

Racial Disparities in Voting Wait Times: Evidence from Smartphone Data

M. Keith Chen, Kareem Haggag, Devin G. Pope, and Ryne Rohla*

October 30, 2020

Abstract

Equal access to voting is a core feature of democratic government. Using data from hundreds of thousands of smartphone users, we quantify a racial disparity in voting wait times across a nationwide sample of polling places during the 2016 U.S. presidential election. Relative to entirely-white neighborhoods, residents of entirely-black neighborhoods waited 29% longer to vote and were 74% more likely to spend more than 30 minutes at their polling place. This disparity holds when comparing predominantly white and black polling places within the same states and counties, and survives numerous robustness and placebo tests. We shed light on the mechanism for these results and discuss how geospatial data can be an effective tool to both measure and monitor these disparities going forward.

JEL Classifications: D70, J15

*Chen: Anderson School of Management, University of California at Los Angeles, keith.chen@anderson.ucla.edu, Haggag: Social and Decision Sciences, Carnegie Mellon University, kareem.haggag@cmu.edu. Pope: Booth School of Business, University of Chicago, devin.pope@chicagobooth.edu. Rohla: Anderson School of Management, University of California at Los Angeles, ryne.rohla@anderson.ucla.edu. The authors thank Stefano DellaVigna, Dean Karlan, Larry Katz, Charles Stewart III, and seminar audiences at Carnegie Mellon University, Stanford University, University of Chicago, and University of Illinois Urbana-Champaign for helpful comments and suggestions, Zachary Goldstein for research assistance, and A. Hoffman, R. Squire, and N. Yonack at SafeGraph for data access and technical assistance.

Providing convenient and equal access to voting is a central component of democratic government. Among other important factors (e.g. barriers to registration, purges from voter rolls, travel times to polling places), long wait times on Election Day are a frequently discussed concern of voters. Long wait times have large opportunity costs (Stewart and Ansolabehere 2015), may lead to line abandonment by discouraged voters (Stein et al. 2019), and can undermine voters' confidence in the political process (Alvarez et al. 2008; Atkeson and Saunders 2007; Bowler et al. 2015). The topic of long wait times has reached the most prominent levels of media and policy attention, with President Obama discussing the issue in his 2012 election victory speech and appointing a presidential commission to investigate it. In their 2014 report, the Presidential Commission on Election Administration concluded that, "as a general rule, no voter should have to wait more than half an hour in order to have an opportunity to vote."

There have also been observations of worrying racial disparities in voter wait times. The Cooperative Congressional Election Study (CCES) finds that black voters report facing significantly longer lines than white voters (Pettigrew 2017; Alvarez et al. 2009; Stewart III 2013). While suggestive, the majority of prior work on racial disparities in wait times has been based on anecdotes and surveys which may face limits due to recall and reporting biases.

In this paper, we use geospatial data generated by smartphones to measure wait times during the 2016 election. For each cellphone user, the data contain "pings" based on the location of the cellphone throughout the day. These rich data allow us to document voter wait times across the entire country and also estimate how these wait times differ based on neighborhood racial composition.

We begin by restricting the set of smartphones to a sample that passes a series of filters to isolate likely voters. This leaves us with a sample of just over 150,000 smartphone users who voted at one of more than 40,000 polling locations across 46 different states. Specifically,

these individuals entered and spent at least one minute within a 60-meter radius of a polling location on Election Day and recorded at least one ping within the convex hull of the polling place building (based on building footprint shapefiles). We eliminate individuals who entered the same 60-meter radius in the week leading up to or the week after Election Day to avoid non-voters who happen to work at or otherwise visit a polling place on non-election days.

We estimate that the median and average times spent at polling locations are 14 and 19 minutes, respectively, and 18% of individuals spent more than 30 minutes voting.¹ We provide descriptive data on how voting varies across the course of Election Day. As expected, voter volume is largest in the morning and in the evening, consistent with voting before and after the workday. We also find that average wait times are longest in the early morning hours of the day. Finally, as a validation of our approach, we show that people show up to the polls at times consistent with the opening and closing hours used in each state.

We next document geographic variation in average wait times using an empirical Bayes adjustment strategy. We find large differences across geographic units – for example, average wait times across congressional districts can vary by a factor of as much as four. We further validate our approach by merging in data from the CCES, which elicits a coarse measure of wait time from respondents. Despite many reasons for why one might discount the CCES measures (e.g. reporting bias and limited sample size), we find a remarkably high correlation with our own measures – a correlation of 0.86 in state-level averages and 0.73 in congressional-district-level averages. This concordance suggests that our wait time measures (and those elicited through the survey) have a high signal-to-noise ratio.

¹The time measure that we estimate in our paper is a combination of wait time in addition to the time it took to cast a ballot. We typically refer to this as just “wait time” in the paper. One may worry that the differences we find are not about wait times, but rather about differences in the amount of time spent casting a ballot. However, there is evidence to suggest this is not the case. For example, we find incredibly strong correlations between our wait time measures and survey responses that ask only about wait times as opposed to total voting time (“Approximately, how long did you have to wait in line to vote?”).

We next explore how wait times vary across areas with different racial compositions. We use Census data to characterize the racial composition of each polling place's corresponding Census block group (as a proxy for its catchment area). We find that the average wait time in a Census block group composed entirely of black residents is approximately 5 minutes longer than average wait time in a block group that has no black residents. We also find longer wait times for areas with a high concentration of Hispanic residents, though this disparity is not as large as the one found for black residents. These racial disparities persist after controlling for population, population density, poverty rates, and state fixed effects. We further decompose these effects into between- and within-county components, with the disparities remaining large even when including county fixed effects. We perform a myriad of robustness checks and placebo specifications and find that the racial disparity exists independent of the many assumptions and restrictions that we have put on the data.

In the Appendix, we consider the potential mechanisms behind the observed racial differences. We ultimately find that a host of plausible candidate explanations do little to explain the disparity in our cross-section, including differences in arrival times of voters, state laws (Voter ID and early voting), the partisan identity of the underlying population or the chief election official, county characteristics (income inequality, segregation, social mobility), and the number of registered voters assigned to a polling place; although we do find larger disparities at higher-volume polling locations. Overall, our results on mechanism suggest that the racial disparities that we find are widespread and unlikely to be isolated to one specific source or phenomenon.

Our paper is related to work in political science that has examined determinants of wait times and also explored racial disparities. Some of the best work uses data from the CCES which provides a broad sample of survey responses on wait times (Pettigrew 2017; Alvarez et al. 2009; Stewart III 2013). For example, Pettigrew (2017) finds that black voters report waiting in line for twice as long as white voters and are three times more likely to wait for

over 30 minutes to vote. Additional studies based on field observations may avoid issues that can arise from self-reported measures, but typically only cover small samples of polling places such as a single city or county (Highton 2006; Spencer and Markovits 2010; Herron and Smith 2016). Stein et al. (2019) collect the largest sample to date, using observers with stopwatches across a convenience sample of 528 polling locations in 19 states. Using a sample of 5,858 voters, they provide results from a regression of the number of people observed in line on an indicator that the polling place is in a majority-minority area. They find no significant effect – although they also control for arrival count in the regression. In a later regression, they find that being in a majority-minority polling location leads to a 12-second increase in the time it takes to check in to vote (although this regression includes a control for the number of poll workers per voter which may be a mechanism for racial disparities in voting times). Overall, we arrive at qualitatively similar results as the political science literature, but do so using much more comprehensive data that avoids the pitfalls of self-reports. Going forward, this approach could produce repeated measures across elections, which would facilitate a richer examination of the causal determinants of the disparities.

Our paper also relates to the broader literature on racial discrimination against black individuals and neighborhoods (for reviews, see Altonji and Blank 1999, Charles and Guryan 2011, and Bertrand and Duflo 2017), including by government officials. For example, Butler and Broockman (2011) find that legislators were less likely to respond to email requests from a putatively black name, even when the email signaled shared partisanship in an attempt to rule out strategic motives. Similarly, White et al. (2015) find that election officials in the U.S. were less likely to respond and provided lower-quality responses to emails sent from constituents with putatively Latino names. Racial bias has also been documented for public officials that are not part of the election process. For example, Giulietti et al. (2019) find that emails sent to local school districts and libraries asking for information were less likely to receive a response when signed with a putatively black name relative to a putatively

white name. As one final example, several studies have documented racial bias by judges in criminal sentencing (Alesina and Ferrara 2014; Glaeser and Sacerdote 2003; Abrams et al. 2012).

1 Data

The three primary datasets we use in this paper include: (1) SafeGraph cell phone location records, (2) Polling locations, and (3) Census demographics.

We use anonymized location data for smartphones provided by SafeGraph, a firm that aggregates location data across a number of smartphone applications (Chen and Rohla 2018). These data cover the days between November 1st and 15th, 2016, and consist of “pings”, which record a phone’s location at a series of points in time. In general, GPS pings are typically accurate to within about a 5-meter radius under open sky, though this varies depending on factors such as weather, obstructions, and satellite positioning (GPS.gov). Pings are recorded anytime an application on a phone requests information about the phone’s location. Depending on the application (e.g. a navigation or weather app), pings may be produced when the application is being used or at regular intervals when it is in the background. The median time between pings in our sample for a given device is 48 seconds (with a mode of 5 minutes).

The geolocation data used in this paper is detailed and expansive, allowing us to estimate wait times around the entire United States. This data, however, naturally raises concerns about representativeness. If we were trying to estimate individual choices, e.g. vote choice, the sample could only produce estimates that are at best representative of the approximately 77% of U.S. adults who owned a smartphone in 2016. While Chen and Pope (2019) show that the data are generally representative of the U.S. along several observable dimensions (with the exception of skewing more wealthy), they may differ on unobservables. However, our goal

is to estimate a property of places rather than individuals. That is, we estimate an outcome of queues that have multiple individuals in them. While the restriction to smartphone users may limit the number of wait times we observe, as long as there is a queueing rule at polling places, we should still observe an unbiased estimate of the wait times faced by voters, both those with and without smartphones.²

Polling place addresses for the 2016 General Election were collected by contacting state and county election authorities. When not available, locations were sourced from local newspapers, public notices, and state voter registration look-up webpages. State election authorities provided statewide locations for 32 states, five of which required supplemental county-level information to complete. Four states were completely collected on a county-by-county basis. In twelve states, not all county election authorities responded to inquiries (e.g. Nassau County, New York).

When complete addresses were provided, the polling locations were geocoded to coordinates using the Google Maps API. When partial or informal addresses were provided, buildings were manually assigned coordinates by identifying buildings through Google Street View, imagery, or local tax assessor maps as available. Additionally, Google Maps API geocodes are less accurate or incomplete in rural locations or areas of very recent development, and approximately 8% of Google geocodes were manually updated.

Of the 116,990 national polling places reported in 2016 by the U.S. Election Assistance Commission, 93,658 polling places (80.1%) were identified and geocoded and comprise the initial sample of polling places in this paper. Appendix Figure A1 illustrates the location of the 93,658 polling places and separately identifies polling places for which we identify likely voters on Election Day and pass various filters that we discuss and impose below.

Demographic characteristics were obtained by matching each polling place location to the

²This is not to dismiss the potential issue of missing polling places or times of day. However, *a priori*, these omissions do not point to systematic bias in a particular direction.

census block group in the 2017 American Community Survey’s five-year estimates. Census block groups were chosen as the level of aggregation because the number of block groups is the census geography that most closely aligns with the number of polling places and because it contains the information of interest (racial characteristics, fraction below poverty line, population, and population density).

2 Methods

In order to calculate voting wait times, we need to identify a set of individuals we are reasonably confident actually voted at a polling place in the 2016 election. To do so, we restrict the sample to phones that record a ping within a certain distance of a polling station on Election Day. This distance is governed by a trade-off – we want the radius of the circle around each polling station to be large enough to capture voters waiting in lines that may spill out of the polling place, but want the circle to be not so large that we introduce a significant number of false positive voters (people who came near a polling place, but did not actually vote).

We take a data-driven approach to determine the optimal size of the radius. In Panel A of Figure 1, we examine whether there are more unique individuals who show up near a polling place on Election Day relative to the week before and the week after the election (using a 100-meter radius around a polling location).³ As can be seen, there appear to be more than 400k additional people on Election Day who come within 100 meters of a polling place relative to the weekdays before and after. In Panel B of Figure 1, we plot the difference in the number of people who show up within a particular radius of the polling place (10 meters to 100 meters) on Election Day relative to the average across all other days. As we

³More precisely, we construct a 100-meter radius around the centroid of the building identified by Microsoft OpenStreetMap as the closest to the polling place coordinates.

increase the size of the radius, we are able to identify more and more potential voters, but also start picking up more and more false positives. By around 60 meters, we are no longer identifying very many additional people on Election Day relative to non-election days, and yet are continuing to pick up false positives. Therefore, we choose 60 meters as the radius for our primary analysis. However, in Section 4.1, we demonstrate robustness of estimates to choosing alternative radii.

For each individual that comes within a 60-meter radius of a polling place, we would like to know the amount of time spent within that radius. Given that we do not receive location information for cell phones continuously (the modal time between pings is 5 minutes), we cannot obtain an exact length of time. Thus, we create upper and lower bounds for the amount of time spent voting by measuring the time between the last ping before entering and the first ping after exiting a polling-place circle (for an upper bound), and the first and last pings within the circle (for a lower bound). For example, pings may indicate a smartphone user was not at a polling location at 8:20am, but then was at the polling location at 8:23, 8:28, 8:29, and 8:37, followed by a ping outside of the polling area at 8:40am; translating to a lower bound of 14 minutes and an upper bound of 20 minutes. We use the midpoint of these bounds as our best guess of a voter's time at a polling place (e.g. 17 minutes in the aforementioned example). In the robustness section, we estimate our effects using values other than the midpoint.

[FIGURE 1 ABOUT HERE]

Another important step in measuring voting times from pings is to isolate people who come within a 60-meter radius of a polling place that we think are likely voters and not simply passing by or people who live or work at a polling location. To avoid including passersby, we restrict the sample to individuals who had an upper bound measure of at least one minute within a polling place circle and for whom that is true at only one polling place

on Election Day. To avoid including people who live or work at the polling location, we exclude individuals who we observe spending time (an upper bound greater than 1 minute) at that location in the week before or the week after Election Day. To further help identify actual voters and reduce both noise and false positives, we restrict the sample to individuals who had at least one ping within the convex hull of the polling place building on Election Day (using Microsoft OpenStreetMap building footprint shapefiles), logged a consistent set of pings on Election Day (posting at least 1 ping every hour for 12 hours), and spent no more than 2 hours at the polling location (to eliminate, for example, poll workers who spend all day at a polling place). In Section 4.1, we provide evidence of robustness to these various sample restrictions.

After these data restrictions, our final sample consists of 154,489 individuals whom we identify as likely voters across 43,413 polling locations. Panel C in Figure 1 shows how many people pass our likely-voter filters on Election Day (154,489), and—as a placebo analysis—how many observations we would have on non-Election (“placebo”) days before and after the 2016 Election that would pass these same filters (modified to be centered around those placebo days). This analysis suggests that more than 87% of our sample are likely voters who would not have been picked up on days other than Election Day. In Appendix Figure A2, we plot the distribution of wait times on each of these placebo non-election days. We find that the wait times of people who would show up in our analysis on non-election days are shorter on average than those that show up on Election Day. Thus, to the degree that we can not completely eliminate false positives in our voter sample, we expect our overall voter wait times to be biased upward. We also would expect the noise introduced by non-voters to bias us towards not finding systematic disparities in wait times by race.

Appendix Table A1 provides summary statistics for our 154,489 likely voters. We find average voting wait times of just over 19 minutes when using our primary wait time measure (the midpoint between the lower and upper bound) and 18% of our sample waited more than

30 minutes to vote. Weighted by the number of voters in our sample, the racial composition of the polling place block groups is, on average, 70% white and 11% black.

3 Results: Overall Voter Wait Times

We plot the distribution of wait times in Panel A of Figure 2. The median and average times spent at polling locations are 14 and 19 minutes, respectively, and 18% of individuals spent more than 30 minutes voting. As the figure illustrates, there is a non-negligible number of individuals who spent 1-5 minutes in the polling location (less time than one might imagine is needed to cast a ballot). These observations might be voters who abandoned after discovering a long wait time. Alternatively, they may be individuals who pass our screening as likely voters, but were not actually voting.

[FIGURE 2 ABOUT HERE]

We next display the number of people who arrive to vote at the polling locations by time of day. This descriptive analysis of when people vote may be of interest in and of itself, but it also serves as a validation of whether people in our sample are indeed likely voters (e.g. if our sample consists primarily of people showing up at the polling locations at 3am, then one should worry about whether our sample is primarily composed of voters). Panel C of Figure 2 shows the distribution of arrival times where the “hour of day” is defined using the “hour of arrival” for a given wait time (i.e. the earliest ping within the polling place radius for a given wait time spell). As expected, people are most likely to vote early in the morning or later in the evening (e.g. before or after work) with nearly twice as many people voting between 7 and 8am as between noon and 1pm. As a consistency check, Appendix Figure A3 repeats this figure separately by state’s opening and closing times – the figures show that likely-voter arrivals match state-by-state poll opening and closing times. Finally, Panel D of

Figure 2 plots the average wait time by time of arrival, showing that the longest averages are early in the morning.

In addition to temporal variation in wait times, we can also explore how voting wait times vary geographically. In online materials, we report average wait times by state, congressional district, and the 100 most populous counties, along with accompanying standard deviations and observation counts, as well as an empirical-Bayes adjustment to account for measurement error.^{4,5} Focusing on the empirical-Bayes adjusted estimates, the states with the longest average wait times are Utah and Indiana (28 and 27 minutes, respectively) and the states with the shortest average wait time are Delaware and Massachusetts (12 minutes each). In Panel B of Figure 2 we map the empirical-Bayes-adjusted average voting wait time for each congressional district across the United States. Average wait times vary from as low as ~ 11 minutes in Massachusetts’s Sixth and Connecticut’s First Congressional District to as high as ~ 40 minutes in Missouri’s Fifth Congressional District. These geographic differences are not simply a result of a noisy measure, but contain actual signal value regarding which areas have longer wait time than others. Evidence for this can be seen by our next analysis correlating our wait time measures with those from a survey.

We correlate our average wait time measures at both the state and congressional district

⁴See Appendix C: <https://www.nber.org/papers/w26487>

⁵Even if all states in the U.S. had the same voter wait time, we would find some dispersion in our measure due to sampling variation. Due to sample size, this measurement error in our estimates would result in the smallest states being the most likely to show evidence of having either very short or very long wait times. Thus, throughout the paper, whenever we discuss voter wait times or racial disparities that have been aggregated up to either the county, congressional district, or state level, we will report estimates that have been adjusted for measurement using a Bayesian shrinkage procedure. This iterative procedure (discussed in detail in Chandra et al. (2016)) shrinks estimates toward the average of the true underlying distribution. The amount of adjustment toward the mean is a function of how far the estimate for each state/county is from the mean and the estimate’s precision. The resulting adjusted estimate is our “best guess” (using Bayesian logic) as to what the actual wait time or disparity is for each geographic unit.

level with the average wait times reported by respondents in the the 2016 wave of the Cooperative Congressional Election Study (Ansolabehere and Schaffner, 2017). The 2016 CCES is a large national online survey of 64,600 people conducted before and after the U.S. general election. The sample is meant to be representative of the U.S. as a whole.⁶ There are several reasons one might be pessimistic that the wait time estimates that we generate using smartphone-data would correlate closely with the wait times reported from the CCES survey. First, given sample sizes at the state and congressional district level, both our wait times and survey wait times may have a fair bit of sampling noise. Second, our wait time measures are a combination of waiting in line and casting a ballot, whereas the survey only asks about wait times. Third, the question in the survey creates additional noise by eliciting wait times that correspond to one of five coarse response options (“not at all”, “less than 10 minutes”, “10 to 30 minutes”, “31 minutes to an hour”, and “more than an hour”).⁷ Lastly, the survey does not necessarily represent truthful reporting. For example, while turnout in the U.S. has hovered between 50 and 60 percent, more than 80% of CCES respondents report voting. Given these reasons for why our wait time results may not correlate well with those from the survey, we find a remarkably strong correlation between the two. Using empirical-Bayes-adjusted estimates for both state-level wait time estimates from the cellphone data and those found in the CCES, we find correlation of 0.86 between the two. We find a similarly strong correlation at the congressional district level (correlation = 0.73). Our wait-time estimates

⁶<https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi%3A10.7910/DVN/GDF6Z0>

⁷There are 34,353 responses to the “wait time” question in the 2016 CCES. We restrict the sample of responses to just use individuals who voted in person on Election Day (24,378 individuals after dropping the 45 who report “Don’t Know”). Following Pettigrew (2017), we translate the responses to minute values by using the midpoints of response categories: 0 minutes (“not at all”), 5 minutes (“less than 10 minutes”), 20 minutes (“10 to 30 minutes”) or 45 minutes (“31 minutes to an hour”). For the 421 individuals who responded as “more than an hour” we code them as waiting 90 minutes (by contrast, Pettigrew (2017) uses their open follow-up text responses.)

are, on average, slightly longer than those in the survey, which is likely a reflection of the fact that our measure includes both wait time and ballot-casting time. Scatter plots of the state and congressional district estimates may be found in Appendix Figure A4. Overall, the strong correlations between the wait times we estimate and those from the CCES survey provide validation for our wait time measure (and for the CCES responses themselves).

4 Results: Racial Disparities in Wait Times

In this section, we provide evidence that wait times are significantly longer for areas with more black residents relative to white residents. We begin with a simple visualization of wait times by race. Figure 3 plots the smoothed distribution of wait times separately for polling places in the top and bottom deciles of the fraction-black distribution. These deciles average 58% and 0% black, respectively. Voters from areas in the top decile spent 19% more time at their polling locations than those in the bottom decile. Further, voters from the top decile were 49% more likely to spend over 30 minutes at their polling locations. Appendix Figures A5 and A6 provide similar density functions of wait-time comparisons for other demographic characteristics.

[FIGURE 3 ABOUT HERE]

[TABLE 1 ABOUT HERE]

Of course, Figure 3 focuses just on polling places that are at the extremes of racial makeup. We provide a regression analysis in Table 1 in order to use all of the variation across polling places' racial compositions and to provide exact estimates and standard errors. Panel A uses wait time as the dependent variable. In column 1, we estimate the bivariate regression which shows that moving from a census block group with no black residents to one that is entirely composed of black residents is associated with a 5.23 minute longer wait time. In

column 2, we broaden our focus by adding additional racial categories, revealing longer wait times for block groups with higher fractions of Hispanic and other non-white groups (Native American, other, multiracial) relative to entirely white neighborhoods. Column 3 examines whether these associations are robust to controlling for population, population density, and fraction below poverty line of the block group (see Appendix Tables A2 and A3 for the full set of omitted coefficients). The coefficient on fraction black is stable when adding in these additional covariates. Column 4 adds state fixed effects and the coefficient on fraction black only slightly decreases, suggesting that racial disparities in voting wait times are just as strong within state as they are between state.

In column 5, we present the results within county. We find that the disparity is mitigated, but it continues to be large and statistically significant. This suggests that there are racial disparities occurring both within and between county. Understanding the level at which discrimination occurs (state, county, within-county, etc.) is helpful when thinking about the mechanism. Further, the fact that we find evidence of racial disparities within county allows us to rule out what one may consider spurious explanations such as differences in ballot length between counties that could create backlogs at other points of service (Pettigrew 2017; Edelstein and Edelstein 2010; Gross et al. 2013).

Panel B of Table 1 is analogous to Panel A, but changes the outcome to a binary variable indicating a wait time longer than 30 minutes. We choose a threshold of 30 minutes as this was the standard used by the Presidential Commission on Election Administration in their 2014 report, which concluded that, “as a general rule, no voter should have to wait more than half an hour in order to have an opportunity to vote” (Bauer et al. 2014). We find that entirely black areas are 12 percentage points more likely to wait more than 30 minutes than entirely white areas, a 74% increase in that likelihood. This remains at 10 percentage points with polling-area controls and 7 percentage points within county.

4.1 Robustness

We have made several data restrictions and assumptions throughout the analysis. In this section, we document the robustness of the racial disparity estimate to using alternative restrictions and assumptions.

In our primary analysis we use the midpoint between the lower and upper bound of time spent near the polling location as the primary measure of wait time. In Panel A of Figure 4, we vary the wait time measure from the lower bound to the upper bound in 10 percent increments, finding that it has little impact on the significance or magnitude of our estimates. We further vary the wait time trimming thresholds in Panel B and the radius around a building centroid used to identify the polling location in Panel C. While these do move the average wait times around, and the corresponding differences, we find that the difference remains significant even across fairly implausible adjustments (e.g. a tight radius of 20 meters around a polling place centroid). We show the associated regression output for this figure in Appendix Table A4.

Another set of assumptions was in limiting the sample to individuals who (a) spent at least one minute at a polling place, (b) did so at only one polling place on Election Day, and (c) did not spend more than one minute at that polling location in the week before or the week after Election Day. As a robustness check, we make (c) stricter by dropping anyone who visited *any other polling place* on any day in the week before or after Election Day, e.g. we would thus exclude a person who only visited a school polling place on Election Day, but who visited a church (that later serves a polling place) on the prior Sunday. This drops our primary analysis sample from 154,489 voters down to 68,812 voters, but arguably does a better job of eliminating false positives. In Appendix Table A5 and Appendix Figure A7 we replicate our primary analysis using this more restricted sample and find results that are very similar to our preferred estimates.

As a placebo check, we perform our primary regression analysis using the same sample

construction methods on the non-Election days leading up to and after the actual Election Day. Specifically, we repeat the regression used in Table 2, Panel A, Column 1 for each of these days. Appendix Figure A8 shows the coefficients for each date. We find that none of these alternative dates produces a positive coefficient, suggesting that our approach likely identifies a lower bound on the racial gap in wait times.

As a final robustness/validation, we correlate the racial disparities in wait times that we identify using the smartphone data with the racial disparities in wait times found using the CCES survey (discussed in the previous section). As we found when correlating our overall wait time measure with the CCES, there is a strong correlation at the state level (0.72). The correlation at the congressional district level is much more modest (0.07).

[FIGURE 4 ABOUT HERE]

5 Discussion and Conclusion

Exploiting a large geospatial dataset, we provide new, nationwide estimates for the wait times of voters during the 2016 U.S. presidential election. In addition to describing wait times overall, we document a persistent racial disparity in voting wait times: areas with a higher proportion of black (and to a lesser extent Hispanic) residents are more likely to face long wait times than areas that are predominantly white. These effects survive a host of robustness and placebo tests and are also validated by being strongly correlated with survey data on voter wait times.

While the primary contribution of our paper is to carefully document voting wait times and disparities at the national level, it is natural to ask why these disparities exist. In the Appendix, we explore the mechanism and do not find conclusive evidence in favor of arrival bunching, partisan bias, early voting, or strict ID laws. We find suggestive evidence that the effects could be driven by fewer resources that leads to congestion especially in high-volume

polling places. We are left with the fact that these racial disparities are not limited to just a few states or areas with particular laws or party affiliations that might reflect strategic motivations. Rather, there is work to be done in a diverse set of areas to correct these inequities. A simple explanation is that government officials in general tend to focus more attention on areas with white constituents at the expense of those with black constituents. For example, this could be due to politicians being more responsive to white voters' complaints about voting administration than those from black voters (and relatedly, white voters lodging more complaints), in line with prior work demonstrating lower responsiveness to black constituents across a variety of policy dimensions (e.g. Butler and Broockman 2011; Giulietti et al. 2019; White et al. 2015).

Our results also demonstrate that smartphone data may be a relatively cheap and effective way to monitor and measure progress in both overall wait times and racial disparities in wait times across various geographic areas. The analysis that we conduct in this paper can be easily replicated after the 2020 election and thereby generate a panel dataset of wait times across areas. Creating a panel dataset across the country may be useful to help pin down the mechanism for disparities (e.g. using difference-in-differences designs to test if disparities in voter wait times change when different laws or election officials take over in a state). We hope that future work can build on the results in this paper to provide a deeper understanding of disparities in voting wait times and their causes.

References

- ABRAMS, D. S., M. BERTRAND, AND S. MULLAINATHAN (2012): “Do judges vary in their treatment of race?” *Journal of Legal Studies*, 41, 347–384.
- ALESINA, A. AND E. L. FERRARA (2014): “A test of racial bias in capital sentencing,” *American Economic Review*, 104, 3397–3433.
- ALTONJI, J. G. AND R. M. BLANK (1999): “Race and Gender in the Labor Market,” in *Handbook of Labor Economics*, ed. by O. Ashenfelter and D. Card, vol. 3, 3143–3259.
- ALVAREZ, R. M., S. ANSOLABEHERE, A. BERINSKY, G. LENZ, C. S. III, AND T. HALL (2009): “2008 Survey of the Performance of American Elections Final Report,” *tech. rep. Caltech/MIT Voting Technology Project*.
- ALVAREZ, R. M., T. E. HALL, AND M. H. LLEWELLYN (2008): “Are Americans confident their ballots are counted?” *Journal of Politics*, 70, 754–766.
- ANSOLABEHERE, S. AND B. F. SCHAFFNER (2017): “CCES Common Content, 2016,” .
- ATKESON, L. R. AND K. L. SAUNDERS (2007): “The effect of election administration on voter confidence: A local matter?” in *PS - Political Science and Politics*, vol. 40, 655–660.
- BAUER, R. F., B. L. GINSBERG, B. BRITTON, J. ECHEVARRIA, T. GRAYSON, L. LOMAX, M. C. MAYES, A. MCGEEHAN, T. PATRICK, AND C. THOMAS (2014): “The American Voting Experience: Report and Recommendations of the Presidential Commission on Election Administration,” 1–71.
- BERTRAND, M. AND E. DUFLO (2017): “Field Experiments on Discrimination,” in *Handbook of Economic Field Experiments*, vol. 1, chap. 8, 309–393.

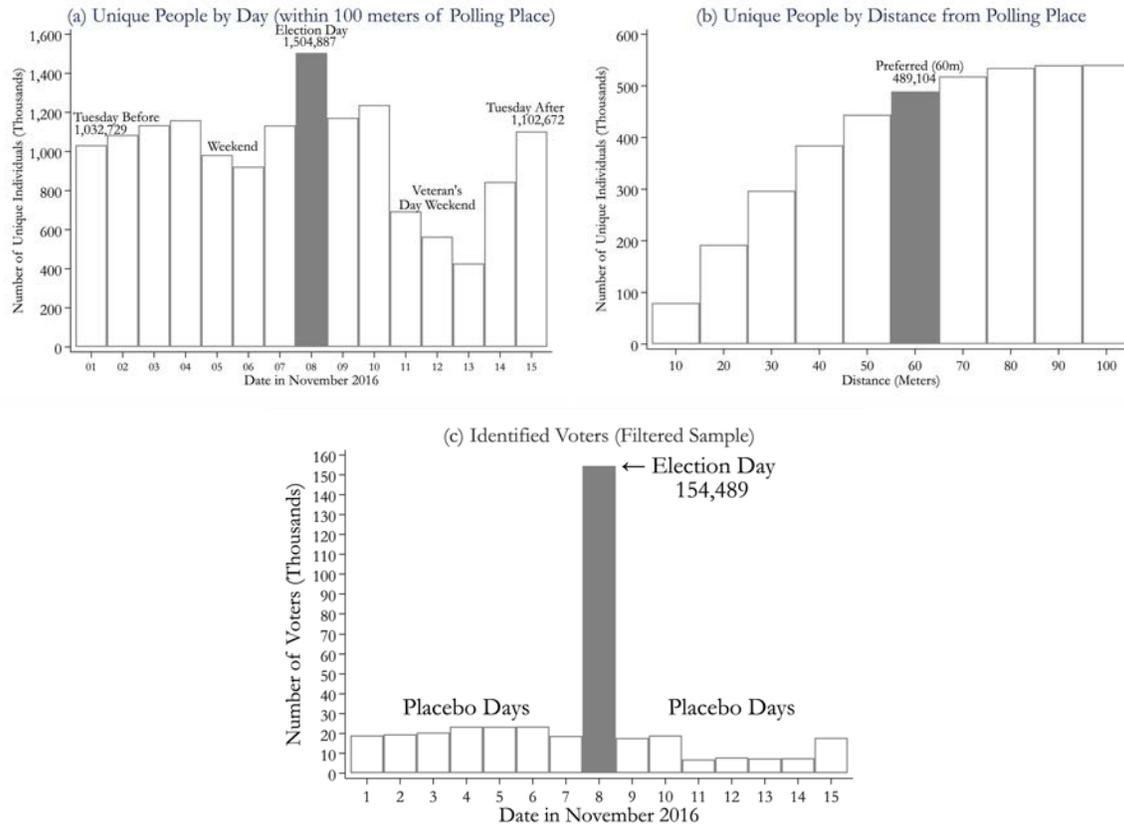
- BOWLER, S., T. BRUNELL, T. DONOVAN, AND P. GRONKE (2015): “Election administration and perceptions of fair elections,” *Electoral Studies*, 38, 1–9.
- BUTLER, D. M. AND D. E. BROCKMAN (2011): “Do politicians racially discriminate against constituents? A field experiment on state legislators,” *American Journal of Political Science*, 55, 463–477.
- CANTONI, E. AND V. PONS (2019): “Strict ID Laws Don’t Stop Voters: Evidence from a U.S. Nationwide Panel, 2008 - 2016,” *NBER Working Paper 25522*, 48.
- CHANDRA, A., A. FINKELSTEIN, A. SACARNY, AND C. SYVERSON (2016): “Health care exceptionalism? Performance and allocation in the US health care sector,” *American Economic Review*, 106, 2110–2144.
- CHARLES, K. K. AND J. GURRYAN (2011): “Studying Discrimination: Fundamental Challenges and Recent Progress,” *Annual Review of Economics*, 3, 479–511.
- CHEN, K. AND D. POPE (2019): “Geographic Mobility in America: Evidence from Cell Phone Data,” *Working Paper*.
- CHEN, M. K. AND R. ROHLA (2018): “The Effect of Partisanship and Political Advertising on Close Family Ties,” *Science*, 360, 1020–1024.
- CHETTY, R. AND N. HENDREN (2018): “The Impacts of Neighborhoods on Intergenerational Mobility II: County-level Estimates,” *The Quarterly Journal of Economics*, 133, 1163–1228.
- EDELSTEIN, W. A. AND A. D. EDELSTEIN (2010): “Queuing and Elections: Long Lines, DREs and Paper Ballots,” *Proceedings of the 2010 Electronic Voting Technology Workshop*.
- ELECTION ASSISTANCE COMMISSION (2017): “EAVS Deep Dive: Poll Workers and Polling Places,” 1–5.

- FAMIGHETTI, C., A. MELILLO, AND M. PÉREZ (2014): “Election Day Long Lines: Resource Allocation,” Tech. rep.
- GIULIETTI, C., M. TONIN, AND M. VLASSOPOULOS (2019): “Racial discrimination in local public services: A field experiment in the United States,” *Journal of the European Economic Association*, 17, 165–204.
- GLAESER, E. L. AND B. SACERDOTE (2003): “Sentencing in Homicide Cases and the Role of Vengeance,” *Journal of Legal Studies*, 32, 363–382.
- GRIMMER, J. AND J. YODER (2019): “The Durable Deterrent Effects of Strict Photo Identification Laws,” *Working Paper*.
- GROSS, D., J. F. SHORTIE, J. M. THOMPSON, AND C. M. HARRIS (2013): *Fundamentals of Queueing Theory: Fourth Edition*, John Wiley and Sons Inc.
- HERRON, M. C. AND D. A. SMITH (2014): “Race, Party, and the Consequences of Restricting Early Voting in Florida in the 2012 General Election,” *Political Research Quarterly*, 67, 646–665.
- (2016): “Precinct resources and voter wait times,” *Electoral Studies*, 42, 249–263.
- HIGHTON, B. (2006): “Long lines, voting machine availability, and turnout: The case of Franklin County, Ohio in the 2004 presidential election,” *PS - Political Science and Politics*, 39, 65–68.
- KAPLAN, E. AND H. YUAN (2019): “Early Voting Laws, Voter Turnout, and Partisan Vote Composition: Evidence from Ohio,” *American Economic Journal: Applied Economics* (Forthcoming).
- PETTIGREW, S. (2017): “The Racial Gap in Wait Times : Why Minority Precincts Are Underserved by Local Election Officials,” *Political Science Quarterly*, 132, 527–547.

- SPENCER, D. M. AND Z. S. MARKOVITS (2010): “Long Lines at Polling Stations? Observations from an Election Day Field Study,” *Election Law Journal: Rules, Politics, and Policy*, 9, 3–17.
- STEIN, R. M., C. MANN, C. S. III, Z. BIRENBAUM, A. FUNG, J. GREENBERG, F. KAWSAR, G. ALBERDA, R. M. ALVAREZ, L. ATKESON, E. BEAULIEU, N. A. BIRKHEAD, F. J. BOEHMKE, J. BOSTON, B. C. BURDEN, F. CANTU, R. COBB, D. DARMOFAL, T. C. ELLINGTON, T. S. FINE, C. J. FINOCCHIARO, M. D. GILBERT, V. HAYNES, B. JANSSEN, D. KIMBALL, C. KROMKOWSKI, E. LLAUDET, K. R. MAYER, M. R. MILES, D. MILLER, L. NIELSON, AND Y. OUYANG (2019): “Waiting to Vote in the 2016 Presidential Election: Evidence from a Multi-county Study,” *Political Research Quarterly*, 1–15.
- STEWART, C. AND S. ANSOLABEHERE (2015): “Waiting to Vote,” *Election Law Journal: Rules, Politics, and Policy*, 14, 47–53.
- STEWART III, C. (2013): “Waiting to Vote in 2012,” *Journal of Law & Politics*, 28, 439–464.
- WHITE, A. R., N. L. NATHAN, AND J. K. FALLER (2015): “What Do I Need to Vote? Bureaucratic Discretion and Discrimination by Local Election Officials,” *American Political Science Review*, 109, 129–142.

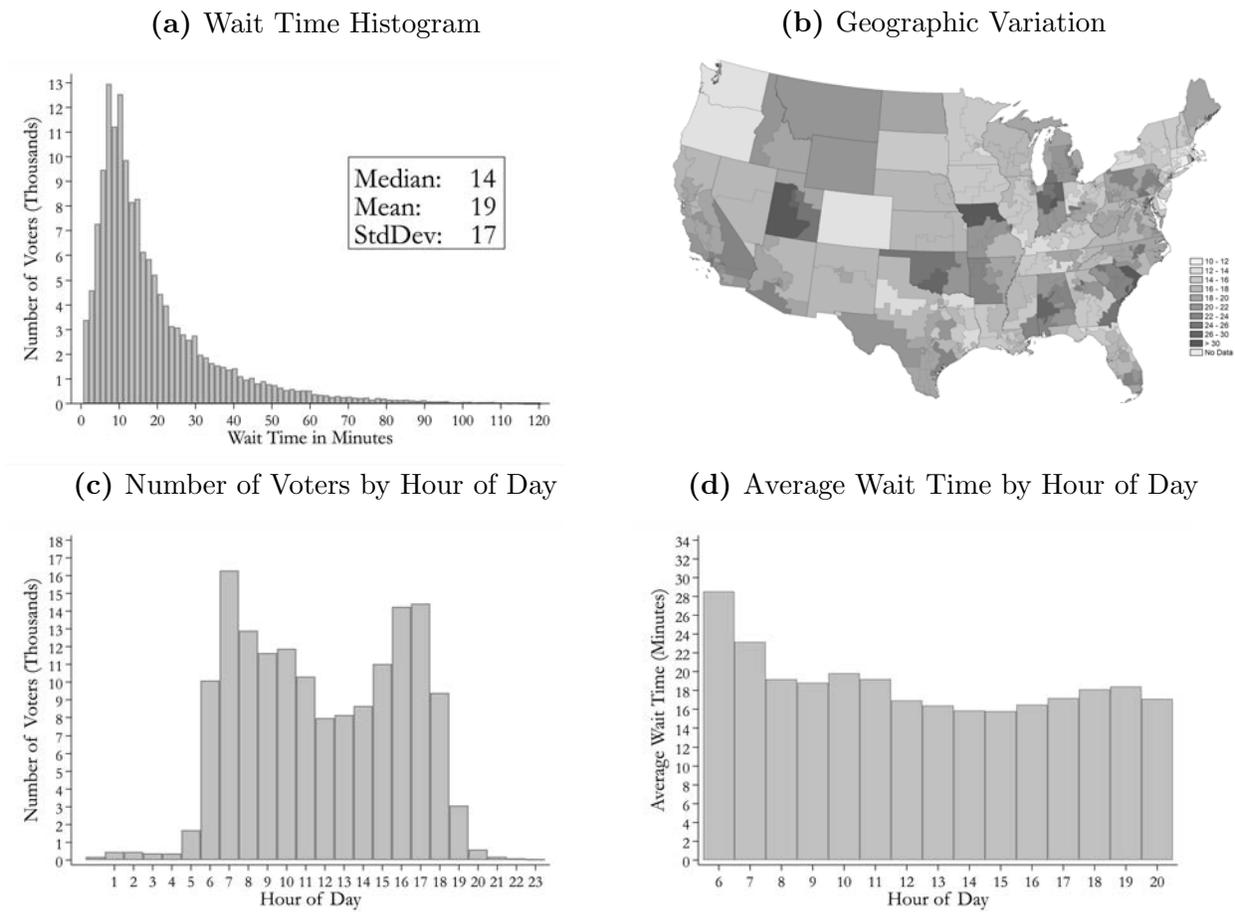
Tables and Figures

Figure 1: Defining the Radius



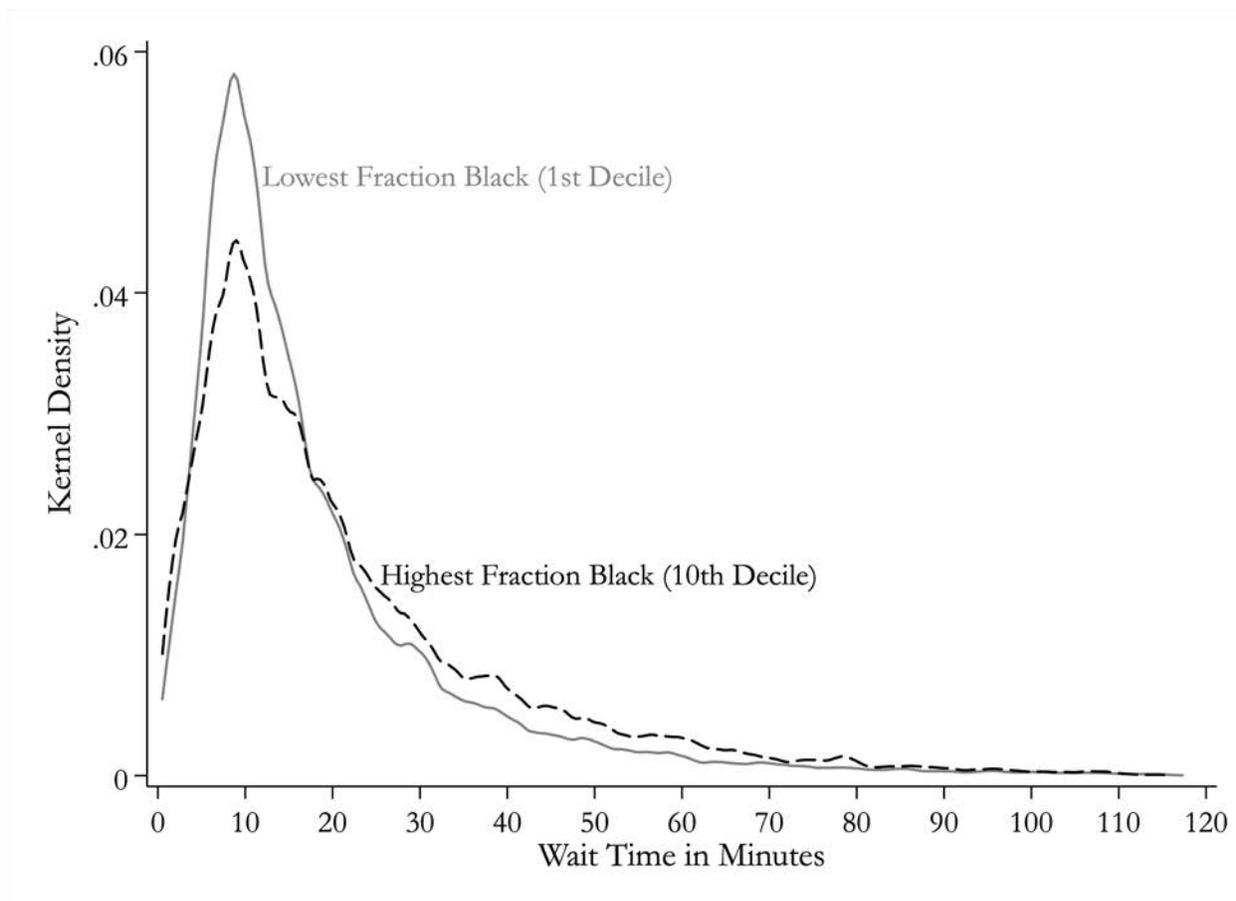
Notes: Panel A plots the number of unique device IDs observed within 100 meters of polling place building centroids on each day from November 1 to November 15 – Election Day (November 8) is shaded. Panel B plots the *difference* in the number of unique devices that are within a particular radius of the polling place (10 meters to 100 meters) on Election Day *relative to the average across all other days* – our final radius of 60 meters is highlighted in orange. Panel C shows the sample of unique devices that are observed within 60 meters of a polling place building centroid after applying the full set of filters – Election Day is shaded. Note that the Y-axes change across subfigures and that Veteran’s Day was on Friday, November 11 in 2016. The initial sample of smartphones that recorded at least one ping on Election Day (November 8, 2016) consisted of 5.2 million unique devices. As Panel A shows, there are 1.5 million devices once we limit to those that recorded at least one ping within 100 meters of a polling place on that date. Limiting to those within 60 meters of a polling place (the final radius used in Panel C) drops this to 1.0 million devices. Further limiting to phones that recorded at least one ping in the convex hull of the polling place building drops this to 406k devices, and limiting to phones that recorded a consistent set of pings on Election Day (1 per hour for 12 hours) drops to 307k devices. Imposing the remaining filters discussed in the text drops to the final sample of 155k observed in the shaded bar of Panel C.

Figure 2: Overall Wait Times



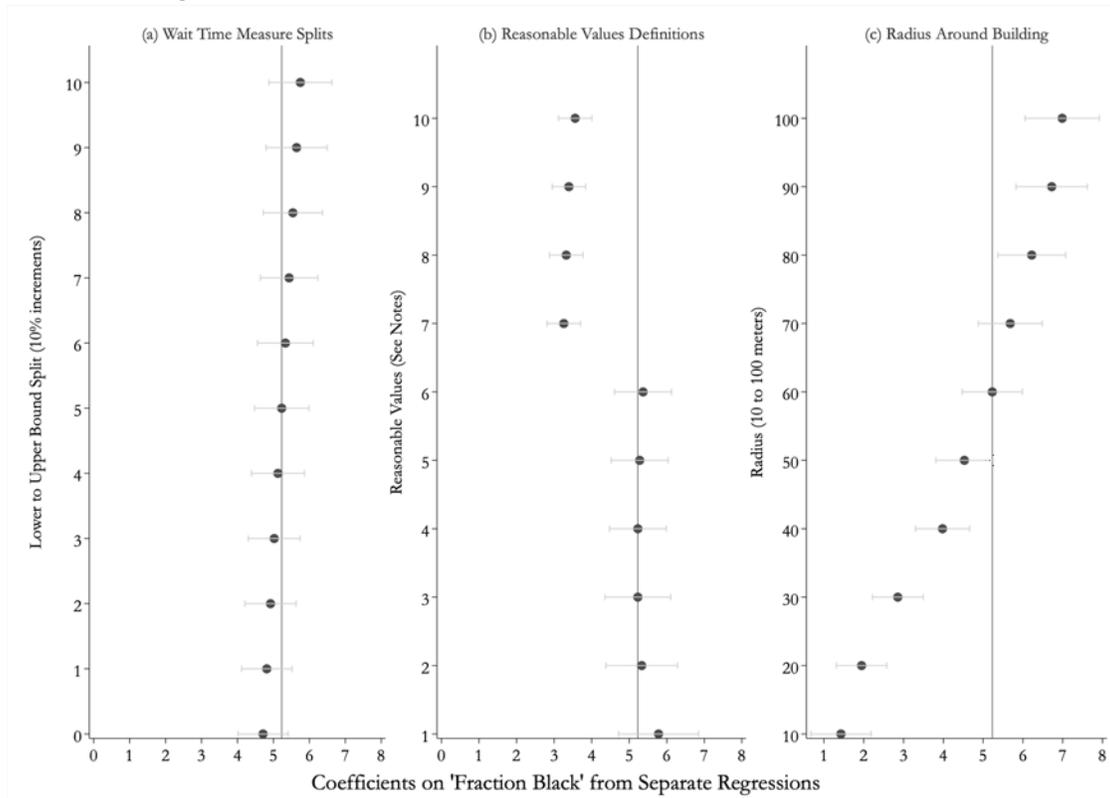
Notes: Panel A plots a histogram corresponding to the 154,495 cell phones who pass the filters used to identify likely voters (using 1.5 minute bins). Panel B shows variation in (empirical-Bayes-adjusted) average wait times by congressional district (115th Congress). Panel C plots the total number of voters (volume) by hour of arrival. Panel D plots the average wait time for each hour of arrival.

Figure 3: Wait Time: Fraction Black 1st vs. 10th Decile



Notes: Kernel densities are estimated using 1 minute half-widths. The 1st decile corresponds to the 34,420 voters across 10,319 polling places with the lowest percent of black residents (mean = 0%). The 10th decile corresponds to the 15,439 voters across the 5,262 polling places with the highest percent of black residents (mean = 58%).

Figure 4: Robustness to Different Data Construction Choices



Notes: Points correspond to coefficients on “Fraction Black” from separate regressions (± 1.96 robust standard errors, clustered at the polling place level). Unit of observation is a cellphone identifier on Election Day. All specifications are of the form used in Column 1 of Panel A, Table 1. Panel A varies the dependent variable across splits between the lower and upper bounds for our wait time measure (as described in Data and Methods); the first point ($y = 0$) corresponds to the lower bound, the last point ($y = 10$) corresponds to the upper bound measure, and all other points are intermediate deciles of the split (e.g. $y = 5$ corresponds to the midpoint of the two measures). Panel B varies the “reasonable values” (RV) filter, as follows: [RV1] Upper Bound under 5 hours ($N = 159,046$; Mean of Dependent Variable = 22.92) [RV2] Upper Bound under 4 hours ($N = 158,167$; Mean = 21.79) [RV3] Upper Bound under 3 hours ($N = 156,937$; Mean = 20.63) [RV4] Upper Bound under 2 hours ($N = 154,411$; Mean = 19.13) [RV5] Upper Bound under 2 hours and over 1.5 minutes ($N = 154,014$; Mean = 19.17) [RV6] Upper Bound under 2 hours and over 2 minutes ($N = 153,433$; Mean = 19.24) [RV7] Upper Bound under 1 hour and over 2 minutes ($N = 141,170$; Mean = 15.64) [RV8] Upper Bound under 1 hour and over 2.5 minutes ($N = 140,470$; Mean = 15.71) [RV9] Upper Bound under 1 hour and over 3 minutes ($N = 139,788$; Mean = 15.78) [RV10] Upper Bound under 1 hour and over 4 minutes ($N = 138,452$; Mean = 15.91). Panel C varies the bounding radius around the polling station centroid from 10 meters ($N = 60,821$; Mean = 12.09) up to 100 meters ($N = 113,797$; Mean = 21.81). The solid vertical line on each figure corresponds to the coefficient from the choice we use in our primary analysis, i.e. the midpoint wait time measure (Panel A), a filter of upper bounds under 2 hours (Panel B), and a radius of 60 meters (Panel C).

Table 1: Fraction Black and Voter Wait Time

	(1)	(2)	(3)	(4)	(5)
Panel A: Ordinary Least Squares (Y = Wait Time)					
Fraction Black	5.23*** (0.39)	5.22*** (0.39)	4.96*** (0.42)	4.84*** (0.42)	3.27*** (0.45)
Fraction Asian		-0.79 (0.72)	-2.48*** (0.74)	1.30* (0.76)	-1.10 (0.81)
Fraction Hispanic		1.15*** (0.37)	0.43 (0.40)	3.90*** (0.46)	1.50*** (0.50)
Fraction Other Non-White		12.01*** (1.94)	11.76*** (1.95)	1.66 (1.89)	2.04 (1.93)
N	154,411	154,411	154,260	154,260	154,260
R^2	0.00	0.00	0.01	0.06	0.13
DepVarMean	19.13	19.13	19.12	19.12	19.12
Polling Area Controls?	No	No	Yes	Yes	Yes
State FE?	No	No	No	Yes	Yes
County FE?	No	No	No	No	Yes
Panel B: Linear Probability Model (Y = Wait Time > 30min)					
Fraction Black	0.12*** (0.01)	0.12*** (0.01)	0.11*** (0.01)	0.10*** (0.01)	0.07*** (0.01)
Fraction Asian		-0.00 (0.02)	-0.04** (0.02)	0.04** (0.02)	-0.02 (0.02)
Fraction Hispanic		0.03*** (0.01)	0.01 (0.01)	0.08*** (0.01)	0.03*** (0.01)
Fraction Other Non-White		0.21*** (0.04)	0.21*** (0.04)	0.03 (0.04)	0.05 (0.04)
N	154,411	154,411	154,260	154,260	154,260
R^2	0.00	0.00	0.01	0.04	0.10
DepVarMean	0.18	0.18	0.18	0.18	0.18
Polling Area Controls?	No	No	Yes	Yes	Yes
State FE?	No	No	No	Yes	Yes
County FE?	No	No	No	No	Yes

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Robust standard errors, clustered at the polling place level, are in parentheses. Unit of observation is a cellphone identifier on Election Day. *DepVarMean* is the mean of the dependent variable. The dependent variable in Panel B is a binary variable equal to 1 if the wait time is greater than 30 minutes. *Polling Area Controls* includes the population, population per square mile, and fraction below poverty line for the block group of the polling station. “Asian” includes “Pacific Islander.” “Other Non-White” includes the “Other,” “Native American,” and “Multiracial” Census race categories.