Getting the Race Wrong: A Case Study of Sampling Bias and Black Voters in Online, Opt-In Polls

Daniel J. Hopkins[1] ⬤, William Halm[1], Melissa Huerta[2] and Josearmando Torres[2]

[1]Department of Political Science, University of Pennsylvania, Perelman Center for Political Science and Economics, Philadelphia, PA, USA and [2]University of Pennsylvania, Perelman Center for Political Science and Economics, Philadelphia, PA, USA
Corresponding author: Daniel J. Hopkins; Email: danhop@sas.upenn.edu

Abstract

Researchers are increasingly reliant on online, opt-in surveys. But prior benchmarking exercises employ national samples, making it unclear whether such surveys can effectively represent Black respondents and other minorities nationwide. This paper presents the results of uncompensated online and in-person surveys administered chiefly in one racially diverse American city—Philadelphia—during its 2023 mayoral primary. The participation rate for online surveys promoted via Facebook and Instagram was .4%, with White residents and those with college degrees more likely to respond. Such biases help explain why neither our surveys nor public polls correctly identified the Democratic primary's winner, an establishment-backed Black Democrat. Even weighted, geographically stratified online surveys typically underestimate the winner's support, although an in-person exit poll does not. We identify some similar patterns in Chicago. These results indicate important gaps in the populations represented in contemporary opt-in surveys and suggest that alternative survey modes help reduce them.

Keywords: Public opinion; online surveys; exit polls; black voters; response bias

Introduction

In recent years, surveys administered online with opt-in samples have come to dominate American survey research (Tourangeau, Conrad and Couper, 2013). Between 2018 and 2022, four leading political science journals together published 364 studies employing such samples.[1] Online, opt-in surveys offer key advantages: at comparatively affordable costs, they enable researchers to conduct survey experiments whose results often generalize well (Berinsky, Huber and Lenz, 2012; Mullinix *et al.*, 2015; Coppock, Leeper and Mullinix, 2018).

Compared to student samples, online opt-in samples are more representative (Druckman and Kam, 2011), and they facilitate panels and experiments, including those with complex designs (Broockman, Kalla and Sekhon, 2017; Yan, Kalla and Broockman, 2018; Hopkins and Gorton 2024a). However, they are subject to various sampling biases. Here, we use data from two U.S. cities to inquire about the ability of online, uncompensated surveys to represent Black voters and other residents of racially diverse cities. Of the journal articles mentioned above, 78 reported results for non-White ethnic/racial groups, including 47 specifically for Black Americans. Already, key studies have employed online, opt-in samples to make inferences about Black and other non-White groups (e.g. Anoll, 2018; Banks, White and McKenzie, 2019; Phoenix, 2019; Bejarano *et al.*, 2021; Burch, 2022; Carter, Wong and Guerrero, 2022).

However, previous benchmarking exercises have focused on the population overall, not subgroups (Vavreck and Rivers 2008; Rivers and Bailey 2009; Keeter 2018; Kennedy *et al.* 2016; MacInnis *et al.* 2018; Enns and Rothschild 2022; but see Broockman, Kalla and Sekhon 2017; Barreto *et al.* 2018; Spry 2021). As a result, they have been driven primarily by biases among the more numerous White respondents. Also, prior research benchmarking online, opt-in surveys typically considers accuracy with respect to marginal estimates (what percentage of the population is Black?) rather than cross-tabs or conditional estimates (what percentage of the Black population has a college degree or backs centrist Democrats?). With respect to political polling specifically, Black voters' strong support for Democrats in national elections (White and Laird, 2020) may further reduce pollsters' need to accurately reflect within-group differences. As a consequence, relatively little is known about the representativeness of Black or other non-White groups in online, opt-in samples.[2] The image of Black public opinion that emerges from surveys with online, opt-in samples—and from most survey experiments—may be skewed.

Mayoral elections provide an opportunity to benchmark the performance of different survey methods in cases where there are large Black and other non-white populations—and where those populations may have heterogeneous preferences. Philadelphia, Pennsylvania is an excellent case in which to study these issues. It has the largest Black population share among America's ten largest cities (U.S. Census Bureau, 2023). Philadelphia is at once large enough to host highly competitive mayoral elections and small enough to permit extensive in-person surveying. Accordingly, this paper reports the results of efforts to survey prospective Philadelphia voters in the run-up to and during the May 2023 Democratic Primary. We employed two survey modes—in-person and online—and multiple sampling strategies (including Facebook and Civiqs) to gauge Philadelphia Democrats' candidate preferences and views on key issues pre-primary. We successfully administered 1,591 uncompensated surveys with 1,359 unique respondents between February and May 2023.

Although the surveys focus primarily on Philadelphia, and in some cases on specific Philadelphia neighborhoods, our results provide important lessons for survey research generally. They highlight how both these polls and all public polls understated support for Cherelle Parker, the Black, establishment-backed mayoral candidate who ultimately won by approximately 10 percentage points. They also

underscore the remarkably limited sample who will take uncompensated, online polls, and the resulting sample biases. Our Facebook polls, for example, had a participation rate of .4%. That was the same rate we found in separate surveys of Chicago's 2023 mayoral election.

The results offer a cautionary tale: both polling modes over-represented White voters, and the online poll respondents in Philadelphia are overwhelmingly college-educated, even in a city where most adults and voters are not. Similar biases are evident in Chicago. Social scientists are increasingly employing online, opt-in surveys—including to study Black Americans and other non-White groups. These sampling biases have the potential to influence the conclusions and generalizability of various research projects. Additional analyses demonstrate that in Philadelphia, the central problem is not the under-representation of certain neighborhoods but the under-representation of certain demographic groups across neighborhoods. Our exit poll provides reasonably close approximations of the in-person votes cast at surveyed precincts, but it is less accurate when weighted and projected to the city as a whole. This reduced accuracy is likely a product of the non-random precincts which we were able to survey as well as the exit poll's inability to access voters who cast ballots by mail.

These results are limited in important ways which we detail below—they come from two large cities and from uncompensated surveys, meaning that they may not generalize straightforwardly to nationwide surveys and to incentivized surveys. In Philadelphia, they come from a closed party primary and so represent only Democrats. Moreover, the focus of this paper is on who participates in different types of uncompensated polls, not on weighting procedures for handling such data. However, the paper also tests the capacity of weighting to address the sample biases we document. Even after weighting based on demographics including race, gender, age, and neighborhood characteristics, we find that our online surveys significantly underestimated support for Cherelle Parker. Such polling biases were especially pronounced in heavily Black areas, suggesting that our surveys' inability to reach some groups of Black voters is one reason for their biases. These findings also demonstrate the potential advantages of supplementing online polls with in-person polls of hard-to-reach populations.

Methods and Context

In the spring of 2023, with the campaign for the Democratic nomination for Philadelphia Mayor underway, we conducted surveys using multiple modes and sampling frames to investigate vote choice and public opinion among Philadelphia Democrats. We employed two modes (face-to-face and online) using two sampling procedures in each case. In the in-person surveys, members of the research team first identified accessible, high-traffic sites (such as grocery stores, parks, and farmers' markets) to survey beginning on February 25th, 2023 and ending on May 6th, 2023 (n = 296). In addition, our research team surveyed respondents at six separate polling locations during the primary itself as an exit poll, netting a total of 196 additional surveys.[3] Note that while the online surveys were designed to be representative of voters citywide, that was not possible with the in-person surveys. See the Appendix for question wording.

We also conducted surveys online, recruiting respondents via two methods. The first method consisted of posting advertisements on Facebook targeting individuals living in randomly selected ZIP codes[4] and then inviting them to complete a brief, anonymous survey via a Qualtrics link (n=484).[5] In total, across our two waves of Facebook data collection, 108,220 unique individuals identified as Philadelphia residents by Facebook saw a total of 319,570 advertisements, making the participation rate .4%. Additionally, we administered a parallel Facebook/ Instagram poll in Chicago's 2023 mayoral run-off, with surveys administered between March 18th and 29th, 2023. There, too, the participation rate was .4%, suggesting that the University of Pennsylvania branding did not meaningfully change the overall participation rate in its home city.

In both Philadelphia and Chicago, our Facebook/Instagram surveys were administered to targeted ZIP codes which were randomly selected. As of 2021, Black Americans (74%) and Hispanic Americans (72%) reported using Facebook at somewhat higher rates than White Americans (67%) (Pew Research Center, 2021), making it a promising avenue through which to survey these groups. Given that both cities have high levels of racial segregation (Loh, Coes and Buthe, 2020), such stratified sampling ensures that the online advertisements inviting respondents to take the surveys were seen by many Black residents in the two jurisdictions. In Philadelphia, for example, the advertisements reached 49,795 unique individuals in majority Black ZIP codes, while in Chicago the number was 4,342.

In a second online polling effort, we partnered with online survey firm Civiqs to survey Philadelphia Democrats. Civiqs empanels respondents to take free, online surveys via advertisements placed on Google, Facebook, and other sites and subsequently emails invitations to take online surveys (Hopkins and Gorton 2024b). These surveys were conducted in two waves from March 17th, 2023 to March 27th, 2023 (n = 298) and then May 8th, 2023 to May 15th, 2023 (n = 317).[6]

On May 16th, 2023, Philadelphia Democrats nominated Cherelle Parker as their mayoral candidate in a closed primary that had featured several credible candidates and more than 30 million dollars in campaign spending (Walsh, Orso and Shukla, 2023). That Parker won was unsurprising: she was a veteran politician with significant support from local Democratic leaders and select unions, and she was the most credible Black candidate in a primary where at least 50% of the votes were likely to be cast in heavily Black precincts (Tannen, 2023). However, Parker bested her nearest competitor, City Controller Rebecca Rhynhart, by almost 10 percentage points. That took observers by surprise, as no public poll had shown Parker ahead (see Appendix Table A1). Chicago's 2023 mayoral election had some similarities in that regard. Brandon Johnson, a progressive Black Democrat, won despite polls underestimating his support in both rounds of that non-partisan election (see Appendix Table A3).

Results

Table 1 reports the unweighted fraction of respondents reporting different racial backgrounds by survey type for Philadelphia. Given that votes from heavily Black neighborhoods accounted for 47% of all 2019 Democratic primary votes (Shukla and Terruso, 2023), Black voters are under-represented across all survey modes.

Table 1. Respondents' racial self-identification by survey type. 2023 Census data derived from census.gov and provided for all Philadelphia County adults, not registered Democratic voters. The fractions for racial classification do not add up to one, as Hispanic people may identify with different races. Fractions of White and Black people in Philadelphia refer to those identifying as White or Black alone. L2 data are modeled from the Pennsylvania voter file for Democrats who voted in the 2023 primary

	Civiqs	Facebook (Eng.)	Facebook (Sp.)	In-Person	2023 Census	L2
Black	.172	.134	.000	.256	.430	.45
Hispanic	.016	.009	.450	.047	.161	.05
Other	.093	.317	.550	.130		
White	.719	.541	.000	.567	.446	.37

Table 2. Fraction college educated by survey type

	CPS	Facebook	Civics	In-Person
Asian Am.	.326	.750	1.000	.960
Black	.155	.725	.885	.581
Hispanic	.132	.750	1.000	.833
White	.379	.861	.924	.967

Although Black voters are overwhelmingly Democratic in national elections (White and Laird, 2020), that homogeneity can obscure important subgroup variation on other questions (Dawson, 2001; Banks, White and McKenzie, 2019; Phoenix, 2019; Jefferson, 2020). Table 2 provides the unweighted fraction of respondents from each broad ethnic/racial group who have at least a Bachelor's degree, with data from the 2018 to 2022 Current Population Survey (CPS; acquired via IPUMS) as a baseline.[7] The results make clear that the online survey modes over-represent college-educated respondents. While the levels of college attainment for in-person respondents are also very high relative to the population, they better reflect the targeted neighborhoods. That's because it was feasible to survey primarily in neighborhoods that were close to the University of Pennsylvania, and so we observe patterns that mirror those neighborhoods' demographics, with variation in educational attainment primarily among Black respondents. Nonetheless, in-person surveys can meaningfully increase the representation of Black respondents without college degrees relative to online surveys.

Table 3 presents respondent demographics for three types of surveys—Facebook, Civiqs, and in-person—compared to the estimated demographics of a random sample of 25,000 Philadelphia 2019 Democratic primary voters obtained via L2. (We use 2019 data because it was available before the primary took place.) The table allows us to observe not just the under-representation of Black and Hispanic respondents but also the sizeable over-representation of younger voters via in-person surveys.

Table 3. Means for various demographic categories for a random sample of Philadelphia Democratic 2019 primary voters as well as via multiple survey modes

	Voter File	Facebook	Civiqs	In-Person
Black	.517	.134	.179	.256
Hispanic	.062	.011	.016	.049
White	.318	.550	.719	.585
Male	.391	.384	.507	.388
Under 30	.046	.082	.039	.236
−39	.139	.071	.111	.209
−49	.139	.082	.072	.181
−59	.161	.108	.146	.108
−69	.230	.205	.242	.106

Participation Rates and ZIP demographics

To what extent is the under-representation of Black voters—especially those with lower educational attainment—a product of geographic biases in who takes surveys? If differing participation rates are simply a product of geography, survey researchers could over-sample people from neighborhoods likely to have lower participation rates.

Appendix Table A5 begins to answer this question by presenting the average fraction of respondents' ZIP codes that fall into each of six clusters of Philadelphia Democratic precinct types (Shukla and Terruso, 2023). Eighteen percent of all votes cast in the 2019 Democratic primary came from precincts labeled as heavily Black and establishment-oriented in their voting, yet such precincts account for only between 5% (in-person) and 8% (via Facebook) of surveyed respondents' ZIP codes. Notice that the Facebook (23%) and Civiqs (21%) surveys had somewhat lower averages in the "Black/not establishment" cluster than the city's electorate (29%), but the in-person surveys actually over-sampled from people in such areas (51%). (Our results below address whether such biases were consequential.) The in-person surveys targeted accessible locations, so it is unsurprising that they under-represent three precinct types.

We can also address the question of neighborhood-level response biases via our Facebook data, which was stratified by ZIP code. (See Appendix Figures A2 and especially A3 for the corresponding maps.) Table 4 presents linear regressions of the number of completed surveys and completed surveys divided by unique Facebook users who saw at least one advertisement for the survey. The independent variables are various demographic measures obtained from the American Community Survey (ACS) via Social Explorer. ZIP codes with more respondents with a bachelor's degree generate more surveys ($\beta = 57$) and have higher participation rates ($\beta = .024$), although neither relationship is significant at $p < .05$. Similarly, while ZIP codes' percentage of Black residents is negatively associated with completed surveys ($\beta = -58$) and lower participation rates ($\beta = -.034$), neither relationship is significant. Overall, the number of surveys by ZIP code and the participation rate

Table 4. For the targeted Philadelphia ZIP codes, this table displays linear regressions predicting the number of surveys and participation rates

	# of Surveys	# Surveys/Reach
Intercept	116.292	.066
	(137.155)	(.061)
Pct. White	−35.494	−.014
	(140.970)	(.063)
Pct. Black	−58.154	−.034
	(145.201)	(.065)
Pct. Hispanic	−12.506	−.013
	(111.039)	(.050)
Pct. Asian Am.	−70.202	−.058
	(157.349)	(.070)
Pct. with BA	56.699	.024
	(35.579)	(.016)
Pct. Moved In	−.823	−.024
	(205.368)	(.092)
Pct. Poverty	−115.988	−.059
	(66.058)	(.030)
Med. Hsh. Income	−96.739*	−.051*
	(38.087)	(.017)
R^2	.279	.357
Num. obs.	44	44

*$p < .05$.

are not very well predicted by available demographic measures. Across all kinds of ZIP codes, survey participation rates are low—the highest observed participation rate was 3%. That leaves limited variation for ZIP code-level predictors to explain. In Philadelphia, the challenge for survey researchers isn't so much getting some ZIP codes to respond, but getting people from most demographic groups to respond.

In Chicago, our Facebook/Instagram advertisements reached 17,617 individuals, 57 of whom completed the survey and gave a valid ZIP Chicago code. Of these respondents, 45 identified as white (79%), 6 identified as Hispanic (11%), and 5 identified as Black/African American (9%). The actual weighted average of Black residents in the selected ZIP codes was 25%, while for Hispanic residents it was 41%, meaning that these groups were meaningfully under-represented among survey respondents. As Appendix Table A4 illustrates, a series of bivariate regressions uncovers that within Chicago, response rates are higher in more heavily White ZIP codes.

Table 5. This table reports estimates of support for leading candidates in Philadelphia's 2023 Democratic mayoral primary weighted to 2019 Democratic primary participants still on the voter file. The actual vote shares include all vote methods (e.g. voting by mail, voting in-person)

	N	Parker	Rhynhart	Gym	Brown	Domb	*Sum(|eror|)*
All Surveys	1591	.120	.137	.188	.180	.106	.427
All Wave 1	883	.084	.082	.130	.233	.125	.642
All Wave 2	512	.171	.210	.210	.101	.064	.247
In-Person	296	.177	.040	.041	.242	.033	.752
Facebook Surveys	464	.095	.100	.104	.203	.283	.767
Facebook Wave 2	94	.219	.270	.062	.124	.053	.403
All Civiqs Surveys	615	.112	.175	.198	.156	.087	.384
Civiqs Wave 2	317	.187	.228	.236	.112	.068	.224
Exit Polls	196	.502	.233	.231	.022	.009	.356
Actual		.327	.228	.220	.088	.113	

Estimating Candidate Vote Shares

Even deliberate efforts to target Black voters will not always result in representative survey samples. But what are the consequences of these biases for surveys' capacity to accurately reflect public opinion? To answer that question, we consider a case in which there is a real-world benchmark—voter preferences in Philadelphia's May 2023 Democratic mayoral primary.

Appendix Table A7 provides unweighted results by candidate for each survey mode/wave—and demonstrates that even in our wave 2 surveys in April and May (the month of the primary election), our surveys significantly under-represented support for Cherelle Parker, a relatively moderate Black Democrat with significant support from the local Democratic Party.

Table 5 reports the results from various subsets of our surveyed population when weighting each subset to the 2019 Democratic primary voters who remained on the voter file in 2023.[8] For example, when we weight all respondents across the different modes, we find that Helen Gym (.188) and Jeff Brown (.180) had the highest levels of support overall in our surveys, followed by Rebecca Rhynhart (.137) and then eventual winner Cherelle Parker (.120). The table's final column reports the sum of the absolute error, and so illustrates which of our surveys more accurately reproduced the primary results.

The table's second and third rows demonstrate that some underestimation is the product of over-time changes in preferences: Parker's support more than doubled between waves 1 (2/20–4/5/23) and 2 (4/18–5/15/23), suggesting sizeable campaign effects.[9] Overall and looking at Facebook and Civiqs, we see that in every case, the total absolute error is lower for second-wave data. In fact, the Civiqs wave-two data has the best performance, in part because it more effectively estimated support for Allan Domb and Jeff Brown, whose backers were concentrated in more moderate White areas that were not well covered by our in-person surveys. Meanwhile, the

in-person surveys estimate Parker's support to be meaningfully higher (.177) than the Facebook surveys (.095) or Civiqs surveys (.112). The strong performance of all the wave-two polls jointly—which were pre-election polls in-person, via Civiqs, and via Facebook—suggests the value of combining multiple polling methods.

However, even the second-wave data underestimates Parker's primary-day support. According to Table 5, the only surveys that overestimate Parker's support are the weighted exit polls; see also Appendix Figure A4 and Appendix Table A10, which reports raw, unweighted results only for in-person voters at select, non-random precincts. (In fact, Appendix Table A10 shows that for in-person voters at a set of non-randomly selected precincts, the exit poll measured Parker's support quite well.[10]) Jointly, these results indicate that in-person exit polling was relatively effective at capturing Parker's support, meaning that the underestimation of Parker's support across other polling methods was not because her supporters were wholly unwilling to take surveys.

The exit polls' over-estimation of Parker's support in Table 5—that is, when weighting to the demographics of primary voters citywide—is instructive in another way. It indicates that even though our in-person Black respondents were less likely to come from areas that traditionally backed establishment candidates like Parker, that fact alone cannot explain why most of our surveys underestimate Parker's support; the exit polls were also concentrated in less pro-establishment Black neighborhoods and yet do not show the same bias after weighting. Our exit polls and to a lesser extent our in-person surveys seem to have accounted for some respondents missed by online sampling, again suggesting the benefit of combining online and in-person survey administration.

One limitation of this paper is that the sampling frame for the in-person surveys is not citywide. However, in Table A11, we provide a comparison of candidate support in our in-person surveys to expected primary support given in-person respondents' ZIP codes.[11] It demonstrates the continued under-representation of Cherelle Parker supporters, who account for only 15 percent of exit poll respondents but 30 percent of all votes cast in survey respondents' ZIP codes.[12] Nonetheless, the exit poll proves more accurate than either wave 1 or wave 2 in-person polling even accounting for the geographic distribution of respondents.

Voters' Racial Backgrounds and Survey Accuracy

It is clear that our surveys underestimated support for Cherelle Parker and that in-person exit polls did better on that score. In our final empirical analysis, we turn to a more difficult problem: what evidence is there that the 2023 Philadelphia surveys were biased because they missed Black voters specifically? To address that question, we employed logistic regression models with the same independent variables as in Appendix Table A6 along with data from all survey respondents to predict voters' average support for each of the five major candidates. In this case, we use an L2 Pennsylvania voter file from October 2023. These data were not available before the primary, but they enable us to perform a retrospective analysis in which we identify the (modeled) demographics of precisely the 248,675 registered Democrats who voted in the May 2023 Mayoral primary and weight accordingly. We predicted each Democratic voter's support for the five major candidates using their demographics

Table 6. This table reports predicted and actual vote shares and their differences for the five major candidates in all ZIP codes (columns 1–3), majority non-Hispanic Black ZIP codes (columns 4–6), and other ZIP codes (columns 7–9). The fractions are weighted by the number of votes cast for these candidates in each ZIP code

	All			Majority black			Other		
	Predict	Actual	Pred-Act	Predict	Actual	Pred-Act	Predict	Actual	Pred-Act
Parker	.159	.335	.176	.246	.525	.279	.101	.211	.109
Rhynhart	.218	.234	.015	.088	.118	.030	.303	.309	.006
Gym	.252	.226	.026	.274	.178	.096	.237	.257	.020
Domb	.144	.116	.029	.085	.063	.023	.183	.150	.033
J. Brown	.227	.090	.137	.306	.116	.190	.176	.073	.102

and five logistic regression models fit to the full survey data set. This gives us a probability of supporting each major candidate for every Democratic primary voter in Philadelphia and in essence serves to reweight our survey data to the actual voting population. We next aggregate these predictions to the ZIP code level to estimate the weighted fraction of voters who supported each of the five major candidates by ZIP code. Similarly, we also aggregated the actual precinct-level results by ZIP code as a point of comparison.

As Table 6's first three columns demonstrate, this procedure also underestimates the fraction voting for Parker, this time by an extraordinary 17.6 percentage points. What's more, the error is especially pronounced when we separate the data set into majority-Black ZIP codes (where 95,781 votes were cast for the five major candidates; columns 4–6) versus other ZIP codes (where 146,332 votes were cast for the same five candidates; columns 7–9). Similarly, Jeff Brown's support is over-estimated by 13.7 percentage points overall, an estimate which is 19 percentage points in majority-Black ZIP codes and just 10 percentage points elsewhere. Such evidence is not dispositive due to ecological inference problems. But it is strongly suggestive that the primary problem with the polls was not just their failure to measure Cherelle Parker's support, but their failure to measure her support among Black voters specifically.

Conclusion

No public poll before Philadelphia's primary had Cherelle Parker ahead (see Appendix Table A1); she won by 10 percentage points, garnering 43% more votes than her nearest competitor. The pre-election polls reported here did not have Cherelle Parker ahead, either. This evidence demonstrates that in part, the understatement of Parker's support was due to campaign dynamics, as it grew in the campaign's closing weeks. Still, that is an incomplete explanation. Both online and in-person polls just before the primary under-represented Black voters, especially those without college degrees, and consistently underestimated support for Parker. Even weighting does not fully remove the biases. These results also indicate that such biases are not simply a product of local geography, as they do not appear to

the same extent at the ZIP code level in Philadelphia. In Chicago, we have a limited sample of ZIP codes and only polling via Facebook/Instagram, but do see some evidence that response rates track neighborhood racial and educational demographics.

In recent years, survey researchers have labored to accurately estimate candidate support from online, opt-in polls, and have had considerable success. But most such efforts seek to estimate general-election support in partisan contexts where Black voters are strongly Democratic (White and Laird, 2020). As a result, many online, opt-in polls are not optimized to estimate differences of opinion *among* Black voters (Dawson, 2001; Jefferson, 2020), not to mention opinions among other non-White groups where rates of English proficiency are sometimes lower. It is also quite plausible that the balance between online, in-person, mail, and phone surveys may vary with geographic scope. Future work should consider whether mail survey invitations (see esp. Yan, Kalla and Broockman, 2018) can help address biases in sampling Black citizens and voters, and also investigate how combining surveys from multiple samples can reduce biases and better estimate heterogeneity in Black Americans' views. Our in-person exit poll performed well at estimating in-person voting in the precincts where it was deployed. Where feasible, future research should also employ the same geographic sampling frame for in-person and online modes which would facilitate comparisons (and the acquisition of baseline data). Subsequent work might also examine how the institution backing the survey influences respondents' willingness to take polls. Given that the University of Pennsylvania is well known in Philadelphia, polls with its name may influence the composition of respondents, improving participation rates among those connected to the university while possibly reducing them among residents most distrusting of it.

Notes

1 The journals include the *American Political Science Review*, *American Journal of Political Science*, *The Journal of Politics*, and *Public Opinion Quarterly*.
2 Research which has looked systematically at these questions for mail-to-online and phone surveys finds lower response rates for Black and Latino respondents via both modes (Broockman, Kalla and Sekhon, 2017).

3 Exit polling at polling places does not provide access to voters by mail; as Appendix Table A9 shows, 24.5% of 2023 primary voters in targeted precincts voted by mail. To address this, Appendix Tables A10 and A11 use only in-person votes as points of comparison.

4 The ZIP codes were randomly drawn from the population of registered Philadelphia voters. See Figure A1 for a sample Facebook advertisement.

5 We supplemented our English-language advertising with Spanish-language advertisements directing respondents to a Spanish-language survey (n = 20).

6 Of these respondents, 232 participated in both waves. Appendix Table A8 presents the transitions in candidate support across waves.

7 The CPS data is citywide; sub-city measures are not available.

8 Specifically, we drew a sample of 25,000 Philadelphia Democratic primary voters in the 2019 primary from L2 data, which was available at the time of the primary, and then fit logistic regression models to predict which demographic groups were over-represented in each subset. The model included both the demographic categories in Table 3 and the fraction of each ZIP code classified into each of the six clusters detailed in Shukla and Terruso (2023). See Appendix Table A6 for an example.

9 See also Appendix Table A8, which uses the two-wave Civiqs respondents to chart changes during the campaign.

10 At the same time, the exit polls under-estimate Brown and Domb's support, likely because the targeted precincts were not in the neighborhoods where they did best.

11 Specifically, we estimate the expected total share of the primary vote going to each candidate by aggregating precinct-level returns to the ZIP-code level and then weighting by the number of respondents per ZIP code.

12 This may partly reflect mode of voting, as the exit poll did not capture respondents who cast mail ballots. As Appendix Table A9 shows, between 16% and 32% of voters at the sampled precincts cast mail ballots.

References

Anoll AP (2018) What makes a good neighbor? Race, place, and norms of political participation. *American Political Science Review* **112**, 494–508.

Banks AJ, White IK and McKenzie BD (2019) Black politics: how anger influences the political actions blacks pursue to reduce racial inequality. *Political Behavior* **41**, 917–943.

Barreto MA *et al.* (2018) Best practices in collecting online data with Asian, Black, Latino, and White respondents: evidence from the 2016 collaborative multiracial post-election survey. *Politics, Groups, and Identities* **6**, 171–180.

Bejarano C *et al.* (2021) Shared identities: intersectionality, linked fate, and perceptions of political candidates. *Political Research Quarterly* **74**, 970–985.

Berinsky AJ, Huber GA and Lenz GS (2012) Evaluating online labor markets for experimental research: amazon. com's mechanical turk. *Political Analysis* **20**, 351–368.

Broockman DE, Kalla JL and Sekhon JS (2017) The design of field experiments with survey outcomes: a framework for selecting more efficient, robust, and ethical designs. *Political Analysis* **25**, 435–464.

Burch T (2022) Not all black lives matter: officer-involved deaths and the role of victim characteristics in shaping political interest and voter turnout. *Perspectives on Politics* **20**, 1174–1190.

Carter N, Wong J and Guerrero LG (2022) Reconsidering group interests: why Black Americans exhibit more progressive attitudes toward immigration than Asian Americans. *Du Bois Review: Social Science Research on Race* **19**, 257–274.

Coppock A, Leeper TJ and Mullinix KJ (2018) Generalizability of heterogeneous treatment effect estimates across samples. *Proceedings of the National Academy of Sciences* **115**, 12441–12446.

Dawson MC (2001) *Black Visions: The Roots of Contemporary African-American Political Ideologies.* Chicago: University of Chicago Press.

Druckman JN and Kam CD (2011) Students as experimental participants. *Cambridge Handbook of Experimental Political Science* **1**, 41–57.

Enns PK and Rothschild J (2022) Do you know where your survey data come from?. https://medium.com/3streams/surveys-3ec95995dde2 (accessed 20 June 2024).

Hopkins DJ and Gorton T (2024a) Unsubscribed and undemanding: partisanship and the minimal effects of a field experiment encouraging local news consumption. *American Journal of Political Science*. Forthcoming.

Hopkins DJ and Gorton T (2024b) On the internet, no one knows you're an activist: patterns of participation and response in an online, opt-in survey panel political research quarterly. Forthcoming.

Jefferson H (2020) The curious case of black conservatives: construct validity and the 7-point liberal-conservative scale. *Available at SSRN 3602209*. http://dx.doi.org/10.2139/ssrn.3602209

Keeter S (2018) The impact of survey non-response on survey accuracy. In Vannette DL and Krosnick JA (eds), *The Palgrave Handbook of Survey Research*. Cham: Springer International Publishing.

Kennedy C *et al.* (2016) *Evaluating Online Nonprobability Surveys*. Washington, DC: Pew Research Center.

Loh TH, Coes C and Buthe B (2020) *The Great Real Estate Reset*. Washington, DC: Smart Growth America.

MacInnis B *et al.* (2018) The accuracy of measurements with probability and nonprobability survey samples: replication and extension. *Public Opinion Quarterly* **82**, 707–744.

Mullinix KJ *et al.* (2015) The generalizability of survey experiments. *Journal of Experimental Political Science* **2**, 109–138.

Pew Research Center (2021) *News Consumption Across Social Media in 2021*. Washington, DC: Pew Research Center.

Phoenix DL (2019) *The Anger Gap: How Race Shapes Emotion in Politics*. Cambridge: Cambridge University Press.

Rivers D and Bailey D (2009) Inference from matched samples in the 2008 US national elections. *Proceedings of the Joint Statistical Meetings* **1**, 627–639.

Shukla A and Terruso J (2023) Philadelphia Inquirer. https://www.inquirer.com/politics/election/philadelphia-democratic-primary-2023-voting-groups-polls-20230403.html (accessed 20 June 2024).

Spry AD (2021) Design challenges and opportunities for studying race and ethnic politics. In Druckman JN and Green DP (eds), *Advances in Experimental Political Science*. Cambridge: Cambridge University Press.

Tannen J (2023) Turnout didn't decide the election. Preferences did. https://sixtysixwards.com/home/turnout-didnt-decide-the-election-preferences-did/ (accessed 20 June 2024).

Tourangeau R, Conrad FG and Couper MP (2013) *The Science of Web Surveys*. Oxford, UK: Oxford University Press.

U.S. Census Bureau (2023) U.S. Census Quick Facts. https://www.census.gov/quickfacts/ (accessed 25 June 2024).

Vavreck L and Rivers D (2008) The 2006 cooperative congressional election study. *Journal of Elections, Public Opinion and Parties* **18**, 355–366.

Walsh SC, Orso A and Shukla A (2023) The most expensive Philly mayor's race in history has now topped $31 million. Philadelphia Inquirer. https://www.inquirer.com/politics/election/mayors-race-campaign-finance-most-expensive-election-domb-rhynhart-20230506.html (accessed 20 June 2024).

White IK and Laird CN (2020) *Steadfast Democrats: How Social Forces Shape Black Political Behavior*. Princeton, New Jersey: Princeton University Press.

Yan A, Kalla J and Broockman DE (2018) Increasing response rates and representativeness of online panels recruited by mail: evidence from experiments in 12 original surveys. *SSRN Electronic Journal*. https://doi.org/10.31219/osf.io/t4msy