Optimal Timing of Policy Announcements in Dynamic Election Campaigns

Yuichiro Kamada and Takuo Sugaya*

July 6, 2018

Abstract

We construct a dynamic model of election campaigns. In the model, opportunities for candidates to refine/clarify their policy positions are limited and arrive stochastically along the course of the campaign until the predetermined election date. We show that this simple friction leads to rich and subtle campaign dynamics. We first demonstrate these effects in a series of canonical static models of elections that we extend to dynamic settings, including models with valence, a multi-dimensional policy space, policy motivated candidates, campaign spending, and incomplete information. We then present general principles that underlie the results from those examples. In particular, we establish that candidates spend a long time using ambiguous language during the election campaign in equilibrium.

*Kamada: Haas School of Business, University of California, Berkeley, Berkeley, CA 94720, e-mail: y.cam.24@gmail.com. Sugaya: Stanford Graduate School of Business, Stanford, CA, 94305, e-mail: tsugaya@stanford.edu; We thank Ernesto Dal Bo, Itay Fainmesser, Drew Fudenberg, Michihiro Kandori, Akitada Kasahara, Kei Kawai, Fuhito Kojima, David Laibson, John B. Londregan, Shih-En Lu, Francisco Martínez-Mora, Adam H. Meirowitz, Stephen Morris, Kristopher Ramsay, Andrzej Skrzypacz, Ken Shepsle, Yuki Takagi, Satoru Takahashi, and the seminar participants at the Harvard University Department of Economics and Department of Government, the Princeton University Department of Economics and Department of Political Science, University of California, Berkeley, Yale University, the 21st Stony Brook International Conference on Game Theory, and the Urrutia Elejalde Workshop on Information, Dynamics and Political Decision Making for helpful comments. In addition, Steven Callander and Josh Gottlieb read through the previous versions of this paper and gave us very detailed comments, which significantly improved the paper. Judith Levi provided a superb professional editorial assistance, which again significantly improved the paper. We thank Emily Her, Douglas Hong, Omair Butt, and especially Lingxuan Wu for their excellent research assistantship. This paper subsumes and supersedes part of our unpublished working paper "Valence Candidates and Ambiguous Platforms in Policy Announcement Games" and its earlier version circulated as "Policy Announcement Game: Valence Candidates and Ambiguous Policies," which Sugaya first presented at the 21st Stony Brook International Conference on Game Theory in 2010.

"Timing is Everything"

—Joe Slade White, the 2014 National Democratic Strategist of the Year

1 Introduction

Election campaigns are inherently dynamic. Joe Slade White, the 2014 National Democratic Strategist of the Year, states "timing makes the difference between winning and losing," in one of his "9 Principles of Winning Campaigns" (White, 2012). Despite the apparent importance of election campaigns on the electoral outcomes and the fact that the campaigns are dynamic in nature, there seem to be no theoretical models of dynamic campaigns in the literature, to the best of our knowledge.[1] One possible reason is that there is no obvious way to model campaigns in a way that would give rise to nontrivial dynamic strategic considerations.[2] The objective of this paper is to fill this gap by proposing a model in which candidates face nontrivial dynamic strategic considerations.

The paper proposes a "policy announcement timing game" in which candidates strategically choose the optimal timing of their policy announcements over a campaign period. Each announcement corresponds to restricting the set of available policies—that is, each candidate may clarify a policy to implement, while she cannot go back to a policy that she has ruled out before—and the final policy announcements before the predetermined election day determine the result of the election.

In our model, opportunities for policy announcements are limited and stochastic. Specifically, we assume that opportunities arrive according to a Poisson process over a campaign period. The assumption of Poisson opportunities is a simple way to represent frictions present in the communication process between candidates and voters. For example, administrative procedures to obtain internal approval for a change of how candidates announce their policies may not always be successful, or candidates may not always be able to communicate with the voters about such a change even if these procedures go through. Moreover, voters may not be convinced that the candidate has changed her policy position.[3] Those frictions cause uncertainty regarding the availability of

[1] By a model of dynamic election campaigns, we mean a model with a single election; in particular, when we speak of "models of dynamic election campaigns," we are excluding models that have primaries and the general election.

[2] The empirical research shows that candidates do react to each other during election campaign (cf. Banda (2013, 2015)).

[3] A richer modeling of administrative procedure or dynamics of voter beliefs would generate a more accurate prediction, but we assume these away and try to concentrate on key effects by investigating what we can say in our

future opportunities, and we capture such uncertainty by using Poisson processes.[4]

After laying out the general model, we present a number of applications to demonstrate that introducing this simple friction to the model generates interesting dynamic strategic considerations and equilibrium dynamics consistent with election dynamics in reality. The first issue we consider is ambiguous policy announcements, which we often see in real election campaigns. For example, in the context of a US presidential election, Nicholas Biddle, the manager of William Henry Harrison's campaign for the US presidency in 1840-1841, advised Harrison in these words: "Let him say not one single word about his principles, or his creed - let him say nothing - promise nothing. Let no Committee, no convention—no town meeting ever extract from him a single word, about what he thinks now, or what he will do hereafter."[5]

We find that our model leads to a new interpretation of ambiguous policy announcements: In our applications, if candidates are purely office-motivated and there is no Condorcet winner in the set of available policies, then the candidates have tentative preferences for the ambiguous policy statement during the course of the election campaign and, in fact, spend most of the campaign time keeping their policy statements ambiguous—not announcing a specific policy.[6] The incentive comes from a dynamic consideration. The candidates' announcements remain ambiguous because the absence of a Condorcet winner implies that it is less favorable to be the first mover than to be the second mover.

There are two leading examples with such "first mover disadvantage": Candidates have valence, or the policy space is multi-dimensional. When there is valence (Section 3.1), we show that, in equilibrium, the weak candidate will not make his policy clear in the early stages of the election campaign, and his policy announcement can possibly occur only close to the election date. This is because if he clarifies his policy too early, then the strong candidate will have enough time to simply copy that policy afterward, so that the weak candidate will certainly lose. The result may

---

simplest framework. As it turns out, the results from our simple setting are quite rich.

[4]Calvo (1983) uses a Poisson process to model uncertainty about future opportunities of changing prices. This approach offers a tractable way of modeling sticky prices and analyzing the effect of fiscal and monetary policies. At the same time, the literature goes forward to offer a micro-foundation of Calvo (1983) by fixed costs of changing prices, rational inattention, and so on (see Klenow and Malin (2010)). Here, we also show that this Poisson approach is useful to analyze campaign dynamics, and leave micro-foundation to the future research.

[5]McGrane, Reginald Charles C. (1919). The quote appears in Shepsle (1972).

[6]In contrast, we show that if candidates are purely office-motivated and there is a Condorcet winner, each candidate announces it as soon as possible. The formal definition of Condorcet winner for this result will be provided in Section 4.2.

explain the dynamics of the election campaign in the 2014 gubernatorial election for Tokyo, Japan, in which Yoichi Masuzoe and Morihiro Hosokawa fought a close campaign. Although Masuzoe had been seen as the strongest candidate from the outset of the campaign, Hosokawa became popular near the election day when he clarified his stance by announcing opposition to the restart of nuclear power generation. Then Masuzoe, who originally had not specified his policy about nuclear power generation, clarified his position to aim for a gradual phase-out of nuclear power. As a result, Masuzoe won against Hosokawa.[7]

When the policy space is multi-dimensional (Section 3.2), generically there does not exist a Condorcet winner (so there does not exist a pure-strategy Nash equilibrium in the static environment). In our policy announcement timing game, however, we can pin down both the equilibrium probability distribution of times at which candidates make policy announcements and winning probabilities even in such a setting. In the absence of a Condorcet winner, once a candidate commits to any policy platform and then the opponent optimally responds to it, the former candidate will lose. Hence, each candidate, upon making an announcement, "becomes a weak candidate" in that being best-responded afterward will bring the worst outcome. In contrast, the payoff structure is such that, if candidates knew that the current opportunity is the last one and there will be no opportunity for either candidate in the future, they would prefer to clarify their policy. Hence, in equilibrium, when a candidate obtains an opportunity near the election day, they clarify their policies.

This analysis on the multi-dimensional policy space, however, does not give us a precise prediction regarding the policies that candidates announce due to its excessive simplicity of pure office motivation. To show that such indeterminacy is not a consequence of the way our general dynamic model is specified, Section 3.3 introduces policy motivation to the model with a multi-dimensional policy space. Again, we show that ambiguous language is used for a long time in equilibrium, and pin down the policies that candidates announce. Interestingly, in equilibrium, a candidate may announce a policy that is Pareto inefficient among both candidates with positive probability. The reason is that announcing such a policy will make it incentive compatible for the other candidate to announce a policy that is not too unfavorable for the candidate in the event that the other can-

[7]Sankei News (2013) argued on December 24, 2013 that Masuzoe was seen as the strongest among the candidates, Asahi Shimbun Digital (2014a) reported on January 9, 2014 that Hosokawa clarified his policy about nuclear power, and Asahi Shimbun (2014b) reported on January 15, 2014 that Masuzoe showed support to the opposition to nuclear power.

didate obtains a chance of a policy announcement afterward. Announcing a policy position with a motive to influence the opposition's policy, though possibly sounding unrealistic, did actually happen in real campaigns. During the Democratic Party presidential primaries in 2016, for example, the far-left Bernie Sanders called for a $15-an-hour minimum wage (more than twice as much as the $7.25 standard back then) and Medicare-for-all health care, and proposed to end TPP. After losing Pennsylvania, Maryland, Delaware, and Connecticut in a row, Sanders declared in a town hall meeting: "But if we do not win, we intend to win every delegate that we can so that when we go to Philadelphia in July, we are going to have the votes to put together the strongest progressive agenda that any political party has ever seen."[8] An article in Vox (Stein, 2016) writes: "Bernie Sanders moved Democrats to the left. The platform is proof. [...] Hillary Clinton may have won the Democratic Party's presidential nomination, but Bernie Sanders has still left an outsize mark on its future."[9]

Another topic which attracts much attention with regard to campaign dynamics is political campaign advertisements (Section 3.4). We provide a simple model in which we reinterpret our policy announcement timing game to encompass dynamic spending in election campaigns. To make the reinterpretation work, we notice that the cumulative spending for advertisements cannot decrease over time; hence any spending is only restricting the set of possible cumulative spending. Supposing that the probability of winning the election only depends on the ratio of the cumulative spending of the two candidates and money can be used for purposes other than the campaign as well and is sufficiently important, we show that, in equilibrium, candidates do not spend as much money as they can in the early stages of the campaign, and make additional spending close to the election date if they can. The intuition is as follows. If two candidates spend as much as possible, both candidates will have a 50% chance of winning (given symmetry). There is little incentive to spend a lot at the early stages because that means that the opponent can cancel out its effect by later spending equally much with a high probability. Rather, candidates would save their money in the early stages and try to spend them later, leaving only a small probability for the opponent to cancel it out later. Our prediction provides a novel explanation for the empirical evidence, which

[8] The facts and the quote appear in Strauss (2016) and Gurciullo and Debenedetti (2016).

[9] The specific interpretation we give to the policy space we study may not be consistent with this episode of Sanders vs. Clinton. We provide this example to make a point that an entry to a policy platform can happen with a motive to influence the opposition's platform.

suggests that candidates often spend monetary resources gradually over time.[10]

Although the general model is set up as a complete information game, one can extend it to include a wider class of settings. To illustrate, our final application allows for incomplete information by considering a model in which candidates' types can be either normal or extreme, and their preferences are such that it is a dominant strategy for the normal types to always announce a median position, while each extremist prefers a non-median policy and dislikes the bliss policy of the extreme opposition (Section 3.5). We first show that, if it is common knowledge that both candidates are extreme, then candidates would be indifferent between announcing their policy as the median and using ambiguous language. Then we show that, if there is a possibility of the candidates being normal, then in any symmetric equilibrium the extreme candidates keep being ambiguous for a long time over the course of the election campaign although the belief that the opposition is extreme can become very close to (but less than) one when the campaign is long. An example of the situation where this model can potentially fit is the Japanese House of Councillors election in 2014, in which Prime Minister Shinzo Abe avoided making the constitutional reform the main issue of the election and did not specify his plan of how to reform the constitution.[11] Nonetheless, the press argued that he had a particular preference for reforming the constitution such as specifying a foundation of the Self Defense Forces.[12] In fact, he started the process of summarizing the issues about the constitutional reform centered around the Self Defense Forces, once he won the election (Ota, 2016).

After discussing the applications, we present general principles that underlie the results from those applications (Section 4). The general results are concerned with the main topics in the applications, i.e., ambiguity, Condorcet winner, and office-motivated candidates, and show that the implications of our model are more robust than being valid in specific examples. First, Section 4.1 presents a result that we call the *long ambiguity theorem*. We formally define a "first-mover disadvantage" condition and show that, under that condition (together with some genericity conditions),

[10]Gerber et al. (2011) empirically show that the effect of campaign spending declines over time (earlier spending has a weaker effect). Although it might be reasonable to assume such depreciation, and it might result in gradual spending, we do not assume depreciation. Our main point is that the candidates have incentives to spend gradually even if the effect of depreciation is absent. Depreciation would strengthen the incentive to spend gradually.

[11]To reform the constitution, no less than two thirds of the members of the parliament have to agree on the reform at both the House of Representatives and the House of Councillors.

[12]For example, an article titled "Shinzo Abe's Constitution Quest" is published in Wall Street Journal (Harris, 2013).

candidates spend most of the time keeping their policy statements ambiguous in equilibrium, provided the campaign period is long enough. This result generalizes the results for valence candidates and a multi-dimensional policy space.

Second, Section 4.2 examines the robustness of the celebrated median-voter theorem to an extension to our dynamic setting. We show the *dynamic median-voter theorem*. Specifically, we consider a general model in which there is a Condorcet winner in the static version of the model, and show that each candidate makes an announcement corresponding to a Condorcet winner as soon as possible. This result generalizes results under some parameter specifications in our applications where the long ambiguity theorem fails. For example, we show that candidates announce their policies as soon as possible in the absence of valence, and spend their money as soon as possible if there are only two levels of spending. These results are implied by the dynamic-median voter theorem. The third general result presented in Section 4.3 pertains to the cases where candidates are purely office-motivated. Specifically, we analyze a general model with constant-sum payoffs, and prove that any perfect Bayesian equilibrium has the Markov property. We call this the *constant-sum Markov theorem*. The theorem generalizes the results from the model with valence and the one with multi-dimensional policy space. It is also used in proving a result for the model with incomplete information.

Section 5 concludes. The Appendix provides main proofs for the general results presented in Section 4. All the proofs not provided in the main text or in the Appendix are provided in the Online Appendix.

1.1 Literature Review

Ambiguity:

Ambiguous policy announcements have long been discussed in the politics and economics literature. The mechanism that generates ambiguity in our model is starkly different from those presented in the existing literature. For example, Shepsle (1972) and Aragonès and Postlewaite (2002) assume that candidates choose their policy positions simultaneously and once and for all. In their models, ambiguity occurs because voters are assumed to possess convex utility functions and therefore prefer ambiguity.[13] In our model, in contrast, ambiguity arises from dynamic strategic in-

[13]Callander and Wilson (2008) also consider a simultaneous-move voting game, and show that candidates' policy

teractions in an election campaign: Each candidate's strategic concern about the opponent's future play causes ambiguity. In particular, we do not assume convexity; rather, in one of the variants of our model discussed in the Online Appendix, we show that ambiguity occurs even when voters have concave utility functions.

Page (1976, 1978) proposes a theory that attributes ambiguity to the fact that candidates have limited resources to make their policy positions precise, and to voters' limited capacity to understand these positions. In our model, however, voters are capable of understanding what the candidates are announcing. Candidates do have a positive probability of not being able to have any chance to make a policy announcement, but we obtain ambiguity even in the limit as this probability shrinks to zero.

Glazer (1990) argues that ambiguity may occur if candidates are uncertain about the median voter's preferences. In his model, fixing a candidate's opponent's announcement, the candidate would prefer ambiguity. In our model, in contrast, our applications include cases where ambiguity is a suboptimal static response for any fixed announcement by the opponent. In other words, we obtain ambiguity due to dynamic strategic consideration.

Alesina and Cukierman (1990) and Aragonès and Neeman (2000) show that ambiguity occurs in elections if candidates prefer to keep the freedom to choose their policies after being elected, even though voters would prefer their candidates to commit themselves to precise policies before the election. That is, the driving force of ambiguity is different from office motivation. In contrast, the long ambiguity theorem in our model can be obtained with pure office motivation.

When the selection of candidates consists of more than one step, as is true for the US presidential election with its primaries and general elections, Meirowitz (2005) shows that candidates announce ambiguous policies in earlier stages if voter preferences are unknown at the beginning but are revealed by the result of the earlier stages. In our model, no new information arrives about voter preferences, and ambiguous policies are purely the result of strategic interactions between candidates.[14]

statements are ambiguous in equilibrium if voters' voting behavior is based on preferences exhibiting a taste for ambiguity due to their context dependence.

[14] Alesina and Holden (2008) show that candidates announce ambiguous policies even without primaries if (i) candidates have policy motivation, (ii) the policy motivation is their private information unknown to the voters, and (iii) campaign contributions from the voters to the candidates affect the electoral outcomes. In contrast, none of these assumptions are necessary to obtain the long ambiguity theorem in our model.

Valence: In the standard simultaneous-move Hotelling-Downs model with valence candidates, there exists no pure-strategy equilibrium: The strong candidate always wants to copy the weak candidate's policy, while the weak candidate does not want to be copied, just as in the "matching pennies" game. There are two approaches to addressing this issue in the literature. The first approach is to assume that the strong candidate is the incumbent and the weak candidate is the entrant (Bernhardt and Ingberman (1985), Berger et al. (2000), and Carter and Patty (2015)). In this approach, a typical result is that the strong candidate positions her policy close to the median voter and the weak candidate positions his policy at a slight distance from the strong candidate's policy, where the distance between the two policies is determined by the degree of asymmetry between candidates' valences.[15] The second approach is that of Aragonès and Palfrey (2002), who consider the simultaneous-move game seriously and characterize a mixed equilibrium.[16] They show that the strong candidate assigns high probabilities to the platforms which are close to the location of the median voter with high probabilities while the weak candidate assigns small probabilities to such platforms. Although these two approaches give us an understanding of what the equilibrium behavior looks like in an electoral situation with valence candidates, in both these models the order of policy announcements is exogenously given by the modelers. In contrast, we view our model with valence candidates as *endogenizing* the order of policy announcements.[17]

Multi-dimensional policy spaces: It is well known that the Downsian model with a multi-dimensional policy space does not have a pure-strategy Nash equilibrium unless a strong assumption about symmetry of the distribution of voters over the policy space is satisfied. As in the case with valence, one way to respond to the nonexistence is to consider a sequential game where the incumbent moves first and the challenger moves second. However, as Roemer (2001) argues, there may be no natural order, and we again view our approach as endogenizing the order of moves. Other approaches to deal with the nonexistence include that of Lindback and Weibull (1987), who allow the voters' behavior to be probabilistic and derive a sufficient condition for the existence of a pure-strategy equilibrium in a one-shot simultaneous-move game (see also Coughlin (1992)), and

[15]See also Ansolabehere and Snyder (2000) and Groseclose (2001) who consider pure-strategy equilibria in models with valence candidates.

[16]More specifically, Aragonès and Palfrey (2002) characterize the unique equilibrium in a discrete policy space and consider a limit as the discrete space approximates the standard continuous policy space. See also Hummel (2010).

[17]This provides a possible answer to the question posed by Aragonès and Palfrey (2002), who ask "What is the correct sequential model."

that of Roemer (2001), who obtains existence using a weaker equilibrium concept in which the set of feasible deviations for each candidate is restricted. In contrast to these papers, our analysis in Section 3.2 keeps the basic structure of the Downsian model.

Campaign spending: Most of the models in the theoretical literature about campaign spending specify how candidates use campaign funds in order to affect voters' behavior.[18] For example, Potters et al. (1997), Prat (2002a,b), and Coate (2004) consider models in which political campaigns can signal the candidate's private information. Also, Bailey (2002) assumes that one candidate chooses the policy position prior to the other, and that contributions can be used to target the campaign at selected people. In the current paper, in contrast, we are agnostic about why campaign spending helps, and focus on the timing of spending. There is also a large strand of empirical literature that analyzes the timing of the campaign spending. We refer interested readers to Gerber et al. (2011) and the reference therein.

Incomplete information: If candidates announce their policies simultaneously and the median voter exists, then it is a unique Nash equilibrium that both candidates announce the policy corresponding to the median voter given that the policy announcement is a credible commitment, regardless of policy preferences or knowledge about them. In models with incomplete information about the candidates' policy preferences, Banks (1990) and Harrington (1992) consider the case in which the policy announcement is not a credible commitment, while there is a cost of implementing a policy different from the announcement. Such a setting is later used by Kartik and McAfee (2007) and Callander and Wilke (2007) to analyze the incentive of telling a lie in elections. In our model, the policy announcement is a credible commitment, while its timing is endogenous (all the papers cited here assume exogenous (simultaneous) timing).

Dynamic games: To formally model the dynamics of policy announcements, we employ a framework with continuous time, a finite horizon, and a Poisson revision process. This modeling device has been extensively explored recently. The revision games in Kamada and Kandori (2017a,b) and Calcagno et al. (2014) consider settings in which players obtain opportunities to revise their preparation of actions according to Poisson processes, and the finally-revised action profile is implemented at the predetermined deadline.[19] In the models of these papers, revisions

[18] In addition, there are papers about the interaction between the lobbyists and politicians. See, for example, Austen-Smith (1987), Baron (1994), and Grossman and Helpman (1996).

[19] Ambrus and Lu (2015) consider a bargaining model in a similar fashion.

of actions are not restricted, in the sense that players can freely choose their actions from a fixed action space at each opportunity to move, as opposed to our assumption that once candidates make their policy platform clear, they cannot change it afterward.

Given the nature of the game analyzed—an election campaign where there is a clear winner and loser—some of our analysis pertains to constant-sum games. While the aforementioned papers mainly consider the situations where cooperation or coordination is important, Gensbittel et al. (2017) analyze general zero-sum revision games in which revisions are not restricted. Our long ambiguity theorem is similar to the "wait and wrestle" property that they find. The difference is that we do not restrict ourselves to constant-sum games, and we consider the case where revisions are restricted.[20]

The policy announcement timing game can be regarded as a stochastic game. Lovo and Tomala (2015) analyze general revision games with payoff-relevant states and show existence of Markov perfect equilibria.[21] In contrast, our focus is on the unique prediction of players' behavior in perfect Bayesian equilibria in the context of election campaigns.

We use a Poisson process to model frictions in the election campaign. Another way to model such frictions is to introduce switching costs. In general, switching costs result in different implications on equilibrium behavior from a Poisson process. See Lipman and Wang (2000) and Caruana and Einav (2008) for models with switching costs in finite-horizon games.

As for the idea of using ambiguous language or not spending their funds in expectation of future events, Gale's (1995, 2001) model of "monotone games" also considers a similar problem. In his model, at each period, players can only (weakly) increase their actions. In effect, players commit to a smaller and smaller subset of their action spaces as time passes, and they will never be able to "expand" that subset (thus, the revisions are restricted). The main difference is that he analyzes "games with positive spillover" played over an infinite horizon and shows that collusive outcomes can be achieved, while we analyze a game with a conflict of interests played over a finite horizon and are interested in uniqueness of an equilibrium outcome.

Another related theoretical literature is about commitment games, where each player simulta-

[20]Note that, although restricted revisions imply that there is no cycling choice of actions as in Gensbittel et al. (2017), it is still not trivial that candidates wait for a long time. Gensbittel et al. (2017) also discuss a comparison between the two models.

[21]Moroni (2018) also provides an existence proof for revision games, allowing for imperfect and incomplete information.

neously commits to a subset of the entire set of actions at the first stage and then plays the game with the restricted set of actions at the second stage (see, for example, Hamilton and Slutsky (1990, 1993); van Damme and Hurkens (1996); Romano and Yildirim (2005); Renou (2009)). These models sometimes yield multiple equilibria. In our analysis, in contrast, the revision opportunities arrive stochastically and asynchronously, and as a result, we obtain an (essentially) unique equilibrium prediction.

2 The Model – Policy Announcement Timing Game

There are two candidates, A and B. Whenever we say candidates i and j, we assume $i \neq j$.[22] There is a set of policies, X. For each candidate $i = A, B$, there is a collection of nonempty subsets of X, denoted $\mathcal{X}_i \subseteq 2^X \setminus \{\emptyset\}$, with a property that $X \in \mathcal{X}_i$ for each $i = A, B$. Each element in $\mathcal{X}_i$ is called i's "policy set." Here, we interpret announcing X as announcing the "ambiguous policy" while announcing other sets in $\mathcal{X}_i$ is seen as (at least partially) specifying a policy platform. Given a profile of policy sets $(X_A, X_B) \in \mathcal{X}_A \times \mathcal{X}_B$, let $v_i(X_i, X_j)$ be candidate i's payoff for each $i = A, B$.

In our policy announcement timing game, time flows continuously from $-T < 0$ to 0. Imagine that 0 is the fixed election date and the campaign starts at $-T$. For each $-t \in [-T, 0]$, according to the Poisson process with arrival rate $\lambda_i > 0$, each candidate $i = A, B$ obtains opportunities to announce her policy set. We assume that the Poisson processes are independent between the candidates. In particular, this implies that policy announcements are asynchronous with probability one. To simplify the exposition, we often use "enter" to denote the act of announcing a singleton set. The result of the election only depends on (X_A, X_B), where X_i with $i \in \{A, B\}$ is candidate i's most recently announced policy set at time 0 (the election date).

In what follows, we analyze perfect Bayesian equilibria of this game. To formally define strategies in our setting, we first define history. A **history** for candidate i is denoted by

$$\left(\left(t_i^k, X_i^k\right)_{k=1}^{k_i}, \left(t_j^l, X_j^l\right)_{l=1}^{l_j}, t, z_i\right),$$

where $-T \leq -t_i^1 < \ldots < -t_i^{k_i} < -t$; $X_i^k \in \mathcal{X}_i$ for all k; $-T \leq -t_j^1 < \ldots < -t_j^{l_j} < -t$; $X_j^l \in \mathcal{X}_j$ for all l; and $z_i \in \{yes, no\}$. The interpretation is that $-t_i^k$ is the time at which candidate i

[22] For ease of exposition, we use feminine pronouns to refer to A and i and masculine pronouns to refer to B and j.

receives his or her k'th revision opportunity, and X_i^k is the policy set that i has chosen at time $-t_i^k$. We assume that candidate i cannot observe whether candidate j receives an opportunity, but can observe candidate j's choice of a policy set whenever it changes.[23] That is, t_j^l is the l'th time that candidate j changes his or her policy set from the previous one, and X_j^l is the policy set that j has chosen at time $-t_j^l$. We let $X_i^0 = X_j^0 = X$, that is, the policy set at time $-T$ is exogenously given to be X. The third element t denotes the current remaining time, and the indicator z_i expresses whether there is an opportunity for candidate i at time $-t$. By $H_i^{k_i,l_j}$, we denote the set of histories in which candidate i for $i = A, B$ has received k_i opportunities in the past and in which candidate j has changed policy sets l_j times. The set of all histories for candidate i is $H_i := \bigcup_{k_i=0}^{\infty} \bigcup_{l_j=0}^{\infty} H_i^{k_i,l_j}$.

A **strategy** for candidate i is denoted by $\sigma_i : H_i \to \Delta(\mathcal{X}_i)$, with three restrictions: First, $\sigma_i(h_i) = X_i^{k_i}$ where k_i is specified in the first element of h_i if the fourth element in h_i specifies $z_i = no$. That is, if there is no opportunity at $-t$ for i, then for notational convenience, we specify that the candidate takes the same policy set as specified in the last opportunity. Second, if $z_i = yes$, then the strategy $\sigma_i(h_i)$ must assign probability zero to $X_i \in \mathcal{X}_i$ if $X_i \not\subseteq X_i^{k_i}$. Thus, the set of candidate i's possible announcements at time $-t$ depends on i's own past policy announcement: If i has already announced $X_i \in \mathcal{X}_i$ in the past, then i can only announce a (weak) subset of X_i. Thus, once a candidate rules out some of the potential platforms, then she cannot go back to them later. The third requirement is technical. To guarantee that candidates' payoffs are integrable with respect to the distribution of the final outcome given the strategy profile, we require that $\sigma_i(h_i)$ puts a positive probability only on a countable subset of $\mathcal{X}_i$.

Let Σ_i be the set of all strategies of candidate i. Let $u_i(\sigma|h_i, h_j)$ be candidate i's continuation payoff given history profile $(h_i, h_j) \in H_i \times H_j$ and the strategy profile $\sigma \in \Sigma_A \times \Sigma_B$.[24] Let $H_j(h_i)$ be the set of candidate j's feasible histories given h_i, and let $\beta(\cdot|\cdot) : H_i \to \Delta(H_j)$ be candidate i's belief about candidate j's history such that $\int_{h_j \in H_j(h_i)} d\beta(h_j|h_i) = 1$ for each $h_i \in H_i$ for each $i = A, B$. Given a belief β, let $u_i^\beta(\sigma|h_i) = \int_{h_j \in H_j} u_i(\sigma|h_i, h_j) d\beta(h_j|h_i)$ be candidate i's expected continuation payoff given h_i. A strategy profile (σ_A^*, σ_B^*) is a **perfect Bayesian equilibrium (PBE)** if there exists a belief β such that, for each $i \in \{A, B\}$, (i) $\sigma_i^* \in \arg\max_{\sigma_i \in \Sigma_i} u_i^\beta(\sigma_i, \sigma_j^*|h_i)$

[23] Except for the incomplete information model in Section 3.5, the prediction of the model will be the same even if candidate i can observe all of candidate j's opportunities, in the sense formalized in the Constant-Sum Markov Theorem (Theorem 3) and Remark 5.

[24] This is well defined because H_i is a countable union of subsets of a finite-dimensional space.

holds for every $h_i \in H_i$ and (ii) β is derived from Bayes rule whenever possible.[25]

3 Examples

We first offer various examples to show that the model of policy announcement timing game enables us to analyze rich strategic considerations when it is applied to otherwise well-known and canonical models of elections.

3.1 Valence Election Campaign

We consider the case in which one candidate is stronger than the other, in the sense that if two of them choose the same policy set, then the former candidate wins. Section 3.1.1 introduces the model. In Section 3.1.2, we establish that if two candidates are perfectly symmetric, then both candidates would want to be clear as soon as possible. In Section 3.1.3, we show that if one candidate is slightly stronger than the other, then there are rich strategic considerations driving the incentive for each candidate to make an ambiguous policy announcement. The incentive for ambiguity follows from the "first-mover disadvantage": The strong candidate wants to copy the weak candidate's policy after the weak candidate enters, while the weak candidate does not want to be the first mover as being copied is the worst outcome. This result presents a novel connection between ambiguity and valence.

3.1.1 The Model

In the language of the general model, candidate A is the strong candidate S, and candidate B is the weak candidate W. We keep the model simple, so as to highlight the complexity introduced by the campaign phase into an election model with valence candidates. In particular, the policy space is assumed to be $X = \{0, 1\}$, and each candidate i's available policy sets are $\mathcal{X}_i = \{X, \{0\}, \{1\}\}$. The Online Appendix presents a general version of the model that involves many other cases, such as a continuous policy space.

If a candidate enters at 0 (or 1) and the other enters at 1 (or 0) or does not enter, then

[25] Although each information set at any time after $-T$ has probability zero, one cay apply Bayes rule to calculate relevant conditional probabilities because any Poisson process has a countable number of arrivals with probability one. We formally define Bayes rule for our context in the Online Appendix.

the former wins with probability p (or $1-p$); if the two candidates enter at the same policy or neither of them enters, then the strong candidate wins with probability one. Candidates are purely office-motivated. That is, we have $v_S(\{0\},\{0\}) = v_S(\{1\},\{1\}) = v_S(X,X) = 1$, $v_S(\{0\},\{1\}) = v_S(\{0\},X) = v_S(X,\{1\}) = p$, $v_S(\{1\},\{0\}) = v_S(\{1\},X) = v_S(X,\{0\}) = 1-p$, and $v_W(X_W,X_S) = 1 - v_S(X_S,X_W)$ for each (X_S,X_W). We assume $p \in (0,\frac{1}{2})$.

This utility function can be micro-founded in the following manner. Suppose that there are a continuum of voters, located at policy 0 and policy 1. The distribution of the voters' locations is stochastic, and with probability p, policy 0 has more voters. During the campaign, the locations of the voters are unknown.

If a candidate $i \in \{S,W\}$ wins the election and implements policy $x \in \{0,1\}$, then a voter with position $y \in \{0,1\}$ obtains a payoff of

$$u(|x-y|) + \delta \cdot \mathbb{I}_{i=S},$$

where $u(0) > u(1)$ and $0 \leq \delta < (u(0)-u(1))/2$, with δ representing the advantage of candidate S due to her charisma or other asymmetries between candidates' characteristics that are unrelated to the policy choices.[26] The voters believe that, if candidate i has specified a policy $x \in \{0,1\}$ and wins, then x will be implemented. If candidate i with the ambiguous policy $X_i = \{0,1\}$ wins, then the voters believe that the policies $\{0\}$ and $\{1\}$ will be implemented with equal probability $\frac{1}{2}$.[27] The voters are sincere, that is, they each vote for the candidate who, if elected, maximizes their expected payoff. The candidate with more votes wins. In the case of a tie, each candidate wins with probability 1/2.

The candidate who obtains more votes wins, and obtains a payoff of 1, while the other candidate obtains a payoff of 0; these are the only payoffs that they receive in this model, i.e., candidates are purely office-motivated. Each candidate's objective is to maximize the expected payoff, that is, their objective is to maximize their probability of winning. The payoff function (v_S, v_W) that we

[26] One way to interpret δ in a "policy related" manner would be to consider a model as in Krasa and Polborn (2010), in which candidates choose one policy out of two for each of multiple policy issues. If candidates make policy announcements for some issues first, they then would compete by choosing policies on remaining issues, where asymmetry between candidates may exist depending on the relative popularity of the policies that each candidate has chosen already. We note that, if $\delta > (u(0)-u(1))/2$, it will be straightforward to show that S wins the election with probability 1 in any PBE.

[27] It is not crucial that the probability is exactly $\frac{1}{2}$. For an open set of probabilities for tie-breaking, our main results are unchanged.

(X_S, X_W) at the deadline	Voters at 0 vote for	Voters at 1 vote for	S's expected utility	W's expected utility
$(\{0,1\},\{0,1\})$	S	S	1	0
$(\{0,1\},\{0\})$	W	S	$1-p$	p
$(\{0,1\},\{1\})$	S	W	p	$1-p$
$(\{0\},\{0,1\})$	S	W	p	$1-p$
$(\{0\},\{0\})$	S	S	1	0
$(\{0\},\{1\})$	S	W	p	$1-p$
$(\{1\},\{0,1\})$	W	S	$1-p$	p
$(\{1\},\{0\})$	W	S	$1-p$	p
$(\{1\},\{1\})$	S	S	1	0

Table 1: Voter behaviors and the expected payoffs for the valence election campaign.

specified above can be obtained by assuming $\delta > 0$. We summarize in Table 1 the voters' behaviors and the resulting expected payoffs for the candidates, given these specifications and $\delta > 0$. Note that, without valence ($\delta = 0$), the environment just specified is the one in which we can apply the median voter theorem in the static version of the model; that is, it is each candidate's dominant action to announce $\{1\}$.

We let $\lambda_S = \lambda_W =: \lambda$.[28] We call this dynamic game a *valence election campaign*. It is characterized by a tuple (p, T, λ).

3.1.2 The Benchmark Case: Perfectly Symmetric Candidates

Before analyzing the model with valence, we analyze the model with symmetric candidates as a benchmark case. The only difference from the model with valence is that, if two candidates end up announcing the same policy set, both of them win with probability $\frac{1}{2}$ (this corresponds to setting $\delta = 0$ in the micro-foundation). Call this game a *no-valence election campaign*. It turns out that there are no incentives to announce the ambiguous policy $\{0, 1\}$.

The following proposition gives us a stark result:

Proposition 1 *In any no-valence election campaign, in any PBE, each candidate announces* $\{1\}$ *as soon as possible.*

To see why this holds, fix time $-t$ and suppose that at any time $-s > -t$, if each candidate has

[28] This assumption is generalized in the Online Appendix.

an opportunity to enter, then he/she enters at 1. Then, at time $-t$, if no candidate has entered, entering at 1 gives the payoff strictly greater than $\frac{1}{2}$, entering at 0 gives $p < \frac{1}{2}$, and not entering gives a payoff of $\frac{1}{2}$ by symmetry of the supposed continuation strategies. Thus, entering at 1 is a unique best response. Therefore, by continuity of the probability in time which implies continuity of the continuation payoff in time, for sufficiently small $\varepsilon > 0$, it is uniquely optimal to enter at 1 for all time in $(-t-\varepsilon, -t]$ if no one has entered. Under the history at which the opponent has entered, an analogous argument shows that entering at 1 is uniquely optimal. This establishes the desired result.[29]

In the next subsection, we demonstrate that (i) the above simple argument breaks down once we introduce asymmetry with respect to candidates' valence ($\delta > 0$ in the micro-foundation), and (ii) candidates face complicated dynamic incentive problems, which involve ambiguous policy announcements. Therefore, a small valence (or small $\delta > 0$) matters and is the key for ambiguous policy announcements.

3.1.3 The Cases with Valence Candidates

Let us start with the following lemma. It states that, if S has an opportunity to enter after W has entered at $x \in \{0, 1\}$, then she enters at x and wins for sure. In contrast, if W has an opportunity to enter after S has entered at $x \in \{0, 1\}$, then he is indifferent between announcing $\{0, 1\}$ and entering at $x' \in \{0, 1\} \setminus \{x\}$. These two conclusions imply that, since the median is more likely to be at policy 1 ($p < \frac{1}{2}$), if a candidate enters before the opponent, he/she enters at $\{1\}$.

Lemma 1 *In any valence election campaign with (p, T, λ), in any PBE, the following are true at any time $-t$:*

1. *Given that W has already entered, S enters at the same platform as soon as possible.*

2. *Given that S has already entered, W is indifferent between announcing $\{0, 1\}$ and entering at the platform different from S's.*

3. *If a candidate i enters before the opponent, then i enters at policy 1.*

[29] This last part follows from the "continuous-time backward induction" that we formally present in Appendix B.

The above lemma characterizes the equilibrium behaviors on and off the equilibrium path except when no candidates have yet entered. It also says that if both are still using ambiguous language and a candidate i enters, then i enters at policy 1. Hence, in the following analysis, we consider the incentives to enter at policy 1 when both are still using ambiguous language.

Before presenting the characterization of the behavior in a PBE in such a situation, we first provide the basic intuition, which exploits the idea that being the first mover is disadvantageous. For the time being, consider the case with $p = \frac{1}{2}$.[30] Suppose that at time $-t$, both S and W have previously announced $\{0, 1\}$. On the one hand, if there is no further revision, W's payoff is 0. So W needs to specify his policy to obtain a positive payoff. Thus, W announces $\{0\}$ or $\{1\}$ at some point in $[-t, 0]$, if he can. Since $\{0\}$ and $\{1\}$ are symmetric with $p = \frac{1}{2}$, assume without loss of generality that W announces $\{1\}$ when he clarifies his policy.

On the other hand, S does not have an incentive to specify her policy until W specifies his policy. This is because she gets $\frac{1}{2}$ for sure by specifying her policy, while using ambiguous language at all times in $[-t, 0]$ gives her either $\frac{1}{2}$ or 1 with the latter taking place with positive probability (when W does not enter afterward and when W enters and S copies his policy).

If W announces $\{1\}$ in the early stages of the campaign, then the probability with which S enters afterward is high. So W wants to postpone announcing. But waiting too much is not optimal for W either, since if he does not have a chance to revise his policy set, W gets a payoff of 0. So there should exist a "cutoff," $-t^*$, until which W announces $\{0, 1\}$ and after which W announces $\{1\}$ when he gets an opportunity of a policy announcement.

Recall that we do not have this type of strategic consideration in the no-valence election campaign ($\delta = 0$), even if we extend the model to include the case with $p = \frac{1}{2}$. The simple argument we provided for Proposition 1 breaks down since the continuation payoff after taking each action is different once we introduce valence. For example, W expects a payoff close to zero if he specifies some policy when the deadline is far away in the valence election campaign, as opposed to a payoff of $\frac{1}{2}$ that he gets in the no-valence election campaign.

Next, consider the case with $p = 0$. In this case, S would want to commit to $\{1\}$ as soon as possible, because she can then obtain a payoff of 1, which is the highest possible payoff. Since W can win if and only if he enters at $\{1\}$ and S does not have an opportunity, W also enters at $\{1\}$

[30] Strictly speaking, since $p < \frac{1}{2}$, this is actually outside of the model, but we consider such a case to provide the intuition. The same comment applies to the case $p = 0$ that we consider next.

as soon as possible.

The next proposition fully characterizes the form of PBE for each $p \in (0, \frac{1}{2}) \setminus \{\frac{1}{1+e}\}$. Suppose that the current policy set of each candidate is $\{0, 1\}$. The equilibrium strategy of W is to wait until a finite cutoff time and to enter as soon as possible after that cutoff. In contrast to the case of $p = 0$, the cutoff is finite for any strictly positive p because the probability that the median voter is at 0 is strictly positive. The equilibrium strategy for S depends on the value of p, and the value $p = \frac{1}{1+e}$ corresponds to the cutoff at which S's incentive changes. If p is close to $\frac{1}{2}$ ($p > \frac{1}{1+e}$, considered in part 1 of Proposition 2), S does not enter until W enters for the same reason as in the case of $p = \frac{1}{2}$. In contrast, for small p ($p < \frac{1}{1+e}$, considered in part 2 of Proposition 2), S enters when the deadline is far away as when $p = 0$, but does not do so when the deadline is close. The intuition for the ambiguity near the deadline is as follows: If S obtains an opportunity when the deadline is close, then the probability with which W has a chance to announce his policy afterward is small. So it is likely that W uses ambiguous language at the deadline. Thus, keeping ambiguous language is profitable for S, because by doing so, S gets a payoff of 1 with a high probability.

Proposition 2 *Consider the valence election campaign with (p, T, λ). There exists a PBE. Moreover, there exist $t^* := \frac{1}{\lambda}$, t_S, and t_W (the latter two depend on p) that are independent of T such that, for any PBE, the following are satisfied if the previous policy sets are both $\{0, 1\}$.*[31]

1. *If $p > \frac{1}{1+e}$, the following hold:*[32]

   (a) *S announces $\{0, 1\}$ for all $-t \in (-\infty, 0]$.*

   (b) *W announces $\{0, 1\}$ for all $-t \in (-\infty, -t^*)$ and $\{1\}$ for all $-t \in (-t^*, 0]$.*

2. *If $p < \frac{1}{1+e}$, then the following hold:*

   (a) *S announces $\{1\}$ for all $-t \in (-\infty, -t_S)$ and $\{0, 1\}$ for all $-t \in (-t_S, 0]$.*

   (b) *W announces $\{0, 1\}$ for all $-t \in (-\infty, -t_W)$ and $\{1\}$ for all $-t \in (-t_W, 0]$.*

   (c) *Moreover, $-t_W < -t_S$, $\frac{dt_S}{dp} > 0$ and $\frac{dt_W}{dp} < 0$.*

[31] If $p = \frac{1}{1+e}$, then there is indeterminacy about S's equilibrium strategy at all $-t < -t^*$ since she is indifferent.

[32] Although the entire game lasts for the time interval $[-T, 0]$, we state results for all times in $(-\infty, 0]$ as the results do not rely on whether the cutoff times (at which equilibrium actions change) are earlier or later than $-T$. Any statement about time interval $K \subseteq (-\infty, 0]$ should be interpreted as a statement about the time interval $K \cap [-T, 0]$. The same caution applies to all other formal statements involving time intervals.

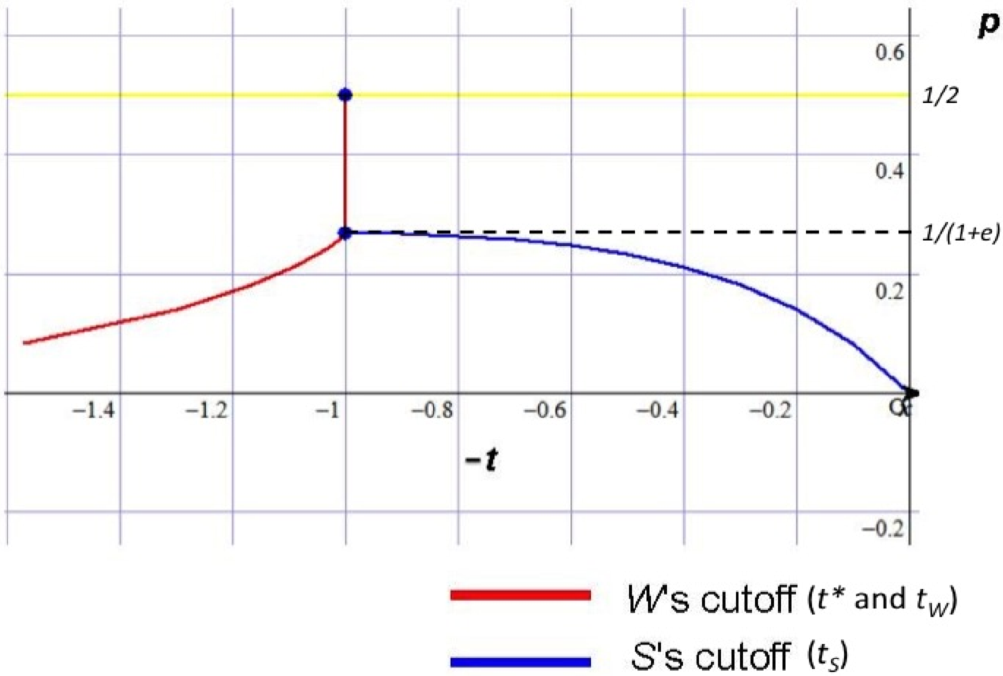

Figure 1: Cutoff times for the valence election campaign ($\lambda = 1$).

Note that the cutoffs are independent of T. Hence, when T and p are large, we expect that candidates use ambiguous language for most of the campaign period. Note also that stretching T and enlarging λ with the same ratio are equivalent. Hence, this also implies that for a fixed length of campaign period T, if we consider the situation in which the opportunities arrive frequently, candidates spend most of the time in $[-T, 0]$ using ambiguous language.

In Figure 1, we depict the times t^*, t_S, and t_W that appear in Proposition 2, for different values of p for the case of $\lambda = 1$. For example, $p = .4$ ($> \frac{1}{1+e}$) corresponds to part 1 of the proposition. In this case, there is one point at which the graph in the figure intersects with the $p = .4$ line, so as a result, the time spectrum is divided into two regions: In the left region, no candidate enters. In the right region, S does not enter while W enters. When $p = .2$ ($< \frac{1}{1+e}$), there are two intersections, and as a result the time spectrum is divided into three regions: In the left-most region, S enters while W does not enter. In the middle region, both candidates enter. Finally, in the right-most region, S does not enter while W enters.

Notice that this particular model predicts that when p is small ($p < \frac{1}{1+e}$) and $T > t_S$, S enters as soon as possible, so if T is large, then there would be almost no ambiguity in equilibrium. This hinges on our assumption that even if W enters after S, S does not incur any loss. In the Online Appendix, we show that if there is even a small loss, S prefers to use ambiguous language until some point in time that does not depend on the horizon length T, and so the modified model is consistent with ambiguity even if p is low. Despite this feature, we believe that the simple model in this section provides a basic intuition about the dynamic incentives that candidates face. The basic takeaway is that the nature of the election game with valence leads candidates to strategically "time" their announcements, since the benefit and cost of maintaining flexibility of choice vary over time. Consider W's incentive, for example. On the one hand, the benefit comes from the fact that the election game is constant-sum, so avoiding being the first mover is a good thing. On the other hand, the cost comes from the difference in valence. He does not want to end up making the same choice as S (that is, taking $\{0, 1\}$). This is the general trade-off of timing strategies faced by electoral candidates, and our model succinctly captures such a trade-off.

Remark 1 (Empirical implication) Note that Proposition 2 applies to any PBE. This uniqueness property enables us to conduct meaningful comparative statics, which one can potentially test empirically. The analysis shows that ambiguity is likely when the probability distribution of the median voter's position is close to uniform (p is close to $1/2$). This is consistent with Campbell (1983) who suggests that opinion dispersion has a strong positive effect on the ambiguity in candidates' language.[33] Also, a researcher would be able to infer which candidate is stronger, given the information about the timing of entry or the final policy profiles announced. More detailed accounts of these claims are in the Online Appendix.

Remark 2 (Robustness of the prediction) The basic structure of the equilibrium is robust even if the two candidates have different arrival rates, although the fine details change. One can show that a relatively higher arrival rate makes the candidate better off. This is due to the fact that the underlying game is constant-sum, and is in a stark contrast to the results for coordination games in Calcagno et al. (2014) that having a higher arrival rate makes the player worse off since it

[33] Specifically, we have in mind a situation where n voters are independently distributed over $\{0, 1\}$ where the probability on the policy 0 is $q < \frac{1}{2}$. A higher q suggests more option dispersion (a higher standard deviation of the preferred policies among the voters. Campbell (1983) also considers standard deviation), and corresponds to a higher p.

decreases his/her commitment power. More detailed discussions about heterogeneous arrival rates and a general model with heterogeneous arrival rates and a general class of payoff functions are provided in the Online Appendix.

Remark 3 (Welfare implications) One may be tempted to conduct a welfare analysis resorting to the micro-foundation we provided, but there is a caveat in doing so: The distribution of the median voter does not necessarily pin down the voter distribution at each realized state of the world. With additional assumptions about the voter distribution, one can conduct welfare analysis. For example, suppose that there is a single voter. It is then necessary that this single voter's ideal policy is 0 with probability p and 1 with probability $1-p$. Then, one can show by a calculation that the voter's expected payoff in our model is smaller than under a unique mixed Nash equilibrium model in which each candidate chooses between 0 and 1 as in Aragonès and Palfrey (2002) when our model predicts long ambiguity ($p > \frac{1}{1+e}$), the valence term $\delta > 0$ is sufficiently small, and T is sufficiently large.[34] Although we acknowledge that welfare analysis is an interesting direction of research, we do not explore it under other assumptions about the voter distribution in the valence election campaign, or in other examples we present. This is because such an exercise necessities additional assumptions about voter distributions which are not necessary when considering our main focus on candidates' timing problems and resulting policies.

3.2 Multi-dimensional Policy Space – the Case with Purely Office-Motivated Candidates

When the policy space is multi-dimensional, there does not generally exist a Condorcet winner, and a pure-strategy Nash equilibrium does not exist in a static model. In this section, in contrast, we show that our policy announcement timing game admits existence of a PBE. The ambiguity again results from a disadvantage of being the first-mover, which follows from the nonexistence of a Condorcet winner.

Suppose that the policy space X is a full-dimensional connected subset of $\mathbb{R}^n$ for some $n \in \mathbb{N}$. The voters are distributed according to a measure $\mu \in \Delta(X)$. As a normalization, let $\mu(X) = 1$. We assume that $\mu(Y) = 0$ for any zero-Lebesgue measure set $Y \subseteq \mathbb{R}^n$. Given a policy profile

[34] The calculation is given in the Online Appendix.

$(x_A, x_B) \in X \times X$, we define the set of supporters for each candidate as:

$$S_A(x_A, x_B) = \{x \in X | \, |x - x_A| < |x - x_B|\} \text{ and } S_B(x_B, x_A) = \{x \in X | \, |x - x_A| > |x - x_B|\},$$

where $|\cdot|$ denotes the Euclidian distance. We define the probability of A's winning, $P_A(x_A, x_B)$, to be 1 if $\mu(S_A(x_A, x_B)) > \mu(S_B(x_B, x_A))$, $\frac{1}{2}$ if $\mu(S_A(x_A, x_B)) = \mu(S_B(x_B, x_A))$, and 0 otherwise. Let the probability of B's winning be $P_B(x_B, x_A) = 1 - P_A(x_A, x_B)$. An interpretation is that each voter receives a utility that is strictly decreasing in the Euclidian distance between her bliss point and a policy, and supports the candidate with the policy that would give rise to a strictly higher utility.

Each candidate is purely office-motivated: She receives payoff 1 if elected, and 0 otherwise. Each candidate's objective is to maximize the expected payoff.

Each problem is characterized by a pair (X, μ). Let the set of all problems be $\mathcal{P}$. Define

$$\mathcal{M} = \left\{ (X, \mu) \in \mathcal{P} \mid \exists x^* \in X \text{ s.t. } \forall y \in X \setminus \{x^*\}, \; \mu\left(\{z \in X | (y - x^*) \cdot (z - x^*) > 0\}\right) = \frac{1}{2} \right\}.$$

Notice first that if $X \subseteq \mathbb{R}$, that is, when X is uni-dimensional, then $(X, \mu) \in \mathcal{M}$ holds for any μ. Notice second that for multi-dimensional X, $\mathcal{M}$ is imposing a severe symmetry condition. For example, if μ is the uniform distribution over X, then $(X, \mu) \in \mathcal{M}$ is equivalent to X being point symmetric. Also, for a given multi-dimensional X, $(X, \mu) \notin \mathcal{M}$ holds generically in the space of μ.[35] Third, when $(X, \mu) \in \mathcal{M}$, the x^* satisfying the condition in the definition of $\mathcal{M}$ is uniquely determined because X is connected. We denote this unique x^* by $x^*(X, \mu)$.

In the policy announcement timing game, for each i, the available policy sets are all the singletons and the entire set X, so $\mathcal{X}_i = \{X\} \cup \left(\bigcup_{x \in X} \{\{x\}\}\right)$ for each $i = A, B$. To define the vote share and winning probabilities for the case in which some candidate does not enter, we expand the domain of S_i with a restriction that $X \setminus (S_i(X_i, X_j) \cup S_j(X_j, X_i))$ has measure zero for any (X_i, X_j) such that $X_i \neq X_j$, and $S_i(X, X) = \emptyset$ for each $i = A, B$. That is, unless the two candidates specify the same policy set, the set of indifferent voters has measure zero. The domain of P_i is expanded accordingly for each $i = A, B$.[36] The payoffs are the same as in the static model:

[35] See Theorem 7.2 of Roemer (2001) for the detail.

[36] That is, $P_A(X_A, X_B)$ is 1 if $\mu(S_A(X_A, X_B)) > \mu(S_B(X_B, X_A))$, it is $\frac{1}{2}$ if $\mu(S_A(X_A, X_B)) = \mu(S_B(X_B, X_A))$, and is 0 otherwise. We let the probability of B's winning be $P_B(X_B, X_A) = 1 - P_A(X_A, X_B)$.

$v_i(X_i, X_j) = P_i(X_i, X_j)$ for each $X_i \in \mathcal{X}_i$, $X_j \in \mathcal{X}_j$, and $i = A, B$. We assume that there exists $\bar{x} \in X$ such that $v_i(\{\bar{x}\}, X) = 1$ for each $i = A, B$, which implies that clarifying a policy position is better when the other candidate is ambiguous. Moreover, once a candidate enters, she prefers the opponent not to enter, that is, $\mu(S_i(x_i, X)) > \inf_{x' \in X} \mu(S_i(x_i, x'))$ holds for each $x_i \in X$, $i = A, B$. This assumption is satisfied if voters believe that candidates, without specifying a policy, take a policy randomly upon being elected, and the voter utility functions are strictly concave. Finally, $S_i(X, X) = \emptyset$ implies that if no one has entered, then the winning probabilities are half-half. We call this dynamic game a *symmetric office-motivated election campaign*. It is characterized by a tuple $(X, \mu, T, \lambda_A, \lambda_B)$.

Proposition 3 *Consider a symmetric office-motivated election campaign with* $(X, \mu, T, \lambda_A, \lambda_B)$. *There exists a PBE, and the following are true.*

1. *Suppose that* $(X, \mu) \in \mathcal{M}$. *Then, in any PBE, conditional on any history, each candidate* i *announces* $x^*(X, \mu)$.

2. *Suppose that* $(X, \mu) \notin \mathcal{M}$. *Then, there exist* $t_A^*, t_B^* \in (0, \infty)$ *such that, in any PBE, if no one has entered at time* $-t$, *candidate* i *does not enter if* $-t \in (-\infty, t_i^*)$, *and does enter at some policy if* $-t \in (t_i^*, 0]$. *It must be the case that* $\operatorname{sign}(\lambda_A - \lambda_B) = \operatorname{sign}(t_A^* - t_B^*)$.

Remark 4 (Existence of a pure-strategy Nash equilibrium) Note that $(X, \mu) \in \mathcal{M}$ if and only if there exists a pure-strategy Nash equilibrium in the static game in which each candidate chooses a policy in X.[37] Hence, the proposition shows that ambiguity emerges in a PBE if and only if there is no pure-strategy Nash equilibrium in such a static game.

Part 1 implies that, if there exists a Condorcet winner, then it is optimal to announce that policy as soon as possible. In part 2, intuitively, each candidate's strategic situation is similar to that of the weak candidate in the valence election campaign (Section 3.1): If the other candidate cannot enter, she prefers entering to not entering (the former gives a payoff of 1 while the latter gives $\frac{1}{2}$). However, if the other candidate can enter, then she prefers not entering to entering (the

[37]The reason for the "if" direction is that there always exists at least one candidate, say i, who receives no more than half of the entire vote share in a Nash equilibrium, and $(X, \mu) \notin \mathcal{M}$ implies that there exists a policy close to j's policy such that i always has an incentive to deviate to it to receive a vote share strictly higher than $1/2$ (see, for example, Theorem 7.1 of Roemer (2001) for a related result).

former gives a positive expected payoff while the latter gives 0). As a result, it is optimal not to enter if the deadline is far because the probability that the other candidate can enter afterward is large. If the deadline is close, however, since the probability of such an event is small, it is optimal to enter.

If candidate B can only move slower ($\lambda_B < \lambda_A$), then the proposition predicts that he is more likely to be ambiguous at the election date, and conditional on entering, the expected entry time is later. For the entry time, there are two opposing forces: On the one hand, since candidate B cannot move fast, the risk of him not being able to enter afterward is substantial. This force would make him willing to enter early. On the other hand, candidate B knows that candidate A is likely to obtain an opportunity later, and this would make him willing to wait until the last moment. Since the loss from the latter is particularly large, he does not want to enter until the last moment ($t_B^* < t_A^*$).

An implication of Proposition 3 is that the faster candidate is more likely to win:

Proposition 4 *Consider a symmetric office-motivated election campaign with* $(X, \mu, T, \lambda_A, \lambda_B)$*. If* $\lambda_A > \lambda_B$*, then for any PBE, candidate* A*'s expected payoff is strictly greater than that of candidate* B*.*

The proposition is straightforward if $(X, \mu) \in \mathcal{M}$. The case of $(X, \mu) \notin \mathcal{M}$ may seem subtle, but there is a simple intuition: Given the previous proposition, candidate B does not enter until $-t_B^*$. Candidate A can obtain a higher payoff than candidate B by simply waiting until time $-t_B^*$ because A receives opportunities more frequently than B does after $-t_B^*$. Such a strategy is suboptimal but provides a lower bound of candidate A's PBE payoff.

Finally, we state an implication of Proposition 3 on the relationship between the dynamics in PBE and the dimensionality of the policy set.

Corollary 1 *Fix* X*. The following are true:*

1. *If* X *is uni-dimensional* ($n = 1$)*, then for any* $(\mu, T, \lambda_A, \lambda_B)$*, the following is true: In the symmetric office-motivated election campaign with* $(X, \mu, T, \lambda_A, \lambda_B)$*, there exists a PBE. Moreover, in any PBE, conditional on any history, each candidate* i *announces* $x^*(X, \mu)$*.*

2. *If X is multi-dimensional ($n \geq 2$), then for generic μ, for any $(T, \lambda_A, \lambda_B)$, the following are true: There exists a PBE. Moreover, there exist $t_A^*, t_B^* \in (0, \infty)$ such that, in any PBE, if no one has entered at time $-t$, candidate i does not enter if $-t \in (-\infty, t_i^*)$, and does enter at some policy if $-t \in (t_i^*, 0]$.*

Notice that the results in this section show the uniqueness of a distribution of entry times in any PBE, while it does not show the uniqueness of the policies to which the candidates enter. In fact, there may exist multiple PBE due to the fact that there may exist multiple policies at which each candidate can win if the opponent does not enter and multiple policies at which each candidate can win when she enters after her opponent. In the next section, we consider the case with a multi-dimensional policy space with policy-motivated candidates and show that, with policy motivation, uniqueness of policies may obtain.

3.3 Multi-dimensional Policy Space – the Case with Policy-Motivated Candidates

We again consider the policy announcement timing game with a multi-dimensional policy space, but now with policy-motivated candidates. We show that, in a PBE, if a candidate cares about the policy implemented by the winner of the election, then she may announce a Pareto-inefficient policy to influence a later announcement by the opposition party. By announcing such a policy, she can induce the opponent to implement a policy that is not too undesirable even in the event that she loses.

Specifically, we consider the following setting of Persson and Tabellini (2000): $X = \{(x_1, x_2) \in [0,1]^2 : x_1 + x_2 \leq 1\}$. Here, a higher x_1 is interpreted as a more conservative economic policy and a higher x_2 is interpreted as a more aggressive military policy. There are three voters: Voter 1's ideal policy is $(1, 0)$ and her utility from policy x is $-(1 - x_1)$. That is, she is right-wing and only cares about the economic policy. Voter 2's ideal policy is $(0, 1)$ and her utility from policy x is $-(1 - x_2)$. That is, she is also right-wing and only cares about the military policy. Finally. voter 3's ideal policy is $(0, 0)$ and her utility from policy x is $-x_1 - x_2$. That is, she generally likes a left-wing policy.

There are two candidates L and R, whose ideal policies are $(0, 0)$ and $\left(\frac{1}{2}, \frac{1}{2}\right)$, respectively.[38]

[38] When we need to distinguish between the two candidates, we use a masculine pronoun for L and a feminine

Their ideal policies are common knowledge, and the voters correctly believe that the candidate who wins without specifying a policy will implement her ideal policy. If a candidate wins with a specified policy x, then she must implement x. The voters vote for the candidate who brings the higher utility, with a tie broken equally in favor of the entrant if there is only one candidate who enters, and in favor of the candidate who enters later if both enter.[39] This in particular implies that R collects two votes when no candidate enters. The candidate collecting two or three votes wins. Since R has an ideal policy that is preferred by two voters (voters 1 and 2), she has a chance of winning with probability 1 if no candidate specifies a policy. In this sense, candidate R is similar to the "strong candidate" in the valence election campaign analyzed in Section 3.1. We will show, however, that the distribution of entry times differs from the one for that model because the payoff from the entry is specified differently.

If a candidate $k \in \{L, R\}$ wins the election and implements policy x, the payoff of candidate $i \in \{L, R\}$ is

$$\mathbb{I}_{i=k} + \varepsilon u_i(x),$$

where $u_L(x) = -\max_{n\in\{1,2\}} x_n$ and $u_R(x) = \min_{n\in\{1,2\}} x_n$ are the utility functions to represent candidates' policy preferences, and $\varepsilon > 0$.[40] The payoff function v_i for each $i = L, R$ is specified accordingly. Persson and Tabellini (2000) show that there is no Condorcet winner (no median voter) and there is no pure-strategy Nash equilibrium in the simultaneous-move game in which choosing X is not allowed.

In the policy announcement timing game, for each i, the available policy sets are again all the singletons and the entire set X, so $\mathcal{X}_i = \{X\} \cup \left(\bigcup_{x\in X} \{\{x\}\}\right)$. As a tie-breaking rule, we assume that if it is optimal for a candidate to enter and $\bar{X}$ is the set of all policies such that entering at any policy in $\bar{X}$ generates the maximum continuation payoff, then she enters at a policy in

pronoun for R.

[39] This tie-breaking rule is consistent with considering a limit of unique PBEs in models with discrete policy spaces. Palfrey (1984) conducts the same exercise of taking a limit in a game where best responses do not exist with a continuous policy space. To be precise, this tie-breaking rule violates the assumption that the payoffs to the candidates are determined solely by the functions $(v_i)_{i=L,R}$ which only depend on the profile of policy sets. One could define a value function that depends both on the policy-set profile and on the times at which they are announced, but we do not write out such formalization in light of the justification due to discrete policy spaces and for the sake or readability. The same comment applies to later sections where we introduce tie-breaking rules (Section 3.5 and the Online Appendix).

[40] To avoid confusion, we use n for the index of a dimension of the policy space; and i, j and k for the indices of the candidates.

$\arg\min_{(x_1,x_2)\in\bar{X}} |x_1 - x_2|$. That is, each candidate enters at a policy that is the most equally right-wing in both dimensions. Call this game a *policy-motivated election campaign.* It is characterized by a tuple $(\varepsilon, T, \lambda_L, \lambda_R)$.

Suppose a candidate has entered at x. Since the tie is broken in favor of the last mover and there is no Condorcet winner, there exists a closed set $X(i,x)$ such that the remaining candidate i wins if she enters at a policy in $X(i,x)$. Let $y_i(x)$ be the unique minimizer of $|x'_1 - x'_2|$ among all $x' \in \arg\max_{x''\in X(i,x)} u_i(x'')$, that is, it is the policy that candidate i enters.[41]

Proposition 5 *Fix λ_L and λ_R such that $\lambda_L \neq 2\lambda_R$. There exists $\bar{\varepsilon} > 0$ such that, for any $T < \infty$ and $\varepsilon \in (0, \bar{\varepsilon})$, any PBE of the policy-motivated election campaign with $(\varepsilon, T, \lambda_L, \lambda_R)$ satisfies the following: Each candidate enters at $y_i(x)$ as soon as possible, once the other candidate enters at x. If the other candidate has not entered, the following hold:*

1. *Candidate R does not enter at any $-t \in (-\infty, 0]$ for any (λ_L, λ_R).*

2. *Candidate L's strategy depends on the parameters (λ_L, λ_R).*

   (a) *If $\frac{\lambda_L}{\lambda_R} > 2$, then there exists $t_L \in (0,\infty)$ such that L does not enter at $-t \in (-\infty, -t_L)$ and does enter at $\left(\frac{1}{2}, \frac{1}{2}\right)$ for $-t \in (-t_L, 0]$.*

   (b) *If $\frac{\lambda_L}{\lambda_R} < 2$, then there exist $t^*_L, t^{**}_L \in (0,\infty)$ such that L does not enter at $-t \in (-\infty, -t^{**}_L)$, enters at either $\left(\frac{2}{3}, 0\right)$ or $\left(0, \frac{2}{3}\right)$ at $-t \in (-t^{**}_L, -t^*_L)$, and enters at $\left(\frac{1}{2}, \frac{1}{2}\right)$ at $-t \in (-t^*_L, 0]$.*

The proof in the Online Appendix provides an explicit expression of $\bar{\varepsilon}$. The bound ensures that it is a dominated strategy for candidate i to enter at a policy x such that i loses at a policy set profile $(\{x\}, X)$.

On the one hand, since both voters $(1,0)$ and $(0,1)$ prefer candidate R's ideal policy, R wins with probability 1 if no candidate specifies a policy. Moreover, if candidate R enters and then candidate L can enter, R will lose for sure. These facts turn out to imply that candidate R does not have an incentive to enter unless candidate L enters.

On the other hand, candidate L has to enter at some point to receive a positive payoff. If the deadline is very far, then since candidate R will enter with a very high probability once L enters,

[41]The proof of Proposition 5 shows uniqueness of the minimizer.

it is optimal for him not to enter. If the deadline is very close, then the probability that candidate R will enter is very small. Therefore, L enters at the policy he prefers the most among those with which he can win, namely, $\left(\frac{1}{2}, \frac{1}{2}\right)$. In the middle, his optimal policy depends on the relative arrival rates of opportunities. If candidate L is a relatively fast mover ($\frac{\lambda_L}{\lambda_R} > 2$), then the risk of not being able to enter at all is small. Hence, he waits until the probability of candidate R entering after L becomes sufficiently small, and then enters at $\left(\frac{1}{2}, \frac{1}{2}\right)$. If L is relatively slow ($\frac{\lambda_L}{\lambda_R} < 2$), it is too risky for him to wait until the probability of candidate R entering becomes small. Hence, he enters even when there is a significant probability of candidate R entering after L. Taking this event into account, he does not enter at the policy he prefers the most among those with which he can win, but at $\left(\frac{2}{3}, 0\right)$ or $\left(0, \frac{2}{3}\right)$. This narrows down the set of policies with which candidate R can win after L's entry, so L can make R's policy more left-wing.

We note that the consideration in this last part (leading L to entering at $\left(\frac{2}{3}, 0\right)$ or $\left(0, \frac{2}{3}\right)$) does not occur if L does not care about what policy R picks when R wins. For example, candidates may care about the utility from being in the office and the cost of persuading the voters that they implement a policy far from their bliss points, while they do not derive any utility from the implemented policy per se. In the Online Appendix, we formalize such a model, and show that the equilibrium dynamics in such a model are simpler.

Remark 5 (Outcome-equivalence for a public-monitoring model) The PBE we characterize in this section (as well as the PBE characterized in the Online Appendix) is Markov-perfect (except for measure-zero sets of times). Hence this equilibrium is outcome-equivalent to a Markov perfect equilibrium in the "public monitoring" model where candidates observe the other candidate receiving opportunities even when the policy set does not change.[42] Moreover, we solve the equilibrium by backward induction. This means that any SPE under public monitoring is outcome-equivalent to a PBE in our main model where the opponent's opportunities are not observable. The same remark applies to Section 3.4.[43]

Remark 6 (Flexibility in office) We believe that there are various reasons for ambiguous announcements in real election campaigns. It is not our intention to capture all of those reasons in our general model, but to focus on those that relate to candidates' dynamic incentives. In

[42]See Section 4.3 for the formal description of such a model.

[43]In Sections 3.1 and 3.2, those claims are a consequence of Theorem 3 in Section 4.3.

the valence election campaign (Section 3.1) and the symmetric office-motivated election campaign (Section 3.2), ambiguity is present because each candidate does not want to be the first mover. In the policy-motivated election campaign in this section, this effect is still present, while there is another reason to be ambiguous: Not specifying a policy gives a flexibility in choosing a preferred policy after being elected. This same reason will be present in the analysis of the situation with incomplete information in Section 3.5.

3.4 Dynamic Campaign Spending Model

The empirical evidence suggests that campaign spending has nontrivial effects on the election outcome, and candidates often spend monetary resources gradually over time.[44] Since the spending can only increase over time, we can represent such a situation using our policy announcement timing game. We specify $X = \{0, L, H\}$ with $0 < L < H$. The interpretation is that there are two levels of positive campaign spending, where L is the lower level of spending and H is the higher level. The available policy sets are $\mathcal{X}_i = \{X, \{L, H\}, \{H\}\}$ for each $i = A, B$. The interpretation is that $\{L, H\}$ implies that candidate i has spent L by the current time and the total spending at the deadline can be either L (if she does not spend more) or H (if she spends more). To focus on issues regarding campaign spending, assume that the election outcome depends solely on the amount of campaign spending.[45] In particular, for each candidate $i = A, B$, the probability of i winning the election under the policy set profile (X_i, X_j) is $w_i(X_i, X_j) := \frac{\min_{x \in X_i} x}{\left(\min_{x \in X_i} x\right) + \left(\min_{x \in X_j} x\right)}$ with a convention that $\frac{0}{0+0} = \frac{1}{2}$. Note that this implies that we are assuming no depreciation of the effect of campaign spending over time. For simplicity, we assume that the arrival rates of opportunities are symmetric: $\lambda_A = \lambda_B := \lambda$.

Assuming that the remaining fund is useful to a candidate, candidate i's payoff is

$$v_i(X_i, X_j) := \alpha w_i(X_i, X_j) + (1 - \alpha) \left(H - \min_{x \in X_i} x \right),$$

where

$$\alpha \in \left(\max \left\{ \frac{L}{L + \frac{1}{2}}, \frac{H + L}{H + L + \frac{1}{2}} \right\}, 1 \right). \tag{1}$$

[44] See, for example, Gerber et al. (2011).

[45] With this definition, we capture both positive and negative advertising.

Here, $H - \min_{x \in X_i} x$ in the second term of $v_i(X_i, X_j)$ is the remaining cash holdings. The restriction on α implies that α is sufficiently large so that a static best response to $X_j = X$ is $\{L, H\}$ and that to $X_j = \{L, H\}$ or $\{H\}$ is $\{H\}$.[46] We call this game a *dynamic campaign spending game*. It is characterized by a tuple $(H, L, \alpha, T, \lambda)$.

A typical pattern found in the empirical literature (Gerber et al., 2011) is that candidates spend more money near the deadline than far from it. With depreciation, such a pattern follows from a mechanical reason: Earlier spending is less effective, holding the other candidate's spending fixed. We show that, even without depreciation, it is possible that the candidates spend more near the deadline. This follows from a strategic reason: Suppose candidate B has spent L and the deadline is far. If candidate A spends H now, then candidate B will match up with a very high probability. If she spends L, then the race of matching up will start immediately as well, resulting in both spending H with a very high probability. In contrast, if she waits at 0, then the opponent will wait at L as well because both of them know that, once the former spends at least L or the latter spends H, then the race of matching up will start immediately. Hence they can avoid a wasteful competition of matching up with the opponent's high spending until the deadline becomes near.

Depending on how important winning is relative to keeping the money, such an incentive may be alleviated due to the risk that the winning probability will be 0 if candidate A does not have an opportunity to spend more later. Specifically, if α is large, that is, if winning the election is sufficiently important compared to keeping the money, then the risk is prominent and thus they spend H as soon as possible far from the deadline. In contrast, if α is small, then the benefit of avoiding escalation when the deadline is far is sufficiently large. Hence, candidates stay at a spending profile $(L, 0)$ or $(0, L)$ for a long time.

Proposition 6 *Fix the dynamic campaign spending game with* $(H, L, \alpha, T, \lambda)$.

1. *If*
$$\alpha > \frac{H + L}{H + L + \frac{1}{4}} \tag{2}$$
holds, then there exists a PBE. In any PBE, the following hold:

[46] Formally, these conditions are expressed as: $\alpha + (1-\alpha)(H-L) > \max\{\alpha, \alpha\frac{1}{2} + (1-\alpha)H\}$, $\alpha\frac{H}{H+L} > \max\{\alpha\frac{L}{H+L} + (1-\alpha)(H-L), (1-\alpha)H\}$, and $\alpha\frac{1}{2} > \max\{\alpha\frac{L}{H+L} + (1-\alpha)L, (1-\alpha)H\}$, which are equivalent to $\alpha > \max\left\{\frac{L}{L+\frac{1}{2}}, \frac{H+L}{H+L+\frac{1}{2}}\right\}$.

(a) *For each* $-t < -\frac{1}{\frac{\alpha}{2(1-\alpha)(H+L)}-1}\frac{1}{\lambda}$*, each candidate spends* H *after each history.*

(b) *For each* $-t > -\frac{1}{\frac{\alpha}{2(1-\alpha)(H+L)}-1}\frac{1}{\lambda}$*,*[47] *each candidate spends* H *if the other candidate has spent* L *or* H*, and spends* L *otherwise.*

2. *If*

$$\alpha < \frac{H+L}{H+L+\frac{1}{4}} \tag{3}$$

holds, then there exists a PBE. In any PBE, the following hold:

(a) *For each* $-t < -\frac{1}{\lambda}$*, each candidate spends* L *if both candidates have spent* 0*; each candidate does not increase the spending if one candidate has spent* L *and the other has spent* 0*; and each candidate spends* H *otherwise.*

(b) *For each* $-t > -\frac{1}{\lambda}$*, each candidate spends* H *if the other candidate has spent* L *or* H*, and spends* L *otherwise.*

Note that it is straightforward to show that the set of parameters (H, L, α) satisfying both (1) and (2) is nonempty, and the one satisfying both (1) and (3) is also nonempty. When the parameters satisfy (1) and (3), the proposition shows that candidates spend a long time not spending the highest possible amount for the campaign.

3.5 Incomplete Information Model

The general model setup in Section 2 features a complete information game. To show the potential of the model to include a wider class of settings, here we allow for incomplete information. We find that ambiguity still prevails in our incomplete-information game.

Suppose that the set of policies is $X = \mathbb{R}$, and there are two candidates L (he) and R (she). Each of the candidates has two types, Normal or Extreme. The normal type has the ideal policy of 0, while the extreme type has the ideal policy of x^i with $i \in \{L, H\}$, depending on the index of the candidate. We assume $x^L < 0 < x^R$ and $x^L = -x^R$. Let $p \in (0, 1]$ be the probability that a candidate is extreme, and we assume that types are independently distributed between candidates. We extend the perfect Bayesian equilibrium defined in Section 2 in a straightforward way, where the associated belief specifies the belief about the other candidate's type.

[47] Given $\alpha > \frac{H+L}{H+L+\frac{1}{4}}$, we have $\frac{\alpha}{2(1-\alpha)(H+L)} - 1 > 0$.

There is a continuum of voters whose ideal policies are distributed on $\mathbb{R}$, with the median position being 0. Given the ideal policy y and policy x, the voter's utility is $-|x-y|$. A candidate's utility given her ideal policy y and implemented policy x is $-|x-y|$. That is, candidates are purely policy-motivated.

In the policy announcement timing game, we let $\mathcal{X}_i = \{X\} \cup \left(\bigcup_{x \in X} \{\{x\}\}\right)$ for each $i = L, R$, and assume that the arrival rates of opportunities are symmetric: $\lambda_L = \lambda_H = \lambda$. The policy commitment is irreversible and credible: Even if a candidate enters at a platform different from her ideal point, she must implement it. However, if a candidate does not enter, then the voters and the other candidate believe that she will implement her ideal policy. Given this commitment and belief, each voter votes for a candidate who brings the higher expected utility (ties are broken randomly with equal probability). The candidate with a higher vote share will win. The payoff function v_i for each type of each $i = L, R$ is specified accordingly. We break ties in favor of a candidate who enters if only one candidate enters; in favor of the second entrant if both candidates enter; randomly with equal probability in all other cases. We call this dynamic game an *election with incomplete information.* It is characterized by a tuple (p, T, λ). We focus on symmetric PBE in this game.[48]

3.5.1 Benchmark: Complete Information

Before fully analyzing the case with $p < 1$, we analyze the case where both candidates are extreme for sure ($p = 1$). The following proposition says that, with $p = 1$, there is a continuum of equilibria when the horizon is sufficiently long.

To state the result formally, given the opponent's policy x, we define $BR_R(x)$ and $BR_L(x)$ to be candidate R's and L's static best responses, respectively, when they are extreme:

$$BR_R(x) := \begin{cases} x^R & \text{if } |x| > x^R \\ x & \text{if } x \in (0, x^R] \\ -x & \text{if } x \in [-x^R, 0] \end{cases}, \qquad BR_L(x) := \begin{cases} -x^R & \text{if } |x| > x^R \\ -x & \text{if } x \in [0, x^R] \\ x & \text{if } x \in [-x^R, 0) \end{cases}. \tag{4}$$

[48]Some results in this section apply to any (possibly asymmetric) PBE, and for those results we do not restrict ourselves to symmetric PBE in stating them.

Proposition 7 *In any election with incomplete information with $(1,T,\lambda)$, σ is a pure PBE if and only if the following hold under σ.*

1. *If the opponent has not entered, then, the following hold.*

   (a) *If i enters at time $-t \in (-\infty, -\frac{1}{\lambda}\ln 2)$, she enters at 0.*

   (b) *If L enters at time $-t = -\frac{1}{\lambda}\ln 2$, he enters at a policy in $[x^L, 0]$. If R enters at time $-t = -\frac{1}{\lambda}\ln 2$, she enters at a policy in $[0, x^R]$.*

   (c) *Each candidate $i \in \{L,R\}$ enters at x^i for $-t \in (-\frac{1}{\lambda}\ln 2, 0]$.*

2. *If the opponent has entered at x, then each candidate $i \in \{L,R\}$ enters at $BR_i(x)$ as soon as possible.*

Intuitively, if the deadline is sufficiently far, it is likely that the opponent will have an opportunity later. Hence, if candidate R enters at $x \geq 0$ then $-x \leq 0$ will be implemented with a high probability, and if L enters at $x \leq 0$ then $-x \geq 0$ will be implemented with a high probability. Thus, it is better for each candidate to enter at 0 if she ever enters.

If she skips an opportunity, then by symmetry and the constant-sum nature of payoff functions, the expected payoff is the same between the candidates. Since each candidate enters at some policy in $\left[-x^R, x^R\right]$ when entering (if a candidate enters outside of this interval, then she will certainly lose since the median voter will prefer the opponent's ideal policy), this symmetry, together with piecewise-linearity of the utility function, means that each candidate's expected payoff is the same as the one from entering at 0.

In total, each candidate is indifferent between entering at 0 and not entering when the deadline is far. When the deadline is close, since it is likely that the opponent cannot enter afterward, it is optimal to enter at her own ideal policy. The cutoff time turns out to be $-\frac{1}{\lambda}\ln 2$.

As will be seen, in the model where $p \in (0,1)$, for sufficiently large T, the extreme candidates do not enter until $-t$ near $-\frac{1}{\lambda}\ln 2$. Intuitively, there is an option value of not entering and figuring out the opponent's type. More precisely, the normal type enters at 0 as soon as possible, and hence waiting allows a candidate to learn about the opponent's type.

3.5.2 Strategy of the Normal Type

The next lemma formally pins down the strategy of the normal type:

Lemma 2 *In any PBE of the election with incomplete information with* (p, T, λ) *with* $p \in (0,1)$*, each normal-type candidate enters at* 0 *as soon as possible at any history.*

The intuition is that, if a normal-type candidate enters at 0, then with probability 1 the winning policy is 0, which is her ideal policy.

3.5.3 Strategy of the Extreme Candidate

Since the candidates are symmetric and we focus on symmetric PBE, without loss, we consider candidate R's incentive.

We first analyze what each extreme candidate does once the opponent enters. Given the definition of the best response, in any PBE, once the opponent enters at x, the extreme type of each candidate $i \in \{L, R\}$ enters at $BR_i(x)$ as soon as possible.

Hence, given an arbitrary conditional probability $\tilde{p}$ of candidate L being extreme, the expected payoff of extreme candidate R entering at x at $-t$ when candidate L has not entered and is extreme depends only on $(t, x, \tilde{p})$ in any PBE. Let $v_t(\tilde{p}, x)$ be this expected payoff of extreme candidate R and let $v_t(\tilde{p}, \text{enter}) = \max_x v_t(\tilde{p}, x)$ be the expected payoff of entering.[49] The next lemma characterizes the optimal policy to enter:

Lemma 3 *For each* $\tilde{p} \in (0,1)$ *and* $t \geq 0$*, we have*

$$\arg\max_x v_t(\tilde{p}, x) = \begin{cases} \{\tilde{p}x^R\} & \textit{if } \frac{e^{-\lambda t}}{1-e^{-\lambda t}} > \tilde{p} \\ [0, \tilde{p}x^R] & \textit{if } \frac{e^{-\lambda t}}{1-e^{-\lambda t}} = \tilde{p} \\ \{0\} & \textit{if } \frac{e^{-\lambda t}}{1-e^{-\lambda t}} < \tilde{p} \end{cases}.^{50}$$

Intuitively, if the deadline is far (that is, $\frac{e^{-\lambda t}}{1-e^{-\lambda t}} < \tilde{p}$), then it is likely that the opponent L will have an opportunity to enter and flip the policy if he is extreme: If extreme candidate R enters at x, then $-x$ will be implemented with a high probability if the opponent is extreme. Hence it is optimal to enter at 0 (if she ever enters). On the other hand, if the deadline is near (that is, $\frac{e^{-\lambda t}}{1-e^{-\lambda t}} > \tilde{p}$), then it is unlikely that the opponent will have an opportunity. Hence she enters

[49] As will be seen, the maximizer always exists.
[50] We use the convention that $\frac{1}{0} = +\infty$ (which applies when $t = 0$).

at a policy close to her ideal policy. The value of entering, $v_t(\tilde{p}, \text{enter}) = \max_x v_t(\tilde{p}, x)$, can be computed by using this lemma.[51]

Since we have pinned down the strategy of the normal type and the continuation strategy of the extreme type after the opponent has entered, we are left to specify extreme candidate R's strategy at the histories where candidate L has not entered. Let $\bar{H}_t^R$ be the set of candidate R's histories such that no candidate has entered by $-t$.

Fix any symmetric PBE σ. Given the history $h_R^t \in \bar{H}_t^R$, let $p\left(h_R^t\right)$ be the posterior probability that candidate L is extreme. Since candidate R's opportunity and candidate L's opportunity are independently distributed and it is possible that candidate L has not obtained any opportunity by the current time, the posterior $p\left(h_R^t\right)$ depends only on the public history—the event that neither candidate has entered by $-t$. Hence, we write it as a function of t, by setting $p\left(h_R^t\right) = p(t)$. Moreover, candidate R and the voters share the same posterior about candidate L being extreme. Thus, $p(t)$ is the voters' posterior about L's type as well.

The next lemma states that candidate R does not enter at $-t$ with $\frac{e^{-\lambda t}}{1-e^{-\lambda t}} \leq p(t)$:

Lemma 4 *In any election with incomplete information with $(p.T, \lambda)$ with $p \in (0,1)$, under any symmetric PBE, for each time $-t$ and $h_R^t \in \bar{H}_R^t$, candidate R does not enter if $\frac{e^{-\lambda t}}{1-e^{-\lambda t}} \leq p(t)$.*

To see the intuition, suppose candidate R has an opportunity at $-t$. Since candidate L and the voters cannot observe if candidate R received an opportunity, if candidate R does not enter, then the situation is the same as the case in which no candidate receives an opportunity at $-t$. Since we focus on symmetric equilibria and the implemented policy is in $\left[-x^R, x^R\right]$, if the opponent is extreme, then candidate R obtains a payoff of $-x^R$ (corresponding to policy 0) if she does not enter. If the opponent is normal, then she obtains a payoff greater than $-x^R$ by not entering. This is because, since the opponent enters at 0, she can obtain a payoff that is at least $-x^R$ at any history; and if no candidate receives further opportunities, then she obtains a payoff of $-\frac{1}{2}x^R$. In contrast, Lemma 3 implies that, if candidate R enters at $-t$ with $\frac{e^{-\lambda t}}{1-e^{-\lambda t}} \leq p(t)$, she obtains $-x^R$. Hence, it is uniquely optimal for her not to enter.

Given this lemma, we now characterize the equilibrium dynamics. Fix any prior $p > 0$. If T is sufficiently large, then we have $\frac{e^{-\lambda t}}{1-e^{-\lambda t}} \leq p \leq p(t)$ for all $-t \in \left[-\frac{T}{2}, -T\right]$. Hence, the extreme candidates do not enter for any $-t \in \left[-\frac{T}{2}, -T\right]$.

[51]The calculation is provided in the proof of Lemma 3.

Since normal types enter as soon as they obtain an opportunity and extreme types do not enter in any symmetric PBE, by a standard argument $p(t)$ can be shown to evolve as follows for $-t \in \left[-\frac{T}{2}, -T\right]$:

$$\frac{d}{dt}p(t) = -\lambda p(t)(1-p(t)).$$

Given this evolution, for large T, we have $p\left(\frac{T}{2}\right) \approx 1$. Since the normal candidates always enter, we have $p(t) \geq p\left(\frac{T}{2}\right) \approx 1$ for all $-t \geq -\frac{T}{2}$.

At $-t = -\frac{T}{2}$, candidate R does not enter. In contrast, for $-t$ sufficiently close to 0, she enters. To see this, consider the following two scenarios: If she enters at $p(t)x^R \approx x^R$, she will win and obtain a payoff near 0 if and only if candidate L cannot enter, which happens with probability close to 1. If candidate R does not enter, then again with a probability close to 1, no candidate obtains a further opportunity. Hence her payoff is near $-x^R$. Thus, for $-t$ sufficiently close to 0, it is optimal for R to enter at $p(t)x^R$.

In fact, we can show that there exists a unique cutoff time $-t^*$ such that candidate R enters at $p(t)x^R$ if $-t > -t^*$ and does not enter if $-t < -t^*$, where $p(t)$ evolves according to $\frac{d}{dt}p(t) = -\lambda p(t)(1-p(t))$ for $-t \in [-T, -t^*]$ (the extreme candidate does not enter for $-t \in [-T, -t^*)$) and $p(t) = p(t^*)$ for $-t \geq -t^*$ (the extreme candidate enters for $-t \in (-t^*, 0]$ and so there is no belief update). Moreover, we can show that t^* converges to the maximum cutoff for the complete-information case as the horizon becomes long ($T \to \infty$), providing one possible refinement of the set of PBE in the case of complete information (cf. Proposition 7). We summarize our results as follows.

Proposition 8 *For each $p \in (0,1)$ and $\lambda > 0$, there exists $\bar{T}_{p,\lambda} < \infty$ such that, for each $T \geq \bar{T}_{p,\lambda}$, in any election with incomplete information with (p, λ, T), there exists a symmetric PBE, and there exists $t^*(p, \lambda, T)$ such that any symmetric PBE satisfies the following equilibrium dynamics: For each $-t < -t^*(p, \lambda, T)$, extreme candidates do not enter and $p(t)$ evolves according to $\frac{d}{dt}p(t) = -\lambda p(t)(1-p(t))$ over $-t = [-T, -t^*]$; and for each $-t > -t^*(p, \lambda, T)$, extreme candidate $i = L, R$ enters at $p(t^*(p, \lambda, T))x^i$ and $p(t) = p(t^*(p, \lambda, T))$. Moreover, for any $p \in (0,1)$ and $\lambda > 0$,*

$-\lambda t^*(p,\lambda,T)$ converges to $-\ln 2$ as $T \to \infty$:

$$\lim_{T\to\infty} |\lambda t^*(p,\lambda,T) - \ln 2| = 0.$$

The proposition shows that candidates use ambiguous language at times before $-t^*(p,\lambda,T)$ in any symmetric PBE. This implies that extreme candidates spend most of the campaign time keeping their policies ambiguous, provided the campaign is sufficiently long.

4 General Predictions

In Section 3, we have seen that the policy announcement timing game can be applied to analyses of various examples. In those examples, we showed results that match observations in real election campaigns (cf. discussions in the Introduction). Now we present general principles that underlie those results. This helps us understand the logic behind various results in Section 3, as well as shows the robustness of those results to wider classes of environments.

To recap, our discussion of the applications have the following in common: Candidates use ambiguous language (or do not use up all the campaign funds) when the election date is not close if entering before the opponent is disadvantageous, while they enter as soon as possible if a Condorcet winner exists. Moreover, we obtained uniqueness of the entry times in many results, and in particular, we obtain uniqueness in the models in which candidates are purely office-motivated. In this section, we aim to generalize those results.

In Section 4.1, we offer a general condition for candidates to use ambiguous language. The key condition is what we call the "first-mover disadvantage," which roughly corresponds to the non-existence of a Condorcet winner. In contrast, Section 4.2 shows that if there is a Condorcet winner, then candidates announce the policy corresponding to the Condorcet winner as soon as possible. Finally, Section 4.3 offers a general implication of the candidates being purely office-motivated.

For each application, Table 2 represents which general prediction is applicable to which application. In some applications, the corresponding general theorem only applies to part of the claims made there. In the subsections that follow, we explain which part of each of those applications is covered by each general theorem.

Model	Long ambiguity	Dynamic median-voter	Constant-sum Markov
Section 3.1: Valence Candidates	Yes	Yes	Yes
Section 3.2: Multi-dimensional policy space, purely office-motivated	Yes	Yes	Yes
Section 3.3: Multi-dimensional policy space, policy-motivated	No	No	No
Section 3.4: Spending	No	Yes	No
Section 3.5: Incomplete information	No	No	Yes

Table 2: General Predictions and Applications: "Yes" means that the corresponding result is used in a proof for the corresponding section, while "No" means it is not.

In Sections 4.1 and 4.2, we assume that, for each candidate $i = A, B$,

$$\mathcal{X}_i = \{\{x\} | x \in X\} \cup \{X\}. \tag{5}$$

That is, the choice of a policy set is either to specify a single policy or not to specify any policy at all.

4.1 The Long Ambiguity Theorem

In this section, we are going to prove the following claim:

Long Ambiguity Theorem: *Under certain conditions, for each candidate i, there exists t_i such that i does not enter for any $-t \in [-T, -t_i)$.*

The actual statement of this result is rather complicated and thus is provided at the end of this section after a presentation of the analysis (Section 4.1.4).

To start the analysis, let (x, X) denote the set of histories at which candidate A has entered at x and candidate B has not entered. Other sets of histories are denoted in an analogous manner. Abuse notation to write "x_i" to mean $\{x_i\}$ as part of the argument of v_i. For each $x_i \in X$, let $BR_j(x_i)$ be the set of candidate j's best responses against candidate i's policy x_i:

$$BR_j(x_i) = \arg\max_{X_j \in \mathcal{X}_j} v_j(X_j, x_i),$$

and suppose that it is non-empty. To simplify the notation, we sometimes write $x_j \in BR_j(x_i)$ to mean $\{x_j\} \in BR_j(x_i)$.

We say that $X_i^* \subseteq X$ is candidate i's **optimal set** if the following hold for each $x_i^* \in X_i^*$.

1. $v_i^{BR_j} := \sup_{x_j \in BR_j(x_i^*)} v_i(x_i^*, x_j) \geq \sup_{x_i \notin X_i^*, x_j \in BR_j(x_i)} v_i(x_i, x_j)$.

2. $v_i(x_i^*, X) = \sup_{x_i \in \mathcal{X}_i} v_i(x_i, X)$.

Note that the equality in part 1 holds if (v_i, v_j) is constant-sum, which happens for example if candidates are purely office-motivated. For general (v_i, v_j), it is straightforward that the definition of the optimal set ensures that there exits a unique largest optimal set (the optimal set that is a superset of all other optimal sets). Hereafter, let X_i^* be the largest optimal set for candidate i.[52]

Assumption 1 *For each candidate i, the largest optimal set X_i^* is non-empty, and satisfies the following properties.*

1. *For any $x_i^* \in X^*$ and $x_j, x_j' \in BR_j(x_i^*)$, $v_i(x_i^*, x_j) = v_i\left(x_i^*, x_j'\right)$ holds.*

2. *For any $x_i^*, x_i^{**} \in X_i^*$, $v_j(x_i^*, X) = v_j(x_i^{**}, X)$ and $\max_{\{x_j\} \in \mathcal{X}_j} v_j(x_i^*, x_j) = \max_{\{x_j\} \in \mathcal{X}_j} v_j(x_i^{**}, x_j)$.*

Assumption 1 ensures that X_i^* is non-empty. Note that once i enters at x_i, j enters at some $x_j \in BR_j(x_i)$ in any PBE. Hence, conditional on any history such that i's opponent has not entered, if i enters, then she enters at some $x_i^* \in X_i^*$. In addition, i's expected payoff when she enters is uniquely pinned down. Moreover, fixing j's strategy, if i enters, then j's payoff is also pinned down uniquely. Assumption 1 thus implies that any $x_i \in X_i^*$ gives the same continuation payoff to both candidates i and j in any PBE.

Assumption 2 *For each candidate i and any $x_i^* \in X_i^*$, $v_i(x_i^*, X) \geq v_i^{BR_j}$.*

This assumption implies that, after i's entry, i cannot be better off by the opponent's subsequent entry. Define:

$$v_{i,t}(\text{enter}) = e^{-\lambda_j t} v_i(x_i^*, X) + \left(1 - e^{-\lambda_j t}\right) v_i^{BR_j}.$$

Assumption 1 implies that this is candidate i's expected payoff at time $-t$ when she enters (in any PBE), and Assumption 2 implies that $v_{i,t}(\text{enter})$ is weakly decreasing in t.

We consider the following three cases, depending on the incentives at the deadline.

[52] For each $x_i \notin X_i^*$, we have either "$v_i^{BR_j} > v_i(x_i, x_j)$ for all $x_j \in BR_j(x_i)$" or "$v_i(x_i^*, X) > v_i(x_i, X)$."

- Case 1: $v_i(X,X) > v_i(x_i^*,X)$ for each i.

- Case 2: $v_i(X,X) < v_i(x_i^*,X)$ for each i.

  - More generally, there exist $t_0 \geq 0$ and a number $v_{i,t_0}(X,X)$ such that the continuation payoff at time $-t_0$ given any history in (X,X) is equal to $v_{i,t_0}(X,X)$ in any PBE, and that $v_{i,t_0}(\text{enter}) > v_{i,t_0}(X,X)$ holds for each i.[53]

- Case 3: $v_A(X,X) > v_A(x_A^*,X)$ and $v_B(X,X) < v_B(x_B^*,X)$.[54]

4.1.1 Case 1: No Candidate Enters at the Deadline

In this case, uniqueness and long ambiguity hold without additional assumptions, as follows.

Proposition 9 *Consider Case 1. Under Assumptions 1 and 2, there exists a PBE. In any PBE, at histories in (X,X), candidate i does not enter at any $-t \in (-\infty, 0]$.*

The intuition is simple: Candidate i's entry at time $-t$ results in either $v_i(x_i^*,X)$ if the opponent j does not enter afterward, or $v_i^{BR_j}$ if j does. Given that no candidate enters at histories in (X,X) after time $-t$, the former payoff is lower than the payoff from not entering, $v_i(X,X)$, by the definition of Case 1, and the latter is weakly lower due to Assumption 2.

4.1.2 Case 2: Both Candidates Enter at the Deadline

Fix t_0 that defines Case 2. For $t > t_0$, define $\bar{v}_{i,t}(\text{not})$ as candidate i's expected continuation payoff at time $-t$ when she does not enter, assuming that each candidate will enter at times in $(-t,-t_0)$ upon receiving an opportunity. Such a payoff is well defined due to Assumption 1.[55] Let

$$t_i^* \equiv \inf\{t > t_0 : \bar{v}_{i,t}(\text{not}) \geq v_{i,t}(\text{enter})\}.$$

Given the continuity of the continuation payoffs in time, $-t_i^*$ is the time closest to the deadline at which candidate i is indifferent between entering and not entering.

[53] "$v_i(X,X) < v_i(x_i^*,X)$ for each i" corresponds to taking $t_0 = 0$.

[54] The case with $v_A(X,X) < v_A(x_A^*,X)$ and $v_B(X,X) > v_B(x_B^*,X)$ is symmetric.

[55] The formal expression of this payoff is complicated, so we relegate it to the Online Appendix.

Assumption 3 (Genericity) *At least one of the following holds:* $v_i^{BR_j} < \sup_{\{x_i\} \in \mathcal{X}_i} v_i (x_i, X)$ *for each* i, *or* $t_A^* \neq t_B^*$, *or* $t_A^* = t_B^* = \infty$.

This assumption is a genericity assumption in the sense that the environment in which it is violated constitutes a degenerate (non-full-dimensional) space in the space of payoff functions.[56]

Proposition 10 *Consider Case 2. Under Assumptions 1, 2, and 3, there exists a PBE. There exists a profile* $(t_A, t_B) \in (\mathbb{R}_{++} \cup \{\infty\})^2$ *such that, for any PBE, at any histories in* (X, X), *candidate* i *does not enter at any* $-t \in (-\infty, -t_i)$, *and enters at every time* $-t \in (-t_i, 0]$. *Moreover, if* $t_i^* \leq t_j^*$, *then* $t_i \leq t_j$ *and* $t_i = t_i^*$.

If $t_i < \infty$, then candidate i does not enter when the deadline is sufficiently far. The following condition, which is stronger than the condition for candidate i in Assumption 2, is a sufficient condition for $t_i < \infty$:

$$\text{First-mover disadvantage for } i \quad \begin{cases} v_i\left(X, x_j^*\right), v_i\left(x_i^*, X\right), v_i\left(X, X\right) \geq v_i^{BR_j} \\ \sup_{x_i \in X} v_i\left(x_i, x_j^*\right) > v_i^{BR_j} \end{cases}. \tag{6}$$

The second line of this condition states that, if both candidates have to enter and the order of the moves is known, then being the first mover is worse than being the second mover. The first line further requires that the disadvantage of being the first mover is so large, that it is the worst option even if we include the possibility of some candidates not specifying a policy. Intuitively, when it is the worst for candidate i to be best-responded by her opponent, i has little incentive to enter when the election day is far away. This is because when the election day is far away, the probability of the opponent best-responding in the future is high. In Section 4.1.4, we explain that this condition holds in various applications.

Proposition 11 *For each* i, *Proposition 10 holds with* $t_i < \infty$ *if we additionally require first-mover disadvantage for* i *to hold.*[57]

[56] To see why $t_A^* \neq t_B^*$ or $t_A^* = t_B^* = \infty$ holds generically, notice that, for each $i = A, B$ and $t < \infty$, $\bar{v}_{i,t}(\text{enter})$ is independent of $v_i(X, X)$, while $\bar{v}_{i,t}(\text{not})$ is strictly increasing in it. Hence, if there exists $w \in \mathbb{R}$ such that $t_A^* = t_B^* < \infty$ holds for some payoff function (v_A, v_B) such that $v_A(X, X) = w$, then $t_A^* \neq t_B^*$ holds for any payoff function that is the same as (v_A, v_B) except that $v_A(X, X) \neq w$.

[57] This result is not inconsistent with the case with $t_A^* = t_B^* = \infty$ which is allowed in Assumption 3 because the proof shows that if first-mover disadvantage for i holds then $t_i^* < \infty$.

4.1.3 Case 3: Only One Candidate Enters at the Deadline

We define $\bar{v}^A_{i,t}(\text{not})$ as candidate i's expected payoff at time $-t$ when she does not enter, assuming that only candidate B will enter at times in $(-t, 0]$ upon receiving an opportunity.[58] Such a payoff is well defined due to Assumption 1.

Let

$$\begin{aligned}\hat{t}_A &\equiv \inf\left\{t > 0 : \bar{v}^A_{A,t}(\text{not}) \leq v_{A,t}(\text{enter})\right\};\\ \hat{t}_B &\equiv \inf\left\{t > 0 : \bar{v}^A_{B,t}(\text{not}) \geq v_{B,t}(\text{enter})\right\}.\end{aligned}$$

Given the continuity of the continuation payoffs in time, $\hat{t}_i$ is the time closest to the deadline at which i is indifferent between entering and not entering, respectively, assuming that only candidate B will enter afterward.

Assumption 4 (Genericity) *$\hat{t}_A \neq \hat{t}_B$ or $\hat{t}_A = \hat{t}_B = \infty$ holds.*

Like Assumption 3, this assumption is again a genericity assumption. If $\hat{t}_A = \hat{t}_B = \infty$, then for each time $-t$ in any PBE, candidate A does not enter and candidate B enters. Hence we focus on the case in which $\hat{t}_A \neq \hat{t}_B$.

Proposition 12 *Consider Case 3. Under Assumptions 1, 2, and 4, there exists a PBE, and the following hold.*

1. *If $\hat{t}_A < \hat{t}_B$, then there exists $\bar{\varepsilon} > 0$ such that for all $\varepsilon \in (0, \bar{\varepsilon})$, in any PBE σ and its associated belief β, at any history at time $-(\hat{t}_A + \varepsilon)$ in (X, X), each candidate strictly prefers to enter under the continuation strategy given by σ and the belief β.*

2. *If $\hat{t}_A > \hat{t}_B$, then for any PBE, at any history in (X, X), no candidate enters at any $-t \in \left(-\infty, -\hat{t}_B\right)$.*

If $\hat{t}_A < \hat{t}_B$, we can use Proposition 10 in Case 2 to characterize any PBE, with a substitution that time t_0 is set to be equal to $\hat{t}_A + \varepsilon$ where $\varepsilon > 0$ is sufficiently small (and with an additional requirement of a genericity assumption (Assumption 3)). If $\hat{t}_A > \hat{t}_B$, in contrast, no candidate enters at any $-t \in \left(-\infty, -\hat{t}_B\right)$.

[58] The superscript denotes the candidate who does not enter close to the deadline in Case 3.

Corresponding to (6), define:

$$\text{Strong first-mover disadvantage for } i \begin{cases} v_i\left(x_i^*, X\right) > v_i^{BR_j} \\ (6) \text{ holds if } \hat{t}_A < \hat{t}_B \end{cases}.$$

Putting the two parts of Proposition 12 together, we can show the following result:[59]

Proposition 13 *Consider Case 3 and suppose that strong first-mover disadvantage for i holds. Under Assumptions 1, 2, and 4, there exists a PBE, and for any PBE, there exists $t_i < \infty$ such that candidate i does not enter at any $-t \in (-\infty, -t_i)$.*

4.1.4 Summary

We are now ready to state our first general prediction:

Theorem 1 (Long Ambiguity) *Under Assumptions 1 and 2, the following claims are true.*

1. *Suppose $v_i(X, X) > v_i(x_i^*, X)$ for each i. Then, there exists a PBE, and in any PBE, candidate i does not enter at any history in (X, X) at any $-t \in (-\infty, 0]$.*

2. *Suppose $v_i(X, X) < v_i(x_i^*, X)$ for each i. Then, with additionally requiring Assumption 3, there exists a PBE. Moreover, if first-mover disadvantage for i holds, then there exists $t_i < \infty$ such that, for any PBE, candidate i does not enter at any history in (X, X) at any $-t \in (-\infty, -t_i)$.*

3. *Suppose that $v_i(X, X) > v_i(x_i^*, X)$ and $v_j(X, X) < v_j\left(x_j^*, X\right)$ for $i \neq j$. Then, with additionally requiring Assumption 4, there exists a PBE. Moreover, fix an arbitrary $k \in \{i, j\}$ and suppose that $\hat{t}_i > \hat{t}_j$ or strong first-mover disadvantage for candidate k holds. Then, there exists $t_k < \infty$ such that, for any PBE, candidate k does not enter at any history in (X, X) at any $-t \in (-\infty, -t_k)$.*

Although the conditions referred to in the theorem involve evaluation of variables that are endogenously determined in equilibrium (such as $\hat{t}_i$), they are fairly easy to check. For example, in the valence election campaign, the environment of Proposition 2 corresponds to part 3 of Theorem

[59] The proof shows that, if strong first-mover disadvantage for B holds, then $\hat{t}_B < \infty$ must hold.

1, where $i = W$ and $j = S$. It satisfies Assumptions 1, 2, and 4, and first-mover disadvantage for W. In addition, for Proposition 3, any symmetric office-motivated election campaign model with $(X, \mu) \notin \mathcal{M}$ satisfies Assumptions 1, 2, 3, and first-mover disadvantage for each candidate. Hence, part 2 of Theorem 1 applies.[60]

Since some of the assumptions in the above theorem are genericity conditions, we can also restate part of the theorem in a way that is easier to interpret, as follows.

Corollary 2 *Under Assumptions 1 and 2, the following claims are true.*

1. *Suppose $v_i(X, X) > v_i(x_i^*, X)$ for each i. Then, there exists a PBE. In any PBE, candidate i does not enter at any history in (X, X) at any $-t \in (-\infty, 0]$.*

2. *Suppose $v_i(X, X) < v_i(x_i^*, X)$ for each i. Then, generically in the space of payoff functions, the following holds. There exists a PBE, and if first-mover disadvantage for i holds, then there exists $t_i < \infty$ such that, for any PBE, candidate i does not enter at any history in (X, X) at any $-t \in (-\infty, -t_i)$.*

Note that the corollary states that we expect long ambiguity in many cases, but does not identify conditions under which we expect it. Theorem 1, in contrast, pins down the sufficient condition for when we expect long ambiguity.

4.2 The Dynamic Median-Voter Theorem

In this section, we consider an extension of the median voter theorem, which has an implication on several of our examples in Section 3. To this end, we focus on symmetric elections, that is, given any $\bar{X}, \bar{X}' \in \mathcal{X}_i = \mathcal{X}_j$, $v_i(\bar{X}, \bar{X}') = v_j(\bar{X}, \bar{X}')$.[61] A policy $x^* \in X$ is a **Condorcet winner** if, for each i, $x^* \in BR_i(x^*)$, $x^* \in BR_i(X)$, $v_i(x^*, X) > v_i(X, X)$, and for each $X_i \neq x^*$, $v_i(x^*, x_j) > v_i(X_i, x_j')$ for some $x_j \in BR_j(x^*)$ and $x_j' \in BR_j(X_i)$.[62] For example, in a uni-dimensional Downsian model

[60]One might think that part 2 of Theorem 1 can be applied to the analysis of the policy-motivated election campaign. However, Assumption 1 fails because the optimal set is empty in that example. Specifically, for candidate L, the intersection of the set of best responses to X, $\{(\frac{1}{2}, \frac{1}{2})\}$, and the set of best responses assuming the opponent's subsequent best response, $\{(\frac{2}{3}, 0), (0, \frac{2}{3})\}$, is disjoint. The example hence demonstrates that even outside the environment in which our assumptions hold, long ambiguity can be an equilibrium phenomenon, showing the robustness of the result.

[61]An extension to non-symmetric elections is straightforward, where one would define a profile of policies as Condorcet winners. Analogous results to the ones we present here as well as in Appendix D would go through.

[62]A policy is in the largest optimal set if it is a Condorcet winner and the stage game is constant sum.

in which (i) a candidate wins with probability one if the vote share is strictly greater than 1/2 and with probability 1/2 if the vote share is 1/2, and (ii) entering at the median voter ensures winning when the opponent does not enter (for example, the voters are risk averse and think that there is uncertainty about what policy a candidate announcing X would implement), the policy corresponding to the median voter is the unique Condorcet winner. In addition, the policy 1 with $\delta = 0$ in Proposition 1 and the policy $x^*(X, \mu)$ in the symmetric office-motivated election campaign with $(X, \mu) \in \mathcal{M}$ (part 1 of Proposition 3) are Condorcet winners. Note that, by definition, there is at most one Condorcet winner.

The following theorem extends the median voter theorem to a dynamic environment.

Theorem 2 (Dynamic Median-Voter) *Suppose that X is finite, (v_A, v_B) is a symmetric constant-sum game, and there exists a Condorcet winner. Then, there exists a PBE, and in any PBE, at any time $-t$, conditional on any history, (i) if the opponent's action is X, each candidate i announces the Condorcet winner and (ii) otherwise, each candidate i best-responds to the opponent's current policy.*

The theorem can be applied to prove that, in the examples mentioned above, candidates enter at the Condorcet winner specified above as soon as possible.

To see the intuition, note first that if a candidate obtains an opportunity at the deadline, then the assumption on the payoff function implies that she enters at the Condorcet winner. To show that this holds for all time $-t$, we resort to the continuous-time backward induction (formally presented in Appendix B), which in particular shows that it is not possible in any PBE that candidates keep using ambiguous language for a long time and try to enter to win at the last moment.

In Appendix D, we generalize the theorem to cover the case with non-constant-sum games. We show the existence of a PBE in which each candidate announces the Condorcet winner. We also show the uniqueness of a PBE when we further require that there is a policy that is strictly dominant for each i, while not requiring condition (5). This last result in particular implies that, in the dynamic campaign spending game, in subgames in which each candidate has already spent L, each candidate spends H as soon as possible.

4.3 The Constant-Sum Markov Theorem

In some of the examples we consider in Section 3, candidates are purely office-motivated, and thus their utility functions are constant-sum since the winning probabilities must add up to one. In this section, we provide a characterization of the equilibrium dynamics for constant-sum elections by showing that, in constant-sum elections, candidates' continuation payoffs at any history is determined only by the remaining time and the current policy set profile. Moreover, we show that it is irrelevant whether each candidate observes the arrival of the opponent's opportunities. More specifically, as specified in Section 2, we assume throughout the paper that each candidate cannot observe the arrivals of opportunities to the opponent but only the changes of the policy set. We compare such a setting with the model in which each candidate can observe the arrivals of the opponent's opportunities, including those that do not involve changes in the policy set. We call the former and the latter setups "private monitoring" and "public monitoring," respectively.

To define the setup of "public monitoring" formally, let $h^t = \left(\left(t_A^k, X_A^k\right)_{k=1}^{k_A}, \left(t_B^k, X_B^k\right)_{k=1}^{k_B}, t\right)$ be the entire history at $-t$, where $-t_j^k < -t$ is the time at which candidate j receives his or her k'th revision opportunity; X_j^k is the policy set that j has chosen at time $-t_j^k$; and t denotes the current remaining time. Let H be the set of all histories. We say that a history for candidate i at time $-t$, denoted h_i^t, is consistent with h^t if the former is given by deleting information about j's opportunities at which j did not change the policy set. Let $\theta\left(h^t\right) = \left(X_A^{k_A}, X_B^{k_B}\right)$ be the most recent policy profile at time $-t$; and $\theta_i\left(h^t\right) = X_i^{k_i}$ be candidate i's most recent policy at $-t$. Note that $\theta\left(h^t\right) = \theta\left(h_i^t\right)$ for each i and t. We allow the available policy set to depend on the current policy sets. Formally, let $\mathcal{X}_i(\theta) \subseteq 2^{\theta_i} \setminus \{\emptyset\}$ be the collection of available policy sets under $\theta \in \mathcal{X}_A \times \mathcal{X}_B$. Candidate i's strategy is a map $\sigma_i : H \to \cup_{\theta \in \mathcal{X}_A \times \mathcal{X}_B} \Delta\left(\mathcal{X}_i(\theta)\right)$, with a restriction that $\sigma_i\left(h^t\right) \in \Delta\left(\mathcal{X}_i\left(\theta\left(h^t\right)\right)\right)$. Let Σ_i be the space of i's strategies, and $\Sigma = \Sigma_A \times \Sigma_B$. With σ_i, candidate i takes $\sigma_i\left(h^t\right)$ if she has an opportunity at time $-t$ and takes $\theta_i\left(h^t\right)$ otherwise. A subgame-perfect equilibrium (SPE) can be defined in the standard manner. We call this setup "public monitoring."

In the "private monitoring" setup, the definition of PBE naturally extends to the case in which each candidate i's feasible announcements depend on the current announcement through the $\mathcal{X}_i(\cdot)$ function. Recall that a strategy profile (σ_A^*, σ_B^*) is a PBE if there exists a belief β such that, for each $i \in \{A, B\}$, (i) $\sigma_i^* \in \arg\max_{\sigma_i \in \Sigma_i} u_i^\beta(\sigma_i, \sigma_j^* | h_i^t)$ holds for every $h_i^t \in H_i$ and (ii) β is derived from Bayes rule whenever possible.

First, take an arbitrary PBE σ in private monitoring, and let $w_t^i\left(\sigma, h_i^t, X_i\right)$ be candidate i's continuation payoff of taking $X_i \in \mathcal{X}_i\left(\theta\left(h_i^t\right)\right)$ when her private history is h_i^t and she receives an opportunity at $-t$. Similarly, let $\hat{w}_t^i\left(\sigma, h_i^t, X_j\right)$ be candidate i's continuation payoff when her private history is h_i^t and candidate j receives an opportunity and takes $X_j \in \mathcal{X}_j\left(\theta\left(h_j^t\right)\right)$; and let $w_t^i\left(\sigma, h_i^t, no\right)$ be candidate i's continuation payoff when her private history is h_i^t and no candidate receives an opportunity at $-t$.

Second, take an arbitrary SPE $\bar{\sigma}$ in public monitoring, and let $W_t^i\left(\bar{\sigma}, h^t, X_i\right)$ be candidate i's continuation payoff of taking policy $X_i \in \mathcal{X}_i\left(\theta\left(h^t\right)\right)$ when the public history is h^t and she receives an opportunity at $-t$. Similarly, let $\hat{W}_t^i\left(\bar{\sigma}, h^t, X_j\right)$ be candidate i's continuation payoff when the public history is h^t and candidate j receives an opportunity and takes $X_j \in \mathcal{X}_j\left(\theta\left(h^t\right)\right)$; and let $W_t^i\left(\bar{\sigma}, h^t, no\right)$ be candidate i's continuation payoff when the public history is h^t and no candidate receives an opportunity at $-t$.

We can show that the continuation payoff of choosing policy X_i and not receiving an opportunity depends only on the current time $-t$ and the current policy set of the opponent $\theta_j\left(h^t\right)$ (recall that $\theta_j\left(h^t\right) = \theta_j\left(h_i^t\right) = \theta_j\left(h_j^t\right)$).

Theorem 3 (Constant-Sum Markov) *Suppose $v_A(X_A, X_B) + v_B(X_A, X_B) = 1$ for each $(X_A, X_B) \in \mathcal{X}_A \times \mathcal{X}_B$. Then, there exists $v_{i,t} : \mathcal{X}_i \times \mathcal{X}_j \to \mathbb{R}$ such that, for any PBE σ under private monitoring, SPE $\bar{\sigma}$ under public monitoring, public history h^t, private history h_i^t consistent with h^t, and $(X_i, X_j) \in \mathcal{X}_i \times \mathcal{X}_j$, we have*

$$w_t^i\left(\sigma, h_i^t, X_i\right) = W_t^i\left(\bar{\sigma}, h^t, X_i\right) = v_{i,t}\left(X_i, \theta_j\left(h^t\right)\right); \tag{7}$$

$$\hat{w}_t^i\left(\sigma, h_i^t, X_j\right) = \hat{W}_t^i\left(\bar{\sigma}, h^t, X_j\right) = v_{i,t}\left(\theta_i\left(h^t\right), X_j\right) \tag{8}$$

and

$$w_t^i\left(\sigma, h_i^t, \theta_i\left(h_i^t\right)\right) = \hat{w}_t^i\left(\sigma, h_i^t, \theta_j\left(h_i^t\right)\right) = w_t^i\left(\sigma, h_i^t, no\right) \tag{9}$$

$$= W_t^i\left(\bar{\sigma}, h^t, \theta_i\left(h^t\right)\right) = \hat{W}_t^i\left(\bar{\sigma}, h^t, \theta_j\left(h^t\right)\right) = W_t^i\left(\bar{\sigma}, h^t, no\right) = v_{i,t}\left(\theta\left(h^t\right)\right).$$

In the revision games with public monitoring, Gensbittel et al. (2017) show that the minimax theorem holds. In addition, Lovo and Tomala (2016) show the existence of Markov perfect equilibrium (MPE) with a finite equilibrium payoff after each profile of current policy sets. Putting them together, we obtain the results for public monitoring.

The theorem's contribution is to show that the equilibrium continuation payoff under private monitoring is the same as the one under public monitoring, and hence depends only on the current policy set profile. In the private monitoring case, for any Markov strategy of candidate i (where i's Markov strategies refer to those that depend only on the current policy set profile, the current time, and whether i receives an opportunity), there exists a best response by j that is Markov. This implies that i can guarantee her minmax value as in the public monitoring case. Since the symmetric argument implies that candidate j can guarantee his minmax value too, the equilibrium continuation payoff is uniquely determined.[63]

In the valence election campaign and symmetric office-motivated election campaign, candidates are office-motivated, so the payoffs are constant-sum. Also, the case with $p = 1$ in the election with incomplete information (the case with no normal types presented in Proposition 7) can be thought of as a constant-sum game after an elimination of strictly dominated strategies. Moreover, these models satisfy (5). Hence, in each of those models, the outcome characterized under private monitoring is outcome-equivalent to the one under public monitoring, and the continuation payoffs are uniquely determined for each time $-t$.

5 Conclusion

We have introduced the first model of dynamic campaigns into the literature on elections, which we call "policy announcement timing game." In the model, candidates cannot always announce their policies, but stochastically obtain opportunities to announce their policies or spend their funds. We applied the model to various examples, demonstrating that the introduction of such a simple friction to the model generates interesting dynamic strategic considerations and equilibrium dynamics consistent with election dynamics in reality. In particular, we showed that it is useful to analyze the candidates' motivations to defer a clear announcement of policies, depending on the opponent's latest announcement and the time left until the election; and to keep the budget for later use, depending on the opponent's cumulative spending and the time left. Depending on the

[63] Although it is intuitive, we do not know whether the result extends to the case with non-constant sum games. The (generic) uniqueness of continuation payoffs is an open question in the revision-games literature.

environment that the candidates face, they may or may not have such incentives for ambiguity. The insights from the examples are generalized in the Long Ambiguity Theorem, the Dynamic Median Voter Theorem, and the Constant-Sum Markov Theorem. Our work raises a wide range of new questions.

First, except for Sections 3.4 and 4.3, we restricted ourselves to the case in which policies are either perfectly ambiguous or perfectly precise. One could allow for "intermediate language" and analyze how gradually candidates shift from ambiguous to clear language over the course of the campaign. For example, in a uni-dimensional Downsian model, one could let the candidates choose any subintervals of $[0, 1]$ for the initial opportunity, and from the next opportunity, let them choose any subintervals included in their most recent announcements. In the multi-dimensional case, we can consider a model where a candidate commits to a policy in one dimension first, and then commits to a policy in another dimension.

Second, it would be more realistic to assume that policy announcements are sometimes synchronous and sometimes asynchronous. Although this problem seems nontrivial as Ishii and Kamada (2011) show in their analysis of revision games with synchronous and asynchronous revisions, we conjecture that there should remain the incentive to announce an ambiguous policy when the deadline is far.

Third, we restricted ourselves to the case in which, once a candidate commits to a particular policy, he or she cannot overturn it later. Although we believe that this is a reasonable starting point for analysis, one could also assume that candidates can change their policies if they are willing to incur a "reputational cost" for announcing "inconsistent" policies. The idea is that if a candidate overturns his or her opinion, voters would infer that it is likely that the candidate would change policies even after the election.

Fourth, it would be interesting to enrich the model by assuming that the median voter's position gets gradually revealed over the course of the campaign (for example, because of polls), so that candidates have an additional reason to wait. Our analysis in Section 3.5 shows that our general model can be extended to cover the cases involving incomplete information.

Fifth, the model of dynamic campaign spending could be enriched to test hypotheses for the "June Puzzle." This is a puzzle that asks why the Obama campaign significantly outspent the Romney campaign in June 2012, even though the election was in November and the effect of

TV advertisements on voter's preferences is known to be short-lived.[64] An explanation for this puzzle argues that popularity in the early stages may help with gathering more donations. Another explanation claims that if the opponent's popularity is below a certain level, then that opponent will "never come back." It will be interesting to enrich the model of Section 3.4 to analyze these hypotheses.

Sixth, we have considered two-candidate elections, but it would be interesting to consider more than two candidates.[65] In such an environment, there is no pure-strategy equilibrium in a static election game, while we can hope for the existence of an (essentially) unique pure-strategy PBE in a corresponding election campaign game, just as in the case with the multi-dimensional policy space.

Finally, our work raises empirical questions as well. For example, first, our model predicts different patterns of the timing of policy clarification/campaign spending for different parameter values. For example, in the valence election campaign, p, which measures how much uncertainty candidates face with respect to the position of the median voter, affects the timing of policy announcements. One may want to test whether this prediction is supported by the data.[66] The second example is about the case with a multi-dimensional policy space. In that model, we obtained a unique prediction about the entry timing and announced policies (when candidates are policy-motivated). The uniqueness may be useful in empirically testing the model.

References

[1] **Alesina, Alberto, and Alex Cukierman (1990)**, "The Politics of Ambiguity," *Quarterly Journal of Economics*, 4, 829–50.

[2] **Alesina, Alberto, and Richard Holden (2008)**, "Ambiguity and Extremism in Elections," unpublished.

[3] **Ambrus, Attila, and Shih-En Lu (2015)**, "A Continuous-Time Model of Multilateral Bargaining," *AEJ: Microeconomics*, 7, 208–249.

[64] See footnote 10 for the discussion on the depreciation of the effect of earlier spending. We thank Avidit Acharya for sharing the story of June Puzzle, who attributes the story to Seth Hill, Brett Gordon, and Michael Peress.

[65] We thank Alessandro Lizzeri for pointing this out.

[66] As mentioned in Remark 1, this pattern is roughly consistent with the empirical finding in Campbell (1983).

[4] **Ansolabehere, Stephen, and James M. Snyder Jr. (2000)**, "Valence Politics and Equilibrium in Spatial Election Models," *Public Choice*, 103, 327–336.

[5] **Aragonès, Enriqueta, and Zvika Neeman (2000)**, "Strategic Ambiguity in Electoral Competition," *Journal of Theoretical Politics*, 12, 183–204.

[6] **Aragonès, Enriqueta, and Thomas R. Palfrey (2002)**, "Mixed Equilibrium in a Downsian Model with a Favored Candidate," *Journal of Economic Theory*, 103, 131–161.

[7] **Aragonès, Enriqueta, and Andrew Postlewaite (2002)**, "Ambiguity in Election Games," *Review of Economic Design*, 7, 233–255.

[8] **Asahi Shimbun Digital (2014a)**, "Hosokawa-Koizumi 'Ex Prime Minister Union,' If Realized, May Change the Scene for Tokyo Gubernatorial Election" (in Japanese), January 14, https://web.archive.org/web/20140121185951/http://www.asahi.com/articles/ASG187WCQG18UTIL04R.html?ref=reca, accessed on 04/23/2018.

[9] **Asahi Shimbun Digital (2014b)** "The Focus is on Anti-Nuclear Power in the Tokyo Gubernatorial Election: Hosokawa and Masuzoe Declared to Run for Election" (in Japanese), January 15, https://web.archive.org/web/20140228074009/http://www.asahi.com/articles/ASG1G5R5SG1GUTIL03Y.html?ref=reca, accessed 04/23/2018.

[10] **Austen-Smith, David (1987)**, "Interest Groups, Campaign Contributions, and Probabilistic Voting," *Public Choice*, 54, 123–139.

[11] **Bailey, Michael A. (2002)**, "Money and Representation: An Exploration in Multiple Dimensions with Informative Campaigns," unpublished.

[12] **Banda, Kevin K. (2013)**, "The Dynamics of Campaign Issue Agendas," *State Politics & Policy Quarterly*, 13, 446–470.

[13] **Banda, Kevin K. (2015)**, "Competition and the Dynamics of Issue Convergence," *American Politics Research*, 43, 821–845.

[14] **Banks, Jeffrey S. (1990)**, "A Model of Electoral Competition with Incomplete Information," *Journal of Economic Theory*, 50, 309–325.

[15] **Baron, David P. (1994)**, "Electoral Competition with Informed and Uninformed Voters," *American Political Science Review*, 88, 33–47.

[16] **Berger, Mark M., Michael C. Munger, and Richard F. Potthoff (2000)**, "With Uncertainty, the Downsian Model Predicts Divergence," *Journal of Theoretical Politics*, 12, 262–268.

[17] **Bernhardt, Daniel, and Daniel Ingberman (1985)**, "Candidate Reputations and the Incumbency Effect," *Journal of Public Economics*, 27, 47–67.

[18] **Calcagno, Riccardo, Yuichiro Kamada, Stefano Lovo, and Takuo Sugaya (2014)**, "Asynchronicity and Coordination in Common and Opposing Interest Games," *Theoretical Economics*, 9, 409–434.

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

[20] **Callander, Steven, and Catherine H. Wilson (2008)**, "Context-Dependent Voting and Political Ambiguity," *Journal of Public Economics*, 92, 565–581.

[21] **Calvo, Guillermo A. (1983)**, "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics*, 12, 383–398.

[22] **Campbell, James E. (1983)**, "Ambiguity in the Issue Positions of Presidential Candidates: A Causal Analysis," *American Journal of Political Science*, 27, 284–293.

[23] **Carter, Jennifer and John W. Patty (2015)**, "Valence and Campaigns." *American Journal of Political Science* 59, 825–840.

[24] **Caruana, Guillermo and Liran Einav (2008)**, "A Theory of Endogenous Commitment," *Review of Economic Studies*, 75, 99–116.

[25] **Coate, Stephen (2004)**, "Political Competition with Campaign Contributions and Informative Advertising," *Journal of the European Economic Association*, 2, 772–804.

[26] **Coughlin, Peter J. (1992)**, *Probabilistic Voting Theory*, New York: Cambridge University Press.

[27] **van Damme, Eric, and Sjaak Hurkens (1996)**, "Commitment Robust Equilibria and Endogenous Timing," *Games and Economic Behavior*, 15, 290–311.

[28] **Gale, Douglas (1995)**, "Dynamic Coordination Games," *Economic Theory*, 5, 1–18.

[29] **Gale, Douglas (2001)**, "Monotone Games with Positive Spillovers," *Games and Economic Behaviors*, 37, 295–320.

[30] **Gensbittel, Fabien, Stefano Lovo, Jérôme Renault, and Tristan Tomala (2017)**, "Zero-Sum Revision Games, *Games and Economic Behavior*, Forthcoming.

[31] **Glazer, Amihai (1990)**, "The Strategy of Candidate Ambiguity," *American Political Science Review*, 84, 237–41.

[32] **Gerber, Alan, James G. Gimpel, Donald P. Green, and Daron R. Shaw (2011)**, "How Large and Long-lasting Are the Persuasive Effects of Televised Campaign Ads? Results from a Randomized Field Experiment," *American Political Science Review*, 105, 135–150.

[33] **Groseclose, Tim (2001)** "A Model of Candidate Location When One Candidate Has a Valence Advantage," *American Journal of Political Science*, 45, 862–886.

[34] **Grossman Gene M., and Elhanan Helpman (1996)**, "Electoral Competition and Special Interest Politics," *Review of Economic Studies*, 63, 256–286.

[35] **Gurciullo, Brianna, and Gabriel Debenedetti (2016)** "Sanders: 'We are in this campaign to win' " in *POLITICO*, https://www.politico.com/blogs/2016-dem-primary-live-updates-and-results/2016/04/bernie-sanders-jeff-weaver-victory-222536, accessed on 04/19/2018.

[36] **Hamilton, Jonathan H., and Steven M. Slutsky (1990)**, "Endogenous Timing in Duopoly Games: Stackelberg or Cournot Equilibria," *Games and Economic Behavior*, 2, 29–46.

[37] **Hamilton, Jonathan H., and Steven M. Slutsky (1993)**, "Endogenizing the Order of Moves in Matrix Games," *Theory and Decision*, 34, 47–62.

[38] **Harrington, Jr. , Joseph E. (1992)**, "The Revelation of Information Through the Electoral Process: An Exploratory Analysis," *Economics and Politics*, 4, 255–275.

[39] **Harris, Tobias (2013)** "Shinzo Abe's Constitution Quest: Instead of focusing on economics, the Japanese leader risks it all to change the rules of governance" in *Wall Street Journal*, http://jp.wsj.com/articles/SB10001424127887324582304578488313061482222, accessed 05/26/2018.

[40] **Hummel, Patrick (2010)**, "On the Nature of Equilibria in a Downsian Model with Candidate Valence," *Games and Economic Behavior*, 70, 425–445.

[41] **Ishii, Yuhta, and Yuichiro Kamada (2011)**, "The Effect of Correlated Inertia on Coordination," unpublished.

[42] **Kamada, Yuichiro, and Michihiro Kandori (2017a)**, "Revision Games, Part I: Theory," unpublished.

[43] **Kamada, Yuichiro, and Michihiro Kandori (2017b)**, "Revision Games, Part II: Applications and Robustness" unpublished.

[44] **Karazas, Ionnis, and Steven E. Shreve (1988)**, *Brownian Motion and Stochastic Calculus*, New York: Springer-Verlag.

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

[46] **Klenow, Peter J. and Benjamin A. Malin (2010)**, "Microeconomic Evidence on Price-Setting," *Handbook of Monetary Economics*, 3, 231–284.

[47] **Krasa, Stefan and Mattias Polborn (2010)**, "The Binary Policy Model," *Journal of Economic Theory*, 145, 661–688.

[48] **Lindbeck, Assar, and Jorgen Weibull (1987)**, "Balanced Budget Redistribution as the Outcome of Political Competition," *Political Choice*, 52, 273-297.

[49] **Lipman, Barton L. and Ruqu Wang (2000)**, "Switching Costs in Frequently Repeated Games," *Journal of Economic Theory*, 93, 149–190.

[50] **Lovo, Stefano, and Tristan Tomala (2015)**, "Markov Perfect Equilibria in Stochastic Revision Games," unpublished.

[51] **McGrane, Reginald C. (ed.) (1919)**, *The Correspondence of Nicholas Biddle*. Boston, New York: The Houghton Mifflin Company.

[52] **Meirowitz, Adam (2005)**, "Informational Party Primaries and Strategic Ambiguity," *Journal of Theoretical Politics*, 17, 107–136.

[53] **Moroni, Sofia (2018)**, "Existence of Trembling Hand Perfect and Sequential Equilibrium in Games with Stochastic Timing of Moves," unpublished.

[54] **Ota, Shinji (2016)** "2016 House of Councillors Election, Prospect of Constitutional Reform and Politics" (in Japanese), *Jiji-Koron*, http://www.nhk.or.jp/kaisetsu-blog/100/248861.html, accessed on 04/19/2018.

[55] **Page, Benjamin I. (1976)**, "The Theory of Political Ambiguity," *American Political Science Review*, 70, 742–752.

[56] **Palfrey, Thomas R. (1984)**, "Spatial Equilibrium with Entry," *Review of Economic Studies*, 51, 139–156.

[57] **Page, Benjamin I. (1978)**, *Choices and Echoes in Presidential Elections*, Chicago: University of Chicago Press.

[58] **Persson, Torsten, and Guido Tabellini (2000)**, *Political Economics: Explaining Economic Policy*, Cambridge: MIT Press.

[59] **Potters, Jan, Randolph Sloof, and Frans van Winden (1997)**, "Campaign Expenditures, Contributions and Direct Endorsements: The Strategic Use of Information and Money to Influence Voter Behavior," *European Journal of Political Economy*, 13, 1–31.

[60] **Prat, Andrea (2002a)**, "Campaign Spending with Office-Seeking Politicians, Rational Voters, and Multiple Lobbies," *Journal of Economic Theory*, 103, 162–189.

[61] **Prat, Andrea (2002b)**, "Campaign Advertising and Voter Welfare," *Review of Economic Studies*, 69, 999–1018.

[62] **Renou, Ludovic (2009)**, "Commitment games," *Games and Economic Behavior*, 66, 488–505.

[63] **Roemer, John E. (2001)**, *Political Competition: Theory and Applications*, Cambridge: Harvard University Press.

[64] **Romano, Richard, and Huseyin Yildirim (2005)**, "On the endogeneity of Cournot–Nash and Stackelberg equilibria: games of accumulation," *Journal of Economic Theory*, 120, 73–107.

[65] **Sankei News (2013)**, "Survey by the Liberal Democratic Party Reveals Strong Supports for Masuzoe and Ikegami for Tokyo Gubernatorial Election" (in Japanese), December 24, https://web.archive.org/web/20140629090111/http://sankei.jp.msn.com/politics/news/131224/elc13122421110000-n1.htm, accessed 04/23/2018.

[66] **Strauss, Daniel (2016)** "Sanders vows to go all the way, but signals a new focus on influencing Clinton" in *POLITICO*, https://www.politico.com/blogs/2016-dem-primary-live-updates-and-results/2016/04/bernie-sanders-election-results-convention-222506, accessed on 04/19/2018.

[67] **Shepsle, Kenneth (1972)**, "The Strategy of Ambiguity: Uncertainty and Electoral Competition," *American Political Science Review*, 66, 555–568.

[68] **Stein, Jeff (2016)** "Bernie Sanders moved Democrats to the left. The platform is proof." in *Vox*, https://www.vox.com/2016/7/25/12281022/the-democratic-party-platform, accessed on 05/07/2018.

[69] **White, Joe Slade (2012)**, "The 9 Principles of Winning Campaigns," http://joesladewhite.com/wp-content/uploads/2015/01/9-principles-blue-jsw.pdf, accessed on 05/07/2018.

A Structure of the Appendix

We first state and prove the continuous-time backward induction, which turns out to be useful in many proofs. Second, we offer the proofs of the results. Although we present the applications before the general theorems in the main text to highlight the applicability of the model, since the general theorems are useful for proving the results in the applications, here we prove the general theorems. The proofs of the results in the applications can be found in the Online Appendix.

B Continuous-Time Backward Induction

The following result, which we call continuous-time backward induction, is due to Calcagno et al. (2014), and is repeatedly used in the proofs of this paper. We reproduce its statement and proof for reader's convenience.

Lemma 5 *Suppose that for any $t \in [0, \infty)$, there exists $\varepsilon > 0$ such that statement $A_{t'}$ is true for all $t' \in [t, t+\varepsilon)$ if statement $A_{t''}$ is true for any $t'' < t$. Then, for any $t \in [0, \infty)$, statement A_t is true.*

Proof. Suppose that the premise of the lemma holds. Let $-t^*$ be the supremum of $-t$ such that A_t is false. If $t^* = \infty$, we are done. So suppose that $t^* < \infty$. Then it must be the case that for any $\varepsilon > 0$, there exists $-\tau \in (-t^* - \varepsilon, -t^*]$ such that A_τ is false. But by the definition of t^*, there exists $\tilde{\varepsilon} > 0$ such that statement A_τ is true for all $-\tau \in (-t^* - \tilde{\varepsilon}, -t^*]$ because the premise of the lemma is true. This is a contradiction. ■

C Proof of Theorem 1

The most important proofs for Theorem 1 are here. Other proofs for the theorem can be found in the Online Appendix.

Proof of Proposition 11.

Suppose $t_j = \infty$. On the one hand, since candidate j enters whenever she has an opportunity, if candidate i does not enter until j enters, the payoff i obtains at time $-t$ converges to $\sup_{x_i} v_i\left(x_i, x_j^*\right)$ as $t \to \infty$. On the other hand, we have $v_{i,t}(\text{enter}) \to v_i^{BR_j}$ as $t \to \infty$ since candidate j will have an opportunity afterward with a probability converging to 1 as $t \to \infty$. Hence, for sufficiently large t, not entering is the unique best response at time $-t$ in any PBE given that first-mover disadvantage for i holds.

Suppose next $t_j < \infty$. On the one hand, for each $-t < -t_j$, if candidate i does not enter until $-t_j$, i obtains the payoff $\bar{v}_{i,t_j}(\text{not})$, which is a convex combination of $v_i\left(X, x_j^*\right)$, $v_i\left(x_i^*, X\right)$,

$v_i\left(X,X\right)$, $\sup_{x_i} v_i\left(x_i, x_j^*\right)$, and $v_i^{BR_j}$. Moreover, since j enters at times sufficiently close to time $-t_0$ (given the continuity of the continuation payoff in time), we have a strictly positive weight on $\sup_{x_i} v_i\left(x_i, x_j^*\right)$ in the convex combination. On the other hand, we have $v_{i,t}(\text{enter}) \to v_i^{BR_j}$ as $t \to \infty$. Hence, for sufficiently large t, not entering is the unique best response at time $-t$ in any PBE given that first-mover disadvantage for i holds. ■

Proof of Proposition 13.

Suppose first that strong first-mover disadvantage for A holds. In this case, if $\hat{t}_A = \hat{t}_B = \infty$, then by the definition of $\hat{t}_A$, A never has an incentive to enter at any time. Hence the conclusion of the proposition holds with $t_A = 0$. If $\hat{t}_A > \hat{t}_B$, then part 2 of Proposition 12 implies that A never has an incentive to enter at all times strictly before $-\hat{t}_B$. Hence the conclusion of the proposition holds with $t_A = \hat{t}_B$. If $\hat{t}_A < \hat{t}_B$, then part 1 of Proposition 12 implies that we can use the argument for Case 2. Hence, Proposition 11 implies that A never has an incentive to enter at all times strictly before $-\tilde{t}$ where $\tilde{t}$ is equal to t_A that we take in Proposition 10. Hence the conclusion of the proposition holds with $t_A = \tilde{t}$.

Next, suppose that strong first-mover disadvantage for B holds. First, we show $\hat{t}_B < \infty$. To see why this holds, observe the following. On the one hand, the payoff from entering at $-t$ converges to $v_B^{BR_A}$ as $t \to \infty$. On the other hand, for an arbitrary fixed $\bar{t} \in (0, \infty)$, not entering until $-\bar{t}$ (and entering when an opportunity arrives at time $-t \geq -\bar{t}$) gives the payoff that is a convex combination of $v_B\left(x_B^*, X\right)$, $v_B^{BR_A}$, and $v_B\left(X, X\right)$ with a strictly positive weight on $v_B\left(x_B^*, X\right)$ (recall that $\hat{t}_B$ is calculated assuming that candidate A never enters unless B has entered). Hence, for sufficiently large t, not entering is better at time $-t$ in any PBE given strong first-mover disadvantage for B. Thus, we have $\hat{t}_B < \infty$. With this condition, we obtain the desired result as in the case of strong first-mover disadvantage for A, using part 2 of Proposition 12 for the case of $\hat{t}_A > \hat{t}_B$ and Proposition 11 for the case of $\hat{t}_A < \hat{t}_B$. ■

D Proof of a Generalized Version of Theorem 2

In in part 3 of the following theorem, we consider general games in which it is not necessarily the case that $\mathcal{X}_i = \{\{x\} | x \in X\} \cup \{X\}$ holds. In such a game, when $\{x_i^*\} \in \mathcal{X}_i$, we say that x_i^* is a **strictly dominant policy** if for all $X_i \in \mathcal{X}_i \setminus \{x_i^*\}$, $v_i(\{x_i^*\}, X_j) > v_i(X_i, X_j)$ for all $X_j \in \mathcal{X}_j$.

We note that, even though we require (v_A, v_B) to be symmetric in the following theorem, it is straightforward to extend the result to the case with non-symmetric cases, as described in footnote 61.

Theorem 4 (General Dynamic Median-Voter) *1. Suppose that (v_A, v_B) is symmetric and, for each $i = A, B$, x^* is a Condorcet winner. Then, there exists a PBE in which, at any time $-t$, conditional on any history, (i) if the opponent's current policy set is X, each candidate i announces $\{x^*\}$, and (ii) otherwise, each candidate i chooses a static best response to the opponent's current policy.*

2. *Suppose that X is finite, (v_A, v_B) is a symmetric constant-sum game, and there exists a Condorcet winner. Then, there exists a PBE, and in any PBE, at any time $-t$, conditional on any history, (i) if the opponent's current policy set is X, each candidate i announces the Condorcet winner and (ii) otherwise, each candidate i chooses a static best response to the opponent's current policy.*

3. *Suppose that X is finite and (v_A, v_B) is symmetric. Consider an environment with a general $(\mathcal{X}_A, \mathcal{X}_B)$. Suppose that x_i^* is a strictly dominant policy for each $i = A, B$. Then, in any PBE, at any time $-t$, conditional on any history, each candidate i announces $\{x^*\}$.*

As will be seen in the Online Appendix, part 3 of Theorem 4 does not hold if we replace "strictly dominant policy" with "weakly dominant policy."

Proof. Part 1: Let x^* be a Condorcet winner. First, note that after the opponent enters at x^*, the given strategy specifies a best response. Second, we show that it is optimal for each candidate to announce x^* at time $-t$ if j's current policy set is X. Take an arbitrary $\bar{x}_j(x_i) \in \max_{x_j \in X} v_i(x^*, x_j)$. By the assumption that x^* is a Condorcet winner, it follows that for any $x_i \in X$, there exists $\bar{x}_j(x_i) \in BR_j(x_i)$ such that $v_i(x^*, \bar{x}_i(x^*)) > v_i(x_i, \bar{x}_j(x_i))$. We consider a strategy profile in which, once i enters at $x_i \in X$, j enters at $\bar{x}_j(x_i)$; and also if no one has entered, each i enters at x^*.

Take any time $-t \in [-T, 0]$. Suppose that for every time in $(-t, 0]$, conditional on any history, each candidate i announces x^*. Under the strategy profile specified above, the following are true.

If i announces x^*, her payoff is

$$e^{-\lambda_j t} v_i(x^*, X) + (1 - e^{-\lambda_j t}) v_i(x^*, \bar{x}_j(x^*)). \tag{10}$$

If i announces $x_i \neq x^*$, her payoff is

$$e^{-\lambda_j t} v_i(x_i, X) + (1 - e^{-\lambda_j t}) v_i(x_i, \bar{x}_j(x_i)). \tag{11}$$

If i announces X, her payoff is

$$\begin{aligned} &e^{-\lambda_j t} \left(e^{-\lambda_i t} v_i(X, X) + (1 - e^{-\lambda_i t}) v_i(x^*, X) \right) + (1 - e^{-\lambda_j t}) \\ &\left(e^{-\lambda_i t} v_i(X, x^*) + (1 - e^{-\lambda_i t}) \left(w_i v_i(x^*, \bar{x}_j(x^*)) + (1 - w_i) \, v_i(\bar{x}_i(x^*), x^*) \right) \right), \end{aligned} \tag{12}$$

where $w_i \in (0, 1)$.

Note that the first term in (10) is weakly larger than the first terms in (11) and (12) due to the assumption that x^* is a Condorcet winner. In addition, the second term in (10) is weakly larger than the second term in (11) by the construction of the function $\bar{x}_j(\cdot)$. Finally, the second term in (10) is weakly larger than the second term in (12) due to $x^* \in BR_i(x^*)$ (implied by x^* being a Condorcet winner), and its implication that $v_i(\bar{x}_i(x^*), x^*) = v_i(x^*, x^*)$. This implies that, for any $t \geq 0$, (10) is weakly larger than both (11) and (12). Hence, it is optimal for each candidate to announce x^* at time $-t$ if j's current policy set is X.

Part 2: Let x^* be the Condorcet winner. For each $X_i \in \mathcal{X}_i$, we have $v_i(x_i^*, x_j) > v_i(X_i, x_j')$ for some $x_j \in BR_j(x^*)$ and $x_j' \in BR_j(X_i)$. Since (v_A, v_B) is constant-sum, we have $v_i(x^*, x_j) > v_i(X_i, x_j')$ for each $x_j \in BR_j(x^*)$ and $x_j' \in BR_j(X_i)$. This implies that, since $x^* \in BR_j(x^*)$, we have $v_i(x^*, x^*) = \min_{X_j \in \mathcal{X}_j} v_i(x^*, X_j) > \min_{X_j \in \mathcal{X}_j} v_i(X_i, X_j)$ for each $X_i \neq \{x^*\}$.

Take any time $-t \in [-T, 0]$. Suppose that, for every time in $[-t, 0]$, conditional on any history, (i) if the opponent's current policy set is X, each candidate i announces x^* and (ii) otherwise, each candidate i takes a static best response to the opponent's current policy. We show that, conditional on any history, if the opponent's current policy set is X at time $-t$, each candidate i has a strict incentive to announce x^* over announcing X or any singleton policy set $\{x_i\} \neq \{x^*\}$. Given the constant-sum assumption, if i announces x^*, her payoff is

$$e^{-\lambda_j t} v_i(x^*, X) + (1 - e^{-\lambda_j t}) v_i(x^*, x^*). \tag{13}$$

If i announces $x_i \neq x^*$, her payoff is

$$e^{-\lambda_j t} v_i(x_i, X) + (1 - e^{-\lambda_j t}) \min_{X_j \in \mathcal{X}_j} v_i(x_i, X_j). \tag{14}$$

If i announces X, her payoff is

$$e^{-\lambda_j t} \left(e^{-\lambda_i t} v_i(X, X) + (1 - e^{-\lambda_i t}) v_i(x^*, X) \right) + (1 - e^{-\lambda_j t}) \left(e^{-\lambda_i t} v_i(X, x^*) + (1 - e^{-\lambda_i t}) v_i(x^*, x^*) \right). \tag{15}$$

Given that x^* is the Condorcet winner and $x_i \neq x^*$, (13) is strictly larger than (14) and (15). Moreover, by the assumption that X is finite, there exits $\varepsilon > 0$ such that, for all $x_i \neq x^*$, the value in (13) is no less than the sum of ε and the value in (14), and also no less than the sum of ε and the value in (15). By continuity of the continuation payoff in time, this implies that there exists $\varepsilon' > 0$ such that i strictly prefers announcing x^* to announcing X or any singleton policy set $\{x_i\} \neq \{x^*\}$ at times in $(-t - \varepsilon', -t]$ if j's current policy set is X. Therefore, by the continuous-time backward induction, in any PBE, at any time $-t$, conditional on any history, each candidate i announces x^* if j's current announcement is X.

Part 3: Fix time $-t$, suppose that at all time strictly after $-t$, each candidate i enters at x_i^* conditional on any history. Then, if i announces x_i^* when the current policy set is (X_i, X_j), then her payoff is

$$e^{-\lambda_j t} v_i(x_i^*, X_j) + (1 - e^{-\lambda_j t}) v_i(x_i^*, x_j^*).$$

If i announces $X_i \neq \{x_i^*\}$ when the current policy set is (X_i, X_j), then her payoff is

$$e^{-\lambda_j t} \left(e^{-\lambda_i t} v_i(X_i, X_j) + (1 - e^{-\lambda_i t}) \bar{v}_i \right) + (1 - e^{-\lambda_j t}) \left(e^{-\lambda_i t} v_i(X_i, x_j^*) + (1 - e^{-\lambda_i t}) \bar{\bar{v}}_i \right),$$

where $\bar{v}_i, \bar{\bar{v}}_i \leq v_i(x_i^*, X_j)$. Note that $\bar{v}_i$ and $\bar{\bar{v}}_i$ are equal to $v_i(x_i^*, X_j)$ and $v_i(x_i^*, x_j^*)$, respectively, if $x_i^* \in X_i$, but they are respectively strictly less than those values otherwise, due to the definition of a strictly dominant policy.

Since $v_i(X_i, X_j) < v_i(x_i^*, X_j)$ and $v_i(X_i, x_j^*) < v_i(x_i^*, x_j^*)$ by the definition of a strictly dominant policy, the payoff from announcing $\{x_i^*\}$ is strictly greater than the payoff from announcing $X_i \neq \{x_i^*\}$. Hence, by the continuous-time backward induction, we obtain the desired result. ■

E Proof of Theorem 3

We first prove that conditions (7)–(9) hold for public monitoring. Using this result, we prove that conditions (7)–(9) hold for private monitoring.

E.1 Public Monitoring

We prove that conditions (7)–(9) hold for public monitoring:

Lemma 6 *Suppose $v_A(X_A, X_B) + v_B(X_B, X_A) = 1$ for each $(X_A, X_B) \in \mathcal{X}_A \times \mathcal{X}_B$. There exists $v_{i,t}(\theta)$ for each $\theta \in \mathcal{X}_i \times \mathcal{X}_j$ such that, for any SPE $\bar{\sigma}$ and h^t, we have*

$$W_t^i\left(\bar{\sigma}, h^t, X_i\right) = v_{i,t}\left(X_i, \theta_j\left(h^t\right)\right);$$

$$\hat{W}_t^i\left(\bar{\sigma}, h^t, X_j\right) = v_{i,t}\left(\theta_i\left(h^t\right), X_j\right);$$

and

$$W_t^i\left(\bar{\sigma}, h^t, \theta_i\left(h^t\right)\right) = \hat{W}_t^i\left(\bar{\sigma}, h^t, \theta_j\left(h^t\right)\right) = W_t^i\left(\bar{\sigma}, h^t, no\right) = v_{i,t}\left(\theta\left(h^t\right)\right).$$

Proof. In the revision games with public monitoring, the minimax theorem holds if there exists an equilibrium with a finite equilibrium payoff after any history (see Gensbittel et al. (2017)). Hence, we are left to show that there exists an equilibrium with a finite equilibrium payoff after each profile of policy sets, θ. Lovo and Tomala (2016) show the existence of Markov perfect equilibrium (MPE) with a finite equilibrium payoff after each θ, as desired. ■

E.2 Private Monitoring

Fix any σ_j (not necessarily an equilibrium strategy). Given h_i^t and $\theta\left(h_i^t\right) = (X, X)$, calculate candidate i's belief about h_j^t. For any time $-t$ and history h_i^t, $\theta_j\left(h_i^t\right) = X$ happens with a positive probability given (σ_i, σ_j) for any strategy σ_i of candidate i since it is possible that no candidate receives any opportunity in the time interval $[-T, -t)$. Hence the belief given $\theta_j\left(h_i^t\right) = X$ should satisfy Bayes rule. In particular, we can show that, since the arrivals of opportunities are independent between candidates, for any two histories of candidate i such that no commitment has been made at $-t$, candidate i's belief about h_j^t is the same. Denote by β^{σ_j} a belief to be explicit about the fact that the belief is solely determined by σ_j:

Lemma 7 *For any σ_j, there exists β^{σ_j} such that, for each h_i^t and $\tilde{h}_i^t$ with $\theta\left(h_i^t\right)=\theta\left(\tilde{h}_i^t\right)=(X,X)$, we have $\beta^{\sigma_j}\left(h_j^t|\tilde{h}_i^t\right)=\beta^{\sigma_j}\left(h_j^t|h_i^t\right)=:\beta^{\sigma_j}\left(h_j^t\right)$.*

Proof. Note that, if $\theta\left(h_i^t\right)=(X,X)$, then we have $\left(t_i^l,X_i^l\right)_{l=1}^{l_i}=\{\emptyset\}$,[67] that is, candidate i never changes her policy announcement. Let $H_j^{\sigma_j}\left(h_i^t\right)$ be the set of candidate j's histories compatible with h_i^t and σ_j.[68] Define $H_j^{\sigma_j}\left(\tilde{h}_i^t\right)$ analogously. Note that $H_j^{\sigma_j}\left(h_i^t\right)$ and $f_i(h_j^t|h_i^t)$ depend on h_i^t only through $\left(t_i^l,X_i^l\right)_{l=1}^{l_i}=\{\emptyset\}$. Hence, $H_j^{\sigma_j}\left(h_i^t\right)=H_j^{\sigma_j}\left(\tilde{h}_i^t\right)$ and $f_i(h_j^t|h_i^t)=f_i(h_j^t|\tilde{h}_i^t)$ for each h_i^t and $\tilde{h}_i^t$ with $\theta\left(h_i^t\right)=\theta\left(\tilde{h}_i^t\right)=(X,X)$. Thus, the result follows from (20). ∎

Using this independence of the belief, we can show that candidate i's continuation payoff does not depend on h_i^t. Take any strategy profile σ (not necessarily an equilibrium). Let

$$\tilde{w}_t^i\left(\sigma_i,\sigma_j,h_i^t,X\right)=\int_{h_j\in H_j^{\sigma_j}\left(h_i^t\right)}u_i\left(\sigma_i,\sigma_j|\left(h_i^t,X\right),h_j^t\right)d\beta^{\sigma_j}\left(h_j^t|h_i^t\right)$$

be candidate i's payoff when she takes X given h_i^t, given that (i) candidate i takes a continuation strategy determined by σ_i and history $\left(h_i^t,X\right)$ for $(-t,0]$,[69] and (ii) if candidate j has never received an opportunity before time $-t$ in h_j^t, she takes a continuation play determined by σ_j and history h_j^t for $(-t,0]$. Note that (i) candidate i's decision X does not affect candidate i's belief $\beta^{\sigma_j}\left(\cdot|h_i^t\right)$; and (ii) the belief $\beta^{\sigma_j}\left(\cdot|h_i^t\right)$ does not depend on whether candidate i obtains an opportunity at time $-t$ by the independence of the Poisson processes.

By Lemma 7, for each h_i^t and $\tilde{h}_i^t$ with $\theta\left(h_i^t\right)=\theta(\tilde{h}_i^t)=(X,X)$, we have

$$\begin{aligned}\sup_{\sigma_i}\tilde{w}_t^i\left(\sigma_i,\sigma_j,h_i^t,X\right) &= \sup_{\sigma_i}\int_{h_j}u_i(\sigma_i,\sigma_j|\left(h_i^t,X\right),h_j^t)d\beta^{\sigma_j}\left(h_j^t\right)\\ &= \sup_{\sigma_i}\int_{h_j}u_i(\sigma_i,\sigma_j|(\tilde{h}_i^t,X),h_j^t)d\beta^{\sigma_j}\left(h_j^t\right)\\ &= \sup_{\sigma_i}\tilde{w}_t^i\left(\sigma_i,\sigma_j,h_i^t,X\right).\end{aligned}$$

The second last line follows since the distribution of the final outcome that candidate i can induce

[67] We follow the convention that, with $l_j<1$, we define $\left(t_i^l,X_i^l\right)_{l=1}^{l_i}=\{\emptyset\}$.

[68] The formal definition of $H_j^{\sigma_j}\left(h_i^t\right)$ is provided in the Online Appendix when we formally define Bayes rule.

[69] Recall that, in Section 2, we define $u_i(\sigma_i,\sigma_j|h_i^t,h_j^t)$. Here, we define $u_i(\sigma_i,\sigma_j|(h_i^t,X),h_j^t)$ analogously, conditional on the event that candidate i takes X at time $-t$.

depends only on $\beta^{\sigma_j}\left(h_j^t\right)$ and $\theta\left(h_i^t\right)$. Hence, we can write

$$\tilde{w}_t^i\left(\sigma_j, X\right) = \sup_{\sigma_i} \tilde{w}_t^i\left(\sigma_i, \sigma_j, h_i^t, X\right)$$

for each h_i^t with $\theta\left(h_i^t\right) = (X, X)$.

Similarly, let $\tilde{w}_t^i\left(\sigma_i, \sigma_j, h_i^t, no\right)$ be candidate i's payoff given that she does not receive an opportunity at time $-t$. We also have

$$\tilde{w}_t^i\left(\sigma_j, X\right) = \sup_{\sigma_i} \tilde{w}_t^i\left(\sigma_i, \sigma_j, h_i^t, X\right) = \sup_{\sigma_i} \tilde{w}_t^i\left(\sigma_i, \sigma_j, h_i^t, no\right) \tag{16}$$

since, given h_j^t, candidate j's continuation play after (h_i^t, X) and that after (h_i^t, no) are the same (candidate j's history will be the same after (h_i^t, X) and after (h_i^t, no)).

Together with the constant-sum assumption, we can show that $\tilde{w}_t^i\left(\sigma_j, X\right) + \tilde{w}_t^j\left(\sigma_i, X\right) = 1$ for any PBE σ.

Lemma 8 *Suppose $v_A(X_A, X_B) + v_B(X_B, X_A) = 1$ for each $(X_A, X_B) \in \mathcal{X}_A \times \mathcal{X}_B$. For any PBE σ, the following holds: Fix $v_i \in [0, 1]$ and $t \geq 0$. Then, the following two claims hold:*

1. *If we have $\tilde{w}_t^i\left(\sigma_j, X\right) > v_i$, then we have $\tilde{w}_t^j\left(\sigma_i, X\right) < 1 - v_i$.*

2. *If we have $\tilde{w}_t^i\left(\sigma_j, X\right) < v_i$, then we have $\tilde{w}_t^j\left(\sigma_i, X\right) > 1 - v_i$.*

Proof. By symmetry, we only prove Claim 1. The ex-ante continuation payoff for candidate i from period t given $\theta\left(h_i^t\right) = (X, X)$ is, by Bayes rule,

$$\begin{aligned}
\frac{\int_{h_i^t:\theta\left(h_i^t\right)=(X,X)} \tilde{w}_t^i\left(\sigma, h_i^t, no\right) d\beta\left(h_i^t\right)}{\int_{h_i^t:\theta\left(h_i^t\right)=(X,X)} d\beta\left(h_i^t\right)} &= \frac{\int_{h_i^t:\theta\left(h_i^t\right)=(X,X)} \tilde{w}_t^i\left(\sigma_j, X\right) d\beta\left(h_i^t\right)}{\int_{h_i^t:\theta\left(h_i^t\right)=(X,X)} d\beta\left(h_i^t\right)} \quad \text{(by the equilibrium condition)} \\
&= \tilde{w}_t^i\left(\sigma_j, X\right) > v_i.
\end{aligned}$$

Similarly, the ex-ante continuation payoff for candidate j from period t given $\theta\left(h_j^t\right) = (X, X)$ is $\tilde{w}_t^j\left(\sigma_i, X\right)$. Since the ex ante continuation payoffs should add up to one, we have $\tilde{w}_t^j\left(\sigma_i, X\right) < 1 - v_i$. ■

Given Lemma 6, together with Lemmas 7 and 8, proving the following lemma will be sufficient for conditions (7)–(9) to hold in private monitoring:

Lemma 9 *Suppose $v_A(X_A, X_B) + v_B(X_B, X_A) = 1$ for each $(X_A, X_B) \in \mathcal{X}_A \times \mathcal{X}_B$. Take $v_{i,t}(\theta)$ that satisfies conditions stated in Lemma 6. Then, for any h_i^t, we have*

$$w_t^i\left(\sigma, h_i^t, X_i\right) = v_{i,t}\left(X_i, \theta_j\left(h_i^t\right)\right);$$

$$\hat{w}_t^i\left(\sigma, h_i^t, X_j\right) = v_{i,t}\left(\theta_i\left(h_i^t\right), X_i\right);$$

and

$$w_t^i\left(\sigma, h_i^t, \theta_i\left(h^t\right)\right) = \hat{w}_t^i\left(\sigma, h_i^t, \theta_j\left(h^t\right)\right) = w_t^i\left(\sigma, h_i^t, no\right) = v_{i,t}\left(\theta\left(h^t\right)\right).$$

Proof. Once a candidate takes $x \in X$, then the other candidate takes a static best response against x whenever he receives an opportunity. In the constant-sum game, this continuation strategy uniquely pins down the equilibrium payoff once a candidate takes $x \in X$, and the payoff does not depend on whether candidates observe the opponent's arrivals of the Poisson process. Hence, we will focus on the case $\theta\left(h_i^t\right) = \theta\left(h_j^t\right) = (X, X)$. By (16) and the equilibrium condition, we can write $w_t^i\left(\sigma, h_i^t, X\right) = \hat{w}_t^i\left(\sigma, h_i^t, X\right) = w_t^i\left(\sigma, h_i^t, no\right) = \tilde{w}_t^i\left(\sigma_j, X\right)$.

Suppose that there exists a PBE $\tilde{\sigma} \in \Sigma$ such that, for some $i \in \{A, B\}$ and h_i^t, we have $\tilde{w}_t^i\left(\tilde{\sigma}_j, X\right) \neq v_{i,t}(X, X)$. Without loss,[70] we can assume

$$\tilde{w}_t^i\left(\tilde{\sigma}_j, X\right) > v_{i,t}(X, X). \tag{17}$$

From Lemma 8, for each $\tilde{h}_j^t$ with $\theta\left(\tilde{h}_j^t\right) = (X, X)$, candidate j's expected payoff is less than $1 - v_{i,t}(X, X)$.

First, candidate i's Markov strategy is a map $\sigma_i : \mathcal{X}_i \times \mathcal{X}_j \times [0, T] \to \Delta(\mathcal{X}_i)$. Let M_i be the the space of i's Markov strategies. Note that the space for Markov strategies in public monitoring is the same as the space for Markov strategies in private monitoring. Since Markov strategies are constant with respect to the part of the histories other than the current policy sets and the current time, we write $M_i \subseteq \Sigma_i$ for each i.

Note that, in the model with public monitoring, there exists a Markov perfect equilibrium (MPE), where each candidate's strategy depends only on t, $\theta\left(h^t\right)$, and whether he or she receives an opportunity at the current time (see Gensbittel et al. (2017) for the proof). Fix a MPE

[70] If $\tilde{w}_t^i\left(\tilde{\sigma}_j, X\right) < v_{i,t}(X, X)$, then since the game is constant-sum, we have $v_{j,t}(X, X) = 1 - v_{i,t}(X, X)$. From Lemma 8, we have $\tilde{w}_t^j\left(\tilde{\sigma}_i, X\right) > v_{j,t}(X, X)$. The following lemma goes through with indices i and j being reversed.

$(\sigma_i, \sigma_j) \in M_i \times M_j$. We have

$$W_t^i \left(\sigma_i', \sigma_j, h^t, X\right) = \hat{W}_t^i \left(\sigma_i', \sigma_j, h^t, X\right) \leq v_{i,t} \left(X, X\right) \tag{18}$$

for each $\sigma_i' \in \Sigma_i$ since σ_i must designate a best response at every h^t in public monitoring.

Since this strategy σ_j is Markov, candidate j can take this strategy in private monitoring. We will show $w_t^i \left(\tilde{\sigma}_i, \sigma_j, \hat{h}_i^t, X\right) > v_{i,t} \left(X, X\right)$ for some $\hat{h}_i^t$ with $\theta\left(\hat{h}_i^t\right) = (X, X)$ by (17) since otherwise candidate j would like to deviate to σ_j from $\tilde{\sigma}_j$ given each h_j^t with $\theta\left(h_j^t\right) = (X, X)$ in private monitoring and obtain the expected payoff no less than $1 - v_{i,t} \left(X, X\right)$. Then, by (16), we have $\tilde{w}_t^i \left(\sigma_j, X\right) > v_{i,t} \left(X, X\right)$.

Thus, for each $\tilde{h}_i^t$ with $\theta\left(\tilde{h}_i^t\right) = (X, X)$, we have

$$\begin{aligned} \tilde{w}_t^i \left(\sigma_j, X\right) &= \sup_{\sigma_i} \tilde{w}_t^i \left(\sigma_i, \sigma_j, \tilde{h}_i^t, X\right) \\ &= \sup_{\sigma_i} \int_{h_j^t} u_i \left(\sigma_i, \sigma_j | \left(\tilde{h}_i^t, X\right), h_j^t\right) d\beta^{\sigma_j} \left(h_j^t\right). \end{aligned}$$

Since $\sigma_j \in M_j$, candidate j's continuation strategy depends only on $\theta_t = (X, X)$. Hence,[71] for each h_j^t, we can write

$$\sup_{\sigma_i} u_i \left(\sigma_i, \sigma_j | \left(\tilde{h}_i^t, X\right), h_j^t\right) = \sup_{\sigma_i \in M_i} u_i(\sigma_i, \sigma_j | \theta_t = (X, X)).$$

Therefore,

$$\tilde{w}_t^i \left(\sigma_j, X\right) = \sup_{\sigma_i \in M_i} u_i(\sigma_i, \sigma_j | \theta_t = X) = v_{i,t} \left(X, X\right).$$

This is a contradiction. Thus, for each PBE σ, we have $w_t^i \left(\sigma, X\right) = v_{i,t} \left(X, X\right)$. ■

[71] This is a standard result in dynamic programming. See Gensbittel, et al. (2017) for the application of this result to the game with Poisson arrivals. Although their paper assumes public monitoring, since σ_j is Markov, the observability of the Poisson arrivals does not affect the formulation of Bellman equations.