

Optimal Timing of Policy Announcements in Dynamic Election Campaigns

Yuichiro Kamada and Takuo Sugaya*

February 23, 2020

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 and a multi-dimensional policy space. We then present general principles that underlie the results from those models. 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. We also thank the careful reading and suggestions by the Editor and the referees. 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 This paper subsumes and supersedes part of our unpublished working paper “Valence Candidates and Ambiguous Platforms in Policy Announcement Games” and its earlier version circulated as “Policy Announcement Game: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.

1 Introduction

For every election, there is a campaign. The campaigns last from a few weeks to well over a year in the United States and they seem to be an important determinant of the election outcome. The campaigns themselves are inherently dynamic, with candidates trying to time their policy positioning to outmaneuver their opponent. Joe Slade White, a National Democratic Strategist of the Year, states that “timing makes the difference between winning and losing,” in one of his “9 Principles of Winning Campaigns” (White, 2012).

Despite the apparent importance of campaigns on election outcomes, and the fact that the campaigns are dynamic in nature, the literature has so far not provided theoretical models to understand the incentives and forces at work in dynamic election campaigns.¹ The objective of this paper is to fill this gap by proposing a model in which candidates face dynamic strategic considerations and obtain predictions about timing of policy announcements.²

The paper proposes a “policy announcement timing game” in which candidates strategically choose the optimal timing of their policy announcements over a campaign period. To capture why timing matters, we introduce a novel yet simple friction: opportunities for policy announcements are limited and arrive stochastically. Specifically, we assume that opportunities arrive according to a Poisson process over a continuous-time campaign period.³ The process continues until the predetermined deadline (the election day) is reached, and the final policy announcements before the deadline determine the result of the election. When an opportunity arrives in a campaign, a candidate must decide whether to take a position (and thereby commit to it) now or to wait for a better time, at the risk that another opportunity may not arise. This stochastic arrival of opportunities may reflect the inattention of voters, the vagaries of the media spotlight, or any other

¹By a model of a dynamic election campaign, 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.

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

³The time being continuous is not essential to our results because any “limit” of the equilibria as the discrete-time models “approach” the continuous-time model must be an equilibrium in the continuous-time model (Moroni [2019]). A discrete-time setting, however, is intractable and the resulting comparative statics would be messier, so the empirical usefulness of the framework would be questionable.

number of campaign variables that candidates cannot control.

It turns out that this framework allows us to obtain predictions about the timing of policy announcements, which would not be possible in the standard static election models. Through a number of applications, we demonstrate that such predictions can be made not only for standard Hotelling models, but also for a wide variety of election models that the literature often rules out.

Our first such application of the framework is to an election model with valence candidates (Section 2.1). The underlying election model is the standard Hotelling one, except that one candidate is stronger than the other, so the stronger wins if the two candidates take the same policy position. The strong candidate always wants to copy the weak candidate’s policy—as her valence advantage will determine the election—while the weak candidate does not want to be copied, just as in the “matching pennies” game. If there is no asymmetry across candidates, that is, if the dynamic framework is applied to the standard Hotelling election model, then such “matching pennies”-like feature disappears, and we show that candidates position themselves at the Condorcet winner as soon as they obtain an opportunity to do so. With the asymmetric valence, however, the incentives are more complicated and we obtain a distinct prediction about the timing of policy announcement. Specifically, it is suboptimal for the weak candidate to announce his policy too early since the strong candidate would then have enough time to copy that policy afterward. At the same time, if both candidates do not announce their policies, then the weak candidate will lose the election due to his valence disadvantage. Hence, waiting until the last moment is too risky given the possibility of not having another opportunity. This tradeoff leads him to taking a simple cutoff strategy in equilibrium: He does not clarify his policy position until a time threshold is passed, after which the risk of not obtaining another opportunity outweighs the risk of being outmaneuvered, and he announces his policy when an opportunity arrives. The key logic here is that the weak candidate has what we call the “first-mover disadvantage”: If he fixes his location, the other candidate can always find a position that beats it (in this specific setting with valence candidates, such a position is the same position as the weak candidate’s).⁴

The result may explain the dynamics of the election campaign in the 2014 gubernatorial election for Tokyo, Japan, in which Yoichi Masuzoe won against Morihiro Hosokawa. Although Masuzoe had been seen as the stronger candidate from the outset of the campaign, Hosokawa became popular

⁴The strong candidate’s strategy depends on the parameter value. For some parameter region, she would not clarify her policy until the weak candidate does so.

near the election day when he clarified his stance by announcing opposition to the restart of nuclear power generation in the wake of Fukushima. After that, 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. He then became more popular again and won the election.^{5,6}

Our second application of the framework shows that the first-mover disadvantage is not limited to the weak candidate in the election with valence. It concerns election campaigns with a multi-dimensional policy space (Section 2.2), for which it has long been known that a Condorcet winner generically does not exist. In contrast to the case of the standard uni-dimensional Hotelling model, the nonexistence of a Condorcet winner implies that there does not exist a pure-strategy Nash equilibrium in the static environment, which has been a major obstacle in conducting an analysis with a multi-dimensional policy space in the static world.⁷ In our dynamic framework, however, we can derive an empirically testable prediction in such an environment. More specifically, when there is no Condorcet winner, if one candidate fixes her location, then the other can always find a position to beat it. Therefore, each candidate, upon making an announcement, “becomes a weak candidate” in the sense that being best-responded afterward will bring about the worst outcome. Hence, no candidate clarifies their policy when the election day is far ahead. At the same time, if a candidate knew that the current opportunity is the last one, she would prefer to clarify her policy (we assume voters prefer candidates with a clear policy announcement). As in the weak candidate of the model with valence, this leads each candidate to taking a simple cutoff strategy in equilibrium.

In these two applications, we assume that a candidate either clarifies all details of the policy position or clarifies nothing. Such a model specification rules out the possibility of progressive commitment. To obtain predictions about the timing of progressive commitment, we introduce a third application (Section 2.3). We extend the model with valence candidates to a model with two policy issues and allow candidates to announce a policy for each issue. We show that the weak candidate has an incentive to partially commit, i.e., clarify a policy for one issue while not

⁵Sankei 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.

⁶See Remark 5 in Section 2.1 for a further real-life example to which our model fits.

⁷The past approach in the literature of multi-dimensional policy space has been to change the rule of the election game in the static context. See, e.g., Lindback and Weibull (1987), Coughlin (1992) and Roemer (2001).

announcing a policy for the other issue. The intuition for such partial commitments can be explained by the following “option value” argument: A partial commitment enables the weak candidate to differentiate from the strong candidate, yet at the same time retain the flexibility to maneuver around the strong candidate should she copy his policy. We show that the weak candidate’s equilibrium strategy has a simple cutoff structure when the strong candidate has yet to clarify her policy positions: he does not clarify his policy position in any issue before the cutoff is passed, while he clarifies his position for one issue after the cutoff. Once the strong candidate copies the weak candidate’s policy for one issue, the situation becomes the same as the one-issue case: The weak candidate waits until another threshold is passed and then clarifies his policy for the remaining issue. This progressive commitment shows more generally how the first-mover disadvantage affects incentives over a long campaign, and corresponds to the reality that candidates convey their policy intent about some issues to voters but not about all issues at once.

In each of the three applications we present, we obtain a unique prediction about the timing of policy announcement, which we show to be characterized by cutoff strategies. This tractability allows us to pin down both the equilibrium probability distribution of times at which candidates make policy announcements and winning probabilities. This enables us to conduct comparative statics, which would be helpful in empirically testing the model.⁸ Given that our policy announcement timing game is the first to analyze the dynamics of election campaigns, we view those comparative statics as novel kinds in the literature.

After discussing the three applications, we consider a general model of election campaign and present general principles that underlie the results from those settings. One of the robust predictions about announcement timing in all three applications of our framework is that candidates spend a long time not clarifying their policy position—using what we call ambiguous language—in face of first-mover disadvantage. We present a general result that formalizes the intuition that the first-mover disadvantage incentivizes a candidate to delay their policy announcement. 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, candidates spend most of the time not announcing any specific policy in equilibrium, provided the campaign period is long enough.

When there is a Condorcet winner, candidates do not face first-mover disadvantage. In the

⁸Those comparative statics are discussed in Remark 2 in Section 2.1, Remark 8 in Section 2.2, and Remark 10 in Section 2.3.

applications, we show that in such an environment, candidates clarify their policies early in the campaign. Section 4.2 generalizes this insight. We show the *dynamic median-voter theorem*, which generalizes the celebrated median-voter theorem to our dynamic setting. Specifically, we suppose that there is a Condorcet winner in the static version of the model and show that each candidate makes an announcement corresponding to the Condorcet winner as soon as possible. This theorem implies that candidates announce their policies as soon as possible in the absence of valence and when the policy space satisfies a certain symmetry condition.

Our election models feature purely office-motivated candidates, hence they are constant-sum games. We present a general result in Section 4.3 for dynamic campaigns for such elections. Specifically, we prove that any perfect Bayesian equilibrium has the Markov property when the election features constant-sum (as in all of our three applications). We call this the *constant-sum Markov theorem*. This theorem guarantees that, in order to calculate the candidate’s payoff—and hence to calculate the probability of winning, it suffices to focus on Markov perfect equilibria, which is a simpler task than considering all (possibly non-Markov) equilibria.

Our long ambiguity theorem serves as a novel explanation for ambiguity in campaigns. Ambiguity in our model is due to dynamic strategic consideration: Candidates seek to move after their opponent does. In contrast, the literature on ambiguity has focused on the standard static model of elections and seeks to explain why candidates may choose not to take a precise policy position. For example, it is due to voters’ preferences such as risk-loving preferences (Shepsle [1972] and Aragonès and Postlewaite [2002]), context dependent preferences (Callander and Wilson [2008]), or due to the fact that ambiguous policy statements make the candidate appealing to many voters (Glazer [1990]). Another explanation is candidates’ limited resources or voters’ limited capacity (Page [1976, 1978], Polborn and Yi [2006], Egorov [2015], and Dragu and Fan [2016]), candidates’ incentive to reserve freedom to choose a policy after being elected or in a next election (Alesina and Cukierman [1990], Aragonès and Neeman [2000], and Meirowitz [2005]), or interaction with policy-motivated donors (Alesina and Holden [2008]).

We assume that opportunities to take policy positions arrive only intermittently, and this assumption is substantive. Election campaigns are buffeted and upended by all sort of forces and actors. For example, administrative procedures to obtain internal approval for a change of what policy a candidate announces may not always be successful. Adopting a policy position requires

that the voting public hears the announcement and is able to digest it. The difficulty of these tasks, and the necessity of conveying the information through media, suggests that opportunities to successfully transmit the messages are limited and arrive at unpredictable times. It is this controlled chaos that our model hopes to capture in a simple way.

Ultimately, our interest is not in the source of these communication difficulties but in their impact on strategic candidates and election outcomes. As such, we take as given that opportunities to take policy positions are limited and arrive stochastically, and explore the consequences. An analogy to the macroeconomics literature may help: 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).⁹ In the present paper, we also show that this Poisson approach is useful to analyze campaign dynamics. We hope a micro-foundation is investigated in future research.^{10,11}

Moreover, the specific way by which we model frictions—using Poisson processes—is proven to be fruitful in other contexts. The assumption of Poisson opportunities has been extensively studied recently in the model called “revision games” (Kamada and Kandori [2019] and Calcagno et al. [2014]).¹² Recent applications using Poisson processes to analyze dynamics include bargaining (Ambrus and Lu [2015]), eBay-like auctions (Ambrus et al. [2014], Kapor and Moroni [2016], and Hopenhayn and Saeedi [2016]). Its validity is examined in experiments by Roy (2014) and Avoyan and Ramos (2018).

The Appendix provides the main proofs for the general results presented in Section 4 as well as additional discussions. All the proofs not provided in the main text or the Appendix are provided in the Online Appendix.

⁹See Klenow and Malin (2010).

¹⁰A possible micro-foundation of the Poisson-type announcement possibility appears in García-Jimeno and Yildirim (2017) in which, in equilibrium, the media stochastically follow a candidate due to its interest in reporting the candidate’s controversial statements.

¹¹Another way to model frictions is to introduce switching costs. See Lipman and Wang (2000) and Caruana and Einav (2008) for models with switching costs in finite-horizon games. Caruana and Einav’s (2008) “grid invariance” results do not hold in our constant-sum setup, so directly applying their framework to our context may not be fruitful. We could instead assume that at each stochastic arrival of an opportunity to announce a policy, there is a cost of announcing a policy that is inconsistent with the previous announcement. We checked that our “long ambiguity” result is robust to such a modification, at least in the model of Section 2.2.

¹²See Ishii and Kamada (2011) and Romm (2014) for extensions.

2 Three Models of Dynamic Election Campaign

This section offers three models to show that the model of a policy announcement timing game enables us to analyze rich strategic considerations when it is applied to otherwise well-known and canonical models of elections. In Section 3, we define a policy announcement timing game, which unifies the three models, and in Section 4 we present general principles that underlie those findings.

All the models that we present here have timing and payoff structures in common: Two candidates receive stochastic opportunities to announce their policies, at which they can either keep announcing the ambiguous policy or clarify their position. Formally, 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 obtains opportunities to announce her “policy set.” A policy set describes a policy announcement by a candidate, and it is a subset of the entire policy space, which we denote by X . The collection of available policy sets varies depending on the model. At each opportunity, each candidate can only announce a subset of the policy set that she has previously announced. Intuitively, this assumption implies that each candidate cannot “flip-flop” her policy.¹³ The policy set at time $-T$ is exogenously given to be the entire policy set X .

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 policy set.

The candidates’ payoffs are determined according to the finally announced policy set profiles as of the time of the election, which is time 0. Candidates are purely office-motivated. Formally, the winning candidate obtains a payoff of 1, while the losing candidate obtains a payoff of 0. Each candidate’s objective is to maximize the expected payoff, that is, their objective is to maximize their probability of winning.

¹³This assumption, which we impose throughout the paper, is a key difference from most of the papers in the revision games literature in which players can costlessly choose their actions from a fixed action space at each opportunity to move. Settings similar to the current model appear in the “monotone games” in which players can only (weakly) increase their actions (Gale, [1995, 2001]), and the “commitment games” in which each player simultaneously commits to a subset of the entire action space and then plays a game (see, for example, Hamilton and Slutsky [1990, 1993]; van Damme and Hurkens [1996]; Romano and Yildirim [2005]; Renou [2009]). These papers are sufficiently different from ours in their stage games as well as the timing of moves.

2.1 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 2.1.1 introduces the model. In Section 2.1.2, we establish that if two candidates are perfectly symmetric, then both candidates clarify their policy position as soon as possible in equilibrium. In Section 2.1.3, 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 clarifies his policy, 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.

2.1.1 The Model

The two candidates are the strong candidate, S , and the weak candidate, W .¹⁴ The policy space is kept simple, so as to highlight the complexity introduced by the campaign phase into an election model with valence candidates. Specifically, the policy space is assumed to be $X = \{0, 1\}$, and the available policy sets are $\{0\}$, $\{1\}$, and $\{0, 1\}$. This is the minimal environment in which ambiguity is a possibility.¹⁵

If a candidate enters at 0 (or 1) and the other enters at 1 (or 0) or does not enter, then 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. We assume $p \in (0, \frac{1}{2})$.

These winning probabilities can be micro-founded by considering a model of sincere voting in which voters derive utility mostly from the implemented policy but they derive small extra utility from the winning candidate being strong. More specifically, suppose that there are a continuum of voters, located at policy 0 and policy 1. The distribution of the voters’ locations is stochastic, and policy 0 has more voters with probability p , and policy 1 has more voters with probability $1 - p$. During the campaign, the locations of the voters are unknown.

¹⁴For ease of exposition, we use feminine pronouns to refer to S , A and i and masculine pronouns to refer to W , B and j .

¹⁵A working paper version of this paper (Kamada and Sugaya [2019]) presents a more general model that involves many other cases, such as a continuous policy space.

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.¹⁶ 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 $\frac{1}{2}$.¹⁷ The voters are sincere, that is, they each vote for the candidate who, if elected, maximizes their expected payoff. The candidate that attracts a larger share of votes wins. In the case of a tie, each candidate wins with probability $1/2$.

The payoff function of the candidates that we specified earlier can be obtained by assuming $\delta > 0$. For example, if the policy sets of the two candidates are the same at time 0, then S 's utility is 1 and W 's is 0. If S 's announcement at time 0 is $\{0\}$ and W 's is $\{0, 1\}$, then S 's utility is p and W 's is $1 - p$.

We call the dynamic game with arrival rates $\lambda_S = \lambda_W =: \lambda$ with the above specification a *valence election campaign*. It is characterized by a tuple (p, T, λ) . We will consider perfect Bayesian equilibria (PBE) of this game.¹⁸

2.1.2 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

¹⁶One 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.

¹⁷It 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.

¹⁸The formal definitions of histories, strategies, and equilibrium are given in the Online Appendix.

$\delta = 0$ 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 $\{0, 1\}$. Note that, in the static version of this game, we can apply the median voter theorem: It is each candidate's dominant action to announce $\{1\}$.

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 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 the symmetry of the supposed continuation strategies. Thus, entering at 1 is a unique best response. Therefore, by the continuity of probabilities in time and boundedness of the payoffs which imply the 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. These observations imply the desired result.¹⁹

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.

2.1.3 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 $x' \in \{0, 1\} \setminus \{x\}$. Since the median is more likely to be at policy 1 ($p < \frac{1}{2}$), these two conclusions imply that, if a candidate enters before the opponent, he/she enters at policy 1.

¹⁹That these observations imply the result follows from the "continuous-time backward induction" that we formally present in Appendix B.

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.*

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 $p = \frac{1}{2}$.²⁰ Suppose that at time $-t$, both S and W have previously announced $\{0, 1\}$. On the one 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 (it happens when W does not enter afterward and when W enters and S copies his policy).

On the other hand, if there is no further revision, W 's payoff is 0 (because S does not enter). 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.

If W announces $\{1\}$ in the early stages of the campaign, then the probability with which S enters afterward is high. Therefore, 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. It turns out that there exists a ‘‘cutoff,’’ $-t^*$, until which W announces $\{0, 1\}$ and after which W announces $\{1\}$ when he gets an opportunity for a policy announcement.

²⁰Strictly speaking, since the model assumes $p < \frac{1}{2}$, 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.

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 for sure, 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 $\{1\}$ as soon as possible.

The next proposition fully characterizes the form of PBE for each $p \in (0, \frac{1}{2}) \setminus \{\frac{1}{1+e}\}$.²¹ 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 $p > \frac{1}{1+e}$ (considered in part 1 of Proposition 2), S does not enter until W enters for a similar reason to in the case of $p = \frac{1}{2}$. In contrast, for $p < \frac{1}{1+e}$ (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 $\{0, 1\}$.²²*

1. If $p > \frac{1}{1+e}$, the following hold.²³

²¹The notation e stands for the base of the natural logarithm.

²²If $p = \frac{1}{1+e}$, then there is indeterminacy about S 's equilibrium strategy at all $-t < -t^*$ since she is indifferent. The Online Appendix presents a characterization of the equilibrium dynamics in this case as well.

²³Although the entire game lasts for the time interval $[-T, 0]$, we state results for all times in $(-\infty, 0]$. Any

- (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 $\frac{dt_W}{dp} < 0$.

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.²⁴ 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 I, 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. 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.

[INSERT FIGURE I HERE]

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.

²⁴Gensbittel et al. (2017) analyze general constant-sum revision games. The result we present here, as well as the long ambiguity results in other sections, is similar to their “wait and wrestle” property. One difference is that we consider the case where announcing a singleton policy set is irreversible, while players can freely choose their action from a fixed action set in the model of Gensbittel et al. (2017). Note that, although our specification implies that there is no cycling choice of actions, it is still not trivial that candidates wait for a long time. Gensbittel et al. (2017) also discuss a comparison between the two models. Another difference is that we obtain long ambiguity in non-constant-sum games as well (see the working paper version of this paper, Kamada and Sugaya [2019]).

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 the probability that S uses ambiguous language is close to 0 in equilibrium. This hinges on our assumption that even if W enters after S , S does not incur any loss. One can 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.²⁵ 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 (Contribution to the literature on valence candidates). 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.²⁶ The second approach is that of Aragonès and Palfrey (2002), who consider the simultaneous-move game and

²⁵A working paper version of this paper (Kamada and Sugaya [2019]) shows this formally.

²⁶See also Ansolabehere and Snyder (2000) and Groseclose (2001) who consider pure-strategy equilibria in models with valence candidates.

characterize a mixed equilibrium.²⁷ 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 about electoral situations with valence candidates under the timing structure observed in many real elections, in both these models the order of policy announcements is exogenously given by the modelers. In contrast, we view our result in this section as *endogenizing* the order of policy announcements.²⁸ □

Remark 2 (Empirical implications). 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.²⁹ 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 3 (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 result is due to the fact that the underlying game is constant-sum.³⁰ It 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 decreases his/her commitment power.³¹ □

²⁷More 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).

²⁸This provides a possible answer to the question posed by Aragonès and Palfrey (2002), who ask “What is the correct sequential model.”

²⁹Specifically, 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 .

³⁰Gensbittel et al. (2017) also prove that the equilibrium payoff is weakly increasing in the relative arrival rate for constant-sum revision games (unlike in our model, players can revise actions back and forth).

³¹More detailed discussions about heterogeneous arrival rates and a general model with heterogeneous arrival rates and a general class of payoff functions are provided in a working paper version of of this paper (Kamada and Sugaya [2019]). One notable insight arising from such an extension is that, in one of the variants of our model discussed in the working paper, we show that ambiguity occurs even when voters have concave utility functions. This is in a stark contrast with the result in Shepsle (1972) and Aragonès and Postlewaite (2002) that we discussed in the introduction, in which convex voter utility functions are assumed to obtain ambiguity results.

Remark 4 (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 (or there are multiple voters whose preferences are homogeneous and are subject to aggregate shocks). 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 ($p > \frac{1}{1+\epsilon}$), the valence term $\delta > 0$ is sufficiently small, and T is sufficiently large. The Online Appendix formally proves this claim, and explains that this is due to the fact that the policy announcement timing game results in a positive correlation between candidates’ positions, which is ex ante not desirable for the median voter (because the probability that there is a candidate at the median voter’s bliss point is small in the presence of such correlation). \square

Remark 5 (Copying of another candidate’s policy). The heart of the incentives the candidates face in this application is that one candidate wants to copy the policy of the other candidate. Such copying behavior may not be realistic when the main political issue for the campaign pertains to the identifying characteristics of the candidates’ parties, such as right wing’s opposition to abortion and gun control. They are, however, more realistic for new issues to which candidates are not necessarily associated with a specific agenda. The issue regarding nuclear power in the wake of Fukushima discussed in the Introduction was such an example. Another example along this line is the promise by the Conservative Party to have a Brexit referendum. In this context, the Conservative Party was an incumbent, and faced UKIP (UK Independence Party), a small right wing party. UKIP’s eurosceptic idea was popular among right wing voters, and the Conservative Party had been losing voters to UKIP. This rise of UKIP is widely seen as the reason why David Cameron in the Conservative Party later promised that the party would hold a referendum if their party won the election in 2013 (see e.g., BBC News [2013] and Mance and Pickard [2016]).

Remark 6 (Incumbent’s commitment). In some elections, there is an incumbent who has already committed to a policy position. Indeed, some media mention that a candidate with a long career

has a disadvantage since he/she has made many policy decisions.³² To understand this problem, suppose that S is an incumbent and W is a challenger (as in Bernhardt and Ingberman [1985], Berger et al. [2000], and Carter and Patty [2015] mentioned in Remark 1) and S has committed to policy 1. In this game, the payoffs for the candidates are $1 - p$ for S and p for W . Candidate S 's payoff in the game without such a commitment is higher than this when $p > \frac{1}{1+e}$ and lower when $p < \frac{1}{1+e}$. The reason is as follows. Without the commitment, when $p > \frac{1}{1+e}$ and the election day is far away, S at her move has a choice to commit to policy 1 but she strictly prefers not to. This means that her payoff is strictly higher than the payoff of committing, which is $1 - p$. In contrast, again without commitment, when $p < \frac{1}{1+e}$ and the election day is far away, S enters at policy 1 as soon as possible. This means that the payoff in the game without the commitment is a convex combination of $1 - p$ and the payoff from being unable to enter, and it is lower than $1 - p$ because S strictly prefers to enter at policy 1 as soon as possible. Moreover, when $p > \frac{1}{1+e}$, the cost of the commitment, which we define to be the difference of S 's payoffs in the two scenarios, can be shown to be strictly increasing in p .³³ This makes sense because S 's payoff in the election campaign game is strictly increasing in p as W 's payoff when W enters and S cannot copy is decreasing in p , while S 's payoff from the commitment strictly decreases in p as policy 1 is less attractive for a higher p . The formal analysis is provided in the Online Appendix. \square

2.2 Multi-dimensional Policy Space

When the policy space is multi-dimensional, generally there does not exist a Condorcet winner, and a pure-strategy Nash equilibrium does not exist in a static two-candidate Hotelling model. The literature has struggled with this nonexistence issue. 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. The ambiguity again results from a disadvantage of being the first-mover, which follows from the nonexistence of a Condorcet winner.

³²For example, Anita Dunn, a top aide to Obama's campaign in 2008, discusses disadvantages of serving in the Senate before running for presidency as a candidate with such a career has cast votes in the Senate (Cadigan and Struyk [2018]).

³³It is $p - (1 - p)e^{-1}$.

The two candidates are A and B . Suppose that the voters are distributed according to a measure $\mu \in \Delta(\mathbb{R}^n)$ for some $n \in \mathbb{N}$. Let the policy space X be the support of μ .³⁴ We assume that X is a full-dimensional connected subset of \mathbb{R}^n . The collection of available policy sets consists of the entire policy set X and all the singleton sets of the form $\{x\}$ for $x \in X$.

We assume sincere voting, where 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. Formally, given a policy profile $(x_A, x_B) \in X \times X$, we define the set of supporters for each candidate as:

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

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 $P_B(x_B, x_A) = 1 - P_A(x_A, x_B)$.

Each problem is characterized by a pair (X, μ) . Let the set of all problems be \mathcal{P} . Define \mathcal{M} to be the set of problems such that the so-called ‘‘Plott conditions’’ are met. Formally, we define:

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

Notice that, if $(X, \mu) \in \mathcal{M}$ and no Y that has zero-Lebesgue measure satisfies $\mu(Y) > 0$,³⁵ then there exists x^* such that any hyperplane that passes x^* divides X into two regions which both have measure $\frac{1}{2}$ according to μ (the weak inequality in the definition of \mathcal{M} can be replaced with an equality in this case). The definition of \mathcal{M} also deals with the case in which some hyperplane has positive measure according to μ . To understand the property of \mathcal{M} , 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 restriction. 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 .³⁶ Third, when $(X, \mu) \in \mathcal{M}$,

³⁴That is, $x \in X$ if and only if $\mu(U) > 0$ for any open neighborhood U of x .

³⁵Together with the definition of X , this means that for any $Y \subseteq X$, $\mu(Y) > 0$ if and only if Y has positive Lebesgue measure.

³⁶See Theorem 7.2 of Roemer (2001) for the detail.

the x^* satisfying the condition in the definition of \mathcal{M} is uniquely determined because X is connected. We denote this unique x^* by $x^*(X, \mu)$.

To define the vote share and winning probabilities for the case in which some candidate does not enter in the policy announcement timing game, we expand the domain of S_i with a restriction that the winning probabilities are half-half if no one has entered. Formally, we assume $S_i(X, X) = \emptyset$ for each $i = A, B$. The domain of P_i is expanded accordingly for each $i = A, B$.³⁷ The payoffs are $P_i(X_i, X_j)$ for each announcement profile (X_i, X_j) and $i = A, B$. We assume that there exists $\bar{x} \in X$ such that $P_i(\{\bar{x}\}, X) = 1$ for each $i = A, B$, which implies that clarifying a policy position is better than being ambiguous when the other candidate is ambiguous. If $(X, \mu) \in \mathcal{M}$, we further assume $P_i(\{x^*\}, X) = 1$ for each $i = A, B$. Moreover, (whether or not $(X, \mu) \in \mathcal{M}$,) 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. We call the dynamic game with the above specification 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 $\text{sign}(\lambda_A - \lambda_B) = \text{sign}(t_A^* - t_B^*)$.*

Remark 7 (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 .³⁸ Hence, the proposition shows that ambiguity emerges in a PBE if and only if there

³⁷That 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 $P_B(X_B, X_A) = 1 - P_A(X_A, X_B)$.

³⁸The 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).

is no pure-strategy Nash equilibrium in such a static game. □

[INSERT FIGURE II HERE]

Part 1 of Proposition 3 implies that, if there exists a Condorcet winner, then it is optimal to announce that policy as soon as possible. In Figure II, we depict the times t_A^* and t_B^* that appear in part 2. Intuitively, each candidate's strategic situation is similar to that of the weak candidate in the valence election campaign (Section 2.1): 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 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, it is optimal to enter since the probability of such an event is small.

Remark 8 (Empirical implications). 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 ($t_B^* < t_A^*$).

Another 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. To prove the claim, we consider the following deviation by candidate A , while fixing B 's strategy: As long as B 's current announcement is ambiguous, A does not enter until $-t_B^*$, and tries to enter with probability $\frac{\lambda_B}{\lambda_A}$ whenever she receives an opportunity afterward. She, however, enters as soon as possible after B 's entry. This strategy is symmetric to B 's strategy except that A can move “more quickly” after the opponent’s entry than B can. Due to this asymmetry, A 's payoff after the deviation is strictly higher than B 's payoff. Such a deviation would not raise A 's payoff, so provides a lower bound of her PBE payoff. We combine this lower bound with the assumption that the election is constant-sum and prove the proposition. \square

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 $-t \in (-t_i^*, 0]$.*

2.3 Partial Commitment

The models presented so far assumes that either the candidate stays ambiguous or completely specifies a policy. This is a natural generalization of the standard Hotelling model which usually models candidates’ move as a choice of a single policy platform. However, once we introduce the possibility of ambiguity, one may want to further ask whether the possibility of partial commitments changes the prediction. This section considers such a question and shows that the main insight—that the ambiguous language is used for a long time—is still valid, although partial commitments may occur on the path of equilibrium play. More specifically, we consider an extension of the

valence election campaign model (Section 2.1) to the case with a two-dimensional policy space $\{0, 1\}^2$. Each candidate can specify a policy in one or two policy issues (dimensions), where for example $\{1\} \times \{0, 1\}$ corresponds to specifying a policy in the first issue while being ambiguous in the second issue. It is shown that, on the path of equilibrium play for some parameter regions, the candidates spend a long time keeping the policy announcement ambiguous, while the weak candidate has an incentive to partially specify a policy close to the election date. The Democratic primary election for the 2008 US presidential election may serve as an example in which such a prediction applies. Obama had a clear position on health care on September 2007 at the latest, which is when Mitt Romney used the term “Obama-care,” while it was October 2007 that Obama clarified his position on LGBT rights.³⁹

The two candidates are S and W . The policy space is $X = \bar{X}^2 = \{0, 1\}^2$. There is a single voter who is located at $(x_1, x_2) \in \{0, 1\}^2$ with the following probability:

	$x_2 = 0$	$x_2 = 1$
$x_1 = 0$	$p_1 p_2$	$p_1 (1 - p_2)$
$x_1 = 1$	$(1 - p_1) p_2$	$(1 - p_1) (1 - p_2)$

We assume that the uncertainty about the voter’s location is higher in the second issue, i.e., $p_1 < p_2 < \frac{1}{2}$. When the voter is located at (x_1, x_2) , her utility from candidate i with policy (y_1, y_2) winning is

$$-\sqrt{(y_1 - x_1)^2 + (y_2 - x_2)^2} + \delta \cdot \mathbb{I}_{i=S},$$

where $\delta \in (0, \frac{\sqrt{2}}{4})$ represents valence. We allow candidates to gradually clarify their policy announcements (policy sets) issue by issue. Formally, when a candidate’s policy set is $X_{i,1} \times X_{i,2}$, we say that her policy set is $(X_{i,1}, X_{i,2})$. We assume that when a candidate’s most recent policy set is $\theta_i = (X_{i,1}, X_{i,2})$, she can announce a policy set $\tilde{\theta}_i = (\tilde{X}_{i,1}, \tilde{X}_{i,2})$ with $\tilde{X}_{i,1} \subseteq X_{i,1}$ and $\tilde{X}_{i,2} \subseteq X_{i,2}$ such that $\tilde{X}_{i,1} \neq \emptyset$ and $\tilde{X}_{i,2} \neq \emptyset$. Since the initial policy set (at time $-T$) is X , this determines the collection of available policy sets. We write $\tilde{\theta}_i \subseteq \theta_i$ for simplicity in this case. Oftentimes we also abuse notation to write, for example, $(1, \bar{X})$ to mean the policy set $(\{1\}, \bar{X})$.

We assume that, given the policy set θ_i announced at the last opportunity before time 0,

³⁹Reeve (2011) and CNN Policy Ticker (2007).

the voter believes that the candidate takes a policy according to the uniform distribution over $\{(x_{i,1}, x_{i,2}) : (\{x_{i,1}\}, \{x_{i,2}\}) \subseteq \theta_i\}$. She votes for the candidate who brings a higher expected utility, while she votes for each candidate with probability 1/2 in the case of a tie.

This determines the payoff function.⁴⁰ For example, suppose that candidate S 's policy set is (\bar{X}, \bar{X}) , candidate W 's policy set is $(1, 1)$, and the voter is located at $(0, 0)$. The voter's payoff from candidate S is

$$-\frac{1}{4} \left(0 + 1 + 1 + \sqrt{2} \right) + \delta = -0.85355 + \delta,$$

and her payoff from W is $-\sqrt{2} = -1.4142\dots$. Hence, she votes for S . The similar calculation shows that candidate S wins if and only if the voter is located at $(0, 0)$, $(0, 1)$, or $(1, 0)$, which happens with probability $1 - (1 - p_1)(1 - p_2)$. Thus, S 's payoff is $1 - (1 - p_1)(1 - p_2)$ and W 's is $(1 - p_1)(1 - p_2)$.

We call the dynamic game with arrival rates $\lambda_S = \lambda_W =: \lambda$ with the above specification a *multi-issue election campaign with valence*. It is characterized by a tuple (p_1, p_2, T, λ) .

We always write candidate S 's policy first when we list both candidates' policies. For example, when we say the current policy announcement is $((\bar{X}, \bar{X}), (1, \bar{X}))$, it means that S 's current policy announcement is (\bar{X}, \bar{X}) while W 's is $(1, \bar{X})$.

We focus on the case where $p_k \geq \frac{1}{1+e}$ for each $k \in \{1, 2\}$. Given the analysis of the one-issue case in Section 2.1, this condition implies that, once both candidates enter at (x_k, \bar{X}) , candidate S never moves until candidate W enters at $(x_k, 1)$, and candidate W enters at $(x_k, 1)$ for time $-t \in (-\frac{1}{\lambda}, 0]$ and stays at (x_k, \bar{X}) for time $-t \in (-\infty, -\frac{1}{\lambda})$. The following proposition characterizes the candidates' behavior under the events that happen on the path of equilibrium play:

Proposition 5. *Consider a multi-issue election campaign with valence with (p_1, p_2, T, λ) such that $p_k \geq \frac{1}{1+e}$ for each $k \in \{1, 2\}$. There exists a PBE, and the following hold for any PBE, where $t_W^* := \frac{1}{\lambda} \frac{5-6p_1-p_2}{3(1-2p_1+p_2)} > \frac{1}{\lambda}$.*

1. *Given $((\bar{X}, \bar{X}), (\bar{X}, \bar{X}))$, (i) candidate S stays at (\bar{X}, \bar{X}) for each $-t \in (-\infty, 0]$ and (ii) candidate W stays at (\bar{X}, \bar{X}) for each $-t \in (-\infty, -t_W^*)$ and enters at $(1, \bar{X})$ for each $-t \in (-t_W^*, 0]$.*

⁴⁰We put the full payoff matrix in the Online Appendix.

2. Given $((\bar{X}, \bar{X}), (1, \bar{X}))$, (i) candidate S enters at $(1, \bar{X})$ for each $-t \in (-\infty, 0]$ and (ii) candidate W stays at $(1, \bar{X})$ for each $-t \in (-\infty, 0]$.
3. Given $((1, \bar{X}), (1, \bar{X}))$, (i) candidate S stays at $(1, \bar{X})$ for each $-t \in (-\infty, 0]$ and (ii) candidate W stays at $(1, \bar{X})$ for each $-t \in (-\infty, -\frac{1}{\lambda})$ and enters at $(1, 1)$ for each $-t \in (-\frac{1}{\lambda}, 0]$.
4. Given $((1, \bar{X}), (1, 1))$, candidate S enters at $(1, 1)$ for each $-t \in (-\infty, 0]$.

[INSERT FIGURE III HERE]

In Figure III, we depict the times t_W^* and $\frac{1}{\lambda}$ that appear in Proposition 5. As in the one-issue case, candidate S moves only to match her state with candidate W . Candidate W also stays at (\bar{X}, \bar{X}) until some time cutoff, since moving too early increases the risk of being copied by candidate S . Candidate W first enters at $(1, \bar{X})$, keeping the possibility of entering at $(1, 1)$ afterward even if candidate S successfully copies $(1, \bar{X})$. Since candidate W leaves this possibility, the cutoff to enter at $(1, \bar{X})$ comes earlier compared to the cutoff to enter at 1 in the one-issue problem (that is, $t_W^* > \frac{1}{\lambda}$).

Remark 9 (Empirical implication). The multi-issue model in this section has an empirical implication that we did not have in the one-issue model (Section 2.1): Candidates clarify their policy positions for the issue k for which the voters' preferences are "more certain," in the sense that p_k is small. The intuition for why this happens is roughly as follows: W 's benefit from delaying the clarification of a position for issue 1 until a position for issue 2 is clarified would arise after the first clarification (clarifying a policy for issue 2) is copied by S . But W can time the first clarification of his policy position to lessen the probability of it being copied, so this benefit is not so large. \square

Another difference from the one-issue case is that the cutoff t_W^* depends on p_1 and p_2 while the cutoff $\frac{1}{\lambda}$ in the one-issue model is independent of p . In the one-issue model, given that S never enters unless W enters, the only event in which W can obtain a positive payoff is that he enters at

1 and S cannot enter afterwards. Since the probability of this event is independent of p , his cutoff is also independent of p . In the two-issue case, since W has multiple cases in which he can win with positive probability and their payoffs are different and depend on (p_1, p_2) , the cutoff t_W^* depends on (p_1, p_2) .

Remark 10 (Comparative statics). Given the dependence of t_W^* on (p_1, p_2) , we obtain the following comparative statics: *The cutoff t_W^* in Proposition 5 is strictly increasing in p_1 and strictly decreasing in p_2 .* This in particular implies that it is more likely for us to observe the situation in which both candidates are completely ambiguous when the popular policy issue is clearer in one issue (with a constraint that it is not too clear—we need $p_1 \geq \frac{1}{1+e}$) while it is highly uncertain in the other issue.

To understand why t_W^* is increasing in p_1 , note that candidate W 's payoff when the realized policy profile is $((\bar{X}, \bar{X}), (1, \bar{X}))$ is $1 - p_1$. This implies that an increase of p_1 reduces W 's continuation payoff from entering at $(1, \bar{X})$ as well as staying at (\bar{X}, \bar{X}) . The payoff of $1 - p_1$ realizes when S cannot enter afterward in the former case, and when W enters and then S cannot enter afterward in the latter case. Since t_W^* can be shown to be not too close to the deadline, at time $-t_W^*$ for a given p_1 , the probability of W 's receiving $1 - p_1$ is larger in the latter case, and hence the reduction of the continuation payoff is larger in that case. Recalling that $-t_W^*$ is the time at which the two continuation payoffs are equal, this implies that the increase of p_1 pushes the cutoff t_W^* further from the deadline. Once $((1, \bar{X}), (1, \bar{X}))$ happens, then W 's best scenario is that the election campaign ends with $((1, \bar{X}), (1, 1))$ and he obtains $1 - p_2$. An increase in p_2 implies that this scenario (S copies $(1, \bar{X})$ and so W needs to differentiate himself in the second issue) becomes less attractive, so candidate W is more afraid of candidate S copying $(1, \bar{X})$. This incentivizes candidate W to skip opportunities until later. \square

Note that, in the proposition, we focus on the case where $p_k \geq \frac{1}{1+e}$ for each $k \in \{1, 2\}$ because this is the most interesting case that yields the result consistent with the reality of election campaigns provided in the Introduction. In the one-issue case, if $p < \frac{1}{1+e}$, there exists $t_S < \infty$ such that S enters at policy 1 as soon as possible for each $-t \in (-\infty, -t_S)$ since the uncertainty of the median voter location is sufficiently low and S is willing to ensure the payoff from policy 1 given that W does not enter early. In the multi-issue model, we can show a similar result. Specifically, for sufficiently small p_1 and p_2 , there exists t_S^{**} such that S enters at policy $(1, 1)$ as soon as possible

for each $-t \in (-\infty, -t_S^{**})$.

Remark 11 (Incumbent’s partial commitment). In the multi-issue election campaign, it is interesting to see the effect of an incumbent’s commitment as in Remark 6 in Section 2.1. Specifically, suppose that S has already committed to $(1, \bar{X})$. In this modified game, we can show that there is a time cutoff such that before the cutoff, W keeps announcing (\bar{X}, \bar{X}) , while W announces $(1, 1)$ after the cutoff. It might be counterintuitive that W announces a policy that may be later copied by S , while it might make more sense for W to choose, say, $(\bar{X}, 1)$. However, under the resulting policy profile $((1, \bar{X}), (\bar{X}, 1))$, W will lose if the voter is at $(1, 1)$, which is the most likely event. This is the key reason why W chooses $(1, 1)$. The cost of the commitment (defined as the difference in S ’s payoffs) can be shown to be strictly increasing in p_1 and strictly decreasing in p_2 .⁴¹ The first comparative statics make sense because a larger p_1 implies that W ’s advantage of committing to policy 1 in the first issue is reduced in the election campaign game, while the commitment to policy 1 for the first issue is less attractive for S . The second comparative statics follow because W ’s payoff is negatively affected if p_2 increases, but the effect is larger in the commitment case in which W does not time his entry to the first issue before contemplating the entry time to the second issue. This means that W ’s payoff difference decreases in p_2 , so S ’s payoff increases in p_2 . The formal analysis is provided in the Online Appendix. \square

3 A Unified Framework for Dynamic Election Campaign—Policy Announcement Timing Game

This section presents a general model that encompasses the applications presented in Section 2. There are two candidates, A and B . Whenever we say candidates i and j , we assume $i \neq j$. There is a set of policies, X . We define a collection of “policy sets,” denoted $\mathcal{X} \in 2^X \setminus \{\emptyset\}$. In the course of election campaign, each candidate announces an element of \mathcal{X} (the available policy sets).⁴² Given a profile of policy sets $(X_A, X_B) \in \mathcal{X} \times \mathcal{X}$, let $v_i(X_i, X_j)$ be candidate i ’s payoff for each $i = A, B$.

This specification allows for virtually any static model of elections. For example, this allows

⁴¹It is $\frac{e^{-\frac{6p_1+p_2-5}{-6p_1+3p_2+3}}(-8p_1p_2+12(p_1-1)p_1-3(p_2)^2+10p_2+1)}{12p_1-6(p_2+1)} + \frac{p_2-1}{e} + 1 - \left(1 - (1-p_2)e^{-1+\frac{p_1}{1-p_2}}\right)$.

⁴²Here we implicitly assume that the available policy sets are common across the two candidates, but it is merely for simplicity and does not affect the essence of our results.

probabilistic voting in which the winning probability is determined once we fix the policy set profile (X_A, X_B) . The following is how the general setup maps to the applications presented in Section 2.

1. **Valence election campaign:** $v_S(\{0\}, \{0\}) = v_S(\{1\}, \{1\}) = v_S(X, X) = 1$, $v_S(\{0\}, \{1\}) = v_S(\{0\}, X) = v_S(X, \{1\}) = p$, $v_S(\{1\}, \{0\}) = v_S(\{1\}, X) = v_S(X, \{0\}) = 1-p$, and $v_W(X_W, X_S) = 1 - v_S(X_S, X_W)$ for each (X_S, X_W) .
2. **Symmetric office-motivated election campaign:** $v_i(X_i, X_j) = P_i(X_i, X_j)$ for each $X_i, X_j \in \mathcal{X}$ and $i = A, B$.
3. **Partial commitment:** $v_S(X', X') = 1$ for any $X' \in \mathcal{X}$, $v_S((\bar{X}, \bar{X}), (1, 1)) = 1 - (1 - p_1)(1 - p_2)$ as explained in Section 2.3, and so forth. The full payoff matrix is in the Online Appendix.

In our policy announcement timing game, time flows continuously from $-T < 0$ to 0. Time 0 is the fixed election date and the campaign starts at $-T$. During the time interval $-t \in [-T, 0]$, 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. The candidates' payoffs are determined according to the finally announced policy set profiles as of time 0.

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 records all of i 's announcements and their times, all of j 's change of announcements and their times, the current time, and whether there is an opportunity at the current time. Formally, it 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 $-T \leq -t_i^1 < \dots < -t_i^{k_i} < -t$; $X_i^k \in \mathcal{X}$ for all k ; $-T \leq -t_j^1 < \dots < -t_j^{l_j} < -t$; $X_j^l \in \mathcal{X}$ for all l ; and $z_i \in \{yes, no\}$. The interpretation is that $-t_i^k$ is the time at which candidate i 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.⁴³ That is, t_j^l is the l 'th time that

⁴³The 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 14. Such invariance would

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 $\sigma_i : H_i \rightarrow \Delta(\mathcal{X})$, 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 $X_i \in \mathcal{X}$ 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 ruled out some policy platform, then she cannot come back to it afterward. 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} .

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 $\sigma \in \Sigma_A \times \Sigma_B$.⁴⁴ Let $H_j(h_i)$ be the set of candidate j 's feasible histories given h_i , and let $\beta(\cdot|\cdot) : H_i \rightarrow \Delta(H_j)$ 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) $\sigma_i^* \in \arg \max_{\sigma_i \in \Sigma_i} u_i^\beta(\sigma_i, \sigma_j^*|h_i)$ holds for every $h_i \in H_i$ and (ii) β is derived from Bayes rule whenever possible.⁴⁵

Remark 12 (Existence of an equilibrium). Our theorems in the next section prove properties of

not obtain if we introduce incomplete information to the model because the past observation may provide additional information about the opponent's type. This possibility is explored in the working paper version of this paper (Kamada and Sugaya [2019]).

⁴⁴This is well defined because H_i is a countable union of subsets of a finite-dimensional space.

⁴⁵Although each information set at any time after $-T$ has probability zero, one can 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.

PBE but do not prove existence of a PBE.⁴⁶ More specifically, Theorems 1 and 2 provide properties that hold for any PBE, and Theorem 3 assumes the existence of an MPE. Moroni (2019) and Lovo and Tomala (2015) show that a PBE and a MPE exists if X is finite, respectively. This latter result implies that the equilibrium value is unique if X is finite. When X is infinite (for example, X is a policy set $[-1, 1]$ as in the Downsian model), our theorems will be useful for guaranteeing the uniqueness of the value once we solve for a particular MPE by backward induction, as we did in Section 2. □

4 General Predictions

In Section 2, we have seen that the policy announcement timing game can be applied to analyses of various settings. In those applications, 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 2, 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 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 all the applications. 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.

⁴⁶The reason for this choice of exposition is that we did not want to rule out continuous action spaces, which the standard Hotelling model features. Although existence could be proven for each application, generally proving existence when there are infinitely many actions would be difficult.

4.1 The Long Ambiguity Theorem

In this section, we identify conditions under which candidates spend a long time keeping their policies ambiguous. For simplicity, here we assume that (v_i, v_j) is constant-sum. The case with non-constant-sum payoffs is discussed in Appendix F.

Before stating the first-mover disadvantage condition, we introduce a few notations. Given $(\theta_i, \theta_j) \in \mathcal{X}^2$, we define i 's first-mover payoff as the payoff that i receives in a deterministic-move game with the initial policy set profile (θ_i, θ_j) in which i moves first and j moves next, with a constraint that i can only announce a singleton policy set. Formally,

$$v_i^{FM}(\theta_i, \theta_j) = \sup_{x_i \in \theta_i} \left(\inf_{X_j \in \mathcal{X}, X_j \subseteq \theta_j} v_i(x_i, X_j) \right).$$

Similarly, i 's second-mover payoff is the payoff that i receives in a deterministic-move game with the initial policy set profile (θ_i, θ_j) in which j moves first and i moves next:

$$v_i^{SM}(\theta_i, \theta_j) = \inf_{X_j \in \mathcal{X}, X_j \subseteq \theta_j} \left(\sup_{X_i \in \mathcal{X}, X_i \subseteq \theta_i} v_i(X_i, X_j) \right).$$

Note that there is slight asymmetry in the two definitions, namely that i is assumed to announce a singleton policy set in defining the first-mover payoff. This is because the theorem below will provide a condition under which i will not announce a singleton policy set.

Normalize the worst possible payoff of the game at 0. We say that, the **first-mover disadvantage** holds for candidate i given (θ_i, θ_j) if the first-mover payoff is relatively small compared to the second-mover payoff:⁴⁷

$$v_i^{FM}(\theta_i, \theta_j) < \begin{cases} e^{-1} v_i^{SM}(\theta_i, \theta_j) & \text{if } \lambda_j = \lambda_i, \\ \left(\frac{\lambda_j}{\lambda_i} \right)^{\frac{\lambda_j}{\lambda_i - \lambda_j}} v_i^{SM}(\theta_i, \theta_j) & \text{if } \lambda_j \neq \lambda_i. \end{cases} \quad (1)$$

Theorem 1 (Long Ambiguity). *Suppose (v_A, v_B) is constant-sum and is bounded, i.e., there exists $\bar{v} < \infty$ such that $0 \leq v_i(X, X') \leq \bar{v}$ for all $X, X' \in \mathcal{X}$ and $i = A, B$. For each (θ_i, θ_j) , there exists $\bar{T} < \infty$ such that, for any PBE, if the first-mover disadvantage holds for candidate i , then she never*

⁴⁷The coefficient $\left(\frac{\lambda_j}{\lambda_i} \right)^{\frac{\lambda_j}{\lambda_i - \lambda_j}}$ converges to e^{-1} as $\lambda_j \rightarrow \lambda_i$.

enters at any $x_i \in X$ for each $-t \in (-\infty, -\bar{T})$ such that the current policy set profile is (θ_i, θ_j) .

The idea of the proof can be explained as follows. Suppose that at time $-t$, the current policy set profile is (θ_i, θ_j) . On the one hand, the payoff from choosing a singleton policy set approaches the first-mover payoff ($v_i^{FM}(\theta_i, \theta_j)$) as the time $-t$ tends to infinity. On the other hand, the payoff from not doing so can be bounded from below by considering the payoff from the strategy in which i keeps choosing θ_i until some (t -independent) time-cutoff $-a$ and then chooses some policy set. A lower bound of the expected payoff under such a strategy can be computed by assuming that (i) i receives the second-mover payoff when i moves after time $-a$ and then j has no chance to react, while (ii) i receives the worst feasible payoff in all other cases. The first-mover disadvantage condition ensures that the comparison of the two payoffs suggests that i does not choose a singleton policy set when t is sufficiently far away. Since we can further show that the convergence speed of the payoff from choosing a singleton policy set does not depend on the choice of the policy to enter, there is a uniform bound $\bar{T} < \infty$ such that for all $-t < -\bar{T}$, i does not enter at any singleton policy set.

Let us explain how we use this theorem. In the valence election campaign (Section 2.1), given $\theta_S = \{0, 1\}$ and $\theta_W = \{0, 1\}$, the first-mover disadvantage holds for candidate W . To see why, notice that on the one hand, we have $\min_{x_S \in \{0, 1\}} v_W(x_W, x_S) = 0$ for $x_W \in \{0, 1\}$ since S can win for sure by copying W 's policy, and hence $v_W^{FM}(X, X) = 0$. On the other hand,

$$v_W^{SM}(\{0, 1\}, \{0, 1\}) = \min_{X_S \in \{0, 1, \{0, 1\}\}} \max_{X_W \in \{0, 1, \{0, 1\}\}} v_W(X_W, X_S) = p > 0.$$

Therefore, there exists $\bar{T} < \infty$ such that, for any PBE, W does not enter for each $-t \in (-\infty, -\bar{T})$ given $\theta_S = \{0, 1\}$ and $\theta_W = \{0, 1\}$. To obtain Proposition 2 in Section 2.1, we separately verify that if $p > 1/(1 + e)$, then for any PBE, for each $-t \in [-\bar{T}, 0]$, S does not enter unless W entered before. Given these results, for each $-t < -\bar{T}$, if S skips all the opportunities until $-\bar{T}$, the policy announcement profile at $-\bar{T}$ is $(\{0, 1\}, \{0, 1\})$. If S enters at $\{1\}$ (or $\{0\}$), her payoff is $1 - p$ (or p). Hence, if S prefers not entering at $-\bar{T}$ in any PBE, then she does not enter until $-\bar{T}$ in any PBE. Therefore, we conclude that $\theta_S = \{0, 1\}$ and $\theta_W = \{0, 1\}$ are realized on the equilibrium path for each $-t \in (-\infty, -\bar{T})$ in any PBE.

In the symmetric office-motivated election campaign (Section 2.2), if the Condorcet winner does

not exist, then for each candidate i , we have (i) $\min_{x_j \in X} v_i(x_i, x_j) = 0$ for each $x_i \in X$ and hence $v_i^{FM}(X, X) = 0$, and (ii) $v_i^{SM}(X, X) = \min_{X_j \in \mathcal{X}} \max_{X_i \in \mathcal{X}} v_j(X_j, X_i) = 1$. Hence, again, there exists $\bar{T} < \infty$ such that for any PBE, each candidate i does not enter for each $-t \in (-\infty, -\bar{T}]$.

In the model of Section 2.3, given (θ_S, θ_W) satisfying $\theta_W \subseteq \theta_S$, the first-mover disadvantage holds for W (the proof is the same as in the valence election campaign). Separately, we verify that if $p_1, p_2 > 1/(1+e)$, then the strong candidate always copies W 's policy on the equilibrium path in any PBE (that is, denoting the current policy set profile by (θ_S, θ_W) , we have $\theta_W \subseteq \theta_S$ all the time on the equilibrium path). Hence, candidate W does not enter for a long time on the equilibrium path. Since $\theta_W \subseteq \theta_S$ on the equilibrium path, this long ambiguity of W implies that S employs ambiguous policies for a long time.

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 models in Section 2. A policy set $X^* \in \mathcal{X}$ is a **Condorcet winner** if, for each i , X^* is a unique best response to X^* and $v_i(X^*, X) \geq v_i(X_i, X)$ for each $X_i \in \mathcal{X}$. Notice that, in a symmetric election ($v_i(\bar{X}, \bar{X}') = v_j(\bar{X}, \bar{X}')$ for any $\bar{X}, \bar{X}' \in \mathcal{X}$) that is constant-sum such that $v_i(\bar{X}, \bar{X}') + v_j(\bar{X}, \bar{X}') = 1$ for any $\bar{X}, \bar{X}' \in \mathcal{X}$, “ X^* is a unique best response to X^* ” is equivalent to “ $v_i(X^*, \bar{X}') > \frac{1}{2}$ for any $\bar{X}, \bar{X}' \in \mathcal{X}$,” hence the name “Condorcet winner.”⁴⁸ Note, however, that a Condorcet winner need not be a weakly dominant strategy under our definition, so our theorem will apply to a wider class of elections.⁴⁹

For example, in a uni-dimensional Downsian model 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, consider a model in which the voters are risk averse and think that there is uncertainty about what policy a candidate announcing X would implement), there is a unique Condorcet winner $\{x^*\}$ if the policy corresponding to the median voter is x^* . In addition, the policy set $\{1\}$ with $\delta = 0$ in

⁴⁸To see that the equivalence holds, consider a symmetric constant-sum election. Symmetry and constant-sum imply $v_i(X^*, X^*) = \frac{1}{2}$. Take an arbitrary $\bar{X}' \neq X^*$. Note that symmetry implies that $v_i(X^*, \bar{X}') > \frac{1}{2}$ is equivalent to $v_i(\bar{X}', X^*) < \frac{1}{2}$. Hence, $v_i(X^*, \bar{X}') > \frac{1}{2}$ is equivalent to $v_i(X^*, X^*) > v_i(\bar{X}', X^*)$. Since this holds for any choice of $\bar{X}' \neq X^*$, we have established the desired equivalence.

⁴⁹Thus, our theorem that candidates enter at the Condorcet winner as soon as possible may not be as straightforward as one might think. Moreover, even if there is a dominant strategy, the revision-games literature shows that it is not necessarily the case that it must be played right away (see, e.g., Kamada and Kandori [2019]).

Proposition 1 and the policy set $\{x^*(X, \mu)\}$ in the symmetric office-motivated election campaign with $(X, \mu) \in \mathcal{M}$ (part 1 of Proposition 3) are Condorcet winners. Note that the definition implies that there is at most one Condorcet winner.⁵⁰

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

Theorem 2 (Dynamic Median-Voter). *Suppose that (v_A, v_B) is constant-sum and there exists a Condorcet winner X^* . Then, in any PBE, at any time $-t \in (-\infty, 0)$, conditional on any history of candidate i such that the current policy set profile is (X', X'') with $X^* \subseteq X'$ and $X'' \in \{X, X^*\}$, i announces X^* .*

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.

The heart of the proof shows that, if candidate i chooses some $X' \neq X^*$ and the opponent j has a chance to make an announcement afterward, j can secure the payoff of $v_j(X^*, X')$ in the event i receives no further move after such j 's move. By the constant-sum assumption, this implies that $v_i(X', X^*)$ is an upper bound of i 's payoff in this case, which is strictly lower than $v_i(X^*, X^*)$ by the assumption that X^* is a unique best response to X^* . This enables us to bound i 's payoff from choosing $X' \neq X^*$ from above. With a similar argument, we can bound the payoff from choosing X^* from below, and show that such a lower bound is larger than the upper bound of the payoff from choosing $X' \neq X^*$.

In a working paper version of this paper (Kamada and Sugaya [2019]), we generalize the theorem to cover the case with non-constant-sum games. In particular, we show the uniqueness of a PBE when we further require that there is a policy that is strictly dominant for each i .

4.3 The Constant-Sum Markov Theorem

In all of the applications we consider in Section 2, 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 are

⁵⁰To prove this, suppose that X^* and X^{**} are both Condorcet winners. Then, $v_i(X^*, X^*) < v_i(X^*, X^{**}) < v_i(X^{**}, X^{**})$, where the first inequality follows because X^* is j 's unique best response against i 's X^* and (v_i, v_j) is constant-sum, and the second inequality follows because X^{**} is i 's unique best response against j 's X^{**} . A symmetric argument shows that $v_i(X^{**}, X^{**}) < v_i(X^*, X^*)$. This is a contradiction.

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 3, 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 $h^t = \left((t_A^k, X_A^k)_{k=1}^{k_A}, (t_B^k, X_B^k)_{k=1}^{k_B}, t \right)$ 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 $\theta(h^t) = (X_A^{k_A}, X_B^{k_B})$ be the most recent policy profile at time $-t$; and $\theta_i(h^t) = X_i^{k_i}$ be candidate i 's most recent policy at $-t$. Note that $\theta(h^t) = \theta(h_i^t)$ for each i and t . Candidate i 's strategy is a map $\sigma_i : H \rightarrow \Delta(\mathcal{X})$, with a restriction that $\theta(h^t) = (X_A, X_B)$ implies that $\sigma_i(h_t) \subseteq X_i$. Let $\check{\Sigma}_i$ be the space of i 's strategies, and $\check{\Sigma} = \check{\Sigma}_A \times \check{\Sigma}_B$. With σ_i , candidate i takes $\sigma_i(h^t)$ if she has an opportunity at time $-t$ and takes $\theta_i(h^t)$ otherwise. A subgame-perfect equilibrium (SPE) can be defined in the standard manner. We call this setup "public monitoring."

In the "private monitoring" setup, 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 $\sigma \in \Sigma$ in private monitoring, and let $w_t^i(\sigma, h_i^t, X_i)$ be candidate i 's continuation payoff of taking $X_i \in \mathcal{X}$ when her private history is h_i^t and she receives an opportunity at $-t$. Similarly, let $\hat{w}_t^i(\sigma, h_i^t, X_j)$ be candidate i 's continuation payoff when her private history is h_i^t and candidate j receives an opportunity and takes $X_j \in \mathcal{X}$; and let $w_t^i(\sigma, h_i^t, no)$ 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 $\bar{\sigma} \in \check{\Sigma}$ in public monitoring, and let $W_t^i(\bar{\sigma}, h^t, X_i)$ be candidate i 's continuation payoff of taking policy $X_i \in \mathcal{X}$ when the public history is h^t and she receives an

opportunity at $-t$. Similarly, let $\hat{W}_t^i(\bar{\sigma}, h^t, X_j)$ be candidate i 's continuation payoff when the public history is h^t and candidate j receives an opportunity and takes $X_j \in \mathcal{X}$; and let $W_t^i(\bar{\sigma}, h^t, no)$ be candidate i 's continuation payoff when the public history is h^t and no candidate receives an opportunity at $-t$.

Third, candidate i 's Markov strategy is a map that is constant with respect to the part of the histories other than the state, where the state is defined as the profile of the current remaining time t and the policy set profile at $-t$. Markov-perfect equilibrium (MPE) is a subset of the set of SPE where candidates take Markov strategies.

We can show that, if an MPE exists in public monitoring, then its value function corresponds to a value function of any SPE under public monitoring and that of any PBE under private monitoring, where the value $v_{i,t}(X_i, X_j)$ represents all three of the following: (a) i 's continuation payoff of i 's taking X_i when only i receives an opportunity at $-t$ in which the current policy set of j is X_j , (b) i 's continuation payoff of j 's taking X_j when only j receives an opportunity at $-t$ in which the current policy set of i is X_i , and (b) i 's continuation payoff when no candidate receives an opportunity at $-t$ in which the current policy set profile is (X_i, X_j) (recall that $\theta_j(h^t) = \theta_j(h_i^t) = \theta_j(h_j^t)$).

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

$$w_t^i(\sigma, h_i^t, X_i) = W_t^i(\bar{\sigma}, h^t, X_i) = v_{i,t}(X_i, \theta_j(h^t)); \quad (2)$$

$$\hat{w}_t^i(\sigma, h_i^t, X_j) = \hat{W}_t^i(\bar{\sigma}, h^t, X_j) = v_{i,t}(\theta_i(h^t), X_j) \quad (3)$$

and

$$\begin{aligned} w_t^i(\sigma, h_i^t, \theta_i(h_i^t)) &= \hat{w}_t^i(\sigma, h_i^t, \theta_j(h_i^t)) = w_t^i(\sigma, h_i^t, no) \\ &= W_t^i(\bar{\sigma}, h^t, \theta_i(h^t)) = \hat{W}_t^i(\bar{\sigma}, h^t, \theta_j(h^t)) = W_t^i(\bar{\sigma}, h^t, no) = v_{i,t}(\theta(h^t)). \end{aligned} \quad (4)$$

First, we explain the result under public monitoring. Take an MPE $\bar{\sigma} \in \check{\Sigma}$ under public monitoring, and let $v_{i,t}(X_i, X_j)$ be the equilibrium value of this MPE when the state is (t, X_i, X_j) . Suppose there exists another SPE $\hat{\sigma} \in \check{\Sigma}$ which achieves the value $\hat{v}_{i,t}(X_i, X_j) \neq v_{i,t}(X_i, X_j)$ for

some t , X_i , and X_j . Without loss, assume $\hat{v}_{i,t}(X_i, X_j) > v_{i,t}(X_i, X_j)$.⁵¹ Since the game is constant-sum, this inequality implies that if candidate i deviates from $\bar{\sigma}_i$ to $\hat{\sigma}_i$ for the time interval $(-t, 0]$ when $\theta_j(h_i^t) = X_i$ and $\theta_j(h^t) = X_j$, then no matter what candidate j will do in the continuation play, candidate i 's payoff improves upon $v_{i,t}(X_i, X_j)$. In particular, her payoff strictly increases when j follows $\bar{\sigma}_j$. Therefore, if $\bar{\sigma}$ is MPE, such $\hat{\sigma}$ cannot exist. This establishes the results for public monitoring.

Second, for private monitoring, suppose there exists a PBE $\tilde{\sigma} \in \Sigma$ which achieves $\tilde{v}_{i,t}(X_i, X_j) \neq v_{i,t}(X_i, X_j)$ for some t , X_i , and X_j . Again, without loss, assume $\tilde{v}_{i,t}(X_i, X_j) > v_{i,t}(X_i, X_j)$. By the same argument as above, if candidate i takes $\tilde{\sigma}_i$, then no matter what candidate j will do under private monitoring, candidate i 's payoff improves upon $v_{i,t}(X_i, X_j)$. In particular, candidate i obtains the payoff no less than $\tilde{v}_{i,t}(X_i, X_j)$ if candidate j takes $\bar{\sigma}_j$ under private monitoring. Here, we use the fact that the space for Markov strategies under public monitoring can be seen as the same as the space for Markov strategies under private monitoring: Since Markov strategies depend only on the current policy sets and the current time, candidate j “can take” $\bar{\sigma}_j$ under private monitoring as well.

However, whether monitoring is private or public, candidate i 's strategy that is a best response against $\bar{\sigma}_j$ after all histories contains a Markov strategy. Intuitively, if the opponent uses a strategy that only depends on t , $\theta_i(h^t)$, and $\theta_j(h^t)$, then it is sufficient for candidate i to take a strategy which depends only on those variables. Together with this result, the conclusion of the last paragraph implies that there exists a Markov strategy of candidate i that achieves the value no less than $\tilde{v}_{i,t}(X_i, X_j) > v_{i,t}(X_i, X_j)$ against $\bar{\sigma}_j$. This contradicts the fact that $v_{i,t}(X_i, X_j)$ is the value of MPE $\bar{\sigma}$. Therefore, under private monitoring, the equilibrium value must be equal to $v_{i,t}(X_i, X_j)$.⁵²

As we mentioned, all of the applications in Section 2, candidates are purely office-motivated, so the payoffs are constant-sum. Thus, Theorem 3 implies that the continuation payoffs are the same between the public monitoring and private monitoring models. This implies that the outcome characterized under private monitoring is outcome-equivalent to the one under public monitoring. Moreover, the equivalence between public and private monitoring is used to prove Theorem 2.

⁵¹This is without loss because, since the game is constant-sum, given $\hat{v}_{i,t}(X_i, X_j) < v_{i,t}(X_i, X_j)$, we have $\hat{v}_{j,t}(X_j, X_i) > v_{j,t}(X_j, X_i)$. The rest of the explanation goes through with i and j reversed.

⁵²Although 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.

5 Discussions

We provide two more examples of settings in which our policy announcement timing game is applicable. The full analysis of those can be found in the Appendix. A working paper version of this paper (Kamada and Sugaya [2019]) contains more topics, such as campaign spending, incomplete information, and synchronous announcement opportunities. The point of this section is to illustrate the wide applicability of the model: There are many ways in which one can use our framework to analyze topics that are relevant for election campaigns.

Multi-dimensional policy space with policy motivation: The analysis on the multi-dimensional policy space in Section 2.2 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, we introduce policy motivation to the model with a multi-dimensional policy space in Appendix G. 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. 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 candidate 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.”⁵³ 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.”

Different reasons for ambiguity: In our model, candidates use ambiguous language due to the

⁵³The facts and the quote appear in Strauss (2016) and Gurciullo and Debenedetti (2016).

first-mover disadvantage. As we discussed in the Introduction, the literature has pointed out other reasons for candidates to be ambiguous. One of those is that candidates have limited resources for making precise policy announcements and voters have limited capacity to understand those announcements (Page [1976, 1978], Polborn and Yi [2006], Egorov [2015] and Dragu and Fan [2016]). In Appendix H, we extend the model of Dragu and Fan (2016), in which candidates spend their funds to various policy issues to make them more prominent, to our policy announcement timing game and demonstrate that our dynamic extension brings in new insights.⁵⁴ For example, Dragu and Fan (2016) show that there is no policy issue for which two candidates spend money on. This prediction is robust to our setting in the sense that there is no history of play at time $-t$ at which both candidates would spend money at time $-t$, but the prediction about the total spending on the election day would change: the identity of the potential candidate who would spend money may change over time, and it is possible that one candidate spends early on in the campaign while the other candidate spends later.

6 Conclusion

We have introduced the first model of dynamic campaigns into the literature on elections, which we call the “policy announcement timing game.” In the model, candidates cannot always announce their policies but stochastically obtain opportunities to do so. We applied the model to various settings, 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 the candidates may or may not have motivations to defer a clear announcement of policies, depending on the opponent’s latest announcement and the time left until the election. The insights from the analysis of our three applications are generalized in the Long Ambiguity Theorem, the Dynamic Median Voter Theorem, and the Constant-Sum Markov Theorem. Our model by no means captures all relevant incentives in election campaigns,

⁵⁴We chose Dragu and Fan (2016) for this analysis because their game has perfect information. Other papers cited here involve incomplete information, so in our dynamic extension we would have to deal with evolving beliefs. Such an analysis is possible (indeed, the working paper version of this paper (Kamada and Sugaya [2019]) considers a model involving incomplete information), but it is outside the scope of this paper. Note that the dynamic extension of the model of Dragu and Fan (2016) can be regarded as a special case of the policy announcement timing game by interpreting the policy set to be the available cumulative spending. This is because the cumulative spending can never decrease over time.

nor do we claim that it captures the most relevant ones. Which types of incentives are most powerful depends on the nature of the given campaign, and hence in future work, it would be important to incorporate other realistic features into the model. Our work shows that the incentives are complicated even in our rather simple setup. It raises a wide range of new questions as follows.

First, 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.

Second, 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.

Third, there may be a positive effect of getting a candidate’s name salient in the early stage of a campaign. For example, the “June Puzzle” 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.⁵⁵ 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 our model to analyze these hypotheses.

Fourth, we have considered two-candidate elections, but it would be interesting to consider more than two candidates.⁵⁶ In such an environment, there is no pure-strategy equilibrium in a static election game, while we can hope for the existence of a unique pure-strategy PBE outcome 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, 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

⁵⁵Gerber et al. (2011) empirically show that the effect of campaign spending declines over time (earlier spending has a weaker effect). We thank Avidit Acharya for sharing the story of June Puzzle, who attributes the story to Seth Hill, Brett Gordon, and Michael Peress.

⁵⁶We thank Alessandro Lizzeri for pointing this out.

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.⁵⁷ The uniqueness of the equilibrium that we obtain in our analysis would be useful in empirically testing the theoretical implications of the model.

References

- [1] **Alesina, Alberto, and Alex Cukierman (1990)**, “The Politics of Ambiguity,” *Quarterly Journal of Economics*, 4, 829–50.
- [2] ——— **and Richard Holden (2008)**, “Ambiguity and Extremism in Elections,” unpublished.
- [3] **Ambrus, Attila, James Burnes, and Yuhta Ishii (2014)**, “Gradual bidding in eBay-like auctions,” unpublished.
- [4] ——— **and Shih-En Lu (2015)**, “A Continuous-Time Model of Multilateral Bargaining,” *AJ: Microeconomics*, 7, 208–249.
- [5] **Ansolabehere, Stephen, and James M. Snyder Jr. (2000)**, “Valence Politics and Equilibrium in Spatial Election Models,” *Public Choice*, 103, 327–336.
- [6] **Aragonès, Enriqueta, and Zvika Neeman (2000)**, “Strategic Ambiguity in Electoral Competition,” *Journal of Theoretical Politics*, 12, 183–204.
- [7] ——— **and Thomas R. Palfrey (2002)**, “Mixed Equilibrium in a Downsian Model with a Favored Candidate,” *Journal of Economic Theory*, 103, 131–161.
- [8] ——— **and Andrew Postlewaite (2002)**, “Ambiguity in Election Games,” *Review of Economic Design*, 7, 233–255.
- [9] **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.
- [10] ——— **(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.
- [11] **Austen-Smith, David (1987)**, “Interest Groups, Campaign Contributions, and Probabilistic Voting,” *Public Choice*, 54, 123–139.
- [12] **Avoyan, Ala and Joan Ramos (2018)**, “A Road to Efficiency through Communication and Commitment,” unpublished.

⁵⁷As mentioned in Remark 2, this pattern is roughly consistent with the empirical finding in Campbell (1983).

- [13] **BBC News (2013)**, “David Cameron promises in/out referendum on EU,” <https://www.bbc.com/news/uk-politics-21148282>, accessed 01/31/2019.
- [14] **Bailey, Michael A. (2002)**, “Money and Representation: An Exploration in Multiple Dimensions with Informative Campaigns,” unpublished.
- [15] **Banda, Kevin K. (2013)**, “The Dynamics of Campaign Issue Agendas,” *State Politics & Policy Quarterly*, 13, 446–470.
- [16] ——— (2015), “Competition and the Dynamics of Issue Convergence,” *American Politics Research*, 43, 821–845.
- [17] **Banks, Jeffrey S. (1990)**, “A Model of Electoral Competition with Incomplete Information,” *Journal of Economic Theory*, 50, 309–325.
- [18] **Baron, David P. (1994)**, “Electoral Competition with Informed and Uninformed Voters,” *American Political Science Review*, 88, 33–47.
- [19] **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.
- [20] **Bernhardt, Daniel, and Daniel Ingberman (1985)**, “Candidate Reputations and the Incumbency Effect,” *Journal of Public Economics*, 27, 47–67.
- [21] **Cadigan, Will and Ryan Struyk (2018)**, “A lot of Democratic senators want to be President. History shows it will be hard.” in *CNN Politics* <https://www.cnn.com/2018/04/26/politics/senator-president-history-analysis/index.html>, accessed 10/08/2019.
- [22] **Calcagno, Riccardo, Yuichiro Kamada, Stefano Lovo, and Takuo Sugaya (2014)**, “Asynchronicity and Coordination in Common and Opposing Interest Games,” *Theoretical Economics*, 9, 409–434.
- [23] **Callander, Steven, and Simon Wilkie (2007)**, “Lies, Damned Lies, and Political Campaigns,” *Games and Economic Behavior*, 60, 262–286.
- [24] ——— and **Catherine H. Wilson (2008)**, “Context-Dependent Voting and Political Ambiguity,” *Journal of Public Economics*, 92, 565–581.
- [25] **Calvo, Guillermo A. (1983)**, “Staggered Prices in a Utility-Maximizing Framework,” *Journal of Monetary Economics*, 12, 383–398.
- [26] **Campbell, James E. (1983)**, “Ambiguity in the Issue Positions of Presidential Candidates: A Causal Analysis,” *American Journal of Political Science*, 27, 284–293.
- [27] **Carter, Jennifer and John W. Patty (2015)**, “Valence and Campaigns.” *American Journal of Political Science* 59, 825–840.
- [28] **Caruana, Guillermo and Liran Einav (2008)**, “A Theory of Endogenous Commitment,” *Review of Economic Studies*, 75, 99–116.

- [29] **CNN Political Ticker (2007)**, “Obama takes heat for link too McClurkin,” <https://web.archive.org/web/20071130064038/http://politicalticker.blogs.cnn.com/2007/10/23/obama-takes-heat-for-link-to-mcclurkin/>, accessed 10/03/2019.
- [30] **Coate, Stephen (2004)**, “Political Competition with Campaign Contributions and Informative Advertising,” *Journal of the European Economic Association*, 2, 772–804.
- [31] **Coughlin, Peter J. (1992)**, *Probabilistic Voting Theory*, New York: Cambridge University Press.
- [32] **van Damme, Eric, and Sjaak Hurkens (1996)**, “Commitment Robust Equilibria and Endogenous Timing,” *Games and Economic Behavior*, 15, 290–311.
- [33] **Dragu, Tiberiu and Xiaochen Fan (2016)**, “An Agenda-Setting Theory of Electoral Competition,” *Journal of Politics*, 78, 1170–1183.
- [34] **Egorov, Georgy (2015)**, “Single-Issue Campaigns and Multidimensional Politics” unpublished.
- [35] **Fudenberg, Drew and Jean Tirole (1991)**, “Perfect Bayesian Equilibrium and Sequential Equilibrium,” *Journal of Economic Theory*, 53: 236–260.
- [36] **Gale, Douglas (1995)**, “Dynamic Coordination Games,” *Economic Theory*, 5, 1–18.
- [37] ——— (2001), “Monotone Games with Positive Spillovers,” *Games and Economic Behaviors*, 37, 295–320.
- [38] **Gensbittel, Fabien, Stefano Lovo, Jérôme Renault, and Tristan Tomala (2017)**, “Zero-Sum Revision Games, *Games and Economic Behavior*, Forthcoming.
- [39] **Glazer, Amihai (1990)**, “The Strategy of Candidate Ambiguity,” *American Political Science Review*, 84, 237–41.
- [40] **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.
- [41] **Groseclose, Tim (2001)** “A Model of Candidate Location When One Candidate Has a Valence Advantage,” *American Journal of Political Science*, 45, 862–886.
- [42] **Grossman Gene M., and Elhanan Helpman (1996)**, “Electoral Competition and Special Interest Politics,” *Review of Economic Studies*, 63, 256–286.
- [43] **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.
- [44] **Hamilton, Jonathan H., and Steven M. Slutsky (1990)**, “Endogenous Timing in Duopoly Games: Stackelberg or Cournot Equilibria,” *Games and Economic Behavior*, 2, 29–46.

- [45] ——— and ——— (1993), “Endogenizing the Order of Moves in Matrix Games,” *Theory and Decision*, 34, 47–62.
- [46] **Harrington, Jr. , Joseph E. (1992)**, “The Revelation of Information Through the Electoral Process: An Exploratory Analysis,” *Economics and Politics*, 4, 255–275.
- [47] **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.
- [48] **Hopenhyan, Hugo and Maryam Saoedi (2016)**, “Bidding Dynamics in Auctions,” unpublished.
- [49] **Hummel, Patrick (2010)**, “On the Nature of Equilibria in a Downsian Model with Candidate Valence,” *Games and Economic Behavior*, 70, 425–445.
- [50] **Kapor, Adam and Sofia Moroni (2016)**, “Sniping in Proxy Auctions with Deadlines,” unpublished.
- [51] **Ishii, Yuhta, and Yuichiro Kamada (2011)**, “The Effect of Correlated Inertia on Coordination,” unpublished.
- [52] **Kamada, Yuichiro, and Michihiro Kandori (2017a)**, “Revision Games, Part I: Theory,” unpublished.
- [53] ——— and **Michihiro Kandori (2017b)**, “Revision Games, Part II: Applications and Robustness” unpublished.
- [54] ——— and **Nozomu Muto (2015)**, “Multi-Agent Search with Deadline” unpublished.
- [55] ——— and **Takuo Sugaya (2019)**, “Optimal Timing of Policy Announcements in Dynamic Election Campaigns” unpublished.
- [56] **Karzas, Ionnis, and Steven E. Shreve (1988)**, *Brownian Motion and Stochastic Calculus*, New York: Springer-Verlag.
- [57] **Kartik, Navin, and R. Preston McAfee (2007)**, “Signaling Character in Electoral Competition,” *American Economic Review*, 97, 852–870.
- [58] **Klenow, Peter J. and Benjamin A. Malin (2010)**, “Microeconomic Evidence on Price-Setting,” *Handbook of Monetary Economics*, 3, 231–284.
- [59] **Krasa, Stefan and Mattias Polborn (2010)**, “The Binary Policy Model,” *Journal of Economic Theory*, 145, 661–688.
- [60] **Kreps, David M. and Robert Wilson (1982)**, “Sequential Equilibria,” *Econometrica*, 50, 863–894.
- [61] **Lindbeck, Assar, and Jorgen Weibull (1987)**, “Balanced Budget Redistribution as the Outcome of Political Competition,” *Political Choice*, 52, 273–297.

- [62] **Lipman, Barton L. and Ruqu Wang (2000)**, “Switching Costs in Frequently Repeated Games,” *Journal of Economic Theory*, 93, 149–190.
- [63] **Lovo, Stefano, and Tristan Tomala (2015)**, “Markov Perfect Equilibria in Stochastic Revision Games,” unpublished.
- [64] **Mance, Henry and Jim Pickard (2016)**, “How David Cameron could have avoided an EU Leave vote,” *Financial Times*, <https://www.ft.com/content/4f8634c6-3873-11e6-9a05-82a9b15a8ee7>, accessed 01/31/2019.
- [65] **McGrane, Reginald C. (ed.) (1919)**, *The Correspondence of Nicholas Biddle*. Boston, New York: The Houghton Mifflin Company.
- [66] **Meirowitz, Adam (2005)**, “Informational Party Primaries and Strategic Ambiguity,” *Journal of Theoretical Politics*, 17, 107–136.
- [67] **Moroni, Sofia (2019)**, “Existence of Trembling Hand Perfect and Sequential Equilibrium in Stochastic Games,” unpublished.
- [68] **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.
- [69] **Page, Benjamin I. (1976)**, “The Theory of Political Ambiguity,” *American Political Science Review*, 70, 742–752.
- [70] ——— (1978), *Choices and Echoes in Presidential Elections*, Chicago: University of Chicago Press.
- [71] **Palfrey, Thomas R. (1984)**, “Spatial Equilibrium with Entry,” *Review of Economic Studies*, 51, 139–156.
- [72] **Persson, Torsten, and Guido Tabellini (2000)**, *Political Economics: Explaining Economic Policy*, Cambridge: MIT Press.
- [73] **Polborn, Mattias K. and David T. Yi (2006)**, “Informative Positive and Negative Campaigning,” *Quarterly Journal of Political Science*, 103, 351–371.
- [74] **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.
- [75] **Prat, Andrea (2002a)**, “Campaign Spending with Office-Seeking Politicians, Rational Voters, and Multiple Lobbies,” *Journal of Economic Theory*, 103, 162–189.
- [76] ——— (2002b), “Campaign Advertising and Voter Welfare,” *Review of Economic Studies*, 69, 999–1018.
- [77] **Reeve, Elspeth (2011)** “Who Coined ‘Obamacare?’” in *The Atlantic*, <https://www.theatlantic.com/politics/archive/2011/10/who-coined-obamacare/335745/>, accessed on 10/03/2019.

- [78] **Renou, Ludovic (2009)**, “Commitment games,” *Games and Economic Behavior*, 66, 488–505.
- [79] **Roemer, John E. (2001)**, *Political Competition: Theory and Applications*, Cambridge: Harvard University Press.
- [80] **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.
- [81] **Romm, Assaf (2014)**, “Building Reputation at the Edge of the Cliff,” unpublished.
- [82] **Roy, Nilanjan (2014)**, “Revision and Cooperation: Evidence from Cournot Duopoly Experiments,” unpublished.
- [83] **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.
- [84] **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.
- [85] **Shepsle, Kenneth (1972)**, “The Strategy of Ambiguity: Uncertainty and Electoral Competition,” *American Political Science Review*, 66, 555–568.
- [86] **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.
- [87] **Watson, Joel (2017)**, “A General, Practicable Definition of Perfect Bayesian Equilibrium,” working paper.
- [88] **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 of policy announcement timing game, 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 the reader's convenience.

Lemma 2. *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. \square

C Proof of Theorem 1

Without loss, we assume that $v_i(X_i, X_j) \in [0, 1]$ for each i and $X_i, X_j \in \mathcal{X}$. Fix a PBE σ and (θ_i, θ_j) , and take \bar{T} such that

$$e^{-\lambda_j \bar{T}} < \begin{cases} e^{-1} v_i^{SM}(\theta_i, \theta_j) - v_i^{FM}(\theta_i, \theta_j) & \text{if } \lambda_j = \lambda_i \\ \left(\frac{\lambda_j}{\lambda_i}\right)^{\frac{\lambda_j}{\lambda_i - \lambda_j}} v_i^{SM}(\theta_i, \theta_j) - v_i^{FM}(\theta_i, \theta_j) & \text{if } \lambda_j \neq \lambda_i \end{cases} \quad (5)$$

and

$$\bar{T} \geq \begin{cases} \frac{1}{\lambda} & \text{if } \lambda_j = \lambda_i = \lambda \\ \frac{\log \lambda_i - \log \lambda_j}{\lambda_i - \lambda_j} & \text{if } \lambda_j \neq \lambda_i \end{cases}.$$

Suppose that candidate i receives an opportunity at time $-t$ when the current policy set profile is (θ_i, θ_j) , and let her history be h_i .

First, suppose for contradiction that there is $x_i \in \theta_i$ such that $\sigma_i(h_i)(\{x_i\}) > 0$. Since σ_j must specify that a static best response to x_i is taken after i 's entry to x_i , an upper bound of candidate

i 's payoff of entering at x_i is given by

$$\left(1 - e^{-\lambda_j t}\right) \inf_{X_j \in \mathcal{X}, X_j \subseteq \theta_j} v_i(x_i, X_j) + e^{-\lambda_j t}, \quad (6)$$

which is no greater than $v_i^{FM}(\theta_i, \theta_j) + e^{-\lambda_j t}$ by the definition of $v_i^{FM}(\theta_i, \theta_j)$.

Second, suppose candidate i deviates to the following strategy denoted by $\sigma_i^{(a)}$ with $a \in [0, t]$ being a parameter: Under $\sigma_i^{(a)}$, i takes $X_i = \theta_i$ for each $-t' \in (-t, -a)$ and takes $\max_{X_i \in \mathcal{X}, X_i \subseteq \theta_i} v_i(X_i, \theta_{j,t'})$ for each $-t' \in (-a, 0)$, where $\theta_{j,t'}$ represents the policy set of candidate j at time $-t'$. Candidate i 's payoff of taking $\sigma_i^{(a)}$ is bounded from below by

$$\begin{aligned} & \int_{\tau=0}^a \lambda_i e^{-\lambda_i \tau} \times e^{-\lambda_j(a-\tau)} \times \inf_{X_j \in \mathcal{X}, X_j \subseteq \theta_j} \sup_{X_i \in \mathcal{X}, X_i \subseteq \theta_i} v_i(X_i, X_j) d\tau \\ = & \begin{cases} \lambda a e^{-\lambda a} v_i^{SM}(\theta_i, \theta_j) & \text{if } \lambda_j = \lambda_i = \lambda, \\ \frac{\lambda_i (\lambda_i e^{-\lambda_i a} - \lambda_j e^{-\lambda_j a})}{\lambda_i - \lambda_j} v_i^{SM}(\theta_i, \theta_j) & \text{if } \lambda_j \neq \lambda_i. \end{cases} \end{aligned}$$

Here, $\lambda_i e^{-\lambda_i \tau}$ represents the probability density that i has the first opportunity after $-a$ at time $-(a - \tau)$, and $e^{-\lambda_j(a-\tau)}$ is the probability that j cannot move after i moves at time $-(a - \tau)$.

Note that (i) $\frac{\lambda_i (\lambda_i e^{-\lambda_i a} - \lambda_j e^{-\lambda_j a})}{\lambda_i - \lambda_j}$ (or $\lambda a e^{-\lambda a}$) is maximized at $a = a^* := \frac{\log \lambda_i - \log \lambda_j}{\lambda_i - \lambda_j}$ (or $a = a^* := \frac{1}{\lambda}$) and (ii) the maximized value is $\left(\frac{\lambda_j}{\lambda_i}\right)^{\frac{\lambda_j}{\lambda_i - \lambda_j}}$ (or e^{-1}). Candidate i 's payoff from taking $\sigma_i^{(a^*)}$ is:

$$\begin{cases} e^{-1} v_i^{SM}(\theta_i, \theta_j) & \text{if } \lambda_j = \lambda_i, \\ \left(\frac{\lambda_j}{\lambda_i}\right)^{\frac{\lambda_j}{\lambda_i - \lambda_j}} v_i^{SM}(\theta_i, \theta_j) & \text{if } \lambda_j \neq \lambda_i. \end{cases} \quad (7)$$

Given the definition of \bar{T} in (5), this implies that, when j uses σ_j , i 's payoff from taking $\sigma_i^{(a)}$ is strictly greater than her payoff from taking σ_i . This is a contradiction to the assumption that σ is a PBE. Hence, for any PBE, for all time $-t < -\bar{T}$, it is optimal not to enter at x_i .

D Proof of Theorem 2

Fix a PBE σ . Take any time $-t \in [-T, 0)$ and a history of candidate i such that j 's current policy set is X or X^* . Suppose that under σ , strictly after time $-t$, candidates announce X^* if the opponent's current policy set is either X or X^* .

First, suppose that j 's current policy set of the opponent is X and compare the payoff of announcing X^* and the payoff of announcing $X_i \neq X^*$. The former payoff is at least

$$e^{-\lambda_j t} v_i(X^*, X) + (1 - e^{-\lambda_j t}) v_i(X^*, X^*) \quad (8)$$

and the latter payoff is at most

$$e^{-\lambda_j t} v_i(X_i, X) + (1 - e^{-\lambda_j t}) \bar{v}, \quad (9)$$

where $\bar{v} < v_i(X^*, X^*)$. The second term of (8) and the second term of (9) need some explanations. Notice that Theorem 3 implies that the continuation payoff can be computed by assuming the model of public monitoring in which candidates observe all the Poisson arrivals in the past, including those of the opponent. Thus we consider the model of public monitoring to compute payoffs.

- For the second term of (8), since σ is a PBE, it must be a Nash equilibrium in the subgame starting from time $-t$. This implies that i 's payoff under σ in that subgame is at least as good as her payoff from always announcing X^* against σ_j . This payoff is minimized when, if j receives an opportunity after time $-t$, he plays X^* because j 's best response against i 's X^* is to announce X^* . This is why the expression in (8) provides a lower bound of the payoff from i 's announcing X^* at time $-t$.
- For the second term of (9), suppose that, after i 's announcement of X_i at time $-t$, j receives an opportunity. Since σ is a PBE, it must be a Nash equilibrium in the subgame starting from such j 's opportunity. This implies that j 's payoff under σ in that subgame is at least as good as his payoff from always announcing X^* against σ_i . Since (v_i, v_j) is constant-sum, we then have that i 's payoff under σ in the game is at most her payoff from j 's always announcing X^* against σ_i .

This upper-bound payoff, which we denote by \bar{v} , is a convex combination of the values in $\{v_i(X'_i, X^*) | X'_i \in \mathcal{X}\}$, with a strictly positive weight on $v_i(X_i, X^*)$. Since (i) the maximum value in $\{v_i(X'_i, X^*) | X'_i \in \mathcal{X}\}$ is $v_i(X^*, X^*)$ because X^* is a best response to X^* and (ii) $v_i(X_i, X^*) < v_i(X^*, X^*)$ because X^* is a unique best response to X^* , \bar{v} is strictly less than $v_i(X^*, X^*)$.

Given $\bar{v} < v_i(X^*, X^*)$ and $v_i(x^*, X) \geq v_i(X_i, X)$ for all $X_i \in \mathcal{X}$, we conclude that the payoff in (8) is strictly larger than the one in (9).

Second, suppose that j 's current policy set of the opponent is X^* and compare the payoff of announcing X^* and the payoff of announcing $X_i \neq X^*$. The former payoff is at most

$$v_i(X^*, X^*) \tag{10}$$

and the latter payoff is

$$e^{-\lambda_j t} v_i(X_i, X) + (1 - e^{-\lambda_j t}) \bar{v}', \tag{11}$$

where $\bar{v}' < v_i(X^*, X^*)$ for the same reasons as above.

Given this and our assumption that X^* is a best response to X^* , the payoff in (10) is strictly larger than the one in (11).

Overall, under σ , at any $-t$ such that j 's current policy set profile is X or X^* , i 's unique best response is to announce X^* , so she assigns probability one to X^* . This proves the theorem.

E Proof of Theorem 3

We first prove that conditions (2)–(4) hold for public monitoring. Using this result, we prove that conditions (2)–(4) hold for private monitoring.

E.1 Public Monitoring

We prove that conditions (2)–(4) hold for public monitoring:

Lemma 3. *Suppose $v_A(X_A, X_B) + v_B(X_B, X_A) = 1$ for each $(X_A, X_B) \in \mathcal{X} \times \mathcal{X}$. If an MPE exists, then there exists $v_{i,t}(\theta)$ for each $\theta \in \mathcal{X} \times \mathcal{X}$ such that, for any SPE $\bar{\sigma} \in \check{\Sigma}$ and h^t , we have*

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

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

and

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

The intuition for this result is simple: Since the game is constant-sum, the value is uniquely determined by the minimax theorem.⁵⁸

Proof. Fix an MPE $\bar{\sigma}^*$, and let $v_{i,t}(\theta)$ be candidate i 's value under $\bar{\sigma}^*$ when the current remaining time is t and the policy set profile is θ . By definition,

$$v_{i,t}(\theta(h^t)) = \sup_{\sigma_i \in \Sigma_i} u_i(\sigma_i, \bar{\sigma}_j^* | h^t) = u_i(\bar{\sigma}^* | h^t), \quad (12)$$

where $u_i(\sigma | h^t)$ is candidate i 's continuation payoff given strategy profile σ .⁵⁹

Suppose there exists h^t such that $W_t^i(\bar{\sigma}, h^t, no) \neq v_{i,t}(\theta(h^t))$. Without loss, assume $W_t^i(\bar{\sigma}, h^t, no) > v_{i,t}(\theta(h^t))$.⁶⁰

$$v_{i,t}(\theta(h^t)) < W_t^i(\bar{\sigma}, h^t, no) = \inf_{\sigma_j \in \Sigma_j} u_i(\bar{\sigma}_i, \sigma_j | h^t) \leq u_i(\bar{\sigma}_i, \bar{\sigma}_j^* | h^t) \leq u_i(\bar{\sigma}^* | h^t),$$

which contradicts (12). The proofs for $W_t^i(\bar{\sigma}, h^t, X_i)$, $\hat{W}_t^i(\bar{\sigma}, h^t, X_j)$, $W_t^i(\bar{\sigma}, h^t, \theta_i(h^t))$, and $\hat{W}_t^i(\bar{\sigma}, h^t, \theta_j(h^t))$ are analogous, so are omitted. \square

E.2 Private Monitoring

Consider the “private monitoring” setup. Fix any $\sigma_j \in \Sigma_j$ (not necessarily an equilibrium strategy). Given h_i^t , let $\vec{\theta}(h_i^t) = \left((t_i^l, X_i^l)_{l=1}^{l_i}, (t_j^l, X_j^l)_{l=1}^{l_j} \right)$ denote the part of h_i^t that is public information across both candidates.⁶¹ By (19) in the Online Appendix, candidate i 's belief about h_j^t does not depend on candidate i 's private history. Denote by $\beta^{\sigma_j}(h_j^t | \vec{\theta}(h_i^t))$ a belief to be explicit about the fact that the belief is solely determined by σ_j and $\vec{\theta}(h_i^t)$.

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

$$w_t^i(\sigma_i, \sigma_j, h_i^t, X_i) = \int_{h_j^t} u_i(\sigma_i, \sigma_j | (h_i^t, X_i), h_j^t) d\beta^{\sigma_j}(h_j^t | h_i^t)$$

⁵⁸Gensbittel et al. (2017) show that the minimax theorem extends to revision games with finite actions and payoffs. Here, since we allow the action set to be infinite, we offer the formal proof of the lemma below.

⁵⁹To simplify the notation, we use the same u for both private and public monitoring.

⁶⁰Since the game is constant-sum, given $W_t^i(\bar{\sigma}, h^t, no) < v_{i,t}(\theta(h^t))$, we have $W_t^j(\bar{\sigma}, h^t, no) > v_{j,t}(\theta(h^t))$. The rest of the proof goes through with i and j reversed.

⁶¹The precise definition is provided in the Online Appendix.

be candidate i 's payoff when she takes X_i given h_i^t , given that (i) candidate i takes a continuation strategy determined by σ_i and history (h_i^t, X_i) for $(-t, 0]$,⁶² and (ii) candidate j takes a continuation play determined by σ_j and history h_j^t for $(-t, 0]$. Note that (i) candidate i 's decision X_i does not affect candidate i 's belief $\beta^{\sigma_j}(\cdot|h_i^t)$; and (ii) the belief $\beta^{\sigma_j}(\cdot|h_i^t)$ does not depend on whether candidate i obtains an opportunity at time $-t$ by the independence of the Poisson processes.

Since $d\beta^{\sigma_j}(h_j^t|h_i^t)$ is independent of h_i^t given $\vec{\theta}(h_i^t)$,⁶³ for each h_i^t, \tilde{h}_i^t satisfying $\vec{\theta}(h_i^t) = \vec{\theta}(\tilde{h}_i^t)$, we have

$$\begin{aligned} \sup_{\sigma_i \in \Sigma_i} w_t^i(\sigma_i, \sigma_j, h_i^t, X_i) &= \sup_{\sigma_i \in \Sigma_i} \int_{h_j} u_i(\sigma_i, \sigma_j | (h_i^t, X_i), h_j^t) d\beta^{\sigma_j}(h_j^t | \vec{\theta}(h_i^t)) \\ &= \sup_{\sigma_i \in \Sigma_i} \int_{h_j} u_i(\sigma_i, \sigma_j | (\tilde{h}_i^t, X_i), h_j^t) d\beta^{\sigma_j}(h_j^t | \vec{\theta}(\tilde{h}_i^t)) \\ &= \sup_{\sigma_i \in \Sigma_i} w_t^i(\sigma_i, \sigma_j, \tilde{h}_i^t, X_i). \end{aligned}$$

The second-to-last line follows since the distribution of the final profiles of policy sets that candidate i can induce given σ_j depends only on $\beta^{\sigma_j}(h_j^t | \vec{\theta}(h_i^t))$ and $\theta(h_i^t)$. Hence, we can write

$$w_t^i(\sigma_j, \vec{\theta}, X_i) = \sup_{\sigma_i \in \Sigma_i} w_t^i(\sigma_i, \sigma_j, h_i^t, X_i) \quad (13)$$

for each $h_i^t \in H_i$ with $\vec{\theta}(h_i^t) = \vec{\theta}$.

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

$$w_t^i(\sigma_j, \vec{\theta}, \theta_i(h_i^t)) = \sup_{\sigma_i \in \Sigma_i} w_t^i(\sigma_i, \sigma_j, h_i^t, \theta_i(h_i^t)) = \sup_{\sigma_i \in \Sigma_i} w_t^i(\sigma_i, \sigma_j, h_i^t, no) \quad (14)$$

since, by (20) imposed in the Online Appendix, (i) candidate i 's belief about h_j^t is the same between $(h_i^t, \theta_i(h_i^t))$ and (h_i^t, no) and (ii) given h_j^t , candidate j 's continuation play after $(h_i^t, \theta_i(h_i^t))$ and that after (h_i^t, no) are the same. In addition, let $\hat{w}_t^j(\sigma_i, \vec{\theta}, X_i)$ be candidate j 's value when candidate i takes X_i at time $-t$ given $\vec{\theta}$.

Together with the constant-sum assumption, we can show that $w_t^i(\sigma_j, \vec{\theta}, X_i) + \hat{w}_t^j(\sigma_i, \vec{\theta}, X_i) = 1$

⁶²Recall that, in Section 3, we define $u_i(\sigma_i, \sigma_j | h_i^t, h_j^t)$. Here, we define $u_i(\sigma_i, \sigma_j | (h_i^t, X_i), h_j^t)$ analogously, conditional on the event that candidate i takes X_i at time $-t$.

⁶³This assumption is formalized by (19) in the Online Appendix.

for any PBE σ .

Lemma 4. *Suppose $v_A(X_A, X_B) + v_B(X_B, X_A) = 1$ for each $(X_A, X_B) \in \mathcal{X} \times \mathcal{X}$. For any PBE σ , the following holds: Fix $v_i \in [0, 1]$, $t \geq 0$, $\vec{\theta}$, and X_i . Then, the following two claims hold:*

1. *If we have $w_t^i(\sigma_j, \vec{\theta}, X_i) > v_i$, then we have $\hat{w}_t^j(\sigma_i, \vec{\theta}, X_i) < 1 - v_i$.*
2. *If we have $w_t^i(\sigma_j, \vec{\theta}, X_i) < v_i$, then we have $\hat{w}_t^j(\sigma_i, \vec{\theta}, X_i) > 1 - v_i$.*

Proof. By symmetry, we only prove part 1. Given PBE $\sigma \in \Sigma$, if $w_t^i(\sigma_j, \vec{\theta}, X_i) > v_i$, then (13) implies that candidate i obtains the payoff strictly greater than v_i after any h_i^t satisfying $\vec{\theta}(h_i^t) = \vec{\theta}$. Since the election is constant-sum, there must exist h_j^t satisfying $w_t^j(\sigma_j, \sigma_i, h_j^t, X_i) < 1 - v_i$ and $\vec{\theta}(h_j^t) = \vec{\theta}$. Since j takes a best response under σ , it implies that $\sup_{\sigma'_j} w_t^j(\sigma'_j, \sigma_i, h_j^t, X_i) < 1 - v_i$. Recalling that we have $\sup_{\sigma'_j} w_t^j(\sigma'_j, \sigma_i, h_j^t, X_i) = w_t^j(\sigma_i, \vec{\theta}, X_i)$, we have $w_t^j(\sigma_i, \vec{\theta}, X_i) < 1 - v_i$. \square

Proving the following lemma will be sufficient for conditions (2)–(4) to hold in private monitoring:

Lemma 5. *Suppose $v_A(X_A, X_B) + v_B(X_B, X_A) = 1$ for each $(X_A, X_B) \in \mathcal{X} \times \mathcal{X}$, and there exists an MPE in public monitoring. Take $v_{i,t}(\theta)$ that satisfies conditions stated in Lemma 3. Then, for any h_i^t , we have*

$$\begin{aligned} w_t^i(\sigma, h_i^t, X_i) &= v_{i,t}(X_i, \theta_j(h_i^t)); \\ \hat{w}_t^i(\sigma, h_i^t, X_j) &= v_{i,t}(\theta_i(h_i^t), X_j); \end{aligned}$$

and

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

Proof. Suppose that there exists a PBE $\tilde{\sigma} \in \Sigma$ such that, for some $i \in \{A, B\}$ and h_i^t , we have $w_t^i(\tilde{\sigma}_j, h_i^t, X_i) \neq v_{i,t}(X_i, \theta_j(h_i^t))$. Without loss,⁶⁴ we can assume

$$w_t^i(\tilde{\sigma}_j, h_i^t, X_i) > v_{i,t}(X_i, \theta_j(h_i^t)). \quad (15)$$

⁶⁴If $w_t^i(\tilde{\sigma}_j, h_i^t, X_i) < v_{i,t}(X_i, \theta_j(h_i^t))$, then since the game is constant-sum, we have $v_{j,t}(X_i, \theta_j(h_i^t)) = 1 - v_{i,t}(X_i, \theta_j(h_i^t))$. From Lemma 4, we have $w_t^j(\tilde{\sigma}_i, h_j^t, X) > v_{j,t}(X_i, \theta_j(h_i^t))$. The following lemma goes through with indices i and j being reversed.

Fix $\bar{\theta}_j = \theta_j(h_i^t)$. From Lemma 4, for each \tilde{h}_j^t with $\theta(\tilde{h}_j^t) = (X_i, \bar{\theta}_j)$, candidate j 's expected payoff is less than $1 - v_{i,t}(X_i, \bar{\theta}_j)$.

First, recall that candidate i 's Markov strategy can be represented as a map $\sigma_i : \mathcal{X} \times \mathcal{X} \times [0, T] \rightarrow \Delta(\mathcal{X})$. Let M_i be 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 . Fix an MPE $(\sigma_i, \sigma_j) \in M_i \times M_j$. We have

$$\sup_{\sigma'_i \in M_i} \hat{W}_t^i(\sigma'_i, \sigma_j, h^t, X_i) = v_{i,t}(X_i, \bar{\theta}_j) \quad (16)$$

for each h^t with $\theta_j(h^t) = \bar{\theta}_j$.

Since this strategy σ_j is Markov, candidate j can take this strategy in private monitoring. We have $w_t^i(\tilde{\sigma}_i, \sigma_j, \hat{h}_i^t, X_i) > v_{i,t}(X_i, \bar{\theta}_j)$ for some \hat{h}_i^t with $\vec{\theta}(\hat{h}_i^t) = \vec{\theta}$ by (15) since otherwise candidate j would like to deviate to σ_j from $\tilde{\sigma}_j$ given each h_j^t with $\vec{\theta}(h_j^t) = \vec{\theta}$ in private monitoring and obtain the expected payoff no less than $1 - v_{i,t}(X_i, \bar{\theta}_j)$. Hence, by (14), we have $w_t^i(\sigma_j, \vec{\theta}, X_i) > v_{i,t}(X_i, \theta_j(h^t))$.

By (14), for each \tilde{h}_i^t with $\vec{\theta}(\tilde{h}_i^t) = \vec{\theta}$, we have

$$\begin{aligned} w_t^i(\sigma_j, \tilde{h}_i^t, X_i) &= \sup_{\sigma_i \in \Sigma_i} w_t^i(\sigma_i, \sigma_j, \tilde{h}_i^t, X_i) \\ &= \sup_{\sigma_i \in \Sigma_i} \int_{h_j^t} u_i(\sigma_i, \sigma_j | (\tilde{h}_i^t, X_i), h_j^t) d\beta^{\sigma_j}(h_j^t | \vec{\theta}). \end{aligned}$$

Since $\sigma_j \in M_j$, candidate j 's continuation strategy depends only on $\theta_t = (X_i, \bar{\theta}_j)$. Hence,⁶⁵ for each h_j^t , we can write

$$\sup_{\sigma_i \in \Sigma_i} u_i(\sigma_i, \sigma_j | (\tilde{h}_i^t, X_i), h_j^t) = \sup_{\sigma_i \in M_i} u_i(\sigma_i, \sigma_j | \theta_t = (X_i, \bar{\theta}_j)).$$

⁶⁵This is a standard result in dynamic programming: Fix candidate j 's Markov strategy σ_j , and see $\sup_{\sigma_i \in \Sigma_i} u_i(\sigma_i, \sigma_j | (\tilde{h}_i^t, X_i), h_j^t)$ as the maximization problem of a single decision maker (candidate i). Since the environment is Markov, the value depends only on t and θ_t , and there exists a solution which only depends on t and θ_t .

Therefore,

$$w_t^i(\sigma_j, \tilde{h}_i^t, X_i) = \sup_{\sigma_i \in M_i} u_i(\sigma_i, \sigma_j | \theta_t = (X_i, \bar{\theta}_j)) = v_{i,t}(X_i, \bar{\theta}_j) \text{ by (16).}$$

This is a contradiction. Thus, for each PBE $\sigma \in \Sigma$, we have $w_t^i(\sigma_j, h_i^t, X_i) = v_{i,t}(X_i, \theta_j(h^t))$.

The proofs for the other equalities can be done similarly. \square

F Long Ambiguity Theorem in Non-Constant-Sum Elections

In this section, we present a version of the long ambiguity theorem that is applicable to elections that are not constant-sum. In order to carry out a proof in such a general environment, we restrict attention to elections in which candidates can either be completely ambiguous or specify a precise policy. Formally, we assume

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

This class of elections does not include the multi-issue election campaign with valence (Section 2.3), but is general enough to encompass the valence election campaign and the symmetric office-motivated election campaign.

Before starting the analysis, let us introduce some notation and assumptions. 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 $X_i \in \mathcal{X}$, let $BR_j(X_i)$ be the set of candidate j ’s best responses against candidate i ’s policy set X_i :

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

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

In the main text, we show that, under the constant-sum and first-mover disadvantage assumptions, long ambiguity prevails. We will show the same result without the constant-sum assumption. Another difference from the result in the main text is that we fully characterize the equilibrium play, in addition to examining whether long ambiguity holds or not.

We say that $X_i^* \subseteq X$ is candidate i 's **optimal set** if the following hold for each $x_i^* \in X_i^*$.⁶⁶

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. $x_i^* \in BR_i(X)$.

If candidate j has a unique best response to each policy of candidate i , the first part implies that taking any action in the optimal set is the best for candidate i if she knows that j will have a move afterwards. The condition is a generalization of it, taking into account the possibility of multiple best responses by j . The second part further requires that if the opponent is using ambiguous language, taking any policy in the optimal set is a static best response. For general (v_i, v_j) , it is straightforward that the definition of the optimal set ensures that there exists 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 .⁶⁷

Assumption 1. For each candidate i , the largest optimal set X_i^* is non-empty, and satisfies the following properties.⁶⁸

1. For any $x_i^* \in X_i^*$ and $X_j, X_j' \in BR_j(x_i^*)$, $v_i(x_i^*, X_j) = v_i(x_i^*, X_j')$ holds.
2. For any $x_i^*, x_i^{**} \in X_i^*$, $v_j(x_i^*, X) = v_j(x_i^{**}, X)$ and $\max_{x_j \in X} v_j(x_i^*, x_j) = \max_{x_j \in X} v_j(x_i^{**}, x_j)$.

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. Note that once i enters at x_i , j chooses some $X_j \in BR_j(x_i)$ when he receives an opportunity in any PBE. Hence, part 1 implies that 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, part 2 requires that, 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.

⁶⁶A policy is in the largest optimal set if it is a Condorcet winner and the stage game is constant-sum. The reason is that the first condition of an optimal set is implied by $\min_{X_j \in \mathcal{X}} v_i(X_i, X_j) \leq v_i(X_i, x^*) < v_i(x^*, x^*)$ and $x^* \in BR_j(x^*)$, and the second condition is implied by $v_i(x^*, X) \geq v_i(X_i, X)$ for each $X_i \in \mathcal{X}$.

⁶⁷For 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 " $v_i(x_i^*, X) > v_i(x_i, X)$."

⁶⁸There does not exist an optimal set (so there is no largest optimal set) in the environment considered in Remark 11 in Section 2.3. There, it is shown that W , when he enters, would choose policy $(1, 1)$, which is in the policy set that S has committed to. But if W expects that S will have an opportunity, then W would prefer to enter at $(0, 1)$, which is outside of the policy set S has committed to.

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. For an arbitrary $x_i^* \in X_i^*$, define

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

Assumption 1 implies that this is candidate i 's expected payoff at time $-t$ when she enters (in any PBE). Moreover, 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.

- 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 .
- Case 3: $v_A(X, X) > v_A(x_A^*, X)$ and $v_B(X, X) < v_B(x_B^*, X)$.⁶⁹

That is, in Case 1, given that the opponent's policy is ambiguous, each candidate strictly prefers to stay ambiguous at the deadline. In Case 2, in contrast, given that the opponent's policy is ambiguous, each candidate strictly prefers to clarify their policies at the deadline. Case 3 is an asymmetric situation in which, given that the opponent's policy is ambiguous, A strictly prefers to stay ambiguous while B strictly prefers to clarify his policy at the deadline.

F.1 Case 1: No Candidate Enters at the Deadline

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

Proposition 6. *Consider Case 1. Under Assumptions 1 and 2, there exists a PBE. In any PBE, at histories in (X, X) , each 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.

⁶⁹The case with $v_A(X, X) < v_A(x_A^*, X)$ and $v_B(X, X) > v_B(x_B^*, X)$ is symmetric.

F.2 Case 2: Both Candidates Enter at the Deadline

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 $(-t, 0]$ upon receiving an opportunity. Such a payoff is well defined due to Assumption 1.⁷⁰ Let

$$t_i^* \equiv \inf \{t > 0 : \bar{v}_{i,t}(\text{not}) \geq v_{i,t}(\text{enter})\}.$$

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.

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

Note that the first part of this assumption is a strengthening of Assumption 2. Since t_A^* and t_B^* are generically different if they are not both infinity, 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.⁷¹

Proposition 7. *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 $t_i = t_i^*$.*

This proposition offers a characterization of the equilibrium dynamics in any PBE in Case 2. It implies that each candidate i 's policy announcement changes before and after a cutoff time t_i : She keeps using ambiguous language before $-t_i$, and clarifies after $-t_i$. In particular, if $t_i < \infty$, then candidate i does not enter when the deadline is sufficiently far. The following condition, which is

⁷⁰The formal expression of this payoff is complicated, so we relegate it to the Online Appendix.

⁷¹To 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 $v_A(X, X) \neq w$.

stronger than the condition for candidate i in Assumption 2, is a sufficient condition for $t_i < \infty$:⁷²

$$\text{First-mover disadvantage}^* \text{ for } i \left\{ \begin{array}{l} v_i(X, x_j^*), v_i(x_i^*, X), v_i(X, X) \geq v_i^{BR_j} \\ \max_{X_i \in \mathcal{X}} v_i(X_i, x_j^*) > v_i^{BR_j} \end{array} \right. . \quad (17)$$

The second line of this condition states that, if the order of the moves is known, then being the first mover is strictly 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 F.4, we explain that this condition holds under some cases in our applications.

Proposition 8. *For each i , Proposition 7 holds with $t_i < \infty$ if we additionally require first-mover disadvantage* for i to hold.*⁷³

Remark 13. The same conclusion as in Propositions 7 and 8 hold in the more general case in which there exist $t_0 \geq 0$ and a number $v_{i,t_0}(X, X)$ such that (i) there exists a PBE if the horizon length is t_0 , (ii) 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 (iii) $v_{i,t_0}(\text{enter}) > v_{i,t_0}(X, X)$ holds for each i .⁷⁴ The proofs provided in the Online Appendix are conducted for this general case. The generalization will be useful in analyzing Case 3. □

F.3 Case 3: Only One Candidate Enters at the Deadline

We define $\bar{v}_{i,t}^A(\text{not})$ 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.⁷⁵ Such a payoff is well defined due to Assumption 1.

⁷²When restricted to the environment in which $\mathcal{X} = \{X\} \cup \{\{x\} | x \in X\}$ holds, this is neither stronger or weaker than the first-mover disadvantage condition defined in (1) in Section 4.1. In particular, first-mover disadvantage* refers to the values of $v_i(X, x_j^*)$, $v_i(x_i^*, X)$, and $v_i(X, X)$, which first-mover disadvantage does not.

⁷³This 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 $t_i^* < \infty$.

⁷⁴" $v_i(X, X) < v_i(x_i^*, X)$ for each i " as stated in Case 2 corresponds to taking $t_0 = 0$.

⁷⁵The superscript denotes the candidate who does not enter close to the deadline in Case 3.

Let

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

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 9. *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 $-t \in (-\infty, -\hat{t}_B)$.*

The first part of this proposition implies that, if $\hat{t}_A < \hat{t}_B$, we can use the generalized version of Proposition 7 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) (see Remark 13). The second part states that, if $\hat{t}_A > \hat{t}_B$, no candidate enters at any $-t \in (-\infty, -\hat{t}_B)$.

Corresponding to (17), define:

$$\text{Strong first-mover disadvantage* for } i \begin{cases} v_i(x_i^*, X) > v_i^{BR_j} \\ (17) \text{ holds if } \hat{t}_A < \hat{t}_B \end{cases}.$$

That is, strong first-mover disadvantage* requires, in addition to first-mover disadvantage*, that it is strictly worse for i who has entered at x_i^* if j best-responds afterward than if j keeps being ambiguous afterward. Putting the two parts of Proposition 9 together, we can show that, under the assumptions imposed in that proposition, in any PBE, candidate i spends a long time using ambiguous language if strong first-mover disadvantage* for i holds.⁷⁶

Proposition 10. *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 $-t \in (-\infty, -t_i)$.*

In proving this theorem, we use strong first-mover disadvantage* for B to show that $\hat{t}_B < \infty$, which is necessary for B not to enter when the deadline is far away. More precisely, suppose that $\hat{t}_B = \infty$, which means that B would be better off entering at any time $-t$ if, after $-t$, A never enters before B enters. This implies that B 's PBE continuation payoff at time $-t$ approaches v_B^{BRA} as $t \rightarrow \infty$ because the probability of A 's receiving an opportunity and best-responding after B 's entry approaches one. But then, there exists a finite $t' < 0$ and $T' < \infty$ such that, for $t > T'$, B has a profitable deviation from time $-t$ onward: It is to keep using ambiguous language until time $-t'$ and then entering afterward. This is a profitable deviation because the expected payoff from such a deviation is a strict convex combination of $v_B(x_i^*, X)$, $v_B(X, X)$, and v_B^{BRA} , all of which are weakly greater than v_B^{BRA} and $v_B(x_B^*, X) > v_B^{BRA}$ by strong first-mover disadvantage* for B where the weight on $v_i(x_i^*, X)$ is bounded away from 0. Since B has a profitable deviation, $\hat{t}_B = \infty$ cannot hold in any PBE. Thus, we have $\hat{t}_B < \infty$. The intuition is simple: If B enters too early, then it is highly likely that A receives an opportunity after B 's entry. Strong first-mover disadvantage* ensures that such an event is not favorable for B , and hence B would rather wait until close to the deadline than enter too early.

F.4 Summary

We are now ready to state our first general prediction:

Theorem 4 (Non-Constant-Sum Long Ambiguity). *Under Assumptions 1 and 2, the following claims are true.*

⁷⁶The proof shows that, if strong first-mover disadvantage* for B holds, then $\hat{t}_B < \infty$ must hold.

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 $v_j(X, X) < v_j(x_j^*, X)$ 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 $-t \in (-\infty, -t_k)$.

That is, under the assumptions imposed, in any PBE, candidates keep announcing X for a long time when the horizon is sufficiently long. 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 4, 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 4 applies.⁷⁷

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 $-t \in (-\infty, 0]$.

⁷⁷One might think that part 2 of Theorem 4 can be applied to the analysis of the policy-motivated election campaign that we present in Appendix G. However, Assumption 1 fails because the optimal set is empty in that application. 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 empty. The application 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.

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 4, in contrast, pins down the sufficient condition for when we expect long ambiguity.

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

We 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): $X = \{(x_1, x_2) \in [0, 1]^2 : x_1 + x_2 \leq 1\}$.⁷⁸ 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 $(\frac{1}{2}, \frac{1}{2})$, respectively.⁷⁹ Their ideal policies are common knowledge, and the voters correctly believe that the candidate

⁷⁸The specific interpretation we give to the policy space we study may not be consistent with this episode of Sanders vs. Clinton that we discussed in Section 5. 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.

⁷⁹When we need to distinguish between the two candidates, we use a masculine pronoun for L and a feminine pronoun for R .

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 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.⁸⁰ 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 2.1. 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 $\varepsilon > 0$.⁸¹ 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, 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 $\arg \min_{(x_1, x_2) \in \bar{X}} |x_1 - x_2|$. 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

⁸⁰ 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 of readability.

⁸¹To 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.

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.⁸²

Proposition 11. *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 $(\frac{1}{2}, \frac{1}{2})$ 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 $(\frac{2}{3}, 0)$ or $(0, \frac{2}{3})$ at $-t \in (-t_L^{**}, -t_L^*)$, and enters at $(\frac{1}{2}, \frac{1}{2})$ 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 achieve a positive probability of winning. If the deadline is very far, then since candidate R will enter with a very high probability once L enters, 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, $(\frac{1}{2}, \frac{1}{2})$. 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

⁸²The proof of Proposition 11 shows uniqueness of the minimizer.

candidate R entering after L becomes sufficiently small, and then enters at $(\frac{1}{2}, \frac{1}{2})$. 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 $(\frac{2}{3}, 0)$ or $(0, \frac{2}{3})$. 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 $(\frac{2}{3}, 0)$ or $(0, \frac{2}{3})$) 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 a model with such preferences of candidates that we call the ‘‘persuasion cost election campaign,’’ and show that the equilibrium dynamics in such a model are simpler.

Remark 14 (Outcome-equivalence for a public-monitoring model). The PBE we characterize in this section (as well as the PBE characterized in the model of persuasion-cost election campaign) is Markov-perfect where the state consists of the current remaining time t and the policy sets at $-t$ (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.⁸³ Moreover, we solve the equilibrium by backward induction, which 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. This in particular implies that the continuation payoffs are ‘‘identical’’ in those two models.⁸⁴ □

Remark 15 (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 the valence election campaign (Section 2.1) and the symmetric office-motivated election campaign (Section 2.2), ambiguity is present because each candidate does not want to be the first mover.

⁸³See Section 4.3 for the formal description of such a model.

⁸⁴In Sections 2.1 and 2.2, such identity is a consequence of Theorem 3 in Section 4.3. Hence, the prediction in Theorem 3 is robust in the current setting which has non-zero-sum payoffs.

In the policy-motivated election campaign in this section, this effect is still present, while there is potentially another reason to be ambiguous: Not specifying a policy gives a flexibility in choosing a preferred policy after being elected. \square

H Dynamic Extension of the Model of Dragu and Fan (2016)

Let us first explain the model by Dragu and Fan (2016). Candidate $i \in \{1, 2\}$ has an ideal policy vector $p^i = (p_1^i, \dots, p_n^i) \in \mathbb{R}^n$. There is a continuum of voters with measure one, and their ideal policy $x \in \mathbb{R}^n$ is distributed according to $x \sim N(\mu, \Sigma)$ where Σ is a diagonal matrix. The voter with ideal policy x derives utility U from a candidate with policy p , defined as follows:

$$U = - \sum_{k=1}^n w_k(a) (p_k - x_k)^2,$$

where $w_k(a)$ is “awareness” of the k^{th} policy issue, which depends on the candidates’ advertising profile $a = (a^1, a^2) = ((a_1^1, \dots, a_n^1), (a_1^2, \dots, a_n^2)) \in \mathbb{R}^{2n}$ with a_k^i being candidate i ’s spending on issue k . Specifically, $w_k(a) := f(a_k^1 + a_k^2) / \sum_{k'=1, \dots, n} f(a_{k'}^1 + a_{k'}^2)$, where $f : \mathbb{R}_+ \rightarrow \mathbb{R}_{++}$ is twice continuously differentiable and weakly concave. Each voter votes for the candidate that generates the higher U . This determines the vote share of each candidate i , denoted $W_i(a)$. If candidate i advertises a_k^i for each issue k , then she pays the cost $\sum_{k=1, \dots, n} c^i(a_k^i)$, where c^i is twice continuously differentiable, increasing, weakly convex, $c^i(0) = c^{i'}(0) = 0$, and $\lim_{a \rightarrow \infty} c^i(a) = \lim_{a \rightarrow \infty} c^{i'}(a) = \infty$. Each candidate maximizes her vote share minus the advertisement cost.

Dragu and Fan (2016) prove that, in any pure-strategy equilibrium, there is no policy issue k in which both candidates spend positive advertising resources.⁸⁵ The intuition is that, since the payoffs from the election (that is, the payoffs except for the advertisement cost) are constant-sum, whenever one candidate has a positive marginal value of advertising in issue k , the other candidate has a negative marginal value.

We now explain the dynamic model, and will see how the conclusion changes in the dynamic case by numerical examples. We assume that each candidate obtains opportunities to advertise according to an independent Poisson process with arrival rate λ . Whenever candidate i has an

⁸⁵This is their Proposition 1. They have other results, but we only focus on a single proposition to keep the comparison simple.

opportunity, she can spend $a^i \in \mathbb{R}^n$ with the cost $\sum_k c^i(a_k^i)$. Let $\theta_t \in \mathbb{R}_+^n$ be the vector of the summations of the total spending up to time $-t$. Candidate i 's payoff from the dynamic game is $W_i(\theta_0) - \sum_{t \in \mathbb{T}_i} \sum_k c^i(a_{k,t}^i)$, where \mathbb{T}_i is the set of times in which candidate i has an opportunity, and $a_{k,t}^i$ denotes the amount i spent on dimension k at time $-t$.

We focus on Markov perfect equilibrium (MPE), where the continuation strategy from $-t$ depends only on θ_t . Since $c^i(0) = 0$ and $\lim_{a \rightarrow \infty} c^i(a) = \lim_{a \rightarrow \infty} c^{i'}(a) = \infty$ while the payoff from winning the election is bounded by one, there exists $\bar{a} \in \mathbb{R}$ such that we can assume that $a^i \in [0, \bar{a}]^n$ for any a^i played after any history in any MPE. Since the feasible choice set is compact, the objective is continuous in θ_0 , and players obtain finitely many Poisson opportunities with probability one, a proof analogous to the one in Lovo and Tomala (2015) shows that a MPE exists.

For a fixed MPE, the Bellman equation for candidate i 's value function $V_{i,t}(\theta)$ is given by

$$\dot{V}_{i,t}(\theta) = \lambda \max_{a^i \in \mathbb{R}_+^n} [V_{i,t}(\theta + a^i) - V_{i,t}(\theta) - C^i(a^i)] + \lambda \left(V_{i,t}(\theta + a_t^j(\theta)) - V_{j,t}(\theta) \right),$$

where the opponent best-responds:

$$a_t^j(\theta) \in \arg \max_{a^j \in \mathbb{R}_+^n} V_{j,t}(a^j) - C(a^j)$$

where the choice of $a_t^j(\theta)$ only depends on t and θ (but not on histories), and the initial condition satisfies

$$V_{i,0}(\theta_0) = W_i(\theta_0).$$

In principle, for any parameter specification of the model, a solution to the above differential equation provides an equilibrium characterization. In what follows, we apply the same principle to numerically solve the dynamic extension of an example provided in Dragu and Fan (2016).⁸⁶

Example 5 (Example 1 of Dragu and Fan [2016]). Let $k = 3$ and suppose that the feasible spending level is binary: $a^i \in \{0, 1\}^3$. Candidate 1 has an advantage in issue 1, while candidate 2 has an advantage in issue 3: $v_1^1 = 10$, $v_2^1 = 0$, $v_3^1 = -2$, $\lambda_1 = 1$, $\lambda_2 = 100$, $\lambda_3 = 1$, $f(\tilde{a}) = \tilde{a} + 1$, and $C(a) = \sum_k 0.05 \times a_k$ for each candidate.

In the static case, the unique pure-strategy Nash equilibrium is $a^1 = (1, 0, 0)$ and $a^2 = (0, 1, 0)$.

⁸⁶Details are available upon request.

In the dynamic case, for each on-path θ_t , we have $a_1^i > 0$ only if $i = 1$, $a_2^i > 0$ only if $i = 2$, and $a_3^i = 0$ for each $i \in \{1, 2\}$. Hence, there is no issue for which both candidates advertise.

Example 6. Consider the same environment as in Example 5, except that we now set $v_3^1 = -10$.

In the static case, the unique pure-strategy Nash equilibrium is $a^1 = (1, 0, 0)$ and $a^2 = (0, 0, 1)$.

In contrast, in the dynamic case, there exists a realization of the Poisson process such that candidate 1 spends on issue 2 at one point in time and candidate 2 spends on issue 2 at another point in time.

To see why this difference shows up, fix a MPE. Suppose candidate i obtains an opportunity at time $-t$. Let $v_k^i(\theta_t)$ be the marginal increase of the probability of winning by increasing the spending on the k^{th} issue at time $-t$ when the state at time $-t$ is θ_t . The dynamic analogue of Proposition 1 of Dragu and Fan (2016), that there is no policy issue k in which both candidates spend positive advertising resources in any pure-strategy equilibrium, would hold if the sign of $v_k^i(\theta_t)$ is constant for each t and each realization of on-path θ_t . In such a case, $v_k^i(\theta_t)$ is positive for one candidate for each t and θ_t and is negative for the other candidate for each t and θ_t . Hence, only the former candidate is willing to spend on issue k .

In the first example, since candidates are sufficiently asymmetric, if $v_k^i(\theta_t)$ is positive for one candidate for some t and θ_t that happens on the equilibrium path, then it is always positive for that candidate for each t and θ_t that happens on the equilibrium path. However, in the second example, candidates are symmetric and there exists a path of the realization of $\{\theta_t\}_t$ such that (i) at some time τ , $v_2^i(\theta_\tau)$ is positive for candidate 1 and negative for candidate 2, and (ii) for another τ' , $v_2^i(\theta_{\tau'})$ is negative for candidate 1 and positive for candidate 2.

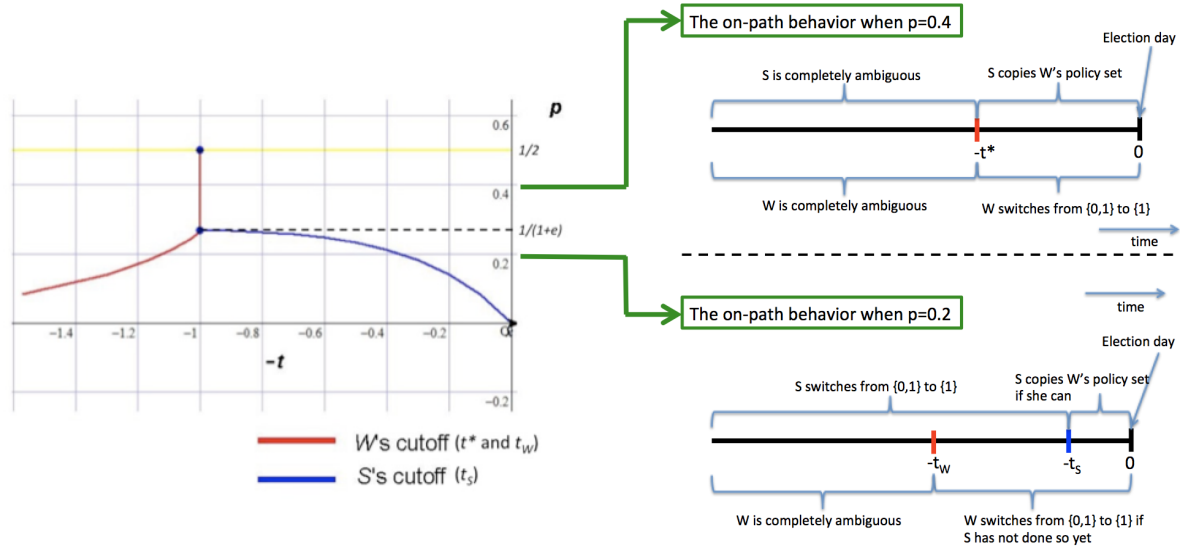


Figure I: Cutoff times and the on-path behavior for the valence election campaign ($\lambda = 1$).

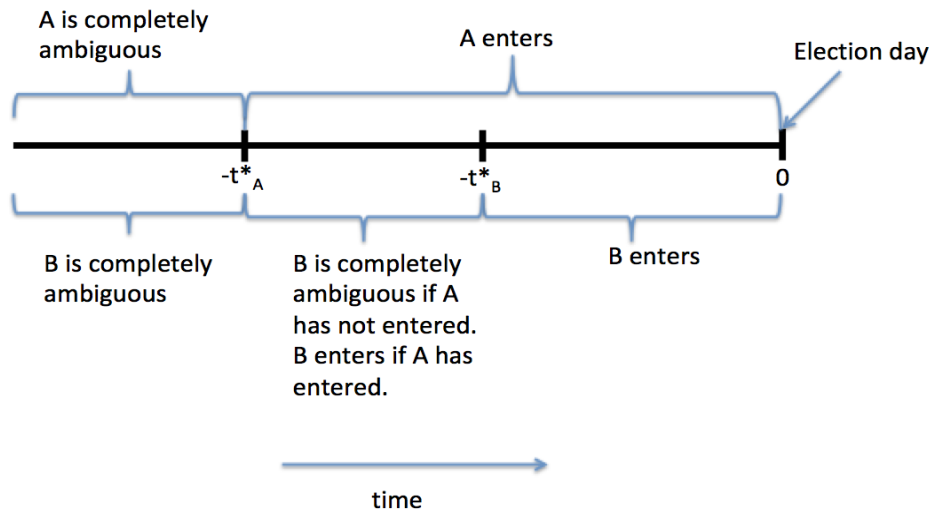


Figure II: The on-path behavior for the symmetric office-motivated election campaign.

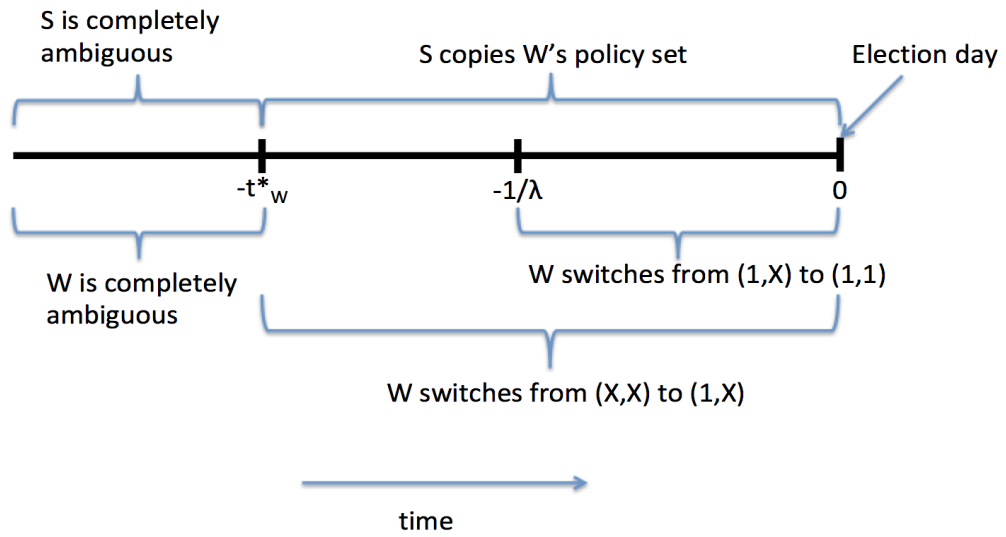


Figure III: The on-path behavior for the multi-issue election campaign with valence.