

Voter polarization and extremism*

Jon Eguia[†] Tai-Wei Hu[‡]

July 21, 2023

Abstract

We present a theory of endogenous policy preferences and electoral competition with boundedly rational voters who find it costly to process detailed information. Voters are otherwise fully rational, and they strategically choose how much memory to devote to processing political information. We find that even if all voters start with a common prior such that they all prefer a moderate policy over extreme alternatives to the left or the right, and even if voters observe only common signals that in the limit would assure a perfectly rational agent that the moderate policy is indeed best for everyone, a majority of voters eventually become extreme and the electorate becomes polarized: some voters support the left policy, and some support the right policy. Two fully rational parties respond by proposing extreme platforms, and thereafter, the policy outcome in every period is extreme.

Keywords: Polarization, extremism, costly information processing, bounded rationality, rational inattention, electoral competition.

*We are grateful to David Austen-Smith, Renee Bowen, Ben Bushong, Carlo Prato, David Miller, Georgy Egorov, Dan Bernhardt and audiences at seminars at the U. Michigan, CERGEI, Hong Kong U., U. Birmingham, U. York, U. Toronto, the 2023 North American Meeting of the Econometric Society, the 2023 John Morgan Alumni conference, the 2023 Great Lakes Political Economy Theory conference, and the 2023 SAET conference for useful comments and feedback.

[†]Michigan State University, email: eguia@msu.edu.

[‡]University of Bristol, email: taiwei.hu@bristol.ac.uk.

1 Introduction

We study how voters' cost of processing political information relates to the polarization of the electorate. We find conditions under which voters who would embrace a policy consensus if they were fully informed or fully uninformed, polarize into two opposed extremes due to the constrained optimal way in which they process common information. To explain this phenomenon, we microfound voters' behavior with explicit formulations of voters' motives to vote and the information-processing costs they bear.

Consider an electorate facing a set of different policies, and a list of candidates running for office. Voters have preferences over certain socio-economic outcomes (their individual wealth, society's wealth and inequality, pollution, etc.) as a primitive. Whereas, a voter's preferences over policies are endogenously derived from the voter's preferences over outcomes, and from her information about how each policy would affect the outcomes over which her primitive preferences are defined.

Given their endogenous preferences over policies, voters derive their preferences over their voting alternatives (i.e. over each candidate in the ballot, and abstention) by combining an outcome-oriented motivation with an "expressive" motivation. The outcome-oriented motivation depends on the effect of their vote on society's policy choice; in a large election, this motivation vanishes as the probability that an individual vote has any effect becomes negligible (Ledyard 1984). The expressive motivation to vote for a given candidate is that voters enjoy supporting good causes: by supporting a cause or a party, a voter *becomes* a supporter of this party or cause (Schuessler 2000), and voters enjoy identifying as members of a group that champions good policies.¹

If a voter's reason to vote is that she enjoys supporting good policies, in order to vote for a candidate, the voter needs to believe that the policies that this candidate would implement are good.² Since voters are uncertain about the effect of policies over downstream socio-economic outcomes, they need to process information to determine

¹ "*Prosperity has many parents; adversity only one*" (Tacitus 2014 [94 AD], page 53, in the original language: "*Prospera omnes sibi vindicant; adversa uni imputatur*").

²This motivation to vote contingent on "getting it right" is unlike, say, an intrinsic motivation to vote regardless of who or what one votes for (Riker and Ordershook 1968) or the motivation to vote for a candidate whose identity is exogenously given as part of the voter's type (Coate and Conlin 2004; Feddersen and Sandroni 2005).

which are the best policies, and thus which candidates (if any) are worth supporting by voting for them.

In the information age, voters are flooded with information, freely available across multiple media platforms. If processing information were costless, a rational voter would use all this freely available information and Bayes rule to formulate a precise posterior belief about the mapping from policies to the outcomes of interest, and she would vote to maximize her expected utility according to that posterior belief. Alas, processing information is costly. Voters need to weigh this cost against the benefit of being better informed (Downs 1957; Davis, Hinich and Ordeshook 1970). Strategic voters exposed to an over-abundance of political news, and with limited capacity to correctly process and store all this information need a simpler, (constrained) optimal rule to determine how to process information, which important pieces of information to keep in mind, and which ones to discard and to forget about.

We construct a theory of political participation under the following two premises. First, voters enjoy supporting a policy in proportion to their expected utility if this policy were implemented. Further, each voter enjoys voting for candidates who support the policy the voter thinks is best for her, while she dislikes voting for candidates who support policies that are very different from the one she thinks best. These expressive preferences over voters' own actions may be weighed arbitrarily little relative to the weight on standard preferences over outcomes, but as long as their importance is not zero, they will influence voters' behavior. Second, voters decide how much cost to incur assimilating and processing political information by weighing how much this information helps them determine which policy is best for them, against the difficulty of processing this information.

We relax the assumption of perfect information processing, and assume instead that it is costly for voters to keep track of precise details about the information they have observed. We formalize this cost of processing information by assuming that voters have a limited set of "memory states" that aggregate past information. A fully rational agent would be one with unlimited memory states, who can record any minutely different history of informative signals (or any sufficient statistic such as a posterior belief recorded to any degree of precision) into its own memory state. A more realistic voter finds each additional memory state costly, so she resorts to a

limited memory capacity with only finitely many different memory states, to deal with all the information. This voter must lump sufficiently similar information histories into the same memory state, resulting on a blurred belief about the state of affairs. Voters trade off this cost of processing information, with the benefit of a more precise understanding about how alternative policies would affect outcomes of interest, and thus deriving more satisfaction from supporting the alternative they think would have the best effect.

We model this endogenously imperfect way to process information using finite automata, as in Wilson (2014). An automaton consists of a collection of finitely many memory states, and a transition rule taking the process from a given memory state to another depending on the information received in that period. Limited memory capacity implies that agents' beliefs are categorized discretely instead of being represented by arbitrarily precise posteriors. This departs from the rational inattention literature (Maćkowiak, Matějka and Wiederholt 2021), because, under rational inattention, agents strategically choose which information to attend to, and then update in a fully rational way with a precise posterior belief. In contrast, under our formulation agents strategically choose their memory capacity and, constrained by this capacity, they choose how best to process all the information they observe.

Given the number of memory states an agent chooses, each memory state represents the agent's state of mind or her set of thoughts about the uncertain state of the world relevant for her voting behavior. Each memory state corresponds to a category of partial histories and a qualitative "belief". This categorization is analogous to the one made by an agent who evaluates sovereign default risk solely based on credit ratings "A", "B", "C", etc., and who then forgets all the detailed information that fed into the rating, including all information that make some countries rated "B" less likely to default than others with the same rating. A distinct feature of this updating process is the discreteness in information processing; there is no "*straw that breaks the camel's back*" because sufficiently small bits of information do not induce a sufficient update to change from one memory state to another, and are thereafter forgotten; rather, it takes quite a substantial bit of information to trigger a transition across memory states.

Agents who use different updating processes may process the same news very differently. To avoid infinite costs, each agent will choose only a finite number of distinct

memory states, accommodating only finitely many different views of the world. A voter's optimal number of memory states and the optimal rule to transition across memory states are determined endogenously by her preference, the informational environment, and the cost structure. We assume homogeneous costs of memory across voters but we allow some minimal heterogeneity in preferences, and this heterogeneity generates endogenously the difference in the constrained optimal updating processes.

We consider the following environment. There is a set of three policies: a moderate policy and two extreme ones, one on each side (left and right) of the ideological divide. Each policy matches one state of the world, and is Pareto superior to the other two in that state: every voter strictly prefers the socio-economic outcome if the state-matching policy is implemented than the outcome if any other policy is implemented. Signals that the state is "moderate" (which we take it to be the normal, expected state) are abundant in every state and hence are commonplace (say a day with no news headlines about street violence), while the signals that shift preferences toward extreme policies (say a shocking case of either coordinated police violence against peaceful demonstrators, or of coordinated violence by armed rioters against peaceful bystanders and police) will make big news as they are rare but very informative.

Our main result is that in this environment, sufficiently impatient voters with costly memory capacity polarize once they observe an extreme signal. Even if all voters start with a common prior about the state of the world and under this prior all voters prefer the moderate policy, even if voters only observe common signals, and even if these common signals are such that any voter who processed information perfectly would formulate a posterior that the moderate policy is indeed best for everyone, given their limited memory capacity, a majority of voters end up favoring extreme policy alternatives, and diverging in their preferences: some prefer the left policy and others prefer the right policy. Chasing these extreme voters, two office-motivated parties offer extreme policy proposals, and the implemented policy becomes extreme.

The underlying mechanism that explains our result is that after seeing any big news that suggest an extreme state of the world, voters ignore any subsequent moderate signal. This uneven updating process, in which commonplace moderate signals are ignored while rare extreme ones are heeded, drive voters away toward the political extremes. Even under common signals and with a common prior, voters polarize at

opposite extremes. Why? Say the common prior is a belief that the state is likely “normal”, and that all voters prefer a “moderate” policy given this prior belief. Once voters commonly observe a signal that suggests the state of the world is extraordinary (which is rare but will happen eventually), all voters agree that the state is indeed more likely to be an extraordinary one that calls for an extreme policy solution. Having concluded that a moderate policy would be unsuitable, voters disagree about which extreme policy is appropriate: some support left policies, while others support right policies. The policy disagreement stems from the heterogeneity in the relative distaste over policy mistakes in one or the other direction under an extreme state of nature.

In this theory, polarization is micro-founded by the individual decisions of each voter to simplify her information environment, by coarsening the partition of possible beliefs under consideration. Polarization is an aggregate phenomenon that can be decomposed as a large number of independent (and disparate) decisions to become extreme made by each individual voter in isolation. According to our theory, polarization is not elite-driven, and it is not driven by the electorate’s network interactions, nor by biased media that reinforces the beliefs of like-minded voters in their own informational bubble. Rather, we show that a large society of Robinson Crusoes, each isolated on their own island, all endowed with a common prior at the time of arrival to their own island, and observing common signals in the night sky each night, would also polarize. Our theory thus provides an account for episodes of ground-up political polarization like the one experienced in the United States after 2009 (Skocpol and Tervor 2019).

We next discuss the related literature. Thereafter, in Section 2 we present a model on information processing and preference formation for impatient voters who face memory costs; in Section 3 we show how an electorate composed of such voters polarizes; and in Section 4 we show that party platforms and implemented policies are extreme in any equilibrium of a stylized electoral competition theory. We discuss our theoretical results in light of new empirical evidence on polarization in Section 5. The main steps of the proofs are in Section 6, with details in the (Online) Appendix.

Related Literature

Our model of learning with finite automata follows the approach in Cover and Hellman (1970) and Wilson (2014).³ Our results contribute to this literature on three fronts. First, while the literature focus on only two states of nature, our model features three with moderation as the status quo and two well-defined extreme states. Polarization means two sets of people moving toward the two extreme states. On the technical front, we also employ completely different techniques to handle three-states as the belief-space is no longer one-dimensional. Second, while Wilson (2014) also obtains belief polarization of agents with different priors, we obtain polarization by heterogeneity in preferences. Moreover, Wilson (2014) only shows that it is *possible* to have divergent beliefs for some signal realization. In our model, we show that under all states of nature and *almost all* signal realizations, polarization inevitably happens and persists. Finally, the literature typically assumes an exogenously given number of memory states, and we endogenize it by introducing a cost.

Our modeling of how voters process information also relates in its motivation and substance to theories of rationally inattentive voters (Prato and Wolton 2016; Matějka and Tabellini 2021). In these models of rational inattention, typically voters can choose from a menu of signals for which more informative ones are more costly. However, agents have perfect abilities to process signals obtained and update them according the Bayes rule. In our model all signals are commonly observed, and the differences arise in what voters do with the signals by introducing costly information processing. In contrast to the rational-inattention models, the constrained optimal rule in our model features a categorization of beliefs according to which small signals, although free to obtain, do not change the voter’s state-of-the-mind, a feature that drives our extremism and polarization result.

Our theoretical finding that voters’ costs of processing information leads to preferences for extremism contributes to a literature documenting political polarization (Abramowitz and Saunders 2008; McCarty, Poole and Rosenthal 2016; Gentzkow 2016), studying its consequences (Gordon and Landa 2017; Buisseret and van Weelden 2022),

³The automata approach to model imperfect processing of information has been recently evaluated by Oprea (2020), and is supported by experimental evidence in Banovetz and Oprea (2023).

suggesting ways to mitigate it (Axelrod, Daymude and Forrest 2021), or explaining some of its causes. Among the latter, Glaeser, Ponzetto and Shapiro (2005); Serra (2010); Bol, Matakos, Troumpounis and Xefteris (2018); Tolvanen, Tremewan and Wagner (2021); and McMurray (2022) focus on candidates' polarization. With regard to voter polarization, it can be triggered by economic growth if contentious moral issues gain salience once basic economic needs are met (Enke, Polborn and Wu 2022); or it can arise if voters choose to follow different sources of information (Nimark and Sundaresan 2019; Che and Mierendorff 2019; Perego and Yuksel 2022); or if they remain uninformed and vote to cancel each other out (Ely 2023); or if they pay disproportionate attention to the issues they care more about (Yuksel, 2022) or to the issues in which the candidates' proposals differ most (Nunnari and Zapal 2020); if they share news with their connections (Bowen, Dmitriev and Galperti 2023); or, even under commonly-observed signals, if voters face ambiguity and are averse to it (Baliga, Hanany and Klibanoff 2013).

Perhaps closest to our work in their linking of voters' memory constraints to polarization are two theories in which voters observe common information, but they process it in a boundedly rational way that leads to polarization. Fryer, Harms and Jackson (2019) assume that voters coarsen the space of signals about the state of the world. In their model, voters reinterpret each uninformative signal as an informative one that conforms with their prior. Voters then update this prior as if the signal had been truly informative; this self-confirming miss-processing of signals, together with heterogeneous priors, leads to polarization. In fact, if some signals are equivocal rather than uninformative, upon observing equivocal signals, fully rational agents with different priors about the meaning of equivocal signals also polarize toward their priors (Benoit and Dubra 2019). In either case, if agents shared a common prior, they would not polarize.

Levy and Razin (2021), like us, present a dynamic theory of electoral competition and polarization with two parties and three possible policies (a moderate one, and an extreme one to each side), in which the driver of polarization is voters' limited temporal memory: voters remember all information for a fixed time, and after this lapse of time they forget. Polarization is candidate-driven, as in Wittman (1983) or Calvert (1985), with cycles of platform polarization and convergence. Meanwhile, voters' beliefs

never polarize, as all voters share commonly updated beliefs. In contrast, our voters’ memory constraint is one of capacity, like computers’ memories, and we complement their account with a theory of voter polarization.

2 The Model

Consider a large democratic society, represented by a set I of voters, with unit mass. The voters are faced with a choice over three policy alternatives in each of infinitely many periods. Let $\mathcal{A} \equiv \{a^L, a^M, a^R\}$ denote the set of alternatives, where a^L denotes a “left” alternative, a^M a “moderate alternative”, and a^R a “right” alternative. Let $\Theta \equiv \{L, M, R\}$ denote the set of possible states of the world (where L , M , and R again respectively denote Left, Moderate, and Right), and let $\theta \in \Theta$ denote a state of the world. All voters share a common prior probability distribution \mathbf{P}_0 over Θ about the state of the world. We envision an environment in which the moderate state M is the most likely, and states L or R represent an extraordinary event or shock. Formally, we assume that the common prior among the agents is such that

$$\mathbf{P}_0(L) = \mathbf{P}_0(R) = p_0 < \frac{1}{4}. \quad (1)$$

In each period $t = 0, 1, 2, \dots$, each voter i chooses which policy alternative to support. Let $a_t^i \in \mathcal{A}$ denote the alternative that voter i supports in period t (more generally we denote individual agent labels as superscripts, and period labels as subscripts). Let $a_t^I \in \mathcal{A}$ denote the policy alternative collectively chosen by society in period t ; in a direct democracy (Section 3), the collective choice a_t^I is the alternative supported by the largest subset of agents; in a representative democracy (Section 4) the collective choice a_t^I is made by the party that wins a competitive election.

Each voter i cares about the policy outcome a_t^I and about the policy a_t^i that she supports, in each period. Let $\lambda \in (0, 1)$ denote the relative weight assigned to the policy outcome, and $1 - \lambda$ the weight assigned to the expressive component of her political preferences, so that in each period t , each voter i derives instantaneous utility

$$\lambda u(a_t^I, \theta, b^i) + (1 - \lambda)u(a_t^i, \theta, b^i), \quad (2)$$

where b^i is voter i 's type as described below.⁴ We assume that voters' intertemporal patience is captured by a discount factor $\delta \in (0, 1)$ across periods, so that the total utility for voter i for an infinite sequence of individual and collective choices is

$$\lambda \sum_{t=0}^{\infty} (\delta)^t u(a_t^I, \theta, b^i) + (1 - \lambda) \sum_{t=0}^{\infty} (\delta)^t u(a_t^i, \theta, b^i), \quad (3)$$

where the first term is the utility from the sequence of policy outcomes, and the second term is the expressive utility from the sequence of individual choices to express support for an alternative.

We assume that in each state $\theta \in \Theta$, every voter derives highest period utility from the policy alternative a^θ that matches the state, so we refer to alternative a^θ as the “correct” alternative in state θ . Each cell in the left matrix in Table 1 shows the utility function $u(a, \theta, b^i)$ as a function of the action a in each column, and of the state of the world θ in each row, and type $b^i \in [-\bar{b}, \bar{b}]$ for each voter i , where $c \in [0, 1]$ and $\bar{b} \in (-1, 1)$. The correct policies are on the diagonal of the matrix. For each given state, utilities are single-peaked with respect to the standard left-to-right order, and, given Assumption (1) on the prior, states Left and Right are ex-ante much less likely than the Moderate state. We thus refer to policy alternatives a^L and a^R , and to states L and R as “extreme.”

Preference parameter $c \in [0, 1]$ captures the disutility of choosing the wrong extreme action in any extreme state, relative to the utility (normalized to zero) of choosing moderation in that same extreme state. It is therefore a convexity parameter of the preferences over the left-to-right order under any given extreme state. We interpret c as a societal “taste for compromise”: if c is high, under a mixed belief about the extreme states, moderation is relatively more appealing than if c is low.

We assume that the distribution over voter types has full support over $[-\bar{b}, \bar{b}]$. Subject to a Moderate state of the world, voters have common-value symmetric preferences over policies, with ideal alternative a^M . Voters also have a common preference order over alternatives in either of the two extreme states of the world, and they all agree

⁴For an overview of citizen's motivations for voting, see Brennan and Lomasky (1993) or a survey by Hamlin and Jennings (2018); and Glazer (1987) for an early theory of elections under expressive voting.

$\theta \backslash a$	a^L	a^M	a^R
L	$1 - b^i$	0	$-c$
M	0	1	0
R	$-c$	0	$1 + b^i$

$\theta \backslash s$	ℓ	m	r
L	$\mu - \epsilon$	$1 - \mu$	ϵ
M	ϵ	$1 - 2\epsilon$	ϵ
R	ϵ	$1 - \mu$	$\mu - \epsilon$

Table 1: Left: payoff matrix; right: signal structure

that alternative a^M is ex-ante the best given their common prior \mathbf{P}_0 .

Type b^i captures the following asymmetry, or **lean** on voter i 's preferences over policy alternatives: voters vary on how much they gain from the correct extreme alternative in each extreme state. We say that a voter “leans left” if she has a stronger preference for the left action in the Left state, than for the right action in the Right state; and that she “leans right” if she has a stronger preference for the right action in the Right state, than for the left action in the Left state. Type b^i is then a measure of this lean, with $b^i < 0$ implying that voter i leans left, and $b^i > 0$ that she leans right. But such leanings only relate to preferences if the state of the world is extreme.

The state of the world, however, is not observable. In period 0 voters rely on the prior to make their decisions, and in each period $t \geq 1$, voters observe a common signal s_t drawn from the set $\mathcal{S} \equiv \{\ell, m, r\}$ independently in each period before making their choices. Conditional on the state of the world $\theta \in \{L, M, R\}$, signal $s \in \{\ell, m, r\}$ is drawn with probability μ_s^θ . Each cell of the right matrix in Table 1 denotes this probability μ_s^θ , as a function of the state θ in each row and the signal s in each column, with $\mu \in (0, 1)$ and $\epsilon \in (0, \mu/2)$.

In each period $t \geq 1$, all voters observe the common signal s_t , and each voter chooses a policy alternative to support (once we introduce an election game in Section 4, we will expand this timeline to let voters also observe the platforms chosen by political parties). For any voter $i \in I$, her choice of an alternative in period t can in principle depend on all relevant past information observed by the voter, which is the sequence $(s_1, \dots, s_t) \in (\mathcal{S})^t$ of common signals up to period t , and the sequence $(a_0^i, \dots, a_{t-1}^i) \in (\mathcal{A})^t$ of voter i 's own actions up to period $t - 1$.

A decision rule for agent i is a function $D^i : \bigcup_{t=0}^{\infty} (\mathcal{S} \times \mathcal{A})^t \rightarrow \mathcal{A}$, where $D^i(\mathcal{S} \times \mathcal{A})^0 = D^i(\emptyset)$ denotes the action taken by voter i in period 0, with no information besides the common prior \mathbf{P}_0 . The decision rule maps each possible sequence of signals and own

actions observed by agent i to the set of alternatives. Since all agents observe the same sequence of signals and this sequence is the only source of information about the state of the world θ , with perfect Bayesian updating, all agents would share the same posterior belief about θ , and this posterior belief, updated in each period according to the latest realization of the common signal, would be the only relevant state variable. The unconstrained optimal rule for each agent can be fully characterized by the agent’s type b^i and by the posterior, p over Θ : for each voter i , the optimal decision rule is to choose alternative a that maximizes $\sum_{\theta \in \Theta} p(\theta)u(a, \theta, b^i)$. We use $\Delta(\Theta)$ to denote the set of posteriors.

Finite automata and implementation

Storing and remembering precise, detailed information is costly for voters. Boundedly rational voters that optimize their choices must take into account this cost in making their decisions. In order for the unconstrained rule described above to be optimal, processing information to compute a precise posterior belief must be costless. Rational agents for which information processing is costly—and these include any voter in any real-world application—will seek to find a constrained optimal rule that is less costly to use. We assume that our agents summarize all observed information using finitely many “memory states,” and update their memory only using the most relevant information, namely, the commonly observed signal. More precisely, we formulate this process as a finite automaton, to be described below.

A **stochastic finite-state automaton (SFSA)** is a list $\langle Q, q_0, \tau, d \rangle$, where Q is a finite set of *memory states*; $q_0 \in Q$ is the initial memory state; $\tau : Q \times \mathcal{S} \rightarrow \Delta(Q)$ is the *transition rule*; and $d : Q \rightarrow \Delta(\mathcal{A})$ is the *decision rule*. Each memory state $q \in Q$ can be interpreted as a state of mind; the transition rule τ then specifies how observing a signal $s \in \mathcal{S}$ triggers a (probabilistic) change in such state of mind; and the decision rule d specifies a probabilistic choice of alternative for each state of mind or, in our jargon, for each “memory state.” Using results from Kalai and Solan (2004), with no loss of generality we can restrict attention to deterministic decision rules, so that $d : Q \rightarrow \mathcal{A}$ and $d(q)$ denotes the alternative chosen whenever the current memory

state is q .⁵ Let \mathcal{Q} denote the set of all such stochastic finite-state automata.

A voter using one of these finite automata only needs to keep track of the memory state $q \in \mathcal{Q}$, to observe the latest signal s , and to remember her transition rule τ and her decision rule d . With just that, she can transition to a new memory state according to her transition function and the signal she observes; and she can take a decision over alternatives according to her decision rule. If the set of memory states \mathcal{Q} is small, this is a simple enough exercise. More complex automata, with more memory states, are costlier to operate.

We refer to the number of memory states in \mathcal{Q} as the “memory capacity” of automaton $\langle \mathcal{Q}, q_0, \tau, d \rangle$ and we assume that using a automaton with finite memory capacity $|\mathcal{Q}|$ costs $\kappa \cdot |\mathcal{Q}|$, for some $\kappa \in \mathbb{R}_{++}$.

Each voter i chooses her automaton to maximize her discounted, total expected utility, by solving the optimization problem

$$\max_{\langle \mathcal{Q}, q_0, \tau, d \rangle \in \mathcal{Q}} \left(\mathbf{E} \left[\lambda \sum_{t=0}^{\infty} (\delta)^t u(a_t^I, \theta, b^i) + (1 - \lambda) \sum_{t=0}^{\infty} (\delta)^t u(d(q_t), \theta, b^i) \right] - \kappa \cdot |\mathcal{Q}| \right). \quad (4)$$

In a large society, the probability that an individual agent’s choice affects the collective choice is negligible—with a unit mass of agents, each agent is infinitesimal, and this probability is exactly zero—and therefore the first summation in the expectation drops out of each voter’s optimization problem (Brennan and Hamlin 1998), which simplifies to finding the automaton that maximizes ex-ante expected expressive utility, net of costs of running the automaton. Formally, an optimal automaton for voter i is one that solves

$$\max_{\langle \mathcal{Q}, \tau, d, q^o \rangle \in \mathcal{Q}} \left(\mathbf{E} \left[(1 - \lambda) \sum_{t=0}^{\infty} (\delta)^t u(d(q_t), \theta, b^i) \right] - \kappa \cdot |\mathcal{Q}| \right). \quad (5)$$

In Section 3 we find the optimal automaton for each voter, and we describe the

⁵We could conceive of an automaton with a transition rule from $\mathcal{Q} \times \mathcal{A} \times \mathcal{S}$ to \mathcal{Q} , according to which a voter’s choice of an alternative together with the observed signal jointly drive the transition to a new memory state. However, since a voter’s own choice does not convey any information to the voter about the state of the world, we simplify the class of automata under consideration to be ones that only transition to a new memory state based on the signals observed by the voter, and not based on the choices she makes.

resulting polarization of voters. In Section 4 we show how the electoral process in a representative democracy shapes the collective choice a_t^I .

3 Voter Polarization

Throughout this section, we assume that the alternative with greatest support in the electorate becomes the collective policy choice a^I , with ties broken randomly.

We define “*voter polarization*” as the phenomenon in which a positive mass of voters support extreme alternative a^L and a positive mass support extreme a^R . We present sufficient conditions under which voters polarize. Polarization is a matter of degree, in proportion to the size of the masses of voters opposing each other at each extreme. We define a degree of polarization that ranges from zero if one extreme has no support, to one if each extreme is supported by exactly half the population.

For each $\theta \in \{L, M, R\}$, let $\nu_t(a^\theta)$ denote the share of voters who support action a^θ in period t .

Definition 3.1. *Society’s degree of polarization in period t is $\Psi_t \equiv 4\nu_t(a^L)\nu_t(a^R)$.*

We say that voters are polarized in period t if $\Psi_t > 0$.

Given that ex-ante —before the first signal is revealed— all voters support a^M ; that they all agree on the best alternative for any state of the world; and that they share the same prior and observe a common sequence of signals that (asymptotically) reveals the true state of the world, it might seem that optimizing agents could not polarize for long.

Indeed, rational agents who process information perfectly, endowed with a common prior and observing a common sequence of signals will agree on their posterior over the state of the world, and once this posterior becomes sufficiently close to degenerate, given their near common-value preferences, they will also agree on which alternative to support. Whereas, we shall show that agents with limited memory capacity have divergent posteriors and polarize, despite their common priors and common sequence of signals.

We start by considering an environment in which voters *do not* polarize: the special case in which extreme signals (ℓ or r) reveal the state of the world (formally, extreme

$q_t \setminus s_t$	ℓ	m	r
q^L	$(1, 0, 0)$	$(1, 0, 0)$	$(1 - \alpha_L - \beta_L, \beta_L, \alpha_L)$
q^M	$(1, 0, 0)$	$(0, 1, 0)$	$(0, 0, 1)$
q^R	$(\alpha_R, \beta_R, 1 - \alpha_R - \beta_R)$	$(0, 0, 1)$	$(0, 0, 1)$

Table 2: Transition probabilities to (q^L, q^M, q^R) under automata in FA_3 .

signals are fully informative if ϵ in Table 1 is zero). In this case, the unconstrained optimal decision rule is straightforward: the posterior on M increases as long as all signals are m , and hence alternative a^M continues to be optimal at period t as long as $s_k = m$ for any $k \in \{1, \dots, t\}$. In contrast, a single ℓ -signal, if $\epsilon = 0$, drives the posterior of L to one and reveals that a^L is the correct alternative. A symmetric argument holds for signal r and action a^R as well.

Such a simple rule is accessible for a voter with very little memory capacity, using any stochastic finite state automaton (SFSA) in the following class.

Definition 3.2. A SFSA is in **class FA_3** if it contains three memory states $\{q^L, q^M, q^R\}$; its initial state is q^M ; its decision rule is $d(q^\theta) = a^\theta$ for each $\theta \in \Theta$; and for some $\alpha_\theta, \beta_\theta \in [0, 1]$ such that $\alpha_\theta + \beta_\theta \in [0, 1]$ for each $\theta \in \{L, R\}$, the transition rule probabilities are given by Table 2, where each cell (q_t, s_t) lists the probability of transitioning, respectively, to q^L , to q^M , and to q^R .

For each extreme state $\theta \in \{L, R\}$, parameters α_θ and β_θ represent, respectively, the probability of transitioning from q^θ to the opposite extreme (α_θ) and of transitioning to moderation (β_θ). For any $(\alpha_L, \beta_L, \alpha_R, \beta_R) \in [0, 1]^4$ satisfying $\alpha_\theta + \beta_\theta \in [0, 1]$ for each $\theta \in \{L, R\}$, let $FA_3(\alpha_L, \beta_L, \alpha_R, \beta_R) \in FA_3$ denote the specific automaton with vector of transition parameters $(\alpha_L, \beta_L, \alpha_R, \beta_R)$, as depicted in Figure 1.

Any automaton in class FA_3 stays in moderate memory state q^M choosing the moderate alternative a^M as long as all signals it observes are moderate (m). But as soon as it observes an extreme (ℓ or r) signal, it transitions to the corresponding extreme memory state (q^L or q^R), and it chooses the corresponding extreme alternative (a^L or a^R). Once it arrives at an extreme memory state, it stays there until it observes a contradictory extreme signal. If $\epsilon = 0$, the first extreme signal revealed the correct state of nature, and a contradictory extreme signal never emerges, so in this case, any




Figure 1: Optimal 3-state automaton $FA_3(\alpha_R, \beta_R, \alpha_L, \beta_L)$

automaton in class FA_3 executes the unconstrained optimal decision rule (Lemma 6.2 in Section 6).

Whereas, if ϵ is positive, a perfectly rational agent with unlimited memory capacity continues to update her belief. For instance, under state of the world M , although the agent would occasionally receive the strong signals ℓ or r indicating L or R , she would also receive many more m -signals and, in the long run, she would conclude the state of the world is M .

We reach our case of interest: $\epsilon > 0$ but small, so that extreme signals are rare and hence very (but not perfectly) informative when they arise. In this environment, any automaton in class FA_3 features some attractive qualities for voters who find memory capacity costly ($\kappa > 0$). First, any 3-state automata is cheap, as the voter only needs to keep track of which of the three memory states she is in. Second, any automaton in FA_3 follows the optimal decision rule (namely, to choose the moderate action) as long as all signals are moderate, which is likely to be for a long time if ϵ is small. Third, if the first extreme signal happens soon enough, an automaton in FA_3 again follows the unconstrained optimal rule in following the signal to the corresponding extreme memory state and choice of alternative. So far, so good.

Automata in FA_3 only make two kinds of mistakes, relative to the unconstrained optimum attainable only with unlimited memory capacity. First, they disregard the small evidence provided by moderate signals. Moderate signals are quite likely in every state, so one such signal does not shift a perfectly-computed posterior much. Automata

in FA_3 regard the very little information contained in a moderate signal as negligible, and do not budge in any way upon observing it. But even if one moderate signal does not mean much, an abundance of them does. So these automata err in not returning to the moderate memory state after observing a sufficiently long history of signals in which moderate signals are overwhelmingly preponderant. The key to the appeal of FA_3 automata is that for ϵ small, it takes a long time to accumulate a large number of moderate signals. So by not returning to moderation when they should, these automata depart from the unconstrained optimal decision rule only far into the future; and an impatient voter finds choices consigned to a distant future to be of little relevance, and not worth incurring a higher cost of memory capacity.

The second problem for any automation in this class is that, since it does not keep track of how many ℓ or r signals it has observed, the automaton finds itself at a bit of a quandary when it is at an extreme memory state (say q^R) and it observes the opposite extreme signal (say ℓ). Should it ignore the signal, or should it switch memory states? It turns out that in such situation, randomization is useful: the optimal automaton among those in class FA_3 features a vector of transition probabilities $(\alpha_L, \beta_L, \alpha_R, \beta_R)$ that depends on all parameters, including the agent's lean; if the cost κ of processing information is sufficiently low, and extreme signals are sufficiently informative (i.e. if ϵ is sufficiently close to zero), this automaton that is optimal among those in class FA_3 is in fact optimal among all automata (Lemma 6.3).

Our first main result establishes sufficient conditions such that voters with lean of high magnitude ($|b^i|$ large), once they get to the extreme memory state congruent with their lean, will not transition away from it, and will support the extreme alternative congruent with this lean, forever. For instance, an agent with a sufficiently high lean to the right (b^i sufficiently close to \bar{b}), upon first seeing a right signal r , will immediately transition to a state in which she supports alternative a^R , and will not be swayed away from this memory state and this support by any subsequent sequence of signals. Likewise, under these conditions, an agent with a sufficiently high lean to the left (b^i sufficiently close to $-\bar{b}$), will never be swayed away from supporting alternative a^L after seeing a single ℓ signal. It follows that after extreme signals ℓ and r have both been observed, these agents polarize.

Proposition 3.1. (Voter polarization) For any taste for compromise $c \in (0, 1)$ and any patience parameter $\delta \in (0, 1)$ that satisfy

$$c < 2 \left(\frac{1 - \delta}{1 - \delta + \delta\mu} \right)^2, \quad (6)$$

there exist \bar{b} sufficiently close to one and sufficiently small memory capacity costs (κ), such that if the probability ϵ of an incorrect extreme signal is sufficiently small, in all states of the world, with probability converging to one in t , and for any sequence $(s_k)_{k=t+1}^\infty$ of signals after period t , votes polarize in all periods $t' > t$ as follows: for some $b^* \in (0, \bar{b})$,

- i. agents with a sufficiently left lean (namely, $\{i \in I : b^i \in [-\bar{b}, -b^*]\}$) support a^L ; whereas,
- ii. agents with sufficiently right lean (namely, $\{i \in I : b^i \in (b^*, \bar{b}]\}$) support a^R .

The intuition for these sufficient conditions is as follows. With regard to societal parameters, the cost κ of processing information must be small enough for voters to be willing to distinguish between circumstances in which they would support different alternatives,⁶ and extreme signals must be informative enough (ϵ small enough) to sway voters when they observe one; these two conditions suffice for the optimal automaton to be one in class FA_3 . An additional condition suffices for polarization to arise: voters' taste for compromise c and/or patience δ must be low as made precise by Condition (6), which makes a return to moderation unappealing.⁷

So, a combination of credible extreme signals, low taste for compromise, and impatience, leads to voter polarization. Under these conditions, voters with lean of highest magnitude polarize, and settle on permanently supporting the alternative that is congruent with their preference lean, regardless of any information they receive after seeing just one signal that supports their extreme choice.

⁶If processing information is prohibitively costly, voters would ignore all signals and always choose the ex-ante preferred action.

⁷Condition (6) relating c , δ and μ can be equivalently restated as $\delta < \frac{1 - \sqrt{0.5c}}{1 - (1 - \mu)\sqrt{0.5c}}$ for any $c \in (0, 1)$. For higher δ 's, we can show that for any $c \in (0, 1)$, there exists $\bar{\delta}(c) < 1$ such that if $\delta > \bar{\delta}$, then for the range κ and ϵ under which the optimal SFSA belongs to the class FA_3 , all voters return to moderation with positive probability.

While the proofs of our results are long and cumbersome, their intuition should be clear. If under the moderate state of the world, extreme signals are sufficiently rare, then when voters see one such extreme signal, they perceive it as very informative, and they all agree to treat it as if it were correct, and to support the extreme policy that corresponds to this extreme signal. As long as subsequent signals are either additional realizations of the same extreme signal, or hardly informative moderate signals that all voters ignore, all voters continue to agree and to support the extreme alternative congruent with the extreme signals they've seen.

Disagreement arises only when voters first receive a contradictory extreme signal; that is, either the first extreme signal was ℓ and now they see an r , or the first extreme signal was r and now they see an ℓ . This is a surprise, and it generates greater uncertainty as to whether the state is L or R (or it could also be M). Given this uncertainty as to which of the two extreme policies is more likely to be correct, if the moderate compromise is insufficiently appealing (if c is low, or if voters are impatient), voters part ways according to their lean: those who lean sufficiently left choose the left alternative, and those who lean sufficiently right choose the right alternative, while with those with lean close to zero randomize between the two extremes, with greater probability of supporting the extreme congruent with their lean.

The option of returning to moderation involves an inter-temporal trade-off that only appeals to patient agents with lean near zero. Namely, because extreme signals are very informative (and because voters forget how many moderate signals they have also observed), conditional on knowing that they have observed two extreme signals, all agents discount the possibility that the state of Nature is moderate, and prefer an extreme action, or a lottery over the two extremes, better than moderation. Returning to moderation thus involves a loss of expected payoff derived from the choice in the current period. This loss is intolerable for impatient agents. Patient agents, on the other hand, value the downstream informational upside of returning to moderation: once the next extreme signal arrives, a voter at a moderate state of memory will follow this signal with renewed confidence that it likely reveals the true state of Nature. This informational benefit is more valuable to patient agents with lean close to zero, who returned to moderation with maximal uncertainty about which extreme is correct, and are thus more willing to wait it out for more information. Whereas, agents with an

extreme lean only hesitate between the two extreme actions if they are almost certain that the state of Nature is the one counter to their lean, so they have less to gain by returning to moderation to gather more information... so they do not.

The full proof proceeds thus. Lemma 6.1 establish optimality conditions on the automata that we use to model a voter’s choice, namely: at each memory state, the voter must take the action that the voter believes gives her the highest payoff, and after each signal the voter must transition to the memory state with the highest continuation payoff. However, the beliefs are not updated according to Bayes rule due to the coarseness of any finite automaton, so they deviate from the Bayesian benchmark. Lemma 6.2 shows that FA_3 is the class of optimal automata for the special case in which extreme signals are fully revealing, and Lemma 6.3 partially extends this result to show that the optimal automata belongs to FA_3 if extreme signals are very (though no longer perfectly) informative. Lemma 6.4 then pins down the optimal automaton for agents with high lean, and from this characterization we infer Proposition 3.1.

Since the distribution over types has full support, Proposition 3.1 implies that, with probability converging to one, the degree of polarization Ψ_t is strictly positive, and is proportional to the product of the mass of agents with leans in $[-\bar{b}, -b^*)$ and the mass of agents with leans in $(b^*, \bar{b}]$.

While the exact value of the degree of polarization depends on parameter values and on functional assumptions on the distribution of types, we show that if voters’ taste for compromise (c) or patience (δ) are low, the polarized voters are not just a small fringe on each side; on the contrary, a majority of the electorate radicalizes into one or the other of the extremes, and in the long run, moderation never attains plurality support.⁸

Define an “automaton profile” to be a measurable mapping $\phi : [-\bar{b}, \bar{b}] \rightarrow \mathcal{Q}$ such that for each $b^i \in [-\bar{b}, \bar{b}]$, voter i with type b^i follows automaton $\phi(b^i) \in \mathcal{Q}$. For any subset $\bar{\mathcal{Q}} \subset \mathcal{Q}$, we say that automaton profile ϕ is “in $\bar{\mathcal{Q}}$ ” if $\phi(b^i) \in \bar{\mathcal{Q}}$ for any $b^i \in [-\bar{b}, \bar{b}]$.

Definition 3.3. *An automaton profile ϕ in FA_3 is **symmetric** if for any $b \in [-\bar{b}, \bar{b}]$,*

⁸To compare the sizes of the mass of voters who support each alternative, we prove (in the Online Appendix) that these subsets of voters are measurable; that is, that there exists a countably additive function that assigns a non-negative value —a size— to each of these subsets.

according to automata $\phi(b)$ and $\phi(-b)$,

$$\alpha_R(b) = \alpha_L(-b), \alpha_L(b) = \alpha_R(-b), \beta_R(b) = \beta_L(-b), \beta_L(b) = \beta_R(-b). \quad (7)$$

By symmetry of the payoff functions and the information structure, there exists an optimal automaton that is symmetric. And if voters follow one such automaton, we can compare the sizes of the support across alternatives, as in the next proposition.

Proposition 3.2. (*Society's extremism*) *Assume the distribution of types is symmetric. For any \bar{b} sufficiently close to one, and for any $(c, \delta) \in (0, 1)^2$ such that*

$$c < \left(\frac{1 - \delta}{1 - \delta + \delta\mu} \right)^2, \quad (8)$$

there exist a range of memory capacity costs (κ) such that that if ϵ is small enough, in all states of the world, and for any symmetric optimal automaton profile that voters may follow, with probability converging to one in t , in any period $t' > t$, any alternative with greatest support in the electorate and thus the collective policy choice a^t are extreme.

In other words: if voters' taste for compromise or patience are low, the probability that moderation sustains a plurality of support among voters and is chosen vanishes.⁹ A partial intuition for this result is as follows: because voters, after observing the first extreme signal, only return to support moderation as a temporary state of indecision triggered by observing conflicting extreme signals, and they leave again as soon as they observe another extreme signal, voters supporting an extreme alternative accumulate over time, but those in support of moderation do not. The proof relies on the symmetry assumption on the distribution of types by noting that for any positive $b \in \mathbb{R}_{++}$, if a majority of voters with lean b return to moderation after observing an extreme signal that contradicts their current extreme state of memory, then after observing the same signal, an at least as large majority of voters with the opposite lean $-b$ ignore the signal and stay put (Lemma 6.6), so aggregating across all types, more voters stay in their current extreme than move to moderation, and moderation never gains a plurality of support.

⁹Like Condition (6), Condition (8) can be restated as $\delta < \frac{1 - \sqrt{c}}{1 - (1 - \mu)\sqrt{c}}$.

4 Parties' Extremism in an Indirect Democracy

We now close our theory of polarization and extremism by introducing a stylized model of political competition and policy implementation, in which in each period, two parties P^1 and P^2 announce policy platforms, voters vote, and the party that obtains most votes wins and implements its announced platform.

Players. We consider an electoral competition game played by the continuum of voters introduced in Section 2, and by two political parties P^1 and P^2 .

Let F_b denote the cumulative distribution function of voter types over $[-\bar{b}, \bar{b}]$. We had assumed that this distribution of types has full support and (for Proposition 3.2) that it is symmetric; we now also assume that it has a continuous density function f_b .

We treat each party P^1 and P^2 as a fully rational unitary actor that follows Bayes rule perfectly to update beliefs. Parties' sole strategic decisions are to choose policy platforms in each period, as a function of the observed past history of play in the game. For each $j \in \{1, 2\}$, let $a_t^{P^j} \in \mathcal{A}$ denote the platform chosen by Party P^j in period t , where for any $a \in \mathcal{A}$, committing to policy platform $a_t^{P^j} = a$ implies that if P^j wins the period t election, then the alternative a_t^I collectively chosen by society through the political process in period t is action a .

Timing and information. The timing in each period $t = 0, 1, 2, \dots$ is as follows:

1. Parties simultaneously commit to their individual platforms, $a_t^{P^j} \in \mathcal{A}$, $j = 1, 2$.
2. If $t = 1, 2, \dots$, the common signal $s_t \in \mathcal{S}$ about the state of the world is commonly observed (in period $t = 0$ agents observe no signal).
3. Each voter i chooses which policy alternative $a_t^i \in \mathcal{A}$ to support.
4. Each voter i chooses one of three voting alternatives: vote for P^1 , vote for P^2 , or abstain.
5. The party that obtains a greater share of votes, denoted by w_t , wins and implements its platform, with ties broken randomly. Hence $a_t^I = a_t^{w_t}$.

We assume that parties observe the pair of platforms, the common signal about the state of the world, the total mass of votes for each party, and the winning party in

each period. In contrast, in each period, each voter observes only the pair of platforms, the common signal about the state of the world, and her own private choices of which policy alternative to support and how to vote.

We model the voters' cost of increasing their memory capacity to process information as a choice of an automaton with costly memory states (as described in sections 2 and 3). Thus, the observed partial history up to period t enters voter i 's decision-making in period t only partially and indirectly through its effect on the memory state q_t^i of the finite automaton $\langle Q^i, q_0^i, \tau^i, d^i \rangle$ that voter i uses to guide her decisions. In particular, the finite automaton employed by voter i enters period t at a given memory state q_{t-1}^i ; it observes the public signal s_t released at the second step; and between the second and the third step in the timing of the strategic environment above, it transitions to a new memory state q_t^i according to τ^i (which can be stochastic), and it produces a recommended alternative to support, $d^i(q_t^i) \in \mathcal{A}$, a recommendation that voter i then follows in Step 3.

Parties' motivations. Parties are office-motivated. They obtain a period payoff of 1 if they win, and 0 otherwise, with time discount factor $\delta_P \in (0, 1)$ across periods. For ease of reading, we drop one parameter from the model by assuming the same time discount for voters and parties, i.e. $\delta_P = \delta$.

Voters' motivations. We assume that in each period, each individual voter solves two individual choice problems sequentially.

In each period $t = 0, 1, 2, \dots$, each voter i first chooses which alternative $a_t^i \in \mathcal{A}$ to support, trading off the desire to choose to support the right alternative, with the cost of processing information. The solution is as detailed in Lemma 6.4: Agent i following optimal automaton $\langle Q^i, q_0^i, \tau^i, d^i \rangle$ chooses to support action $a_t^i = d^i(q_t^i) \in \mathcal{A}$, where q_t^i is the memory state i 's automaton reaches in period t , and d^i is the automaton's decision rule.¹⁰

Once voter i has identified which alternative in $a_t^i \in \mathcal{A}$ she supports, voter i faces

¹⁰Notice that there is a restriction here. Namely, voters' automata only process information about the state of the world directly obtained through the sequence of signals $(s_t)_{t=1}^\infty$, but they do not recognize the potentially informative indirect signaling content of the parties' announced equilibrium platforms. This restriction is supported by evidence that agents overweigh their own experience and their own private signals, over the information that can be inferred from the behavior of other agents (Kogan 2008; Kaustia and Knüpfer 2008).

a second choice problem; namely, whether to vote for Party 1, to vote for Party 2, or to abstain. Let $v_t^i \in \{P^1, P^2, \emptyset\}$ denote the voting decision of agent i in period t , with $v_t^i = \emptyset$ representing abstention.

We assume that voters obtain an expressive payoff from voting, additive in each period to Expression (3), so that the expression of agent i 's overall utility in the democratic environment with an election in each period is

$$\lambda \underbrace{\sum_{t=0}^{\infty} (\delta)^t u(a_t^I, \theta, b^i)}_{\text{Instrumental utility}} + (1 - \lambda) \underbrace{\sum_{t=0}^{\infty} (\delta)^t u(a_t^i, \theta, b^i)}_{\text{Support expressive utility}} + \underbrace{\sum_{t=0}^{\infty} (\delta)^t u_v(v_t^i)}_{\text{Vote expressive u.}}, \quad (9)$$

where u_v is the period expressive utility derived from voting.

Since a voter's individual vote cannot have any influence over the election outcome, nor over future play (it cannot even be individually observed by other agents), the instrumental utility component drops out of the summation, and each voter's behavior is driven exclusively by the expressive payoffs.

Expression (9) decouples the act of supporting a policy alternative, from the act of voting. A citizen can support any policy alternative by advocating for it in conversation, in writing, or in civic activism, deriving an expressive payoff from any of these activities. A citizen can only vote by casting a ballot for one of the two competing parties, and it is this specific act that delivers the additional expressive utility term $u_v(v_t^i)$. The expressive utility from the choice of alternative to support is the one maximized, net of the cost of memory capacity, by the optimal automaton in the optimization problem (5). The expressive utility from voting is determined by the voter by her voting choice in each period.

We normalize the expressive utility from abstaining to zero, so $u_v(\emptyset) = 0$ for any agent i , for any period t . We assume that the expressive utility from voting depends on whether the vote aligns or not with the policy alternative that the voter has determined is best, according to her optimal automaton. Namely, if voter i votes sincerely for a party that commits to the alternative chosen by voter i 's optimal automaton, then voter i obtains a positive expressive payoff; whereas, if voter i votes for a party committed to an alternative that is not the one chosen by voter i 's optimal automaton, then voter i incurs a disutility from such vote.

Formally, there exists a parameter $\bar{u}_v > 0$ such that, for each $j \in \{1, 2\}$, for each voter i and for each period t , if $d^i(q_t^i)$ is the alternative chosen in period t by the optimal automaton chosen by voter i to solve her optimization problem (5), then

$$u_v(v_t^i) = \begin{cases} \bar{u}_v & \text{if } v_t^i = P^j \text{ and } a_t^{P^j} = d^i(q_t^i); \\ -\bar{u}_v & \text{if } v_t^i = P^j \text{ and } a_t^{P^j} \neq d^i(q_t^i); \text{ and} \\ 0 & \text{if } v_t^i = \emptyset. \end{cases} \quad (10)$$

We think of $d^i(q_t^i)$ as the optimal automaton's recommendation, so that voter i has agency over the choice of alternative $a_t^i \in \mathcal{A}$ given this recommendation. If voters follow their optimal automaton, they support the alternative chosen by their automaton (that is, if $a_t^i = d^i(q_t^i)$ for every voter i and period t). If so, agents derive expressive utility from voting for the alternative they support.

Whereas, if voter i deviates and chooses to support an alternative $a \neq d^i(q_t^i)$ that is not the one recommended by the voter's optimal automaton, then the voter enjoys a positive expressive payoff of voting if she votes for a party that commits to alternative $d^i(q_t^i)$, not for voting for a party that commits to a .¹¹

An agent's vote has no effect over the agent's current period instrumental payoff, no effect over the current period expressive payoff from the choice of an alternative to support, and no effect on future play. Therefore, the agent's voting problem in each period t reduces to the static optimization problem

$$\max_{v_t^i \in \{P^1, P^2, \emptyset\}} u_v(v_t^i). \quad (11)$$

Equilibrium concept. In our model, an equilibrium is a profile in which each voter chooses an optimal automaton and takes actions aligned with her automaton's recommendations, and in which parties' strategies are sequentially rational given voters' behavior and given beliefs updated by Bayes rule. We next formally define this concept.

At the beginning of the game, each voter i chooses a stochastic finite state au-

¹¹Alternative $d^i(q_t^i)$ is the one voter i thinks is best, while we can interpret alternative a_t^i as the one voter i claims to support in public. Vote v_t^i is cast in a secret ballot, where the positive expressive utility from voting comes from voting one's conscience sincerely for what one thinks best.

tomaton $\langle Q^i, q_0^i, \tau^i, d^i \rangle \in \mathcal{Q}$. At the time she chooses an action in period t , voter i has observed the following: the party platforms, the public signals, her automaton’s memory states and recommended actions, and her own chosen action and voting decision in every period up to $t - 1$, plus the party platforms, the public signal and her automaton’s memory state and recommended action in period t . Let voter i “support function” refer to a mapping from the set of all these observables for any period t , to the set of actions \mathcal{A} . Similarly, the set of all observables at the time voter i chooses her vote in period t includes all of the above, plus her own choice of which action to support in the current period. Let voter i ’s “voting function” refer to a mapping from the set of all such observables for any period t , to the set of probability distributions over voting options $\{P^1, P^2, \emptyset\}$.

For each party, a pure strategy is a standard object: a mapping from information sets to the set of actions \mathcal{A} , and a mixed strategy is a probability distribution over pure strategies.

Definition 4.1. *An **equilibrium** is a support function and a voting function for each voter, and a mixed strategy profile for candidates that satisfy the following.*

1. **(Voters optimize)** *There exists a symmetric automaton profile $\phi : [-\bar{b}, \bar{b}] \rightarrow \text{FA}_3$ such that for each voter i , $\phi(b^i) \equiv \langle Q^i, q_0^i, \tau^i, d^i \rangle \in \text{FA}_3$ solves optimization problem (5) and is such that for any period t ,*
 - (a) **(Sincere support)** *The action a_t^i that i chooses in period t is $d^i(q_t^i)$ for any realization of all the observables observed by i up to her choice of which action to support in period t .*
 - (b) **(Sincere voting)** *For any realization of observables observed by i up to her period t vote, any vote v_t^i that agent i casts with positive probability is a solution to*

$$\max_{v_t^i \in \{P^1, P^2, \emptyset\}} u_v(v_t^i).$$

2. **(Parties optimize)** *The parties’ strategy profile is a sequentially rational profile given the voters’ support function and voting function, and given that parties update beliefs according to Bayes rule.*

The intuition behind this formal notion of equilibrium is as follows. Voters want to learn which alternative is best, but it is costly for them to keep track of all the informative signals in detail, so they resort to a cost-efficient automata. An optimal automaton makes the best possible recommendation to maximize the expressive utility from supporting an action, based on the available signals and on the constraints induced by the cost of memory capacity. Equilibrium condition 1(a) requires voters to follow this recommendation: each voter supports the alternative that is recommended by an automaton that is optimal for her, given her type.

Equilibrium condition 1(b) requires each voter to vote optimally, given what she thinks best. Equilibrium condition 2 is that parties best respond at every information set, given standard Bayesian-updating beliefs. Equivalently, taking voters' optimal behavior as given, parties play a Weak Perfect Bayes Nash equilibrium (Mas Colell, Whinston and Green 1995) of the 2-player electoral competition game induced by voters' behavior.

We say an equilibrium is “neutral” if voters, when indifferent between the two parties, vote for each party with equal probability. Formally, in a neutral equilibrium $\operatorname{argmax}_{v_t^i \in \{P^1, P^2, \emptyset\}} u_v(v_t^i) = \{P^1, P^2\}$ implies $\Pr[v_t^i = P^1] = \Pr[v_t^i = P^2] = \frac{1}{2}$.

Results

A neutral equilibrium exists (Lemma A.2 in the Appendix). The intuition for existence is as follows. First, an optimal automaton profile exists for voters; for each voter i , following her automaton's recommendation for which alternative to choose, and then voting for any party that announces this alternative (and abstaining if none do) constitute a profile with sincere support and sincere voting. With regard to the two parties, the two-player, one-period game that takes voters optimal behavior as given and ignores the future is a finite game, and thus it has an equilibrium. A collection of such one-period equilibria of the two-player game, one for each possible one-period game that the parties might face, constitutes an equilibrium of the full game.

If voters are sufficiently impatient and averse to compromise so that the most supported policy is surely to be an extreme one (as shown in Proposition 3.2), then both parties choose extreme policy platforms, and the implemented policy is also extreme.

Proposition 4.1. (*Parties’ extremism*) *Given a symmetric distribution of types with full support over $[-\bar{b}, \bar{b}]$ and \bar{b} sufficiently close to one, and for any preference parameter $c \in (0, 1)$ and discount $\delta \in (0, 1)$ that satisfy*

$$c < \left(\frac{1 - \delta}{1 - \delta + \delta\mu} \right)^2,$$

there exist a range of costs of processing information (κ) such that if the probability ϵ of an incorrect extreme signal is sufficiently small, in all states of the world, in all neutral equilibria, with probability converging to one in t , in every period $t' \geq t$, both parties propose extreme policy platforms.

Society’s radicalization to the extremes (Proposition 3.2) drives candidates’ to the extremes as well. Expressive voters are uncompromising in the sense that they only vote for a candidate who embraces a policy position that they sincerely support. Candidates, thus, must react as in the quote attributed to French Minister Ledru-Rollin: “*there go the people, and I must follow them, for I am their leader.*”

In the Moderate state of the world (M), this candidate behavior constitutes pandering: over time, parties accumulate sufficient moderate signals to be arbitrarily close to certainty that the moderate policy a^M is best for every citizen, as it is in fact the case. If voters were fully rational and could process information costlessly (i.e. $\kappa = 0$), they too would learn that moderate policies are best, and at least one party would offer moderation and would get elected. As it is, voters who find processing information costly only react to extreme signals, and a majority end up supporting extreme alternatives. Candidates know better, but chasing votes, they too locate at the extremes, where the voters are found.¹²

If parties posterior belief at a given period t given the sequence of observed signals is that the state of Nature is most likely Moderate, they anticipate that the signal s_t will most likely be moderate (m), and that the distribution of support across alternatives in period t will be most likely the same as in the previous period. In particular, if one extreme alternative, say a^L , gathered the most support in period $t - 1$, the same

¹²The rationale for such pandering was perhaps best articulated by then Prime Minister of Luxembourg, Jean-Paul Juncker: “*We all know what to do, but we don’t know how to get re-elected once we have done it*”, quoted in The Economist in “The quest for prosperity”, March 15, 2007.

extreme alternative will likely gather the most support once again. Parties mutually best respond by both announcing this alternative as their platform.

Extension: Differentiated candidates. We consider an extension that generates platform divergence: policy-differentiated parties, such that in the eyes of the voter, each party has an exogenously given advantage on a different extreme policy, as in Krasa and Polborn (2010, 2012, 2014). Formally, instead of labeling parties neutrally as P^1 and P^2 , let them be meaningfully labeled P^L and P^R , such that for each $\theta \in \{L, R\}$, party P^θ announcing platform a^θ delivers additional utility $u^+ \in (0, \bar{u}_v)$, additive to the payoffs indicated in Expression (10) to any voter voting for the party. Under such specialization, parties cannot effectively compete on the extreme policy for which they have a comparative disadvantage. Regardless of which of the two extreme alternatives gathered greater support in the previous period, say a^L , as long as parties believe that a^L is most likely to gather the most support again, but that if $s_t = r$, then a^R would gather the most support, then parties would polarize as well, each of them announcing the extreme platform in which they have a comparative advantage.¹³

Differentiated parties thus polarize even after they learn that the moderate policy a^M is Pareto superior. In this environment, electoral majorities and implemented policies alternate between the two extremes according to the realization of the sequence of signals, and aggregate welfare suffers.

5 Discussion

In our theory of polarization, costs of processing information drive voters to polarize (Proposition 3.1). The moderate political center loses support, and most of the electorate radicalizes to the extremes (Proposition 3.2). Fully rational office-motivated parties pander to this radicalized electorate and their platforms become extreme as well (Proposition 4.1).

This theory of voter-driven polarization is consistent with the evidence from its motivating phenomenon: the political polarization in the United States around the

¹³This result holds as well if one of the two parties also has a small advantage on the moderate policy: in circumstances in which choosing moderation is dominated by extremism, it does not matter who would win if both parties chose moderation.

	L	M	R	Polarization Ψ_t
2004	28%	57%	15%	0.168
2011	27%	50%	23%	0.248
2014	32%	44%	24%	0.307
2017	42%	37%	21%	0.353
2019	42%	38%	20%	0.336

Table 3: Polarization of the US electorate, 2004-2019.

election and presidencies of Obama and Trump.

Longitudinal survey data collected by the Pew Center records a sharp polarization of the US electorate from 2004 to 2019, as shown in Table 3.¹⁴

Longitudinal data from the American National Election Studies (ANES) shows that respondents perceive the two main political parties to have polarized as well. As shown in Figure 2, the gap in the perceived position of the Democratic and Republican parties on a 0 to 10 ideological scale has more than doubled from 2008 to 2020; while until 2012 both parties were perceived to locate within the middle third of the scale, in 2020 the parties were perceived to be located one at each extreme third of the scale.

The standard wisdom is that the GOP moved to the right under pressure from its voters and activists: *“In this telling, GOP candidates and elected officials are merely reflecting what their most fervent popular supporters demand”* (Skocpol 2019). Skocpol and Williamson (2016, Ch. 5) explain that in this period, *“the Tea Party boosts the GOP and prods it rightward.”* Drawing from rich qualitative evidence and quantitative analysis, Skocpol and Tervor (2019) document how grassroots movements —aided by elite political actors— on the right and the left pulled not only the GOP but also the Democratic Party toward more polarized positions.

Our theory provides an explanation for the root cause of this process: the ideological polarization of the electorate.

Going forward, our theory predicts that, absent a shock outside our model, a state

¹⁴We use data from the 2004, 2011, 2014 and 2017 Political Typology surveys (December 1-16, 2004, n=2,000; February 22 to March 14, 2011, n=3,030; January 23 to March 16, 2014, n=10,013; and June 8-18, 2017, n=2,504); and the 2019 Political survey (September 5-16, 2019, n=2,004) to place respondents on a left/right scale from -10 to $+10$ based on their answers to ten invariant ideological questions, and we partition the scale into three equal-length intervals: L for values from -10 to -4 , M for values from -3 to $+3$, and R for values from $+4$ to $+10$.

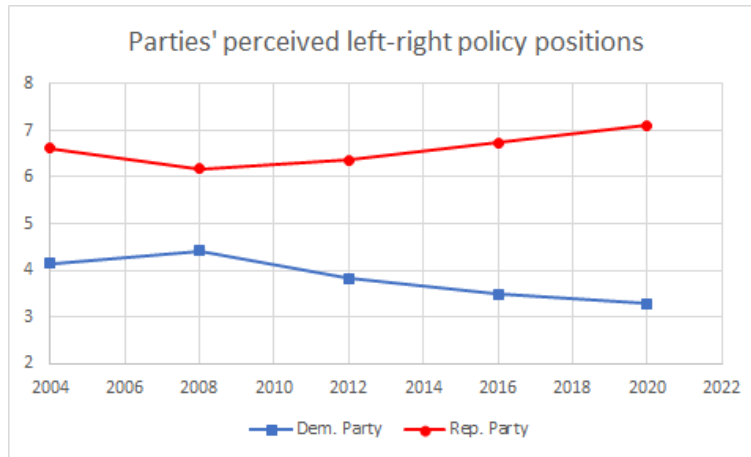


Figure 2: Party polarization in the US, 2004-2020.

of high support for extremist policies, and a polarized electorate, are irreversible: in an environment with frequent (and thus weak) moderate signals and ex-ante rare extreme signals, voters with limited capacity to process information will never learn enough to be convinced that the state is moderate, and will always be subject to the pull of extremism.

Whether polarization is indeed “permanent” in practice depends on the time scale under consideration. In the United States, polarization was permanently high for decades during the Gilded Age and the Progressive Era (1880s-1910s) (Putnam 2020), and it is once again high a little over a century later... but in between these two crests, during World War II and in its wake, there was a period of a broad national consensus and low polarization that lasted through the Eisenhower Presidency in the 1950s (Hofstadter 1964). Over long historical time-spans, polarization ebbs and flows. While our theory is dynamic with an infinite horizon, we interpret it as applicable to explain only part of this historical evolution: the rise of polarization over a relatively short period of time. The return to a national consensus responds to factors outside our model. Chief among such potential factors are external security threats (Desch 1995) that trigger a “rally round the flag” unifying effect (Groeling and Baum 2008).

One could, if desired, embed our theory into a more general account of the rise and fall of polarization over history: suppose that at the conclusion of each period, with some small probability, the entire political environment suffers a common-knowledge

shock that resets the state of the world, drawing it anew from its common prior distribution. Formally, this is but a small departure from our model: it suffices to interpret the shock as “ending” the game and starting a new one,¹⁵ and to reinterpret the discount factor δ as including both the time discount and the probability that the game restarts after each period. Voters’ behavior over the infinite sequence of such games is cyclical: agents abruptly return to a consensus on moderation at the start of each new game, and then in this new game they eventually polarize at the extremes, before returning to moderation at the next game-ending shock.

If we are correct in our prediction that —until the next unifying national crisis, threat or shock— polarization is permanent, social interventions that seek to nudge the electorate back to moderation (such as regulating online content, deplatforming extreme speakers, or seeking to break information bubbles by exposing audiences to both sides of an argument) are likely to have limited success. Rather, fostering norms of tolerance of dissent and view-point diversity, and enshrining civil discourse and the democratic process as the means to channel ideological disagreement, may better help to manage ongoing political polarization.

6 Proofs

The proofs are organized as follows. We first present a technical lemma, Lemma 6.1, which we subsequently use to prove the results in Section 3. We prove Proposition 3.1 by first proving two preliminary lemmas, Lemma 6.2 and Lemma 6.3, and then establishing Lemma 6.4, from which Proposition 3.1 follows as a corollary. Proposition 3.2 is proved by two technical lemmas, Lemmas 6.5 and 6.6. Finally, we prove Proposition 4.1.

We first characterize the optimal SFSA for a given number of memory states, $|Q| = K$, and will later take the cost κ of the memory states into account. This generalizes the result for two states of nature in Wilson (2014) to a our setting with three states of nature. The key to this characterization is a definition of surrogate “beliefs” at each memory state $q \in Q$. Given a state of nature θ and a memory state $q \in Q$, for an

¹⁵Just like the Algiers crisis of 1958 triggered the end of the Fourth Republic and the advent of the Fifth Republic in France.

agent i with type b^i , the expected payoff accumulated from q conditional on θ is then

$$\begin{aligned}
& \mathbf{1}_{q=q_0} u[d(q_0), \theta, b^i] + \delta \sum_{s_1 \in \mathcal{S}} \tau(q_0, s_1; q) \mu_{s_1}^\theta u[d(q), \theta, b^i] \\
& + (\delta)^2 \sum_{q_1 \in Q, s_1, s_2 \in \mathcal{S}} \tau(q^o, s_1; q_1) \mu_{s_1}^\theta \tau(q_1, s_2; q) \mu_{s_2}^\theta u[d(q), \theta, b^i] + \dots \\
& = \frac{1}{1 - \delta} f_q(\theta) u[d(q), \theta, b^i],
\end{aligned} \tag{12}$$

where the auxiliary function $f_q(\theta)$ is defined as follows:

$$f_q(\theta) \equiv \sum_{T=1}^{\infty} (1 - \delta)(\delta)^{T-1} \left[\sum_{(q_1, \dots, q_{T-1}), (s_1, \dots, s_{T-1}), q_T=q} \mathbf{1}_{q_1=q_0} \prod_{t=1}^{T-1} \mu_{s_t}^\theta \tau(q_t, s_t; q_{t+1}) \right]. \tag{13}$$

This then motivates the following definition of the surrogate beliefs:

$$p(q)(\theta) \equiv \frac{\mathbf{P}_0(\theta) f_q(\theta)}{\sum_{\theta'} \mathbf{P}_0(\theta') f_q(\theta')} \text{ and } p(q, s)(\theta) \equiv \frac{\mathbf{P}_0(\theta) f_q(\theta) \mu_s^\theta}{\sum_{\theta'} \mathbf{P}_0(\theta') f_q(\theta') \mu_s^{\theta'}}. \tag{14}$$

Indeed, at memory state q , optimality requires

$$d(q) \in \arg \max_{a \in \mathcal{A}} \sum_{\theta \in \Theta} p(q)(\theta) u(a, \theta, b^i). \tag{15}$$

Note that $p(q, s)$ is updated using Bayes rule from $p(q)$ and signal s . The following characterization result extends this optimality requirement to transitions. We use $V_q(\theta, b^i)$ to denote the continuation value for agent i with type b^i at memory state q conditional on the state of nature being θ . We say that two memory states are *equivalent* if they share the same transition rules to any other states or their equivalents, and have the same decision rule. With these definitions, we have the following necessary conditions for optimality.

Lemma 6.1. *Let $K \in \mathbb{N}$ and assume $\langle Q, q_0, \tau, d \rangle$ is a SFSA without equivalent states that is optimal among those of size $|Q| = K$. For each $q \in Q$, and for each type*

$b^i \in [-\bar{b}, \bar{b}]$, define

$$\Pi_q(b^i) \equiv \left\{ p \in \Delta(\Theta) : \sum_{\theta \in \Theta} p(\theta) V_q(\theta, b^i) \geq \sum_{\theta \in \Theta} p(\theta) V_{q'}(\theta, b^i) \text{ for all } q' \in Q \right\}. \quad (16)$$

Then, for each $q \in Q$ with $p(q)$ and $p(q, s)$ defined by the pair of expressions (14),

$$\tau(q, s; q') > 0 \Rightarrow p(q, s) \in \Pi_{q'}(b^i). \quad (17)$$

The proof of Lemma 6.1 follows similar steps as those in Wilson (2014); for completeness we give a proof in the Appendix.

We next prove the claim that if extreme signals are fully informative (i.e., $\epsilon = 0$ in Table 1), class FA_3 implements the unconstrained optimal rule and hence is uniquely optimal among all SFSA with 3 memory states.

Lemma 6.2. *Under the preferences and the information structure given by Table 1, for any $b^i \in [-\bar{b}, \bar{b}]$, $c \in (0, 1)$ and $\epsilon = 0$, and for any $(\alpha_L, \beta_L, \alpha_R, \beta_R) \in [0, 1]^4$ satisfying $\alpha_\theta + \beta_\theta \in [0, 1]$ for each $\theta \in \{L, R\}$, the unconstrained optimal rule can be implemented by automation $FA_3(\alpha_L, \beta_L, \alpha_R, \beta_R)$. Further, among automata with 3 memory states, only those in class FA_3 implement this unconstrained optimal rule.*

Proof. The optimal decision rule, denoted by D^* , is such that $D^*(\emptyset) = a^M$ as the prior has $\mathbf{P}_0(M) > 1/2$, and, for any $t \geq 1$, and for any sequence $(a^\nu)_{\nu=0}^{t-1}$,

$$D^*(s_1, \dots, s_t; a_0, \dots, a_{t-1}) = a^M \text{ if } s_\tau = m \text{ for all } \tau = 1, \dots, t; \quad (18)$$

$$D^*(s_1, \dots, s_t; a_0, \dots, a_{t-1}) = a^L \text{ if } s_\tau = \ell \text{ for some } \tau \text{ such that } s_{\tau'} = m \text{ for all } \tau' < \tau;$$

$$D^*(s_1, \dots, s_t; a_0, \dots, a_{t-1}) = a^R \text{ if } s_\tau = r \text{ for some } \tau \text{ such that } s_{\tau'} = m \text{ for all } \tau' < \tau.$$

Under the information structure in Table 1 with $\epsilon = 0$, it is impossible to have an r -signal followed by an ℓ -signal or vice versa. Thus, the transition probabilities $\alpha_L, \alpha_R, \beta_L, \beta_R$ do not matter, and $FA_3(\alpha_R, \beta_R, \alpha_L, \beta_L)$ implements D^* for any $\alpha_R, \beta_R, \alpha_L, \beta_L \in [0, 1]$. Moreover, any optimal SFSA with $|Q| = 3$ must take this form. \square

Now we take the cost of processing information, κ , into account, and the following lemma then shows that, for sufficiently small κ , and sufficiently small ϵ , the class FA_3

includes all the optimal SFSA.

Lemma 6.3. *There exist $\bar{\kappa} \in \mathbb{R}_{++}$ and a function $\bar{\epsilon} : (0, \bar{\kappa}) \rightarrow \mathbb{R}_{++}$ such that for any $b^i \in [-\bar{b}, \bar{b}]$, $c \in (0, 1)$, and for any $\kappa \in (0, \bar{\kappa})$ and for any $\epsilon \in (0, \bar{\epsilon}(\kappa))$, the optimal SFSA belongs to class FA_3 .*

Proof. For any $K \in \mathbb{N}$, let $\bar{V}_K(\epsilon, b^i, c)$ be the optimal payoff from K -memory-state finite automata under $\epsilon \geq 0$ for $b^i \in [-\bar{b}, \bar{b}]$ and $c \in [0, 1]$. Note that for any (b^i, c) , $\bar{V}_2(0, b^i, c) > \bar{V}_1(0, b^i, c)$, $\bar{V}_K(0, b^i, c) > \bar{V}_2(0, b^i, c)$ for all $K \geq 3$, and $\bar{V}_K(0, b^i, c) = \bar{V}_{K'}(0, b^i, c)$ for all $K \geq 3$ and all $K' \geq K$.

Define

$$\kappa(b^i, c) \equiv \min \left\{ \frac{\bar{V}_3(0, b^i, c) - \bar{V}_1(0, b^i, c)}{2}, \bar{V}_3(0, b^i, c) - \bar{V}_2(0, b^i, c) \right\},$$

and $\bar{\kappa} \equiv \inf_{b^i \in [-\bar{b}, \bar{b}], c \in (0, 1)} \kappa(b^i, c)$, and notice that $\kappa(b^i, c) > 0$ for any parameters in that range, with values bounded away from zero over there, so $\bar{\kappa} > 0$. Assume $\kappa \in (0, \bar{\kappa})$; then every agent prefers the optimal three-state automaton over any automata with fewer states. Given κ , by continuity of $\bar{V}_k(\epsilon, b^i, c)$ with respect to ϵ , there exists $\bar{\epsilon}(\kappa) \in \mathbb{R}_{++}$ sufficiently small that $\bar{V}_k(\bar{\epsilon}(\kappa), b^i, c) - \bar{V}_3(\bar{\epsilon}(\kappa), b^i, c) < k\kappa$ for all $k \leq \bar{V}_3(0, b^i, c)/\kappa$ and for any $b^i \in [-\bar{b}, \bar{b}]$, $c \in (0, 1)$. Then for any $\epsilon \in (0, \bar{\epsilon}(\kappa))$, the optimal automaton has three memory states.

Given that optimal SFSA has $K = 3$, now we prove that any optimal SFSA has the form of $FA_3(\alpha_L, \beta_L, \alpha_R, \beta_R)$. To do so, first we compute the value functions, whose detailed derivation can be found in the Appendix. Second, we rely on Lemma 6.1 to conclude that optimal automata belong to FA_3 .

To show the optimality of $FA_3(\alpha_R, \beta_R, \alpha_L, \beta_L)$, we divide the set of SFSA into two groups. Automata in the first group have transition probabilities close to those in $FA_3(\alpha_R, \beta_R, \alpha_L, \beta_L)$, while the second group consist of all others. We then show that $FA_3(\alpha_R, \beta_R, \alpha_L, \beta_L)$ with optimal α 's and β 's is the unique optimal SFSA within the first group, and outperforms those in the second group. The first claim is proved using Lemma 6.1, while the second follows from the uniqueness in Lemma 6.2 and continuity of the optimal value for ϵ close to zero.

We first define a distance between SFSA. For any two SFSA $\langle Q, q_0, \tau, d \rangle$ and

$\langle Q, q_0, \tau', d \rangle$, define the distance between them as $\max_{q \in Q} \|\tau(q, s) - \tau'(q, s)\|$, where $\|\cdot\|$ is the Euclidean distance over $\Delta(Q)$.

Now, let $\tau_{(\alpha, \beta)}(q, s)$ and d_3 denote the transition probabilities given $(q, s) \in Q \times \mathcal{S}$ and the decision rule of automaton $FA_3(\alpha_R, \beta_R, \alpha_L, \beta_L)$, and define

$$\mathcal{FA}(\rho) \equiv \{ \langle (q^L, q^M, q^R), q^M, \tau, d_3 \rangle : \|\tau(q, s) - \tau_{(\alpha, \beta)}(q, s)\| < \rho \text{ for all } (q, s) \neq (q_R, \ell), (q_L, r) \}. \quad (19)$$

That is, $\mathcal{FA}(\rho)$ consists of SFSA within distance of ρ to some SFSA in class FA_3 . Let $\mathcal{FA}^c(\rho)$ denote the set of all SFSA with $|Q| = 3$ not in $\mathcal{FA}(\rho)$.

Now, for any $\rho, \epsilon \in \mathbb{R}_{++}$, and any (b^i, c) , define

$$W(\rho, \epsilon, b^i, c) \equiv \max_{FA \in \mathcal{FA}^c(\rho)} V(FA, \epsilon, b^i, c),$$

where $V(FA, \epsilon, b^i, c)$ is the expected ex ante payoff under (b^i, c) from an arbitrary SFSA FA under ϵ . Notice that since $\mathcal{FA}^c(\rho)$ is compact, and $V(FA, \epsilon, b^i, c)$ is continuous in FA , the maximum exists and $W(\rho, \epsilon, b^i, c)$ is well defined.

For any (b^i, c) in the range and for any $\alpha_L, \beta_L, \alpha_R, \beta_R \in [0, 1]$,

$$W(\rho, 0, b^i, c) < V[FA_3(\alpha_R, \beta_R, \alpha_L, \beta_L), 0, b^i, c].$$

By continuity and the Theorem of the Maximum, for any (b^i, c) , there exists $\tilde{\epsilon}(b^i, c) \in (0, \bar{\epsilon}(\kappa)]$ (where $\bar{\epsilon}(\kappa)$ is as defined above), such that

$$W(\rho, \epsilon, b^i, c) < V[FA_3(\alpha_R, \beta_R, \alpha_L, \beta_L), 0, b^i, c]$$

for all $\epsilon \leq \tilde{\epsilon}(b^i, c)$. Further, $\inf_{b^i, c} \tilde{\epsilon}(b^i, c) > 0$, so there also exists a common $\tilde{\epsilon}$ such that

$$W(\rho, \epsilon, b^i, c) < V[FA_3(\alpha_R, \beta_R, \alpha_L, \beta_L), \epsilon, b^i, c]$$

for all $\epsilon \leq \tilde{\epsilon}$. Therefore, for sufficiently small ϵ , the optimal automaton in $\mathcal{FA}(\rho)$ (if there is one) is also strictly better than any automaton in $\mathcal{FA}^c(\rho)$, and thus if it exists, it is the optimal 3 state automaton.

To show that there exists a three-state automaton of the form $FA_3(\alpha_R, \beta_R, \alpha_L, \beta_L)$

that is optimal among those in $\mathcal{FA}(\rho)$, we appeal to Lemma 6.1 to show that optimal transitions follow those in $FA_3(\alpha_R, \beta_R, \alpha_L, \beta_L)$. Recall that the surrogate belief at q^M after seeing signal r for state R , $p(q^M, r)(R)$, according to Definition (14) is given by

$$\frac{\mathbf{P}_0(R)f_{q^M}(R)\mu_r^R}{\mathbf{P}_0(L)f_{q^M}(L)\mu_r^L + \mathbf{P}_0(M)f_{q^M}(M)\mu_r^M + \mathbf{P}_0(R)f_{q^M}(R)\mu_r^R},$$

which is arbitrarily close to one for sufficiently small ϵ ; see detailed computation of the beliefs $f_q(\theta)$'s in the Appendix. Similarly, if ϵ is sufficiently small, then $p(q^R, r)(R)$ and $p(q^R, m)(R)$ are arbitrarily close to one. Symmetrically, $p(q^M, l)(L)$, $p(q^L, l)(L)$ and $p(q^L, m)(L)$ are arbitrarily close to one. Note as well that $V_{q^R}(R) > V_{q^M}(R) > V_{q^L}(R)$ and $V_{q^L}(L) > V_{q^M}(L) > V_{q^R}(L)$ for ϵ sufficiently small (again, see detailed computations in the Appendix). These two observations, combined together and with Assumption (1), imply that if ϵ is sufficiently small, then under any automaton in FA_3 , for any (q, s) other than (q_R, ℓ) and (q_L, r) , if $\tau_{(\alpha, \beta)}(q, s) = q'$, then $p(q, s) \in \text{INT}(\Pi_{q'})$.

Since the continuation values and beliefs given by the pair of expressions (14) are continuous in both ϵ and in the transition probabilities, there exist $\rho_0 > 0$ and $\epsilon_0 > 0$ such that for any (q, s) other than (q_R, ℓ) and (q_L, r) , if $\tau_{(\alpha, \beta)}(q, s) = q'$, then $p(q, s) \in \text{INT}(\Pi_{q'})$ as well for all $\epsilon \leq \epsilon_0$ and for all SFSA in $\mathcal{FA}(\rho_0)$. Lemma 6.1 then implies that for all $\epsilon \leq \epsilon_0$, among SFSA in $\mathcal{FA}(\rho_0)$, any optimal automaton must be of the form $FA_3(\alpha_R, \beta_R, \alpha_L, \beta_L)$. Further, since the class of automata $FA_3(\alpha_R, \beta_R, \alpha_L, \beta_L)$ is compact, and utilities are continuous in transition probabilities, a solution to the voter's optimization problem (5) exists, and thus $FA_3(\alpha_R, \beta_R, \alpha_L, \beta_L)$ with optimal α 's and β 's is the optimal SFSA among those in $\mathcal{FA}(\rho_0)$. \square

The following lemma identifies sufficient conditions for an agent to transition to an absorbing extreme memory state such that she never transitions out of such state and it is key to the proof of Proposition 3.1. We use u_R to denote $1 + b^i$ and u_L for $1 - b^i$, $b^i \in [-\bar{b}, \bar{b}]$, and for each $\theta \in \{L, R\}$ we let $\gamma_\theta \equiv 1 - \alpha_\theta - \beta_\theta$ denote the probability that the SFSA stays at q_θ when receiving an extreme signal (ℓ or r) that is not equal to θ .

Lemma 6.4. *For any $\delta \in (0, 1)$ and any sufficiently small cost κ , there exist $\bar{\epsilon}(\kappa, \delta) > 0$ such that for all $\epsilon \leq \bar{\epsilon}$, and for any $(u_R, u_L) \in [1, 2]^2$ and $c \in (0, 1)$ any optimal SFSA takes the form of $FA_3(\alpha_L, \beta_L, \alpha_R, \beta_R)$. Moreover, we have the following characteriza-*

tion.

- If

$$u_R \geq \frac{\delta\mu(1-\delta+\delta\mu)u_L + (1-\delta+\mu\delta)^2c}{(1-\delta)^2} + \frac{2\delta\epsilon}{\mu(1-\delta)^3}, \quad (20)$$

then any optimal SFSA features $\beta_L = \beta_R = 0$.

- Further, if Inequality (20) holds and

$$\frac{(u_R + c)}{(u_L + c)} > \left(\frac{1 - \delta(1 - \mu + \epsilon)}{1 - \delta(1 - \epsilon)} \right)^2, \quad (21)$$

then optimal $\gamma_R = 1 - \alpha_R - \beta_R = 1$ and optimal $\alpha_L = 1$.

Proof. Lemma 6.3 shows that optimal SFSA takes the form of $FA_3(\alpha_L, \beta_L, \alpha_R, \beta_R)$. Now we show that under condition (20), the optimal SFSA features no transition to moderation, i.e. $\beta_L = \beta_R = 0$. Denote $\rho_0 \equiv \mathbf{P}_0(M)/\mathbf{P}_0(L)$. Let

$$G(\alpha_R, \beta_R, \alpha_L, \beta_L) = V_{qM}(L) + V_{qM}(R) + \rho_0 V_{qM}(M).$$

Then optimal α 's and β 's optimize G . For our purpose, it is easier to work with β 's and γ 's, taking $\alpha_L = 1 - \beta_L - \gamma_L$ and $\alpha_R = 1 - \beta_R - \gamma_R$. We derive the expressions of the value functions according to this substitution in the Appendix. First we claim that, for ϵ small, optimal $\beta_R = \beta_L = 0$ under Condition (20). To prove the claim, we show that, under Condition (20), for all $\beta_L, \beta_R, \gamma_L, \gamma_R$, $\frac{\partial G}{\partial \beta_L} < 0$. This follows from straightforward algebraic computations and comparisons of the conditions, with details in the Appendix. This implies that optimal $\beta_L = 0$. Similarly, we claim that $\frac{\partial G}{\partial \gamma_R} > \frac{\partial G}{\partial \beta_R}$ for all $\beta_L, \beta_R, \gamma_L, \gamma_R$. Again, this follows from Condition (20) with straightforward computations, with details in the Appendix. This implies that optimal $\beta_R = 0$.

Finally, we claim that, if both conditions (20) and (21) hold, then optimal $1 - \gamma_R = 0 = \gamma_L$. We show this by computing $\frac{\partial G}{\partial \gamma_R}$ with $\beta_L = \beta_R = 0$ and argue that it is positive for all γ_R and γ_L if (21) holds. Similarly, $\frac{\partial G}{\partial \gamma_L} < 0$ for all $\gamma_R, \gamma_L \in [0, 1]$ if Condition (21) holds. This implies that optimal $1 - \gamma_R = 0 = \gamma_L$. Again, details are in the Appendix. \square

We next use Lemma 6.4 to prove Proposition 3.1.

Proof of Proposition 3.1.

Proof. Define

$$b^{*1} \equiv \frac{\delta\mu(1 - \delta + \delta\mu) - (1 - \delta)^2 + (1 - \delta + \delta\mu)^2c + \frac{2\delta\epsilon}{\mu(1-\delta)}}{\delta\mu(1 - \delta + \delta\mu) + (1 - \delta)^2}. \quad (22)$$

Plugging $u_R = 1 + b^i$ and $u_L = 1 - b^i$, we find that Condition (20) holds if and only if $b^i \geq b^{*1}$. Now we show that, for ϵ sufficiently small, Condition (6) implies that $b^{*1} < 1$. To see this, $b^{*1} < 1$ if and only if

$$(1 - \delta + \delta\mu)^2c + \frac{2\delta\epsilon}{\mu(1 - \delta)} < 2(1 - \delta)^2,$$

which holds under Condition (6) if ϵ is small. Similarly, plugging in $u_R = 1 + b^i$ and $u_L = 1 - b^i$, we find that Condition (21) holds if and only if

$$b^i \geq b^{*2} \equiv \frac{[(1 - \delta + \delta(\mu - \epsilon))^2 - (1 - \delta(1 - \epsilon))^2](1 + c)}{(1 - \delta + \delta(\mu - \epsilon))^2 + (1 - \delta(1 - \epsilon))^2}.$$

Moreover, under Condition (6), $b^{*2} < 1$ because

$$c < \frac{2(1 - \delta)^2}{(1 - \delta + \delta\mu)^2} < \frac{2(1 - \delta + \delta\epsilon)^2}{(1 - \delta + \delta(\mu - \epsilon))^2 - (1 - \delta + \delta\epsilon)^2},$$

where the last inequality holds if ϵ is small.

Thus, for $b^i \geq b^* \equiv \max\{b^{*1}, b^{*2}\}$, by Lemma 6.4, $\alpha_L = 1$ and $\alpha_R = \beta_R = 0$. By a symmetric argument, up to relabeling of left and right, $b^i \leq -b^*$, $\alpha_R = 1$ and $\alpha_L = \beta_L = 0$. \square

The next lemma resolves some technical questions about measurability (with the proof in the Appendix). Given a measurable automaton profile ϕ , for each $b \in [-\bar{b}, \bar{b}]$, each $\theta \in \{L, M, R\}$, and each $t \in \mathbb{N}$, let $\nu_t^b(a^\theta)$ denote the fraction of voters of type b who support action a^θ at time t given that they follow automaton $\phi(b)$.

Lemma 6.5. *A measurable optimal automaton profile ϕ in FA_3 exists, and if voters*

follow such ϕ , then for any $\theta \in \{L, M, R\}$, the size of support for alternative a^θ is

$$\nu_t(a^\theta) = \int_{b \in (-\bar{b}, \bar{b})} \nu_t^b(a^\theta) f(b) db. \quad (23)$$

We next give a result from which Proposition 3.2 follows as a corollary. We relegate the proof of Lemma 6.6 to the Appendix.

Lemma 6.6. *Assume voters follow a symmetric measurable automaton profile ϕ in FA_3 . For any path of signal realizations and any period t such that at least one extreme signal has realized in periods 1 through $t - 1$,*

$$\nu_t(a^L) > \nu_t(a^M) \text{ if } s_t = r; \quad \nu_t(a^R) > \nu_t(a^M) \text{ if } s_t = \ell. \quad (24)$$

Now, Proposition 3.2 follows as a corollary of Lemma 6.6: note that since all voters employ a SFSA of the form $FA_3(\alpha_R, \beta_R, \alpha_L, \beta_L)$, voters do not change their actions when seeing an m -signal. Thus, by the inequalities (24), once one extreme signal is realized, moderation never gains plurality. We next prove Proposition 4.1.

Proof of Proposition 4.1.

Proof. First, an equilibrium as defined in Definition 4.1 exists (Lemma A.2 in the Appendix). Because parties' chosen platforms $a_t^{P_1}$ and $a_t^{P_2}$ in period t have no effect over the sequence of signals in periods after t , nor on voters' voting behavior in any subsequent period, taking the voters' behavior as given, and normalizing each party's period payoff from winning with probability 0.5 to zero, parties P^1 and P^2 face an infinite sequence of two-player, symmetric zero-sum one-period games. In any equilibrium of such an infinite horizon two-player zero-sum game, each party must obtain a zero period-payoff in each period, obtained by playing an equilibrium of the one-period game. We solve for the equilibrium of this two-player symmetric zero-sum one-period game.

Let $T \in \mathbb{N}$ denote the first period in which $s_t \in \{\ell, r\}$ (if the entire infinite sequence of public signals consists of m signals, then let $T = \infty$, but this occurs with probability zero). For any $t \in \mathbb{N}$, let $\pi(s_1, s_2, \dots, s_{t-1})$ denote the posterior about the state of Nature given the sequence of signals $(s_1, s_2, \dots, s_{t-1})$, according to Bayes rule. Consider

$a_t^{P_1} \backslash a_t^{P_2}$	a^L	a^M	a^R
a^L	$(\frac{1}{2}, \frac{1}{2})$	$(1, 0)$	$(z_1, 1 - z_1)$
a^M	$(0, 1)$	$(\frac{1}{2}, \frac{1}{2})$	$(z_2, 1 - z_2)$
a^R	$(1 - z_1, z_1)$	$(1 - z_2, z_2)$	$(\frac{1}{2}, \frac{1}{2})$

Table 4: Period-t game probability of winning matrix.

any period $t > T$ such that (without loss of generality),

$$\nu_{t-1}(a^L) \geq \nu_{t-1}(a^R). \quad (25)$$

Consider two cases.

Case 1. The most recent extreme signal is an r signal. By Lemma 6.6, $\nu_{t-1}(a^M) < \nu_{t-1}(a^L)$. We consider three subcases. (a) If $s_t = \ell$, then $\nu_t(a^M) < \nu_t(a^R) \leq \nu_{t-1}(a^R) \leq \nu_{t-1}(a^L) \leq \nu_t(a^L)$, where the first inequality is by Lemma 6.6, the second because $s_t = \ell$, the third by Assumption (25), and the fourth because $s_t = \ell$. So $\nu_t(a^M) < \nu_t(a^L)$. (b) If $s_t = m$, then $\nu_t(a^M) = \nu_{t-1}(a^M)$, $\nu_t(a^L) = \nu_{t-1}(a^L)$, where the last inequality follows from the two equalities and $\nu_{t-1}(a^M) < \nu_{t-1}(a^L)$. So $\nu_t(a^M) < \nu_t(a^L)$. (c) If $s_t = r$, then $\nu_t(a^M) < \nu_t(a^L)$ follows directly from Lemma 6.6. Together, we find that for any signal realization s_t , $\nu_t(a^M) < \nu_t(a^L)$.

We next construct the partial matrix of parties' probabilities of winning (Table 4) in the two-player one-period electoral competition game in period t , for some $z_1 \in (0, 1]$ and $z_2 \in [0, 1)$ that are left undetermined. Note that along the diagonal probabilities of winning are equal because the equilibrium is neutral. Since, subject to observing signal $s_t = \ell$, $\nu_t(a^M) < \nu_t(a^R) \leq \nu_t(a^L)$, while following any signal $\nu_t(a^M) < \nu_t(a^L)$, and since according to posterior π , the probability that $s_t = \ell$ is strictly positive, it follows that $z_1 > z_2$, which implies that a^M is strictly dominated and in equilibrium parties choose only extreme platforms.

Case 2. The most recent extreme signal is an l -signal. Then, by Lemma 6.6, $\nu_{t-1}(a^M) < \nu_{t-1}(a^R)$. We have three subcases. (a) If $s_t = \ell$, then $\nu_t(a^M) < \nu_t(a^R) \leq \nu_{t-1}(a^R) \leq \nu_{t-1}(a^L) \leq \nu_t(a^L)$, where the first inequality is by Lemma 6.6 the second because $s_t = \ell$, the third by Assumption (25), and the fourth because $s_t = \ell$. So $\nu_t(a^M) < \nu_t(a^L)$. (b) If $s_t = m$, then $\nu_t(a^M) = \nu_{t-1}(a^M) < \nu_{t-1}(a^R) \leq \nu_{t-1}(a^L) = \nu_t(a^L)$, so $\nu_t(a^M) < \nu_t(a^L)$.

(c) If $s_t = r$, then $\nu_t(a^M) < \nu_t(a^L)$ follows from Lemma 6.6.

Together, we find that for any signal realization s_t , $\nu_t(a^M) < \nu_t(a^L)$. The resulting matrix of probabilities of winning is again Table 4, as in Case 1, and the argument thus proceeds identically as in Case 1 to conclude that a^M is strictly dominated and thus both parties choose extreme platforms. \square

References

- [1] Abramowitz, A. I., and K. L. Saunders. “Is polarization a myth?” *Journal of Politics* 70.2 (2008): 542-555.
- [2] Axelrod, R., J. J. Daymude, and S. Forrest. “Preventing extreme polarization of political attitudes.” *Proceedings of the National Academy of Sciences* 118.50 (2021).
- [3] Baliga, S., E. Hanany, and P. Klibanoff. “Polarization and ambiguity.” *American Economic Review* 103.7 (2013): 3071-83.
- [4] Banovetz, J., and R. Oprea. ”Complexity and procedural choice.” *American Economic Journal: Microeconomics* 15.2 (2023): 384-413.
- [5] Benoît, J., and J. Dubra. “Apparent bias: What does attitude polarization show?” *International Economic Review* 60.4 (2019): 1675-1703.
- [6] Bol, D., K. Matakos, O. Troumpounis, and D. Xefteris. “Electoral rules, strategic entry and polarization.” *Journal of Public Economics* 178 (2019): 104065.
- [7] Bowen, R., D. Dmitriev, and S. Galperti. “Learning from shared news: when abundant information leads to belief polarization.” *Quarterly Journal of Economics*, 138.2 (2023): 955-1000.
- [8] Brennan, G., and A. Hamlin. “Expressive voting and electoral equilibrium.” *Public Choice* 95 (1998): 149–175.
- [9] Brennan, G., and L. Lomasky. *Democracy and Decision* (1993). Cambridge: Cambridge University Press.
- [10] Buisseret, P., and R. Van Weelden. “Polarization, valence, and policy competition.” *American Economic Review: Insights* 4.3 (2022): 341-52.
- [11] Calvert, R. L. “Robustness of the multidimensional voting model: Candidate motivations, uncertainty, and convergence.” *American Journal of Political Science* (1985): 69-95.

- [12] Che, Y.-K., and K. Mierendorff. “Optimal dynamic allocation of attention.” *American Economic Review* 109.8 (2019): 2993-3029.
- [13] Coate, S., and M. Conlin. “A group rule-utilitarian approach to voter turnout: Theory and evidence.” *American Economic Review* 94.5 (2004): 1476-1504.
- [14] Cover, T. M., and M. E. Hellman. “The two-armed-bandit problem with time-invariant finite memory.” *IEEE Transactions on Information Theory* 16.2 (1970): 185-195.
- [15] Davis, O. A., M. J. Hinich, and P. Ordeshook. “An expository development of a mathematical model of the electoral process.” *American Political Science Review* 64.2 (1970): 426-448.
- [16] Desch, M. C. “War and strong states, peace and weak states?” *International Organization* 50.2 (1996): 237-268.
- [17] Downs, A. “An economic theory of political action in a democracy ” *Journal of Political Economy* 65.2 (1957): 135-150.
- [18] Ely, J. “Ruth, Anthony and Clarence. ” Working paper, May 2023.
- [19] Enke, B., M. Polborn, and A. Wu. “Values as Luxury Goods and Political Polarization.” NBER No. w30001 (2022).
- [20] Feddersen, T., and A. Sandroni. “A theory of participation in elections.” *American Economic Review* 96.4 (2006): 1271-1282.
- [21] Fryer Jr, R. G., P. Harms, and M. O. Jackson. “Updating beliefs when evidence is open to interpretation: Implications for bias and polarization.” *Journal of the European Economic Association* 17.5 (2019): 1470-1501.
- [22] Gentzkow, M. “Polarization in 2016.” Toulouse Network for Information Technology white paper (2016).
- [23] Gordon, S. C., and D. Landa. “Common problems (or, what’s missing from the conventional wisdom on polarization and gridlock).” *Journal of Politics* 79.4 (2017): 1433-1437.
- [24] Glazer, A. “A new theory of voting: Why vote when millions of others do.” *Theory and Decision* 22.3 (1987): 257–270.
- [25] Groeling, T., and M. A. Baum. “Crossing the water’s edge: Elite rhetoric, media coverage, and the rally-round-the-flag phenomenon.” *Journal of Politics* 70.4 (2008): 1065-1085.
- [26] Hamlin, A., and C. Jennings. “Expressive voting.” *The Oxford Handbook of Public Choice* 1 (2018): 333.

- [27] Hofstadter, R. “A long view: Goldwater in history.” *New York Review of Books*, October 6th, 1964 Issue.
- [28] Kaustia, M., and S. Knüpfer. “Do investors overweight personal experience? Evidence from IPO subscriptions.” *The Journal of Finance* 63 (2008): 2679-2702.
- [29] Kogan, S. “Distinguishing the effect of overconfidence from rational best-response on information aggregation.” *The Review of Financial Studies* 22.5 (2009): 1889-1914.
- [30] Krasa, S., and M. K. Polborn. “Competition between specialized candidates.” *American Political Science Review* 104.4 (2010): 745-65.
- [31] Krasa, S., and M. K. Polborn. “Political competition between differentiated candidates.” *Games and Economic Behavior* 76.1 (2012): 249-71.
- [32] Krasa, S., and M. K. Polborn. “Social ideology and taxes in a differentiated candidates framework.” *American Economic Review* 104.1 (2014): 308–22.
- [33] Ledyard, J. O. “The pure theory of large two-candidate elections.” *Public Choice* 44.1 (1984): 7-41.
- [34] Levy, G., and R. Razin. “Short-term political memory and the inevitability of polarisation.” CEPR DP16737 (2021).
- [35] Maćkowiak, B., F. Matějka, and M. Wiederholt. “Rational inattention: A review.” *Journal of Economic Literature*, forthcoming (2021).
- [36] Matějka, F., and G. Tabellini. “Electoral competition with rationally inattentive voters.” *Journal of the European Economic Association* 19.3 (2021): 1899-1935.
- [37] McCarty, N., K. T. Poole, and H. Rosenthal. *Polarized America: The Dance of Ideology and Unequal Riches*. MIT Press, 2016.
- [38] McMurray, J. “Polarization and pandering in common-interest elections.” *Games and Economic Behavior* 133 (2022): 150-161.
- [39] Nimark, K. P., and S. Sundaresan. “Inattention and belief polarization.” *Journal of Economic Theory* 180 (2019): 203-228.
- [40] Nunnari, S., and J. Zapal. “A model of focusing in political choice.” CEPR DP 12407 (2020).
- [41] Oprea, R. “What makes a rule complex?” *American Economic Review* 110.12 (2020): 3913-51.
- [42] Perego, J., and S. Yuksel. “Media competition and social disagreement.” *Econometrica* 90.1 (2022): 223-265.

- [43] Prato, C., and S. Wolton. “The voters’ curses: why we need Goldilocks voters.” *American Journal of Political Science* 60.3 (2016): 726-737.
- [44] Putnam, R. D. *The Upswing: How America Came Together a Century Ago and How We Can Do It Again*. Simon and Schuster, 2020.
- [45] Riker, W. H., and P. C. Ordeshook. “A theory of the calculus of voting.” *American Political Science Review* 62.1 (1968): 25-42.
- [46] Schuessler, A. A. “Expressive voting.” *Rationality and Society* 12.1 (2000): 87-119.
- [47] Skocpol, T. “The elite and popular roots of contemporary Republican extremism.” In *Upending American politics*, Skocpol, T., and C. Tervo, eds (2019). Ch. 1, 3-28.
- [48] Skocpol, T., and C. Tervo, eds. *Upending American politics*. Oxford University Press, 2019.
- [49] Skocpol, T., and V. Williamson. *The Tea Party and the remaking of Republican conservatism*. Oxford University Press, 2016.
- [50] Serra, G. “Polarization of what? A model of elections with endogenous valence.” *Journal of Politics* 72.2 (2010): 426-437.
- [51] Tolvanen, J., J. Tremewan, and A. K. Wagner. “Ambiguous Platforms and Correlated Preferences: Experimental Evidence.” *American Political Science Review* (2021): 1-17.
- [52] Tacitus, C. *Tacitus: Agricola*. Cambridge University Press, 2014 [94AD].
- [53] Wilson, A. “Bounded memory and biases in information processing.” *Econometrica* 82.6 (2014): 2257-2294.
- [54] Wittman, D. “Candidate motivation: A synthesis of alternative theories.” *American Political Science Review* 77.1 (1983): 142-157.
- [55] Yuksel, S. “Specialized learning and political polarization.” *International Economic Review* 63.1 (2022): 457-474.