Targeted Advertising and Candidate Entry[*,†]

Gregory J. Martin[‡] Saikun Shi[§]

August 2026

Abstract

We examine how the introduction of targeted digital advertising influences the pool of candidates who run for Congressional and state legislative offices in the United States. We develop a model of candidate entry with an advertising stage, and use the model to generate predictions about two consequences of digital advertising: 1) the large reduction in candidates' advertising costs, relative to television ads, and 2) the large increase in the ability to precisely target subsets of the electorate. We show that the first effect generally increases political competition and encourages the entry of candidates closer to the district median. But the second effect depends on the structure of targeting features available, and can in some cases undo the moderating effect of price reduction. We take the predictions to data on candidate entry and advertising decisions in Congressional and state legislative races. We find asymmetric consequences across the two major parties, with the competition-enhancing price effect dominating on the Democratic side but the targeting effect dominating on the Republican side. We connect this result to the advantageous structure of targeting features available to candidates on the right relative to those on the left of the ideological spectrum.

[*]Researcher(s)' own analyses calculated (or derived) based in part on data from The Nielsen Company (US), LLC and marketing databases provided through the Nielsen Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the researcher(s) and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

[†]Meta Platforms, Inc. provided researchers access to the Ad Targeting and Ad Library Datasets, https://doi.org/10.48680/meta.adtargetingdataset. All analyses and conclusions are researchers' own.

[‡]Stanford GSB. gjmartin@stanford.edu.

[§]Stanford GSB. saikun@stanford.edu.

Most existing research on campaign advertising focuses on its effects (or lack thereof) on voter behavior, such as candidate choices, turnout decisions, or attention to politics.[1] But regardless of the magnitude of such effects on voters, the potential to advertise can also affect *candidates*' behavior, in particular the decision to run for office at all. As elected officials were once candidates, changes in the candidate pool are passed through to representation and policy outcomes.

A relatively recent literature (Hall, 2019; Fouirnaies, 2021; Avis et al., 2022; Cox, 2026, e.g.) studies this side of the equation, demonstrating that changes in candidates' need to raise money to spend on campaign expenses affect the pool of potential entrants to electoral contests. In this paper we abstract from the fundraising question and focus instead on advertising technology, specifically the precision with which political advertising messages can be targeted to sub-segments of the electorate. Existing work has shown that such targeting can enhance campaigns' persuasive capacity (Tappin et al., 2023). Our aim is to understand how technological change in media—such as the early-21st century shift from broadcast advertising to more precisely targeted digital advertising—can affect which kinds of candidates are willing to run for office as well as how they reach voters when they run.

The nature of such effects is not obvious ex ante. On the one hand, studies have connected the rise of digital media to polarization; some scholars argue that the easier and more rapid exchange of information is at the root of increasing disinformation and misperception, which in turn contribute to polarization (e.g. Tucker et al. (2018), Bail et al. (2018)). In an unruly and chaotic digital marketplace of ideas, perhaps those candidates willing to make the wildest claims or stake out the most extreme positions are the ones who stand out enough to draw voters' attention. And perhaps effective data-mining operations to identify and target likely supporters can make ideologically extreme candidates viable despite their unpopular positions (Hindman, 2018).

[1] See e.g. Aggarwal et al. (2023); Spenkuch and Toniatti (2018); Gordon and Hartmann (2013); Ashworth and Clinton (2007); Huber and Arceneaux (2007), among many others.

On the other hand, as Fowler et al. (2021) point out, a more mundane consequence of improved targeting through digital advertising is that it enables candidates to better target voters *in their district* as opposed to outside of it. Digital ads reduce the effective cost of political advertising compared to broadcast ads, due to reductions in wasted ad impressions shown to viewers ineligible to vote for the candidate. The associated reduction in cost can be extremely large, due to the typically poor congruence between electoral districts and broadcast media markets.[2] By making advertising feasible in a much broader set of races, more-precise targeting might increase political competition, allowing new challengers to enter and take on entrenched incumbents out-of-step with their districts.

Model To organize ideas and make precise our empirical expectations, we develop a model of candidate entry with an advertising stage. The model extends the classic setup of Osborne and Slivinski (1996) to a setting in which a candidate is not in a voter's consideration set at the voting booth unless informed of the candidate's existence through advertising, which is costly to candidates. We model the introduction of digital forms of advertising as decreasing the average cost of advertising (reflecting the effect of better geographic precision) and also allowing for messages to be targeted to subsets of voters within the district.

Our model of advertising is purely informative, in the sense of Bagwell (2007); seeing a candidate's ad perfectly informs the voter about the candidate's ideal policy position. It is the starkest possible model of informative advertising, in that there is exactly one piece of information that is relevant to all voters and candidates have only the binary choice of revealing it, or not. In this sense our assumptions are starker than the related theoretical paper by Titova (2026), who allows candidates to send ambiguous messages about their policy preference. We intentionally abstract from the possibility of lying, obfuscation, misrepresentation, or emotional appeals that affect candidate choice through non-informational channels, in order

[2] This expansion of advertising feasibility is particularly evident in down-ballot contests; Fowler et al. (2021) show that about 1% of state legislative candidates advertise on television, whereas roughly half do so on Facebook.

to focus attention on the effect of targeting precision per se.[3]

We show that in this model, reducing advertising costs while maintaining the requirement that candidates advertise to the whole district leads to higher political competition and a more moderate (relative to the district) candidate pool. The set of one-candidate equilibria shrinks when advertising prices fall, and the set of candidates who enter in some equilibrium is closer to the district median. Digital advertising thus helps to break down incumbency or name-recognition advantages that insulate out-of-step incumbents from electoral challenges.

The introduction of within-district targeting, however, maintains the effect on competition but can undo the effect on moderation. In spite of the fact that advertising in our model is purely informative, precise within-district targeting by itself can advantage extreme entrants. The reason is that targeting can give extreme candidates a relative efficiency advantage—a higher marginal return in votes, per dollar spent—in advertising, by creating segments in which they are much more popular than they are district-wide and where the returns to advertising are very high.[4] A candidate who is a minority choice district-wide but enjoys concentrated support within some segment *that can be targeted* gains a substantial advantage, compared to the "broadcast" setting where she must either advertise to the whole district or not at all.

Data We take the model predictions to data from campaigns for the US House and for US state legislative chambers. Our analysis draws on a range of data sources. We build a comprehensive view of the candidate pool from campaign contribution records at federal and state levels, augmented with information from Ballotpedia's database of candidate contact information. We collect granular information on campaigns' digital advertising choices, including targeting information, at the advertisement level from Meta's Ad Targeting and Ad

[3]It is possible that targeting precision may help campaigns identify "persuadable" or easily manipulable voters, as Cambridge Analytica claimed; we abstract from this possibility here. See Hewitt et al. (2024) for some reasons to be skeptical of campaigns' ability to reliably identify the most persuadable groups.

[4]Note that our use of "moderate" and "extreme" here need not necessarily align with a left-right ideological dimension. By "extreme" we simply mean far from the district median along whatever preference dimension is salient in voters' choices. The efficiency mechanism we describe applies on any such dimension.

Library datasets (Meta Platforms, Inc., 2026a,b). And we measure variation in the relative attractiveness of digital versus broadcast advertising across districts by collecting effective TV advertising prices at the district level, taking advantage of variation in alignment between district boundaries and TV markets.

We use these datasets to establish several empirical patterns. First, we show that candidates in districts where the effective price of TV ads is higher allocate higher percentages of their campaign expenditures to digital ads. Digital advertising thus appears to be a viable substitute for TV ads: candidates resort more to digital ads when TV advertising costs are high and vice versa. Second, we show that in the period before digital advertising became widespread in campaigns, fewer candidates entered electoral contests on average in districts where the effective cost of TV advertising was high. The negative relationship between TV ad costs and contestation diminishes over time and ultimately disappears by the most recent election cycles. The relative cost-effectiveness of digital advertising does appear to have increased the level of competition in these formerly difficult-to-contest races.

We also find some evidence that high TV advertising costs are associated with a more extreme candidate pool. Prior to the availability of digital advertising, candidates in high TV-cost districts tended to have more extreme CFScores (Bonica, 2023) compared to what would be expected given the district's electorate. That relationship also disappears once digital advertising becomes widely available, at least for Democratic candidates; for Republican candidates the correlation is preserved or even strengthened over time.

Last, we bring in detailed targeting data from Meta that covers congressional and state legislative races from 2020 to 2024 to describe how candidates take advantage of the ability to target subsets of the electorate. We show that even at the state legislative level, and even in partisan primaries, there is very little overlap in the segments of the electorate targeted by opposing candidates in the same race. One reason for this is that candidates rely heavily on what are known as "custom" or "lookalike" audiences: lists of users supplied to the platform by the campaign for ad delivery, either to the user list or to other users with similar attributes to the users identified by the campaign. These patterns are consistent with our model's efficiency mechanism, in which candidates concentrate spending in segments where their marginal return per dollar is highest and rivals' returns are lowest.

The advantage of targeting to ex-ante receptive groups is not evenly distributed across candidate types, however. We find that of the segments most advertised-to by candidates in our sample, the large majority are right-leaning, and targeted overwhelmingly by Republican candidates. This pattern is a result of the asymmetry in the US media market: a large fraction of the most popular ad targets on the right are users who follow conservative media outlets and conservative media personalities. There simply do not exist parallel targetable features on the left with similarly wide reach and tight ideological focus.

This distribution of targeting features gives an efficiency advantage to candidates of the right. Relatedly, we find that while the introduction of digital ads seems to have moderated the Democratic candidate pool in high-TV-cost districts, consistent with our model's price mechanism, it has if anything made Republican candidates in those districts more extreme. This asymmetric pattern is especially evident at the state legislative level, where voters' knowledge of candidates is more thin. The impact of digital advertising thus depends critically on the alignment of targetable features on social media platforms. The asymmetry in the US political media sphere has an unanticipated consequence of favoring the entry of right-wing candidates, not through any form of voter persuasion but simply through making it easier for candidates of the right to find receptive audiences online.

Relationship to existing literature On the theoretical side, our model directly extends Osborne and Slivinski (1996), but our motivating question is closest to that in Titova (2026), who also models a game of targeted advertising in elections. Like Titova (2026), we assume candidates have no ability to commit to policy, and model advertising as revealing verifiable information to voters. We differ in restricting the set of messages that candidates can send to a point rather than a range; in allowing for variation in voters' pre-advertising knowledge of candidates rather than assuming that one (incumbent) candidate's position is perfectly known; and in taking the sub-groups that can be targeted as a model primitive rather than a choice variable of the candidates. We think these differences are appropriate to our empirical setting, but they also show that Titova's (2026) conclusion—that targeted advertising can allow candidates who would lose under full information to win—does not depend on these specific modeling choices.

There is also a broader theoretical literature on (broadcast) advertising in elections. Much of this literature (e.g. Ashworth (2006) or Desai and Duggan (2021)) focuses on the interaction with donors, who seek to influence candidate policy platforms through promises to provide or withhold contributions. We abstract from this interaction and take candidates' ability to spend as given, focusing instead on how advertising technology influences different candidates' chances of winning for a given budget. We also abstract from the possibility of non-informational influences of advertising, which is explored in a large literature following Baron (1994), e.g. Buisseret and Van Weelden (2025); our aim is to show that precise targeting can be consequential even if campaigns are not able to use it to identify and reach manipulable voters. Our advertising model is one of verifiable or "hard" information, as in e.g. Polborn and Yi (2006).

On the empirical side, several large-scale field experiments (Allcott et al., 2026; Aggarwal et al., 2023; Coppock et al., 2022; Enríquez et al., 2024, e.g.) have estimated effects of digital advertising on voting, candidate perceptions, and turnout. The general pattern of results in US federal elections is one of small-to-zero treatment effects, consistent with experimental evidence on broadcast advertising (Kalla and Broockman, 2018; Coppock et al., 2020). We note that our hard-information theory of advertising is entirely consistent with zero—or even negative—average treatment effects in designs where exposure to a single advertising campaign is randomized across voters, like Aggarwal et al. (2023) and Coppock et al. (2022).

The design in Allcott et al. (2026) is the most relevant to our theory, as those authors randomly suppressed Facebook users' ability to see political advertising *targeted to them*. But the setting in Allcott et al. (2026) is a US presidential general election, most of the ads suppressed relate to the presidential campaign, and the main outcomes measured are presidential candidate evaluations. The central mechanism for advertising influence in our theory is through adding a new candidate to the voter's consideration set, a mechanism which is not in play at the presidential general election level. As Enríquez et al. (2024) show, social media advertising campaigns can have sizable electoral effects in lower-information contexts.

Finally, several authors have investigated empirical patterns in targeting. Kim et al. (2018) and Kim et al. (2026) track all ads served to a panel of users, focusing on divisive issue campaigns and vote-suppression messages, respectively. We capture the universe of ads run by a defined set of campaigns, rather than the universe of ads received by a defined set of users. Votta et al. (2024) and Bär et al. (2024), closer to our design, use advertiser-level data from the Meta Ad Targeting dataset; we differ in our focus on candidate-sponsored advertising in the US. Waldfogel (2026) examines the differences in message content between ads targeted to different audiences by the same campaign in US House elections. Our focus is on the choice of targets itself rather than the choice of message conditional on the target.

In the remainder of the paper, we set up a model of candidate entry with an advertising stage and derive several empirical predictions. We then take the predictions to data from the US Congress and state legislatures.

A model of advertising and candidate entry

We characterize a game between citizen-candidates extending the classic Osborne and Slivinski (1996) model. The game involves a set of potential candidates $\mathbb{J}$ who can opt to run for office or not, and a set of voters $\mathbb{N}$ who choose between the set of candidates who decide to run. Both candidates and voters are defined by their ideal policy x on the real line; candidates cannot commit to implement any policy other than their ideal when in office. We assume the support of ideal points for both candidates and voters is contained in the same bounded subset of $\mathbb{R}$.

Our extension is to relax the assumption of complete information about candidate types, and to add an additional stage of the game where candidates choose how much to inform the public about their preferred policy, at some cost. The game proceeds in three stages:

1. **Entry**. Candidates choose whether or not to enter the election. The entry choice of candidate j is denoted $e_j \in \{0,1\}$. A candidate who enters pays a fixed cost c, just as in the standard Osborne and Slivinski (1996) model. We denote the set of candidates who choose to enter ($e_j = 1$) as $\mathbb{J}_1$.

2. **Advertising**. Once candidates have made entry decisions, candidates who chose to enter choose how much money to spend on advertising $\mathbf{a}_j$. We will allow $\mathbf{a}_j$ to be vector valued, allowing for the possibility of targeting ads to some segments of the population. Advertising can be purchased at constant and common cost $\mathbf{p}$, which we assume to take the form $p_m = p\mu_m$, where μ_m is the size of segment m as a fraction of the total population.

3. **Voting**. All citizen-candidates vote for the candidate whose perceived platform is closest to their own; the candidate receiving the most votes wins. The winning candidate, denoted x_w, gets an office benefit b and implements her ideal policy.

At the end of the game, payoffs for all players are:

$$u_j(e_j, \mathbf{a}_j) = \begin{cases} -|x_j - x_w| & e_j = 0 \\ b - \mathbf{p} \cdot \mathbf{a}_j - c & e_j = 1, j = w \\ -|x_j - x_w| - \mathbf{p} \cdot \mathbf{a}_j - c & e_j = 1, j \neq w \end{cases}$$

Details of each stage

Voting

We assume the set of citizen ideal points $\mathbb{N}$ is distributed over the real line according to the continuous, known distribution function G, which has bounded support. At the voting stage, all voters vote non-strategically, selecting the candidate whose perceived position is closer to their own.

We model knowledge as binary: a voter i is either informed or uninformed about the type of a candidate j for office. If informed, she learns the type exactly and with certainty, i.e. her belief is a degenerate distribution at the true value x_j. If uninformed, she believes the candidate's policy could lie anywhere on the real line, with a belief such that $E[|x_j - x_i|] = \infty$.[5] This implies that a voter always strictly prefers a candidate with known policy position to one about whom she is uninformed. We assume that if the voter's most-preferred candidate is not unique, she splits her vote with equal probability between the set of candidates who are tied for most preferred.

The share of voters who vote for candidate j is denoted s_j, and the vector of shares for all candidates is $\mathbf{s}$. We assume that the winner of the election is assigned according to a probabilistic plurality rule. Namely:

$$\mathbb{P}(j \text{ wins} \mid \mathbf{s}) = \Phi_j(\mathbf{s}) \equiv \frac{\Phi(s_j - \max_{j' \neq j} s_{j'})}{\sum_{j'' \in \mathbb{J}_1} \Phi(s_{j''} - \max_{j' \neq j''} s_{j'})}$$

where the function $\Phi(\cdot)$ is any valid CDF with support on $(-1, 1)$ that admits a bounded and continuous PDF $\phi \equiv \frac{\partial}{\partial s}\Phi(s)$, and is symmetric about 0, $\Phi(a) = 1 - \Phi(-a)$.

[5]This is equivalent to assuming that a candidate is only in a citizen's "consideration set" if the citizen is aware of her position.

Advertising

At the advertising stage, all candidates who have entered the race decide how much to spend on advertising. We assume there is some finite and pre-determined partition $\mathbb{A}$ of sets of citizens such that $\bigcup_m^{|\mathbb{A}|} \mathbb{A}_m = \mathbb{N}$, and $\mathbb{A}_m \bigcap \mathbb{A}_{m'} = \emptyset$, $\forall\ m \neq m'$. The fraction of citizens in each segment m is denoted μ_m, with $\sum_m \mu_m = 1$. Candidates choose the level of advertising for each segment in the partition, $a_{jm} \geq 0\ \forall\ m$, at a cost that scales with the segment size: $C(a_{jm}) = pa_{jm}\mu_m$.

The advertising technology is characterized by a continuous and twice differentiable function $f_j(a_{jm})$ that describes the fraction of citizens in the segment m who will become informed about j's ideal position if j purchases a quantity of advertising a_{jm} on advertising in segment m. The remaining fraction in the segment are uninformed. We make the following assumptions about f_j:

$$
\begin{aligned}
0 < f_j(\cdot) &< 1 \\
f_j'(\cdot) &> 0 \\
f_j''(\cdot) &< 0
\end{aligned}
$$

In words, f_j is bounded between 0 and 1, strictly increasing in its argument, and strictly concave in its argument.

We define the parameter $\alpha_j \equiv f_j(0), \alpha_j \in (0,1)$, which determines the level of knowledge voters have about candidate j's position in the absence of advertising. High α_j captures the idea that some candidates may already be more familiar to the electorate than others before any advertising occurs, due e.g. to incumbency status. We will further assume that if $\alpha_{j'} \geq \alpha_{j''}$

for two candidates j', j'', then $f_{j'}(a) \geq f_{j''}(a) \; \forall \; a \in [0, \infty)$,[6] and that the inequality is strict if $\alpha_{j'} > \alpha_{j''}$.

Entry

At the entry stage, candidates simultaneously choose to enter or not, given expectations about which other candidate(s) will enter and the cost of campaigning against those candidates in the advertising stage. A candidate who enters pays a fixed and constant cost c which is sunk before the advertising stage. Following Osborne and Slivinski (1996), we assume that all candidates get utility of $-\infty$ if no candidate enters.

Equilibrium Analysis and Comparative Statics

We consider subgame-perfect equilibria of the game defined above. Given our empirical application, our primary interest will be in the conditions necessary to sustain a one- versus a two-candidate equilibrium. It is clear that in a one-candidate equilibrium, the single entering candidate does not advertise (since she wins for sure regardless, and advertising is costly).

The model at first glance looks identical to Osborne-Slivinski in this case. However, entry of the second candidate may induce the first to want to advertise, and sustaining the equilibrium requires that she want to stay in the race despite both the chance of losing and the increase in costs. A potential second entrant similarly needs to consider not only the fixed entry cost but also the additional expense of advertising she will incur.

An immediate result on ideological convergence Suppose there is a one-candidate equilibrium where only candidate x_0 enters. A necessary condition for such an equilibrium is that no other candidate wants to enter. This implies that:

[6] Note that this implies that the functions $f_{j'}$ and $f_{j''}$ are the same if $\alpha_{j'} = \alpha_{j''}$.

$$-|x_j - x_0| \geq \mathbb{P}(j \text{ wins } |\mathbf{a}_j^*, \mathbf{a}_0^*)(b - c - \mathbf{p} \cdot \mathbf{a}_j^*) + (1 - \mathbb{P}(j \text{ wins } |\mathbf{a}_j^*, \mathbf{a}_0^*))(-|x_j - x_0| - c - \mathbf{p} \cdot \mathbf{a}_j^*) \; \forall j$$

$$\Rightarrow \quad \mathbb{P}(j \text{ wins } |\mathbf{a}_j^*, \mathbf{a}_0^*) \left[|x_j - x_0| + b\right] \leq c + \mathbf{p} \cdot \mathbf{a}_j^* \; \forall j$$

$$\Rightarrow \quad \Phi_j(\mathbf{s}(\mathbf{a}_j^*, \mathbf{a}_0^*)) \left[|x_j - x_0| + b\right] \leq c + \mathbf{p} \cdot \mathbf{a}_j^* \; \forall j$$

Note that the last inequality above is identical to the Osborne and Slivinski (1996) condition except for the addition of the (weakly positive) term $\mathbf{p} \cdot \mathbf{a}_j^*$ to the right-hand side and the substitution of $\Phi_j(\mathbf{s}(\mathbf{a}_j^*, \mathbf{a}_0^*)) \leq 1$ for 1 on the left-hand side. The necessity of advertising thus implies that the ideological range of candidates who can sustain a one-candidate equilibrium (weakly) widens compared to the case of perfect information: any x_0 that satisfies the Osborne-Slivinski condition for all j will also satisfy the above for all j.

Two comparative statics We will consider two further comparative statics, corresponding to the effects of the introduction of digital advertising on candidate competition. First, we consider the effect of reducing the effective price of advertising (p) down from ∞. The $p = \infty$ case approximates the situation faced by candidates in a district which makes up only a small fraction of the containing TV market; in such a district the effective cost of advertising can be extremely high, as from the candidate's perspective a large majority of advertising impressions may be wasted on TV viewers who cannot vote in the candidate's district. The introduction of digital advertising, which can be targeted to smaller geographical areas, allows candidates to cost-effectively reach their specific electorate. In analyzing this change, we abstract from the possibility of targeting within the district; that is, we assume that the targeting partition $\mathbb{A}$ has cardinality one.

Second, we consider the effects of increasing the fineness of the partition $\mathbb{A}$, holding price constant. This change corresponds to the ability of digital advertising to better target voters on ideological grounds, compared to broadcast advertising.

For both of these comparative static analyses, we consider the effects on both the number of candidates who enter a race in equilibrium, and their ideological distribution relative to the distribution of voters in the district.

Price Decline

In the extreme case $\mathbf{p} \to 0$ and for sufficiently steep Φ, our model reduces to that of Osborne and Slivinski (1996).[7] Their results show that one-candidate equilibria with non-median candidates exist only when $b < c$, and one-candidate equilibria of any kind exist only when $b \leq 2c$. Both of these properties will fail to hold when p is sufficiently large: in fact, there exists at least one configuration of the α_j's that sustains a one-candidate equilibrium for *any* candidate, regardless of their position relative to the median. The most extreme candidates require a very asymmetric configuration of α's to sustain a one-candidate equilibrium, while more moderate candidates have a broader set of such configurations that sustain them as the only entrant.

To see this, consider what happens when $\mathbf{p} \to \infty$. When price is sufficiently high, no candidate will be willing to advertise (since the cost of this is infinite, and the benefit is bounded). Vote shares are therefore given by $\mathbf{s}(\mathbf{0})$. In a potential two-candidate race, the share of the vote obtained by candidate x_j against x_0 when neither candidate advertises is:

$$\begin{aligned} s_j(0,0) &= \alpha_j\alpha_0\tilde{G}\left(\frac{x_j + x_0}{2}\right) + \alpha_j(1-\alpha_0) + (1-\alpha_j)(1-\alpha_0)\frac{1}{2} \\ &= \alpha_j\alpha_0\left(\tilde{G}\left(\frac{x_j + x_0}{2}\right) - \frac{1}{2}\right) + \frac{1}{2}(\alpha_j - \alpha_0) + \frac{1}{2} \end{aligned}$$

Where $\tilde{G}(\cdot) \equiv 1(x_j < x_0)G(\cdot) + 1(x_j \geq x_0)(1 - G(\cdot))$. Obviously if we set $\alpha_0 = 1, \alpha_j \to 0\ \forall j \neq 0$, then $s_j \to 0$ for any entrant against x_0, every such j loses (nearly) for sure, and therefore no other candidate will be willing to enter. So there is at least one configuration of

[7] The same obviously holds if, instead of $p \to 0$, all α_j's approach one; both of these limiting cases are equivalent to the baseline full-information model.

the α's that sustains a one-candidate equilibrium where only x_0 enters, for any x_0 and for any choice of c and b. In general, the set of such α's takes the form

$$\left\{\alpha \;:\; \max_j \Phi(2s_j(0,0)-1)(b+|x_j-x_0|)<c\right\}$$

which depends on the shape of the function Φ as well as the entry cost and office benefit parameters. In the limiting case where Φ is the step function at one-half, a necessary condition for entry is $s_j(0,0) \geq \frac{1}{2}$, which simplifies the condition on α that sustains x_0 as the only entrant to the more easily interpretable expression:

$$\left\{\alpha \;:\; \alpha_0 > \max_{j\neq 0} \frac{\alpha_j}{1-\alpha_j\left(2\tilde{G}\left(\frac{x_j+x_0}{2}\right)-1\right)}\right\}$$

So, a candidate x_0 at the median (where $\tilde{G}\left(\frac{x_j+x_0}{2}\right)$ is at most one-half) can sustain a one-candidate equilibrium regardless of c and b whenever $\alpha_0 > \max_{j\neq 0}\alpha_j$. A very extreme candidate (where $\tilde{G}\left(\frac{x_j+x_0}{2}\right)$ approaches one for all j) can sustain a one-candidate equilibrium if $\alpha_0 > \max_{j\neq 0}\frac{\alpha_j}{(1-\alpha_j)} > \max_{j\neq 0}\alpha_j$.[8] In general, the farther from the median is x_0, the more stringent is the condition on α_0 relative to the other candidates' α's required to sustain x_0 as the only entrant.

We have shown that one-candidate equilibria exist for a much wider range of cost and office benefit parameters when $p \to \infty$ compared to the baseline case with complete information. These one-candidate equilibria furthermore involve a much wider range of candidate positions entering in equilibrium than in the baseline case. Because of the continuity in p of the participation constraint on entry, for p sufficiently close to 0, the Osborne-Slivinski equilibrium characterization obtains in our model. Hence, we can state our first comparative static

[8]Note that this is possible only if the second-largest α_j is less than one-half.

predictions:

Hypothesis 1. *As the effective price of advertising falls away from* ∞*, we expect the number of one-candidate races in equilibrium to fall, and the number of candidates entering in equilibrium to rise.*

Hypothesis 2. *As the effective price of advertising falls away from* ∞*, we expect the average ideological distance of entering candidates from the median of the voter distribution to fall.*

Targeting Precision Increase

We next consider what happens to candidate entry when, holding price fixed at some finite level p, we increase the fineness of the partition of the district with which the candidates can target advertising. To understand this, we first need to consider the advertising stage of the game. For tractability, we will consider the two-candidate case; that is, we assume two candidates x_j and x_k have decided to enter. At this stage, the entry cost c is sunk. The candidates solve:

$$\max_{\mathbf{a}_j} \Phi_j(2s_j(\mathbf{a}_j, \mathbf{a}_k) - 1)[b + |x_j - x_k|] - \mathbf{a}_j \cdot \mathbf{p}$$
$$\max_{\mathbf{a}_k} \Phi_k(1 - 2s_j(\mathbf{a}_j, \mathbf{a}_k))[b + |x_j - x_k|] - \mathbf{a}_k \cdot \mathbf{p}$$

The symmetry of the game implies that for any segment m in which both candidates advertise,

$$\frac{\partial s_{jm}}{\partial a_{jm}} = -\frac{\partial s_{jm}}{\partial a_{km}}$$

in equilibrium. The share of candidate j in segment m is:

$$s_{jm} = f_j(a_{jm}) f_k(a_{km}) \left(\tilde{G}_m \left(\frac{x_j + x_k}{2} \right) - \frac{1}{2} \right) + \frac{1}{2}(f_j(a_{jm}) - f_k(a_{km})) + \frac{1}{2}$$

where G_m is the conditional distribution of voter ideal points formed by restricting G to the segment m, and $\tilde{G}_m$ is defined as above. To simplify the notation we define $\gamma_m^{j,k} \equiv \tilde{G}_m \left(\frac{x_j + x_k}{2} \right)$ to be the full-information vote share that j obtains against k in segment m. Thus:

$$\frac{f_j'(a_{jm}^*)}{f_k'(a_{km}^*)} = \frac{1 - f_j(a_{jm}^*) \left(2\gamma_m^{j,k} - 1 \right)}{1 + f_k(a_{km}^*) \left(2\gamma_m^{j,k} - 1 \right)} \tag{1}$$

By inspection, the ratio above is less than 1 if and only if j is majority-preferred to k in segment m ($\gamma_m^{j,k} > \frac{1}{2}$). Because of our assumptions about the curvature of f_j, this implies that a greater share of voters must be informed about j than k in segment m when j is majority-preferred there; if $\alpha_j \leq \alpha_k$ then j necessarily spends more on advertising in such a segment than does k.

This is our third hypothesis:

Hypothesis 3. *The fraction of total spending in a given targeting segment made by one candidate in a race is increasing in that candidate's fraction of supporters in the segment under full information.*

When $\mathbb{A} = \{\mathbb{N}\}$, that is, when the candidates must advertise to the entire district, then a majority-preferred candidate with equal initial voter awareness (i.e. $\alpha_j = \alpha_k$) always earns a higher vote share district-wide than her opponent (and a majority-preferred candidate disadvantaged in initial voter awareness always closes the gap with her advantaged opponent, relative to the no-advertising case). Both of these properties will fail to hold, in general, when advertising can be targeted ideologically.

To see this, note that the first-order condition on a candidate's maximization problem implies that for any segment in which candidate j spends a positive amount,

$$\phi(2s_j - 1)[b + |x_j - x_k|]\mu_m \frac{\partial s_{jm}}{\partial a_{jm}} = \mu_m p$$

from which the μ_m appearing on both sides cancels. Advertising levels are thus determined only by the candidate's relative advantage in the segment and the amount that the opponent advertises in the segment:

$$f_j'(a^*_{jm}) = \frac{p}{\phi(2s_j - 1)[b + |x_j - x_k|]} \left(\frac{1}{1 + f_k(a^*_{km}) \left(2\gamma_m^{j,k} - 1\right)} \right)$$

We showed previously that post-advertising vote shares increase, relative to the no-advertising baseline, for the candidate who is majority-preferred by voters *in that segment*. This does not guarantee that the district-wide majority-preferred candidate under full information will necessarily gain share of the district-wide vote relative to the no-advertising case unless he is also majority-preferred in every available segment $\left(\gamma_m^{j,k} > \frac{1}{2}\ \forall m\right)$.

The candidate's total vote share is a linear combination of segment shares with weights μ_m ($s_j = \sum_m \mu_m s_{jm}$). μ_m appears nowhere in the expression for candidates' optimal advertising strategy, implying that segment-level shifts induced by advertising are unrelated to μ_m and hence the effect of advertising on total vote share is ambiguous. Depending on the specific structure of the partition $\mathbb{A}$, either candidate may be advantaged relative to the benchmark of broadcast ($\mathbb{A} = \{\mathbb{N}\}$) advertising.

Marginal efficiency To make precise the notion of advantage due to targeting, we define a candidate's *marginal advertising efficiency* e_j as the marginal effect of increasing the candidate's advertising budget, W_j, on her vote share, s_j. Because $s_j = \sum_m \mu_m s_{jm}$, for any given budget level W, there is generically one segment $m^*(W) = \arg\max_m \frac{\partial s_{jm}}{\partial a_{jm}}\big|_{a^*_{jm}(W)}$ which yields the maximum marginal gain in votes for j. Efficiency is just the marginal gain in votes that j gets from infinitesimally increasing her budget allocation to m^*, holding expenditures in all other segments equal:

$$e_j(W) \equiv \left.\frac{ds_j}{dW_j}\right|_W = \mu_{m^*(W)} \frac{ds_{jm^*(W)}}{da_{jm^*(W)}} \frac{da_{jm^*(W)}}{dW_j}$$

Given the assumption that advertising can be purchased at constant marginal cost $\mu_m p$, $\frac{da_{jm^*(W)}}{dW_j} = \frac{1}{\mu_{m^*(W)} p}$. Multiplying by $\frac{ds_{jm^*(W)}}{da_{jm^*(W)}}$ gives:

$$e_j(W) = \frac{1}{p} f'_j\left(a^*_{jm^*(W)}\right)\left[f_k\left(a^*_{km^*(W)}\right)\left[\gamma^{j,k}_{m^*(W)} - \frac{1}{2}\right] + \frac{1}{2}\right] \tag{2}$$

Note that j's marginal advertising efficiency is strictly (because $f_k(\cdot) > 0$) increasing in $\gamma^{j,k}_{m^*(W)}$, which is the share of j's supporters in the segment in $\mathbb{A}$ where the marginal impact of advertising is at its maximum. If $\mathbb{A}$ is a singleton, then, all else equal, a candidate closer to the district-wide median (with higher $\tilde{G}\left(\frac{x_j+x_k}{2}\right)$) has higher marginal advertising efficiency than a candidate farther away. In any other case this need not hold, as what matters is the candidate's share in the segment of the partition *in which her underlying support is at its maximum.* The maximum share of j against k across the cells of the partition can be greater than the maximum share of k against j, even if k is preferred to j district-wide.

An illustrative example An example of how the alignment of the partition $\mathbb{A}$ with candidates' positions can advantage majority-dispreferred candidates is given in Figure 1. In this example, candidates j and k located at x_j and x_k respectively compete in a district whose median voter is x_m and in which there are two targetable segments available, A and B, depicted as regions between the dashed lines. Although k is preferred by the majority district-wide, in segment A

all voters prefer j. In segment B, j's share is $\frac{1}{5}$.

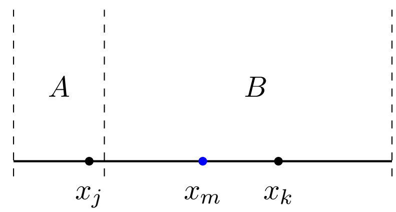

Figure 1: A pathological targeting example.

We suppose $f_j = f_k = f$ (which implies $\alpha_j = \alpha_k = \alpha$), and that initially neither candidate advertises ($a_{jA} = a_{jB} = a_{kA} = a_{kB} = 0$). At this point, the candidates' marginal efficiencies of advertising are:

$$e_j(0) = f'(0)\left(\frac{1+\alpha}{2}\right) > f'(0)\left(\frac{1+0.6\alpha}{2}\right) = e_k(0)$$

In other words, j initially has an efficiency advantage: the same fundraising effort goes farther for her than it does for k, in spite of the fact that a majority of voters in the district prefer k. In contrast, k would have an efficiency advantage if the candidates were required to advertise to the whole electorate $A \bigcup B$.

Average efficiency Of course, at the entry stage, what matters is not the marginal but the average advertising efficiency. The candidate must consider the total return, in votes, of the funds she expects to spend in the campaign stage. This is just the integral of (2) from 0 to the amount that the candidate expects to spend in equilibrium:

$$\mathcal{E}_j(W) = \int_0^W \frac{1}{p} f_j'\left(a^*_{jm^*(W')}\right)\left[f_k\left(a^*_{km^*(W')}\right)\left[\gamma^{j,k}_{m^*(W')} - \frac{1}{2}\right] + \frac{1}{2}\right] dW' \tag{3}$$

To evaluate the integral in (3), recall that at any spending level there is generically one maximally-efficient segment in which to increase spending; this fact implies that there is a sequence of segments $\{m_1, m_2, \ldots, m_{L(W)}\}$ in which j will increase her spending as her total budget increases, given the distribution of support across segments and her opponent's spending. Combined with the substitution $dW' = p\mu_{m^*(W')} da_{jm^*(W')}$, we have:

$$\mathcal{E}_j(W) = \sum_{l=1}^{L(W)} \mu_{m_l} \left[f_j\left(a_{jm_l}^{l}\right) - f_j\left(a_{jm_l}^{l-1}\right)\right] \left[f_k\left(a_{km_l}^{*}\right) \left[\gamma_{m_l}^{j,k} - \frac{1}{2}\right] + \frac{1}{2}\right] \tag{4}$$

Where $a_{jm_l}^{l}$ is j's equilibrium advertising level in segment m_l in the event that W is large enough to reach the lth element of the sequence $\{m_1, m_2, \ldots, m_{L(\infty)}\}$.

The expression for j's efficiency in (4) contains equilibrium quantities: it depends on the behavior of the opponent k as well as the equilibrium total expenditure W. However, it is increasing in the product $\mu_m \gamma_m^{j,k}$ for segments m that appear early in the optimal sequence for j, that is, the segments with the highest $\gamma_m^{j,k}$.[9]

Higher efficiency $\mathcal{E}_j$ directly translates into greater probability of entry, as the cost of achieving any given vote share, and hence the equilibrium entry cost, is decreasing in the candidate's efficiency. Hence, we have our last hypothesis:

Hypothesis 4. *The introduction of targeted advertising encourages the entry of candidates with larger* maximum *values of $\mu_m \gamma_m$ (the product of segment size and the candidate's share of supporters in the segment) across the targetable segments of the electorate $\mathbb{A}_m$.*

Data

Effective Price of TV Ads

To measure the difficulty of broadcast advertising across districts, we construct a measure of the effective price of TV advertising at the congressional district (CD) level or state legislative

[9] Note that the incentive to advertise in the candidate's strongest segments *strengthens* if the opponent increases his advertising in those segments, so the rank-order is not affected by opponent behavior.

district (SLD) level. TV advertising is sold in units of 1000 "impressions," or viewer-exposures. Regulations on television carriage in the US enforce that all cable system operators in a given media market, referred to as a Designated Market Area, or DMA, carry the set of broadcast TV channels operating within that DMA. The result is that the impressions purchased by an advertiser on a broadcast TV station could be seen by anyone within the geographic boundaries of its DMA.[10] Because media market boundaries generally do not align with political boundaries, the result for political advertisers is unintended spillovers of advertising outside of the area to which the candidate would like to advertise.[11] These wasted impressions effectively raise the cost of advertising in districts with poor congruence with their containing media market(s).

We use per-impression prices at the media market level, collected from two sources: pre-2010 years come from the data provider Spot Quotations and Data (SQAD), while years 2010 on come from the Nielsen Ad Intel dataset. We reweight these prices into an effective district-level price that politicians need to pay to reach voters in their district if they are to advertise on TV. The effective price P_{ij} that a candidate in district i needs to pay to reach voters in the overlapping area between DMA j and her district is:

$$P_{ij} = \frac{\pi_j}{\pi_{ij}} P_j$$

Where π_{ij} is the number of voters in the intersection, and π_j is the total number of voters in DMA j. Aggregating over all DMAs that overlap with her district gives us the effective district-level TV price:

$$P_i = \sum_j P_{ij} \frac{\pi_{ij}}{\pi_i} = \sum_j \frac{\pi_j}{\pi_i} P_j \tag{5}$$

The effective price is thus higher, when the district i makes up a small fraction of the

[10]DMAs typically encompass multiple counties, including a central metropolitan or micropolitan area and its surrounding counties.

[11]These spillovers have been used in multiple papers (Spenkuch and Toniatti, 2018; Huber and Arceneaux, 2007; Ashworth and Clinton, 2007, among others) to estimate effects of advertising on voters.

population of the DMA(s) with which it overlaps. The ratios π_j/π_i can in practice be quite high, creating wide variation in effective cost across districts that ranges from close to parity with commercial advertisers to 10 or 20 times as high.

Meta (Facebook / Instagram) Ads

To measure targeting of digital ads, we rely on two datasets produced by Meta (the parent company of the Facebook and Instagram social platforms) for use by academic researchers. First, the Ad Targeting dataset, which has advertisement-level information on advertiser-specified targeting characteristics such as gender, age, geographical location, interests or behavior of users whom the advertiser wants to target with the ad. The Ad Targeting data also contains a variable with the estimated potential reach of the ad, which is the number of users who meet the targeting criteria. Second, the Ad Library dataset, which contains information on the content of the ad, the number of impressions that the advertiser purchased, and the estimated cost of placing the ad. The Ad Targeting and Ad Library datasets can be matched together using a unique ID, allowing us to measure the amount that candidates spent to reach different audiences. The Ad Library covers political advertisements beginning in May of 2018, while Ad Targeting is newer, covering ads published on or after August 3, 2020.[12] Our data will thus focus on candidates in election cycles from 2020 to 2024.[13]

Our goal is to construct a dataset containing total spending and impressions by candidate, election, and targeting cell (the set of features specified in the ad's targeting criteria). To do this, we first collect all candidates' campaign page IDs, and then query the database for ads run by that set of pages. To generate the set of page ID's to query, we used the following procedures for federal and state-legislative candidates.

[12]See https://developers.facebook.com/docs/ad-targeting-dataset/overview.

[13]By the constraint of publishing date, our cycle 2020 Primary measures have a truncated window that apply to all candidates.

Congressional Candidates We start with all political candidates from campaign contribution records and election returns, and match them to the list of all Meta advertisers who have run ads on social issues, elections or politics[14] using candidate and page/disclaimer names. The matched pages are validated in two ways. First, we use a Facebook API scraper to retrieve page information such as text in the intro box, contact information, external websites linked to the page, etc., and then input this information to GPT-4o and Llama 3.1 to validate whether the page corresponds to the matched candidate. Second, we pull the page's advertisements during the election window from the Ad Library and verify that the ad delivery period is consistent with the period in which the candidate was active. The second step is especially useful when the candidate's page is inactive.

State Legislative Candidates We start with a list of candidates compiled by BallotPedia, which also provides the candidates' campaign Facebook page URLs. Because these URLs differ from page IDs, we first use a Facebook API scraper to map URLs to page IDs. This step works only for pages that are still active at the time of scraping. For inactive pages, we follow the same matching and validation procedure described above for congressional candidates.

Other Data

To measure entry and ideological positioning of candidates, we rely on several data sources. We use campaign registration records from the Federal Elections Commission (FEC) to measure the candidate pool in House elections. The database maintained by FollowTheMoney allows us to include state legislative candidates in our analysis as well. The advantage of using campaign finance related data rather than election results is that we are able to capture the comprehensive candidate pool — including candidates who expressed an initial intention to run but dropped out of the race before the primaries. In addition, we use Bonica's (2023)

[14] The full report can be downloaded from https://www.facebook.com/ads/library/report/?source=onboarding. Our version was time stamped with June 2, 2025.

Database on Ideology, Money in Politics, and Elections (DIME) for estimates of candidates' ideological position.

Finally, we use data on candidate expenditures from Sheingate et al. (2022) to measure candidates' total resource allocation to digital versus broadcast advertising. This data covers the period 2004-2020.

Empirical Patterns

Digital vs TV

We first show that digital advertising is a substitute for television advertising, used disproportionately by candidates running in districts whose geography makes television advertising prohibitively expensive. Figure B1 illustrates that there is a strong positive relationship between district-level effective TV ad prices and the percentage of campaign media expenditures spent on digital advertising, as classified by Sheingate et al. (2022). This relationship is also evident in regression form, in column (1) of Table 1. In the lowest-cost districts, candidates spend the large majority of their media budgets on television, while in the highest-cost districts the fractions are reversed. The slope of this relationship flattens somewhat over the 2004-2020 period due to a general increase in the use of digital advertising by all candidates. The cross-district pattern indicates that the introduction of digital advertising can be thought of as in effect a large reduction in advertising costs in the highest-TV-cost districts and a relatively smaller reduction in the lowest-TV-cost districts.

Second, digital and TV ads have different audiences, and consequently the viewers that are easy to reach differ across the two media.[15] Figure B3 examines one of the largest digital platforms used by political advertisers—Meta—and shows that Meta political ads generally reach a younger audience compared to political ads on television. The median campaign ad on TV has 46% of its audience in the 65+ age group, compared to 28% for Facebook

[15] See Gentzkow et al. (2024) for evidence on the impacts of this difference for ad pricing.

/ Instagram ads. This difference implies that in addition to the price decline, digital ads change the available segmentation of the audience. Our model implies that such changes to segmentation should be consequential in encouraging entry by candidates who find themselves with efficiently-reachable supporter groups.

Table 1: The relationship of TV ad prices to digital media use and candidate entry in US House elections.

	Digital Media Share	N. Candidates			Uncontested		
		Total	Rep.	Dem.	General	Dem. Pri	Rep. Pri
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log TV Ad Price	0.0462***	-0.2053***	-0.0750**	-0.1162***	0.0091	0.0425***	0.0127
	(0.0074)	(0.0539)	(0.0370)	(0.0319)	(0.0095)	(0.0120)	(0.0119)
Log TV Ad Price × Post-2012	-0.0075	0.2755***	0.0896**	0.1701***	-0.0030	-0.0569***	-0.0123
	(0.0091)	(0.0645)	(0.0443)	(0.0382)	(0.0114)	(0.0144)	(0.0142)
Pop. Density Controls	✓	✓	✓	✓	✓	✓	✓
Year FEs	✓	✓	✓	✓	✓	✓	✓
Observations	3,819	5,112	5,112	5,112	5,112	5,112	5,112
R^2	0.14	0.05	0.09	0.04	0.08	0.04	0.07

Notes: All regressions include year fixed effects and a cubic spline in log district population density. Digital media share is from Sheingate et al. (2022) and covers 2004-2020. US House candidate data covers 2002-2024. "Post-2012" is an indicator variable equal to one for years 2012 and later. "Log TV Ad Price" is the logarithm of the district-level *effective* ad price per 1000 impressions, calculated as in Equation 5. Standard errors in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 2: The relationship of TV ad prices to candidate entry in state legislative elections.

	N. Candidates			Uncontested		
	Total	Rep.	Dem.	General	Dem. Pri	Rep. Pri
	(1)	(2)	(3)	(4)	(5)	(6)
Log TV Ad Price	-0.0675***	-0.0121**	-0.0519***	-0.0018	0.0265***	0.0100***
	(0.0080)	(0.0054)	(0.0053)	(0.0028)	(0.0022)	(0.0023)
Log TV Ad Price × Post-2012	0.0164*	-0.0169***	0.0495***	-0.0052	-0.0203***	0.0028
	(0.0095)	(0.0063)	(0.0063)	(0.0033)	(0.0026)	(0.0028)
Pop. Density Controls	✓	✓	✓	✓	✓	✓
Year FEs	✓	✓	✓	✓	✓	✓
Observations	58,681	58,681	58,681	58,681	58,681	58,681
R^2	0.010	0.07	0.05	0.03	0.02	0.03

Notes: All regressions include year fixed effects and a cubic spline in log district population density. State legislative candidate data covers 2002-2024. "Post-2012" is an indicator variable equal to one for years 2012 and later. "Log TV Ad Price" is the logarithm of the district-level *effective* ad price per 1000 impressions, calculated as in Equation 5. Standard errors in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Candidate Pool

We now present some results on the effects of digital advertising on candidate entry. We first examine the effect on the size of the candidate pool. We use candidate registration data from the Federal Elections Commission (FEC) to count the number of candidates entering a given congressional or state legislative race, and estimate the relationship between the size of the pool and the effective cost of television advertising in the district. Our model suggests that as advertising costs decline, the number of candidates should increase. We expect that as digital advertising over time becomes a more widely used substitute for TV advertising, the relationship between TV ad prices and candidate entry should weaken.

One complication with this analysis is that effective prices tend to be highest in urban areas, as urban districts can be much smaller in population than their containing metro area. Figure B2 shows that there is a consistent positive relationship between population density and effective price of television ads at the district level. There is still considerable residual variation driven by differences in alignment between district and containing media market, even for districts with similar density, but we want to avoid conflating differential trends between urban and rural places with the effect of high-cost TV markets. We will therefore always control for a flexible function of district population density in estimating the relationship between prices and competitiveness outcomes.

Table 1 shows the estimated relationship between effective TV ad price and competitiveness in the US House. We regress the number of candidates who enter (columns 2-4) and indicators for a race being uncontested (columns 5-7) on our measure of effective TV ad prices, controlling for population density and year fixed effects. We show results for the total number of candidates and entrants within each party. We interact the price measure with an indicator for the period from 2012-on when digital advertising becomes mainstream, to allow the importance of TV ads to decline as digital ads become available.

The results are broadly consistent with effective TV ad prices being a strong predictor of competition in the earlier years when digital advertising is less common or not available to political candidates. That relationship largely disappears after 2012.[16] This pattern implies that digital ads provide a lower-cost medium for candidates to reach voters, equalizing political competition across high- and low-TV-cost districts.

Table 2 presents the same regressions as Table 1, for state legislative candidates.[17] The magnitudes of the coefficients on effective TV ad prices here are smaller, likely because advertising is prohibitively expensive for most candidates in these races. The pattern of high TV ad prices being associated with lower competition in the early years is nonetheless present here as well. Unlike in the House, however, the post-2012 attenuation is concentrated among Democratic candidates.[18]

Our second set of outcome variables captures candidate characteristics. We examine candidates' positioning on the left-right axis, as measured by campaign finance scores (CFscores, Bonica (2023)). Our model implies that the price effect of digital advertising's introduction should lead to greater congruence between district and candidate ideology. The effect of more precise targeting, however, is ambiguous, and could help candidates at the ideological extremes if ads are easily targeted to viewers on ideological grounds.

In Table 3, we examine the relationship of the average ideology of candidates running for the House and for state legislatures to television prices over time, as in our previous tables. Because the meaning of "extreme" on the CFScore axis differs by party, we show these separately for Democratic and Republican candidates. Figures B8 and B9 show the corresponding patterns in the data visually, by year. Note that the Bonica CFScore measure that we use here is constant within candidate, and so change in district averages necessarily must come from entry of new candidates. We also include a measure of district-level preferences from Warshaw and Tausanovitch (2022), such that what we capture in the regressions is deviation from expected

[16]We cut the sample into 2002-2010 and 2012-2024 periods for parsimony in the regression specification. Figures B4-B5 show the relationship each year.

[17]With the exception of the share of digital spending, which is not available for state legislative candidates.

[18]Figures B6 and B7 show year-by-year visual analogs of the regression in Table 2.

CFScore given the district's electorate.

For Democratic candidates, the pattern looks similar to the entry regressions. In the earlier years of the sample, high television prices are associated with relatively extreme (left) positions among Democratic candidates, but this pattern flattens over time.

For Republicans, the baseline relationship is different between state legislative and US House races. In both cases, however, the slope of the relationship actually increases over time. This increase is especially evident for state legislative candidates; see Figure B9.

Table 4 reports the same regressions using an alternative measure, DW-DIME scores (Bonica, 2018). Because DW-DIME scores map contribution patterns directly onto predicted roll-call voting behavior, they might better capture the policy positions candidates take in office than CFScores. The results are similar to those in Table 3, with the exception of column (4): the baseline relationship for Republicans in the state legislature is now negative, which is consistent with the pattern for Republicans in US House.

While the pattern for Democratic candidates looks consistent with the prediction for price reduction (as high-TV-cost districts get less extreme candidates more in line with district preferences after digital ads become available), more-extreme Republican candidates seem to be entering differentially *more* in places where digital advertising is more attractive. This pattern is suggestive that the balance between the price effect and the precision effect is different between the two parties, a possibility we take up in the next section.

Table 3: The relationship of TV ad prices to candidate CFScores in US House and state legislative elections.

	US House		State Legislatures	
	Mean CFScore			
	Dem. (1)	Rep. (2)	Dem. (3)	Rep. (4)
Log TV Ad Price	-0.0805***	-0.0521***	-0.0313***	0.0060**
	(0.0079)	(0.0086)	(0.0025)	(0.0028)
Log TV Ad Price × Post-2012	0.0675***	0.0146	0.0300***	0.0080**
	(0.0093)	(0.0105)	(0.0029)	(0.0034)
Pop. Density Controls	✓	✓	✓	✓
Tausanovitch-Warshaw District Ideology	✓	✓	✓	✓
Year FEs	✓	✓	✓	✓
Observations	4,767	4,509	36,896	36,752
R^2	0.12	0.11	0.08	0.03

Notes: All regressions include year fixed effects and a cubic spline in log district population density. DIME candidate ideology data covers 2002-2022. "Post-2012" is an indicator variable equal to one for years 2012 and later. "Log TV Ad Price" is the logarithm of the district-level *effective* ad price per 1000 impressions, calculated as in Equation 5. Standard errors in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Ad Targeting on Facebook and Instagram

In this section, we describe how candidates advertise on Facebook and Instagram, where within-district targeting is much easier than on television. Meta's Ad Targeting dataset (Meta Platforms, Inc., 2026a) provides advertisement-level information on the targeting features that candidates select when placing ads, including geographic areas, age and gender ranges, interests, behaviors, and custom audience lists. Because these features are observed at the individual ad level, we can characterize in detail how candidates segment the electorate, i.e. which subgroups they choose to reach, how narrowly they define their audiences, and whether rivals in the same race target overlapping or distinct sets of voters. The targeting choices that candidates make reveal which features of the electorate they believe give them an advantage: a candidate who narrows her ads to a particular age group or geographic area is implicitly

Table 4: The relationship of TV ad prices to candidate DW DIME scores in US House and state legislative elections.

	US House		State Legislatures	
	Mean DW DIME Score			
	Dem. (1)	Rep. (2)	Dem. (3)	Rep. (4)
Log TV Ad Price	-0.0586***	-0.0146	-0.0351***	-0.0125***
	(0.0055)	(0.0106)	(0.0027)	(0.0040)
Log TV Ad Price × Post-2012	0.0645***	0.0002	0.0351***	0.0341***
	(0.0065)	(0.0131)	(0.0031)	(0.0048)
Pop. Density Controls	✓	✓	✓	✓
Tausanovitch-Warshaw District Ideology	✓	✓	✓	✓
Year FEs	✓	✓	✓	✓
Observations	4,312	3,848	25,530	25,120
R^2	0.36	0.16	0.22	0.19

Notes: All regressions include year fixed effects and a cubic spline in log district population density. DIME candidate ideology data covers 2002-2022. "Post-2012" is an indicator variable equal to one for years 2012 and later. "Log TV Ad Price" is the logarithm of the district-level *effective* ad price per 1000 impressions, calculated as in Equation 5. Standard errors in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

betting that her message resonates more strongly with that segment than with the electorate as a whole. This is precisely the logic captured in our model, in which the ability to direct advertising to segments of the population creates differential incentives for candidates depending on the distribution of their support across those segments.

Table 5 summarizes our data coverage. Across the 2020–2024 cycles, we identify over 30,000 unique state legislative candidates and over 3,600 congressional candidates. Roughly two-thirds of state legislative candidates have a Facebook page. Among candidates we were able to match to a valid Facebook page ID, around 34% advertised on Facebook or Instagram during primaries, and around 50% did so during generals. Among the Congressional candidates, we were able to match around 58% of candidates to a valid Facebook page ID. Out of the matched sample, 59% advertised on Facebook or Instagram during primaries, and around 85% did so during generals.

Table 5: Meta Platforms Advertising by State Legislature and Congressional Candidates 2020-2024

	State Legislature		**US House**		**US Senate**	
	Primary	General	Primary	General	Primary	General
Unique candidates	30,664	23,325	3,614	1,636	510	238
Candidates with Facebook pages	19,691	15,273	–	–	–	–
Candidates with matched Facebook pages	14,508	13,321	2,087	1,022	305	160
Candidates who advertised on Facebook	4,897	6,625	1,230	854	191	149

Notes: Counts are based on unique candidates. We exclude special elections or elections that were canceled. A Facebook page is considered *matched* if the Facebook URL can be successfully linked to a valid Facebook page ID. This row is specific to state legislative candidates since we use a compiled list from BallotPedia. This row is not filled for US House and Senate candidates because we constructed a Facebook account list directly via page ID searches. Ads run by candidates between Dec. 1st of the year before the election year and the primary election date are counted as primary election advertising. Ads run by candidates between the day after the primary election date and the general election date are counted as general election advertising.

Tables 6 and 7 report the share of ad impressions run by candidates in either group employing each targeting method and the associated cost per impression at the aggregate level. Geographic targeting is the most prevalent: 65–69% of congressional impressions and 76–81% of state legislative impressions use sub-national geographic targeting (e.g., targeting impressions to users residing in a specific state). Sub-state targeting, which narrows the target audience to a region smaller than the state level (e.g., a specific city or collection of zip codes), is used for 34–38% of congressional ad impressions and 65-67% of state legislative ad impressions. The substantially higher rate among state legislative candidates is natural: these candidates represent smaller districts and are less likely to be seeking out-of-state donations than congressional elections, making geographic precision particularly valuable for eliminating wasted impressions. What we call custom audiences, which includes Meta's categories of custom, lookalike, and connection-based user lists,[19] account for 47–57% of impressions, while

[19]A Meta "custom audience" is a list of user accounts or email addresses provided to Meta by the advertiser, which the advertiser specifies as the audience for the ad. Meta delivers these ads only to users in the advertiser's provided list. "Lookalike" audiences begin from a custom audience but allow Meta to deliver the ads to other users similar to those in the provided list. "Connection" audiences similarly begin from a custom audience but allow Meta to deliver the ad to other users who are connected through the friend or follow relationships to those in the provided list.

interest-based targeting (interest, behaviors, occupation, education, etc.) is less common at 21–28%.[20] In general, ads involving custom audiences are more expensive. Facebook ads also appear to be more expensive during general elections than during primary elections.

Table 6: Meta Platforms Advertising by Congressional Candidates, 2020-2024

	Primary		**General**	
	% Impressions	Avg. $ / Impression	% Impressions	Avg. $ / Impression
Any custom audience	56.8	0.0272	57.2	0.0300
Any sub-national geographic targeting	64.7	0.0178	69.0	0.0213
Any sub-state geographic targeting	34.0	0.0156	37.5	0.0198
Any interest targeting	25.8	0.0160	28.3	0.0229

Notes: Ads run by candidates between Dec. 1st of the year before the election year and the primary election date are counted as primary advertising. Ads run by candidates between the day after the primary election date and the general election date are counted as general advertising. Custom audience includes any use of custom, lookalike or connection Facebook user list. Interest targeting includes any targeting based on interests, behaviors, occupation, and education background.

Table 7: Meta Platforms Advertising by State Legislative Candidates, 2020-2024

	Primary		**General**	
	% Impressions	Avg. $ / Impression	% Impressions	Avg. $ / Impression
Any custom audience	47.4	0.0201	48.4	0.0202
Any sub-national geographic targeting	76.4	0.0146	81.3	0.0162
Any sub-state geographic targeting	64.6	0.0136	67.3	0.0152
Any interest targeting	20.9	0.0132	21.4	0.0160

Notes: Ads run by candidates between Dec. 1st of the year before the election year and the primary election date are counted as primary advertising. Ads run by candidates between the day after the primary election date and the general election date are counted as general advertising. Custom audience refers to ads that include any use of custom, lookalike or connection Facebook user list. Interest targeting includes any targeting based on interests, behaviors, occupation, and education background.

A central prediction of our model (Hypothesis 3) is that targeted advertising leads candidates to concentrate messages in segments that are likely to be favorable, which is corroborated by descriptive patterns in the Meta data. Figure B10 shows that a significant portion of candidates primarily or even only advertise to custom audiences. Custom audiences are based on a list

[20] Allcott et al. (2026) report an even higher share of custom-audience targeting, at 62% for Facebook and 71% for Instagram ads. The higher share is likely driven by their inclusion of presidential advertising, where the share of custom-audience targeting is especially high: 73% and 90% for Facebook and Instagram respectively.

of users collected by the campaign itself (often through the user making a contribution or participating in campaign events or email lists). Given how campaigns gather these lists, they are very likely to be made up of people inclined to support the campaign: they either have already supported the campaign in the past, or are similar to or have friend relationships with people who do. Custom audiences have grown notably more important in recent cycles.

A second prediction of the model is that candidates will only overlap (both candidates advertising to the same cell of the audience partition) when they are well-funded. This is because candidates sort elements of the partition in order of favorability, and proceed from opposite ends as their budget grows. Figure B11 shows that nearly all candidates in our sample advertise only to exclusive targeting cells, effectively partitioning the electorate into non-overlapping segments. Here, we define a targeting cell as the full combination of targeting elements specified in an ad: age, gender, geographic areas, interests, behaviors, occupation, and education. A cell is exclusive if, in a multi-candidate race, only one candidate delivers ads to it. We treat custom audiences as exclusive by definition. Table B1 aggregates this pattern to the race level. Among races in which at least two candidates advertised on Facebook or Instagram, few feature any targeting cell reached by more than one candidate, and the share is even lower in state legislature races.

Figure 2 examines the relationship between candidates' financial resources and targeting precision. In this figure, we plot campaign contribution received by candidates against the average potential reach of Meta ads rescaled by district size. The potential reach of an ad is the size of the population that could be reached by that targeting combination, and hence can be considered as a measure for the narrowness of the audience in the target cell.

Congressional candidates who raise more money in total target narrower audiences. The pattern is even stronger when we focus attention on ads whose primary goal is not fundraising.[21] The slope is correspondingly much flatter for fundraising ads. When candidates are

[21] We define fundraising ads as ads with a call-to-action link containing the words "donate," "contribute," or "give."

fundraising, they cast a wider net, but when they aim to persuade or turn out voters, they prefer narrower targeting cells as long as their budgets allow. For state legislative candidates, the negative relationship holds for both fundraising and non-fundraising ads, perhaps because these races are less nationalized and the donor pool is also localized. These patterns connect to our model's efficiency concept. The optimal strategy for a given candidate is to allocate first to the segment where the marginal vote gain per advertising dollar is highest, and then move to the next-best options. Well-resourced candidates are able to target more precise (hence lower potential reach) audiences.

Figure 3 shows a similar story but in extensive margins: well-funded candidates target more. In this figure, we plot the share of a candidate's non-targeted impressions against campaign donation. We define a non-targeted ads as those that use neither a custom audience nor any interest or behavior based targeting; such ads may still have criteria on age, gender, and geographical location. Candidates who raise more money deliver a smaller share of their impressions through non-targeted ads.

Finally, Table 8 lists the top 50 most common targets, in terms of dollars spent on ads targeted to users who have that feature, by candidates in Congressional elections.[22] The second and third columns in this table report the net share by spending or unique candidates, respectively, of ads that target this feature that are run by Republican candidates. Cells where the net Republican share exceeds 0.25 are colored red, and cells where the net Republican share is less than -0.25 are colored blue.

Notably, there are only three targets in the top 50 by dollars that are not overwhelmingly targeted by candidates from one party: engagement with moderate political content, being interested in Politics, and being interested in Golf. The vast majority of targeted interests are highly partisan in the sense that candidates targeting that feature mostly come from one party or the other.

[22]Note that, per Tables 6 and 7, ads that target *any* interest or behavior are a relatively small minority, roughly 21% to 28%, of ads run by candidates for office.

Figure 2: Well-funded candidates target narrower audiences.

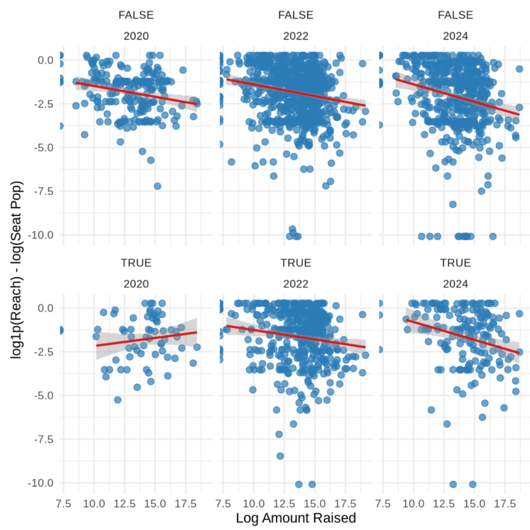

(a) Congressional primary elections

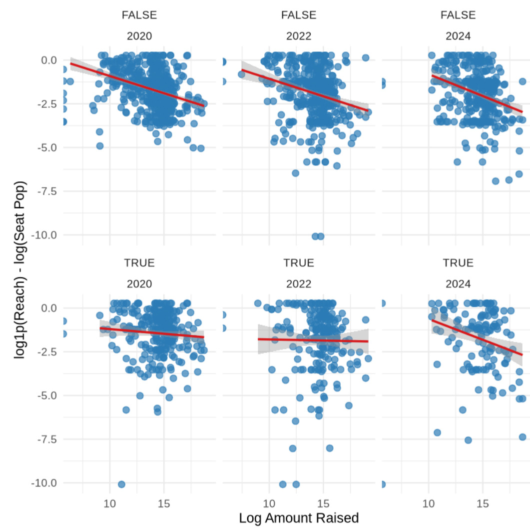

(b) Congressional general elections

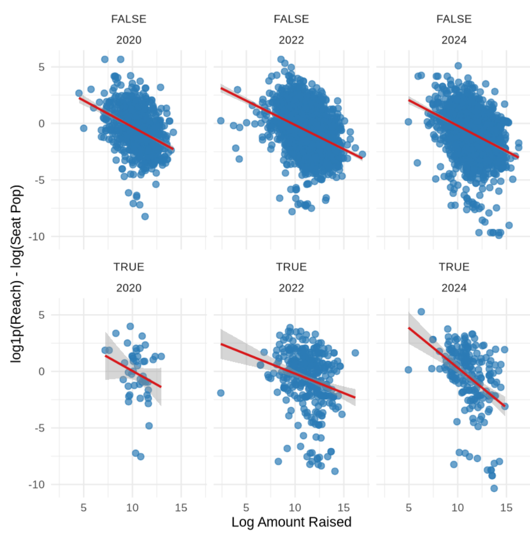

(c) State legislative primary elections

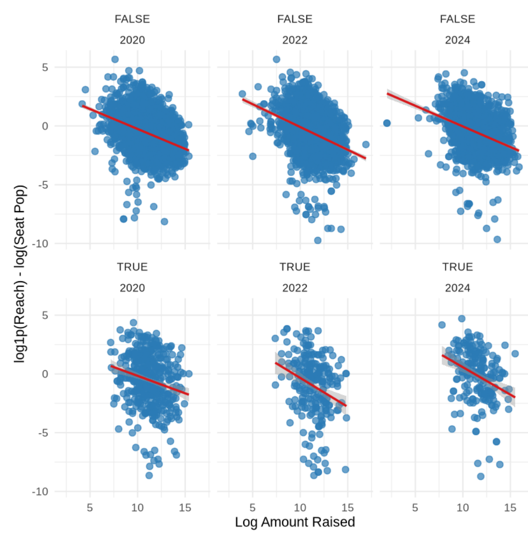

(d) State legislative general elections

Notes: Each panel shows the relationship between a candidate's total amount raised (in log terms) and the average potential reach of their Meta ads, as a fraction of the district population, with a linear fit and 95% confidence band. Panels are faceted by election cycle and whether the ad contains fundraising content. Potential reach is measured as the midpoint of the lower and upper bounds reported by Meta, averaged across a candidate's ads using impression-weighted means. Data source: Meta Ad Library, Ad Targeting Dataset, and DIME.

Figure 3: Well-funded candidates target more.

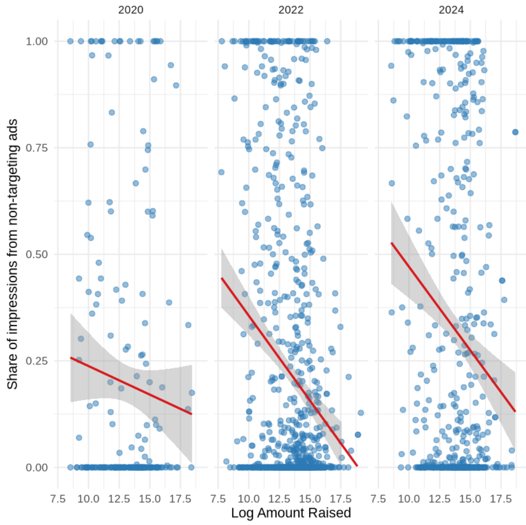

(a) Congressional primary elections

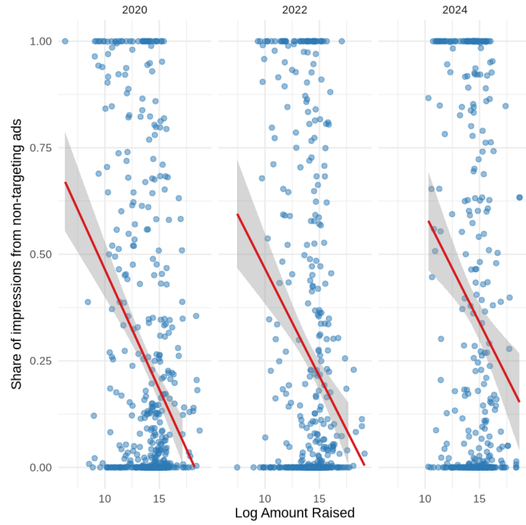

(b) Congressional general elections

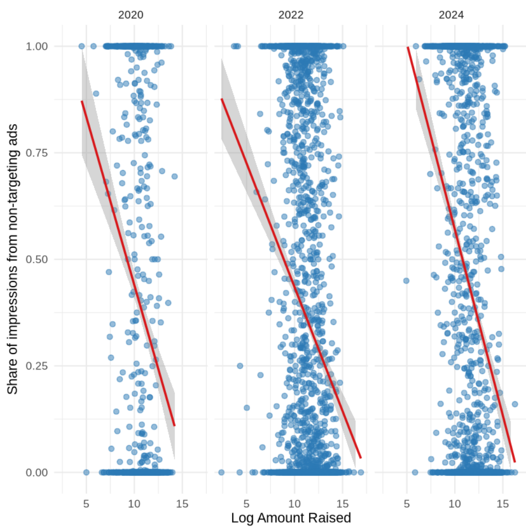

(c) State legislative primary elections

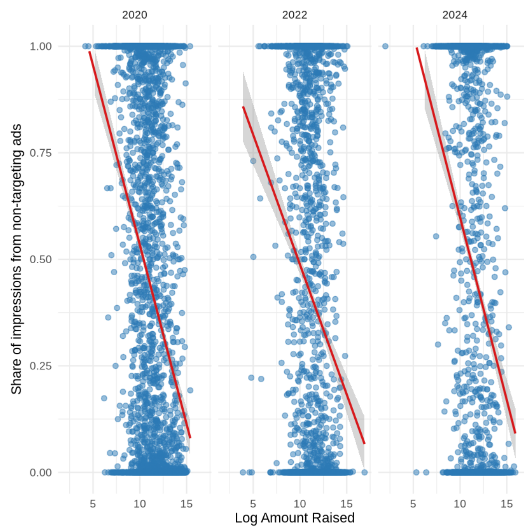

(d) State legislative general elections

Notes: Each panel shows the relationship between a candidate's total amount raised (in log terms) and the share of their Meta ads that target neither a custom list nor interests and behaviors, with a linear fit and 95% confidence band. Panels are faceted by election cycle. Data source: Meta Ad Library, Ad Targeting Dataset, and DIME.

Table 8: Targeting criteria ranked by ad spending, by partisan lean

Target	Net Republican $	Net Republican Cand.	N. Cands.	Total Spending
Likely engagement with US political content (conservative)	0.931	0.860	2851	26.3M
Donald Trump	0.894	0.935	1814	23.2M
Fox News Channel	0.860	0.951	1704	18.9M
Sean Hannity	0.861	0.943	1289	18.9M
Likely engagement with US political content (liberal)	-0.924	-0.838	1694	18.1M
Ben Shapiro	0.973	0.966	1154	13.4M
Likely engagement with US political content (moderate)	-0.205	0.005	1404	13.3M
Politics	-0.023	0.170	2670	12.9M
Rush Limbaugh	0.793	0.928	999	12.4M
The Sean Hannity Show	0.792	0.965	724	12.1M
Community issues	-0.461	-0.258	1654	11.6M
Fox & Friends	0.789	0.955	751	11.2M
Glenn Beck	0.969	0.967	826	10.9M
Donald Trump Jr.	0.996	0.926	366	10.7M
Hunting	0.962	0.890	1534	10.4M
TheBlaze	0.988	0.978	650	10.2M
Eric Trump	0.990	0.955	308	9.6M
The Rush Limbaugh Show	0.731	0.953	661	9.4M
Elizabeth Warren	-0.993	-0.970	539	8.8M
Mike Pence	1.000	0.952	187	8.5M
The Daily Caller	0.973	0.973	483	8.4M
Republican Party (United States)	0.968	0.855	1282	8.3M
Fishing	0.959	0.835	1023	8.3M
Tucker Carlson	0.715	0.951	571	8.1M
Bill O'Reilly (political commentator)	0.967	0.967	512	8.0M
Donald Trump for President	0.994	0.933	595	7.9M
Kamala Harris	-0.982	-0.948	699	7.9M
RedState	0.994	0.990	496	7.9M
Cory Booker	-0.998	-0.971	349	7.8M
Republican National Committee	0.999	0.962	447	7.7M
Barack Obama	-0.988	-0.946	447	7.7M
Mark Levin	0.994	0.983	526	7.6M
Make America Great Again	0.999	0.985	271	7.6M
NPR	-0.832	-0.830	700	7.4M
Current events	-0.642	-0.263	963	7.3M
Voting	-0.549	-0.136	906	7.1M
Golf	0.213	0.756	373	7.1M
Bernie Sanders	-0.929	-0.952	583	7.0M
Judge Jeanine Pirro	0.986	0.979	527	6.9M
Rachel Maddow	-0.986	-0.951	405	6.9M
Michelle Obama	-0.987	-0.921	228	6.8M
Chick-fil-A	0.979	0.979	373	6.7M
Social change	-0.935	-0.668	780	6.7M
The Heritage Foundation	0.993	0.981	411	6.5M
National Republican Senatorial Committee	0.999	0.984	258	6.5M
Charity and causes	-0.339	-0.246	769	6.4M
Activism	-0.854	-0.689	813	6.4M
The Rachel Maddow Show	-0.986	-0.959	392	6.2M
Democratic Party (United States)	-0.961	-0.821	823	6.2M
TheBlaze Radio Network	0.999	0.994	329	6.2M

The table also reveals an asymmetry: the most popular targets are dominated by right-leaning features targeted by Republican candidates. The only Democratic-leaning target to crack the top 10 is "Likely engagement with US political content (liberal)."

Examining the most-targeted features reveals that this asymmetry is downstream of the asymmetry in the media landscape in the US. Of the top fifty targets, sixteen are conservative media outlets or personalities. Only three (NPR, plus two versions of Rachel Maddow) are liberal media figures or outlets.

The existence of a right-wing media sphere in the US produces highly ideologically sorted features to target on the right, in a way that has no parallel on the left. Candidates of the ideological right are thus advantaged in the sense described in our model: right-wing candidates have available targeting cells that are both large in potential reach and contain a very high fraction of potential supporters. This asymmetry may explain the fact that the introduction of digital ads seems to have moderated the Democratic candidate pool in previously cost-prohibitive districts, but made the Republican candidate pool if anything more extreme in those same districts.

Conclusion

Fears about the introduction of targeted advertising in campaigns have tended to focus on the idea that campaigns might identify persuadable or vulnerable voter groups and target them with misinformation, misleading claims, or messages aimed at vote suppression. For example, Kim et al. (2026) reports that non-white voters were disproportionately targeted with ads discouraging voting in the 2016 presidential election. A Brookings Institution report in the leadup to the 2022 midterms (Lai, 2022) claimed that "Microtargeting has also allowed the spread of disinformation, allowing both political entities and individuals to disseminate ads to targeted groups with great precision." And even in the absence of any manipulative or false claims, discussion of micro-targeting tends to focus on its potential to enhance campaigns' persuasive ability. Tappin et al. (2023), for example, find that tailoring campaign messages to individual voters can substantially increase the persuasive effectiveness of campaign messaging compared to broadcasting a single best message to all voters.

We show in this paper that persuasion of any kind—whether with misinformation, true information, or any other stimulus—is not necessary for targeted advertising to be consequential for the set of candidates who contest and win elections. Even in the best-case world of perfectly informative, policy-focused advertising, the existence of targeting features that correlate with ideological preferences can advantage some candidates over others. Candidates whose likely supporters are easily identified by targeting features available on a platform gain a cost-efficiency advantage in reaching those supporters and informing them of the candidate's position and presence in the election. The efficiency advantage is especially important in down-ballot contests where initial voter awareness is low, and in primary elections where party labels are not present to guide voters' choices.

We find that candidates use targeted advertising in ways consistent with our theory, and not with theories of disinformation-spreading. They overwhelmingly use digital ads to target their likely supporters (a finding consistent with Allcott et al.'s (2026) randomized experiment in the 2020 presidential election). Only the best-funded campaigns advertise to targets that are likely to contain even a mixture of their own and their opponents' supporters, much less targeting their opponent's base.

Theoretically, the consequences of this kind of targeted advertising for candidate entry are ambiguous, depending on the specific structure of the targeting cells available to political advertisers on a given platform at a given time. It appears that in the US, on Facebook and Instagram, the structure of targets mirrors that of the media landscape. The ideological asymmetry in US media means that it is easy to find user attributes that identify large, highly concentrated groups of right-wing users in most electoral districts. It is much harder to do the same for left-wing or moderate users. It appears that the efficiency advantage this structure gives to right-wing candidates is large enough to outweigh what would otherwise be the moderating influence on the candidate pool of digital ads' cost advantage over broadcast media.

References

Aggarwal, Minali, Jennifer Allen, Alexander Coppock, Dan Frankowski, Solomon Messing, Kelly Zhang, James Barnes, Andrew Beasley, Harry Hantman, and Sylvan Zheng, "A 2 million-person, campaign-wide field experiment shows how digital advertising affects voter turnout," *Nature Human Behaviour*, March 2023, *7* (3), 332–341.

Allcott, Hunt, Matthew Gentzkow, Ro'ee Levy, Adriana Crespo-Tenorio, Natasha Dumas, Winter Mason, Devra Moehler, Pablo Barbera, Taylor Brown, and Juan Carlos Cisneros, "The effects of political advertising on Facebook and Instagram before the 2020 US election," *Nature Human Behaviour*, 2026, *10* (5), 884–895.

Ashworth, S. and J.D. Clinton, "Does Advertising Exposure Affect Turnout?," *Quarterly Journal of Political Science*, 2007, *2* (1), 27–41.

Ashworth, Scott, "Campaign finance and voter welfare with entrenched incumbents," *American Political Science Review*, 2006, *100* (01), 55–68.

Avis, Eric, Claudio Ferraz, Frederico Finan, and Carlos Varjão, "Money and Politics: The Effects of Campaign Spending Limits on Political Entry and Competition," *American Economic Journal: Applied Economics*, 2022, *14* (4), 167–199.

Bagwell, Kyle, "The Economic Analysis of Advertising," in M. Armstrong and R. Porter, eds., *Handbook of Industrial Organization*, Vol. 3, Elsevier, 2007, pp. 1701 – 1844.

Bail, Christopher A, Lisa P Argyle, Taylor W Brown, John P Bumpus, Haohan Chen, MB Fallin Hunzaker, Jaemin Lee, Marcus Mann, Friedolin Merhout, and Alexander Volfovsky, "Exposure to opposing views on social media can increase political polarization," *Proceedings of the National Academy of Sciences*, 2018, *115* (37), 9216–9221.

Bär, Dominik, Francesco Pierri, Gianmarco De Francisci Morales, and Stefan Feuerriegel, "Systematic discrepancies in the delivery of political ads on Facebook and Instagram," *PNAS nexus*, 2024, *3* (7).

Baron, D.P., "Electoral competition with informed and uniformed voters," *American Political Science Review*, 1994, pp. 33–47.

Bonica, Adam, "Inferring roll-call scores from campaign contributions using supervised machine learning," *American Journal of Political Science*, 2018, *62* (4), 830–848.

_ , "Database on Ideology, Money in Politics, and Elections," 2023.

Buisseret, Peter and Richard Van Weelden, "Inequality, Polarization, and Culture Wars," Technical Report, working paper 2025.

Coppock, Alexander, Donald P. Green, and Ethan Porter, "Does digital advertising affect vote choice? Evidence from a randomized field experiment," *Research & Politics*, 2022, *9* (1).

_ **, Seth J. Hill, and Lynn Vavreck**, "The small effects of political advertising are small regardless of context, message, sender, or receiver: Evidence from 59 real-time randomized experiments," *Science Advances*, September 2020, *6* (36).

Cox, Christian, "The Equilibrium Effects of Campaign Finance Deregulation on US Elections," *Econometrica*, 2026, *Forthcoming.*

Desai, Zuheir and John Duggan, "A Model of Interest Group Influence and Campaign Advertising," *Quarterly Journal of Political Science*, January 2021, *16* (1), 105–137.

Enríquez, José Ramón, Horacio Larreguy, John Marshall, and Alberto Simpser, "Mass Political Information on Social Media: Facebook Ads, Electorate Saturation, and Electoral Accountability in Mexico," *Journal of the European Economic Association*, August 2024, *22* (4), 1678–1722.

Fouirnaies, Alexander, "How Do Campaign Spending Limits Affect Electoral Competition?: Evidence From Great Britain 1885-2010," *American Political Science Review*, 2021, *115* (2).

Fowler, Erika Franklin, Michael M. Franz, Gregory J. Martin, Zachary Peskowitz, and Travis N. Ridout, "Political advertising online and offline," *American Political Science Review*, 2021, *115* (1), 130–149.

Gentzkow, Matthew, Jesse M. Shapiro, Frank Yang, and Ali Yurukoglu, "Pricing power in advertising markets: Theory and evidence," *American Economic Review*, 2024, *114* (2), 500–533.

Gordon, Brett R. and Wesley R. Hartmann, "Advertising effects in presidential elections," *Marketing Science*, 2013, *32* (1), 19–35.

Hall, Andrew B, *Who wants to run?: How the devaluing of political office drives polarization*, University of Chicago Press, 2019.

Hewitt, Luke, David Broockman, Alexander Coppock, Ben M. Tappin, James Slezak, Valerie Coffman, Nathaniel Lubin, and Mohammad Hamidian, "How Experiments Help Campaigns Persuade Voters: Evidence from a Large Archive of Campaigns' Own Experiments," *American Political Science Review*, February 2024, pp. 1–19.

Hindman, Matthew, *The Internet Trap*, Princeton: Princeton University Press, 2018.

Huber, G.A. and K. Arceneaux, "Identifying the persuasive effects of presidential advertising," *American Journal of Political Science*, 2007, *51* (4), 957–977.

Kalla, Joshua L. and David E. Broockman, "The minimal persuasive effects of campaign contact in general elections: Evidence from 49 field experiments," *American Political Science Review*, 2018, *112* (1), 148–166.

Kim, Young Mie, Jordan Hsu, David Neiman, Colin Kou, Levi Bankston, Soo Yun Kim, Richard Heinrich, Robyn Baragwanath, and Garvesh Raskutti, "The Stealth

Media? Groups and Targets behind Divisive Issue Campaigns on Facebook," *Political Communication*, 2018, *35* (4), 515–541.

— **, Ross Dahlke, Hyebin Song, and Richard Heinrich**, "Targeted digital voter suppression efforts likely decrease voter turnout," *Proceedings of the National Academy of Sciences*, 2026, *123* (5).

Lai, Samantha, "Data misuse and disinformation: Technology and the 2022 elections," Technical Report, Brookings Institution 2022.

Meta Platforms, Inc., "Meta Ad Targeting Dataset," 2026.

— , "Meta Content Library API version v6.0," 2026.

Osborne, Martin J. and Al Slivinski, "A Model of Political Competition with Citizen-Candidates," *The Quarterly Journal of Economics*, 1996, *111* (1), 65–96.

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

Sheingate, Adam, James Scharf, and Conner Delahanty, "Digital Advertising in U.S. Federal Elections, 2004-2020," March 2022.

Spenkuch, Jörg L. and David Toniatti, "Political Advertising and Election Results," *The Quarterly Journal of Economics*, 2018, *133* (4), 1981–2036.

Tappin, Ben M., Chloe Wittenberg, Luke B. Hewitt, Adam J. Berinsky, and David G. Rand, "Quantifying the potential persuasive returns to political microtargeting," *Proceedings of the National Academy of Sciences*, 2023, *120* (25).

Titova, Maria, "Targeted Advertising in Elections," *American Economic Journal: Microeconomics*, 2026, *Forthcoming.*

Tucker, Joshua A, Andrew Guess, Pablo Barberá, Cristian Vaccari, Alexandra Siegel, Sergey Sanovich, Denis Stukal, and Brendan Nyhan, "Social media, political polarization, and political disinformation: A review of the scientific literature," *Political polarization, and political disinformation: a review of the scientific literature (March 19, 2018)*, 2018.

Votta, Fabio, Simon Kruschinski, Mads Hove, Natali Helberger, Tom Dobber, and Claes de Vreese, "Who Does(n't) Target You? Mapping the Worldwide Usage of Online Political Microtargeting," *Journal of Quantitative Description: Digital Media*, May 2024, *4.*

Waldfogel, Sarah, "Two-faced on facebook: Ad tailoring in U.S. house campaigns," *Journal of Public Economics*, July 2026, *259.*

Warshaw, Christopher and Chris Tausanovitch, "Subnational ideology and presidential vote estimates," 2022.

Appendices

A Data and Measurement

Candidate Pool

This section describes how we construct candidate counts and characteristics. For the US House races, we build the district-by-cycle (i.e. race level) candidate counts primarily from Bonica (2023). Not all persons with DIME records should be considered as candidates. We consider a DIME person to be a candidate in a primary race if they meet at least one of the three conditions: 1. they are flagged as having won or lost in the primary; 2. they are flagged as having won or lost in the general; 3. they have obtained non-missing votes in either the primary or the general. Similarly, we consider a DIME person to be a candidate in a general race if they meet at least one of the three conditions: 1. they are flagged as having won the primary; 2. they are flagged as having won or lost in the general; 3. this person has obtained non-missing votes in either the general. For the 2024 cycle, we supplement the DIME with CQ's 2024 House vote returns. Candidate characteristics such as party label, incumbency status are from DIME.

For the state legislature races, we construct counts using a similar logic but from two different sources: DIME for election cycles before 2020 and Ballotpedia for 2020-2024. We apply the same rules as above for DIME-based counts, and take all records reported by Ballotpedia as valid candidates unless they are flagged as "disqualified" or "withdrawn" and with missing votes. Party and incumbency come from DIME before 2020 and from Ballotpedia from 2020 on. DIME records are matched to Ballotpedia based on candidate names and district such that for the post 2020 data we can assign ideology measures accordingly.

The candidate count regressions and figures control for population densities. For the congressional races, this is straightforward and pulled directly from the Census. For the state legislature level, we aggregate both population and land area from Census tract level to get state legislature district level measure. All measures are adjusted when redistricting happens.

Candidate Facebook Ads

We use Meta's Ad Library and Ad Targeting datasets for Ad level information such as impression, potential reach, targeting features, spend; information is available for ads run beginning August 3, 2020. The two datasets are linked together via ad IDs and table partition data stamps that are periodically updated by Meta. For a complete list of variables available and their ranges, see https://developers.facebook.com/docs/ad-targeting-dataset/table-schema. In order to use these two datasets, we also need to obtain Facebook page IDs for all congressional and state legislature candidates.

For congressional candidates, we match candidates to a lifelong Facebook advertiser list by names. We only consider advertisers who have run at least one ad with a political disclaimer. We validate each of the matches rigorously. First, we scrape the matched page's information such as the page introduction, the texts in the "About" box, page category and contact information etc., and feed all these information to two LLMs — Chat GPT-4o and Llama 3.1-8b — to check whether the information obtained resembles the candidate matched to that page; this check is possible for any Facebook page that is currently active and public. Second, we retrieve each page's ad delivery regions from the Ad Library API, and check the top regions against the candidate's state; this check is primarily for pages that are no longer active. Any remaining ambiguous case is hand checked.

For state legislature candidates, we start with Ballotpedia's candidate campaign Facebook URLs — this is different from the page IDs that identifies pages in Meta's datasets. We first recover the page IDs by scraping the URLs. For those not recoverable — most of the time it was because the page is no longer active or public — we go back to a similar process described above for congressional candidates. That is, we match candidates to the lifelong Facebook advertiser list by names, and then validate the match via Ad Library API pulled information.

Finally, we use these candidate-page crosswalks to pull ads run by these candidates. We define the primary window to be December 1st of the year before the election year until the date of primary election, and the general window to be the day right after the primary election until the date of general election. We only consider ads delivered within these windows.

B Additional Figures and Tables

Figure B1 shows the relationship between effective television ad price and the fraction of US House candidates' media spending on digital ads in each year from 2004 to 2020.

Figure B1: Percentage of campaign media expenditures by US House candidates on digital ads, by effective TV ad price.

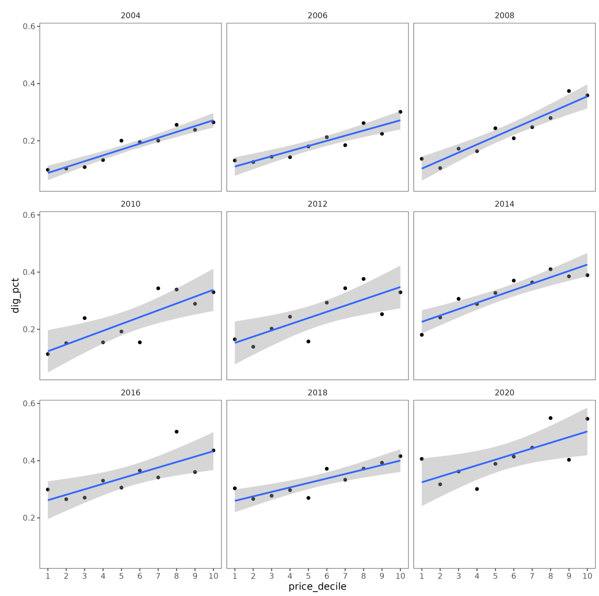

Notes: The x-axis is the decile rank of the Congressional district in effective price of television ads, computed as in equation 5. The y-axis is the fraction of total media expenditures by Congressional campaigns going to digital platforms (primarily Facebook/Instagram and Google/Youtube ads but including other social media platforms as well). Points are binned means, and the blue line is the OLS fit (separately by year). Sources: Prices from SQAD and Nielsen Ad Intel, expenditure data from Sheingate et al. (2022).

Figure B2 shows that there is a strong relationship between population density and effective TV ad price. But, there remains substantial residual variation in effective prices even for districts with similar density, because of variation in alignment between district boundaries and media market boundaries.

Figure B2: Effective TV ad price against Congressional district log population density.

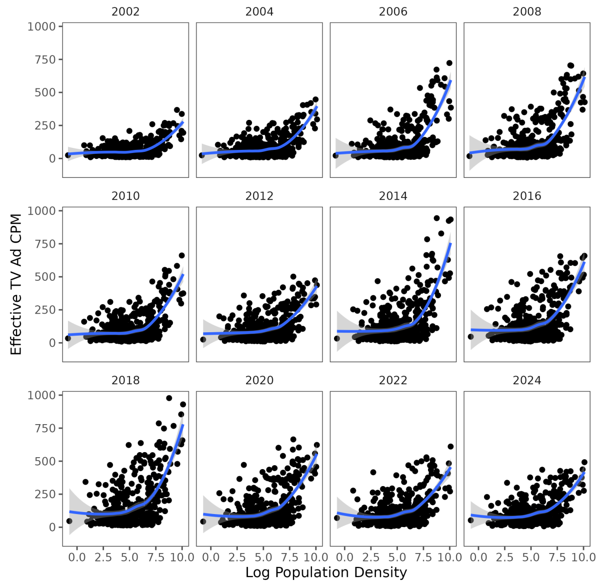

Notes: Each point is a congressional district-year. The blue curve is a LOESS smoother.

Figure B3 shows the empirical distribution of the audience across age and gender cells on TV and on digital ads run on Meta platforms. The sample here is ads run by candidates in US Congressional races.

Figures B4 and B5 show visual analogs of the regressions in Table 1. Figures B6 and B7 give the same analogs for Table 2.

Figure B3: Audience distribution of political advertising on TV and on Meta platforms.

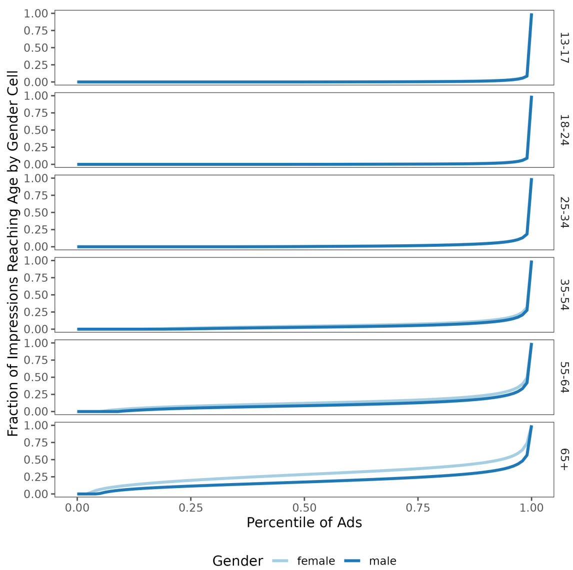

(a) TV

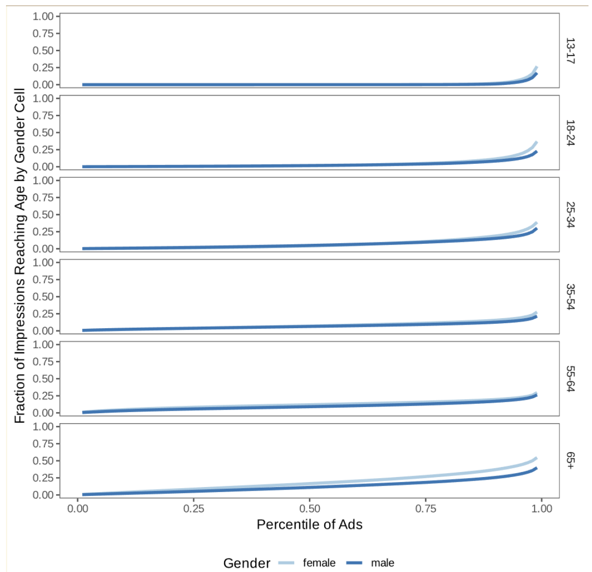

(b) Meta

Notes: Each panel shows the empirical cumulative distribution function (CDF) of audience characteristics of political advertisements on television and on Facebook. The variable whose CDF is plotted in each panel is the share of the ad's audience falling into the relevant age and gender bin, e.g., the median TV ad has roughly 20% of its viewers in the 65+ Men group. Figures include all ads run by US Congressional campaigns between 2020 and 2024 (the period for which the Meta Ad Targeting dataset has coverage). Sources: Facebook Open Research and Transparency (FORT) Ad Targeting and Ad Library datasets; Nielsen Ad Intel.

Figure B4: Number of US House candidates per district by TV ad price decile, by year.

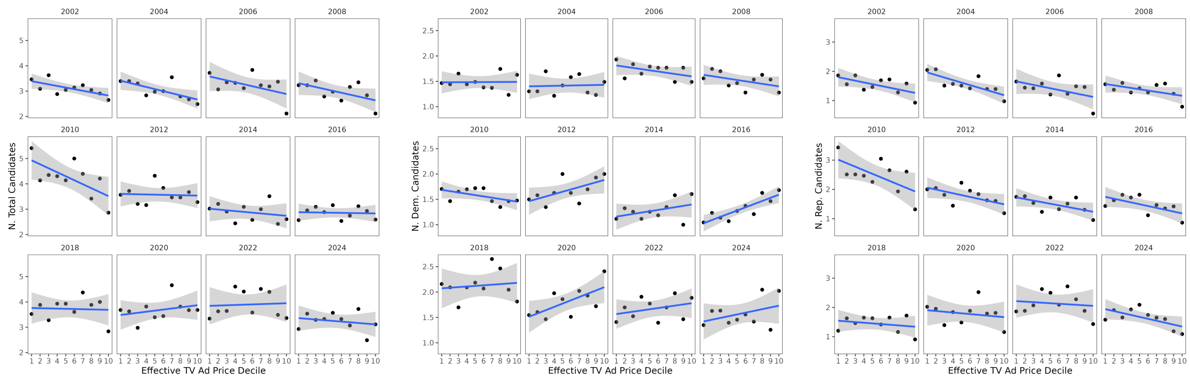

(a) Total Candidates (b) Democratic Candidates (c) Republican Candidates

Notes: Each panel shows binned means of the number of candidates per House district, for each decile of the distribution of district-level effective TV ad price. The blue line is the OLS fit.

Table B1: Percent of races where candidates targeted overlapping audiences

	Congress		State Legislature	
	Primary	General	Primary	General
2020	6.5	14.7	4.0	7.4
2022	8.1	2.3	2.3	2.1
2024	15.9	8.2	4.3	2.5
All	10.2	8.6	3.3	4.7

Each entry is the percentage of races in which at least one targeting cell received ads from more than one candidate in the race. A targeting cell is defined as the combination of all elements of targeting on age, gender, geographic areas, interests, behaviors, occupation, and education background. The denominator in each cell is the set of races in which at least two candidates advertised on Facebook or Instagram during the relevant election window. Pooled across cycles, the denominators are 342 congressional primary, 349 congressional general, 1,293 state legislative primary, and 2,060 state legislative general races. Ads run by candidates between Dec. 1st of the year before the election year and the primary election date are counted as primary election advertising. Ads run by candidates between the day after the primary election date and the general election date are counted as general election advertising. State legislative races held in odd-numbered years are assigned to the following even-year cycle. Data source: Meta Ad Library and Ad Targeting Datasets.

Figure B5: Uncontested US House general and primary elections by TV ad price decile, by year.

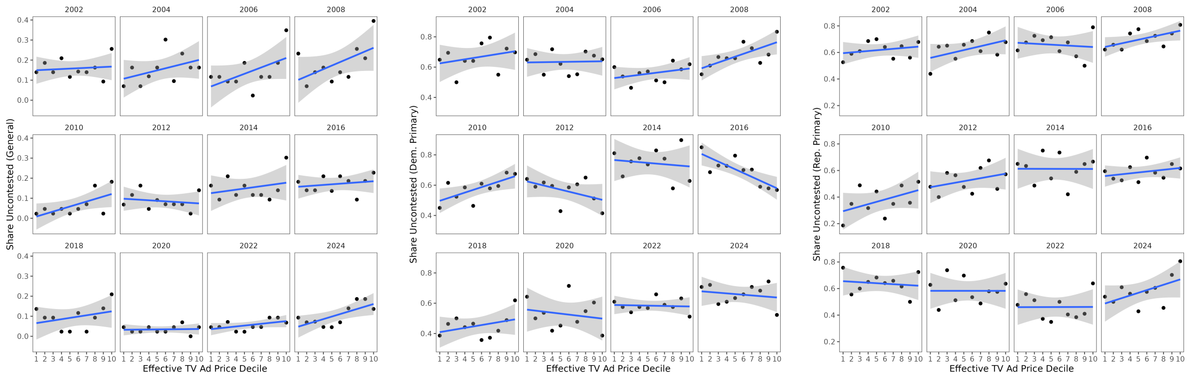

(a) General (b) Democratic Primary (c) Republican Primary

Notes: Each panel shows binned means of the fraction of House races that are uncontested, for each decile of the distribution of district-level effective TV ad price. The blue line is the OLS fit.

Figure B6: Number of state legislative candidates per district by TV ad price decile, by year.

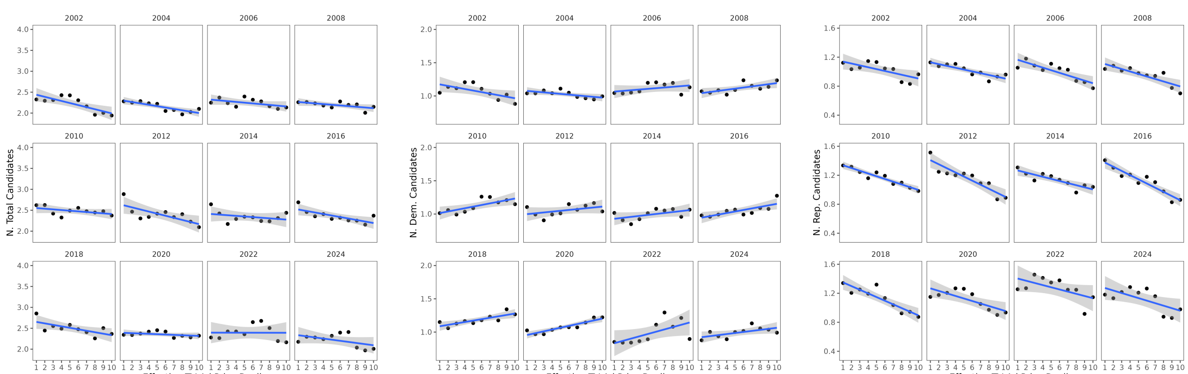

(a) Total Candidates (b) Democratic Candidates (c) Republican Candidates

Notes: Each panel shows binned means of the number of candidates per state legislative district, for each decile of the distribution of district-level effective TV ad price. The blue line is the OLS fit.

Figure B7: Uncontested state legislative general and primary elections by TV ad price decile, by year.

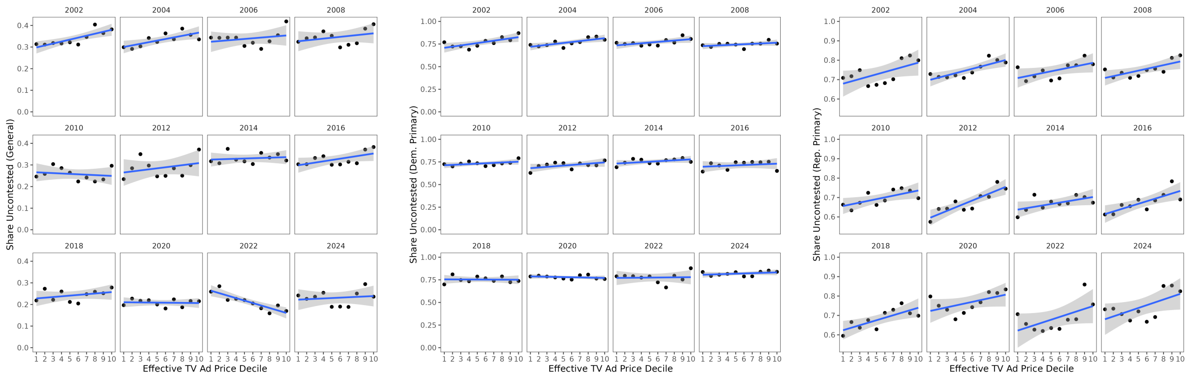

(a) General (b) Democratic Primary (c) Republican Primary

Notes: Each panel shows binned means of the fraction of state legislative races that are uncontested, for each decile of the distribution of district-level effective TV ad price. The blue line is the OLS fit.

Figure B8: Candidate CFScores in US House races by TV ad price decile, by year.

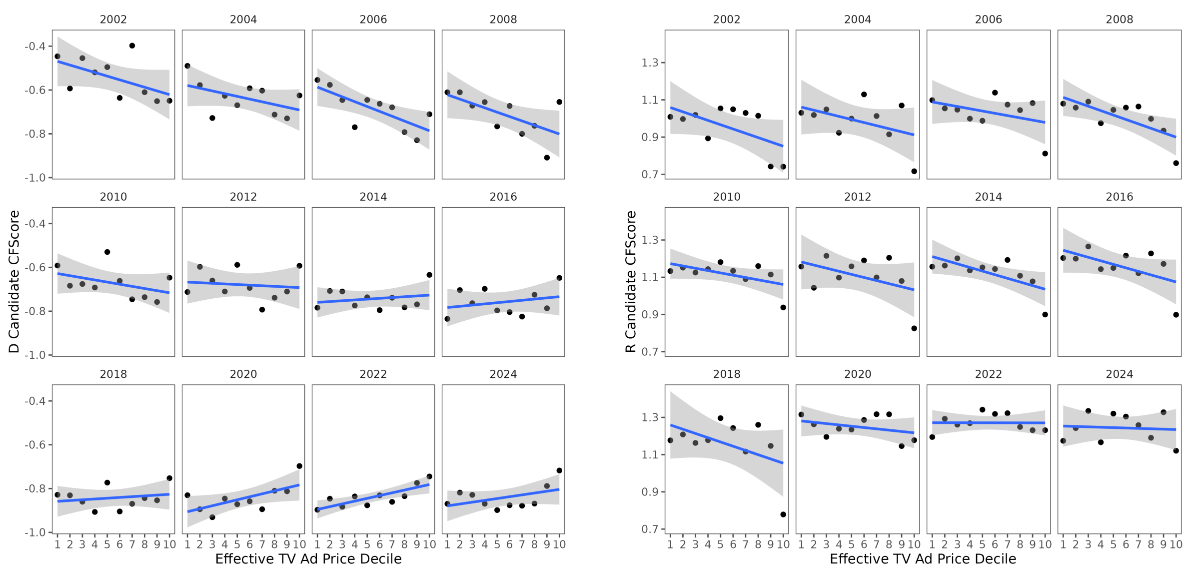

(a) Democratic Candidates (b) Republican Candidates

Notes: Each panel shows binned means of candidates' CFScores (Bonica, 2023) in US House races, for each decile of the distribution of district-level effective TV ad price. The blue line is the OLS fit.

Figure B9: Candidate CFScores in state legislative races by TV ad price decile, by year.

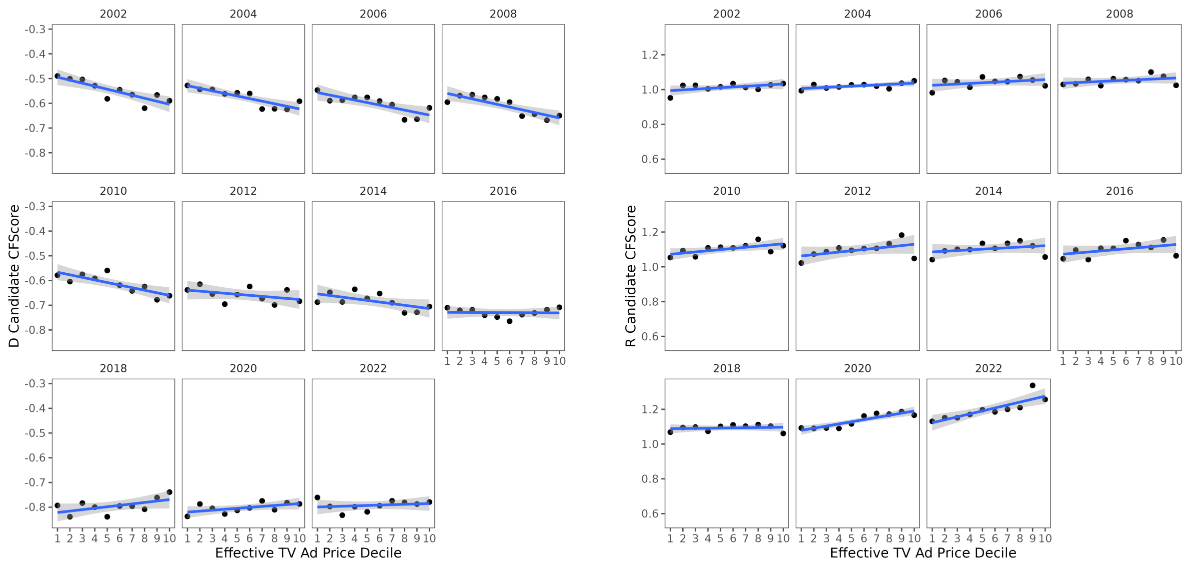

(a) Democratic Candidates

(b) Republican Candidates

Notes: Each panel shows binned means of candidates' CFScores (Bonica, 2023) in state legislative races, for each decile of the distribution of district-level effective TV ad price. The blue line is the OLS fit.

Figure B10: Candidates increasingly deliver Meta ads to custom audiences.

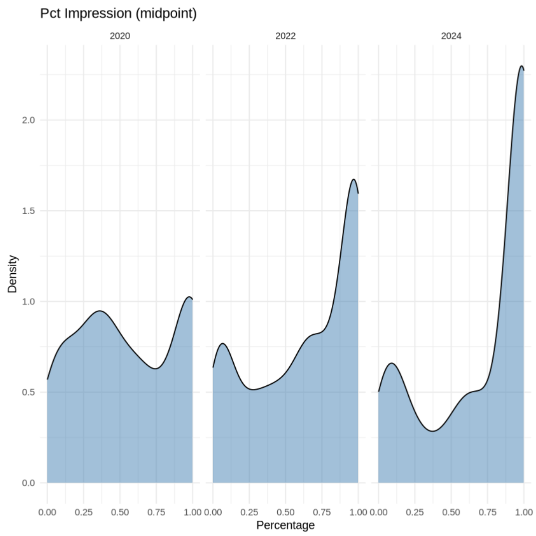

(a) Congressional primary elections

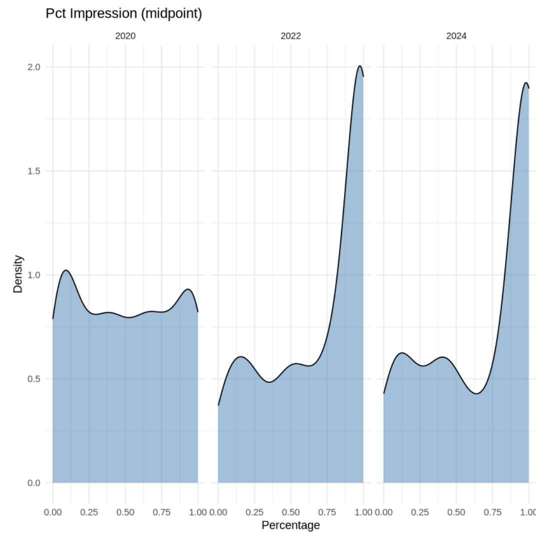

(b) Congressional general elections

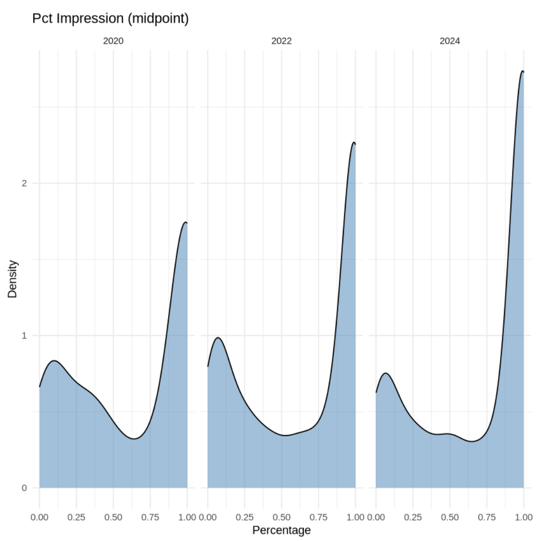

(c) State legislative primary elections

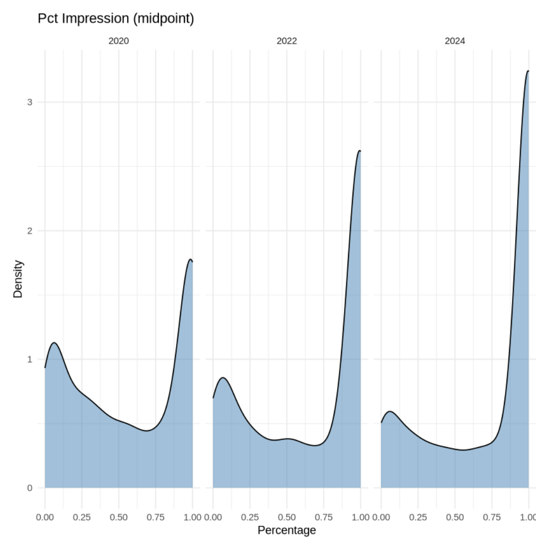

(d) State legislative general elections

Notes: Each density plot shows the distribution of candidates' share of impressions delivered to a custom audience. Each point in the underlying data is a candidate; the variable whose distribution is plotted is the candidate's total impressions of ads delivered to custom audiences, divided by the candidate's total impressions of all ads. Candidate totals are aggregated from the ad-level underlying data. Meta reports ad impressions in ranges defined by a lower and upper bound for each ad; we use the midpoint of the bounds and sum the midpoints over all ads run by a given candidate. Data source: Meta's Ad Library and Ad Targeting Datasets.

Figure B11: Candidates almost entirely advertise in exclusive targeting cells.

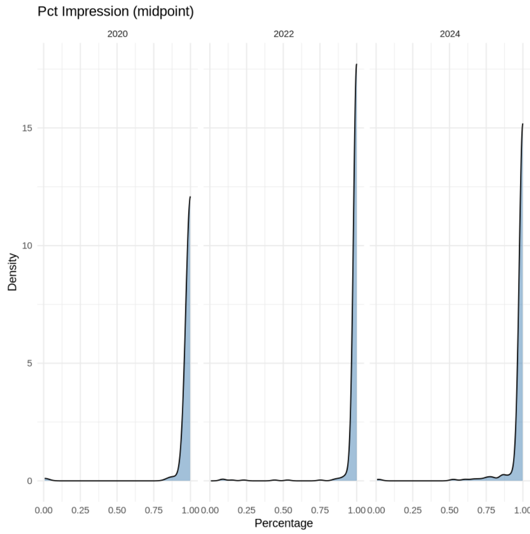

(a) Congressional primary elections

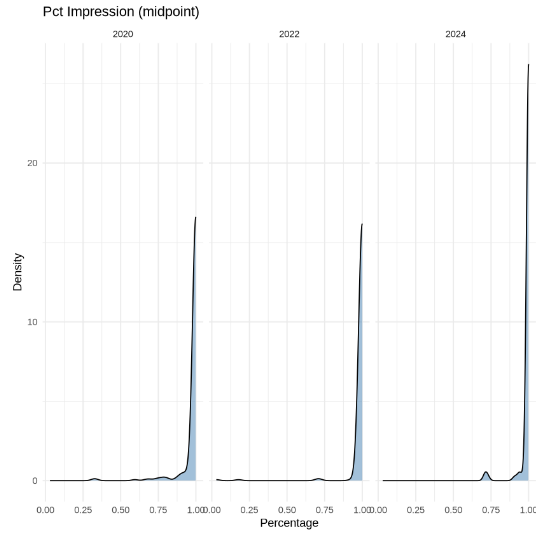

(b) Congressional general elections

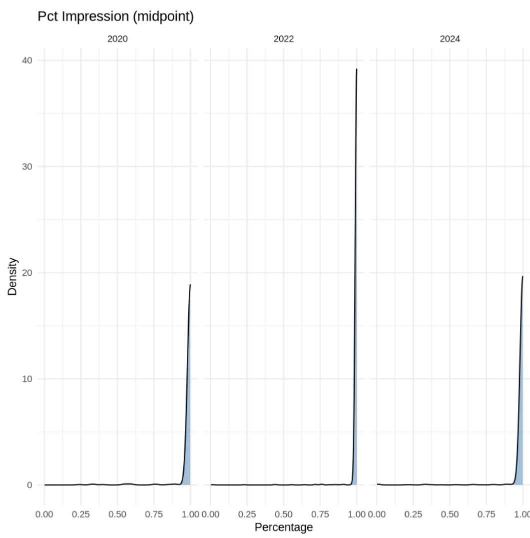

(c) State legislative primary elections

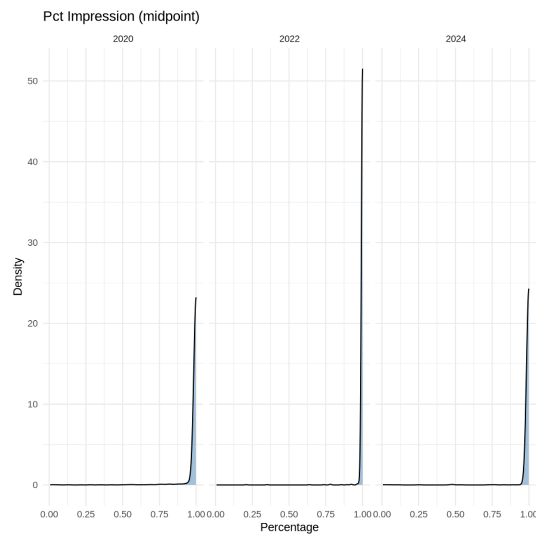

(d) State legislative general elections

Notes: Each density plot shows the distribution of candidates' share of impressions delivered to an exclusive targeting cell. A targeting cell is defined as a combination of all elements of targeting on age, gender, geographic areas, interests, behaviors, occupation, education background. A targeting cell is considered exclusive if in a race with more than one candidate, only one candidate delivers ads to this cell. Each point in the underlying data is a candidate; the variable whose distribution is plotted is the candidate's total impressions of ads delivered to exclusive targeting cells, divided by the candidate's total impressions of all ads. Candidate totals are aggregated from the ad-level underlying data. Meta reports ad impressions in ranges defined by a lower and upper bound for each ad; we use the midpoint of the bounds and sum the midpoints over all ads run by a given candidate. Data source: Meta Ad Library and Ad Targeting Datasets.