

Selective Exposure and Electoral Competition*

Avidit Acharya,[†] Peter Buisseret,[‡] Adam Meirowitz,[§] Floyd Zhang[¶]

January 20, 2023

Abstract

We theoretically study how selective information exposure to voters impacts electoral competition between two policy motivated candidates – a ‘liberal’ versus a ‘conservative.’ Each candidate has stochastic valence that is realized after the candidates choose platforms. In our model of selective exposure voters receive information about the candidates’ valences that is slanted to reflect their ideological preferences. Existing work predicts that selective exposure intensifies platform polarization. We instead show that it can reduce platform polarization relative to a context where voters all learn the candidates’ valences from a neutral source.

Keywords: electoral competition, media consumption, polarization

*We are grateful to seminar audiences at Kellogg School of Management, and to John Duggan, for helpful suggestions and feedback.

[†]Professor, Department of Political Science, Stanford University, *Email:* avidit.acharya@gmail.com.

[‡]Assistant Professor, Department of Government, Harvard University, *Email:* pbuisseret@fas.harvard.edu.

[§]Corresponding Author. Professor, Department of Political Science, Yale University, *Email:* adam.meirowitz@yale.edu.

[¶]PhD Candidate, Stanford GSB, *Email:* floydjz@stanford.edu.

1 Introduction

In the United States, and elsewhere, voters acquire political knowledge from information sources that reflect their ideological predispositions, exposing them to only selective information. This selective exposure manifests in voters’ choice of partisan media (Iyengar and Hahn, 2009), or their interpersonal and social networks (Mutz and Martin, 2001). These information sources shape electoral assessments; for example, a partisan media outlet may buttress its audience’s favorable perception of one candidate by suppressing news that highlights her weaknesses, while simultaneously “inflaming negative views” about her opponent (Pierson and Schickler, 2020). As a consequence, liberal and conservative voters may base their electoral choices on starkly different information.

In this paper, we study how voters’ selective exposure shapes the policies candidates propose in elections. We ask: Should selective exposure intensify or instead moderate platform polarization in elections, relative to an environment in which all voters are equally informed about candidates?

To address this question, we pursue a novel extension of the classical Calvert-Wittman model of electoral competition between policy-motivated candidates (Wittman, 1983, Calvert, 1985). Two candidates—a ‘liberal’ versus a ‘conservative’—simultaneously announce platforms. A valence shock is then realized for each candidate.¹ Voters learn about the candidates’ valences, and cast ballots for one of the two candidates. The majority-winning candidate implements her announced platform.

The critical premise of our framework is that voters with different ideologies learn different information about the candidates over the course of the election. To capture this idea in a stark way, we assume a liberal voter only learns the liberal candidate’s valence when it is better than expected (“good news”), and only learns the conservative candidate’s valence when it is worse than expected (“bad news”). Similarly, conservatives only learn good news about the conservative’s valence, and bad news about the liberal’s valence. The most natural interpretation is that liberals and conservatives receive news from different and biased media sources. These sources cannot fabricate stories, but they can suppress or highlight stories about each of the candidates in ways that cater to their audiences’ political views.

Our analysis focuses on the consequences of these discrepancies in voter information for platform polarization. We do not examine the demand-and-supply channels through

¹These shocks capture the unfolding of features that would generally impact voter evaluations in similar ways—for example, evidence that speaks to a leader’s competence or corruption.

which these discrepancies arise. However, Bernhardt, Krasa and Polborn (2008) show how profit-maximizing media firms slant their reporting by suppressing news stories that might be seen as unfavorable by their audiences, generating the pattern of information segmentation by voter ideology that our analysis presumes. Broockman and Kalla (2022) study “partisan coverage filtering,” whereby “a media outlet is more likely to report information flattering to politicians and causes on their ideological or partisan side, and not to report information unflattering to the same.” They highlight its consequence for voters’ knowledge about politics in a novel experimental setting.

We ask how selective information exposure that can arise in these settings shapes candidates’ platform choices. Our main result shows that selective information exposure may *reduce* platform polarization. It highlights that segmentation of voters’ information according to ideology need not, in and of itself, lead to more polarized elections.

To understand why, consider a full-learning benchmark in which candidates stake out policy positions anticipating that all voters will learn both candidates’ valences over the course of the campaign. As in Wittman (1983) and Calvert (1985), the candidates adopt differentiated platforms that trade-off two opposing incentives: moderation increases the probability of winning, but extremism increases the value of winning. This trade-off arises because voters care about both policy and valence. When the realized valence difference is large policy positions have less impact on voter behavior, encouraging each candidate to locate closer to her ideal policy. Small realized valence differences make the vote count more dependent on policy, encouraging each candidate to moderate her platform. Because candidates choose platforms before valences are realized, their interior platform choices reflect their uncertainty about what valences voters will see and how they will decide the election.

Conversely, in a no-learning benchmark where no voter ever learns either candidate’s valence, voters choose between candidates exclusively on policy. This intensifies the candidates’ incentives to moderate in pursuit of the median voter’s preference: the unique equilibrium yields full convergence at the median voter’s preferred policy. The contrast with the full-learning benchmark is that a softening of competition on one dimension (valence) intensifies competition on the other dimension (policy).

Our analysis highlights that selective exposure is an intermediate case. On average, selective exposure reduces voters’ perceived differences in candidate quality. Relative to the full-learning benchmark, but not to the extent of no learning, selective exposure about valence makes voters’ choices more sensitive to differences in the candidates’ poli-

cies. We correspondingly show that candidates respond by competing more aggressively on policy, thereby converging more than in the full-information setting.

In related work, Glaeser, Ponzetto and Shapiro (2005) and Virág (2008) characterize equilibria in spatial electoral competition when different voters have different information about the candidates. Similar to our framework, these papers assume that liberal voters are better-informed about the liberal candidate and conservatives are better-informed about the conservative candidate. However, unlike our paper, both of these papers predict that platforms are more polarized under selective information exposure.

Why do we obtain the opposite result? The reason is that these papers differ from ours in two ways: (i) voters are certain about candidates' valences but uncertain about their policy choices, and (ii) voters face a cost of voting and may abstain.² In our framework, voters know the candidates' policies, but are uncertain about their valence. Our analysis unearths that selective exposure's implications for policy extremism depend on which dimension—valence or policy—is selectively learned.³ Chan and Suen (2008), analyze a model with uncertainty about both valence and a policy relevant state but since in their model the media only provides information about the policy-relevant state the equilibrium forces are similar to those in the above cited papers and distinct from ours. They find that a smaller share of uninformed voters makes platforms more responsive to the state. This spurs convergence on the median voter's preferred policy. Our results highlight that more information about a dimension that is orthogonal to platforms—valence—softens policy competition, yielding greater policy differentiation.

Polborn and Yi (2006) and Bernhardt, Krasa and Polborn (2008) are closest to our work in one important sense; they also assume voters are uncertain about valence. However, these papers presume politicians' policy platforms are fixed, whereas our explicit focus is on candidates' policy choices. Duggan and Martinelli (2011) study how a single (possibly biased) media outlet optimally filters information to voters about a challenger's fixed but unknown platform.

In concurrent work, Hu, Li and Segal (2019) examine platform polarization under common versus personalized (i.e., targeted) information that is consumed by rationally

²In Glaeser, Ponzetto and Shapiro (2005), abstention is driven by alienation; in Virág (2008), abstention is driven by indifference.

³Further, the choice by Glaeser, Ponzetto and Shapiro (2005) and Virág (2008) to model utility as strictly concave and allow abstention due to costs of voting necessarily introduces a strong force for policy extremism. The reason is that, amongst those voters who selectively learn a candidate's platform, more ideologically extreme voters reward extreme platforms more than moderate voters punish them. This is not necessarily true in our framework.

inattentive voters. While personalized media increases polarization in leading cases, the authors unearth circumstances in which that finding reverses. Nonetheless, the mechanism is different: their purely office-seeking candidates trade off support from moderate versus extreme voters that learn and respond differently to distinct information. Our policy-motivated candidates trade off their intrinsic preference for more extreme policy with moderation’s benefit to their odds of winning. Further, in the authors’ setting with three voter ideal points all centrists obtain the same information. Our model of selective exposure instead allows moderate left- and right-of-centre voters to learn distinct information. This feature might appear to stack the deck in favor of a finding that selective exposure increases polarization—yet we find the opposite.

2 Model

Two candidates, L and R , simultaneously choose campaign platforms, x_L and x_R , prior to an election. The policy space is $X = [-1, 1]$. There is a unit mass of citizens, and their ideal points are uniformly distributed on X . If candidate $j \in \{L, R\}$ implements platform x_j , voter i with ideal point y_i derives payoff

$$U^j(x_j, v_j, y_i) = v_j - |x_j - y_i|,$$

where v_j is candidate j ’s valence. The valences v_L and v_R are independently and uniformly distributed on the interval $[0, 1]$. Candidates care exclusively about policy. Like voters, they have linear policy losses. We assume that they have extreme preferences: L ’s ideal policy is -1 and R ’s ideal policy is $+1$.

The interaction proceeds as follows.

1. The candidates simultaneously announce platforms x_L and x_R .
2. Valences are realized, and revealed according to a process we describe, below.
3. Each voter votes for one of the candidates. Ties are broken with a fair coin-toss.
4. The winner implements her policy announcement, and payoffs are realized.

We study Nash equilibria in which voting is sincere and so each voter casts her ballot for the candidate that provides a higher expected utility.⁴ Without loss of generality, we restrict L to choose $x_L \in [-1, 0]$, and R to choose $x_R \in [0, 1]$.

⁴Because there are a continuum of voters, no voter is ever pivotal and thus any voting strategy can be supported in Nash (or perfect) equilibria. As voting stages involve a binary choice the appeal of sincere voting is unambiguous and natural.

3 Benchmark Cases

3.1 Full Learning

We start with a setting in which voters learn the realization of each candidate's valence before casting their ballots. Our first proposition characterizes the unique equilibrium, in which the parties partly differentiate their platforms.

Proposition 1. *In the full learning benchmark, there is a unique equilibrium, in which $x_L^* = -\frac{1}{4}$ and $x_R^* = \frac{1}{4}$.*

In the setting with full learning, each candidate faces a trade-off: moving towards her ideal policy raises the value of winning, but moving towards the expected ideal point of the median voter (policy 0) raises the probability of winning.

Let us illustrate this trade-off. Note that a voter with ideal policy y_i prefers party R if and only if:

$$|x_L - y_i| - |x_R - y_i| + v_R - v_L \geq 0. \tag{1}$$

Let y^* denote the unique ideal point of a voter who is indifferent between the two candidates, i.e., the value of y_i such that (1) holds with equality. R wins the election if and only if $y^* < 0$ (with a tie occurring at the 0-probability event that the relationship holds with equality). For $x_L \leq 0 \leq x_R$, we have

$$y^* < 0 \iff v_R - v_L > x_R + x_L.$$

Now suppose the candidates locate at a symmetric pair $x_R = +x$ and $x_L = -x$, and consider a local deviation by candidate R to $x + \Delta$, for $\Delta > 0$. The above expression implies that R wins if and only if $v_R > v_L + \Delta$. Figure 1 plots L 's valence $v_L \in [0, 1]$ on the horizontal axis, and R 's valence $v_R \in [0, 1]$ on the vertical axis. When $\Delta = 0$, the 45-degree line demarcates the two key events, $v_R > v_L$ in which case R wins, and $v_L > v_R$ in which case L wins. Since the candidates' valences are uniformly distributed, R wins with probability 1/2.

The dashed line in that figure shows the new region in which R wins, after locating at $x + \Delta$ for $\Delta > 0$. R 's probability of winning after the shift $\Delta > 0$ in policy is the area above the diagonal line in Figure 1. We denote this probability as

$$\pi^{FL}(x, \Delta) \equiv \frac{(1 - \Delta)^2}{2}. \tag{2}$$




Figure 1: Election outcomes for valence pairs $(v_L, v_R) \in [0, 1]^2$, when $x_L = -x$ and $x_R = x + \Delta$. The thick diagonal line represents $\Delta = 0$, and the dashed diagonal line represents $\Delta > 0$.

This probability of winning decreases in Δ ; however, R 's value from winning increases in Δ . Leaving out terms that do not depend on Δ , we write R 's objective:

$$\Pi_R^{FL}(\Delta, x) \equiv \pi^{FL}(x, \Delta)(2x + \Delta),$$

from which we obtain

$$\frac{\partial \Pi_R^{FL}(0, x)}{\partial \Delta} = 0 \iff x = \frac{1}{4}.$$

The trade-off between the probability of winning and the value of winning also appears in settings with uncertainty about the median voter's preferred policy (Wittman, 1983, Calvert, 1985). In those contexts, a candidate's choice of more extreme policy reflects her gamble that the median voter will align with her on policy grounds. In our context, instead, a more extreme policy reflects a candidate's gamble that her net valence advantage will be large enough to sway the known median voter's preference. The

consequence in either setting is to soften policy competition, precluding full convergence on the median (or expected median) voter’s preferred policy.

3.2 No Learning

We next consider a setting in which voters do not learn the realization of each candidate’s valence. Analogous to the full learning case, a voter with ideal policy y_i prefers candidate R if and only if

$$|x_L - y_i| - |x_R - y_i| + \mathbb{E}[v_R] - \mathbb{E}[v_L] \geq 0.$$

With no learning $\mathbb{E}[v_R] = \mathbb{E}[v_L]$, and so R wins the election if and only if $|x_R| < |x_L|$ (again the 0-probability event that this holds with equality results in a tie). The immediate consequence of this is full convergence of platforms to the median voter’s preferred policy.

Proposition 2. *In the no-learning benchmark, there is a unique equilibrium, in which $x_L^* = x_R^* = 0$.*

Despite their policy motivation, the candidates fully converge. The reason is simple: a candidate must win in order to implement her policy, and a candidate wins if and only if she is preferred by the median. With no learning about valence, valence is strategically irrelevant and we are in the classic case of the median voter theorem.

4 Main Analysis: Selective Exposure

In reality, different voters receive different information about the candidates over the course of an election. The most natural explanation is that voters learn about the candidates from media sources, and different media sources cover different stories. The decision of what stories to cover likely derives from the preferences of an outlet’s consumers—outlets that cater to liberal voters may focus on stories that buttress the liberal candidate’s image, while highlighting the conservative’s gaffes or scandals.

To capture these ideas in the simplest possible way, we assume that a voter’s information about the candidates’ valences is mediated by her ideology. A liberal voter with ideal policy $y_i \leq 0$ receives a pair of signals $s_L^{\text{lib}} \in \{\emptyset\} \cup [0, 1]$ and $s_R^{\text{lib}} \in \{\emptyset\} \cup [0, 1]$, such that

$$s_L^{\text{lib}} = \begin{cases} v_L & \text{if } v_L \geq \mathbb{E}[v_L] \\ \emptyset & \text{if } v_L < \mathbb{E}[v_L], \end{cases} \quad \text{and} \quad s_R^{\text{lib}} = \begin{cases} \emptyset & \text{if } v_R > \mathbb{E}[v_R] \\ v_R & \text{if } v_R \leq \mathbb{E}[v_R]. \end{cases} \quad (3)$$

In words, a liberal voter learns L 's valence if and only if it is better-than-expected, and learns R 's valence if and only if it is worse-than-expected. Analogously, a conservative voter learns R 's valence if and only if it is better-than-expected and learns L 's valence if and only if it is worse-than-expected. That is, a voter with ideal policy $y_i \geq 0$ receives a pair of signals $(s_L^{\text{con}}, s_R^{\text{con}})$:

$$s_L^{\text{con}} = \begin{cases} \emptyset & \text{if } v_L > \mathbb{E}[v_L] \\ v_L & \text{if } v_L \leq \mathbb{E}[v_L], \end{cases} \quad \text{and} \quad s_R^{\text{con}} = \begin{cases} v_R & \text{if } v_R \geq \mathbb{E}[v_R] \\ \emptyset & \text{if } v_R < \mathbb{E}[v_R]. \end{cases} \quad (4)$$

We assume that voters are fully rational and understand this information environment. Our main result characterizes the unique symmetric equilibrium.

Theorem 1. *With selective exposure, there exists a unique symmetric pure strategy Nash equilibrium, in which $x_L^* = -\frac{1}{6}$ and $x_R = \frac{1}{6}$. Policies therefore diverge more than in the no-learning benchmark, but less relative to a setting with full learning.*

Comparing the theorem with previous lemmas, selective exposure leads to less polarization than the full-learning context, and more polarization than the no-learning context. In the appendix we relax the assumption that voters fully account for the informational environment by allowing them to have misspecified beliefs about the informational environment. Here, we focus on developing the intuition by way of local analysis that leads to necessary conditions for platform locations which are uniquely satisfied by the locations in Theorem 1.⁵

A voter with preferred policy y_i prefers R if and only if

$$|x_L - y_i| - |x_R - y_i| + \mathbb{E}[v_R | s_R^j] - \mathbb{E}[v_L | s_L^j] \geq 0. \quad (5)$$

for $j \in \{\text{lib}, \text{con}\}$. Conjecture, as before, that $x_L = -x$ and $x_R = x + \Delta$, for $\Delta \geq 0$. Figure 2 plots L 's valence on the horizontal axis, and R 's valence on the vertical axis.

We begin with the bottom left-hand corner, which corresponds to both candidates having worse-than-expected valence. This implies that conservatives learn L 's valence, but do not learn R 's valence, while liberals learn R 's valence, but do not learn L 's valence. Bayesian updating yields:

$$\mathbb{E}[v_R | s_R^{\text{lib}}] = v_R; \quad \mathbb{E}[v_L | s_L^{\text{lib}}] = \frac{1}{4}; \quad \mathbb{E}[v_R | s_L^{\text{cons}}] = \frac{1}{4}; \quad \mathbb{E}[v_L | s_L^{\text{cons}}] = v_L.$$

⁵The proof presented in the Appendix considers both small and large deviations.




Figure 2: Election outcomes for valence pairs $(v_L, v_R) \in [0, 1]^2$, when $x_L = -x$ and $x_R = x + \Delta$; the thick lines represent $\Delta = 0$, and the dashed lines represent $\Delta > 0$.

If the liberal's valence is sufficiently favorable—i.e., if $v_L > \frac{1}{4} - \Delta$ —then there exists an (interior) indifferent conservative ‘swing’ voter with ideal policy:

$$y_R = \frac{\Delta + v_L - \frac{1}{4}}{2} > 0,$$

such that conservatives with ideal policies in $[0, y_R]$ favor L , while those with ideal policies in $(y_R, 1]$ favor R . Otherwise, all conservatives favor R . Similarly, if the conservative's valence is sufficiently favorable—i.e., if $v_R > \frac{1}{4} + \Delta$ —then there exists an indifferent liberal swing voter with ideal policy:

$$y_L = \frac{\Delta - v_R + \frac{1}{4}}{2} < 0,$$

such that liberals with ideal policies in $[-1, y_L]$ favor L , while those with ideal policies in $(y_L, 0]$ favor R . Otherwise, all liberals favor L .

When $v_R < \frac{1}{4} + \Delta$ and $v_L < \frac{1}{4} - \Delta$, liberals uniformly support L and conservatives uniformly support R , implying that the candidates each win with probability one half. When $v_R < \frac{1}{4} + \Delta$ and $v_L > \frac{1}{4} - \Delta$, liberals and some conservatives support L , implying that L wins. Finally, if $v_R > \frac{1}{4} + \Delta$ and $v_L > \frac{1}{4} - \Delta$, there is both an interior liberal swing voter and an interior conservative swing voter. With uniformly distributed ideal policies, R wins the election if and only if

$$0 - y_L + 1 - y_R > y_L - (-1) + y_R - 0 \iff v_R > v_L + 2\Delta.$$

Comparing with the previous Figure 1 for the full-information benchmark, Figure 2 shows that R 's shift in platform to $x + \Delta$ changes her support by a factor of 2Δ instead of Δ , since her platform shift affects the location of two swing voters, as opposed to one.

The area between the thick and dashed lines in the bottom left quadrant of Figure 2 shows how the regions of victory and ties (in expectation) shift with Δ . We conclude that R 's probability of winning when $v_L, v_R < 1/2$, when $x_L = -x$ and $x_R = x + \Delta$ for small $\Delta > 0$ is

$$\frac{1}{2} \times \underbrace{\left(\frac{1}{4} - \Delta\right) \times \left(\frac{1}{4} + \Delta\right)}_{\text{Probability of a tie}} + \underbrace{\left(\frac{1}{4} - \Delta\right)^2 \frac{3}{2}}_{\text{Probability } R \text{ wins}} = \Delta^2 - \frac{3\Delta}{4} + \frac{1}{8}. \quad (6)$$

The top right quadrant in Figure 2 corresponds to the setting where both candidates' valences are instead *above* their expectations, i.e., a setting where $v_R > \frac{1}{2}$ and $v_L > \frac{1}{2}$. It constitutes a symmetric case to the bottom left quadrant, and R 's probability of winning is again given by (6). The top left-hand quadrant corresponds to $v_R > \frac{1}{2} > v_L$, in which case for sufficiently small $\Delta > 0$ R *always* wins. Notice that, in this context, segmentation *weakens* the responsiveness of R 's victory prospects to small platform changes. The bottom right-hand panel is the opposite case in which $v_R < \frac{1}{2} < v_L$, so that small enough $\Delta > 0$ ensures L always wins.

Putting all this together, Right's probability of winning after the small deviation from x to $x + \Delta > x$ is

$$\pi^{SE}(x, \Delta) \equiv 2 \left(\Delta^2 - \frac{3\Delta}{4} + \frac{1}{8} \right) + \frac{1}{4}, \quad (7)$$

where the first term accounts for the bottom left and top right quadrants in Figure 2, and the second term of one quarter accounts for the top left quadrant in the Figure.

For Δ small, $\pi_R^{SE}(x, \Delta)$ in expression (7) decreases in Δ *faster* than the corresponding probability from the full-learning benchmark, $\pi_R^{FL}(x, \Delta)$. In other words: Right’s winning prospects are *more* sensitive to local platform choices in the context of selective exposure than in the benchmark. Right’s choice of $\Delta \geq 0$ therefore maximizes:

$$\Pi_R^{SE}(\Delta, x) = \pi^{SE}(x, \Delta)(2x + \Delta),$$

and it is easy to verify:

$$\frac{\partial \Pi_R^{SE}(0, x)}{\partial \Delta} = 0 \iff x = \frac{1}{6}.$$

5 Discussion

Departing from a canonical model of two-party competition with policy-motivated candidates and uncertainty about voter preferences, we examine how selective information exposure shapes candidate incentives. Relative to a full-learning benchmark, we unearth the result that selective exposure need not be a force for platform polarization. We show that it can in fact reduce polarization.

Our result offers a counter-point to a received wisdom that takes the presence of echo-chambers and polarized information as a mechanism that necessarily induces greater elite polarization. It highlights that selective exposure can encourage voters to behave more sensitively to elite policy choices by reducing their responsiveness to differences in candidate quality. This intensifies competitive incentives to moderate platforms.

In the Appendix, we pursue an extension in which voters are imperfectly rational: when they do not learn a candidate’s valence, their posterior beliefs imperfectly account for the fact that information is selectively censored. Platform divergence increases smoothly with the degree of imperfect information processing. Consistent with our main presentation, however, polarization is still less than in the full-learning benchmark. Why does even *less* learning about valence further reduce platform competition? The reason is that even though imperfectly Bayesian voters learn less accurate information from “no news,” their failure to internalize their imperfect information processing makes their induced preferences over candidates even more responsive to no news. This net increases valence competition, thereby softening policy competition.

We close by emphasizing that our analysis does not preclude political contexts in which selective exposure may drive elite polarization. On the contrary, comparison of Figures 1 and 2 suggests that this prospect depends on the joint distribution of candidate

valences—in particular, the probability that these valences are realized in distinct subsets of the state space. By re-casting these questions in the canonical Calvert-Wittman framework, however, we hope that our framework complements important empirical studies of selective exposure’s effects on voter preferences (recently, Broockman and Kalla, 2022), and spurs future insights into its effects on politicians’ strategies.

References

- Bernhardt, Dan, Stefan Krasa and Mattias Polborn. 2008. “Political polarization and the electoral effects of media bias.” *Journal of Public Economics* 92(5-6):1092–1104.
- Broockman, David and Joshua Kalla. 2022. “The manifold effects of partisan media on viewers’ beliefs and attitudes: A field experiment with Fox News viewers.” Working Paper.
- Calvert, Randall L. 1985. “Robustness of the multidimensional voting model: Candidate motivations, uncertainty, and convergence.” *American Journal of Political Science* pp. 69–95.
- Chan, Jimmy and Wing Suen. 2008. “A spatial theory of news consumption and electoral competition.” *The Review of Economic Studies* 75(3):699–728.
- Duggan, John and Cesar Martinelli. 2011. “A spatial theory of media slant and voter choice.” *The Review of Economic Studies* 78(2):640–666.
- Glaeser, Edward, Giacomo Ponzetto and Jesse Shapiro. 2005. “Strategic extremism: Why Republicans and Democrats divide on religious values.” *Quarterly Journal of Economics* pp. 1283–1330.
- Hu, Lin, Anqi Li and Ilya Segal. 2019. “The politics of personalized news aggregation.” *arXiv preprint arXiv:1910.11405* .
- Iyengar, Shanto and Kyu S Hahn. 2009. “Red media, blue media: Evidence of ideological selectivity in media use.” *Journal of communication* 59(1):19–39.
- Mutz, Diana and Paul Martin. 2001. “Facilitating communication across lines of political difference: The role of mass media.” *American political science review* 95(1):97–114.

- Pierson, Paul and Eric Schickler. 2020. "Madison's constitution under stress: A developmental analysis of political polarization." *Annual Review of Political Science* 23:37–58.
- Polborn, Mattias K and David T Yi. 2006. "Informative positive and negative campaigning." *Quarterly Journal of Political Science* 1(4):351–371.
- Virág, Gábor. 2008. "Playing for your own audience: Extremism in two-party elections." *Journal of Public Economic Theory* 10(5):891–922.
- Wittman, Donald. 1983. "Candidate Motivation: A Synthesis of Alternative Theories." *American Political Science Review* 77:142–157.

Supplementary Online Appendix: Proofs of Propositions

Proof of Proposition 1. Given $x_L \leq 0$, R chooses $x_R \in [0, 1]$ to maximize

$$V_R^{FL}(x_L, x_R) = \pi_R^{FL}(x_L, x_R)(x_R - x_L),$$

where

$$\pi_R^{FL}(x_L, x_R) \equiv \int_0^{1-x_R-x_L} \left(\int_{v_L+x_R+x_L}^1 dv_R \right) dv_L.$$

An interior best response $\hat{x}_R(x_L)$ by candidate R satisfies the FONC

$$\frac{\partial \pi_R^{FL}(x_L, \hat{x}_R(x_L))}{\partial x_R} (\hat{x}_R(x_L) - x_L) + \pi_R^{FL}(x_L, \hat{x}_R(x_L)) = 0. \quad (8)$$

Substituting primitives into (8) yields two possible solutions: $x_R = 1 - x_L$, and $x_R = \frac{1+x_L}{3}$, and direct computation yields $V_R^{FL}(x_L, \frac{1+x_L}{3}) > V_R^{FL}(x_L, 1)$. Finally, the corresponding SOC is

$$\frac{\partial^2 V_R^{FL}(x_L, x_R)}{\partial x_R^2} = \frac{\partial^2 \pi_R^{FL}(x_L, x_R)}{\partial x_R^2} (x_R - x_L) + 2 \frac{\partial \pi_R^{FL}(x_L, x_R)}{\partial x_R} < 0. \quad (9)$$

Substituting primitives reveals (9) is strictly negative if and only if $x_R < \frac{2-x_L}{3}$, which is true for any $x_R \in [0, 1)$, for all $x_L \leq 0$. We conclude that $\hat{x}_R(x_L) = \frac{1+x_L}{3}$ is R 's unique best response to $x_L \leq 0$. The problem for candidate L is analogous. Solving the mutual best responses yields the unique pair $x_R^* = \frac{1}{4}$ and $x_L^* = -\frac{1}{4}$. Finally, observe that if $x_L = 0$ the corner solution of x_R yields a strictly lower payoff than any policy $0 < x_R < 1$ and thus we have also ruled out $(0, 0)$. \square

Proof of Proposition 2. The result is immediate from the discussion immediately preceding the Proposition's statement in text. \square

Proof of Theorem 1. We prove the Theorem in a more general framework that parameterizes voter bias in Bayesian inference. Specifically, we assume that for $\beta \leq 1$,

$$\begin{aligned}\mathbb{E}[v_L | s_L^{\text{con}} = \emptyset] &= \mathbb{E}[v_R | s_R^{\text{lib}} = \emptyset] = \beta \frac{3}{4} + (1 - \beta) \frac{1}{2} \\ \mathbb{E}[v_R | s_R^{\text{con}} = \emptyset] &= \mathbb{E}[v_L | s_L^{\text{lib}} = \emptyset] = \beta \frac{1}{4} + (1 - \beta) \frac{1}{2}.\end{aligned}$$

The interpretation is that when a voter receives no news, her posterior is a convex combination of (1) the “correct” Bayesian posterior, with weight β , and (2) a “naive” posterior—the prior belief—that fails to account for how the outlet censors information about valence. Our benchmark model corresponds to the special case of $\beta = 1$. All other aspects of the model are unchanged.

We now show that a symmetric pure strategy equilibrium exists only for sufficiently high values of β , and platform divergence in a symmetric equilibrium decreases in β .

Theorem 1*. *There is a threshold $\bar{\beta} \approx 0.456612$ such that in the generalized model:*

1. *For all $\beta < \bar{\beta}$, there is no symmetric pure strategy Nash equilibrium.*
2. *For all $\beta > \bar{\beta}$, there is a unique symmetric pure strategy Nash equilibrium with platforms $x_L = -\frac{1}{2(2+\beta)}$ and $x_R = \frac{1}{2(2+\beta)}$.*

Recognize that the case of $\beta = 1$ corresponds to Theorem 1 in the main text.

Proof of Theorem 1*, part 1. For convenience, we work with a re-parameterized posterior expectation given no news, namely $b(\beta) = \frac{\beta}{4}$. The posterior expectation is then either $\frac{1}{2} + b$ or $\frac{1}{2} - b$.

Given the media strategies and consumption choices, we can partition realizations of v_L, v_R into four qualitatively distinct patterns of inference for citizens on the left and right halves of the policy space. In particular, each of the following four events occurs with probability $\frac{1}{4}$.

- Event A: $\max\{v_L, v_R\} < \frac{1}{2}$ so citizens with $y_i > 0$ evaluate the candidates at expected valences of $(v_L, \frac{1}{2} - b)$ and citizens with $y_i < 0$ evaluate the candidates at expected valences of $(\frac{1}{2} - b, v_R)$.
- Event B: $\min\{v_L, v_R\} > \frac{1}{2}$ so citizens with $y_i > 0$ evaluate the candidates at expected valences of $(\frac{1}{2} + b, v_R)$ and citizens with $y_i < 0$ evaluate the candidates at expected valences of $(v_L, \frac{1}{2} + b)$.

- Event C: $v_L < \frac{1}{2} < v_R$ so citizens with $y_i > 0$ evaluate the candidates at expected valences of (v_L, v_R) and citizens with $y_i < 0$ evaluate the candidates at expected valences of $(\frac{1}{2} - b, \frac{1}{2} + b)$.
- Event D: $v_R < \frac{1}{2} < v_L$ so citizens with $y_i > 0$ evaluate the candidates at expected valences of $(\frac{1}{2} + b, \frac{1}{2} - b)$ and citizens with $y_i < 0$ evaluate the candidates at expected valences of (v_L, v_R) .

We refer to voters with $y_i \leq 0$ that are tuning in to the left wing media as L voters and voters with $y_i \geq 0$ tuning in to the right wing outlet as R voters.

Suppose that the candidates locate at platforms $(-x, x)$. Fixing $x_L = -x$, we derive necessary conditions on x for candidate R to not have a profitable deviation to $x + \Delta$ for $\Delta \neq 0$. Because the expected utility to R captures different effects to moderating or becoming more extreme, this step requires considering separately the cases of deviations with $\Delta > 0$ and $\Delta < 0$.

A sufficient condition for candidate i obtaining a majority is then that both of the following conditions hold: voters with $y = 0$ who consume the left outlet prefer candidate i and those with $y = 0$ who consume the right outlet also prefer candidate i . Candidate i can also win if voters with $y = 0$ who consume the left outlet prefer candidate R and those with $y = 0$ who consume the right outlet prefer candidate L and the indifferent voter consuming i 's outlet is more moderate than the indifferent voter consuming the other candidate's outlet.

Right-side deviations: Consider the case of a deviation to $x + \Delta$ with $\Delta \geq 0$. The payoffs to R for the three possible outcomes (R wins, L wins, and the vote is tied) are:

$$V_R(R \text{ wins}) = x + \Delta - 1, \quad V_R(L \text{ wins}) = -x - 1, \quad V_R(\text{tie}) = -1 + \frac{\Delta}{2}.$$

Now consider the four events that partition realizations of v .

Event A: With $\Delta \in (0, \frac{1}{2} - b)$, R voters with ideal $y = 0$ vote L if $v_L - x > \frac{1}{2} - b - (x + \Delta)$ which is equivalent to $v_L > \frac{1}{2} - b - \Delta$ and so the probability that R voters with ideal $y = 0$ vote L in Event A is $2(b + \Delta)$. Similarly, L voters with ideal $y = 0$ vote R if $v_R > \frac{1}{2} - b + \Delta$ so the probability of this is $2(b - \Delta)$. Observe that v_R and v_L are independent conditional on Event A. If both L and R voters with ideal point $y = 0$ vote for the the *opposite* party (which happens with probability $4(b^2 - \Delta^2)$ conditional on Event A) then R wins if $v_R - v_L > 2\Delta$. Otherwise L wins. Thus Conditional on this sub

event the probability that R (resp. L) wins is $\frac{b-\Delta}{2(b+\Delta)}$ (resp. $\frac{b+3\Delta}{2(b+\Delta)}$). Notice now that if $\Delta \in [b, \frac{1}{2} - b]$ then R cannot win and L will win as long as an R voter with $y = 0$ votes L , which happens with probability $2(b + \Delta)$. Finally, if $\Delta > \frac{1}{2} - b$ then the probability of this last event is actually 1 so L wins.

Combining these conclusions we obtain

$$\text{prob}(R \text{ wins}|A) = \begin{cases} 2(b - \Delta)(1 - 2(b + \Delta)) + 2(b - \Delta)^2 & \Delta \leq b \\ 0 & b < \Delta \end{cases}$$

$$\text{prob}(L \text{ wins}|A) = \begin{cases} (1 - 2(b - \Delta))2(b + \Delta) + 2(b - \Delta)(b + 3\Delta) & \Delta \leq b \\ 2(b + \Delta) & b < \Delta \leq \frac{1}{2} - b \\ 1 & \frac{1}{2} - b < \Delta \end{cases}$$

and the probability of a tie is $1 - \text{prob}(L \text{ wins}|A) - \text{prob}(R \text{ wins}|A)$.

Event B: R voters with ideal $y = 0$ vote for L if $\frac{1}{2} + b - x > v_R - x - \Delta$. This occurs if $v_R < \frac{1}{2} + b + \Delta$ and so the probability that R voters with $y = 0$ vote for L in Event B is $2(b + \Delta)$. L voters with $y = 0$ vote for R if $v_L - x < \frac{1}{2} + b - x - \Delta$. This occurs if $v_L < \frac{1}{2} + b - \Delta$ and so the probability that L voters with $y = 0$ vote for R is $2(b - \Delta)$. We see that the probabilities of the three payoff relevant events conditional on Event B are identical to the corresponding probabilities conditional on Event A.

Event C: For sufficiently small $\Delta > 0$ (namely $b > \Delta > 0$, R always wins. But with larger deviations, the outcome depends on the realization of v_L and v_R . Recall that R voters condition on the realizations of v_L, v_R and L voters condition on neither taking the valence advantage for R as $2b$. If $b < \Delta < 2b$, a measure of $b - \frac{1}{2}\Delta$ (resp. $\min\{0, \frac{1}{2}(\Delta - (v_R - v_L))\}$) of L (resp. R) voters vote for the opposite party, and R 's winning probability is $8(\Delta - b)^2$. If $\Delta > 2b$, no L voter will vote for R , so the election will be tied if no R voter votes for L either (which happens if $v^R - v^L > \Delta$). Otherwise, L wins. Thus we obtain:

$$\text{prob}(R \text{ wins}|C) = \begin{cases} 1 & \Delta \leq b \\ 1 - 8(\Delta - b)^2 & b < \Delta \leq 2b \\ 0 & 2b < \Delta \end{cases}$$

$$\text{prob}(L \text{ wins}|C) = \begin{cases} 0 & \Delta \leq b \\ 8(\Delta - b)^2 & b < \Delta \leq 2b \\ 2\Delta^2 & 2b < \Delta \leq \frac{1}{2} \\ 1 - 2(1 - \Delta)^2 & \frac{1}{2} < \Delta \end{cases}$$

Event D: As long as $\Delta > 0$, L voters with $y = 0$ vote for L and because $b > 0$ R voters with $y = 0$ vote for L . Thus the outcome in Event D is a certain L win.

Now, combining the above observations yields an expected utility for R . We normalize the expected utility by multiplying it by 2 in all cases below.

If $\Delta < b$:

$$\begin{aligned} \mathbb{E}V_R^+ \propto & \underbrace{[2(b - \Delta)(1 - 2(b + \Delta)) + 2(b - \Delta)^2 + \frac{1}{2}]}_{\text{probability that R wins}}(x + \Delta) + \\ & \underbrace{[2(b + \Delta)(1 - 2(b - \Delta)) + 2(b - \Delta)(b + 3\Delta) + \frac{1}{2}]}_{\text{probability that L wins}}(-x) + \\ & \underbrace{[(1 - 2(b - \Delta))(1 - 2(b + \Delta))]}_{\text{probability of a tie}}\left(\frac{\Delta}{2}\right) \end{aligned} \quad (10)$$

If $b < \Delta \leq \min\{2b, \frac{1}{2} - b\}$:

$$\begin{aligned} \mathbb{E}V_R^+ \propto & \frac{1}{2}[1 - 8(\Delta - b)^2](x + \Delta) + [2(b + \Delta) + \frac{1}{2}8(\Delta - b)^2 + \frac{1}{2}](-x) + \\ & [1 - 2(b + \Delta)]\left(\frac{\Delta}{2}\right) \end{aligned} \quad (11)$$

Assuming $b \leq \frac{1}{6}$, if $2b \leq \Delta < \frac{1}{2} - b$:

$$\mathbb{E}V_R^+ \propto [2(b + \Delta) + \frac{1}{2}2\Delta^2 + \frac{1}{2}](-x) + [(1 - 2(b + \Delta)) + \frac{1}{2}(1 - 2\Delta^2)]\left(\frac{\Delta}{2}\right) \quad (12)$$

Assuming $b > \frac{1}{6}$, if $\frac{1}{2} - b \leq \Delta < 2b$:

$$\mathbb{E}V_R^+ \propto \frac{1}{2}[1 - 8(\Delta - b)^2](x + \Delta) + [1 + \frac{1}{2}8(\Delta - b)^2 + \frac{1}{2}](-x) \quad (13)$$

If $\max\{2b, \frac{1}{2} - b\} < \Delta \leq \frac{1}{2}$:

$$\mathbb{E}V_R^+ \propto [1 + \frac{1}{2}(2\Delta^2) + \frac{1}{2}](-x) + [\frac{1}{2}(1 - 2\Delta^2)]\left(\frac{\Delta}{2}\right) \quad (14)$$

And finally, if $\frac{1}{2} < \Delta$:

$$\mathbb{E}V_R^+ \propto [1 + \frac{1}{2}(1 - 2(1 - \Delta)^2) + \frac{1}{2}](-x) + [\frac{1}{2}(2(1 - \Delta)^2)]\left(\frac{\Delta}{2}\right) \quad (15)$$

Left-side deviations: Consider now the case of a deviation to $x + \Delta$ with $\Delta \leq 0$, and note that the parties' win probabilities for Events A and B are the same as in the case of $\Delta \geq 0$ as long as $\Delta > -\frac{b}{3}$. Thus for either event $\mathcal{E} = A, B$,

$$\text{prob}(R \text{ wins}|\mathcal{E}) = \begin{cases} 2(b - \Delta)(1 - 2(b + \Delta)) + 2(b - \Delta)^2 & -\frac{b}{3} \leq \Delta \\ 2(b - \Delta)(1 - 2(b + \Delta)) + 4(b^2 - \Delta^2) & -b \leq \Delta < -\frac{b}{3} \\ 2(b - \Delta) & b - \frac{1}{2} \leq \Delta < -b \\ 1 & \Delta < b - \frac{1}{2} \end{cases}$$

$$\text{prob}(L \text{ wins}|\mathcal{E}) = \begin{cases} (1 - 2(b - \Delta))2(b + \Delta) + 2(b - \Delta)(b + 3\Delta) & -\frac{b}{3} \leq \Delta \\ (1 - 2(b - \Delta))2(b + \Delta) & -b \leq \Delta < -\frac{b}{3} \\ 0 & \Delta < -b \end{cases}$$

Event C: As long as $\Delta < 0$, R voters vote for R and because $b > 0$, L voters votes for R . Thus the outcome is that R wins with certainty.

Event D: For $\Delta \geq -b$, L wins for sure. When $\Delta < -b$, R wins if the valence difference is small. The winning probabilities are given by:

$$\text{prob}(R \text{ wins}|A) = \begin{cases} 0 & -b \leq \Delta \\ 8(\Delta + b)^2 & -b - \frac{1}{4} < \Delta \leq -b \\ 1 - 2(1 - 2(\Delta + b))^2 & -b - \frac{1}{2} \leq \Delta < -b - \frac{1}{4} \end{cases}$$

$$\text{prob}(L \text{ wins}|A) = \begin{cases} 1 & -b \leq \Delta \\ 1 - 8(\Delta + b)^2 & -b - \frac{1}{4} < \Delta \leq -b \\ 2(1 - 2(\Delta + b))^2 & -b - \frac{1}{2} \leq \Delta < -b - \frac{1}{4} \end{cases}$$

$$\text{prob}(R \text{ wins}|D) = 0, \quad \text{prob}(L \text{ wins}|D) = 1 - 2\Delta^2, \quad \text{prob}(\text{tie}|D) = 2\Delta^2$$

First Order Conditions: Differentiating the right side of Equation (10) with respect to Δ yields the following first-order necessary condition for an optimal deviation:

$$1 + 12\Delta^2 - (4 + 8b)x - (4 + 8b - 16x)\Delta = 0 \tag{16}$$

which holds at $\Delta = 0$ if and only if $x = \frac{1}{4(1+2b)}$. This conjectured symmetric Nash equilibrium $(-\frac{1}{4(1+2b)}, \frac{1}{4(1+2b)})$ yields 0 expected utility for both parties by symmetry. Recall that utility was normalized above. Consider a deviation $b < \Delta \leq \min\{2b, \frac{1}{2} - b\}$

and let

$$\Delta' = \min\left\{\frac{1}{12}\left(-1 + 8b - \frac{2 - \sqrt{15 + 36b - 12b^2 - 48b^3 + 64b^4}}{1 + 2b}\right), 2b, \frac{1}{2} - b\right\} \quad (17)$$

Substituting $\Delta = \Delta'$ into the expression on the right of (11) shows that the value is strictly positive if and only if $b < \bar{b} \approx 0.114153$ or equivalently $\beta < \bar{\beta} = 4\bar{b} \approx 0.456612$. Since equation (16) holding at $\Delta = 0$ is necessary, this profitable deviation implies a symmetric pure strategy Nash equilibrium does not exist if $\beta < \bar{\beta}$.

Proof of Theorem 1*, part 2. Again, let $b = \frac{\beta}{4}$ and assume now that $b \geq \bar{b} = \frac{\bar{\beta}}{4} \approx 0.114153$. Without loss and invoking symmetry, we continue to focus on possible deviations by R . Fixing $x = \frac{1}{4(1+2b)}$, we argue that R 's expected utility, as expressed in the proof of Part 1, is weakly negative for any $\Delta \neq 0$. (Recall that expected utility is 0 at the conjectured symmetric equilibrium, where $\Delta = 0$.)

Right-side deviations: Suppose that $\Delta < b$, and observe that (10) is cubic in Δ and $\Delta = 0$ is a local maximum. To make sure $\Delta = 0$ is also the global maximum over $[0, b]$, it suffices to compare $\Delta = 0$ with $\Delta = b$. Plugging in $x = \frac{1}{4(1+2b)}$, we see that the expression on the right of (10) equals 0 at $\Delta = 0$ and is negative at $\Delta = b$ for $b > 0$.

Now suppose that $\Delta < \min\{2b, 1 - \frac{b}{2}\}$, and notice that the expression on the right of (11) is cubic in Δ and goes to $-\infty$ as $\Delta \rightarrow \infty$. Its first-order condition, after plugging in $x = \frac{1}{4(1+2b)}$, is:

$$1 - 16b^3 - 12\Delta - 24\Delta^2 + 4b^2(-3 + 16\Delta) + b(10 + 24\Delta - 48\Delta^2) = 0$$

The greater root $\frac{1}{12}\left(-1 + 8b - \frac{2 - \sqrt{15 + 36b - 12b^2 - 48b^3 + 64b^4}}{1 + 2b}\right)$ is the local maximum. Since the local minimum $\frac{1}{12}\left(-1 + 8b - \frac{2 + \sqrt{15 + 36b - 12b^2 - 48b^3 + 64b^4}}{1 + 2b}\right)$ is always less than $2b$, it suffices to check that R 's expected payoff is weakly negative at Δ' if $b > \bar{b}$ where Δ' is defined in (17). Simple numerical calculations shows that this is the case.

Now suppose that $b \leq \frac{1}{6}$ and $2b \leq \Delta < \frac{1}{2} - b$, and note that the expression in (12) is cubic in Δ and goes to $-\infty$ as $\Delta \rightarrow \infty$. Its first-order condition, after plugging in $x = \frac{1}{4(1+2b)}$, is:

$$1 - 8b^2 - 10\Delta - 6\Delta^2 + 2b(1 - 8\Delta - 6\Delta^2) = 0$$

The local maximum (minimum) is the greater (smaller) root and always (never) in $[2b, \frac{1}{2} - b]$ if $b \geq \bar{b}$. It thus suffices to check that at the local maximum R 's expected utility is weakly negative. This is indeed the case.

Now suppose that $b > \frac{1}{6}$ and $\frac{1}{2} - b \leq \Delta < 2b$. Expression (13) is cubic in Δ and goes to $-\infty$ as $\Delta \rightarrow \infty$. Its first-order condition, after plugging in $x = \frac{1}{4(1+2b)}$, is:

$$2 - 16b^3 - 8\Delta - 24\Delta^2 + 8b^2(-1 + 8\Delta) + 2b(5 + 16\Delta - 24\Delta^2) = 0$$

The local minimum is the smaller root and always less than $\frac{1}{2} - b$ for any $b \in [\bar{b}, \frac{1}{4}]$. It thus suffices to show the local maximum value is no greater than 0 for any $b \in [\bar{b}, \frac{1}{4}]$. Simple calculations show it is weakly negative.

Now suppose $\max\{2b, \frac{1}{2} - b\} < \Delta \leq \frac{1}{2}$. Expression (14) is cubic in Δ , and goes to $-\infty$ as $\Delta \rightarrow \infty$. Its first-order condition, after plugging in $x = \frac{1}{4(1+2b)}$, is:

$$\frac{1}{2} - \frac{\Delta}{2 + 4b} - 3\Delta^2 = 0$$

Since the interval $[\max\{2b, \frac{1}{2} - b\}, \frac{1}{2}]$ is always strictly between the two roots of the first-order condition for any $b \in [\bar{b}, \frac{1}{4}]$, one only needs to consider R 's expected utility at the outer limit $\frac{1}{2}$. Simple numerical calculations shows it is negative for any $b \in [\bar{b}, \frac{1}{4}]$.

Finally, suppose $\Delta > \frac{1}{2}$. Expression (15) is cubic in Δ and goes to ∞ as $\Delta \rightarrow \infty$. Its first-order condition, after plugging in $x = \frac{1}{4(1+2b)}$, is:

$$\frac{1}{2}(-3 + 3\Delta)\Delta + b(1 - 4\Delta + 3\Delta^2) = 0$$

The local maximum is always less than $\frac{1}{2}$, and the local minimum exceeds the policy space. It thus suffices to check the limit of (15) as $\Delta \rightarrow \frac{1}{2}^+$. Simple numerical calculations show it is negative for any $b \in [\bar{b}, \frac{1}{4}]$.

Left-side deviations: First we argue that $\Delta < 0$ cannot be a profitable deviation. When $x = \frac{1}{4(1+2b)}$ as conjectured, equation (10) is locally maximized at $\Delta = 0$ and monotonically increasing in Δ for $\Delta < 0$. Inspecting the wining probabilities for $\Delta < 0$ in the proof of Part 1 shows that for any $-b \leq \Delta < 0$, R 's wining (losing) probability is weakly lower (higher) than the corresponding probability incorporated in (10) given the same Δ value. This means such deviations would not be profitable for R .

Since parties cannot move across 0, we can disregard the cases where $\Delta < b - \frac{1}{2}$ or $\Delta < -b - \frac{1}{4}$ as $x = \frac{1}{4(1+2b)} \leq \frac{1}{4}$. So the only case left to consider is $-b - \frac{1}{4} \leq \Delta < -b$. In this case,

$$\begin{aligned} \mathbb{E}V_R^- \propto & (2(b - \Delta) + \frac{1}{2}8(\Delta + b)^2)(x + \Delta) + \frac{1}{2}(1 - 8(\Delta + b)^2)(-x) \\ & + (1 - 2(b - \Delta))\frac{\Delta}{2} \end{aligned} \quad (18)$$

This expression is cubic in Δ and goes to ∞ in Δ . Its first-order condition, after plugging in $x = \frac{1}{4(1+2b)}$, is:

$$1 + 16b^3 + 4\Delta + 24\Delta^2 + 4b^2(3 + 16\Delta) + 2b(7 + 12\Delta + 24\Delta^2) = 0$$

The left side has no real root for $b \in [\bar{b}, \frac{1}{4}]$. This implies (18) is monotonically increasing in Δ for $b \in [\bar{b}, \frac{1}{4}]$, and it suffices to check its limit as $\Delta \rightarrow -b - \frac{1}{4}^-$, which is negative.

Uniqueness: We have so far proved that $(-\frac{1}{4(1+2b)}, \frac{1}{4(1+2b)})$ is a symmetric pure strategy Nash equilibrium. To establish uniqueness in the family of symmetric pure strategy Nash equilibria, observe that (16) holding at $\Delta = 0$ is necessary for any symmetric pure strategy Nash equilibrium. \square