Misinformation

Li, Hao
University of British Columbia & University of Toronto

Wei Li
University of British Columbia & University of California, Riverside

March 2010

Abstract

We model political campaigns as a costly information structure in a two-type signaling game. A candidate's type consists of private information about his own qualifications and that about his rival, and chooses both the target and the level of an information campaign, with a higher positive (negative) campaign generating a more accurate public signal about his (the rival's) qualifications. We consider both the expected-margin payoff structure, where the candidate's payoff depends on the difference in the median voter's posterior beliefs about himself and about the rival, and the winner-take-all payoff structure. Under both payoff structures, the high-type candidate has favorable information about his own qualifications and/or unfavorable information about his rival. Under expected-margin, the high type chooses the kind of campaign that he has relative advantage in to separate from the low type. The required separating level is higher when the relative advantage is smaller. Under winner-take-all, separation can occur only through different kinds of campaign. The required separating level is independent of the high type's private information, and is higher when the low type has a weaker preference between positive and negative campaigns.

JEL classification: D72, D82, D83

Keywords: Positive and negative campaign, campaign level, least cost separation

1 Introduction

Much attention has been paid to the issue of positive versus negative advertising in political campaigns in political science and in popular press (Arceneaux 2006, Lau and Rovner 2007, Sides and Grossman forthcoming). The focus of existing literature is on which type of campaign—praising one's own qualifications or discrediting that of the rival's—is more effective in influencing the voters. Less attention, however, has been paid to the informativeness of campaign advertisements, either positive or negative. Some campaign feature "feel-good" ads, such as the 1984 "Morning in America" presidential election advertising campaign, which reveal little about a candidate's qualifications, while others give more specific information about a candidate's records and character. If campaign choices, both the target and the level of informativeness of a campaign, reflect how much a candidate is willing to let the voters learn about either himself or his rival, they may influence the voters' evaluation of the candidates.

Viewing through the lens of how much information the candidate wants to reveal or to conceal in his campaign choices, we can address several important questions. When the candidate has private information about his own qualifications and about those of his rival, should a strong candidate who has favorable information about himself and unfavorable information about his rival run positive or negative campaigns? How does the level of the campaign depend on the private beliefs of the candidate? How do these campaign choices depend on the election formats? And finally, can it happen that in equilibrium voters are misinformed by the candidate's campaign choices?

In this paper, we study the campaign choices of such a candidate. The candidate may be of two types, each type consisting of a noisy signal about his own qualifications and another about his rival's qualifications. He chooses one campaign, modeled here as a costly public information structure, but not its realization, to convey his private information. He controls both the target of the campaign, which can be either his own qualifications (*positive campaign*) or his rival's (*negative campaign*), as well as the informativeness of the campaign, which is the accuracy of the public signal the voter observes. Since the candidate has private information, his private beliefs about the types of both candidates are generally different from the *interim beliefs* that the voter forms after observing his campaign choices and then uses to evaluate the realized signal generated by the campaign. In this case, we say that *misinformation* occurs.

In our first model, the candidate maximizes the expected margin of the voter's belief about his own qualifications over his rival's. In this model, information per se has no value in the sense that under complete information the candidate has no incentive to engage in costly information campaigns. In contrast, the value of misinformation to the candidate arises from misleading the voter into forming a higher belief about his own qualifications than his private belief, or a lower belief about his rival's qualifications. However, since the candidate does not control the realization of the campaign signal, misinformation is necessarily imperfect as the voter's interim beliefs after observing the candidate's campaign choice are partially corrected by the realized campaign signal. In particular, if the voter has a higher interim belief about the candidate's qualifications given his strategy than his private belief, then he prefers to "obfuscate the negative" by running an uninformative campaign because in expectation, any informative public signal is likely to lower the voter's ex post belief of his qualifications. In this case, the marginal value of misinformation is negative, because the realized campaign signal agrees with his private information in expectation, leading to a higher opinion of the rival's qualifications. In the opposite case when the voter has a lower interim belief about the candidate's qualifications given his strategy than his private belief, the candidate prefers to "accentuate the positive" by running an informative campaign about his own qualifications. The marginal value of misinformation is positive.

The high type, whose private beliefs about his own qualifications and those of his rival have a greater difference, has incentives to use the appropriate target and the level of the information campaign to separate from the low type. There is always a separating equilibrium in which the low type runs a costless, uninformative campaign, because for at least one kind of campaign the high type expects a greater payoff from the same information campaign than the low type. Misinformation does not happen in any separating equilibrium, but the high type's equilibrium campaign choices, both the target and the level, are driven by the low type's incentives to misinform the voter. If the high type has both more favorable information about his own qualifications and about his rival, separation can occur only with a positive campaign. This is because the voter is more likely to learn the rival's qualifications from a negative campaign, and the candidate loses the chance of convincing the voter of his own qualifications. Similarly, if the high type has both less favorable information about his own qualifications and about his rival, separation can occur only with a negative campaign. If the high type has more favorable information about his own qualifications and less favorable information about his rival, the less costly choice for separation depends on which kind of campaign gives him a relative advantage to separate from the low type. The relative advantage is determined by the private information of the high type: for the same difference in the private beliefs about his own qualifications and those of his rival, the high type has a relative advantage in positive campaigns if both beliefs are high and in negative campaigns if both beliefs are low.[1] Furthermore, the relative advantage in positive (negative) campaigns is stronger if both beliefs are higher (lower), and as a result, the required campaign level for separation is lower. Finally, conditional on the same relative advantage, the separating campaign level increases if the high type becomes more attractive for imitation by the low type, that is, if the high type has both more favorable information about himself and less favorable information about his rival. However, it is possible that to separate from the same low type, a high type with a greater difference in his private beliefs about his own qualifications and those of the rival runs a lower level of campaign than another high type, because the former has a strong relative advantage in one kind of campaign while the latter has little relative advantage.

In our second model, the candidate cares only about the event that the voter holds a higher posterior belief about his qualifications than about his rival. One difference in this winner-take-all model from the expected-margin model is that here information per se has value when the voter's prior belief about the candidate's qualifications is lower than that about the rival. In particular, if information campaigns are not too costly, under complete information each type will engage in the minimum level of information campaign so that with a positive probability the voter's posterior belief about his own qualifications is just equal to that about his rival. The same minimum level is required for a positive campaign to generate accurate enough evidence supporting the candidate and for a negative campaign to do the same against the rival. Whether the positive or the negative campaign yields a greater probability of winning depends on whether the voter has more favorable prior belief about the candidate's qualifications or less favorable belief about his rival. Another important difference is that in the signaling game under winner-take-all, the interim beliefs of the voter enter the candidate's payoff in a type-independent way. Thus, unlike the expected-margin model, misinformation is unchecked by the realization of the campaign signal. As a

[1] Equivalently, the relative advantage of the high type in running positive versus negative campaigns to separate from the low type is determined by latter's information: for the same difference in the private beliefs of the low type, the high type has a relative advantage in positive campaigns if both beliefs of the low type are low and in negative campaigns if both beliefs are high.

result, the two types cannot be separated using the same kind of campaign, for a lower level of campaign would both raise the chance of winning and reduce the campaign cost. Separation is possible only through different kinds of campaigns.

Similar to the expected-margin model, the high type has a more favorable private belief about his own qualifications or a less favorable private belief about his rival, or both, and would thus run a lower level of campaign under complete information that would generally invite imitation by the low type. The high type can be separated from the low type only if the former has a relative preference for the opposite kind campaign that the low type prefers under complete information; if this condition fails, pooling and equilibrium misinformation will occur. For example, if the low type prefers negative campaigns under complete information, separation can occur if the high type has a weaker preference for negative campaigns under complete information even though the high type would rather run a negative campaign under complete information. Furthermore, since the way interim beliefs of the voter affects the candidate's payoff does not depend on his type, the separating level is independent of the high type's private beliefs. Instead, the separating level is a function of the low type's private beliefs. If the low type has a stronger preference for negative campaigns under complete information, then the separating level of the positive campaign run by the high type is lower, but fewer high types can be separated. Conversely, in the limit when the low type has no preference between the two kinds of campaigns, the separating level run by the high type is maximized, but any high type may be separated from the low type by having them run opposite kinds of campaign.

In existing papers, advertising, political or otherwise, is either modeled as directly informative because they contain hard information about the candidate's own (Nelson 1974, Coate 2004), or indirectly informative (Milgrom and Roberts 1986, Prat 2002). Milgrom and Roberts (1986) shows that the amount of money a firm spends on advertising can signal its product quality. Prat (2002) considers an electoral campaign setting where voters do not learn useful information from the campaign advertising itself. Rather, they use the amount of money spent on the campaign, in terms of campaign contributions, as a costly signal of the interest groups who have private information about the candidate's qualifications. Kotowitz and Mathewson (1979) study a setting where consumers' purchase decision based on their expectations of the quality of an experience good. They show that a monopolist can use advertising to mislead the consumers into having higher expectations than the actual quality of the good, at least in the short run.

Mostly closely related to our paper, Polborn and Yi (2006) considers a model in which each candidate has two characteristics, of which he can only reveal one. He runs an informative and truthful campaign, which can be a positive one about one good characteristic of his own or a negative one about one bad characteristic of his rival. The voters rationally estimate the characteristic not revealed. Thus a candidate is more likely to choose a positive campaign when his own characteristic is good and/or his rival's is good too. Our model differs from Polborn and Yi (2006) in that the candidate chooses the informativeness of the campaign signal the voter receives, which influences the voters' beliefs about his own qualifications and his rival's qualifications.

2 The Setup

Consider the following signaling model. There are two political candidates for office, a and b, where only a is a player (the sender) in the game described below. Each candidate is either "qualified" or "unqualified." To model private information that candidate a has, we assume that there are two types of a, with each type a pair of beliefs about the qualifications of a and b. Denote the two types as (α_L, β_L) and (α_H, β_H), where the first component of each type represents a's private belief that he himself is qualified and the second component represents the belief that the rival candidate is qualified. There is a non-degenerate common prior about the type of candidate a.

Define an *information campaign* as an observable choice of information structure—a distribution of a publicly verifiable signal about the qualifications of candidate a or b. An information structure is a "positive campaign" if it generates a signal about a's qualifications, and a "negative campaign" if it is about b. For simplification, we assume that regardless of whether its target is a or b, each information structure generates a "campaign signal" that is either $\overline{s}$ or $\underline{s}$ about a candidate's qualification. Further, each information structure is represented by a number k between $\frac{1}{2}$ and 1, called the "level" of the campaign. Campaign level is both the probability of the signal being $\overline{s}$ conditional on that the targeted candidate is qualified and the probability of the signal being $\underline{s}$ conditional on that the target is unqualified.

A representative voter (the receiver) first observes candidate a's choice of campaign, which includes both the target and the level of the campaign, and then observes the realized campaign signal. To focus on information provision, we assume that the voter is not a strategic player of this game: she simply uses Bayes' rule to form a pair of posterior beliefs about the qualifications of the candidates. These beliefs, denoted as π^a and π^b, together with a campaign cost function $F(k)$, determine the payoff to candidate a in the game. Two payoff structures are considered: in the expected-margin model, candidate a maximizes the difference of the voter's posterior beliefs about himself over b. The payoff to candidate a is

$$\pi^a - \pi^b - F(k).$$

In the winner-take-all model, candidate a maximizes the chance of winning, or the probability that the voter's posterior belief is higher about himself than about b. We adopt the convention that the tie-breaking rule favors candidate a. Candidate a's payoff is

$$\begin{cases} 1 - F(k), & \text{if } \pi^a \geq \pi^b \\ -F(k), & \text{otherwise.} \end{cases}$$

We assume that F is continuous and strictly increasing, with $F\left(\frac{1}{2}\right) = 0$.

A few remarks about the setup are in order. First, we assume that candidate a cannot directly control the realization of the campaign signal, consistent with the idea of information provision. Instead, a chooses both the target and the level of the campaign, while the voter updates her beliefs about the qualifications of the two candidates based on the campaign choice and the realized public campaign signal. The idea that the voter directly observes the level, i.e., informativeness, of a campaign is of course an abstraction, but in practice the public may be able to indirectly infer it from, for example, campaign spending disclosures.

Second, we have implicitly assumed that candidate a cannot simultaneously run both a positive and a negative campaign in order to focus on the choice of the target (positive or negative) of the campaign. Correspondingly, we have also assumed that the cost function of an informative campaign is the same for positive and for negative campaigns to highlight the choice of the target. We will extend the analysis to "double campaigns" below; in the extension we assume that the total campaign cost is simply the sum of the two costs, and the public signals from the two campaigns are independent conditional on the respective candidate qualifications.

Third, as a signaling game, the above model is special because the signal is not some action as in a typical game, but an information structure. This feature creates an important analytical role for the "interim belief" held by the voter regarding the qualifications of the two candidates after observing the campaign choices made by a but before the campaign signal ($\overline{s}$ or $\underline{s}$) is realized. In a positive campaign, for example, the posterior belief π^a about the qualifications of candidate a is determined according to Bayes' rule by the interim belief $\tilde{\alpha}$ and the realized campaign signal. Of course in any (Perfect Bayesian) equilibrium the interim belief is endogenous, or, part of the equilibrium. Another distinguishing feature of the signaling game considered here is that the sender's type is multi-dimensional, and so is the sender's strategy space.[2] In our model, candidate a has both private information about his own qualifications and the rival's qualifications; and he can choose both the target and the level of the campaign.

Finally, the candidates' payoffs are modeled in a reduced form, which can be made more general without affecting the results qualitatively. The chosen payoff specifications have natural interpretations in the context of political campaigns: the expected-margin model is suitable in a proportional representation system, while the winner-take-all model is more appropriate in a plurality system. These payoff specifications may also be appropriate in the context of marketing and advertising campaigns, for instance whether a firm aims to expand its market shares or to capture a market completely.

3 Expected Margins

In this model, candidate a maximizes the difference between the voter's posterior belief about him over that about candidate b. We begin by showing how the voter's posterior belief depends on her interim belief and the kind and the level of candidate a's campaign. In a positive campaign of level k, given a's private belief about himself α and the voter's interim belief $\tilde{\alpha}$ about his qualifications, a's expected posterior belief $\Pi(\alpha, \tilde{\alpha}; k)$ is given by

$$\begin{aligned}\Pi(\alpha, \tilde{\alpha}; k) &= (\alpha k + (1-\alpha)(1-k))\frac{\tilde{\alpha}k}{\tilde{\alpha}k + (1-\tilde{\alpha})(1-k)} \\ &\quad + (\alpha(1-k) + (1-\alpha)k)\frac{\tilde{\alpha}(1-k)}{\tilde{\alpha}(1-k) + (1-\tilde{\alpha})k}, \qquad (1)\end{aligned}$$

[2]There is a small literature on multi-dimensional signaling. In Quinzii and Rochet (1985) and Engers (1987), the sender's type is multi-dimensional but the signal is one-dimensional and the focus is on the conditions for separation to occur; in Damiano, Li, and Suen (2008) the type is one-dimensional but the signal is multi-dimensional and the focus is on the equilibrium structure of pooling.

where the first fraction is the voter's posterior belief about a's qualifications after observing $\overline{s}$, and the second fraction her posterior belief after observing $\underline{s}$. For notational convenience, define the difference in these two posterior beliefs as a function of the interim belief and the campaign level:

$$\Delta(\tilde{\alpha};k) \equiv \frac{\tilde{\alpha}k}{\tilde{\alpha}k+(1-\tilde{\alpha})(1-k)} - \frac{\tilde{\alpha}(1-k)}{\tilde{\alpha}(1-k)+(1-\tilde{\alpha})k}.$$

Then, we can rewrite expression (1) as

$$\Pi(\alpha,\tilde{\alpha};k) = \tilde{\alpha} + (\alpha-\tilde{\alpha})(2k-1)\Delta(\tilde{\alpha};k).$$

Note that $\Pi(\alpha,\tilde{\alpha};k)$ is linear in a's private belief α, with a positive slope less than 1. The expression for the expected posterior belief $\Pi(\beta,\tilde{\beta};k)$ for candidate b after a negative campaign of level k, given the private belief β and the interim belief $\tilde{\beta}$, can be similarly derived. It is straightforward to establish the following properties of function Π.

Lemma 1 *(i)* $\Pi(\alpha,\tilde{\alpha};k) = \alpha$ *if* $\tilde{\alpha} = \alpha$*; (ii)* $\Pi(\alpha,\tilde{\alpha};k)$ *decreases in* k *if* $\alpha < \tilde{\alpha}$*; and (iii)* $\Pi(\alpha,\tilde{\alpha};k)$ *increases in* k *if* $\alpha > \tilde{\alpha}$.

Part (i) of Lemma 1 shows that the value of information is zero in the present expected-margin model, in the sense that under complete information there is no benefit in running any information campaign. This also implies that on the equilibrium path of any separating equilibrium, the "value of misinformation" is also zero in the expected-margin model. In general, Lemma 1 can be thought of as a characterization of the *marginal value of misinformation* through changing campaign levels. From part (iii), if candidate a is privately more confident about his qualifications than the voter is, then he can increase his average perceived type by choosing a more informative positive campaign, which highlights the good news about herself. That is, the marginal value of misinformation is positive. From part (ii), candidate a can also hide the bad news by reducing the informativeness of his campaign signal if he is privately less confident about his own qualifications than the voter. In this case, the marginal value of misinformation is negative. Naturally, the opposite holds for a negative campaign: the candidate lowers the voter's perception about his rival by running an informative campaign if he has worse news about the rival than the voter. Note that the voter's posterior beliefs are bounded by $\Pi\left(\alpha,\tilde{\alpha};\frac{1}{2}\right) = \tilde{\alpha}$, and $\Pi(\alpha,\tilde{\alpha};1) = \alpha$.

3.1 Equilibrium Campaign Choices

Given the lemma, we look for the least cost separating equilibrium, in which the "low type" runs an uninformative campaign with $k_L = \frac{1}{2}$ (it does not matter whether it is positive or negative) and the "high type" runs an informative campaign of some level $k_H > \frac{1}{2}$, either positive or negative, such that the low type is just indifferent between running k_L and k_H. The identity of the low type in such an equilibrium can be determined as follows; since we have already denoted the two types as (α_L, β_L) and (α_H, β_H), we just need to provide a necessary condition on these values for the low type in any least cost separating equilibrium to be (α_L, β_L). In any such equilibrium, type (α_H, β_H) is weakly better off running k_H than running k_L:

$$\alpha_H - \beta_H - F(k_H) \geq \alpha_L - \beta_L,$$

where the left-hand-side follows from part (i) of Lemma 1. By assumption, $F(k_H) > 0$, so a necessary condition for the low type in any least cost separating equilibrium to be (α_L, β_L) is to have lower private belief about himself relative to that of candidate b:

$$\alpha_H - \beta_H > \alpha_L - \beta_L. \tag{2}$$

Throughout this section, we assume that condition (2) holds. Under this assumption, there are three cases regarding the location of type (α_H, β_H) relative to (α_L, β_L) as shown in Figure 1. In the first case, labeled "P" in the figure, we have

$$\beta_H \geq \beta_L; \tag{3}$$

in the second case, labeled "N,"

$$\alpha_H \leq \alpha_L; \tag{4}$$

and in the third case, labeled "P/N," both (3) and (4) hold with strict inequalities.

We claim that in the P-region, where type (α_H, β_H) candidate a has both higher private belief about himself *and* about candidate b, the least cost separating equilibrium calls for a positive campaign. Suppose that instead type (α_H, β_H) runs a negative campaign of some level $k^n > \frac{1}{2}$. Since type (α_H, β_H) weakly prefers k^n to $k_L = \frac{1}{2}$ and the reverse is true for type (α_L, β_L),

$$\alpha_H - \beta_H - F(k^n) \geq \alpha_L - \beta_L \geq \alpha_H - \Pi(\beta_L, \beta_H; k^n) - F(k^n).$$

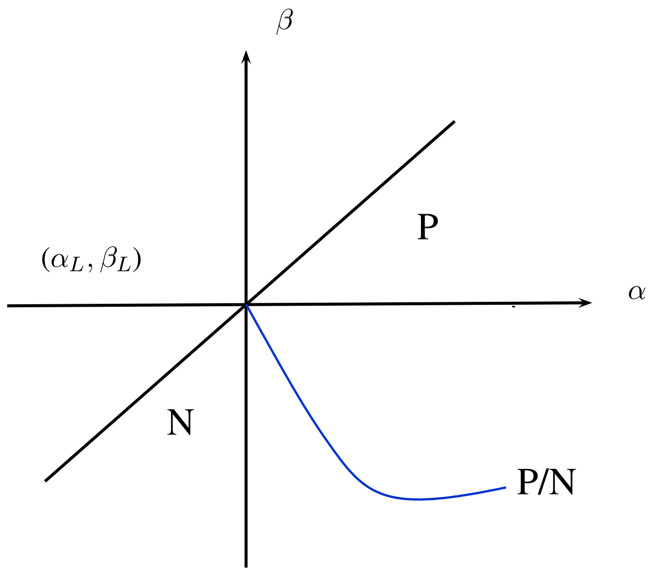

Figure 1: Expected-margin Model

This implies

$$\Pi(\beta_L, \beta_H; k^n) \geq \beta_H,$$

contradicting (3). Thus, in the P-region, separation of type (α_H, β_H) from type (α_L, β_L) can only be accomplished through a positive campaign. This is an extreme case of type (α_H, β_H) having an advantage in positive campaigns *relative* to negative campaigns in separating from type (α_L, β_L). Intuitively, because $\beta_H \geq \beta_L$, for any given level of negative campaign, type (α_L, β_L) is more successful in lowering the voter's expected posterior belief of candidate b than type (α_H, β_H). In comparison, type (α_L, β_L) candidate a is less successful in raising the voter's expected posterior belief of himself than type (α_H, β_H) in a positive campaign. A symmetric argument establishes that in the N-region, where type (α_H, β_H) candidate a has both lower private belief about himself and about candidate b, type (α_H, β_H) runs a negative campaign in the least cost separating equilibrium, which is an extreme case of type (α_H, β_H) having a relative advantage in running negative campaigns. Finally, in the P/N-region, both kinds of campaigns may be used in separation. In contrast to the other two regions, type (α_H, β_H) does not have extreme relative advantage in running one kind of campaigns over the other kind.

How informative a campaign does type (α_H, β_H) candidate a need to run to separate from type

(α_L, β_L)? In the P-region, the level k_H is given by k_H^p that satisfies the indifference condition of type (α_L, β_L):

$$\alpha_L - \beta_L = \Pi(\alpha_L, \alpha_H; k_H^p) - \beta_H - F(k_H^p). \tag{5}$$

The above is well-defined for all (α_H, β_H) in the P-region, because at $k_H^p = \frac{1}{2}$ by condition (2), the left-hand-side is strictly smaller than the right-hand-side, while at $k_H^p = 1$, the opposite is true. Symmetrically, in the N-region, the level k_H is given by k_H^n that satisfies

$$\alpha_L - \beta_L = \alpha_H - \Pi(\beta_L, \beta_H; k_H^n) - F(k_H^n). \tag{6}$$

Finally, in the P/N-region, the least cost separation level k_H is simply the minimum of k_H^p and k_H^n defined by (5) and (6) respectively. We have the main result for the expected-margin model.

Proposition 1 *In the least cost separating equilibrium, type (α_L, β_L) candidate runs an uninformative campaign. Type (α_H, β_H) runs an informative campaign of level k_H, with $k_H = k_H^p$ in the P-region, $k_H = k_H^n$ in the N-region, and $k_H = \min\{k_H^p, k_H^n\}$ in the P/N-region.*

Given the preceding analysis, the above proposition follows immediately from Lemma 1. First, any level k lower than k_H defined above is insufficient to deter type (α_L, β_L) from imitating type (α_H, β_H) in each of the three regions in Figure 1. Second, the indifference condition of type (α_L, β_L) implies that type (α_H, β_H) strictly prefers running a campaign of level k_H to running an uninformative campaign. For instance, in the case of positive campaign, equation (5) and Lemma 1 implies that

$$\alpha_H - \beta_H - F(k_H^p) > \alpha_L - \beta_L.$$

To rule out all other deviations, we can specify the out-of-equilibrium belief as follows. In the P-region, for any deviation to a positive campaign of some level k^p, the interim belief is (α_L, β_L) if $k^p < k_H^p$, and (α_H, β_H) if $k^p > k_H$; for any deviation to a negative campaign of any level, the out-of-equilibrium belief is (α_L, β_L). In the N-region, the out-of-equilibrium belief can be symmetrically specified. In the P/N-region, the out-of-equilibrium belief is the same as that in the P-region if the deviation is in a positive campaign, and the same as that in the N-region if the deviation is in a negative campaign.

As in standard signaling games, there is some degree of freedom in specifying the out-of-equilibrium beliefs. However, Lemma 1 implies that the beliefs specified above are the only ones satisfying the Intuitive Criterion of Cho and Kreps (1987). Further, the same refinement rules out other separation equilibria with type (α_H, β_H) running a higher level of campaign k than k_H given in the proposition.[3]

Comparative statics of the least cost separation level k_H in the P and N regions of Figure 1 is straightforward from equations (5) and (6), which measures the strength of type (α_L, β_L)'s misinformation incentives. If type (α_H, β_H) candidate a becomes a stronger candidate himself, he must run a more expensive campaign to deter type (α_L, β_L): both k_H^p and k_H^n increase in α_H. Conversely, as β_H increases, both k_H^p and k_H^n decrease. As α_L increases, k_H^n decreases because there is less incentive for type (α_L, β_L) to imitate type (α_H, β_H) with a negative campaign. Since $\Pi(\alpha_L, \alpha_H; k_H^p)$ increases in α_L at a rate less than 1, an increase in α_L will still reduce the gain from imitation with a positive campaign for type (α_L, β_L), leading to a decrease in the least cost separation level k_H^p as well. Symmetrically, as β_L increases, both k_H^p and k_H^n increase. Further, if α_L and β_L change at the same rate such that $\alpha_L - \beta_L$ stays constant, then the least cost separation level k_H^p in a positive campaign increases and that in a negative campaign k_H^n decreases.

Comparative statics of the least cost separation equilibrium for type (α_H, β_H) in the P/N-region of Figure 1 is more involved because we need to consider the possibilities of both kinds of campaigns. Fix any type (α_H, β_H) in the P/N-region and rewrite (5) and (6) as

$$\alpha_H - \beta_H - (\alpha_L - \beta_L) = \alpha_H - \Pi(\alpha_L, \alpha_H; k_H^p) + F(k_H^p) = \Pi(\beta_L, \beta_H; k_H^n) - \beta_H + F(k_H^n). \quad (7)$$

Note that by Lemma 1, the second and the third expressions above are increasing in k_H^p and k_H^n respectively. Using the definition of $\Delta(\tilde{\alpha}; k)$, we have

$$\alpha_H - \Pi(\alpha_L, \alpha_H; k_H^p) = (\alpha_H - \alpha_L)(2k_H^p - 1)\Delta(\alpha_H; k_H^p),$$

and

$$\Pi(\beta_L, \beta_H; k_H^n) - \beta_H = (\beta_L - \beta_H)(2k_H^n - 1)\Delta(\beta_H; k_H^n).$$

[3] To see this, note that in any such equilibrium, there exists a level $\tilde{k}$ strictly between k_H and k such that type (α_L, β_L) strictly prefers k_L to $\tilde{k}$ even if the interim belief after the deviation to $\tilde{k}$ is (α_H, β_H), while type (α_H, β_H) strictly prefers $\tilde{k}$ to k under the same belief. Thus, the equilibrium fails the Intuitive Criterion.

Simple algebra shows that $\Delta(\alpha_H; k)$ increases in α_H if $\alpha_H < \frac{1}{2}$ and decreases in α_H if $\alpha_H > \frac{1}{2}$. This is because the spread in the posterior beliefs after $\overline{s}$ versus after $\underline{s}$ is greater if the interim belief is closer to $\frac{1}{2}$, when the posterior belief is most responsive to the realized campaign signals. Consider another type (α'_H, β'_H) such that $\alpha'_H - \beta'_H = \alpha_H - \beta_H$, with $\alpha_H < \alpha'_H < \frac{1}{2}$ and $\beta'_H > \beta_H > \frac{1}{2}$. Let ${k^p_H}'$ and ${k^n_H}'$ be the corresponding least cost separating positive and negative campaign levels respectively. Then, it follows from (7) that ${k^p_H}' < k^p_H$ and ${k^n_H}' > k^n_H$. We have the following corollary.

Corollary 1 *Fix any type (α_L, β_L), and consider the least cost separating equilibrium for any type (α_H, β_H) satisfying $\alpha_H < \frac{1}{2} < \beta_H$. (i) If some type (α_H, β_H) runs a positive campaign, then any type (α'_H, β'_H) with $\alpha'_H - \beta'_H = \alpha_H - \beta_H$ and $\alpha'_H > \alpha_H$ also runs a positive campaign at a strictly lower level; and (ii) if (α_H, β_H) runs a negative campaign, then any (α'_H, β'_H) with $\alpha'_H - \beta'_H = \alpha_H - \beta_H$ and $\alpha'_H < \alpha_H$ also runs a negative campaign at a strictly lower level.*

Suppose that $\alpha_L < \frac{1}{2} < \beta_L$. Corollary 1 implies that there exists a boundary in the P/N-region of Figure 1, such that any type (α_H, β_H) on the boundary is indifferent between running a positive campaign and a negative campaign of the same level in the least cost separating equilibrium. This boundary, connecting (α_L, β_L) to $\left(\frac{1}{2}, \frac{1}{2}\right)$, is defined by equations (7) after setting $k^p_H = k^n_H = k_H$ and eliminating k_H from the two equations. To understand the nature of the boundary further, imagine a 45-degree line that goes through such an indifferent type (α_H, β_H), representing all types (α'_H, β'_H) with $\alpha'_H - \beta'_H = \alpha_H - \beta_H$. By Corollary 1, the indifferent type (α_H, β_H) is the unique intersection of this 45-degree line with the boundary, with each type (α'_H, β'_H) above the intersection running a positive campaign of a lower level than type (α_H, β_H) in the corresponding least cost separating equilibrium, and each type below the intersection running a negative campaign of a lower level than type (α_H, β_H). Further, moving down the boundary from type (α_L, β_L) to $\left(\frac{1}{2}, \frac{1}{2}\right)$, it is easy to see from the comparative statics in the P and N regions that the equilibrium level of least cost separation increases.[4]

Corollary 1 is also useful in refining the comparative statics analysis of the least cost separating equilibrium level when both α_H and β_H change for fixed type (α_L, β_L). Consider any two types (α_H, β_H) and (α'_H, β'_H) who choose to run a positive campaign of levels k^p_H and ${k^p_H}'$ respectively. From the

[4]Note that Corollary 1 does not imply that the boundary is monotonically decreasing.

comparative statics in the P-region, we already know that ${k^p_H}' \geq k^p_H$ if $\alpha'_H \geq \alpha_H$ and $\beta'_H \leq \beta_H$. Under the same additional assumption as in the corollary, equations (7) imply that ${k^p_H}' \geq k^p_H$ if $\alpha'_H - \beta'_H \geq \alpha_H - \beta_H$ and $\alpha'_H \leq \alpha_H$. Thus, ${k^p_H}' \geq k^p_H$ so long as $\alpha'_H - \beta'_H \geq \alpha_H - \beta_H$ and $\beta'_H \leq \beta_H$.[5] Symmetrically, conditional on types (α_H, β_H) and (α'_H, β'_H) both running negative campaigns, the corresponding separating levels k^n_H and ${k^n_H}'$ satisfy ${k^n_H}' \geq k^n_H$ so long as $\alpha'_H - \beta'_H \geq \alpha_H - \beta_H$ and $\alpha'_H \geq \alpha_H$. With the above comparative statics results, we can develop one unambiguous measure of the advantage of type (α_H, β_H) in running positive campaigns relative to negative campaigns to separate from type (α_L, β_L). For all (α_H, β_H) such that $\alpha_H - \beta_H$ is constant, a smaller α_H (and correspondingly a smaller β_H) means a smaller relative advantage in positive campaigns (and correspondingly a greater relative advantage in negative campaigns). This results in a higher level of campaign for type (α_H, β_H) when the least cost separating campaign is positive, and a lower level when it is negative.

Finally, since the boundary depends on the position of type (α_L, β_L), if α_L and β_L change at the same rate such that $\alpha_L - \beta_L$ remains constant, then as they both increase, the comparative statics result above suggests that any type (α_H, β_H) on the old boundary needs to run a higher level positive campaign or a lower level negative campaign than before. Since he always chooses the least costly campaign, type (α_H, β_H) runs a negative campaign. This also suggests if (α_L, β_L) moves up (down) along the 45-degree line in Figure 1, the new boundary shifts up (down).

3.2 Discussion

3.2.1 Pooling equilibrium

The same additional assumptions on the types as in Corollary 1 turn out to be sufficient to rule out pooling equilibrium in the P/N-region under the equilibrium refinement of D1 of Banks and Sobel (1987). Consider a pooling equilibrium in a positive campaign of level $\hat{k}^p$. Denote the equilibrium interim belief as (α_m, β_m), which is the weighted average of (α_L, β_L) and (α_H, β_H) using the common prior. We have $\alpha_m < \frac{1}{2}$ under the assumption that $\alpha_L, \alpha_H < \frac{1}{2}$. Further, since in any equilibrium type (α_L, β_L) gets at least the complete information payoff, we have $\hat{k}^p < k^p_H$. We argue that for any deviation $k^p \in (\hat{k}^p, k^p_H)$

[5] On the other hand, ${k^p_H}' < k^p_H$ if $\alpha'_H - \beta'_H = \alpha_H - \beta_H$ and $\alpha'_H > \alpha_H$. By continuity, we may have ${k_H}' < k_H$ even if $\alpha'_H - \beta'_H > \alpha_H - \beta_H$.

in a positive campaign, if the interim belief (α, β) is such that type (α_L, β_L) weakly benefits from the deviation, then under the same interim belief type (α_H, β_H) strictly benefits. To see this, note since $k^p > \hat{k}^p$, for type (α_L, β_L) to benefit from such deviation, we need $\alpha > \alpha_m$. Further,

$$\Pi(\alpha_L, \alpha; k^p) - \beta - F(k^p) \geq \Pi(\alpha_L, \alpha_m; \hat{k}^p) - \beta_m - F(\hat{k}^p).$$

For type (α_H, β_H) to strictly benefit from the deviation, we need

$$\Pi(\alpha_H, \alpha; k^p) - \beta - F(k^p) > \Pi(\alpha_H, \alpha_m; \hat{k}^p) - \beta_m - F(\hat{k}^p).$$

The desired inequality is

$$\Pi(\alpha_H, \alpha; k^p) - \Pi(\alpha_H, \alpha_m; \hat{k}^p) > \Pi(\alpha_L, \alpha; k^p) - \Pi(\alpha_L, \alpha_m; \hat{k}^p).$$

Using the definition of Δ (equation 3) and rearranging the terms, we find that the above is equivalent to

$$(2k^p - 1)\Delta(\alpha; k^p) > (2\hat{k}^p - 1)\Delta(\alpha_m; \hat{k}^p),$$

which holds because $\Delta(\alpha; k)$ increases in k and in α for any $\alpha < \frac{1}{2}$. Thus, applying the refinement D1, we set the out-of-equilibrium belief to (α_H, β_H) after the deviation to k^p. Clearly, the pooling equilibrium fails D1 because type (α_H, β_H) benefits by deviating to a positive campaign of level k^p just above $\hat{k}^p$.

3.2.2 Banning negative campaigns

Given the prevalent and increasing use of negative campaigns, there have been many proposals of imposing bans on negative campaigns. The perpetual discussion about the impact of negative or "attack" ads on politics gained force during and after the 2004 US Presidential Election, which saw a surprising amount of venom on both sides. This followed some similarly negative campaigns for the 2002 Senate elections. Some politicians also tried to reach an agreement with their rival to ban negative campaigns.[6] In the current model, banning of negative campaigns may lead to pooling of types (α_L, β_L) and (α_H, β_H). We already know that if type (α_H, β_H) is in the N-region, it cannot be separated from type (α_L, β_L) through a positive campaign. If negative campaigns are banned, then there is a pooling equilibrium in which both

[6] "Van Hollen Urges Ban On Negative Advertising," *Washington Post*, October 1, 2002.

types run an uninformative campaign. Of course, if type (α_H, β_H) is in the P-region or in the P/N-region above the boundary implied by Corollary 1, banning negative campaigns does not affect the least cost separating equilibrium. Finally, if type (α_H, β_H) is in the P/N-region below the boundary, then banning negative campaigns will result in a higher level of positive campaign used by type (α_H, β_H) to separate from type (α_L, β_L).

3.2.3 Double campaigns

So far candidate a can run only a single campaign, we now turn to the case of "double-campaigning" to see whether type (α_H, β_H) can reduce the cost of separation by running both a positive campaign and a negative campaign. We claim that "double campaigning" cannot reduce the cost of separation if type (α_H, β_H) is in either P or N-region of Figure 1. Suppose that type (α_H, β_H), with $\alpha_H > \alpha_L$ and $\beta_H \geq \beta_L$, runs a positive campaign of level k^p and a negative campaign of level k^n. For the two campaigns to be separating, we need

$$\alpha_L - \beta_L \geq \Pi(\alpha_L, \alpha_H; k^p) - \Pi(\beta_L, \beta_H; k^n) - F(k^p) - F(k^n). \tag{8}$$

From equation (5), the total campaign cost in running both k^p and k^n is smaller than the cost of running just k^p_H only if $k^p < k^p_H$ and

$$\Pi(\alpha_L, \alpha_H; k^p) - \Pi(\beta_L, \beta_H; k^n) < \pi(\alpha_L, \alpha_H; k^p_H) - \beta_H.$$

The above is impossible by Lemma 1 because $\alpha_L < \alpha_H$ and $\beta_H \geq \beta_L$. Recall from Proposition 1 that in the P-region, type (α_H, β_H) cannot separate from type (α_L, β_L) by running a negative campaign because type (α_L, β_L) faces a less qualified candidate b. Thus if type (α_H, β_H) must run an informative negative campaign, he has to run a higher level positive campaign to separate, which only increases the total cost of campaigning. The argument for why double campaigning cannot reduce the cost of separation in the N-region is symmetric.

In the P/N-region, for type (α_H, β_H) to separate from type (α_L, β_L) with two campaigns k^p and k^n, the necessary total cost of campaigns $F(k^p) + F(k^n)$ is such that condition (8) is satisfied with equality. By definition, condition (8) is satisfied with equality by $k^p = \frac{1}{2}$ and $k^n = k^n_H$, and by $k^p = k^p_H$ and $k^n = \frac{1}{2}$.

Consider increasing k^p and simultaneously decreasing k^n, starting from $k^p = \frac{1}{2}$ and $k^n = k^n_H$, such that (8) is satisfied with equality. Then, the change in the required total cost of campaigns $F(k^p) + F(k^n)$, given by $F'(k^p)dk^p + F'(k^n)dk^n$, has the same sign as

$$\frac{F'(k^p)}{F'(k^n)} + \frac{\partial\Pi(\alpha_L, \alpha_H; k^p)/\partial k^p}{\partial\Pi(\beta_L, \beta_H; k^n)/\partial k^n}. \tag{9}$$

Note that at $k^p = \frac{1}{2}$ and $k^n = k^n_H$, we have $\partial\Pi(\alpha_L, \alpha_H; k^p)/\partial k^p = 0$ and thus expression (9) is positive. We claim that when F is concave, expression (9) can change sign at most once. To see this, note that when F is concave, the first ratio in (9) is decreasing as k^p increases and k^n decreases. For the second ratio, by taking derivatives we can verify that $\Delta(\tilde{\alpha}; k)$ is convex in k for all $\tilde{\alpha}$: a greater k increases the spread in the posterior beliefs at a greater rate regardless of the interim belief. Using equation (3), we then find that $\partial\Pi(\alpha_L, \alpha_H; k^p)/\partial k^p$ is negative and decreasing as k^p increases because $\alpha_L < \alpha_H$, while $\partial\Pi(\beta_L, \beta_H; k^n)/\partial k^n$ is positive and decreasing as k^n decreases.

As a result, the second ratio in the expression (9) is negative and decreasing as k^p increases and k^n decreases. It follows that (9) can change sign only from positive to negative. This implies that if at $k^n = \frac{1}{2}, k^p = k^p_H$, the total cost is above $k^p = \frac{1}{2}, k^n = k^n_H$, type (α_H, β_H) cannot reduce cost by running two campaigns; and otherwise, $k^n = \frac{1}{2}, k^p = k^p_H$ is the total cost minimizing campaigns. In either case, double campaigning cannot reduce the total cost of separation for type (α_H, β_H).

Clearly, the restriction of one informative campaign does not affect type (α_H, β_H) candidate a in either the P-region or the N-region. This result also suggests that double campaigns, if allowed, cannot hurt the strong type (α_H, β_H) candidate a in the P/N-region. Intuitively, type (α_H, β_H) can better convince the voter of his high qualifications and of candidate b's low qualifications than type (α_L, β_L) candidate a. Therefore if the campaign cost is sufficiently convex, it may reduce cost to run two informative campaigns at lower levels to separate than to run one highly informative campaign.

3.2.4 Multiple types

The additional assumptions about types in Corollary 1 are important because it gives rise to a natural single crossing property which allows us to generalize the construction of the least cost separating equilibrium to more than two types. Let there be $n > 2$ types, denoted as (α_i, β_i), $i = 1, \ldots, n$. Assume that

$\alpha_{i-1} \leq \alpha_i < \frac{1}{2}$ and $\beta_{i-1} \geq \beta_i$ with at least one strict inequality for all $i = 2, \ldots, n$. We focus on the analysis of positive campaigns; the case of negative campaigns is similar.

The least cost separating equilibrium levels of positive campaigns, k_i^p are defined iteratively by the indifference condition of type $(\alpha_{i-1}, \beta_{i-1})$ between its own equilibrium campaign of level k_{i-1}^p and type (α_i, β_i)'s level k_i^p, starting with $k_1^p = \frac{1}{2}$:

$$\alpha_{i-1} - \beta_{i-1} - F(k_{i-1}^p) = \Pi(\alpha_{i-1}, \alpha_i; k_i^p) - \beta_i - F(k_i^p). \tag{10}$$

Since $\alpha_{i-1} \leq \alpha_i$ and $\beta_{i-1} \geq \beta_i$ with at least one strict inequality, the above condition implies that $k_{i-1}^p < k_i^p$.[7] Consider first "upward" deviations; downward deviations can be symmetrically analyzed. We claim that type (α_i, β_i) strictly prefers k_i^p to any k_j^p with $j \geq i+2$. To see this, note that type $(\alpha_{j-1}, \beta_{j-1})$ is indifferent between k_{j-1}^p and k_j^p. From equation (10) for type $(\alpha_{j-1}, \beta_{j-1})$, we have that type (α_i, β_i) strictly prefers k_{j-1}^p to k_j^p if

$$\Pi(\alpha_{j-1}, \alpha_{j-1}; k_{j-1}^p) - \Pi(\alpha_i, \alpha_{j-1}; k_{j-1}^p) < \Pi(\alpha_{j-1}, \alpha_j; k_j^p) - \Pi(\alpha_i, \alpha_j; k_j^p),$$

which is true because $\alpha_{j-1} \leq \alpha_j < \frac{1}{2}$ and $k_{j-1}^p < k_j^p$. Since type (α_i, β_i) is indifferent between k_i^p and k_{i+1}^p, an iteration of the above argument establishes that type (α_i, β_i) strictly prefers k_i^p to k_j^p.

4 Winner Takes All

In this section we consider the winner-take-all model. Unlike the expected-margin model, under complete information, each type of candidate may want to run informative campaigns to increase the chance of winning. To simplify, we assume throughout this section that

$$\alpha_L, \alpha_H \leq \frac{1}{2};\ \beta_L, \beta_H \geq \frac{1}{2}, \tag{11}$$

so that under complete information both types of candidate a have to engage in sufficiently informative campaigns to have a chance of winning, which occurs if the voter observes a favorable realized campaign signal. We also assume that the campaign cost $F(k)$ is small for relevant levels, so that both types can

[7] We assume that $F(1)$ is sufficiently great, or alternatively, the differences $\beta_i - \beta_{i-1}$ are sufficiently small, so that all levels k_i^p are well defined.

afford the necessary campaigns. Then, under complete information, for type (α_L, β_L) candidate a to have a chance of winning, the minimum level k_L^c of a positive campaign satisfies

$$\frac{\alpha_L k_L^c}{\alpha_L k_L^c + (1-\alpha_L)(1-k_L^c)} = \beta_L, \tag{12}$$

so that candidate a wins if and only if the realized campaign signal is $\overline{s}$. For any higher positive campaign level $k > k_L^c$, the payoff to type (α_L, β_L) is

$$\alpha_L k + (1-\alpha_L)(1-k) - F(k), \tag{13}$$

which is decreasing in k because $\alpha_L < \frac{1}{2}$, and linear in the private belief α_L, as in the expected-margin model. It is easy to verify that the level k_L^c is the same level required of a negative campaign for type (α_L, β_L) to win (when the realized campaign signal about candidate b's qualifications is $\underline{s}$).

Since $\alpha_L < \frac{1}{2} < \beta_L$, type (α_L, β_L)'s optimal campaign level under complete information is k_L^c, and the optimal target is the candidate a himself if and only if

$$\alpha_L k_L^c + (1-\alpha_L)(1-k_L^c) > (1-\beta_L)k_L^c + \beta_L(1-k_L^c),$$

or $\alpha_L + \beta_L > 1$, i.e., if and only if α_L is closer to $\frac{1}{2}$ than β_L is. In this case, we say that type (α_L, β_L) has a preference for positive campaigns over negative campaigns under complete information. If $\alpha_L + \beta_L < 1$, type (α_L, β_L) has a relative advantage in negative campaigns. The same analysis applies to type (α_H, β_H). Intuitively, the voter's posterior belief only matters after one realized campaign signal: $\overline{s}$ in a positive campaign or $\underline{s}$ in a negative campaign; and the relevant posterior belief is more responsive to the realized campaign signal if her prior belief, which is also the interim belief here, is closer to $\frac{1}{2}$.

For any type α, β, his complete information optimal campaign is positive if $\alpha + \beta > 1$; negative if $\alpha + \beta < 1$; either if $\alpha + \beta = 1$. Under assumption (11), in either positive or negative campaigns, the marginal value of information for each type (α, β) is zero for any information campaign at a level below k^c defined by (12), and is negative at a level above it. Comparative statics of the complete information campaign level is straightforward. For each type (α, β), the level k^c, whether in a positive or negative campaign, is decreasing in α and increasing in β.

4.1 Equilibrium Campaign Choices

The equilibrium analysis for the winner-take-all model differs from that in the expected-margin model in an important aspect. In the expected-margin model, the interim belief affects a's payoff continuously by changing the posterior beliefs for the realized campaign signals (equation (1)). In contrast, in the present model, the outcome (who wins) and hence the payoff to candidate a depends on the interim belief discontinuously: the interim belief matters only when it makes the outcome depend on the realized campaign signal. Moreover, the interim belief enters the different types of candidates' payoff calculation in the same way: if an interim belief is such that some level of a positive campaign is sufficient for one type of candidate to win when the realized signal is $\overline{s}$, then the same interim belief is also sufficient for the other type to win.

Consequently, unlike in the expected-margin model, here it is impossible to have separation through different levels of the same kind of campaign. In any equilibrium of the present model, an informative and thus costly positive campaign run by either type must lead to an equilibrium interim belief such that candidate a wins if the realized campaign signal is $\overline{s}$; and symmetrically, an informative negative campaign run in equilibrium must lead to a win by a if the realized signal is $\underline{s}$. Given this, and given the assumption of $\alpha_L, \alpha_H \leq \frac{1}{2}$ and $\beta_L, \beta_H \geq \frac{1}{2}$, if there were a separating equilibrium through different levels in either positive or negative campaign, then the type running a higher level of campaign prefers to deviate to the lower level run by the other type, which strictly increases his chance of winning at a lower cost from equation (13).

Lemma 2 *In the winner-take-all model, there is no separating equilibrium in which types* (α_H, β_H) *and* (α_L, β_L) *run the same kind of campaign.*

Let the low type be the one with a higher optimal campaign level under complete information, say type (α_L, β_L), since this is the type that may have incentives to imitate the other type by getting a more favorable interim belief through a lower campaign level. Type (α_H, β_H) has no incentive to imitate type (α_L, β_L) if they both run the same kind of campaign, in the sense that the former's complete information payoff is always strictly greater than the payoff that type (α_H, β_H) can obtain by masquerading as the low type and getting the same interim belief as type (α_L, β_L). Throughout this section, we assume that

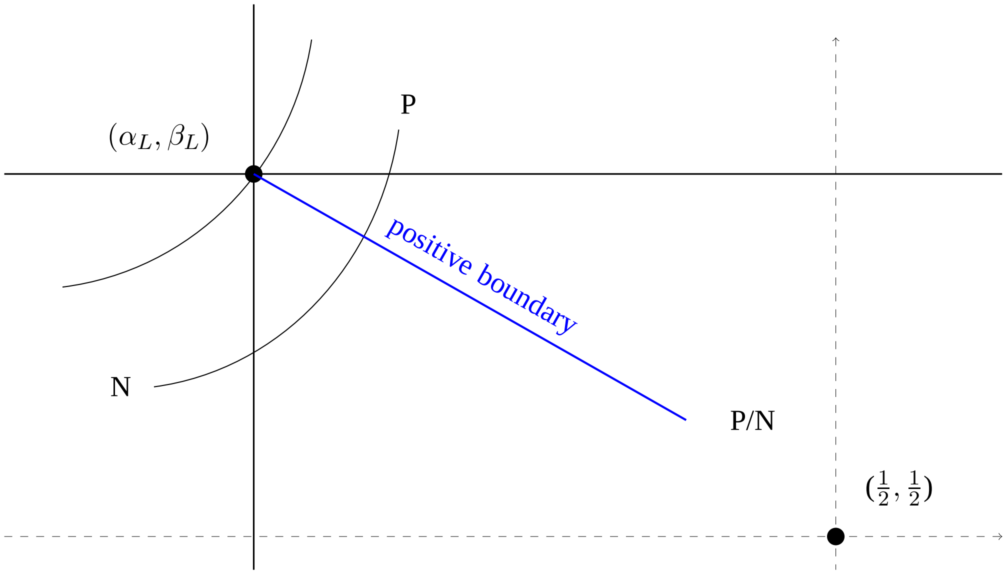

Figure 2: Winner-take-all Model

$k_L^c > k_H^c$, or

$$\frac{(1-\alpha_L)\beta_L}{\alpha_L(1-\beta_L)+(1-\alpha_L)\beta_L} > \frac{(1-\alpha_H)\beta_H}{\alpha_H(1-\beta_H)+(1-\alpha_H)\beta_H}. \tag{14}$$

As in the expected-margin model, the types (α_H, β_H) satisfying (14) are located to the right of all types (α_H, β_H) such that $k_L^c = k_H^c$. The condition $k_L^c = k_H^c$ is equivalent to[8]

$$\frac{(1-\alpha_L)\beta_L}{\alpha_L(1-\beta_L)} = \frac{(1-\alpha_H)\beta_H}{\alpha_H(1-\beta_H)}. \tag{15}$$

As in the expected-margin model, we construct different equilibria depending on the location of type (α_H, β_H) relative to (α_L, β_L). Suppose first that $\beta_H \geq \beta_L$, then $\alpha_H > \alpha_L$ by condition (14), and so type (α_H, β_H) is located in the P-region in Figure 2. Imagine that type (α_L, β_L) benefits from switching from a negative campaign of any level k^n, which leads to a win for candidate a if and only if the realized signal is $\underline{s}$, to a positive campaign of any level k^p, which leads to a win if and only if the realized signal is $\overline{s}$, i.e.

$$\beta_L(1-k^n) + (1-\beta_L)k^n - F(k^n) \leq \alpha_L k^p + (1-\alpha_L)(1-k^p) - F(k^p). \tag{16}$$

[8] By taking derivatives, we can easily show that for a fixed type (α_L, β_L), the value of β_H that satisfies the following equality increases at an increasing rate as α_H increases.

If inequality (16) holds, then type (α_H, β_H) would strictly benefit from such a switch:

$$\beta_H(1-k^n) + (1-\beta_H)k^n - F(k^n) < \alpha_H k^p + (1-\alpha_H)(1-k^p) - F(k^p). \tag{17}$$

In this case, type (α_H, β_H) has a preference for positive campaigns *relative* to type (α_L, β_L).

As a result, if type (α_L, β_L)'s complete information optimal campaign is negative, the two types can be separated by type (α_L, β_L) running a negative campaign and type (α_H, β_H) running a positive campaign. For this type of separation to be a least cost separating equilibrium, type (α_L, β_L)'s negative campaign level must be at the complete information level k_L^c, while type (α_H, β_H)'s positive campaign level must be between k_H^c and k_L^c such that type (α_L, β_L) is just indifferent. This least cost separating equilibrium level k_H^p satisfies:

$$(1-\beta_L)k_L^c + \beta_L(1-k_L^c) - F(k_L^c) = \alpha_L k_H^p + (1-\alpha_L)(1-k_H^p) - F(k_H^p). \tag{18}$$

Type (α_H, β_H) has to run a strictly higher level of positive campaign than his complete information level ($k_H^p > k_H^c$) if type (α_L, β_L)'s preference for negative campaigns is not too strong. That is, if

$$(1-\beta_L)k_L^c + \beta_L(1-k_L^c) - F(k_L^c) \leq \alpha_L k_H^c + (1-\alpha_L)(1-k_H^c) - F(k_H^c), \tag{19}$$

then type (α_L, β_L) prefers deviating to the positive campaign of k_H^c to pretending to be type (α_H, β_H).[9] If type (α_L, β_L) has a sufficiently strong preference for negative campaigns so that inequality (19) is violated, then the least cost separating equilibrium level $k_H^p = k_H^c$, so long as type (α_H, β_H) does not have a strong preference for negative campaigns, or,

$$(1-\beta_H)k_L^c + \beta_H(1-k_L^c) - F(k_L^c) \leq \alpha_H k_H^c + (1-\alpha_H)(1-k_H^c) - F(k_H^c). \tag{20}$$

Note that (20) is implied by (19) because type (α_H, β_H) has a relative preference for positive campaigns compared to type (α_L, β_L). If (20) is not satisfied, then separation between the two types is impossible. The analysis for the N-region is symmetric: the counterpart for (18) is

$$\alpha_L k_L^c + (1-\alpha_L)(1-k_L^c) - F(k_L^c) = (1-\beta_L)k_H^n + \beta_L(1-k_H^n) - F(k_H^n), \tag{21}$$

[9] This follows because the right-hand-side of (18) is decreasing in k_H^p: inequality (19) implies that the right-hand-side is greater than the left-hand-side at $k_H^p = k_H^c$, while the assumption $\alpha_L + \beta_L < 1$ implies the right-hand-side is strictly less than the left-hand-side at $k_H^p = k_L^c$.

and the counterpart for (20) is

$$\alpha_H k_L^c + (1-\alpha_H)(1-k_L^c) - F(k_L^c) \leq (1-\beta_H)k_H^c + \beta_H(1-k_H^c) - F(k_H^c). \tag{22}$$

We have the following proposition.

Proposition 2 *(i) For any (α_H, β_H) in the P-region, there is no separating equilibrium if $\alpha_L + \beta_L > 1$ or if $\alpha_L + \beta_L < 1$ but (20) is violated, otherwise in the least cost separating equilibrium, (α_L, β_L) runs a negative campaign of level k_L^c and (α_H, β_H) runs a positive campaign of level $k_H^p \in [k_H^c, k_L^c)$; and (ii) for any (α_H, β_H) in the N-region, there is no separating equilibrium if $\alpha_L + \beta_L < 1$ or if $\alpha_L + \beta_L > 1$ but (22) is violated, otherwise in the least cost separating equilibrium, (α_L, β_L) runs a positive campaign of level k_L^c and (α_H, β_H) runs a negative campaign of level $k_H^n \in [k_H^c, k_L^c)$.*

The most interesting case in the above result is the least cost separating equilibrium in which type (α_H, β_H) runs a different kind of campaign to separate from type (α_L, β_L). To establish this equilibrium, note that by conditions (18) and (19), type (α_L, β_L) weakly prefers the positive campaign of level k_H^p to the negative campaign of level k_L^c. Since (α_H, β_H) has a preference for positive campaigns relative to (α_L, β_L), the former strictly prefers k_H^p to k_L^c. To prevent deviations, we specify the out-of-equilibrium interim belief in the P-region to be (α_H, β_H) after any positive campaign of level higher than k_H^p, (α_L, β_L) after any positive campaign of level lower than k_H^p, and always (α_L, β_L) after any negative campaign. The beliefs in the N-region are similarly specified. Finally, this separating equilibrium is the only one that satisfies the Intuitive Criterion.

The above result suggests that in this model, candidate a's equilibrium choice is often driven by their relative preference for the kind of campaign rather than their complete information choices. This is the case if both types prefer the same kind of campaign under complete information, but type (α_H, β_H) is willing to run the other kind of campaign (which involves a higher cost) to separate from type (α_L, β_L). For example, in the P-region, suppose that both $\alpha_L + \beta_L < 1$ and $\alpha_H + \beta_H < 1$, but inequality (20) is satisfied. Then type (α_H, β_H) has a preference for positive campaigns relative to type (α_L, β_L) even though both types prefer negative campaigns under complete information, and we obtain the least cost separating equilibrium which type (α_H, β_H) runs a negative campaign to separate from type (α_L, β_L).

Comparative statics of the least cost separating levels k_H^p and k_H^n, defined respectively by (18) and (21), is straightforward. Recall that the payoff to candidate a depends on the interim belief only when it makes the outcome depend on the realized campaign signal. Thus, in a separating equilibrium the separation levels k_H^p and k_H^n, if they are above k_H^c, are unaffected when there are small changes to (α_H, β_H), in contrast with the comparative statics in the expected-margin model. Next, take the case of type (α_H, β_H) in the P-region. When β_L increases, type (α_L, β_L) receives a smaller complete information payoff from the negative campaign and at the same time the complete information negative campaign level k_L^c increases, which further reduces type (α_L, β_L)'s equilibrium payoff. As a result, type (α_H, β_H) must run a higher level positive campaign to separate in the least cost separating equilibrium. In contrast, as α_L increases, there are two opposing effects on type (α_L, β_L): the deviation payoff from running the positive campaign of level k_H^p (the right-hand-side of (18)) increases, but at the same time the equilibrium payoff (the left-hand-side of (18)) also increases because the complete information level k_L^c decreases. To gain insights about which effect dominates, consider what happens when α_L and β_L move in the same direction so that the complete information level k_L^c of type (α_L, β_L) is unaffected. In Figure 2, this happens when we move along the boundary through (α_L, β_L). Consider the least cost separating level k_H^p of positive campaigns. Assuming that $\alpha_L + \beta_L < 1$, then as we increase both α_L and β_L along the boundary, the left-hand-side of (18) decreases while the right-hand-side increases, implying that k_H^p must increase to restore the indifference of type (α_L, β_L). Symmetrically, starting with $\alpha_L + \beta_L > 1$, when we decrease both α_L and β_L along the boundary, the least cost separating level k_H^n of negative campaigns given by (21) also increases. Thus, as type (α_L, β_L)'s preference for either kind of campaign becomes weaker, the least cost separating level increases.

Separation is impossible if type (α_L, β_L) prefers positive campaigns under complete information and type (α_H, β_H) prefers positive campaigns relative to type (α_L, β_L). Lemma 2 showed that the two types can only be separated through different kinds of campaigns. Since type (α_L, β_L) prefers positive campaigns under complete information, in any separating equilibrium, type (α_L, β_L) must run a positive campaign. The only possibility for separation is then type (α_H, β_H) running a negative campaign and type (α_L, β_L) running a positive campaign. However, since type (α_H, β_H) has a relative preference for positive campaigns, the equilibrium conditions for the two types cannot be simultaneously satisfied.

Separation is also impossible in the remaining case of Proposition 2, where type (α_L, β_L) prefers negative campaigns under complete information and type (α_H, β_H) has a strong preference for negative campaigns, even though the latter prefers positive campaigns relative to the former. This can occur only if type (α_H, β_H) prefers negative campaigns under complete information, i.e., $\alpha_H + \beta_H < 1$, and if k_H^c is sufficiently close to k_L^c. For separation to occur in this case, type (α_H, β_H) must run a positive campaign while type (α_L, β_L) runs the negative campaign of level k_L^c. This is impossible, because in any separating equilibrium type (α_H, β_H) has to run a positive campaign of level that is at least k_H^c, but then type (α_H, β_H) would deviate to the negative campaign run by (α_L, β_L) when (20) is violated.

Proposition 2 can be extended to the P/N-region of Figure 2, where $\alpha_H > \alpha_L$ and $\beta_H < \beta_L$. Suppose for now $\alpha_L + \beta_L < 1$ so that type (α_L, β_L) prefers negative campaigns under complete information. Furthermore, suppose that (19) holds so that type (α_L, β_L) weakly prefers type (α_H, β_H)'s complete information positive campaign to its own complete information negative campaign. Then, the level k_H^p given by (18) is the unique value of k_H^c that satisfies (19) with equality, and the assumption of $\alpha_L + \beta_L < 1$ implies that $k_H^p < k_L^c$. Condition (19) then requires type (α_H, β_H) to be such that the corresponding complete information level k_H^c is below k_H^p, or equivalently, type (α_H, β_H) to fall to the right of the boundary given by $k_H^c = k_H^p$, which is a parallel shift of the boundary through (α_L, β_L) defined by (15), which can be seen in Figure 2. Since k_H^p increases as α_L and β_L both increase to keep the complete information level k_L^c constant, the new boundary given by $k_H^c = k_H^p$ moves towards the one through (α_L, β_L), as we move up along the latter. Note that the two boundaries coincide when $\alpha_L + \beta_L = 1$, i.e., when type (α_L, β_L) has no preference between the two kinds of campaigns under complete information.[10]

Now, consider first any separating equilibrium. To begin, note that type (α_H, β_H) being in the P-region is sufficient but not necessary for having a relative preference for positive campaigns. It is easy to see that (16) implies (17) if

$$(\alpha_H - \beta_L)(2k^p - 1) > (\beta_L - \beta_H)(2k^n - 1), \tag{23}$$

which is always satisfied if type (α_H, β_H) is in the P-region but this is not necessary. Since in any

[10]The case of type (α_H, β_H) located above and to the left of the boundary given by $k_H^c = k_H^p$ is less interesting; whether or not there is pooling or separation of the two types depends entirely on whether type (α_H, β_H) has a strong preference for negative campaigns, as in Proposition 2.

separating equilibrium type (α_L, β_L)'s campaign is always at level k_L^c, and the kind of campaign is the same as it prefers under complete information, we can substitute k_L^c for k^n on the right-hand-side of (23). Further, under the assumption of $\alpha_L + \beta_L < 1$ and (19) there is a unique $k_H^p \in [k_H^c, k_L^c)$ that satisfies (18), we can substitute k_H^p for k^p on the left-hand-side of (23). This suggests a "positive boundary" in the P/N-region in the case where type (α_L, β_L) prefers negative campaigns under complete information, given by

$$(\alpha_H - \alpha_L)(2k_H^p - 1) = (\beta_L - \beta_H)(2k_L^c - 1). \tag{24}$$

Above the positive boundary, type (α_H, β_H) may be said to have a preference for positive campaigns relative to type (α_L, β_L). Note that since k_H^p is independent of α_H and β_H, the positive boundary is a downward sloping line in the P/N-region. Further, since $k_H^p < k_L^c$, the positive boundary has a slope less than 1. In the opposite case where type (α_L, β_L) prefers positive campaigns under complete information, with the least cost separating level k_H^n defined by (21), the corresponding "negative boundary" is given by

$$(\beta_L - \beta_H)(2k_H^n - 1) = (\alpha_H - \alpha_L)(2k_L^c - 1), \tag{25}$$

which is below the equal preference line. When type (α_H, β_H) is located below the negative boundary, we may say that it has a preference for negative campaigns relative to type (α_L, β_L). The counterpart of (19) is

$$\alpha_L k_L^c + (1 - \alpha_L)(1 - k_L^c) - F(k_L^c) \leq (1 - \beta_L)\ k_H^c + \beta_L(1 - k_H^c) - F(k_H^c). \tag{26}$$

We have the following corollary.

Corollary 2 *(i) Suppose $\alpha_L + \beta_L < 1$ and (19) holds. There is a separating equilibrium in which type (α_L, β_L) runs the negative campaign of level k_L^c and (α_H, β_H) runs the positive campaign of level k_H^p if (α_H, β_H) is above the positive boundary, and no separating equilibrium otherwise. (ii) Suppose $\alpha_L + \beta_L > 1$ and (26) holds. There is a separating equilibrium in which type (α_L, β_L) runs the positive campaign of level k_L^c and (α_H, β_H) runs the negative campaign of level k_H^n if (α_H, β_H) is below the negative boundary, and no separating equilibrium otherwise.*

The argument to establish the separating equilibrium is the same as in Proposition 2. We only need to use the definitions of positive and negative boundaries to show that separation is impossible if type (α_L, β_L) prefers one kind of campaigns under complete information and type (α_H, β_H) has a preference for the other kind relative to type (α_L, β_L). Consider the case where $\alpha_L + \beta_L < 1$ and (α_H, β_H) is located below the positive boundary. The only way to separate is for type (α_L, β_L) to run its complete information negative campaign of level k_L^c and type (α_H, β_H) to run a positive campaign of some level $k^p \in [k_H^c, k_L^c)$. Summing up the incentive conditions for the two types, we need

$$(\alpha_H - \alpha_L)(2k^p - 1) \geq (\beta_L - \beta_H)(2k_L^c - 1).$$

Since type (α_H, β_H) is below the positive boundary, the above requires that $k^p > k_H^p$. By the definition of the positive boundary, type (α_H, β_H) prefers the negative campaign of level k_L^c to the positive campaign of level k_H^p, so a positive campaign at a higher level would make it even less attractive than the negative campaign. As a result, separation is impossible.

Corollary 2 implies that separation of the two types occurs in equilibrium when type (α_H, β_H) has a relative preference for the opposite kind of campaign that type (α_L, β_L) prefers under complete information, and otherwise there will be pooling in equilibrium. Comparative statics regarding the separating equilibrium is straightforward. As in Proposition 2, the separating level depends on type (α_L, β_L) and not on type (α_H, β_H). The boundaries for separation to occur depend on how changes in α_L and β_L affect both k_L^c and k_H^p, or both k_L^c and k_H^n. As we have learned from the comparative statics of k_H^p and k_H^n with respect to (α_L, β_L), starting with $\alpha_L + \beta_L < 1$, as we increase both α_L and β_L so that its complete information level k_L^c stays the same, the least cost separating level k_H^p increases. As a result, when type (α_L, β_L)'s preference for either campaign becomes weaker under complete information, the positive boundary becomes steeper and the negative boundary becomes flatter, and the separating region for type (α_H, β_H) becomes larger: there are more types (α_H, β_H) that can be separated from type (α_L, β_L), even as the least cost separating level becomes higher. When $\alpha_L + \beta_L$ approaches 1 and type (α_L, β_L) becomes indifferent between the two kinds campaigns under complete information, the boundaries approach each other and approach the equal preference line. All types (α_H, β_H) can be separated in the limit.[11]

[11] Of course, there also exist pooling equilibria for the same types because type (α_L, β_L) is indifferent between the two kinds of campaigns under complete information.

4.2 Discussion

4.2.1 Double campaigns

There is no benefit to running both a positive campaign and a negative campaign under complete information. This can be established by showing that at least one of the two campaigns can be eliminated without affecting the outcome. Suppose that type (α, β) runs both a positive of level k^p and a negative campaign of level k^n. Consider first the case where the posterior belief about candidate a's qualifications after $\underline{s}$ from k^p is greater than or equal to the posterior belief about b after $\underline{s}$ from k^n:

$$\frac{\alpha(1-k^p)}{\alpha(1-k^p)+(1-\alpha)k^p} \geq \frac{\beta(1-k^n)}{\beta(1-k^n)+(1-\beta)k^n}.$$

Since $\alpha < \beta$ by assumption, the above implies that $k^p < k^n$, which in turn implies that the posterior belief about a after $\overline{s}$ from k^p is smaller than the posterior belief about b after $\overline{s}$ from k^n. Therefore, the outcome is determined by the realized campaign signal of the negative campaign: candidate a wins if and only if the realized signal is $\underline{s}$ from k^n. The positive campaign can be eliminated. Next, suppose that

$$\frac{\alpha(1-k^p)}{\alpha(1-k^p)+(1-\alpha)k^p} < \frac{\beta(1-k^n)}{\beta(1-k^n)+(1-\beta)k^n} \leq \frac{\alpha k^p}{\alpha k^p+(1-\alpha)(1-k^p)}.$$

In this case, if furthermore

$$\frac{\beta k^n}{\beta k^n+(1-\beta)(1-k^n)} \leq \frac{\alpha k^p}{\alpha k^p+(1-\alpha)(1-k^p)},$$

then the outcome is determined by the realized signal of the positive campaign, and thus the negative campaign is redundant. If instead the opposite is true, then candidate a wins if and only if the realized campaign signal is $\overline{s}$ from k^p and $\underline{s}$ from k^n. This cannot be optimal, since candidate a can increase the probability of winning by reducing k^n (and thus increasing the probability of $\underline{s}$ from k^n). Finally, if

$$\frac{\beta(1-k^n)}{\beta(1-k^n)+(1-\beta)k^n} > \frac{\alpha k^p}{\alpha k^p+(1-\alpha)(1-k^p)},$$

then a can never win and both campaigns can be eliminated.

The above analysis establishes that in any separating equilibrium type (α_L, β_L) will run only one campaign. The remaining question is whether type (α_H, β_H) can reduce the cost of separation by running two campaigns simultaneously. To consider this question, suppose that $\alpha_L+\beta_L < 1$ so that in a separating equilibrium with single campaigns type (α_L, β_L) runs the complete information negative campaign of level k^c_L and type (α_H, β_H) runs the positive campaign of level k^p_H. We claim that type (α_H, β_H) cannot reduce the cost of separation by adding a negative campaign of level k^n and simultaneously appropriately reducing the positive campaign to some k^p, so long as k^n is sufficiently low that the campaign outcome remains decided by the realization of the positive campaign alone, that is, so long as

$$\frac{\beta_H k^n}{\beta_H k^n + (1-\beta_H)(1-k^n)} \leq \frac{\alpha_H k^p}{\alpha_H k^p + (1-\alpha_H)(1-k^p)}.$$

To see this, note that for separation to obtain with the two campaigns, we need

$$(1-\beta_L)k^c_L + \beta_L(1-k^c_L) - F(k^c_L) \geq \alpha_L k^p + (1-\alpha_L)(1-k^p) - F(k^p) - F(k^n).$$

The change in the total required cost required for separation, $F'(k^p)dk^p + F'(k^n)dk^n$, is then given by $-(1-2\alpha_L)dk^p$, which is positive. Similarly, type (α_H, β_H) cannot reduce the cost of separation by adding a negative campaign of level k^n and simultaneously reducing the positive campaign to some k^p, if k^n is sufficiently high that the campaign outcome is decided by the realization of the negative campaign alone, that is, if

$$\frac{\beta_H(1-k^n)}{\beta_H(1-k^n) + (1-\beta_H)k^n} \leq \frac{\alpha_H(1-k^p)}{\alpha_H(1-k^p) + (1-\alpha_H)k^p}.$$

The remaining possibility is that the levels of k^p and k^n are such that candidate a wins when the realized signal is $\overline{s}$ from the positive campaign and is $\underline{s}$ from the negative campaign, or

$$\begin{aligned}\frac{\alpha_H(1-k^p)}{\alpha_H(1-k^p) + (1-\alpha_H)k^p} &< \frac{\beta_H(1-k^n)}{\beta_H(1-k^n) + (1-\beta_H)k^n} \\ &\leq \frac{\alpha_H k^p}{\alpha_H k^p + (1-\alpha_H)(1-k^p)} < \frac{\beta_H(1-k^n)}{\beta_H(1-k^n) + (1-\beta_H)k^n}.\end{aligned}$$

In this case, the required total cost of separation $F(k^p) + F(k^n)$ satisfies[12]

$$(1-\beta_L)k^c_L + \beta_L(1-k^c_L) - F(k^c_L) = (\alpha_L k^p + (1-\alpha_L)(1-k^p))(\beta_L(1-k^n) + (1-\beta_L)k^n) - F(k^p) - F(k^n).$$

[12] The equation below is necessary for separation but not sufficient. Unlike in the expected-margin model, the indifference of type (α_L, β_L) between its equilibrium choice and the double campaigns does not imply that type (α_H, β_H) prefers the latter. However, our argument is that the double campaigns do not reduce the total cost even if they are separating under suitable conditions.

The change $F'(k^p)dk^p + F'(k^n)dk^n$ can be shown to have the same sign as

$$\frac{F'(k^n)}{F'(k^p)} - \frac{(2\beta_L - 1)(\alpha_L k^p + (1-\alpha_L)(1-k^p))}{(1-2\alpha_L)(\beta_L(1-k^n) + (1-\beta_L)k^n)}.$$

As k^p decreases and k^n increases, the second ratio in the above expression increases, and when F is concave, the first ratio decreases. As a result, the change in the total cost of separation can only change sign at most once, from positive to negative. This implies that the cost of separation is minimized when k^p and k^n are such that the posterior beliefs of the two candidates are equal, either after observing $\overline{s}$ from both campaigns, or after observing $\overline{s}$ from the positive campaign and $\underline{s}$ from the negative campaign, or after observing $\underline{s}$ from both campaigns. Double campaigning will not help reduce the cost of separation, if the cost is higher in these corner cases than the equilibrium cost with the single positive campaign.

4.2.2 Equilibrium pooling

When it is impossible to separate type (α_H, β_H) from (α_L, β_L), there is a continuum of pooling equilibria. For example, suppose that $\alpha_L + \beta_L > 1$ and (α_H, β_H) is in the P-region. Let k^c_m be the minimum campaign level for the average type (α_m, β_m) to win after a realized campaign signal $\overline{s}$ under complete information. Clearly $k^c_m \in (k^c_H, k^c_L)$. For any level $\hat{k}^p \in [k^c_m, k^c_L]$, we can construct a pooling equilibrium in which both types run a positive campaign of level $\hat{k}^p$. To support this equilibrium, we specify the out-of-equilibrium belief to be (α_L, β_L) after any deviation to a positive campaign of level below $\hat{k}^p$, (α_H, β_H) after any deviation to a positive campaign of level above $\hat{k}^p$, and (α_L, β_L) after any deviation to a negative campaign. At any such pooling equilibrium except when $\hat{k}^p = k^c_L$, with a positive probability type (α_L, β_L) candidate a wins even though he should lose if there is no private information about his type. In this case, misinformation occurs in equilibrium.

As mentioned in the argument leading towards Lemma 2, in the present winner-take-all model, the interim belief enters the calculation of each type's payoff in the same way. As a result, the set of out-of-equilibrium beliefs that would make it profitable to deviate from the pooling equilibrium is the same for the two types. Any pooling equilibrium thus survives the standard belief refinement.

Since separation is impossible within the same kind of campaign, pooling equilibria may also occur if negative campaigns are banned. Again, there is a continuum of pooling equilibria, at any level below both complete information optimal levels that is sufficient for the average type (α_m, β_m) to win with $\overline{s}$ in positive campaigns, and any such equilibrium survives the standard belief refinement.

Finally, since separation is impossible within the same kind of campaign and since there are only two targets in the model, when there are more than two types, only the lowest type can be possibly separated from the rest. In this sense, pooling and equilibrium misinformation are robust features of the winner-take-all model.

5 Concluding Remarks

In our model of information campaigns, both the kind and the level of a campaign reveal information about a candidate's private information about herself and her rival. The candidate may have an incentive to misinform the voter by affecting the interim beliefs that the voter uses to evaluate the public campaign signals. A crucial element of the model is that the candidate does not control the realized campaign signal by making the campaign choices. That is, our model features a signaling game in which the signal is an information structure. This allows us to distinguish interim beliefs from ex post beliefs, in order to study the value of misinformation.

A related idea in principal-agent models is that the private type of the agent is a signal structure. This appears in the price discrimination model of Courty and Li (2000), in which a consumer knows only the distribution of his valuation for the good and has to report the private realized valuation later in a sequential mechanism by the seller. Perhaps more related, in the optimal auction model of Bergemann and Pesendorfer (2007), a bidder's private type is a signal structure that is optimally designed together with the auction. See also related models in an industrial organization context, for example in monopoly pricing by Ottaviani and Moscarini (2001) and price competition by Damiano and Li (2007). Unlike these models, in our paper the realized campaign signal is publicly observed. Indeed, all we need in our model is that the campaign signal is verifiable; the standard unraveling argument implies that the voter can infer the signal in equilibrium given the candidate's disclosure policy. How to generalize this to richer type and signal space is a topic for further research.

References

Arceneaux, K. (2006): “Do Campaigns Help Voters Learn? A Cross-National Analysis,” *British Journal of Political Science*, 36, 159–173.

Banks, J. S., and J. Sobel (1987): “Equilibrium Selection in Signaling Games,” *Econometrica*, 55(3), 647–61.

Bergemann, D., and M. Pesendorfer (2007): “Information Structures in Optimal Auctions,” *Journal of Eonomic Theory*, 137(1), 580–609.

Cho, I.-K., and D. Kreps (1987): “Signaling Games and Stable Equilibria,” *Quarterly Journal of Economics*, 102(2), 179–222.

Coate, S. (2004): “Pareto-Improving Campaign Finance Policy,” *American Economic Review*, 94(3), 628–655.

Courty, P., and H. Li (2000): “Sequential Screening,” *Review of Economic Studies*, 67, 697–717.

Damiano, E., and H. Li (2007): “Information Provision and Price Competition,” mimeo, University of Toronto.

Damiano, E., H. Li, and W. Suen (2008): “Credible Ratings,” *Theoretical Economics*, 3, 325–365.

Engers, M. (1987): “Signalling with Many Signals,” *Econometrica*, 55, 663–674.

Kotowitz, Y., and F. Mathewson (1979): “Advertising, Consumer Information, and Product Quality,” *Bell Journal of Economics*, 10(2), 566–588.

Lau, Richard R, L. S., and I. B. Rovner (2007): “The Effects of Negative Political Campaigns: A Meta-Analytic Reassessment,” *Journal of Politics*, 69, 1176–1209.

Milgrom, P., and J. Roberts (1986): “Price and Advertising Signals of Product Quality,” *Journal of Political Economy*, 94(4), 796–821.

Nelson, P. (1974): “Advertising as Information,” *Journal of Political Economy*, 82(4), 729–754.

Ottaviani, M., and G. Moscarini (2001): "Price Competition for an Informed Buyer," *Journal of Economic Theory*, 101(2), 457–493.

Polborn, M. K., and D. T. Yi (2006): "Informative Positive and Negative Campaigning," *Quarterly Journal of Political Science*, 1(4), 351–371.

Prat, A. (2002): "Campaign Advertising and Voter Welfare," *Review of Economic Studies*, 69(4), 999–1017.

Quinzii, M., and J. C. Rochet (1985): "Multidimensional Signalling," *Journal of Mathematical Economics*, 14, 261–284.

Sides, John, K. L., and M. Grossman (forthcoming): "Do Voters Perceive Negative Campaigns as Informative Campaigns?," *American Politics Research*.