Political Influence Through Microtargeting^{*}

Michael Eldar

Nuffield College, University of Oxford, Department of Economics michael.eldar@economics.ox.ac.uk

Sinem Hidir

University of Warwick, Department of Economics s.hidir@warwick.ac.uk

February 6, 2025

Abstract

Political actors routinely target custom audiences on social media in order to influence elections. We model this process, focusing on the way in which it induces voters to learn about their own preferences. This differs from past literature, which has focused on party platforms and the effects of bias. We find that the optimal strategy based on some empirically estimated parameters is to target groups favoring one's opponents, providing a rational explanation for negative campaigning. More generally, log-concave cost of voting distributions can give rise to a nonconvex set being targeted—weak supporters of the politician and strong supporters of their opponent. We make use of this setup to provide a novel analysis of the effects of micro-targeting on turnout, and find a sense in which lower costs of voting encourage negative campaigning.

Keywords: Microtargeting; Negative Campaigning; Mobilisation; Demobilisation; Political Economy; Social Media; and Political Influence

^{*}We would like to thank D. Carlos Akkar, Jean-Paul Carvalho, Ian Jewitt, Dan Bernhardt, Mehdi Shadmehr, Alvaro Delgado Vega, Lukas Leucht, Meg Meyer, Inés Moreno de Barreda, Aidan Smith, Hubert Wu, participants of the University of Oxford's Student Theory Workshop, and the University of Oxford's Political Economy seminar series for valuable feedback and comments.

1 Introduction

Micro-targeting is the use of a range of observable characteristics to send particular messages to particular groups of people. For example, a supermarket might choose to send an advert about back to school supplies to families with young children just before the beginning of the academic year. Micro-targeting is most effectively used on social media platforms, where advertisers choose how much to spend on ads sent to particular types of audiences¹. Facebook allows advertisers to choose audiences based on age, gender, ethnicity, interests, occupation, and even past engagement history.²

Since Barack Obama's 2012 presidential campaign, micro-targeting has been a prominent tool used by political campaigns to improve the effectiveness of their messaging by concentrating on groups which they think they can sway in a useful direction. Microtargeting is seen by the political class as so effective that political influence firm Cambridge Analytica, which has been employed by a range of right-wing political movements, was famously at the center of a scandal in 2018 when it was revealed to have illegally harvested personal profiles of up to 87 million people so that they could improve their micro-targeting precision (Kang 2018).

Much of the attention that has been paid to political micro-targeting has focused on persuasion through fake and sensationalized content by enabling messages to be tailored to the different behavioral biases of different parts of society. According to this, if one could effectively moderate political advertising so that it does not include fake or highly sensationalized content, there need not be worry about the implications of political micro-targeting.³

However, whilst knowing what made an electoral campaign a success or not is difficult, there is evidence that successful micro-targeting campaigns do not always use fake news. In 2010, Cambridge Analytica orchestrated the "Do So!" campaign in Trinidad and Tobago, which they claim reduced turnout among young black people enough to swing an election by critiquing the political system itself, encouraging them to simply disengage (Amer and Noujaim 2019). While this campaign has most likely affected the turnout and hence the outcome of the election, it was not found to rely on fake news. Instead, it seems that Cambridge Analytica targeted a youth that was unsure about the value of participating in electoral politics.

Recent research shows the 'fake' part of fake news is not necessarily what drives its political influence. Nyhan et al. (2020) and Barrera et al. (2020) find that correcting

¹Throughout this paper we discuss micro-targeting as a method used in social media because it is the most natural application, but there are many instances of political communication outside of social media in which our formal notion of microtargeting can be applied. This is briefly discussed in section 7.

 $^{^{2}} See \ www.facebook.com/business/ads/ad-targeting$

³Of course, the will to engage in fact-checking that aims to eliminate fake content has waned as the political interests of the owners of tech companies have shifted (Bradshaw and Craggs Mersinoglu 2025).

people's misconceptions does not significantly alter their decision whether to vote or not for political candidates. In particular, Barrera et al. (2020) find that fake news is effective, but its effectiveness does not seem to stem from voters believing the falsehoods contained in it. In a more recent experimental analysis of sharing of fake news, Guriev et al. (2023) find that people are relatively good at identifying falsehoods when they are simply nudged to think about fake news. These findings together suggest that political influence online does not in general work by persuading people to believe in fake news.

We aim to give insight into how campaigns might function via a theory of political influence through micro-targeting that does not rely on behavioral biases or fake news. Instead, we study how politicians can use micro-targeting to selectively improve how much particular groups learn about *their own* preferences towards the proposals of the opposing parties.

For example, consider a policy proposal to open more routes for legal immigration into a country. One possible approach to thinking about online influence regarding this policy proposal is to look at the extent to which false information about immigrants is spread online. This paper, on the other hand, considers the effect of people gauging their own emotional reactions to news related to immigrants in their country. If every time they see an article about immigrants they find themselves supportive, then they learn that they probably do want to support more routes to legal immigration. If on the other hand they have a negative reaction when they read about immigrants, then they learn that they do not want to support more legal routes. Assuming this issue dominates an election, then their voting (or abstaining) decision will depend on their inclination on this issue in relation to what the political parties propose. If this mechanism is important, what specific groups would politicians want to send additional content to?

Another possible situation is voters deciding whether their country should leave a political union, such as in the Brexit referendum. This is a complex decision, based on a range of considerations, many of which are emotional. An unsure voter consumes news content on this issue. Every piece of content elicits some emotional reaction that makes the voter feel either like they belong in the union or not, which might shape their voting decision. Voters learn about how they feel on the political issue through engaging with it. This emotional reaction follows a different distribution for each voter, which is correlated with their observable characteristics. This means that by targeting which audiences will receive more content about the union, politicians can shift opinions in predictable ways by simply prompting certain people to reflect on a topic.

Using this intuition, we build a model where politicians identify groups using their observable characteristics and prompt them to reflect on their views through sending content online. In this setting, the cost of voting distribution, which in our model can be interpreted as how strong a voter's preference for one candidate over another needs to be before they go out to vote, becomes crucially important.

We find that based on the cost of voting distribution estimated by Blais et al. (2019) (monotonically decreasing PDF), targeting groups that are supportive of one's opponents is the optimal strategy. This could be interpreted as an explanation for the phenomenon of 'negative campaigning' without relying on the commonly used assumptions of behavioral biases (Skaperdas and Grofman 1995), or lack of information about policy positions (Titova 2024; Schipper and Woo 2018; Harrington and Hess 1996). This is driven by the fact that given the estimated cost of voting distribution, in expectation more voters are likely to be demobilized from voting for the opponent after learning about their own preferences compared to those being mobilized to vote who were previously inclined to abstain.

With a monotonically increasing PDF, the logic reverses and instead it becomes profitable to target groups supportive of oneself. This is because in this case the mobilization effect outweighs the demobilization effect, in expectation voters are more likely to switch to vote from abstaining as opposed to switch to abstain from voting.

When the cost of voting PDF is not monotonically increasing or decreasing but still satisfies log-concavity, then for initial targeting each politician find it profitable to target two distinct region of voters: a region of voters that is mildly supportive of themselves and another region that is strongly supportive of their opponent. The dynamics of targeting may eventually result in a state like we see with monotone costs of voting where a unique region of voters, on different sides of the median voter, is targeted by each politician, but will never result in more than two distinct regions of voters targeted by each politician. The proof of this finding relies partly on generalizing the insights from the monotone voting case and partly on the utilization of tools from Total Positivity theory.

For robustness, we investigate the impacts of common valence shocks, and small numbers of voters. We find that although these add additional nuance, they do not qualitatively change the nature of our results.

Lastly, we explore the comparative statics relating to different cost of voting distributions and to more/less exogenously informed voters. We find that there is a sense in which higher costs of voting encourage targeting of one's own supporters, and that although in the limit more exogenously informed voters are less profitable to target, the story can be more complicated.

The structure of the paper is as follows. Section 2 presents a formal motivating example to fix ideas and build intuition for what will follow. Section 3 summarizes the related literature within Economics and Political Science. Section 4 outlines the formal model. Section 5 presents the main analysis of the paper, where the benefits of microtargeting are characterized as a weighted measure of the convexity of the politician's valuation of beliefs of voters. This is then used to derive the incentives for micro-targeting when the cost of voting distribution is log-concave. The characterization of the benefits of micro-targeting is relatively simple, meaning that the model presented here can be applied to empirical work. Extensions investigate the extent to which the insights presented carry over to different specifications.

Section 6 presents some comparative statics, finding that higher costs of voting tend to encourage politicians to target groups that are already supportive of them more in a mobilization effort, whereas lower costs of voting tend to encourage targeting of groups that are not supportive on average in a demobilization effort. Section 7 concludes.

2 Motivating Example

Consider an election between two candidates: A and B. Candidate A proposes policy 0 whereas candidate B proposes policy 1. There are two voters in this election—Lucy and Ryan. Lucy's publicly known prior over her most preferred policy x_{Lucy} is represented by the distribution Beta[2, 8]. This distribution (PDF illustrated in green in Figure 1) is single-peaked and implies the expected bliss point 0.2. This means she leans towards the policy platform of A. Ryan's publicly known prior over his most preferred policy x_{Ryan} is Beta[8, 2] (PDF illustrated in red in figure 1), implying his expected bliss point is 0.8. He leans towards the policy platform of B.



Figure 1: Prior over bliss points of Lucy (green) and Ryan (orange).

Voter utility is decreasing in the distance from the bliss point to the policy of the politician elected. That is:

$$U_i(A) = -x_i,$$
$$U_i(B) = -(1 - x_i).$$

We assume costly voting based on utility difference. That is, Lucy and Ryan draw a cost of voting c_i from some distribution f_c with positive support, then decide whether they

want to vote comparing this cost to their perceived utility difference from the candidates. That is, voter i votes for politician A iff:

$$\mathbb{E}[U_i(A) - U_i(B)] > c_i.$$

Similarly, voter i votes for politician B iff:

$$\mathbb{E}[U_i(B) - U_i(A)] > c_i.$$

When indifferent between two or more options, voters randomise equally. We assume that $f_c \sim Beta[5,5]$, so that costs of voting are distributed around 0.5 as illustrated in Figure 2.



Figure 2: Plot of $f_c = Beta(5,5)$

Now suppose that there exists a private binary signal which realises as 1 for voter i with probability x_i , and as 0 otherwise. This signal is informative about the voter's bliss point, and it represents the movement in opinion following some internal reflection prompted by seeing political content. Suppose that before the vote, politician A may pay some price p to send one of these signals to either Lucy, Ryan, or both. Assuming a politician shares the same prior over the distribution of bliss points as Lucy and Ryan do, what strategy maximizes their expected vote share?

Before any signal, Lucy's expected utility difference from politician A is $\mathbb{E}[U_{Lucy}(A)] = -0.2 + (1 - 0.2) = 0.6$. This means that Lucy will not vote for Politician B, and her probability of voting for A is $1 - F_c(0.6) \approx 0.73$, corresponding to the probability that the cost of voting is low enough. Similarly, Ryan will vote for Politician B with probability ≈ 0.73 .

If Lucy observes a private signal, with probability 0.2 it realises as 1 and with probability 0.8 it realises as 0. Using the conjugate prior property of the Beta distribution, we know that the posterior belief for Lucy following a signal realization 1 is Beta[3,8] leading to $\hat{x}_{Lucy} \approx 0.27$, and the belief following the signal realization 0 is Beta[2,9]leading to $\hat{x}_{Lucy} \approx 0.18$. Lucy will still never vote for politician B, but her probability of voting for A falls to ≈ 0.40 with probability 0.2 and rises to ≈ 0.81 with probability 0.8. This means that in expectation, sending the signal to Lucy will decrease the probability she votes for A by $\approx 1\%$. The same calculations reveal that sending the signal to Ryan decreases his probability of voting for B by $\approx 1\%$, so that politician A would like to send the signal to Ryan only.

Note that politician A chooses to target the voter supportive of their opponent. This is because following a signal realisation, a voter may move to the left or to the right, but the effect this has on the election is mediated by the cost of voting distribution. Most distributions mean that expected net votes are non-linear in beliefs, so that logic similar to that applied in Bayesian Persuasion can be applied. In the case of the above example, the relationship between beliefs and expected vote share happens to be sufficiently concave at the space around 0.2 and sufficiently convex at the space around 0.8, meaning that targeting has negative value for Lucy at 0.2 and positive value for Ryan at 0.8.

The example above will be generalized to think about a situation in which instead of Lucy and Ryan, a politician can pick a segment of the population to target. This segment may have a degenerate distribution of expectations (as in the example—Lucy and Ryan each had one expected belief), or it may have a non-degenerate distribution, representing the more general case where different people targeted together have different expectations about their own bliss points. The incentives of a politician to send content to one group rather than another depend crucially on the distribution of costs of voting. We find that given the distribution estimated by Blais et al. (2019), voter demobilization is particularly a tempting strategy. Echoing the example above, sending content to segments of the population that are already likely to be at least somewhat supportive of one's cause are predicted to aid one's opponent.

3 Literature Review

This paper is most closely related to recent literature which has focused on modeling micro-targeted political advertising. To our knowledge, this literature has so far only been concerned with advertising that selectively reveals party platforms to voters, as opposed to our setup which selectively induces learning about voters' own political preferences. Further, our use of imperfectly revealing signals and explicit incorporation of cost of voting distributions allow us to explore effects on turnout, unlike past literature.

Schipper and Woo (2018) assume that politicians have full knowledge about voters' preferences when targeting and can fully and truthfully reveal their platform to individually chosen voters. All information related to party platforms is revealed to voters; as exactly one of the two candidates will find it in their interest to reveal their platform to a given voter.

Titova (2024) allows an entrant politician whose platform is initially unknown to send verifiable but private messages about their own platform to individual voters. Using this strategy, a candidate can win elections that are ex-ante unwinnable via public signals by simultaneously attracting left and right voters. Similarly to Schipper and Woo (2018), the politician perfectly knows each individual voter's bliss point.

Prummer (2020) models media networks explicitly, giving rise to notions of 'media centrality' and fragmentation. Politicians advertise their own policy platforms through specific media outlets which have specific audiences. She finds that media fragmentation, allowing for increased precision of targeting, leads to increased polarization. She distinguishes between attached and unattached voters, where only unattached voters may change their mind. Among unattached voters, the politician targets those whose current belief about the party platform is far from their bliss point. Prummer also allows politicians to change their policy platform without informing unattached voters. This means policy platforms can be catered to targeted voters.

Schultz (2007) finds that the parties would target advertisement spending on those who are most receptive to news, whose uncertainty is largest, and who are relatively uninformed. Vaeth (2024) explores the implications or rational voter learning on issue dimensionality and party positioning.

The theoretical literature on the political economy of mass media is well-developed. David Strömberg and Andrea Prat are key contributors (Strömberg 2001; Stromberg 2004; Prat 2018, and reviews by DellaVigna and Gentzkow 2010; Prat and Strömberg 2013).⁴

We relate to the part of this literature concerning the endogenous provision of political information. Much of this literature focuses on biased news provision. In Mullainathan and Shleifer (2005), the belief confirmation bias of readers leads to media slant towards extreme positions. In Bernhardt, Krasa, and Polborn (2008), due to confirmation bias of voters, competing media firms bias/suppress news to pander to certain voters. This leads to polarization and increased probability of electoral mistakes. In J. Duggan and Martinelli (2011), slanting is an outcome of competition among media, and optimal media slanting depends on whom the media favors.⁵

Our findings also relate to the literature on the question of targeting swing versus

⁴Strömberg (2001, 2004) studies the effects of mass media on policy choice, where more informed groups are offered more favorable policies by politicians.

⁵There is a wider literature on competition among information providers. Endogenous provision of information with partisan providers is usually welfare enhancing (Baron 2006; Gentzkow and Shapiro 2006; Burke 2008; Chan and Suen 2009; Anderson and McLaren 2012; Gentzkow, Shapiro, and Stone 2015). For an exception, see Perego and Yuksel (2022), who find that competition leads to specialization among news providers, and polarized views among voters, resulting in lower welfare. In Chen and Suen (2023) competition among news providers leads to a downward spiral on quality of news as attention is divided. In Matějka and Tabellini (2021) increased access to information and ability to avoid information lead to dispersion in how informed voters are.

core voters (Cox and McCubbins 1986; Lindbeck and Weibull 1987; Dixit and Londregan 1995; Dixit and Londregan 1996). In Glaeser, Ponzetto, and Shapiro (2005), extreme policies are effective in energizing core supporters with low risk of alienating other voters, as party affiliates are relatively more likely to learn about the policy platforms. Hence, the party can increase turnout among core supporters by deviating from the median towards a more extreme platform. Bernhardt, Buisseret, and Hidir (2020) show that core voters can be targeted via policies in a setup where parties' objective is to maximize their vote shares.

This paper also speaks to the literature on campaign contributions and spending, as we implicitly assume. Part of this literature assumes the ability to attract uninformed voters purely using spending, i.e. vote buying (Baron 1994; Dekel, Jackson, and Wolinsky 2008).

Some work considers informative campaigns where politicians consider weight up the implications for voters and for donors (Austen-Smith 1987; Coate 2004b). Coate (2004a) shows that limiting campaign contributions can restrain politicians from implementing extreme policies to appeal to donors.⁶

Empirical work has also investigated the use of targeted messages in elections, with findings that suggest some effect but which do not cohere together to a clear overall view. Liberini et al. (2020) show that micro-targeted ads are effective in influencing turnout. Beknazar-Yuzbashev and Stalinski (2022) found a negative but insignificant effect of micro-targeting on turnout, and no effect on polarization. Kendall, Nannicini, and Trebbi (2015) conduct an experiment to find that campaign messages on valence are effective in boosting votes in a setting where voting is compulsory. Spenkuch and Toniatti (2018) find no relationship between political advertising and turnout, but they do find higher vote shares for candidates with more advertising. This suggests that theorizing in this field is valuable for building the bigger picture.

Finally, this paper relates to work on the political effects of social media. See Zhuravskaya, Petrova, and Enikolopov (2020) and Tucker et al. (2018) for reviews of this field.⁷ This literature is generally concerned with fake news and disinformation as tools of political influence. We see this paper as providing an additional way to view political influence, especially in light of the findings of Nyhan et al. (2020); Barrera et al. (2020); and Guriev et al. (2023), that falsehoods do not seem to be driving political influence online.

⁶There are some papers where campaign contributions act as signals; indirectly informative, where interest groups only care about favors and not ideology. (Potters, Sloof, and Van Winden 1997; Prat 2002b; Prat 2002a).

⁷Also of interest may be work done on the changing nature of news provision(Kavanagh et al. 2019).

4 Framework

We analyze a model in which there are two politicians who are committed to certain policies, and voters can learn further about their own preferences from micro-targeted political ads. Following the Banks and John Duggan (2005) presentation of the Downsian election model, we assume that the politicians' objective is to maximise their expected plurality. Voters draw a random cost of voting, and decide to vote based on utility difference.

We assume that there is a population of voters so that the outcome of the election is deterministic—expected vote share is actual vote share. In section 4.2 we explore the implications of purely office-motivated candidates and non-deterministic elections.

There is a unit mass of voters with bliss points distributed uniformly on the [0, 1] interval, and two politicians labelled politician A and politician B, where A's policy platform is 0 and B's policy platform is 1. The assumption is that the politicians' policy platforms are fixed, while they can choose which segments of voters to send political adverts to. We refer to individual voters by i, and to a position on the political space as x, so that a voter's bliss point is x_i , where:

$$x_i \sim U[0,1]$$

This setup can be interpreted in two ways:

- There is one policy being voted on, which only politician B will enact. The bliss point of voters is equal to the probability that they will benefit from the policy, initially unknown to them.
- Each politician has a whole set of policy profile which we normalise respectively to 0 and 1. The position of voters is the probability that they will benefit from the platform of B more than that of A, given enough regularity.

The politicians' objective is to maximize expected plurality, i.e. the difference between their vote and that of the other party.

We assume that voting is costly in the sense that voters draw a cost of voting from an iid distribution and vote for their preferred candidate if and only if the cost of voting is lower than their expected utility difference from the candidates, which is equivalent to comparing the distance of policies from the voter's estimated bliss point in this setting. We denote the cost of voting by f_c , which will be assumed log-concave with support on [0, 1]. Denote a voter *i*'s cost of voting as c_i . We assume when indifferent voters do not vote.

Therefore, a voter votes for A if:

$$\hat{x}_i - (1 - \hat{x}_i) < c_i,$$

and votes for B if:

$$(1 - \hat{x}_i) - \hat{x}_i < c_i,$$

and a voter abstains otherwise.

4.1 Learning and Segmentation

We assume that the bliss points of voters are initially **not observed by either the politicians or voters**. Instead, voters learn about their bliss point by consuming content and observing their own reaction to it. This is formalised as a series of signals. These signals are Bernoulli trials drawn with a probability of realising as 1 equal to the underlying underlying bliss point x_i . This represents the notion that reaction to content is informative about one's bliss point. Note that we need not assume that these signals are actually informative about a real underlying bliss point, just that voters become increasingly certain about their position and that the distribution of signal-generating processes within a segment is that implied by the specified prior.

The set of realisations for voter i is denoted s_i . We refer to a voter's estimated bliss point as:

$$\hat{x}_i := \mathbb{E}(x_i | s_i)$$

as a function of all observed signals. These binary signals will be split into two groups exogenous signals s_{ex} and endogenous signals s_{en} . All signals are unbiased.

Exogenous signals represent the information that people's observable characteristics contain, such as how old a person is, their occupation, their interests, their gender, etc. These characteristics determine which party a voter has an inclination at the outset by giving an estimate of the distribution of their underlying bliss point. Political actors also observe these so that they can estimate a voter's distribution of bliss points based on their observable characteristics.⁸

This underlies our notion of micro-targeting—we define a set of people sharing a same set of observable characteristics as a *segment*. This maps directly to Facebook's notion of an audience⁹. Politicians and voters correctly estimate the distribution of bliss points within a segment, and voters know which segment they are in. For simplicity, we model

⁸It could be that voters have more characteristics than what the parties can observe, which will be considered as endogenous signals, but this doesn't affect the targeting strategies as long as they are not observable to politicians.

⁹See Facebook, 2024, for the notion of an audience. It is a set of people sharing a set of characteristics, and micro-targeting is carried out by constructing specific audiences to send ads to.

exogenous signals similarly to endogenous signals so that the distribution of bliss points within a segment is a Beta distribution.

We allow fractions of exogenous signals so that a segment can contain voters that realized 3.2 positive signals and 1.5 negative signals. This allows our analysis to touch on bliss point distributions within segments using any combination of mean and variance within the Beta distribution.

Endogenous signals are interpreted as micro-targeted political ads. Unlike exogenous signals, endogenous signals are those whose amount and target audience are chosen by politicians. All voters in a segment receive the same endogenous signals but each voter i will interpret them as a function of their underlying bliss point x_i . For simplicity, we do not allow fractions of endogenous signals, so that each one must be a proper Bernoulli trial. We use the term endogenous signal instead of political ad to emphasise the fact that this can be interpreted as an ad, an item of news, or any other piece of content that is promoted and which prompts reflection on the relevant political topic.

Using Bernoulli trials to represent signals is of course not without loss, as it requires that the posterior of voters given this information is a Beta distribution. Note though that Beta distributions can represent a wide range of single-peaked distributions. The marginal effect of additional political ads can be higher or lower as a result of different priors within segments to begin with, and the belief dispersion within a segment can be achieved by assuming people in a segment have already been subjected to endogenous signals. In particular, we model the characteristics of voters that are not observable to politicians as additional endogenous signals that are realised before politicians decide whether to send more signals. Further, using Bernoulli signals simplifies the analysis and provides a simple and clear interpretation.

There is a population of voters in each segment. We denote the set of all segments as Ψ , so that each segment is denoted as $\psi \in \Psi$. Then, adjusting the preceding notation, we use $s_{ex,\psi}$ to denote exogenous signals and $s_{en,\psi}$ to denote endogenous signals in segment ψ . We also use the following notation to keep track of the number of signals in segments:

$$n_{\psi} := |s_{ex,\psi}|$$

$$u_{\psi} := |s_{en,\psi}|$$

We define α_{ψ} as the number of positive realizations in $s_{x,\psi}$.

Remark 1. The distribution of bliss points in a segment ψ is:

$$x_{\psi} \sim Beta(\alpha_{\psi}, n_{\psi} - \alpha_{\psi}).$$

Denote the distribution of estimated bliss points within a segment by $g_{\psi}(\hat{x})$. When

no endogenous signals have yet been realized in a segment, this will be a degenerate distribution with mass at $\frac{\alpha_{\psi}}{n_{\psi}}$ as exogenous signal realizations in a segment are identical. When endogenous signals have been sent to the segment, given $Beta_z(\cdot)$ denotes the PDF of the Beta distribution at value z, this will be a probability mass function:

$$g_{\psi}(\hat{x}) = \int_{[0,1]} Beta_z(\alpha_{\psi}, n_{\psi} - \alpha_{\psi}) \binom{u_{\psi}}{u\hat{x}} z^{u\hat{x}}(1-z)^{n-u\hat{x}} dz.$$

Unpacking the above, $g_{\psi}(\hat{x})$ gives the proportion of people in the segment with belief \hat{x} . We take every possible bliss point z, working out what proportion of the population has that bliss point $(Beta_z(\cdot))$. Then, for every z, we work out what proportion of voters with that bliss point would end up with belief \hat{x} given the set of exogenous signals common to the segment and the u endogenous signals that have been privately realised. This is given by $\binom{u_{\psi}}{u\hat{x}} z^{u\hat{x}} (1-z)^{n-u\hat{x}}$. Added up, this gives us the proportion of people with belief \hat{x} in segment ψ .

Sometimes it will be useful to refer to segments by the average belief within the segment, which is equivalent to the only belief within the segment if no endogenous signals have yet been realised. One can think of this as the expected belief given observable characteristics or the belief around which the segment is formed. For convenience we will simply say "segment x" or "segment $\psi = x$ " in place of "segment with average belief x".

4.2 Relationship to Bayesian Persuasion

This model clearly has close ties to Bayesian Persuasion models. Restricted to one individual segment, the problem of a politician is to decide whether a particular signal is profitable in exactly the same way that agents decide whether an experiment is profitable in a Bayesian Persuasion game. Indeed, it turns out that the insights from Bayesian Persuasion about the convexity of the value function being integral to the profitability of experiments play a major role in our analysis.

However, there are important differences in the formal structure of the model, and in the interpretation when applied to problems of electoral politics.

Let us now discuss the formal differences. First, in our micro-targeting model, the politician's ability to design a signal is restricted, while we introduce and focus on a dimension not often included in Bayesian Persuasion models: the choice of which type of receivers to target. Recently, Curello and Sinander (2024), contribute in this direction: they compare the optimal signal across different environments, but they do not model the decision of picking which groups to target. Second, our model deals with the problem of persuading a set of people, each of whom is interested in a different unknown variable. The unknown variables are correlated in a very particular way, but we are not aware of any model that allows for this.

These differences mean that the application to electoral politics has a different interpretation. Models that use Bayesian Persuasion to analyze electoral politics mostly assume that politicians can commit to public experiments. This is clearly a strong assumption that limits the possible applications of these models. On the other hand, in our setup, signals are realized privately, corresponding to a notion of self-discovery. This is arguably a more prominent aspect of electoral politics, allowing for a broader application.

5 Analysis

In this section we present the value of sending an endogenous signal to a segment in a form that makes clear its dependence on the convexity of a value function, then we go on to explore what kinds of segments politicians benefit from targeting. We find that in cases that most closely mirror realistic estimates, a demobilization tactic consisting of targeting core voters of an opponent, is profitable. When the cost of voting distribution is log-concave with an interior peak, politicians find it useful to target two distinct and separate kinds of segments.

Before beginning the analysis, we introduce two assumptions:

A 1. Signals are not too informative: $\frac{1}{n+u} < \frac{1}{2}$.

A 2. f_c is log-concave.

Assumption A 1 simply tells us that following the realization of one additional signal, two voters that were previously identical in terms of signal realizations they had observed must end up with less than 0.5 distance between them where 0.5 is the distance between being completely indifferent and being 100% sure about one of the options. Therefore, we believe this to be a very weak assumption.

The assumption that f_c is log-concave is relatively standard. It allows most commonly used distributions such as the normal distribution, the exponential distribution, etc. However, it does rule out more complicated distributions such as distributions with two or more peaks, or distributions with discontinuities of certain kinds. This might be thought to be problematic in the current setup because one could imagine that in some populations the cost of voting is either very high or very low and very rarely anything in between. In cases where this is observed, for example when the high cost of voting stems from lack of access to a car or from having a particular job, then we may consider populations likely to have a high cost of voting and populations likely to have a low cost of voting as different segments, and the analysis below can be applied separately.

5.1 Value of Targeting

An immediate question to ask is what kind of segments are profitable to target in this setting for the politician? Proposition 1 below formulates the value of sending one endogenous signal to a particular segment ψ in a way that sheds some light on this. Specifically, we find that a coarse measure of convexity around the mass of beliefs in a segment is what determines the value of targeting.

Before turning to Proposition 1, it is useful to define a value function, which is the net value to politician A of a voter holding a particular belief \hat{x} just before voting.

Definition 1. A value function for politician A of a voter with belief \hat{x} is defined as:

$$v_A(\hat{x}) := F_c(1 - 2\hat{x}) - F_c(2\hat{x} - 1)$$

which is the probability of a voter with bliss point \hat{x} of voting for politician A minus their probability of voting for politician B, implied by the objective of plurality maximizing. Note that v_A takes the argument \hat{x} rather than x, because voters will vote based on estimated, not actual bliss points. Then, the expected net value of a segment ψ for politician A is given by:

$$\Pi_A(\psi) := \int_0^1 v_A(\hat{x}) g_\psi(\hat{x}) dx.$$

Where we integrate over different bliss points that may exist in a segment, allowing for the possibility that the voters in the segment may already have received some endogenous signals. If a segment has only exogenous signals (i.e. when considering sending the first piece of endogenous news), the above simplifies to:

$$v_A(\hat{x})$$

as all the voters in the segment hold the same belief.

What politicians are interested in is how the above value changes following a microtargeted endogenous signal. Proposition 1 demonstrates this.

Proposition 1. The value of sending one endogenous signal to segment ψ for politician A is given by:

$$\pi_A(\psi) := \int_0^1 \varphi_A(\hat{x}) g_\psi(\hat{x}) d\hat{x},$$

Where

$$\varphi_A(\hat{x}) := (1 - \hat{x})v_A(L) + \hat{x}v_A(R) - v_A(\hat{x})$$

$$L := \hat{x} - \frac{\hat{x}}{n_{\psi} + u_{\psi} + 1} \qquad R := \hat{x} + \frac{1 - \hat{x}}{n_{\psi} + u_{\psi} + 1}$$

 $\varphi_A(\hat{x})$ denotes the change in $v_A(\hat{x})$ following an additional signal, and hence it is a measure of convexity of $v_A(\hat{x})$ around \hat{x} . This implies that the convexity (concavity) of the weighting function is what determines whether targeting value is positive (negative) for a politician, as well as its magnitude. This is reassuringly in line with the insights of Bayesian Persuasion.

Note that the measure of convexity is over an interval. This is because when a voter observes a signal, they may move a little to the right or a little to the left. These will be jumps because the signals are binary. For politician A (left), π_A compares how many net votes are gained through voters moving a little to the left with how many net votes are lost through voters moving a little to the right.

We use Proposition 1 to explore the optimal targeting behaviour of politicians in this set-up. First, we show that a monotonic f_c gives rise to a simple if somewhat counterintuitive behaviour. Then we show that f_c with an interior peak may lead politicians to target two distinct groups on different sides of the political spectrum.

5.1.1 Monotonic cost of voting distributions

At first glance it might not seem intuitive to investigate the implications of increasing or decreasing cost of voting PDFs. Without the mathematical formalism, it is not clear how this would affect the incentives of targeting, and we generally think about central and symmetric distributions in the first instance. However, there is evidence that a monotonically decreasing f_c may be a good approximation of reality (Blais et al. 2019), and within the model presented, a monotonic f_c is useful for generating a sharp result that allows us to say something about equilibrium.

An implication of a strictly increasing (decreasing) f_c is that v_A is always convex (concave) below 0.5 and always concave (convex) above 0.5. This suggests that there is a sense in which an increasing f_c encourages targeting segments with low average \hat{x} , while a decreasing f_c encourages targeting segments with high average \hat{x} . Below, we present an example demonstrating this. Following the example, Proposition 2 demonstrates that this principle can be generalized to all monotonic costs of voting, and for any number of endogenous signals.

Example 2

Recall the motivating example above where we had an election between politicians A and B, and two voters Lucy and Ryan. As before, Lucy's publicly known prior over her most preferred policy is Beta[2, 8], and Ryan's is Beta[8, 2]. Both are illustrated in Figure 3.



Figure 3: Prior over bliss points of Lucy (green) and Ryan (orange).

Deviating from the motivating example, we assume that f_c is monotonically decreasing. To pin ideas we use $f_c = Beta[1,3]$, illustrated in Figure 4.



Figure 4: Plot of $f_c = Beta(1,3)$

From Proposition 1, we have the value of targeting as:

$$\pi_A(\psi) := \int_0^1 \varphi_A(\hat{x}) g_\psi(\hat{x}) d\hat{x},$$

Where

$$\varphi_A(\hat{x}) := (1 - \hat{x})v_A(L) + \hat{x}v_A(R) - v_A(\hat{x}).$$

Since we are only concerned with Lucy, $\pi_A(Lucy)$ simplifies to $\varphi_A(0.2)$.

If Lucy observes an endogenous signal, with probability 0.2 it realizes as 1 and with probability 0.8 it realizes as 0. As before, we know that the posterior belief for Lucy following a signal realization 1 is Beta[3, 8] leading to $\hat{x}_{Lucy} \approx 0.27$, so that R = 0.27. The belief following the signal realization 0 is Beta[2,9] leading to $\hat{x}_{Lucy} \approx 0.18$, so that L = 0.18.

This means that we have:

$$\pi_A(Lucy) = 0.8v_A(0.18) + 0.2v_A(0.27) - v_A(0.2).$$

Now we note the following (immediate from the definition of v_A):

$$v_A''(z) = \begin{cases} f_c' & if \ z < 0.5 \\ -f_c' & if \ z > 0.5 \end{cases}$$

Therefore, we know that v_A is concave below 0.5, so that $\pi_A(Lucy)$ must be negative. By similar arguments, we also know that $\pi_A(Ryan)$ must be positive.

We may be worried that this may not hold when \hat{x} is too close to 0.5 or when the segment contains many different priors, possibly on different sides of 0.5. Proposition 2 assures us that the qualitative nature of the logic is invariant to these changes.

Proposition 2. If the probability density function of the cost of voting is strictly increasing (decreasing), then politician A will only benefit from sending signals to segments $\psi < 0.5$ ($\psi > 0.5$). The inverse is true for politician B.

Corollary 1. In equilibrium:

- If f_c is increasing, then politicians only spend resources targeting segments that are on average supportive.
- If f_c is decreasing, then politicians only spend resources targeting segments that are on average not supportive.

The above cannot be interpreted as a direct prediction of an actual election because there are many other mechanisms at play, but it illustrates the way in which we expect this mechanism to determine behavior by political actors.

This means that in the simplest cases we have a relatively straightforward prediction. Namely, if a higher cost of voting is always more likely than a lower cost of voting, politicians find it profitable to increase the political content consumed by people who are already likely to be supporters. We can view this as a *mobilization* strategy because all of the benefit comes from moving people who are already supporters a little to the left so that they are sufficiently supportive of A over B that they are willing to pay the cost of voting.

The mobilisation strategy also has the effect of always increasing turnout:

Remark 2. When f_c is strictly increasing, turnout always increases in the amount of targeting.

To see this note that the turnout function is:

$$Turnout(\hat{x}) = F_c(1 - 2\hat{x}) + F_c(2\hat{x} - 1).$$

When f_c is increasing, the above is everywhere convex so that by Proposition 1, turnout always increases with targeting.

On the other hand, if f_c is decreasing (meaning a lower cost of voting is more likely), then politicians find it profitable to target those who are likely to be supporters of the opponent. This is because these voters are already likely to vote for the opponent, and giving them more information is more likely to moderate them enough that they switch to abstaining than it is to move them enough in the other direction to mobilise them. This maps to a *demobilisation* strategy.

One might think that the demobilisation strategy will always decrease turnout. This is not true. To see this, consider the case where f_c is strictly decreasing and no endogenous signals have yet been sent. Then, there exists $\varepsilon > 0$ such that the segment $\psi = 0.5 + \varepsilon$ would have higher turnout following targeting. To see this note that it is sufficient for L (the estimated bliss point following a negative realization) to have the property $L < 0.5 - \varepsilon$. This is equivalent to $\frac{n(0.5+\varepsilon)}{n+1} < 0.5 - \varepsilon \Leftrightarrow \varepsilon < 0.5 \frac{1-\frac{n}{n+1}}{1+\frac{n}{n+1}}$. $\varepsilon > 0$ satisfying this clearly exists.

Through this example one can see that for segments sufficiently close to 0.5, the demobilization strategy actually does make people switch the candidate they end up voting for. This is a consequence of the concavity of the positive part of the value function meaning that the gain in terms of mobilizing some supporters never outweighs the loss from pushing some voters in the opposite direction.

The assumptions needed to produce the mobilization and demobilization strategies may seem strong, but there is evidence in Blais et al. (2019) pointing to a cost of voting distribution that is in general monotonically decreasing. This means our model's key prediction is that voter demobilization should be a key tactic in elections.

5.1.2 Log-concave cost of voting distributions

If we assume that the PDF of the cost of voting is neither monotonically increasing nor decreasing, but still satisfies log-concavity, then we see that the politicians may value two disjoint sets of voters as targets. These two sets will be voters who are moderately supportive of themselves and voters who are strongly supportive of the opponent.

Proposition 3. Given a strengthened version of A 1:

A' 1. Signals are not too informative: $\frac{1}{n+1} < \frac{mode(f_c)}{2}$.¹⁰

¹⁰This is a strengthening of A 1 needed for technical reasons. In general we do not expect this assumption to be problematic because the informational value of one signal is likely to be small. Further,

Then for some $0 < \alpha_1 < 0.5 < \alpha_2 < 1$ *:*

$$\varphi_A(\hat{x}) \begin{cases} \leq 0 & if \ \hat{x} \in (0, \alpha_1) \cup (0.5, \alpha_2) \\ \geq 0 & if \ \hat{x} \in (\alpha_1, 0.5) \cup (\alpha_2, 1) \\ 0 & Otherwise \end{cases}$$

This means that for the first endogenous signal:

- It is profitable for politician A to target segments which are mildly supportive of A and segments which are very supportive of B.
- It is harmful for politician A to target segments which are very supportive of A and segments which are mildly supportive of B.

See Figure 5 for an illustration of $\varphi_A(\hat{x})$.



Value of targeting by segment

Figure 5

Proposition 3 applies only to the first endogenous signal. However, the corollary below indicates that this is the most complicated form π_A can take for any additional signals.

in many cases where this assumption would be violated, it would be reasonable to approximate f_c as a monotone function, in which case Proposition 2 can be used to discern the shape of π_A .

This means that either π_A is positive-negative or negative-positive (as in Proposition 2), or it is the shape expounded on in Proposition 3 and illustrated in Figure 3.

Corollary 2. One of the following is true for the value of additional endogenous signals following the first one:

- 1. π_A takes the shape described in Proposition 3.
- 2. π_A changes sign once at $\psi = 0.5$.
- 3. π_A is zero everywhere.

5.1.3 Mandatory Voting

We can extend our analysis to the case of mandatory voting. In this case we have a modified value function:

$$v_A(\hat{x}) = \begin{cases} 1 & if \, \hat{x} < 0.5 \\ -1 & if \, \hat{x} > 0.5 \\ 0 & if \, \hat{x} = 0.5 \end{cases}$$

This represents voters picking whichever side looks better however close to the center they may be. For voters at exactly 0.5, we are agnostic about whether they randomize or spoil their ballot. In this case we immediately see that when we do not have to worry about a distribution of beliefs within segments, politicians find it most profitable to target segments that are supportive of their opponent, but as close to indifferent as possible.

Remark 3. With no endogenous signals:

When voting is mandatory, politician A finds it most profitable to target the segment with $\psi > 0.5$ that is closest to the center.

This is due to the fact that targeting a segment disperses the views within that segment. If that segment already supports a politician, then dispersing views can only serve to turn some voters to the other side. If, however, they disperse the views of those that are supportive of their opponent, then they can only turn voters to their own side. Picking voters closest to the center maximisers the chance that voters will change their mind about which politician they prefer.

We can also see that mandatory voting is a limit case of a decreasing f_c . Everybody has a cost of voting equal to 0. Indeed, in this limit, the insights from Proposition 2 still hold:

Remark 4. If voting is mandatory, then politician A will only benefit from sending signals to segments $\psi > 0.5$. The inverse is true for politician B.

The proof for this is a slight modification to the proof used for Proposition 2. This implies that in cases where voting is mandatory, negative campaigning that targets moderates is predicted to be especially effective.

5.2 Office-Motivated Candidates and Non-Deterministic Elections

The above analysis assumed that the objective of politicians is plurality maximization in a deterministic election setting. This allowed us to isolate the mechanism we are interested in, but some readers may be concerned that excluding stochastic election outcomes and office-motivated candidates may have distorted the analysis. To address these concerns, we present below the following:

- A comparison of purely office-motivated (maximizing winning probability) and plurality-maximizing (maximizing vote difference) politicians in the baseline setup.
- A setup where voters experience a common valence shock after targeting but before voting, making the outcome of the election stochastic. We show that our results from above still hold in this setup for the plurality-maximizing candidate.
- A setup where each segment is small enough, making the outcome of the election stochastic. We show that for plurality-maximizing candidates, this setup is equivalent to our benchmark. Further, we show that under some restrictions the optimal strategy of a purely office-motivated candidate is the same as the optimal strategy of a plurality-maximizing candidate, and discuss what the difference in strategies is when there is one.

5.2.1 Office-Motivated Candidates in Original Setting

Since the election in our original setting is deterministic, comparing the strategies of office-motivated (maximizing winning probability) and plurality-maximizing candidates is relatively straightforward. If there is no opportunity cost to using up a whole budget, then the optimal strategy set of a plurality-maximizer is a subset of the optimal strategy set of the office-motivated candidate, as there are many ways to ensure winning the election while there is only one way to maximize plurality. If, however, there is an opportunity cost to use up the budget, then an office motivated candidate will refrain from targeting in order to further increase votes once wining is guaranteed, and they will refrain from any targeting altogether when winning cannot be guaranteed.

5.2.2 Common Valence Shock

A common valence shock is such that that after observing all exogenous and endogenous signals, but before voting, voters experience a common opinion shock. This means that all voters move a little to the left or a little to the right, introducing randomness to their voting strategy and hence to the outcome of an election.

In this section we demonstrate that although the introduction of a common valence shock may make the results of the paper less sharp, it does not qualitatively change them.

Formally, a common valence shock is a random variable $\epsilon \sim f_{\epsilon}$, which equates to a shift in the expected bliss point of each voter (observed by the voter). Politician A's valuation of a voter with belief \hat{x} is hence transformed to $v_A(x + \epsilon)$ after the valence shock.

We assume that the distribution of beliefs before the common valence shock is realized and the support of the common valence shock are such that $x + \epsilon \in [0, 1]$. This is to avoid having to extend the definition of v_A outside of [0, 1]. We also assume that the common valence shock is part of the exponential family and symmetric around 0.

We can now present a technical lemma that will be used to show how a common valence shock does not significantly alter our results.

Lemma 1. The number of sign changes of $\pi_A(\psi)$ does not increase when a common valence shock is introduced.

Proposition 4. If the probability density function of the cost of voting is strictly increasing (decreasing), then a politician will send signals only to segments that on average support (oppose) them even if there is a common valence shock before the election.

It is not necessarily true that the shape characterised for the first endogenous signal in Proposition 3 is maintained when a common valence shock is introduced. This is for the same reason that we cannot guarantee that the shape is maintained after more endogenous signals are sent: the distribution of views within segments becomes too diffuse to guarantee that enough movements within each phase occur for a particular segment. From Lemma 1 there also follows a modified version of Corollary 2:

Corollary 2'. When there is a common valence shock, one of the following is true:

- 1. π_A takes the shape described in Proposition 3.
- 2. π_A changes sign once at $\psi = 0.5$.
- 3. π_A is zero everywhere.

5.2.3 Small Segments

In this section we assume that each segment has only one voter, so that the outcome of the election is stochastic. In this setup, the strategies that maximize plurality coincide with the strategies that maximize plurality in the deterministic benchmark. However, the strategies that maximize the probability of winning may be different. Below we find that the difference, when there is any, comes down to the ways in which different segments are likely to be pivotal.

Denote by $M_w(\neg\psi)$ the probability of winning the election by one vote if segment ψ doesn't participate, by $M_l(\neg\psi)$ the probability of losing the election by one vote if segment ψ doesn't participate, and by $M_d(\neg\psi)$ the probability of a draw if segment ψ doesn't participate. These are the three ways in which segment ψ can be pivotal, because they are the cases where this segment (which consists of one voter) can change the outcome of the election. We say we have *pivot symmetry* if $M_w(\neg\psi) = M_l(\neg\psi) \ \forall \psi \in \Psi$. We say we have *pivot equality* if for $\psi \neq \psi'$, $M_w(\neg\psi) = M_w(\neg\psi')$ and $M_l(\neg\psi) = M_l(\neg\psi') \ \forall \psi, \psi' \in \Psi$.

Remark 5. If we have pivot symmetry and pivot equality, the optimal strategies of plurality-maximizing and purely office-motivated (winning maximizing) candidates are identical.

Remark 3 assures us that the difference between office-motivated and plurality-maximizing candidates comes down to office-motivated candidates paying attention to the probability and way in which a segment is pivotal. In some settings it might be reasonable to assume that the probability of different segments being pivotal is roughly equivalent. An example of such a setting would be a national referendum or targeting different groups within states for presidential elections, in which there are a lot of segments due to the diverse population but there is no concern about local constituencies.

However, in case segments are in different constituencies in a First Past the Post system (or other systems with local representatives for national elections such as the Electoral College in USA), then segments in swing constituencies are clearly more likely to be pivotal. In these cases, it makes sense to simply control for how likely the segment is to be pivotal, as shown in the following corollary.

Corollary 3. If we have pivot symmetry but not pivot equality, then all else being equal, a segment with higher probability of being pivotal will have a higher absolute value of being targeted.

If we do not have pivot symmetry (so that the probability of winning by one vote without segment ψ is not the same as losing by one vote without segment ψ), then it is no longer the case that increasing the probability of ψ voting for A is worth as much as decreasing the probability that ψ votes for B. This is because for a politician that only cares about winning the election, the probability of a voter voting for A is only relevant in cases where the politician would not have won if the voter abstained. This means that negative demobilization is more profitable in cases where M_w is higher, and positive mobilization is more profitable in cases where M_l is higher.

6 Comparative Statics

In this section we explore how targeting incentives change following some changes in the parameters of our model.

First we look at how incentives change following a change in the cost of voting distribution, and find a sense in which higher costs of voting encourage the targeting of segments which are on average supportive of a party. We then provide some interpretation of why this is and present some examples where higher costs of voting also lead to higher incentives to target more extreme segments.

Next we explore how targeting incentives change when the number of exogenous signals already received by segments increases. That is, when more information is available through observable voter characteristics of the segments. Here we make use of some insights from Bayesian Persuasion to help our understanding.

6.1 Different costs of voting

How do the incentives for targeting change if people are more likely to have higher costs of voting? Intuitively, we can say that higher costs of voting means that the mass of cost of voting distribution f_c leans more towards 1 and less towards 0. In a single-peaked distribution this means that more of the convexity and concavity of $v_A(x)$ moves towards more extreme voters, so that the value of targeting (positive and negative) should become more concentrated on these. This is due to the fact that π_A is an integral over a measure of convexity. This logic serves us well in the following example:

Example 3:

Suppose f_c is as defined below:

$$f_c(z) = \begin{cases} 0 & if \ z < b \\ \frac{2(z-b)}{(1-b)^2} & if \ z \in [b,1] \end{cases}$$

where b > 0.5. This cost of voting distribution is illustrated in Figure 4.



 f_c is a linear increasing PDF starting at b, meaning that higher b corresponds to higher costs of voting. Suppose also that u = 0 (no endogenous news yet been sent) and $n \ge \frac{1}{1-b}$ (exogenous signals).

In this case, the only segments that can have $\pi(\psi) > 0$ satisfy $\frac{n\hat{x}}{n+1} < \frac{b}{2}$. This is because by Proposition 2 only segments with $\hat{x} \leq 0.5$ can have $\pi(\psi) > 0$ for a weakly increasing f_c and further, because for segments satisfying $\frac{n\hat{x}}{n+1} \geq \frac{b}{2}$ and $x \leq 0.5$, no voter can be convinced to vote as their costs of voting will be higher than whatever utility difference they end up with following one additional signal. This logic works in the exact same way for segments with $\pi(\psi) < 0$, so that the set of segments that some politician wants to target become more concentrated around more extreme segments as b rises.

Further, the maximizer of π must be close to $\frac{b}{2}$ while the minimizer must be close to $1 - \frac{b}{2}$, so that increasing the cost of voting pushes the most profitable segments towards the edges. To see this note the following: first, $v_a(\hat{x})$ is linear or concave on $\hat{x} \in [\frac{b}{2}, 1]$, so that a segment centered around $\hat{x} > \frac{b}{2}(1 - \frac{1}{n+1})^{-1}$ cannot have $\pi(\hat{x}) > 0$ (because the posterior following one more negative signal in these segments is above $\frac{b}{2}$). In fact this upper bound is tighter because $\pi(\frac{b}{2}) > \pi(\omega)$ for $\omega > \frac{b}{2}$. This is because the people who observe a positive realisation in segments $\hat{x} \in [\frac{b}{2}, \frac{b}{2}(1 - \frac{1}{n+1})^{-1}]$ will not vote. Within this interval, choosing a segment that is closer to $\frac{b}{2}$ increases the amount of people within that segment who will observe a negative signal and thus move left, and it also increases the gain from moving these people to the left because they will move to a more extreme opinion.

We also know that the maximiser will not be a segment that is too close to $\hat{x} = 0$. This is because on $\hat{x} \in [0, \frac{b}{2} - \frac{1-\hat{x}}{n+1}]$, we have:

$$v_A(z) = \int \frac{2(z-b)}{(1-b)^2} = \frac{z^2 - 2zb}{(1-b)^2}.$$

Plugging this into φ (defined in Proposition 1) , taking a derivative, and simplifying, we find:

$$\frac{\partial \varphi(\hat{x})}{\partial \hat{x}} = \frac{4 - 8\hat{x}}{(n+1)^2(1-b)^2}.$$

This is positive for $\hat{x} < 0.5$, so that the maximiser \hat{x}^+ lies in $\left[\frac{b}{2} - \frac{1-\hat{x}}{n+1}, \frac{b}{2}\right]$. A similar procedure can be used to find that the minimiser \hat{x}^- lies in $\left[\left(1-\frac{b}{2}\right), \left(1-\frac{b}{2}\right) + \frac{1-\hat{x}}{n+1}\right]$.

To formalize the notions explored above, we need to formally define what is meant by people being more likely to have a higher cost of voting. In the literature, the most common notions that are used are first order stochastic dominance and dominance on the likelihood ratio order. Indeed, the likelihood ratio order is useful for presenting a result that gestures towards the intuition developed in the example above:

Proposition 5. As the cost of voting distribution increases in the likelihood ratio order, the area below 0.5 for which φ (targeting value for party A) is positive weakly increases while the area above 0.5 for which φ is positive weakly decreases.

This result illustrates a sense in which higher costs of voting increase the incentives of politicians to target groups that support them. The interval of supportive segments that you would like to target weakly grows as costs of voting increase, while the interval of non-supportive segments that you would like to target weakly shrinks with growing costs of voting. This is a reflection of the fact that convex parts of F_c indicate increasing f_c while concave parts indicate decreasing f_c , so that increasing costs of voting, which has the effect of shifting the peak of f_c towards 1, creates a bigger convex area in F_c . The proof directly uses the sense in which mass moves rightwards in the likelihood ratio order.

The inverse here is also true—lower costs of voting in the log-likelihood order encourage the targeting of groups supportive of the opponent. This means that lower costs of voting are may encourage more negative campaigning. This builds on the intuition built in Proposition 2 that decreasing f_c (indicating generally low costs of voting) produces an incentive for only targeting opposing groups.

6.2 More exogenous signals

We have defined above our notion of exogenous signals, relating to the already fixed observable characteristics of voters. When we speak about a segment with more exogenous signals, we are referring to a segment where the average belief is the same, but the number of exogenous signals is higher. This is the same as thinking of it as a segment where the observable characteristics are more informative, though they lead to the same expected bliss point.

The effect of more exogenous signals then is to decrease the shift in belief caused by additional endogenous signals. One might think that this would have the straightforward effect of reducing the targeting value, π . Indeed, when F_c is continuous, the value of targeting vanishes as the number of exogenous signals becomes high. However, signals inducing smaller movement can actually increase the targeting value π if previously the endogenous signal induced *too much* movement, compared to the optimal signal.

The reason for this is linked to the same logic as in the concavification approach in Bayesian persuasion—signals are useful as long as they move beliefs within an interval where the value function is convex, but further movement into a concave region is harmful. In the context of this model, the concern with moving into concave sections is particularly salient because politicians do not pick how informative their signals are, and they may end up targeting segments in which there are voters both among convex and concave parts of the value function.

In general, it is not clear whether a segment with more exogenous news is more profitable to target than a segment with less, keeping the average belief in the two segments same. We next demonstrate in example 4 a case of a segment where increasing exogenous signals unambiguously decrease the value of targeting and a segment where increased exogenous signals can at first increase the value of targeting.

Example 4:

In this example we will assume u = 0 and we will explore the implications of increasing n with the following cost of voting distribution:

$$f_c = Beta(10,2)(z)$$

Where Beta(10, 2) corresponds to the Beta distribution PDF with shape parameters 10 and 2. This distribution is illustrated in Figure 3.



Figure 7: Plot of $f_c(z) = Beta(10, 2)(z)$

Note that $mode(f_c) = 0.1$, meaning that $v_A(\hat{x})$ is concave on $\hat{x} \in [0, 0.45] \cup [0.5, 0.55]$ and convex on $\hat{x} \in [0.45, 0.5] \cup [0.55, 1]$. This means that if we take a segment which satisfies $\frac{n\hat{x}}{n+1} > 0.55$, $\varphi(\hat{x})$ is a measure of convexity over a function that is only convex. Therefore, the value of targeting this segment can only fall as n increases.

Remark 6. If $v_A(\hat{x})$ is convex or concave on $[\frac{n\hat{x}}{n+1}, \frac{n\hat{x}+1}{n+1}]$, then $|\varphi(\hat{x})|$ weakly decreases in n.

If, on the other hand, we take a segment satisfying $\frac{n\hat{x}}{n+1} < 0.45$, $\hat{x} \in (0.45, 0.5)$, $\frac{n\hat{x}+1}{n+1} \in (0.5, 0.55)$, then a marginal increase in n increases $\varphi(\hat{x})$. Once n is large enough (that is, once $\frac{n\hat{x}}{n+1} \in [0.45, \hat{x})$ and $\frac{n\hat{x}+1}{n+1} \in (\hat{x}, 0.5]$), then any increase in n decreases φ . There exists a segment satisfying the above centered in $\hat{x} \in (0.45, 0.5)$ if n starts sufficiently small. This illustrates the sense in which increasing n can have non-monotonic effects on φ —it can first be useful for avoiding sending a signal that is too informative, before it makes the signal uninformative.

Though a characterization of when exactly more informative signals are useful is not presented here, simulations can help to understand the trade-offs in specific applications with an estimated cost of voting distribution. For a deeper theoretical understanding of when exactly more informative signals are useful in a setting where comparison between segments is not central, the reader is directed to Curello and Sinander (2024).

7 Conclusion

We have developed a model of online political influence that relies on prompting selfreflection. In line with evidence from recent research, our model does not rely on misleading voters via fake or sensationalized content. We have found that politicians benefit from selectively targeting specific groups based on their observable characteristics, which maps exactly to the way in which Facebook allows advertisers to pick their 'audiences'.

Through a relatively simple mechanism, we have developed an economic argument for negative campaigning. When the cost of voting function, f_c is decreasing, prompting reflection among parts of the population that support one's opponent is profitable with an aim to demobilize them. This is somewhat worrying given Blais et al. (2019)'s empirical finding that f_c is indeed decreasing. However, the mechanism explored implies that every targeted individual estimates their own bliss point better than they would have otherwise, so that the outcome of this type of negative campaigning is not straightforwardly problematic.

This leads us to think that a positive effect from this kind of microtargeting is possible, but not necessarily likely. In an ideal world, the increased reflection stemming from micro-targeting would lead voters who learn that they are not satisfied with the choices presented to support and pursue better political projects. However, when the barriers to entry are high as they are in countries like the United States, the more likely outcome may well be disengagement. Even when there is an alternative for voters that have been demobilized by microtargeting tactics as in many countries in Europe, it is not clear that this would lead to improved political choices because it is not clear we can say much about the sort of political thinking that would be adopted after this kind of disaffection.

We have further explored the incentives for micro-targeting, finding that log-concave cost of voting distributions with an interior peak may cause politicians to target a nonconvex set of segments. In this case, both segments that are mildly supportive of the politician and those very supportive of their opponent may simultaneously be profitable to target. Furthermore, we found a sense in which higher costs of voting incentivizes more targeting of supportive segments and less targeting of opposing segments.

Of course, as with other kinds of political influence, we should worry when its use is controlled by a wealthy elite. Following the initial increase in interest in political influence on social media in 2015/16, tech firms have taken steps to signal their commitment to democratically-oriented moderation. The credibility of these commitments is in question, most clearly in the case of X (formerly Twitter) with its owner Elon Musk aligning himself firmly with Donald Trump in the run up to and follow up to the 2024 US presidential election. Musk has also sought to align himself with the German, Italian, and British far right (Nöstlinger 2024). A commitment to democratically-oriented moderation is also in question with regards to another US-based tech giant, Meta (Bradshaw and Craggs Mersinoglu 2025).

Micro-targeted political content does not need to be harmful for democracy. It can allow for smaller actors to reach specific audiences, so that a broader range of movements can gather momentum, although there is reasonable worry that the type of smaller movements that would mostly benefit from this are authoritarian and anti-democratic.

This highlights the need for effective regulation. We are not in a position to recommend specific regulatory measures, but we hope that the new tool for thinking about political influence online we have developed might help in this endeavor.

There are several avenues for further research. Theoretically, the analysis above can be extended to provide a fuller characterization of the comparative statics relating to the magnitude of the benefit of targeting certain groups, rather than just the sign, so that we can say something about the groups that a politician would most like to target. The analysis can also be extended to include several policy dimensions, which may be correlated.

Empirically, the model can be tested in several ways. Most immediately, an experiment can be run to investigate whether the mechanism works as expected. Analysis can also be performed to investigate whether past political campaigns have targeted the groups that the model predicts they should, and whether campaigns that target the groups that the model predicts they should end up doing better electorally than those who don't.

Though we have used the interpretation of microtargeting on social media, the logic of this model can be applied to any setting in which the audience of a particular platform cam be assumed to have single-peaked beliefs, arrived at through consuming a similar amount of content. For example, one might think of a segment as an audience of a particular newspaper, or the people likely to see a particular billboard. The most natural setting in which to think about microtargeting is social media because of the explicit ways in which platforms allow advertisers to construct audiences, but all forms of communication have specific audiences.

Bibliography

Amer, Karim and Jehane Noujaim (2019). The Great Hack.

- Anderson, Simon P. and John McLaren (2012). "Media Megers and Media Bias with Rational Consumers". In: Journal of the European Economic Association 10.4, pp. 831–859. ISSN: 15424766. DOI: 10.1111/j.1542-4774.2012.01069.x.
- Austen-Smith, David (1987). "Interest groups, campaign contributions, and probabilistic voting". In: *Public Choice* 54.2, pp. 123–139. ISSN: 0048-5829, 1573-7101. DOI: 10. 1007/BF00123002.
- Banks, Jeffrey S. and John Duggan (2005). "Probabilistic Voting in the Spatial Model of Elections: The Theory of Office-motivated Candidates". In: Social Choice and Strategic Decisions. Ed. by David Austen-Smith and John Duggan. Red. by M. Salles, P. K. Pattanaik, and K. Suzumura. Series Title: Studies in Choice and Welfare. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 15–56. ISBN: 978-3-540-22053-4 978-3-540-27295-3. DOI: 10.1007/3-540-27295-X_2.
- Baron, David P. (1994). "Electoral Competition with Informed and Uninformed Voters".
 In: American Political Science Review 88.1, pp. 33–47. ISSN: 0003-0554, 1537-5943.
 DOI: 10.2307/2944880.
- (2006). "Persistent media bias". In: Journal of Public Economics 90.1, pp. 1–36. ISSN: 00472727. DOI: 10.1016/j.jpubeco.2004.10.006.
- Barrera, Oscar et al. (2020). "Facts, alternative facts, and fact checking in times of posttruth politics". In: *Journal of Public Economics* 182, p. 104123. ISSN: 00472727. DOI: 10.1016/j.jpubeco.2019.104123.
- Beknazar-Yuzbashev, George and Mateusz Stalinski (2022). "Do social media ads matter for political behavior? A field experiment". In: *Journal of Public Economics* 214, p. 104735. ISSN: 00472727. DOI: 10.1016/j.jpubeco.2022.104735.
- Bernhardt, Dan, Peter Buisseret, and Sinem Hidir (2020). "The Race to the Base". In: American Economic Review 110.3, pp. 922–942. ISSN: 0002-8282. DOI: 10.1257/aer. 20181606.
- Bernhardt, Dan, Stefan Krasa, and Mattias Polborn (2008). "Political polarization and the electoral effects of media bias". In: *Journal of Public Economics* 92.5, pp. 1092– 1104. ISSN: 00472727. DOI: 10.1016/j.jpubeco.2008.01.006.
- Blais, André et al. (2019). "What is the cost of voting?" In: *Electoral Studies* 59, pp. 145–157. ISSN: 02613794. DOI: 10.1016/j.electstud.2019.02.011.
- Bradshaw, Tim and Yasemin Craggs Mersinoglu (2025). "Meta ends third-party factchecking programme". In: *Financial Times*.
- Burke, Jeremy (2008). "Primetime Spin: Media Bias and Belief Confirming Information".
 In: Journal of Economics & Management Strategy 17.3, pp. 633-665. ISSN: 1058-6407, 1530-9134. DOI: 10.1111/j.1530-9134.2008.00189.x.

- Chan, Jimmy and Wing Suen (2009). "Media as watchdogs: The role of news media in electoral competition". In: *European Economic Review* 53.7, pp. 799–814. ISSN: 00142921. DOI: 10.1016/j.euroecorev.2009.01.002.
- Chen, Heng and Wing Suen (2023). "Competition for Attention and News Quality". In: American Economic Journal: Microeconomics 15.3, pp. 1–32. ISSN: 1945-7669, 1945-7685. DOI: 10.1257/mic.20210259.
- Coate, Stephen (2004a). "Pareto-Improving Campaign Finance Policy". In: *American Economic Review* 94.3, pp. 628–655. ISSN: 0002-8282. DOI: 10.1257/0002828041464443.
- (2004b). "Political Competition with Campaign Contributions and Informative Advertising". In: Journal of the European Economic Association 2.5, pp. 772–804. ISSN: 1542-4766, 1542-4774. DOI: 10.1162/1542476042782251.
- Cox, Gary W. and Mathew D. McCubbins (1986). "Electoral Politics as a Redistributive Game". In: *The Journal of Politics* 48.2, pp. 370–389. ISSN: 0022-3816, 1468-2508.
 DOI: 10.2307/2131098.
- Curello, Gregorio and Ludvig Sinander (2024). The comparative statics of persuasion. Version Number: 6. DOI: 10.48550/ARXIV.2204.07474.
- Dekel, Eddie, Matthew O. Jackson, and Asher Wolinsky (2008). "Vote Buying: General Elections". In: Journal of Political Economy 116.2, pp. 351–380. ISSN: 0022-3808, 1537-534X. DOI: 10.1086/587624.
- DellaVigna, Stefano and Matthew Gentzkow (2010). "Persuasion: Empirical Evidence".
 In: Annual Review of Economics 2.1, pp. 643-669. ISSN: 1941-1383, 1941-1391. DOI: 10.1146/annurev.economics.102308.124309.
- Dixit, Avinash and John Londregan (1995). "Redistributive Politics and Economic Efficiency". In: American Political Science Review 89.4, pp. 856–866. ISSN: 0003-0554, 1537-5943. DOI: 10.2307/2082513.
- (1996). "The Determinants of Success of Special Interests in Redistributive Politics". In: *The Journal of Politics* 58.4, pp. 1132–1155. ISSN: 0022-3816, 1468-2508. DOI: 10.2307/2960152.
- Duggan, J. and C. Martinelli (2011). "A Spatial Theory of Media Slant and Voter Choice".
 In: The Review of Economic Studies 78.2, pp. 640–666. ISSN: 0034-6527, 1467-937X.
 DOI: 10.1093/restud/rdq009.
- Gentzkow, Matthew and Jesse M. Shapiro (2006). "Media Bias and Reputation". In: Journal of Political Economy 114.2, pp. 280–316. ISSN: 0022-3808, 1537-534X. DOI: 10.1086/499414.
- Gentzkow, Matthew, Jesse M. Shapiro, and Daniel F. Stone (2015). "Media Bias in the Marketplace". In: *Handbook of Media Economics*. Vol. 1. Elsevier, pp. 623–645. ISBN: 978-0-444-63691-1. DOI: 10.1016/B978-0-444-63685-0.00014-0.
- Glaeser, Edward, Giacomo Ponzetto, and Jesse M. Shapiro (2005). "Strategic Extremism: Why Republicans and Democrats Divide on Religious Values". In: *The Quarterly*

Journal of Economics 120.4, pp. 1283–1330. ISSN: 0033-5533, 1531-4650. DOI: 10. 1162/003355305775097533.

- Guriev, Sergei et al. (2023). "Curtailing False News, Amplifying Truth". In: SSRN Electronic Journal. ISSN: 1556-5068. DOI: 10.2139/ssrn.4616553.
- Harrington Jr., Joseph E. and Gregory D. Hess (1996). "A Spatial Theory of Positive and Negative Campaigning". In: *Games and Economic Behavior* 17.2, pp. 209–229. ISSN: 08998256. DOI: 10.1006/game.1996.0103.
- Jewitt, Ian (1987). "Risk Aversion and the Choice Between Risky Prospects: The Preservation of Comparative Statics Results". In: *The Review of Economic Studies* 54.1, p. 73. ISSN: 00346527. DOI: 10.2307/2297447.
- Kang, Cecilia (2018). Facebook Says Cambridge Analytica Harvested Data of Up to 87 Million Users. New York Times.
- Karlin, Samuel (1968). Total Positivity. Stanford University Press.
- Kavanagh, Jennifer et al. (2019). News in a Digital Age. RAND.
- Kendall, Chad, Tommaso Nannicini, and Francesco Trebbi (2015). "How Do Voters Respond to Information? Evidence from a Randomized Campaign". In: American Economic Review 105.1, pp. 322–353. ISSN: 0002-8282. DOI: 10.1257/aer.20131063.
- Liberini, Federica et al. (2020). "Politics in the Facebook Era Evidence from the 2016 US Presidential Elections". In: SSRN Electronic Journal. ISSN: 1556-5068. DOI: 10. 2139/ssrn.3584086.
- Lindbeck, Assar and Jrgen W. Weibull (1987). "Balanced-budget redistribution as the outcome of political competition". In: *Public Choice* 52.3, pp. 273–297. ISSN: 0048-5829, 1573-7101. DOI: 10.1007/BF00116710.
- Matějka, Filip and Guido Tabellini (2021). "Electoral Competition with Rationally Inattentive Voters". In: *Journal of the European Economic Association* 19.3, pp. 1899– 1935. ISSN: 1542-4766, 1542-4774. DOI: 10.1093/jeea/jvaa042.
- Mullainathan, Sendhil and Andrei Shleifer (2005). "The Market for News". In: *American Economic Review* 95.4, pp. 1031–1053. ISSN: 0002-8282. DOI: 10.1257/0002828054825619.
- Nöstlinger, Nette (2024). "Backlash builds as Elon Musk endorses Germany's far right". In: *Politico*.
- Nyhan, Brendan et al. (2020). "Taking Fact-Checks Literally But Not Seriously? The Effects of Journalistic Fact-Checking on Factual Beliefs and Candidate Favorability". In: *Political Behavior* 42.3, pp. 939–960. ISSN: 0190-9320, 1573-6687. DOI: 10.1007/s11109-019-09528-x.
- Perego, Jacopo and Sevgi Yuksel (2022). "Media Competition and Social Disagreement". In: *Econometrica* 90.1, pp. 223–265. ISSN: 0012-9682. DOI: 10.3982/ECTA16417.
- Potters, Jan, Randolph Sloof, and Frans Van Winden (1997). "Campaign expenditures, contributions and direct endorsements: The strategic use of information and money to

influence voter behavior". In: *European Journal of Political Economy* 13.1, pp. 1–31. ISSN: 01762680. DOI: 10.1016/S0176-2680(96)00032-8.

- Prat, Andrea (2002a). "Campaign Advertising and Voter Welfare". In: *Review of Economic Studies* 69.4, pp. 999–1017. ISSN: 0034-6527, 1467-937X. DOI: 10.1111/1467-937X.00234.
- (2002b). "Campaign Spending with Office-Seeking Politicians, Rational Voters, and Multiple Lobbies". In: *Journal of Economic Theory* 103.1, pp. 162–189. ISSN: 00220531.
 DOI: 10.1006/jeth.2001.2793.
- (2018). "Media Power". In: Journal of Political Economy 126.4, pp. 1747–1783. ISSN: 0022-3808, 1537-534X. DOI: 10.1086/698107.
- Prat, Andrea and David Strömberg (2013). "The Political Economy of Mass Media". In: Advances in Economics and Econometrics. Ed. by Daron Acemoglu, Manuel Arellano, and Eddie Dekel. 1st ed. Cambridge University Press, pp. 135–187. ISBN: 978-1-107-01605-7 978-1-107-67416-5 978-1-139-06002-8. DOI: 10.1017/CB09781139060028.004.
- Prummer, Anja (2020). "Micro-targeting and polarization". In: Journal of Public Economics 188, p. 104210. ISSN: 00472727. DOI: 10.1016/j.jpubeco.2020.104210.
- Schipper, Burkhard C. and Hee Yeul Woo (2018). "Political Awareness and Microtargeting of Voters in Electoral Competition". In: SSRN Electronic Journal. ISSN: 1556-5068. DOI: 10.2139/ssrn.2039122.
- Schultz, Christian (2007). "Strategic Campaigns and Redistributive Politics". In: *The Economic Journal* 117.522, pp. 936–963. ISSN: 0013-0133, 1468-0297. DOI: 10.1111/j.1468-0297.2007.02073.x.
- Shaked, Moshe and J. George Shanthikumar (2007). *Stochastic orders*. Springer series in statistics. New York: Springer. ISBN: 978-0-387-34675-5.
- Skaperdas, Stergios and Bernard Grofman (1995). "Modeling Negative Campaigning". In: American Political Science Review 89.1, pp. 49–61. ISSN: 0003-0554, 1537-5943. DOI: 10.2307/2083074.
- Spenkuch, Jörg L and David Toniatti (2018). "Political Advertising and Election Results*". In: The Quarterly Journal of Economics 133.4, pp. 1981–2036. ISSN: 0033-5533, 1531-4650. DOI: 10.1093/qje/qjy010.
- Stromberg, D. (2004). "Radio's Impact on Public Spending". In: The Quarterly Journal of Economics 119.1, pp. 189–221. ISSN: 0033-5533, 1531-4650. DOI: 10.1162/ 003355304772839560.
- Strömberg, David (2001). "Mass media and public policy". In: *European Economic Review* 45.4, pp. 652–663. ISSN: 00142921. DOI: 10.1016/S0014-2921(01)00106-4.
- Titova, Maria (2024). "Targeted Advertising in Elections". In: Working Paper.
- Tucker, Joshua et al. (2018). "Social Media, Political Polarization, and Political Disinformation: A Review of the Scientific Literature". In: SSRN Electronic Journal. ISSN: 1556-5068. DOI: 10.2139/ssrn.3144139.

- Vaeth, Martin (2024). "Rational Voter Learning, Issue Alignment, and Polarization". In: *Working Paper*.
- Zhuravskaya, Ekaterina, Maria Petrova, and Ruben Enikolopov (2020). "Political Effects of the Internet and Social Media". In: *Annual Review of Economics* 12.1, pp. 415–438. ISSN: 1941-1383, 1941-1391. DOI: 10.1146/annurev-economics-081919-050239.

Appendix

Proof of Remark 1. This follows directly from the observation that the Beta distribution is the conjugate prior probability distribution for the Bernoulli distribution. \Box

Proof of Proposition 1. Suppose an exogenous signal is sent only to voters in segment ψ with belief \hat{x}_0 . Before targeting, their value to the politician is:

$$v_A(\hat{x}_0)g_\psi(\hat{x}_0)$$

Note that the probability of a randomly chosen individual observing the additional exogenous signal as 1 is \hat{x}_0 . This is because voters update rationally, so that \hat{x}_0 is the expected bliss point of the people with that belief. To see this more explicitly, note that for set of voters q, we have the probability that an additional signal realises as positive given as: $pr(signal = 1 | \hat{x}_i = \hat{x}_0) = \frac{1}{|q|} \sum_{i \in q} x_i = \mathbb{E}[x_i] = \hat{x}_0$.

This means that in expectation a proportion \hat{x}_0 of the individuals in the segment with belief \hat{x}_0 will observe the additional signal as 1, while $(1 - \hat{x}_0)$ of those will observe it as 0. Our assumption that each segment contains a population of individuals means we can treat this as deterministic, rather than as an expectation.

Therefore, the new valuation in the segment for those who started with belief \hat{x}_0 is:

$$(1 - \hat{x}_0)v_A\left(\hat{x}_0 - \frac{\hat{x}_0}{n_{\psi} + u_{\psi}}\right) + \hat{x}_0v_A\left(\hat{x} + \frac{1 - \hat{x}}{n_{\psi} + u_{\psi} - 1}\right).$$

Taking away the initial valuation we have:

$$(1 - \hat{x}_0)v_A\left(\hat{x}_0 - \frac{\hat{x}_0}{n_{\psi} + u_{\psi}}\right) + \hat{x}_0v_A\left(\hat{x} + \frac{1 - \hat{x}}{n_{\psi} + u_{\psi} - 1}\right) - v_A(\hat{x}_0) = \varphi_A(\hat{x}_0).$$

And summing over all beliefs in the segment we get:

$$\int_0^1 \varphi_A(\hat{x}) g_\psi(\hat{x}) d\hat{x}.$$

In the case of no endogenous signal received, hence the politician considering sending the first pieces of news, there is a single belief \hat{x} in the segment and the above simplifies to:

$$\varphi_A(\hat{x}).$$

Proof of Proposition 2. This proof proceeds by pinning down the sign of φ_A and arguing that this pins down the sign of $\int_0^1 \varphi_A(\hat{x}) g_{\psi}(\hat{x}) d\hat{x}$.

Note that by symmetry $\varphi_A(0.5) = 0$. Further, note that a monotonically increasing cost of voting PDF implies that v_A is convex below 0.5 and concave above 0.5. Similarly, a monotonically increasing cost of voting PDF implies v_A is concave below 0.5 and convex above 0.5. Therefore, in this proof we use the convexity or concavity of v_A directly, instead of using the properties of f_c .

Below we use L and R as defined in Proposition 1. These are beliefs that voters move to after observing a negative and positive signal respectively. This means that L and Rare always functions of some initial belief \hat{x} , but this initial belief is obvious from the context so we omit it from the notation. It will be useful to remember that we always have $L < \hat{x} < R$.

Case 1: $v_A(\hat{x})$ is concave on $\hat{x} \in [0, 0.5)$ and convex on $\hat{x} \in (0, 0.5]$. We will show that:

$$\varphi_A(\hat{x}) < 0 \Leftrightarrow \hat{x} < 0.5$$

$$\varphi_A(\hat{x}) > 0 \Leftrightarrow \hat{x} > 0.5$$

Case 1.1a: $\hat{x} < 0.5$ and R < 0.5.

Clearly $\varphi_A(\hat{x}) < 0$ by the concavity of $v_A(\hat{x})$ on $\hat{x} \in [0, 0.5)$.

Case 1.1b: $\hat{x} > 0.5$ and L > 0.5.

Clearly $\varphi_A(\hat{x}) > 0$ by the convexity of $v_A(\hat{x})$ on $\hat{x} \in (0.5, 1]$.

Case 1.2a: $\hat{x} < 0.5$ and R > 0.5.

 $(1-\hat{x})v_A(L) + \hat{x}v_A(R)$ is the height of the secant line between $v_A(L)$ and $v_A(R)$ at \hat{x} . This means that it is sufficient to show that the secant line lies below $v_A(\hat{x})$.

By symmetry a line going through $v_A(0.5 - \gamma)$ and $v_A(0.5 + \gamma)$ must go through the point (0.5, 0), and it must therefore be below v_A on $\hat{x} \in (0.5 - \gamma, 0.5)$ and above v_A on $\hat{x} \in (0, \gamma)$. This means that for any $0.5 \ge \gamma' > \gamma \ge 0$, the line going through $v_A(0.5 - \gamma')$ and $v_A(0.5 + \gamma)$ must lie below v_A on $x \in (0.5 - \gamma', 0.5 + \gamma)$.

Further, for $n_{\psi} + u_{\psi} \ge 1$, it must be that $|L - 0.5| \ge |R - 0.5|$, so that by the logic above v_A must lie above the secant line joining $v_A(L)$ and $v_A(R)$ on $x \in (L, 0.5)$. This is because we can plug in $L = 0.5 - \gamma'$ and $R = 0.5 + \gamma$ for some $0.5 \ge \gamma' > \gamma \ge 0$.

Case 1.2b: x > 0.5 and L < 0.5.

A mirror image of the argument in 1.2a yields the result that $\varphi(x) \ge 0$.

Case 2: $v_A(x)$ is convex on $x \in [0, 0.5)$ and concave on $x \in (0, 0.5]$.

Steps of the proof are a mirror image of Case 1.

Finally, we show that this property extends to $\int_0^1 \varphi_A(x) g_{\psi}(x) dx$. By the variation diminishing property and Remark 5, $\int_0^1 \varphi_A(x) g_{\psi}(x) dx$ cannot change sign more than once on [0, 1]. By symmetry $\int_0^1 \varphi_A(x) g_{0.5}(x) dx = 0$.

Now all that is left to do is demonstrate that:

- 1. For a convex-concave v_A :
 - a. $\pi_A(\psi)$ is positive somewhere on $\psi \in (0, 0.5)$.
 - b. $\pi_A(\psi)$ is negative somewhere on $\psi \in (0.5, 1)$.

2. For a concave-convex v_A :

- a. $\pi_A(\psi)$ is negative somewhere on $\psi \in (0, 0.5)$.
- b. $\pi_A(\psi)$ is positive somewhere on $\psi \in (0.5, 1)$.

We will demonstrate 1.a. and the rest will follow by symmetry.

Since signals are binary and everyone within the same segment observes the same number of signals, estimated bliss points within a segment are spaced $\frac{1}{n+u}$ apart. Further, the mass of g_{ψ} is concentrated around ψ . Therefore, a segment with $\psi < 0.5$ for which $g_{\psi}(0.5 - \frac{1}{n+u}) > 0$ will have the following property:

$$g_{\psi}(0.5+\alpha) > 0 \Rightarrow g_{\psi}(0.5-\alpha) > g_{\psi}(0.5+\alpha) \quad \forall \alpha > 0$$

By symmetry of φ_A , this means that this segment will have $\pi_A(\psi) > 0$. This segment exists. We show this by showing that a segment with $\psi < 0.5$ containing mass on any point $\gamma < 0.5$ always exists.

A condition that is sufficient is that for some ψ the following holds:

$$f(\psi) := \frac{\psi n + \lceil \gamma u \rceil}{n + u} = \gamma$$

Clearly $f(0) < \gamma$, $f(0.5) > \gamma$, and $f(\psi)$ is continuous in ψ .

Proof of Proposition 3. Note that **A' 1** guarantees that there exists an increasing sequence x_1, x_2, x_3, x_4 such that $\varphi(x_1) < 0$, $\varphi(x_2) > 0$, $\varphi(x_3) < 0$, $\varphi(x_4) > 0$. To see this is true, consider the following:

- Choose x_1 such that $\frac{nx_1+1}{n+u+1} < \frac{mode(f_c)}{2}$. Then $\varphi(x_1)$ is a measure of convexity over a wholly concave interval, so that $\varphi(x_1) < 0$.
- Choose x_2 such that $\frac{nx_2+1}{n+u+1} = 0.5$. Then $\varphi(x_1)$ is a measure of convexity over a wholly convex interval, so that $\varphi(x_2) > 0$.
- Choose x_3 such that $\frac{nx_3}{n+u+1} = 0.5$. Then $\varphi(x_1)$ is a measure of convexity over a wholly concave interval, so that $\varphi(x_2) < 0$.
- Choose x_4 such that $\frac{nx_4}{n+u+1} > 1 \frac{mode(f_c)}{2}$. Then $\varphi(x_1)$ is a measure of convexity over a wholly convex interval, so that $\varphi(x_1) > 0$.

Next we will show that $\varphi(\hat{x})$ changes sign at most three times on $\hat{x} \in [0, 1]$. We introduce a modified version of φ :

$$\begin{split} \tilde{\varphi}(z) &:= (1-\xi)v(z) + \xi v(z + \frac{1}{n+u+1}) - v(\xi+z) \\ z &\in (0, 1 - \frac{1}{n+u+1}); \quad n > 0; \quad \xi \in (0, \frac{1}{n+u+1}) \end{split}$$

This modified function may need some explanation. z plays the role L did previously, but here we choose L instead of \hat{x} , and the distance between L and \hat{x} is set by ξ . φ then is a special case of $\tilde{\varphi}$ where $\xi = \frac{\hat{x}}{n+u+1}$.

Rewrite the function as an integral:

$$\tilde{\varphi}(z) = \int_{[-1,1]} v(\eta + z) d(\mu_1(\eta) - \mu_2(\eta))$$

Where:

$$d\mu_1(\eta) = \begin{cases} (1-\xi) & if \ \eta = 0\\ \xi & if \ \eta = \frac{1}{n+u+1}\\ 0 & Otherwise \end{cases}$$
$$d\mu_2(\eta) = \begin{cases} 1 & if \ \eta = \xi\\ 0 & Otherwise \end{cases}$$

Using integration by parts:

$$\int_{[-1,1]} v(\eta+z)d(\mu_1(\eta)-\mu_2(\eta)) = [v(\eta+z)(\mu_1(\eta)-\mu_2(\eta))]_{-1}^1 - \int_{[-1,1]} (\mu_1(\eta)-\mu_2(\eta))v'(\eta+z)d\eta$$

The first term on the RHS is 0, so we are left with:

$$-\int_{[0,1]} (\mu_1(\eta) - \mu_2(\eta)) v'(\eta + z) d\eta$$

Where $\mu_1(\eta) - \mu_2(\eta)$ changes sign exactly once (at ξ). If $v'(\cdot)$ is TP_2 over some interval then, holding ξ constant, $\tilde{\varphi}(z)$ can change sign at most once over that interval.

The definition of v is given below:

$$v(\eta + z) = F_c(1 - 2(\eta + z)) - F_c(2(\eta + z) - 1)$$

$$v'(\eta + z) = \frac{\partial v(\eta + z)}{\partial \eta} = -2(f_c(1 - 2(\eta + z)) + f_c(2(\eta + z) - 1))$$

For $z \in [0, 0.5 - \frac{1}{n+1}] \cup [0.5, 1 - \frac{1}{n+1}]$, $v'(\eta + z)$ is only ever evaluated on one side of 0.5 at a time, so that is TP_2 by the log-concavity of f_c . This means that for the above intervals, $\tilde{\varphi}$ changes sign at most once at each interval, but we still need to show that this implies $\varphi(x)$ changes sign at most once on those intervals.

To show this, first note that the only difference between $\tilde{\varphi}$ and φ is that $\tilde{\varphi}$ is a version of φ in which the distance between L and x is pinned down by a choice of ξ , whereas in φ this distance increases as L increases. This means that we need to ensure that as Lincreases in this interval, the movement of ξ does not introduce new sign changes.

To argue this we note that φ is a measure of distance between v(x) and the secant line joining v(L) and v(R), so that given the following two properties:

- The secant line crosses v(x) at most once on (R, L).
- This crossing moves closer to L as L increases (if it does not disappear).

Then we could guarantee that φ does not have more sign changes than $\tilde{\varphi}$ on the intervals $L \in [0, 0.5 - \frac{1}{n+1}]$ and $L \in [0.5, 1 - \frac{1}{n+1}]$. This is because these conditions would mean that as ξ moves to the right, the crossing moves to the left, so that the movement of ξ does not introduce additional sign changes.

For the first bullet point—v'' changes sign at most once on [L, R], so that v(x) can cross a straight line at most three times on this interval. The secant line is a straight line that crosses v(x) once at v(L) and once at v(R), so that it can cross v at most once on (L, R).

For the second bullet point, suppose this does not occur. Now consider a < a' where the secant line crosses v three times for all $L \in [a, a']$. This means that there must be a change in the sign of convexity within this segment for all $L \in [a, a']$. By assumption **A' 1**, there is exactly one such change. In the concave-convex case, the secant line is below v_A before the middle crossing and above it after the middle crossing. If the crossing is further from L for L = a' than for L = a, then for the concave-convex case this means that the secant line of a' must lie below the secant line of a at all points to the right of the middle crossing of a (because v_A is decreasing). This means that the rightmost point at which the secant line crosses v_A given L = a' is to the left of the rightmost point at which the secant line crosses v_A given L = a. This implies $a' + \frac{1}{n+u+1} < a + \frac{1}{n+u+1}$ because the R intercept must be a distance $\frac{1}{n+u+1}$ from L. This is a contradiction. A similar argument applies to the convex-concave case.

Lastly, we examine the case where $L \in (0.5 - \frac{1}{n+1}, 0.5)$:

If $mode f_c < 1 - \frac{1}{n+u+1}$, then $v_A(x)$ is convex on $x \in (0.5 - \frac{1}{n+1}, 0.5]$ and concave on $x \in [0.5, 0.5 + \frac{1}{n+1})$. Further v(0.5 - q) = -v(0.5 + q) for $q \in [0, 0.5]$ and for $L \in (0.5 - \frac{1}{n+1}, 0.5)$, $L, R \in (0.5 - \frac{1}{n+1}, 0.5 + \frac{1}{n+1})$. Therefore, the proof from Proposition 2 can be used to show that $L \in (0.5 - \frac{1}{n+1}, 0.5 - \frac{1}{2(n+1)}) \Rightarrow \varphi > 0$, $L = 0.5 - \frac{1}{2(n+1)} \Rightarrow \varphi = 0$, and $L \in (0.5 - \frac{1}{2(n+1)}, 0.5) \Rightarrow \varphi > 0$. Therefore, we have that φ changes at most once over $L \in [0, 0.5 - \frac{1}{n+1}]$, exactly once over $L \in (0.5 - \frac{1}{n+1}, 0.5)$, and at most once over $L \in [0.5, 1 - \frac{1}{n+1}]$. This is equivalent to three times over $x \in (0, 1)$.

Remark 7. $g_{\psi}(x)$ is TP_{∞} .

Proof of Remark 7. The Beta distribution and Binomial distribution are TP_{∞} (Karlin 1968). By Theorem 2 from Jewitt (1987) (first presented in Karlin 1968), this means that $g_{\psi}(x)$ is also TP_{∞} .

Proof of Corollary 2. By Proposition 3, $\varphi(x)$ changes sign at most three times over $x \in (0, 1)$. The value of targeting in a segment is:

$$\int_0^1 \varphi_A(x) g_\psi(x) dx,$$

Where $g(\psi)$ is TP_{∞} by Remark 3, so that the value of targeting a segment changes sign at most three times.

Note that $\pi(0.5) = 0$ always, so that by symmetry it is impossible for π to change sign twice. This means that if targeting is ever profitable, then π either has the shape described in Proposition 2 or the shape described in Proposition 3.

Proof of Remark 2. Define the following:

$$L := x - \frac{x}{n+1},$$
$$R := x + \frac{1-x}{n+1}.$$

Rewriting $\varphi(x)$:

$$\varphi(x) = xv_A(R) + (1-x)v_A(L) - v_A(x).$$

For any given x, if L and R are either both below or both above 0.5, then clearly $\varphi(x) = 0$. If one of L and R is below 0.5 while the other is above 0.5, then targeting the segment will be profitable *iff* x > 0.5. Note that x < 0.5 means some people go from voting for A to voting for B while nobody goes from voting for B to not voting or to voting for A. The reverse is true for x > 0.5. For x = 0.5, the movement to the left and to the right cancel each other out.

For the case that L and R are on different sides of 0.5, the proportion of people who move left is (1 - x). Within the set bigger than 0.5, this is maximised for x approaching 0.5 from above.

Proof of Remark 4. Clearly $\varphi(\hat{x}) > 0 \Rightarrow \hat{x} > 0.5$ and $\varphi(\hat{x}) < 0 \Rightarrow \hat{x} < 0.5$. This is simply because if $\hat{x} < 0.5$, then voters with belief \hat{x} will already definitely vote for politician A so that there is nothing to be gained by targeting. Similarly, if $\hat{x} > 0.5$, then voters with belief \hat{x} will already definitely vote for politician B so that there is nothing to be lost by targeting. Further, by total positivity of g_{ψ} and the variation diminishing property, $\pi_A(\hat{x})$ changes sign at most once. By symmetry, $\pi_A(0.5) = 0$. Therefore, all that is left to show is that $\pi_A(\hat{x}) < 0$ for some $\hat{x} < 0.5$ and $\pi_A(\hat{x}) > 0$ for some $\hat{x} > 0.5$. We use the same construction as in the proof of Proposition 2 to show this:

Since signals are binary and everyone within the same segment observes the same number of signals, estimated bliss points within a segment are spaced $\frac{1}{n+u}$ apart. Further, the mass of g_{ψ} is concentrated around ψ . Therefore, a segment with $\psi < 0.5$ for which $g_{\psi}(0.5 - \frac{1}{n+u}) > 0$ will have the following property:

$$g_{\psi}(0.5+\alpha) > 0 \Rightarrow g_{\psi}(0.5-\alpha) > g_{\psi}(0.5+\alpha) \quad \forall \alpha > 0$$

By symmetry of φ_A , this means that this segment will have $\pi_A(\psi) < 0$. This segment exists. We show this by showing that a segment with $\psi < 0.5$ containing mass on any point $\gamma < 0.5$ always exists.

A condition that is sufficient is that for some ψ the following holds:

$$f(\psi) := \frac{\psi n + \lceil \gamma u \rceil}{n+u} = \gamma$$

Clearly $f(0) < \gamma$, $f(0.5) > \gamma$, and $f(\psi)$ is continuous in ψ .

Proof of Lemma 1. Suppose that the common shock has some distribution such that the posterior distribution of positions, given belief x before the shock realises, is f_x .

A common shock transforms the value of a segment:

$$\int v_A(x)g_\psi(x)dx \to \int \int (v_A(z)f_x(z)dz) g_\psi(x)dx$$
$$= \int \int (v_A(z)f_x(z)g_\psi(x)dz) dx.$$

By Fubini:

$$= \int \int (v_A(z)f_x(z)g_\psi(x)dx)\,dz.$$

Then by Proposition 1, the marginal value of a signal is:

$$\int \int (\varphi_A(z) f_x(z) g_\psi(x) dx) dz.$$
$$= \int \int (\varphi_A(z) g_\psi(x) dx) f_x(z) dz$$

Note that f_x is TP_{∞} because it is part of the exponential family. Therefore, by the variation diminishing property, there are no more sign changes with a common shock than without a common shock.

Proof of Proposition 4. Using Lemma 1 and Proposition 2, we need only demonstrate that:

- 1. For a convex-concave v_A :
 - a. $\pi_A(\psi)$ is positive somewhere on $\psi \in (0, 0.5)$.
 - b. $\pi_A(\psi)$ is negative somewhere on $\psi \in (0.5, 1)$.
- 2. For a concave-convex v_A :
 - a. $\pi_A(\psi)$ is negative somewhere on $\psi \in (0, 0.5)$.
 - b. $\pi_A(\psi)$ is positive somewhere on $\psi \in (0.5, 1)$.

We proceed as in Proposition 2. We prove case 1.a and the rest follow by symmetry. We will argue that the construction used in Proposition 2 is still valid in the presence of a common valence shock.

As per Proposition 2, we pick any segment with $\psi < 0.5$ for which $g_{\psi}(0.5 - \frac{1}{n+u})$. We know it will have the following property:

$$g_{\psi}(0.5+\alpha) > 0 \Rightarrow g_{\psi}(0.5-\alpha) > g_{\psi}(0.5+\alpha) \quad \forall \alpha > 0$$

Further, by the symmetry of the common valence shock and v_A , we know that $\int (v_A(z)f_{0.5-\alpha}) = -\int (v_A(z)f_{0.5-\alpha})$

Proof of Remark 5. Suppose every segment has one person in it, so that whether the election is won is stochastic. We only consider the marginal effect of news on winning. Denote by $\nu(\psi, X)$ the probability of segment ψ voting for party X. We state the increased probability of winning when we change probabilities in one segment ψ from $\nu(\psi, A), \nu(\psi, B)$ to $\nu'(\psi, A), \nu'(\psi, B)$. Then the gain from changing the probabilities as above is:

$$(\nu'(A) - \nu(A))(M_d + M_l) - (\nu'(B) - \nu(B))(M_d + M_w).$$

This is because a change in the probability of one segment voting for A or B only matters when that segment is pivotal. Specifically, the probability of the segment voting for A only matters when that segment is pivotal and A would not have won if ψ abstained (because abstaining and voting for A here leads to the same electoral outcome). Similarly, the probability of the segment voting for B only matters when B would not have won if ψ abstained.

From here we can see that if we have pivot symmetry and pivot equality, then maximising the chance of winning is equivalent to maximising net expected plurality.

Proof of Proposition 5.

Case 1: R < 0.5

We construct $\tilde{\varphi}$ as in Proposition 3 which holds the distance between L and x constant for this proof, then we argue that because the proposition is true for any distance between L and x, then it is true for φ .

$$\alpha := \frac{1}{n+u}$$
$$\tilde{\varphi} := \xi v_A(L) + (1-\xi)v_A(R) - v_A(L+\xi\alpha)$$
$$= \xi F_c(1-2L) + (1-\xi)F_c(1-2(L+\alpha)) - F_c(1-2(L+\xi\alpha))$$

$$= \xi \left(\int_{0}^{1-2(L+\alpha)} f_{c}(z)dz + \int_{1-2(L+\alpha)}^{1-2L} f_{c}(z)dz \right) + (1-\xi) \left(\int_{0}^{1-2(L+\alpha)} f_{c}(z)dz \right) \\ - \left(\int_{0}^{1-2(L+\alpha)} f_{c}(z)dz + \int_{1-2(L+\alpha)}^{1-2(L+\xi\alpha)} f_{c}(z)dz \right) \\ = \xi \left(\int_{1-2(L+\alpha)}^{1-2L} f_{c}(z)dz \right) - \left(\int_{1-2(L+\alpha)}^{1-2(L+\xi\alpha)} f_{c}(z)dz \right).$$

For convenience we define a measure on F_c that takes sets as arguments as $\mu(\cdot)$. We define the following sets:

$$a := [1 - 2(L + \alpha), 1 - 2(L + \xi \alpha)]; \qquad b := [1 - 2(L + \alpha), 1 - 2L].$$

So that we have:

$$\tilde{\varphi} = \xi \mu(b) - \mu(a).$$

Note that b > a. Now suppose that for some L and cost of voting distribution F_{c1} corresponding to measure μ_1 , we have:

$$\xi \mu_1(b) - \mu_1(a) \ge 0$$
$$\Leftrightarrow \frac{\mu_1(a)}{\mu_1(b)} \le \xi. \tag{1}$$

From the definition of the LR order, we have that for $\mu_2 \succ_{LR} \mu_1^{11}$:

$$\frac{\mu_1(a)}{\mu_1(b)} \ge \frac{\mu_1(a)}{\mu_1(b)}.$$

So that if (1) is true for some μ_1 , it must be true for any μ_2 such that $\mu_2 \succ_{LR} \mu_1$. Further, since this is true for any ξ , it holds for φ .

Case 2: L > 0.5.

Same proof as above but flipped.

Case 3: L < 0.5 and R > 0.5.

We know from Proposition 3 that for this area we have:

$$x < 0.5 \Rightarrow \varphi \ge 0; \quad x > 0.5 \Rightarrow \varphi \le 0.$$

| . 1 | - | - | - | - | - | ٦ |
|-----|---|---|---|---|---|---|
| | | | | | | |
| | | | | | | |
| | | | | | | |
| | | | | | | |
| | | | | | | |

¹¹Modification of definition 1.C.3 from Shaked and Shanthikumar (2007).