

Razor-Thin Elections

David K. Levine¹, Cesar Martinelli²

Abstract

We model head-to-head elections as a competition between incentive schemes to turn out voters. We show that elections are either heavily contested, and decided by thin margins, or safe, meaning that voters in one of the two sides effectively give in, possibly leading to a landslide in favor of the larger side. In equilibrium, as the quality of polling improves, contested elections with razor-thin margins become prevalent.

Email addresses: david@dklevine.com (David K. Levine), cmarti33@gmu.edu (Cesar Martinelli)

¹Department of Economics, EUI and WUSTL

²Department of Economics, George Mason University.

Acknowledgements: We would like to thank Andrea Mattozzi. We gratefully acknowledge support from the MIUR PRIN 2017 n. 2017H5KPLL_01.

1. Introduction

Head-to-head mass elections decided by a few thousand voters are frequent enough as to raise the question of whether they are the deliberate result of political parties' choices. For instance, in the last two presidential runoff elections in Peru, in 2016 and 2021, the margin between the two candidates was 41,057 and 44,263 voters, respectively, out of 17.1 and 17.6 million valid votes.³ In 2016, preelection polls reported close to a tie throughout the campaign.⁴ In 2021, however, the first preelection polls initially 41.5% to Castillo and 21.5% to Fujimori. During the following two months, up to the election, the voting intention for Castillo hovered around 41% while the voting intention for Fujimori, remarkably, crawled to tie that of Castillo, who ended up winning the election.⁵ To give an idea of how close this margin is, if voters were drawn independently with a chance of 49.8% or less of voting for the less popular candidate, the chances of a margin this thin are less than one in one hundred thousand.⁶

Razor-thin margins have occurred with some frequency elsewhere, including, albeit under a different electoral system (the electoral college) several states of the US in the 2021 presidential election. The bottom line is we find it hard to see how such close outcomes could occur unless the parties intentionally try to arrange that the election be close.

Why would a political party prefer a close election to turning out a few extra voters at minimal cost to win the election? We provide an answer by explicitly modeling elections as competition between incentive schemes. There are two ways to turn out voters: one is through costly monitoring and peer pressure, as discussed in Levine and Mattozzi (2020). Alternatively voters may turn out because they are pivotal, as in Palfrey and Rosenthal (1985).⁷ Our central contention is that these are substitutes. While turning out a few extra voters may guarantee a win it also assures that voters are not pivotal and increases costs substantially by requiring costly monitoring and peer pressure: hence—depending on circumstances in a way which we elucidate—a close election may be preferable.

³The official tally is available at <https://www.onpe.gob.pe/elecciones/historico-elecciones/>

⁴See e.g. <https://www.ipsos.com/es-pe/opinion-data-29-de-mayo-de-2016>

⁵IEP Informes de Opinión, <https://iep.org.pe/noticias/informes-de-opinion/>

⁶To a good approximation with independent probabilities p near 50% the standard error of the percent turnout is approximately $(1/2)/\sqrt{N} \approx .01\%$. The observed vote for the less popular candidate was 49.874% of valid turnout. Hence any mean below 49.8% is more than 4.5 standard deviations from 49.874% and the chances of observing a normal more than 4.5 standard deviations from the mean are less than 1 in 100,000.

⁷There is a large literature in economics and political science that we will not attempt to summarize in here. It is common in this literature to contrast peer pressure and other social motivations to strategic voter motives as explanation of turnout; see e.g. Blais (2000). In our model, both peer pressure and strategic voter motivations are possible.

2. The Model

Voters are divided into two parties $k = \{S, L\}$, the small and the large. There are a large number N of voters and party k has $\eta_k N$ members with $\eta_L > \eta_S > 0$. The parties compete in an election for a prize worth N to the party that produces the greatest number of votes. Voting is costly and we order voters so that higher numbered voters have higher participation cost. Specifically, we assume that voter n_k in party k faces a participation cost of $C(n_k/(\eta_k N))$. We assume that this marginal participation cost increases with the fraction of voters participating and does so symmetrically for both parties. This is the *direct cost* of votings net of social pressure: it consists of costs such as those of time, inconvenience, and transportation.

Our core assumption is that each party is able to design a mechanism providing incentives to individual voters in the form of social pressure. We model this as an ability to impose penalties that are costly to individual party members. The mechanism has two parts. The first part is a target fraction ϕ_k together with a rule for party members prescribing voting if $n_k/(\eta_k N) \leq \phi_k$. This means that voters with sufficiently low costs of voting are expected to vote. We refer to this as the *social norm* for the party. The second part of the mechanism is a punishment $P_k \geq 0$ representing social disapproval for failing to comply with the social norm. This punishment may also be costly to party members who carry out the punishment: if Tim is punished by David refusing to have a beer with him this may be costly to David as well as to Tim. Hence the social cost of the punishment is taken to be ψP_k where $\psi > 0$ (possibly greater than one). This punishment takes place after the results of the election are known.

The ability to apply punishments is limited by imperfect information. While there is no difficulty in determining whether or not a party member voted, individual costs of voting are not transparent and the party can only observe a noisy binary signal about whether or not a non-voter followed the social norm. In particular, among voters that failed to vote there are two types: high cost voters who were “excused” according to the social norm and low cost voters who were not. If the non-voting member violated the social norm, that is, $n_k/(\eta_k N) \leq \phi_k$, a negative signal is received with probability $\pi_1 \in [0, 1]$ and a positive signal with corresponding probability $1 - \pi_1$. A non-voting member with $n_k/(\eta_k N) > \phi_k + \mu$ where $\mu > 0$ clearly has a good excuse so always receives a positive signal.⁸ The remainder who did not violate the social norm, that is, $\eta_k N(\phi_k + \mu) \geq n_k > \eta_k N\phi_k$, receive a negative signal with probability π .

A positive signal should be thought of as “producing a good excuse for not voting” so that a non-voter who is genuinely excused is more likely to be able

⁸Note that this simplifies the formulation in Levine and Mattozzi (2020) who have an upper bound on the number of voters and assume that all non-voters must be punished. That results in eventually decreasing monitoring costs: here we limit attention to the region of increasing monitoring costs in order to focus on the substitutability between pivotality and peer punishment.

to produce a good excuse than one who is not. Punishment can then be based based on whether or not the party member voted, on the signal received and may also be based on whether or not the election was won or lost. For notational simplicity we assume that only bad signal are punished and that when the probability of losing the election is strictly between zero and one punishment occurs only when there is a loss. We will subsequently show that this is without loss of generality as it is the optimal punishment scheme.

In addition it is possible that a party might wish to prevent members for voting. Again for notational simplicity we assume that it can do so costlessly, for example because it can perfectly observe and punish voters who were not supposed to vote. As in fact it is never optimal to prevent members from voting, we will show that this assumption also is without loss of generality.

The choice of turnout ϕ_k takes place in two stages. In the first stage each party announces a commitment of ϕ_k^0 voters. One party K always by accident or design always attains commitment first, with each party having an equal chance of being that *first mover*. The party that fails to commit first, the *second mover* then chooses a level of turnout ϕ_{-K}^1 . Each party has an equal chance of being first mover.

We want to reflect the fact that depending on the quality of polling and other data it is difficult to know exactly what the turnout of the other party is. We model this by assuming that actual party turnout is a random function of intended party turnout: specifically we assume that in addition to the $\phi_k \eta_k N$ intended voters each party costlessly receives a random number ζ_k of additional votes. These shocks are independently drawn from the same distribution. The actual vote differential is then equal to the intended vote differential $h_k \equiv \phi_k \eta_k N - \phi_{-k} \eta_{-k} N$ plus an error $\xi_k = \zeta_k - \zeta_{-k}$. Let $\sigma^2 = \text{var} \xi_k$ and let F denote the cdf for the normalized random variable ξ_k / σ . To cleanly distinguish between *close elections* and *sure elections* we assume that there is a cutoff σL such that $F(L) = 1$ and $F(h_k / \sigma) < 1$ for $h_k < \sigma L$. In this setup by choosing h_k , the second mover determines whether the election is close or sure. We assume, moreover, that the second mover cannot choose $h_k = \sigma L$ exactly but may choose either σL^- slightly below σL or σL^+ slightly above σL . The meaning of this we explain in the next paragraph.

In principle the cdf F should be discrete but we approximate it as continuous except at the cutoff σL . Let f be the density corresponding to F for $\xi_k < \sigma L$. Then the normalized density $(1/\sigma)f(h_k/\sigma)$ is an approximate measure of pivotality: the probability of a tie, the probability of a one vote loss, and the probability of a one vote win. Specifically, in a close election, an individual who shifts from voting to not voting increases the probability of a loss by half the probability of a tie or a one vote win, which is to say, (approximately) $(1/\sigma)f(h_k/\sigma)$. In case the second mover chooses σL^- we take the probability of winning to be $F^- \equiv \lim_{h_k \uparrow \sigma L} F(h_k/\sigma)$ and the probability of being pivotal to be $(1/\sigma)f^- \equiv (1/\sigma) \lim_{h_k \uparrow \sigma L} f(h_k/\sigma)$. In case $h_k > \sigma L$, including $h_k = \sigma L^+$, the probability of winning is 1 and the probability of being pivotal is 0. Similarly, in case $h_k < -\sigma L$, including $h_k = -\sigma L^+$, the probability of winning is 0 and the probability of being pivotal is 0.

We assume that individual voting decisions are made after the second mover determines h_{-K} and punishment schemes must satisfy the interim incentive compatibility constraint that conditional on knowing ϕ_K^0 and h_{-k} party members who are supposed to vote are willing to do so and those who are not supposed to vote are willing not to do so. Given this, parties attempt to maximize the expected value of winning the prize less the expected costs of turning out voters. For convenience we will normalize this utility dividing by N .

Functional Form and Parameter Restrictions

We make two assumptions about functional form. First, we assume that the direct cost of voting $C(n_k/(\eta_k N)) = cn_k/(\eta_k N)$ is linear. This implies that the total cost of voting is quadratic and given by $\eta_k N \phi_k^2/2$. Second, we assume that the additional voters ζ_k are drawn from an exponential distribution with mean σ . The latter implies that ξ_k follows a Laplace distribution with mean 0 and scale parameter σ so that for positive x we have $f(x) = (1/2)e^{-x}$ and $F(x) = 1 - (1/2)e^{-x}$.

Cost of Voting

We limit attention to the case in which direct costs are high enough that the constraint $\phi_k \leq 1$ does not bind and that the larger party does not overwhelm the smaller party. Specifically, we assume that

$$c > 2/((1 - \mu)^2 \eta_S),$$

$$c > 2\eta_L/\eta_S^2.$$

The former insures that the direct cost of turning out a fraction $1 - \mu$ voters $\eta_k N c (1 - \mu)^2/2$ is strictly greater than the value N of the prize for both parties so that the fraction of non-voters will always exceed μ . The latter assures that the most voters L is willing to turn out $N\sqrt{2/(\eta_L c)}$ does not exceed the greatest number of voters $\eta_S N$ that the small party is able to turnout.

Electoral Shock

We assume that size of the random shock to the number of voters as measured by σL is small relative to the the size of the population but reasonably large in absolute terms. Specifically, we assume that

$$\sigma L/N \approx 0,$$

$$e^{-\sigma L} \approx 0,$$

$$\sigma \geq 1.$$

The first assumption is that close elections involve a small fraction of voters, while the second is that at the truncation point the probability of a loss is negligible. The final assumption assures that the probability the second mover loses when a single voter fails to vote $(1/2)e^{-h_k/\sigma} + (1/\sigma)(1/2)e^{-h_k/\sigma} \leq 1$ for all $h_k \geq 0$.

Approximation

We adopt the approximation that $(1/\sigma)f(h_k/\sigma)$ is the probability of being pivotal.

3. Objective Functions

We now analyze the objective functions and incentive constraints, and introduce a useful approximation for the case of close elections that we will adopt throughout the remainder of the paper. Define $\theta = \mu\psi\pi/\pi_1$, and let the probability that party k wins be Q_k .

Theorem 3.1. *For sure elections with $N\eta_{-K}\phi_{-K} \in [0, N\eta_K\phi_K - \sigma L^+]$, or $N\eta_{-K}\phi_{-K} \in [N\eta_K\phi_K + \sigma L^+, N\eta_{-K}]$ with optimal punishment the normalized objective function of party k is*

$$U_k^s = Q_k - (1/2)c\eta_k(\phi_k)^2 - \eta_k\theta c\phi_k.$$

Define

$$R = \frac{1 - Q_K}{1 - Q_K + (1/\sigma)f(e_{-K}/\sigma)} \max\{0, c\phi_K - (1/\sigma)f(h_{-K}/\sigma)/\eta_K\}.$$

For close elections with $\eta_{-K}N\phi_{-K} = \eta_KN\phi_K + h_{-K}$ and $h_{-K} \in [-\sigma L^-, \sigma L^-]$, with optimal punishment the normalized objective functions are

$$U_K^c = Q_K - (1/2)c\eta_K(\phi_K)^2 - \eta_K\theta R$$

and

$$U_{-K}^c = Q_{-K} - (1/2)c(\eta_K/\eta_{-K})\eta_K(\phi_K)^2 - \eta_K\theta R + O(\sigma L/N).$$

Throughout the remainder of the paper we will adopt the approximation that $U_{-K}^c = Q_{-K} - (1/2)c(\eta_K/\eta_{-K})\eta_K(\phi_K)^2 - \eta_K\theta R$.

Proof. Consider first sure elections. In a sure election either $N\eta_{-K}\phi_{-K} \in [0, N\eta_K\phi_K - \sigma L)$ (including $N\eta_K\phi_K - \sigma L^+$) in which case the probabilities of winning are $Q_K = 1, Q_{-K} = 0$ or $N\eta_{-K}\phi_{-K} \in (N\eta_K\phi_K + \sigma L, \infty)$ (including $N\eta_K\phi_K + \sigma L^+$) in which case the probabilities of winning are $Q_K = 0, Q_{-K} = 1$. There is no pivotality and punishment P_k is for bad signals. Notice that the cost of punishment ψP_k must be paid for the fraction of the non-voting population for whom bad signals can be received, μ , times the probability π they get bad signals. Hence the total social cost of punishment is $\eta_k N\pi\mu\psi P_k$. The normalized objective function of party k is therefore

$$Q_k - (1/2)c\eta_k(\phi_k)^2 - \eta_k\pi\mu\psi P_k.$$

The marginal voter n_k receives $-cn_k/(N\eta_k) = -c\phi_k$ for voting and $-\pi_1 P$ for not voting. Hence the incentive constraint is

$$P_k \geq c\phi_k/\pi_1.$$

Since punishment is socially costly, optimality requires that it be chosen as small as possible subject to the incentive constraints. Solving the incentive constraints with equality and substituting into (approximate) objective functions gives

$$Q_k - (1/2)c\eta_k(\phi_k)^2 - \eta_k(\pi/\pi_1)\mu\psi c\phi_k$$

for sure elections. Here the final term is the *monitoring cost* due to the need to provide incentives.

Now consider close elections. In a close election, the intended turnout of the first mover is $\eta_K N\phi_K$ and the turnout of the second mover is $\eta_K N\phi_K + h_{-K}$ where $h_{-K} \in [-\sigma L^-, \sigma L^-]$. The probability the second mover wins is $Q_{-K} = F(h_{-K}/\sigma)$, the probability the first mover wins is $1 - Q_{-K}$, and the probability of being pivotal is $(1/\sigma)f(h_{-K}/\sigma)$ for both parties. Suppose that for close elections punishment is conditional on losing the election; we will show later on that this is without loss of generality since it reduces the expected monitoring cost. Conditioning on losing the election, the expected monitoring cost is $(1 - Q_k)\pi\mu\psi P_k$. The normalized objective function of party K is therefore

$$Q_K - (1/2)c\eta_K(\phi_K)^2 - \eta_K(1 - Q_K)\pi\mu\psi P_K.$$

That of the second mover is

$$Q_{-K} - (1/2)c\eta_{-K} \left(\frac{\eta_K N\phi_K + h_{-K}}{\eta_{-K} N} \right)^2 - \eta_{-K}(1 - Q_{-K})\pi\mu\psi P_{-K}.$$

By assumption $|h_{-K}| \leq \sigma L \ll N$ we can approximate h_{-K}/N as being zero, and we have the approximation

$$Q_{-K} - (1/2)c(\eta_K/\eta_{-K})\eta_K(\phi_K)^2 - \eta_{-K}(1 - Q_{-K})\psi\pi\mu P_{-K}.$$

The first mover's marginal voter receives $-cn_K/(N\eta_K) + (1/\sigma)f(h_{-K}/\sigma)/\eta_k$ for voting and $-\pi_1(1 - Q_K + (1/\sigma)f(h_{-K}/\sigma))P_K$ for not voting, so the incentive constraint is

$$P_K \geq \frac{c\phi_K - (1/\sigma)f(h_{-K}/\sigma)/\eta_K}{\pi_1(1 - Q_K + (1/\sigma)f(h_{-K}/\sigma))}.$$

Notice that if the punishment were not given conditional on losing the election, then the incentive constraint would be $P_K \geq (c\phi_K - (1/\sigma)f(h_{-K}/\sigma)/\eta_K)/\pi_1$, but the expected cost of monitoring would be proportional to P_K rather than to $(1 - Q_K)P_K$, and therefore would be at least as large and strictly larger if the expected cost of punishment were positive.

The second mover's marginal voter n_{-K} receives

$$-cn_{-K}/(N\eta_{-K}) + (1/\sigma)f(h_{-K}/\sigma)/\eta_{-K} = \\ -c\frac{\eta_K N\phi_K + h_{-K}}{N\eta_{-K}} + (1/\sigma)f(h_{-K}/\sigma)/\eta_{-K}$$

for voting and $-\pi_1(1 - Q_{-K} + (1/\sigma)f(h_{-K}/\sigma))P_{-K}$ for not voting. Again we propose to approximate by ignoring the small term h_{-K}/N , so the incentive constraint is

$$P_{-K} \geq \frac{c(\eta_K/\eta_{-K})\phi_K - (1/\sigma)f(h_{-K}/\sigma)/\eta_{-K}}{\pi_1(1 - Q_{-K} + (1/\sigma)f(h_{-K}/\sigma))}.$$

By a similar argument to the first mover, if the punishment were not conditional on losing the election, the expected cost of monitoring would be at least as large and strictly larger if the expected cost of punishment were positive.

For close elections we must take account of the fact that the lower bound on punishment might be negative due to pivotality. Hence we solve the incentive constraints with equality to get P_k and substituting $\max\{0, P_k\}$ into (approximate) objective functions. For the first mover we have

$$Q_K - (1/2)c\eta_K(\phi_K)^2 \\ -\eta_K(\pi/\pi_1)\mu\psi\frac{1 - Q_K}{1 - Q_K + (1/\sigma)f(h_{-K}/\sigma)} \max\{0, c\phi_K - (1/\sigma)f(h_{-K}/\sigma)/\eta_K\}$$

and for the second

$$Q_{-K} - (1/2)c(\eta_K/\eta_{-K})\eta_K(\phi_K)^2 \\ -\eta_K(\pi/\pi_1)\mu\psi\frac{1 - Q_{-K}}{1 - Q_{-K} + (1/\sigma)f(h_{-K}/\sigma)} \max\{0, c\phi_K - (1/\sigma)f(h_{-K}/\sigma)/\eta_K\},$$

as given in the statement of the theorem. \square

4. Equilibrium

Our main theorem is

Theorem 4.1. *In any equilibrium there are three possible type of commitment: preemptive in which $\phi_L > 0$ and $h_S = 0$, concession in which $\phi_K = 0$ and $h_{-K} = \sigma L$, or contested in which commitment is $\eta_K\phi_K = (1/2)(1/\sigma)$ and the second mover responds with $h_{-K} = 0$. The small group never preempts, and for all $c, \mu, \pi, \pi_1, \eta_L, \eta_S, \sigma$ there is some $\bar{\psi} > 0$ such that for $\psi > \bar{\psi}$ the large group does not preempt and for $\psi < \bar{\psi}$ the large group does preempt. In addition for fixed $c, \mu, \pi, \pi_1, \sigma$ we have $\bar{\psi} \rightarrow 0$ as $\eta_S/\eta_L \rightarrow 1$ and in particular if $\eta_S/\eta_L > 1/3$ the election is contested whenever $\psi > \pi_1(\sigma + 1)/(\mu\pi)$.*

This follows from Theorems 4.2 and 4.3 below. A key point of the theorem is that decreasing σ , meaning better targeting of voters by the parties, results in closer vote totals and higher turnout as long as groups are not too asymmetric.

Backwards Induction

To solve for subgame perfect equilibrium and prove the theorem, we proceed by backward induction. We first calculate the best responses of the second mover and then we turn to the optimal behavior for the first mover.

Theorem 4.2. *The best response of the second mover is either $\phi_{-K} = 0$ or as follows if the solution yields utility bigger than zero:*

- (i) if either $\phi_K < (1/c)(1/\sigma)(1/2)e^{-L}/\eta_K$ or $\theta/(\sigma+1) < 1$, it is $h_{-K} = \sigma L^-$.
- (ii) if $\phi_K > (1/c)(1/\sigma)(1/2)e^{-L}/\eta_K$ and $\theta/(\sigma+1) > 1$, it is $h_{-K} = -\sigma \log(2\sigma c\eta_K\phi_K)$ if $\sigma c\eta_K\phi_K < 1/2$, and it is $h_{-K} = 0$ otherwise.

Proof. The second mover can guarantee a payoff of zero by choosing $\phi_{-K} = 0$. Therefore the second mover will only choose a close election if it can guarantee itself a positive payoff. The second mover would never choose $h_{-K} < 0$ since given the probability of winning the election would be higher choosing $|h_{-K}|$, the monitoring cost would be the same, given the Laplace distribution assumption, and the cost of voting would be approximately the same. Using the Laplace distribution, in a close election the second mover's objective function for $h_{-K} \geq 0$ is

$$1 - (1/2)e^{-h_{-K}/\sigma} - (1/2)c(\eta_K/\eta_{-K})\eta_K(\phi_K)^2 - \eta_K(\pi_1/\pi)\mu\psi \frac{\sigma}{\sigma+1} \max\{0, c\phi_K - (1/\sigma)(1/2)e^{-h_{-K}/\sigma}/\eta_K\}.$$

For $c\phi_K - (1/\sigma)(1/2)e^{-h_{-K}/\sigma}/\eta_K < 0$ this is strictly increasing in h_{-K} , so either $\phi_K < (1/c)(1/\sigma)(1/2)e^{-L}/\eta_K$ in which case the optimum is σL^- or the derivative of the objective function with respect to h_{-K} is

$$\begin{aligned} & (1/\sigma)(1/2)e^{-h_{-K}/\sigma} - \eta_K(\pi/\pi_1)\mu\psi \frac{\sigma}{\sigma+1} (1/\sigma^2)(1/2)e^{-h_{-K}/\sigma}/\eta_K \\ &= \left(1 - (\pi/\pi_1)\mu\psi \frac{1}{\sigma+1}\right) (1/\sigma)(1/2)e^{-h_{-K}/\sigma}. \end{aligned}$$

Hence for $(\pi/\pi_1)\mu\psi/(\sigma+1) < 1$ the solution is $h_{-K} = \sigma L^-$ and for $(\pi/\pi_1)\mu\psi/(\sigma+1) > 1$ the solution entails to reduce h_{-K} until either $e^{-h_{-K}/\sigma} = 2\sigma c\eta_K\phi_K$, so that $c\phi_K - (1/\sigma)(1/2)e^{-h_{-K}/\sigma}/\eta_K = 0$ and the cost of monitoring is zero, or it is given by $h_{-K} = 0$. This finishes the proof of the theorem.

Notice that in order to gain the election with probability near one the second mover can choose $h_{-K} = \sigma L^-$. We are assuming e^{-L} negligible and $\sigma \geq 1$ so that $(1/\sigma)e^{-L}$ is also negligible; this means that we attribute to σL^- no chance of a loss and no chance of being pivotal, but the multiplier

$$\frac{1 - Q_{-K}}{1 - Q_{-K} + (1/\sigma)f(e_{-K}/\sigma)}$$

is still equal to $\sigma/(\sigma + 1)$ so that σL^- is better than σL^+ . That is, the second mover will only induce a sure election if it concedes with $\phi_{-K} = 0$. \square

We can revisit now the assumption that voters with costs above the marginal voter can be costlessly deterred from voting. For sure elections, since the marginal voter $n_k = \phi_k N \eta_k$ is indifferent between voting or not, voters with larger costs will not be tempted to vote even if there is no punishment or monitoring for them. For close elections, given the best response behavior of the second mover, both the first mover and the second mover's marginal voters are indifferent between voting or not, so again voters with larger costs than the marginal voter are not tempted to vote in equilibrium. Thus, the assumption is without loss of generality.

Our main Theorem now follows from

Theorem 4.3. *If $\theta < \sigma + 1$, the first mover commit to $\phi_K = 0$ and the second mover responds with $h_{-K} = \sigma L^-$, or the first mover preempts. If $\theta > \sigma + 1$, the first mover commits to $\eta_K \phi_K = (1/2)(1/\sigma)$ and the second mover responds with $\eta_{-K} \phi_{-K} = (1/2)(1/\sigma)$ and $h_{-K} = 0$, or the first mover preempts. The small group never preempts, and for all $c, \eta_L, \eta_S, \sigma$ there is some $\bar{\theta}$ such that for $\theta > \bar{\theta}$ the large group does not preempt and for $\theta < \bar{\theta}$ the large group does preempt. Moreover, for all c, σ , we have $\bar{\theta} < \sigma + 1$ if $\eta_S/\eta_L > 1/3$ and $\bar{\theta} \rightarrow 0$ as $\eta_S/\eta_L \rightarrow 1$.*

Proof. Consider first the case $\theta/(\sigma + 1) < 1$. Given the best response of the second mover in Theorem 4.2, the second mover wins the election with probability one as long as it can guarantee itself a positive payoff. Thus, the first mover either preempts by choosing ϕ_K large enough so that the second mover is better off conceding and choosing $\phi_{-K} = 0$, or it concedes by choosing $\phi_K = 0$, and the second mover chooses $h_{-K} = \sigma L^-$. Preemption by the first mover requires that the group's turnout choice, say ϕ^* , is large enough to keep the second mover out, that is

$$1 - (1/2)c(\eta_K/\eta_{-K})\eta_K(\phi^*)^2 - \frac{\sigma}{1 + \sigma}\eta_K\theta c\phi^* = 0, \quad (4.1)$$

but small enough so that the first mover has a positive gain by not conceding, that is

$$1 - (1/2)c\eta_K(\phi^*)^2 - \eta_K\theta c\phi^* \geq 0.$$

The inequality cannot be satisfied if the first mover is the small group, since for $\eta_K/\eta_{-K} < 1$ the LHS of the inequality is smaller than the LHS of equation 4.1. Hence the small group never preempts for $\theta/(\sigma + 1) < 1$. Consider that the large group is the first mover. As $\theta \rightarrow 0$ as $0 \leq \phi^* \leq 1$ we see that $(1/2)c\eta_L(\phi^*)^2 \rightarrow \eta_S/\eta_L$. Hence the LHS of the inequality approaches $1 - \eta_S/\eta_L > 0$. Thus, the large group will choose to preempt for small enough θ .

We claim that if for some $0 < \bar{\theta} < \sigma + 1$ the large group is indifferent or prefers not to preempt, then the large group prefers not to preempt for all

$\bar{\theta} < \theta \leq \sigma + 1$. To see this, notice that the first mover prefers not to preempt if ϕ^* is larger than $\bar{\phi}^*$ given by the unique solution for $\bar{\phi}^* \geq 0$ of

$$1 - (1/2)c\eta_L(\bar{\phi}^*)^2 - \eta_L\theta c\bar{\phi}^* = 0.$$

Applying the implicit function theorem to this equation and equation 4.1 for ϕ^* we have

$$\partial\bar{\phi}^*/\partial\theta = -\bar{\phi}^*/(\bar{\phi}^* + \theta)$$

and

$$\partial\phi^*/\partial\theta = -\phi^*/(((1 + \sigma)/\sigma)(\eta_L/\eta_S)\phi^* + \theta).$$

Hence for $\bar{\phi}^* < ((1 + \sigma)/\sigma)(\eta_L/\eta_S)\phi^*$ it must be that $\partial\bar{\phi}^*/\partial\theta < \partial\phi^*/\partial\theta$. Suppose that for some $0 < \bar{\theta} < \sigma + 1$ the large group is indifferent or prefers not to preempt: so it must be that $\phi^* \geq \bar{\phi}^*$. So certainly $((1 + \sigma)/\sigma)(\eta_L/\eta_S)\phi^* > \bar{\phi}^*$ so $\partial\bar{\phi}^*/\partial\theta < \partial\phi^*/\partial\theta$. Hence for $\theta > \bar{\theta}$ we must have $\phi^* \geq \bar{\phi}^*$ implying that the large group strictly prefers not to preempt.

Now consider the case $\theta/(\sigma + 1) > 1$. The first mover loses for sure if $\phi_K < (1/c)(1/\sigma)(1/2)e^{-L}/\eta_K$ so any such choice is dominated by $\phi_K = 0$. In the range $(1/c)(1/\sigma)(1/2)e^{-L}/\eta_K \leq \phi_K \leq (1/c)(1/\sigma)(1/2)/\eta_K$, given the best response of the second mover, the payoff of the first mover is

$$U_K^c = \sigma c\eta_K\phi_K - (1/2)c\eta_K(\phi_K)^2,$$

which is strictly increasing in ϕ_K since $\phi_K < 1 \leq \sigma$. Hence, if choosing a close election, the first mover chooses

$$\phi_K = (1/2)(1/\sigma)(1/c)/\eta_K,$$

which is the minimum turnout guaranteeing $h_{-K} = 0$; any higher choice in a close election leads to the same probability of winning, larger turnout costs, and positive monitoring costs. The first mover utility of choosing a close election is therefore

$$\hat{U}_K^c = (1/2) - (1/2)c\eta_K((1/2)(1/\sigma)(1/c)/\eta_K)^2 = 1/2 - (1/8)(1/\sigma^2)(1/(c\eta_K)).$$

Notice that the first mover utility of choosing a close election is always better than conceding since $\sigma \geq 1$ and the lower bound $c > 2\eta_L/\eta_S^2$ implies $c\eta_S > 2$.

Now we need to consider whether it would be better for the first mover to preempt. If it is profitable for the first mover to preempt, it must be profitable for them to keep the second mover out, that is, the turnout choice of the first mover $\phi^{**} > (1/2)(1/\sigma)(1/c)/\eta_K$ must be large enough so that the payoff of $h_{-K} = 0$ is at most zero,

$$1 - (1/2)c(\eta_K/\eta_{-K})\eta_K(\phi^{**})^2 - \eta_K\theta(c\phi^{**}\sigma/(1 + \sigma) - 1/(2\sigma\eta_K)) = 0, \quad (4.2)$$

but small enough so that the first mover has a positive gain by not accepting a

close election, that is

$$1 - (1/2)c\eta_K(\phi^{**})^2 - \eta_K\theta c\phi^{**} \geq 1/2 - (1/8)(1/\sigma^2)(1/(c\eta_K)).$$

The inequality cannot be satisfied if the first mover is the small group, since for $\eta_K/\eta_{-K} < 1$ the LHS of the inequality is smaller than the LHS of the equation 4.2, but the RHS of the inequality is larger than the RHS of equation 4.2. Hence the small group never preempts for $\theta/(\sigma+1) < 1$. Consider that the large group is the first mover. As $0 \leq \phi^{**} \leq 1$ for $\theta \rightarrow \infty$ the first equation implies that $c\phi^{**}\sigma/(1+\sigma) - 1/(2\sigma\eta_L) \rightarrow 0$, in particular ϕ^{**} is bounded away from zero so the RHS of the inequality must approach $-\infty$ and the inequality must be violated. Thus, the large group will choose not to preempt for large enough θ .

Finally, we claim that if for some $\sigma+1 \leq \bar{\theta}$ the large group is indifferent or prefers not to preempt, then the large group prefers not to preempt for all $\bar{\theta} < \theta$. To see this, notice that the first mover prefers not to preempt if ϕ^{**} is larger than $\bar{\phi}^{**}$ given by

$$1 - (1/2)c\eta_L(\bar{\phi}^{**})^2 - \eta_L\theta c\bar{\phi}^{**} = 1/2 - (1/8)(1/\sigma^2)(1/(c\eta_L)).$$

Applying the implicit function theorem to this and equation 4.2 for ϕ^{**} we have

$$\partial\bar{\phi}^{**}/\partial\theta = -\bar{\phi}^{**}/(\bar{\phi}^{**} + \theta)$$

and

$$\partial\phi^{**}/\partial\theta > -\phi^{**}/(((1+\sigma)/\sigma)(\eta_L/\eta_S)\phi^{**} + \theta).$$

Hence for $\bar{\phi}^{**} < ((1+\sigma)/\sigma)(\eta_L/\eta_S)\phi^{**}$ we have $\partial\bar{\phi}^{**}/\partial\theta < \partial\phi^{**}/\partial\theta$. It follows that $\bar{\phi}^{**} \leq \phi^{**}$ for $\bar{\theta} \geq \sigma+1$ implies $\bar{\phi}^{**} < \phi^{**}$ for all $\bar{\theta} < \theta$.

From previous steps, for all $c, \eta_L, \eta_S, \sigma$ there is some $\bar{\theta}$ such that for $\theta > \bar{\theta}$ the large group does not preempt and for $\theta < \bar{\theta}$ the large group does preempt. We claim that for all c, σ , we have $\bar{\theta} < \sigma+1$ if $\eta_S/\eta_L \geq 1/3$. To prove this, we seek to establish that for $\eta_S/\eta_L > 1/3$, the LHS of equation 4.2 is larger than the LHS of the positive gain inequality, but the RHS of the inequality is larger than the RHS of equation 4.2 for all $\theta \geq \sigma+1$. This is the case if

$$\begin{aligned} - (1/2)c(\eta_L/\eta_S)\eta_L(\phi^{**})^2 - \eta_L\theta(c\phi^{**}\sigma/(1+\sigma) - 1/(2\sigma\eta_L)) \\ > - (1/2)c\eta_L(\phi^{**})^2 - \eta_L\theta c\phi^{**}, \end{aligned}$$

or equivalently

$$\eta_L\theta c\phi^{**}(1 - \sigma/(1+\sigma)) + \theta/(2\sigma) > (1/2)(\eta_L/\eta_S - 1)c\eta_L(\phi^{**})^2.$$

Using $(\phi^{**})^2 < \phi^{**}$, it is enough to show that $\theta(1 - \sigma/(1+\sigma)) \geq (1/2)(\eta_L/\eta_S - 1)$, which verifies for $\theta \geq \sigma+1$ as long as $\eta_S/\eta_L \geq 1/3$. Similarly, fixing any $1+\sigma > \bar{\theta} > 0$, for η_S/η_L close enough to one, the LHS of equation 4.1 is larger than the LHS of the positive gain inequality, but the RHS of the inequality is larger than the RHS of equation 4.1 for all $\theta \geq \bar{\theta}$, so that $\bar{\theta} \rightarrow 0$ as

$\eta_S/\eta_L \rightarrow 1$. □

From the statement of Theorem 4.3, for η_S/η_L close enough to one, we have $\bar{\theta} < \mu\pi/\pi_1$. Hence, for enough symmetry in the support of the parties, there are no preemption equilibria for $\psi \geq 1$.

5. Conclusion

We propose a dynamic model of electoral competition to explain the apparent abundance of knife-edge elections. If the cost of monitoring voters is low, and the small group commits to a turnout level first, it concedes the election, and if the large group commits to a turnout level first, it either concedes, or (if there is enough asymmetry in the electoral support) it preempts and wins by a landslide. If instead the cost of monitoring voters is high, there is a knife-edge election regardless of which group moves first. Moreover, the cost threshold moves down as targeting of voters becomes more precise, which implies that better targeting of voters will make head-to-head elections increasingly knife-edge.

References

Blais, André (2000), *To Vote or Not to Vote: The Merits and Limits of Rational Choice Theory*. University of Pittsburgh Press.

Levine, David K. and Andrea Mattozzi (2020), Voter Turnout with Peer Punishment. *American Economic Review* 110: pp. 3298–3314.

Palfrey, Thomas R., and Howard Rosenthal (1985), Voter Participation and Strategic Uncertainty. *American Political Science Review* 79: 62–78.