

Political social-learning: short-term memory and cycles of polarisation

Gilat Levy and Ronny Razin, LSE¹

Abstract: In this paper we explore the effect of *short-term memory* on political outcomes in a model in which politics is viewed as a collective learning process. We analyse a dynamic model in which voters use past observations to make inferences about the true data generating process, and political parties are self-interested with polarised ideal policies. Voters balance party loyalty with a desire to vote for the party whose policy is based on a better interpretation of past observations. We show that short-term memory leads to political cycles of polarisation and consensus. A short-term history involving only periods of consensus implies little variation in voters' data, and hence less precise knowledge about the true state of the world; this allows parties to push their self interests. Alternatively, periods of polarisation imply sufficient variation which at some point allows voters to be confident about what is the true model; this forces parties to converge on the policy that fits that model. Our framework also sheds light on the relation between policy uncertainty and political polarisation, and on the effects of crises on political competition.

1 Introduction

“Social ideologies usually evolve in response to historical experience...Each nation’s political and ideological trajectory can be seen as a vast process of collective learning and historical experimentation. Conflict is inherent in the process because different social and political groups have not only different interests and aspirations but also different memories. Hence they interpret past events differently and draw from them different implications regarding the future. From such learning experiences, national consensus on certain points can nevertheless emerge, at least for a time. Though partly rational, these collective learning processes nevertheless have their limits. Nations tend to have short memories (people often forget their own country’s experiences after a few decades or else remember only scattered bits, seldom chosen at random).” Thomas Piketty, *Capital and Ideology*, page 10.

¹Department of Economics, Houghton St. WC2A 2AE, UK. Email addresses: g.levy1@lse.ac.uk, r.razin@lse.ac.uk. For useful comments we thank Steve Callander, Andrew Little, and many seminar participants and discussants. Roberto Maura-Rivero and Enrico Turri provided helpful research assistance.

To study the evolution of policy it is instructive to view politics as a process of social learning and experimentation. In his book, *Capital and Ideology*, Thomas Piketty uses social learning as a prism to analyse the different historical paths taken by different countries when constructing their property rights and tax regimes. The history of regulation in the US can similarly be viewed as a process by which politicians, bureaucrats and the electorate learn about the effects of regulation on business and consumer welfare and about which regulatory tools are more suitable.² In both of these examples, policies are initially chosen, in some cases with little information, and later experiences influence future opinions and agendas of politicians, voters and other stakeholders.

How well does the political process enable society to learn which policies are most suitable? To learn, a society needs to experiment with different policies, compare outcomes over time, and agree on what is best. Some aspects of the political process may support or alternatively might obstruct learning. Heterogenous preferences may hinder agreement about which policy to implement. But political turnover can also be useful as it encourages experimentation with different policies. Short-term incentives of voters and politicians may also hinder the ability to learn. Is social learning viable in the context of the political process? Will society reach a consensus, or will polarised world views persist in the long run? How is this affected by voters' short memories mentioned in the quote above?

The question of whether consensus arises or not is particularly pertinent given the recent wave of polarisation in Western societies. As political scientists already noticed, polarisation is not a new phenomenon; taking a step back from current polarised positions, and zooming out to a longer time perspective, a somewhat cyclical pattern between consensus and polarisation is observed. For example, polarisation of policy positions of Senators and Congress persons in the US was high in the beginning of the 20th century, declined in the 1930s, remained low until the late 1970s and has been rising ever since (see Figure 1).



Figure 1: Historical polarisation in the US Senate and Congress (McCarty 2019).

Similar cyclical patterns can be observed in political parties' stated ideology by looking

²See Goldin and Libecap (2008).

at manifestoes over time. The Manifesto Project decodes each policy dimension into a unidimensional score and tracks how it changes over time.³ Looking at the manifestos of the two US parties we see how parties oscillate between polarisation and convergence on economic issues. For example, Figure 2 plots the Democratic and Republican party positions on market regulation for the period 1948-2020. Periods of relative consensus arise around the late 40s, late 80s to early 2000s and more recently in 2020, while there is relative polarisation on this issue in all other times. Interestingly, John Williamson’s famous “Washington Consensus” paper about the standard reform package for developing countries was published in 1989, and Dani Rodrik’s response, entitled “Good bye Washington Consensus, Hello Washington confusion?...”, was published in 2006.⁴




Figure 2: US Democratic and Republican parties’ regulation policies (Democratic party in red with higher levels of regulation).

In this paper we show how a political social-learning process, coupled with short-term memory, can give rise to similar patterns of consensus and polarisation. Specifically, we show that when voters have an unbounded memory, their collective learning must give rise to a long-term consensus between parties. In contrast, when voters’ memory is short, cycles of consensus and polarisation will arise. These cycles are manifested in public opinion, in the narratives espoused by politicians and in the policies that are implemented. Crucially, it is a phase of consensus which brings about periods of polarisation and vice versa; after periods of polarisation, a consensus forms affecting voting behaviour and parties’ platforms.

We analyse a dynamic model with uncertainty about the true model generating observable outcomes. For example, society might be uncertain about whether growth is better achieved

³see <https://manifesto-project.wzb.eu/>.

⁴See Rodrik (2006). Discussing a World Bank report about the experience in the 1990s, Rodrik remarks that "In fact, it is a rather extraordinary document insofar as it shows how far we have come from the original Washington Consensus. There are no confident assertions here of what works and what doesn't-and no blueprints for policy makers to adopt. The emphasis is on the need for humility, for policy diversity, for selective and modest reforms, and for experimentation"

with policies that increase inequality in order to facilitate risky investment and job creation by the rich (a position pushed by one party), or with a more redistributive policy which aims to achieve more consumption power (a position pushed by another). Policy is determined endogenously following political competition between ideologically motivated parties each pushing a different interest. Each party’s platform includes a policy choice, and a model that justifies it by explaining the historical data. Voters observe the historical experiences and compare the different narratives that parties put forward. They then balance off their motivation to vote for a policy that is based on a model that has a better fit to the data, with their (political or economic) affinity to one of the parties. Thus, throughout time, the implemented policies are affected by voters’ estimation of the likelihood of the different political narratives, and these likelihoods are updated given the experiences from these policies.

Our results build on two key mechanisms. The first (political) mechanism relates to the effect of voters’ uncertainty on polarisation, when parties have policy preferences.⁵ In our model, when voters are not sure which model fits the data better, it is harder for them to assess which policy is better. Parties, who trade-off the probability of winning with the policies they will implement upon winning, take advantage of this and push their interests more easily. Alternatively, when voters’ knowledge about the correct policy is relatively precise, parties are “disciplined” to choose this policy and find it harder to push their own interests.

In the benchmark case when voters have unbounded memory, the above implies that society reaches a consensus: The unique long-term outcome is for both parties to offer the same policy. If parties do not offer the same policy in the long term, political turnover implies that voters’ historical data contains large variation in policies. A high level of “experimentation” implies however that the likelihood of the true model must converge to be infinitely larger than that of other models. As voters’ knowledge about the true outcome generating process becomes very precise, parties will be forced to choose the same -and optimal- policy. It therefore cannot be that parties’ platforms do not coincide in the long run.

A second (statistical) mechanism in our model implies that cycles of polarisation and consensus arise when voters have short-term memory. If there is a sequence of periods in which parties offer similar platforms, there will be little variation in policy in the short history that voters observe. Low policy variation implies a relatively uninformative history, and so voters cannot sufficiently discriminate between the different models parties may offer. As a result, at some point, parties will find it easier to pursue different models and platforms, closer to their ideal ones. In contrast, higher variation in policy over the years, due to periods of

⁵This mechanism goes back to Calvert (1985) and more recently is explored in Callander, Izzo and Martin (2021).

polarisation and political turnover, yields histories with more information about what effect policies have on outcomes. Following such periods of polarisation, society will be more in agreement about what is the best course of action. This implies that self-interested parties find it harder to push their own agendas and have to settle on a policy that accords better with the emerging consensus. The above two mechanisms, together with the endogeneity of policies, imply that cycles must arise; polarisation phases induce consensus, and vice versa.

In our model, political polarisation is closely linked to policy and economic uncertainty in two ways. First, political polarisation arises exactly when there is model uncertainty, and so little information in the economy with regard to what should be the right course of action. As we show in this paper, this uncertainty is then exacerbated by political polarisation; which policy will be implemented next is uncertain compared to when there is consensus. Our model can then shed light on recent empirical work looking at policy uncertainty and political polarisation. Specifically, Baker et al (2020) and Bloom et al (2014) show that economic policy uncertainty rises in the month leading to elections, and specifically, in the month leading to presidential elections that are close and polarized. In addition, Bloom et al (2014) show that political polarisation is a key reason for economic uncertainty.

Our framework can also be valuable to consider the effect of crises on the political system. While in the main analysis we assume a fixed outcome generating process, in Section 4.2 we discuss how polarisation increases following a change in the environment. Intuitively, such a change implies that voters' historical knowledge is less valuable, increasing the inclination of parties to polarise.⁶

Finally, while in our main analysis we assume that all voters are exposed to the same history, in Section 5 we discuss other possibilities. For example, cohort effects may imply that voters experience different events in their formative years, the years on which they base their knowledge and beliefs.⁷ We also illustrate how echo chambers, that expose different groups of voters to different selective memories and political narratives, increase political polarisation.

⁶Some types of crises indeed arise when the state of the world changes; for example, the 2008 financial crisis was partially driven by new technologies of financial derivatives and trades. Our result that polarisation follows crisis is consistent with the results of Mian et al (2014) and Funke et al (2016) who both show that voters become more polarised following a financial crisis.

⁷Malmendier 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.

2 Related literature

In this paper we provide a theoretical model of politics as a process of social learning. Piketty (2020) provides a comprehensive 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. Little (2019) studies voter learning problems in which motivated reasoning distorts beliefs.⁸ Strulovici (2010) and Messner and Polborn (2004) analyse 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. Levy, Razin and Young (forthcoming) analyse a political social-learning model in which groups in society differ in their subjective model of the true data generating process.

In our model parties compete over votes by offering different models/narratives to interpret historical data. Eliaz and Spiegler (2021) present a static model of political competition based on competing narratives that draw voters' attention to different causal variables and mechanisms. Callander, Izzo and Martin (2021) analyse a static model of political competition over narratives. They show that in equilibrium parties always propose different narratives. Our analysis differs as we focus on the dynamic implications of electoral competition over narratives.

Our paper proposes a new theory of political cycles. In our theory, cycles arise due to the nature of politics as a collective learning process, coupled with the short-term memory of the public.⁹ Rogoff (1990) proposes a theory of political cycles that arise as a result of the incentives of politicians in election years.¹⁰ Alesina and Rosenthal (1995) focus on the US political system and show how political cycles can arise when voters use midterm elections to tame polarised presidents by splitting their votes. Battaglini and Coate (2008) show

⁸In his model a voter trades off the belief that is most likely to explain her (exogenous) set of observations, with the belief that justifies her preferred ideological outcome. In our model a similar feature arises as voters compare the narratives offered by ideologically polarised parties on the basis of the likelihood of these models given the observed history. See also Little (2021), and Little, Schankenberg and Turner (2020) who show how motivated reasoning weakens politicians' accountability.

⁹Since Schlesinger (1949), scholars have considered cycles of different policies, e.g., conservative and liberal, whereas we focus on cycles of consensus and polarisation. See also Schlesinger (1999).

¹⁰Relatedly, most of the empirical literature on voters' short term memory has focused on the question of whether voters respond more to outcomes that arise in election year or consider a longer set of previous outcomes. See for example Bechtel and Hainmueller (2011), Healy and Lenz (2014) and Achen and Bartels (2014).

how policy making can cycle between a regime in which legislators accumulate debt by over redistributing at the expense of future budgets, and a regime in which policies maximize the collective good. In Levy, Razin and Young (forthcoming), the polity converges to have cycles between groups that hold a complex view of the world to those that hold a simple world view. The intuition is that perpetual rule by one party 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.¹¹ Wolitzky and Acemoglu (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.

3 The model

We first describe the economic environment, which is a simple mapping between policies and outcomes. There are two policy dimensions, l and r . We assume a discrete set of policies, to simplify the exposition. In particular, at every period t , parties can offer one of the following three policy vectors (denoting the direction of policy in dimensions l and r):

$$L = (1, 0), \quad M = \left(\frac{1}{2}, \frac{1}{2}\right), \quad R = (0, 1).$$

The above represents three policy regimes, policy L is biased towards policy dimension l , policy R is biased towards r and a compromise policy, M , that invests in both dimensions.¹²

The common outcome, y_t , given some policy $p \in \{L, M, R\}$ is determined by,

$$y_t = \begin{cases} \beta_l^* + \varepsilon_t & \text{if } p = L \\ \beta_r^* + \varepsilon_t & \text{if } p = R \\ \frac{1}{2}\beta_l^* + \frac{1}{2}\beta_r^* + \varepsilon_t & \text{if } p = M \end{cases}$$

where ε_t is iid across time and normally distributed with zero mean and variance σ^2 .

Voters understand how the data generating process depends on parameters $\beta = (\beta_l, \beta_r)$, in a set $B = [0, \bar{\beta}]^2$, for some $\bar{\beta} > 0$, but do not know the true value of these parameters, $\beta^* = (\beta_l^*, \beta_r^*) \in B$.¹³

¹¹Azzimonti and Fernandez (2018) and Bohren and Hauser (forthcoming) are two additional examples of social learning models in which convergence need not arise; in the former because of bots that provide misinformation, and in the latter due to individuals having misspecified models and hence not able to fully learn under some conditions.

¹²Our results can be extended to consider continuous policies, see footnote 21.

¹³In the main part of the analysis we consider a fixed β^* , while in Section 4.2 we discuss how polarisation is exacerbated when the environment shifts (that is, when β^* changes over time).

To make the model interesting, we consider a rich enough set B , so that each of the three policies $p \in \{L, M, R\}$ can be “justified” by some model β . Let $E[y|p, \beta]$ denote the expected outcome y given some policy $p \in \{L, M, R\}$ and a model $\beta = (\beta_l, \beta_r)$ (that is, absent the noise ε). Specifically, and in accordance with the outcome function specified above, $E(y|L, \beta) = \beta_l$, $E(y|R, \beta) = \beta_r$, and $E(y|M, \beta) = \frac{1}{2}(\beta_l + \beta_r)$. We assume that policies are costly, with $c(M) < c(L) = c(R)$.¹⁴ Let then B_p denote the subset of models β for which policy p is optimal. Thus, for example for policy M :

$$B_M = \{\beta \in B | E(y|M, \beta) - c(M) \geq \max\{E(y|L, \beta) - c(L), E(y|R, \beta) - c(R)\}\}.$$

It is then easy to see that given some cost function as above, when $\bar{\beta}$ is large enough, $B_p \neq \emptyset$ for any $p \in \{L, M, R\}$. Specifically, as detailed in Figure 3 below, models with a high enough β_l and low enough β_r rationalise L , those with a high enough β_r and a low enough β_l rationalise R , and those where β_l and β_r are sufficiently similar rationalise M . For concreteness, we assume that $\beta^* \in \text{interior}(B_M)$, and so M is the optimal policy.¹⁵

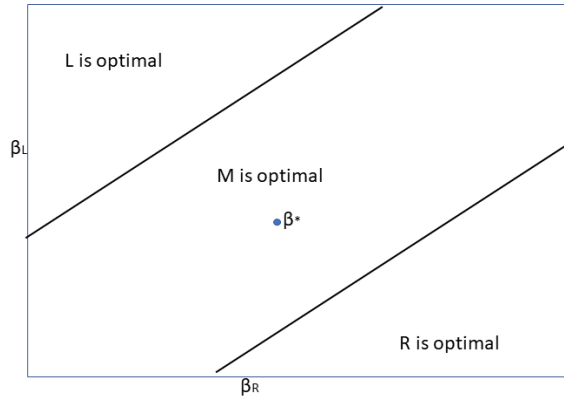


Figure 3: The set of models and their respective optimal policies

3.1 Political parties and electoral competition

There are two parties, each identified with a special interest on a different policy dimension. Party L prefers investment in policy dimension l and party R prefers investment in policy dimension r . The utilities of party L and R from some policy $p \in \{L, M, R\}$, $U_L(p)$ and

¹⁴This is a standard assumption of convex costs. Note that otherwise, given the simple linear outcome function, M cannot maximise y_t .

¹⁵Our results generalise to the cases when $\beta^* \in \text{interior}(B_p)$ for $p \in \{L, R\}$.

$U_L(p)$, are therefore:

$$\begin{aligned} U_R(R) &= 1, U_R(M) = \frac{1}{2}, U_R(L) = 0; \\ U_L(p) &= 1 - U_R(p). \end{aligned}$$

In addition, parties enjoy some small office-rents α when they win the election. Thus, given an election with a policy outcome p , party J 's utility is $U_J(p) + I_J\alpha$, where $I_J = 1$ if party J won the election and 0 otherwise.

During the campaign, at any period t , each party $J \in \{L, R\}$ chooses a policy vector $p_t^J \in \{L, M, R\}$ and a model $\beta_t^J \in B_{p_t^J}$ to justify their choice of policy.

3.2 Histories and Voting

At each period t , the voters observe data from only the last K periods. In particular denote the history observed by voters at period t by $H_t = (p_\tau, y_\tau)_{\tau=t-K}^{\tau=t-1}$ where $p_\tau \in \{L, M, R\}$ is the implemented policy in period τ and y_τ is the policy outcome in that period.

Given the information they have, voters are inclined to vote for the party that offers models that are more likely to explain past observations. This can arise from some moral obligation, a need to justify decisions to oneself or others, or a simple motivation to seek the correct policy. Different models β will rationalise the implementation of different policies. Thus, voters will use historical observations to calculate the likelihoods of the different models espoused by parties, β_t^J .

In our main model all voters observe the same information and hence agree on the likelihood of the models that parties offer.¹⁶ Given a history $H_t = (p_\tau, y_\tau)_{\tau=t-K}^{\tau=t-1}$ and a model β , the likelihood of β is given by

$$\log \mathcal{L}(\beta|H_t) = \log \prod_{\tau=t-K}^{t-1} f(y_\tau - E[y|p_\tau, \beta])$$

where $f(\cdot)$ is the (normal) density of the shock ε .

Voters may differ on additional dimensions though. For example, voters may have economic interests in line with those of the parties, or they may be attached to particular parties. We summarise all these by assuming a median voter, with a bias towards L denoted by ϕ . The median voter votes for party L if

$$\log \mathcal{L}(\beta_t^L|H_t) - \log \mathcal{L}(\beta_t^R|H_t) + \phi > 0,$$

where ϕ is uniformly distributed on $[-\frac{1}{2\zeta}, \frac{1}{2\zeta}]$.¹⁷ Voters balance then their bias in favour of

¹⁶We discuss other possibilities in Section 5.

¹⁷In case of an equality, WLOG we assume that the voter votes for party L with probability 0.5.

one of the parties with their desire to elect the party that pushes forwards a model that has a better fit with the data.¹⁸

3.3 Learning from a no-variation history

We make one key assumption: When a *very long* history contains no variation in policy (and so learning is diffused), any policy can be rationalised by some model that can interpret past events. Specifically, for any $p \in \{L, M, R\}$:

$$E[y|p, \beta^*] = E[y|p, \beta] \text{ for some } \beta \in B_{p'} \text{ for any } p' \in \{L, M, R\}.$$

Thus, if some policy vector p is implemented repeatedly, then it is possible to explain the average output generated by it, $E[y|p, \beta^*]$, using a model β that will rationalise any rival policy p' .¹⁹ For example, given $E[y|M, \beta^*] = \frac{1}{2}\beta_l^* + \frac{1}{2}\beta_r^*$, we can find $\beta \in B_R, \beta \in B_M$ and $\beta \in B_L$ that satisfy $\frac{1}{2}\beta_l^* + \frac{1}{2}\beta_r^* = \frac{1}{2}\beta_l + \frac{1}{2}\beta_r$. Or, given that $E[y|L, \beta^*] = \beta_l^*$, we can find vectors with $\beta_l = \beta_l^*$ in B_L, B_M and B_R . This is illustrated in Figure 4 below. This assumption is readily derived if we let the components of β^* be greater than the value $c(L) - c(M)$, together, as above, with a high enough $\bar{\beta}$.

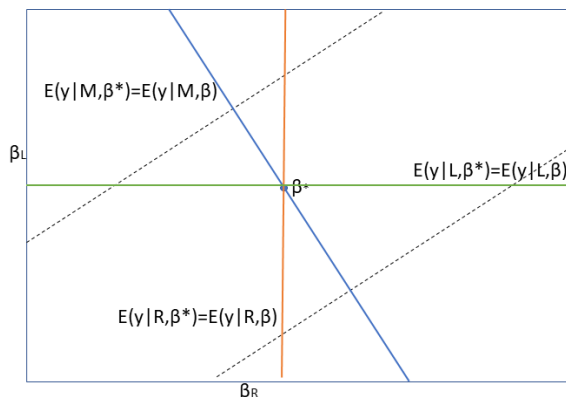


Figure 4: Long term data with no policy variation implies that each policy can be optimal

¹⁸In our model voters consider which policy is rationalised by a more likely explanation, as we focus on social learning with political narratives. Alternatively, a fully rational voter (albeit with a bounded memory), may instead compute her expected utility from each possible policy, given her full set of beliefs. Our result is more stark with the behavioural model presented above (see Callander et al 2021 for a similar assumption), but can also be derived in a rational model with some additional assumptions on voters' utilities.

¹⁹The condition is somewhat stronger than what we need but presented as such for clarity.

3.4 Dynamics

We are now ready to put all the ingredients of the model together. The dynamic model is defined as follows:

1. There is some initial history H_0 .
2. In period t , party J that won the election implements $p_t^J \in \{L, M, R\}$.
3. Given y_t , history evolves from $H_t = \{p_\tau, y_\tau\}_{\tau=t-K}^{t-1}$ to $H_{t+1} = \{p_\tau, y_\tau\}_{\tau=t-K+1}^t$.
4. The two parties offer $\{p_{t+1}^J, \beta_{t+1}^J\}$, where $\beta_{t+1}^J \in B_{p_{t+1}^J}$.
5. In period $t + 1$, ϕ is drawn and party L wins the election if

$$\log \mathcal{L}(\beta_{t+1}^L | H_{t+1}) - \log \mathcal{L}(\beta_{t+1}^R | H_{t+1}) + \phi > 0$$

or with probability 0.5 if the above is satisfied with equality.

In equilibrium, in any period, voters calculate the likelihood of each model espoused by the parties given the observed K -period history and vote accordingly. Parties are best responding to each other, while anticipating the voters' behaviour. It is easy to see that given the voters' choice rule, and given policy choices p_t^J , each party J will choose a model that maximises the likelihood of H_t within the set of models that justify p_t^J , $B_{p_t^J}$. That is,

$$\beta_t^J \in \arg \max_{\beta \in B_{p_t^J}} \mathcal{L}(\beta | H_t).$$

4 Cycles of consensus and polarisation

In this Section we analyse the dynamic model presented above. Our focus is on whether parties offer the same policy or different ones. In particular, we distinguish between two per-period electoral competition outcomes: A polarisation outcome in which parties offer different policies, and a consensus outcome in which both parties offer the same policy. Our first, preliminary, result highlights the mechanism by which the level of information in historical data affects electoral competition. Fix a history H_t that voters observe at time t and consider the one-period political competition game that ensues. For Lemma 1 below, let \hat{p} be a policy that can be justified by a model that maximises the likelihood of H_t , that is, there exists $\hat{\beta} \in B_{\hat{p}}$, such that $\hat{\beta} \in \arg \max_{\beta \in \mathbf{B}} L(\beta | H)$.

Lemma 1 (Consensus vs Polarisation): *At period t , (i) If \hat{p} is an ideal policy of party J , in equilibrium party J offers \hat{p} . (ii) High relative likelihood implies convergence: If*

$$(*) \log \frac{\mathcal{L}(\hat{\beta} | H_t)}{\mathcal{L}(\beta' | H_t)} > \frac{1}{2\zeta} - 2\zeta \frac{\alpha}{1 + \alpha} \text{ for all } \beta' \in B_{p'}, \text{ for all } p' \neq \hat{p},$$

then the unique equilibrium at period t involves consensus on \hat{p} . (iii) Similar likelihoods

induce polarisation: If

$$(**) \log \frac{\mathcal{L}(\hat{\beta}|H_t)}{\mathcal{L}(\beta'|H_t)} < \frac{1}{2\zeta} - 2\zeta \frac{\alpha}{0.5 + \alpha} \text{ for some } \beta' \in B_{p'}, \text{ for some } p' \neq \hat{p},$$

then parties polarise (e.g., choose different policies).

The lemma connects the likelihoods of different models, gleaned from historical data, with the incentives of parties to polarise or reach a consensus. Generally speaking, parties prefer to pursue their own interests, but to successfully do so, must be elected with some probability. This implies that they may be disciplined by voters to choose the policy that is more likely to be supported by the model that best explains the historical data. If the historical data makes voters sufficiently confident about a potential model as the correct one, i.e., the log-likelihood ratio of this model compared to others is sufficiently high, then parties have to offer this policy, and hence reach a consensus. If they offer any other policy, they will neither serve their policy interest nor their office motivation, as they will face only a slim probability of being elected. Alternatively, if the historical data does not sufficiently discriminate between different models, then parties can afford to offer platforms that better serve their own policy interests. As a result, parties polarise.

Given the above Lemma, we can move to the analysis of the dynamic model, in which history is endogenous. In the following, let $\hat{\eta}_t(p)$ be the fraction of time in the history up to time t that policy p was implemented. In addition, we denote by $\eta_t(\textit{polarisation})$ the fraction of time in the history up to time t that the two parties offered different platforms and by $\eta_t(\textit{consensus}) = 1 - \eta_t(\textit{polarisation})$ the fraction of time in the history up to time t that the two parties offered the same platform. Note that the dynamic evolution of policies involves some randomness, given the attachment shock, ϕ , and the policy shock, ε (through the latter's effect on beliefs). This then induces a probability distribution P over the set of possible infinite histories \mathbb{H} . Thus, when we write “almost surely”, here and in the Appendix, we mean P -almost surely on \mathbb{H} .

As a benchmark we first consider the case in which the history that voters remember is unlimited, i.e., when $K = \infty$. Our result shows that with full memory, the two parties will converge to offer the same platform.

Proposition 1: *Assume that $K = \infty$. Then almost surely: (i) $\eta_t(\textit{consensus}) \rightarrow 1$. (ii) There exists a $p \in \{L, M, R\}$ such that $\hat{\eta}_t(p) \rightarrow 1$.*

To see the intuition for Proposition 1, note that polarisation cannot be part of a long-term equilibrium. If this were the case, given political turnover, polarisation implies experimentation with at least two different policies that will allow the voters to learn the truth. For

example, if M and L are observed in the long term, voters' beliefs will concentrate on models β that satisfy both $E[y|M, \beta^*] = E[y|M, \beta]$ and $E[y|L, \beta^*] = E[y|L, \beta]$. As can be seen from Figure 4, this then implies that the (true) model β^* will become infinitely more likely compared to other models. This, by Lemma 1(ii), induces parties to converge on M , in contradiction to our initial supposition of long-term polarisation. Thus, when $K = \infty$, voters' beliefs in the long term must allow for only one policy to be properly justified, and parties must converge to reach a consensus on this policy.

We note that convergence is not guaranteed to be on the optimal policy M ; as in any learning problem with myopic agents (or more generally with discount factors strictly smaller than one), learning can sometimes be wrong due to insufficient experimentation. For example, a rare series of very good shocks early on, let's say when L is implemented, may convince voters that β_l is very high, and parties may then stick to offering L .

4.1 Short-term memory and cycles

We now turn to consider short-term memory, i.e., finite K . With short-term memory, voters will never fully learn the state of the world. Moreover, 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 very little variation in policies. Alternatively when the history involves a high frequency of changes in policies, voters' data will be relatively informative.

Short-term memory implies therefore a potential for cycles of policy change. Consider first the possibility that parties polarise. Following some history of political polarisation (and turnover), if voters' memory is long enough, their data will contain a lot of information and will therefore discriminate between models, putting relatively higher likelihoods on models that are more in line with the true model. In this case, at some point, parties will be drawn to put forward models that are closer to the truth and hence reach a consensus. But now short-term memory kicks in; at some point voters may observe a recent history entailing only consensus, and hence with little variation in policy. This history is not very informative about the true data generating process. As a result different models will yield comparable likelihoods, implying a lower electoral cost for parties when they choose models that support their extreme ideal policies. This then paves the way for polarisation, and so on.

Our next result formalises the intuition above.

Proposition 2: *For low enough ζ and a large enough K , almost surely:*

- (i) $\liminf_{t \rightarrow \infty} \eta_t(\text{polarisation}) > 0$ and $\liminf_{t \rightarrow \infty} \eta_t(\text{consensus}) > 0$.
- (ii) $\liminf_{t \rightarrow \infty} \hat{\eta}_t(p) > 0$ and $\liminf_{t \rightarrow \infty} \hat{\eta}_t(p') > 0$ for at least two policies $p \neq p'$.

Our key result is that the endogeneity of policies, together with short-term memory, leads

to cycles of polarisation and consensus. As stated in (i), society will forever cycle between periods of consensus and periods of polarisation. In periods in which parties offer different policies, as each must be elected with some positive probability, over time there will be variation in policies. When K is large enough this insures that voters become confident about what is the correct model, fostering consensus.²⁰ However, following periods in which there is consensus, short-term memory implies that voters at some point will only remember consensus periods. In this case historical data has little policy experimentation allowing parties to push their agendas. Part (ii) illustrates the implication of our result to implemented policies; society never converges on one policy, in contrast with Proposition 1.²¹

Remark 1 (Policy uncertainty and polarisation): In our model, political polarisation is closely linked to policy and economic uncertainty in two ways. First, political polarisation arises exactly when there is model uncertainty, and so little information in the economy with regard to what should be the right course of action. This uncertainty is then exacerbated by political polarisation as which policy will be implemented is uncertain in a phase of polarisation, compared to a consensus phase. Our model can then rationalise recent empirical work looking at policy uncertainty and political polarisation. Specifically, Baker et al (2020) and Bloom et al (2014) show that economic policy uncertainty (EPU) rises in the month leading to presidential elections that are close and polarised, compared to elections that are neither. In addition, Bloom et al (2014) show that political polarisation is a key reason for higher EPU.

Proposition 2 formally establishes that the shares of time that the polity is in a polarisation phase and a consensus phase are substantial. We now explore in more detail the dynamics of political cycles and the length of such phases. First, we fully characterise the equilibria in the model when the variance of noise is small. Naturally, short memory and the stochastic nature of the model (the noise ε and the affinity shock ϕ) imply that many patterns of behaviour may arise in equilibrium. As we consider the case where the variance of the shock ε , σ^2 , goes to zero, we illustrate how cycles arise even without the randomness induced by policy noise:

Proposition 3: *Let $\sigma^2 \rightarrow 0$. Then, for $K \geq 2$, equilibrium dynamics in the limit are*

²⁰Convergence is more likely to be on the right policy, but as before, parties may wrongly converge to offer a different policy, depending on the previous sequence of shocks.

²¹In this paper we simplified our model by focusing on a discrete set of policies. Our result generalises to the case of continuous policies. With continuous policies, parties will never choose the same policy; they always have an incentive to polarise slightly. Still, when the history is very informative, parties will have to choose policies that are sufficiently close to each other. Close policies and finite memory will imply little variation and little learning as we have in our simple three-policy model.

characterised by: (i) A consensus phase, lasting exactly K periods, in which both parties choose M . (ii) A polarisation phase, lasting exactly one period, where each party offers its ideal policy.

The case of low variance allows us to identify the systemic reason for cycles. Once parties moderate, the lack of variation in policies after K periods implies a complete “reset” in learning. Given the low variance, beliefs will concentrate on all models β satisfying $E[y|M, \beta^*] = E[y|M, \beta]$, which, as Figure 4 illustrates, can justify all policies. This will allow parties to polarise. However, when the variance is very small, learning is very fast: In the limit observing two different policies is enough to fully learn the state of the world. Thus polarisation is short-lived as voters can remember at least two different policies that were implemented, M and either L or R . This will entail perfect learning on β^* and therefore convergence of both parties to offering the correct policy M . And so on.

While this version of the model is stark, it highlights that the key assumption in our analysis is the short-term memory. It also shows that the length of the polarisation and convergence phases depend on the interaction between K and σ^2 ; cycles arise also for a small K , if the variance of the shock is small enough.

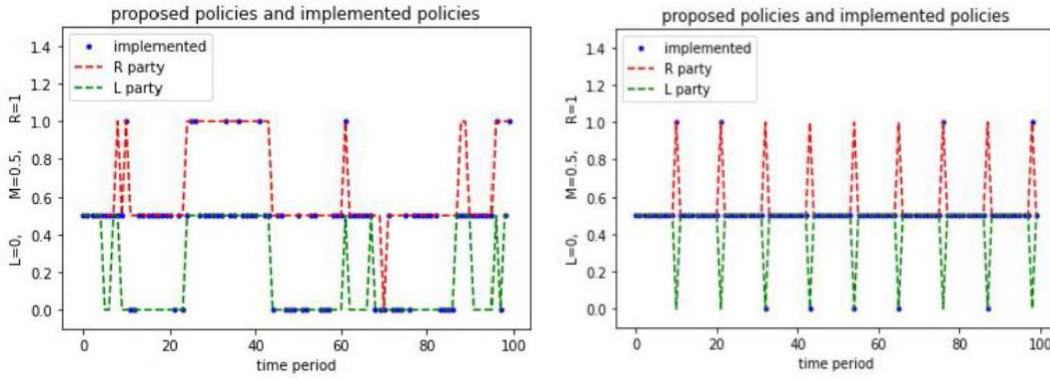
To further illustrate the relation of the parameters of our model to the length of the different polarisation and consensus phases, we report the results of a simulation that explores the effect of both K and σ^2 .²² Table 1 summarises averages of statistics from a simulation of hundred-period models for ten different draws of initial history, and for different values of the noise variance, σ^2 , and length of memory, K . For each configuration of parameter values the table presents the average proportion of periods in which there was consensus, $\eta(\text{consensus})$, the average length of the consensus phase, and the proportion of time in which consensus was actually on the correct policy. As can be seen in the table, all three statistics are inversely related to σ^2 and are increasing in K .

In addition, Figures 5a and 5b show the dynamics of party behaviour; in Figure 5a, we plot outcomes for $K = 10$ and $\sigma^2 = 1.3$, and in Figure 5b we illustrate outcomes of polarisation and convergence for $K = 10$ and $\sigma^2 = 0.3$. As can be seen in the figures, as the variance of the shock decreases, the cycles of polarisation and consensus become more regulated and predictable, with the length of polarisation phases shrinking to one and the length of consensus phases increasing to K .

²²In the simulation we use $\zeta = 1, \beta^* = (3, 3), \bar{\beta} = 6, c(M) = 0, c(L) = 1$.

σ^2 / K	10	3
1.8	$\eta(\text{consensus}) = 23\%$ Av length consensus = 2.7 consensus=M 90.42%	$\eta(\text{consensus}) = 7.8\%$ Av length consensus = 1.29 consensus=M 0%
0.8	$\eta(\text{consensus}) = 45.3\%$ Av length consensus = 4.14 consensus=M 97.8%	$\eta(\text{consensus}) = 15.2\%$ Av length consensus = 1.39 consensus=M 82.91%
0.3	$\eta(\text{consensus}) = 88.5\%$ Av length consensus = 9.26 consensus=M 100%	$\eta(\text{consensus}) = 69.8\%$ Av length consensus = 2.81 consensus=M 100%

Table 1: Comparative statics with respect to σ^2 and K



On the left, Figure 5a simulates parties' platform choices over a 100 periods for $K = 10$ and $\sigma = 1.3$. Figure 5b on the right uses $K = 10$ and $\sigma = 0.3$.

Remark 2 (Different policy dimensions): Our political competition model focused on one policy dimension. Different policy dimensions might have different attributes that affect the nature of cycles that will arise, as manifested by our parameters σ^2 and/or K . One example is how fast information about policy outcomes is observed. For example while statistics about crime rates are easy to understand and can be published relatively quickly following policy reforms, the effect of redistribution schemes on inequality of opportunities in the economy has a slower clock. Another example is the ease with which policy changes can be made. While some policy reforms are easy to implement, on some dimensions policies may be more sticky.

To conclude this subsection, note that information about the state of the world is only gathered in our model from endogenous outcomes. In some cases, the public may receive external signals about the true data generating process for example by observing other polities.

Although learning from others is more nuanced as there are of course differences across polities, one would expect some learning in these cases. In our model, such exogenous learning will enhance the force of moderation and consensus in society as it increases the knowledge of voters. For example, the recent Covid-19 global pandemic allowed the public and politicians to learn from the experiences of those in other countries. Another example is how US states learn from each other. In 1932, US Supreme Court justice Louis Brandeis coined the term “Laboratories of democracies”, when he wrote: “A single courageous State may, if its citizens choose, serve as a laboratory; and try novel social and economic experiments without risk to the rest of the country.”²³ The history of industry regulation in the US is a good example of this; it started in the late 1800s at the state level and the experiences of different states affected both the legislation in other states as well as later legislation at the federal level.²⁴

4.2 Crises and changing worlds

Our analysis so far focused on a fixed technology that the polity is learning about. However, sometimes polities might experience a change in such technologies which can be triggered by both external factors (e.g., a war, a pandemic) as well as by endogenous factors such as technological changes. Changes to the true state of the world can also arise from implementation of wrong policies, as in the case of climate change. In such situations, and especially with suboptimal policies that are not tailored to the new technology, a crisis can arise.

Take for example the 2008 financial crisis; following the onset of the crisis, investors and governments realised that the effects of financial innovations have not been fully understood and hence policies were not properly tailored to the evolving technologies. As a result, and after some time, investment banks, governments and economists had to change their old models in favour of new ones. Old models and empirical analysis that relied on many years of data were deemed less relevant.

Our model can be easily adapted to better understand the effects of such crises on politics. A simple way to accommodate a different data generating process is to assume that at each period t , with probability λ , nature draws a new set of parameters β_t^* according to some distribution on the set B . For simplicity, let us consider the situation in which voters are aware once the state of the world has changed at some period t (or alternatively one can consider some delay in understanding such change). After a change happened, the history up to that period becomes uninformative about the future. In the language of our model, this means that $K = 0$ at period t . Once memory is reset, parties will then polarise. Therefore

²³In *New State Ice Co. v. Liebmann*, 285 U.S. 262 (1932).

²⁴See Goldin, C. and G. D. Libecap (2008).

when voters understand that the world has changed, political polarisation will follow.²⁵

The implications of the discussion above can shed light on recent empirical literature on crises and economic and political uncertainty. Mian et al (2014) show that financial crises lead to political polarisation. Funke et al (2016) show that financial crises create more political polarisation (mainly right-wing extremism) compared to normal recessions, and that these effects diminish after around ten years. Bloom (2014) details how financial crises and recessions lead to uncertainty, and that political uncertainty following the great recession has amplified and propagated the slow recovery. In our model, a change in the state of the world has a direct effect on learning and naturally increases uncertainty.²⁶ Moreover, this uncertainty implies polarisation, which in itself increases policy uncertainty as party and policy turnover is relatively high. Finally, such uncertainty indeed subsides at some point when individuals learn sufficiently.

5 Discussion: Different and selective memories

In this Section we discuss an additional feature of our initial quote in the introduction, the possibility that groups in societies have different memories and its implications to the prevalence of consensus and polarisation. We now discuss two such possibilities. One relies on a natural process of different memories, and another on a strategic process in which politicians instill selective memories in different groups.

5.1 Echo chambers and selective narratives

We now consider the effect of echo chambers on polarisation cycles. Within our model, if voters observe information within echo chambers, then it is possible that each group is exclusively exposed to information from one party and therefore develops a unique narrative of history. In general, when groups of voters are exposed only to information from their own respective parties, polarisation is more likely to arise (see for example Levy et al 2021). We consider a simple extension in which voters observe information through the filter of parties. Mainly, we allow different parties to highlight different periods and so select the history it offers to its voters.

²⁵ Another option to analyze is one in which voters are not aware when technology has changed. To analyze this situation there are two ways to proceed. First, one can assume that voters are aware of the fact that the state of the world might change and try to fit a model that includes the timing of these changes. The second option would be to assume that voters are not aware at all that the state might change and so always try to fit one model to analyze the history they observe. These are two interesting research questions that are beyond the scope of this paper.

²⁶ See Orlik and Veldkamp (2014) for a model looking at the relation between model uncertainty and crises.

Assume that there are two groups of voters, one attached to party L and one to R . Assume for simplicity that the two groups of voters are equally sized, and that each voter either votes for the party that represents her or does not vote at all. In addition, assume that party J is able to select a set of k^J periods, $H_t^{k^J}$, a subset of the recent H_t , to show its voters (and has no communication with the opposite set of voters). Using some costly voting model, a voter in group L is more likely to vote for party L the more she is able to justify its policies using the history $H_t^{k^L}$:

$$\log \mathcal{L}(\beta_t^L | H_t^{k^L}) + \gamma_L,$$

where $\gamma_L > 0$ is some attachment parameter of the voter to party L . Similarly, a voter in group R is more likely to vote for party R the more she is able to justify its policies using the history $H_t^{k^R}$:

$$\log \mathcal{L}(\beta_t^R | H_t^{k^R}) + \gamma_R,$$

where $\gamma_R > 0$ is some attachment parameter of the voter to party R . And so now party L wins the election if its share of voters is larger, which means whenever:

$$\log \mathcal{L}(\beta_t^L | H_t^{k^L}) - \log \mathcal{L}(\beta_t^R | H_t^{k^R}) + \phi > 0$$

Letting $\phi \equiv \gamma_L - \gamma_R$, this allows us to easily extend our framework to consider selective histories.²⁷ It is easy to see that the above modification of our model implies that polarisation will be more pronounced the more flexibility parties have in their “pick and mix” of historical periods. Thus, in line with other research and common wisdom, here too echo chambers increase the prevalence of polarisation.

Remark 3: (Nostalgia and anti-nostalgia narratives): Politicians who strategically manipulate memory may find it easier to do so when they use familiar tropes about the “good old times” or alternatively about the fact that “the world has changed”. Nostalgia is often used in political debates and is prevalent amongst voters.²⁸ Politicians can highlight particular periods and rekindle memory of “good times”. For example, one party can look back at history and shed light on periods in which they implemented their ideal policy and had obtained good outcomes. Within our model this arises when a party puts large weight on periods in which the shock ε was sufficiently high and attributes these events to a high parameter value relating to their desired policy. Alternatively, there are also narratives that are used to ignore specific histories. For example, if recent histories are not favourable to one party’s desirable outcomes, then a possible narrative is to state that

²⁷Note that in this way we also introduce turnout decisions into our model (with or without echo chambers).

²⁸There is a recent literature that studies nostalgic memory in politics. See Kenny (2017) for a theoretical discussion of the use of nostalgia by politicians and Elçi (2021) and Stefaniak et al (2021) for empirical papers showing evidence for nostalgic narratives as predictors of political attitudes.

the data generating process is changing and hence voters need to forget these periods. As opposed to our discussion in Section 4.2, such narratives can be attempted by politicians even if the state of the world had not changed.

5.2 Cohort effects

Voters' beliefs may not necessarily be shaped by the most recent K periods, but by the periods that consist their formative years. Such periods may affect them disproportionately compared to others. For example, 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. This can then translate into policy making; Malmendier et al (2021) show how personal experiences of inflation strongly influence the hawkish or dovish leanings of central bankers. Aksoy et al (2020) find that epidemic exposure in an individual's impressionable years (ages 18 to 25) has a persistent negative effect on confidence in political institutions, leaders, and public health systems, suggesting that the Covid-19 pandemic may leave behind a long-lasting political scar on the current young generation.

An interesting extension of our model, within an overlapping generation framework, is to see how cohort effects alter the beliefs of different generations who experience different formative events. Cohort effects might imply then different patterns of cycles of consensus and polarisation among younger and older generations.

6 Conclusion

We analyse a model of a collective learning process in which we show how short-term memory implies cycles of polarisation and consensus in party platforms and public attitudes. Our model is easily extended to analyse how politicians can use selective histories to justify their policies such as the use of nostalgia and anti nostalgia in political discourse. Alternatively, the model can be extended to analyze which features of the political system will hasten or delay the start of the polarisation phase.

Our paper contributes to the current literature that focuses on the polarisation of politics in recent decades. In particular, the analysis shines a light on an inherent feature of democratic political systems that implies the recurrence of polarisation phases. In this way we hope to complement other theories that have focused on more current trends as explanations for the recent polarisation in politics.

7 Appendix

Proof of Lemma 1: For (i) the proof is straightforward and is similar to the intuition in Calvert (1985). It arises as offering one's policy maximises the probability of winning with no compromise over policy. To see (ii) assume both parties offer \hat{p} . Recall that parties' utility from the election and an outcome p is $U_J(p) + I_J\alpha$, where $I_J = 1$ if J wins the election. Suppose that $\hat{p} = R$. Party R will always offer it by (i). Party L will not deviate to offer L , if:

$$0.5\alpha > \Pr(L \text{ is elected})(1 + \alpha)$$

and not to M if

$$0.5\alpha > \Pr(M \text{ is elected})(0.5 + \alpha)$$

both of which are satisfied by condition (*) as $\Pr(J \text{ is elected}) = 1 - F(\log \frac{\mathcal{L}(\hat{\beta}|H_i)}{\mathcal{L}(\beta'|H_i)})$, where β' justifies J , and β justifies R . The same argument applies when $\hat{p} = L$.

If $\hat{p} = M$ and one party offers M , the other party does not deviate to offering its ideal policy $J \in \{L, R\}$ if

$$0.5\alpha > \Pr(J \text{ is elected})(0.5 + \alpha)$$

which is satisfied by condition (*) and therefore consensus on \hat{p} is an equilibrium. We need to make sure that in this case there is no other equilibrium, that is an equilibrium in which each party J offers its ideal policy. In this case one of the parties, say L , wins with at most probability half. As a result party L expected utility from offering L is at most $\frac{1}{2}(1 + \alpha)$. If on the other hand it switches to M , its expected utility is at least $[1 - F(-\frac{1}{2\zeta} + 2\zeta\frac{\alpha}{1+\alpha})](\frac{1}{2} + \alpha)$, given (*). But given the uniform distribution on ϕ , the condition holds for any α .

To see (iii), note that if (**) holds, then given the above, no convergence on \hat{p} can arise. Moreover, there cannot be a convergence on p' , as in this case surely some party will prefer policy \hat{p} in terms of its preferences and \hat{p} provides a higher probability of being elected given the premise of the Lemma. ■

For the following proofs, as defined in the text, denote the expected outcome when policy p is played and beliefs centre on some parameters β , as $E[y|p, \beta]$. let $\delta(p, \beta) = E[y|p, \beta^*] - E[y|p, \beta]$. That is, given a policy p , $\delta(p, \beta)$ measures the average mistake that a model β yields. Furthermore, let $B(p)$ denote the set of vectors β that solve $\delta(p, \beta) = 0$. When $p = L$ (R) this set corresponds to all vectors where $\beta_l = \beta_l$ ($\beta_r = \beta_r$). $B(M)$ is the line of all vectors β satisfying $\frac{1}{2}\beta_l + \frac{1}{2}\beta_r = \frac{1}{2}\beta_l + \frac{1}{2}\beta_r$, i.e. a line $\beta_r = \beta_r + \beta_l - \beta_l$ with slope -1 and intercept given by the sum of the components of β^* . Note that when $p \neq p'$, $B(p) \cap B(p')$ is a singleton and includes only β^* .

The random policy function (that arises given the randomness in the election and the randomness of the shock ε , through its effect on beliefs), induces a probability distribution

P over the set of possible histories \mathbb{H} . Thus, when we write "almost surely" we mean P -almost surely on \mathbb{H} .

Remember that for history H_t we define the associated distribution over implemented actions at time t , $\hat{\eta}_t$, as the share of time each policy was implemented in H_t and we let η_t (*polarisation*) be the fraction of time in the history up to time t that the two parties offered different platforms.

Proof of Proposition 1:

Although the voters in our model are not Bayesian, they do use likelihood ratios to compare between different models. In this proof it will be useful to think of a Bayesian agent with a uniform prior on B that uses standard Bayesian updating to update her beliefs about the true model β^* . The uniform prior implies that any inferences we make about likelihood ratios in the beliefs of the Bayesian agent will exactly coincide with those our voters will end up calculating.

A Bayesian updater who observes the history H_t satisfies the conditions of the martingale convergence theorem. Therefore, beliefs converge almost surely to some limiting probability distribution μ_∞ on B . Let the support of these beliefs be denoted by B_∞ . This implies that the likelihood ratio satisfies $\log \frac{\mathcal{L}(\beta|H_t)}{\mathcal{L}(\beta'|H_t)} \rightarrow \infty$ for all $\beta \in B_\infty$ and $\beta' \notin B_\infty$.

Step 1: In this step we show that the posterior distribution of the beliefs over β , denoted by μ_t , will concentrate on the set B_∞^η whose elements β have a Kulbeck-Liebler value that is close to the minimiser of the Kulbeck-Liebler value.

Our model satisfies Assumptions 1-3 in Esponda et al (forthcoming), henceforth EPY.²⁹ The Kulbeck-Liebler (KL) divergence value of some vector of parameters β given some $\hat{\eta}_t$ is defined as

$$KL(\beta|\hat{\eta}_t, \beta^*) = \sum_{p \in \{L, M, R\}} \hat{\eta}_t(p) \int_{\mathbb{R}} f(\varepsilon) \ln \frac{f(\varepsilon)}{f(E[y|p, \beta^*] + \varepsilon - E[y|p, \beta])} d\varepsilon$$

where $f(\varepsilon)$ is the density over ε , assumed here to be normal with mean zero. The KL divergence is always non-negative, given Gibbs inequality, and β^* is a minimizer of KL for which $KL(\beta^*|\hat{\eta}_t, \beta^*) = 0$ regardless of $\hat{\eta}_t$. Following EPY, $\beta_{\min}(\hat{\eta}_t) = \beta^*$ is a minimizer and $K(\hat{\eta}_t) \equiv KL(\beta_{\min}(\hat{\eta}_t)|\hat{\eta}_t, \beta^*) = 0$. Theorem 1 in EPY implies that there exists a set $E \subset \mathbb{H}$ such that $P(E) = 0$ and that for all $H \in \mathbb{H} \setminus E$,

$$\lim_{t \rightarrow \infty} \int_B KL(\beta|\hat{\eta}_t, \beta^*) d\mu_{t+1}(\beta) = 0.$$

We now show that this implies that if for some β the KL value is strictly positive that a ball around β must have zero measure in the limit beliefs. So suppose $\beta \in B$ is such that

²⁹In our model the policy function is random at every period but this has no bearing on the proof of Theorem 1 in EPY.

$KL(\beta|\hat{\eta}_t, \beta^*)$ does not converge to zero. Then there exists $\psi > 0$ and a subsequence t_n such that for all t_n , $KL(\beta|\hat{\eta}_{t_n}, \beta^*) \geq \psi$. By continuity there is an epsilon ball $B^\epsilon(\beta)$, around β , such that

$$\int_B KL(\beta|\hat{\eta}_{t_n}, \beta^*) d\mu_{t_k+1} \geq \psi \int_{B^\epsilon(\beta)} \mu_{t_n+1}(\beta') d\beta',$$

and then we must have $\int_{B^\epsilon(\beta)} \mu_{t_n+1}(\beta') d\beta' \rightarrow 0$ and so $\mu_{t_n}(\beta') \rightarrow 0$ almost surely in $B^\epsilon(\beta)$, as $t_n \rightarrow \infty$.

Step 2: If $B_\infty \subseteq B_p$ for some p , then Lemma 1 implies that parties converge to both offering p . Consider now paths on which B_∞ includes beliefs both in B_p and $B_{p'}$ where $p \neq p'$. Let p have a model $\beta \in B_p$ that maximises the likelihood of the history. If the limit beliefs satisfy condition (*) strictly, then we have convergence far enough along the sequence. Assume then that condition (*) is not satisfied strictly. In this case, if the model that maximises the likelihood of history supports $p = M$, and condition (**) is not satisfied, the unique equilibrium involves convergence on M . If condition (**) is satisfied, then the unique equilibrium has polarisation. If instead the model that maximises the likelihood of history supports L or R then the unique equilibrium is again polarisation.

So we focus on these cases in which the unique equilibria involve polarisation. Each of the parties is elected in equilibrium with a strictly positive probability due to $\alpha > 0$. We then have a positive measure of paths for which $\hat{\eta}_t(L)$ and $\hat{\eta}_t(R)$ are bounded away from zero. This means that $\beta^* = \arg \min_{\beta'} KL(\beta'|\hat{\eta}_{t_n}, \beta^*)$, i.e. β^* is the unique minimizer of the KL divergence. By continuity and Step 1, beliefs can only concentrate on an ball around β^* . By the convergence of beliefs we must have that beliefs along these paths have $B_\infty = \{\beta^*\} \subseteq B_M$. This contradicts the supposition at the beginning of this step. ■

Proof of Proposition 2:

We start with part (ii). Suppose not, therefore there is a positive measure of paths along which there is one policy p and a subsequence t_n such that $\hat{\eta}_{t_n}(p) \rightarrow 1$. For any t denote the preceding K periods of history as the K – window at t . For any t_n , consider the set of all K – windows for any t' up to t_n .

Claim 1: (a) After each K – window in which only one policy $p \in \{L, R\}$ was implemented, the next period must involve each party choosing a different policy. (b) When K is large and $t_n \rightarrow \infty$, on a strictly positive measure of paths, almost surely after each K – window in which only one policy $p = M$ was implemented, each party will choose its ideal policy in the next period.

Proof of Claim 1: First note that given any K observations in which only one policy p

was implemented, $(y_s = E[y|p, \beta^*] + \varepsilon_s)_{s=t'-K}^{t'-1}$, the log-likelihood of some belief β will be

$$\begin{aligned} \log \prod_{s=t'-K}^{t'-1} f(y_s - E[y|p, \beta]) &= \frac{1}{\sigma\sqrt{2\pi}} \sum_{s=t'-K}^{t'-1} \log e^{-\frac{1}{2}\left(\frac{\delta(p, \beta) + \varepsilon_s}{\sigma}\right)^2} = \\ &= -\frac{1}{2\sigma^3\sqrt{2\pi}} \sum_{s=t'-K}^{t'-1} (\delta(p, \beta) + \varepsilon_s)^2. \end{aligned}$$

It follows that β maximizes likelihood if and only if $\delta(p, \beta)$ minimizes $\sum_{s=t'-K}^{t'-1} (\delta(p, \beta) + \varepsilon_s)^2$, which implies $\delta(p, \beta) = -\frac{1}{K} \sum_{s=t'-K}^{t'-1} \varepsilon_s$. Then we see that the set of $\beta \in \mathbb{R}^2$ that solve the equation is a line parallel to $B(p)$.

(a) In the case $p \in \{L, R\}$, the line parallel to $B(p)$ would be vertical or horizontal. If the realization of $(\varepsilon_{t'-K}, \dots, \varepsilon_{t'-1})$ is such that the set of solutions intersects B , at least two different policies maximise likelihood. Thus the two parties cannot offer the same platform at time t' . If the set of solutions does not intersect B , since the log-likelihood is strictly concave, then one of the edges of B is the set of maximizers of the log-likelihood and the same argument applies since there are beliefs on any edge that support both M and either L or R .

(b) Consider now the case $p = M$. Here polarisation happens in finite time only (almost surely). To see this, observe that if the realization of the shocks is such that

$$\beta_r + \beta_l + \frac{1}{K} \sum_{s=t'-K}^{t'-1} \varepsilon_s \in [c(R) - c(M), 2\beta - c(L) + c(M)],$$

then each party polarizes on their favourite policy (otherwise, consensus continues). This follows from the fact that a line with slope -1 that intersects B_R must also intersect B_L and vice versa. Thus, when K is large and t_n goes to infinity, then on a strictly positive measure of paths, each K -window with only $p = M$ will almost surely imply complete polarisation in the next period. \square

We can now use Claim 1 to establish (ii). As $\hat{\eta}_{t_n}(p) \rightarrow 1$, as t_n goes to infinity, the fraction of these K -windows with only p implemented within the window must be going to one. By Claim 1, each of these will lead to polarisation for sure in the next period (in the case where $p = L, R$) or to polarisation almost surely (in the case where $p = M$). This however contradicts our assumption that $\hat{\eta}_{t_n}(p) \rightarrow 1$.

Part (i): We start with showing that $\liminf \eta_t(\text{polarisation}) > 0$. Suppose not, and so there is a positive measure of paths along which there is a subsequence t_n such that $\eta_{t_n}(\text{polarisation}) \rightarrow 0$. This implies that if we look at all the K -windows almost all of them include no polarisation. Similar to (ii), it cannot be that there is strictly positive measure of K -windows with only one policy implemented as then we would have $\eta_{t_n}(\text{polarisation})$

bounded from zero. Thus the only possibility that remains is that in almost all K -windows, at least two policies p and p' are implemented, and that parties will shift from a consensus on one policy p to a consensus on another policy p' (a “consensus-switch”).

Claim 2: *At any time t when there is a consensus-switch from p to p' , there is a $D > 0$ such that $\max\{|\varepsilon_{t-K-1}|, |\varepsilon_{t-1}|\} > D$.*

Proof of Claim 2: To see this note that the log likelihood of any β at period t is given by:

$$\begin{aligned} \log \mathcal{L}(\beta|H_t) &= \log \prod_{s=t-K}^{t-1} f(y_s - E[y|p_s, \beta]) = \frac{1}{\sigma\sqrt{2\pi}} \sum_{s=t-K}^{t-1} \log e^{-\frac{1}{2}\left(\frac{\delta(p_s, \beta) + \varepsilon_s}{\sigma}\right)^2} = \\ &= -\frac{1}{2\sigma^3\sqrt{2\pi}} \sum_{s=t-K}^{t-1} (\delta(p_s, \beta) + \varepsilon_s)^2 = \end{aligned}$$

and therefore the difference between periods t and $t-1$ log likelihoods is:

$$\log \mathcal{L}(\beta|H_t) - \log \mathcal{L}(\beta|H_{t-1}) = -\frac{1}{2\sigma^3\sqrt{2\pi}} [(\delta(p_{t-1}, \beta) + \varepsilon_{t-1})^2 - (\delta(p_{t-1-K}, \beta) + \varepsilon_{t-1-K})^2]$$

By Lemma 1, and the fact that the consensus switched at t from some policy p to another policy p' we must have that there is a $\beta \in B_p$ and a $\beta' \in B_{p'}$ such that

$$\log \frac{\mathcal{L}(\beta|H_{t-1})}{\mathcal{L}(\beta'|H_{t-1})} > \frac{1}{2\zeta} - 2\zeta \frac{\alpha}{0.5 + \alpha} \text{ and } \log \frac{\mathcal{L}(\beta'|H_t)}{\mathcal{L}(\beta|H_t)} > \frac{1}{2\zeta} - 2\zeta \frac{\alpha}{0.5 + \alpha}$$

But this implies that

$$\begin{aligned} \log \frac{\mathcal{L}(\beta|H_{t-1})}{\mathcal{L}(\beta'|H_{t-1})} + \log \frac{\mathcal{L}(\beta'|H_t)}{\mathcal{L}(\beta|H_t)} &> \frac{1}{\zeta} - 4\zeta \frac{\alpha}{0.5 + \alpha} \Leftrightarrow \\ \log \mathcal{L}(\beta|H_{t-1}) - \log \mathcal{L}(\beta|H_t) + \log \mathcal{L}(\beta'|H_t) - \log \mathcal{L}(\beta'|H_{t-1}) &> \frac{1}{\zeta} - 4\zeta \frac{\alpha}{0.5 + \alpha} \Leftrightarrow \\ \frac{(\delta(p_{t-1}, \beta) + \varepsilon_{t-1})^2 - (\delta(p_{t-1}, \beta') + \varepsilon_{t-1})^2 + (\delta(p_{t-1-K}, \beta') + \varepsilon_{t-1-K})^2 - (\delta(p_{t-1-K}, \beta) + \varepsilon_{t-1-K})^2}{2\sigma^3\sqrt{2\pi}} &> \\ > \frac{1}{\zeta} - 4\zeta \frac{\alpha}{0.5 + \alpha} \end{aligned}$$

But since B is compact, for low enough ζ we need to have a $D > 0$ such that $\max\{|\varepsilon_{t-k-1}|, |\varepsilon_{t-1}|\} > D$. \square

Claim 2 will lead to a contradiction to our supposition that in almost all K -windows, a consensus-switch occurs. To see this, note that for any D , given the normal distribution over ε , we can compute the probability that one of these shocks is higher in magnitude than D . As D grows large, this probability goes down, and becomes lower than $\frac{1}{K}$. This implies for a

K large enough, we can find a small enough ζ such that the probability of a consensus-switch goes to zero and contradicts the supposition that in almost all K – windows where such a switch occurs. Thus, there is no positive measure of paths along which there is a subsequence t_n such that $\eta_{t_n}(\text{polarisation}) \rightarrow 0$.

We now show that $\limsup_t \eta(\text{polarisation}) < 1$. Suppose not, and so there is a positive measure of paths along which there is a subsequence t_n such that $\eta_{t_n}(\text{polarisation}) \rightarrow 1$. This implies that if we look at all the K – windows almost all of them include polarisation at every period, implying that for all windows there exist at least two different policies p and p' implemented with a strictly positive probability. As a result, for a large enough K and $t_n \rightarrow \infty$, after almost all the K – windows we have that, as in Proposition 1, beliefs almost surely concentrate on a ball around β^* . This implies that both parties must choose M after almost all these K – windows, a contradiction to $\eta_{t_n}(\text{polarisation}) \rightarrow 1$. ■

We now consider the case of $\sigma^2 \rightarrow 0$ and introduce two helpful Lemmata.

Lemma A1: (i) Assume that a policy p was implemented for one period. Then $E[\log \frac{f(\delta(p, \beta') + \varepsilon)}{f(\delta(p, \beta'') + \varepsilon)}] = 0$ for all $\beta', \beta'' \in B(p)$, and when $\sigma^2 \rightarrow 0$, $E[\log \frac{f(\delta(p, \beta') + \varepsilon)}{f(\delta(p, \beta'') + \varepsilon)}] \rightarrow \infty$ for all $\beta' \in B(p)$ and $\beta'' \notin B(p)$. (ii) Assume that two different policies p, p' were implemented across periods 1 and 2. Then when $\sigma^2 \rightarrow 0$, $E[\log \frac{f(\delta(p, \beta^*) + \varepsilon_1)}{f(\delta(p, \beta') + \varepsilon_1)} \frac{f(\delta(p, \beta^*) + \varepsilon_2)}{f(\delta(p, \beta') + \varepsilon_2)}] \rightarrow \infty$ for all $\beta' \neq \beta^*$.

Proof of Lemma A1: To see (i), note that by definition of $B(p)$, $E[\log \frac{f(\delta(p, \beta') + \varepsilon)}{f(\delta(p, \beta'') + \varepsilon)}] = E[\log \frac{f(\varepsilon)}{f(\varepsilon)}] = 0$ for all $\beta', \beta'' \in B(p)$. Also, $E[\log \frac{f(\delta(p, \beta') + \varepsilon)}{f(\delta(p, \beta'') + \varepsilon)}] = E[\log \frac{f(\varepsilon)}{f(\delta(p, \beta'') + \varepsilon)}]$ where $\delta(p, \beta'') \neq 0$ as $\beta'' \notin B(p)$. This implies that $E[\log \frac{f(\delta(p, \beta') + \varepsilon)}{f(\delta(p, \beta'') + \varepsilon)}] = \frac{\delta(p, \beta'')^2}{2\sigma^2} \rightarrow_{\sigma^2 \rightarrow 0} \infty$ for any $\beta' \in B(p)$ and $\beta'' \notin B(p)$. To see (ii), note again that $E[\log \frac{f(\delta(p, \beta^*) + \varepsilon_1)}{f(\delta(p, \beta') + \varepsilon_1)} \frac{f(\delta(p, \beta^*) + \varepsilon_2)}{f(\delta(p, \beta') + \varepsilon_2)}] = \frac{\delta(p, \beta')^2 + \delta(p, \beta'')^2}{2\sigma^2} \rightarrow_{\sigma^2 \rightarrow 0} \infty$ for any $\beta' \neq \beta^*$. ■

Lemma A2: When $\sigma^2 \rightarrow 0$, (i) if only one policy was implemented throughout the observed K periods of H_t , then parties polarise; (ii) if there are two different policies implemented throughout the observed K periods of the history H_t then with probability one the unique equilibrium is convergence on M .

Proof of Lemma A2: (i) Given Lemma A1, when one policy p is implemented for K periods, in expectation all beliefs in $B(p)$ maximise the likelihood and all models outside of $B(p)$ have infinitely smaller likelihoods. By our assumptions about the state space and by Lemma 1(iii) parties polarise.

(ii) When p and p' have been implemented, then given Lemma A1, β^* is the unique maximum likelihood model and all other models are infinitely less likely. As $\beta^* \in B_M$, by Lemma 1(ii) the unique equilibrium is (M, M) . ■

Proof of Proposition 3: Following Lemma A2, if both parties offered M in the last K periods, then they both polarise on their preferred policy. Moreover, if two different policies are observed subsequently, a cycle of converge on M lasting K periods begins. Then we are left to prove that at some point in time, two different policies need to be implemented one following the other. To see this, consider what happens in period t . If H_t includes two different policies implemented in periods $t - 1$ and $t - 2$ we are done. Assume then that this is not the case. If the last K periods include no variation in policies, then parties fully polarise at period t offering L and R . As each is elected with probability half in this case, in finite time at some period periods $t + s$ and $t + s - 1$ will include two different policies implemented. If the history includes some variation (but not across periods $t - 1$ and $t - 2$ as assumed), it must be that H_{t-1} included variation in policies, which implies that the policy implemented in period $t - 1$ (and hence $t - 2$) is M , and will be so until M until the first history sequence in which M is implemented K times. Following that, parties must polarise and we then have two consecutive periods in which two different policies are implemented.

■

References

- [1] Abiad, A. and A. Mody (2005), “Financial Reform: What Shakes It? What Shapes It?”, *The American Economic Review*, Vol. 95 (1), pp. 66-88.
- [2] Acemoglu, D. and A. Wolitzky (2014), “Cycles of Conflict: An Economic Model”, *American Economic Review* 104, 1350-1367.
- [3] Aksoy, C.G, B. Eichengreen and O. Saka (2020), “The Political Scar of Epidemics”, mimeo.
- [4] Alesina, A., S. Ardagna and F. Trebbi (2006), “Who Adjusts and When? The Political Economy of Reforms”, *IMF Staff Papers*, Vol. 53, pp. 1-29
- [5] Alesina, A. and Rosenthal, H. (1995). *Partisan Politics, Divided Government, and the Economy (Political Economy of Institutions and Decisions)*. Cambridge: Cambridge University Press.
- [6] Azzimonti, M. and M. Fenandes (2018), “Social Media Networks, Fake News, and Polarisation”, NBER working paper 24462.
- [7] Baker, S.R, A. Baksy, N. Bloom, S. David and J.A. Rodden (2020), “Elections, political polarisation and economic uncertainty”, NBER working paper 27961.

- [8] Baker, S.R., N. Bloom, B. Canes-Wrone, S.J. Davis, and J. Rodden (2014), “Why Has US Policy Uncertainty Risen Since 1960?”, *American Economic Review: Papers and Proceedings*, 104(5): 56–60.
- [9] Baker, S.R., N. Bloom, and S.J. Davis, (2014), “Measuring Economic Policy Uncertainty”, *The Quarterly Journal of Economics*, Volume 131, Issue 4, Pages 1593–1636
- [10] Barber, M. and N. McCarty (2013), “The Causes and Consequences of polarisation.” in *Negotiating Agreement in Politics*, eds. Cathie Jo Martin and Jane Mansbridge. APSA task force report.
- [11] Battaglini, Marco, and Stephen Coate. 2008. "A Dynamic Theory of Public Spending, Taxation, and Debt." *American Economic Review*, 98 (1): 201-36.
- [12] Bechtel, M. M. and J. Hainmueller (2011), “How Lasting Is Voter Gratitude? An Analysis of the Short- and Long-Term Electoral Returns to Beneficial Policy”, *AJPS* 55, 4, Pages 852-868.
- [13] Becker, S., T. Fetzer and D. Novy (2017), “Who Voted for Brexit? A Comprehensive District-Level Analysis”, *Economic Policy*, 32(92), 601–650, October 2017.
- [14] Bloom, N, (2014), “Fluctuations in Uncertainty”, *Journal of Economic Perspectives*, Vol. 28, No 2, 153-176.
- [15] Callander, S. (2011), “Searching for good policies”, *American Political Science Review*, 105 (04), 643–662.
- [16] Callander, S., F. Izzo and G. Martin (2021), “Ideological Competition”, mimeo.
- [17] Calvert, R. L. (1985), “Robustness of the Multidimensional Voting Model: Candidate Motivations, Uncertainty, and Convergence,” *American Journal of Political Science*, 29 (1), 69–95.
- [18] Drazen, A. and V. Grilli (1993), “The Benefit of Crises for Economic Reforms”, *The American Economic Review*, Vol. 83 (3), pp. 598-607.
- [19] Eliaz, K. and R. Spiegler (2020), “A Model of Competing Narratives”, *American Economic Review*. Vol. 110, No. 12, pp. 3786–3816.
- [20] Esponda, I., D. Pouzo and Y. Yamamoto (forthcoming), "Asymptotic Behavior of Bayesian Learners with Misspecified Models", *Journal of Economic Theory*.

- [21] Funke, M., M. Schularick, and C. Trebesch (2016), "Going to extremes: Politics after financial crises, 1870–2014", *European Economic Review*, Volume 88, Pages 227-260.
- [22] Goldin, C. and G. D. Libecap, eds. *The regulated economy: A historical approach to political economy*. University of Chicago Press, 2008.
- [23] Healy, Andrew and Lenz Gabriel S. (2014), "Substituting the End for the Whole: Why Voters Respond Primarily to the Election-Year Economy", *AJPS*, Volume 58, Pages 31-47.
- [24] Kenny, M. (2017) *Back to the populist future?: understanding nostalgia in contemporary ideological discourse*, *Journal of Political Ideologies*, 22:3, 256-273.
- [25] Levy, G., R. Razin and A. Young (forthcoming), "Misspecified Politics and the Recurrence of Populism", *American Economic Review*, March 2022.
- [26] Levy, G., Moreno de Barreda, I., and R. Razin (2021), "Polarized Extremes and the Confused Centre: Campaign Targeting of Voters with Correlation Neglect", *Quarterly Journal of Political Science*: Vol. 16: No. 2, pp 139-155.
- [27] Little, A.T. (2019), "The Distortion of Related Beliefs", *American Journal of Political Science*, Volume 63, Issue 3, Pages 675-689.
- [28] Little, A. (2021), *Directional motives and different priors are observationally equivalent*, mimeo.
- [29] Little, A., K. Schnakenberg and I. Turner (2020), "Motivated Reasoning and Democratic Accountability", mimeo.
- [30] Malmendier, U., S. Nagel, and Z. Yang (2021), "The Making of Hawks and Doves", *Journal of Monetary Economics* 117, 19-42.
- [31] Malmendier, U. and S. Nagel (2016), "Learning from Inflation Expectations", *Quarterly Journal of Economics*, 53-87.
- [32] McCarty, N. (2019). *Polarization and the Changing American Constitutional System*. In F. Lee & N. McCarty (Eds.), *Can America Govern Itself? (SSRC Anxieties of Democracy, pp. 301-328)*. Cambridge: Cambridge University Press.
- [33] McCarty, N., Poole, K., & Rosenthal, H. (2013). *Political Bubbles: Financial Crises and the Failure of American Democracy*. Princeton; Oxford: Princeton University Press.

- [34] Mian, A., A. Sufi and F. Trebbi (2014), “Resolving Debt Overhang: Political Constraints in the Aftermath of Financial Crises”, *American Economic Journal: Macroeconomics*, Vol. 6, No. 2, pp. 1-28.
- [35] Myrick R. (2020), “Do External Threats Unite or Divide? Security Crises, Rivalries, and Polarization in American Foreign Policy”, *International Organisation*, 1-38. doi:10.1017/S0020818321000175.
- [36] Orlik, A. and L. Veldkamp (2014), “Understanding Uncertainty Shocks and the Role of Black Swans”, NBER working paper 20445.
- [37] Piketty, Thomas, *Capital and Ideology* (Cambridge, MA: Harvard University Press, 2020).
- [38] Rogoff, K. (1990). “Equilibrium Political Budget Cycles.” *American Economic Review* 80: 21-36.
- [39] Schlesinger, Arthur Sr. (1949). *Paths to the Present*. Macmillan.
- [40] Schlesinger, Arthur Jr. (1999). *The Cycles of American History*. Houghton Mifflin Harcourt.
- [41] Strulovici, B. (2010), “Learning While Voting: Determinants of Collective Experimentation”, *Econometrica*, May 2010.