

Political Social Learning: Short-Term Memory and Cycles of Polarization[†]

By GILAT LEVY AND RONNY RAZIN*

This paper investigates the effect of voters' short-term memory on political outcomes by considering politics as a collective learning process. We find that short-term memory may lead to cycles of polarization and consensus across parties' platforms. Following periods of party consensus, short-term memory implies that there is little variation in voters' data and therefore limited information about the true state of the world. This in turn allows parties to further their own interests and hence polarize by offering different policies. In contrast, periods of polarization and turnover involve sufficient variation in the data that allows voters to be confident about what the correct policy is, forcing both parties to offer this policy. (JEL D72, D83, E61)

The rise of polarization in democratic societies in the last few decades has garnered significant attention in recent academic literature.¹ However, polarization is not a new phenomenon. By looking at longer time periods, a cyclical pattern of consensus and polarization can be observed. For instance, in the United States, policy positions of senators and congress members show that polarization was high at the beginning of the twentieth century, decreased in the 1930s, remained low until the late 1970s, and has been increasing since (see Figure 1).

Similar cyclical patterns can be seen in political parties' stated ideologies by examining their manifestos over time. The Manifesto Project decodes policy dimensions into a unidimensional score and tracks changes in these scores over time (Manifesto Project Dataset 2022). The manifestos of the two major parties in the United States clearly oscillate between polarization and relative consensus on economic issues. For example, Figure 2 plots the positions of the Democratic and Republican parties on market regulation, demonstrating periods of little difference in positions from the 1980s to the early 2000s, with relative polarization before and after these periods.

We investigate the question of whether cycles of polarization and consensus in party platforms are an inherent feature of political systems. By studying a dynamic

*Levy: LSE (email: g.levy1@lse.ac.uk); Razin: LSE (email: r.razin@lse.ac.uk). Pietro Ortoleva was the coeditor for this article. We thank the editor and four anonymous referees for their very helpful advice. For useful comments we also thank Steve Callander, Andrew Little, and many seminar participants and discussants. Roberto Maura-Rivero, Raffaele Blasone, and Enrico Turri provided helpful research assistance.

[†]Go to <https://doi.org/10.1257/aer.20220226> to visit the article page for additional materials and author disclosure statement(s).

¹See Barber and McCarty (2013) and McCarty (2019).



FIGURE 1. HISTORICAL POLARIZATION IN THE US SENATE AND CONGRESS

Source: Barber and McCarthy (2013), Figure 2.1




FIGURE 2. US DEMOCRATIC (THE HIGHER CURVE) AND REPUBLICAN PARTIES' REGULATION POLICIES

Source: The Manifesto Project Dataset (version 2024a)

model in which voters learn from history about the optimal policy, we demonstrate that such cycles can emerge when voters have limited or short-term memory. Our findings suggest that short-term memory can imply a systemic cyclical force in politics.

The examination of how short-term memory affects the evolution of policies in society is important for both positive and normative reasons. Previous research has lead to a broad consensus indicating that voters in presidential and congressional elections in the United States tend to “myopically ignore any information beyond

the recent past” (Peltzman 1990), a conclusion that is shared by other studies (Achen and Bartels 2008; Bartels 2008; Gerber and Green 1998; Lenz 2010; Healy and Lenz 2014; Angelucci and Prat 2023). Short-termism can then hinder learning from past mistakes. One example of repeated mistakes is the cyclical pattern of over-regulation and under-regulation in financial markets. Rajan (2009) notes that “once memories of the current crisis fade and the ideological cycle turns, the political pressure to soften capital requirements or their enforcement will be enormous.”

Motivated by the above empirical research, we incorporate the assumption of short-term memory into a dynamic political framework in which voters learn about the true data generating process that influences observable outcomes. In our model, there are two ideologically motivated parties, each with its own political interests, such as a preference for high levels of regulation or deregulation. Voters observe historical experiences and compare the expected utility they will gain from the policies that parties propose. To analyze the electoral competition between these parties, we use a probabilistic voting model and so parties’ choice of policies takes into account their partisan interests as well as voters’ uncertain preferences. In this model beliefs and policies evolve together over time.

Our model embeds a classic mechanism from the political economy literature, where party positions can diverge when parties are policy motivated and uncertain about voter preferences (as in Calvert 1985). The extent of this polarization depends in our model on how much voters are able to differentiate between the policies and understand which is better. When their knowledge is relatively weak their vote is determined by other stochastic factors such as party or candidate attachment and this allows parties to polarize. When voters’ knowledge allows them to understand well which policy is optimal, their voting behavior will be less noisy and parties will be disciplined to offer this policy.

What is novel in our paper is that we study this electoral competition between parties in a dynamic model in which both parties’ policies as well as voters’ knowledge are endogenous. We show how this can lead to cycles of party polarization and consensus. To see these dynamics, assume first that voters’ memory includes only periods in which parties were polarized and so offered different policies. Voters’ historical data will then include sufficient variation in policies (resulting from political turnover). This variation provides voters with sufficient information about the effectiveness of different policy options, enabling them to identify the optimal policy. Parties will then be incentivized to offer this policy and will therefore veer towards more consensus. However, after extended periods in which parties are in consensus, voters’ memories will consist only of experiences with the same policy. This lack of variation in the observed data will imply that while voters will understand well the benefits of this policy, they might not know how good or bad they have it as compared with an alternative. This weak ability of voters to distinguish among the alternatives will enable parties to credibly offer policies more in line with their partisan inclination, resulting in party polarization. Ultimately, it is the phase of consensus that can lead to polarization, and the phase of polarization that can eventually give way to consensus.

We identify a simple sufficient condition that gives rise to such dynamics. Specifically, cycles emerge when voters are able to differentiate the expected utility of different policies more effectively when they have access to an informative

history with policy variation, as compared to a less informative history. For any given prior distribution over the state of the world, there always exist states of the world for which this is satisfied. Additionally, we show how for some prior distributions (featuring “scale-free” learning) this condition is satisfied for almost all states of the world.

In our model, transitions between the phases of the cycle result from the effects of both endogenous policies and exogenous shocks on voter learning. Absent exogenous shocks, voters can learn relatively quickly which policy is optimal, and the cyclical dynamics are such that phases of policy consensus are relatively long compared to phases of party polarization. This highlights the role of endogenous policies as a systemic force driving the emergence of cycles in political systems. Voters’ learning can also be affected by exogenous shocks, such as natural disasters, wars, or pandemics, and our analysis illustrates how such shocks together with endogenous policies combine to affect the cycle dynamics (see Section IIB).

The implication of the forgetfulness of voters is that society will make repeated mistakes; short-term memory is “costly” as voters’ knowledge becomes limited ever so often, and so a polity that implements the correct policy in a consensus phase will necessarily desert it after a while. However, in some environments this can actually be beneficial. To see this, note that another way to view polarization phases is that once in a while, the polity inadvertently “experiments” with different policies. If the true state of the world changes sometimes, such experimentation allows voters to detect it. We show in Section III that this feature allows polities with short-term memory to sometimes outperform polities with unbounded memory.

In Section IV we extend the model to consider an environment in which different voters are exposed to different histories. Such heterogeneity among voters can arise from echo chambers or potentially from cohort effects where voters are affected by different experiences from their youth. We show how such different memories affect the nature of cycles. In Section V we consider other extensions, Section VI discusses the related literature, and all proofs are in a Supplemental Appendix.

I. The Model and Preliminary Results

A polity is considering a choice between two policies l and r . We first describe the economic environment, which is a simple mapping between policies and outcomes. Specifically, let the observable economic outcome y_t at period t be

$$(1) \quad y_t = \begin{cases} \beta_l^* + \varepsilon_t, & \text{if policy } l \text{ is chosen} \\ \beta_r^* + \varepsilon_t, & \text{if policy } r \text{ is chosen} \end{cases},$$

where ε_t is i.i.d across time and Normally distributed with zero mean and variance σ^2 . Our model and analysis are simplified by the assumption that the set of policies is discrete. Our results about cycles can be generalized to the case of continuous policies (see the discussion in Section VC).

Voters understand how the data generating process depends on parameters $\beta = (\beta_l, \beta_r)$, but do not know the true value of these parameters, $\beta^* = (\beta_l^*, \beta_r^*)$. They are endowed with a continuous and symmetric prior $G(\beta_l, \beta_r)$ on some compact set $B \in \mathbb{R}^2$, which determines how β^* is generated. In the main part of the

analysis we consider a fixed β^* . We discuss the possibility of changing states later on.

The outcome y_t is a common element in the voters' preferences, and so everything else equal, at any period t , all voters prefer policy l if given their information at period t , Ω_t , $E[\beta_l|\Omega_t] > E[\beta_r|\Omega_t]$. This will feed into their voting behavior, which we will describe below.

A. Political Parties and Electoral Competition

There are two parties, each identified with a special interest on a different policy. Party L prefers policy l and party R prefers policy r . The utilities of party L and R from policy $p \in \{l, r\}$, $U^R(p)$, and $U^L(p)$, satisfy:

$$(2) \quad U^R(r) = U^L(l) = 1, \quad U^R(l) = U^L(r) = 0.$$

In addition, parties also enjoy office rents, denoted by $\alpha > 0$. Thus, given an election at period t , and an implemented policy p_t , party J 's utility, for $J \in \{L, R\}$, is $U^J(p_t) + I_t^J \alpha$, where $I_t^J = 1$ if party J won the election and 0 otherwise. Our results also hold for other types of policy motives, e.g., if parties have a legacy motive and so enjoy a higher utility from their ideal policy if they are also the party that implements this policy.

In an election, at any period t , each party offers a policy $p_t^J \in \{l, r\}$. We say that parties polarize when $p_t^R \neq p_t^L$; naturally, this will imply that $p_t^R = r$ and $p_t^L = l$. When $p_t^R = p_t^L$, we say that parties are in consensus.

B. Histories and Memory

At each period t voters observe data from only the last $K < \infty$ periods. In particular, we start the model at period 0 with an initial history $H_0 = (p_\tau, y_\tau)_{\tau=-K}^{\tau=-1}$ and denote the history observed by voters at period $t \geq 0$ by $H_t = (p_\tau, y_\tau)_{\tau=t-K}^{\tau=t-1}$ where for any τ , $p_\tau \in \{l, r\}$ is the implemented policy in period τ and y_τ is the policy outcome in that period. Thus, at every period t , the information Ω_t held by voters is composed of the prior G , and the K -period history H_t . We model short-term memory by assuming that voters are not able to decode past learning from the behavior of parties and past electorates. To do so, we use two assumptions: Each period voters use the same original prior G , and treat the history as exogenous. Thus, voters do not ask themselves "why did the parties choose those platforms in the last few periods."

Voters use Bayes rule to compute their posterior distribution G_t on the vector β , given $\{G, H_t\}$. Specifically, the posterior density distribution $g_t(\cdot)$ on vectors $\beta = (\beta_l, \beta_r)$ is

$$g_t(\beta) = \frac{g(\beta) \prod_{\tau=t-K}^{\tau=t-1} f(y_\tau - E[y_\tau | p_\tau, \beta])}{\int_{\beta'} g(\beta') \prod_{\tau=t-K}^{\tau=t-1} f(y_\tau - E[y_\tau | p_\tau, \beta']) d\beta'},$$

where $E[y_\tau | l, \beta] = \beta_l$, $E[y_\tau | r, \beta] = \beta_r$, and $f(\cdot)$ is the (normal) density of the shock ε_t .

One can consider alternative assumptions on memory. For example, that individuals put different weights on different past periods. All our results hold if instead of short term memory, the weight on events in the past declines fast enough with time.

C. Electoral Competition

As already shown in Wittman (1983) and Calvert (1985), party polarization can arise in an electoral competition when parties are ideological (as we assume) and when they face some uncertainty with regards to voting behavior. Above we described the element of preferences which is common to voters (the outcome y_t , and observing the same history H_t). Naturally voters may differ on additional aspects which parties may have uncertainty about. For example, voters may react differently to the personal attributes of candidates. We summarize all these other factors, as in the probabilistic voting model (Lindbeck and Weibull 1987), by a random variable ϕ_t which represents the bias of the median voter towards party L . In period t , the median voter votes for party L if

$$(3) \quad E[y(p_t^L) - y(p_t^R)|H_t] + \phi_t > 0,$$

where ϕ_t is i.i.d and uniformly distributed on $[-1/(2\zeta), 1/(2\zeta)]$, where $\zeta \in (0, \infty)$, and

$$\begin{aligned} E[y(p_t^L) - y(p_t^R)|H_t] \\ = \begin{cases} E[(\beta_l - \beta_r)|H_t] = \int_{\beta} (\beta_l - \beta_r) g_t(\beta) d\beta, & \text{if parties polarize} \\ 0, & \text{otherwise.} \end{cases} \end{aligned}$$

In the case of equality in (3), without loss of generality, we assume that the voter votes for party L with probability $1/2$.

We assume that voters and parties are myopic. Myopia is a standard assumption in models of electoral competition; we discuss the effects of forward-looking parties and voters in Section VB.

Given the above, in period t , the probability that party L wins the election is

$$\begin{aligned} \Pr(L \text{ wins} | p_t^L, p_t^R) &= \Pr(\phi_t > -E[y(p_t^L) - y(p_t^R)|H_t]) \\ &= \begin{cases} 1, & \text{if } \frac{1}{2} + \zeta E[y(p_t^L) - y(p_t^R)|H_t] > 1 \\ 0, & \text{if } \frac{1}{2} + \zeta E[y(p_t^L) - y(p_t^R)|H_t] < 0 \\ \frac{1}{2} + \zeta E[y(p_t^L) - y(p_t^R)|H_t], & \text{otherwise} \end{cases} \end{aligned}$$

Given the platform of party R , p_t^R , party L maximizes her expected utility,

$$\Pr(L \text{ wins} | p_t^L, p_t^R) (U^L(p_t^L) + \alpha) + [1 - \Pr(L \text{ wins} | p_t^L, p_t^R)] U^L(p_t^R),$$

by choosing p_t^L in equilibrium. Similarly, given p_t^L , party R chooses p_t^R to maximize an analogous expression, where $U^J(p)$ is defined in (2) and parties know G, H_t . For

completeness we assume that if a party is in equilibrium indifferent between offering l or r then it offers its ideal policy (this assumption plays no role in the analysis as this case will generically not arise).

At any period t , the equilibrium in the electoral competition model is a Nash equilibrium between the parties, where parties choose the platforms p_t^L and p_t^R simultaneously to maximize their utilities detailed above. Following the choice of platforms, ϕ_t is drawn and the election result is in accordance with (3).

D. Dynamics

The only dynamic link between the periods is the evolving memory of voters. The dynamic model is defined therefore, as follows:

- (i) There is some initial history H_0 of size K .
- (ii) In period t the party that won the election, party $J \in \{L, R\}$, implements $p_t^J \in \{l, r\}$.
- (iii) Given y_t , history evolves from $H_t = \{p_\tau, y_\tau\}_{\tau=t-K}^{t-1}$ to $H_{t+1} = \{p_\tau, y_\tau\}_{\tau=t-K+1}^t$.
- (iv) A new electoral competition equilibrium arises in period $t+1$: the two parties offer their equilibrium policy platforms p_{t+1}^J , ϕ_{t+1} is drawn and party L wins the election if

$$E[y(p_{t+1}^L) - y(p_{t+1}^R) | H_{t+1}] + \phi_{t+1} > 0,$$

or with probability $1/2$ if the above is satisfied with equality.

E. The Incentives to Polarize: A Useful Lemma

Our first result highlights the mechanism by which the level of information in historical data affects electoral competition. Fix a history H_t and consider the one-period political competition game that ensues:

LEMMA 1 (Consensus vs Polarization): *In period t :*

- (i) *If*

$$|E[(\beta_l - \beta_r) | H_t]| > 1/[2\zeta(1 + \alpha)],$$

then the unique equilibrium involves both parties choosing the same policy (consensus);

- (ii) *If*

$$|E[(\beta_l - \beta_r) | H_t]| < 1/[2\zeta(1 + \alpha)],$$

then the unique equilibrium involves each party choosing its ideal policy (polarization).²

Generally speaking, parties prefer to pursue their own interests. But they also care about being elected. This implies that sometimes they may be disciplined by voters to choose the policy that voters think is more likely to generate a higher outcome given the historical data. If, given the historical data, voters sufficiently prefer one of the policies, say l , then parties have to offer l and are hence in consensus. In this case, if a party offers r , it will serve neither its own policy interest nor its office motivation, as it will face only a slim probability of being elected. Alternatively, if the historical data does not allow voters to sufficiently discriminate between the two policies, then parties can afford to take the risk and offer platforms that better serve their own policy interests and hence polarize.

Note that a higher office rent α pushes parties to offer the same policy, as they are more eager to get elected and therefore to satisfy the voters' will. This also implies that when parties do polarize, each party must win with a probability that is bounded away from zero; if a party polarizes and has only a negligible chance of being elected, it can deviate to the same platform as the other party, gain $1/(2\alpha)$, and be better off. A higher ζ implies a smaller variance for the shock ϕ_t : This means that parties are more certain about how voters vote, and again this pushes parties to reach consensus more often. As both ζ and α work in the same direction, we can denote the right-hand-side of the expressions in the Lemma as

$$\rho \equiv 1/[2\zeta(1 + \alpha)],$$

where a lower value for ρ is more conducive to party consensus.

To make learning meaningful in the model, it is reasonable to consider the case where parties are in consensus when voters know the true state. With this in mind, given Lemma 1, we will assume that $|\beta_l^* - \beta_r^*| > \rho$. This is a reasonable requirement in a social learning model, to allow voters to reach the optimal policy. For concreteness and without loss of generality, let us assume that the state of the world, a primitive of the model, is such that l is the optimal policy so that $\beta_l^* > \beta_r^*$.

II. Cycles of Consensus and Polarization

In what follows we are interested in whether and how a polity transitions between periods of party polarization and party consensus. To this end the following notation will be useful. Let $\eta_t(\text{polarization})$ denote the fraction of time in the full history up to time t that the two parties offered different platforms and by $\eta_t(\text{consensus}) = 1 - \eta_t(\text{polarization})$ the fraction of time in the full history up to time t that the two parties offered the same platform. Denote by $\hat{\eta}_t(p)$ the fraction of time in the full history up to time t that policy p was implemented.

²Both types of equilibria can hold when $|E[(\beta_l - \beta_r)|H_t]| = 1/[2\zeta(1 + \alpha)]$.

Note that the dynamic evolution of policies involves some randomness, given the voting shock ϕ_t and the policy shock ε_t (through the latter's effect on beliefs). This then induces a probability distribution P over the set of infinite paths of history $\mathbb{H}$. Thus, when we write “almost surely” below, here and in the Appendix, we mean P -almost surely on $\mathbb{H}$.

A. A Benchmark: Unbounded Memory

As a benchmark we first consider the case in which the history that voters remember is unlimited, that is, when $K = \infty$. The model and specifically our definitions of the evolution of histories can be easily extended to this infinite memory case.³ Our result below shows that with full memory, the two parties will reach a consensus on the same platform.

PROPOSITION 1: *Assume that $K = \infty$. Then almost surely the polity experiences long-term consensus, i.e., $\eta_t(\text{consensus}) \rightarrow 1$ and $\hat{\eta}_t(p) \rightarrow 1$ for some $p \in \{l, r\}$.*

To see the intuition for Proposition 1, note first that long-term polarization cannot arise in equilibrium. When $K = \infty$, the beliefs of voters will converge in the long term as they form a martingale. If beliefs converge so that voters do not sufficiently distinguish between the utilities of the two policies, parties polarize in line with Lemma 1. As parties also have office motivations, it implies that when parties polarize, each has a probability of being elected that is bounded from zero. Given the stochastic political turnover, polarization implies that voters (unintentionally) “experiment” in the long term with two different policies: This then allows voters to learn the true state, β^* . As we assumed that $\beta_l^* - \beta_r^* > \rho$, this will lead to a contradiction as by Lemma 1 the two parties must reach a consensus on l . When $K = \infty$, voters' beliefs in the long term must converge then to beliefs that induce parties to be in consensus.⁴

Consensus however is *not* guaranteed to be on the optimal policy; as in any learning problem with myopic agents, learning can sometimes lead voters to believe that the wrong policy is optimal due to insufficient “experimentation.” For example, a rare and finite series of very bad shocks when l is implemented, may convince voters that r is more likely to be better and therefore induce parties to reach a consensus on r . This consensus can be sustained in the long run; voters will learn β_r^* but might maintain beliefs that the difference between the expected utilities of the two policies is large enough, as they never forget the initial series of bad shocks on l . Thus, long-term memory also implies that long-term outcomes can be history dependent.

³Specifically, we start with some finite initial history H_0 , and at any period t , given y_t , history evolves from H_t to $\{H_t(p_t, y_t)\}$.

⁴The condition $|\beta_l^* - \beta_r^*| > \rho$ is only a sufficient condition for long-term consensus. If $|\beta_l^* - \beta_r^*| < \rho$, then beliefs can converge to either allow long-term party polarization or long-term party consensus.

B. Short-Term Memory and Cycles

We now consider short-term memory, that is, finite K . In this case the nature of voters' data can change over time. If, for example, power did not change hands or parties' platforms are the same, history contains little variation in policies and voters' data is rather uninformative. Alternatively when the history involves party polarization and a high political turnover, voters' data will be relatively informative.

We already know by Lemma 1 that if voters' history is null, then parties will polarize, as according to the (symmetric) prior, $|E[\beta_l - \beta_r]| = 0$. At the other extreme, if voters' history is fully informative and voters learn the state, parties will reach a consensus on l as $\beta_l^* - \beta_r^* > \rho$. What happens on the equilibrium path with short-term memory will depend then on whether voters' knowledge allows them to sufficiently distinguish the expected utilities of the two policies. The evolution of the short-term knowledge of voters is further complicated by the fact that policies are themselves endogenous, and by the stochastic noise in the learning process.

To illustrate the mechanism of cycles it is instructive to shut down this stochastic noise and so we first consider the case in which the variance of the outcome shock is infinitesimal. When policy noise vanishes, learning is very fast. Thus, after any period, in the limit as $\sigma \rightarrow 0$, voters learn the true effectiveness of the policy that was implemented in that period. As a result, observing two periods with different implemented policies allows voters to learn the state in full.

Consider any $K \geq 2$ and the case in which voters' observed history contains only one implemented policy $p \in \{l, r\}$. As $\sigma \rightarrow 0$, voters will learn the true effectiveness of p , β_p^* . The condition below relates to what voters will learn about the other policy (that was not implemented), denoted by $-p$. Let $E[\beta_{-p} | \beta_p^*]$ denote the expected effectiveness of policy $-p$ when voters know β_p^* (and nothing else). The following assumption will allow us to characterize when cycles arise. For completeness, Assumption 1 includes the assumption that we made at the end of Section I, that $\beta_l^* - \beta_r^* > \rho$.

ASSUMPTION 1: *The state β^* satisfies*

$$\max\{|\beta_l^* - E[\beta_r | \beta_l^*]|, |\beta_r^* - E[\beta_l | \beta_r^*]|\} < \rho < \beta_l^* - \beta_r^*.$$

Assumption 1 implies that when voters learn the benefit of one policy only, they are less sure about the relative benefit of the two policies, compared to the case in which they learn the benefit of both. We postpone for later the discussion of when Assumption 1 is likely to hold. We henceforth focus on the parameters that satisfy Assumption 1.

In Proposition 2 we derive systemic cycles of polarization due to endogenous learning.

PROPOSITION 2: *Let $\sigma \rightarrow 0$ and $K \geq 2$. Then the polity experiences perpetual cycles of polarization and consensus. The unique limit equilibrium has two phases:*

- (i) *A consensus phase lasting exactly K periods in which both parties espouse the optimal policy l .*

- (ii) *A polarization phase in which each party espouses its ideal policy. This stage lasts until the party espousing r wins the election. The expected length of the polarization phase is*

$$\frac{1}{1/2 - \zeta(\beta_l^* - E[\beta_r|\beta_l^*])},$$

and so, in the limit equilibrium, the fraction of time the polity implements the correct policy is larger than $1 - 1/K$.

The proposition uncovers systemic polarization cycles. To see the intuition, assume that the polity starts with both parties offering the same policy p . As σ is vanishing, after K periods, the knowledge of voters is perfect about β_p^* , whereas their knowledge on β_{-p} arises from the prior G and their knowledge of β_p^* , with an expectation equal to $E[\beta_{-p}|\beta_p^*]$. By Assumption 1 parties will then polarize. Then, at some point in this polarization phase, $-p$ will be implemented, allowing the polity to learn from the most recent period and the one before that l is optimal. A consensus phase then arises in which both parties espouse l , and this lasts as long as voters remember their benefit from each policy. Following the K th period of the consensus phase, voters only remember periods in which l was implemented, which by Assumption 1 means that parties polarize. Thus, society is in an absorbing state of cycles between phases of consensus on l , and polarization until r is implemented.⁵

Assumption 1 is a sufficient condition that guarantees that cycles can arise: $\beta_l^* - \beta_r^* > \rho$ implies that periods of polarization will result in consensus on l and $\rho > \max\{|\beta_l^* - E[\beta_r|\beta_l^*]|, |\beta_r^* - E[\beta_l|\beta_r^*]|\}$ implies that periods of consensus on l (and more generally also on r) will result in polarization. In the Supplemental Appendix we characterize equilibria when Assumption 1 does not hold. We show that the equilibria have either cycles as above, or perpetual consensus or perpetual polarization.

A Discussion of Assumption 1: When would we expect Assumption 1 to be satisfied? To answer this, note that we can decompose the assumption into two. The first states that $\beta_l^* - \beta_r^* > \max\{|\beta_l^* - E[\beta_r|\beta_l^*]|, |\beta_r^* - E[\beta_l|\beta_r^*]|\}$. This is a condition on the state β^* and the prior G . The second part requires that ρ , which captures parameters relating to electoral competition, falls within the range of the above two values.

The condition $\beta_l^* - \beta_r^* > \max\{|\beta_l^* - E[\beta_r|\beta_l^*]|, |\beta_r^* - E[\beta_l|\beta_r^*]|\}$ highlights an intuitive feature of policy evaluation between observed and unobserved policies. In practice, relative policy evaluation between two policies that have both been experienced in the past is easier compared with the case in which only one policy has been experienced. The condition boils down to requiring that when only one policy is experienced it is hard to tell the difference between the two policies. While voters understand well the benefits of this policy, they might not know how good or bad they have it compared with an alternative. This condition is intuitive yet in contrast

⁵ While here we present the results with a stark model of a finite memory, our results hold more generally with infinite weighted-memory, as long as weights on the past histories decay fast enough.

to standard models of voters with ideal policies, which assume that voters know what is their satiation level.

One way to generate the condition above is in the case in which voters' learning is "scale-free."⁶ In particular, suppose that the scale of utility voters experience from one policy is not too informative about the other; voters might not be sure what is the highest utility scale they can gain, and whether they had reached this level. In the Supplemental Appendix we show that the condition can then be easily satisfied; we provide an example of a "scale-free" learning environment, for which almost all states satisfy it. In the simple example that we present, $\max\{|\beta_l^* - E[\beta_r|\beta_l^*]|, |\beta_r^* - E[\beta_l|\beta_r^*]|\} \approx 0$, and so learning the utility scale of one policy provides no knowledge about which one is better.

The second part of Assumption 1 involves the level of ρ . In our analysis we assume that ρ is fixed throughout but more generally one can entertain that the parameters that affect the level of ρ , namely ζ and α , might change over time. The level of uncertainty about voters' intentions, ζ , might depend on particular aspects of the election such as political engagement and media consumption. The utility of parties from winning office, α , might also change from year to year depending on politicians' ability to capture rents once in office. As long as the distribution of ρ is such that with strictly positive probability it is in the desired range, the qualitative nature of our results will hold.

The Length of Consensus and Polarization Phases: Note that the parameters discussed above affect the length of the polarization phase in a natural way. Proposition 2 readily implies that the consensus phase is higher when K is larger as voters retain for longer the knowledge of the effectiveness of both policies.

The average length of the polarization phase is in turn affected by the probability that policy l is elected in this phase, which increases in $\zeta(\beta_l^* - E[\beta_r|\beta_l^*])$. When $\beta_l^* - E[\beta_r|\beta_l^*]$ is higher (but still lower than ρ), voters are more keen on l , and when ζ is higher, their vote is more certain in this direction, implying that l is more likely to be elected. The more likely l is to be elected in a polarization phase, the longer is the polarization phase on average, as it only ends once r is elected for the first time. Note that a longer phase of polarization (when $\sigma \rightarrow 0$) is beneficial for voters as it implies that on average, the correct policy l is implemented more often throughout the cycle.

We have so far abstracted away from exogenous shocks as we have assumed above that $\sigma \rightarrow 0$. Naturally these have an important role to play in the dynamics of politics. Our next result establishes that the cycles uncovered above also arise when $\sigma > 0$ and K is large enough.

PROPOSITION 3: *For a large enough K , the polity almost surely experiences cycles of polarization and consensus. In particular, (i) there exists an $\eta_K > 0$ such that for any $\sigma > 0$, $\liminf_{t \rightarrow \infty} \eta_t(\text{polarization}) > \eta_K$ and $\liminf_{t \rightarrow \infty} \eta_t(\text{consensus}) > \eta_K$ almost surely. (ii) There is a strictly positive probability that when the polity is in consensus it is on the wrong policy.*

⁶We thank a referee for suggesting the idea of scale-free learning.

TABLE 1—EMPIRICAL MOMENTS FROM SIMULATION

	$K = 5$			$K = 10$		
	$\sigma = 0.2$	$\sigma = 1.2$	$\sigma = 2.5$	$\sigma = 0.2$	$\sigma = 1.2$	$\sigma = 2.5$
Optimal policy (%)	80.22 (3.31)	57.51 (10.12)	48.25 (9.53)	87.76 (3.20)	65.34 (14.68)	53.14 (14.74)
Consensus (%)	73.78 (4.05)	59.12 (8.87)	61.71 (6.99)	83.96 (3.85)	64.95 (13.28)	63.08 (11.05)
Consensus on optimal policy (%)	100.00 (0.05)	88.88 (6.92)	72.08 (10.72)	100.00 (0.04)	93.11 (8.31)	77.51 (14.06)
Length of consensus phases	4.69 (0.26)	4.31 (0.92)	3.98 (0.89)	9.22 (0.56)	7.43 (2.17)	6.08 (2.05)

By comparing the result above to the one in Proposition 2, note that the stochastic element of exogenous outcome shocks affects the nature of cycles in the following ways. First, a large level of noise implies that during a polarization phase voters may not necessarily learn the true optimal outcome. This implies that parties might reach a consensus on the wrong policy. For example, a finite series of positive shocks when implementing policy r and/or negative shocks when implementing l might convince voters that policy r is substantially better than policy l . While this series of shocks might not happen too frequently, it can arise with a strictly positive probability.

Second, during a polarization phase, it may take longer to converge to some knowledge (correct or not) that some policy is better than the other. This implies that polarization phases might be longer and involve more political turnover. Moreover, large policy shocks can trigger abrupt switches from party polarization to party consensus and vice versa. A large one-off shock during a phase of polarization can convince voters that policy p is very beneficial and can induce a consensus on p , whereas in such a consensus phase, another shock can dissuade voters from implementing p altogether.

To illustrate further how the systemic element of the cycles works in conjunction with the exogenous shocks to affect the length of the different phases of the cycles, we now present results from simulations of the model. These also allow us to shed some light on the normative implications of the model (e.g., with regard to how often the optimal policy is implemented). For the simulation, we use $B = [0, 6]^2$ with a uniform prior. We set $\beta_l^* = 3.5$ and $\beta_r^* = 2.5$ (so the optimal policy is l), and $\zeta = \alpha = 0.5$ so that $\rho = 2/3$. The parameters are chosen to satisfy Assumption 1. We randomize the initial history H_0 and then run the simulation for a hundred periods.

Table 1 reports positive and normative implications of the model using averages of 20,000 such simulations. For each vector of values for σ and K we measure the average length of a consensus phase, the fraction among the periods of consensus on which consensus is on the optimal policy, and the proportion of time the polity implements the optimal policy (either in a consensus or a polarization phase).

The variables σ and K affect the results in similar ways as they are both related to how much the voters can learn; when σ decreases, learning is faster and so for any K , the polity is able to implement the optimal policy in more periods and also a greater

proportion of consensus periods are on the true optimal policy. As K increases, the polity has a longer memory to learn from, implying similar results. As we can see, the length of the consensus phase is decreasing in σ and gets closer to its limit of K as σ gets smaller, in line with Proposition 2.

While, as shown above, both a high K and a small σ increase the information voters have, the comparative statics of these parameters in the limit are more nuanced. Specifically we now explore the cycle outcomes when $\sigma > 0$ but $K \rightarrow \infty$.

PROPOSITION 4: *Let $\sigma > 0$. As $K \rightarrow \infty$, almost surely:*

- (i) *Consensus arises, i.e., $\lim_{K \rightarrow \infty} \limsup_{t \rightarrow \infty} \eta_t(\text{polarization}) / \eta_t(\text{consensus}) = 0$.*
- (ii) *The wrong policy is implemented in a strictly positive fraction of time, that is, $\lim_{K \rightarrow \infty} \liminf_{t \rightarrow \infty} \hat{\eta}_t(r)$ is bounded from below.*

The result indicates that while cycles always exist, the share of time the polity spends in a polarization phase shrinks as $K \rightarrow \infty$. But, as in Proposition 1, due to the myopia of voters and parties, consensus may arise on the wrong outcome when $\sigma > 0$. This arises no matter how high K is and in contrast to Proposition 2. Thus, while both a higher K and a lower σ imply better learning on average, they affect learning and hence outcomes in different ways.

Proposition 4 together with the simulations illustrated some of the normative properties of the outcomes of the cycle in terms of the share of time the optimal policy is implemented. We now elaborate more on the welfare properties of short-term memory.

III. Short-Term Memory, Changing States, and Welfare

In this section we evaluate the welfare properties of a political system with voters who have short-term memory. Our welfare criterion is long-run information aggregation, that is, the long-run share of time in which the implemented policy accords with the optimal policy. We consider the welfare cost of short-term memory (compared to a polity with $K = \infty$), and then show that, in some environments, short-term memory can actually be beneficial. We focus throughout on the case of $\sigma \rightarrow 0$, which implies in our model that a polity with a long-term memory converges to implement the optimal policy almost surely.

To see the cost of short-term memory as compared to a long-term memory, recall that as we had seen in Proposition 2, the optimal policy is implemented for at least $1 - (1/K)$ of the time. The cost of short-term memory is that every now and then polarization arises and the polity necessarily implements the suboptimal policy r . However, if short-term memory is not too short, the share of time in which this happens is small.

Short-term memory may however have benefits in some environments. Specifically, forgetfulness can be useful if the state of the world might change over time. Indeed sometimes polities might experience a change in technology which can be triggered by both external factors (e.g., a war, a pandemic) as well

as by endogenous factors (e.g., climate change and technological breakthroughs). Remember that voters are myopic in our model and so do not intentionally experiment. But with short-term memory the polity drifts into polarization phases now and then which, when the state of the world changes, might serve as periods of “unintended experimentation.” This can allow the polity to detect changes in the state. Below we identify environments in which voters with short-term memory and who are not aware that the state of the world may change, may gain a higher welfare compared with voters with unbounded memory that are aware that the state of the world may change.

To illustrate this, consider an environment in which the state of the world β^* might change over time. Some changes to the fundamentals are easily detected, but some changes may remain undetected at least for a while. For example, if the polity is currently in consensus on l , a change to β_r^* —which may imply that now r is the optimal policy—will not be detected. Voters with short-term memory will detect these changes once they reach a polarization phase and will switch to the new optimal policy. How will voters with long-term memory fare in this environment? As they are aware of such possible changes to the state they will have a belief that a change has occurred before time t ; denote it by $\phi(t)$. Using this, voters can compute the expected utility difference between policies l and r , which generally can be written as

$$\Delta(t) \equiv \phi(t) E_{t,\text{post-change}}[\beta_l - \beta_r] + [1 - \phi(t)] E_{t,\text{pre-change}}[\beta_l - \beta_r].$$

By Lemma 1, for any t such that $\Delta(t) > \rho$, parties will remain in consensus on l , and after the first t' for which $\Delta(t') < \rho$, parties will polarize. Once parties polarize, the polity will eventually detect the change and will return to implementing an optimal policy.

Once a change does occur and goes undetected as above, the question is how quickly consensus unravels to allow voters to learn the new optimal policy. As can be seen from the expression $\Delta(t)$, the timing of the unravelling of consensus does not depend only on the level of $\phi(t)$ but also on the expectations of $(\beta_l - \beta_r)$ before and after a change; this may imply that voters may be stuck on the wrong consensus for a very long time even though they believe that there is a high probability that a change of the state has already occurred. Example 1 formalizes the intuition above.

EXAMPLE 1 (A Simple Model with Changing States): Assume that there are only two possible states of the world, a and b , with β^a and β^b , where $\beta_l^{a*} - \beta_r^{a*} > \rho$ and $\beta_r^{b*} - \beta_l^{b*} > \rho$. The prior, $\pi \in (0, 1)$, is the probability that the state is a . In any period, with Poisson arrival rate λ , a new state is drawn from the prior distribution. Assume further that $\beta_l^{a*} = \beta_l^{b*} = \beta_l^*$; this will imply that a change in the state cannot be observed when policy l is implemented (but can be detected if r is implemented). In the spirit of Assumption 1 assume that $0 < \beta_l^* - [\pi \beta_r^{a*} + (1 - \pi) \beta_r^{b*}] < \rho$. This insures that when voters are certain that there was a change, parties will polarize. We then have the following.

PROPOSITION 5: Consider Example 1. Let $\sigma \rightarrow 0$ and $K > 1/(1 - \pi)$. Then, when $\lambda \rightarrow 0$, there are values of $\{\beta_l^*, \beta_r^{a*}, \beta_r^{b*}\}$ for which voters with short-term

memory (who are not aware of the possibility of changing states) will implement the optimal policy for a higher fraction of time compared to voters with long-term memory (who are aware that the state may change).

Intuitively, when voters have long-term memory their myopia implies that they do not engage in intentional experimentation; considering their current payoff, even a strong possibility that a change of state had occurred may still not lead to polarization. This is because, in expectation, they may view the current consensus policy as still optimal. In contrast, voters with short-term memory will always unintentionally experiment as they forget why they have settled on the current consensus policy.

A more general analysis of changing states may yield other implications. For example, if naive voters in our model do realize that the state of the world can change, they can attach less weight to periods in which they believe the state was less “similar” in some sense to the current one, or alternatively, completely discard historical information which is not relevant to the current period. To a degree, this implies that even when voters have full length memory, then the possibility of changing states may imply that short-term memory can arise endogenously. We leave this for future research.

IV. Extension: Different Memories

The recollection of history might be different for different voters or groups. One example of this is highlighted in a recent literature which suggests that voters’ beliefs are shaped by experiences mostly accumulated during their formative years. Malmendier and Nagel (2016) show that life-time experiences of inflation significantly affect beliefs about future inflation, and that this channel explains the substantial disagreement between young and old individuals in periods of highly volatile inflation, such as the 1970s. Alternatively, a large literature studies and documents the effect of echo chambers on polarization, and so different groups that are exposed to different sources of information can end up with different memories and recollection of histories.

As we show below our model is well equipped to analyze the dynamics of political systems when different voter groups have different memories. To see this, we adjust the model above and assume that the voting population is divided into m groups. All the voters in group j observe the same data which constitutes their memory. In period t , the history of each group j is the pairs of policies and outcomes in some selected history with K^j periods, denoted by H_t^j . This allows the voters in group j to compute their posterior distribution G_t^j on the vector β in a similar fashion to our model above.

Consider the electoral competition between the two parties at some period in which the groups are exogenously endowed with their respective histories. Below we will omit the subscript t when no confusion arises. We assume that within group j and across the groups, voters may differ on additional dimensions. Thus, voter i in group j votes for party L if

$$E[y(p^L) - y(p^R) | H^j] + v^j + \phi > 0,$$

where ϕ is the aggregate shock uniformly distributed on $[-1/(2\zeta), 1/(2\zeta)]$ and v^{ij} is an idiosyncratic group-specific shock distributed on $[-1/(2\xi^j), 1/(2\xi^j)]$. The indifferent voter in each group satisfies

$$\hat{v}^{ij} = E[y(p^L) - y(p^R)|H^j] + \phi,$$

and so party L 's overall vote, given each group's share in the population γ^j , where $\sum_{j \in \{1, \dots, m\}} \gamma^j = 1$, is

$$\frac{1}{2} + \sum_{j \in \{1, \dots, m\}} \gamma^j \xi^j (E[y(p^L) - y(p^R)|H^j] + \phi).$$

Party L wins therefore if

$$\sum_j \gamma^j \xi^j E[y(p^L) - y(p^R)|H^j] > -\phi \sum_j \gamma^j \xi^j$$

which arises with probability

$$\Pr(L|p^L, p^R)$$

$$= \begin{cases} 1, & \text{if } \frac{1}{2} + \zeta \sum_j w^j E[y(p^L) - y(p^R)|H^j] > 1 \\ 0, & \text{if } \frac{1}{2} + \zeta \sum_j w^j E[y(p^L) - y(p^R)|H^j] < 0 \\ \frac{1}{2} + \zeta \sum_j w^j E[y(p^L) - y(p^R)|H^j], & \text{otherwise} \end{cases}$$

where $w^j = \gamma^j \xi^j / \sum_j \gamma^j \xi^j$ denotes the political weight of group j in the electoral competition. This weight is increasing in group size as is intuitive and decreases in the variance of the distribution of idiosyncratic shocks of group j . The latter effect is due to the fact that as the variance decreases, this group of voters is more sensitive to policy utility differences.

Anticipating the above, and with the knowledge of H_t^j , parties choose policies in equilibrium to maximize their expected utility as before; we then have an analogous result to Lemma 1.

LEMMA 2: *In period t , a consensus on policy $p \in \{l, r\}$ arises if*

$$\sum_{j \in \{1, \dots, m\}} |w^j E[(\beta_p - \beta_{-p})|H_t^j]| > \rho,$$

and polarization arises if for any policy $p \in \{l, r\}$,

$$\sum_{j \in \{1, \dots, m\}} |w^j E[(\beta_p - \beta_{-p})|H_t^j]| < \rho.$$

When a group has strong preferences towards one policy given their memory, consensus is more likely to arise when this group's political weight is larger; similarly, when a group has weak preferences for policy given their memory, it would be easier for parties to polarize if this group's political weight is larger.

To see the implications of Lemma 2, we now concentrate on a specific case of different memories. Assume that each group j remembers the recent K^j periods. Without loss of generality order groups in terms of the length of their history, that is, $K^m > K^{m-1} > \dots > K^1 \geq 2$. An example of such a polity could be age groups, so that older voters have a longer observable history.⁷ We maintain Assumption 1 and consider generic distributions so that $w^j > 0$ and $w^j \neq w^{j'}$ for any j, j' . We also consider for simplicity $\sigma \rightarrow 0$:

PROPOSITION 6: *The equilibrium in the model with heterogeneous (and nested) length of memories is characterized by cycles of polarization and consensus. Specifically, for all distribution of political weights,*

- (i) *There is a pivotal group $j^* \in \{1, \dots, m\}$ such that the length of the consensus stage is K^{j^*} , the length of memory of j^* .*
- (ii) *The expected length of the polarization phase depends on the memory length of all m groups.*

When policy noise is vanishing, any group's knowledge centres on either $|\beta_p^* - E[\beta_p | \beta_{-p}^*]|$ or on $\beta_l^* - \beta_r^*$ as in Proposition 2. Consider just two groups, 1 and 2. Assume a consensus phase, let's say on l , and consider the polity after K^1 such periods. Group 1's knowledge has just switched from $\beta_l^* - \beta_r^* > \rho$ to $|\beta_l^* - E[\beta_r | \beta_l^*]| < \rho$, whereas group 2's knowledge constitutes of more periods and is therefore still at $\beta_l^* - \beta_r^* > \rho$. From Lemma 2 we know that if group 1's political weight is sufficiently large, parties will now polarize and this group will be the pivotal one. If its political weight is too small, parties will be in consensus for K^2 periods, that is, until the knowledge of group 2 also switches from $\beta_l^* - \beta_r^*$ to $|\beta_l^* - E[\beta_r | \beta_l^*]|$. The Proposition shows more generally that there is always one unique such pivotal group.

While the length of the consensus phase only depends on the memory capacity of the pivotal group, the length of the polarization phase depends on the length of the memory of all groups. To see why, note that once parties polarize, the polarization phase ends once the two policies had been implemented one after the other. The probability of each policy to be elected depends on $\sum_{j \in \{1, \dots, m\}} w^j E[(\beta_p - \beta_{-p}) | H_t^j]$, and so on the knowledge of all groups.

To illustrate what this implies, consider again the case of two groups. Suppose that group 1 is the pivotal one, and now increase K^2 to infinity. As we do this we do not change the length of the consensus phase in equilibrium which is exactly

⁷ In Levy and Razin (2023) we study an overlapping generations model and focus on the generational divide in terms of the information of different cohorts of voters. We establish a similar cyclical pattern to that we uncovered in Propositions 2 and 3 in which the polity cycles between periods of polarization and consensus. For recent literature on polarization and extremism in different demographic groups see Ortoleva and Snowberg (2015) and Boxell, Gentzkow, and Shapiro (2017).

K^1 . But by increasing K^2 we guarantee that group 2 will almost surely know the state at any period. This will make group 2 quite “frustrated” about politics as they know what should be done, but because of their de facto weak political power they are not able to “save” the polity from choosing r every now and then in the polarization phase. Still, all is not lost as group 2’s superior information affects the outcome even though it does not affect the strategies of parties; specifically, group 2’s knowledge increases the probability that the correct policy l wins the election in a polarization phase. While this means that the polarization phase is longer, this is welfare enhancing: It will take longer to sample policy r during a polarization phase and so the correct policy l will be implemented for longer in expectation.

Note that the above illustrates that it can be welfare-improving to have heterogeneity of groups’ memories. If we start with a society where memory length is homogenous at some level K , all in society can be better-off if some parts of society will have less memory and some more. Assume that we create such heterogeneity while maintaining the pivotal group as the one with memory length K . As policy noise is vanishing, this implies that the length of the consensus phase is as before, while the length of the polarization phase increases. This is welfare-improving as then policy l is implemented in equilibrium more often in expectation.

V. Discussion

In this section we discuss our results and assumptions and suggest some extensions of the main model.

A. *Rationally Inattentive Voters*

In our model, voters’ learning is characterized by one key behavioral feature, which is short-term memory. Beyond that, voters are fully attentive to policies and political outcomes, and update their beliefs rationally given the information they have. However, it is reasonable to consider the case that voters are not always fully attentive to politics. And in the context of our model, it is reasonable that voters’ attention may differ in times of consensus and in times of polarization: In times of political polarization voters might pay more attention in order to understand whom to vote for; the benefits of doing so are more apparent as parties offer different policies. In periods in which parties are in consensus, voters might pay little attention to politics as they see little difference between the parties and will therefore be ill informed. One way to think about this is to consider news content providers that in times of consensus may focus on entertainment news to attract readers, while in times of polarization, due to voters’ heightened demand for information, may put more emphasis on political content.

For simplicity, suppose that we extend our model and assume that in times of consensus voters do not observe the political outcome y , while in time of polarization all remains as before. Following a consensus on some policy, after K periods, voters’ history is null as voters did not pay attention to any outcome in these K periods. Thus the only information voters have is the prior G . As G is symmetric, $E[\beta_l - \beta_r] = 0$, and parties will polarize. Once a polarization phase starts, at some point voters will acquire sufficient knowledge that will induce a consensus,

as before. In such a model, an analogous assumption to Assumption 1 facilitates cycles. Thus, adding rational inattention to the model does not change the qualitative nature of our results.⁸

B. *Nonmyopic Voters and Parties*

In our model voters and parties are fully myopic. The role of myopia was highlighted in Section III. We now discuss the possibility of having preferences with longer horizons.

First note that the effect of preferences with long-term horizons is potentially restricted due to the short-term memory of future actors. This implies that current voters and parties can only directly influence a finite number of periods in the future. The only effect that they can have on longer horizons is through the indirect effects of their behavior today, such as changing the pattern and timing of future cycles.

Consider voters first. If voters are not myopic, they may have some incentive to actively experiment and especially so when they are least informed. This can happen during a phase of consensus in which, as time goes by, voters' information about the difference between the two policies becomes weaker. To see this, consider $\sigma \rightarrow 0$. Assume voters' history contains only instances of policy l being implemented. At this point voters compare the current period's benefit of voting for l which is $\beta_l^* - E[\beta_r|\beta_l^*]$ to the future benefit of learning what is the true state tomorrow. As long as $\beta_l^* - E[\beta_r|\beta_l^*] > 0$ our results are robust to voters having a small discount factor. For higher discount factors voters might be tempted to experiment by preferring to vote for policy r ; if this motivation is very strong, parties may then switch to both offering r and so sometimes, rather than polarization, we may have a switch to a different consensus.

For parties, myopia is a standard assumption in models of electoral competition: From their perspective, the stakes are so high that it is reasonable that they concentrate on winning the current election. If parties are not myopic, there are a few things to consider. First, they might want to affect the learning of voters in the future. One example could be that during periods when voters are quite sure that l is the optimal policy, party R might be tempted to offer policy r , rather than offer l . Party R may be elected with some small probability and potentially change the beliefs of voters. Thus, polarization may be hastened. When parties have longer horizons we might also consider collusion between parties in the long run. This analysis is beyond the scope of this paper and left for future research.

C. *Continuous Policies*

We have used a discrete model with two policies for simplicity. When policies are continuous, it is possible that parties do not fully converge to offer the same policy (unless their office or ego rents, measured by α , are sufficiently large). However, even if they do not fully converge, the same mechanisms that drive our results are

⁸We thank a referee for suggesting this alternative version of the model.

still in play with continuous policies and the polity can cycle between periods of substantial party polarization and periods of low party polarization.

Let us consider a simple example of continuous policies to illustrate one key mechanism that is behind our results. Suppose that parties choose policies $p \in [0, 1]$, and then the outcome at period t from a policy p is given by

$$(4) \quad y_t = (1 - p)\beta_l^* + p\beta_r^* + \varepsilon_t.$$

The utilities of party L and R from policy $p \in \{l, r\}$, $U^R(p)$ and $U^L(p)$, are now $U^R(p) = p$, $U^L(p) = 1 - p$. At period t , we measure polarization by the distance between parties platforms, that is, $|p_t^R - p_t^L|$. Everything else remains as in our main model. It is easy to derive from the above that the probability voters vote for L increases the higher is

$$(p_t^L - p_t^R)E[\beta_r - \beta_l|H_t].$$

From this it can already be seen, as in Lemma 1, that the harder voters find it to substantially differentiate between β_r and β_l , that is, the lower is $|E[\beta_r - \beta_l|H_t]|$, the more their vote will be based on the random shocks, ε , allowing parties to polarize more. Thus, less informed voters allow parties to substantially polarize and more informed voters induce less polarization.

To complete the intuition for cycles in our model we also need to show that the political equilibrium affects voters' knowledge. In particular, that a larger degree of party polarization $|p_t^R - p_t^L|$ will imply that voters become more informed whereas a small distance between p_t^R and p_t^L will imply that voters are less informed. In our discrete model this relation is simplified as consensus means a complete lack of variation in the data. In a model in which there is always some degree of polarization (but the degree of polarization changes), we can show a similar result.

For illustration, suppose that voters observe a political outcome in one period t^R in which some p^R is implemented, and the outcome in one period t^L in which some p^L is implemented. Then they can use

$$\begin{aligned} \Delta y &= y_{t^L}(p^L) - y_{t^R}(p^R) \\ &= (p^L - p^R)(\beta_r^* - \beta_l^*) + \varepsilon_{t^L} - \varepsilon_{t^R} = (p^L - p^R)\Delta\beta^* + \Delta\varepsilon \end{aligned}$$

to form beliefs on $\Delta\beta^* \equiv (\beta_r^* - \beta_l^*)$. Their signal extraction problem is about trying to filter out the stochastic effects of the shocks ε ; but note that these shocks obfuscate the true parameters more when $(p^L - p^R)$ is smaller. As a result, periods of low polarization will involve less learning than periods of high polarization.

VI. Related Literature

Our paper contributes to the current literature that focuses on the polarization of politics in recent decades (for a recent example see Callander and Carbajal 2022). In particular, the analysis shines a light on an inherent feature of democratic political systems that implies the recurrence of polarization phases. In this way we

complement other theories that have focused on more current trends as explanations for the recent polarization in politics.

Previous literature in political economy that analyzed political cycles focused on cycles between two types of policies rather than cycles of polarization and consensus as we have here. For example, Battaglini and Coate (2008) show how policy making in legislatures can cycle between a regime in which debt is accumulated by over-redistributing at the expense of future budgets, and a regime in which policies maximize the collective good. Transitions between these regimes arise due to dynamic equilibrium considerations; an incumbent finds it optimal running deficits when it expects future incumbents to be more prudent and vice versa. The polity cycles between “good” and “bad” policies, and the phases of the cycle coincide with electoral cycles. In our analysis we have cycles of polarization and consensus, and the cycle phases relate to voters’ knowledge and last longer than election cycles.

We provide a theoretical model of politics as a process of collective learning. Hall (1993) surveys the literature in political science that views the political process as a learning endeavor. Piketty (2020) provides a historical overview of inequality regimes and ideologies in different countries through the prism of a collective learning process. Piketty (1995) analyses a model in which individuals learn about the true data generating process, but only from their own actions and thus there is no social learning element. Strulovici (2010) and Messner and Polborn (2004) analyze group strategic experimentation and show that under-experimentation arises as individuals worry about losing their position as the median voter in society. Callander (2011) analyses a political social learning model with a focus on the dynamics of learnings when the mapping between policies and outcomes is complex.

Several recent papers analyze collective learning in political processes, when voters are behavioral to some degree. In Callander, Izzo, and Martin (2023) voters vote for parties according to their beliefs about the effect of a policy; these beliefs are derived in a non-Bayesian manner, as voters adopt one party’s interpretation of historical data that has a higher likelihood. Little (2019) studies voter learning problems in which motivated reasoning distorts beliefs.⁹ Levy, Razin, and Young (2022) analyze a dynamic political social learning model in which groups in society differ in their subjective model of the true data generating process and so one group has a misspecified model. In that model, power in society changes hands between a group that holds a complex view of the world to one that holds a simple world view. The reason is that perpetual rule by one group implies that the party in opposition becomes more intense in its preferences to win the election due to their subjective interpretation of the outcomes implemented by the ruling party. Eliaz and Spiegler (2020) and Eliaz, Spiegler, and Galperti (2022) analyze political competition when voters are not able to understand the true correlation structure between a political action and a political outcome. Hence voters may consider “false correlations” in their models of the world between some nonrelevant variable to the policy outcome. Different groups of politicians offer narratives (models) to voters where in equilibrium these models have to be consistent with the data generated when the winning groups implement policies according to them. A key result in both these papers is

⁹ See also Little (2021) and Little, Schankenberg, and Turner (2022) who show how motivated reasoning weakens politicians’ accountability.

that in some cases, the (static) equilibrium outcomes are such that groups share power, as any distribution of outcomes generated by a model advocated by a particular group, can also be explained by another model advocated by another group.¹⁰

Closer to our focus on short-term memory, Eguia and Hu (2022) assume that voters have bounded memory, and strategically design what information (signals about the state of the world) to retain in their memory (a finite automaton, as in the work of Wilson 2014). They characterize environments in which such memory choices will imply that voters will always polarize even though they observe the same information. In contrast we assume that voters naively treat past history as exogenous. Jehiel and Newman (2010) and Bhaskar and Thomas (2019) analyze social learning with bounded memory. Acemoglu and Wolitzky (2014) analyze dynamic conflicts between groups with limited memory of previous history. A sufficiently long history of a conflict allows the groups to realize that a conflict has started by mistake, and revert to a coordination phase.

VII. Conclusion

This paper investigates politics as a collective learning process. We demonstrate how short-term memory of the political players leads to cycles of polarization and consensus in party platforms. In the model it is both the endogenous policies that are chosen and exogenous shocks that impact the observed outcomes and therefore provide information about optimal policies.

In practice there might be other political institutions, internal as well as external, that might provide other sources of information to the polity. Internally, political systems often have institutions such as the bureaucracy that have both longer horizons and perhaps a longer memory. In many countries tensions between these different political players, with different incentives and memories, have been rising in the last decade. An interesting extension of the analysis would be to understand how these different institutions affect the ability of the polity to aggregate information and the dynamics of policy making.

Other information sources are more external. Countries can learn from the experiences of other countries. A good example is the COVID-19 pandemic in which both the public and politicians were able to observe and learn from the experiences of countries that were affected by the virus earlier. Similarly, in large countries or unions such as the United States, the federal government can learn from policy experiments of member states, as exemplified by the concept of “laboratories of democracies” proposed by Supreme Court Justice Louis Brandeis. The regulation of industry in the United States serves as another example, with state-level experiences influencing both legislation in other states and federal legislation (see Goldin and Libecap 2008). This paper serves as a starting point for further exploration of these issues and a deeper understanding of the study of politics as a collective learning process.

¹⁰ Azzimonti and Fernandez (2018) and Bohren and Hauser (2021) are two additional examples of social learning models in which convergence need not arise; in the former because of bots that provide misinformation that prevents learning, and in the latter due to individuals having misspecified models and hence not able to fully learn under some conditions.

REFERENCES

- Acemoglu, Daron, and Alex Wolitzky.** 2014. "Cycles of Conflict: An Economic Model." *American Economic Review* 104 (4): 1350–67.
- Achen, Charles, and Larry Bartels.** 2008. "Myopic Retrospection and Party Realignment in the Great Depression." Unpublished.
- Angelucci, Charles, and Andrea Prat.** 2023. "Is Journalistic Truth Dead? Measuring How Informed Voters Are About Political News?" Unpublished.
- Azzimonti, Marina and Marcos Fenandes.** 2018. "Social Media Networks, Fake News, and Polarisation." NBER Working Paper 24462.
- Barber, Michael, and Nolan McCarty.** 2013. "The Causes and Consequences of Polarisation." In *Negotiating Agreement in Politics*, edited by Cathie Jo Martin and Jane Mansbridge, 19–53. Washington, DC: American Political Science Association.
- Bartels, Larry.** 2008. *Unequal Democracy: The Political Economy of the New Gilded Age*. Princeton University Press.
- Bhaskar, V. and Caroline Thomas.** 2019. "Community Enforcement of Trust with Bounded Memory." *Review of Economic Studies* 86 (3): 1010–32.
- Battaglini, Marco, and Steve Coate.** 2008. "A Dynamic Theory of Public Spending, Taxation, and Debt." *American Economic Review* 98 (1): 201–36.
- Bohren, Aislinn, and Daniel Hauser.** 2021. "Learning with Heterogeneous Misspecified Models: Characterization and Robustness." *Econometrica* 89 (6): 3025–77.
- Boxell, Levy, Matthew Gentzkow, and Jesse Shapiro.** 2017. "Greater Internet Use is Not Associated with Faster Growth in Political Polarisation Among US Demographic Groups." *Proceeding of the National Academy of Science* 114 (40): 10612–17.
- Callander, Steve.** 2011. "Searching for Good Policies." *American Political Science Review* 105 (4): 643–62.
- Callander, Steve, Federica Izzo, and Greg Martin.** 2023. "Ideological Competition." *American Journal of Political Science* 67 (3): 687–700.
- Callander, Steve, and Juan Carbajal.** 2022. "Cause and Effect in Political Polarisation: A Dynamic Analysis." *Journal of Political Economy* 130 (4): 825–80.
- Calvert, Randall.** 1985. "Robustness of the Multidimensional Voting Model: Candidate Motivations, Uncertainty, and Convergence." *American Journal of Political Science* 29 (1): 69–95.
- Eguia, Jon, and Tai-Wei Hu.** 2022. "Voter Polarisation and Extremism." Unpublished.
- Eliaz, Kfir, and Ran Spiegler.** 2020. "A Model of Competing Narratives." *American Economic Review* 110 (12): 3786–3816.
- Eliaz, Kfir, Ran Spiegler, and Simone Galperti.** 2022. "False Narratives and Political Mobilization". Unpublished.
- Esponda, Ignacio, Damien Pouzo, and Yuichi Yamamoto.** 2021. "Asymptotic Behavior of Bayesian Learners with Misspecified Models." *Journal of Economic Theory* 195: 105260.
- Gerber, Alan, and Donald Green.** 1998. "Rational Learning and Partisan Attitudes." *American Journal of Political Science* 42 (3): 794–818.
- Goldin, Claudia, and Gary D. Libecap.** 2008. *The Regulated Economy: A Historical Approach to Political Economy*. University of Chicago Press.
- Hall, Peter.** 1993. "Policy Paradigms, Social Learning, and the State: The Case of Economic Policy-making in Britain." *Comparative Politics* 25 (3): 275–96.
- Healy, Andrew, and Gabriel Lenz.** 2014. "Substituting the End for the Whole: Why Voters Respond Primarily to the Election-Year Economy." *American Journal of Political Science* 58 (1): 31–47.
- Jehiel, Philip, and Andrew Newman.** 2010. "Loopholes: Social Learning and the Evolution of Contract Form." Unpublished.
- Levy, Gilat, and Ronny Razin.** 2023. "The Dynamics of Cohort Effects in Politics." Unpublished.
- Levy, Gilat, and Ronny Razin.** 2025. *Data and Code for: "Political Social Learning: Short-Term Memory and Cycles of Polarization"*. Nashville, TN: American Economic Association; distributed by Inter-university Consortium for Political and Social Research, Ann Arbor, MI. <https://doi.org/10.3886/E208145V1>.
- Levy, Gilat, Ronny Razin, and Alwyn Young.** 2022. "Misspecified Politics and the Recurrence of Populism." *American Economic Review* 112 (3): 928–62.
- Lehmann, Pola, Simon Franzmann, Denise Al-Gaddooa, Tobias Burst, Christoph Ivanusch, Sven Regel, Felicia Riethmüller, Andrea Volkens, Bernhard Weßels, and Lisa Zehnter.** 2024. The Manifesto Data Collection. Manifesto Project (MRG/CMP/MARPOR). Version 2024a. Berlin: Wissenschaftszentrum. Berlin für Sozialforschung (WZB), Göttingen: Institut für Demokratieforschung (IfDem). <https://doi.org/10.25522/manifesto.mpsds.2024a>.

- Lenz, Gabriel.** 2010. "(Mis)Remembering Incumbent Performance." Presented at the annual meeting of the Midwest Political Science Association.
- Lindbeck, Asa, and John Weibull.** 1987. "Balanced-Budget Redistribution as the Outcome of Political Competition." *Public Choice* 52: 273–97.
- Little, Andrew.** 2019. "The Distortion of Related Beliefs." *American Journal of Political Science* 63 (3): 675–89.
- Little, Andrew.** 2021. "Directional Motives and Different Priors are Observationally Equivalent." Unpublished.
- Little, Andrew, Keith Schnakenberg, and Ian Turner.** 2022. "Motivated Reasoning and Democratic Accountability." *American Political Science Review* 116 (2): 751–67.
- Malmendier, Ulrike, and Stefan Nagel.** 2016. "Learning from Inflation Expectations." *Quarterly Journal of Economics* 131 (1): 53–87.
- Manifesto Project Dataset.** 2024. "Manifesto Project Dataset (version 2024a)." Manifesto Project. <https://manifesto-project.wzb.eu/datasets> (accessed August 27, 2024).
- McCarty, Nolan.** 2019. *Polarization: What Everyone Needs to Know*. Oxford University Press.
- Messner, Matthias, and Mattias Polborn.** 2004. "Voting on Majority Rules." *Review of Economic Studies* 71 (1): 115–32.
- Ortoleva, Pietro, and Erik Snowberg.** 2015. "Overconfidence in Political Behavior." *American Economic Review* 105 (2): 504–35.
- Peltzman, S.** 1990. "How Efficient Is the Voting Market?" *Journal of Law and Economics* 33 (1): 27–63.
- Piketty, Thomas.** 2020. *Capital and Ideology*. Harvard University Press.
- Piketty, Thomas.** 1995. "Social Mobility and Redistributive Politics." *Quarterly Journal of Economics* 110 (3): 551–84.
- Rajan, G., and Raghuram.** 2009. "The Credit Crisis and Cycle-Proof Regulation." *Federal Reserve Bank of St. Louis Review* 91 (5): 397–402.
- Strulovici, Bruno.** 2010. "Learning While Voting: Determinants of Collective Experimentation." *Econometrica* 78 (3): 933–71.
- Wilson, Andrea.** 2014. "Bounded Memory and Biases in Information Processing." *Econometrica* 82 (6): 2257–94.
- Wittman, Donald.** 1983. "Candidate Motivation: A Synthesis of Alternatives." *American Political Science Review* 77 (1): 142–57.