Journal of Public Economics 247 (2025) 105397

Contents lists available at ScienceDirect

Journal of Public Economics

journal homepage: www.elsevier.com/locate/pubec

The dynamics of cohort effect in politics☆

Gilat Levy*, Ronny Razin

Department of Economics, LSE, Houghton St., WC2A 2AE, London, UK

HIGHLIGHTS

- We propose a new theoretical framework that encompasses a dynamic social-learning model of politics, where cohort effects are endogenously derived from preceding generations' political decisions.
- We highlight the role of political experiences in shaping the beliefs of younger cohorts, which subsequently influence policy decisions as these individuals mature.
- We demonstrate how these dynamic intergenerational linkages lead to cyclical patterns of polarised and cohesive cohorts, as well as a cyclical pattern of parties' behaviour, oscillating between polarisation and consensus.
- We show how the phases of the cycles depend on demography parameters, and how transitory shocks can exert persistent influence on politics due to these dynamic linkages.
- We also present some suggestive evidence, using ANES surveys, showing that different cohorts' opinions exhibit different levels of variance, consistent with our theoretical framework.

ARTICLE INFO

Keywords:
Cohort effects
Polarisation

ABSTRACT

This paper investigates the dynamic ramifications of cohort effects on politics. We propose a theoretical framework that encompasses a dynamic social-learning model of politics, where cohort effects are endogenously derived from preceding generations' political decisions. This process underscores the role of political experiences in shaping the beliefs of younger cohorts, which subsequently influence policy decisions as these individuals mature. We demonstrate how these dynamic intergenerational linkages lead to cyclical patterns of polarised and cohesive cohorts. In the proposed model, cohorts emerging during periods of political consensus display less familiarity with optimal policies, resulting, due to random external shocks, in high variance of public opinions. Conversely, cohorts maturing amidst polarisation and political turnover demonstrate greater knowledge about optimal policies, leading to more cohesive public opinions. Notably, our model suggests that transitory shocks can exert persistent influence on politics due to these dynamic linkages. We also present some suggestive evidence, using ANES surveys, showing that different cohorts' opinions exhibit distinct levels of variance.

1. Introduction

Intergenerational conflicts are an important determinant of political change and understanding them is one of the challenges facing Political Economy today. Young and old cohorts differ in how they vote and whether they vote. Conflicts between young and old voters are manifested in a broad range of topics ranging from the traditional ones such as education or pension schemes to more recent debates such as how to deal with the climate crisis. Indeed, age was one of the best predictors of election results in the US in 2016 and the Brexit referendum in the UK.[1]

The study of generational differences has a long tradition in social sciences, mainly focusing on age and life-cycle events. Alwin and Krosnick (1991) analyse NES data and provide evidence for a life-cycle process in which attitudes start forming in the formative years (18–24), and gradually stabilise by the time individuals are 40. Other recent examples are Ortoleva and Snowberg (2015) who show how individual

☆ We thank the Editor and two anonymous referees for useful comments. Lachi Singh, Emiliano Rinaldi and Raffaele Blasone provided helpful research assistance.
* Corresponding author.
Email addresses: g.levy1@lse.ac.uk (G. Levy), r.razin@lse.ac.uk (R. Razin).

[1] 73 % of people aged between 18 and 24 voted to Remain in the European Union, compared with just 40 % of people aged over 65. In fact, the propensity to have voted Leave increases with age, with the three oldest age groups here voting leave and the three youngest voting to Remain (source: Statista, Department of statistics, the London School of Economics).

https://doi.org/10.1016/j.jpubeco.2025.105397
Received 28 May 2024; Received in revised form 3 May 2025; Accepted 6 May 2025
Available online 15 May 2025

opinions strengthen and polarise with age and Boxell et al. (2017) who document different levels of polarisation across age groups.

Beyond age and life-cycle effects, a recent strand of the empirical literature has focused on how events that happen during a cohort's impressionable and formative years affect their political attitudes. Malmendier and Nagel (2016) show how living with inflation in early adulthood years affects preferences and Malmendier et al. (2021) show how this affects old and young committee members at the FOMC. Ghitza et al. (2022) study generational voting in American presidential elections and show that voters' partisan preferences are weighted heavily by events that occur in a voter's teenage and early adult years.

If political memories and attitudes are shaped in a particular period in an individual's life, this implies that events that occur in the impressionable years of one generation, such as wars or pandemics, will affect its political attitudes. Naturally this is not restricted to exogenous events; the outcomes of implemented policies will also affect the attitudes of the younger voters who experience them in their formative years. But these policies are typically shaped by the political landscape of previous generations. This then implies dynamic linkages between cohorts' experiences and hence attitudes. In this paper we introduce a new theoretical framework to study these linkages and their effects on politics.

We assume that throughout the political process there is uncertainty about the true model generating observable political outcomes as a function of the implemented policies. For example, voters might attempt to learn whether growth is better achieved 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 that aims to achieve more consumption power (a position pushed by another). In our model voters learn from political outcomes in their formative years, inducing potential cohort effects. We use an overlapping generation environment where voters go through three phases: A learning phase in their formative years in which they accumulate their experiences and form their beliefs about the right course of action (but do not vote), and two voting phases, in which young and old cohorts vote. Policy is determined endogenously through the political competition between two ideologically-motivated parties. Voting decisions are based on voters' knowledge, acquired in their formative years, along with stochastic elements that influence voting. In this framework, voters from different cohorts may accumulate different levels of knowledge and beliefs; they might be exposed to different implemented policies as well as to different random shocks to political outcomes. Moreover, cohorts are linked dynamically through their experiences.

The first insight we highlight is that cohorts differ in their level of information (which will affect the strength of their beliefs), depending on the experiences they go through. A cohort that experienced "turbulent" times, with polarised parties and high political turnover, "samples" a rich set of different policies and can therefore become quite informed about the optimal policy. The shared informative history they observe will create a consensus among its members, and a relatively homogeneous voting behaviour. A good example is the war generation, as observed by Alwin and Krosnick (1991): "This exception may represent a unique cohort effect on attitude stability, reflected in its departure from the general pattern of partisan stability over the life course. It exists in the partisan attitudes of the cohort born 1939–46, who were 26–33 in the early 1970s. This cohort came of age during the 1960s, one of the most turbulent periods in recent political history, experiencing their 18th through 20th years between 1957 and 1966. It seems likely that this cohort may have entered the 1970s with highly crystallized political orientations, with attitude stabilities that, somewhat prematurely, achieved the highest possible level when they were 26–33 years old during the early 1970s."

In contrast, consider a cohort that had lived their formative years in an era during which parties were relatively in consensus (as for example the generation growing up around and after WWII). The set of policies this cohort observes is quite limited in scope and that might put some constraints on individuals' levels of information about the optimal policy: While these voters might have good information about the policies they have experienced, they might not necessarily know how good or bad those policies are compared to other policies they have not observed. This cohort will seem to be more polarised in their voting behaviour, as their votes are not based on a strong shared belief or experience, but potentially on some other random influences or less predictable factors (e.g., their identity or cultural upbringing).

In other words, experiences of political outcomes may affect not only the direction of the ideologies of cohorts, but also their distribution. While the literature typically focuses on the first moment of the distributions of cohorts' preferences or beliefs, we identify the possibility that second moments can differ, and one cohort can be more polarised than another. In Section 2 below we show such differences in the opinions of cohorts in the ANES data. We provide suggestive evidence that such differences are related to the level of information that each cohort holds and to the level of political polarisation in the impressionable years of the cohort.

Our analysis illustrates that the heterogeneity among cohorts that we identify above can indeed arise in equilibrium through the dynamic linkages across cohorts. We demonstrate a cyclical pattern in political outcomes and cohorts' distribution of opinions: Cycles of polarisation and consensus of party platforms, and as a result, cycles of cohorts that are more and less informed respectively. To see how this arises, note that once parties face cohorts that are relatively informed, parties will be disciplined to move towards the policy that voters believe is best, forming a consensus. But this phase of consensus in party positions will imply that the younger cohorts who now go through their formative years will have less variation in their data. These cohorts will grow up to potentially have less clear knowledge about the desired policies, which will in turn make it easier for parties in future periods to push their own different interests and offer different policies. Such party polarisation will then allow the following generations to gather more information, and so on.

We explore the determinants of differences in opinion between contemporaneous cohorts. We show that when older cohorts are more politically powerful (for example, when the population declines over time), the phases of the cycle (phases of party polarisation and of party consensus) are longer. This implies that consecutive cohorts are more likely to go through similar experiences of policies, which implies that such cohorts are more likely to have similar distributions of attitudes. When the younger cohorts are more politically powerful and parties cater more to their attitudes, the cycle phases are shorter and so contemporaneous cohorts are more likely to have different distributions of opinions.

The dynamic linkages between cohorts highlighted above have some important implications for the effects of one-off exogenous shocks on future polities. Economic crises, or other transitory shocks to political and economic outcomes, will have lasting effects: The shocks will influence the beliefs of impressionable voters, these voters will shape policies in the future, and these policies together with future shocks will influence future generations' beliefs. We illustrate these effects in several ways. We study the case of the baby boom, in which a transitory population shock affects the political outcomes of all future generations. We also use a simulation to illustrate how transitory shocks to political or economic outcomes can reverberate through to future generations and affect politics many periods later.

We contribute to the literature in several ways. First, we theoretically and empirically formalise a way to consider how cohorts can differ in the second moments of their distribution of beliefs or preferences rather than in the first moment, which was the focus of the literature cited above. Second, we provide a dynamic model of political linkages between cohorts, and how cohorts' attitudes are shaped by previous generations'

attitudes. The results illustrate how differences in second moments can arise endogenously and fluctuate over time. The theoretical framework can more generally inform empirical work on cohort effects.[2]

Our theoretical result on the cyclical nature of party politics (that is, polarisation or consensus across parties) relates to Levy and Razin (2025), in which we analyse a dynamic model without cohorts, and show how short-term memory of voters can imply such cycles.[3] In this paper we focus instead on cohorts' memories. Crucially, this implies that different voters have different memories, and in particular that these different memories take a particular form as we go through the generations. Looking at cohorts also provides the perfect platform to tap into the important and broader question of intergenerational conflicts, how these are resolved in the political arena, and how they are affected by demography and other parameters.

Finally, an additional theoretical contribution is to provide a model of political competition with groups that have different memories.[4] This allows us to connect the differences in knowledge of the different groups to the propensity of parties to polarise or converge.

2. Stylised facts: second moment cohort effects

To motivate our analysis in the paper, we use survey data from the ANES cumulative data set (see ANES, 2022) to show that there are differences between cohorts in terms of the second moments of the distribution of political opinions. Moreover, we show some suggestive evidence that these differences may be related to differences in political engagement and knowledge about politics as well as the level of polarisation and turnover in cohorts' impressionable years.

From the ANES set of questions we have selected all the questions that have some relevance to political opinions. This resulted in a total of 33 questions. To maximise the number of questions that we can include, we focused on the questions that have been asked consistently in all the surveys and at least from the 1970s onwards. We excluded questions that seemed less relevant (e.g., questions on attitudes towards Jews, Blacks, Catholics, Whites) as well as questions such as whether one likes anything about the Democratic party, and focused on clear questions about political preferences over policies.[5] This process resulted in ten questions on which the respondents provide their opinions. The questions are opinions on whether: The Government benefits all, The Government should provide health, The Government should provide jobs, The Government wastes taxes, The Government should invest in military, The Government should provide welfare, the US is better off if unconcerned about the rest of the world, Thermometer question about big business, Thermometer question about labor unions, Thermometer question about the Military, and finally a Thermometer question about trust in Government.[6] We have bundled individuals according to their birth year into different cohorts (we report below on cohorts that are composed of a ten-year band. In additional empirical exercises reported in Appendix B, we also checked five-year, eight-year and twelve-year bands which provided similar results).

We use the cumulative dataset which includes all answers throughout the survey years. Our aim is to check how the second moments of the distribution of opinion may vary across cohorts. To do so, separately for each of the ten questions, we aggregate the answers of a particular cohort across all survey years, and compute the Allison–Foster polarisation index for these answers.[7] The Allison–Foster index is the difference between the score for respondents with scores above the median and the score for respondents with scores below the median (see Allison and Foster, 2004). Focusing on differences from the median, rather than from the average, measures variance in a way that does not depend on the numerical scaling assigned to the categories, which is useful as it is not clear how respondents in surveys interpret the scale.

The Allison–Foster index is reported in Fig. 1; the first 10 figures report the index for the 10 questions. The last figure reports the polarisation index for all the thermometer questions aggregated.[8] As can be seen in Fig. 1, cohorts clearly differ in the second moments of the distribution of their opinions. Also, for some questions, a U-shaped pattern emerges across cohorts so that younger and older cohorts exhibit higher dispersion in survey answers.

To illustrate the potential mechanism highlighted in our model, we next look at data about the level of knowledge that different cohorts have about politics. It would be useful to understand what knowledge these voters have about the specific effectiveness of different policies, however such data is not available. Instead, as a proxy for such political knowledge, we chose from the ANES data several questions pertaining to general political engagement and knowledge of politics. These questions that can potentially reflect how informed cohorts are. We report in Fig. 2 the share of respondents that state that they "don't know" the answer, or report "no identification", by birth cohorts, for the following questions. The first three questions attempt to see if individuals know what positions named people hold. For example, "What job or office is held by Nancy Pelosi?" Generally then, as names change throughout the years, we denote the question: What job or office is held by [named Speaker of the House]. Similarly two other questions are: What job or office is held by [named Vice-President], and What job or office is held by [named Chief Justice]. Other questions are: Do you Approve/Disapprove of U.S. Congress, Do you Approve/Disapprove of Running U.S. House Incumbent, Thermometer question on the Supreme Court, and Do you know the difference between the major parties. The eighth question that completes the list is a question about how much one is interested in elections and we report in the last panel of Fig. 2 the share of those who are "not much interested."[9]

Again, the trend, which is consistent with the mechanism in our paper and its results, is such that older cohorts and the very young cohorts are less informed, with cohorts born around the 1940s being the most informed. In our theoretical framework, it is indeed the cohorts that are less politically informed that should have larger dispersion in their views, as they would form their opinions based on other, idiosyncratic elements, such as features related to identity or personality. This makes their opinions harder to predict and hence less cohesive.

Finally, we also consider how the above may be related to the level of polarisation of different cohorts experienced in their formative years. Recall that previous work in Political Science had identified how political polarisation levels were relatively low in the 1930s–1960s, and started to increase at the end of the 1960s (see Barber and McCarty, 2013). Thus, we expect generations that came of age around the 1940s

[2] Current theoretical work on cohort effects, e.g., the model in Malmendier and Nagel (2016), does not take into consideration dynamic linkages or more generally learning from endogenous policies.

[3] In our work we focus on cycles of party polarisation and party consensus, which translate into cycles of uninformed (and more polarised) or informed (and less polarised) public opinions. Other work in political cycles typically focused on cycles of "good" and "bad" policies or of different ideologies (see for example Battaglini and Coate, 2008) or cycles of simple and complex policies (see Levy et al., 2022).

[4] The original probabilistic voting model in Lindbeck and Weibul (1987) focuses on voters with different preferences. Strömberg (2004) analyses a probabilistic voting model with voters with different information.

[5] The full list of the 33 political questions along with additional 18 questions some of which we used as controls is provided in Appendix B (Table A.1).

[6] The questions' codes (in ANES, 2022) are VCF0605, VCF0806, VCF0809, VCF0606, VCF0208, VCF0210, VCF0213, VCF0220, VCF0604, VCF0823.

[7] Answers to the questions are coded in the survey from 0–100 in Thermometer questions, 1–7 in questions 1–4, and 1–2 for question 10.

[8] Such questions are easy to aggregate as the answers to them are on the same scale.

[9] These questions' codes in the ANES data set are VCF0992, VCF0991, VCF9005, VCF9260, VCF9261, VCF9262, VCF0501, VCF0310.

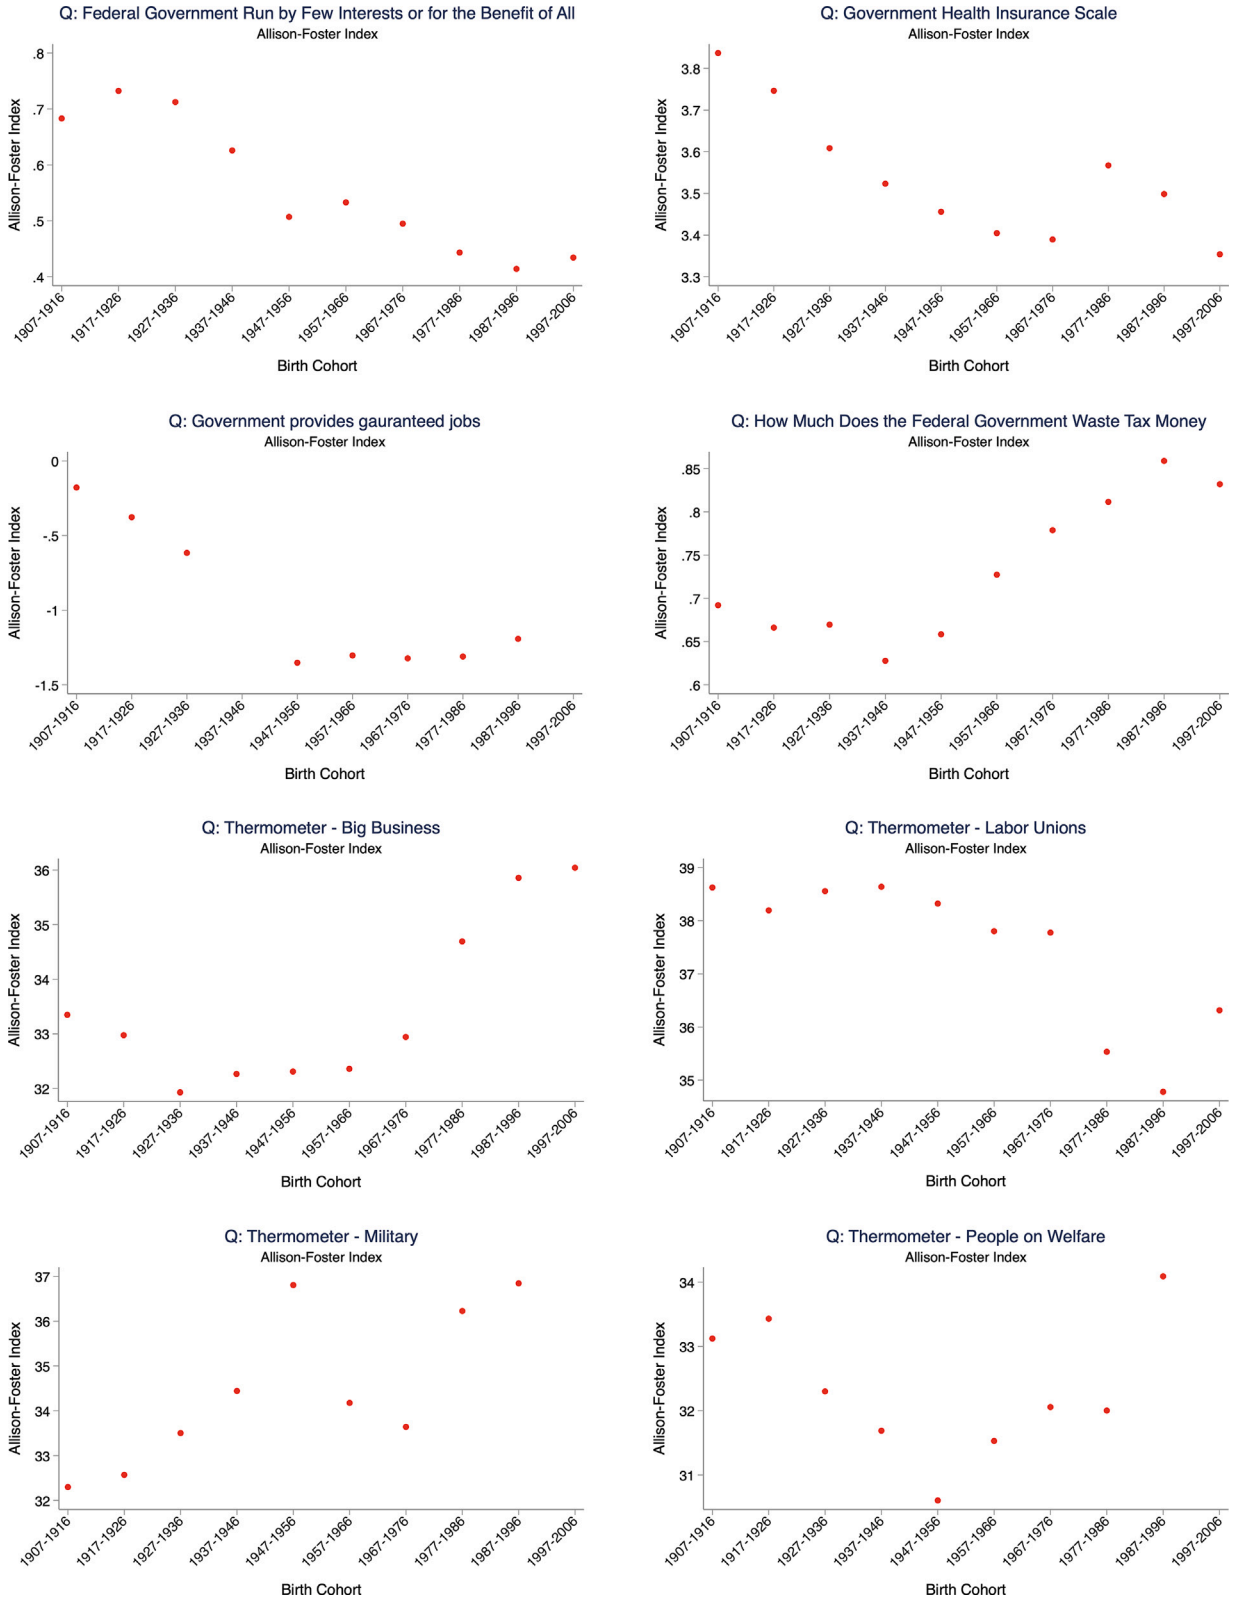

Fig. 1. Allison–Foster index for survey questions.

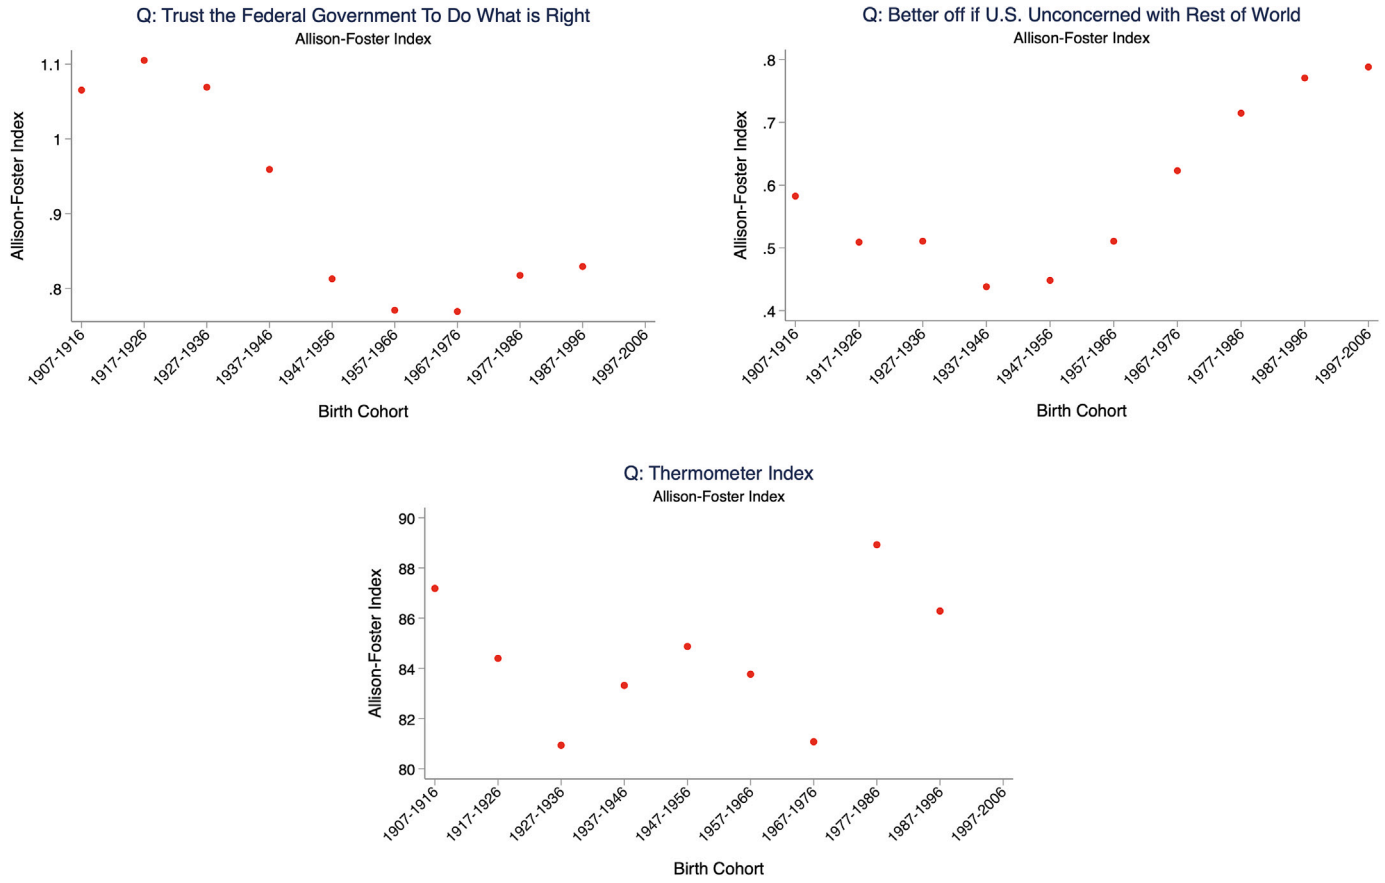

Fig. 1. *(Continued)*

to be exposed to little variation in their data, and so to be less knowledgeable and more polarised in their views. Alternatively, those who came of age at the end of the 60s (e.g., born in the 40s), have experienced high level of polarisation and are thus more knowledgeable and have strong shared experiences and hence views. This is consistent both with the level of information identified in the graphs above, and the level of the dispersion in opinions of cohorts captured by some of the trends of the Allison–Foster Index above.[10]

In Appendix B we employ another approach to check for differences in second moments across cohorts. Specifically we do a regression analysis on opinions (for the same questions reported above) and a heteroskedasticity test to measure how different cohorts "contribute" to the unexplained residuals. Similar trends arise in this approach as well, as we find that differences in cohorts' second moments are statistically significant.

3. The model

The model has two main parts. The first is a dynamic, overlapping generations structure in which voters learn exclusively in their impressionable years and vote when they mature. The second describes each period's political competition and consists of a probabilistic voting model with two ideological parties that compete over voters from different cohorts, with potentially different histories/memories.

[10] High levels of polarisation may sometimes be associated with gridlock and hence little variation in the policies that voters actually observe. This may be a feature of the political arena from the end of the 1990s onwards, but is less relevant for the polarisation level of the 1960s and 1970s. We thank a referee for raising this point.

It will suffice to focus on a simple economic environment which we present first: A polity is considering, every period, a choice between two policies l and r. The economic environment is a simple mapping between policies and outcomes. Specifically, the observable economic outcome x_t at period t is:

$$(II.1)x_t = \begin{cases} \beta_l^* + \varepsilon_t \text{ if policy } l \text{ is chosen} \\ \beta_r^* + \varepsilon_t \text{ if policy } r \text{ is chosen} \end{cases}$$

where ε_t is iid across time, Normally distributed with zero mean and variance σ^2. The results in the paper can be generalised to continuous policies and other distributions of the noise.

Voters understand how the data generating process depends on parameters $\beta = (\beta_l, \beta_r)$, but do not know the true value of these parameters, $\beta^* = (\beta_l^*, \beta_r^*)$. They are endowed with a continuous and symmetric prior $G(\beta_1, \beta_2)$ on some compact set $B \in \mathbb{R}^2$, which determines how β^* is generated. Without loss of generality, assume that the true state satisfies $\beta_l^* > \beta_r^*$, hence l is the optimal policy.

We think of the outcome x_t as a common element in the voters' preferences. In particular, at any period t, all voters will be generally inclined to vote for policy l if, given their information at period t, Ω_t, $E[\beta_l|\Omega_t] > E[\beta_r|\Omega_t]$. Below we describe the (potentially different) information held by different groups and how this will feed into their voting behaviour.

3.1. Overlapping generations and voters' beliefs

For simplicity we divide time into different phases; these could be thought of as decades in which the young form their political beliefs. Each phase $\kappa \in \{1, 2, ...\}$ has $K < \infty$ periods. In phase κ, we have two cohorts of voters, young and old, Y^κ and O^κ, and a generation of young

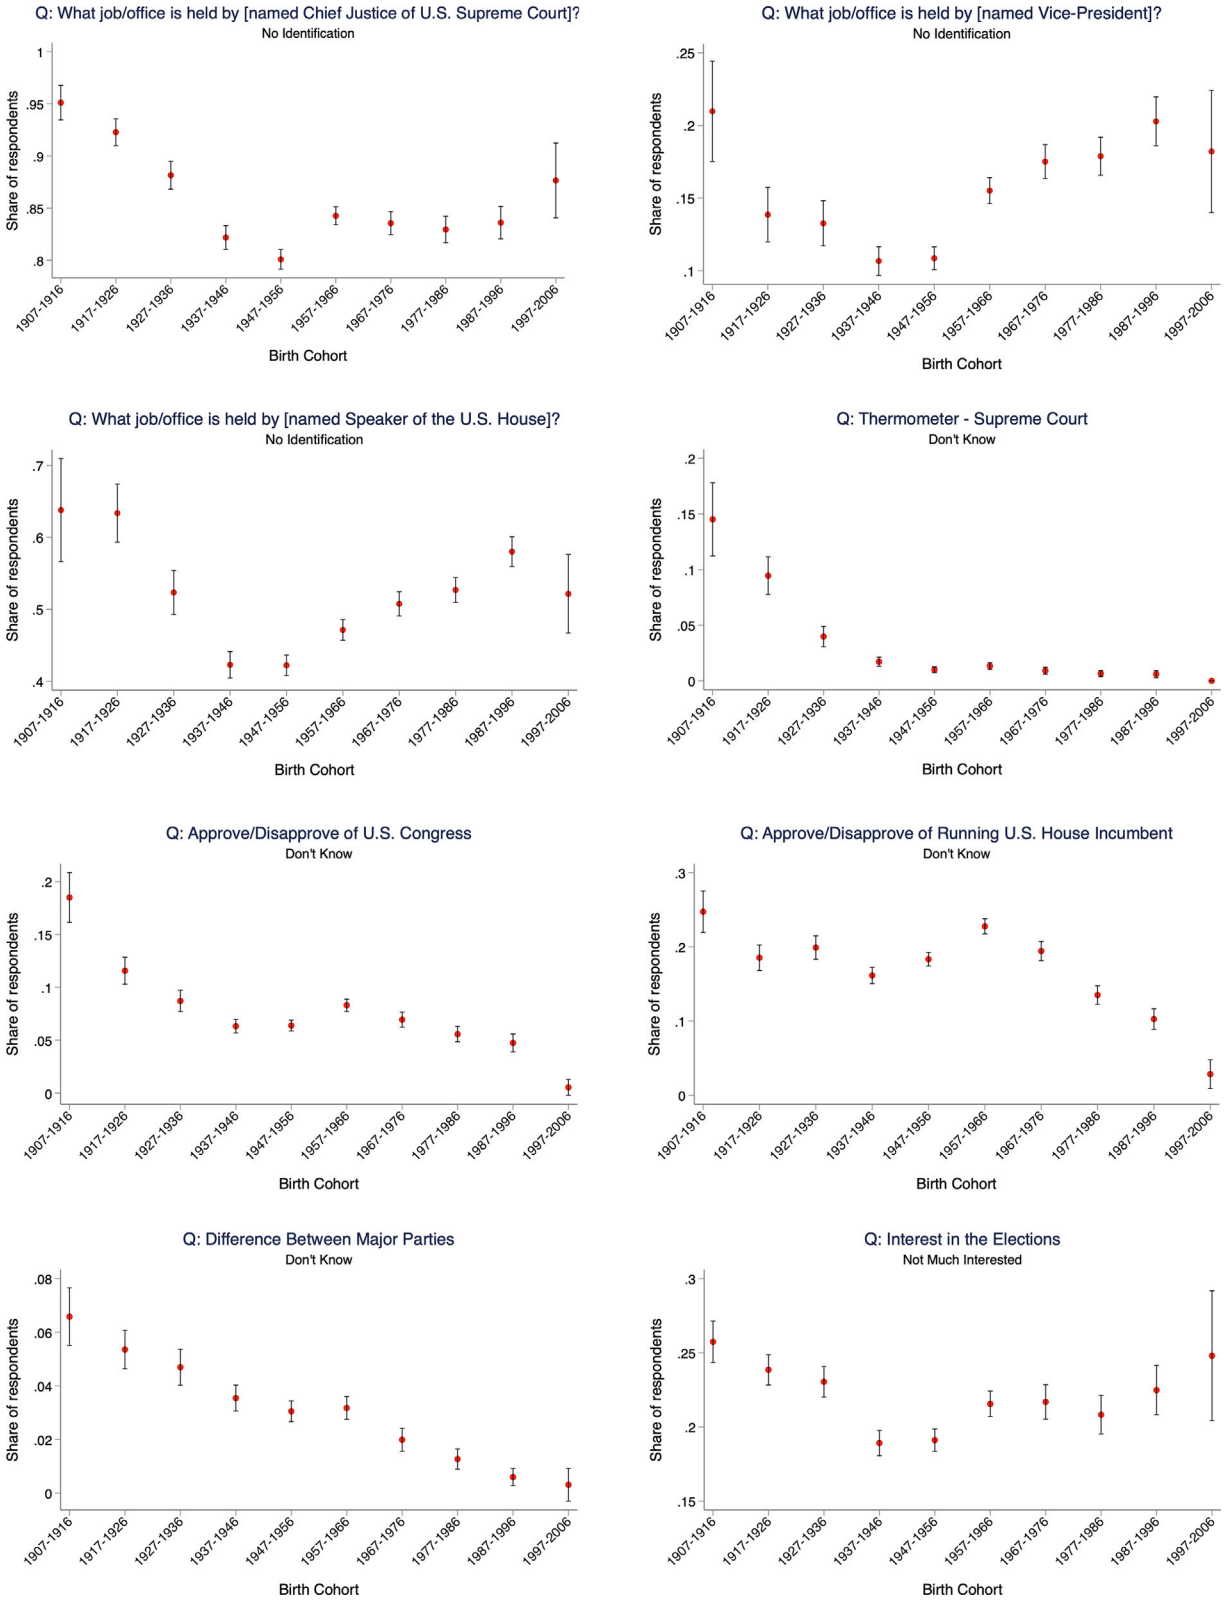

Fig. 2. How informed are the different cohorts?

individuals, I^κ, in their *impressionable years* (e.g., ages 14–24); they do not vote, but accumulate experiences in this phase and learn about politics. When moving to phase $\kappa+1$, O^κ perishes, Y^κ turns into $O^{\kappa+1}$, I^κ becomes $Y^{\kappa+1}$, and a new generation of young individuals, $I^{\kappa+1}$, arrives.

The information held by a voting cohort $c \in \{Y^\kappa, O^\kappa\}$ at some phase κ and period t is composed of a prior G, and a history of political outcomes they have observed. Each cohort has its history which is accumulated during its impressionable years (and only then), and so the two voting cohorts had observed different histories. Cohort Y^κ had observed the history of policies and outcomes in phase $\kappa-1$, $H^{Y^\kappa} = \{p_\tau, x_\tau\}_{\tau=t_{(\kappa-1),1}}^{\tau=t_{(\kappa-1),K}}$, and cohort O^κ had observed the political outcomes at phase $\kappa-2$, $H^{O^\kappa} = \{p_\tau, x_\tau\}_{\tau=t_{(\kappa-2),1}}^{\tau=t_{(\kappa-2),K}}$, where $p_\tau \in \{l, r\}$ is the implemented policy in some past period τ and x_τ is the policy outcome in that period. The history of each cohort of voters is attained endogenously through previous generations' experiences, preferences, and the political process, which we describe below.

Histories are then translated into beliefs. Specifically, the impressionable individuals of phase κ, I^κ, treat the history they observe as exogenous, and compute their posterior distribution on $\beta \in B$, which will constitute their knowledge for their two voting phases. Specifically, the posterior density distribution $g^{I^\kappa}(\cdot)$ on vectors $\beta = (\beta_l, \beta_r)$ is:

$$(II.2)\ g^{I^\kappa}(\beta) = \frac{g(\beta)\prod_{\tau=t_{\kappa,1}}^{\tau=t_{\kappa,K}} f(x_\tau - E[x_\tau|p_\tau, \beta])}{\int_{\beta'} g(\beta')\prod_{\tau=t_{\kappa,1}}^{\tau=t_{\kappa,K}} f(x_\tau - E[x_\tau|p_\tau, \beta'])d\beta'}$$

where $E[x_\tau|l, \beta] = \beta_l$ and $E[x_\tau|r, \beta] = \beta_r$ and $f(\cdot)$ is the (Normal) density of the shock ε_t. Thus, at each period t in phases $\kappa+1$ and $\kappa+2$, this cohort will use $g^{I^\kappa}(\beta) = g^{Y^{\kappa+1}}(\beta) = g^{O^{\kappa+2}}(\beta)$ as their information Ω_t to feed into their voting decisions.

Note that in our OLG model we make several simplifying assumptions. First, voters vote based only on their beliefs formed in their impressionable years and so their current experiences do not matter. Second, while generations could be defined continuously, here we have a discrete notion of generations with completely separate memories. Finally, to simplify, we focus on only two generations of voting cohorts, and also assume that those in the impressionable years age group do not vote. All the above assumptions are made to allow for clarity and simplicity and to focus on the effect of impressionable years and the possibility of an intergenerational conflict that arises due to different memories. The results are robust to variations of these assumptions, as we discuss later on in Section 4.5.

3.2. Electoral competition

We now describe how political outcomes are determined at each generic period t. There are two parties, each identified with a special interest on a different policy. Party L prefers policy l and party R prefers policy r. The utilities of party L and R from policy $p \in \{l, r\}$, $U^R(p)$ and $U^L(p)$, satisfy:

$$(II.3)\ U^R(r) = U^L(l) = 1, U^R(l) = U^L(r) = 0.$$

In addition, parties enjoy office rents denoted by $\nu > 0$. Thus, given an election at period t, and an implemented policy p_t, party $J's$ utility, for $J \in \{L, R\}$, is

$$U_t^J(p) + I_t^J \nu$$

where $I_t^J = 1$ if party J won the election at period t and 0 otherwise.

In the election at period t, each party offers a policy $p_t^J \in \{l, r\}$. We say that parties polarise when $p_t^R \neq p_t^L$; naturally, this will imply that $p_t^R = r$ and $p_t^L = l$. When $p_t^R = p_t^L$, we say that parties are in consensus. Polarisation can only arise in electoral competition when parties are ideological (as we assume) and when that they face some uncertainty in the election with regards to voting behaviour.

To introduce uncertain voting behaviour we adopt the probabilistic voting model, following Lindbeck and Weibul (1987). Beyond their histories/memories, voters may differ on additional aspects of the voting choice which parties may have uncertainty about. For example, voters may be affected by personal attributes of candidates. Thus, at period t, given the policies espoused by the different parties, a voter i in cohort c, $c \in \{Y^\kappa, O^\kappa\}$, votes for L if:

$$E[x_t|p_t^L, H^c] - E[x_t|p_t^R, H^c] + v_t^{ic} + \phi_t > 0,$$

where ϕ_t is an aggregate shock uniformly distributed on $[-\frac{1}{2\zeta}, \frac{1}{2\zeta}]$ and v_t^{ic} is an idiosyncratic cohort-specific shock distributed on $[-\frac{1}{2\xi^c}, \frac{1}{2\xi^c}]$, where $\zeta, \xi^c \in (0, \infty)$.[11] Both types of shocks capture the uncertainty of parties over voters in general as well as over the different groups. Note that $E[x_t|p_t^L, H^c] - E[x_t|p_t^R, H^c] = 0$ if parties are in consensus, whereas if they polarise, then:

$$E[x_t|p_t^L, H^c] - E[x_t|p_t^R, H^c] = E[\beta_l - \beta_r|H^c] = \int_\beta (\beta_l - \beta_r)g^c(\beta)d\beta$$

Let each cohort $c's$ share in the voter population be denoted by γ^c, where $\sum_{c\in\{Y^\kappa, O^\kappa\}} \gamma^c = 1$. After standard manipulations as in probabilistic voting models, we get that party L wins the election at period t with probability

$$(II.4)\ \Pr(L|p_t^L, p_t^R, \{H^c\}_c) = \begin{cases} 1 \text{ if } \frac{1}{2} + \zeta\sum_c \omega^c(E[x_t|p_t^L, H^c] - E[x_t|p_t^R, H^c]) > 1 \\ 0 \text{ if } \frac{1}{2} + \zeta\sum_c \omega^c(E[x_t|p_t^L, H^c] - E[x_t|p_t^R, H^c]) < 0 \\ \frac{1}{2} + \zeta\sum_c \omega^c(E[x_t|p_t^L, H^c] - E[x_t|p_t^R, H^c]) \text{ otherwise} \end{cases},$$

where $\omega^c = \frac{\gamma^c\xi^c}{\sum_{c'}\gamma^{c'}\xi^{c'}}$ denotes the *political power* of cohort c in the electoral competition. Note that this weight is increasing in group size as is intuitive and decreases in the variance of the distribution of idiosyncratic shocks of cohort c. The latter effect is due to the fact that as the variance decreases, this group of voters is more sensitive to changes in what different parties offer.

In a Nash equilibrium, at any period t at phase κ, given p_t^R, party L chooses p_t^L to maximise:

$$\Pr\left(L|p_t^L, p_t^R, \{H^c\}_{c\in\{Y^\kappa, O^\kappa\}}\right)(U^L(p_t^L) + \nu) + \left(1 - \Pr\left(L|p_t^L, p_t^R, \{H^c\}_{c\in\{Y^\kappa, O^\kappa\}}\right)\right)U^L(p_t^R),$$

with an analogous expression for party R, where parties know $G, \{H^c\}_{c\in\{Y^\kappa, O^\kappa\}}$. Following the choice of platforms, ϕ_t and v_t^{ic} are drawn and the election result is in accordance with $(II.4)$.

The above is a full description of the political competition at period t in some phase κ. Note that at any phase κ, voters' knowledge is fixed (as for simplicity they are not influenced by current events) and thus the same equilibrium arises throughout the phase; either a consensus on some policy $p \in \{l, r\}$ or a polarisation equilibrium. The election's outcome though is stochastic and so a different party may win the election at each period at phase κ, and moreover the political outcome is also subject to a per period shock ε. While having the same political equilibrium throughout a phase is a stark feature, it allows us to describe the impressionable years of generations as either those that had lived through consensus years, or those that had lived through polarised times.

3.3. Dynamics

Putting all the above ingredients together, the dynamic model is defined as follows. The model starts with some initial history. In each phase

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

κ, and any period t in this phase, a Nash equilibrium in the party competition game arises with the platforms p_t^J. After the cohort-specific and the aggregate shocks are drawn, the winning party is determined by (*II*.4) and this party implements its platform. After the policy shock ε_t is drawn, histories are updated. In phase $\kappa+1$, cohort Y^κ becomes $O^{\kappa+1}$ (but retains its historical knowledge), cohort I^κ become $y^{\kappa+1}$, and a new generation $I^{\kappa+1}$ with a null history and prior G is "born".

Note that the behaviour of voters and of parties described in the model implicitly assumes that voters and parties are myopic. For voters this means that when they choose whom to vote for they do not take into account the experimentation value of voting for a policy and how it affects the beliefs of future voters. Parties also do not think about manipulating voters' beliefs to affect their winning chances in the future. Myopia is a standard assumption in models of electoral competition, yet our results can also hold more generally.

3.4. A preliminary lemma

We now highlight the mechanism by which the level of information in historical data affects electoral competition. Fix the histories H^c, $c \in \{Y^\kappa, O^\kappa\}$, that voters observe, and consider the political game that arises in (any period in) phase κ. The result below is reminiscent of the result in Calvert (1985):

Lemma 1 (Consensus vs Polarisation). *Let* $\rho = \frac{1}{2\zeta(1+\nu)}$. *In phase* κ, *if*

$$\left|\sum_{c\in\{Y^\kappa,O^\kappa\}} \omega^c[E[\beta_l|H^c] - E[\beta_r|H^c]]\right| > \rho$$

then both parties choose the same platform (consensus), and otherwise each party chooses its ideal policy (polarisation).

Generally speaking, parties prefer to pursue their own interests, but to successfully do so, they must be elected with some probability. This implies that they may be disciplined by voters to choose the policy that is more likely to generate a higher outcome given the historical data. If the historical data makes voters sufficiently confident that one of the policies, say l, is more likely to generate a higher outcome, then parties have to offer this policy and hence reach a consensus on l. If a party offers r instead, it will neither serve its own policy interest nor its office motivation, as it will face only a slim probability of being elected. Alternatively, if voters' historical data does not sufficiently discriminate between the different policies, then parties can afford to offer platforms that better serve their own policy interests and hence polarise.

To see the role that the parameters play, note that a higher office rent ν pushes parties to be in consensus, as they are more eager to get elected and therefore to satisfy the voters' will. Similarly, a higher ζ implies a smaller interval for the shock ϕ_t : This means that parties are more certain about how voters vote, and again this pushes parties to be more in consensus. As both ζ and ν work in the same direction to affect the possibility of consensus, we summarised them with the parameter ρ, so that a lower ρ is more conducive to consensus. Finally, as can be seen in Lemma 1, parties average the preferences of the cohorts, which differ in their history H^c. How they do this averaging depends on the cohorts' political weight ω^c, which relates to the population share of group c, γ^c, and the variance of the idiosyncratic component of the cohorts, ξ^c. We will discuss the effects of these parameters later on.

Note that absent any information to both cohorts, $E(\beta_l - \beta_r) = 0$, as G is symmetric, and so parties will polarise. To make learning meaningful in the model, it is reasonable to consider the case where parties reach a consensus when both cohorts of voters know the true state. We will henceforth assume that $\beta_l^* - \beta_r^* > \rho$, that is, the true state is such that if voters knew the state, consensus will arise on l. In contrast, we assume that if all voters, young and old, have fully learnt the benefit of only one policy, then this knowledge is not sufficient to distinguish the two policies. Specifically, let $E[\beta_{-p}|\beta_p^*]$ denote the expected effectiveness of policy $-p$ when voters knows β_p^* (and nothing else). We then assume:

Assumption 1. The state of the world β^* satisfies

$$\beta_l^* - \beta_r^* > \rho > \max\{|\beta_l^* - E(\beta_r|\beta_l^*)|, |\beta_r^* - E(\beta_l|\beta_r^*)|\}$$

Note that the important assumption here is that $\beta_l^* - \beta_r^* > \max\{|\beta_l^* - E(\beta_r|\beta_l^*)|, |\beta_r^* - E(\beta_l|\beta_r^*)|\}$, which implies that voters are able to better differentiate the two policies when they had experienced both. In contrast, when voters only experience one policy, they are less able to tell how good or bad they have it compared to the other option. This is a reasonable assumption in a political context (and more generally potentially in other learning environments).

4. Cycles, intergenerational conflicts, and demography

Given their historical experiences in their formative years, both exogenous policy shocks and endogenous policies will have an effect on a cohort's knowledge. To see this note that when two cohorts were exposed to the same policies, the different historical experiences due to the random shocks (ε_t) will shape opinions in different ways. Alternatively, even when exposed to the same shocks, cohorts will have different experiences if different policies were implemented in their formative years. For example, experiencing only policy l is very different to being exposed to both policies l and r.

Below we mostly focus on the case of $\sigma \to 0$. This allows us to highlight the effect of endogenous policies on cyclicality of politics and cohort differences. In the Appendices we show how the results hold for any $\sigma > 0$ and a large enough K, as well as illustrate this by simulations in the main text.

4.1. Generational differences in polarisation levels

The impressionable cohort either experiences a history with the same policy p implemented, or a history of polarisation and turnover in which they experienced both policies. As we focus on an arbitrarily small σ, learning is very fast, and so a cohort's expected difference in utility from the two policies is either $\beta_l^* - \beta_r^*$ in the latter case and $\beta_l^* - E(\beta_r|\beta_l^*)$ or $\beta_r^* - E(\beta_l|\beta_r^*)$ in the former case. An impressionable cohort that goes through a consensus phase, or a polarisation phase with no turnover, is then less informed about the differential benefit of policies compared to a cohort that has its impressionable years at a polarisation phase with turnover. Such a cohort will know exactly what the policy they observe delivers, but not necessarily how good or bad they have it compared to the other policy they haven't experienced.

Given our assumptions about idiosyncratic noise in voting, one way to compare cohorts is in terms of how predictable their voting intentions are. In general, a cohort that has better knowledge is more predictable, as their beliefs about the optimal policy crowd out the stochastic idiosyncratic noise element. In contrast, the voting behaviour of the cohort that is less knowledgeable will be more unpredictable in terms of how they vote: Their beliefs imply a smaller expected difference between policies, leaving room for their voting to be more affected by other, more random, influences.

To formalise this define α_l^j as the probability a cohort c would vote for Party L if the two parties polarise:

$$\alpha_l^c = \begin{cases} 1 \text{ if } \frac{1}{2} + \xi^c[E(\beta_l|H^c) - E(\beta_r|H^c)] > 1 \\ 0 \text{ if } \frac{1}{2} + \xi^c[E(\beta_l|H^c) - E(\beta_r|H^c)] < 0 \\ \frac{1}{2} + \xi^c[E(\beta_l|H^c) - E(\beta_r|H^c)] \text{ otherwise} \end{cases},$$

where we have averaged over the distribution of ϕ. We now define the level of polarisation of cohort c as the variance of this (potential) voting behaviour:

Definition 1 *The level of polarisation of cohort* c *is given by* $\alpha_l^c(1-\alpha_l^c)$.

Note that above definition of polarisation is based on how individuals think about the two policies, l and r, and not about the current platforms of the parties. In this sense the definition is useful when we consider data on public opinions, as for example in our analysis in Section 2 of the ANES opinion polls.

From the definition above and Assumption 1 we then derive Observation 1:

Observation 1: *A cohort that had its impressionable years during times of consensus will be more polarised compared to a cohort that had its impressionable years during times of polarisation.*

A cohort that experienced polarisation and turnover will know the true state and so will have an expected difference between policies equal to $\beta_l^* - \beta_r^*$. The expected difference between policies for a cohort that was impressionable in times of consensus is lower by Assumption 1, resulting in a more polarised distribution of votes or opinions.

Motivated by the analysis above, the stylized facts presented in Section 2 explore how polarisation of opinions differs across cohorts.[12] Cohorts may have more spread in their opinions over different policies and will therefore differ in the level of second moments of their distribution of opinions. The mechanism that creates this in our model relates to how much information they had gained on the difference between the effectiveness of these policies. If they gained little information, they are more likely to form opinions on the basis of other, less predictable, idiosyncratic elements such as features of identity or personality. In other words, if $|E[x_t|l, H^c] - E[x_t|r, H^c]|$ is small, the noise element, υ_t^{ic}, will loom larger. If alternatively the common shared knowledge on policies is large, then we should expect less variance.

4.2. Cycles of polarisation and consensus

Above we highlighted the potential for generational differences in their level of polarisation of attitudes. We now show that these differences arise in equilibrium when policies are determined endogenously in our dynamic model. Specifically, our first result illustrates that the polity cycles forever between party polarisation and party consensus, and hence polarised cohorts and cohesive cohorts.

Proposition 1. *For all vectors of political power ω, the polity cycles forever between: (i) Party polarisation and party consensus. (ii) Polarised and cohesive cohorts.*

To see the intuition, recall first that when $\sigma \to 0$, learning is very fast, and so it is enough for a cohort to observe a policy implemented once to assess its benefit. Consider now the polity when it goes through a polarisation stage. Specifically, consider the situation, which will arise with a strictly positive probability, that in two consecutive periods two different policies were implemented (implying a different party winning the election in each period).[13] The impressionable cohort in this phase will then learn β^*. And if this were to happen for two consecutive phases, then in the next phase, these two consecutive cohorts become the young and old voting population. As they know what is the optimal policy, parties must reach consensus on l by Assumption 1 and Lemma 1. Thus, eventually, polarisation must give way to consensus. Intuitively, polarisation amounts to an unintended "experimentation" which facilitates learning of the optimal policy. This clear knowledge of the voters then disciplines the parties to reach a consensus on it.

Let us consider now the case in which parties offer the same policy and are in consensus for two consecutive phases (on the same policy). But a consensus cannot last more than two phases; if at some point two consecutive impressionable cohorts have only experienced one policy, then in the next phase the young and the old voters that comprise the voting population will have relatively little knowledge and little variation in their data. For example, suppose that both cohorts have had experienced only policy l, and so they both learn β_l^*. The only information that they have about β_r though is the knowledge of β_l^* (if the prior G allows for some correlation in the values of the βs). As a result their expected utility difference between the two policies is $|\beta_l^* - E(\beta_r|\beta_l^*)|$. They will find it harder to differentiate the benefits of the two policies, and given Assumption 1 and Lemma 1, parties will polarise. Thus, as the knowledge of these cohorts is limited and based only on their impressionable years (and not the full history which may include experiences of r as well), party consensus must bring about polarisation.

A cycle of party consensus and polarisation necessarily brings about a cycle of polarised and cohesive cohorts, and the other way around. We now delve further into understanding the features of these cycles.

4.3. Generational differences and the nature of cycles

We now further explore how the political power of cohorts affects the nature of the cycle in party platforms and intergenerational differences.

We start by addressing how a generational conflict can be resolved at each phase. Proposition 1 above implies heterogeneous generational attitudes whereby cohorts might differ in the direction of their beliefs and their level of polarisation. In these cases, it is important to understand who parties will cater to.

Recall that the political power of cohort c is ω^c, which increases in γ^c (its population share) and ξ^c (the certainty of this cohort's voting). We now focus on demography, the first element of political power. For example, a case of constant population growth (a large enough share γ of the young) implies that, holding other things equal, the young are always more powerful. Similarly, population decline implies that the old are politically more powerful. Expanding on the result in Proposition 1, we then have:

Proposition 2. *When the population grows fast enough compared to the case in which it declines fast enough: (i) Both stages of party consensus and party polarisation are shorter; (ii) Contemporaneous cohorts are less likely to have similar views and similar levels of cohesiveness/polarisation.*

When we have population growth, and so the young cohorts are politically more powerful, polarisation and consensus phases will be short-lived and changes in party politics will be frequent. This arises because the beliefs of the young are determined by their experiences in the previous phase and so if in the previous phase parties were in consensus, the current young voting cohort is polarised and parties will follow suit as they follow the preferences of the young. If the previous phase consisted of party polarisation, then, with political turnover, the current young cohort will be more cohesive and parties will reach consensus. Intuitively, if the young are powerful, parties cater to them, and so recent experiences become more important leading to quick changes in the parties' behaviour, which feed into quick changes in the young's attitudes and so on.

When the old are the politically powerful cohort, and so parties cater to their preferences, cycle stages last longer: Consider the possibility that there are at least two consecutive phases of polarisation (which, as we will show below, will arise). This implies that two consecutive cohorts potentially learn the state (conditional on turnover) and so later on, there will be two consecutive cohorts of old voters that will enforce parties to be in consensus. But these two consecutive phases of consensus will expose two consecutive impressionable cohorts to only one policy. These two cohorts will become polarised, which will then lead to two phases of polarisation, when each of these cohorts becomes the old cohort. Parties catering to the political power of the old implies more persistence in party equilibrium behaviour.

The above also implies that contemporaneous cohorts are more likely to differ in their knowledge and attitudes when the young are more powerful, as then parties' switches between polarisation and consensus are most frequent. Thus at any phase, old and young voters are most likely to have different opinions. When the old are more powerful and the cycle stages are longer, consecutive cohorts have similar experiences and so the whole voting block may be in agreement.

[12] Note that our measure of polarisation, $\alpha_l(1-\alpha_l)$, is suitable for binary choices, whereas in the ANES surveys, choices are on a scale of up to 1–100. Thus our qunatitative measure of polarisation in Section 2 is adapted accordingly.

[13] This will arise in finite expected time as when parties do polarise, each party must win with a probability that is bounded away from zero. This is guaranteed by our assumption that $\nu > 0$. If a party polarises and has only a negligible chance of being elected, it can deviate to be in consensus with the other party, gain $\frac{1}{2}\nu$, and be better off. Thus whenever parties polarise, each must win with a strictly positive probability.

In Proposition 2 the focus was on a constant growth or decline of the population. The results show the implications of these demographic trends on the nature of cycles in the political process. Our next result illustrates how, due to the dynamic linkages in the political system, even transitory shocks reverberate to affect the polity in the long run.

Specifically, we now explore the short and long-term effects of a one-off demographic change on future generations' political outcomes, specifically that of a baby boom. In our intergenerational comparative statics above we looked so far at constant power positions of the old and the young, but of course in reality a relatively large young cohort at some phase can imply a large old cohort later on. A one-off baby boom can be modelled then as follows: At some phase κ, an impressionable cohort is very large. At time $\kappa+1$, the young cohort is very large and hence politically powerful, and so parties follow the polarisation attitudes of the young. At the next phase $\kappa+2$, it is the old generation that is politically powerful and parties follow the attitudes of the old. From period $\kappa+3$ we go back to a constant process. It is then easy to see:

Observation 2: *A one-off baby boom (or any transitory shock to the political power of different generations in one period), can affect the political equilibrium indefinitely.*

To illustrate the above, let us assume that the demographic process, before a baby boom, is such that no cohort is too powerful and that ρ is sufficiently low. These two conditions imply, as is intuitive, that it is sufficient that one cohort (disregarding which one) is cohesive for parties to be in consensus.[14] To facilitate the discussion, suppose that within a phase there are sufficient experiences (as will be with a large K), so that polarisation also implies political turnover and so during such a phase the impressionable cohort can learn the truth.

In the Figure below, we describe the steady state equilibrium for these parameters. The first row describes the party equilibrium during the phase, polarisation (*Pol*) or consensus (*Con*). The second and third rows respectively describe the attitudes of the young and the old voting cohort; whether they are polarised (*Pol*) or are in consensus (*Con*) on the optimal policy l. For example, the young at phase $\kappa-1$ gain their attitudes from the party equilibrium at phase $\kappa-2$; as they have experienced both policies in a polarisation equilibrium they will form a cohesive cohort. As can be seen polarisation lasts for one period only and consensus for two phases, given the case of a relatively low ρ.

	$...\kappa-2$	$\kappa-1$	κ	$\kappa+1$	$\kappa+2$	$\kappa+3$	$\kappa+4...$
Eq.	*Pol*	*Con*	*Con*	*Pol*	*Con*	*Con*	*Pol*
Y	*Pol*	*Con*	*Pol*	*Pol*	*Con*	*Pol*	*Pol*
O	*Pol*	*Pol*	*Con*	*Pol*	*Pol*	*Con*	*Pol*

Now consider a baby boom and so a large impressionable cohort at phase κ that becomes a politically powerful young cohort in $\kappa+1$ and a politically powerful old cohort in period $\kappa+2$. Now the party equilibrium changes as the parties follow the preferences of the politically powerful cohort: the young in $\kappa+1$ and the old in $\kappa+2$. We then have:

	$...\kappa-2$	$\kappa-1$	κ	$\kappa+1$	$\kappa+2$	$\kappa+3$	$\kappa+4...$
Eq.	*Pol*	*Con*	*Con*	*Pol*	**Pol**	*Con*	**Con**
Y	*Pol*	*Con*	*Pol*	*Pol*	*Con*	**Con**	*Pol*
O	*Pol*	*Pol*	*Con*	*Pol*	*Pol*	*Con*	**Con**

As can be seen in the new table, there is no effect on party behaviour in period $\kappa+1$ compared to the previous table as both cohorts have the same attitudes. A direct change arises when the baby boomers constitute the old generation, and the political outcome switches from consensus to polarisation, as now parties follow the attitudes of the old. This means that the polarisation stage lasts longer. But beyond this direct short-term effect that arises from the preferences of the baby boomers, there is also an indirect long-term effect: Party polarisation at phase $\kappa+2$ affects the attitudes of the impressionable voters at that phase, and they are now in consensus instead of being polarised. So from that phase onwards, the timing of the cycle phases change (and so party equilibria are different for the same phases across the two tables); as a result, the beliefs of *all future cohorts are affected*. This feature of the reverberation of a transitory shock again illustrates the importance of considering dynamic linkages across cohorts, through endogenous policies and beliefs.

There are other demographic parameters that we can consider. For example, more recent cohorts may live longer. There are different ways to think about lifespan in the context of our model; perhaps the most natural way is to think about more cohorts of voters. So with increasing lifespan, we can have for example three voting cohorts, let's say young, old, and very old. To understand whom the parties cater to, we need to go back to Lemma 1. Suppose then that no cohort can overcome the political power of the two other cohorts. Then we can have longer phases of polarisation and consensus as a longer life span may imply more persistence and hence more overlap of knowledge across cohorts. Another way to think about increasing life span is to increase K, the number of political experiences a cohorts has in their phase. A higher K implies a higher chance for political turnover in a phase, and in the case in which σ is larger, better learning. This will increase the average instances of consensus (and on the correct policy, in the model with substantial noise).

4.4. The effect of shocks

The analysis above focused on arbitrarily small σ and so focused on systemic effects that arise when policies are endogenous. But exogenous policy shocks also affect political outcomes as well as learning. As with a demographic change, a large policy shock will have direct effect on the political process through the effect it has on the learning of the impressionable cohort. And it will also have an indirect effect as this impressionable cohort will potentially induce parties to implement different policies following this shock once they become active voters; this in turn will affect the beliefs of the future generations.

In Appendix A we extend all our results above to the case of $\sigma > 0$, and a large K. A large K implies that an impressionable cohort has sufficient experiences and so can learn despite the noise. How large K has to be depends on σ, as both affect the possibility of learning.

The simulations below illustrate that adding noise does not change the essence of our results even for small K. We used the following values: $B=[0,6]^2$ with a uniform distribution, $(\beta_l^*,\beta_r^*)=(3.5,2.5)$, and an office rent $\nu=2$. For these values, knowing the exact state of the world implies an expected difference of policy effectiveness of $\beta_l^*-\beta_r^*=1$. When a cohort learns only one of the policy parameters β_p^* then $|E[\beta_{-p}-\beta_p^*|\beta_p^*]|=0.5$. To get the graphs below, we simulated the model over 100 periods.

We first report the level of knowledge of voters over time as a function of two values of σ, $\sigma=0.2$ or $\sigma=2.5$, and for the case of a population decline $(\omega^Y,\omega^O)=(0.2,0.8)$. As can be seen in Fig. 3, for both values of σ, the voters oscillate between periods of being more and less informed (beliefs centred on 1 and half respectively), which is a result of the parties oscillating between polarisation and consensus. When we switch from $\sigma=0.2$ to $\sigma=2.5$, while the same pattern arises, variability of beliefs increases as expected. Moreover, transitory large shocks can influence beliefs substantially, and as a result change the "clock" for future generations' politics.

In the next set of diagrams we plot the differences between the polarisation level of the young and the old contemporaneous cohorts. We compare these differences in the cases in which the young are more powerful and the old are more powerful, $(\omega^Y,\omega^O)=(0.8,0.2)$ or $(\omega^Y,\omega^O)=(0.2,0.8)$, for the case of a variance of $\sigma=0.2$. As can be seen in Fig. 4, the rate at which the generational conflict oscillates is higher with population growth (powerful young) and the proportion of times generations have the same level of polarisation decreases. This is again, in line with our results in Proposition 2 and illustrates how the gist of our results holds also with policy shocks.

[14] We show this formally in the proof of Proposition 2 in Appendix A.

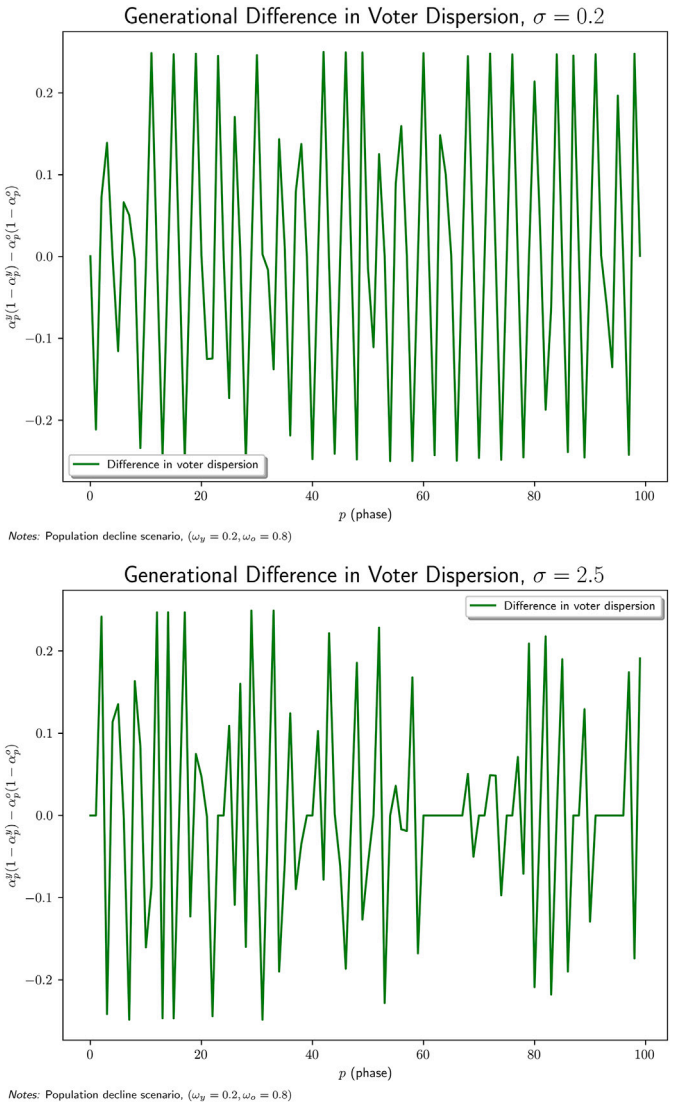

Fig. 3. Changing σ with population decline.

Finally, we used 20,000 repeated simulations of the above to calculate some positive and normative measures of the political cycles that we report below. As can be seen from Table 1, increasing the variance has the effect of reducing the polity's ability to learn and implement the optimal policy (l). Still, the cycles are pretty regular even with high variance, and consensus arises for around 50 % of the time. Note though that for low K or for high σ, consensus is not necessarily on the correct policy, due to wrong learning.

4.5. Extension: allowing voters to learn as they grow older

In our main model we assumed that voters learn only in their impressionable years, get no knowledge from previous generations, and do not accumulate further knowledge during their voting phases as young and old adults.

Note that if voters gain all information from previous generations and continue to fully learn throughout their life time, this means that they observe the full history of events. In this case it is easy to see that consensus is inevitable. Observing the full history nullifies any potential effect for cohorts in the long term, as all cohorts have the same information. It is also easy to see that beliefs must converge to some particular posterior distribution over the value of β as there is no "disruption" in how voters learn. Now if these beliefs converge to be such that parties can polarise, it means that voters must learn the true state at some point. This is a contradiction to the assumption that beliefs converge to be such that parties can polarise.

But whenever some frictions arise across cohorts, so information is not fully transmitted or voters do not fully learn as they grow older, then cycles will arise. To see this, let us assume that voters continue to learn also throughout their young phase. In other words, their beliefs only stabilise when they become older. We can then show (focusing again on $\sigma \to 0$):

Proposition 3. *When voters learn both in their impressionable phase and in their young phase, cycles still arise, with a shorter polarisation phase compared to the case in which they only learn during their impressionable years. As before, the phases of the cycles are shorter when the young are politically powerful.*

To see the intuition, note that the key assumption that changes here is that there is a phase in which individuals simultaneously vote and learn. This means that as they learn in their young phase, the party platform equilibrium can potentially change. If they have been through a polarisation phase in their impressionable years, they have already learnt the true state, and so this does not change their beliefs. But if

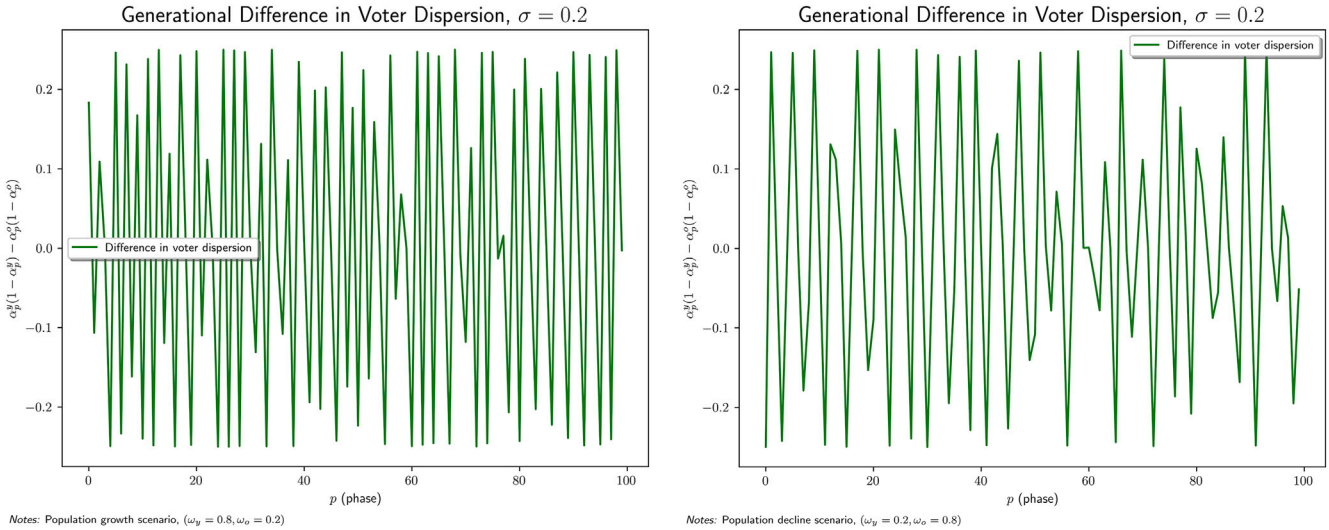

Fig. 4. Changing population scenario with fixed $\sigma = 0.2$.

Table 1
Empirical moments from simulation.

	$K = 5$			$K = 10$		
	$\sigma = 0.2$	$\sigma = 1.2$	$\sigma = 2.5$	$\sigma = 0.2$	$\sigma = 1.2$	$\sigma = 2.5$
Optimal policy	86.89 %	78.63 %	67.93 %	88.32 %	82.52 %	73.42 %
	(1.17 %)	(2.98 %)	(4.06 %)	(0.88 %)	(2.13 %)	(3.54 %)
Consensus	44.81 %	50.01 %	53.82 %	49.69 %	50.47 %	53.14 %
	(2.41 %)	(4.14 %)	(4.48 %)	(1.45 %)	(3.77 %)	(4.20 %)
Optimal policy | Consensus	99.87 %	95.98 %	79.93 %	99.92 %	98.86 %	88.35 %
	(0.51 %)	(2.88 %)	(6.00 %)	(0.39 %)	(1.48 %)	(4.77 %)
Length of consensus phases	8.68	10.59	11.17	18.79	20.91	22.22
	(0.43)	(1.20)	(1.49)	(0.60)	(2.05)	(2.73)

they have been through a consensus phase, more learning allows them to identify the true state while still in their young phase. This means that parties potentially switch to consensus within this phase, affecting the knowledge of the next impressionable cohort.

5. Conclusions: generational conflicts and memories

We have used a relatively simple OLG framework to explore endogenous dynamic linkages between cohorts. We have focused on particular aspects that can result in intergenerational conflicts, namely that cohorts gain their important experiences or memories at different times. In our dynamic model this led to cohorts differing in the second moments of their distribution of attitudes, with a cycle of some being more polarised and some more cohesive.

One feature of political processes that we abstracted from is the ability of parties and other organisations such as media or interest groups to potentially affect beliefs or preferences by highlighting particular past events. In this regard, a potential difference between old and young cohorts is that the old have actually lived through a longer history which parties can allude to. A potential political narrative is a *nostalgic* one, that rekindles some memories (potentially biased) of the "good times" that people might have experienced when they were younger. Naturally this resonates more with old voters who may then prefer to implement policies that are more in line with those of the past. While other papers have considered political narratives, as far as we know, none of them had looked into different narratives directed at different cohorts.[15] In future analysis it would be interesting to consider also the supply side of memories, potentially arising strategically on behalf of parties.

Our empirical investigation has suggested statistically significant differences between the distribution of opinions of different cohorts. While in this paper we provide some stylised facts that suggest that the mechanism in our model is present, further empirical analysis is needed. Our data of public opinion at the national level lacks the power to explain how the trends we uncover in public opinions are related to the level of polarisation in policies and to the level of information that voters have. One can extend the analysis using state level data on policy polarisation and public opinion and hopefully to get a more clear picture of the exact mechanism behind the different distributions of cohorts' public opinions.

6. Appendix

6.1. Appendix A: proofs

6.1.1. Proofs for results in the text

We now prove the following useful claim:

Claim A1: *Assume that $K \geq 2$ and $\sigma^2 \to 0$. (i) Suppose H^c includes only one policy that was implemented throughout the history. Then with probability arbitrarily close to one beliefs concentrate on $(\beta_p^*, E[\beta_{-p}|\beta_p^*])$. (ii) Suppose H^c includes both policy l and policy r that were implemented. Then with probability arbitrarily close to one beliefs concentrate on β^*.*

Proof of Claim A1: (i) Assume one policy p is implemented during H^c. As shocks are distributed normally, as $\sigma \to 0$ the distribution of shocks concentrates on its expectation. Therefore, for any $\gamma', \gamma'' > 0$ there is a $\bar{\sigma} > 0$ such that for all $\sigma < \bar{\sigma}$ with probability $1 - \gamma'$ all the shocks in the K periods are in $[-\gamma'', \gamma'']$. As a result, when $\sigma \to 0$, with probability arbitrarily close to one, the posterior belief after any path will be concentrated on $(\beta_p^*, E[\beta_{-p}|\beta_p^*])$. (ii) Assume that both l and r have been implemented in H^c. Again, for any $\gamma', \gamma'' > 0$ there is a $\bar{\sigma} > 0$ such that for all $\sigma < \bar{\sigma}$ with probability $1 - \gamma'$ all the shocks in the K periods are in $[-\gamma'', \gamma'']$. As a result, when $\sigma \to 0$, with probability arbitrarily close to one, the posterior belief after almost any path will be concentrated on β^*. □

Proof of Lemma 1: Assume that party L offers l. If party R offers l too it attains $\frac{1}{2}\nu$, whereas if it switches to r it attains $(1-\Pr(L \text{ wins}|l,r))(1+\nu)$, where

$$(1 - \Pr(L\text{wins}|l,r)) = \begin{cases} 1 \text{ if } \frac{1}{2} + \sum_c \omega^c [E[\beta_r|H^c] - E[\beta_l|H^c]] > 1 \\ 0 \text{ if } \frac{1}{2} + \sum_c \omega^c [E[\beta_r|H^c] - E[\beta_l|H^c]] < 0 \\ \frac{1}{2} + \sum_c \omega^c [E[\beta_r|H^c] - E[\beta_l|H^c]] \text{ otherwise} \end{cases}.$$

Note that if $\frac{1}{2} + \sum_c \omega^c [E[\beta_r|H^c] - E[\beta_l|H^c]] < 0$, then party R indeed offers l, and when $\frac{1}{2} + \sum_c \omega^c [E[\beta_r|H^c] - E[\beta_l|H^c]] > 1$, party R will best respond by offering r. When $\frac{1}{2} + \sum_c \omega^c [E[\beta_r|H^c] - E[\beta_l|H^c]] \in (0,1)$, then party R will offer l when $(\frac{1}{2} + \sum_c \omega^c [E[\beta_r|H^c] - E[\beta_l|H^c]])(1+\nu) < \frac{1}{2}\nu$, which amounts to $\sum_c \omega^c [E[\beta_l|H^c] - E[\beta_r|H^c]] > \frac{1}{2\zeta(1+\nu)}$. Given the above, whenever $\sum_c \omega^c [E[\beta_l|H^c] - E[\beta_r|H^c]] > \frac{1}{2\zeta(1+\nu)}$, party R offers l when party L offers l. Note that if this is the case, party L for sure offers l. An analogous condition, $\sum_c \omega^c [E[\beta_r|H^c] - E[\beta_l|H^c]] > \frac{1}{2\zeta(1+\nu)}$, guarantees that a consensus on r is the unique equilibrium. In all other cases, polarisation must arise as the unique equilibrium, that is, when

$$|\sum_c \omega^c [E[\beta_r|H^c] - E[\beta_l|H^c]]| < \frac{1}{2\zeta(1+\nu)}.$$

In the non-generic cases in which $|\sum_c \omega^c [E[\beta_r|H^c] - E[\beta_l|H^c]]| = \frac{1}{2\zeta(1+\nu)}$ both polarisation and consensus on one of the policies will be an equilibrium. □

Proof of Proposition 1: Consider (i) first. Assume first that consensus continues perpetually. As long as consensus is on the same policy, after two phases of consensus let's say at κ and $\kappa + 1$, then the voting cohorts at $\kappa+2$ includes two generations with the same polarised beliefs, implying that parties polarise by Assumption 1. Assume that there are no two such phases but also that parties never polarise. Thus consensus switches every other phase to a different policy. Suppose then that at phases $\kappa, \kappa+1$ respectively, we have consensus on l, r. But again given Assumption 1, at phase $\kappa + 2$ both voting cohorts are polarised and so parties must polarise.

Assume now that polarisation continues perpetually. As $\nu > 0$, parties polarise only if there is a strictly positive probability for each to be

[15] For recent papers on political narratives, see Levy et al. (2022), Eliaz and Spiegler (2020) and Eliaz et al. (2022).

elected. Hence in finite time in expectations, there will be two consecutive phases that include political turnover. Suppose these arise at κ and $\kappa+1$, implying that the voting cohorts of $\kappa+2$, are both in consensus on l given $\sigma \to 0$, inducing parties to be in consensus on l. Part (ii) follows then from the above and Claim A1. □

Proof of Proposition 2: Suppose now that at some phase κ, one cohort has full knowledge of the optimal policy, and the other one has only experienced one policy. Using Assumption 1, we can define two cutoffs $\bar{\omega}, \underline{\omega} \in (0,1)$, with $\bar{\omega} > \underline{\omega}$, such that:

$$\bar{\omega}(\beta_l^* - \beta_r^*) + (1-\bar{\omega})\min\{|\beta_l^* - E(\beta_r|\beta_l^*)|, |\beta_r^* - E(\beta_l|\beta_r^*)|\} = \rho$$

$$\underline{\omega}(\beta_l^* - \beta_r^*) + (1-\underline{\omega})\max\{|\beta_l^* - E(\beta_r|\beta_l^*)|, |\beta_r^* - E(\beta_l|\beta_r^*)|\} = \rho$$

Below, for $c \in \{Y^\kappa, O^\kappa\}$ let $-c = \{Y^\kappa, O^\kappa\}/c$. We can then show,

Lemma 2. *Suppose that cohort c has full knowledge of the optimal policy and is therefore cohesive and cohort -c has only full knowledge of the benefit of one policy and is therefore polarised. Then parties cater to cohort c and so are in consensus on l if* $\omega^c > \bar{\omega}$ *or* $\bar{\omega} > \omega^c > \underline{\omega}$ *and* $E[\beta_l|H^{-c}\}] - E[\beta_r|H^{-c}] = \max\{|\beta_l^* - E(\beta_r|\beta_l^*)|, |\beta_r^* - E(\beta_l|\beta_r^*)|\}$. *Otherwise parties cater to cohort* $-c$ *and are therefore polarised.*

Proof: Let c be the cohort for which $E[\beta_l|H^c] - E[\beta_r|H^c] = (\beta_l^* - \beta_r^*)$ and remember that $\omega^c = 1 - \omega^{-c}$. By the definition of $\bar{\omega}$, as $\bar{\omega} > \underline{\omega}$ and by Lemma 1, if $\omega^c > \bar{\omega}$ the parties will be in consensus. By the definition of $\underline{\omega}$ and Lemma 1 if $\bar{\omega} > \omega^c > \underline{\omega}$ and $E[\beta_l|H_t^{-c}] - E[\beta_r|H^{-c}] = \max\{|\beta_l^* - E(\beta_r|\beta_l^*)|, |\beta_r^* - E(\beta_l|\beta_r^*)|\}$ then we will have consensus. By the definition of $\underline{\omega}$ and Lemma 1 if $\bar{\omega} > \omega^c > \underline{\omega}$ and $E[\beta_l|H_t^{-c}] - E[\beta_r|H^{-c}] = \min\{|\beta_l^* - E(\beta_r|\beta_l^*)|, |\beta_r^* - E(\beta_l|\beta_r^*)|\}$ then we will have polarisation. By the definition of $\underline{\omega}$ and Lemma 1 if $\omega^c < \underline{\omega}$ we have polarisation. □

We now prove the Proposition: (i) Consider first the case in which the young are politically powerful. Assume that in phase κ parties polarise. This readily implies that the young at phase $\kappa+1$ are either in consensus, or, if there was no political turnover at phase κ, that they are polarised. And so polarisation ends immediately once there is political turnover in a polarisation stage. Assume that in phase κ parties are in consensus. This implies that the young cohort at $\kappa+1$ is polarised, and so consensus must end after one period. Consider now the case in which the old are politically powerful. Consider the possibility of two phases of polarisation $\kappa, \kappa+1$ in which there was political turnover. In phases $\kappa+2, \kappa+3$ we must have then consensus. This implies that in phases $\kappa+4, \kappa+5$ we will have polarisation, disregarding whether there was political turnover in the first stage. Thus stages of polarisation and consensus are longer compared to the case in which the young are politically powerful.

As for part (ii), note that when the young are politically powerful, consensus arises only in one stage and polarisation stops whenever political turnover arises in one phase. This means that the probability that two consecutive impressionable generations will have the same beliefs is the lowest compared to all other cases, where systematically either consensus lasts more than one phase, or polarisation lasts more than one phase even when this phase included political turnover, or both. □

Proof of Proposition 3: Consider the case where the young are the decisive cohort. Assume that at phase κ, the young arrive with knowledge that allows parties to polarise. Then assuming political turnover, consensus will already arise in phase κ, as beliefs of voters are changing. As in the main model, I^κwill fully learn the state as well. In this case, at phase $\kappa+1$, we have consensus on l. Then $I^{\kappa+1}$ will have beliefs that allow for polarisation and so on. Thus, there is as before quick succession of polarisation and consensus but polarisation once it starts may end during a phase.

Consider now the case when the old are the decisive cohort. Assume that this old cohort has beliefs that allow for polarisation. This means that both I^κ and Y^κ will learn the state with political turnover. And so in phases $\kappa+1$ and phases $\kappa+2$ parties will reach a consensus on l. But note that while $Y^{\kappa+2}$ will become $O^{\kappa+3}$ and will have beliefs that allow for polarisation, with turnover at this phase, $Y^{\kappa+3}$ will learn the true state and consensus will arise at $\kappa+4$. Thus again polarisation is short-lived compared to our main model.

Table A.1
List of all recurring political questions and controls in ANES with survey years (ANES, 2022).

Code	Question	Years
VCF0004	Survey Year	All
VCF0006	ID	All
VCF0050a	Level of political info: pre-election	1968–
VCF0050b	Level of political info: post-election	1966–
VCF0101	Age	1952–
VCF0102	Age Group	1948, 1952
VCF0104	Gender	1948–
VCF0105a/b	Race	1948–
VCF0110	Highest level education	1948–
VCF0111	Urban/suburban/rural	1952–
VCF0114	Income percentile	1948–
VCF0115	Occupation category	1952–
VCF0118	Work Status	1952–
VCF0127	Union Membership	1948–
VCF0128	Religion	1948–
VCF0136	Where grew up - urbanism	1956–1976
VCF0137	Where grew up - urbanism	1978–2000
VCF0155	Worried about losing job	1984–2008
VCF0201/2	Thermometer: Democrats/Republicans	1964–1982
VCF0218/24	Thermometer: Democratic/Republican Party	1978–2020
VCF0204	Thermometer: Catholics	1964–2012
VCF0205	Thermometer: Jews	1964–2020
VCF0206	Thermometer: Blacks	1964–2020
VCF0207	Thermometer: Whites	1964–2020
VCF0208	Thermometer: Big Business	1964–2020
VCF0210	Thermometer: Labor Unions	1964–2020
VCF0211/2	Thermometer: Liberals/Conservatives	1964–2020
VCF0213	Thermometer: Military	1964–2012
VCF0220	Thermometer: People on Welfare	1978–2012
VCF0226	Thermometer: Young People	1972–1980, 2004
VCF0301	Party Identification - Intensive	1952–
VCF0303	Party Identification - Summary	1952–
VCF0306/7	Party Identification Father/Mother	1952–1992
VCF0310	Interest in Elections	1952–
VCF0314	Number of positive mentions - Democratic	1952–2004
VCF0315	Number of negative mentions - Democratic	1952–2004
VCF0318	Number of positive mentions - Republican	1952–2004
VCF0319	Number of negative mentions - Republican	1952–2004
VCF0374	Likes anything about Democratic Party	1952–2020
VCF0380	Dislikes anything about Democratic Party	1972–2020
VCF0386	Likes anything about Republican Party	1952–2020
VCF0392	Dislikes anything about Republican Party	1972–2020
VCF0604	Trust Fed gov to do what is right	1958–2012
VCF0605	Fed Gov run by few interests or benefit of all	1964–2020
VCF0606	How much does fed gov waste tax money	1958–2020
VCF0703	Registered and votes?	1952–2020
VCF0704	Which party presidential vote	1948–2020
VCF0806	Should governments provide health insurance	1970–2020
VCF0808	Government provide guaranteed jobs	1956–1968
VCF0809	Government provide guaranteed jobs	1972–2020
VCF0823	Better if US unconcerned with RoW	1956–2020
VCF0830	Aids to Blacks	1970–2020
VCF0901a	FIPS State code	1952–2020

6.1.2. Proofs for the case of $\sigma > 0$

In this part of the Appendix we provide results and proofs analogous to the ones in the text, but for the case of $\sigma > 0$ and a large K. Denote by $\hat{\eta}_\kappa(p)$ the fraction of time in a phase κ that policy p was implemented. Note that the dynamic evolution of policies involves some randomness, given the voting shock ϕ_t and the policy shock ϵ_t (through the latter's effect on beliefs). This then induces a probability distribution P over the set of infinite paths of history $\mathbb{H}$. Thus, when we write "almost surely", here and in the Appendix, we mean P-almost surely on $\mathbb{H}$.

Claim A2: *Suppose that* $\sigma > 0$, *and that K is sufficiently large. A cohort that has its impressionable years in a phase of party polarisation will almost*

Table A.2
Summary of thermometer – labor unions.

Scale	Birth cohort										All cohorts
	1907–1916	1917–1926	1927–1936	1937–1946	1947–1956	1957–1966	1967–1976	1977–1986	1987–1996	1997–2006	
	%age of respondents										
0 (Unfavourable)	4.38 %	4.02 %	4.82 %	5.63 %	5.17 %	4.82 %	4.33 %	3.24 %	3.25 %	2.13 %	4.65 %
10	3.68 %	3.77 %	4.16 %	5.24 %	4.51 %	4.23 %	4.00 %	2.50 %	1.99 %	0.80 %	4.07 %
20	0.75 %	0.95 %	0.84 %	0.93 %	1.23 %	1.17 %	1.36 %	0.87 %	0.65 %	0.80 %	1.04 %
30	5.98 %	6.13 %	6.79 %	7.36 %	7.32 %	6.05 %	5.20 %	5.22 %	3.49 %	3.20 %	6.28 %
40	8.86 %	9.69 %	11.73 %	12.13 %	11.39 %	10.15 %	8.25 %	7.00 %	6.05 %	2.93 %	10.05 %
50	18.20 %	17.97 %	18.53 %	17.77 %	17.40 %	20.07 %	21.28 %	24.39 %	28.88 %	27.73 %	19.72 %
60	13.68 %	14.08 %	13.14 %	13.56 %	14.60 %	13.58 %	12.66 %	12.58 %	12.47 %	10.13 %	13.53 %
70	12.22 %	13.21 %	12.83 %	11.90 %	12.64 %	12.41 %	13.69 %	13.38 %	10.84 %	10.67 %	12.60 %
80	8.46 %	9.88 %	8.24 %	9.07 %	9.10 %	9.61 %	10.40 %	9.88 %	10.07 %	11.47 %	9.40 %
90 (Favourable)	8.28 %	8.00 %	8.16 %	6.03 %	6.11 %	6.78 %	6.92 %	8.41 %	9.50 %	14.93 %	7.19 %
NA	15.50 %	12.30 %	10.76 %	10.37 %	10.52 %	11.14 %	11.91 %	12.53 %	12.79 %	15.20 %	11.46 %
No. of Respondents	2258	4324	4878	7495	9930	8639	5056	3887	2462	375	49,304

Table A.3
Initial regression coefficients for different cohorts.

	(1) Thermometer – Labor Unions
1917–1926	−0.230
	(0.757)
1927–1936	−2.024**
	(0.993)
1937–1946	−4.108***
	(1.308)
1947–1956	−3.078*
	(1.660)
1957–1966	−2.966
	(2.013)
1967–1976	−1.979
	(2.416)
1977–1986	−0.669
	(2.806)
1987–1996	−1.768
	(3.203)
1997–2006	−2.073
	(3.702)
Constant	64.713***
	(1.695)
Observations	49,304

Standard errors in parentheses. Omitted Birth cohort is 1907–1916. The regression includes added controls for respondents' age group, gender, race, education and year of survey. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

surely have beliefs concentrating on β^*, and a cohort that has its impressionable years in a phase of party consensus will almost surely have beliefs that concentrate on $\beta_l^* - E(\beta_r|\beta_l^*)|$ or on $|\beta_r^* - E(\beta_l|\beta_r^*)|$, depending on which policy was implemented.

Proof of Claim A2: When K is sufficiently large, a history of a polarisation phase implies that the beliefs of the impressionable cohort will almost surely concentrate on β^*, as each of the policies, l and r, is implemented with a strictly positive probability, and so at phase κ we must have $\frac{\hat{\eta}_\kappa(l)}{\hat{\eta}_\kappa(r)} \to c_i$ for some finite, non zero, c_i. On the other hand a consensus phase implies that almost surely beliefs concentrate on $\beta_l^* - E(\beta_r|\beta_l^*)|$ or on $|\beta_r^* - E(\beta_l|\beta_r^*)|$, depending which policy was implemented, as the cohort can learn in the limit the utility from the policy implemented and nothing else on the alternative policy. □

Claim A3: (i) Following any consecutive two phases $\kappa, \kappa+1$, where in both phases only one policy was implemented (potentially different ones in each phase), the next phase $\kappa+2$ will almost surely have parties polarising. (ii) Following any consecutive two phases $\kappa, \kappa+1$, where in both phases parties were polarised, the next phase $\kappa+2$ will almost surely have parties being in consensus on l.

Proof of Claim A3: This follows from Claim A2, Assumption 1, and noting that at phase $\kappa+2$ the two cohorts will either both be in consensus on l or both be polarised. □

Following on from Claims A2 and A3, all the results in the text follow almost surely when instead of $\sigma \to 0$ we consider a large enough K.

6.2. Appendix B: additional empirical analysis

In this Appendix we report on an additional empirical exercise which we conducted to check whether different cohorts (defined by years of birth) have different levels of polarisation in their political preferences or opinions. The data that we use is the same as for the computation of the Allison–Foster index reported in Section 2, with additional data that we use as controls.

Specifically, we use the ANES survey data. We used the cumulative data set, which included all answers from respondents for every edition of the ANES dataset (see ANES, 2022).

Description of Data: The raw data records answers of respondents to the different ANES questions. As explained in Section 2, we selected all the questions that have some relevance to political opinions. This resulted in a total of 33 questions. To maximise the number of questions that we can include, we focused on the questions that have been asked consistently in all the surveys and at least from the 1970s onwards. We excluded questions that seemed less relevant (e.g., questions on attitudes towards Jews, Blacks, Catholics, Whites) as well as questions such as whether one likes anything about the Democratic party, and focused on clear questions about political preferences. This process resulted in ten questions on which the respondents provide their opinions. The full set of initial questions that have been consistently asked is provided in Table A.1.

The ten questions that we focus on are:
1: Government benefits all (VCF0605),
2: Government should provide health (VCF0806),
3: Government should provide jobs (VCF0809),
4. Government wastes taxes (VCF0606),
5: Thermometer big business (VCF0208),
6: Thermometer labor unions (VCF0210),
7: Thermometer military (VCF0213),
8: Thermometer Welfare (VCF0220),
9: Thermometer trust in Government (VCF0604),
10: [US is better off if] Unconcerned about the rest of the world (VCF0823).

Answers to the questions are coded in the survey from 0–100 in Thermometer questions, 1–7 in questions 1–4, and 1–2 for question 10.

Table A.4
Breusch–Pagan regression coefficients for different cohorts.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Economy prediction	Thermometer	Community involvement	Curiousity or manners	Obidient or self-reliant	Considerate or well behaved	Independence or respect elders
1917–1926	0.038	0.005	0.105*	−0.028	−0.016	0.000	0.067
	(0.032)	(0.025)	(0.054)	(0.065)	(0.057)	(0.069)	(0.075)
1927–1936	0.005	−0.007	0.114**	0.026	0.044	0.036	0.131*
	(0.031)	(0.024)	(0.051)	(0.062)	(0.054)	(0.065)	(0.071)
1937–1946	0.031	−0.018	0.083*	0.167***	0.144***	−0.011	0.258***
	(0.029)	(0.023)	(0.049)	(0.059)	(0.052)	(0.063)	(0.068)
1947–1956	0.013	0.019	0.086*	0.188***	0.154***	−0.026	0.287***
	(0.028)	(0.022)	(0.048)	(0.059)	(0.051)	(0.062)	(0.067)
1957–1966	0.016	0.082***	0.114**	0.118**	0.143***	0.018	0.274***
	(0.028)	(0.023)	(0.048)	(0.059)	(0.051)	(0.062)	(0.067)
1967–1976	0.010	0.099***	0.051	0.134**	0.140***	0.077	0.373***
	(0.029)	(0.025)	(0.049)	(0.059)	(0.052)	(0.062)	(0.068)
1977–1986	−0.024	0.089***	0.025	0.195***	0.175***	0.045	0.507***
	(0.030)	(0.027)	(0.049)	(0.059)	(0.052)	(0.062)	(0.068)
1987–1996	−0.024	0.129***	0.040	0.223***	0.198***	0.064	0.561***
	(0.032)	(0.031)	(0.050)	(0.059)	(0.052)	(0.063)	(0.068)
1997–2006	−0.128**	0.014	0.008	0.296***	0.175***	0.110	0.644***
	(0.052)	(0.071)	(0.061)	(0.069)	(0.061)	(0.073)	(0.079)
Constant	0.611***	0.597***	0.609***	0.600***	0.690***	0.707***	0.308***
	(0.027)	(0.020)	(0.047)	(0.058)	(0.051)	(0.061)	(0.066)
Observations	41,576	40,816	25,535	22,457	22,422	22,468	22,448

Standard errors in parentheses. Omitted Birth cohort is 1907–1916. *$p < 0.10$, **$p < 0.05$, ***$p < 0.01$

Variables and Controls: We have grouped individuals into cohorts according to their birthdate. We have used different year bands to define cohorts, 8 year bands, 10 year bands and 12 year bands. We report here the results from 10-year bands (the results were very similar when using other bands). Specifically, the cohort variables were defined as birth_cohort_x for $x \in \{1, \ldots, 9\}$. For example if $x = 4$ this will include individuals who were born in the ten year period of 1947–1956.

Age was used as a control and we try three different types of control, one in which age is grouped in four-year bands (e.g., 18–22 and so on), and one in which ages of respondents are grouped in eight-year bands (e.g., 18–25 and so on) and 12-year bands. The results were very similar with both types of age frames and below we report the results for the age control with the four-year bands. There are additional dummy variables for Presidential election years and for survey years. Crucially, when comparing multiple survey waves over time, there is a need to control for age of respondents and survey wave, however this brings forward the age-period-cohort problem. We use the standard method in the literature of grouping ages together to circumvent this problem.

Beyond age and time dummies described above, we use the following controls, which are the respondents' answers to questions VCF0104 (gender), VCF0105 a (race), VCF0110 (education) and VCF0303 (demorep, that is, whether one's views are represented by the democratic party).

Table A.2 shows an example of the summary statistics of the dependent variables for Question 6, throughout the different survey waves.

Empirical Strategy: We proceed in two steps. We first run an initial regression, using the above explanatory variables (cohort dummies, age dummies, survey year dummies and controls) to explain the dependent variable, which is the responses (throughout the years) to a survey question. Specifically, we first run the following regression,

$$Opinion_{ic} = a + \beta_c I_{ic} + \beta_X X_i + \mu_i$$

where i is the respondent and c the cohort, I_{ic} represents a cohort effect and X_i a set of controls, including age group, year, sex, race, education, and respondent's political identity. We ran different versions of this regression allowing for different sets of controls. Table A.3 is an example of one of the initial regression results, for Question 6.

As can be seen, there are naturally differences in the average responses of each cohort (the first moments). Our main interest though is the second moments; given the above initial regressions, our next step is to run the Breusch-Pagan (BP) Regression. The aim is to check how cohorts "contribute" to our inability to explain the variance in the answers. The Breusch-Pagan test measures whether heteroskedasticity is present in a regression, i.e., whether the residuals are potentially distributed with different variances at each level of the variables. We therefore regress birth cohorts on the *square of the residuals* from the initial regressions. We interpret then the coefficients on the cohorts dummies as representing the variance of the error terms for each cohort. Again, our theory implies that a cohort that has weaker political preferences will imply less ability for us as researchers to predict their opinions.

Results: Table A.4 reports the results of these regressions for all the questions (which were run using the initial regressions both with and without controls). In all the questions differences between some cohort pairs are statistically significant.[16] In most of the questions (see columns 1–3,5, 6, 8, and 9), we find a trend indicating a downward shift in the variance of public opinion as cohorts become younger and a small rise in the youngest cohorts. A good example of this is the Thermometer question on Big Business in column 5. This trend is consistent with an interpretation that the polarisation level of cohorts goes down with time, and potentially then goes up. For some other questions the trend is not as clear.[17]

[16] To check for binary differences of the coefficients of different cohorts we ran the Wald test. Table A.5 provides the matrices of the p-values of coefficient equality test for all questions. As can be seen, there are statistically significant differences across the coefficients of the BP regression.

[17] These questions are Government waste taxes, Government should invest in military and Unconcerned in the rest of the world, in columns 4, 7, 10.

Table A.5
Coefficient equality test *p*-values.

	1907–1916	1917–1926	1927–1936	1937–1946	1947–1956	1957–1966	1967–1976	1977–1986
Panel A: Govt. benefits all								
1917–1926	0.111	.	.	.	.	.	.	.
1927–1936	0.739	0.115	.	.	.	.	.	.
1937–1946	0.060	0.000	0.004	.	.	.	.	.
1947–1956	0.000	0.000	0.000	0.000	.	.	.	.
1957–1966	0.001	0.000	0.000	0.037	0.142	.	.	.
1967–1976	0.000	0.000	0.000	0.000	0.140	0.008	.	.
1977–1986	0.000	0.000	0.000	0.000	0.000	0.000	0.014	.
1987–1996	0.000	0.000	0.000	0.000	0.000	0.000	0.003	0.454
Panel B: Govt. should provide health insurance								
1917–1926	0.157	.	.	.	.	.	.	.
1927–1936	0.001	0.024	.	.	.	.	.	.
1937–1946	0.000	0.000	0.063	.	.	.	.	.
1947–1956	0.000	0.000	0.003	0.219	.	.	.	.
1957–1966	0.000	0.000	0.000	0.017	0.187	.	.	.
1967–1976	0.000	0.000	0.000	0.014	0.133	0.736	.	.
1977–1986	0.000	0.000	0.087	0.997	0.280	0.035	0.027	.
1987–1996	0.000	0.000	0.016	0.338	0.941	0.393	0.291	0.371
Panel C: Govt. should provide jobs								
1917–1926	0.000	.	.	.	.	.	.	.
1927–1936	0.000	0.000	.	.	.	.	.	.
1937–1946	0.000	0.000	0.000	.	.	.	.	.
1947–1956	0.000	0.000	0.000	0.000	.	.	.	.
1957–1966	0.000	0.000	0.000	0.029	0.001	.	.	.
1967–1976	0.000	0.000	0.000	0.029	0.019	0.749	.	.
1977–1986	0.000	0.000	0.000	0.724	0.000	0.034	0.030	.
1987–1996	0.000	0.000	0.001	0.003	0.000	0.000	0.000	0.015
Panel D: Govt. wastes taxes								
1917–1926	0.367	.	.	.	.	.	.	.
1927–1936	0.085	0.332	.	.	.	.	.	.
1937–1946	0.002	0.011	0.124	.	.	.	.	.
1947–1956	0.044	0.201	0.862	0.102	.	.	.	.
1957–1966	0.855	0.323	0.034	0.000	0.006	.	.	.
1967–1976	0.243	0.013	0.000	0.000	0.000	0.070	.	.
1977–1986	0.000	0.000	0.000	0.000	0.000	0.000	0.003	.
1987–1996	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.057
Panel E: Thermometer – Big Businesses								
1917–1926	0.193	.	.	.	.	.	.	.
1927–1936	0.128	0.805	.	.	.	.	.	.
1937–1946	0.009	0.133	0.207	.	.	.	.	.
1947–1956	0.000	0.000	0.000	0.004	.	.	.	.
1957–1966	0.000	0.000	0.000	0.012	0.997	.	.	.
1967–1976	0.001	0.014	0.022	0.164	0.493	0.522	.	.
1977–1986	0.005	0.047	0.066	0.283	0.515	0.535	0.943	.
1987–1996	0.352	0.652	0.710	0.989	0.393	0.399	0.579	0.617
Panel F: Thermometer – Labor Unions								
1917–1926	0.130	.	.	.	.	.	.	.
1927–1936	0.187	0.778	.	.	.	.	.	.
1937–1946	0.140	0.835	0.917	.	.	.	.	.
1947–1956	0.062	0.811	0.555	0.583	.	.	.	.
1957–1966	0.017	0.354	0.193	0.175	0.377	.	.	.
1967–1976	0.007	0.154	0.076	0.064	0.144	0.488	.	.
1977–1986	0.000	0.001	0.000	0.000	0.000	0.002	0.027	.
1987–1996	0.001	0.020	0.009	0.007	0.016	0.070	0.236	0.484
Panel G: Thermometer – Military								
1917–1926	0.693	.	.	.	.	.	.	.
1927–1936	0.830	0.454	.	.	.	.	.	.
1937–1946	0.950	0.549	0.836	.	.	.	.	.
1947–1956	0.002	0.001	0.000	0.000	.	.	.	.
1957–1966	0.447	0.671	0.212	0.259	0.002	.	.	.
1967–1976	0.402	0.585	0.213	0.256	0.017	0.848	.	.
1977–1986	0.000	0.000	0.000	0.000	0.002	0.000	0.000	.
1987–1996	0.000	0.000	0.000	0.000	0.012	0.000	0.000	0.793

(*continued on next page*)

Table A.5 *(continued)*

	1907–1916	1917–1926	1927–1936	1937–1946	1947–1956	1957–1966	1967–1976	1977–1986
Panel H: Thermometer – Welfare								
1917–1926	0.674	.	.	.	.	.	.	.
1927–1936	0.977	0.554	.	.	.	.	.	.
1937–1946	0.529	0.151	0.397	.	.	.	.	.
1947–1956	0.090	0.003	0.015	0.082	.	.	.	.
1957–1966	0.436	0.091	0.276	0.829	0.096	.	.	.
1967–1976	0.645	0.973	0.506	0.114	0.001	0.061	.	.
1977–1986	0.848	0.810	0.784	0.326	0.025	0.241	0.779	.
1987–1996	0.064	0.084	0.031	0.006	0.000	0.004	0.080	0.070
Panel I: Thermometer – Trust in govt.								
1917–1926	0.914	.	.	.	.	.	.	.
1927–1936	0.255	0.138	.	.	.	.	.	.
1937–1946	0.004	0.000	0.039	.	.	.	.	.
1947–1956	0.000	0.000	0.001	0.211	.	.	.	.
1957–1966	0.001	0.000	0.006	0.435	0.691	.	.	.
1967–1976	0.000	0.000	0.002	0.120	0.505	0.351	.	.
1977–1986	0.403	0.413	0.091	0.004	0.000	0.001	0.000	.
1987–1996	0.660	0.686	0.339	0.090	0.035	0.050	0.022	0.918
Panel J: US should be unconcerned about rest of the world								
1917–1926	0.026	.	.	.	.	.	.	.
1927–1936	0.067	0.611	.	.	.	.	.	.
1937–1946	0.000	0.013	0.002	.	.	.	.	.
1947–1956	0.000	0.094	0.020	0.266	.	.	.	.
1957–1966	0.262	0.091	0.255	0.000	0.000	.	.	.
1967–1976	0.001	0.000	0.000	0.000	0.000	0.000	.	.
1977–1986	0.000	0.000	0.000	0.000	0.000	0.000	0.000	.
1987–1996	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.022

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The data are available from ANES.

References

Allison, R.A., Foster, J.E., 2004. Measuring health inequality using qualitative data. J. Health. Econ. 23(3):505–524.

Alwin, D.F., Krosnick, J.A., 1991. Aging, cohorts, and the stability of sociopolitical orientations over the life span. Am. J. Sociol. 97(1):169–195.

American National Election Studies, 2022. ANES time series cumulative data file [dataset and documentation]. September 16, 2022 version. www.electionstudies.org. Available on https://electionstudies.org/data-center/anes-time-series-cumulative-data-file/.

Barber, M., McCarty, N., 2013. The causes and consequences of polarisation. In: Cathie Jo, M., Jane, M., (Eds.), Negotiating Agreement in Politics, APSA task force report.

Battaglini, M., Coate, S., 2008. A dynamic theory of public spending, taxation, and debt. Am. Econ. Rev. 98, 201–236.

Boxell, L., Gentzkow, M., Shapiro, J., 2017. Greater internet use is not associated with faster growth in political polarisation among US demographic groups. Proc. Natl. Acad. Sci. USA 114, 10612–10617.

Calvert, R.L., 1985. Robustness of the multidimensional voting model: candidate motivations, uncertainty, and convergence. Am. J. Pol. Sci. 29, 69–95.

Eliaz, K., Spiegler, R., 2020. A Model of competing narratives. Am. Econ. Rev. 110, 3786–3816.

Eliaz, K., Spiegler, R., Galperti, S., 2022. False narratives and political mobilization. Mimeo.

Ghitza, Y., Gelman, A., Auerbach, J., 2022. The great society, Reagan's revolution, and generations of presidential voting. Am. J. Pol. Sci. doi:https://doi.org/10.1111/ajps.12713

Levy, G., Razin, R., 2025. Political social-learning: short-term memory and cycles of polarisation. Am. Econ. Rev. 115(2):635–659.

Levy, G., Razin, R., Young, A., 2022. Misspecified politics and the recurrence of populism. Am. Econ. Rev. 112, 928–962.

Lindbeck, A., Weibull, J., 1987. Balanced-budget redistribution as the outcome of political competition. Public Choice 52, 273–297.

Malmendier, U., Nagel, S., Yang, Z., 2021. The making of hawks and doves. J. Monet. Econ. 117, 19–42.

Malmendier, U., Nagel, S., 2016. Learning from inflation expectations. Q. J. Econ. 131, 53–87.

Ortoleva, P., Snowberg, E., 2015. Overconfidence in political behavior. Am. Econ. Rev. 105, 504–535.

Strömberg, D., 2004. Radio's impact on public spending. Q. J. Econ. 119(1):189–221.