

A Generational Model of Political Learning

Larry M. Bartels

Department of Political Science, Vanderbilt University
Nashville, TN, USA 37203-5721
larry.bartels@vanderbilt.edu

Simon Jackman

Department of Political Science, Stanford University
Stanford, CA, USA 94305-6044
jackman@stanford.edu

REVISED: 19 February 2013

We propose a mathematical framework for modeling opinion change using large-scale longitudinal data sets. Our framework encompasses two varieties of Bayesian learning theory as well as Mannheim's theory of generational responses to political events. The basic assumptions underlying the model are (1) that historical periods are characterized by shocks to existing political opinions, and (2) that individuals of different ages may attach different weights to those political shocks. Political generations emerge endogenously from these basic assumptions: the political views of identifiable birth cohorts differ, and evolve distinctively through time, due to the interaction of age-specific weights with period-specific shocks. We employ this model to examine generational changes in party identification using survey data from the 1952-2008 American National Election Studies.

Keywords: Bayesian learning theory; party identification; partisan change; running tally; generational imprinting.

Do political experiences in adolescence and early adulthood powerfully shape subsequent attitudes and behavior? Are citizens responsive to political information throughout their lives? Is their incorporation of new information consistent with the precepts of rational political learning? And how do individual responses to political events add up to large-scale shifts in political opinions over decades or generations?

Our aim here is to develop and apply a mathematical model of opinion change that can shed light on these questions. The model is quite simple, with only two “moving parts”—(1) a sequence of parameters characterizing the political events of successive historical periods and (2) a sequence of parameters representing the distinctive weights attached to those events by individuals at various points in the life-cycle. Despite its simplicity, the model is sufficiently flexible to subsume a variety of more specific models, ranging from Bayesian learning theories (Achen 1992; Gerber and Green 1998) to Mannheim’s (1952) generational theory of political change.

Our empirical analysis employs longitudinal data to estimate the two sequences of parameters that interact to produce political change over time and across generations. In this respect, our work is firmly situated in the long tradition of “cohort analysis” exemplified by the work of Campbell et al. (1960, chap. 7), Carlsson and Karlsson (1970), Converse (1969; 1976), Abramson (1975), Markus (1983), and Miller and Shanks (1996), among others. However, our analysis departs in some important ways from the theoretical framework underlying most of this work, the so-called “Age-Period-Cohort” (APC) framework.

The APC framework treats political attitudes or behavior as a function—typically a simple

additive function—of variables indexing each individual's age, year of observation ("period"), and year of birth ("cohort"). One fundamental (and much-remarked) problem with this approach is that the model in its general form is statistically underidentified; the age, period, and cohort variables are collinear. As a result, distinct age, period, and cohort effects cannot be disentangled without additional, more or less arbitrary assumptions.

An even more significant (but less-remarked) limitation of the APC framework is that it lacks a clear theoretical rationale. As Markus (1983, 720) put it, "the APC model is primarily an accounting equation rather than an explanatory one. That is, its purpose is to partition variation into distinct bundles (age, period, cohort); it is rarely purported to represent in mathematical form the underlying process generating the observed data."

In combination, these two limitations of the APC framework seem to produce a pernicious mixture of empirical complexity and conceptual confusion. For example, analysts often overcome the identification problem by imposing constraints on the various age, period, and cohort effects that have odd and presumably unintended implications for observed patterns of generational change. And since no analyst really believes that an individual's birth year has a direct causal impact on her political attitudes or behavior, "cohort" effects in the APC framework are often loosely interpreted as reflecting some mysterious combination of past "period" effects and (unmodeled) generational imprinting—a state of affairs that inspired Converse (1976, 80) to wonder "who is responsible for insisting on the conceptual partition between 'period effects' and 'generational effects,' since it seems to verge on a distinction without a difference."

Our aim here is to develop an alternative to the APC framework that is more explicitly grounded in theories of political learning. Rather than attempting to partition observed variance

into additive “period” and “cohort” components by brute force, we posit a single process of political learning in which the two important elements are (1) period-specific “shocks” reflecting the distinctive political events of a given time period, and (2) age-specific “weights” reflecting the extent to which these shocks are internalized by individuals at various points in the life-cycle. Generational patterns of political change arise endogenously from the interaction of these basic elements—a form of interaction that cannot be captured within the conventional additive APC framework.¹

Our empirical analysis focuses on partisan learning as reflected in changing levels of party identification in the American National Election Studies (ANES) data from 1952 through 2008. Party identification provides a useful empirical focus for a variety of reasons as cogent now as they were for Converse a generation ago, including “the demonstrated importance of these abiding feelings of party attachments in the determination of voting choices,” the fact that “few time series available in sample survey data ... are now as rich or as long,” and the fact that partisanship has been “a workhorse in efforts to develop a proper methodology of cohort

¹ Brady and Elms (1999) proposed an APC model of political participation in which period effects multiply age- and cohort-specific baseline levels of participation. Their model, like ours, transcends the conventional additive APC framework; but from the perspective proposed here it does so in a way that is precisely backward. For Brady and Elms, time periods may be more or less intense but have no distinctive political character, while cohort-specific baseline levels of participation presumably reflect (unspecified) differences in political socialization. In the model proposed here, the defining characteristic of each time period is a distinctive political “shock”; age plays the role of multiplier, causing individuals of different ages to respond to these period effects more or less intensely.

analysis" (Converse 1976, 9). However, our framework is intended to be sufficiently general to apply to a wide variety of political attitudes and beliefs, and we shall conclude by calling for just such comparative analysis.

1. Three Models of Political Learning

We begin by describing three distinct models of political learning through the life-cycle. There are strong family resemblances among these three distinct models, and indeed all three will turn out to be encompassed by a more general framework in which each appears as a special case. The more general framework will serve as the basis for our empirical analysis, so that data can be used to assess the adequacy of the special assumptions underlying each of the distinct models presented here. However, from the standpoint of fixing ideas it seems preferable to proceed from the simplest model to more complex alternatives rather than beginning with the general framework and deriving the distinct models as special cases.

All three models are *generational* in the sense that they differentiate individuals at any point in time solely on the basis of when they were born. It would certainly be possible—and, for some purposes, desirable—to elaborate this framework in order to incorporate other politically relevant individual characteristics. However, in the simple formulation offered here, all the respondents of similar age in a given survey are treated as homogeneous, and their political views are accounted for on the basis of the sequence of political events they experience in common at different points in their shared life-history up to the point at which they are

interviewed.²

In each model, political events are represented as a sequence of “shocks” to existing partisanship. We do not measure these events directly, but infer their effects from observed changes in partisanship. For example, non-southern whites became noticeably more Republican between 1982 and 1984, and we infer that political events in that two-year period were, on balance, perceived as favorable to the Republican Party. Our analysis provides a quantitative estimate of the magnitude of the corresponding partisan shock, but does not attempt to attribute it to specific events such as the robust economic recovery from the recession of 1981-82 or the respective presidential campaigns of Ronald Reagan and Walter Mondale.

The three models differ in how individuals are supposed to incorporate relevant information at various points in their lives. In the first model, an individual’s attitude or belief is based upon a simple arithmetic average of all the relevant political events she has experienced throughout her life. In the second model, recent events receive more weight than those experienced in the distant past. In the third model, events experienced during “formative” years (perhaps during adolescence and early adulthood) receive more weight than those experienced earlier or later in life.

A key simplifying assumption common to all three models is that age has no independent impact on partisan attitudes. If members of a given cohort become more Republican (or

² In our empirical analysis of party identification we render this assumption more plausible than it would otherwise be by excluding African-Americans and southerners—two groups that responded very distinctively to political events in the historical period covered by our analysis.

Democratic) as they grow older, their changing partisanship is assumed to reflect the accumulation of specific political experiences to which they have been exposed rather than the process of aging per se. Moreover, the political experiences to which they have been exposed are assumed to be unaffected by their age at the time of exposure. Thus, a given war, scandal, or economic boom cannot be interpreted by some cohorts as a pro-Republican event but by other cohorts as a pro-Democratic event.³

These simplifying assumptions will obviously be more or less plausible depending on the specific research context. For example, direct effects of aging seem more likely to be consequential in an analysis of political involvement than in an analysis of party identification (Converse 1969; Converse 1976; Brady and Elms 1999), and political events with clear age-specific ramifications (say, the imposition of a military draft or the collapse of a public pension system) seem more likely than other kinds of events to produce distinctive responses in different birth cohorts (Delli Carpini 1989). Limitations of this sort notwithstanding, the framework developed here seems sufficiently flexible to be potentially fruitful in analyzing a variety of specific instances of political learning.

1.1. The Running Tally

The first model we consider is a simple Bayesian learning model (Zeckman 1979; Achen 1992). The model posits a sequence of observations $\{\pi_{c+1}, \dots, \pi_i, \dots, \pi_t\}$, with

³ However, the *impact* of such an event may, and in general will, be age-dependent—a point we shall elaborate in due course.

$$\{1\} \quad \pi_i = \theta + \delta_i,$$

where θ is a constant (unknown) parameter and δ_i is a normally-distributed random variable with mean zero and (known) variance $\psi^2 > 0$.⁴ In Achen's (1992) formulation, θ represents a partisan differential in expected (prospective) benefits for a given citizen under Republicans and Democrats, and the observations $\{\pi_{c+1}, \dots, \pi_i, \dots, \pi_t\}$ represent a sequence of actual benefits experienced by that citizen in successive time periods, beginning with the first period following her birth (period $c+1$) and continuing through the current period t .⁵ Benefits in each period are more or less affected by random factors (reflected by the variance ψ^2); however, favorable or unfavorable experience with either party inclines the citizen to adjust her beliefs about the unknown partisan differential θ accordingly.⁶ For the specific case of party identification, Achen's model represents a mathematical formalization of Fiorina's (1977, 611) notion of party identification as a "running balance sheet" of retrospective evaluations of the competing parties'

⁴ More realistically, the variance ψ^2 may also be treated as an unknown parameter. That complication is of no great importance, and is ignored here.

⁵ Focusing on a single quantity, the partisan differential, entails no loss of generality in the context of a two-party system. Applying the same model to a multi-party system would presumably require separate "running tallies" reflecting benefits experienced under each party in the system. The resulting analysis would be more complex, but not fundamentally altered.

⁶ Achen's formulation includes informative prior beliefs produced by parental socialization in addition to direct experience. Since the available longitudinal data do not include reliable information on parental partisanship, we dispense with that complication here and assume that each cohort begins life without partisan predispositions.

performance in office.

In this simple “running tally” model, the best estimate $\mu_{c,t}$ of the underlying value θ at time t for an individual born in cohort c is a simple average of the observed π values from period $c+1$ through period t :

$$\{2\} \quad \mu_{c,t} = (\sum_{c < i \leq t} \pi_i) / (t - c).$$

Since the distribution of expected benefits is assumed to be constant over time, inferences about the unknown value θ are unaffected by the temporal order of the observations—good or bad experience in the distant past is just as relevant as today’s headlines. However, the *incremental* impact of each new observation does depend on the number of previous observations, as may be seen from the recursive relationship between the current estimate $\mu_{c,t}$ and the previous estimate $\mu_{c,t-1}$:

$$\{3\} \quad \mu_{c,t} = \mu_{c,t-1} (t - c - 1) / (t - c) + \pi_t / (t - c) .$$

For individuals in a given birth cohort c , the marginal impact of each new experience π_t declines over the life-cycle (that is, the weight $1/(t-c)$ attached to π_t declines as t increases) even though experiences at different times enter identically in equation {2}. Alternatively, at any given time t , the marginal impact of a new experience π_t is greater for younger cohorts (since the weight $1/(t-c)$ attached to π_t increases as c increases). As Achen (1992) pointed out, this fact provides a theoretical rationale for a familiar pattern in the empirical literature on political attitudes: the greater stability of observed opinions among the old than among the young.

Achen’s (1992) discussion of the Bayesian model focused primarily on the dynamic relationship between current and past partisanship represented by equation {3} rather than on

the reduced-form impact of the sequence of period shocks highlighted by equation {2}. Thus, he suggested models and methods for capturing the age-dependent relationship between current and past partisanship in panel data—in effect, providing a more explicit theoretical rationale for empirical work along similar lines by Franklin and Jackson (1983), Markus (1983), and others. However, since our approach here is predicated on the use of longitudinal data rather than panel data, the key relationship for our purposes is equation {2}, which expresses the partisan differential $\mu_{c,t}$ as a simple average of the period shocks $\{\pi_{c+1}, \dots, \pi_i, \dots, \pi_t\}$.

1.2. Partisan Change and Temporal Discounting

The simple “running tally” model assumes that the underlying political value of interest, θ , remains constant over a citizen’s entire lifetime. Gerber and Green (1998) proposed a more realistic dynamic model in which θ is subject to temporal drift:

$$\{4\} \quad \theta_t = \gamma \theta_{t-1} + \varepsilon_t,$$

where γ is a constant parameter between zero and one and ε_t is a normally-distributed random variable with mean zero and variance σ^2 . As Gerber and Green (1998) noted, the “running tally” model is a special case of this more general model in which $\gamma=1$ and $\sigma^2=0$.⁷

The model defined by equations {1} and {4} does not fit neatly within the framework proposed here, because the optimal estimate $\mu_{c,t}$ of the current value θ_t for a member of birth cohort c at time t is not a weighted average of the observed π values. However, setting $\gamma=1$

⁷ The more general model is in turn subsumed within a broader class of Dynamic Linear Models analyzed by West and Harrison (1997).

while allowing $\varphi^2 > 0$ produces a somewhat simpler random-walk model in which the optimal estimate $\mu_{c,t}$ of θ_t does turn out to be a weighted average of the period shocks $\{\pi_{c+1}, \dots, \pi_i, \dots, \pi_t\}$:

$$\{5\} \quad \mu_{c,t} = (\sum_{c < i \leq t} \omega_{c,i} \pi_i) / (\sum_{c < i \leq t} \omega_{c,i}).$$

The optimal weight $\omega_{c,i}$ for observation π_i for citizens in birth cohort c is

$$\{6\} \quad \omega_{c,i} = [(v_{c,i-1} + \varphi^2) / v_{c,i-1}] \omega_{c,i-1},$$

where $v_{c,i-1}$ is the variance of the posterior distribution of $\mu_{c,i-1}$.

Since $\varphi^2 > 0$, $[(v_{c,i-1} + \varphi^2) / v_{c,i-1}] > 1$ in equation {6}, which implies $\omega_{c,i} > \omega_{c,i-1}$. Because the underlying partisan differential varies over time, the optimal weights attached to each period's shock increase with the index value i ; viewed from the current period t , recent experience is more relevant than older experience for inferring the current value θ_t .⁸ This pattern of optimal weights distinguishes the "partisan change" model from the simple "running tally" model. When θ is fixed, the optimal weights for observations $\{\pi_{c+1}, \dots, \pi_i, \dots, \pi_t\}$ are identical. But when θ is allowed to change over time, the optimal weights for observations from more recent periods (or, equivalently, for events experienced later in life) are greater than those attached to similar events experienced earlier.

⁸ The loss of generality entailed by setting $\gamma=1$ is not trivial, since it may well be reasonable to expect party identification to erode if it is not bolstered by additional favorable experience. This sort of erosion provides an additional reason for past experience to be discounted, implying an even more marked increase in age-specific weights over the course of the life-cycle. Thus, our empirical test of the random-walk model also provides a "conservative" qualitative test of the more general model.

1.3. Generational Imprinting

In discussing the limitations of their Bayesian model of political learning, Gerber and Green (1998, 815) argued that “any such model must account for the ‘period effects’ that stamp different generations with distinctive partisan coloration. ... Learning models that stress the influence of contemporaneous information have difficulty explaining the persistence of early formative experiences.” Indeed, at least since the publication of Mannheim’s (1952) classic essay on “The Problem of Generations,” the notion that early formative experiences have some special power to stamp different generations with distinctive political attitudes and beliefs has been one of the most familiar and influential ideas in the literature on political socialization.

Within the framework proposed here, a natural way to represent the notion of generational imprinting is to suppose that the weights attached by each citizen to experiences in successive time-periods are a function of when in her life they occur. For a citizen in birth cohort c , the relevant weights are $\{\omega_1, \dots, \omega_i, \dots, \omega_{t-c}\}$, where $(t-c)$ represents the age of birth cohort c at time t . If “early formative experiences” are especially influential, then—in contrast to the “running tally” and “partisan change” models—the earlier elements in this sequence of weights will be larger in magnitude than the later elements.

It is not entirely clear from the literature what, exactly, is supposed to make these formative experiences so important. For example, Markus (1983, 723) noted that generational analysts “posit that the socializing experiences of late adolescence and early adulthood are of crucial importance in forming political outlooks because of the heightened sensitivity of cohort members during this formative life stage.” But to argue that “heightened sensitivity” produces distinctive responsiveness to political stimuli is a not-very-enlightening tautology.

Fortunately, our interest here is not in the specific psychological or sociological mechanisms that may underlie political socialization, but in the patterns of generational opinion change they produce. The literature seems to be a good deal stronger on that score, with careful empirical studies shedding light on the continuity of political attitudes over the life-span and on variations in political attitudes across cohorts or generations.⁹ Nevertheless, these studies fall short of providing a clear empirical test of what we take to be the core hypothesis of Mannheim's generational theory.

On one hand, long-term panel studies of single cohorts over significant portions of the life-span (Newcomb et al. 1967; Jennings and Niemi 1981; Beck and Jennings 1991) provide strong evidence of the persistent effects of early formative experiences; however, the nature of these studies precludes detailed analysis of the timing of socialization or detailed comparison of the experiences and attitudes of different cohorts. On the other hand, short-term panel studies (conducted over periods ranging from a few months to a few years) with respondents of different ages (for example, Carlsson and Karlsson 1970; Glenn 1980; Franklin and Jackson 1983; Alwin and Krosnick 1991) have demonstrated that political attitudes are more labile in early adulthood than at later points in the life-cycle; but that fact does not, in itself, provide evidence of generational imprinting in Mannheim's sense.

In our view, a real test of Mannheim's theory requires us to distinguish clearly between the *weight* associated with each age and the corresponding *incremental impact* of events experienced at that age. This distinction is perhaps clearest in the case of the simple "running

⁹ Delli Carpini (1989) provided a thorough and insightful review of the relevant literature.

tally” model. In that model, the *weight* attached to each period’s experience is, by assumption, constant; however, the *incremental impact* of events declines over the course of the life-cycle, simply because each new experience is integrated with a growing body of previous experience.¹⁰ So, for example, any given political shock will produce only half as much change in the views of a 40-year-old as in the views of a 20-year-old. Nevertheless, it would be a mistake to conclude that political events experienced at age 20 are somehow more powerful or important than those experienced at age 40. From the perspective of understanding the current political beliefs of a 40-year-old (or, for that matter, a 60-year-old or 80-year-old), the events experienced at age 20 would be exactly as important as those experienced at age 40—no more and no less. That does not seem to be what Mannheim and subsequent “generational” theorists had in mind.

A stronger, and correspondingly more interesting, version of the “generational imprinting” hypothesis would hold that the age-weights themselves, and not only the incremental impacts associated with those age-weights, peak during a crucial period of political development in adolescence or early adulthood. Cast in this light, the “generational imprinting” hypothesis fits within the same framework of period-specific shocks and age-specific weights that we have used to represent the Bayesian “running tally” and “partisan change” models. In particular, the political attitudes of an individual in birth cohort c at time t can be represented as a weighted average of previous political experiences,

$$\{7\} \quad \mu_{c,t} = (\sum_{i \leq t-c} \omega_i \pi_{c+i}) / (\sum_{i \leq t-c} \omega_i),$$

where π_{c+i} is the period-specific shock representing the distinctive political events of the period

¹⁰ The former fact is clear from equation {2}, the latter fact from equation {3}.

in which cohort c reached age i and ω_i is the *age-specific* weight representing the distinctive sensitivity to political experience of individuals at age i .

Equation {7} exactly parallels equation {5}, except that the notation has been slightly recast to emphasize the age-dependence of the weights $\{\omega_1, \dots, \omega_i, \dots, \omega_{t-c}\}$. Thus, in our formulation, the crucial distinction between the “generational imprinting” model and the Bayesian learning models is not in *how* political experiences are incorporated, but simply in the *extent* to which political experiences are incorporated at different points in the life-cycle.

The precise pattern of age-weights implied by the hypothesis of “generational imprinting” is somewhat unclear, but hardly so unclear as to render the hypothesis vacuous. In an extensive review of the relevant literature, Delli Carpini (1989, 20) conceded that “there is some disagreement on the specific years involved,” and cited various estimates (including 17 to 25, 18 to 26, 15 to 30, and 20 to 30) for the “critical” years of political development. An attractive feature of equation {7} as a representation of the “generational imprinting” hypothesis is that it is sufficiently flexible to capture any relevant variation in age-weights over the course of the life-cycle. Thus, evidence rather than assumptions can identify which (if any) stage of life is distinctively “formative”—a good thing, given that “almost any pattern can be defended” logically (Delli Carpini 1989, 31).

Another attractive feature of equation {7} is that the *political content* of generational imprinting varies naturally from cohort to cohort with changes in the political environment. As Delli Carpini (1989, 18) summarized Mannheim’s generational hypothesis,

The possibility of a new generation existed with the coming of age of each

new cohort, but the actuality of a new generation, as well as the specific attitudes that developed and the specific ways in which those attitudes were expressed, depended on the external environment.

In our mathematical formulation, the contribution of the external environment is to produce a sequence of political experiences which may or may not depart markedly from those experienced in previous periods. When they do, the new cohort or cohorts most open to being shaped by those distinctive experiences develop markedly different political attitudes from previous cohorts, transforming a "potential generation" into an "actual generation."¹¹

2. Data and Estimation

Our empirical analysis employs data on party identification from the American National Election Study (ANES) Cumulative Data File.¹² The ANES surveys conducted between 1952 and 2008 include 29,008 non-southern white respondents with identifiable birth years between 1862

¹¹ Our formulation is sufficiently flexible to represent either sharp or gradual generational differences. If a crucial political experience occurs suddenly, and if the weight attached to political experiences varies sharply with age, then a relatively narrow birth cohort may develop quite distinctive attitudes by comparison with earlier and later cohorts. On the other hand, a political experience that extends over several years will tend to produce less sharp generational differences, as will a pattern of relatively equal age-weights over a significant fraction of the life-span.

¹² The data are from the August 25, 2011 version of the ANES Time Series Cumulative Data File on the American National Election Studies website, www.electionstudies.org.

and 1990.¹³ We exclude African-Americans and white southerners from our analysis because their distinctive responses to political events in the period covered by our analysis make it implausible to assume that their partisan attachments were based on the same sequence of partisan shocks as those of non-southern whites.¹⁴ We measure party identification using the familiar ANES seven-point scale, recoded to range from zero (for strong Democrats) to +100 (for strong Republicans).

Because the ANES surveys have been conducted biennially (in the fall of each election year, with the exception of 2006), it seems natural to define birth cohorts and estimate period-specific political shocks in two-year increments. So, for example, we categorize respondents who were 18 or 19 at the time of the 2008 ANES survey (born in 1989 or 1990) as members of a 1990 birth cohort. These respondents' partisan loyalties in 2008 are assumed to be the product of a series of ten biennial political shocks: a 1989-90 shock experienced at age 0-1, a 1991-92 shock experienced at age 1-3, and so on through a 2007-08 shock experienced at age 17-19.

The resulting distribution of respondents by birth cohort and survey year is shown in Figure 1. The dark vertical bands in the figure represent ANES surveys with unusually large numbers of respondents, while the darker blocks within each vertical band indicate the best-represented birth cohorts in the corresponding survey. The surveys conducted in the 1950s drew primarily from cohorts born between the "critical" elections of 1896 and 1932; the surveys conducted in

¹³ Our analysis excludes respondents whose age was not ascertained, including all of the respondents in the 1954 ANES survey, in which age was only recorded in ten-year intervals.

¹⁴ Our definition of the South encompasses the eleven former Confederate states.

the 1970s drew primarily from cohorts born between about 1912 and 1958; the 2008 survey drew primarily from cohorts born after World War II. The oldest cohorts in each survey are sparsely represented; so too, in most cases, are the very youngest cohorts.

***** Figure 1 *****

2.1. Model Specification

The general form of our weighted-average model is

$$\{8\} \quad y_{c,t} \sim \text{Normal}(\mu_{c,t}, \sigma^2/n_{c,t}),$$

where $y_{c,t}$ is the observed average level of party identification for survey respondents in birth cohort c interviewed at time t , $n_{c,t}$ is the number of survey respondents in birth cohort c interviewed at time t ,

$$\{9\} \quad \mu_{c,t} = \boldsymbol{\omega}_{c,t}' \boldsymbol{\pi}_{c,t} / \boldsymbol{\omega}_{c,t}' \mathbf{1},$$

where $\boldsymbol{\omega}_{c,t}$ is a $(t-c)$ -vector of unknown weights and $\boldsymbol{\pi}_{c,t}$ is a $(t-c)$ -vector of unknown shocks, and σ^2 is the variance of the distribution of idiosyncratic factors affecting party identification. Up through time t , the cohort born in period c has been exposed to the sequence of shocks $\boldsymbol{\pi}_{c,t} = \{\pi_{c+1}, \dots, \pi_i, \dots, \pi_t\}'$ reflecting the relevant events of its political lifetime. This sequence of shocks is weighted by the vector $\boldsymbol{\omega}_{c,t} = \{\omega_1, \dots, \omega_i, \dots, \omega_{t-c}\}'$, where $(t-c)$ represents the age of cohort c at time t . The quantity $\mathbf{1}$ is a unit vector of length $(t-c)$, so the denominator in equation {9} is simply the sum of the relevant weights. Put differently, the vectors $\boldsymbol{\omega}_{c,t}$ and $\boldsymbol{\pi}_{c,t}$ are selecting the appropriate elements from the complete sequence $\boldsymbol{\omega} = \{\omega_1, \dots, \omega_{46}\}'$ of age-specific weights and the complete sequence $\boldsymbol{\pi} = \{\pi_{1862}, \dots, \pi_{2008}\}'$ of period-specific shocks, respectively, for the cohort

born in period c and interviewed at time t .¹⁵

Equations {2}, {5}, and {7}—the corresponding expressions for $\mu_{c,t}$ in the “running tally,” “partisan change,” and “generational imprinting” models—are simply special cases of this general model with less explicit notation.

It is clearly unrealistic to assume that the sequence of period-specific shocks in equation {9} is identical for all the members of a given cohort, even in our sub-population of non-southern whites. It would not be difficult (though it would be unwieldy) to further sub-divide that sub-population, positing distinct period-specific shocks for different sub-groups; but that complexity is beyond the scope of the present analysis. Instead, we simply rely on the variance term σ^2 in equation {8} to incorporate all of the factors that produce heterogeneity in partisan loyalties among the members of a given birth cohort at a point in time. Given the obvious importance of these factors, it should not be surprising that the generational model defined by equation {8} will leave a great deal of the total variance in observed partisanship unaccounted for.

2.2. Bayesian Inference

We adopt a Bayesian approach to inference for the unknown parameters in the model defined by equations {8} and {9}, $\Theta = (\boldsymbol{\pi}', \boldsymbol{\omega}', \sigma^2)$. For the period-specific shocks, we use independent uniform prior distributions spanning a large range relative to the zero-to-100 scale

¹⁵ We curtail the sequence of age-specific weights at age 90 and count the few respondents over that age in the ANES Cumulative Data File as 90-year-olds.

of our dependent variable: $\pi \sim \text{Uniform}(-200, 300)$.¹⁶ For the age-specific weights, we also use independent uniform prior distributions: $\omega \sim \text{Uniform}(0, 5)$. For σ , we use a uniform prior distribution broad enough to allow our generational model to account very badly (or very well) for observed party identification: $\sigma \sim \text{Uniform}(0, 100)$.

We use Markov chain Monte Carlo methods to explore the posterior density of Θ given the specified model and prior densities. We use JAGS (Plummer 2012) to deploy a MCMC sampling scheme specific to our model and priors. We direct JAGS to generate two, parallel MCMC schemes and save 250,000 MCMC iterates from each chain; we retain every 50-th iteration for inference. Formal run-length and stationarity diagnostics strongly suggest that the MCMC algorithm is sampling from the posterior density of the model parameters. All code is available upon request.

There is good reason to expect that the pattern of age-specific weights in equation {9} will be a more or less smooth function of age. If respondents are strongly affected by political events when they are 20, it seems reasonable to suppose that they are also strongly affected at 18 and 22, though perhaps not at 10 or 30. Thus, it is tempting to impose some structure on the age-specific weights, such as a random-walk specification or a hierarchical model, in order to reduce the complexity of the model and the number of parameters to be estimated. We resist

¹⁶ While the categorical party identification variable derived from ANES survey data is recoded to range from zero to 100, the estimated period-specific shocks for each two-year period are not constrained to fall within that range. We think of underlying partisanship as unbounded, so that period-specific shocks less than zero pull everyone (including strong Democrats) in a Democratic direction, while shocks greater than 100 pull everyone (including strong Republicans) in a Republican direction.

that temptation here, for two reasons. First, given the absence of previous research along these lines, it seems far from clear what specific functional form would be appropriate; even a seemingly unrestrictive choice could unintentionally miss or distort important features of the data. Moreover, the general expectation of continuity in the weights over the course of the life-span provides a useful check on the plausibility of the empirical results: if the estimated weights display a reasonably smooth pattern when nothing in the analysis forces them to do so, their credibility will be correspondingly enhanced.

While our approach to the estimation of age-weights is as unrestrictive as possible, one restriction is imposed by the structure of the model. Since all of the age-weights always appear in both the numerator and the denominator of equation {9}, it should be clear that their absolute scale is indeterminate—multiplying each of the weights by the same positive constant would have no effect on the magnitudes of the period shocks or on the expected value of the dependent variable. Thus, it is necessary to normalize the age-weights. We do so by constraining the 46 distinct weights $\{\omega_1, \dots, \omega_{46}\}$ reflecting the impact of events from birth to age 90 to sum to 46; thus, the average age-weight is 1, with values greater than 1 denoting larger-than-average weights and values less than 1 denoting smaller-than-average weights. Once this constraint is imposed, the prior distribution for any given ω has a 95% credible interval of [0.05, 2] and is approximately uniform over the interval from 0 to 1.7, albeit with a long right tail.

Once this normalization is imposed, Bayesian analysis of the model embodied in equations {8} and {9} is straightforward. The output of our statistical analysis consists of posterior distributions for the 121 parameters in our model—74 period-specific shocks $\{\pi_{1862}, \dots, \pi_{2008}\}$, 46 age-specific weights $\{\omega_1, \dots, \omega_{46}\}$, and the variance σ^2 . The posterior distributions of the period-

specific shocks are summarized in Figure 2, while the posterior distributions of the age-specific weights are summarized in Figure 3.

***** Figure 2 *****

One important limitation of the ANES data in estimating our generational model is immediately evident in Figure 2: those data provide little useful information regarding the period-specific shocks for early years in the sequence $\{\pi_{1862}, \dots, \pi_{2008}\}$. With no respondents interviewed before 1952—and most interviewed after 1976—it is obviously very difficult to disentangle the specific impact of political events in, say, 1882 from those in 1884.¹⁷ Thus, our estimates of the π parameters for the early years of the analysis are quite imprecise.¹⁸ For the later years the survey data are much more informative; for the 29 estimates reflecting political

¹⁷ These parameters are not completely unidentified because, under the assumptions of our generational model, survey respondents in the 1884 birth cohort (who were 68 or 69 years old in 1952, 72 or 73 years old in 1956, etc.) could have been affected by events in 1884, but not by events in 1882. However, there are only 114 non-southern white ANES respondents in this cohort; and even the evidence they can provide is clouded by the fact that political events “experienced” in their infancy may have a very modest impact on their partisan attachments seven decades later. The situation is rather less bleak for later years, not only because there are more ANES respondents in the relevant cohorts but also because, on average, those respondents are observed at younger ages (so that their observed partisan loyalties have to be decomposed into fewer distinct period-specific shocks).

¹⁸ The average width of the posterior 95% credible intervals for the first 23 shocks (1862-1906) is 455.8, covering 91% of the uniform prior interval. The average width of the 95% credible intervals for the next 22 shocks (1908-1950) is 378.4, covering 76% of the uniform prior interval. For the shocks after 1950, the average width of the 95% credible intervals amounts to 27% of the uniform prior interval.

events after 1950, the average width of the posterior 95% credible interval is 134.6—less than one-third the average width of the corresponding credible intervals for earlier years.

Given this degree of statistical imprecision, it would be rash to suppose that the estimated partisan shocks displayed in Figure 2 can shed detailed light on the partisan political history of the United States even over the past half-century, much less over the past 150 years. However, since the π parameters play a crucial role in our generational model, it seems reassuring that the estimated values plausibly track the political events that seem likely to have shaped the partisan attachments of non-southern whites over the period encompassed by the ANES surveys. For example, the estimates suggest that the largest pro-Republican shocks in that period were in 2002 (in the wake of the 9-11 terrorist attacks), 1984 (Ronald Reagan’s landslide reelection), 1954 (the first midterm election of the Eisenhower era), 1972 (Richard Nixon’s landslide reelection), and 1994 (the first time in decades that Republicans regained control of the House of Representatives). However, the frequency of shocks erasing partisan gains made in the preceding period (reflected in the saw-tooth pattern in Figure 2) suggests that the estimates may be capturing some short-term fluctuations in party identification as well as more durable shifts, and that our model might be improved by distinguishing the former from the latter.

3. Partisan Learning through the Life Cycle

Our key findings are presented in Figure 3, which summarizes the posterior distributions of the 46 estimated age-specific weights in our Bayesian model.

*** * * Figure 3 * * ***

The 95% credible intervals displayed in Figure 3 reflect a good deal of statistical uncertainty

about the magnitudes of the age-specific weights, especially after the age of 60 or so. This uncertainty reflects the inevitable fact that the ANES data include many more former 20-year-olds than former 60-year-olds, and very few former 80-year-olds.¹⁹ However, it is important to bear in mind that the credible intervals in Figure 3 are derived from a model in which no smoothing or other structure is imposed on the 46 distinct age-specific weights; each of them is estimated independently (aside from the normalization constraining them to sum to 46). Thus, runs of several consecutive high or low values are quite unlikely to arise by chance, even when the individual credible intervals overlap the average value of 1.0.

Within the limits set by the imprecision of the parameter estimates, four aspects of the results presented in Figure 3 seem noteworthy:

- There is clearly a sustained period of heightened sensitivity to political events during adolescence. This period occurs somewhat earlier than one might guess from the existing literature on generational imprinting: the peak period of sensitivity seems to be from age 7 to 17. (The average posterior estimate of ω for respondents between these ages is 1.42—that is, 42% greater than the lifetime average value.)

- The ANES respondents seem to have been less affected by political events that occurred when they were in their twenties. This period of diminished political sensitivity may reflect personal distractions of early adulthood, including new jobs, homes, and marriages. (The

¹⁹ To a lesser extent, the uncertainty of the estimated age-weights for older survey respondents also reflects the structure of the generational model, in which the current attitudes of 60-year-olds must be apportioned among 31 distinct period shocks, while the attitudes of current 20-year-olds are apportioned among only 11 distinct shocks.

average posterior estimate of ω for respondents between the ages of 19 and 29 is .76—almost 25% below the lifetime average value.)

- The twenties trough is followed by another period of heightened sensitivity during early middle age. (The average posterior estimate of ω for respondents between the ages of 33 and 43 is 1.39—roughly matching the average value for adolescents.)

- Respondents in their sixties and beyond are, once again, somewhat less responsiveness to political events. The estimates for each two-year age group beyond the age of 60 are even more imprecise than those for earlier age groups; however, the average posterior estimate of ω for the 15 oldest age groups is only .84.

These patterns provide remarkably little support for the notion that citizens rationally discount (or, for that matter, simply forget) past political events. The portion of the life span over which the age-weights are clearly increasing ends in early adolescence, and the subsequent (more gradual) increase through much of the twenties and thirties merely reverses a decline through the late teens and early twenties. Simply discounting past political experience by, say, 5% every two years would produce an age-weight for 60-year-olds more than three times as large as that for 10-year-olds, while discounting past experience by 10% every two years would produce an age-weight for 60-year-olds more than *ten* times as large as that for 10-year-olds. Clearly, nothing like this is going on in the ANES party identification data.

If the image of a citizenry acutely sensitive to changes in the partisan landscape fares poorly in this analysis, the hypothesis of generational imprinting receives considerable support, albeit with some important twists. The two sustained periods of heightened sensitivity to political

events evident in Figure 3 occur earlier and later in the life cycle than would be expected on the basis of the existing literature on political generations; the earlier period clearly ends in the mid-teens, while the second only begins in the early thirties. The early- and mid-twenties, which are often portrayed as a crucial phase in the political life-course, seem from our results to be distinctly less important in shaping the partisan loyalties of ANES survey respondents.

While we attach significance to the fluctuations in political sensitivity over the life-cycle evident in Figure 3, our findings also provide surprising support for the simpler “running tally” model, in which partisanship at any point in time represents a simple average of political experiences over an individual’s entire lifetime. Certainly, if our aim was to account for the current political views of a typical 50- or 60-year-old, we would not go too far astray by computing a simple average of the political shocks she had experienced over the preceding five or six decades of her life. One implication of this fact is that successive generations are unlikely to differ sharply in their partisan views, since their views are mostly constructed from the same political events incorporated in similar ways.

Lest a simple “running tally” of political experiences seem unremarkable, it is worth underlining the impressive staying power of political impressions implied by a simple average computed over a period of decades. That staying power is suggested graphically by the top line of Figure 4, which recasts our statistical results from Figure 3 by showing the estimated weight of partisan experience accumulated before the age of 18, expressed as a proportion of the cumulative weight of all the partisan shocks experienced at each point in the life-cycle. The striking implication is that adolescent political experiences play a substantial role in shaping partisan identities throughout the life-course. The party identification of a typical 38-year-old in

the ANES surveys seems to have been shaped as much by what happened before she was 18 as by what happened after she was 18. This implies, for example, that the partisanship of 38-year-olds in 2008 was shaped as much by the presidencies of Richard Nixon, Jimmy Carter, and Ronald Reagan as by those of George H. W. Bush, Bill Clinton, and George W. Bush.

*** * * Figure 4 * * ***

By way of comparison, the bottom line in Figure 4 shows the estimated weight of the most recent partisan shock at each point in the life cycle, likewise expressed as a proportion of the cumulative weight of all the partisan shocks experienced up to that point. (For 18-year-olds—the youngest age group shown in the figure—these two proportions necessarily sum to unity, since their partisan life-experience consists of shocks experienced before the age of 18 plus the most recent shock.) The relative weight of contemporary events in shaping party identification is always less than ten percent, and over most of the life cycle it is less than five percent.

4. Conclusion

The weight of the past implied by the estimates presented in Figure 4 seems hard to square with the notion that partisanship is a product of “rational” political learning. As Fiorina (1977, 611) wrote about his own version of the “running tally” model:

I would hesitate to call the model I have proposed a rational choice model.

Why in the world would a sixty-five year-old union member vote for

McGovern on the basis of what he thought of Roosevelt? Sunk costs are sunk,

our colleagues in economics say.

The subsequent work of Achen (1992) and Gerber and Green (1998) clarified precisely how, and to what extent, it would be rational for memories of Roosevelt to motivate a vote for McGovern—only to the extent that the continuity of Democratic (and Republican) personnel, policy commitments, and performance over the intervening decades produced a rational expectation that the partisan differential experienced in the 1930s would be diagnostic of the parties' relative performance in the 1970s.

The estimated age-specific weights presented in Figure 3 imply that a 65-year-old in 1972 would have given the political events of the Roosevelt administration (when he was in his late 20s and 30s) almost 30% *more* weight in the formulation of his current party identification than the events of the decade he had just lived through—including the civil rights movement, Lyndon Johnson's Great Society, race riots, assassinations, and the war in Vietnam, up through and including McGovern. By this calculation, sunk costs appear to be anything but sunk! Nor is this remarkably long partisan memory attributable to the particular features of the New Deal era or of the 1960s—given the structure of our model, the calculation would be exactly the same for a 65-year-old in 1932 looking back to the era of William Jennings Bryan, or for a 65-year-old Fiorina in 2012 looking back to McGovern, Watergate, and Jimmy Carter.

In our view, the weight of the past implied by these estimates is too heavy to be convincingly accounted for using the tidy Bayesian logic of "rational" political learning. While the American political landscape has displayed a considerable degree of stability over the period covered by our analysis—and we have intentionally side-stepped the most momentous ways in which it has changed by excluding African-Americans and white southerners from our calculations—it is simply implausible to suppose that decades-old partisan experiences are just

as relevant as contemporary experiences in forecasting what the parties will do in the future.

On the other hand, it would be rash to conclude from these results that citizens never forget political events, or that they consciously give events from the distant past as much or even more weight than very recent events in formulating their current partisanship. A more plausible interpretation of our results is along the lines suggested more than 50 years ago by the authors of *The American Voter*, who argued that party identification “raises a perceptual screen through which the individual tends to see what is favorable to his partisan orientation” (Campbell et al. 1960: 133). In this view, partisan shocks from the Carter administration have a strong effect on the party identification of 65-year-olds in 2012, not because those 65-year-olds consciously weigh all of their life experiences in formulating responses to the ANES party identification question, or because they have cumulated those experiences in a constantly-updated “running tally,” but because early experiences color their perceptions of subsequent political events. If partisan loyalties formed during the presidency of Jimmy Carter shape the way people see the economic record of the Reagan administration, deficit reduction under Bill Clinton, and missing weapons of mass destruction in Iraq, then long-ago events will seem to weigh as heavily as contemporary events in shaping current party identifications, even though their impact is largely indirect.²⁰

²⁰ Bartels (2002) used a Bayesian model of opinion change to examine partisan biases in assessments of presidential performance, candidates’ character traits, and “objective” economic conditions, including changes in the unemployment and inflation rates under Ronald Reagan. See also Achen and Bartels (2006) on perceptions of the budget deficit under Bill Clinton and Gaines et al. (2007) on partisan interpretations of the war in Iraq.

As Gerber and Green (1998) have noted, party identification may be a very atypical political attitude in this respect, since both the theory and the operationalization of party identification emphasize long-term psychological attachment rather than short-term political calculation (Campbell et al. 1960, chaps 6-7). Nothing in the results we have presented here is inconsistent with the notion that citizens update *most* political attitudes and beliefs—about candidates, policies, and other features of the political world—with much greater temporal discounting than is evident here. The only way to find out is to look. Thus, a high priority for future research is to apply the generational model proposed here to a variety of other political attitudes and beliefs.

Acknowledgements

This is a revised version of a paper originally presented at the 2001 meeting of the American Political Science Association. We are grateful to participants in the Oxford workshop on New Approaches in Age, Period, Cohort Analysis and (repeatedly) to Christopher Achen for helpful discussion and comments; to Princeton University's Woodrow Wilson School of Public and International Affairs, Stanford University's Department of Political Science, and Vanderbilt University's May Werthan Shayne Chair in Public Policy and Social Science for financial support; and to the American National Election Studies project for providing high-quality survey data spanning six decades.

References

- Abramson, P.R., 1975. *Generational Change in American Politics*. D.C. Heath, Lexington, MA.
- Achen, C.H., 1992. Social Psychology, Demographic Variables, and Linear Regression: Breaking the Iron Triangle in Voting Research. *Political Behavior* 14 (3), 195-211.
- Achen, C.H., Bartels, L.M., 2006. *It Feels Like We're Thinking: The Rationalizing Voter and Electoral Democracy*. Prepared for presentation at the Annual Meeting of the American Political Science Association, Philadelphia.
- Alwin, D.F., Krosnick, J.A., 1991. Aging, Cohorts, and the Stability of Sociopolitical Orientations Over the Life Span. *American Journal of Sociology* 97 (1), 169-195.
- Bartels, L.M., 2002. *Beyond the Running Tally: Partisan Bias in Political Perceptions*. Political

Behavior 24 (2), 117-150.

Beck, P.A., Jennings, M.K., 1991. Family Traditions, Political Periods, and the Development of Partisan Orientations. *Journal of Politics* 53 (3), 742-763.

Brady, H.E., Elms, L., 1999. Age-Period-Cohort Analysis with Noisy, Lumpy Data. Presented at the annual meeting of the Political Methodology Group, College Station, TX.

Campbell, A., Converse, P.E., Miller, W.E., Stokes, D.E., 1960. *The American Voter*. John Wiley & Sons, New York.

Carlsson, G., Karlsson, K., 1970. Age, Cohorts, and the Generation of Generations. *American Sociological Review* 35 (4), 710-718.

Converse, P.E., 1969. Of Time and Partisan Stability. *Comparative Political Studies* 2 (2), 139-171.

Converse, P.E., 1976. *The Dynamics of Party Support: Cohort-Analyzing Party Identification*. Sage Publications, Beverly Hills, CA.

Delli Carpini, M.X., 1989. Age and History: Generations and Sociopolitical Change. In: Sigel, R.S. (Ed.), *Political Learning in Adulthood*. University of Chicago Press, Chicago, pp. 11-55.

Fiorina, M.P., 1977. An Outline for a Model of Party Choice. *American Journal of Political Science* 21 (3), 601-625.

Franklin, C.H., Jackson, J.E., 1983. The Dynamics of Party Identification. *American Political Science Review* 77 (4), 957-973.

Gaines, B.J., Kuklinski, J.H., Quirk, P.J., Petyon, B., Verkuilen, J., 2007. Same Facts, Different Interpretations: Partisan Motivation and Opinion on Iraq. *Journal of Politics* 69 (4), 957-974.

- Gerber, A., Green, D.P., 1998. Rational Learning and Partisan Attitudes. *American Journal of Political Science* 42 (3), 794-818.
- Glenn, N.D., 1980. Values, Attitudes, and Beliefs. In: Brim, O.G. Jr., Kagan, J. (Eds.), *Constancy and Change in Human Development*. Harvard University Press, Cambridge, MA, pp. 596-640.
- Jennings, M.K., Niemi, R.G., 1981. *Generations and Politics: A Panel Study of Young Adults and Their Parents*. Princeton University Press, Princeton, NJ.
- Mannheim, K., 1952. The Problem of Generations. in Kecskemeti, P. (Ed.), *Essays on the Sociology of Knowledge*. Routledge and Kegan Paul, London, pp. 276-320
- Markus, G.B., 1983. Dynamic Modeling of Cohort Change: The Case of Political Partisanship. *American Journal of Political Science* 27 (4), 717-739.
- Miller, W.E., Shanks, J.M., 1996. *The New American Voter*. Harvard University Press, Cambridge, MA.
- Newcomb, T.M., Koenig, K.E., Flacks, R., Warwick, D.P., 1967. *Persistence and Change: Bennington College and Its Students After Twenty-Five Years*. John Wiley & Sons, New York.
- Plummer, Martyn. 2012. JAGS users manual, Version 3.2.0.
<http://sourceforge.net/projects/mcmc-jags>
- West, M., Harrison, J., 1997. *Bayesian Forecasting and Dynamic Models*, 2nd edition. Springer, New York.
- Zeckman, M., 1979. Dynamic Models of the Voter's Decision Calculus. *Public Choice* 34 (3-4), 297-315.

Figure 1: Sample Size by Birth Cohort and Survey Year

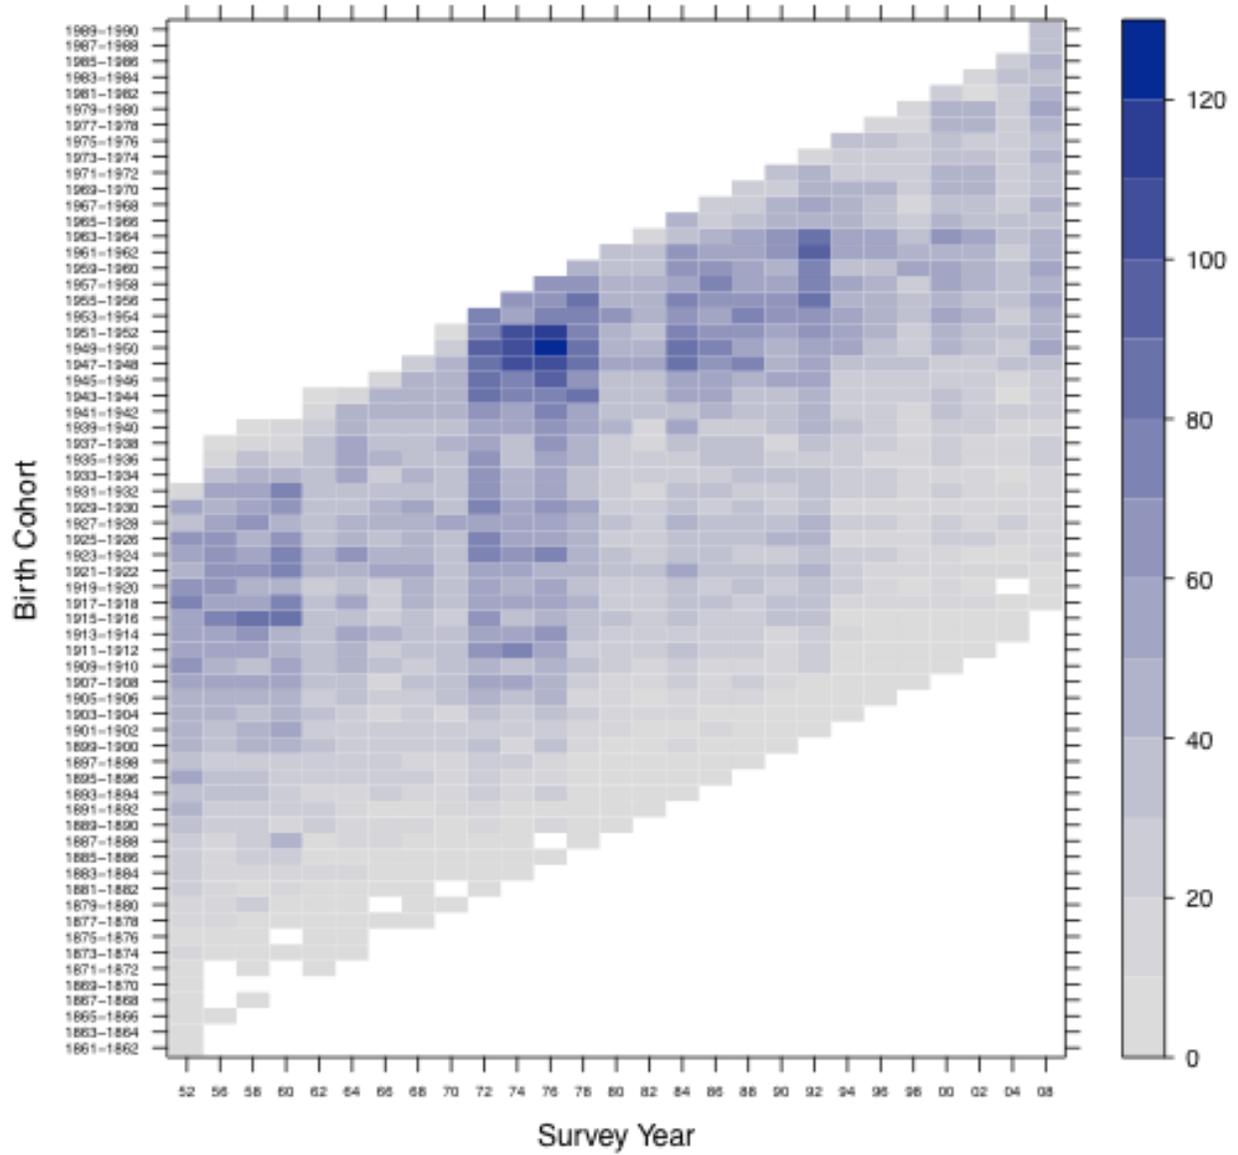


Figure 2: Estimated Partisan Shocks, 1864-2008

Bayesian MCMC parameter estimates with 95% credible intervals
(0=Strong Democratic; 50=Neutral; 100=Strong Republican)

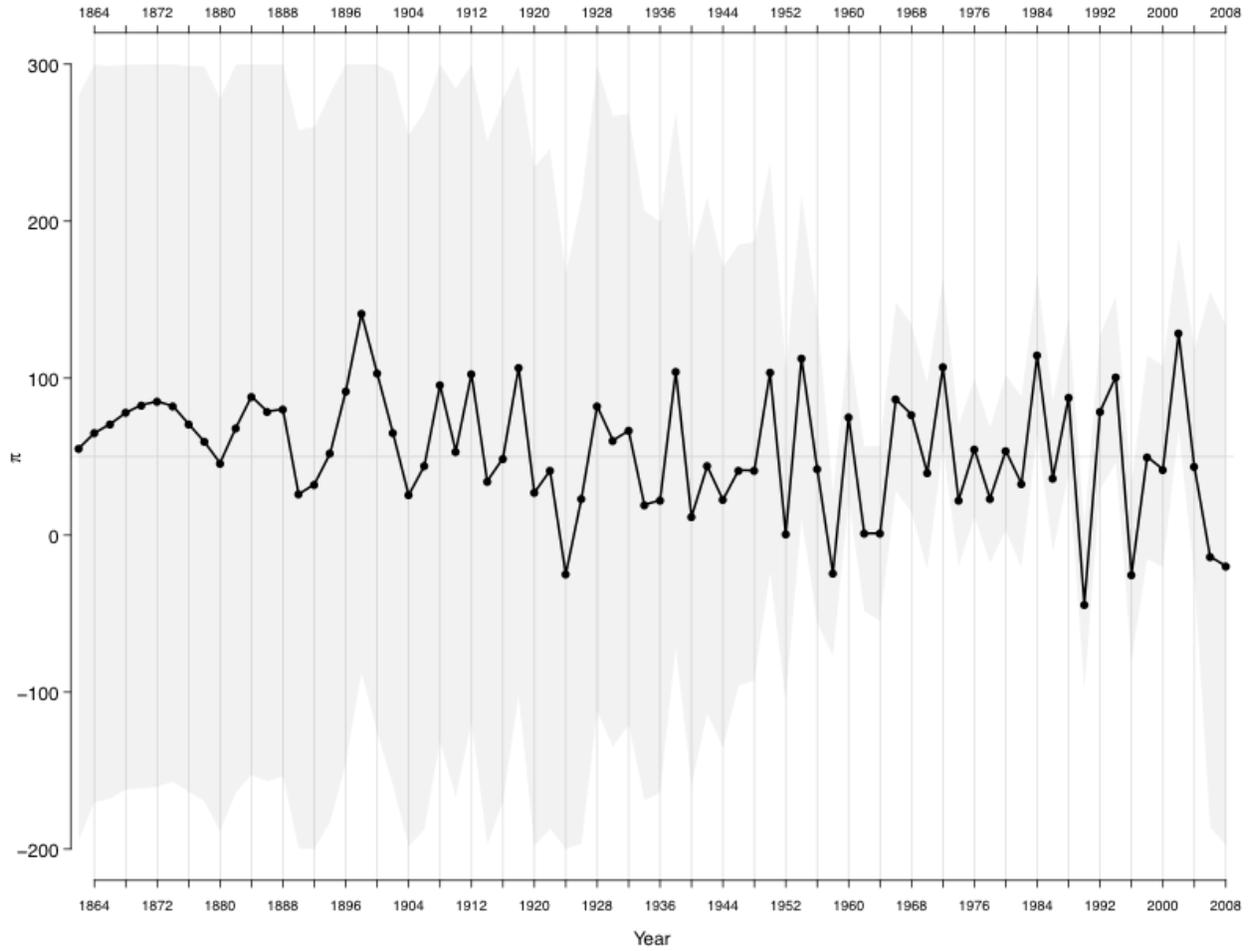


Figure 3: Estimated Age-Specific Weights

Bayesian MCMC parameter estimates with 95% credible intervals

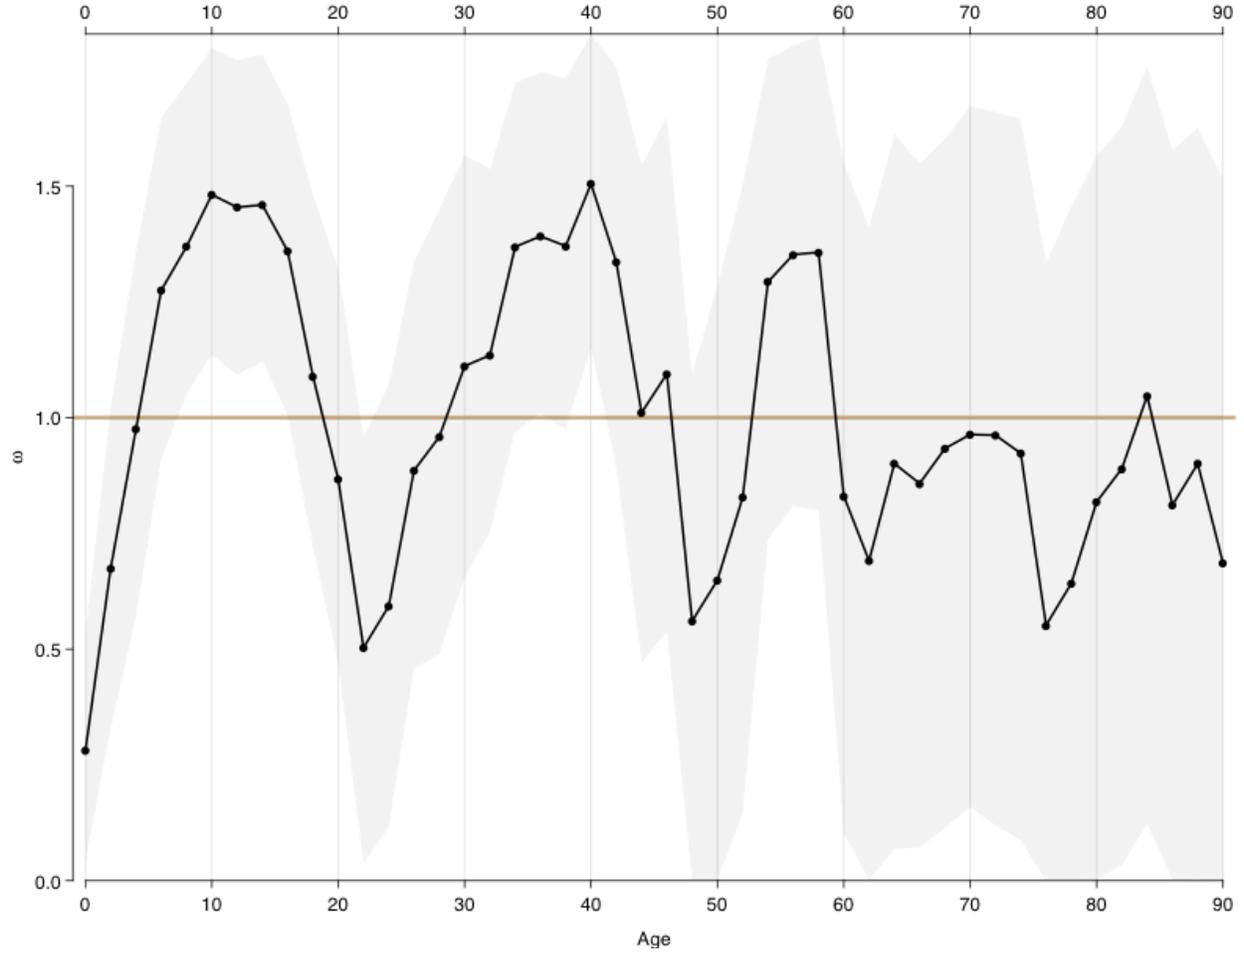


Figure 4: The Persistent Impact of Early Partisan Experiences

Bayesian MCMC parameter estimates with 95% credible intervals

