When Do Voter Files Accurately Measure Turnout? How Transitory Voter File Snapshots Impact Research and Representation

Seo-young Silvia Kim[*‡] and Bernard L. Fraga[†]

[‡]American University
[†]Emory University

September 13, 2022

Abstract

Voter files are an essential tool for both election research and campaigns, but relatively little work has established best practices for using these data. We focus on how the timing of voter file snapshots affects the most commonly cited advantage of voter file data: accurate measures of who votes. Outlining the panel structure inherent in voter file data, we demonstrate that opposing patterns of accretion and attrition in the voter registration list result in temporally-dependent bias in estimates of voter turnout for a given election. This bias impacts samples for surveys, experiments, or campaign activities by skewing estimates of the potential and actual voter populations; low-propensity voters are particularly impacted. We provide an approach that allows researchers to measure the impact of this bias on their inferences. We then outline methods that measurably reduce this bias, including combining multiple snapshots or using commercial files that preserve the turnout histories of dropped voters.

[*]Kim is an Assistant Professor at American University, and is the corresponding author (email: sskim.research@gmail.com. We thank Bradley Spahn and Alan Yan for their work on related projects examining voter file data that shaped the present manuscript.

Voter registration databases or voter files[1] are a rich and important source of data increasingly used for both election campaigns and research (Green and Gerber, 2005; Cooper et al., 2009). Since McDonald (2007) noted the "significant, untapped research potential" of voter files, at least 83 articles using voter file data have been published in the three leading outlets for scholars researching American electoral behavior.[2] Examining the literature on American turnout, 45% of articles over the last 15 years made use of voter files, including at least 44 articles appearing in the top three journals over the previous five years alone.[3] A transition from reliance on survey self-reports to administrative records of turnout[4] is well underway.

That said, little work has established best practices for using voter files for academic research (McDonald, 2007; Nyhan et al., 2017; Igielnik et al., 2018). In particular, to produce replicable analyses, researchers must observe static "snapshots" of fundamentally transitory voter registration records. The *timing* of these snapshots can be crucial in how accurate our estimates are for the key quantity motivating the use of voter files in the first place—an accurate measure of who votes. In this manuscript, we provide both theoretical and empirical evidence that the transitory nature of voter files creates opposing patterns of *accretion bias* and *attrition bias* that skew our estimates of voters' true turnout history. Our data includes weekly voter files from North Carolina and Ohio between 2018–2021, as well as voter file snapshots from Georgia and its voting history records. We show that these biases have real-world implications on both research and representation, particularly impacting measures of behavior for racial/ethnic minorities and other low turnout groups. We conclude by discussing methods for reducing bias in both observational and experimental voter file-dependent research.

1 The Structure and Usage of Voter Files

There are several reasons why researchers and practitioners may draw on voter files. For academic research, voter files are administrative records kept by election administrators, and can provide information for past elections when researchers need to revisit political behavior after time has passed. For campaign practitioners, they provide data on where to perform mobilization and persuasion fieldwork, as they carry information on basic demographics, voters' residential addresses, political district assignments, contact information, and turnout history (Hersh, 2015).

However, the main advantage of voter files over surveys is that they reflect official records of voter registration and turnout. Misreporting of voter turnout is a well-known phenomenon

[1] Sometimes these are also referred to as voter (registration) lists. We will use these interchangeably.

[2] The *American Political Science Review*, the *American Journal of Political Science*, and the *Journal of Politics*.

[3] We identified 84 studies in the aforementioned three journals with voter turnout as a primary dependent variable from 2007–2017 and approximately 100 articles from 2018–2022 that have turnout as a primary dependent variable. For a full list, see Appendix (forthcoming). We thank Alan Yan and Sam England for their work on compiling these lists.

[4] Note that we label studies using aggregate statistics of voter turnout as *not* being based on voter files. The number of articles relying solely on survey self-reports is undoubtedly smaller than suggested here.

plaguing survey work (Silver et al., 1986; Bernstein et al., 2001). Thus initial studies used voter file data to validate self-reported turnout by matching survey respondents to voter registration lists (Clausen, 1968; Traugott and Katosh, 1979; Ansolabehere and Hersh, 2012). Later work extended this logic, using raw voter files to determine the number of voters directly (Sigelman et al., 1985; McDonald, 2007) and using voter files for field experiments and to construct sampling frames for surveys (Green and Gerber, 2005; Green, 2006). Because administrative datasets have the additional virtue of having many more observations than typical surveys, voter files have been regarded as good research sources.

This trust in voter files as an official, validated dataset is certainly warranted. Turnout self-reports from surveys are notoriously unreliable, diverging from population-level turnout rates by as much as 24 percentage points (Burden, 2000). Even if some of it is due to a sampling bias rather than misreporting, several studies have found that about 6% of self-reported voters match to a voter file record showing no vote (Ansolabehere and Hersh, 2012; Berent et al., 2016; Jackman and Spahn, 2019). The average aggregate error rate of voter file-derived turnout numbers as reported in McDonald (2007) is consistently lower. Given this, voter file numbers can be regarded as a more trustworthy indicator of whether an individual has voted than survey self-reports, which are the only viable alternative.

Unfortunately, the above advantage of accuracy is contingent on the *timing* of the voter file acquisition. Voter files' contents vary depending on when the observer chooses to take a 'snapshot' of the data. Primarily, this has to do with the fact voter files are constantly updated and maintained. As the underlying electorate shifts due to events such as new registration, deaths (Ansolabehere and Hersh, 2014), felon dis/re-enfranchisements (Morse, 2021; Morris, 2021), and moves (Kim, 2022), the voter files are changing almost daily (Kim et al., 2020) and may even have coverage, measurement, or processing error (Shino et al., 2020). In addition, many states have use-it-or-lose-it voting laws that will remove registrations of active voters if they have not voted for some period (Rosenstone and Wolfinger, 1978). For states with same-day registration provisions, some voters for a given election are not reflected on the registration data until several weeks have passed after Election Day. All of these issues are complicated by the fact that changes are neither real-time nor smoothly administered across time and space (Cao et al., 2022).

That even a relatively fresh file could be missing votes is startling. After all, they are supposed to be minimally burdened by the churn of voters moving in and out of the state or getting dropped due to not voting. This indicates that while voter files have low bias relative to other measures, their measurement errors are not zero and should be taken seriously as a source of error. In some analyses, this bias might be simple to correct. For instance, in the turnout validation studies cited above, simply attributing about 1.6 percentage points of the over-report bias to errors in the files might be a reasonable correction. But in other research designs that depend on multiple elections or older vote history, like the measure of non-response bias in Jackman and Spahn (2019) or experimental effects on habitual voting (Gerber et al., 2003; Coppock and Green, 2016), these biases can grow, leading to serious measurement issues—and effectively renders the second reason for using voter files impractical.

Given this, it is critical to ask: what are the best practices in using voter files and why? To understand what measurement errors occur in voter files, we need to understand their data generating process, which we detail in the next section.

2 Voter Files as Dynamic Datasets

Voter files are, by nature, dynamic, as we have illustrated above. Voters can be added, deleted, or have their information changed in the database, all of which are affected both by election schedules and internal maintenance (Kim et al., 2020). For example, new voters are added to the database when voters come of age, move into the given district, newly register, or renew their registration through events such as felon re-enfranchisement. Conversely, voters may die, move out of the given district, be restricted from voting, or—in some cases—have not voted nor updated their information in a fixed period, leading to removals from the database.

2.1 Temporal Bias via Accretion and Attrition

We classify two primary types of temporal dependence that skew voter file estimates: *accretion* and *attrition*. We define *accretion* as the phenomenon where the true electorate of a given election will not be fully included in a voter file acquired *before* Election Day. Many last-minute registrations are not reflected in real-time, and so a voter file snapshot misses many young and peripheral voters who have overcome the barrier of registration late into the cycle. *Attrition* refers to the issue that files do not reflect the true electorate of a past election *after* Election Day. This mainly stems from the fact that the file goes under list maintenance according to legal requirements and jurisdictions' specific practices. We expand on these concepts below.

Accretion. Accretion bias means that all of those who will vote in an election may not appear in a voter file collected before Election Day. It is an inherent feature of a voter file, not a bug. As an election intensifies, so do mobilization efforts, leading to registrations of more voters. Some people will utilize convenience voting measures such as election day registration (EDR). Only when such a last-minute flurry of activities has passed can there be a significant chance of having all records of registrants who did vote. In a sense, accretion bias results from processing error, as noted in Shino et al. (2020).

However, accretion bias has implications for quantifying demographic differences in turnout. Prior studies have found that EDR or same-day registration (SDR), a convenience voting measure that makes registration and voting a one-step process (Leighley and Nagler, 2013), helps minorities.[5] Alvarez et al. (2002), using estimated that EDR would increase turnout for periph-

[5] Unlike EDR, SDR is commonly used to refer to early voting period registration. However, depending on the

eral voters and among racial groups, Hispanics in particular. Similarly, Alvarez and Nagler (2011), in a policy brief, estimated that EDR would increase California's turnout for Latinos and newly-naturalized citizens by 5.1%, compared to 4.4% for Whites and 4.0% for Blacks. Bryan Cole (2016) also documents that in North Carolina's 2008 elections, compared to White voters, Black voters were slightly more likely to utilize SDR, both in the primaries and in the general election.

This means that if a campaign procures a voter file well before Election Day to carry out ground operations, they may miss a disproportionately minority registrant population who will register during the early voting period or on Election Day. Given underlying differences in turnout and contact rates by race (Fraga, 2018; Ramírez et al., 2018), accretion bias may further exacerbate inequities in voting participation.

Attrition. Attrition bias means that all of those who have voted in a given election may not appear in a voter file collected after Election Day. Since the primary interest of election administrators is to maintain the most up-to-date voter file instead of preserving past records, voters who are no longer eligible to vote or have not updated their information for a long time will typically be removed from the active voter file. The specific details of how this removal or 'purging' occurs varies by jurisdiction—see Kauffman (2018); Welsh-Huggins (2020) for an example.

Existing academic work has focused on reasons for or potential errors with such deletions (Brater et al., 2018; Feder and Miller, 2020; Huber et al., 2021) or its participatory effects (Biggers and Smith, 2020). In this paper, we are focusing on how this affects our understanding of who really voted.

While some jurisdictions may preserve the best relevant voter file for each election, in many cases, the only data available to researchers making new requests for data is the active, current file. Therefore, attempts such as identifying drop-off voters based on voter files (Igielnik et al., 2018) may be subject to bias from when the data was sampled.

An Illustration. In definition, the two biases look symmetric, as both are about temporal coverage of a given election's true electorate. However, as we discuss above, the problems they cause and their implications are very different.

First, we should note that these biases are structural. These do not arise because there is a foul play from election administrators. Yet, they suggest a precious small window of a voter file reflecting the entire, true electorate of even a recent election. Figure 1 shows a possible scenario of how the voter file may only partially carry the list of voters who vote in a particular election.[6] In fact, there might not be a period where the voter file is 'perfect.' There could be instances where the list maintenance of removing movers and nonvoters is performed earlier than adding

author, SDR can sometimes encompass EDR.

[6]The axes/numbers are hypothetical and not reflective of the true magnitude of the problem.

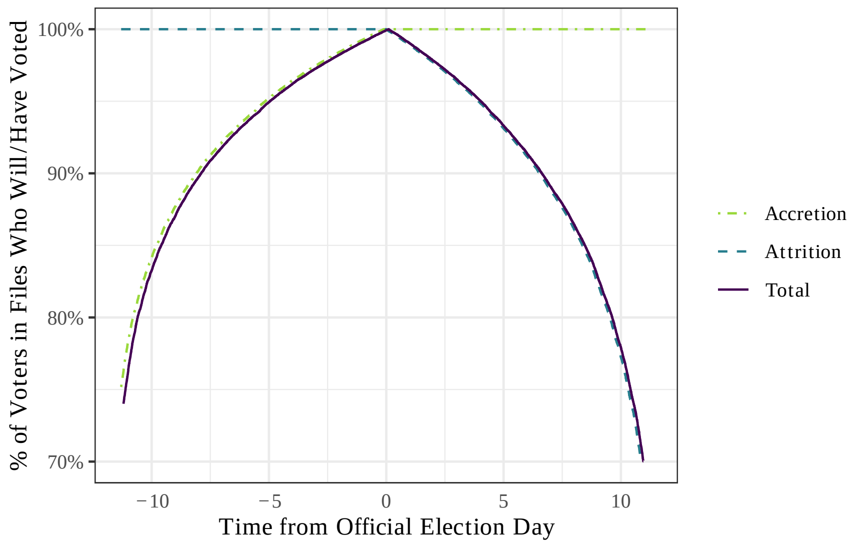

Figure 1: A Hypothetical Reflection of the Temporally Dependent Problems of Voter Files

same-day registrants and voters. The existence and magnitude of accretion and attrition would be available only after gathering regular snapshots of the data and comparing them.

So why are accretion and attrition an issue? Attrition interferes with research. Suppose a researcher takes an interest in using the turnout records of past elections. Many voters' records will be missing for *systematic* reasons, creating an attrition bias in the estimates of the quantities of interest. We show some hypothetical quantities that can be badly skewed, even if these elections were not so long ago. We also show that the biases are stronger for peripheral voting groups, including racial minorities and younger voters. This differential bias poses particular problems for subgroup analyses.

What about accretion? Accretion can create issues for representation because voter files by design do not include non-registrants. Consider an organization working on either get-out-the-vote (GOTV) or persuasion in Nevada during the 2020 presidential election. The operations depend primarily on the state-provided voter file, but late, same-day, and Election Day registrants will be, in ascending order, less likely to appear on a registration list as the age of the file increases. All elements of voting blocs—the size, turnout, and loyalty (Axelrod, 1972)—must be considered for an effective campaign strategy to form a winning coalition. Yet, the law of available data governs grassroots organizations to be constrained to registrants on the voter file at the time of its acquisition. Potentially the most elastic of voters, verging on staying home and turnout, or swing voters that have room to be persuaded, are "unlisted" and cannot be reached (Jackman and Spahn, 2021).

From our read of the literature published in leading political science journals, none of the papers using turnout records in voter files mention the possibility of temporally-dependent biases [7]. This

[7] Kim et al. (2020) discusses the temporal nature of voter file snapshots but focuses on implications for election

lack of recognition is not surprising. For scholars of elections, few have the resources to recognize and rectify potential biases. The solution would require collecting state/national voter files at regular intervals before and after the election and reverse-engineering the change logs. This is because, depending on the research question and context, there could be months to years' worth of gap between Election Day and the voter file obtainment date. Even if a researcher recognizes the temporally-dependent problems, the solution could be costly, depending on the jurisdiction. For example, the Alabama Secretary of State charges $35,854.14 for a single snapshot of all its registered voters.[8] National voter file snapshots from commercial vendors such as Catalist, L2, or Aristotle may also be quite expensive. As a discipline, the "law of available data" will prevail—meaning that individual-level inference with voter files is concentrated around a few states which provide their voter files freely and at regular intervals.

In this paper, therefore, we provide not only the above cautionary tales in using voter files but explore how much of a problem these biases are in practice. In Section 3 we explain how we acquire and use weekly voter files from North Carolina and Ohio between 2018–2021, as well as voter file snapshots from Georgia and its voting history records. Then, in Section 4 we demonstrate the existence of both accretion and attrition bias, quantifying their magnitudes.

3 Data

3.1 Official Turnout Statistics

In order to measure the accuracy of voter file data and the potential for accretion and attrition bias, we need to establish the true number of participants in recent elections. While electoral jurisdictions are generally mandated to provide vote totals by candidate, the aggregate number of *participants* in a given election may exceed the sum of votes for all candidates due to selective participation in certain elections or a myriad of other reasons. That said, most states do report the total number of ballots cast, with the most complete source of such information being the website maintained by McDonald (2022).[9]

3.2 State Voter Files

Our first set of comparison data consists of voter files acquired directly from states. Through requests for information from states and academic researchers, we were able to acquire recent voter files from several states, but for better comparison across the years, we mostly rely on North Carolina, Ohio, and Georgia files.

administration and recordkeeping

[8] Last checked/updated file of March 31, 2021, total 3,585,414 voters.

[9] Some states fail to report total ballot counted numbers, which forces researchers to rely on the number of votes cast for the highest elected office up for election. McDonald (2022) also provides such information, and thus while we prefer to rely on the total ballots cast as our "true" count of the number of voters, we use the votes cast for the highest office as a proxy when necessary.

Rules for removing voters from the voter list vary by state. Glossing over substantial heterogeneity in procedures, some states drop or "purge" voters who have not voted in two consecutive federal elections, while other states mark such voters as "inactive" and keep them on the voter file. Nearly all states remove voters who have died or moved out of state, termed "deadwood" (Ansolabehere and Hersh, 2012). As we will discuss below, this presents a special challenge for researchers because not only do such registrants disappear from the dataset, but if they voted in previous elections, we would completely miss those voters unless we have earlier voter files to compare the current files against.

Which registration status the voter file contains usually has three different flavors: (1) offering only active voters, (2) offering both active and inactive but eligible voters, and (3) offering inactive, active, and some number of removed voters.[10] In addition, how the files are structured is important.

Voter Characteristics and History in a Single File. The first structural type has a single file containing both voter characteristics and history. This means that the voting history is appended to the voter ID and attributes in a wide format, where each row corresponds to a single voter, with several columns indicating turnout for different elections. Those columns would be named, for example, "PRIMARY-08/02/2022." Data does not offer beyond a certain number of turnout variables to keep the distributing database manageable. Both Ohio and Georgia are these types.

Often, the turnout columns have two responses: (1) an indicator for turnout and (2) a missing value that indicates that either the voter was ineligible/unregistered at the time or did not vote despite being registered and eligible. It is up to the researcher to deduce who was eligible at the time, using fields such as voter registration status, first registered date (if available), last registered date, and date of birth.

Voter Data and History in Two Separate Files. The second type has separate voter and history files, with a voter ID that acts as foreign key/crosswalk data. The voter file—defined in the narrower sense—is the same as the first type but without the turnout variables. In our data, North Carolina is this type.

The history file is formed differently, however, and usually offered in a long format so that each row is not a single voter but a single instance of a voter voting in one election. For example, if a voter A has voted in the 2018 and 2020 general elections, she would have two rows in this second type. Naturally, this would not contain instances when the voter was eligible but did not vote, which requires merging with the voter file and deducing who was eligible. This has pros and cons, as we will show later in this paper.

[10] North Carolina is the last type, but for comparison and demonstration, when we show voter file-based figures and tables, we remove the voters who are labeled as "removed" or "declined."

3.3 Third-party Commercial Vendors

There are, of course, commercially compiled datasets in companies such as Catalist, L2, Aristotle, and many others. As of now, there have been no systematic investigations into the differences in quality between these commercial vendors. We leave such explorations to future studies—commercial datasets are outside the scope of this paper.

4 Results

4.1 Evidence of Accretion Bias

Suppose that, as a GOTV organizer, you obtained a Georgia voter file on Jun 10, 2020, to launch campaign operations for the November general election. You would have missed 191,309 late- and same-day registrants, or 2.3% of total voters who actually turned out. If you tried to use the same file for the Senate runoff races that took place on Jan 5, 2021, that is 4.4% of all runoff voters.

Naturally, this gap can be reduced if a later snapshot of the voter file is used. Given an Oct 20, 2020 snapshot, you would only miss 18,127 voters (0.2% of actual voters) of the general 2020 election. However, this leaves very little time frame for mobilization or persuasion efforts. Who are these voters that you would have missed due to the accretion bias?

Racial Gap of Early and Late Registrants. Table 1 shows how late registrants' race distribution differs compared to the all registrants.[11] What is noticeable is that racial minorities registered at higher rates at these late stages of the election especially concentrated close to the election. Asian Americans and Pacific Islanders, in particular, although only 2.6% of the Georgian electorate,[12], formed 11.3% of the last-minute registrants in the two-week period leading up to the general election. Hispanics also showed concentrated registrations around this last-minute period. Black registrants show somewhat subdued registration rates but registered at higher rates right before the runoff election.

Election	Voter File Snapshot	White	Black	Hispanic	AAPI	Other
Gen. 2020	Jun 10, 2020	2.30	-5.83	2.28	2.54	-1.29
Gen. 2020	Oct 20, 2020	-12.19	-2.22	5.54	8.68	0.19
Run. 2020	Jun 10, 2020	1.15	-3.77	1.52	2.43	-1.34
Run. 2020	Oct 20, 2020	-0.73	-3.45	1.79	3.54	-1.15
Run. 2020	Dec 16, 2020	-8.41	3.41	2.80	3.91	-1.72

Table 1: Deviations of Late Registrants' Race Distribution Compared to All Registrants, Georgia, 2020 Elections

[11] Each row, therefore, adds to zero. The voter File of Dec 16, 2020, is used as a reference.
[12] This is conditional on registration.

You may be wondering at this point why we term this as a *bias*. After all, new registrations when there is a contentious race are a natural phenomenon. However, many late registrants—although we used this term for lack of a better word—are *not first-time registrants.* In fact, the more last-minute registrant you are, the likely you are to be a registrant whose record has been (1) restored after a period of having been removed from the file due to the use-it-or-lose-it clause,[13] or (2) for reasons such as administrative backlogs not reflected on the voter roll in time. For simplicity, we will call them *restored voters.*

Election	Voter File Snapshot	Restored Voters (%)	New Voters (%)
Gen. 2020	Jun 10, 2020	17.5	82.5
Gen. 2020	Oct 20, 2020	99.4	0.6
Run. 2020	Jun 10, 2020	14.4	85.6
Run. 2020	Oct 20, 2020	30.3	69.7
Run. 2020	Dec 16, 2020	99.9	0.1

Table 2: Breakdown of Late Registrants: Restored Voters and New Voters

With the June snapshot, a field organizer would have missed 33,035 voters for the 2020 general election whose registration date was before the snapshot, which is 17.5% of all late-registered voters between June and November. With the October snapshot, she would have missed 8,430, 99.4% of those who registered in late October. Similarly, for the runoff, she would have missed 14,383 (30.3%) of late-but-not-new registrants with the October snapshot, and 6,121 (99.9%) with the December snapshot.

Given that the control of the Senate depended on the runoff election that was ultimately decided on the difference of 55,232 votes, these numbers are not trivial. Recognizing that the mileage of a voter file for practical purposes is somewhat limited depending on its snapshot timing and that some of it may be salvaged by keeping track of who is being deleted from the rolls is an important fact that must be acknowledged.

4.2 Attrition Bias

Suppose you are a scholar of elections, and your research design demands data about elections long past. For example, you might be interested in how changes in refugee admissions during the Trump administration changed the political participation of existing naturalized citizens. You might want to, with a similar research agenda, examine whether naturalized citizens' participation fell post the Sept 11 attacks with heterogeneous effects by ethnic groups. But the year is 2022, and you have been informed that the registrar of voters in jurisdictions of interest does not keep historical archives of voter turnout data.

[13] For example, only 1.5% of restored voters in late October voted in the 2018 general election. The number is 2.3% for the 2016 general election.

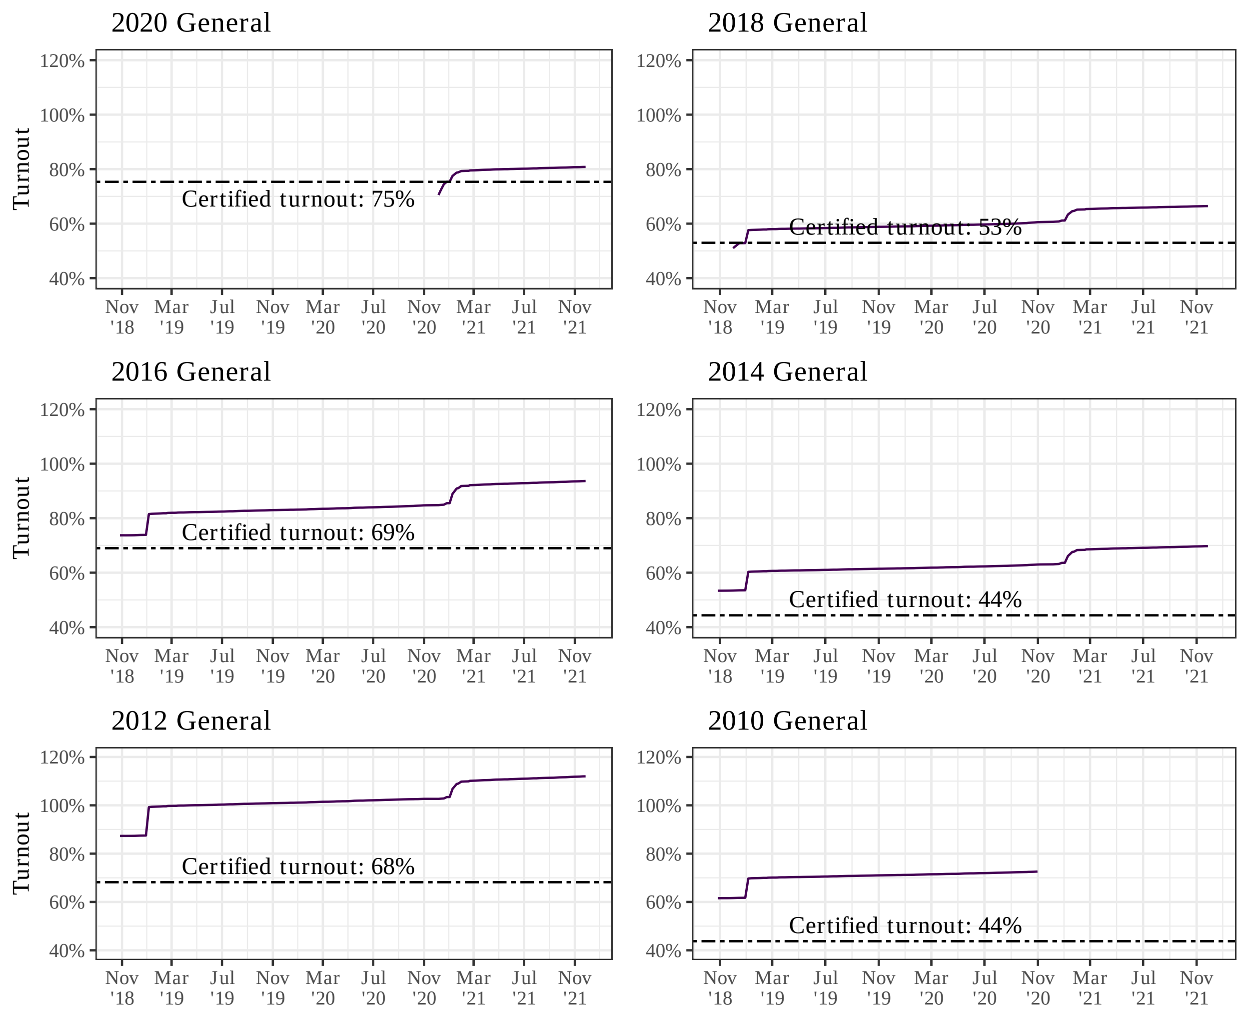

Figure 2: Comparison of Voter File-based Turnout and Certified Turnout, 2010–2020 General Elections, North Carolina

Is it safe to use a recent voter file snapshot to study such elections? Unfortunately, some opportunities 'expire' without securing timely data. First, jurisdictions will likely only provide a few recent elections' worths of data to simplify the publicly distributed data. For example, for a North Carolina file downloaded in the first week of January 2020, the earliest calendar year that comes with individual turnout history is 2010.

But more importantly, even if individual turnout history seems preserved for past elections, we must be wary about accepting them as true records of who voted. Figure 2 shows how voter file-based calculations of turnout may be distorted by using snapshots between 2018–2021. Turnout is calculated as the number of ballots matched to a voter file divided by the number of registrants who are described to be eligible.[14] As can be seen, the further away we are from elections of interest, the starker the turnout gap becomes. In fact, the 2012 general turnout is a preposterous

[14]This includes active, inactive, or temporary (such as UOCAVA voters) voters, excluding those labeled as denied or removed.

number that goes over 100% due to there being records of ineligible voters who voted at the time. Including those labeled as ineligible in the denominator does reduce these gaps, but not close enough to officially certified numbers. For example, with a Nov 27, 2021 snapshot, the 2012 general turnout would be 87.6% for what was a 68.2% turnout.

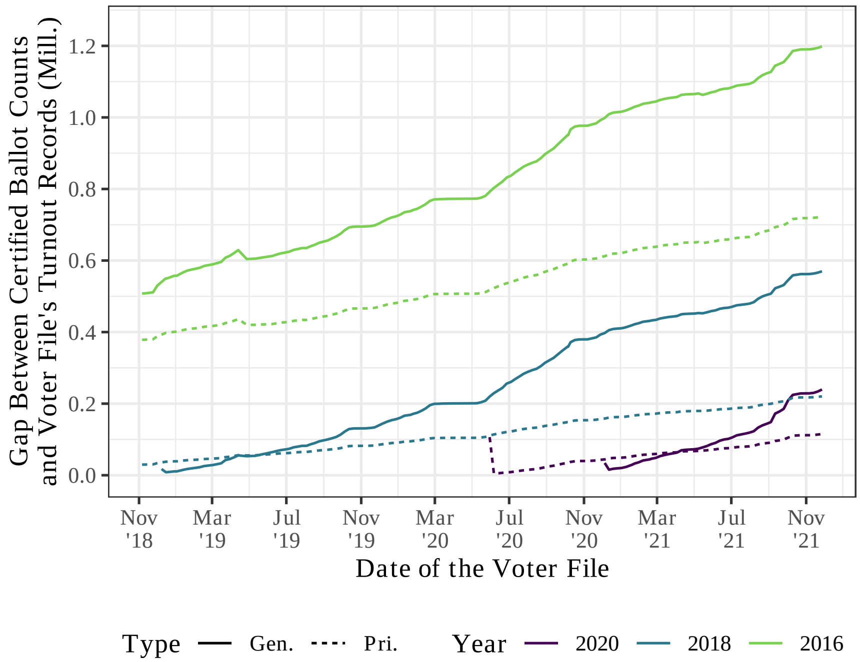

Figure 3: Comparison of Voter File-based Turnout and Certified Turnout, 2016–2020 Elections, Ohio

Figure 3 shows a similar story for Ohio by displaying the gap between certified ballot counts and voter file-based ballot counts. The 2018 general, 2020 primary, and 2020 general elections have all been conducted during the data collection period. As has been shown with North Carolina, earlier elections, such as the 2016 elections, have a significantly larger gap from certified results. In addition, given a particular election, we can see the ballot gap increasing over time as the distance between the election and the voter data snapshot increases.

Partisan Gap of Voters and Non-voters. So we have learned that a later voter file may not carry all voters of a given election because they have been removed for a variety of potential reasons. Does it matter in terms of our understanding of the election? The answer is, unfortunately, yes.

One important statistic that we care about is whether voting is representative—that is to say, whether there are systematic differences between voters and nonvoters (Citrin et al., 2003; Fowler, 2013). In a similar vein, we may be interested in the difference between regular/core and marginal/peripheral voters (Ansolabehere and Schaffner, 2015; Fowler, 2015; Bhatti et al., 2019). Analyzing whether GOTV mobilizations increase representational gaps or inequality (Enos et al.,

2014) is a similar line of research with substantial implications.

In particular, are there partisan gaps between voters and nonvoters? The answer to this, unfortunately, may depend significantly on when the dataset is collected. Figure 4 shows how the partisan gap, calculated as the proportion of Democrats in nonvoters minus the proportion of Democrats in voters, depends on the voter file timing. While Figure 4a shows it over the full data period, the more extreme cases should be discarded, because these are affected by accretion bias. For example, the first dataset after Nov 6, 2018, general election date, collected on Nov 10, 2018, shows a gap close to -0.45. This is because a single Democrat was counted to have voted, and all other registrants' voting histories were not yet reflected on file. Figure 4b truncates those extreme cases strongly affected by accretion bias for better visibility.

We can see an interesting pattern here. First, of course, it is not surprising to see that in most elections, the gap is positive—that is to say, nonvoters are more likely to be Democrats than Republicans. That nonvoters are likely more progressive in their policy preferences has already been suggested in previous works (Lijphart, 1997; Citrin et al., 2003; Martinez and Gill, 2005; Fowler, 2013, 2015), while whether the difference itself is significant enough to swing election results with higher turnout is a different question.

Second, the magnitude of the partisan gap starts to decrease when later voter files are used. In particular, estimates seem to decrease sharply around the 2020 General Election Day. Of course, how much the estimated gap fluctuates vary by election. For example, the 2020 primaries' partisan gap shifts from -0.02 to more than -0.10 when estimated with a later file (approximately a difference of 0.078), while the 2020 general election's gap is, while decreasing, more or less similar to the earlier values.

Suppose that we estimate a structural breakpoint using a simple Bai-Perron approach (Bai, 1997; Bai and Perron, 1998, 2003; Zeileis et al., 2002) that minimizes the sum of squared residuals for all potential number of breakpoints m and the resulting segments $m + 1$.[15] Even if we do not assume that the changes across multiple election's partisan gaps are correlated, common dates emerge as breakpoints when we compute optimal breakpoints for each election's changes.

Election	04/13/2019	04/27/2019	03/14/2020	06/06/2020	11/07/2020	01/23/2021
Pri. 2016	O		O		O	
Gen. 2016	O		O		O	
Pri. 2018			O		O	
Gen. 2018		O	O		O	
Pri. 2020				O	O	
Gen. 2020						O

Table 3: Breakdates Detected, Shifts in Partisan Gap Estimation from Voter Files Over Time, North Carolina, 2016–2020 Elections (From Oct 27, 2018, to Nov 27, 2021, Weekly Updated Files)

[15] Because we do not expect any seasonalities, we do not use other regressions or time-series models to account for temporal regularity.

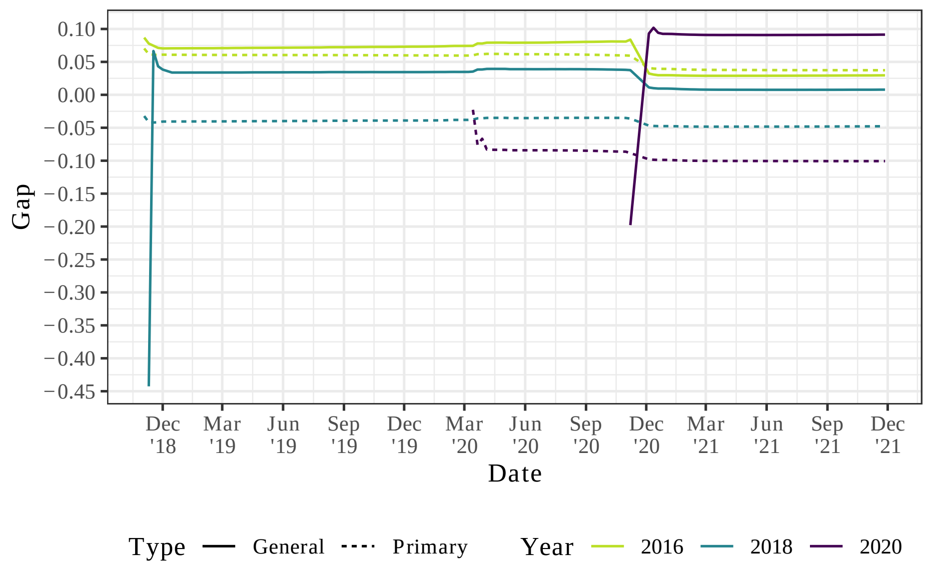

(a) Full Scale, Nov 2018 to Nov 2021

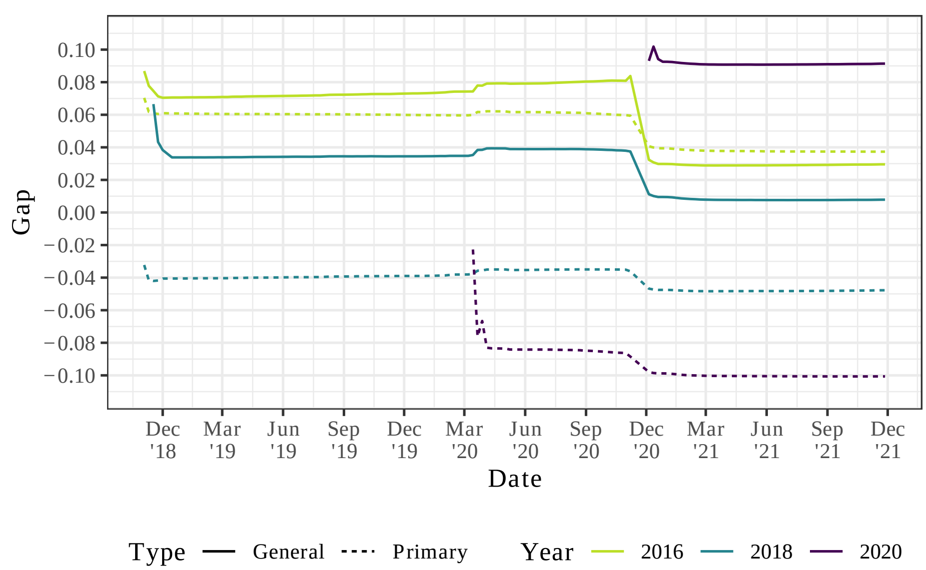

(b) Panel (a) without Periods Affected by Accretion Bias

Figure 4: Partisan Gap between Voters and Non-voters, North Carolina

Table 3 shows the detected change points. We see that some dates emerge as common breakpoints, such as Nov 7, 2020. Because North Carolina updates its voter files every Saturday, and some weeks the data collection was missed, the latent breakpoint might not exactly match the detected dates.[16] But given that Nov 3, 2020, was the general election date, and Mar 3, 2020, was the Democratic and Republican primary date, it seems that post-election list maintenance is bringing about a significant change in who is included in the voter file. Other dates are likely also related to internal list maintenance activities that researchers might not be able to document immediately.

We have shown that the degree to which there are partisan gaps between voters and nonvoters can depend significantly on *when* the voter file is selected to be analyzed. Given that this research question is a vital one in political science that (1) determines whether voting, a bedrock of democracy, is unequal and (2) analyzes whether higher turnout might have shifted election results, we must be watchful of how our conclusions may be a product of the data generating process.

Racial Gap of Voters and Non-voters. We may also care about the descriptive representation of voting (Fraga, 2018). Figure 5 is Figure 4's equivalent for black-white racial gap in North Carolina, and Table 4 corresponds to Table 3. The racial gap is calculated as the proportion of the Black electorate in nonvoters minus the proportion of the Black electorate in voters.[17]

As with the partisan gap, the racial gap in voting may also depend on when the dataset is collected. Looking at Figure 5b for clarity, we can see that for all elections documented, over time, the gap drops significantly—for example, the 2018 general election, which was conducted on Nov 6, 2018. Suppose you procure a snapshot on Nov 24, 2018, which produces a racial gap of 0.08. This will drop to 0.04 in a year and 0.03 in another two years.

As was the partisan gap, the breakpoint estimation gives Mar 14, 2020, and Nov 7, 2020, as major breakpoints. As was with the partisan gap, the post-election list maintenance changes the outlook of who voted in past elections.

Election	03/23/2019	03/07/2020	03/14/2020	05/30/2020	11/07/2020	01/19/2021
Pri. 2016	O		O		O	
Gen. 2016	O		O		O	
Pri. 2018	O	O			O	
Gen. 2018	O					
Pri. 2020				O		
Gen. 2020						O

Table 4: Breakdates Detected, Shifts in Racial Gap Estimation from Voter Files Over Time, North Carolina, 2016–2020 Elections

[16] For example, while the modal period between voter files is a week, three data collections were missed after Nov 7, 2020. Considering the time it takes to update to the most recent data, the likely true breakpoint is after Nov 7, 2020.

[17] We discard those voters who are marked as 'undesignated.'

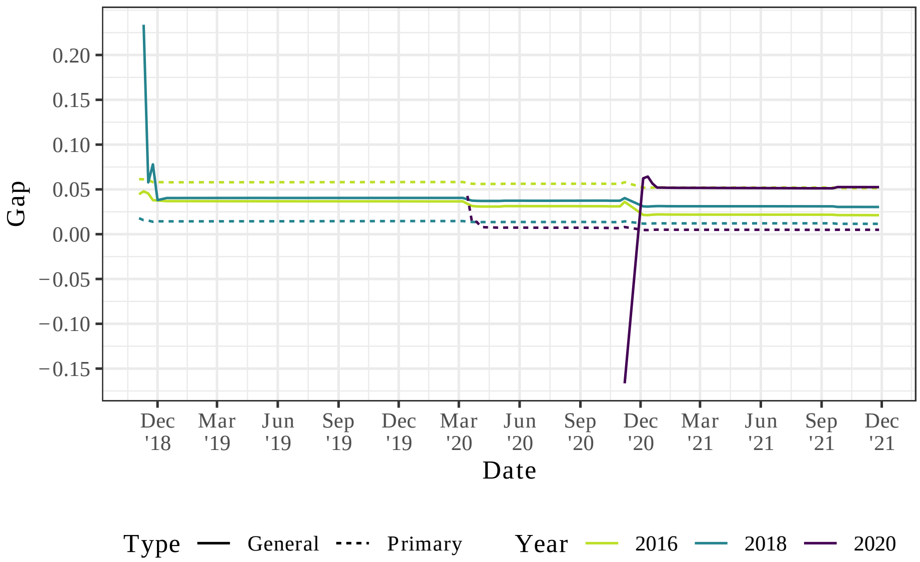

(a) Full Scale, Nov 2018 to Nov 2021

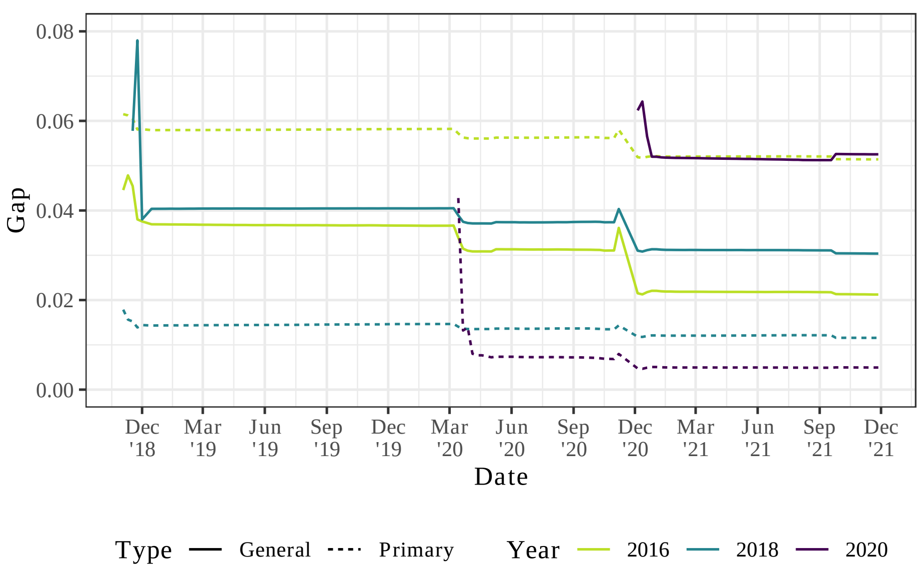

(b) Panel (a) without Periods Affected by Accretion Bias

Figure 5: Black-White Racial Gap between Voters and Non-voters, North Carolina

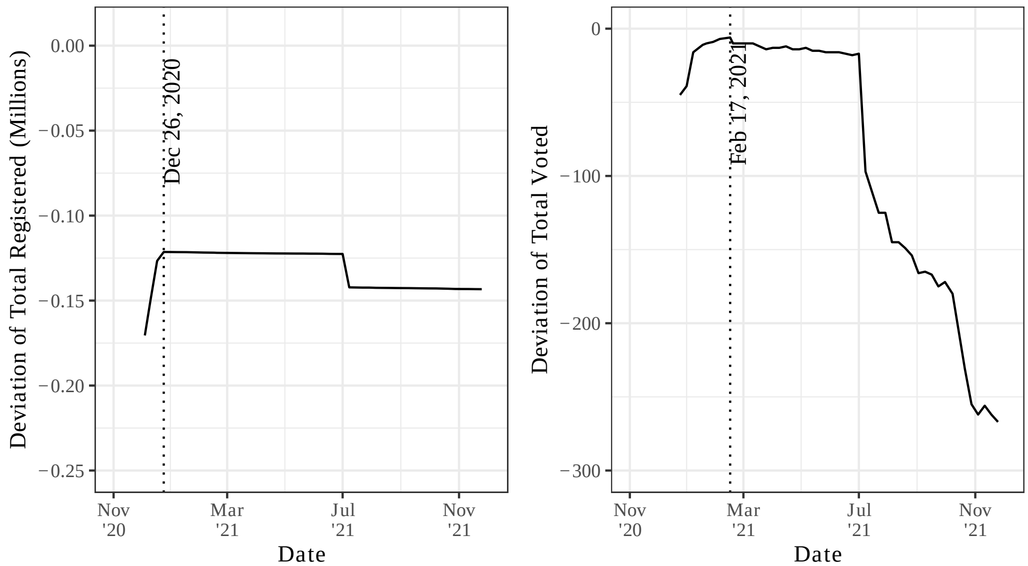

(a) History File Based

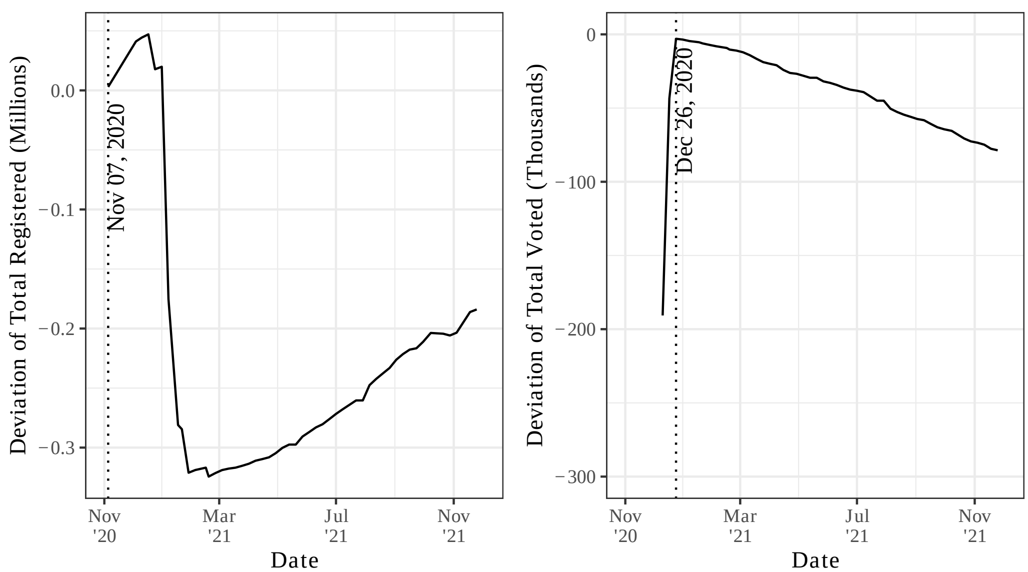

(b) Voter File Based

Figure 6: Absolute Deviation of Total Registered and Total Number of Ballots Compared to Certified Results, North Carolina, 2020 General Elections, Rescaled

4.3 Temporal Bias: A Summary

Given the above, we must ask: is there any time before or after Election Day when the voter file is the best version of itself—that is to say, an accurate record of who has voted in a given election? This allows us to also visualize the accretion and attrition bias in a single framework as to who is missing from the data.

There are two dimensions to this calculation. First, we must identify those who have actually voted. Second, we must identify those who were eligible to vote and registered at the time but have not voted. Given the unique voter IDs for both these groups, in fact, we have three important dimensions of how accurate a given snapshot is, compared to the ground truth: (1) the number of voters—explicitly referring to those who did vote—who are not included, (2) the number of nonvoters who were eligible at the time who are not included, and (3) the difference of turnout calculated from the snapshot and the official, certified turnout.

No voter snapshot may accurately reflect the certified result to begin with. For example, the certified voter turnout statistics for the 2020 general election in North Carolina has 5,544,018 total voters who voted and 7,371,229 total registered voters on the books. Figure 6 shows how compared to these official two numbers, history-file-based and voter-file-based statistics of total registered and total ballots counted differ, rescaled so that the variability unaffected by accretion bias would be better displayed. The vertical lines display the week in which there was a minimum absolute deviation from the benchmark numbers.

First, a voter file collected on Election Day or even a week after will fail to contain much of the updates about the most recent election. Based on voter files (Figure 6), on Nov 7, 2020, there were only five ballots counted as having voted. By Dec 6, 2020, which is a month after the election, there are still 329,339 ballots short. This disparity falls under three digits only by Dec 26, 2020.

For simplicity, we will use the sum of absolute deviations for the number of voters and nonvoters as an accuracy metric. In this specific case, we choose Dec 26, 2020, as the most accurate snapshot for the 2020 general election, which is about 7.5 weeks from Election Day and still has 45 ballots missing and 17,743 more eligible voters than were certified based on the voter file. *On no account does this mean systematic fraud in these datasets.* This means that the cycle at which the publicly available voter files are renewed does not match the certification timing. If we use only the history file, 121,404 eligible voters cannot be found, as the history file will not contain those eligible but who never voted in recent elections.

Given that we have the most accurate snapshot identified, we now have a list of voter IDs and nonvoter IDs for the election. With Figure 7 we can see how they deviate before and after the Election Day based on the voter file, i.e., the realization of Figure 1.

We do see the inverse U-curve that we hypothesized. In addition, we can see that the root of the attrition bias is less about those who voted, but the removal of nonvoters who were eligible at the time. Out of 7.3 million, we see more than 1 million nonvoters (14.9% of registrants) of the election removed from the books within a year. However, the number of voters who did cast a ballot but then are removed in a year—again, likely for valid reasons such as death or residential moves—are also significant at more than 80,000 (1.5% of voters). Depending on the research question and the effect size being estimated, this may break or make the estimation strategy.

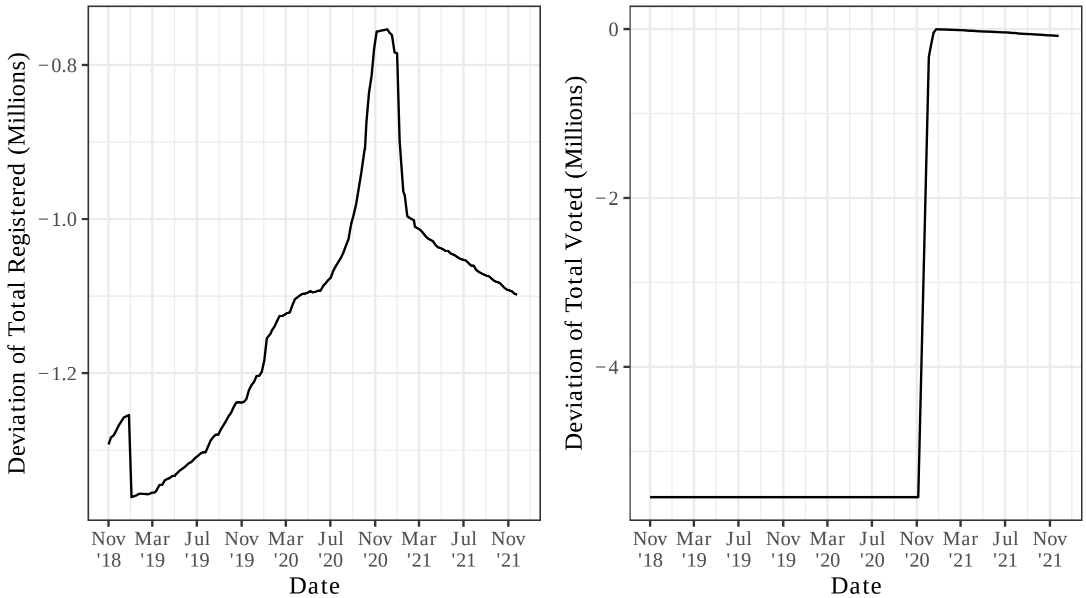

Figure 7: Absolute Deviation of Total Registrants and Total Number of Voters Compared to Most Accurate Snapshot (Dec 26, 2020), North Carolina, 2020 General Elections, Full Scale

5 Toward Best Practices in Voter File Analysis

5.1 Reverse-Engineering the Most Accurate Individual Turnout Records

We have seen above that a voter file snapshot very close to Election Day might not be appropriate to capture turnout. Table 5 shows that the first reasonably good snapshot containing less than 2,500 deviations in ballot counts is around 5 to 9 weeks after Election Day. Again, this timing of publicly distributed, individual-level voting history does not have to match the certification timing of the jurisdiction.

Election	Election Date	Official Ballot Count	First Date, Permissible Diff.	Diff	Days from Election Day
Gen. 2018	2018-11-06	3,755,778	2018-12-15	-1,719	39
Pri. 2020	2020-03-03	2,164,731	2020-05-02	-293	60
Gen. 2020	2020-11-03	5,545,848	2020-12-26	-1,875	53

Table 5: First Date with Permissible Difference Between Official Ballot Count and History File-Based Ballot Count, North Carolina

However, we might be unable to acquire voter files at the ideal time. It could be whether the budget was executable at the time or whether a research agenda was formed later in hindsight. It could be just a human error in forgetting to buy or download the file at the right time. And the longer we wait, the more likely it is that stronger attrition bias kicks in, resulting in many relevant nonvoters being removed. We would also like to minimize the number of false positives in those who were eligible to vote.

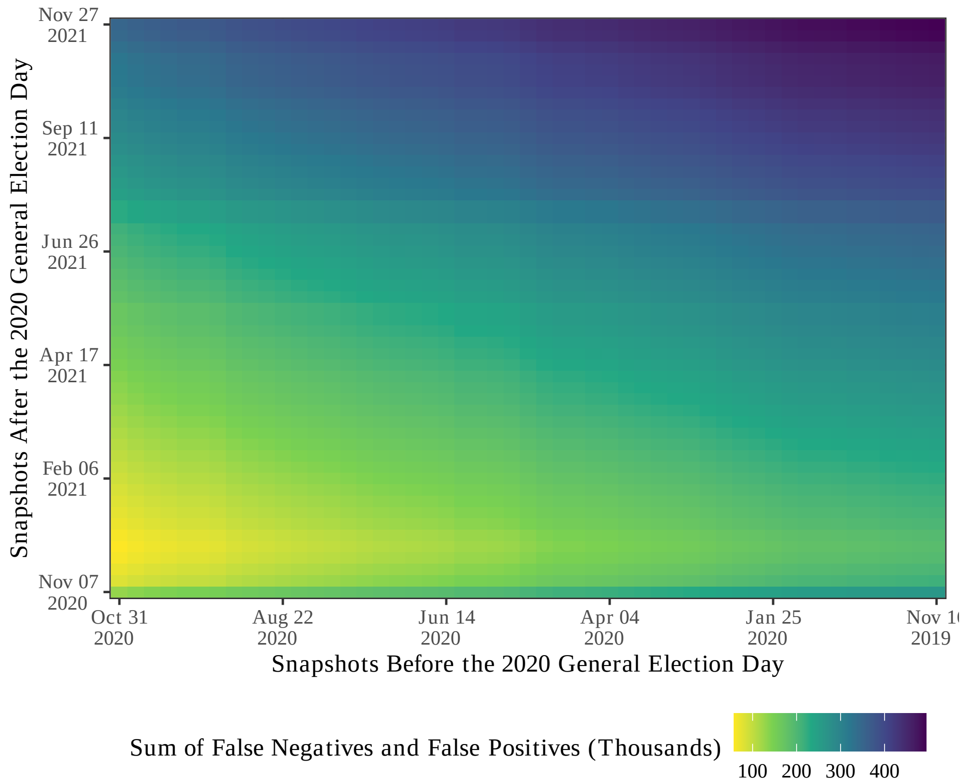

Figure 8: Coverage Error by Using Two Snapshots Before and After Election Day, North Carolina

What would we do in those instances? An intuitive strategy is to secure one snapshot before Election Day and another after Election Day. Figure 8 shows how the sum of false negatives and false positives would be distributed by a combination of snapshots, using snapshots one year before Election Day and one year after. The further one goes to the right side of the x-axis or above the y-axis, the further away your snapshots are from Election Day, and we are looking to minimize this coverage error.

We see that it matters less that the before-snapshot is close to Election Day, but the quality of the combined data quickly deteriorates when the after-snapshot is taken too long after the election. Of course, one must wait a few weeks after Election Day for the best results. But even so, given around a late December to early January after-snapshot, ideally, the before-snapshot occurs after August for the best combination that allows less than 100,000 occurrences of false positives and false negatives.

Is there some way to salvage voter file data that is too long after an election has taken place? This is hard to generalize, as it depends on the specific administrative practice of jurisdictions being studied as well as the research question. While we do not fully explore how to model the uncertainty associated with voter file timing, quantifying the potential magnitude of the problem is an important first step.

5.2 Transparency in Observational Research

Our findings also indicate recommendations for scholars using voter files for observational research. At a minimum, researchers must be transparent about the snapshot dates of voter files they use, as well as when the data was secured by the researcher. As an example, in Clinton et al. (2020) exact snapshot dates of the data sources are supplied.[18] In this article, the researchers indicate that they acquired snapshots on Election Day for each year from 2008–2016. As of 2022, North Carolina has same-day registration (SDR) during the early voting period only; registrants cannot register for the first time and vote on Election Day itself (EDR). It is likely that the Clinton et al. (2020) snapshots accurately capture SDR voters, but if North Carolina had an EDR system, these persons could be missed *despite the fact that the registration list was current as of Election Day.* Individuals who registered via EDR would be excluded from Clinton et al. (2020) 's panel framework in the first election where they voted, potentially biasing estimates of the effect of subsequently changing that voter's polling place downward if we think they are especially susceptible to changes to their polling place.

For researchers using voter file data provided by a third-party vendor, where snapshot date may be difficult to acquire, information about when the data was provided to the researcher could give a broad indication of potential bias. However, the onus should be on the researcher to gather information from the vendor about the snapshot date and any potential list maintenance performed by the vendor since the snapshot was generated. Unfortunately, this is rarely documented. In Ansolabehere and Hersh (2012), for instance, the researchers may be overestimating the rate of overreporting by CCES respondents if they are missing a sizeable share of the EDR and SDR population. Since third-party voter file data is often used to conduct analyses that are more costly than single-state studies, forcing scholars using such data to better document their sources may reduce incentives to turn to more expensive yet less transparent datasets.

5.3 Conducting Unbiased Field Operations and Field Experiments

The dynamic nature of voter file data over relatively short time frames also presents a challenge to work entailing contact with registered voters. Recent evidence suggests potential heterogeneity in the effect of get-out-the-vote (GOTV) campaigns, with the CATE for high-propensity voters being substantially higher than for low-propensity voters (Enos et al., 2014). The potential for accretion bias and attrition bias in voter file data suggests an alternative explanation, at least when examining aggregate effects: voter file snapshots taken well before or well after an election exclude a disproportionately low-propensity voter population from the targeting universe *entirely*.

[18] See their Appendix A, which states as follows: "Snapshots of the North Carolina Voter Roll provided by North Carolina State Board of Election (NCSBE) between 2008 and 2016 was downloaded by the authors from the NCSBE data site http://dl.ncsbe.gov/index.html in November of 2017. Data for the 2016 presidential election comes from the Nov 8, 2016 snapshot, data for the 2012 presidential election comes from the Nov 6, 2012 snapshot, and data for the 2008 presidential election comes from the Nov 4, 2008 snapshot."

Experimental research using voter file data may thus produce less generalizable results than a reliance on administrative records would otherwise suggest. For example, if low propensity voters are systematically excluded from voter file-based samples used in experiments, estimates of the CATE on low propensity voters and ATE will both be skewed away from detecting the true population effect. This could occur even if researchers acquire multiple snapshots of the voter file and are careful to keep voters who are dropped between snapshots, since it is not possible to retroactively apply treatment to registrants appearing due to accretion. It may be impossible to capture the full, in-context effect of any treatment on the types of voters who are likely to register on or close to Election Day.

Given the temporal biases we identify, should practitioners operating in the field collect all possible snapshots to keep track of potential voters? Ultimately this may require a trade-off between false positives and false negatives (see Figure 8). That is to say, while expanding the mobilization and persuasion efforts to all registrants that ever appeared on the voter file could be fruitful in catching re-registering voters, but as the lion's share of dropped registrants could be true "deadwood" who will never re-register (e.g., out-of-state movers or deceased registrants), it may also dilute constrained resources. Yet contact could stimulate not only turnout but also registration; our estimates of accretion may be lower than the number of persons who would register if targeted. While not observable with the given data, in a close, contentious election, perusing a list of deleted voters could be decisive.

6 Conclusion

Voter files have become a key tool for academic researchers seeking to understand who votes in the United States. However, analysts considering voter file data for the first time confront a dearth of practical guidance regarding best practices for using this resource. In this paper, we document how the dynamic nature of voter file data may induce bias in both observational and experimental studies. Separating this bias into two components, which we label *accretion bias* and *attrition bias*, we show how a less biased depiction of who votes generally indicates acquisition of a voter file 5–9 weeks *after* Election Day, but even in this instance some potential voters and actual registrants will be incorrectly excluded from the population of interest. Researchers may instead need to acquire time series voter file data that sufficiently account for voter accretion and attrition, such as snapshots taken close to but a few weeks before *and* a few weeks after Election Day.

On a substantive level, our findings suggest extant voter file-based work may provide a distorted picture of disparities *and remedies for disparities* in political participation. Observational analyses relying on single snapshots, or snapshots of unknown timing from third-party vendors, may miss the late registering or readily purged populations that are less likely to vote. Accretion and attrition bias also means that experiments may underestimate the CATE of treatments for low-propensity voters, and field operatives may be missing a population of supporters who are missing from their lists. Given the increased attention to the dynamics of turnout in recent studies of American elections, academics and practitioners alike should seek to better understand the impact of an inherently dynamic data source on their actions and assumptions.

References

Alvarez, R. Michael, Stephen Ansolabehere, and Catherine H. Wilson (2002). Election Day Voter Registration in the United States: How One-step Voting Can Change the Composition of the American Electorate. Working Paper, Caltech/MIT Voting Technology Project.

Alvarez, R. Michael and Jonathan Nagler (2011). Election Day Voter Registration in California. Technical report, Demos.

Ansolabehere, Stephen and Eitan Hersh (2012). Validation: What Big Data Reveal About Survey Misreporting and the Real Electorate. *Political Analysis 20*(4), 437–459.

Ansolabehere, Stephen and Eitan Hersh (2014). Voter Registration: The Process and Quality of Lists. In B. C. Burden and C. Stewart, III (Eds.), *The Measure of American Elections*. Cambridge: Cambridge University Press.

Ansolabehere, Stephen and Brian Schaffner (2015). Beyond the Core and Periphery: A New Look at Voter Participation Across Elections. pp. 21.

Axelrod, Robert (1972). Where the Votes Come From: An Analysis of Electoral Coalitions, 1952-1968. *The American Political Science Review 66*(1), 11–20.

Bai, Jushan (1997). Estimation of a Change Point in Multiple Regression Models. *The Review of Economics and Statistics 79*(4), 551–563.

Bai, Jushan and Pierre Perron (1998). Estimating and Testing Linear Models with Multiple Structural Changes. *Econometrica 66*(1), 47–78.

Bai, Jushan and Pierre Perron (2003). Computation and analysis of multiple structural change models. *Journal of Applied Econometrics 18*(1), 1–22.

Berent, Matthew K., Jon A. Krosnick, and Arthur Lupia (2016). Measuring Voter Registration and Turnout in Surveys. *Public Opinion Quarterly 80*(3), 597–621.

Bernstein, Robert, Anita Chadha, and Robert Montjoy (2001). Overreporting Voting: Why It Happens and Why It Matters. *Public Opinion Quarterly 65*(1), 22–44.

Bhatti, Yosef, Jens Olav Dahlgaard, Jonas Hedegaard Hansen, and Kasper M Hansen (2019). Core and Peripheral Voters: Predictors of Turnout Across Three Types of Elections. *Political Studies 67*(2), 348–366.

Biggers, Daniel R. and Daniel A. Smith (2020). Does threatening their franchise make registered voters more likely to participate? Evidence from an aborted voter purge. *British Journal of Political Science 50*(3), 933–954.

Brater, Jonathan, Kevin Morris, Myrna Pérez, and Christopher Deluzio (2018). Purges: A Growing Threat to the Right to Vote. Technical report, Brennan Center for Justice.

Bryan Cole, J. (2016). Does Same Day Registration Lead to Repeat Customers at the Ballot Box? *Election Law Journal: Rules, Politics, and Policy 15*(4), 271–284.

Burden, Barry C. (2000). Voter Turnout and the National Election Studies. *Political Analysis 8*(4), 389–398.

Cao, Jian, Seo-young Silvia Kim, and R. Michael Alvarez (2022). Bayesian Analysis of State Voter Registration Database Integrity. *Statistics, Politics and Policy 13*(1), 19–40.

Citrin, Jack, Eric Schickler, and John Sides (2003). What if Everyone Voted? Simulating the Impact of Increased Turnout in Senate Elections. *American Journal of Political Science 47*(1), 75–90.

Clausen, Aage R. (1968). Response Validity: Vote Report. *The Public Opinion Quarterly 32*(4), 588–606.

Clinton, Joshua D., Nick Eubank, Adriane Fresh, and Michael E. Shepherd (2020). Polling place changes and political participation: evidence from North Carolina presidential elections, 2008–2016. *Political Science Research and Methods 9*(4), 800–817.

Cooper, Christopher A., Moshe Haspel, and H. Gibbs Knotts (2009). The Value of Voterfiles for U.S. State Politics Research. *State Politics & Policy Quarterly 9*(1), 102–121.

Coppock, Alexander and Donald P. Green (2016). Is Voting Habit Forming? New Evidence from Experiments and Regression Discontinuities. *American Journal of Political Science 60*(4), 1044–1062.

Enos, Ryan D., Anthony Fowler, and Lynn Vavreck (2014). Increasing Inequality: The Effect of GOTV Mobilization on the Composition of the Electorate. *The Journal of Politics 76*(1), 273–288.

Feder, Catalina and Michael G. Miller (2020). Voter Purges After Shelby: Part of Special Symposium on Election Sciences. *American Politics Research 48*(6), 687–692.

Fowler, Anthony (2013). Electoral and Policy Consequences of Voter Turnout: Evidence from Compulsory Voting in Australia. *Quarterly Journal of Political Science 8*(2), 159–182.

Fowler, Anthony (2015). Regular Voters, Marginal Voters and the Electoral Effects of Turnout. *Political Science Research and Methods 3*(2), 205–219.

Fraga, Bernard L. (2018). *The Turnout Gap: Race, Ethnicity, and Political Inequality in a Diversifying America*. Cambridge University Press.

Gerber, Alan S., Donald P. Green, and Ron Shachar (2003). Voting May Be Habit-Forming: Evidence from a Randomized Field Experiment. *American Journal of Political Science 47*, 11.

Green, D. P. (2006). Can Registration-Based Sampling Improve the Accuracy of Midterm Election Forecasts? *Public Opinion Quarterly 70*(2), 197–223.

Green, Donald P. and Alan S. Gerber (2005). Recent Advances in the Science of Voter Mobilization. *The ANNALS of the American Academy of Political and Social Science 601*(1), 6–9.

Hersh, Eitan D. (2015). *Hacking the Electorate: How Campaigns Perceive Voters*. Cambridge University Press.

Huber, Gregory A., Marc Meredith, Michael Morse, and Katie Steele (2021). The racial burden of voter list maintenance errors: Evidence from Wisconsin's supplemental movers poll books. *Science Advances 7*(8).

Igielnik, Ruth, Scott Keeter, Courtney Kennedy, and Bradley Spahn (2018). Commercial Voter Files and the Study of U.S. Politics. Technical report.

Jackman, Simon and Bradley Spahn (2019). Why Does the American National Election Study Overestimate Voter Turnout? *Political Analysis 27*(2), 193–207.

Jackman, Simon and Bradley Spahn (2021). Politically invisible in america. *PS: Political Science and Politics 54*(4), 623–629.

Kauffman, Johnny (2018). Georgia Law Allows Tens Of Thousands To Be Wiped From Voter Rolls. *NPR*.

Kim, Seo-young Silvia (2022). Automatic Voter Re-registration as a Housewarming Gift: Quantifying Causal Effects on Turnout Using Movers. *American Political Science Review*.

Kim, Seo-young Silvia, Spencer Schneider, and R. Michael Alvarez (2020). Evaluating the Quality of Changes in Voter Registration Databases. *American Politics Research 48*(6), 670–676.

Leighley, Jan E. and Jonathan Nagler (2013). *Who Votes Now?: Demographics, Issues, Inequality, and Turnout in the United States*. Princeton University Press.

Lijphart, Arend (1997). Unequal Participation: Democracy's Unresolved Dilemma Presidential Address, American Political Science Association, 1996. *American Political Science Review 91*(1), 1–14. Publisher: Cambridge University Press.

Martinez, Michael D. and Jeff Gill (2005). The Effects of Turnout on Partisan Outcomes in U.S. Presidential Elections 1960–2000. *The Journal of Politics 67*(4), 1248–1274.

McDonald, Michael P. (2007). The True Electorate: A Cross-Validation of Voter Registration Files and Election Survey Demographics. *Public Opinion Quarterly 71*(4), 588–602.

McDonald, Michael P. (2022). Voter turnout. *United States Elections Project* *http://www.electproject.org/home/voter-turnout/voter-turnout-data. Last Accessed on August 26, 2022.*

Morris, Kevin (2021). Turnout and Amendment Four: Mobilizing Eligible Voters Close to Formerly Incarcerated Floridians. *American Political Science Review 115*(3), 805–820.

Morse, Michael (2021). The Future of Felon Disenfranchisement Reform: Evidence from the Campaign to Restore Voting Rights in Florida Democracy Reform Symposium. *California Law Review 109*(3), 1143–1198.

Nyhan, Brendan, Christopher Skovron, and Rocío Titiunik (2017). Differential Registration Bias in Voter File Data: A Sensitivity Analysis Approach. *American Journal of Political Science 61*(3), 744–760.

Ramírez, Ricardo, Romelia Solano, and Bryan Wilcox-Archuleta (2018). Selective recruitment or voter neglect? race, place, and voter mobilization in 2016. *Journal of Race, Ethnicity, and Politics Forthcoming*.

Rosenstone, Steven J. and Raymond E. Wolfinger (1978). The Effect of Registration Laws on Voter Turnout. *American Political Science Review 72*(1), 22–45.

Shino, Enrijeta, Michael D. Martinez, Michael P. McDonald, and Daniel A. Smith (2020). Verifying Voter Registration Records. *American Politics Research*.

Sigelman, Lee, Philip W. Roeder, Malcolm E. Jewell, and Michael A. Baer (1985). Voting and Nonvoting: A Multi-Election Perspective. *American Journal of Political Science 29*(4), 749.

Silver, Brian D., Barbara A. Anderson, and Paul R. Abramson (1986). Who Overreports Voting? *The American Political Science Review 80*(2), 613–624.

Traugott, Michael W. and John P. Katosh (1979). Response Validity in Surveys of Voting Behavior. *The Public Opinion Quarterly 43*(3), 359–377.

Welsh-Huggins, Andrew (2020). Ohio may purge 120K inactive voters from rolls post-election. *The Associated Press*.

Zeileis, Achim, Friedrich Leisch, Kurt Hornik, and Christian Kleiber (2002). strucchange: An R Package for Testing for Structural Change in Linear Regression Models. *Journal of Statistical Software 7*, 1–38.