-1, -1)$ has support $[-0.25, 0.25]$; and the conditional distribution following $(1, 1)$ has support $[0.75, 1.25]$. When ϵ is small, the conditional probability that candidate A receives signal $i = -1$ is close to one, regardless of B's signal. In contrast, the conditional probability that B receives signal j is close to one when A receives signal $i = j$. Let $x_{-1} = 0$, $x_1 = 1.25$, $y_{-1} = 0$, and $y_1 = -1.25$. The strategy profile (X, Y) so-defined is a Bayesian equilibrium, despite the fact that $x_1 = 1.25 > \overline{m} = 1$ and $y_1 = -1.25 < \underline{m} = -1$, violating the bound given in Theorem 6. To see this, first note that candidate A maximizes probability of winning following signal -1: Moving to the left from $x_{-1} = 0$ only decreases A's probability of winning when B receives signal -1; and moving far enough to increase A's probability of winning when B receives signal 1 means A must position at $x'_{-1} < -0.75$, but then A would win with probability zero in case B receives the more likely signal -1. Similarly, A maximizes probability of winning following signal 1: A already wins with probability one when B receives signal 1; and A cannot increase the probability of winning when B receives signal -1 without moving to the left of $y_{-1} = 0$, but then A would win with probability zero when B receives the more likely signal 1. A symmetric argument for B establishes the claim. □

We conclude this section by deriving the properties of atoms of mixed strategy equilibrium distributions. We prove that with a slight strengthening of condition (C4), the only possible atoms of the distributions G_i are at the conditional medians $m_{i,i}$.

(C4*) There exists a linear ordering $\leqslant$ of I, with asymmetric part $<$, such that: For all signals $i, i' \in I$ with $i < i'$,

$$\text{for all } j \in I, \text{ and for all } z \in M \text{ with } 0 < F_{i',j}(z) < 1, \text{ we have } F_{i',j}(z) < F_{i,j}(z).$$

As with (C4), the stochastic dominance condition in (C4*) arises naturally if higher signals are correlated with higher values of μ. Note that (C4*) implies that for signals $i < i'$, we have $F_{i',j}(z) \leqslant F_{i,j}(z)$ for all $z \in M$.

We now use condition (C4*) to characterize the location of mass points of equilibrium mixed strategies, thereby extending Lemma 1 to mixed strategy equilibria.

Lemma 2. *In the Canonical Model with* (C4*)*, let* (G, H) *be a symmetric mixed strategy Bayesian equilibrium. For all* $z \in M$*, if both candidates place positive probability mass on* z*, that is,* $G_i(z) - G_i(z)^- > 0$ *for some* $i \in I$*, then* $z = m_{i,i}$.

The intuition behind Lemma 2 is simple. Suppose that the candidates choose the same policy z with positive probability following some signal realization i; the proof uses (C4*) to show that the candidates cannot put positive probability on z after any other signal realizations. The argument then proceeds as in Lemma 1. Conditional on signal i, each candidate expects to choose z with positive probability, and if z is not equal to $m_{i,i}$, then a candidate, say A, can shift probability mass from z toward $m_{i,i}$ by an arbitrarily small amount. This increases A's expected payoff discretely when B chooses z, and it affects A's expected payoff continuously otherwise. Hence, a slight deviation increases A's expected payoff.

We now apply Lemma 2 to prove that in the Canonical Model with (C4*), the only possible atom of an equilibrium distribution, G_i, is the conditional median $m_{i,i}$. But for the strengthening of (C4), this result generalizes Theorem 1.

Theorem 7. *In the Canonical Model with* (C4*)*, let* (G, H) *be a mixed strategy Bayesian equilibrium. If* $G_i(z) - G_i(z)^- > 0$ *or* $H_i(z) - H_i(z)^- > 0$ *for some* $i \in I$*, then* $z = m_{i,i}$.

Proof. Let (G, H) be a mixed strategy Bayesian equilibrium, and suppose $G_i(z) - G_i(z)^- > 0$ for some $i \in I$, but $z \neq m_{i,i}$. By symmetry and interchangeability, (G, G) is an equilibrium. By Theorem 6, we must have $z \in M$. But then Lemma 2 implies $z = m_{i,i}$, a contradiction. □

Theorem 7 does not quite allow us to use differentiable methods to analyze mixed strategy equilibria. However, it may be reasonable to restrict attention to a subset of "regular" mixed strategy equilibria for which the distribution following each signal is differentiable at all continuity points. This restriction solves technical problems at the cost of possibly omitting pathological equilibria. We can then decompose the probability measure generated by any equilibrium distribution into a finite number of degenerate measures and an absolutely continuous measure with density defined at all but at most a finite number of policies. The usual first-order condition must then be satisfied at any platform in the support of a candidate's equilibrium density following any signal, generating a system of differential equations. See BDS for this approach.

6. Robustness of the median voter theorem

In this section, we first derive the robustness of the probabilistic version of the median voter theorem as an easy consequence of Theorems 5 and 6. Observe that after decomposing the Probabilistic Voting Model into its component games, Theorem 6 applies to each one separately. Since the set of medians for the component game corresponding to the signal pair (i, i) is just the singleton $\{m_i\}$, it immediately follows that the unique mixed strategy Bayesian equilibrium of the component game must be the point mass on m_i for both candidates. In other words, the unique pure strategy Bayesian equilibrium of the Probabilistic Voting Model is unique among all *mixed strategy* Bayesian equilibria. Then Theorem 5 gives us a strong continuity result: In models close to the Probabilistic Voting Model, mixed strategy equilibria must be close in the sense of weak convergence to the pure strategy equilibrium in which candidates locate at m_i following signal i.

Theorem 8. *Assume γ is a Probabilistic Voting Model. If (G, H) is a mixed strategy Bayesian equilibrium of γ, then the candidates locate at m_i^γ with probability one following signal each $i \in I$, i.e., $G_i(m_i^\gamma) - G_i(m_i^\gamma)^- = H_i(m_i^\gamma) - H_i(m_i^\gamma)^- = 1$ for all $i \in I$. Given any sequence $\{\gamma_n\} \to \gamma$ in Γ and given $(G^n, H^n) \in E(\gamma_n)$ for all n, then for all $i \in I$, $\{G_i^n\}$ and $\{H_i^n\}$ converge weakly to the point mass on m_i^γ.*

We now turn to the robustness of the median voter theorem in the Downsian Model. Because that model is characterized by discontinuous conditional distributions of the median voter's ideal point, the robustness result of Theorem 5 does not apply. As Example 3 shows below, such discontinuities are more than minor technicalities: The introduction of discontinuous conditional distributions into the Canonical Model can actually overturn the robustness results of the previous section. The Downsian Model is distinguished, however, by the fact that candidates have complete information about the median voter's location, and it is this structure that underlies the theorems to follow.

To extend our results to the Downsian Model, we say that a sequence $\{\gamma_n\}$ in Γ is *approximately Downsian* if: (a) for each n, γ_n satisfies condition (C1); (b) for all $i \in I$, $P^{\gamma_n}(i|i)$ converges to one; and (c) for all $i \in I$, $\{F_{i,i}^{\gamma_n}\}$ converges weakly to the point mass on some m_i. While the models indexed by γ_n lie in Γ and satisfy our maintained assumption of continuous conditional distributions, the "limiting model" (which is implicit here) is Downsian and exhibits discontinuous conditional distributions. We now extend the continuity result of Theorem 4 to approximately Downsian sequences, showing that each candidate's ex ante probability of winning goes to one-half.

Theorem 9. *If the sequence $\{\gamma_n\}$ in Γ is approximately Downsian, then $v_A(\gamma_n) \to 1/2$.*

Proof. Let $\{\gamma_n\}$ be approximately Downsian. We will show that $\liminf v_A(\gamma_n) \geqslant 1/2$, and a symmetric argument for candidate B will then imply that $\lim v_A(\gamma_n) = 1/2$. For each n, let $P^n(i, j)$ denote the prior probability of signal pair (i, j) in γ_n; let $F_{i,j}^n$ denote the distribution of μ conditional on signal pair (i, j); and let $m_{i,j}^n$ denote the median of $F_{i,j}^n$. Now let X^n be defined by $x_i^n = m_{i,i}^n$ for all $i \in I$, and let Y^n be an arbitrary pure strategy for B. Then A's expected payoff from (X^n, Y^n) conditional on signal i in game γ_n, denoted $\Pi_A^n(X^n, Y^n|i)$, is equal to

$$\sum_{j \in I: x_i^n < y_j^n} P^n(j|i) F_{i,j}^n\left(\frac{x_i^n + y_j^n}{2}\right) + \sum_{j \in I: y_j^n < x_i^n} P^n(j|i)\left(1 - F_{i,j}^n\left(\frac{x_i^n + y_j^n}{2}\right)\right)$$
$$+\frac{1}{2} \sum_{j \in I: x_i^n = y_j^n} P^n(j|i).$$

Note that $P^n(i|i) \to 1$ and $P^n(j|i) \to 0$ for $j \neq i$, and define the sequence $\{\Phi^n\}$ by

$$\Phi^n = \begin{cases} F^n_{i,i}(\frac{m^n_{i,i}+y^n_i}{2}) & \text{if } m^n_{i,i} < y^n_i, \\ 1 - F^n_{i,i}(\frac{m^n_{i,i}+y^n_i}{2}) & \text{if } y^n_i < m^n_{i,i}, \\ \frac{1}{2} & \text{otherwise.} \end{cases}$$

Then $\liminf \Pi^n_A(X^n, Y^n|j) \geqslant \liminf \Phi^n \geqslant 1/2$, where the last inequality follows from $\Phi^n \geqslant 1/2$ for all n. Thus, A can guarantee an expected payoff arbitrarily close to one half as n goes to infinity, and we conclude that $\liminf v_A(\gamma_n) \geqslant 1/2$, as required. □

Whereas Theorem 4 established continuity of the payoffs for the class of models with continuous conditional distributions, Theorem 9 extends this result to certain "boundary points" of this class exhibiting discontinuous conditional distributions, namely the Downsian Models. The next example shows that our continuity results do not extend when complete information in the Downsian Model is relaxed.

Example 3 *(Continuity of Value Violated with Discontinuous Conditional Distributions).* Let $I = \{0, 1\}$ and $P(i, j) = 1/4$ for all $i, j \in I$, with conditional distributions $F_{i,j}$ equal to the point mass on zero for all $i, j \in I$. For each n, let $F^n_{0,j}$, $j = 0, 1$, be the uniform distribution with density n centered at zero, and let $F^n_{1,j}$, $j = 0, 1$, be the uniform distribution with density n centered at $1/n$. Note that $F^n_{i,j} \to F_{i,j}$ weakly and that the supports of $F^n_{0,j}$ and $F^n_{1,j}$ are contiguous for all $i, j \in I$. Thus, for each n, candidate A has full information about the distribution of the median voter, whereas candidate B receives no information: Conditional on $j = 0$ and $j = 1$, the distribution of the median voter has mean 0 and mean $1/n$ with equal probability. For each n, define $x^n_0 = 0$, $x^n_1 = 1/n$, and $y^n_0 = y^n_1 = 0$, and note that the pure strategy profile (X^n, Y^n) so-defined is a Bayesian equilibrium. In it, candidate A always locates at the median, winning for sure when $i = 1$ and matching B when $i = 0$. Thus, candidate A's equilibrium expected payoff is $3/4$ and the value of the game for A is $v^n_A = 3/4$ for all n. In contrast, the limiting model is symmetric, and payoffs for the candidates must be one half in any equilibrium.

We have shown that the value in models close to Downsian must be close to one half, and the next result pushes this further: The Bayesian equilibrium mixed strategy distributions used following any signal realization must themselves be close to the location of the median in the Downsian Model conditional on that signal. Thus, though Example 1 shows that pure strategy equilibria may fail to exist near the Downsian Model, Theorem 3 shows that mixed strategy equilibria do exist, and these equilibria must be close to the Downsian equilibrium.

Theorem 10. *If the sequence $\{\gamma_n\}$ is approximately Downsian and $(G^n, H^n) \in E(\gamma_n)$ for all n, then for all $i \in I$, $\{G^n_i\}$ and $\{H^n_i\}$ converge weakly to the point mass on m_i.*

Proof. Let $\{\gamma_n\}$ be approximately Downsian, let $(G^n, H^n) \in E(\gamma_n)$ for each n, and suppose that H^n_i does not converge to the point mass on m_i for some $i \in I$. Then there exists an interval (a, b) containing m_i such that either $\limsup H^n_i(a) > 0$ or $\liminf H^n_i(b)^- < 1$. Without loss of generality, assume the latter, and consider a subsequence (still indexed by n) along which this lim inf is achieved, i.e., $\lim H_i^n(b)^- < 1$. Using the notation from the proof of Theorem 9, define X^n so that $x_{i'}^n = m_{i',i}^n$ for each $i' \in I$ and each n. As in the proof of Theorem 9, we can show that

$$\liminf_{n\to\infty} \inf_{Y\in\Re^T} \Pi_A^n\big(X^n, Y|i'\big) \geqslant \frac{1}{2} \tag{2}$$

for all $i' \in I$. Moreover, if we restrict candidate B to strategies such that the candidate locates to the right of b following signal i, then this inequality becomes strict conditional on candidate A also receiving signal i:

$$\liminf_{n\to\infty} \inf_{Y\in\Re^T:\ y_t\geqslant b} \Pi_A^n\big(X^n, Y|i\big) > \frac{1}{2}. \tag{3}$$

To see this, let Y^n be an arbitrary pure strategy for B such that $y_i^n \geqslant b$, and take c such that $m_i < c < (m_i + b)/2$. Note that

$$\liminf \Phi^n = \liminf F_{i,i}^n\left(\frac{m_{i,i}^n + y_t^n}{2}\right) \geqslant \liminf F_{i,i}^n(c)^- \geqslant F_{i,i}(c)^- > \frac{1}{2},$$

where the second-to-last inequality follows from weak convergence of $F_{i,i}^n$ to $F_{i,i}$, and the strict inequality follows from $c > m_i$ and our assumption that $F_{i,i}$ is the point mass on m_i. By (2) and (3), we have

$$\begin{aligned}
&\liminf \Pi_A^n\big(X^n, H^n|i\big) \\
&\quad\geqslant \liminf \int\limits_{(-\infty,b)} \pi_A^n\big(m_{i,i}^n, y|i,i\big) H_i^n(\mathrm{d}y) + \liminf \int\limits_{[b,\infty)} \pi_A^n\big(m_{i,i}^n, y|i,i\big) H_i^n(\mathrm{d}y) \\
&\quad\geqslant \frac{1}{2} \lim H_i^n(b)^- + F_{i,i}(c)^-\big(1 - \lim H_i^n(b)^-\big),
\end{aligned}$$

which, by $F_{i,i}(c)^- > 1/2$ and $\lim H_i^n(b)^- < 1$, is greater than $1/2$. Thus, $\liminf \Pi_A^n(G^n, H^n) > 1/2$, i.e., $\liminf v_A(\gamma_n) > 1/2$. But $v_A(\gamma_n) \to 1/2$ by Theorem 9, a contradiction. □

In Example 3, the equilibrium platforms of the candidates converge to zero, the unique equilibrium of the limiting model. This leaves the possibility that, while our continuity result for the value does not extend, the upper hemicontinuity result of Theorem 10 does extend beyond the Downsian Model. However, Example A.2 in Appendix A shows that this too does not hold.

7. Conclusion

This paper provides theoretical results on elections with privately-informed candidates, where the source of private information may be a candidate's personal experiences or polls conducted by the candidate's campaign organization or political party. We characterize the unique pure strategy Bayesian equilibrium of the electoral game when it exists, and we give results on the existence and continuity properties of mixed strategy Bayesian equilibria, as well as bounds on the supports of mixed strategy equilibria and restrictions on equilibrium atoms. By posing the analysis within a general framework, not only do we strengthen the foundation of these results, but we open the possibility of developing special cases of interest in applications. We do just that for the Downsian Model and the Probabilistic Voting Model, the most commonly used electoral models. In particular, we show that pure strategy equilibria of the Downsian Model are fragile, but we provide robustness results in mixed strategies: If a small amount of private information is added to either model, then there exist mixed strategy Bayesian equilibria, and in all equilibria a candidate must place probability near one close to the median conditional on her signal.

To obtain more explicit characterizations, BDS consider a natural specification of the Canonical Model, where the median μ is decomposed as $\mu = \alpha + \beta$. The component α is uniformly distributed on $[-a, a]$, and β is an independently-distributed discrete random variable. Candidates share the same set of signals realizations. Signals depend stochastically on the realization of β and are independent of α, the latter capturing aspects of voter preferences at the time of the election that are not revealed by the candidates' polls. BDS prove that there is a unique "ordered" mixed strategy Bayesian equilibrium and derive its properties. They obtain an explicit solution for the mixed strategy equilibrium when the pure strategy equilibrium fails to exist: A candidate receiving a "moderate" signal i locates at the median $m_{i,i}$ conditional on both candidates receiving that signal, while candidates who receive "extreme" signals mix, moderating their platforms relative to the pure strategy equilibrium choices. Welfare analysis shows that even though candidates locate more extremely than if they had simply targeted the median voter given their private information, all voters may be better off if candidates chose *even more* extreme platforms.

The model assumes some structural aspects of elections that may be of interest to develop more fully, such as the nature of the polling process and the determinants of the median voter's ideal policy. Several extensions suggest themselves as well. First, because the strategic value of better information is always positive for candidates, it is conceptually straightforward to endogenize the choice of costly polling technologies by candidates. It would also be worthwhile to determine how outcomes are affected if candidates have ideological preferences, and to endogenize contributions by ideologically-motivated lobbies to fund polling by candidates. Finally, as Ledyard (1989) observes, it would be useful to uncover how equilibrium outcomes are altered if candidates choose platforms sequentially. Then, the second candidate can see where the first locates, and hence may be able to unravel the latter's signal, before locating.

Acknowledgments

We thank Jean-Francois Mertens for helpful discussions during his visit to the Wallis Institute of Political Economy at the University of Rochester. The paper also benefited from discussions with Roger Myerson and Tom Palfrey. Dan Bernhardt gratefully acknowledges support from the National Science Foundation, grant number SES-0317700. John Duggan gratefully acknowledges support from the National Science Foundation, grant number SES-0213738.

Appendix A

Lemma A.1. *Let (G, H) be any pair of mixed strategies. For all $s \in S$ and all $w, z \in \Re$, one of three possibilities obtains*: *Either*

$$\sum_{t\in T} P(t|s)\big[H_t(z) - H_t(z)^-\big]\big[F_{s,t}(w)\big]$$
$$= \frac{1}{2}\sum_{t\in T} P(t|s)\big[H_t(z) - H_t(z)^-\big] = \sum_{t\in T} P(t|s)\big[H_t(z) - H_t(z)^-\big]\big[1 - F_{s,t}(w)\big],$$

or

$$\sum_{t\in T} P(t|s)\big[H_t(z) - H_t(z)^-\big]\big[F_{s,t}(w)\big]$$
$$< \frac{1}{2}\sum_{t\in T} P(t|s)\big[H_t(z) - H_t(z)^-\big] < \sum_{t\in T} P(t|s)\big[H_t(z) - H_t(z)^-\big]\big[1 - F_{s,t}(w)\big],$$

or the reverse inequalities hold. Likewise for G and all $t \in T$ and all $w, z \in \Re$.

Proof. Given $s \in S$ and $z \in \Re$, the first possibility obtains if $P(t|s)[H_t(z) - H_t(z)^-] = 0$ for all $t \in T$. Suppose $P(t|s)[H_t(z) - H_t(z)^-] > 0$ for some $t \in T$, and define the function F^* by:

$$F_s^*(w) = \frac{\sum_{t\in T} P(t|s)[H_t(z) - H_t(z)^-]F_{s,t}(w)}{\sum_{t\in T} P(t|s)[H_t(z) - H_t(z)^-]}.$$

Since each $F_{s,t}$ is continuous at w, the three possibilities above correspond to the three possibilities $F_s^*(w) = 1/2$, $F_s^*(w) < 1/2$, and $F_s^*(w) > 1/2$. □

Lemma A.2. *Let (G, H) be a mixed strategy Bayesian equilibrium. For all $z \in \Re$, if $G_{s'}(z) - G_{s'}(z)^- > 0$ for some $s' \in S$ and $H_{t'}(z) - H_{t'}(z)^- > 0$ for some $t' \in T$ with $P(s', t') > 0$, then*

$$\sum_{t\in T} P(t|s')\big[H_t(z) - H_t(z)^-\big]\big[F_{s',t}(z)\big]$$
$$= \frac{1}{2}\sum_{t\in T} P(t|s')\big[H_t(z) - H_t(z)^-\big] = \sum_{t\in T} P(t|s')\big[H_t(z) - H_t(z)^-\big]\big[1 - F_{s',t}(z)\big]$$

and

$$\sum_{s\in S} P(s|t')\big[G_s(z) - G_s(z)^-\big]\big[F_{t',s}(z)\big]$$
$$= \frac{1}{2}\sum_{s\in S} P(s|t')\big[G_s(z) - G_s(z)^-\big] = \sum_{s\in S} P(s|t')\big[G_s(z) - G_s(z)^-\big]\big[1 - F_{t',s}(z)\big].$$

Proof. We prove the first equalities. If they do not hold for some z and some s' and t' with $P(s', t') > 0$, then, by Lemma A.1, we may assume that one of the following holds:

$$\sum_{t\in T} P(t|s')\big[H_t(z) - H_t(z)^-\big]\big[1 - F_{s',t}(z)\big] > \frac{1}{2}\sum_{t\in T} P(t|s')\big[H_t(z) - H_t(z)^-\big], \tag{A.1}$$

$$\sum_{t\in T} P(t|s')\big[H_t(z) - H_t(z)^-\big]\big[F_{s',t}(z)\big] > \frac{1}{2}\sum_{t\in T} P(t|s')\big[H_t(z) - H_t(z)^-\big].$$

We focus on the first inequality, as a symmetric proof addresses the second. For each $t \in T$, let λ_t denote the probability measure generated by the distribution H_t, let μ_t denote the degenerate measure with mass $H_t(z) - H_t(z)^-$ on z, and let $\nu_t = \lambda_t - \mu_t$. Let $\{x^n\}$ be a sequence decreasing to z, and let G^n be the mixed strategy defined by replacing G_s in G with the point mass on x^n. Let $\pi_t(w) = \pi_A(z, w|s', t)$ denote A's probability of winning using z when B receives signal t and chooses platform w, and let $\pi_t^n(w) = \pi_A(x^n, w|s', t)$ denote A's probability of winning using x^n when B receives signal t and chooses platform w. Then

$$\begin{aligned}\Pi_A\big(G^n, H|s'\big) - \Pi_A(G, H|s') &= \sum_{t\in T} P(t|s') \int \big[\pi_t^n(w) - \pi_t(w)\big]\lambda_t(\mathrm{d}w)\\ &= \sum_{t\in T} P(t|s')\big[H_t(z) - H_t(z)^-\big]\left[1 - F_{s',t}\left(\frac{z+x^n}{2}\right) - \frac{1}{2}\right]\\ &\quad + \sum_{t\in T} P(t|s') \int \big[\pi_t^n(w) - \pi_t(w)\big]\nu_t(\mathrm{d}w).\end{aligned}$$

Since $\pi_t^n - \pi_t \to 0$ almost everywhere (ν_t), the corresponding integral terms above converge to zero. Thus, $\lim_{n\to\infty} \Pi_A(G^n, H|s') - \Pi_A(G, H|s')$ equals

$$\sum_{t\in T} P(t|s')\big[H_t(z) - H_t(z)^-\big]\big[1 - F_{s',t}(z)\big] - \frac{1}{2}\sum_{t\in T} P(t|s')\big[H_t(z) - H_t(z)^-\big],$$

which is positive by (A.1). Hence, $\Pi_A(G^n, H|s) > \Pi_A(G, H|s)$ for high enough n, a contradiction. □

Proof of Lemma 1. Let G be a mixed strategy such that for all $i \in I$, $G_i(x_i) - G_i(x_i)^- = 1$. Let $I' = \{i \in I : x_i = z\}$, and take any $i \in I'$. Lemma A.2 implies

$$\sum_{i'\in I'} \frac{P(i'|i)}{P(I'|i)}\big[F_{i,i'}(z)\big] = \frac{1}{2}, \tag{A.2}$$

where we use $G_{i'}(z) - G_{i'}(z)^- = 1$ for $i' \in I'$ and $G_{i'}(z) - G_{i'}(z)^- = 0$ otherwise. Thus, $z = m_{i,I'}$. If there exists $j \neq i$ such that $j \in I'$, then we have $m_{i,I'} = z = m_{j,I'}$, contradicting (C4). Therefore, (A.2) reduces to $F_{i,i}(z) = 1/2$, i.e., $z = m_{i,i}$. □

Proof of Theorem 2. We show that (X, Y) is an equilibrium, where $x_i = y_i = m_{i,i}$ for all $i \in I$. Without loss of generality, we focus on candidate B's best response problem following signal j. Consider a deviation to strategy Y'. There are two cases: $y'_j < m_{j,j}$ and $m_{j,j} < y'_j$. In the first case, let

$$\mathcal{G} = \{i \in I\colon\ m_{i,i} \leqslant y'_j\} \quad \text{and} \quad \mathcal{L} = \{k \in I\colon\ m_{j,j} \leqslant m_{k,k}\}.$$

Note that for all $i \in I \setminus (\mathcal{G} \cup \mathcal{L})$, we have $y'_j < m_{i,i} < m_{j,j}$. Hence, for i with $P(i|j) > 0$,

$$F_{i,j}\left(\frac{y'_j + m_{i,i}}{2}\right) - \left[1 - F_{i,j}\left(\frac{m_{i,i} + m_{j,j}}{2}\right)\right] \leqslant 0,$$

where we use (C5) to deduce that $F_{i,j}(\frac{y'_j+m_{i,i}}{2}) \leqslant 1/2$ and $F_{i,j}(\frac{m_{i,i}+m_{j,j}}{2}) = 1/2$. That is, B's gains from deviating when A receives signal $i \in I \setminus (\mathcal{G} \cup \mathcal{L})$ are non-positive. Therefore, the change in B's interim expected payoff, $\Pi_B(X, Y'|j) - \Pi_B(X, Y|j)$, is less than or equal to

$$\begin{aligned}&\sum_{i\in\mathcal{G}} P(i|j)\left[1 - F_{i,j}\left(\frac{m_{i,i} + y'_j}{2}\right) - \left(1 - F_{i,j}\left(\frac{m_{i,i} + m_{j,j}}{2}\right)\right)\right]\\ &\quad + \sum_{k\in\mathcal{L}} P(k|j)\left[F_{j,k}\left(\frac{y'_j + m_{k,k}}{2}\right) - F_{j,k}\left(\frac{m_{j,j} + m_{k,k}}{2}\right)\right]\\ &= \sum_{i\in\mathcal{G}} P(i|j)\left[\frac{1}{2} - F_{i,j}\left(\frac{m_{i,i} + y'_j}{2}\right)\right] + \sum_{k\in\mathcal{L}} P(k|j)\left[F_{j,k}\left(\frac{y'_j + m_{k,k}}{2}\right) - \frac{1}{2}\right],\end{aligned}$$

which is non-positive as long as

$$\sum_{i\in\mathcal{G}} P(i|j)\left[\frac{1}{2} - F_{i,j}\left(\frac{m_{i,i}+y'_j}{2}\right)\right] \leqslant \sum_{k\in\mathcal{L}} P(k|j)\left[\frac{1}{2} - F_{j,k}\left(\frac{y'_j+m_{k,k}}{2}\right)\right]. \tag{A.3}$$

Let i^* minimize $F_{i,j}(\frac{m_{i,i}+y'_j}{2})$ over $\mathcal{G}$, and let k^* maximize $F_{j,k}(\frac{m_{i,i}+y'_j}{2})$ over $\mathcal{L}$. Then inequality (A.3) holds if

$$\left[\frac{1}{2} - F_{i^*,j}\left(\frac{m_{i^*,i^*}+y'_j}{2}\right)\right]\sum_{i\in\mathcal{G}} P(i|j) \leqslant \left[\frac{1}{2} - F_{j,k^*}\left(\frac{y'_j+m_{k^*,k^*}}{2}\right)\right]\sum_{k\in\mathcal{L}} P(k|j). \tag{A.4}$$

Note that because of (C4), we have $\mathcal{G} \subseteq \{i \in I \mid i < j\}$ and $\mathcal{L} = \{i \in I \mid j \leqslant i\}$, so (C7) implies $\sum_{i\in\mathcal{G}} P(i|j) \leqslant \sum_{i\in\mathcal{L}} P(i|j)$. Furthermore, $i^* < j < k^*$ and $y'_j \in [m_{i^*,i^*}, m_{k^*,k^*}]$, so (C6) implies

$$F_{i^*,j}\left(\frac{m_{i^*,i^*}+y'_j}{2}\right) \geqslant F_{j,k^*}\left(\frac{y'_j+m_{k^*,k^*}}{2}\right).$$

Thus, inequality (A.4) holds, and it is unprofitable for B to deviate to $y'_j < m_{j,j}$. A symmetric argument applies for deviations $y'_j > m_{j,j}$. □

Proof of Theorem 3. We use the existence theorem of Dasgupta and Maskin (1986) for multi-player games with one-dimensional strategy spaces. To apply this result, view the electoral game as a $(|S| + |T|)$-player game in which each type (corresponding to different signal realizations) of each candidate is a separate player. Player s (or t) has strategy space $M \subseteq \Re$, a compact and convex set, with pure strategies x_s (or y_t). Then $(X, Y) = (x_s, y_t)_{s\in S, t\in T}$ is a pure strategy *profile,* one for each type. Let (X_{-s}, Y) denote the result of deleting x_s from (X, Y). The payoff function of player $s \in S$ is $U_s(X, Y) = P(s)\Pi_A(X, Y|s)$, and the payoff function of player $t \in T$ is $U_t(X, Y) = P(t)\Pi_B(X, Y|t)$. The space of mixed strategies for each player type s (or t) is $\mathcal{M}$, the Borel probability measures on M, with mixed strategies denoted G_s (or H_t). Then (G, H) is a mixed strategy *profile,* one for each type. Note that

$$\sum_{s\in S} U_s(X, Y) + \sum_{t\in T} U_t(X, Y) = 1,$$

for all X and Y, so that the total payoff is trivially upper semi-continuous. Furthermore, payoffs are between zero and one, so they are bounded. Note that U_s is discontinuous at (X, Y) only if $x_s = y_t$ for some $t \in T$. Therefore, the discontinuity points of U_s lie in a set that can be written as $A^*(s)$, as in Dasgupta and Maskin's equation (2). The discontinuity points of U_t lie in a similar set. It remains to show that U_s (likewise U_t) is weakly lower semi-continuous in x_s; that is, for all $x_s \in M$, there exists a $\lambda \in [0, 1]$ such that for all (X_{-s}, Y),

$$U_s(X, Y) \leqslant \lambda \liminf_{z\downarrow x_s} U_s(z, X_{-s}, Y) + (1-\lambda) \liminf_{z\uparrow x_s} U_s(z, X_{-s}, Y).$$

It is straightforward to verify that this condition holds with equality for $\lambda = 1/2$. Let $T^+ = \{t \in T\colon x_s < y_t\}$, let $T^- = \{t \in T\colon y_t < x_s\}$, and let $T^0 = \{t \in T\colon x_s = y_t\}$. Since

$$U_s(X, Y) = \sum_{t\in T^-} P(s,t)\left(1 - F_{s,t}\left(\frac{x_s+y_t}{2}\right)\right) + \sum_{t\in T^0} \frac{P(s,t)}{2} + \sum_{t\in T^+} P(s,t)F_{s,t}\left(\frac{x_s+y_t}{2}\right),$$

it follows that $\liminf_{z\uparrow x_s} U_s(z, X_{-s}, Y)$ equals

$$\sum_{t\in T^-} P(s,t)\left(1 - F_{s,t}\left(\frac{x_s + y_t}{2}\right)\right) + \sum_{t\in T^0\cup T^+} P(s,t)F_{s,t}\left(\frac{x_s + y_t}{2}\right)$$

and $\liminf_{z\downarrow x_s} U_s(z, X_{-s}, Y)$ equals

$$\sum_{t\in T^-\cup T^0} P(s,t)\left(1 - F_{s,t}\left(\frac{x_s + y_t}{2}\right)\right) + \sum_{t\in T^+} P(s,t)F_{s,t}\left(\frac{x_s + y_t}{2}\right).$$

The claim

$$U_s(X,Y) = \frac{1}{2}\liminf_{z\downarrow x_s} U_s(z, X_{-s}, Y) + \frac{1}{2}\liminf_{z\uparrow x_s} U_s(z, X_{-s}, Y)$$

then follows immediately. The same argument can be used to verify that U_t is weakly lower semi-continuous. Then by Dasgupta and Maskin's (1986) Theorem 5, there exists a mixed strategy equilibrium of the multi-player game, and hence of the electoral game when strategies are restricted to M. To see that following a signal s, candidate A has no profitable deviations outside M, take any $x > \overline{m}$. For all $t \in T$ and all $y \in M$, we have $\pi_A(\overline{m}, y|s,t) \geqslant \pi_A(x, y|s,t)$. Let G' be any deviation such that G'_s puts probability one on $x_s > \overline{m}$, and let G'' put probability one on $\overline{m}$ instead. Then $\Pi_A(G', H|s) \leqslant \Pi_A(G'', H|s) \leqslant \Pi_A(G, H|s)$. A similar argument applies when $x < \underline{m}$, yielding the claim. Because the electoral game is a two-player, symmetric constant-sum game, existence and interchangeability of equilibria imply that there exists a symmetric mixed strategy equilibrium. □

Example A.1 *(Diagonal Better Reply Security Violated with Discontinuous Conditional Distributions).* Consider a discontinuous version of the Canonical Model in which $I = \{1, 2, 3\}$ and, for all $i, j \in I$, $F_{i,j}$ is the point mass on $m_{i,j}$, given in the table below. We assign priors on $I \times I$ as indicated there.

$j = 3$	$P(1,3) = 0.16\epsilon$ $m_{1,3} = 1.8$	$P(2,3) = 0.41\epsilon$ $m_{2,3} = 2.3$	$P(3,3) = 1 - 2.89\epsilon$ $m_{3,3} = 2.6$
$j = 2$	$P(1,2) = 0.45\epsilon$ $m_{1,2} = 1$	$P(2,2) = 0.05\epsilon$ $m_{2,2} = 1.9$	$P(3,2) = 0.41\epsilon$ $m_{3,2} = 2.3$
$j = 1$	$P(1,1) = 0.8\epsilon$ $m_{1,1} = 0$	$P(2,1) = 0.45\epsilon$ $m_{2,1} = 1$	$P(3,1) = 0.16\epsilon$ $m_{3,1} = 1.8$
	$i = 1$	$i = 2$	$i = 3$

Define the mixed strategy G as follows: $G_1(m_{1,1}) - G_1(m_{1,1})^- = \alpha$, $G_1(m_{1,2}) - G_1(m_{1,2})^- = 1 - \alpha$, $G_2(2) - G_2(2)^- = 1$, and $G_3(m_{3,3}) - G_3(m_{3,3})^- = 1$, and set $\alpha = 0.8$ and $H = G$. That is, after signal 1, the candidates mix between two conditional medians, $m_{1,1}$ and $m_{1,2}$; after signal 2, the candidates adopt the platform 2 (which does not correspond to a conditional median); and after signal 3, the candidates adopt the conditional median $m_{3,3}$. Reny's (1999) diagonal better reply security requires that if (G, H) is not a mixed strategy Bayesian equilibrium, then candidate A has a mixed strategy deviation that is profitable, even if B's mixed strategy is allowed to vary within some open set. Specifically, there must exist a mixed strategy $\hat{G}$ for A and an open set $\mathcal{H}$ of mixed strategies for B such that $H \in \mathcal{H}$ and $\inf_{\hat{H} \in \mathcal{H}} \Pi_A(\hat{G}, \hat{H}) > 1/2$.[8]

Note that this strategy is clearly a best response following signal 3, if we set $\epsilon > 0$ sufficiently small. Following signal 2, candidate A would tie with candidate B in case B received signal 1 and positioned at $m_{1,1} = 0$, would lose to B in case B received signal 1 and positioned at 1, and would tie with B in case B received signals 2 or 3: A's expected payoff would be

$$\begin{aligned} P(2)\Pi_A(G, H|2) &= \alpha P(1,2)(0.5) + (1-\alpha)P(1,2)(0) + P(2,2)(0.5) + P(2,3)(0.5) \\ &= \epsilon(0.225\alpha + 0.23) = 0.41\epsilon, \end{aligned}$$

where we weight the interim payoff by the marginal probability of signal 2. If candidate A moved to the left following signal 2, then the candidate could move a small enough amount to $x_2' \in (1.8, 2)$ to win against candidate B in case B received signal 1 and positioned at $m_{1,1} = 0$ or received signal 2, but would then lose against B in case B received signal 3: A's expected payoff would again be

$$\begin{aligned} P(2)\Pi_A(X', H|2) &= \alpha P(1,2)(1) + (1-\alpha)P(1,2)(0) + P(2,2)(1) + P(2,3)(0) \\ &= \epsilon(0.45\alpha + 0.05) = 0.41\epsilon. \end{aligned}$$

Moving further to the left, A could do no better than locate at $m_{1,2} = 1$, which yields an expected payoff of

$$\begin{aligned} P(2)\Pi_A(X', H|2) &= \alpha P(1,2)(1) + (1-\alpha)P(1,2)(0.5) + P(2,2)(0) + P(2,3)(0) \\ &= \epsilon(0.225\alpha + 0.225) = 0.405\epsilon. \end{aligned}$$

Moving to the right, the best A could do would be to win against B in case B received signal 3 and lose otherwise, which yields an expected payoff of

$$P(2)\Pi_A(X', H|2) = \alpha P(1,2)(0) + (1-\alpha)P(1,2)(0) + P(2,2)(0) + P(2,3)(1) = 0.41\epsilon.$$

Thus, G is a best response to H following signal 2 as well.

Define G' as G, but with $G_1'(m_{1,1}) - G_1'(m_{1,1})^- = 1$; that is, according to G', candidate A plays as in G but chooses the conditional median $m_{1,1}$ with probability one following signal 1. Letting X be the pure strategy with $x_1 = m_{1,2}$, we have

$$\begin{aligned} P(1)\Pi_A(X, H|1) &= P(1,1)\big(\alpha(0) + (1-\alpha)(0.5)\big) + P(1,2)(1) + P(1,3)(0.5) \\ &= \epsilon(0.1\alpha + 0.53) = 0.61\epsilon \end{aligned}$$

and

$$\begin{aligned} P(1)\Pi_A(G', H|1) &= P(1,1)\big(\alpha(0.5) + (1-\alpha)(1)\big) + P(1,2)(0.5) + P(1,3)(0) \\ &= \epsilon(0.6\alpha + 0.225) = 0.705\epsilon. \end{aligned}$$

Since G_1 puts positive probability on $x_1 = m_{1,2}$, we conclude that (G, H) is not a Bayesian equilibrium. Thus, diagonal better reply security requires $\hat{G}$ and $\mathcal{H}$, as described above. Note that $H \in \mathcal{H}$, and that G_2 and G_3 are best responses to H conditional on signals 2 and 3, respectively.

8 Here, we give candidate B's strategy space the product topology, where each factor, the set of distributions over the real line, is given the weak* topology.

Furthermore, we claim that the only position that increases A's expected payoff conditional on signal 1 is, in fact, $x_1 = m_{1,1}$. If A took a position $x_1' \in (m_{1,1}, m_{1,2})$ following signal 1, then A's expected payoff would be

$$\begin{aligned} P(1)\Pi_A(X', H|1) &= P(1,1)\big(\alpha(0) + (1-\alpha)(1)\big) + P(1,2)(1) + P(1,3)(0) \\ &= \epsilon(1.25 - 0.8\alpha) = 0.61\epsilon. \end{aligned}$$

If A took a position $x_1' \in (m_{1,2}, 2)$ following signal 1, then A's expected payoff would again be

$$\begin{aligned} P(1)\Pi_A(X', H|1) &= P(1,1)\big(\alpha(0) + (1-\alpha)(0)\big) + P(1,2)(1) + P(1,3)(1) \\ &= 0.61\epsilon. \end{aligned}$$

And positioning further to the right would yield an even lower expected payoff. Therefore, $\hat{G}$ must involve the transfer of probability mass from $m_{1,2}$ to $m_{1,1}$ following signal 1.

The difficulty for diagonal better reply security is that such a change no longer delivers an *ex ante* expected payoff for A above one half if we perturb H slightly to $\hat{H}$ by specifying that $\hat{H}_2$ put probability one on a point to the left of, and close to, 2: In that case, A loses to B when A receives signal 1 and locates at $m_{1,1}$ and B receives signal 2. To see the claim, consider $\hat{G} = G'$. Then A's ex ante expected payoff in excess of one half, $\Pi_A(\hat{G}, \hat{H}) - 1/2$, is

$$\begin{aligned} &P(1)\Pi_A(G', \hat{H}|1) + P(2)\Pi_A(G', \hat{H}|2) + P(3)\Pi_A(G', \hat{H}|3) - 0.5 \\ &\quad = \big[P(1,1)\big(\alpha(0.5) + (1-\alpha)(1)\big) + P(1,2)(0) + P(1,3)(0)\big] \\ &\qquad + \big[P(1,2)(\alpha(0.5) + \big(1-\alpha\big)(0)) + P(2,2)(0) + P(2,3)(0.5)\big] \\ &\qquad + \big[P(1,3)\big(\alpha(1) + (1-\alpha)(0.5)\big) + P(2,3)(1) + P(3,3)(0.5)\big] - 0.5 \\ &\quad = \alpha\big[(0.5)(0.8\epsilon) - 0.8\epsilon + (0.5)(0.45\epsilon) + 0.16\epsilon - (0.5)(0.16\epsilon)\big] \\ &\qquad + 0.8\epsilon + (0.5)(0.41\epsilon) + (0.5)(0.16\epsilon) + (0.41\epsilon) + (0.5)(1 - 2.89\epsilon) - 0.5 \\ &\quad = \epsilon\big[\alpha(-0.095) + 0.05\big] \\ &\quad = -0.026\epsilon, \end{aligned}$$

which is negative. Thus, diagonal better reply security is not fulfilled by $\hat{G} = G'$.

Reny's diagonal better reply security condition does not require that $\hat{G} = G'$, as in the above calculation: Candidate A could, for example, move probability mass from 2 or $m_{3,3}$ following signals 2 and 3, respectively; as already confirmed, this would not increase A's ex ante expected payoff when B uses H, but it could conceivably mitigate the problem illustrated in the preceding, "protecting" A from B's slight move to the left following signal 2. A closer look shows, however, that no such protection is available. Following signal 3, of course, any change in G_3' will lead to a discontinuous decrease in A's expected payoff, weighted by $1 - 2.89\epsilon$, which can be made arbitrarily close to one. Following signal 2, A might move probability mass from 2 to the left to defeat candidate B in case B also receives signal 2, but such a change means that A would lose to B in case B received signal 3, and we have seen that the two effects cancel. Finally, after signal 1, A might move probability mass from $m_{1,1}$ to the right in order to defeat B when B receives signal 2, but we have seen that such a move does not increase A's interim expected payoff conditional on signal 1. This completes the example. □

Proof of Theorem 4. We prove lower semi-continuity of v_A at γ. A symmetric argument proves lower semi-continuity of $v_B = 1 - v_A$, which, in turn, gives us upper semi-continuity of v_A. Let

$\gamma_n \to \gamma$, and suppose $v_A(\gamma) > \liminf v_A(\gamma_n)$. Let Π_A^n denote A's ex ante expected payoff function corresponding to γ_n, and let Π_A denote the ex ante payoffs corresponding to γ. Let M^n denote the interval $M(\gamma_n)$, let $M = M(\gamma)$, and let $\hat{M}$ be any compact set containing M in its interior. By continuity, therefore, $M^n \subseteq \hat{M}$ for high enough n. For each n, let (G^n, H^n) be an equilibrium with support in M^n for the electoral game indexed by γ_n, so $\Pi_A^n(G^n, H^n) = v_A(\gamma_n)$ and $\Pi_B^n(G^n, H^n) = v_B(\gamma_n)$. By compactness of $\hat{M}$, there exists a weakly convergent subsequence of $\{(G^n, H^n)\}$, also indexed by n, with limit (G, H). Going to a further subsequence if necessary, we may assume $\{v_A(\gamma_n)\}$ converges to limit $v < v_A(\gamma)$. Let (G^*, H^*) be an equilibrium of the electoral game indexed by γ, so G^* is an optimal strategy for A, which guarantees a payoff of at least $v_A(\gamma)$ in game γ. Thus, $\Pi_A(G^*, H) \geqslant v_A(\gamma)$. In particular, there exists a pure strategy X^* such that $\Pi_A(X^*, H) \geqslant v_A(\gamma) > v$. We claim that as a consequence, there exists a pure strategy X' such that

$$\Pi_A^n\big(X', H^n\big) > \frac{\Pi_A(X^*, H) + v}{2},$$

for high enough n. But this, with $v_A(\gamma_n) \to v$, contradicts the assumption that G^n is a best response to H^n for candidate A. We establish the claim in three steps.

Step 1. By Lemma A.1, for every $s \in S$, either

$$\sum_{t\in T} P(t|s)\big[H_t\big(x_s^*\big) - H_t\big(x_s^*\big)^-\big]\big[F_{s,t}\big(x_s^*\big)\big] \geqslant \frac{1}{2}\sum_{t\in T} P(t|s)\big[H_t\big(x_s^*\big) - H_t\big(x_s^*\big)^-\big] \tag{A.5}$$

or

$$\sum_{t\in T} P(t|s)\big[H_t\big(x_s^*\big) - H_t\big(x_s^*\big)^-\big]\big[1 - F_{s,t}\big(x_s^*\big)\big] \geqslant \frac{1}{2}\sum_{t\in T} P(t|s)\big[H_t\big(x_s^*\big) - H_t\big(x_s^*\big)^-\big]. \tag{A.6}$$

Let S^- be the set of $s \in S$ such that (A.5) holds, and let S^+ be the set of $s \in S \setminus S^-$ such that (A.6) holds. For $s \in S^-$, let $\{x_s^k\}$ be a sequence increasing to x_s^*, and for $s \in S^+$, let $\{x_s^k\}$ be a sequence decreasing to x_s^*. In addition, we choose each x_s^k to be a continuity point of H_t for all $t \in T$; this is possible because T is finite and each H_t has a countable number of discontinuity points. Thus, $H_t(x_s^k) - H_t(x_s^k)^- = 0$ for all $t \in T$. For each k, define the strategy $X^k = (x_s^k)$ for candidate A.

Step 2. We now argue that X^k satisfies $\liminf \Pi_A(X^k, H) \geqslant \Pi_A(X^*, H)$. For each $t \in T$, let λ_t denote the probability measure generated by the distribution H_t, let μ_t denote the degenerate measure with mass $H_t(x_s^*) - H_t(x_s^*)^-$ on each x_s^*, and let $\nu_t = \lambda_t - \mu_t$. Let $\pi_{s,t}^*(z) = \pi_A(x_s^*, z|s, t)$ denote A's probability of winning using x_s^* conditional on signal s when B receives signal t and chooses platform z, and let $\pi_{s,t}^k(z) = \pi_A(x_s^k, z|s, t)$ denote A's analogous probability of winning using x_s^k. Note that

$$\begin{aligned}
\Pi_A\big(X^k, H\big) - \Pi_A(X^*, H) &= \sum_{s\in S} P(s) \sum_{t\in T} P(t|s) \int \big[\pi_{s,t}^k(z) - \pi_{s,t}^*(z)\big]\lambda_t(\mathrm{d}z) \\
&= \sum_{s\in S^-}\sum_{t\in T} P(s,t)\big[H_t\big(x_s^*\big) - H_t\big(x_s^*\big)^-\big]\left[F_{s,t}\left(\frac{x_s^* + x_s^k}{2}\right) - \frac{1}{2}\right] \\
&\quad + \sum_{s\in S^+}\sum_{t\in T} P(s,t)\big[H_t\big(x_s^*\big) - H_t\big(x_s^*\big)^-\big]\left[1 - F_{s,t}\left(\frac{x_s^* + x_s^k}{2}\right) - \frac{1}{2}\right] \\
&\quad + \sum_{s\in S}\sum_{t\in T} P(s,t) \int \big[\pi_{s,t}^k(z) - \pi_{s,t}^*(z)\big]\nu_t(\mathrm{d}z).
\end{aligned}$$

Since $\pi^k_{s,t} - \pi^*_{s,t} \to 0$ almost everywhere (ν_t), the corresponding integral terms above converge to zero. Thus, by construction of X^k, $\liminf_{m\to\infty} \Pi_A(X^k, H) \geqslant \Pi_A(X^*, H) > v$, as desired.

Step 3. Choose k such that $\Pi_A(X^k, H) > (\Pi_A(X^*, H) + v)/2$ and set $X' = X^k$. To prove the claim that $\Pi^n_A(X', H^n) > v$ for high enough n, define the functions

$$\phi'_{s,t}(z) = \pi_A\left(x'_s, z|s, t\right) \quad \text{and} \quad \phi^n_{s,t}(z) = \pi^{\gamma_n}_A\left(x'_s, z|s, t\right).$$

Note that

$$\Pi_A(X', H) = \sum_{s\in S}\sum_{t\in T} P(s, t) \int \phi'_{s,t}(z) H_t(\mathrm{d}z),$$

and, letting $P^n = P^{\gamma_n}$,

$$\Pi^n_A\left(X', H^n\right) = \sum_{s\in S}\sum_{t\in T} P^n(s, t) \int \phi^n_{s,t}(z) H^n_t(\mathrm{d}z).$$

Since $\Pi_A(X', H) > v$, it suffices to show that

$$\int \phi^n_{s,t}(z)\, \mathrm{d}H^n_t \to \int \phi'_{s,t}(z)\, \mathrm{d}H_t,$$

for each $s \in S$ and $t \in T$. To prove this, fix $\epsilon > 0$. Because $x'_s = x^k_s$ is not a mass point of H_t, we may specify an interval $Z = [\underline{z}, \bar{z}]$ with $x'_s \in (\underline{z}, \bar{z})$ such that $H_t(\bar{z}) - H_t(\underline{z})^- < \epsilon/4$. By weak convergence, $H^n_t(\bar{z}) - H^n_t(\underline{z})^- < \epsilon/2$ for sufficiently high n. Furthermore, $\{\phi^n_{s,t}\}$ is a sequence of functions that are non-decreasing on $[\underline{m}, \underline{z}]$ and converge pointwise to $\phi'_{s,t}$ on this interval, so they converge uniformly to $\phi'_{s,t}$ on the interval. Similarly, each $\phi^n_{s,t}$ is non-increasing on $[\bar{z}, \overline{m}]$, so the functions converge uniformly to $\phi'_{s,t}$ on this interval. Choosing n high enough that $|\phi^n_{s,t}(z) - \phi'_{s,t}(z)| < \epsilon/2$ for all $z \in [\underline{m}, \underline{z}] \cup [\bar{z}, \overline{m}]$, we have

$$\left| \int \phi^n_{s,t}(z)\, \mathrm{d}H^n_t - \int \phi'_{s,t}(z)\, \mathrm{d}H_t \right| < \epsilon,$$

as required. □

Proof of Theorem 5. Since Γ and $\mathcal{B}^{S\cup T}$ are metrizable, we may restrict attention to sequences, rather than nets. Let $\gamma_n \to \gamma$, let $(G^n, H^n) \in E(\gamma_n)$ for each n, and suppose that $(G^n, H^n) \to (G, H)$. If $(G, H) \notin E(\gamma)$, then, using the notation from the proof of Theorem 4, one candidate, say A, has a pure strategy X such that $\Pi_A(X, H) > v_A(\gamma)$. But then, as in the proof of Theorem 4, we can find a strategy X' satisfying the latter inequality such that no x'_s is a mass point of any H_t, and then we can show that

$$\Pi^n_A\left(X', H^n\right) > \frac{\Pi_A(X, H) + v_A(\gamma)}{2},$$

for high enough n. But $v_A(\gamma_n) \to v_A(\gamma)$ by Theorem 4, so it follows that $\Pi^n_A(X', H^n) > v_A(\gamma_n)$ for high enough n, contradicting the assumption that G^n is a best response to H^n for A in the electoral game indexed by γ_n. □

Proof of Theorem 6. Let (G, H) be a mixed strategy Bayesian equilibrium, let $\underline{x}_i = \sup\{x \in \Re : G_i(x) = 0\}$ be the lower bound of the support of G_i for each $i \in I$, and let $\underline{x} = \min_{i\in I} \underline{x}_i$ be the minimum of these lower bounds. Suppose that $\underline{x} < \underline{m}$, and take i such that $\underline{x}_i = \underline{x}$. By symmetry and interchangeability, (G, G) is also an equilibrium, so we may assume that $H = G$. Consider a sequence of pure strategies $\{X^n\}$ satisfying the following. If G_i puts positive probability on $\underline{x}$, i.e., $G_i(\underline{x}) - G_i(\underline{x})^- > 0$, then let $x_i^n = \underline{x}$ for all n. Otherwise, let $\{x_i^n\}$ be a sequence decreasing to $\underline{x}$ such that each x_i^n is in the support of G_i. Furthermore, choose x_i^n so that $\Pi_A(X^n, H|i) = \Pi_A(G, H|i)$ for all n. To see that this can be done, set $x_i^0 > \underline{x}$ arbitrarily and, if possible, let x_i^n be any continuity point of A's expected payoff function in the support of G_i and in the interval $[\underline{x}, (\underline{x} + x_i^{n-1})/2]$ to satisfy the desired condition. Since there is at most a countable number of discontinuity points of A's payoff function, such a policy can be found unless the support of G_i in $[\underline{x}, (\underline{x} + x_i^{n-1})/2]$ is countable. In that case, however, any policy in the support of G_i in this interval satisfies the desired condition, and there is at least one such policy since $\underline{x}_i = \underline{x}$. In any case, we have $\Pi_A(X^n, H|i) = \Pi_A(G, H|i)$ for all n and $\lim_{n\to\infty} G_i(x_i^n)^- = 0$. Now consider a pure strategy X' satisfying $x_i' = \underline{m}$, and note that for n such that $x_i^n < \underline{m}$,

$$\begin{aligned}
&\Pi_A(X', H|i) - \Pi_A\big(X^n, H|i\big) \\
&\quad= \sum_{j\in I} P(j|i)\Bigg[\int\limits_{[\underline{x}, x_i^n)} \left[F_{i,j}\left(\frac{x_i^n + z}{2}\right) - F_{i,j}\left(\frac{\underline{m} + z}{2}\right)\right] H_j(\mathrm{d}z) \\
&\qquad + \big(H_j\big(x_i^n\big) - H_j\big(x_i^n\big)^-\big)\left[\frac{1}{2} - F_{i,j}\left(\frac{x_i^n + \underline{m}}{2}\right)\right] &\text{(A.7)}\\
&\qquad + \int\limits_{(x_i^n, \underline{m})} \left[1 - F_{i,j}\left(\frac{\underline{m} + z}{2}\right) - F_{i,j}\left(\frac{x_i^n + z}{2}\right)\right] H_j(\mathrm{d}z) &\text{(A.8)}\\
&\qquad + \big(H_j(\underline{m}) - H_j(\underline{m})^-\big)\left[\frac{1}{2} - F_{i,j}\left(\frac{x_i^n + \underline{m}}{2}\right)\right] &\text{(A.9)}\\
&\qquad + \int\limits_{(\underline{m}, \infty)} \left[F_{i,j}\left(\frac{\underline{m} + z}{2}\right) - F_{i,j}\left(\frac{x_i^n + z}{2}\right)\right] H_j(\mathrm{d}z)\Bigg].
\end{aligned}$$

For each $j \in I$, the first integral goes to zero, since $\lim_{n\to\infty} H_j(x_i^n)^- = \lim_{n\to\infty} G_j(x_i^n)^- = 0$. Further, the last integral is clearly non-negative. So, too, the other terms are non-negative, because $F_{i,j}(\frac{w+z}{2}) < 1/2$ for all $w \leqslant \underline{m}$ and all $z < \underline{m}$. This establishes that the right-hand side is non-negative. It is strictly positive because, when $j = i$, the bracketed terms in (A.7)–(A.9) have strictly positive limits; and the total probability mass on these terms is strictly positive since $\lim_{n\to\infty} H_j(\underline{m}) - H_j(x_i^n)^- = G_i(\underline{m}) - G_i(\underline{x})^- = G_i(\underline{m}) > 0$. Because $P(i|i) > 0$ by (C3), we conclude that $\Pi_A(X', H|i) > \Pi_A(X^n, H|i) = \Pi_A(G, H|i)$ for high enough n, contradicting the premise that (G, H) is an equilibrium. An analogous argument establishes that the supports of equilibrium strategies are bounded from above by $\overline{m}$. □

Proof of Lemma 2. Let (G, H) be a symmetric mixed strategy Bayesian equilibrium, and take any $z \in M$. Define the set $I' = \{i : G_i(z) - G_i(z)^- > 0\}$. Lemma A.2 establishes that for all signals $i \in I'$, z must solve

$$\sum_{j\in I} \alpha_j F_{i,j}(z) = \frac{1}{2}, \tag{A.10}$$

where the non-negative weights,

$$\alpha_j = \frac{[G_j(z) - \lim_{w\uparrow z} G_j(z)]P(j|i)}{\sum_{l\in I'}[G_l(z) - \lim_{w\uparrow z} G_l(z)]P(l|i)},$$

sum to one. Under (C4*), signals are ordered. If I' contains multiple signals, then let $\underline{i}$ and $\bar{i}$ denote the minimum and maximum of I', respectively. Accordingly,

$$\sum_{j\in I}\alpha_j F_{\bar{i},j}(z)=\sum_{j\in I}\alpha_j F_{\underline{i},j}(z). \tag{A.11}$$

Suppose in order to show a contradiction that $F_{\underline{i},\bar{i}}(z)=0$. Then, by (C4*), we have $F_{j,\bar{i}}(z)=0$ for all $j>\underline{i}$, which implies $\sum_j \alpha_j F_{\bar{i},j}(z)=0$. This contradicts (A.10) with $i=\bar{i}$, as desired. An analogous argument establishes that $0<F_{\underline{i},\bar{i}}(z)<1$. Since $z\in M$, condition (C4*) then implies that $F_{\underline{i},\bar{i}}(z)$ is both the uniquely smallest element of $\{F_{\underline{i},\ell}(z)\mid \ell\in I'\}$ and the uniquely largest element of $\{F_{\bar{i},\ell}(z)\mid \ell\in I'\}$. Condition (A.11) therefore implies that $\alpha_{\bar{i}}=\alpha_{\underline{i}}=1$, a contradiction. We conclude that I' is a singleton. Taking i such that $I'=\{i\}$, condition (A.10) therefore reduces to $F_{i,i}(z)=1/2$, i.e., $z=m_{i,i}$. □

Example A.2 *(Upper Hemicontinuity Violated with Discontinuous Conditional Distributions).* Let $I=\{-1,0,1\}$, with uniform priors, i.e., $P(i,j)=\frac{1}{9}$ for each $i,j\in I$. For each $i,j\in I$, let $F_{i,j}$ be the point mass on $m_{i,j}$, and let $\{F^n_{i,j}\}$ be a sequence of uniform distributions with density n, where all conditional medians are depicted below.

	$i=-1$	$i=0$	$i=1$
$j=1$	$m_{-1,1}=0$ $m^n_{-1,1}=\frac{1}{n}$	$m_{0,1}=0$ $m^n_{0,1}=0$	$m_{1,1}=0$ $m^n_{1,1}=\frac{1}{n}$
$j=0$	$m_{-1,0}=0$ $m^n_{-1,0}=0$	$m_{0,0}=0$ $m^n_{0,0}=0$	$m_{1,0}=0$ $m^n_{1,0}=0$
$j=-1$	$m_{-1,-1}=-1$ $m^n_{-1,-1}=-1$	$m_{0,-1}=0$ $m^n_{0,-1}=0$	$m_{1,-1}=0$ $m^n_{1,-1}=\frac{1}{n}$

Thus, each n defines an instance of the Canonical Model, and the conditional distributions $F^n_{1,-1}$, $F^n_{-1,1}$, and $F^n_{1,1}$ converge weakly to the degenerate distributions $F_{1,-1}$, $F_{-1,1}$, and $F_{1,1}$, respectively. Furthermore, the supports of $F^n_{-1,0}$ and $F^n_{-1,1}$ are contiguous, as are the supports of $F^n_{0,-1}$ and $F^n_{1,-1}$. For each n, define the pure strategy X^n as follows: $x^n_{-1}=-1$, $x^n_0=0$, and $x^n_1=1/n$, and let $Y^n=X^n$. Then (X^n,Y^n) is a Bayesian equilibrium of the nth game in the sequence. To see this, note that following signal -1, $x^n_{-1}=-1$ produces a tie if B receives signal -1, a loss if B receives signal 0, and a loss if B receives signal 1, yielding a conditional expected payoff for candidate A of 1/6. Moving from $x^n_{-1}=-1$ to the right, A's conditional expected payoff is maximized for $x'_{-1}\in[0,1/n]$, which yields

$$\left(\frac{1}{3}\right)(0)+\left(\frac{1}{3}\right)F_{-1,0}\left(\frac{x'_{-1}-x^n_0}{2}\right)+\left(\frac{1}{3}\right)\left(1-F_{-1,1}\left(\frac{x^n_1-x'_{-1}}{2}\right)\right)=\frac{1}{6},$$

and thus $x^n_{-1}=-1$ is a best response. Following signal 0, $x^n_0=0$ produces a win if B receives signals -1 or 1 and a tie if B receives signal 0, and this is clearly a best response. Following signal 1, $x_1^n = 1/n$ produces a win if B receives signal -1, a loss if B receives signal 0, and a tie if B receives signal 1, and this is also a best response, establishing the claim. Clearly, (X^n, Y^n) converges to (X, Y) defined by $x_{-1} = y_{-1} = -1$, $x_0 = y_0 = 0$, and $x_1 = y_1 = 0$. But this is not an equilibrium of the limiting model, because $x_{-1} = -1$ is not a best response: while $x_{-1} = -1$ produces a tie and two losses, moving to $x'_{-1} = 0$ produces two ties and one loss, increasing A's expected payoff.

References

Aragones, E., Palfrey, T., 2002. Mixed equilibrium in a Downsian model with a favored candidate. J. Econ. Theory 103, 131–161.

Banks, J., Duggan, J., 2005. Probabilistic voting in the spatial model of elections: The theory of office-motivated candidates. In: Social Choice and Strategic Decisions: Essays in Honor of Jeffrey S. Banks. Springer, New York.

Besley, T., Coate, S., 1997. An economic model of representative democracy. Quart. J. Econ. 112, 85–114.

Bernhardt, D., Duggan, J., Squintani, F., 2005. Private polling in elections and voter welfare. Mimeo.

Black, D., 1958. Theory of Committees and Elections. Cambridge Univ. Press, Cambridge.

Budge, I., Klingemann, H., Volkens, A., Bara, J., Tanenbaum, E., 2001. Mapping Policy Preferences: Estimates for Parties, Electors, and Governments, 1945–1998. Oxford Univ. Press, Oxford.

Calvert, R., 1985. Robustness of the multidimensional voting model: Candidate motivations, uncertainty, and convergence. Amer. J. Polit. Sci. 29, 69–95.

Chan, J., 2001. Electoral competition with private information. Mimeo.

Dasgupta, P., Maskin, E., 1986. The existence of equilibrium in discontinuous economic games, I: Theory. Rev. Econ. Stud. 53, 1–26.

Downs, A., 1957. An Economic Theory of Democracy. Harper and Row, New York.

Duggan, J., Fey, M., 2005. Electoral competition with policy-motivated candidates. Games Econ. Behav. 51, 490–522.

Eisinger, R.M., 2003. The Evolution of Presidential Polling. Cambridge Univ. Press, Cambridge.

Groseclose, T., 2001. A model of candidate location when one candidate has a valence advantage. Amer. J. Polit. Sci. 45, 862–886.

Heidhues, P., Lagerlöf, J., 2003. Hiding information in electoral competition. Games Econ. Behav. 42, 48–74.

Hotelling, H., 1929. Stability in competition. Econ. J. 39, 41–57.

Jackson, M., Simon, L., Swinkels, J., Zame, W., 2002. Communication and equilibrium in discontinuous games of incomplete information. Econometrica 70, 1711–1740.

Ledyard, J., 1989. Information aggregation in two-candidate elections. In: Ordeshook, P. (Ed.), Models of Strategic Choice and Politics. Univ. Michigan Press, Ann Arbor.

Martinelli, C., 2001. Elections with privately informed parties and voters. Public Choice 108, 147–167.

Martinelli, C., Matsui, A., 2002. Policy reversals and electoral competition with privately informed parties. J. Public Econ. Theory 4, 39–62.

Medvic, S., 2001. Political Consultants in US Congressional Elections. Ohio State Univ. Press, Columbus.

Merrill, S., Grofman, B., 1999. A Unified Theory of Voting. Cambridge Univ. Press, Cambridge.

Milgrom, P., 1981. Rational expectations, information acquisition, and competitive bidding. Econometrica 49, 921–943.

Osborne, M., Slivinski, A., 1996. A model of political competition with citizen-candidates. Quart. J. Econ. 111, 65–96.

Ottaviani, M., Sorensen, P., 2003. Forecasting and rank order contests. Mimeo.

Page, B., 1978. Choices and Echoes in Presidential Elections. Univ. Chicago Press, Chicago.

Poole, K., Rosenthal, H., 1997. Congress: A Political Economic History of Roll Call Voting. Oxford Univ. Press, New York.

Reny, P., 1999. On the existence of pure and mixed strategy Nash equilibria in discontinuous games. Econometrica 67, 1029–1056.

Simon, L., Zame, W., 1990. Discontinuous games and endogenous sharing rules. Econometrica 58, 861–872.

Wittman, D., 1983. Candidate motivation: A synthesis of alternatives. Amer. Polit. Sci. Rev. 77, 142–157.