2017

Are Voters Cursed When Politicians Conceal Policy Preferences?

Nichole Szembrot
Trinity College, Hartford Connecticut, nichole.szembrot@trincoll.edu

Are Voters Cursed When Politicians Conceal Policy Preferences?

Nichole Szembrot[1]

June 15, 2017

Abstract

In campaigns, candidates often avoid taking positions on issues, concealing the policy preferences that would guide them if elected. This paper describes a novel explanation for ambiguity in political campaigns. It develops a model of candidate competition in which policy-motivated candidates can choose whether or not to announce their policy preferences to voters. It applies Eyster and Rabin's (2005) concept of cursed equilibrium, which allows for varying degrees of understanding of the connection between type (policy preference) and strategy (whether to announce). If voters updated according to Bayes' Rule, they would understand that candidates who do not announce positions are strategically concealing an unpopular policy preference. In equilibrium, only the most extreme candidates, those located furthest from the median voter's position, would choose to take no position. However, if voters do not sufficiently appreciate the informational content of a non-announcement, unraveling will not occur and both extremists and more moderate candidates will not announce positions.

[1] Department of Economics, Trinity College, 300 Summit St., Hartford, CT 06106
Tel: 860-297-4106
E-mail: nichole.szembrot@trincoll.edu

1 Introduction

While Hillary Clinton was seeking the Democratic nomination for president for 2016, a voter asked her a direct question about whether she supported the Keystone XL pipeline. Clinton responded by continuing her policy of not taking a position on the issue, telling the voter that "'if it is undecided when I become president, I will answer your question'" (Merica and Zeleny 2015).

In political campaigns, candidates often avoid taking positions on issues, instead making vague statements that conceal the policy preferences that would guide them if elected. For example, Tomz and Houweling (2009) analyzed statements made regarding tax policy by the presidential and vice-presidential candidates in debates during the 2007-2008 election cycle. They found that over half of such statements were ambiguous. For example, a candidate might say that he would not increase taxes, but would not reveal the magnitude of the tax cut he supported. Given such a statement, voters would be left unsure whether the candidate would enact extreme tax cuts, moderate tax cuts, or leave tax rates as they were.

The existence of such ambiguity is a fundamental puzzle in political economy. This paper proposes a novel explanation for ambiguity in elections: voters do not fully understand the informational content of a non-announcement. If voters were fully rational, they would be able to infer that the fact that a candidate chooses to conceal his preferences actually reveals information about those preferences: they must not be congruent with the voters' preferences. However, if voters do not follow this reasoning, then candidates may actually be able to conceal their preferences by making vague campaign statements.

To elaborate on the mechanism summarized above, this paper focuses on a particular informational asymmetry: politicians have more information about their own policy preferences

than their opponents and the electorate have. Since a candidate's willingness to reveal his true policy position depends on the likelihood that announcing it will garner support from the electorate, a candidate's decision to conceal his preferred policy may actually reveal information about that policy preference to voters. More concretely, consider the following example, in which citizens must elect a representative who will make a decision regarding gun control. Suppose that the median voter prefers a centrist gun control policy: mandatory registration and background checks. There are two political parties, each composed of three factions. Candidates from the Left party generally prefer a stricter level of gun control. The Left party includes centrists, moderate Leftists who prefer banning military-style assault rifles, and extreme Leftists who would ban private ownership of handguns. The Right party, in contrast, fields candidates who are centrists, moderates who would allow some exemptions from background checks, and extremists who would not require gun registration. Suppose that the median voter is equally dissatisfied with moving to the left or to the right of her ideal policy. Figure 1 illustrates this environment graphically.

[Figure 1 here]

A centrist candidate will not want to pool with less desirable candidates and will clearly choose to announce his preferred policy. Suppose, first, that only centrist candidates make announcements. If voters update correctly, they would infer that a Left politician who made no announcement either wants to ban assault rifles or also wants to ban handguns. Therefore, a moderate Leftist would be more appealing to voters if he revealed his preferred policy than if he pooled with extremists by taking no position. Thus, only the extreme candidates could possibly remain silent; regardless of whether they announce their true preferences or try to conceal them, voters would infer that they are extremists.

The standard model of play in games of incomplete information assumes that players are able to go through the train of logic outlined above and update their beliefs about other players' types correctly. However, the unraveling process may break down if voters lack sufficient cognitive ability or are simply unwilling to bear the effort costs that may be required to think through a complex game. If when a candidate takes no position, voters (mistakenly) put enough weight on the candidate being a centrist, they may instead prefer the candidate who took no position to a candidate who revealed himself to be a moderate Leftist or a moderate Rightist. This paper will show conditions under which moderate candidates would respond by remaining ambiguous about their future plans by developing a game-theoretic model that formalizes the ideas described above and solving it using the solution concept of cursed equilibrium (Eyster and Rabin 2005). When a voter is "cursed", she under-appreciates the relationship between a candidate's type (his policy preference) and his campaign strategy.

The cursed equilibrium formulation is intended to be a portable model that applies in many games of incomplete information. A literature based mostly on the beauty contest game has found that limited strategic thinking is correlated with cognitive ability (Burnham et al. 2009), scores on the Cognitive Reflection Test (Brañas-Garza, García-Muñoz and González 2012), and agreeableness and emotional stability (Gill and Prowse 2016), suggesting that fairly stable personal traits may explain behavior in a variety of strategic environments. In addition to the evidence from several experiments using auctions and jury voting cited in Eyster and Rabin (2005), experiments conducted more recently have found additional support for the concept of cursed equilibrium in different contexts. Crawford and Iriberri (2007) use data from several different auction experiments to estimate models that allow for heterogeneity in the cursedness parameter. They estimate that almost half of subjects are fully cursed, about one-fifth are

Bayesians, and the rest are partially cursed to some degree. Charness and Levin (2009) find that behavior in a modified acquire-a-company game is consistent with buyers being fully cursed. While some papers, such as Battaglini, Morton and Palfrey (2010) in an experimental study of the swing voter's curse, have found evidence of strategic sophistication, the evidence as a whole suggests that at least some degree of cursedness is a more realistic assumption than Bayesian updating.

The rest of the paper is organized as follows. Section 2 explains how this paper relates to the existing literature. Section 3 analyzes a model with a discrete policy space in which candidates decide whether or not to announce their preferred policies to voters before the election. In section 4, several extensions or alternative formulations of the main model are discussed: adding uncertainty regarding the location of the median voter, using a continuous policy space rather than a discrete one, and allowing candidates to make binding commitments to policies. Section 5 concludes.

2 Literature review

This paper contributes to the theoretical literature in economics and political science that attempts to explain the prevalence of ambiguity in campaigns.

One class of explanations views ambiguous statements by candidates as responses to true voter preferences. Shepsle (1972) makes an important theoretical contribution with an extension of the Downsian model in which candidate strategies are probability distributions, instead of points on a subset of the real line. He shows that, when voters care intensely about getting their most preferred option and experience diminishing sensitivity further away from their bliss points,

a challenger can defeat an incumbent whose position is known by choosing a non-degenerate lottery, rather than committing to a single point in the policy space. On the other hand, if voters are risk-averse, a candidate should not take actions that would make voters uncertain about his policy position (Bernhardt and Ingberman 1985). Aragonès and Postlewaite (2002) reinforce Shepsle's results in a model similar to his, except for the introduction of asymmetry by requiring each candidate to put some minimum probability on that candidate's "preferred" policy. These authors are also responsible for the framing of the key assumption as preference intensity, rather than risk-loving preferences. Preference intensity will also play a role in the model developed in this paper.

Several authors have noted that ambiguity on a given policy issue may be beneficial for candidates when voters' preferences are multi-dimensional. According to emphasis theory, candidates optimally choose to focus voters' attention on issues on which there is consensus, leaving no attention to divisive issues (Page 1976). Dellas and Koubi (1994) suggest that candidates hide information about their performances from voters in order to win elections based on charisma and similar characteristics. While this paper considers political competition on a single dimension, future work could consider the interaction of the mechanism analyzed here and the additional concerns that arise with competition on several dimensions.

If candidates must simultaneously appeal to moderate voters and extremist campaign contributors, as assumed in Alesina and Holden (2008), candidates with different policy preferences may pool by committing to enact a policy within an interval that contains both of their true preferences. Relatedly, candidates may face initial uncertainty regarding the location of the median voter (Alesina and Cukierman 1990; Meirowitz 2005; Agranov 2012). In these models, policy-motivated candidates do not reveal information about their preferences when

taking no position, because non-extreme candidates also have a desire to remain ambiguous when the preference of the median voter is not known with certainty. Without this assumption, introducing policy-motivated candidates would add a layer of complexity, since actions taken during campaigns may reveal information about what a candidate will do in the future. Given the present ubiquity of polls in the time leading up to elections, explanations that rely on having enough of this uncertainty may not seem as plausible as they did even a few years ago. The model presented below drops the assumption of uncertainty about the location of the median voter, implying that the only motive for keeping quiet about one's policy position is to conceal a preference for an unpopular policy. Somewhat similarly, Aragonès and Neeman (2000) posit that ambiguity exists in elections because candidates would like to have the flexibility to adjust policy after uncertainty about the optimal policy is resolved. They argue that politicians are likely to be ambiguous when the best policy, given preferences, is less certain at the time of the election; this paper suggests that ambiguity can result in the absence of such uncertainty if enough variation exists in candidate preferences.

This paper is not the first work to point out that voters should make inferences about the types of candidates who would choose to be vague; Chappell (1994) also makes that argument. However, his model differs from this paper's model in several important assumptions, and his results are not directly comparable to those that will be discussed here. Two papers are closely related to this paper in the sense that candidates are responding to behaviorally biased voters. Callander and Wilson (2008) develop a model of context-dependent voting, in which voters care about each candidate's platform directly and in comparison to his opponent's platform. Given this setup, and assuming that there is uncertainty about the valence (dis)advantage held by each candidate, each candidate chooses a probabilistic platform that puts some weight on the median

voter's position and some weight on a position closer to his own. In the model of Jensen (2009), voters interpret vague statements made by candidates that they like for non-policy reasons by projecting their own views onto such candidates. His key result is that a well-liked office-motivated candidate can defeat a candidate who commits to the median voter's position by taking an interval position containing the median. It is not clear how this bias would interact with cursedness, because Jensen does not analyze the bias within a signaling game with policy-motivated candidates.

This paper provides a novel view of political competition as a signaling game played by actors with less than complete cognizance of that strategic environment. Future work should do more to connect these disparate explanations, but an understanding of this key mechanism is a prerequisite for such work.

More broadly, this work also contributes to the theoretical literature on disclosure. The announcement game studied in this paper has the same general structure as disclosure games: players with information decide whether to reveal that information to an uninformed agent whose decision also affects the informed player. The most closely related papers consider whether standard unraveling results hold in the presence of unsophisticated consumers, beginning with Milgrom and Roberts (1986). Their "naïvely credulous" decision-makers who are completely unsophisticated with respect to game-theoretic reasoning correspond to fully cursed voters in our setting.

3 Announcement game

3.1 Environment

The purpose of the election is to select a candidate who will implement one of five possible policies. Suppose that, as illustrated in Figure 2, each policy can be represented by a point on the real line, with one unit of distance between each pair of adjacent policies.

[figure 2 here]

There are two policy-motivated candidates, one from each side of the ideological spectrum. The reader may think of them as being chosen randomly (or chosen based on characteristics that are orthogonal to policy preference) from two entrenched political parties. Candidate 1 may prefer policy A, B, or C, and Candidate 2 may prefer policy C, D, or E. There are N voters distributed across the policy space, and the median voter prefers policy C.[2]

The overall distribution of candidate preferences is symmetric around C, and this distribution is common knowledge. Let π_j^i be the probability that Candidate i's preferred policy is j. Since the distribution is symmetric, the notation can be simplified by writing $\pi_A = \pi_A^1 = \pi_E^2$, $\pi_B = \pi_B^1 = \pi_D^2$ and $\pi_C = \pi_C^1 = \pi_C^2$. Both candidates and voters receive utility based on how close the implemented policy is to their preferred policy; payoffs are described in more detail below.

In this game, all candidates will implement their preferred policies. The decision for a candidate in the campaign stage is whether or not to announce this preferred policy to voters.

3.2 Timing

---

[2] The median voter theorem (Black 1958) implies that, since majority rule is assumed and preferences are single-peaked, the outcome of the election will be determined by the preference of the median voter; hence, analysis is greatly simplified by focusing on the preference of the median voter. Additionally, assuming that the median voter prefers C is consistent with work done by Anderson and Meagher (2012), who show that, with endogenous party formation in a continuous environment, the parties' moderate boundaries lie on either side of the median voter.

The timing of the announcement game is as follows:

Stage 0: Nature draws a preferred policy for each candidate, and each candidate learns his own preferred policy.

Stage 1: The candidates simultaneously choose campaign platforms. Each candidate chooses whether or not to announce his preferred policy.

Stage 2: Voters observe the campaign platforms and choose their preferred candidates. If a voter is indifferent, she abstains.[3] The candidate with the most votes wins the election and implements his preferred policy. If both candidates receive the same number of votes, a winner is chosen at random, with each candidate having an equal chance of being chosen.

3.3 Payoffs

Each player chooses his strategy to maximize his expected utility. Each player i receives a utility payoff that depends on his preferred policy, denoted by x_i, and the policy implemented, $\bar{x}$. This implies that candidates are purely policy-motivated and receive no additional reward from holding office.[4] Let z_i denote the distance between player i's preferred policy and the policy

---

[3] Nothing changes if she instead votes randomly when indifferent.
[4] Allowing a reward from holding office does not change the results of the announcement game. As will be shown below, each candidate will choose the strategy that maximizes his chance of winning the election. Adding a reward from holding office would only increase the difference in utility between the strategy that maximizes winning probability and the other strategy.

implemented: $z_i = |x_i - \bar{x}|$. Each player's preferences are represented by $u(z)$, where $u(\cdot)$ is

decreasing in the distance between the preferred policy and the policy implemented. For ease of

notation, define the following values: for $i = 1, \ldots, 5$, $u_i = u(i-1)$. Each player receives u_1 if his

preferred policy is implemented, u_2 if his second-favorite policy is implemented, and so on.

Without loss of generality, u_5 is set to 0.

3.4 Equilibria

Let θ_1 and θ_2 denote the types of Candidate 1 and Candidate 2, respectively. A strategy for

Candidate 1, $\sigma_1(\theta_1)$, is a mapping from preferred policy $\theta_1 \in \{A, B, C\}$ to campaign platform

$a_1 \in \{\theta_1, \varnothing\}$; he chooses between announcing his preferred policy or announcing nothing. The

strategy for Candidate 2, $\sigma_2(\theta_2)$, is defined analogously. A strategy for voter i, $\sigma_{vi}(a_1, a_2)$, is a

mapping from observed candidate platforms a_1 and a_2 to a voting action $a_{vi} \in \{1, 2, 0\}$. She

chooses whether to vote for Candidate 1, vote for Candidate 2, or abstain (denoted by 0) after

viewing the platforms.

The following definition applies the standard perfect Bayes-Nash equilibrium solution

concept to this game.

Definition 1 The perfect Bayes-Nash equilibrium of the announcement game is a strategy profile

σ^{BN} and a set of beliefs $\pi_{\theta|a}$ for voters such that

1. For each candidate j, $\sigma_j^{BN}(\theta_j)$ maximizes his utility, given $\sigma_{-j}^{BN}(\theta_{-j})$ and

   $\sigma_{v1}^{BN}(a_1, a_2), \ldots, \sigma_{vN}^{BN}(a_1, a_2)$.

2. For each voter i, $\sigma_{vi}^{BN}(a_1, a_2)$ maximizes her utility, given $\sigma_1^{BN}(\theta_1)$, $\sigma_2^{BN}(\theta_2)$, and her beliefs $\pi_{\theta|a}$.

3. Each voter forms posterior beliefs $\pi_{\theta|a}$ using Bayes' Rule.

Voters' beliefs about the preferred policy of a candidate who makes an announcement are determined by that announcement, since candidates can only make truthful statements. A voter who did not expect a candidate to announce that he prefers A/E nonetheless believes the statement if she does hear it. When a voter observes that a candidate takes no position, she must form beliefs about the candidate's preferred policy, since that is the policy that he would implement if elected. In a perfect Bayes-Nash equilibrium, all players use Bayes' Rule to update their beliefs about the likelihood that a candidate has each preferred policy conditional on having taken no position.

Alternatively, the game can be solved using Eyster and Rabin's (2005) solution concept of cursed equilibrium. This concept is a Bayes-Nash equilibrium, with a modification of the requirement that all players use Bayes' Rule when updating their beliefs. Eyster and Rabin (2005) model the error as players forming beliefs as if other players play their type-contingent strategies with probability $1 - \chi$ and draw a strategy at random from the average distribution of strategies with probability χ, where $\chi \in [0,1]$ parameterizes the voter's degree of cursedness. Equivalently, a cursed voter's posterior beliefs about the types of the candidates following the campaign stage are a convex combination that puts weight $1 - \chi$ on the correct Bayesian

posteriors and weight χ on the prior beliefs. A voter who is more cursed (larger χ) has less understanding of the link between type and strategy chosen. If $\chi = 0$, this concept is simply Bayes-Nash equilibrium. At the other extreme, if voters are fully cursed ($\chi = 1$), they do not update their beliefs at all based on actions taken by other players. In this case, when a voter sees that a candidate does not make a policy announcement, her belief that the candidate has a given policy preference is simply her prior belief that the candidate has that policy preference, based on the distribution of candidate preferences.

Definition 2 The cursed equilibrium of the announcement game is a strategy profile σ^* and a set of beliefs $\tilde{\pi}_{\theta|a}$ for voters such that

1. For each candidate j, $\sigma_j^*(\theta_j)$ maximizes his utility, given $\sigma_{-j}^*(\theta_{-j})$ and

$$\sigma_{v1}^*(a_1, a_2), \ldots, \sigma_{vN}^*(a_1, a_2).$$

2. For each voter i, $\sigma_{vi}^*(a_1, a_2)$ maximizes her utility, given $\sigma_1^*(\theta_1)$, $\sigma_2^*(\theta_2)$, and her beliefs $\tilde{\pi}_{\theta|a}$.

3. Each voter forms posterior beliefs $\tilde{\pi}_{\theta|a} = (1 - \chi)\pi_{\theta|a} + \chi\pi$, where $\pi_{\theta|a}$ is constructed using Bayes' Rule.

Candidates are assumed to be Bayesians. If candidates were cursed, their beliefs about the type of their opponent conditional on the opponent's strategy would be biased. However, as will be shown below, only the probability of winning the election matters when a candidate chooses his strategy. Since beliefs about the policy that would be implemented if the opponent wins are not relevant to the decision, the assumption that candidates are Bayesians is without loss of generality in the announcement game.

We appeal to Selten's (1975) concept of trembling-hand perfection to rule out perfect Bayes-Nash equilibria that hinge on particular off-path beliefs when no candidate takes no position in equilibrium. In the perturbed game, each player chooses his non-equilibrium strategy with some small probability. This implies that "no position" is on the equilibrium path of the perturbed game, and Bayes' Rule can be used to update beliefs about the preferred policy of a player who chose "no position." Since there is positive probability that a candidate who took no position does not prefer A/E in the perturbed game, a voter would choose a candidate who took no position over a candidate who announced A/E. Therefore, it is a dominant strategy for candidates who prefer A/E to take no position; given that the other candidate may tremble, the candidate has a higher chance of winning by taking no position. On the basis of this argument, we restrict attention to equilibria in which candidates who prefer A/E choose to take no position.

Attention is further restricted to equilibria in which players do not choose weakly dominated strategies. This restriction rules out the possibility that a voter would be indifferent between the candidates because she believes that her vote will not be pivotal. Under this assumption, one can model a voter's decision problem as if she alone determines the outcome of the election. Since only two candidates are on the ballot and voting is costless, voting for her favorite candidate based on the policy she expects him to implement if elected weakly dominates all other voting strategies.

Voters' decisions are straightforward given most combinations of candidate platforms; interesting decisions involve a candidate who took no position. As explained above, the median voter would vote for a candidate who took no position over a candidate who announced that he prefers A/E. Similarly, because trembles imply a positive probability that a candidate who took no position does not prefer C regardless of the equilibrium strategies, the median voter would

also vote for a candidate who announced that he prefers C rather than a candidate who took no position. We will return to the choice between a candidate who announced a preference for B/D and a candidate who took no position below.

For candidates who prefer C, announcing is a weakly dominant strategy, because the median voter would prefer a candidate who announced C to a candidate who chose any other strategy. It is only weakly dominant because a candidate who prefers C would be indifferent between winning the election and having an opponent who also prefers C win the election. As explained above, candidates who prefer A/E have a dominant strategy to take no position. Therefore, discussion of possible equilibria will focus on strategies chosen by candidates who prefer B/D.

The first possible equilibrium is one in which all information about candidates' types is revealed to Bayesian voters in equilibrium. The label attached to the equilibrium defined just below derives from this fact.

Definition 3 *Revelation equilibrium:* Candidates who prefer A/E take no position, and all other types announce their preferred policies. The median voter chooses a candidate who announced B/D over a candidate who took no position.

To understand the decision problem faced by candidates who prefer B/D, one must begin by understanding the voter behavior to which the candidates are responding. Voters' choices depend on their beliefs about the likely preferred policy of a candidate who took no position. The following lemma describes the beliefs held by Bayesian and cursed voters under the assumption that only candidates who prefer A/E choose to take no position.

Lemma 1 *Suppose that in equilibrium, only candidates who prefer A/E choose to take no position.*

1. A Bayesian voter believes the following about a candidate who took no position:

$$\pi_{A|\varnothing} = 1$$
$$\pi_{B|\varnothing} = 0 \quad\quad (1)$$
$$\pi_{C|\varnothing} = 0$$

2. A cursed voter believes the following about a candidate who took no position:

$$\tilde{\pi}_{A|\varnothing} = 1 - \chi(1 - \pi_A)$$
$$\tilde{\pi}_{B|\varnothing} = \chi\pi_B \quad\quad (2)$$
$$\tilde{\pi}_{C|\varnothing} = \chi\pi_C$$

Proof See supplementary materials.

Bayesian voters understand that since only candidates who prefer A/E take no position in equilibrium, a candidate who took no position must prefer A/E. In contrast, a cursed voter believes that a candidate who took no position may prefer B/D or C. The following proposition uses Lemma 1 to show that candidates who prefer B/D are better off announcing their preferred policies than concealing them when facing Bayesian voters, but this is not necessarily true when voters are cursed.

Proposition 1

1. The revelation equilibrium exists as a perfect Bayes-Nash equilibrium.

2. The revelation equilibrium exists as a cursed equilibrium if and only if

$$\chi \leq \chi_R \equiv \frac{u_2 - u_3}{(1 - \pi_A)(u_2 - u_3) + \pi_C(u_1 - u_2)} \quad\quad (3)$$

Proof See supplementary materials.

Given Bayesian voters' behavior, the best response of a candidate who prefers B/D is to announce his preferred policy, since he would be more likely to win the election by doing so than by taking no position. However, if cursed voters prefer a candidate who took no position to one who announced B/D, then candidates who prefer B/D would choose to deviate from their equilibrium strategies of announcing their preferred policies. Given the distribution of candidate preferences and the median voter's preferences over outcomes, a threshold degree of cursedness exists; if voters are sufficiently close to Bayesian, then the revelation equilibrium holds.

Alternatively, candidates who prefer B/D may choose to take no position by making ambiguous campaign statements.

Definition 4 *Ambiguity equilibrium*: Candidates who prefer A/E or B/D take no position, and candidates who prefer C announce their preferred policies. The median voter prefers a candidate who took no position to a candidate who announced B/D.

The following lemma gives the posterior beliefs that voters hold, given that candidates who prefer both A/E and B/D choose to take no position.

Lemma 2 *Suppose that in equilibrium, only candidates who prefer C announce their preferred policies.*

1. A Bayesian voter believes the following about a candidate who took no position:

$$\pi_{A|\varnothing} = \frac{\pi_A}{1 - \pi_C}$$

$$\pi_{B|\varnothing} = \frac{\pi_B}{1 - \pi_C} \tag{4}$$

$$\pi_{C|\varnothing} = 0$$

2. A cursed voter believes the following about a candidate who took no position:

$$\tilde{\pi}_{A|\varnothing} = \frac{\pi_A \left(1 - \chi \pi_C\right)}{1 - \pi_C}$$

$$\tilde{\pi}_{B|\varnothing} = \frac{\pi_B \left(1 - \chi \pi_C\right)}{1 - \pi_C} \tag{5}$$

$$\tilde{\pi}_{C|\varnothing} = \chi \pi_C$$

Proof See supplementary materials.

In this case, voters will be uncertain about the preferred policy of a candidate who took no position, because more than one type chooses that strategy in equilibrium. However, a Bayesian would still know for certain that a candidate who took no position cannot prefer policy C. Cursed voters, on the other hand, perceive more uncertainty than actually exists because they assign some probability on a candidate who prefers C choosing to take no position.

The next proposition establishes that choosing to remain ambiguous cannot benefit candidates who prefer B/D when voters are Bayesians, but ambiguity may exist in equilibrium if voters are cursed.

Proposition 2

1. The ambiguity equilibrium does not exist as a perfect Bayes-Nash equilibrium.

2. The ambiguity equilibrium exists as a cursed equilibrium if and only if

$$\chi \geq \chi_{Amb} \equiv \frac{\pi_A \left(u_2 - u_3 \right)}{\pi_C \left[\pi_A \left(u_2 - u_3 \right) + \left(1 - \pi_C \right) \left(u_1 - u_2 \right) \right]} \tag{6}$$

Proof See supplementary materials.

If Bayesian, the median voter would prefer a candidate who announced B/D to a candidate who offered a lottery over B/D and A/E by taking no position. Based on this voter preference, a candidate who prefers B/D would gain by deviating from taking no position to announcing B/D. Therefore, the ambiguity equilibrium does not exist as a perfect Bayes-Nash equilibrium.

In contrast, if the median voter is cursed, she believes that a candidate who took no position may prefer C, because she does not fully appreciate how candidates' preferences affect their platform choices. These distorted beliefs lead her to prefer a candidate who took no position to a candidate who announced a preference for B/D. In response to the median voter's behavior, a candidate who prefers B/D optimally chooses to take no position.

The next proposition establishes the conditions under which the revelation and ambiguity equilibrium are unique. As is clear from the conditions discussed above, the revelation equilibrium is the unique perfect Bayes-Nash equilibrium. However, depending on the value of χ, there may be multiple cursed equilibria.

Proposition 3 *When restricting attention to equilibria in which players do not choose weakly dominated strategies and that are robust to trembles:*

1. *If $\dfrac{\pi_A}{\pi_C} < \dfrac{u_1 - u_2}{u_2 - u_3}$, then:*

    a. $0 < \chi_{Amb} < \chi_R < 1$

    b. *The revelation equilibrium is unique if and only if*

$$\chi < \chi_{Amb} \equiv \frac{\pi_A \left(u_2 - u_3 \right)}{\pi_C \left[\pi_A \left(u_2 - u_3 \right) + \left(1 - \pi_C \right) \left(u_1 - u_2 \right) \right]}$$

c.   *The ambiguity equilibrium is unique if and only if*

$$\chi > \chi_R \equiv \frac{u_2 - u_3}{\left(1 - \pi_A \right) \left(u_2 - u_3 \right) + \pi_C \left(u_1 - u_2 \right)}$$

d.   *The revelation and ambiguity equilibria both exist (and no other*

equilibria exist) if and only if $\chi_{Amb} \leq \chi \leq \chi_R$

2.   *If* $\dfrac{\pi_A}{\pi_C} = \dfrac{u_1 - u_2}{u_2 - u_3}$, *then the revelation equilibrium is unique if* $\chi < 1$, *and*

both the revelation equilibrium and the ambiguity equilibrium exist if

$\chi = 1$.

3.   *If* $\dfrac{\pi_A}{\pi_C} > \dfrac{u_1 - u_2}{u_2 - u_3}$, *then the revelation equilibrium is unique, for all values of*

χ .

Proof   See supplementary materials.

Note that a set of parameter values exists for which both the revelation and ambiguity equilibria emerge. It is not clear that one of these equilibria is more focal than the other. As shown below, both the median voter and candidates who prefer B/D are better off under the revelation equilibrium, which may lead that equilibrium to be selected. However, given the myriad of other reasons for candidates to be ambiguous discussed in section 2, it may also be natural for players to coordinate on the ambiguity equilibrium.

If $\frac{\pi_A}{\pi_C} > \frac{u_1 - u_2}{u_2 - u_3}$, then $\chi_R \geq 1$, and the revelation equilibrium exists and the ambiguity

equilibrium does not, regardless of the degree of cursedness. This condition holds if and only if

the median voter prefers receiving u_2 for certain to a random draw from the distribution of

candidate preferences. The cursed equilibrium model implies that the most optimistic belief that

a voter may hold about a candidate who took no position is that he is a random draw, rather than

disproportionately unlikely to be an extremist. If the median voter still prefers a candidate who

would implement policy B/D, even if fully cursed, then cursedness has no effect on the

predictions of the model; the revelation equilibrium exists.

In order for it to be true that $\frac{\pi_A}{\pi_C} < \frac{u_1 - u_2}{u_2 - u_3}$ with risk-neutral or risk-averse voters, it must be

the case that the share of extremists is smaller than the share of centrists. It could also hold with a

larger share of extremists if voters care more about moving away from their ideal point than

about differences further from their ideal point, but this would imply that voters are not averse to

risk. While the underlying distribution of candidate preferences is unobservable, the empirical

distribution of individual preferences suggests that the assumption that the share of extremists is

not too large is reasonable. According to 2012 data from the American National Election Studies,

35% of respondents who reported an ideology considered themselves to be middle of the road.

The share who consider themselves moderate, slightly liberal, or slightly conservative is 62%,

while the share who consider themselves extremely liberal or extremely conservative is just 8%.

The next set of results show how these thresholds are affected by changes in the distribution

of candidate preferences and voters' risk preferences. To analyze the effect of changing the

distribution of candidate preferences, the next proposition considers an increase in the share of

extremists, while reducing shares of moderates and centrists in fixed proportions.

Proposition 4 *Suppose that π_A increases by ε, while π_B decreases by $\lambda\varepsilon$ and π_C decreases by $(1-\lambda)\varepsilon$ for $\lambda \in [0,1]$. Then, χ_R and χ_{Amb} increase.*

Proof See supplementary materials.

Recall that an increase in χ_R implies that the revelation equilibrium is more likely to exist; similarly, an increase in χ_{Amb} implies that the ambiguity equilibrium is less likely. If candidates who prefer A make up a larger fraction of the candidate pool, then it is more likely that a candidate who took no position has that preferred position (and less likely that he has a position that is better for the median voter). Thus, a voter must really not understand the link between candidate type and strategy if she is still willing to vote for a candidate who took no position.

The following proposition shows that risk aversion and cursedness work in opposite directions. This paper defines one set of preferences u as more risk-averse than another set of preferences v if and only if the certainty equivalent of any lottery is less under u than under v.

Proposition 5 *Suppose that utility function u represents more risk-averse preferences than utility function v. Then χ_R and χ_{Amb} are higher under u than under v.*

Proof See supplementary materials.

3.5 Welfare

Proposition 3 showed that for a range of parameter values, the ambiguity equilibrium exists if the median voter is sufficiently cursed, while the revelation equilibrium exists if the median voter

holds beliefs that are sufficiently close to Bayesian. Given that cursedness may lead to the ambiguity equilibrium being played instead of the revelation equilibrium, it is natural to consider the welfare consequences.[5]

Proposition 6

1. The expected utility for players who prefer C and B/D is higher in the revelation equilibrium than in the ambiguity equilibrium.

2. The expected utility for players who prefer A/E is higher in the revelation equilibrium than in the ambiguity equilibrium if and only if

$$u_4 - u_5 > u_1 - u_2 \qquad (7)$$

Proof See supplementary materials.

Note that while the analysis focused on the median voter, who prefers C, other voters are distributed across the policy space. Therefore, one may care about the welfare of players who prefer B/D and A/E not only because one cares about the candidates' welfare, but also because many voters also prefer those policies. While players who prefer C and B/D are clearly better off when more moderate candidates have a better chance of being elected, the welfare effects of cursedness on players who prefer A/E are ambiguous and depend on the shape of the utility function. While a player who prefers A is better off if Candidate 1 is able to conceal his preference for A and be elected, it may be a concern that a Candidate 2 who prefers E is able to do the same thing. Therefore, the ambiguity equilibrium is better only for risk-loving extreme

---

[5] A literature in public choice (e.g., Gersbach 1995, 2000) analyzes the conditions under which public information improves collective decisions. In that environment, the outcome delivered by each alternative is uncertain (as are candidate preferences in our model); however, in contrast to our model, the alternatives themselves are not candidates who have strategic incentives to reveal information.

candidates. It is particularly notable that candidates who prefer B/D--the ones who change their behavior when voters are cursed--would prefer to live in a world in which the median voter is a Bayesian. In other environments, firms might take advantage of a behavioral bias to increase profit. To borrow an example from O'Donoghue and Rabin (2003), potato chip manufacturers are better off when consumers are present-biased than when their preferences are time-consistent. In contrast, moderate politicians take no position in response to the voters' behavior, but they would be better off if the voters understood incentives well enough to allow the moderate candidates to announce their preferred policies.

4 Extensions

4.1 Uncertainty about the location of the median voter

Shifting the median voter's preferred policy changes the threshold value of χ that, conditional on the candidates' equilibrium strategies, induces the median voter to switch from voting for a candidate who announced B to voting for a candidate who took no position. Therefore, uncertainty about the location of the median voter translates into uncertainty about that threshold value of χ. Since a candidate's optimal strategy is the one that maximizes his chance of winning the election, the revelation equilibrium exists if and only if the probability that a majority of voters have a degree of cursedness below the χ_R threshold (as defined for each voter based on her preferred policy) is greater than or equal to $\frac{1}{2}$. This implies that if the median voter's preferred policy is drawn from a uniform distribution centered around C, the equilibrium conditions are unchanged from the case in which the median voter prefers C for certain.

4.2 Continuous policy space

An analysis of the announcement game with a continuous policy space can be found in section 2 of the supplementary materials. The main result is that a coarse policy space is necessary in order for cursedness to lead to ambiguity in equilibrium. A cut-off strategy in which candidates who are more extreme than some threshold in the interior of the policy space take no position cannot be supported, even with cursed voters. In many real-world applications, a coarse policy space is the appropriate assumption. Gun control, as discussed in the introduction, is one example. However, when the policy is more continuous, such as the choice of a tax rate, one should not expect that cursedness will lead to ambiguity.

4.3 Binding commitments

While it may be reasonable to assume that candidates will always pursue their preferred policies, it may be possible for a candidate to credibly make a binding commitment to implement a different policy. This section summarizes the robustness of the results from the announcement game to the alternative assumption that candidates may make binding commitments to any policy; a more detailed discussion and proofs of related propositions can be found in section 3 of the supplementary materials. The structure of this game differs from the announcement game only in the campaign stage:

Stage 1: The candidates simultaneously choose campaign platforms. Candidate 1 may choose to make a binding commitment to policy A, B, or C or to take no position; Candidate 2 may choose to commit to policy C, D, or E or to take no position. A candidate who does not make a commitment is free to implement his preferred policy if he wins the election.

Since candidates now have the option of committing to policies that are not their preferred policies, more equilibria are possible. In particular, a trade-off now exists between choosing a strategy that would allow the candidate to implement a more preferred policy and choosing a strategy that would give the candidate a greater chance of winning the election.[6] A candidate in the commitment game may choose to commit to a policy that is not his preferred policy in order to increase his chance of winning the election, thereby preventing his opponent from implementing an even worse alternative.

In the commitment game, taking no position may not be the strategy chosen by candidates who prefer B/D, even if they believe that the median voter prefers candidates who took no position to candidates who committed to B/D. Taking no position is a risky strategy, since it allows him to implement his preferred policy if elected, but he risks receiving a low payoff if his opponent wins. He can choose instead to commit to C, which would guarantee him his second-favorite policy.

In equilibrium, all candidates choose to commit to the median voter's preferred policy if they are risk-averse. Cursedness still makes taking no position more attractive; however, taking no position is too risky, given that the safer strategy of committing to the voter's preferred policy is

---

[6] The assumption that candidates are Bayesians matters in the commitment game, because cursed candidates do not correctly predict the policy that would be implemented by an opponent who does not take a position. However, this has a small effect on incentives, only slightly changing the degree of risk aversion that would lead a candidate to choose one strategy over another.

available. In the announcement game, candidates who prefer B/D or A/E could only choose between risky strategies; in that environment, cursedness plays an important role in determining whether moderate candidates choose to announce their preferences to voters or to take no position. In contrast, when candidates can make binding commitments to any policy, risk-averse candidates choose to commit to policy C, regardless of the voters' degree of cursedness.

5 Conclusion

To summarize, this paper shows how the outcomes of elections can be affected if voters do not understand the strategic incentives faced by candidates who choose whether or not to reveal their preferred policies ahead of Election Day. If a candidate does not take a position, voters should infer something about his preferred policy from this action. However, if voters do not make these inferences, then a candidate may respond by taking no position, even when he prefers a moderate policy.

Given that voters may not make the correct inferences and may reduce their own welfare by doing so, policies that would require disclosure could improve outcomes. In their review of the disclosure literature, Dranove and Jin (2010) find empirical support for the prediction that mandated disclosure leads consumers to choose higher-quality products. Similarly, policies that provided candidates with incentives to disclose their policy preferences can help voters make better choices. For instance, debate moderators may limit the speaking time of candidates who dance around questions about their policy preferences. Whether a news organization that sponsors a given debate would prefer to do this is another question, as the organization may believe that having an extremist politician to cover may boost ratings. However, if an

organization cares enough about the public interest or perceives its own interests as being served by not having an extremist in office, it may choose to try to encourage disclosure if it realizes that voters would not make the correct inferences in the absence of such disclosure.

A key difference between the stylized model and a real-world election is that, in actual elections, it is not always clear whether a candidate has taken no position on an issue or has not discussed a particular issue simply because it is not salient. In the model, a single policy dimension was assumed, and it was common knowledge that candidates would have the opportunity to announce a position. Suppose instead that candidates have made a decision about a position (or lack thereof) on each issue at the beginning of the campaign. With some probability, an issue gains attention and the candidate's strategy of announcing or taking no position is revealed. If not, voters do not observe that the candidate has taken a position on the issue. For example, if voters do not observe a candidate making a statement about Social Security reform, they would be uncertain about whether the candidate has a campaign plank on the issue that has not been revealed because of lack of media attention or whether the candidate intends not to take a position on Social Security reform. The primary difference here is that different types of issues have different probabilities of being made salient. For instance, in modern US elections, the probability that attention will be drawn to a candidate's position on healthcare policy approaches one, while it is much less likely that the media will focus on potential nominees to obscure positions in government.

While candidates in this paper were restricted to either making truthful announcements or binding commitments, actual candidates have a larger set of possible strategies. For instance, they may choose to lie (perhaps with some cost). Extending the theory in these directions would require adjusting the refinements that restrict the beliefs about other players' types that players

may hold after observing an action that is off the equilibrium path. For example, if a voter expected that all candidates would lie and announce that they shared the median voter's position, what should voters believe if a candidate announced a position slightly to the left of the median voter's position? Refinements such as divinity (Banks and Sobel 1987) require that voters consider which types of candidates would be more likely to make a given deviation; this kind of strategic thinking is at odds with the spirit of the cursed equilibrium model. Future research should develop "behavioral" refinements to perfect Bayes-Nash equilibrium that permit beliefs that are consistent with limited strategic thinking.

Additionally, candidates may be heterogeneous with respect to other characteristics that voters care about, such as integrity or decisiveness. Since such traits may also affect a candidate's decision about whether to remain ambiguous, a platform of "no position" can be a signal of both policy preference and other character traits. For example, a candidate who faces internal costs when making decisions about taking action may also face similar costs when deciding which policy to espouse during a campaign; taking no position would then be a signal of indecisiveness. An extension of the model to this more complicated signaling game is left for future work.

References

Agranov, M. (2013). *Flip-flopping, intense primaries and the selection of candidates*. Available at SSRN: http://dx.doi.org/10.2139/ssrn.2412252

Alesina, A. & Cukierman, A. (1990). The politics of ambiguity. *The Quarterly Journal of Economics, 105*(4): 829-850.

Alesina, A.F. & Holden, R.R. (2008). *Ambiguity and extremism in elections*. NBER Working Paper 14143.

The American National Election Studies, Table 3.1

http://www.electionstudies.org/nesguide/toptable/tab3_1.htm

Anderson, S.P. & Meagher, K.J. (2012). *Choosing a champion: party membership and policy platform.* CEPR Discussion Paper 8941.

Aragonès, E. & Postlewaite, A. (2002). Ambiguity in election games. *Review of Economic Design, 7*(3): 233-255.

Aragonès, E. & Neeman, Z. (2000). Strategic ambiguity in electoral competition. *Journal of Theoretical Politics, 12*(2): 183-204.

Banks, J.S. & Sobel, J. (1987). Equilibrium selection in signaling games. *Econometrica, 55*(3): 647-661.

Battaglini, M., Morton, R.B, & Palfrey, T.R. (2010). The swing voter's curse in the laboratory. *Review of Economic Studies, 77*(1): 61-89.

Bernhardt, M.D. & Ingberman, D.E. (1985). Candidate reputations and the "incumbency effect." *Journal of Public Economics, 27*(1): 47-67.

Black, D. (1958). *The Theory of Committees and Elections*. Cambridge: Cambridge University Press.

Brañas-Garza, P., García-Muñoz, T., & González, R.H. (2012). Cognitive effort in the beauty contest game. *Journal of Economic Behavior & Organization, 83*(2): 254-260.

Burnham, T.C., Cesarini, D., Johannesson, M., Lichtenstein, P., & Wallace, B. (2009). Higher cognitive ability is associated with lower entries in a *p*-beauty contest. *Journal of Economic Behavior & Organization, 72*(1): 171-175.

Callander, S. & Wilson, C.H. (2008). Context-dependent voting and political ambiguity. *Journal of Public Economics, 92*(3/4): 565-581.

Chappell, H.W., Jr. (1994). Campaign advertising and political ambiguity. *Public Choice, 79*(3/4): 281-303.

Charness, G. & Levin, D. (2009). The origin of the winner's curse: a laboratory study. *American Economic Journal: Microeconomics, 1*(1): 207-236.

Crawford, V.P. & Iriberri, N. (2007). Level-k auctions: can a nonequilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions? *Econometrica, 75*(6): 1721-1770.

Dellas, H. & Koubi, V. (1994). Smoke screen: a theoretical framework. *Public Choice, 78*(3/4): 351-358.

Dranove, D. & Jin G.Z. (2010). Quality disclosure and certification: theory and practice. *Journal of Economic Literature, 48*(4): 935-963.

Eyster, E. & Rabin, M. (2005). Cursed equilibrium. *Econometrica, 73*(5): 1623-1672.

Gersbach, H. (1995). Information efficiency and majority decisions. *Social Choice and Welfare, 12*(4): 363-370.

Gersbach, H. (2000). Public information and social choice. *Social Choice and Welfare, 17*(1): 25-31.

Gill, D. & Prowse, V. (2016). Cognitive ability, character skills, and learning to play equilibrium: a level-*k* analysis. *Journal of Political Economy, 124*(6): 1619-1676.

Jensen, T. (2009). Projection effects and strategic ambiguity in electoral competition. *Public Choice, 141*(1/2): 213-232.

Meirowitz, A. (2005). Informational party primaries and strategic ambiguity. *Journal of Theoretical Politics, 17*(1): 107-136.

Merica, D & Zeleny, J. (2015). Hillary Clinton won't answer Keystone XL pipeline question. *CNN.com* *http://www.cnn.com/2015/07/28/politics/hillary-clinton-keystone-xl-pipeline/* [Last accessed: 23 May 2017].

Milgrom, P. & Roberts, J. (1986). Relying on the information of interested parties. *The RAND Journal of Economics, 17*(1): 18-32.

O'Donoghue, T & Rabin, M. (2003). Studying optimal paternalism, illustrated by a model of sin taxes. *The American Economic Review: Papers & Proceedings, 93*(2): 186-191.

Page, B.I. (1976). The theory of political ambiguity. *American Political Science Review, 70*(3): 742-752.

Selten, R. (1975). Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory, 4*(1): 25-55.

Shepsle, K.A. (1972). The strategy of ambiguity: uncertainty and electoral competition. *American Political Science Review, 66*(2): 555-568.

Tomz, M. & Van Houweling, R.P. (2009). The electoral implications of candidate ambiguity. *American Political Science Review, 103*(1): 83-98.