Age, Period, and Cohort Analysis with Bounding and Interactions*

Jiwon Lee†

Published in *Sociological Methods & Research*

Abstract

This article uses the example of voter turnout in US presidential elections to compare two new methods for age, period, and cohort (APC) analysis: the APC interaction model and the APC bounding analysis. While discussing the formal, conceptual, and interpretive differences between the two methods, the analysis demonstrates how both methods can be used to generate distinct but complementary findings. Because the two methods take alternative positions on the appropriate cohort-effect estimands, the comparison underscores the importance of well-grounded conceptual foundations in APC analysis.

*A previous version of this manuscript, titled "Age, Period, and Cohort Effects on Voter Turnout for the 1976–2020 Presidential Elections," was presented at the 2022 Annual Meetings of the Population Association of America. I thank the discussant, Andrew Fenelon, and other attendees for their thoughtful comments. Additionally, I would like to thank Julia Burdick-Will, Andrew Cherlin, Robert Moffitt, Steve Morgan, the editor, five anonymous reviewers, and participants of the PSI Seminar at Johns Hopkins University Department of Sociology in November 2021, and the Korea Inequality Research Network Colloquium in September 2023 for their valuable feedback and suggestions on earlier drafts of this article. Any errors that remain are my own.

†Postdoctoral Fellow, Cornell Population Center, Cornell University, NY, USA. Email: jiwon.lee@cornell.edu.

1. Introduction

A record-setting 160 million Americans, consisting of 66.8 percent of the voting-eligible population, cast ballots in the 2020 general election (McDonald 2020). The voting surge was welcomed by commentators worried about declining turnout rates and vanishing voters (e.g., Rosenstone and Hansen 1993; Teixeira 1992; Miller and Shanks 1996; Patterson 2002; cf. McDonald and Popkin 2001), and several explanations have been proposed.[1] Some focus on factors confined to the election, such as the resurgence of interest in politics fueled by the presidency of Donald Trump and the broadening of absentee voting to mitigate the impact of the COVID-19 pandemic. Others suggest cohort-replacement explanations, underscoring the unusually high turnout among young voters who came of age during a period of growing economic insecurity.[2]

At least since the 1960s, with the development of age-period-cohort (APC) models, demographers and social scientists have investigated analogous explanations for voting behavior (e.g., Alwin 1998; Dassonneville 2017; Frenk, Yang, and Land 2013; Glenn and Grimes 1968; Hout and Knoke 1975; Land, Hough, and McMillen 1986; Smets and Neundorf 2014). However, determining how APC processes structure an outcome such as turnout is challenging. One must contend with the identification problem that arises from the linear dependencies among three predictors on the same time scale (age = period - birth year). The problem has prompted some scholars to use methods that require that unverifiable restrictions be placed on the data.[3]

This article has two main purposes. First, more narrowly, it revisits the prior findings on how APC processes have shaped voter turnout trends in the US and extends them to cover the three most recent presidential elections. To this end, two new approaches to APC analysis — namely, APC bounding analysis (Fosse and Winship 2019b) and the APC-interaction (APC-I) model (Luo and Hodges 2020a) — are applied to the 1976-2020 Current Population Survey Voting and Registration Supplements (CPS-VRSs).

Second, more broadly, it offers a methodological comparison of two promising approaches. These two methods are grounded in distinct approaches to APC analysis that have long coexisted within the literature. Nevertheless, direct comparisons and contrasts of these approaches, using the same data for a specific analytical question, have seldom been conducted (but see Keyes, Utz, Robinson, and Li 2010). The examination of voter turnout offers an ideal context for demonstrating and contrasting the new methods, given the extensive existing knowledge on the subject and the large sample sizes available from the CPS-VRSs. Through the motivating analysis, I aim to elucidate the formal and conceptual differences that underlie the two methods. By connecting the methods

[1] See Figure S1 in the Online Supplement for the national turnout rates from 1976 to 2020.

[2] See, for some journalistic examples, McDonald (2020), Corasaniti and Rutenberg (2020), and Hess (2020).

[3] The latest wave of the literature demonstrates that identification claims of the methods employed in prior analyses are overstated, given that they entail arbitrary constraints that are often difficult to justify (for review, see Bell 2019; Fosse and Winship 2019a). For example, the Hierarchical-APC (H-APC) model of Yang and Land (2006) has been used in recent studies of turnout trends (Dassonneville 2017; Frenk et al. 2013; Smets and Neundorf 2014). The H-APC model entails strong assumptions that are insufficiently grounded (see Bell and Jones 2014; Fosse and Winship 2019a; Luo and Hodges 2020b). As a result, the conclusions of the studies may be incorrect.

with the more abstract frameworks they each represent, I aim to assist researchers in choosing the most fitting method for their objectives.

2. How APC Processes Shape Turnout Trends in the US

In this section, I first explain how APC processes can shape turnout trends. The preexisting literature on age differentiation in prior general elections is extensive. A robust finding in this literature is that age has a curvilinear relationship with turnout; turnout climbs incrementally with age, plateaus, and then falls with aging among older individuals. This age curve is observed almost universally in Western democracies, including the US (for review, see Dassonneville 2017; Smets 2021).[4]

Period effects are present when election-specific events and circumstances shift aggregate voting patterns. Typical sources of period effects include the competitiveness of the election, campaign efficiency in voter mobilization, and changes in registration and voting laws (for recent reviews, see Cancela and Geys 2016; Smets and Van Ham 2013; Stockemer 2017). Among them, three factors have been particularly relevant to the 2020 election. First, the election was highly partisan and competitive (Drutman 2021; Pew Research Center 2020). Second, one candidate was especially provocative (McDonald 2020). Third, the expansion of absentee voting made voting much easier (Scherer 2021; cf. Yoder, Handan-Nader, Myers, Nowacki, Thompson, Wu, Yorgason, and Hall 2021).[5]

Cohort differentiation or cohort replacement can also drive turnout trends (Ryder 1965). Formal education increases essential cognitive, social, and civic resources for political engagement (Cancela and Geys 2016). In the US, formal education saw substantial growth over the twentieth century.[6] Thus, the substantial role of education in predicting political engagement and its broad increase across most cohorts suggests an overall increase in voter turnout due to cohort replacement.

In the political socialization literature, it is also argued that life-long political dispositions are disproportionately shaped early in the life course (e.g., Alwin, Cohen, and Newcomb 1991; Bartels and Jackman 2014; Jennings and Niemi 2014; Smets 2021). Because the sociopolitical contexts of the formative years are unique for each cohort but common within a cohort, cohorts may differ in political orientations. Aggregate-level change in turnout occurs when cohorts move through time and cohort compositions change in the electorate.

The most prominent example of relevant cohort differentiation is the comparison of those

[4]Specifically, in the US, prior studies generally suggest two critical points; one in the fifties and early-sixties when turnout peaks and another around seventy when turnout rates start to decline with age (e.g., Campbell, Converse, Miller, and Stokes 1980; Dassonneville 2017; Glenn and Grimes 1968; Hout and Knoke 1975; Land et al. 1986; Smets 2021; Wolfinger, Rosenstone, and Rosenstone 1980).

[5]As supporting evidence, see Pew Research Center (2020) and Scherer (2021) which together shows that significantly more registered voters in 2020 1) perceived the two major-party candidates to differ on various issues, 2) thought the result of the election "really matters" to the nation, 3) cited opposition to the competing candidate (e.g., "he is not Trump") as a reason for deciding support a candidate, and 4) voted early or by mail compared to prior elections.

[6]This educational growth was particularly pronounced in the earlier decades of the century, with bachelor's degree attainment by age twenty-five increasing from 6-7 percent in the 1915 cohort to 28 percent in the 1975 cohort, though this trend plateaued for those born post-1980 (Bailey and Dynarski 2011).

who did (or did not) pass through the Great Depression and the New Deal during their formative years. It is often claimed that the experiences of economic distress, government interventionism, and the sudden reshaping of the party system distinguish these cohorts' socio-political orientations from others (e.g., Abramson 1974; Alwin et al. 1991; Miller and Shanks 1996).[7]

3. Data and Methods of Analysis

I draw data from the 1976-2020 CPS-VRS, disseminated by the Integrated Public Use Microdata Series (Flood, King, Ruggles, and Warren 2021). CPS-VRSs are fielded bi-annually when congressional and presidential elections are held. Because this article focuses on presidential elections, I only use data from the years in 1976 to 2020 when a presidential election was held. I also group age and period variables into 4-year bins to align the time scales, as is the common practice in APC modeling. Additional detail on the data source is provided in Section S1 in the Online Supplement.

3.1. APC Accounting Model Framework

Let y denote the outcome set to 1 if a respondent voted and 0 otherwise. The age-period data array for the CPS-VRS data is indexed $a = 1, \ldots, 15$ and $p = 1, \ldots, 12$ for the age and period groups. Cohorts are then indexed from 1 to 26, yielding 26 groups born from 1899-1902 to 1999-2002. The classical APC accounting model is

$$y = \mu + \alpha_a + \beta_p + \gamma_c + \epsilon \tag{1}$$

where α_a, β_p, and γ_c denote, in this application, the 15 age effect parameters, 12 period effect parameters, and 26 cohort effect parameters, respectively (Mason, Mason, Winsborough, and Poole 1973).[8] The intercept is μ, and ϵ represents all other variation.

Following Holford (1983), it is useful to express the same model in a linearized form by specifying linear components of the effect parameters in Equation (1):

$$y = \mu + \underbrace{\alpha_L A_a^L + \beta_L P_p^L + \gamma_L C_c^L}_{\text{Linear components}} + \underbrace{\tilde{\alpha}_a + \tilde{\beta}_p + \tilde{\gamma}_c}_{\text{Nonlinear components}} + \epsilon \tag{2}$$

[7] For example, Putnam (2000) refers to individuals born between 1910 and 1940 as the "long civic generation." Putnam contrasts these cohorts' impressive levels of civic engagement to the successive cohorts, which he interprets as the impact of experiencing the Depression and World War II when coming of age. He further claims that the decline of civic activities over the late twentieth century is due primarily to the replacement of this civic generation with the "laid-back" Boomers. Similar claims on electoral participation are made in Miller (1992) and Lyons and Alexander (2000).

[8] The seminal work of Mason et al. (1973) uses the word effect to describe the age, period, and cohort parameters in Equation (2) (see their Equation (2)). Fosse and Winship (2019b) do the same for their parameters (see their Equation (1)). I follow this convention in the article but do not intend to imply causal effects based on the notions of counterfactuals (see Morgan and Winship 2015). See Section S3 in the Online Supplement for a brief discussion of the challenges in ascribing causal interpretations to APC estimands.

The slope parameter α_L is a fixed coefficient, and A_a^L is indexed from 1 to 15 for the age groups. The effect parameter denoted by $\tilde{\alpha}_a$ is then the departure from the linear trend determined by $\alpha_L A_a^L$. The period and cohort effect parameters are defined analogously.

The linear dependency of the APC variation renders Equation (2) equivalent to

$$y = \mu + \alpha_L A_a^L \tilde{\alpha}_a + \beta_L P_p^L + \gamma_L(P_p^L - A_a^L) + \tilde{\beta}_p + \tilde{\gamma}_c + \epsilon \tag{3a}$$

$$= \mu + \Big[(\alpha_L - \gamma_L)A_a^L + \tilde{\alpha}_a\Big] + \Big[(\beta_L + \gamma_L)P_p^L + \tilde{\beta}_p\Big] + \tilde{\gamma}_c + \epsilon \tag{3b}$$

The APC identification challenge is that the data alone cannot determine the linear slopes $(\alpha_L, \beta_L, \gamma_L)$, but only lower-dimensional differences and summations, such as $(\alpha_L - \gamma_L)$ and $(\beta_L + \gamma_L)$ in Equation (3b) (see Rodgers 1982). All the nonlinear components $(\tilde{\alpha}_a, \tilde{\beta}_p, \tilde{\gamma}_c)$ are point identified, but only as deviations from the linear components $(\alpha_L A_a^L, \beta_L P_p^L, \gamma_L C_c^L)$ that remain unidentified. As a result, the nonlinear components are identified only in a conditional sense.

3.2. APC-I Model

The APC-I Model of Luo and Hodges (2020a) is

$$y = \mu^* + \alpha_a^* + \beta_p^* + \eta_{ap(c)}^* + \epsilon^* \tag{4}$$

where $\eta_{ap(c)}^*$ denotes cross-product interactions between the a^{th} age group and p^{th} period for the cohort group c.[9] Notably, the stand-alone cohort parameters γ_c from the accounting model in Equation (1) are substituted by these age-period interaction effect parameters $\eta_{ap(c)}^*$. As I discuss in more detail later, this subtle change reflects the key conceptual shift from the accounting model. Following the reasoning of Ryder (1965), cohort effects are themselves assumed to be generated by period effects that vary with age, or vice versa.

Interpreted in terms of an age-period data array, the APC-I model corresponds to a form of the two-factor model with interactions (Fienberg and Mason 1979). The effect parameters α_a^* and β_p^* are marginal age and period effects, respectively. The interaction terms $\eta_{ap(c)}^*$ account for the cell-specific deviations from the values implied by the two marginal effect parameters.[10]

Unlike the accounting model, the APC-I model allows all parameters to be estimated directly from the data, which may then be interpreted as APC effect parameters when appropriately structured. Luo and Hodges suggest taking the arithmetic mean of the age-period interaction terms corresponding to a specific cohort as an overall summary of a cohort effect. This measure represents

[9]For different approaches for modeling cohort effects using age-period interactions, see James and Segal (1982); Keyes and Li (2010); Ohtaki, Kim, and Munaka (1990).

[10]It is important to note that this interpretation of interaction effect parameters depends on the use of effect coding (i.e., sum-to-zero coding), as recommended by Luo and Hodges. While the essence of the interactions — cell-level deviations from marginal effect parameters — remains consistent across several alternative coding schemes, effect coding facilitates the most straightforward interpretation of APC-I model estimates.

the average cohort deviation of the outcome from the level jointly predicted by the marginal age and period effect parameters.

Additionally, building upon Ryder (1965), the APC-I model offers a structured approach for modeling intra-cohort developments throughout the life cycle. This captures systematic over-time outcome deviations from marginal age and period effects within a cohort. Luo and Hodges suggest using orthogonal polynomial contrasts of an appropriate order to analyze the age-period interaction terms for a specific cohort, thereby mapping the trajectories of $\eta^*_{ap(c)}$ across different ages or periods within that cohort. This approach marks a theoretical advancement by allowing the cohort effects to evolve dynamically, thereby helping to align the cohort estimand more closely with its initial formulation. I expand upon these points in subsequent sections, and an illustrative representation of this refined feature is provided in Figure S3 of the Online Supplement.

3.3. APC Bounding Analysis

Embracing the partial identification approach of Manski (1990), Fosse and Winship (2019b) have developed a bounding approach for APC analysis based on the accounting model. Their main idea is to use social, cultural, and/or biological theories to restrict the parameter space of the unknown linear slopes to credible identification intervals.

To illustrate the ideas, first recall that particular combinations of the accounting model's linear slopes are estimable, such as single values for $(\alpha_L - \gamma_L)$ and $(\beta_L + \gamma_L)$ in Equation (3b). These estimable values provide information about the individual slopes embedded in them. For example, assume that the cohort slope is contained in an interval $[\gamma_L^{min}, \gamma_L^{max}]$. Then, the true age slope is also restricted to the interval $[(\alpha_L - \gamma_L) + \gamma_L^{min}, (\alpha_L - \gamma_L) + \gamma_L^{max}]$ because $\gamma_L^{min} \leq \gamma_L \leq \gamma_L^{max}$.[11] When multiple assumptions are introduced, the intersection of the implied intervals can narrow the range of credible values.

Fosse and Winship suggest two strategies for introducing assumptions. First, a researcher may refer to scientific theories stipulating that the outcome changes positively or negatively in relation to a linear APC component. The sign of the linear slope is then restricted (i.e., zero is imposed as an upper or a lower bound of one or more slopes). Secondly, one may assert an assumption of a positive or negative incremental change in the outcome over age, period, or cohort (sub-)intervals. These assumptions can also yield bounds on the linear slopes. I will use both strategies below.

4. Findings

4.1. Visualization and Descriptive Results

Figure 1 presents period-by-age turnout rates from 1976 to 2020 for respondents under the age of 78 (see Section S1 in the Online Supplement for an explanation of this age cutoff). In almost all periods, turnout rates increase monotonically in age until the age category of 58-61. After some leveling-off during the sixties, turnout rates mostly decline from ages 70-73 to 74-77. The drop-off in turnout during the seventies is more pronounced in the periods before 2000.

Age gaps in turnout rates tend to vary across periods, suggesting period-specific age effects. While turnout rates of young adults are lower in most elections, their turnout increases in years when the Democrats have won back the White House (1976, 1992, 2008, and 2020). The rates peaked for all age groups in 2020, increasing by four to eleven percentage points from 2016 to 2020. This widespread increase suggests a strong period effect associated with the 2020 election.

In Figure 2, I present cohort-by-period turnout trends. In all elections, turnout rates trace an inverted U, suggesting that inter-cohort variations could be relatively constant across periods. Nevertheless, I also find that the inflection points shift across elections, suggesting that age variations are confounding these cohort patterns. The figure also shows that turnout rates in the most recent elections, especially the 2020 election, are high relative to the prior elections for each cohort group.

Finally, Figure 3 displays cohort-by-age turnout rates. Turnout rates for each age group generally increase for successive cohort groups. These positive associations between cohort and turnout

[11]The complete sets of bounding formulas are presented in Fosse and Winship (2019b), which also offers more comprehensive expositions of the general method.

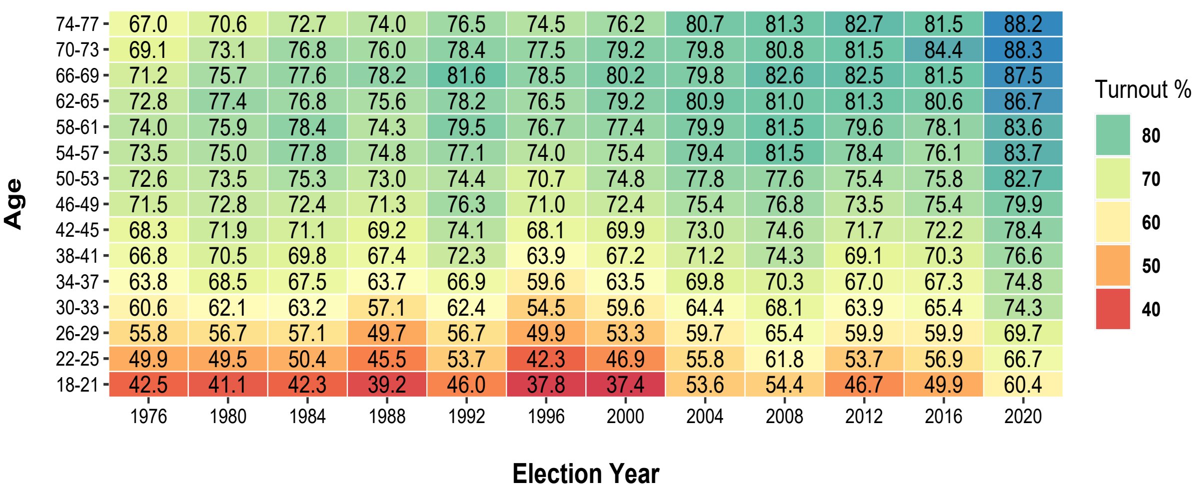

Note: Survey weights augmented by adjustments for non-response (N=981,858).

Figure 1. Turnout Rates in Presidential Elections by Election Year and Age Groups

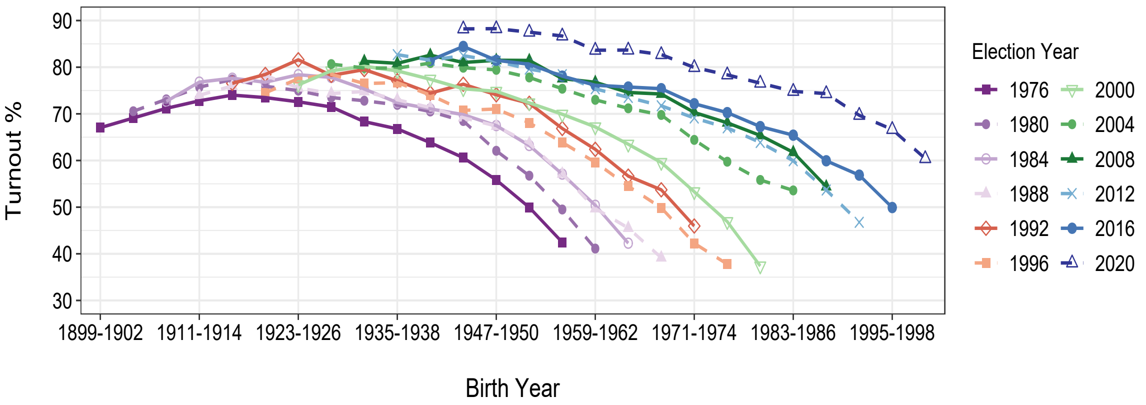

Note: Survey weights augmented by adjustments for non-response (N=981,858).

Figure 2. Turnout Rates in Presidential Elections by Cohort Groups and Election Year

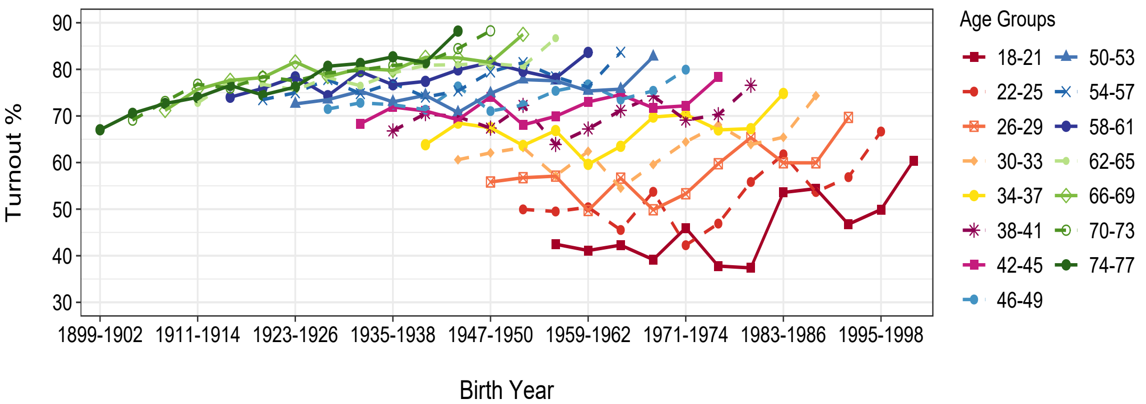

Note: Survey weights augmented by adjustments for non-response (N=981,858).

Figure 3. Turnout Rates in Presidential Elections by Cohort and Age Groups

are more pronounced among age groups below 50. Moreover, when stratified by age groups, turnout rates of cohorts after 1982 are higher than most of the prior cohorts. These patterns may imply significant modification of period effects by age. However, it is also plausible that these observed patterns are confounded by positive period effects for recent elections, as depicted in Figure 1.

Overall, the patterns in Figures 1, 2, and 3 suggest curvilinear age patterns and increasing turnout rates by period, which peak in 2020. Moreover, turnout rates are higher among voters born after 1980, conditional on age at the time of the election. However, the APC variables are confounded in the descriptive analysis. To develop evidence for stronger conclusions, I use the bounding analysis and APC-I model in the next two sections.

4.2. Bounding Analysis

The bounding analysis follows these steps. First, assert assumptions that can be defended as credible. Second, exploit the known results that (a) particular combinations of the component linear slopes are identified and estimable and (b) nonlinear components of the total APC effect parameters are always (conditionally) identified. Then, estimate these components and use the estimates with the assumptions to calculate implied bounds for the unidentified linear components of the total APC effect parameters.

Partial Identification Assumptions. I use the following assumptions to investigate the pattern of age, period, and cohort effects:

A1. Age effects increase monotonically from the age groups 18-21 to 58-61.

A2. Age effects decline monotonically from the age groups 70-73 to 74-77.

C1. The linear cohort slope is non-negative.

C2. The (weighted) average cohort effect of individuals born between 1911 and 1926 is equal to or greater than those born between 1947 and 1966.[12]

These assumptions are substantively "net" of each other. A1 and A2 are supported by the lifecourse aging literature (see Smets 2021). C1 is supported by the resource-based political engagement literature (Brady, Verba, and Schlozman 1995; Verba, Schlozman, Brady, and Nie 1993), which implies that turnout increases with increases in education. Finally, C2 is supported by the literature on the differences between the Great Depression generation and the Baby Boomer generation (e.g., Alwin et al. 1991; Putnam 2000).

Estimates of Bounds. Table 1 presents the estimates of the combinations of linear slopes from linear probability models with the coefficients scaled as percentage points.[13] These estimates are challenging to interpret on their own. The linear age and period slope sum ($\alpha_L + \beta_L$) is estimated to be 0.77. This estimate could be interpreted, for example, as an increase in the linear age component induced by a year increase in age, conditional on the linear period slope being zero.

Figure 4 plots the accounting model's nonlinear components — $(\tilde{\alpha}_a, \tilde{\beta}_p, \tilde{\gamma}_c)$ in Equation (2) — of the total APC effect parameters (added to an estimated intercept $\hat{\mu} = 68.1$; see Table S2 in the

[12]Given that the Depression and New Deal era began in 1929 and persisted through the 1930s, the assumption is that critical formative years are located somewhere between ages 3 to 29, which is consistent with the findings in the literature (see Smets 2021: 288).

[13]I constructed the design matrix using weighted orthogonal polynomial contrasts (Elbers 2020), which is analogous to the standard polynomial contrast coding but adjusts for imbalances in the number of observations between groups. The design matrix is also normalized such that the linear slopes are interpretable as the outcome variation induced by a year increase in age, period, or cohort. I reach substantively similar conclusions using logistic regression models (see Figure S2 in the Online Supplement).

Sums and Differences of Linear Slopes	Estimate	(Standard Error)
Age + Period Slope ($\alpha_L + \beta_L$)	0.77	(0.005)
Period + Cohort Slope ($\beta_L + \gamma_L$)	0.19	(0.004)
Age - Cohort Slope ($\alpha_L - \gamma_L$)	0.59	(0.003)

Table 1. Sums and Differences of Linear Slopes

Online Supplement for the full estimates). Interpretation of the estimates in Figure 4 is difficult without information on the unknown linear slopes. Nevertheless, when combined with the identification assumptions, they inform the parameter space of the linear slopes.

Assumption A1 requires that the total age effect ($\alpha_L A_a^L + \tilde{\alpha_a}$) is nonincreasing from the age categories 18-21 to 58-61. For the data to be consistent with Assumption A1, the linear age slope must be sufficiently large to compensate for the most significant decline in the nonlinear component within the asserted interval (i.e., $\alpha_L \geq \min\{\tilde{\alpha}_{a+1} - \tilde{\alpha}_a\}_{a=1}^{10}/4$). The estimated nonlinear components for these age groups are 59.5, 64.5, 67.4, 70.3, 72.0, 72.8, 72.3, 71.9, 70.5, 69.9, and 68.3. Therefore, their (forward) differences are 5.0, 2.9, 2.9, 1.6, 0.9, -0.6, -0.4, -1.4, -0.6, and -1.6. The most substantial decline in the nonlinear age component is 1.6 percentage points. This decline occurs between age groups 54-57 and 58-61 and is represented by the blue shaded area in the left panel of Figure 4. After normalization, this result suggests that the lower bound of the linear age slope α_L is 0.4 (i.e., 1.6/4 = 0.4) percentage points.[14]

Assumption A2 stipulates a decline in the age effect parameters among the elderly: the age effect parameters for the age group 74–77 are no higher than those for the age group 70–73. This assumption is represented by the red shaded area on the left panel of Figure 4. The nonlinear age component declines between the two groups by 3.3 percentage points. Therefore, the linear age slope is bounded above by 0.82 ($\approx$ 3.3/4) by the assumption because a linear age slope larger than 0.82 would offset the decline in nonlinear components among older individuals, violating the assumption.

The nonlinear cohort components are plotted in the right panel. Assumption C1 asserts that the linear cohort slope is not negative, rendering zero as a lower bound of the linear cohort slope. Assumption C2, represented by the blue areas in the figure, requires that the 1911-1926 cohort group has the same or higher turnout rates than the 1947-1966 cohorts conditional on age and period. When the linear cohort slope is greater than 0.03, the average cohort effect parameters of the latter group exceed those of the 1911-1926 cohorts. Therefore, the assumption bounds the linear cohort slope from above at 0.03 percentage points (see section S2 in the Online Supplement for the derivation).

I summarize the bounds implied by each set of assumptions in Table 2.[15] Notably, combining

[14] Since the identification power of Assumption A1 comes solely from the age groups subject to the most significant decline, other age groups' nonlinear components are redundant in bounding the linear slopes as long as the interval of the most significant decline is unchanged. Assumption A1 is, therefore, effectively weaker than asserted because it only requires that the age effect parameters do not decrease between the age groups 54-57 to 58-61.

[15] Section S2 in the Online Supplement explains, in more detail, how the particular bounds are derived in this study.

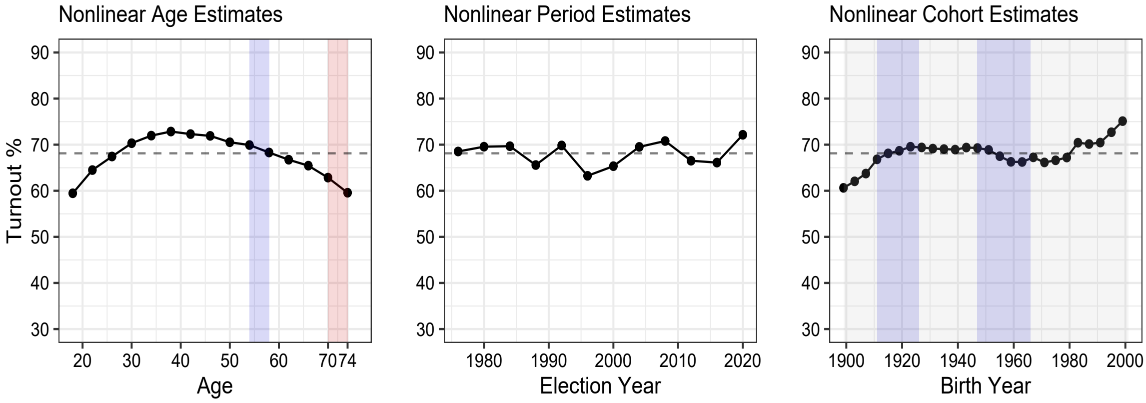

Note: The shaded areas represent the ranges that inform the identification intervals. Survey weights augmented by adjustments for non-response (N=981,858).

Figure 4. Estimates of Nonlinear Age, Period, and Cohort Components

	Age Slope (α_L)	Period Slope (β_L)	Cohort Slope (γ_L)
Assumptions			
A1 + A2	[0.40, 0.82]	[-0.04, 0.37]	[-0.18, 0.23]
A1 + A2 + C1	[0.59, 0.82]	[-0.04, 0.19]	[0.00, 0.23]
A1 + A2 + C2	[0.40, 0.61]	[0.16, 0.37]	[-0.18, 0.03]
A1 + A2 + C1 + C2	[0.59, 0.61]	[0.16, 0.19]	[0.00, 0.03]

Note: The slope estimates represent a change in turnout rates induced by a year increase in age, period, or cohort in terms of percentage points.

Table 2. Upper and Lower Bounds on the Linear Slopes

Assumptions A1, A2, C1, and C2 effectively point identifies all three linear slopes since the three slopes are all restricted to intervals of a length of less than 0.03 percentage points. While it is tempting to embrace only the tightest bounds, I will explain below why caution is a better interpretative modality. One reason can be stated now: when both C1 and C2 are maintained, the bounds appear to be overly sensitive to the regression models' functional form and the precision of the nonlinear component estimates. For example, imposing the logistic function yields implausible bounds (e.g., an upper bound of the linear cohort slope is smaller than its lower bound) when all four assumptions are maintained. To be cautious and remain in line with the general spirit of a partial identification analysis, I will use only Assumption C1 or C2, but not both simultaneously, in subsequent analysis.

4.3. APC-I Model Implementation

As described, in terms of an age–period cross-classification table, the APC-I model follows a form of two-way frequency table analysis with interactions. As such, the estimation of its parameters and statistical inference adhere to standard procedures applicable to this type of analysis. For this study, the R package `APCI` (Xu and Luo 2022) is utilized, which is designed specifically to support the estimation and inferential processes described in Luo and Hodges (2020a). This package also provides additional tools for model evaluation and visualization (for a different estimation approach, refer to Keyes and Li 2010).

4.4. Total APC Effect Parameter Estimates of the Two Methods

In this section, I compare the two methods by presenting APC-I estimates alongside the bounds implied by the combinations of the assumptions introduced in the last section. The two methods' age and period effect estimates align, but the cohort effect estimates differ as expected. I first discuss age and period effect estimates before discussing the differences in the cohort estimates. I conclude with an explanation of the APC-I estimates as if the accounting model were the true model.

Combined Results. In Figure 5, the bounding analysis is represented by the blue-shaded regions. The APC-I estimates are presented in red. All the estimates have been scaled up by their respective intercepts — represented by the dashed horizontal lines of the corresponding colors. The APC-I model's 95-percent confidence intervals are also plotted in red, though their visibility is limited due to their minimal extent (see Online Supplement Table S3 for the detailed estimates).[16]

The APC-I model estimates remain consistent across all three rows as they are unaffected by the bounding assumptions. The age effect parameters display monotonic increases from the 18-21 age group to the 58-61 age group, then decrease from the 70-73 to 74-77 age group. The period effect parameter estimates are plotted in the center column. They display trendless fluctuation from 1976 to 2000, reaching the lowest mark of 65.0 percent in 1996. Between 2004 and 2016, they appear to oscillate around a higher local mean. In 2020, there is a notable surge in the turnout rate by 7.1 percentage points, reaching a peak at 78.8 percent. The raw increase in turnout rates from 2016 to 2020 in the analytical sample is estimated to be 7.8 percentage points. Hence, the increase suggests that more than 90 percent of the increase from 2016 to 2020 is attributable to a period effect (i.e., $7.1/7.8 \approx 0.91$).

Now, consider how the three rows of Figure 5 differ for the bounding analysis of the age and period effects. Figure 5 (a), based on Assumptions A1 and A2, largely align its age and period bounds with the APC-I estimates. The age and period bounds are mostly consistent with the APC-I estimates. Layering Assumption C1 on Assumptions A1 and A2 in Figure 5 (b) narrows the bounds. While the general age and period patterns are similar, the post-1996 period effects in these bounds suggest more modest increases compared to APC-I estimates. For instance, the period effect parameter for 2020 is bounded by [71.1, 76.5], which is lower than the APC-I estimate of 78.8. In Figure 5 (c), I instead add Assumption C2 to Assumptions A1 and A2. For these estimates, the narrower intervals include the APC-I estimates. Notably, the 2020 period effect falls between 75.9 and 80.8, which includes the APC-I estimate of 78.8 for that year.

[16]Standard errors are not reported for the bounding analysis because the large sample size suggests that sampling variations are likely small, and researchers have yet to reach a consensus on the best practices for inference on partially identified parameters. For an overview of inference approaches for such parameters, refer to Tamer (2010).

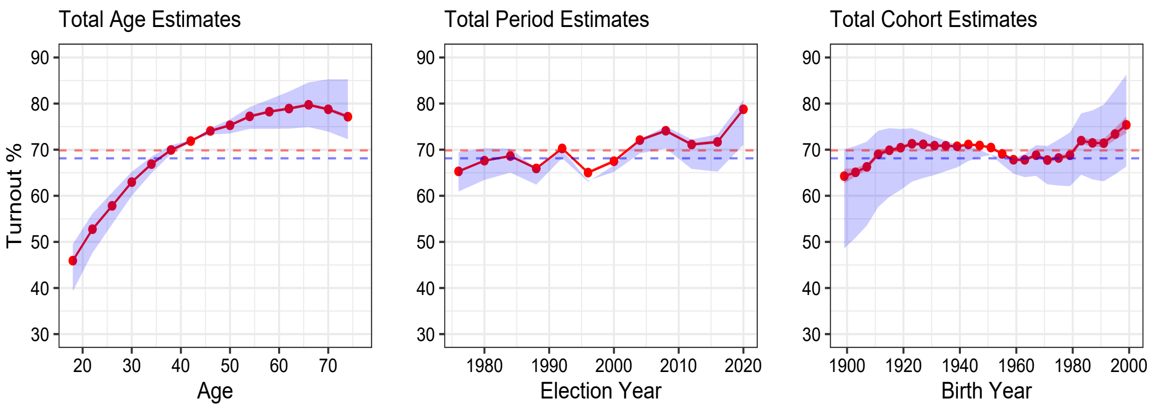

(a) Assumptions A1 + A2

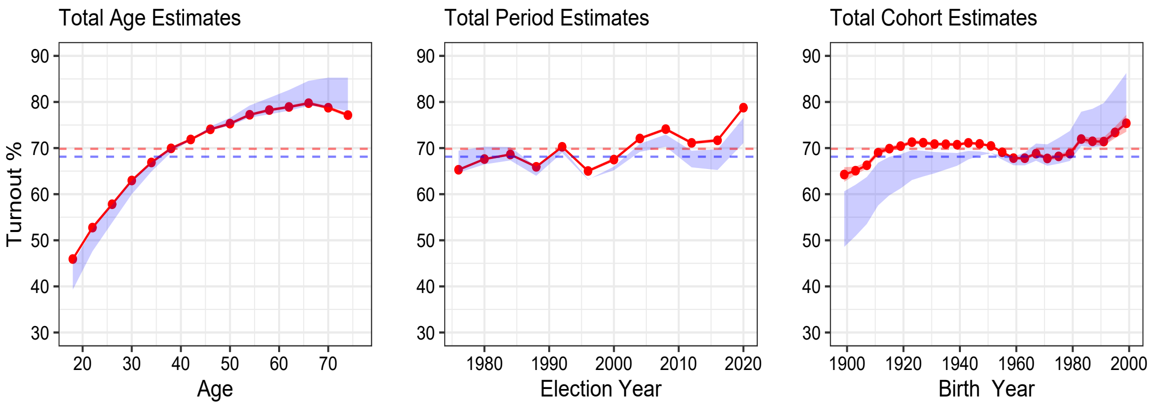

(b) Assumptions A1 + A2 + C1

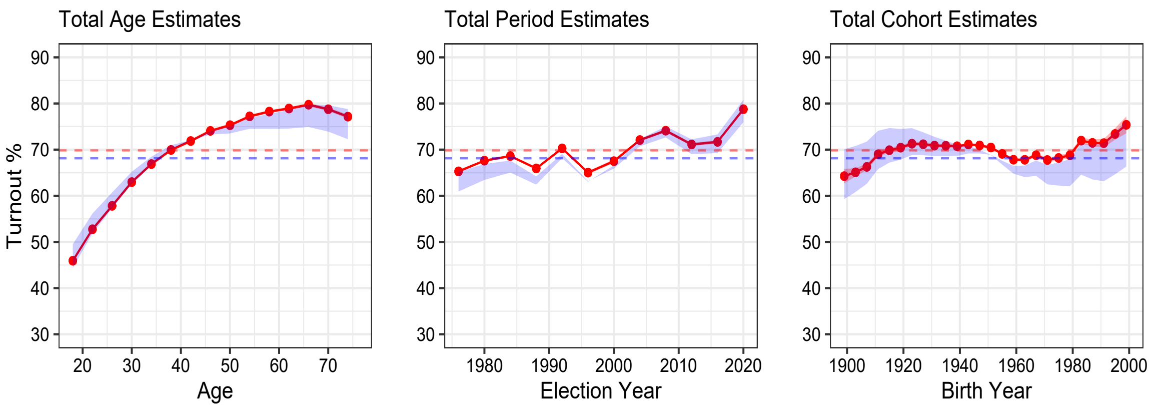

(c) Assumptions A1 + A2 + C2

Note: The blue shaded areas represent the bounding analysis identification intervals. The red circles represent the APC-I model's point estimates. The dashed horizontal lines represent the intercepts. Survey weights augmented by adjustments for non-response (N=981,858).

Figure 5. Total APC Effect Estimates of the APC-I Model and Bounding Analysis

In contrast to the convergent evidence for age and period effect estimates, especially under the A1, A2, and C2 assumptions, the cohort effect estimates display more variation. First, the APC-I estimates indicate that cohort differences are generally modest. The 1910-1950 cohorts yield estimates slightly above the grand mean (i.e., the intercept), while the 1960-1980 cohorts' estimates are below. Estimates for the 1980s cohorts average around two percentage points higher than the grand mean. Notably, the estimates for the two youngest cohorts show a significant increase. However, as these cohorts are only represented in one or two elections within the dataset, these findings should be approached with caution.

Figure 5 (a) shows that little can be inferred about the cohort effects from the bounding analysis based on Assumptions A1 and A2, given the wide implied intervals. Introducing Assumption C1 in Figure 5 (b), the implied non-negative linear cohort slope suggests that the pre-1950 cohort estimates are notably below the APC-I estimates. When Assumption C2 is instead imposed in Figure 5 (c), the cohort estimates appear more aligned with the APC-I estimates. In contrast to the findings in Figure 5 (b), the cohort estimates before 1960 mostly contain the APC-I estimates. However, this bounding analysis suggests lower cohort effect estimates than the APC-I model and the earlier bounding analysis for the cohorts after 1960.

Interpretation of the APC-I Estimates from an Accounting Model Perspective. How can we interpret the differences between the two estimates shown in Figure 5? When adopting the traditional perspective that treats the accounting model as the data-generating model, as detailed in Fosse and Winship (2023), we identify two primary sources of variation in the age and period effect parameters:

1. *Linear cohort slope.* Within the APC-I model, where no explicit cohort term is present, both the age and period effect parameters embed the linear cohort slope γ_L. The APC-I model's linear age component reflects the differential between the accounting model's linear age and cohort slopes, while the linear period component of the APC-I model includes the sum of the accounting model's linear period and cohort slopes.

2. *Nonlinear cohort component.* The age effect parameter of the APC-I model also includes parts of the variation in the accounting model's nonlinear cohort components that correlate with age and are not absorbed by any of the APC-I model parameters. Likewise, the APC-I model's period effect parameter includes the partial covariance with the accounting model's unique nonlinear cohort components.

The second factor, while estimable, plays a minor role in this analysis and likely in many others, as much of the variation in the accounting model's nonlinear cohort components is already integrated and distributed among the APC-I model parameters. Thus, the primary distinction between the age and period effect parameters of the two methods stems from the linear cohort slope. This observation, for example, helps explain why the APC-I model's age and period estimates tend to trace the edges of the bounding analysis intervals in Figure 5 (c). Given that the cohort slope is bounded between -0.18 and 0.03 in this analysis, these APC-I estimates align with a case where the linear cohort slope is near zero.

The differences in the cohort effect parameters are more nuanced. The outcome variation captured by the accounting model's nonlinear cohort component can be decomposed into contributions from age-period interactions and distinctive nonlinear variations in cohorts (Fosse and Winship 2023). The latter encapsulates quite irregular (high-order) outcome variations in cohorts that are unrelated to the variables in the APC-I model. As a result, the outcome variation partitioned by the APC-I model's age-period interactions $\eta^*_{ap(c)}$ and the accounting model's nonlinear cohort component $\tilde{\gamma}_c$ will often overlap significantly. As a result, the APC-I cohort estimates will tend to closely resemble those of the accounting model's nonlinear cohort estimates (Luo and Hodges 2020a:1199), as illustrated by comparing the right panel of Figure 4 against the right panels of Figure 5.

To be sure, this comparison is made entirely from the accounting model's perspective that the constant underlying linear cohort effect is both defined and a parameter of fundamental interest. Now, consider the comparison in the opposite direction: *there is no theoretical basis for considering cohort effects as containing (linear) components that are independent of, and can be separated from, age and period effects.*. This is the implicit position of the APC-I model, which I consider more carefully in the following sections after briefly summarizing the main empirical findings below.

4.5. Summary of the Main Findings

In the wake of the historic 2020 election turnout, one objective of this article was to examine the APC patterns of turnout in the 1976-2020 presidential elections with new data and methods. The main substantive findings can be summarized as follows.

First, age patterns align with past research: 1) age is a strong predictor of turnout; 2) age displays a curvilinear relationship with the turnout; 3) age effects peak when individuals are in their sixties. Second, period effects on turnout have been increasing since the 1996 election. Period effects increased incrementally during the two George W. Bush elections as well as Obama's first-term election. After slight declines in the 2012 and 2016 elections, period effects rose substantially in 2020 to their highest level during the last 44 years. These findings are generally supported by the two methods employed and the sets of substantive assumptions maintained.

Third, the APC-I model suggests that cohort effects are modest compared to the age and period effects. Nevertheless, the most notable cohort effect pertains to those typically identified as Millennials. Voters born between 1983 and 1994 had turnout rates, on average, about two percentage points higher than what their age and period profiles would predict. For those born after 1995, the cohort effects were even more substantial. These voters' turnout rates surpassed predicted levels by 3.6 to 5.5 percentage points, although these estimates are primarily based on these cohorts' participation in the 2016 and 2020 elections. Further analysis of intra-cohort developments through the APC-I model, however, reveals considerable variations in life-course trends among these cohorts' turnout rates, as presented in Figure S3 in the Online Supplement.

5. Discussion

About a decade after Ryder (1965)'s seminal theoretical development, Mason et al. (1973) introduced the APC accounting model, which soon established itself as a prevailing APC analysis approach. From the onset, Mason and colleagues acknowledged the model's conceptual foundations were debatable, hinging critically on the assumption of additive separability of the three effect parameters:

> The appropriateness of a two- rather than a three-way cohort analysis rests on the question of whether we view age, period, and birth cohort as causally distinct in relation to a given dependent variable (Mason et al. 1973:244).

Nonetheless, they firmly supported the need for a "three-way" model, as reflected by:

> [C]ohort analyses which ignore one of the three dimensions of age, birth cohort and time period are often unsatisfactory on substantive grounds. (p. 253)

> Because a distinct causal interpretation can often be applied to age, period and cohort, the failure to control for one of the three variables leaves results open to the possibility of spurious effects. For these cohort problems, a three-way analysis is desirable. (p. 243)

Although initially met with some skepticism (e.g., Glenn 1976), its widespread adoption, particularly within the field of sociology, attests to its broad acceptance among their peers and later researchers.[17] This acceptance led subsequent methodological advancements to primarily focus on addressing the model's inherent identification problem. However, it seems researchers have been less attentive to the theoretical ramifications of this modeling setup, particularly how the derived estimands align with the foundational theories of cohort effects and their interpretive complexities.

As many scholars have previously noted, the accounting model's particular parameterization entails substantial behavioral assumptions. Specifically, within the observation interval (i.e., the range of the periods chosen for analysis), the model assumes: 1) homogeneous marginal age effects for all cohorts and across all periods; 2) period effects that are constant for all age groups; and 3) cohort effects that remain unvarying across all ages and periods (Clogg 1982; Fienberg and Mason 1979; Glenn 1976; Schulhofer-Wohl and Yang 2016). These assumptions are vital as they sustain the general structure of the model, enabling the extrapolation of the effect parameters beyond the joint support of age and period data, thereby facilitating the estimation of APC effects that apply globally to all studied individuals.[18]

Nevertheless, these constraints also introduce significant interpretational challenges. First, they present a substantial conceptual paradox: the model presupposes that any interactions between

[17] For instance, Clogg (1982:460) discusses "intrinsic cohort tendencies" as cohort attributes that are "separate from age or period effects and determined exogenously to the periods of observation."

[18] In this regard, early critics, such as Glenn (1976:902), directly challenged Mason et al. (1973), arguing that the model's assumptions "are rarely realistic in the case of dependent variables of interest to sociologists."

age and period that could give rise to cohort effects are nonexistent within the observation interval, even though it is precisely these interactions that are believed to generate cohort effects. Essentially, for the model to remain internally coherent under these constraints, it necessitates that relevant age-varying period experiences that later manifest as cohort effects be confined to periods before the observation interval. Then, all age groups are, on average, expected to respond identically to period influences during the observation interval, and cohort effects must remain constant for each individual across different periods and ages of observation (Hobcraft, Menken, and Preston 1982). Clearly, the plausibility of such marked behavioral discontinuities may vary by context. However, they do pose limitations as a general framework, especially when considering the often arbitrary selection of the starting period in relation to the outcomes studied.

To illustrate, in this study, the assumptions imply that for individuals born in 1950, the cohort effects are triggered by events from the 1960s to early 1970s, which they experienced, in an age-dependent way, during their formative years. However, once these initial age-period interactions surface as observable cohort effects in the data, these individuals are expected to react to all subsequent period conditions post-1976 — the start of the observation period — in a manner indistinguishable from other age or cohort groups. This expectation suggests that the formative perspectives formed during pivotal moments like the civil rights, anti-war, and liberation movements offer no valuable insights into how these cohort members might engage with later societal events, such as 9/11, the War on Terror, or the Black Lives Matter movement, in terms of their electoral participation (see also Morgan 2023; Morgan and Lee 2024).

Moreover, the notion of unchanging cohort effects seems to diverge from the foundational views presented in the seminal works. For example, while Mannheim (1952:298) clearly underscores the profoundness of "natural worldviews" formed during one's early years, his articulation also implies that these effects progressively evolve via a "dialectic" process between such rooted dispositions and ongoing experiences with emerging societal events.[19] In line with this, Ryder (1965:861) explicitly emphasizes the importance of examining the "intra-cohort temporal development throughout the life cycle" as a crucial element of cohort analysis.

Acknowledging these subtleties, Luo and Hodges (2020a) presents the APC-I model as a methodological and theoretical "critique" of the accounting model framework that more faithfully represents the original theoretical perspectives. Central to this critique is the recognition that defining cohort effects, both conceptually and empirically, must inherently involve age-period interactions, which the constant-effect parameters of the accounting model may not fully capture. Therefore, rather than asserting the existence of pure cohort effects that are separable from age and period influences, the APC-I model proposes that cohort effects are most accurately described as accumulations of age-patterned period effects that occur throughout the life course (see also Schulhofer-Wohl

[19]Specifically, Mannheim (1952:298) writes: "Early impressions tend to coalesce into a natural view of the world. All later experiences then tend to receive their meaning from this original set, whether they appear as that set's verification and fulfillment or as its negation and antithesis. Experiences are not accumulated in the course of a lifetime through a process of summation or agglomerations, but are "dialectically" articulated."

and Yang 2016).[20][21] As a logical consequence of these principles, within this approach, the linear component of a constant cohort effect parameter, as posited by the accounting model, is not only unidentifiable but fundamentally *irrelevant* in modeling cohort-driven patterns of social change.

This conceptual reorientation of the APC-I model achieves three primary objectives. First, it renders unnecessary the APC identification problem. Second, it permits age-heterogeneous period effects and allows for the continuous renewal of cohort differences within the observation interval, as implied by the foundational theoretical pieces. Third, as a result, it enables the APC-I model to directly analyze intra-cohort developments throughout the life cycle, an endeavor that Ryder recognized as critically important. Recent empirical applications of the APC-I model are leveraging this additional functionality to generate potentially important insights in their respective areas (e.g., Lu and Luo 2021; Lu, Luo, and Santos 2022; Phillips 2022).

In this respect, the APC-I model could be viewed as simpler yet substantively more realistic, aligning closer with the original theoretical intent of APC analysis. However, in this approach, too, contemporaneous aging and period effects from before the observation interval are not directly observed, making it difficult to disentangle the earlier marginal age and period effects from the age-period interactions that later surface as cohort effects. Thus, the cohort-effect estimands of the APC-I model reflect a mix of 1) age-varying period effects that accumulate within the observation interval and 2) the intertwined age and period effects (including their interactions) from before the observational interval that are carried forward to the initial period of observation in the data (see Morgan and Lee 2024).[22]

[20]The interpretation of cohort effects as age-period interactions is well-established in epidemiological research, reflecting the field's specific analytical aims. In contrast, the sociological approach, as demonstrated by the accounting model, aims to understand broad societal changes or "structural transformations," portraying cohorts or generations as causal agents of such changes (Mannheim 1952; Ryder 1965). In this sociological framework, belonging to a particular cohort tends to be seen as a significant "exposure" that encompasses the collective historical experiences of that group. Meanwhile, the epidemiological viewpoint treats cohort effects primarily in terms of age-period interactions, focusing on the impact of more local and specific incidents like pandemics, natural disasters, or major public health actions. Here, cohorts serve more as indices that capture age-specific susceptibilities to these critical events or conditions (for a more comprehensive discussion, see Keyes et al. 2010).

[21]While the APC-I model does not inherently suggest it, extending this viewpoint to its logical limit prompts an intriguing conceptual consideration: do cohort effects represent any intrinsic causal substance (i.e., social processes) that is ontologically irreducible to particular configurations of age and period effect? If causal primacy is assigned to age and period and cohort effects are articulated simply as age-structured lagged period effects, one might argue that what we refer to as cohort effects are not independent causal forces but are rooted in the foundational processes of aging and historical context, thus offering limited unique causal contribution from the cohorts themselves. Nevertheless, even within this reductionist view, cohorts may still serve as informative "formal variables" or "surrogate indexes" that help to trace demographic processes influencing social change (for more comprehensive discussions, see Davis 2001; Hobcraft et al. 1982; Morgan 2023).

[22]To be clear, this limitation is not inherent to the APC-I model. In the accounting model, this limitation is dealt with by extrapolations of, say, age effects from the support of the age and period data to those whose aging processes are directly not observed in the data.

6. Conclusion

The broader goal of this article was to elucidate the similarities and differences between APC bounding analysis and the APC-I model by applying them simultaneously in a single study, thereby providing an empirical bridge between these two principled approaches to APC analysis. Each method demonstrates significant promise in furthering their respective facets of APC analysis, both of which have long co-existed. APC bounding analysis provides a robust path within the conventional APC framework in sociology. It directly confronts the identification problem, compelling researchers to clarify and justify the assumptions crucial for extracting meaningful data.

On the other hand, the APC-I model presents an alternative approach to APC analysis, firmly rooted in a distinct theoretical understanding of cohort effects. It regards continuously accumulating age-period interactions as central to cohort effect estimation, thereby prompting a reevaluation of the sometimes tricky interpretative bases of the accounting model. From a purely statistical standpoint within the traditional accounting model framework, the APC-I model could be seen as a variant that assumes a zero cohort slope and refines the accounting model's nonlinear cohort components to derive cohort effect estimates. However, the fundamental distinction between these models lies in their foundational conceptual philosophies, not in their mathematical formulations.

For too long, the APC literature has been ensnared in the allure of purely technical advancements, overshadowing the vital early insight that "atheoretical cohort analysis is a useless exercise" (Glenn 1976:903). The introduction of methodologies that prioritize theoretical rigor marks a progressive step. By contrasting these methodologies, this study emphasizes the importance of a robust theoretical framework to define cohort effects and a statistical approach that complements it. The choice of methodology in APC analysis extends beyond mere statistical considerations; it demands a grounding in behavioral theories and deep substantive knowledge. Thus, it is essential for researchers to precisely define "cohort effects" and their sociological interpretations in their work, moving away from the frequent pitfall of adhering to conventions without solid justification.

In practical terms, researchers aligned with the theoretical foundations of the accounting model who are also seeking a principled approach to address the inherent identification issues may find the bounding analysis particularly valuable. Alternatively, those who view cohort effects as age-period interactions and are keenly interested in exploring intra-cohort variations may discover significant advantages in using the APC-I model. To emphasize, it is crucial to ensure greater transparency about the conceptual and theoretical bases of the methodologies used, regardless of the chosen approach.

Given the entrenched position of the accounting model within the discipline, it appears that sociologists have yet to fully explore the analytical possibilities presented by the broader perspective represented by the APC-I model. As innovative methods like the bounding analysis emerge to address the identification problem, there is also a burgeoning potential to integrate distinctive insights from APC literature that align with alternative conceptual frameworks (for recent examples, see Morgan and Lee 2024; Schulhofer-Wohl and Yang 2016). Moving forward, these developments promise to enhance and diversify the analytical resources available for analyzing social change.

Author's Note

Replication materials, including the code used for the analysis in this study, can be found in the author's GitHub repository: http://github.com/lee-jiwon/apc-turnout-smr.

References

Abramson, P. R. (1974). Generational change in american electoral behavior. *American Political Science Review 68*(1), 93–105.

Alwin, D. F. (1998). The political impact of the baby boom: Are there persistent generational differences in political beliefs and behavior? *Generations: Journal of the American Society on Aging 22*(1), 46–54.

Alwin, D. F., R. L. Cohen, and T. M. Newcomb (1991). *Political attitudes over the life span: The Bennington women after fifty years.* University of Wisconsin Press.

Bailey, M. J. and S. Dynarski (2011). Inequality in postsecondary education. In G. Duncan and R. Murnane (Eds.), *Whither Opportunity? Rising Inequality, Schools, and Children's Life Chances.* New York, New York: Russell Sage.

Bartels, L. M. and S. Jackman (2014). A generational model of political learning. *Electoral Studies 33*, 7–18.

Bell, A. (2019). Age period cohort analysis: A review of what we should and shouldn't do. *Annals of Human Biology 47*(2), 208–217.

Bell, A. and K. Jones (2014). Another 'futile quest'? A simulation study of yang and land's hierarchical age-period-cohort model. *Demographic Research 30*, 333–360.

Brady, H. E., S. Verba, and K. L. Schlozman (1995). Beyond SES: A resource model of political participation. *American Political Science Review 89*(2), 271–294.

Campbell, A., P. E. Converse, W. E. Miller, and D. E. Stokes (1980). *The American voter.* University of Chicago Press.

Cancela, J. and B. Geys (2016). Explaining voter turnout: A meta-analysis of national and subnational elections. *Electoral Studies 42*, 264–275.

Clogg, C. C. (1982). Cohort analysis of recent trends in labor force participation. *Demography 19*(4), 459–479.

Corasaniti, N. and J. Rutenberg (2020). Republicans pushed to restrict voting. Millions of Americans pushed back. *New York Times.*. December 5, 2020 (https://www.nytimes.com/2020/12/05/us/politics/2020-election-turnout.html).

Dassonneville, R. (2017). Age and voting. In K. Arzheimer, J. Evans, and M. S. Lewis-Beck (Eds.), *The SAGE handbook of electoral behaviour*, pp. 137–158. Sage.

Davis, J. A. (2001). Testing the demographic explanation of attitude trends: Secular trends in attitudes among us householders, 1972–1996. *Social Science Research 30*(3), 363–385.

Drutman, L. (2021). The high turnout in 2020 wasn't good for American democracy. *The Washington Post.*. February 10, 2021 (https://www.washingtonpost.com/politics/2021/02/10/high-turnout-2020-wasnt-good-american-democracy/).

Elbers, B. (2020). Orthogonal polynomial contrasts and applications to age-period-cohort models. *SocArXiv*.

Fienberg, S. E. and W. M. Mason (1979). Identification and estimation of age-period-cohort models in the analysis of discrete archival data. *Sociological Methodology 10*, 1–67.

Flood, S., M. King, S. Ruggles, and J. R. Warren (2021). *Integrated public use microdata series, current population survey: Version 8.0.[dataset].*

Fosse, E. and C. Winship (2019a). Analyzing age-period-cohort data: A review and critique. *Annual Review of Sociology 45*, 467–492.

Fosse, E. and C. Winship (2019b). Bounding analyses of age-period-cohort effects. *Demography 56*(5), 1975–2004.

Fosse, E. and C. Winship (2023). The anatomy of cohort analysis: Decomposing comparative cohort careers. *Sociological Methodology 53*(2), 217–268.

Frenk, S. M., Y. C. Yang, and K. C. Land (2013). Assessing the significance of cohort and period effects in hierarchical age-period-cohort models: applications to verbal test scores and voter turnout in us presidential elections. *Social Forces 92*(1), 221–248.

Glenn, N. D. (1976). Cohort analysts' futile quest: Statistical attempts to separate age, period and cohort effects. *American Sociological Review 41*(5), 900–904.

Glenn, N. D. and M. Grimes (1968). Aging, voting, and political interest. *American Sociological Review 33*(4), 563–575.

Hess, A. J. (2020). The 2020 election shows Gen Z's voting power for years to come. *CNBC.*. November 8, 2020 (https://www.cnbc.com/2020/11/18/the-2020-election-shows-gen-zs-voting-power-for-years-to-come.html).

Hobcraft, J., J. Menken, and S. Preston (1982). Age, period, and cohort as sources of variation in demography. *Population Index 48*, 4–43.

Holford, T. R. (1983). The estimation of age, period and cohort effects for vital rates. *Biometrics 39*(2), 311–324.

Hout, M. and D. Knoke (1975). Change in voting turnout, 1952–1972. *Public Opinion Quarterly 39*(1), 52–68.

James, I. R. and M. R. Segal (1982). On a method of mortality analysis incorporating age-year interaction, with application to prostate cancer mortality. *Biometrics*, 433–443.

Jennings, M. K. and R. G. Niemi (2014). *Generations and politics.* Princeton University Press.

Keyes, K. M. and G. Li (2010). A multiphase method for estimating cohort effects in age-period contingency table data. *Annals of epidemiology 20*(10), 779–785.

Keyes, K. M., R. L. Utz, W. Robinson, and G. Li (2010). What is a cohort effect? comparison of three statistical methods for modeling cohort effects in obesity prevalence in the United States, 1971–2006. *Social Science & Medicine 70*(7), 1100–1108.

Land, K. C., G. C. Hough, and M. M. McMillen (1986). Voting status life tables for the united states, 1968–1980. *Demography 23*(3), 381–402.

Lu, Y. and L. Luo (2021). Cohort variation in us violent crime patterns from 1960 to 2014: An age–period–cohort-interaction approach. *Journal of Quantitative Criminology 37*, 1047–1081.

Lu, Y., L. Luo, and M. R. Santos (2022). Social change and race-specific homicide trajectories: An age-period-cohort analysis. *Journal of Research in Crime and Delinquency*, 00224278221129886.

Luo, L. and J. S. Hodges (2020a). The age-period-cohort-interaction model for describing and investigating inter-cohort deviations and intra-cohort life-course dynamics. *Sociological Methods & Research*.

Luo, L. and J. S. Hodges (2020b). Constraints in random effects age-period-cohort models. *Sociological Methodology 50*(1), 276–317.

Lyons, W. and R. Alexander (2000). A tale of two electorates: Generational replacement and the decline of voting in presidential elections. *The Journal of Politics 62*(4), 1014–1034.

Mannheim, K. (1952). The problem of generations. In P. Kecskemeti (Ed.), *Karl Mannheim: Essays.* Routledge.

Manski, C. F. (1990). Nonparametric bounds on treatment effects. *The American Economic Review 80*(2), 319–323.

Mason, K. O., W. M. Mason, H. H. Winsborough, and W. K. Poole (1973). Some methodological issues in cohort analysis of archival data. *American Sociological Review 38*(2), 242–258.

McDonald, M. P. (2020). Trump and polarization drove record turnout. So did mail voting, which should be universal. *USA Today.*. November 4 , 2020 (https://www.usatoday.com/story/opinion/2020/11/04/trump-polarization-mail-voting-drove-record-2020-turnout-column/6153959002/).

McDonald, M. P. and S. L. Popkin (2001). The myth of the vanishing voter. *American Political Science Review 95*(4), 963–974.

Miller, W. E. (1992). The puzzle transformed: Explaining declining turnout. *Political Behavior 14*(1), 1–43.

Miller, W. E. and J. M. Shanks (1996). *The new American voter.* Harvard University Press.

Morgan, S. L. (2023). A double-diamond retrospective on modeling change in attitudes and opinions. *Social Science Research 109*, 102809.

Morgan, S. L. and J. Lee (2024). A rolling panel model of cohort, period, and aging effects for the analysis of the general social survey. *Sociological Methods & Research 53*(1), 369–420.

Morgan, S. L. and C. Winship (2015). *Counterfactuals and causal inference.* Cambridge University Press.

Ohtaki, M., D.-K. Kim, and M. Munaka (1990). A nonparametric method for estimating interaction effect of age and period on mortality. *Environmental Health Perspectives 87*, 115–121.

Patterson, T. E. (2002). *The vanishing voter: Public involvement in an age of uncertainty.* Vintage.

Pew Research Center (2020, August). Voters are highly engaged, but nearly half expect to have difficulties voting. Technical report, Washington, D.C.

Phillips, J. (2022). Affective polarization: Over time, through the generations, and during the lifespan. *Political Behavior 44*(3), 1483–1508.

Putnam, R. D. (2000). *Bowling alone: The collapse and revival of American community.* Simon and schuster.

Rodgers, W. L. (1982). Estimable functions of age, period, and cohort effects. *American Sociological*

Review, 774–787.

Rosenstone, S. J. and J. M. Hansen (1993). *Mobilization, participation, and democracy in America.* Longman Publishing Group.

Ryder, N. B. (1965). The cohort as a concept in the study of social change. *American Sociological Review 47*(6), 843–861.

Scherer, Z. (2021). Majority of voters used nontraditional methods to cast ballots in 2020. *New York Times.*. April 29, 2021 (https://www.census.gov/library/stories/2021/04/what-methods-did-people-use-to-vote-in-2020-election.html).

Schulhofer-Wohl, S. and Y. C. Yang (2016). Modeling the evolution of age and cohort effects. In R. Schoen (Ed.), *Dynamic Demographic Analysis*, pp. 313–335. Springer.

Smets, K. (2021). Aging and politics: age differences in political behavior in comparative perspective. In K. Ferraro and D. Carr (Eds.), *Handbook of Aging and the Social Sciences*, pp. 287–300. Elsevier.

Smets, K. and A. Neundorf (2014). The hierarchies of age-period-cohort research: Political context and the development of generational turnout patterns. *Electoral Studies 33*, 41–51.

Smets, K. and C. Van Ham (2013). The embarrassment of riches? A meta-analysis of individual-level research on voter turnout. *Electoral Studies 32*(2), 344–359.

Stockemer, D. (2017). What affects voter turnout? A review article/meta-analysis of aggregate research. *Government and Opposition 52*(4), 698–722.

Tamer, E. (2010). Partial identification in econometrics. *Annual Review of Economics 2*(1), 167–195.

Teixeira, R. A. (1992). *The disappearing American voter.* Brookings Institution Press.

Verba, S., K. L. Schlozman, H. Brady, and N. H. Nie (1993). Citizen activity: Who participates? What do they say? *American Political Science Review 87*(2), 303–318.

Wolfinger, R. E., S. J. Rosenstone, and S. J. Rosenstone (1980). *Who votes?* Yale University Press.

Xu, J. and L. Luo (2022). Apci: an R and Stata package for visualizing and analyzing age-period-cohort data. *The R Journal 14*(2), 77.

Yang, Y. and K. C. Land (2006). A mixed models approach to the age-period-cohort analysis of repeated cross-section surveys, with an application to data on trends in verbal test scores. *Sociological Methodology 36*(1), 75–97.

Yoder, J., C. Handan-Nader, A. Myers, T. Nowacki, D. M. Thompson, J. A. Wu, C. Yorgason, and A. B. Hall (2021). How did absentee voting affect the 2020 US election? *Science Advances 7*(52), eabk1755.