DISCUSSION PAPER SERIES

DP16737

Short-term political memory and the inevitability of polarisation

Ronny Razin and Gilat Levy

POLITICAL ECONOMY

CEPR

ISSN 0265-8003

Short-term political memory and the inevitability of polarisation

Ronny Razin and Gilat Levy

Discussion Paper DP16737
Published 20 November 2021
Submitted 16 November 2021

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

Short-term political memory and the inevitability of polarisation

Abstract

In this paper we explore the effect of short memories 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 learn about the optimal policy 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 convergence. Historical periods of convergence lead parties to polarise, whereas periods of polarisation imply convergence of platforms. Our framework also allows us to model the strategic use of biased histories and narratives in political competition, such as the use of nostalgia.

JEL Classification: N/A

Keywords: N/A

Ronny Razin - r.razin@lse.ac.uk
London School of Economics and CEPR

Gilat Levy - g.levy1@lse.ac.uk
London School of Economics and CEPR

Acknowledgements
We thank seminar participants and discussants in the ETH 4th workshop on democracy, SITE 21 in Stanford, and the Berkeley Political Economy Conference 2021.

Short-term political memory and the inevitability of polarisation

Gilat Levy and Ronny Razin, LSE[1]

Abstract: In this paper we explore the effect of *short memories* 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 learn about the optimal policy 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 convergence. Historical periods of convergence lead parties to polarise, whereas periods of polarisation imply convergence of platforms. Our framework also allows us to model the strategic use of biased histories and narratives in political competition, such as the use of nostalgia.

1 Introduction

The rise of political polarisation in many western democracies in the last few decades has received much attention recently.[2] But 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, Barber and McCarty (2013) show that 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).

[1]Department of Economics, Houghton St. WC2A 2AE, UK. Email addresses: g.levy1@lse.ac.uk, r.razin@lse.ac.uk. We thank seminar participants and discussants in the ETH 4th workshop on democracy, SITE 21 in Stanford, and the Berkeley Political Economy Conference 2021.

[2]See Barber and McCarty (2013) for a survey of this literature.

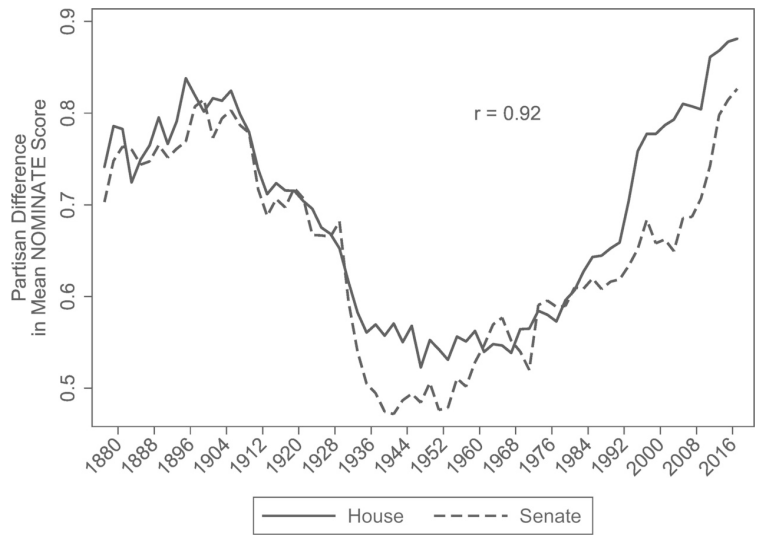

Figure 1: Historical polarisation in the US Senate and Congress (McCarty 2019).

Similar cycles can be observed in political parties' stated ideology by looking at manifestoes over time. The Manifesto project decodes manifestos of political parties into a unidimensional score on many topics and allows for comparisons over time. Looking at the manifestos of the two US parties we see oscillation between polarisation and convergence on economic issues (see Figure 2).

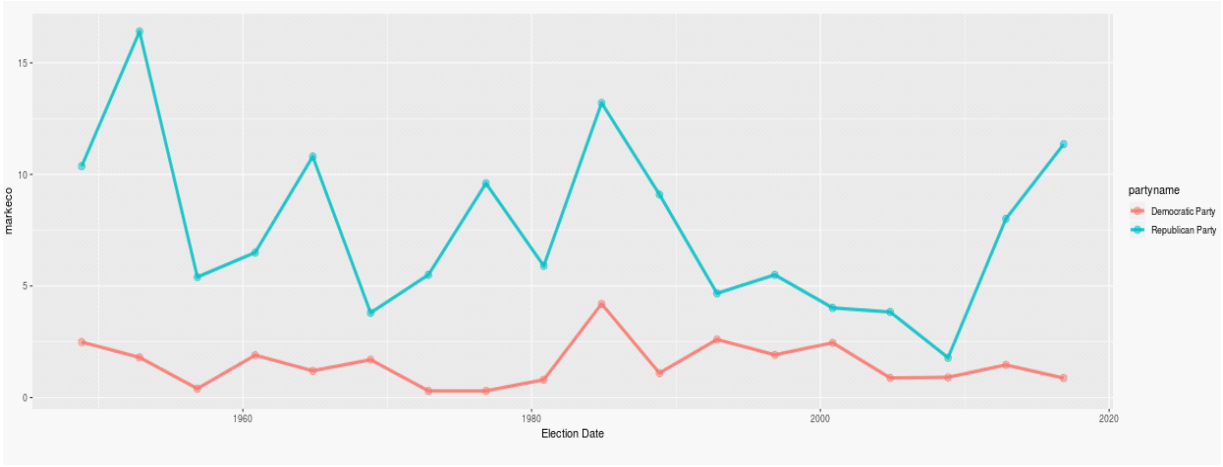

Figure 2: US party platforms on economic issues. Taken from the Manifesto Project.

This cyclical pattern manifests itself in many public opinion polls as well. On tax issues for example, there are many periods in which the population is highly divided on whether they pay too much or too little tax, and periods in which a majority believes it pays too much (see Figure 3). Similar dynamic patterns emerge on many other issues, such as attitudes towards the death penalty.

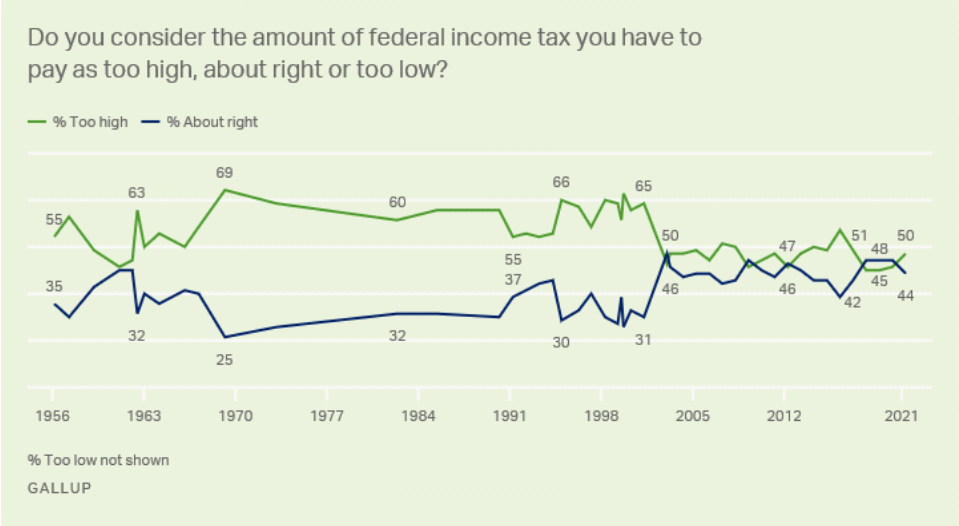

Figure 3: A US Gallup survey on attitudes towards taxation.

Such cycles in public opinion and in politicians' stated preferences and voting behaviour might also have an effect on implemented policy (and vice verse of course). Looking at US tax rate brackets over the years, there is substantial variation of policy leading up to the late 1970s with little variation after 1980 (see Figure 4):

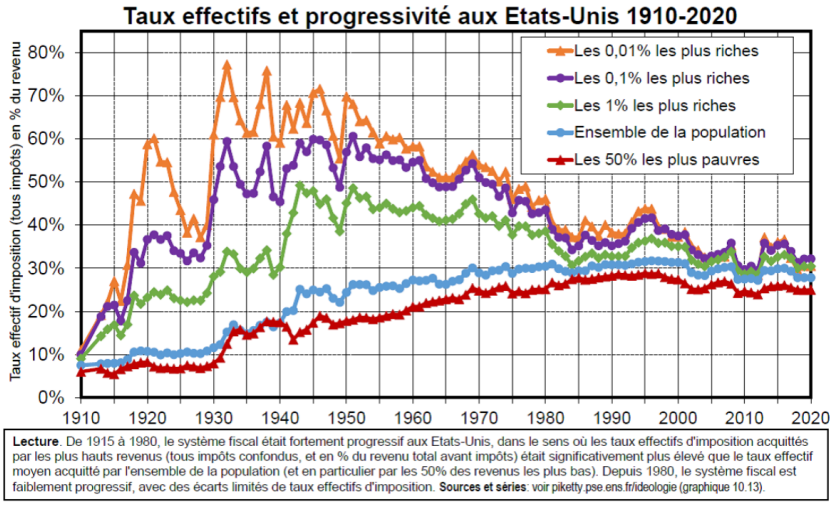

Figure 4: Progressive taxation in the US 1910-2020 (taken from Thomas Piketty's Lecture).

Somewhat unsurprisingly, the little variation in policies in the latter period accords with the convergence of the manifestos of the two US parties; their platforms are indeed very similar on economic issues in the 1990s up to 2010 (see Figure 2 above). As a result, even as power has changed hands between Republican control to Democratic control a few times in this period, there is little policy change.[3]

[3]Specifically, there was a reduction of progressive taxation also during the time of the Clinton regime who followed Reagan and George W. H. Bush, as well as during the time of Obama who followed George

What fosters this cyclical pattern of national consensus versus polarisation in public attitudes and political outcomes? Of course for any particular period in history the literature has found different driving forces for polarisation or consensus, be it economic, technological or institutional changes. In this paper we show however that there is also a more fundamental reason for such political cycles: When societies go through a collective learning process, short memory in the political arena is a key driver for polarisation and consensus cycles. Specifically, we show the effect of short-term memory on the dynamic evolution of *consensus* versus *polarisation* in public attitudes as well as political party platforms. We also show how together with short-term memory, other forces mentioned in the literature (e.g., crises) can then exacerbate or ameliorate the possibility for polarisation. Moreover, selective memories of different groups, will lengthen polarisation and may give a political advantage to some particular demographics, such as old voters.

Our model as well as results are in fact well summarised by a paragraph in Thomas Piketty's recent book, Capital and Ideology, where he writes: "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)."[4]

We analyse a dynamic model in which policy is determined endogenously, following political competition between ideologically motivated parties each pushing a different interest. Voters look at the historical experience of their polity to form beliefs and balance their votes between their affinity to a party and a platform that is considered to be a better fit to the state of the world, given the data. For example, they might consider whether growth is better achieved with policies that increase inequality to increase risky investment and job creation by the rich (a position pushed by one party), or with a more redistributive based policy which aims to achieve more consumption power (a position pushed by another). The optimal policy that fits the true state of the world is a compromise or moderate policy (which we first assume to be fixed over time). Our key assumption is that while voters learn from the observed outcomes, their memory is short.

W. Bush. Overall, the top marginal rate fell to an average of 39 percent from 1980 to 2018, half its level in the period 1932-1980. Similarly, the federal minimum wage fell in real value.

[4]Thomas Piketty, Capital and Ideology, page 10.

Our cycles result builds on two key insights. First, periods of little policy variation will imply uninformative histories which would leave room for alternative and opposing views to be espoused by parties. Specifically, when information about the right course of action is scarce or there is little to discern between competing narratives, special interests will find it easier to push their agendas forward. In contrast, higher variation in policy over the years yields more information about what effect policies have on outcomes. In this latter case Society might 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 agenda and have to settle on a consensus policy.

Second, short memory is an essential feature for transitions between consensus and polarisation. If the voters had perfect memory, they would become more informed as time passes, accumulating more and more experiences of polices and outcomes. On the other hand, a polity that only remembers a number of previous periods in history, oscillates between having more or less information depending on the variation in the observed data. Specifically, periods of party convergence or low political turnover will imply relatively uninformative histories, enabling platforms' polarisation. Alternatively, periods of high polarisation and turnover will imply a greater agreement on what is best, enabling platforms' to convergence. The polarisation phase of the cycle sows then the seeds for the consensus phase and vice verse.

Our framework also allows us to make predictions about how the political system responds to crisis. In our framework we highlight two sources of crises that can affect the polity. One driver of a crisis is a large and unexpected shock to the economy, affecting the way voters interpret historical data. We show that the effect of a large shock on the polity is long lasting, as long as this period is part of the collective memory. In fact, in line with cohort effects identified in the literature, such large events (wars, pandemics) can have disproportional weight in the voters' memory. Moreover, we show that the effect of a crisis depends on the current policy that was implemented when the shock hit. In particular, shocks tend to cause extended periods of polarisation when they hit an economy that was already in a polarisation phase, but they may also induce more moderation when the shock hits at a period in which a moderate policy was implemented.

The second driver for a crisis in our model is a change in technology or in the true data generating process. For example, the 2007 financial crisis was partially driven by new technologies of financial derivatives and trades. Many players in the financial markets did not fully understand this technological change, and continued to use old models and old data to predict where market trends. We show that when the state of the world changes, the polity polarises. When voters are aware of this change, naturally they put less weight on historical

observations and thus have less knowledge. We show that even when voters are unaware of this change, their learning is hampered as they are unable to properly interpret their data. This again allows parties to polarise. This is consistent with the results in Mian et al (2014) that show that voters become more polarised following a financial crisis.

In practice different voters sometimes rely on selected and different historical observations. In this vein we also extend the model to allow politicians to highlight different historical evidence to different voters. Naturally, such echo chambers increase polarisation. However, different groups may have an advantage when history is concerned. For example, the old generation have lived through a longer history and this allows politicians to draw observations from a larger set. Malmendier and Nagel (2016) indeed 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.[5] We analyze the inter-generational conflict between the old and young and show how learning from endogenous policies exacerbates the advantage of the old. Specifically, the young that are bound by more recent history have to moderate more often whereas the old can pursue their self-interest with a higher probability.

The remainder of the paper proceeds as follows. In Section 2 we describe the related literature. In Section 3 we present our dynamic model. Section 4 presents the main cycles result, and also comparative statics of the cycles to very large as well as small shocks. In Section 5 we consider what happens when the data generating process changes and in Section 6 we discuss selective histories and the conflict between the old and the young.

2 Related Literature

Our paper highlights a new interpretation of the term political cycles as inherent to the nature of politics as a collective learning process. The literature has mainly focused so far on the interaction between the economy and the political system.[6] Rogoff (1990) focuses on the incentive of politicians in election years to stimulate the economy in order to improve their prospects of reelection. Relatedly, 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). Alesina and Rosenthal (1995) focus on the US political system and show how political cycles can

[5] This can then translate into policy making; Malmandier et al (2021) show how personal experiences of inflation strongly influence the hawkish or dovish leanings of central bankers.

[6] Since 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).

arise when voters use midterm elections to tame polarised presidents by splitting their votes. McCarty et al (2013) analyse the interaction between politics and the economy from the perspective of the 2007 financial crisis in the US. They suggest that the US political system amplifies economic crises and is ill-equipped to take action in the aftermath of a crisis.

Levy, Razin and Young (2020) is also a social learning model in which society chooses endogenously which party can implement its action. They show that political cycles arise between two groups, when one group has a complex (and correct) model of the world and another group has a simple model, in that it considers fewer policy instruments to be relevant. In their model too voters attempt to learn the true state of the world, but each group has a different model in mind. They show that in the long term the polity converges to have cycles where periods in which complex rule must follow periods of simple rule. The intuition there is different and relates to the fact that perpetual rule by one party implies that the party in opposition becomes more eager in its preferences to win the election due to wrong interpretation of the outcomes implemented by the ruling party.[7] Wolitzky and Acemoglu (2014) analyze a conflict model in which two groups fail to coordinate on the good outcome when they receive a wrong signal about the intentions of the other group (and have limited memory of previous history). Still, a sufficiently long history of a conflict allows the groups to realize that a conflict has started by mistake reverting to a coordination phase thus generating cycles.

Our paper is related to a recent literature that focuses on learning throughout the political process. Piketty (2020) provides a comprehensive historical overview of inequality regimes and ideologies in different countries through the prism of politics as a collective learning process. Our work formalises this learning process and how different ideologies or narratives could be put forward to explain the same historical data. Callander (2011) analyses a model of learning from outcomes when the mapping between policies and outcomes is complex. Callander, Martin and Izzo (2021) analyse a static one-period model in which two ideological parties compete to endow a voter with a belief about the true state of the world; in their model too, an uninformative history leads to polarisation. In our dynamic model the history is determined endogenously, which allows us to derive the cycle results.

Our paper is also related to Little (2019) who studies motivated beliefs in political environments. Little (2019) explores in a one-voter decision problem how motivated reasoning distorts beliefs.[8] In his model a voter trades off the belief that is most likely to explain her set

[7] Azzimonti 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.

[8] See also Little (2021), and Little, Schankenberg and Turner (2020) who show how motivated reasoning

of observations, with the belief that justifies her preferred outcome. In our model a similar feature arises as voters seek the belief that is most likely to explain the state, while parties are polarised ideologically and offer a model to voters that trades off ideological preferences and a preference for seeking the truth. Our results also hold when voters themselves have some ideological preferences, as we show in Section 5. As in our work, in Little (2019) it is also the case that when voters have less information, their directional motives matter more.

3 The model

We consider an environment in which at each period a party is elected to govern and implements her campaign promise. Voters want to maximise some common outcome (e.g., GDP or crime prevention) but are uncertain about the mapping between the implemented policy and these outcomes. Two policy-motivated parties compete in the election. Given their knowledge, which is based on a short term memory of previous policy outcomes, voters balance the optimality of the policy platforms and their own idiosyncratic attachment to each party. Below we explain the model in detail, starting from the common outcome.

3.1 The common outcome and feasible policies

We consider an environment with a simple linear mapping between policies and outcomes. There are two policy dimensions, l and r. Let $(x_{l,t}, x_{r,t})$ be a vector of policies implemented at each period. The common outcome, y_t, is determined by,

$$y_t = \beta_l^* x_{l,t} + \beta_r^* x_{r,t} + \varepsilon_t,$$

where ε_t is iid across time and normally distributed with zero mean and variance σ^2. The parameters $\boldsymbol{\beta}^*$ are unknown to the voters; the voters hold a uniform prior over some compact set of feasible parameters $B = [0, \bar{\beta}]^2$ for some, sufficiently large, $\bar{\beta} > 0$.[9] In the main part of the analysis we consider a fixed $\boldsymbol{\beta}^*$, while in Section 4 we show how polarisation is exacerbated when the environment shifts (that is, when $\boldsymbol{\beta}^*$ changes over time).

We assume a discrete set of policies, to simplify the exposition (our qualitative results are maintained when we consider a continuous set of policies). In particular, at every period t, parties can offer one of the following three fixed policy vectors, (x_l, x_r):

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

weakens politicians' accountability.

[9] The analysis can also be generalised to data generating processes that are not linear.

The above represent three policy regimes, one biased towards policy dimension l, one towards r and a compromise policy that invests in both policy dimensions.

At every period t, given some policy $p_t \in \{L, M, R\}$, voters gain utility from the policy through common outcome y_t, which is simply defined as

$$V(p_t) = y_t - c(p_t),$$

where the cost of the policies satisfy $c(M) < c(L) = c(R)$.[10] We normalize $c(M) = 0$, $c(L) = c(R) = c$. We assume for concreteness that at the true state $\boldsymbol{\beta}^*$, the optimal policy which maximises $V(y)$ is M, but our results can be generalised to the other cases.

3.2 Political competition and voters' short-term memory

There are two parties, each identified with an 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 $p \in \{L, M, R\}$, $U_R(p)$ and $U_L(p)$, are therefore $U_R(R) = 1$, $U_R(M) = \frac{1}{2}$, $U_R(L) = 0$ and $U_L(p) = 1 - U_R(p)$. In addition, holding everything else equal, parties enjoy small office-rents when they win the election. During the campaign, at any period t, each party $J \in \{L, R\}$ chooses a policy vector $p_t^J \in \{L, M, R\}$, which is can commit to.[11]

At each period t, the voters are aware of the history of 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. We define voters' short-term memory at period t to be composed of a uniform prior over B, and the history H_t.

Given the information they have, voters are inclined to vote for the party that is more likely to generate a higher utility given its platforms. In our main model all voters observe the same information and hence agree on what platform provides a higher utility. Voters may differ however in their attachment to a party. We summarise it by focusing on the median voter, with a median bias towards L denoted by ϕ, votes for party L if

$$EV(p_t^L|H_t) - EV(p_t^R|H_t) + \phi > 0$$

where ϕ is uniformly distributed on $[-\frac{1}{2\zeta}, \frac{1}{2\zeta}]$, for some $\zeta > 0$.[12] Voters balance then their bias in favour of one of the parties with their desire to elect the party that pushes forwards

[10] This is a standard assumption of convex costs, so that it is less costly to produce a combination of policies. Note that without such an assumption, M will never be optimal.

[11] Parties may or may not know the true $\boldsymbol{\beta}^*$. Our results will hold under both assumptions.

[12] In case of an equality, we assume that the voter votes for party L with probability 0.5.

a platform that it more likely to generate a higher utility, given their (limited) historical knowledge.

3.3 A preliminary Lemma and assumptions on parameters

Before we introduce the dynamic model, we state a simple Lemma that highlights the static forces that play a role in parties' decisions at every period to polarise or converge.[13]

Lemma 1 (Consensus vs Polarisation): At period t,

(i) If $p^J = J$ maximises $EV(p|H_t)$, *then it is a dominant strategy for party* J *to offer its ideal policy.*

*(ii) I*f for some p and all $p' \neq p$, $EV(p|H_t) - EV(p'|H_t) \geq \frac{1}{2\zeta}$, *then in the unique equilibrium at period* t *involves convergence on* p.

(iii) There is no convergence equilibrium on policy p *if for some policy* $p' \neq p$, $EV(p|H_t) - EV(p'|H_t) < \frac{1}{2\zeta}$.

Part (i) establishes that if a party can maximise the utility of voters using its own favourite policy, then it will surely offer its policy in equilibrium. Part (ii) considers the case in which voters consider p to be sufficiently superior so that no attachment shock can allow another policy to win with any probability. In this case both parties must offer p. Part (iii) shows a converse statement: Convergence on some p cannot arise, if a party can win the election with even a small probability when it switches to another p'.

Given the above, we make two assumptions about ζ, which measures the variance of the attachment shock. First, we assume that ζ is sufficiently large, which implies that the attachment shock, is sufficiently small:

$$E_{\boldsymbol{\beta}^*}V(M) - E_{\boldsymbol{\beta}^*}V(p') > \frac{1}{2\zeta} \text{ for all } p' \in \{L, R\} \rightarrow \zeta > \frac{1}{|\beta_r^* - \beta_l^*| - c}$$

This implies that when the state is known, both parties must converge to offer the optimal policy M, as follows from Lemma 1(ii). This assumption will imply that long term variation in policies will lead to consensus on the correct outcome.

Our second assumption insures that ζ is not too large, so that when knowledge is sparse about the optimal action, parties will find it in their interest to polarise. Let $E[y|p, \boldsymbol{\beta}]$ denote the expected outcome y given some policy $p \in \{L, M, R\}$ and parameters $\boldsymbol{\beta} = (\beta_l, \beta_r)$ (that is, absent the noise ε). Specifically, $E(y|L, \boldsymbol{\beta}) = \beta_l$, $E(y|R, \boldsymbol{\beta}) = \beta_r$, and $E(y|M, \boldsymbol{\beta}) = \frac{1}{2}(\beta_l + \beta_r)$. For a policy p, let $B(p) \subset B$ denote the set of parameters $\hat{\boldsymbol{\beta}}$ that satisfy

$$E[y|p, \boldsymbol{\beta}^*] = E[y|p, \hat{\boldsymbol{\beta}}],$$

[13] These all rely on familiar arguments in the literature, see for example Calvert (1985).

that is, the set of parameters that is consistent with the observations generated after a long time in which p is implemented. We then assume:

$$\begin{aligned} E_{\beta\sim U[B(p)]}V(p') - E_{\beta\sim U[B(p)]}V(p'') &< \frac{1}{2\zeta}, \text{ for all } p, p', p'' \in \{L, M, R\} \rightarrow \\ \zeta &< \frac{1}{\bar{\beta}^2 - (\beta_l^* + \beta_r^*) - c} \end{aligned}$$

which implies that if some policy p has been implemented for a long time, voters' updates beliefs, which are necessarily updated to be uniformly distributed on $B(p)$, allow for all policies brought forward in an election to win at least with some small probability.[14]

3.4 Dynamics

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. Outcome y_t is observed and history is updated so that $H_{t+1} = \{p_\tau, y_\tau\}_{\tau=t-K+1}^{\tau=t}$.
4. The two parties offer p_{t+1}^J.
5. At period $t+1$, ϕ is drawn and party L wins the election if

$$Eu(p_t^L|H_t) - Eu(p_t^R|H_t) + \phi > 0$$

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

In equilibrium, in any period, parties are myopically best responding to each other, while anticipating the voters' behaviour. Moreover, the voters update their beliefs and compute their expected utility from each platform given the observed $K-$ period history. While it is possible to accommodate parties choosing policies strategically rather than myopically to win the current election, this will add little to the key results.

4 Cycles of convergence and polarisation

In this Section we analyse the dynamic model presented above. We show below that the equilibrium involves cycles of consensus and polarisation in the policy choices of parties. But first we start with a benchmark 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 platforms.

Proposition 1: *Assume that $K = \infty$. Then in any equilibrium, almost surely both parties offer the same platform in the long run.*

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

It is instructive to think about our model of political competition in the case of full memory by noting that the beliefs of a rational observer must converge. If they converge to a set that allows only for one policy to be optimal, both parties must offer the same platform by Lemma 1(ii). Thus convergence in policies arises in this case. We note that convergence is not guaranteed to be on the optimal policy M; as in any one-agent decision problem, learning can sometimes be wrong, due to insufficient experimentation arising from myopic behaviour. For example, a rare series of bad shocks early on may convince voters that the correct policy is an extreme one and parties have to conform to that.

Alternatively, If in the long term belief of the voters converges to a set that allows for more than one policy to be optimal, then the self-interested parties will polarise. In the long term however, "experimenting" with two different policies will allow the voters to learn the truth; in this case parties must choose the optimal policy, rendering such an equilibrium infeasible. Thus voters' belief in the long term must allow for only one policy to be optimal and parties must converge to offer this policy.

4.1 The main result

We now turn to consider finite K. Given the attachment shock and the outcome noise, the dynamics of policy on equilibrium path is stochastic and, coupled with short term memory, voters will never fully learn the state of the world. Moreover, for finite histories, the nature of voters' data can change over time. If for example power did not change hands or parties' platforms are very similar, this will imply that history contains very little variation in policies. If alternatively the history involves a high frequency of changes in policies, voters' data will be relatively informative.

These changes in the nature of history that arise with short-term memory imply the potential for cycles of policy change. Following histories with little variation in policy, the 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. Alternatively extended periods of political turnover and polarisation will imply the opposite; histories now 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, parties will be drawn to put forward models that are closer to the truth to avoid electoral defeat.

Our next result formalises the intuition above:

Proposition 2: *For a large enough K, almost surely, the polity experiences perpetual cycles of convergence and polarisation.*

cycles diagrams

$(2)_{P}age_{3}$

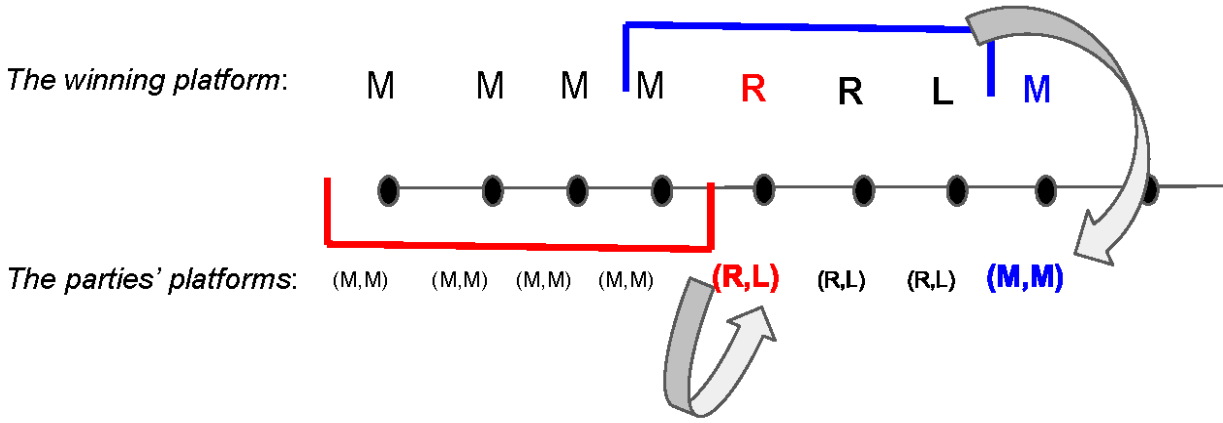

$_{5}.png$

Our key result is that the endogeneity of policies, together with short-term memory, leads to cycles of polarisation and convergence. Again, it is instructive to think about the public's learning process; short-term memory hampers learning, opening up the possibility for more than one policy to be optimal. This is when parties' self-interest kicks in, allowing these to polarise. But polarisation in policies ensures sufficient variation in what the public observes. Together with a memory that is not too short, this will allow voters to learn more precisely, and at some point the parties must converge.[15] From then on, the short-term memory implies that after at most K periods, only the convergence periods are recalled. Voters' history contains then little variation in policies. This dampens previous learning, allowing parties to espouse different policies again, and so on, oscillating between learning and unlearning. Figure 5 illustrates a cycle with $K = 4$:

Note that we simplified our model by focusing on only three policies. Our result generalises to the case of continuous policies (e.g., a one-dimensional space where we move from fully investing in the l dimension to fully investing in the r dimension). 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 due to sufficiently different policies, 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. As a result, after periods of close by policies, parties will start polarising more in the direction of their ideal policies. Therefore, a model with continuous policies will still involve cycles of convergence and polarisation.

[15] Convergence is 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.

Remark 1 (How stagnation breeds polarisation): In our model, dampening of knowledge arises when parties converge and so there is insufficient variation in the historical implementation of policies which allows parties to pursue then their self-interest. Another avenue for such insufficient variation can also arise in the polarisation phase, when one party keeps on winning. For example, if for some reason voters' attachment parameter ϕ is persistent in a particular direction for a few periods, then in periods of polarisation, we are more likely to observe the same party winning and hence the same platform implemented over and over again. This implies less variation in policies and so also after periods of such stagnation, polarisation is likely to persist and the consensus phase of the cycle will be delayed.

Our main result establishes cycles for a large enough K. For lower K cycles will still arise with a strictly positive probability, which depends on the other parameters of the model, namely, the level of political uncertainty (the attachment shock ϕ density) and the variance of the shock ε. To illustrate, consider the case where the variance of the shock, σ^2, goes to zero. This implies that learning is very fast; in the limit observing one period is enough to fully learn the state of the world. Analysing this case is informative about the dynamics of the model and will come useful when we extend the model in later sections.

Proposition 3: *Let $\sigma^2 \to 0$. Then, for $K \geq 2$: (i) The convergence phase of the cycle will have both parties offer M and will last K periods; (ii) The polarisation phase of the cycle where party R (L) offers policy R (L) lasts for only one period.*

Once parties moderate, the lack of variation in policies after K periods implies a complete "reset" in learning. This will cause parties to polarise again. Polarisation however 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 and therefore convergence of both parties to offering the correct policy M. 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 in fact for a small K, if the variance of the shock is small enough.

4.2 The effects of crises

We have identified a fundamental reason for polarisation and consensus cycles, that rely on short-term memory when the public attempts to learn what is the true data generating process. While learning lays the foundation for cycles, there are many features of the political process that can increase or decrease the length of each phase of the cycle. We first model and discuss the effects of crises, and then consider other such features.

4.2.1 Crisis I: The effect of unexpected large shocks

Large random shocks constitute "surprising" events, where the outcome y_t is very large or very small. Following such a rare event, whatever knowledge was accumulated before, voters' beliefs may substantially change.

Indeed, a large (and negative) shock often triggers a crisis in the political and economic realm. A shock to oil or real estate prices, a natural disaster such as an earthquake or a pandemic, reduce output dramatically. There is a large literature that looks at the effect of crises on political decision making and on polarisation among voters. For example, Mian et al (2014) show how voters become more politically polarized and fractionalized following financial crises.[16] In other situations such as sudden external threats, Myrick (2021) shows that a polity may unite or divide, depending on the domestic political environment in which they are introduced.

In our model a large shock will substantially affect the dynamics of the convergence and polarisation phases and will have an effect that can last for K periods. Moreover, whether a crisis triggers polarisation or not, depends on the current status quo policy. In fact, a polarised policy accompanied by a large shock leads to more periods of polarisation, while a moderate policy accompanied by a large shock leads to more periods of moderation:

Proposition 4: *Given a sufficiently large one-off* $|\varepsilon_t|$ *: (i) If* $p_t = M$*, then almost surely for the next* K *periods, the equilibrium policies are* (M, M)*. (ii) If* $\varepsilon_t > (<)0$ *and* $p_t = L$*, then almost surely the equilibrium for the next* K *periods exhibits polarisation, with* L *choosing* $L(M)$ *and party* R *choosing* $M(R)$*. (iii) If* $\varepsilon_t > (<)0$ *and* $p_t = R$*, then almost surely the equilibrium for the next* K *periods exhibits polarisation with* R *choosing* $R(M)$ *and party* L *choosing* $M(L)$*.*

To see the intuition, note that a very large shock implies substantial learning, resulting from the fact that the normal distribution on ε has "thin tails". Thus a large shock increases dramatically the likelihood of extreme models. If M was implemented, then if the shock was positive it indicates that both β_r and β_l are very large, and if the shock was negative it indicates that with a high likelihood both are very small. In either case, the models that justify M become much more likely, and this overrides any previous learning, as well as any future learning for the next $K-1$ periods. Thus, a moderate policy accompanied by a large shock will increase moderation.

If some other policy was implemented, let's say L, then a very large positive shock indicates

[16] The polarisation of voters they identify also sometimes creates a gridlock in some political environments that rely on coalitional governments (see Alesina et al 2006, Drazen and Grillo 1993, and Mody and Abbdul 2005).

that β_l is very large. This means that both L and M can be justified with the highest likelihood (using the highest possible $\hat{\beta}_l = \max \beta_l$). A very large negative shock indicates that β_l is very small, allowing now for both R and M to be considered optimal (using the lowest possible $\hat{\beta}_l = \min \beta_l$). Moreover, again, a large enough shock implies that no matter what finite set of $K-1$ events follow, the likelihood ratio between these extreme models and others persist for the next K periods (as long as no other rare event had occurred). Thus once a large shock hits during the polarisation phase of the cycle its effect is prolonged substantially.

Remark 2 (A cohort effect): The fact that a large shock can have long lasting effects is also related to a "cohort effect" explored in the literature. This is defined as an event that can disproportionately affect a whole generation's preferences and beliefs (e.g., the effect of the Vietnam war on the US young generation at the time). One additional example of the effect of a shock that persists in this way is the tax revolt in California in the late 1970s and the passage of Proposition 13 in the state. The crisis started when sharply rising property prices implied that many households could not afford to pay their property taxes. This had led to tax revolts and substantial tax reductions. Some consider this shock to be instrumental in Ronald Reagan's US election win a few years later. Piketty (2020, pp. 836) remarks on this historical episode: "Apart from its importance in the rise of Reaganism, this episode is interesting because it shows how very short-term phenomena... can combine with longer-term intellectual and ideological failures... to produce major political change."

4.2.2 Crisis II: Changes in the true data generating process

Above we interpreted a crisis as arising from a large negative shock that hits the economy. However, a change in the state of the world, such as a technological change, along with policies that are not tailored to this new state, can often lead to a crisis as well. For example, following the financial crises investors and governments realised that the effects of financial innovations have not been fully understood. As a result, investment banks, government 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 was deemed less relevant (see....). We now accommodate in our model the possibility that the state of the world (the data generating process) changes over time. We show how this increases polarisation, whether voters are aware of such a possible change in the data generating process or not.

An easy 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 $\boldsymbol{\beta}_t^*$ from the set B using the uniform prior distribution. To simplify, assume that these are iid draws over time, but our analysis can also accommodate serial correlation over time.

When considering changing worlds, we need to take a stand about whether the voters are aware of this possibility. If voters are aware once the state of the world has changed at period t, then history becomes irrelevant. In the language of our model, this means that $K = 0$ at period t. Once memory is reset, parties can then polarise. Thus in this simple case, polarisation will be manifested more often.

A more realistic case is that the voters or parties are not aware of this change or only understand it with a lag. Again, even with a lag, once voters realise that the state had changed following a crisis, they will only find the recent history relevant implying that they assess different models according to some $K' < K$, which is at least weakly more likely to lead to polarisation; there will be (weakly) less variation in policies the shorter is the history considered. Note that in fact even if the state had not changed, politicians may argue that it did change, and use this narrative in order to foster polarisation. We discuss such political narratives in Section 5 below.

A final possibility is the case where voters are not aware at all that the state of the world can change. Voters then have a misspecified model: They believe that $\boldsymbol{\beta}_t^* = \boldsymbol{\beta}^*$ for all t or in other words that $\lambda = 0$ (while in reality $\lambda > 0$), and do not change this assessment over time. That is, they simply do not consider the possibility of a change in the fundamental environment. Even in this case, and even when the noise ε is minimal, polarisation can increase:

Proposition 5: *Let $\sigma^2 \to 0$, and $K > 3$. Then there exists a small enough ζ, so that: (i) The higher is λ, on average, the shorter is the moderation phase of the cycle and the longer is polarisation phase of the cycle. (ii) When $\lambda = 1$, polarisation is indefinite.*

To see the intuition, let us consider the extreme $\lambda = 1$ case when the environment changes every period. It then becomes impossible to explain with certainty the observed outcomes; voters have a wrong model in the sense that they do not perceive changes to the environment, and the fact that the noise is minimal implies that their beliefs cannot evolve to fully explain their observations. As a result, even when the noise is minimal, voters' knowledge is sufficiently sparse and parties can take the chance and polarise, when $\frac{1}{2\zeta}$ is sufficiently small. In that case, any policy can have even a small chance of winning and no convergence can arise, in line with Lemma 1(iii).

4.2.3 Discussion

We now consider other features of the political system that increase or decrease polarisation.

Sticky policies: It is often the case that policies on some dimensions are sticky and not easy to change. This contributes further to little variation in the observed data and hence

fosters polarisation. In a richer model that allows for multidimensional policy space, we expect more polarisation to arise on these dimensions where policies are more likely to be harder to change for technical or bureaucratic reasons.

Learning from other countries/states: In our model for example, information about the state of the world is only gathered from endogenous outcomes. In some cases, the public may receive external signal about the true data generating process. For example, in the case of a global pandemic, as other countries face similar environments, the public can learn from the experiences of these other countries. External signals that allow more learning will hasten convergence and moreover, such convergence is more likely to be on the optimal policy. Still, short term memory implies that polarisation will always arise at some stage.

5 Extensions: Selective histories and echo chambers

Our model is built on voters explaining the world according to their observed history, relying on short term memory. However, it is often the case that politicians attempt to manipulate this memory. For example, they might highlight particular historical experiences or alternatively attempt to convince voters to ignore others. Additionally, the presence of echo chambers and the ability to target specific voters implies that parties may highlight different histories to different voters. A party which can communicate with its voters exclusively, which is often the case in an echo chamber, will find it easier to manipulate memory as its voters are not confronted with any conflicted historical accounts.

As short-term memory is concerned, some aspects of memory are not necessarily manipulated strategically. For example, some voters may nostalgically remember "good" periods even without being prompted by politicians. Alternatively, some voters might have experienced a crisis in their formative years and this "cohort effect" event influences them disproportionately compared to other periods. Still, politicians who strategically manipulate memory may find it easier to do so for such events when they use familiar narratives about the "good old times" or alternatively about the fact that "the world has changed". We first discuss informally such narratives and then construct a model of parties with echo chambers and apply it to the intergenerational conflict between the old and the young.

5.1 Narratives for selective histories

We discuss two types of narratives:

Nostalgia: Nostalgia is often used in political debates and is prevalent amongst voters.[17]

[17] There is a recent literature that studies nostalgic memory in politics. See Kenny (2017) for a theoretical

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 attribute these events to a high parameter value relating to their desired policy.

A changing world or anti-nostalgia: An alternative to nostalgic selection of histories, which points at periods in which outcomes were good conditional on some desired policy, 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 the world is evolving; that is, the data generating process is changing. This implies that we need to "forget" these periods as they are not relevant. Politicians can convince voters to ignore periods in which their desired policy was implemented but the overall outcome was disappointingly low. As opposed to our analysis in Section 5, such narratives can be attempted by politicians even if the state of the world had not changed (or at least even when the politicians are not aware of such change).

A good example of the use of selective histories in politics is documented in Kenny (2017) who studies current populism through the prism of populist movements in history. Kenny (2017) documents how in the late 1960s Great Britain, Enoch Powell used both nostalgia and anti-nostalgia to make his case for the future of the country. For the case of anti nostalgia, Kenny (2017) writes, "Powell's signature contribution to these debates was to denounce the illusions and overdeveloped sense of international responsibility associated with nostalgia for empire. The notion that Britain retained enduring responsibilities to the inhabitants of Commonwealth countries, who had long been treated as subjects of the British crown, was dismissed as a soft-headed piece of romanticism, fuelled by nostalgia for imperial glory."

5.2 Echo chambers, selective histories, and the intergenerational conflict

We now consider the effect of echo chambers on polarisation. 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. 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 now consider a simple extension in which voters observe information through the filter of parties. Mainly, we allow different parties to highlight

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.

different periods and so select the history it offers to its voters.

5.2.1 Extending the model to allow for echo chambers and selective histories

Assume that there are two groups of voters, one attached to party L and one to R. The key assumption here is 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). 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. Using some costly voting model, a voter in group L is then more likely to vote for party L the more she is able to justify its policies using the history $H_t^{k^L}$:

$$EV(p_t^L|H_t^{k^L}) + \gamma_L,$$

where $\gamma_L > 0$ is some affinity shock that accords with these voters' preference for 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 $H_t^{k^R}$:

$$EV(p_t^R|H_t^{k^R}) + \gamma_R,$$

where again $\gamma_R > 0$ is an affinity shock that accords with their preference for R. Let $\phi = \gamma_L - \gamma_R$ and so now party L wins the election iff its share of voters is larger, which means whenever:

$$EV(p_t^L|H_t^{k^L}) - EV(p_t^R|H_t^{k^R}) + \phi > 0$$

This allows us to easily extend our previous model to allow for parties to select specific periods to justify their policies and voters are not confronted with the other party's information. The only difference now is that parties can choose different histories to support their model and hence policies.

It is easy to see that the above modification of our model implies that polarisation will be more pronounced the less restrictive is the number of periods that a party needs to explain. Assume for example that voters are satisfied with a short history of two periods. The left party may be able to choose two periods in which it implemented L and y_t was high, and analogously for the right party, which will foster polarisation. Moreover, the larger is the history they can draw from, the more flexibility it allows them to pick and mix periods. We next analyse a specific example in which the two groups are not symmetric in the history they remember and such echo chambers increase one-sided polarisation.

5.2.2 The advantage of being old

Consider the following extension of our basic model to two groups of voters, old and young. Specifically, now party L represents the young voters and party R represents the old voters.

For example, we can think of policy L as increased investment in education and employment opportunities, or alternatively as subsidies for green energy, whereas policy R can be thought of as more resources spent on pensions and social care or alternatively on traditional sources of energy. The conflict between these two groups is relevant for our model for the reason that the old can potentially recall a larger set of histories that they have actually lived through. The young have not experienced all these histories first hand. 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.

Specifically, we assume that each party needs to select $k < K$ periods of history to show to their group of voters. The old can recall K periods and thus the R party can choose any k periods in $\{t-K, t-1\}$, while the young remember only the last k periods and thus the L party is constrained to provide policy that should be a good fit for beliefs that are updated given the observations of the most recent k periods. This feature captures the asymmetry between the young and the old, where the party that represents the old can manipulate the observed history by selectively "reminding" its voters about different periods.

How does the ability of the old to be more selective in terms of the evidence they provide to their voters affect their election platforms? Intuitively, the young are bound by the recent history. If the recent history includes different types of implemented policies, they are then more likely to be bound by the true model and thus moderate. The old on the other hand can pick and choose from a longer history and present a more homogenous set of policies allowing for less learning and thus enabling polarisation on their side, as well as to choose favourable periods that justify their policies. However, histories are also endogenous and so the additional ability of the old to pick and mix histories depends on what the young and the old offer to voters.

In fact, we show that the endogeneity of policies exacerbates the advantage of the old. For any given history, given the discussion above, the young will tend to moderate more often and the old will tend to polarise more often. But the young moderating and the old polarising implies that any recent history will include sufficient variation and thus will necessitate the young moderating even with a higher probability or more often; similarly, the old may be able to find a homogenous history composed of either all moderate policies or all their own policy, which will enable them to polarise even more. Thus the young moderating with a high probability enables the old to polarise, and the old polarising implies that the young are bound to moderate. As a result we have:

Proposition 6: *Let $\sigma^2 \to 0$. (i) For any k, there exists a large enough K', such that for all $K \geq K'$, the old always choose R. (ii) For a high enough k and a high enough K, the*

unique equilibrium has the young always moderating and the old always polarising.

Demography plays an important role in recent political conflicts, be it climate change or Covid-19. One of the key parameters with a high predictive power for voting to leave the EU in the UK Brexit referendum in 2016 was identified to be age.[18] Populations across Western countries have a demographic bias towards older voters, which also tend to vote more. Our model above identifies another reason for the success of parties that are more aligned with old voters in electoral campaigns: Their ability to be more selective in their choice of what histories to shed light on. Counter-intuitively, old voters are manipulated in the sense that the histories chosen allow their parties to deviate from the truly optimal policies but this manipulation is better for these voters in terms of serving their "material" interest.

6 Conclusion

We provide 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. We show how crises can alter the cycle dynamics and can have lasting effects. 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: (i) 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), note that the premise implies that a party that offers p wins for sure unless the other party offers p as well in which case each wins with probability 0.5. Therefore both choosing p is an equilibrium as deviating from p yields the same ideological outcome and no office-rents. This implies that there is no divergent equilibrium in which

[18] See Becker et al (2017).

only one party chooses p. Clearly both offering some $p' \neq p$ is not an equilibrium as p will be preferred ideologically by one of the parties and hence this party will deviate. We remain with the case of a divergent equilibrium in which neither party offers p. Suppose $p = M$. If each party offers its ideal policy then, one wins with a probability of at most $1/2$ and hence has an expected utility of $1/2$. If it deviates to $p = M$, it gains a utility of $1/2$, as it wins for sure, and in addition gets some office rents and hence this cannot arise as an equilibrium. Suppose $p = L$. But then by (i) L must offer L and so such an equilibrium cannot arise and similarly if $p = R$. (iii) Assume both parties offer p. As one of the parties must ideologically prefer some p', if it deviates, it can win with some small probability and gain utility from this outcome and hence no convergence can arise.■

For the following proofs, as defined in the text, denote the expected outcome when policy p is played and parameters β are the effectiveness parameters 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 $\frac{1}{2}\beta_l + \frac{1}{2}\beta_r = \frac{1}{2}\beta_l^* + \frac{1}{2}\beta_r^*$. Note that when $p \neq p'$, $B(p) \cap B(p')$ is a singleton and includes only β^*.

Proof of Proposition 1:

Step 1: Our model satisfies Assumptions 1-3 in Esponda et al (2020), henceforth EPY.[19] Given H_t, let σ_t be the distribution over implemented actions, that is, the share of time L, M, R were implemented. The Kulbeck-Liebler (KL) divergence value of some vector of parameters β given some σ_t is defined as as

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

where $f(\varepsilon)$ is the density over ε, assumed normal with mean zero in our case. Let $\beta_{\min}(\sigma_t)$ be a minimizer of the above and let $K(\sigma_t) \equiv KL(\beta_{\min}(\sigma_t)|\sigma_t, \beta^*)$. Theorem 1 in EPY implies that in the limit, when $t \to \infty$, the posterior distribution over the beliefs, μ_{t+1}, will concentrate on the set B_∞^σ, so that for any $\beta \in B_\infty^\sigma$, for which $\mu_{t+1}(\beta) > 0$, their KL value is close to $K(\sigma_t)$, that is:

$$\lim_{t \to \infty} \int_{\boldsymbol{\beta}} |KL(\beta|\sigma_t, \beta^*) - K(\sigma_t)| d\mu_{t+1}(\beta) = 0$$

This implies that for any β whose KL value is not close to $K(\sigma_t)$, it must be that $\mu_{t+1}(\beta) \to 0$. Note that this result holds even if σ_t does not converge. We will use this step in Step 3 below.

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

Step 2: A Bayesian updater who observes the history H_t satisfies the conditions of the martingale convergence theorem. Therefore, beliefs converge almost surely. Let the support in the limit be denoted by B_∞. This implies that the likelihood ratio satisfies $\frac{\mathcal{L}(\boldsymbol{\beta}|H_t)}{\mathcal{L}(\boldsymbol{\beta}'|H_t)} \to \infty$ for all $\beta \in B_\infty$ and $\beta' \notin B_\infty$.

Step 3: Consider paths on which B_∞ is such that both some unique $J = R, L$ as well as M are in $\arg\max E_{B_\infty} V(p)$. In the limit, for this to constitute an equilibrium, party J must polarise on policy J and the other party will moderate. Each of the parties is elected in equilibrium with a strictly positive probability. Consider a strictly positive measure of paths with a subsequence $\sigma_{\tau^t} \to \sigma$ satisfying $\frac{\sigma_{\tau^t}(\mathbf{J})}{\sigma_{\tau^t}(\mathbf{M})} \to c$ for some finite non zero c. This means that along τ^t, $\beta^* = \arg\min_{\boldsymbol{\beta}'} KL(\beta'|\sigma_{\tau^t}, \beta^*)$, and $KL(\beta^*|\sigma_{\tau^t}, \beta^*) = 0$. 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 $B_\infty = \{\beta^*\}$. This contradicts the supposition of a path with such a subsequence $\sigma_{\tau^t} \to \sigma$ as both parties must moderate. This implies that such a subsequence cannot arise with a strictly positive probability implying that in the limit, parties are not best responding, a contradiction to the equilibrium existence. Similar analysis rules out situations in which B_∞ includes beliefs that allow for more than one policy to maximise $E_{B_\infty} V(p)$. Thus, B_∞ can only include beliefs that give rise to a unique policy to maximise $E_{B_\infty} V(p)$.■

Proof of Proposition 2: Suppose that after some t in any period both parties offer the same policy p. This implies that at any period after period $t + K$, the history will include only observations of the policy p. Given any K observations for such periods t', $(y_s = E[y|p, \beta^*] + \varepsilon_s)_{t'}^{t'+K-1}$, the likelihood of some belief β, will be

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

To maximise the likelihood, we choose $\delta(p, \beta)$ to minimize $\sum_{s=t'}^{t'+K-1}(\delta(p, \beta) + \varepsilon_s)^2$, and hence choose $\delta(p, \beta)$ so that $\delta(p, \beta) = -\frac{1}{K}\sum_{s=t'}^{t'+K-1} \varepsilon_s$. For a large enough K, this implies that $\delta(p, \beta) \simeq 0$ is the minimizer of the KL value almost surely (as $-\frac{1}{K}\sum_{s=t'}^{t'+K-1} \varepsilon_s$ converges in probability to zero). This implies that far enough in the future, beliefs concentrate to be uniformly distributed on $B(p)$. Given our assumptions, there will surely be a period t'' far enough into the future for which polarisation will arise.

Assume that polarisation continues indefinitely, and that K is large enough. Note that throughout the polarisation phase, it must be that each party is elected with a strictly positive probability bounded away from zero (otherwise it will moderate or offer the same

policy as the other party). This implies that along this sequence, H_t is such that $\frac{\sigma_t(\mathbf{x})}{\sigma_t(\mathbf{x}')}$ is non zero and finite. We know from continuity and Proposition 1 Step 1, that for a large enough K, if for any t, this then implies that $\beta^* = \arg\min_{\boldsymbol{\beta}'} KL(\beta'|\sigma_t, \beta^*)$, $KL(\beta^*|\sigma_t, \beta^*) = 0$, and beliefs can only concentrate on an a ball around β^*. As a result, at some point, parties must moderate and so a polarisation phase cannot continue indefinitely for a large enough K. Note that scenarios such as both parties at some point offering the same policy $p \neq M$ cannot be ruled out, so with strictly positive probability we will also have periods of consensus with both parties choosing L or both choosing R. Together with (i), this implies cycles of convergence and polarisation.■

Proof of Proposition 3: Note that almost surely, for any history with $K-1$ periods,

$$\frac{\prod_{s=t}^{t+K-2} f(y_s - E[y|p_s, \beta])}{\prod_{s=t}^{t+K-2} f(y_s - E[y|p_s, \beta'])} < \infty$$

However,

$$\lim_{y\to\infty} \frac{f(y - E[y|p, \beta'])}{f(y - E[y|p, \beta''])} = \infty$$

for all $\beta' \in \arg\max E[y|p, \beta]$ and all other β' that do not belong to this set, and similarly

$$\lim_{y\to-\infty} \frac{f(y - E[y|p, \beta'])}{f(y - E[y|p, \beta''])} = \infty$$

for all $\beta' \in \arg\min E[y|p, \beta]$ and all other β'' that do not belong to this set. As a result, when a crisis occurs at period $t' - K - 1$, then at period t' and for the next $K-1$ periods, then almost surely $EV(p|H_t) \approx E_{\beta'}V(p|H_t)$ for on the respective β' in each case. If a policy M was implemented at period $t' - K - 1$, then $\beta' = (\bar{\beta}, \bar{\beta})$ in the former case and $(0, 0)$ in the latter, implying in both cases that M is the optimal policy. If on the other hand L was implemented, then a positive shock triggers an (L, M) equilibrium β' includes all beliefs with $\beta_l = \bar{\beta}$, and so on as detailed in the Proposition. Moreover, almost surely it lasts for K periods.■

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

Lemma A1: *(i) Assume that a policy p was implemented for one period. Then* $E[\frac{f(\delta(p,\beta')+\boldsymbol{\varepsilon})}{f(\delta(p,\beta'')+\boldsymbol{\varepsilon})}] = 1$ *for all* $\beta', \beta'' \in B(p)$, *and when* $\sigma^2 \to 0$, $E[\frac{f(\delta(p,\beta')+\boldsymbol{\varepsilon})}{f(\delta(p,\beta'')+\boldsymbol{\varepsilon})}] \to \infty$ *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* $\sigma^2 \to 0$, $E[\frac{f(\delta(p,\beta^*)+\boldsymbol{\varepsilon}_1)}{f(\delta(p,\beta')+\boldsymbol{\varepsilon}_1)} \frac{f(\delta(p,\beta^*)+\boldsymbol{\varepsilon}_2)}{f(\delta(p,\beta')+\boldsymbol{\varepsilon}_2)}] \to \infty$ *for all* $\beta' \neq \beta^*$.

Proof of Lemma A1: To see (i), note that by definition of $B(p)$, $E[\frac{f(\delta(p,\beta')+\boldsymbol{\varepsilon})}{f(\delta(p,\beta'')+\boldsymbol{\varepsilon})}] = \frac{f(\varepsilon)}{f(\varepsilon)} = 1$ for all $\beta', \beta'' \in B(p)$. Also, $E[\frac{f(\delta(p,\beta')+\boldsymbol{\varepsilon})}{f(\delta(p,\beta'')+\boldsymbol{\varepsilon})}] = E[\frac{f(\varepsilon)}{f(\delta(p,\beta'')+\varepsilon)}]$ where $\delta(p, \beta'') \neq 0$ as $\beta'' \notin B(p)$.

This implies that $E[\frac{f(\delta(p,\beta')+\boldsymbol{\varepsilon})}{f(\delta(p,\beta'')+\boldsymbol{\varepsilon})}] \rightarrow_{\sigma^2\rightarrow 0} \infty$ for any $\beta' \in B(p)$ and $\beta'' \notin B(p)$. To see this, note that $\frac{f(\varepsilon)}{f(\boldsymbol{\delta}+\varepsilon)} \rightarrow^{p}_{\sigma^2\rightarrow 0} \frac{f(0)}{f(\delta)}$ for some $\delta \neq 0$, and that $\frac{f(0)}{f(\delta)} \rightarrow_{\sigma^2\rightarrow 0} \infty$. To see (ii), note again that $E[\frac{f(\delta(p,\beta^*)+\boldsymbol{\varepsilon}_1)}{f(\delta(p,\beta')+\boldsymbol{\varepsilon}_1)}\frac{f(\delta(p,\beta^*)+\boldsymbol{\varepsilon}_2)}{f(\delta(p,\beta')+\boldsymbol{\varepsilon}_2)}] = E[\frac{f(\varepsilon_1)}{ff(\delta(p,\beta')+\boldsymbol{\varepsilon}_1)}\frac{f(\varepsilon_2)}{f(\delta(p,\beta')+\boldsymbol{\varepsilon}_2)}] \rightarrow_{\sigma^2\rightarrow 0} \infty$ for any $\beta' \neq \beta^*$ using similar arguments to the ones above.■

Lemma A2: *When $\sigma^2 \rightarrow 0$, (i) if only one policy was implemented in throughout the observed K periods of H_t, then a party polarises. (ii) if there are two different policies implemented throughout the observed K periods of the history H_t, then a party moderates.*

Proof of Lemma A2: (i) Given Lemma A1, when one policy p is implemented for K periods, then almost surely $EV(p|H_t) \approx E_{\beta\sim UB(p)}V(p)$. Therefore, by Lemma 1(i) each party polarises. (ii) When p and p' have been implemented, then given Lemma A1(ii), $EV(p|H_t)$ $\approx E_{\beta^*}V(p)$ and so by Lemma 1(ii) the unique equilibrium is (M, M).■

Proof of Proposition 4: Following Lemma A2, assume that both parties choose M for K periods. In the next period each party $J = L, R$, will polarise and choose policy J. Once the parties polarise one of them will win the election. As $K \geq 2$, whoever won the election there will be two consecutive periods with different implemented policies and so we will go back to the moderation phase. Note that the moderation phase will continue for exactly K periods.■

Proof of Proposition 5: We start with $\lambda = 1$. We now augment the definition of δ as the state changes: Let $\delta_{t+i}(\beta) = E[y_{t+i}|p_{t+i}, \beta^*_{t+i}] - E[y_{t+i}|p_{t+i}, \beta]$. Almost surely, for any sequence of policies, $p_t, p_{t+1}, ..., p_{t+K-1}$, and any sequence of $\beta^*_t, .., \beta^*_{t+K-1}$, we have that

$$E(\frac{\prod_{i=0}^{K-1} f(\delta_{t+i}(p_{t+i}, \beta') + \varepsilon_{t+i})}{\prod_{i=0}^{K-1} f(\delta_{t+i}(p_{t+i}, \beta'') + \varepsilon_{t+i})}) < \infty$$

To see this, when $K = 3$, generically, for any β, there exists at least one i for which $\delta_{t+i}(\beta) \neq 0$ and is bounded, and for $K > 3$, there exist at least $K - 2$ such is. Given the compactness of B, there exists then some number $X < \infty$ for which for all β', β'' :

$$\lim_{\sigma^2\rightarrow 0} E[\frac{\prod_{i=0}^{K-1} f(\delta_{t+i}(p_{t+i}, \beta') + \varepsilon_{t+i})}{\prod_{i=0}^{K-1} f(\delta_{t+i}(p_{t+i}, \beta'') + \varepsilon_{t+i})}] \leq X.$$

This implies that there exists a ζ small enough, for which $\lim_{\sigma^2\rightarrow 0}(EV(p|H_t) - EV(p'|H_t)) < |\frac{1}{2\zeta}|$ for any p, p', implying that any policy has a chance to win over any other policy, at least

with some (possibly small) probability bounded away from zero. By Lemma 1(iii) this implies that parties do not converge on the same platform.

Consider now λ that is bounded from one. For any sequence of K policies, with probability $1-(1-\lambda)^K$, there has been a change of state during this phase. Whenever a change occurs at a phase in which history consists both of moderate and polarised policies, the above applies. Whenever a change occurs and following the change history consists of only moderate policies, then consensus will arise. To see this, note for $\beta^*_{t+K-1} \neq \beta^*_{t+K-2}$ that there is a unique value of β that solves

$$\delta_{t+K-2}(M,\beta)=0, \delta_{t+K-1}(M,\beta)=0,$$

which we denote by $\hat{\beta}$. Thus

$$\lim_{\sigma^2\to 0} E[\frac{\prod_{i=0}^{K-1} f(\delta_{t+i}(p_{t+i},\hat{\beta})+\varepsilon_{t+i})}{\prod_{i=0}^{K-1} f(\delta_{t+i}(p_{t+i},\beta')+\varepsilon_{t+i})}] \to \infty$$

for any other β'. However, if consensus implies that $p=L,R$ are the right policies, or if a change occurs at any other phase in the moderation cycle or in the polarisation phase, then polarisation arises. As on average, for a sufficiently high K (e.g., $K>3$), a change is less likely to only occur at this event, then on average polarisation is then more likely. Thus, the higher is λ, on average, the shorter are the moderation phases of the cycle and the longer are the polarisation phases.■

Proof of Proposition 6: (i) Let $K'=3k$. This implies that we have at least one action, L, M, or R that has been implemented at least k times. Thus party R can choose those k periods. Following on from Lemma A2 part (i), this party polarises.

(ii) Assume that k is large, let $K=3k$ and so the old polarise. First note that what we have described is a long term equilibrium, as for a large enough k, it must mean by the law of large numbers that both R and M are observed in any k periods and hence the young must moderate perpetually. We now show that this is the unique one by showing that in the limit, besides paths of measure zero, we must converge to this equilibrium.

Consider some period $4k+1$. Assume by way of contradiction that the young do not moderate. Following on from Lemma A2, this must mean that only one policy was implemented in the previous k periods. By (i) this must be R. This can arise for large k only if: (i) The young offer R for a measure one of the k periods; (ii) The young offer potentially another policy for a positive measure of periods but win with probability close to zero. In case (i), by Lemma A2, at any period in $\{3k+1, 4k\}$, the young should either offer M or L, so this case cannot arise. This is because if the history of k recent policies only include R, then they can offer L, whereas if the history includes two different policies, then they will offer

M. In case (ii), by the proof of Lemma A2 and the argument above, either M or L must garner at least probability $1/2$ of getting elected and so this cannot arise as well (besides on a measure zero of paths).

Thus at period $4k+1$ the only equilibrium must be that the old polarise and the young moderate. This implies that the same logic above applies to the next period and so on.■

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] 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

[4] Alesina, A. and Rosenthal, H. (1995). Partisan Politics, Divided Government, and the Economy (Political Economy of Institutions and Decisions). Cambridge: Cambridge University Press.

[5] Azzimonti, M. and M. Fenandes (2018), "Social Media Networks, Fake News, and Polarisation", NBER working paper 24462.

[6] 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.

[7] 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.

[8] 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.

[9] Callander, S. (2011), "Searching for good policies", American Political Science Review, 105 (04), 643–662.

[10] Callander, S., F. Izzo and G. Martin (2021), "Ideological Competition", mimeo.

[11] Calvert, R. L. (1985), "Robustness of the Multidimensional Voting Model: Candidate Motivations, Uncertainty, and Convergence," American Journal of Political Science, 29 (1), 69–95.

[12] Drazen, A. and V. Grilli (1993), "The Benefit of Crises for Economic Reforms", The American Economic Review, Vol. 83 (3), pp. 598-607.

[13] Eliaz, K. and R. Spiegler (2020), "A Model of Competing Narratives", American Economic Review. Vol. 110, No. 12, pp. 3786–3816.

[14] 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.

[15] Kenny, M. (2017) Back to the populist future?: understanding nostalgia in contemporary ideological discourse, Journal of Political Ideologies, 22:3, 256-273.

[16] Levy, G., R. Razin and A. Young (2021), "Misspecified Politics and the Recurrence of Populism", mimeo.

[17] 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.

[18] Little, A.T. (2019), "The Distortion of Related Beliefs", American Journal of Political Science, Volume 63, Issue 3, Pages 675-689.

[19] Little, A. (2021), Directional motives and different priors are observationally equivalent, mimeo.

[20] Little, A., K. Schnakenberg and I. Turner (2020), "Motivated Reasoning and Democratic Accountability", mimeo.

[21] Malmendier, U., S. Nagel, and Z. Yang (2021), "The Making of Hawks and Doves", Journal of Monetary Economics 117, 19-42.

[22] Malmendier, U. and S. Nagel (2016), "Learning from Inflation Expectations", Quarterly Journal of Economics, 53-87.

[23] 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.

[24] McCarty, N., Poole, K., & Rosenthal, H. (2013). Political Bubbles: Financial Crises and the Failure of American Democracy. Princeton; Oxford: Princeton University Press.

[25] 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.

[26] 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.

[27] Piketty, Thomas, Capital and Ideology (Cambridge, MA: Harvard University Press, 2020).

[28] Rogoff, Kenneth. 1990. "Equilibrium Political Budget Cycles." American Economic Review 80: 21-36.

[29] Schlesinger, Arthur Sr. (1949). Paths to the Present. Macmillan.

[30] Schlesinger, Arthur Jr. (1999). The Cycles of American History. Houghton Mifflin Harcourt.