Data Descriptor

Cast vote records: A database of ballots from the 2020 U.S. Election

Shiro Kuriwaki[1,†,*], Mason Reece[2,†], Samuel Baltz[2], Aleksandra Conevska[3], Joseph R. Loffredo[2], Can E. Mutlu[3], Taran Samarth[1], Kevin E. Acevedo Jetter[2], Zachary Djanogly Garai[2], Kate Murray[2], Shigeo Hirano[4], Jeffrey B. Lewis[5], James M. Snyder, Jr.[3], and Charles H. Stewart, III[2]

[1]Yale University, Institution for Social and Policy Studies, New Haven, CT, 06511, USA
[2]Massachusetts Institute of Technology, Department of Political Science, Cambridge, MA, 02139, USA
[3]Harvard University, Department of Government, Cambridge, MA, 02138, USA
[4]Columbia University, Department of Political Science, New York, NY, 10027, USA
[5]University of California Los Angeles, Department of Political Science, Los Angeles, CA, 90095, USA
[*]Corresponding author: shiro.kuriwaki@yale.edu
[†]These authors contributed equally to this work.

ABSTRACT

Ballots are the core records of elections. Electronic records of actual ballots cast (*cast vote records*) are available to the public in some jurisdictions. However, they have been released in a variety of formats and have not been independently evaluated. Here we introduce a database of cast vote records from the 2020 U.S. general election. We downloaded publicly available unstandardized cast vote records, standardized them into a multi-state database, and extensively compared their totals to certified election results. Our release includes vote records for President, Governor, U.S. Senate and House, and state upper and lower chambers – covering 42.7 million voters in 20 states who voted for more than 2,204 candidates. This database serves as an uniquely granular administrative dataset for studying voting behavior and election administration. Using this data, we show that in battleground states, 1.9 percent of solid Republicans (as defined by their congressional and state legislative voting) in our database split their ticket for Joe Biden, while 1.2 percent of solid Democrats split their ticket for Donald Trump.

Background & Summary

Ballots are the core records of elections. While the *total counts* of ballots for individual candidates are reported regularly[1–4], *individual* ballots are rarely made available.

In paper-based elections, three types of records at the individual ballot level have been available to some researchers and litigators. The first is the actual, marked paper ballot. The second is the electronic scan of the paper ballot, often referred to as a ballot image. The third is an electronic record of the machine's interpretation of that scanned record, called *cast vote records* (CVRs, Figure 1). The National Institute of Standards and Technology describes CVRs as "an electronic record of a voter's selections, with usually one CVR created per sheet (page) of a ballot."[5]. CVRs are not the ultimate basis of an election, but because "election results are [often] produced by tabulating the collection of CVRs"[5], they should directly reproduce vote totals produced by ballot tabulators. Among the three types of records, CVRs lends itself best for analysis.

In this study, we introduce a dataset of CVRs representing 42.7 million voters. The dataset is available on the *Dataverse* at `https://doi.org/10.7910/DVN/PQQ3KV`[6]. Unlike certified election results — typically made available by state election officials — CVRs are rarely centralized at the state level. Instead, CVRs are often produced as a byproduct of the tabulation process conducted at the sub-state level and retained by local election officials.

Following the November 3, 2020 U.S. general election, local election officials saw a surge in requests for cast vote records by anonymous constituents. Some states saw a four to five-

Figure 1. Cast Vote Record Example. An example of an actual ballot image (left) and the authors' representation of the associated cast vote record (right) in Wisconsin. Blank marks are recorded as an undervote.

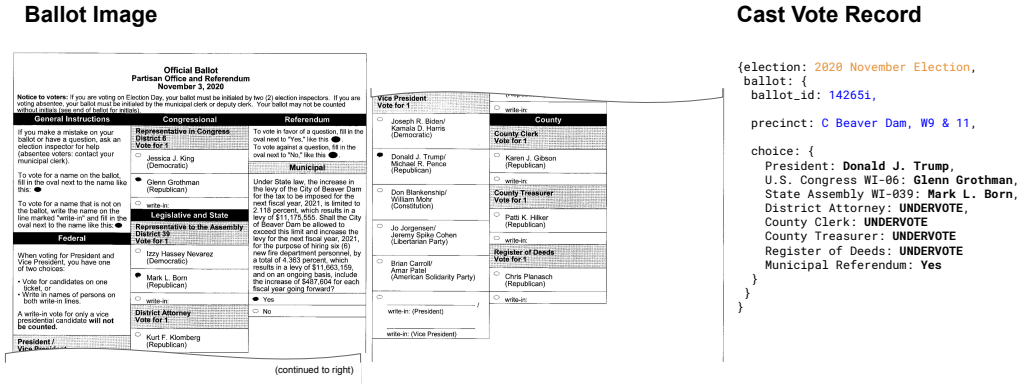

fold increase in records requests between 2020 and 2022[7,8]. Owing to the work of election administrators in responding to these requests, and to O'Donnell for posting those unprocessed records online[9], researchers now have access to an unprecedented quantity of cast vote records. O'Donnell and his collaborators originally collected the CVRs in order to investigate the validity of the 2020 presidential election[10] (to evaluate these claims, we refer readers to work by other researchers[11]). We standardized the publicly available CVRs into a single database, and extensively compared them to official election results.

Our final dataset has two main audiences. The first audience are those in political science, economics, sociology, and other related fields who study electoral behavior. CVR data will allow researchers to study a wide variety of electoral phenomena. In particular, it will allow researchers to measure important aspects of voting behavior much more accurately than is possible using aggregate election returns or surveys. The CVRs are at the individual level and record voters' actual choices across multiple offices. For example, the CVRs include choices for state legislative races. These are rarely included in surveys because researchers typically focus on top-of-the-ticket races. Moreover, survey estimates of vote choices in down-ballot elections can result in especially large measurement errors because respondents are less likely to correctly recall who they supported in down-ballot elections. Such estimates may have especially large sampling errors too, because only a few respondents in a typical nationally representative sample will have participated in any particular down-ballot contest such as those for state legislative seats.

One important electoral phenomenon is split-ticket voting, that is, voting for candidates with different partisan affiliations across offices. For example, aggregate-level election data can tell us how many votes Donald Trump received in a state and how many voted for each of the Republican candidates for offices further down the ballot, but it cannot tell us how many voters who voted for Donald Trump for President *also* voted for Republicans all on the down-ballot offices[12–14]. CVRs allow researchers to measure this type of behavior precisely, and at different levels of government, e.g., national level, state level, and across levels[15–17]. Researchers can also use CVRs to count the number of voters who vote for Democratic candidates and also vote for the progressive position in referendums[18–20].

The data include geographic identifiers: counties and often precincts. One application possible with CVRs is to measure which state legislative or congressional districts have an especially high share of split-ticket voters. To the degree that these voters can be viewed as "swing" or "persuadable" voters, while straight-ticket voters can be viewed as "core" or "loyal" partisan voters, this measure could be used to test theories of resource allocation and campaign strategies[21–23].

Researchers can of course merge the CVR data with other information about the candidates and

contests, such as incumbency status, patterns of campaign spending, and candidate attributes, such as race, ethnicity, gender, ideology, and/or experience[24]. With this extra data in hand, researchers can begin to study a wide variety of phenomena, including which voters split their tickets and in which ways as a function of the available choices. They can also study the degree to which split-ticket voting favors incumbents or the candidate who has a fundraising advantage[17]. Researchers can also merge the CVR data with precinct-level demographic and socio-economic information to explore relationships between split-ticket voting and these types of variables. For example, does ticket splitting vary with the types and amounts of media (e.g., local newspapers) available in an area[25]?

The CVR data could be used to study other electoral phenomena as well. For example, CVRs for ranked choice elections allows researchers to analyze voter's preferences across all candidates[26,27]. Or, consider roll-off, where voters cast ballots but choose to abstain in particular races (also called undervoting). Using the CVRs, researchers can measure this behavior accurately. They can then investigate the types of races where this behavior is more prevalent, and the types of voters (e.g., straight-ticket Democrats, straight-ticket Republicans, or ticket-splitters) that are more likely to roll-off, as well as the interaction between contest and voter types. One final example is voting for candidates of minor parties, such as the Libertarian Party and the Green Party. Since the CVR database has millions of records, it contains many thousands of records of individuals who supported minor parties that received a small fraction of the vote[28]. The typical survey is less useful for studying this type of behavior due to sample size.

The second audience are those in election law, forensics, and administration, who seek to study the integrity of the electoral process[29,30]. Cast vote records are of interest to scholars of attitudes towards election integrity because voters' mistrust of the vote counting process, when it exists, often revolves around the ballots[31–33]. Ballot-level data have been used in studies of election administration to help explain seemingly surprising election results[34,35]. Our dataset can also contribute to debates around the tradeoff between transparency and privacy[36–39].

Finally, our data has implications that go beyond the particular states in this release, or even the 2020 election in particular. The November 2020 election was a turning point in U.S. politics, where the administration of elections became an overtly partisan issue. The conduct and administration of any other future elections now risks being politicized. Self-enforcing the legitimacy of elections requires election administrators, data scientists, and social scientists to work together upon a common understanding of election technology. We hope that our data release serves as a standard for future work in this area of growing interest.

Methods

Out of 3143 counties in the U.S., `votedatabase.com` releases data from 464 individual counties and 3 statewide data sets (Alaska, Delaware, and Rhode Island).As described by O'Donnell[9], this website provides access to the raw electronic file containing CVRs that were acquired by numerous citizens via open records requests. Of the 467 files, we validate and release 362 counties in 20 states. We additionally modify precinct information from 5391 vote choices in the data (0.003%) for privacy protection. Privacy and our technical validation are discussed in more detail below.

We approached the raw data cautiously. CVRs may not include all of the ballots cast in an election for a number of reasons. Even within a single county, it is possible that some valid ballots are cast and tabulated in a way that creates CVRs while others are not. Sometimes ballots that are held aside for manual adjudication, such as provisional or damaged ballots, are not scanned through tabulators, and therefore a cast vote record is not created for such ballots. We can increase our confidence that CVR files posted are complete, uncorrupted, and genuine by comparing vote tallies produced using the downloaded CVR data to the official reports of vote totals from the same jurisdictions. Note that our goal is *not* an audit of the election. CVRs are neither sufficient nor necessary to rule out election malfeasance. While CVRs are convenient representations of the paper ballot, they are simply not intended to be the official results of an election. On the other hand,

even if the cast vote record were complete, ballots could have been tampered prior to tabulation.

We start with data that O'Donnell and his collaborators obtained and subsequently uploaded to `votedatabase.com`[9]. O'Donnell reports that citizens requested CVRs according to their state's open record law guidelines, noting that all records were "obtained through these valid public records requests"[9]. He reports that requests were sent to "nearly all counties in all states," but 23 states responded that they did not have any records relevant to the request that could be provided under the state's open records law (Alabama, Connecticut, Hawaii, Indiana, Kansas, Louisiana, Maine, Massachusetts, Mississippi, Missouri, Montana, Nebraska, New Hampshire, New York, North Carolina, North Dakota, Oklahoma, South Carolina, South Dakota, Utah, Virginia, Washington, and Wyoming). Indeed, the availability of CVRs is limited by state law and executive order, with states including North Carolina, South Carolina, and New York shielding the CVR from open records requests[39]. In all, the CVRs in the data presented here originate from counties in 27 states and D.C. that were collected and made available by O'Donnell[9]. The database contains data from swing states like Wisconsin, Michigan, and Georgia, as well as solidly Democratic states, including New Jersey, and solidly Republican states, including Texas.

Our standardization and validation proceeded in five steps. First, we downloaded the data files and standardized them so that their values were comparable across different voting machines and jurisdictions. While the website does provide some normalized versions of the raw data, we have chosen not to rely on these and instead independently process the raw data. We only considered files that were clearly CVRs in machine-readable form covering multiple offices and more than a handful of precincts. Standardization here entails that we identify the party affiliation of each candidate, code invalid votes consistently across jurisdictions, and standardize the formatting of candidate names. In the initial phase of the study, two groups conducted these pursuits independently, without awareness of each other's work. This gave us nearly independent measures of inter-coder reliability and reduced the possibility that a single coding error propagated to all counties.

About 15 percent of the counties had CVRs in non-rectangular formats such as JSON or XML. The remaining files were tabular files such as CSV or Microsoft Excel. Formats differed by the vendor of the machine (the main three being Dominion, Hart Intercivic, and ES&S), and the make of each machine. For information about the particular machine used in each county, see `https://verifiedvoting.org/verifier`. We parsed these data with our own R and Python scripts, eventually normalizing all data into a long format where one row represents a single vote choice by a voter. This article releases the full codebase we developed.

Second, we assigned a cast vote record identifier to each voter, within a county or jurisdiction. Most times, this number indicates an individual voter – one anonymous voter gets one identifier for all their choices. In about 30 counties with ballots spanning multiple double-sided pages, each page was separated before it was scanned. (Note that from the administrator's perspective, once the ballot itself is deemed valid, the identity of the voter is irrelevant in the counting process.) In about 20 of these counties, we used metadata to link pages into a single ID (see Supplementary Information B). In the remaining 10, the records were irreversibly separated. This includes several counties in California (Los Angeles, San Francisco, San Bernardino, Ventura Counties). However, in the remaining counties, there is a one-to-one correspondence between the ID we assign and a single voter. Even in many counties with long ballots such as Maricopa, Arizona (with around 60 contests per ballot), the ballot record for each individual voter is preserved.

Third, we extensively checked the CVRs against other official sources of data, mainly the MIT Election Data and Science Lab (MEDSL) 2020 precinct-level returns[3,40]. This was an ideal dataset to use as validation because it is at the precinct-level, it has standardized its formatting across states, it includes district-level as well as statewide contests, and features extensive documentation. We limited our attention to six offices: U.S. President, Governor, U.S. Senate, U.S. House, State Senate, and State House. The CVRs include votes for many other offices, including local administrative offices, school boards, and referendums. Other work analyzes these offices[15,16] but we exclude them in this dataset because we lack fully standardized official data to extensively validate against.

Name	Description
`state`	The name of the state that the ballot is from
`county_name`	The name of the county the ballot is from
`cvr_id`	A unique ID given to each ballot within a state-county
`precinct_medsl`	One field that indicates the precinct that can be matched to Baltz *et al.* (2022).[3]
`precinct_cvr`	Concatenated precinct values as recorded in the raw data, for each `precinct_medsl`.
`office`	The name of the office the voter is choosing a candidate in
`district`	The district of the office the voter is choosing from
`candidate`	The name of the candidate the voter has selected in the office-district
`party`	A simplified name of the party of the candidate on the ballot
`party_detailed`	The full name of the party of the candidate on the ballot
`magnitude`	The number of candidates a voter could have chosen in this particular contest

Table 1. Dataset Variables

After attempting to find the best matching official result for all counties, we only release counties with candidate-level discrepancies of 1 percent or less at the candidate-county level. The section "Technical Validation" discusses the process in more depth.

Fifth, we extracted precinct information in the cast vote records and standardized them to match the MIT Election Data and Science Lab database of standardized precinct names. Precinct identifiers were often either names (e.g., "City of Madison Ward 1") or numeric codes (e.g., Precinct 35-001). These classifications are known to vary widely across jurisdictions and machines, with no national standard. We used fuzzy string matching and triangulation of vote counts to link the precincts, which we detail in Supplementary Information A. There are 312 counties which we matched to the precinct level database.

Data Records

Our dataset includes 166,470,734 rows, with each row indicating a choice of a voter for a given contest. Our data is deposited in Dataverse at `https://doi.org/10.7910/DVN/PQQ3KV`[6].

Variables Table 1 describes the variables in the data. Voters are uniquely identified by the state, county name, and `cvr_id` assigned by ourselves, with the exception of unmerged fragmented ballots described in the methods section. Contests are uniquely identified by the state, office, and district. The names and values for our variables generally follow the naming conventions of the MIT Election Data and Science Lab[3].

The `candidate` value is the name of the candidate that ballot was cast for, or it can be an undervote (`UNDERVOTE`), overvote (`OVERVOTE`), or write-in (`WRITEIN`). Undervotes refer to a blank choice for that contest or mark that the tabulator could not interpret, and overvotes refer to a voter marking more candidates than they could vote for in a given contest. Even though both types of votes are invalid, undervoting in particular is seen as a form of contest-specific abstention and is of interest to election scholars.

In all offices but the State House in Arizona and West Virginia, the contests represented here are single-choice elections (a `magnitude` of one). Arizona voters can cast up to 2 valid votes to elect their Representatives for the state's lower chamber, with 2 winners per district. Until 2020, West Virginia voters in some districts could cast up to 2 or 3 valid votes for their lower chamber.

Third-party candidates and write-in candidates rarely win themselves, but the ideological orientation of voters who vote for them is of interest to researchers[28]. When a candidate is listed on the ballot with a registered party, they are listed with their party affiliation in the `party_detailed` variable (such as Libertarian, Green, or "No Party Affiliation" in Florida). The `party_detailed` variable is set to missing for undervotes and overvotes. Jurisdictions vary in whether write-ins are reported separately or grouped together simply as write-in votes. Ballot access for third parties

also varies by state. When the candidate is not listed on the ballot with a party, we record them as write-in candidates with no party affiliation (For more details, see Supplementary Information C).

Our dataset also records the precinct of the cast vote record, as discussed in the Methods section. We provide two variables for precincts (Table 1). `precinct_medsl` is our matched precinct name, formatted to correspond exactly with those described by Baltz *et al.*[3]. `precinct_cvr` is a concatenation of the original precinct or precinct portion name as it appears in the cast vote record by the pipe character `|`. The concatenation occurs for every set of precincts for a given `precinct_medsl`. For example, if precinct `001` as defined in the MEDSL precinct data[3] contains two precinct portions `001A` and `001B`, and the cast vote records record the precinct portion, we give all voters in precinct `001` the value of `001A | 001B`.

Privacy Considerations For certain vote records, we further aggregate the values of `precinct_cvr` and `precinct_medsel` by concatenating additional precinct portions and by concatenating two more MEDSL precincts. In those records, the `precinct_cvr` and `precinct_medsl` fields indicate that the vote was cast in one of the listed precincts, but not which one. We apply this additional aggregation to avoid facilitating the linkage of any particular voter to a vote that they cast.

To understand how highly-disaggregated election results can reveal the vote choices of particular voters, suppose there is a precinct in which only 3 voters shared a particular ballot style. *Who* voted in U.S. elections is public via lists of voter rolls produced by election officials. Those voter rolls contain sufficient geographic information to infer the identity of each of these voters from their precinct and ballot style. Here we refer to the ballot style as the subset of contests that the voter could vote on that are among the six we report. If all of the voters in that precinct who used that ballot style supported the same candidate, the vote choice for those 3 voters would be revealed by the CVR data to anyone who has access to an accurate voter list[39]. While the release of CVRs might seem to increase opportunities for unraveling the secret ballot, nearly all of the information that can be used to reveal vote choices from the CVR data is already contained in the official precinct aggregates that counties commonly made public[39].

If we reported the `precinct_cvr` and `precinct_medsl` fields without additional aggregation, as many as 5,391 votes choices out of 166.7 million vote choices (0.003%) contained in the database are considered *revealed* as defined by Kuriwaki, Lewis, and Morse[39], i.e., can theoretically be linked to the voters who cast them, using the precinct and ballot style information contained in the database. We note that each of these 5,391 linkages is only theoretical because of practical data limitations that can sometimes make it impossible to infer the ballot style from the CVR or to construct a complete and accurate enumeration of every voter in a precinct who used a particular ballot style.

To address this potential linkage of votes to voters, we first compiled a list of precinct and ballot style combinations in which every voter supported the same candidate in at least one contest. In each of those instances, knowing who voted in the precinct using that ballot style is sufficient to learn one or more of the vote choices made by each of those voters. To avoid this revelation, we then pooled those revealed voters with voters from another precinct in the same county who used the same ballot style. In particular, each revealed precinct-ballot style voter combination was concatenated with the smallest precinct-ballot style voter combination in the same county that shared the same ballot style. For example, suppose the voters using ballot style `A` in precinct `001` had their votes revealed because all of them supported Donald Trump for President, and suppose precinct `002` had the smallest number of voters in the county among all non-revealed precinct-ballot style combinations using ballot style `A`. Then, we report the precinct value for voters using style `A` in both precincts as `001 | 002`. Note that for other ballot styles used in precincts `001` and `002` (if any), there may be no vote revelations. For this reason, the database may contain records in which precincts `001` and `002` are concatenated (those associated with voters using ballot style `A`) and others in which they appear alone (those associated with voters using other ballot styles).

Geographic Coverage A total of 20 states are covered by our release. Table 5 lists all counties in our released dataset. Nevada, Rhode Island, and Delaware had all its counties released, while in

	Percent Biden			Biden, Trump, and Jorgensen Voters		
State	CVR	Pop.	Diff.	CVR	Pop.	%
Arizona	52.7	50.2	+3	2,677,978	3,385,294	79
California	65.3	64.9	0	14,245,273	17,305,067	82
Colorado	54.5	56.9	-2	2,640,797	3,221,419	82
Delaware	59.6	59.6	0	500,227	501,871	100
Florida	50.4	48.3	+2	5,694,485	11,036,075	52
Georgia	48.8	50.1	-1	3,623,410	4,998,482	72
Illinois	41.4	58.7	-17	247,034	5,985,350	4
Iowa	33.0	45.8	-13	11,191	1,676,370	1
Maryland	64.9	67.0	-2	2,382,347	2,994,925	80
Michigan	59.3	51.4	+8	400,193	5,511,059	7
Nevada	51.2	51.2	0	1,387,901	1,387,957	100
New Jersey	61.5	58.1	+3	2,209,816	4,523,390	49
Ohio	43.7	45.9	-2	1,917,001	5,901,568	32
Oregon	51.4	58.3	-7	938,860	2,340,413	40
Rhode Island	60.6	60.6	0	512,238	512,461	100
Tennessee	31.5	38.2	-7	212,808	2,998,363	7
Texas	42.5	47.2	-5	1,787,191	11,272,753	16
Utah	46.8	39.3	+8	6,747	1,463,868	<1
West Virginia	27.0	30.2	-3	49,261	792,053	6
Wisconsin	43.0	50.3	-7	520,911	3,279,229	16
All 20 States	56.6	53.4	+3	41,965,669	91,087,967	46
All 50 States	—	52.2	—	—	155,291,327	—

Table 2. Comparison of Data Coverage to Entire State. *The first set of columns compares the two-party Biden vote share in our data (CVR) and the entire state (Pop.). The second set of columns shows the total number of Biden, Trump, and Jorgensen votes.*

other states, only a portion of the state's counties is released. Table 2 compares the characteristics of our collected samples with the entire state. A convenient cross-state metric to indicate the representativeness of our release is the percentage of voters that vote for the Democratic presidential candidate, Joseph R. Biden. Table 2 compares the percentage of Biden voters as a share of Biden and Trump voters in the dataset and the state(s) as a whole. Overall, 56.6 percent of our collection's Presidential voters are Biden voters (as a percentage of the two-party vote), while his two-party vote share in all 50 states was 52.2 percent[41].

Table 3 compares the set of counties based on demographics. We take demographic data from the 2020 decennial Census and compare our counties with the entire United States. The Table reports the average and median value of each demographic variable, as well as the overall (population-weighted) value and the standard deviation. The population in our set of counties does not differ from the average U.S. county in terms of age or homeownership, but our counties are less White (52.8% vs. 61.2% nationwide), more Hispanic, and more urban.

Technical Validation

We extensively analyzed each county's files and compared them with official results, as described in the previous section.

Validation In general terms, we define a discrepancy as any difference between the summed total of CVRs cast for candidates in the CVRs we processed from a county and the certified results published by that county. More precisely, we compute the discrepancy in a county k by the

	Overall		Average		Median		Standard Dev.	
	CVR	Nation	CVR	Nation	CVR	Nation	CVR	Nation
Percent White	52.8	61.2	71.3	75.3	74.4	82.1	17.0	19.9
Percent Black	11.7	12.3	10.0	8.7	3.8	2.2	13.4	14.0
Percent Hispanic	27.5	19.5	14.8	11.9	9.6	4.8	14.7	19.2
Percent Under 18	22.0	22.0	21.7	21.9	22.0	21.9	3.4	3.4
Percent Over 65	16.3	16.9	19.4	20.2	18.5	19.9	5.3	4.7
Percent Urban	89.4	80.1	51.3	37.1	56.4	34.8	34.3	34.2
Percent Homeowning	22.2	24.2	27.0	28.5	27.0	29.0	4.7	4.4

Table 3. Characteristics of Counties Used. *Comparison of the counties in our sample (CVR) with all counties in the United States (Nation). Data come from the 2020 Decennial Census.*

percentage

$$\text{discrepancy}_k = \max_{j \in \mathbb{I}_k} \left\{ \left| \frac{N_j^c - N_j^v}{N_j^v} \right| \right\}, \tag{1}$$

where N^c is the CVR vote count, N^v is the official vote count, j indexes candidates in a county, $|\cdot|$ is the absolute value function, and $\mathbb{I}_k$ is the set of all Republican, Democratic, and Libertarian candidates in the six offices in county k. We limited our validation to these three parties because affiliations for other choices can be reported in different ways by the county or voting machine, leading to mismatches even when the count is correct. For example, all our released counties (except for Santa Cruz, Arizona and Cumberland, New Jersey) include undervotes, but some counties do not report the number of undervotes in their official reports so we cannot verify all these numbers. We relied on MEDSL's precinct data[3] as our official vote count, and complemented it with official statements of the vote from counties directly where necessary.

We then release counties where the discrepancy is 1 percent or less. Figure 2 shows the distribution of discrepancies. In panel 2a, we grouped the discrepancy rates in six categories, showing that 258 counties had no discrepancy across all six offices, and 104 had a candidate with non-zero discrepancies within 1 percent. Panel 2b shows the candidate-level discrepancies, focusing on the range of -5 to 1 percent. We see that almost all the discrepancies are within 0.2 percent (and most are exactly zero). Some data points with 0 discrepancy are nevertheless not released and shown in light gray. This is because those counties have zero discrepancies in some offices but a discrepancy larger than 1 percent in other offices. Panel 2c shows the correlates of discrepancy by population. The discrepancies tend to be dispersed across large and small jurisdictions.

In addition to computing discrepancies at the county-level, we also conducted a precinct-level validation. Our matching procedure described previously produced 23,139 precincts in 312 counties that could be matched to MEDSL's standardized precinct names[3]. Figure 3 shows the alignment between the total votes for Presidential candidates in the CVR, per precinct, and the independently reported[3] votes in the precinct we matched to. The match is not perfect due to the inclusion of counties with up to 1 percent discrepancies, and possible incorrect matches. Nevertheless, 20,139 precincts (out of 23,139 assigned) matched exactly, and 21,773 matched within 3 votes.

Reasons for Discrepancy In the rare cases of discrepancies, we conducted an extensive, county-by-county investigation to reconcile the numbers, using state and county certification reports. Our public Github repository's *Issue* feature documents our diagnosis for each problematic county that contains up-to-date diagnoses (`https://github.com/kuriwaki/cvr_harvard-mit_scripts/issues`). We found that the discrepancies tended to fall in one one of several categories.

Figure 2. Discrepancy Rates. *(a) number of counties by discrepancy. (b) distribution of discrepancies at the candidate level, limited to the neighborhood of 1 percent. (c) relationship between differences in total votes.*

a

	Released	N
0 discrepancy	1	257
any <1% discrepancy	1	102
any < 5% discrepancy	0	28
any < 10% discrepancy	0	9
any > 10% discrepancy	0	44
candidate missing	0	22

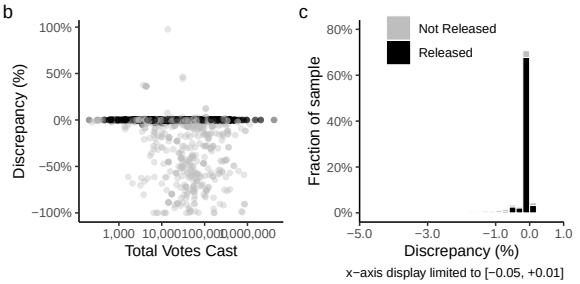

Figure 3. Precinct Level Validation. *Comparisons of vote totals for each Presidential candidate at the precinct level, with those from a precinct-level database[3] on the x-axis and those from our cast vote record database (with approximately matched precinct) on the y-axis. Axes are shown on a square root scale.*

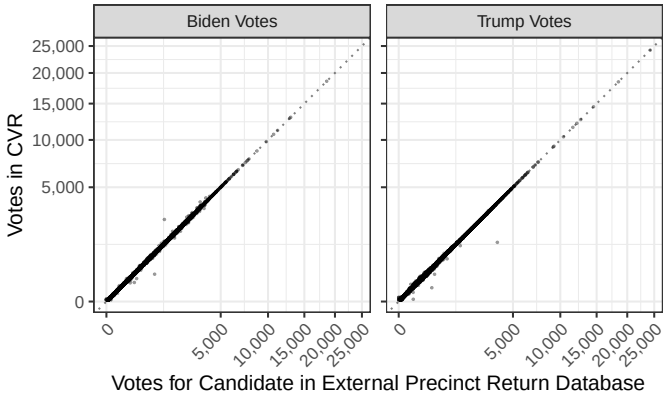

First, in some counties, entire precincts were missing from the cast vote record data. This can happen if a county chooses to hand-count a batch of its ballots, or if the precinct's ballots are processed differently.

Second, in some counties the cast vote record data did not include votes cast by certain methods (vote by mail, in-person, provisional). For example, in Santa Rosa County, Florida and Cuyahoga County, Ohio, it appears that the mailed-in votes have not been included in the cast vote record export uploaded to O'Donnell's[9] website. We verified this by comparing our counts with the county's published vote totals broken out by vote method. This may occur if mailed and in-person votes are handled by a different tabulator.

Another type of known discrepancy is due to redactions done by counties to protect the integrity of the secret ballot, as we previewed in the Data Description section. Thirteen counties in Colorado and California provided CVRs where the vote choices of ballots in small ballot styles were removed. This practice is one way in which jurisdictions have tried to balance their dual roles of transparency and privacy. However, perfect redaction is known to be difficult because counties also need to report the total number of votes cast with perfect fidelity, and the redacted information can be backed out by triangulating the total election results. We do not attempt to unredact the CVRs.

There were many other counties with smaller discrepancies that could not be resolved with publicly available data. We believe one possible explanation for these minute discrepancies is the designation of provisional and disputed votes and ballot that were hand-tabulated. Our cast vote record may either exclude or include votes that were disputed but were later counted towards a major party candidate in the certified tally. For example, Dodge County, Wisconsin's website cautions that

> "Cast Vote Record (CVR) Reports are unofficial results from election night. These are the results the voting equipment tabulated on Election Day. The final, official canvass results posted on the Wisconsin Elections Commission's website for any state/federal races also include counted provisional ballots and other small adjustments. These adjustments are not tallied by, or in, the voting equipment, [but] rather through the County Board of Canvass process. The Cast Vote Record (CVR) Reports contain all data fields available in the ES&S Election Software. Also, please note that if a Municipal Clerk has accidentally corrupted their election data after printing their results tapes and electronically transferring the results into the County for a specific election, that data will not be able to be archived and therefore, would have no ballots to be read and included in the CVR Report." (`https://bit.ly/46eISNX`, accessed July 1, 2024)

While the dispute resolution process is publicly recorded on video in many counties, we cannot link each resolved ballot to a record in our database.

The case of Dane County, Wisconsin illustrates how several of these discrepancies can manifest. Figure 4 shows a precinct map of the county, which includes the City of Madison, using shapefiles[4] and each precinct colored by type of discrepancy compared to official reported results. Two precincts were missing from the Dane CVR because they were held by the neighboring county. Wisconsin cities and villages, which conduct their own election tabulation, may be located near a county border and straddle two counties. In three other instances, we found that a cast vote record from a town actually belonged to the jurisdiction of a different county. We learned that in two other precincts in Dane County, scanning machine failures prompted a hand-count of the votes, with no ballot image or cast vote record present. In 25 other precincts, the cast vote record was only 1 vote short of the reported election result. We contacted the Dane County clerk to resolve these issues, and obtained a letter written by the clerk's office in 2021 that provided these explanations for the discrepancies[42].

We chose to limit our inquiry to county clerks to a minimum, given that many are already at or above capacity in their day-to-day duties[43]. We believe it is unlikely that the original collectors of

Figure 4. Precinct-level Validation. An example from precinct-level matches in Dane County, Wisconsin

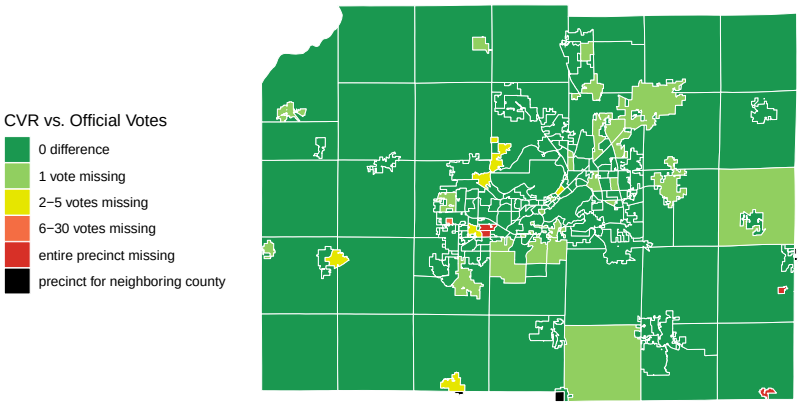

the data tampered with these remaining cases before posting because discrepancies are small and inconsequential for the apparent winner of the contest.

Usage Notes

In the remaining section, we illustrate how users can read in the data and analyze it. Although we use R as the running example, Python or any database-friendly programming language can read the dataset. We conclude with an example that studies the party loyalty of Republican and Democratic voters in their choice for President.

Reading in the Data We store our dataset in a parquet format. Parquet is a modern file storage format optimized for querying large datasets. It is partitioned by grouping variables, and it is columnar (so that users do not need to read in an entire row to extract a value from one column). Our dataset is prohibitively large to read and write in a plain-text format (20 Gb), but is compact and easy to read from in parquet (700 Mb). In R, we use the `arrow` package to query parquet files. For more information on how to read and write parquet files in R, see https://r4ds.hadley.nz/arrow[44]. Parquet is also designed for usage in Python (https://arrow.apache.org/docs/python/parquet.html) and several other programming languages.

The following command opens the dataset.

```
library(tidyverse)
library(arrow)

ds <- open_dataset("cvrs")
```

Here, `"cvrs"` indicates the path to the top-level folder containing the parquet files downloaded from Dataverse. Our data is organized by county, nested within states. After unzipping the zip file or by previewing on Dataverse, we see that `cvrs` has the following structure:

```
├── state=ARIZONA
│   ├── county_name=MARICOPA
│   │   └── part-0.parquet
│   ├── county_name=PIMA
│   │   └── part-0.parquet
```

```
│   ├── county_name=SANTA%20CRUZ
│   │   └── part-0.parquet
│   └── county_name=YUMA
│       └── part-0.parquet
...
── state=UTAH
│   └── county_name=SAN%20JUAN
│       └── part-0.parquet
└── state=WISCONSIN
    ├── county_name=BROWN
    │   └── part-0.parquet
    ├── county_name=KENOSHA
    │   └── part-0.parquet
    ├── county_name=PIERCE
    │   └── part-0.parquet
    └── county_name=WAUKESHA
        └── part-0.parquet
```

Because parquet is columnar, users will find it much faster to produce summary statistics of the data. Even though the code below counts some 166 million rows, it performs the count in one second on a personal laptop.

```
ds |> count(office) |> collect()
```

```
# A tibble: 6 x 2
  office              n
  <chr>           <int>
1 US PRESIDENT 42710448
2 US SENATE    19667514
3 US HOUSE     42613049
4 STATE SENATE 22146498
5 STATE HOUSE  38770601
6 GOVERNOR       562624
```

To perform this count, we used `count()` from `dplyr`, which totals the number of occurrences of each unique value in our `office` variable. We make use of R's pipe operator, `|>`, to pass our data objects forward onto subsequent operations we want to perform.

Finally, we must use the `collect()` command from `arrow` to extract the summary. All transformations before `count()` are *lazily-loaded*, meaning that they are not executed until needed. The arrow program combines the transformations internally in a way that avoids duplicative operations.

Extracting Summaries Users should use the combination of `state`, `office`, and `party` variable to identify candidates. The code below first limits to vote choices for President in Wisconsin ballots using the `filter()` command, and counts the number of records for each candidate-party collection, sorted from most frequent to least.

```
ds |>
  filter(state == "WISCONSIN", office == "US PRESIDENT") |>
  count(candidate, party, sort = TRUE) |>
  collect()
```

```
# A tibble: 8 x 3
  candidate       party      n
  <chr>           <chr>  <int>
1 DONALD J TRUMP  REP   293283
2 JOSEPH R BIDEN  DEM   221356
3 JO JORGENSEN    LBT     6272
4 WRITEIN         <NA>    1412
5 UNDERVOTE       <NA>    1337
6 BRIAN T CARROLL OTH      874
7 DON BLANKENSHIP OTH      724
8 OVERVOTE        <NA>     493
```

For individual voters, use the `cvr_id` variable within a state and county. This ID is a numeric variable that is defined within counties. These numbers do not in any way indicate the time in which the ballot was cast, or the personal identity of the voter. The following code extracts the vote from the voter marked with the `cvr_id` of 1.

```
ds |>
  filter(state == "ARIZONA", county_name == "MARICOPA") |>
  filter(cvr_id == 1) |>
  select(county_name, cvr_id, office, district, candidate, party) |>
  collect()
```

```
# A tibble: 6 x 6
  county_name cvr_id office       district candidate        party
  <chr>        <int> <chr>        <chr>    <chr>            <chr>
1 MARICOPA         1 US PRESIDENT FEDERAL  JOSEPH R BIDEN   DEM
2 MARICOPA         1 US SENATE    ARIZONA  KELLY MARK       DEM
3 MARICOPA         1 US HOUSE     008      MUSCATO MICHAEL  DEM
4 MARICOPA         1 STATE SENATE 013      KERR SINE        REP
5 MARICOPA         1 STATE HOUSE  013      DUNN TIMOTHY TIM REP
6 MARICOPA         1 STATE HOUSE  013      SANDOVAL MARIANA DEM
```

This example shows that this voter split their ticket, voting for Democrats in the Presidential and Congressional race, while voting for one Republican candidate in State Senate. However, further investigation into this voter's State Senate district shows that it was uncontested. That is, with the following query,

```
ds |>
  filter(state == "ARIZONA", office == "STATE SENATE", district == "013") |>
  count(candidate, party_detailed) |>
  collect()
```

```
# A tibble: 7 x 3
  candidate       party_detailed     n
  <chr>           <chr>          <int>
1 "KERR SINE"     REPUBLICAN     93388
2 "NOT QUALIFIED" <NA>            1852
3 "UNDERVOTE"     <NA>           34391
4 "WRITEIN"       WRITEIN          531
5 ""              <NA>             119
6 "BACKUS BRENT"  <NA>             145
7 "OVERVOTE"      <NA>              17
```

we see that none of the ballots in State Senate district 13 were for a Democrat candidate, indicating that no Democrat ran in this district.

Application: Biden and Trump's Party Loyalty As our main exercise, we ask whether partisans — defined by their votes for Congress and state legislature — vote for their party's presidential candidate. Donald Trump was a polarizing candidate. Election observers have wondered if Trump drew less support from Republican-leaning voters compared to the support his opponent, Joe Biden, drew from Democratic-leaning voters. Some referred to these group of voters as "Never Trump Republicans."

For this analysis, we study the counties in five battleground states which together decided the election: Wisconsin, Michigan, Georgia, Arizona, and Nevada. Biden won Georgia, Arizona, and Wisconsin by less than a percentage point, and won Nevada and Michigan by less than 3 percentage points.

```
ds_states <- ds |>
  filter(state %in% c("WISCONSIN", "MICHIGAN", "GEORGIA", "ARIZONA", "NEVADA"))
```

Recall that aggregate election results report how many votes Biden and Trump received, but they do not reveal which of those votes came from Republican and Democratic voters. Only cast vote records can classify individual voters into partisan types based on how they voted in multiple contests.

We first need to narrow down our data so that we only use voter-contest pairs in contests contested by a Democrat and a Republican. In other words, the voter needed to have a choice to vote for a Republican or Democrat.

```
ds_contested <- ds_states |>
  collect() |>
  # Contested contests
  filter(any(party == "REP") & any(party == "DEM"),
         .by = c(state, office, district)) |>
  # Ballots with Presidential vote
  filter(any(office == "US PRESIDENT"),
         .by = c(state, county_name, cvr_id))
```

The first `filter()` command in this code limits to vote choices for contested offices. For each state-office-district combination, we examine if there are any Republican candidates *and* any Democrats. Contests that do not meet this criteria are dropped. The second `filter()` command limits to ballots with a Presidential choice. This excludes fragmented ballots where the President and the rest of the ballot is separated. Both commands are done after `collect()` because the `arrow` package does not support group-specific filter commands as of version 16.1.0.

We now construct a dataset where each row is a single voter. We first create a dataset of Presidential votes:

```
## Voters based on President
ds_pres <- ds_contested |>
  filter(office == "US PRESIDENT") |>
  select(
    state, county_name,
    cvr_id, candidate,
    pres_party = party) |>
  mutate(pres = case_when(
    pres_party == "REP" ~ "Trump",
    pres_party == "DEM" ~ "Biden",
```

```
    pres_party == "LBT" ~ "Libertarian",
    candidate == "UNDERVOTE" ~ "Undervote",
    .default = "Other"))
```

Separately, we construct a dataset that classifies the same voters based on their non-Presidential vote choice. The variable `nonpres_party` is `Down-ballot Democrat` if the voter only votes for Democrats down-ballot (using the `all()` command) and it is `Down-ballot Republican` if the voter only votes for Republicans down-ballot.

```
## subset to all-Dem voters based on everything except President
ds_D <- ds_contested |>
  filter(office != "US PRESIDENT") |>
  filter(all(party == "DEM"), .by = c(state, county_name, cvr_id)) |>
  distinct(state, county_name, cvr_id) |>
  mutate(nonpres_party = "Down-ballot Democrat")

## same subset, but for all-Rep voters
ds_R <- ds_contested |>
  filter(office != "US PRESIDENT") |>
  filter(all(party == "REP"), .by = c(state, county_name, cvr_id)) |>
  distinct(state, county_name, cvr_id) |>
  mutate(nonpres_party = "Down-ballot Republican")
```

Now we join voter's choices for President with their down-ballot choices. Because each row is now a single voter, we join one-to-one using `state`, `county_name`, and `cvr_id`. Voters who were not classified into Democrats or Republican, are, by construction, those who voted for some Democratic down-ballot candidates and Republican down-ballot candidates, or undervoted in some of these offices. We label them `Mixed`.

```
ds_analysis <- ds_pres |>
  left_join(
    bind_rows(ds_D, ds_R),
    by = c("state", "county_name", "cvr_id"), relationship = "one-to-one") |>
  mutate(nonpres_party = replace_na(nonpres_party, "Mixed"))
```

Finally, we construct a cross-tabulation of this dataset using the base-R `xtabs()` function.

```
xtabs(~ nonpres_party + pres, ds_analysis) |>
  addmargins()
```

```
                      pres
nonpres_party            Biden Libertarian   Other   Trump Undervote     Sum
  Down-ballot Democrat   3430497       15087   12012   43239      4478 3505313
  Down-ballot Republican   70136       26723   12396 3540120      9166 3658541
  Mixed                   793905       71431   21379  619255     16688 1522658
  Sum                    4294538      113241   45787 4202614     30332 8686512
```

This table shows for example that among 3,505,313 solidly Democratic voters, 3,430,497 voted for Joe Biden. We can show cell counts in terms of proportions of the entire row, with the following base-R operation:

```r
xtprop <- xtabs(~ nonpres_party + pres, ds_analysis) |>
  prop.table(margin = 1) |>
  round(3)

## add margins
N <- xtabs(~ nonpres_party, ds_analysis) |>
  format(big.mark = ",")

## reorder columns and append totals
xtprop[, c("Biden", "Trump", "Libertarian", "Undervote")] |>
  cbind(N) |>
  kableExtra::kbl(format = "latex", booktabs = TRUE)
```

	Biden	Trump	Libertarian	Undervote	N
Down-ballot Democrat	0.979	0.012	0.004	0.001	3,505,313
Down-ballot Republican	0.019	0.968	0.007	0.003	3,658,541
Mixed	0.521	0.407	0.047	0.011	1,522,658

Table 4. Party Loyalty in Five Battle Ground States

This formatted table shows more clearly that the ticket splitting rate among solid partisans was on the order of 1 percent in this sample. Such small subgroups are almost impossible to detect in a survey. In contrast, 97 percent of solid Republicans stuck with their party's nominee, Trump, and 98 percent of solid Democrats stuck with Biden. Trump's party loyalty was a percentage point smaller than Biden's.

A starker difference arises in the mixed group (those who vote for some Republicans and some Democrats, or undervoted, down-ballot). Biden won this group of weak partisans by more than 10 points. Close to 5 percent of this group voted for the third-party Libertarian candidate for President, instead of picking either Biden or Trump. Undervoting for President was low, less than 1 percent, in these group of voters.

More can be done to examine if these results vary by state, county, or precinct. Future versions of this dataset can also include ballot measures and local candidates that give more context for these patterns.

Arizona
4 counties, 2,706,032 voters
Maricopa
Pima
Santa Cruz
Yuma

California
34 counties, 14,642,275 voters
Alameda
Amador
Contra Costa
Del Norte
El Dorado
Fresno
Glenn
Humboldt
Kern
Kings
Lake
Los Angeles
Marin
Mendocino
Merced
Nevada
Orange
Placer
Riverside
San Benito
San Bernardino
San Diego
San Francisco
San Luis Obispo
San Mateo
Santa Barbara
Santa Cruz
Shasta
Sonoma
Sutter
Tehama
Tuolumne
Ventura
Yuba

Colorado
57 counties, 2,698,653 voters
Adams
Alamosa
Archuleta
Bent
Broomfield
Chaffee
Cheyenne
Clear Creek
Conejos
Costilla
Crowley
Custer
Delta
Denver
Dolores
Douglas
Eagle
El Paso
Elbert
Fremont
Garfield
Gilpin
Grand
Gunnison
Hinsdale
Huerfano
Jefferson
Kiowa
Kit Carson
La Plata
Lake
Larimer
Lincoln
Logan
Mesa
Mineral
Moffat
Montezuma
Montrose
Morgan
Otero
Ouray
Park
Phillips
Pitkin
Prowers
Pueblo
Rio Blanco
Rio Grande
Routt
Saguache
San Miguel
Sedgwick
Teller
Washington
Weld
Yuma

Delaware
507,773 voters
(Statewide)

Florida
26 counties, 5,738,195 voters
Bradford
Broward
Calhoun
Citrus
Collier
Duval
Escambia
Flagler
Gulf
Hamilton
Hillsborough
Holmes
Lafayette
Lake
Manatee
Marion
Martin
Nassau
Okaloosa
Orange
Palm Beach
Pasco
St Johns
Sumter
Wakulla
Walton

Georgia
86 counties, 3,643,546 voters
Bacon
Barrow
Bartow
Ben Hill
Berrien
Bibb
Bleckley
Brantley
Bryan
Burke
Camden
Carroll
Catoosa
Charlton
Chatham
Chattooga
Cherokee
Clarke
Cobb
Colquitt
Columbia
Cook
Crisp
Dade
Dawson
Dekalb
Douglas
Early
Echols
Elbert
Emanuel
Fayette
Forsyth
Franklin
Gilmer
Glascock
Glynn
Gordon
Grady
Greene
Gwinnett
Hall
Haralson
Harris
Hart
Heard
Henry
Houston
Irwin
Jackson
Jasper
Lanier
Laurens
Lee
Lowndes
Lumpkin
Madison
Mcduffie
Mitchell
Morgan
Murray
Muscogee
Newton
Oconee
Paulding
Pierce
Pike
Polk
Pulaski
Putnam
Rabun
Richmond
Rockdale
Schley
Spalding
Talbot
Tattnall
Terrell
Thomas
Tift
Upson
Walker
Walton
Warren
Whitfield
Wilcox

Illinois
9 counties, 250,778 voters
Brown
Clinton
Hamilton
Jo Daviess
Mchenry
Monroe
Pike
Union
Wayne

Iowa
1 county, 11,306 voters
Dickinson

Maryland
21 counties, 2,420,605 voters
Allegany
Anne Arundel
Baltimore
Baltimore City
Calvert
Caroline
Carroll
Cecil
Charles
Dorchester
Frederick
Garrett
Harford
Howard
Kent
Prince George's
Queen Anne's
Somerset
Talbot
Washington
Wicomico

Michigan
10 counties, 403,687 voters
Alcona
Alger
Charlevoix
Clinton
Eaton
Gladwin
Iron
Mason
Missaukee
Washtenaw

Nevada
17 counties, 1,407,496 voters
Carson City
Churchill
Clark
Douglas
Elko
Esmeralda
Eureka
Humboldt
Lander
Lincoln
Lyon
Mineral
Nye
Pershing
Storey
Washoe
White Pine

New Jersey
9 counties, 2,262,720 voters
Bergen
Camden
Cumberland
Essex
Gloucester
Hudson
Monmouth
Passaic
Salem

Ohio
33 counties, 1,941,167 voters
Allen
Ashtabula
Auglaize
Belmont
Butler
Champaign
Clark
Clinton
Darke
Erie
Fayette
Franklin
Gallia
Greene
Hancock
Harrison
Highland
Hocking
Logan
Miami
Pickaway
Preble
Putnam
Richland
Ross
Seneca
Shelby
Trumbull
Tuscarawas
Van Wert
Wayne
Wood
Wyandot

Oregon
14 counties, 958,676 voters
Columbia
Coos
Douglas
Harney
Josephine
Klamath
Lincoln
Linn
Marion
Polk
Union
Wasco
Washington
Yamhill

Rhode Island
518,200 voters
(Statewide)

Tennessee
4 counties, 216,215 voters
Loudon
Pickett
Sevier
Williamson

Texas
28 counties, 1,800,681 voters
Andrews
Bosque
Collin
Comal
Cooke
Cottle
Crane
Ellis
Erath
Fort Bend
Grayson
Guadalupe
Hidalgo
Kendall
Lee
Montague
Navarro
Oldham
Orange
Parmer
Potter
Scurry
Smith
Stephens
Taylor
Walker
Washington
Wharton

Utah
1 county, 6,938 voters
San Juan

West Virginia
2 counties, 49,754 voters
Nicholas
Wood

Wisconsin
4 counties, 525,751 voters
Brown
Kenosha
Pierce
Waukesha

Table 5. List of Counties Included. For a spreadsheet version of this list, see file `county_info.xlsx` in our Dataverse repository.

Code availability

The code we use to construct the dataset is available at https://github.com/kuriwaki/cvr_harvard-mit_scripts.

References

1. King, G. *et al.* The Record of American Democracy, 1984-1990 (1997). https://road.hmdc.harvard.edu/pages/road-documentation.
2. Ansolabehere, S., Palmer, M. & Lee, A. Precinct-Level Election Data, 2002-2012, https://doi.org/10.7910/DVN/YN4TLR (2014).
3. Baltz, S. *et al.* American election results at the precinct level. *Nature Scientific Data* https://doi.org/10.1038/s41597-022-01745-0 (2022).
4. Voting and Election Science Team. Precinct-Level Election Results, https://doi.org/10.7910/DVN/K7760H (2024). Last accessed 2024.
5. Wack, J. P. Cast vote records common data format specification version 1.0. *National Institute of Standards and Technology* https://doi.org/10.6028/NIST.SP.1500-103 (2019).
6. Kuriwaki, S., Reece, M. *et al.* Cast vote records: A database of ballots from the 2020 U.S. Election, Harvard Dataverse, https://doi.org/10.7910/DVN/PQQ3KV (2024).
7. Green, R. FOIA-Flooded Elections. *Ohio State Law Journal* 255–306 (2024).
8. Leingang, R. Election activists are seeking the "cast vote record" from 2020. here's what it is and why they want it. *Votebeat, September 7, 2022* (2022). https://perma.cc/55F4-NVKW.
9. O'Donnell, J. The fingerprints of fraud: Evidence and analysis of multi-state conspiracy to defraud the 2020 general election, vol. 1. *Unpublished Manuscript Available Online* (2023). https://perma.cc/6KY3-3GFK.
10. Bloomberg Technology. 'Raccoon Army' Swamps Election Officials in Dubious Campaign to Disprove Results (2022). October 25, 2022, https://perma.cc/J2H6-TAUD.
11. Grimmer, J., Herron, M. C. & Tyler, M. Evaluating a new generation of expansive claims about vote manipulation. *Election Law Journal* https://doi.org/elj.2022.0070 (2024).
12. Burden, B. C. & Kimball, D. C. *Why Americans Split their Tickets: Campaigns, Competition, and Divided Government* (University of Michigan Press, 2009).
13. Rogers, S. National forces in state legislative elections. *The Annals of the American Academy of Political and Social Science* **667**, 207–225 (2016).
14. Malzhan, J. & Hall, A. B. Election-denying republican candidates underperformed in the 2022 midterms. *Forthcoming, American Political Science Review* (2024).
15. Reece, M. *et al.* Hidden Partisanship in American Elections. *Preprint* (2024). https://ssrn.com/abstract=4721012.
16. Conevska, A. *et al.* How Partisan are U.S. Local Elections? Evidence from 2020 Cast Vote Records. *Preprint* (2024).
17. Kuriwaki, S. Ticket Splitting in a Nationalized Era. *Preprint* https://doi.org/10.31235/osf.io/bvgz3 (2023).
18. Dubin, J. A. & Gerber, E. R. Patterns of voting on ballot propositions: A mixture model of voter types. Tech. Rep., California Institute of Technology, Division of the Humanities and Social Sciences (1992). https://doi.org/10.7907/xprd7-7jt08.
19. Gerber, E. R. & Lewis, J. B. Beyond the Median: Voter Preferences, District Heterogeneity, and Political Representation. *Journal of Political Economy* **112**, 1364–1383, https://doi.org/10.1086/424737 (2004).

20. Besley, T. & Coate, S. Issue Unbundling via Citizens' Initiatives. *Quarterly Journal of Political Science* **3**, 379–397 (2008).

21. Strömberg, D. How the Electoral College influences campaigns and policy: the probability of being Florida. *American Economic Review* **98**, 769–807 (2008).

22. Larcinese, V., Snyder, J. M. & Testa, C. Testing models of distributive politics using exit polls to measure voters' preferences and partisanship. *British Journal of Political Science* **43**, 845–875 (2013).

23. Hill, S. J. Changing votes or changing voters? How candidates and election context swing voters and mobilize the base. *Electoral Studies* **48**, 131–148 (2017).

24. Dowling, C. M., Miller, M. G. & Morris, K. Can Voters Locate Copartisan Candidates in Nonpartisan Elections? Evidence from Cast Vote Records. *Working Paper* (2024).

25. DeLuca, K. Editor's choice: Measuring candidate quality using local newspaper endorsements. *APSA Preprints* https://doi.org/10.33774/apsa-2023-3qdmj-v3 (2024).

26. Alvarez, R. M., Hall, T. E. & Levin, I. Low-information voting: Evidence from instant-runoff elections. *American Politics Research* **46**, 1012–1038 (2018).

27. Atsusaka, Y. Analyzing ballot order effects when voters rank candidates. *Political Analysis* 1–9 (2024).

28. Herron, M. C. & Lewis, J. B. Did Ralph Nader Spoil a Gore Presidency? A Ballot-level Study of Green and Reform Party Voters in the 2000 Presidential Election. *Quarterly Journal of Political Science* **2**, 205–226, https://doi.org/10.1561/100.00005039 (2007).

29. Adler, E. S. & Hall, T. E. Ballots, transparency, and democracy. *Election Law Journal* **12**, 146–161, https://doi.org/10.1089/elj.2012.0179 (2013).

30. Bernhard, M. *et al.* Public evidence from secret ballots. In *Electronic Voting: Second International Joint Conference, E-Vote-ID 2017, Bregenz, Austria, October 24-27, 2017, Proceedings 2*, 84–109, https://doi.org/10.1007/978-3-319-68687-5_6 (Springer, 2017).

31. Gerber, A. S., Huber, G. A., Doherty, D. & Dowling, C. M. Is there a secret ballot? ballot secrecy perceptions and their implications for voting behaviour. *British Journal of Political Science* **43**, 77–102, https://doi.org/10.1017/S000712341200021X (2013).

32. Atkeson, L. R., McKown-Dawson, E., Hood III, M. & Stein, R. Voter perceptions of secrecy in the 2020 election. *Election Law Journal* **22**, 234–253, https://doi.org/10.1089/elj.2022.0064 (2023).

33. Jaffe, J., Loffredo, J., Baltz, S., Flores, A. & Stewart III, C. Trust in the count: Improving voter confidence with post-election audits. *Public Opinion Quarterly* **88**, 585–607 (2024).

34. Wand, J. N. *et al.* The butterfly did it: The aberrant vote for Buchanan in Palm Beach County, Florida. *American Political Science Review* **95**, 793–810 (2001).

35. Bafumi, J., Herron, M. C., Hill, S. J. & Lewis, J. B. Alvin Greene? Who? How Did He Win the United States Senate Nomination in South Carolina? *Election Law Journal* **11**, 358–379, https://doi.org/10.1089/elj.2011.0137 (2012).

36. Biggers, D. R. *et al.* Can addressing integrity concerns about mail balloting increase turnout? results from a large-scale field experiment in the 2020 presidential election. *Journal of Experimental Political Science* **10**, 413–425, https://doi.org/10.1017/XPS.2022.31 (2023).

37. Gerber, A. S., Huber, G. A., Doherty, D., Dowling, C. M. & Hill, S. J. Do perceptions of ballot secrecy influence turnout? results from a field experiment. *American Journal of Political Science* **57**, 537–551, https://doi.org/10.1111/ajps.12019 (2013).

38. Williams, J. R., Baltz, S. & Stewart, C. Votes can be confidently bought in some ranked ballot elections, and what to do about it. *Political Analysis* 1–13, https://doi.org/10.1017/pan.2024.4 (2024).

39. Kuriwaki, S., Lewis, J. B. & Morse, M. The still secret ballot: The limited privacy cost of transparent election results. *Preprint* https://doi.org/10.48550/arXiv.2308.04100 (2023).

40. MIT Election Data and Science Lab. Precinct-level returns 2020 by individual state. Harvard Dataverse, https://doi.org/10.7910/DVN/NT66Z3 (2022).

41. Federal Election Commission. Federal Elections 2020: Election Results for the U.S. President, the U.S. Senate and the U.S. House of Representatives. Tech. Rep., United States Government (2022). https://perma.cc/DE4X-SQFC.

42. Dane County Clerk. Re: November 2020 general election ballot images. Letter to Brian McGrath and Kyle Koenen on August 13, 2021 (2021).

43. Ferrer, J., Thompson, D. M. & Orey, R. Election official turnover rates from 2000-2024. *Bipartisan Policy Center* (2024). https://bipartisanpolicy.org/report/election-official-turnover-rates-from-2000-2024/.

44. Wickham, H., Çetinkaya-Rundel, M. & Grolemund, G. *R for Data Science (2e)* (O'Reilly, 2023), second edn. https://r4ds.hadley.nz/.

Acknowledgements

We thank Luka Bulić Bračulj for additional work analyzing the data. We also thank Scott McDonnell of Dane County for answering our questions about Dane County, as well as other local election officials that answered our questions. We also thank Larry Moore for inspiration on the usage example. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-2139841 (to T.S.).

Author contributions statement

S.K. drafted the manuscript; M.R., J.S., and J.B.L. were the principal maintainers of the database construction; A.C., C.S., J.B.L., J.R.L., J.S., M.R., S.H., S.K. contributed to the manuscript; A.C., C.M., J.B.L., J.R.L., J.S., K.A., K.M., S.H., S.K., M.R., T.S., Z.G., performed data processing, validation, investigation; J.B.L., J.R.L., J.S., M.R., S.K., S.B., T.S., wrote software; C.S., J.S., M.R., S.B., S.K. supervised the project. All authors reviewed the paper.

Competing interests

Authors declare no conflicting interests that might be perceived to influence the results and/or discussion reported in this paper. This project was determined as not human subjects research by the MIT and Yale Institutional Review Board (IRB).

Supplementary Information

This document is supplementary information for the Cast Vote Records Data Descriptor[1].

A. Precinct Name Standardization

Linking CVR data to other precinct-level data sources is not straightforward. The released CVR data typically contains information about the voting precinct in which each ballot was cast. However, there is no agreed-upon standard for the labeling of precincts. The precinct labels used in the CVR data often do not match those used in the official precinct-level results published by the county or state. Further, CVRs sometimes provide more detailed geographic information (commonly called "subprecinct" information) that is not included in the official precinct-level tallies.

In order to make the CVR data more useful to researchers, we constructed a crosswalk from the precinct labels used in the CVR data to the precinct labels in a widely-used freely-available national database of certified precinct-level 2020 General Election results published by the MIT Election Data and Science Lab (see the descriptor by Baltz *et al.*[2]). Using that crosswalk, we have appended the MEDSL precinct label to our dataset.

We created this crosswalk by the following steps:

1. First, we aggregated the CVR data to the (sub)precinct level based on the precinct labels that they include. (Note, for counties such as Los Angeles, California for which the CVRs do not provide precinct identifiers, we have not attempted to place the CVRs in MEDSL precincts. Such a linkage is not in general possible, though it might be accomplished in some cases by using the set of contests included on each ballot – sometimes called the "ballot style" – to identify the precinct.) We then attempted to match our CVR-based precinct records to the MEDSL precinct data. Working county-by-county, we matched CVR precincts to MEDSL precincts based upon the reported number of votes cast for Democratic and Republican candidates for US President, US Senate, US House, State Senate, and State House. For many counties, unique exact matches for every CVR precinct could be found among the MEDSL precincts. For those counties, this set of unique exact matches was used as the crosswalk between CVR and MEDSL precincts labels.

2. For counties in which unique exact candidate vote matches could not be found for every precinct, we transformed the CVR precinct names to more closely match the format of the MEDSL precinct names. We then disambiguated cases in which a given CVR precinct's candidate vote total exactly matched more than one MEDSL precinct using the edit distance between the transformed CVR precinct name and the MEDSL precinct name by selecting the potential match having the smallest edit distance. In cases in which no exact match based on the candidate vote totals existed, we established a linkage if the transformed CVR precinct name exactly matched the MEDSL precinct name.

3. For those counties containing precincts for which neither the candidate vote totals nor the transformed precinct names could be exactly matched across the CVR and MEDSL data, potential matches were identified by the smallest sum of absolute deviations in candidate votes and confirmed by manual comparison of the CVR and the MEDSL precinct labels. Where this was not possible, no linkage was established.

4. For some counties, we determined that the CVR precinct labels included subprecinct information meaning that the CVR-based precinct-level data was at a lower level of aggregation than that reported in the MEDSL data. In those counties, we transformed to CVR precinct names to remove the subprecinct information and re-aggregated to produce CVR precinct records at the same level of aggregation as the MEDSL precincts before constructing the crosswalk using the methods described above.

MEDSL precinct labels were not applied to records for which a precinct identifier was not present in the CVR data nor to CVR precincts for which neither the names nor the vote totals established a convincing linkage between the two data sources.

In four Michigan Counties (Alcona, Clinton, Gladwin, and Missaukee) few or none of the CVR precincts could be matched to MEDSL precincts using the method described. In these four counties, the inability to link CVR precincts to MEDSL precincts appeared to be the result of irregularities in the MEDSL data. Those issues may be resolved in the future.

Fifteen other counties contained between one and five CVR precincts that could not be linked to MEDSL precincts. All of the precincts in the remaining 296 counties for which CVR precinct information was available could be linked to MEDSL precincts. In total, 28,540 CVR unique (sub)precincts were linked to 23,467 MEDSL precincts.

B. Fixing Fragmented Ballots

The counties in Table S1 did not have a clear cast vote record identifier. We nevertheless paired records as one voter (one CVR ID) using a matching algorithm described in Algorithm 1. The algorithm relies on the user to supply two arguments, (1) a grouping column, `args.groupcol`, that is a condition for pages to belong to the same voter and (2) a target column, `args.targetcol`, in the CVR that can be used to match pages together. For example, if the column is the President, if two rows actually belong to the same voter, the pair must have an identical value for the `args.groupcol`, and one row will have a value for President and the other row will be blank for President.

Algorithm 1 Process Missing Values in Dataset

Initialize an empty set `used_rows` to keep track of processed rows.
Determine the total number of rows in the dataset, `num_rows`.
for each row `current_row` at index i in the dataset **do**
    if i is in `used_rows` **then**
        continue
    end if
    if `current_row[args.targetcol]` is missing a value **then**
        for each index j in expanding spiral around i **do**
            if $j \geq$ `num_rows` **then**
                break
            end if
            if j is not in `used_rows` and `current_row[args.groupcol]` equals `data[j, args.groupcol]` and `data[j, args.targetcol]` is not missing **then**
                Merge `current_row` with the complementary row at index j to form `merged_row`.
                Update `data[i]` with `merged_row`.
                Add j to `used_rows`.
                break
            end if
        end for
    end if
end for

C. Third Party Candidates and Write-ins

Designations for third parties vary widely by state. This is compounded by the different ways counties print party on the ballot, and the types of voting machines store and record information about third party candidates and write-ins.

Table S1. Counties Where Algorithm 1 was used to re-connect pages

County	Group Column
Alameda, California	Ballot Type
Contra Costa, California	Ballot Type
Kings, California	Ballot Type
Merced, California	Ballot Style
Riverside, California	Ballot Type
San Benito, California	Ballot Type
San Mateo, California	Ballot Type
Sonoma, California	Ballot Type
Yuba, California	Ballot Type
Denver, Colorado	Ballot Type
Eagle, Colorado	Ballot Type
Routt, Colorado	Ballot Type
Gwinnett, Georgia	Ballot Type
Baltimore, Maryland	Ballot Style
Baltimore City, Maryland	Ballot Style
Montgomery, Maryland	Ballot Style
Prince George's Maryland	Ballot Style
Butler, Ohio	Ballot Type
Champaign, Ohio	BallotStyleID
Cuyahoga, Ohio	Ballot Style
Greene, Ohio	Ballot Type
Rhode Island	Ballot Style

Howie Hawkins was the Green Party's nominee for U.S. President but the Green Party did not get ballot access in some states. Among the states that we examined, the Green Party appears to not have had ballot access in Arizona, Georgia, Pennsylvania, and Wisconsin. Some jurisdictions, e.g., in Wisconsin, do not report out Hawkins as a specific candidate in either the cast vote record or the election returns; he is lumped into simply `WRITEIN`. In Georgia, the cast vote record and the election returns report Hawkins as its own candidate, but they are still write-ins. (See `https://ballotpedia.org/Georgia_official_sample_ballots,_2020` for a sample ballot, and `https://bit.ly/3RYhL3A` for an example of how Hawkins is reported in an election return.) In all these cases, the `party` value for Hawkins is a Write-in, not the Green Party. In contrast, New Jersey's CVRs records each and every write-in choice as a valid candidate.

It is worth noting that some of the data on `votedatabase.com` lost information about write-ins due to information loss from Excel to CSV. In some machines, notably the ES&S DS200 scanner, write-ins are scanned and can produce a cast vote record in Microsoft Excel. In these sheets, the write-in candidates are not transcribed into text but stored as an image file. However, O'Donnell instructed his collaborators to upload these data as a plain-text CSV file, which loses information in the image. In these cases, write-in votes are tracked as blank cells and excluded from our dataset. Our data, therefore, systematically misses write-in votes in many counties that use this format of cast vote record.

D. Sample Characteristics

See Table 4 in the main text for the list of counties included in our data release. In Delaware and Rhode Island, we count the whole state as one county given the format of the data.

Figure S1 shows the states in our dataset and the coverage of its population for each office. The percentage values assigned for each state indicate the total number of ballots relative to the

Figure S1. Coverage of CVR data. *The percentage shows the fraction of total voters contained in the CVR for that state and office. States are not shown on the map if the office was not on the ballot. States have no percentage if CVRs were not available for that state.*

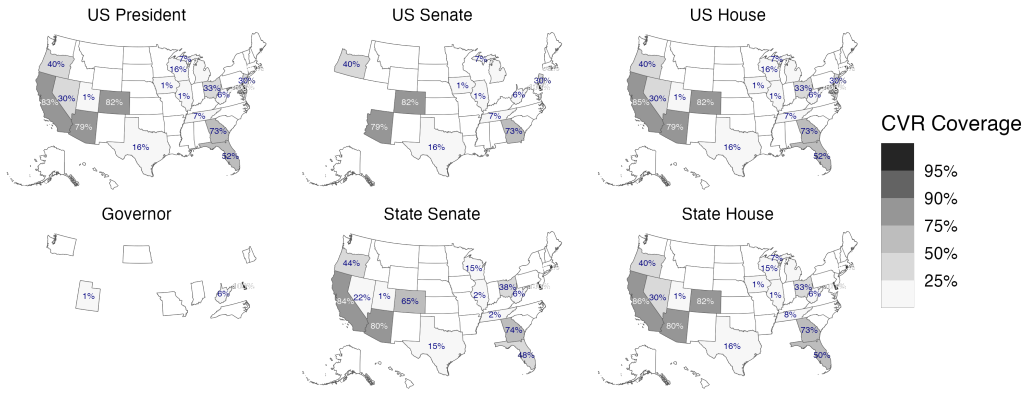

total count of votes reported in the entire state. For example, in Texas, we have around 15 percent of the votes for the President, Congress, and state legislature. There was no election for Governor in the state.

References

1. Kuriwaki, S., Reece, M. *et al.* Cast vote records: A database of ballots from the 2020 U.S. Election, Harvard Dataverse, 10.7910/DVN/PQQ3KV (2024).
2. Baltz, S. *et al.* American election results at the precinct level. *Nature Scientific Data* 10.1038/s41597-022-01745-0 (2022).