Games and Economic Behavior 78 (2013) 44–63

Contents lists available at SciVerse ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

How to win a large election ☆

Michael Mandler

Department of Economics, Royal Holloway College, University of London, Egham, Surrey TW20 0EX, United Kingdom

ARTICLE INFO

Article history:
Received 17 December 2010
Available online 13 November 2012

JEL classification:
D72
D81

Keywords:
Elections
Expected margin of victory
Law of large numbers
Local limit theorem

ABSTRACT

We consider the optimization problem of a campaign trying to win an election when facing aggregate uncertainty, where agents' voting probabilities are uncertain. Even a small amount of uncertainty will in a large electorate eliminate many of counterintuitive results that arise when voting probabilities are known. In particular, a campaign that can affect the voting probabilities of a fraction of the electorate should maximize the expected difference between its candidate's and the opposing candidate's share of the fraction's potential vote. When a campaign can target only finitely many voters, maximization of the same objective function remains optimal if a convergence condition is satisfied. When voting probabilities are certain, this convergence condition obtains only at knife-edge combinations of parameters, but when voting probabilities are uncertain the condition is necessarily satisfied.

1. Introduction

How should a political campaign maximize its chances of winning an election? If agents have the option of staying home or abstaining, campaigns face hard decisions. A campaign may have to choose between aggressive advertisements that make its supporters more likely to show up at the polls and more sober ads that gently try to sway the undecided. A campaign's uncertainty about agents' voting probabilities will be the centerpiece of our attack on this problem. In classical approaches, voters' actions – voting for one of the candidates or abstaining – are independent draws from known probability distributions and a campaign's policy decision is to choose a profile of distributions from some set of available profiles. The classical starting point leads to conceptual trouble however: a campaign's ranking of policies can be counterintuitive.

Example 1. The simplest model with no aggregate uncertainty has n i.i.d. voters, each with the probabilities of voting left L, voting right R, and abstaining A. One might guess that the likelihood of victory for the right would increase in $R - L$ (the expected difference between the right and left's share of the potential vote) or perhaps $\frac{R}{L+R}$ (the right's expected share of the actual vote). But suppose (L, A, R) can take one of three values, $(.67, 0, .33)$, $(.5, .3, .2)$, $(.3, .6, .1)$, a sequence that increases in $R - L$ and decreases in $\frac{R}{L+R}$. It turns out that, for all n sufficiently large, the probability of victory for the right is highest with $(.3, .6, .1)$ followed by $(.67, 0, .33)$ and then $(.5, .3, .2)$. So it is not optimal for the right always to choose the bigger $R - L$ or the bigger $\frac{R}{L+R}$. Whether or not the above ordering of (L, A, R)'s by their probability of victory seems counterintuitive, the example shows that even the simplest independent-action models do not deliver transparent

☆ This paper was originally motivated by discussions with political campaigns regarding how different advertisements should be evaluated, on the assumption that polling data could measure changes in voting likelihoods. It became clear that any practical evaluation would not be able to draw on detailed information about the electorate as a whole, a principle that I have tried to follow in this paper. Let me thank Sophie Bade, Roger Myerson, Phil Steitz, an Associate Editor, and two referees for helpful suggestions and advice. This paper has also circulated under the title 'Large Bayesian elections'.

E-mail address: m.mandler@rhul.ac.uk.

0899-8256/$ – see front matter
http://dx.doi.org/10.1016/j.geb.2012.09.005

recommendations. Patty (2002) offers a related example, credited to Duggan. In Section 4.2 we will show how to order i.i.d. electorates by their large-population probability of victory.

We will show in models like Example 1 that rankings where the right's probability of victory varies inversely with $R - L$ are always fragile: we can move to a nearby model with a little uncertainty about (L, A, R) that will overturn the ranking. So unless one has information of absolute precision about how policies affect (L, A, R), it will be hopeless in an i.i.d. model – and indeed in any model of independently distributed voter actions – to offer reliable advice to a campaign. A model with uncertain voting probabilities in contrast is robust and generates more plausible rankings. A second difference between the approaches concerns the information needed to make a decision that targets only a fraction of the electorate. When voting probabilities are known, optimal decision-making will turn on detailed information about the entire electorate but when voting probabilities are uncertain a policy choice that affects just one group can be decided solely by that group's characteristics.

In our framework, there are the two candidates and a typical voter i has voting probabilities $\pi_i = (L_i, A_i, R_i) \geqslant 0$ where $L_i + A_i + R_i = 1$. Let π denote the sequence $(\pi_1, \pi_2, \ldots)$. When π is known and the realizations of agent actions are independent (as in Example 1), the model exhibits *aggregate certainty*: the proportion of votes garnered by a candidate will under mild restrictions converge to a specific quantity with probability 1. We allow for *aggregate uncertainty* by letting π be uncertain. Uncertainty about π allows a plausible weakening of independence. Campaigns rely on polling data to estimate voting probabilities and these data are too coarse to pin down voting probabilities with precision. This imprecision makes violations of full independence inevitable but our framework lets the aggregate uncertainty take a mild form, where for example all voters or all voters of a specific type share common voting probabilities. To keep the model tractable we assume that, conditional on π, the realizations of agent actions remain independent. That voters are described by probabilities does not mean that any agent employs a randomization device; all of the uncertainty could be a reflection of a campaign's ignorance.

A campaign's policy decisions – its choice of platform, advertisements, etc. – matter only insofar as they affect the likelihood of various π's. So we view a campaign policy as a probability measure P that gives the likelihood of sets of π's and the resulting likelihoods of voter actions. For concreteness, we always take the right campaign's point of view.

Uncertainty about π may seem to add complications but in fact it makes campaign decisions simpler. Fix some particular sequence π for the moment and let $\rho(\pi)$ denote the right's *expected margin of victory* (EMV), the expected difference in a large electorate between the right and left's share of the pool of potential votes:

$$\rho(\pi) = \lim_{n\to\infty} \frac{1}{n}\sum_{i=1}^{n}(R_i - L_i),$$

where $\pi_i = (L_i, A_i, R_i)$ is the ith coordinate of π. The law of large numbers (LLN) implies that for a given π the difference between the right and the left's share of the vote out of a population of size n will (with probability 1) converge to $\rho(\pi)$ as n increases. So in a large population the right wins with near certainty if $\rho(\pi) > 0$ and the right loses with near certainty if $\rho(\pi) < 0$. Taking into account that π is uncertain, we conclude that a policy that increases the probability that $\rho(\pi) > 0$ will raise the right's probability of victory when n is sufficiently large (Theorem 1). While this result does not validate full-scale EMV maximization, it does lead to some EMV-maximizing rules of thumb for subpopulations of voters. A policy that targets some nonzero fraction of agents and raises that subset's expected margin[1] will under mild conditions increase the probability that $\rho(\pi)$ is positive; hence the probability of victory will increase (Theorem 2). This conclusion allows a campaign to ignore the voters it cannot influence, but curiously the irrelevance of unaffected voters depends on their voting probabilities being uncertain.

When π and hence ρ is certain, counterintuitive cases like Example 1 can easily arise. If ρ is certain and has the same sign before and after a policy change, then $P(\rho > 0)$ will be unaffected and hence the LLN classifies the election as virtually sure to be won (if $\rho > 0$) or virtually sure to be lost (if $\rho < 0$) whether or not the policy change occurs. The effect of the policy change then depends on second-order factors that are necessarily delicate. We will see how to make such policy decisions (Theorem 3) but the important message is the fragility of these rankings: a little uncertainty about π can convert a case where ρ is known and a policy decision leaves the sign of ρ unchanged to a case where the policy decision has some small impact on $P(\rho > 0)$ and the latter effect, no matter how small, is always decisive (see Corollary 1 to Theorem 4).

To maximize its probability of victory, a campaign has to consider not only policies that affect a nonzero fraction of the electorate but also policies that affect finite sets of agents, whose proportion of the electorate goes to 0 as $n \to \infty$. For illustration purposes, suppose a campaign considers a policy decision that affects just the vote of agent 1. Let $P(i)$ denote the probability that the election excluding agent 1 results in an i vote lead for the right. We will see that maximization of agent 1's expected margin, $R_1 - L_1$, is optimal only if the ratios $\frac{P(1)}{P(0)}$ and $\frac{P(-1)}{P(0)}$ both converge to 1 as the population size n increases.[2] If we fix some π and by chance $\rho(\pi) = 0$ then these ratios will in fact converge to 1. But outside of

[1] The expected margin of a subset of agents J is $\lim_{n\to\infty} \frac{1}{\#\{i\in J:\ i\leqslant n\}} \sum_{i\in J:\ i\leqslant n}(R_i - L_i)$.

[2] For $|i| > 1$, $P(i)$ is irrelevant to any decision that affects only agent 1: agent 1 can change the election's outcome only if there is a tie or a one vote difference among the remainder of the electorate.