Unrepresentative Big Surveys Significantly Overestimate US Vaccine Uptake

Valerie C. Bradley[1*], Shiro Kuriwaki[2*], Michael Isakov[3],
Dino Sejdinovic[1], Xiao-Li Meng[4], Seth Flaxman[5,†]

[1]Department of Statistics, University of Oxford, Oxford, UK
[2]Department of Political Science, Stanford University, Stanford, CA, USA
[3]Harvard College, Harvard University, Cambridge, MA, USA
[4]Department of Statistics, Harvard University, Cambridge, MA, USA
[5]Department of Computer Science, University of Oxford, Oxford, UK
[*]These authors contributed equally to this work.
[†]Correspondence to: seth.flaxman@cs.ox.ac.uk.

October 2021

Surveys are a crucial tool for understanding public opinion and behavior, and their accuracy depends on maintaining statistical representativeness of their target populations by minimizing biases from all sources. Increasing data size shrinks confidence intervals but magnifies the impact of survey bias -- an instance of the Big Data Paradox[1]. Here we demonstrate this paradox in estimates of first-dose COVID-19 vaccine uptake in US adults: Delphi-Facebook[2,3] (about 250,000 responses per week) and Census Household Pulse[4] (about 75,000 per week). By May 2021, Delphi-Facebook overestimated uptake by 17 percentage points and Census Household Pulse by 14, compared to a benchmark from the Centers for Disease Control and Prevention (CDC). Moreover, their large data sizes led to negligible statistical uncertainty on the incorrect estimates. In contrast, an Axios-Ipsos online panel[5] with about 1,000 responses following survey research best practices[6] provided reliable estimates and uncertainty. We decompose observed error using a recent analytic framework[1] to explain the inaccuracy in the three surveys. We then analyze the implications for vaccine hesitancy and willingness. We show how a survey of 250,000 respondents can produce an estimate of the population mean that is no more accurate than an estimate from a simple random sample of size 10. Our central message is that data quality matters far more than data quantity, and compensating the former with the latter is a mathematically provable losing proposition.

Governments, businesses, and researchers rely on survey data to inform the provision of government services[7], steer business strategy, and guide response to the COVID-19 pandemic[8,9]. With the ever-increasing volume and accessibility of online surveys and organically-collected data, the line between traditional survey research and Big Data is becoming increasingly blurred[10]. Large datasets enable analysis of fine-grained subgroups, which are in high-demand

for designing targeted policy interventions[11]. However, counter to common intuition[12], larger sample sizes alone do not ensure lower error. Instead, small biases are *compounded* as sample size increases[1].

We see initial evidence of this in the discrepancies in estimates of first-dose COVID-19 vaccine uptake, willingness, and hesitancy from three online surveys in the US. Two of them — Delphi-Facebook's COVID-19 symptom tracker[2,3] ($n \approx 250{,}000$ per week and with over 4.5 million responses from January to May 2021) and the Census Bureau's Household Pulse survey[4] ($n \approx 75{,}000$ per survey wave and with over 600,000 responses from January to May 2021) — have large enough sample sizes to render standard uncertainty intervals negligible, yet report significantly different estimates of vaccination behavior with nearly identically-worded questions (Table 1). For example, Delphi-Facebook's state-level estimates for willingness to receive a vaccine from the end of March 2021 are 8.5 percentage points lower on average than those from the Census Household Pulse (Extended Data Fig. 1A), with differences as large as 16 percentage points.

The US Centers for Disease Control and Prevention (CDC) compiles and reports vaccine uptake statistics from state and local offices[13]. These figures serve as a rare external benchmark, permitting us to compare survey estimates of vaccine uptake to those from the CDC. The CDC has noted the discrepancies between their own reported vaccine uptake and that of the Census Household Pulse[14,15], and we find even larger discrepancies between the CDC and Delphi-Facebook data (Fig 1a). In contrast, the Axios-Ipsos' Coronavirus Tracker[5] ($n \approx 1{,}000$ responses per wave, and over 10,000 responses from January to May 2021) tracks the CDC benchmark well. None of these surveys use the CDC benchmark to adjust or assess their estimates of vaccine uptake, thus by examining patterns in these discrepancies, we can infer each survey's

accuracy and *statistical representativeness*, a nuanced concept that is critical for the reliability of survey findings[16–19].

The Big Data Paradox in vaccine uptake

We focus on the Delphi-Facebook and Census Household Pulse surveys because their large sample sizes (each greater than 10,000 respondents[20]) present the opportunity to examine the Big Data Paradox[1] in surveys. The Census Household Pulse is an experimental product designed to rapidly measure pandemic-related behavior. Delphi-Facebook has stated that the intent of their survey is to make comparisons over space, time, and subgroups and that point estimates should be interpreted with caution[3]. However, despite these intentions, Delphi-Facebook has reported point estimates of vaccine uptake in its own publications[11,21].

Delphi-Facebook and Census Household Pulse surveys persistently overestimate vaccine uptake relative to the CDC's benchmark (Fig. 1a). Despite being the smallest survey by an order of magnitude, Axios-Ipsos' estimates track well with the CDC rates (see Fig 1a), and their 95% confidence intervals contain the benchmark estimate from the CDC in 10 out of 11 surveys (an empirical coverage probability of 91%).

One might hope that estimates of changes in first-dose vaccine uptake are correct, even if each snapshot is biased. Unfortunately, errors have increased over time, from just a few percentage points in January 2021 to 4.2 percentage points (Axios-Ipsos), 14 percentage points (Census Household Pulse), and 17 percentage points (Delphi-Facebook) by mid-May 2021 (Fig. 1b). For context, for a state near the herd immunity threshold (70-80% based on recent

estimates[22]), a discrepancy of 10 percentage points in vaccination rates could be the difference between containment and uncontrolled exponential growth in new SARS-CoV-2 infections.

Conventional statistical formulas for uncertainty further mislead when applied to biased big surveys because as sample size increases, bias (rather than variance) dominates estimator error. Fig. 1a shows 95% confidence intervals for vaccine uptake based on each survey's reported sampling standard errors and weighting design effects[23]. Axios-Ipsos has the widest confidence intervals, but also the smallest design effects (1.1-1.2), suggesting that its accuracy is driven more by minimizing bias in data collection rather than post-survey adjustment. Census Household Pulse's 95% confidence intervals are widened by large design effects (4.4-4.8) but they are still too narrow to include the true rate of vaccine uptake in almost all survey waves. The confidence intervals for Delphi-Facebook are vanishingly small, driven by large sample size and moderate design effects (1.4-1.5), and give us essentially zero chance of being even close to the truth.

One benefit of such large surveys might be to compare estimates of spatial and demographic subgroups[24–26]. However, in March of 2021, Delphi-Facebook and Census Household Pulse over-estimated CDC state-level vaccine uptake by 16 and 9 percentage points, respectively (Extended Data Fig. 1G-H), and by equal or larger amounts by May 2021 (Extended Data Fig. 2G-H). Relative estimates were no better than absolute estimates in March of 2021: there is barely any agreement in a survey's estimated state-level rankings with the CDC (a Kendall rank correlation of 0.31 for Delphi-Facebook in Extended Data Fig. 1I, 0.26 for Census Household Pulse in Extended Data Fig. 1J) but they improved in May of 2021 (correlations 0.78 and 0.74 in Extended Data Fig. 2I-J). Among 18-64 year-olds, both Delphi-Facebook and Census Household Pulse overestimate uptake, with errors increasing over time (Extended Data Fig. 6).

These examples illustrate a mathematical fact. That is, when biased samples are large, they are doubly misleading: they produce confidence intervals with incorrect centers and substantially underestimated widths. This is the Big Data Paradox[1]: *the larger the data size, the surer we fool ourselves* when we fail to account for bias in data collection.

A framework for quantifying data quality

While it is well-understood that traditional confidence intervals capture only survey sampling errors[27] (and not total error), the traditional survey framework lacks analytic tools for quantifying nonsampling errors separately from sampling errors. A recently formulated statistical framework[1] permits us to exactly decompose total error of a survey estimate into three components:

$$\text{Total Error} = \textbf{Data Quality Defect} \times \textbf{Data Scarcity} \times \textbf{Inherent Problem Difficulty}. \quad (1)$$

This framework has been applied to COVID-19 case counts[28] and election forecasting[29]. Its full application requires ground-truth benchmarks or their estimates from independent sources[1].

Specifically, **Total Error** is the difference between the observed sample mean $\bar{Y}_n$ as an estimator of the ground truth, the population mean $\bar{Y}_N$. The **Data Quality Defect** is measured using $\hat{\rho}_{Y,R}$, called *data defect correlation* (*ddc*)[1], which quantifies total bias (from any source), measured by the correlation between the event that an individual's response is recorded and its value, Y. The impact of data quantity is captured by **Data Scarcity**, which is a function of the sample size n and the population size N, measured as $\sqrt{\frac{N-n}{n}}$, and hence what matters for error is

the relative sample size, i.e., how close n is to N, rather the absolute sample size n. The third factor is **Inherent Problem Difficulty**, which measures the population heterogeneity (via standard deviation σ_Y of Y), because the more heterogeneous a population is, the harder it is to estimate its average well. Mathematically, Equation 1 is given by $\bar{Y}_n - \bar{Y}_N = \hat{\rho}_{Y,R} \times \sqrt{\frac{N-n}{n}} \times \sigma_Y$. Incidentally, this expression was inspired by the Hartley-Ross inequality for biases in ratio estimators published in *Nature* in 1954.[30]. More details on the decomposition are provided in the Methods ("Calculation and interpretation of *ddc*"), where we also present a generalization for weighted estimators.

Decomposing error in COVID surveys

While the data defect correlation *ddc* is not directly observed, COVID-19 surveys present a rare case in which it can be deduced because all other terms in Equation 1 are known (see "Calculation and interpretation of *ddc*" in the Methods for an in-depth explanation). We apply this framework to the aggregate error shown in Fig. 1b, and the resulting components of error from the right-hand side of Equation 1 are shown in Fig. 1c-e.

We use the CDC's report of the cumulative count of first doses administered to US adults as the benchmark[8,13], $\bar{Y}_N$. This benchmark may suffer from administrative delays and slippage in how the CDC centralizes information from states[31–34]. As a sensitivity analysis to check the robustness of our findings to further misreporting, we present our results with sensitivity intervals under the assumption that CDC's reported numbers suffer from $\pm 5\%$ and $\pm 10\%$ error. These scenarios were chosen based on analysis of the magnitude by which the CDC's initial estimate for vaccine uptake by a particular day increases as the CDC receives delayed reports of

vaccinations that occurred on that day (Extended Data Fig. 3 and Supplementary Information A.2). That said, these scenarios may not capture latent systemic issues affecting CDC vaccination reporting.

The **Total Error** of each survey's estimate of vaccine uptake (Fig. 1b) increases over time for all studies, most markedly for Delphi-Facebook. The **Data Quality Defect**, measured by the *ddc*, also increases over time for Census Household Pulse and for Delphi-Facebook (Fig. 1c). The *ddc* for Axios-Ipsos is much smaller and steady over time, consistent with what one would expect from a representative sample. The **Data Scarcity** ($\sqrt{\frac{N-n}{n}}$) for each survey is roughly constant across time (Fig. 1d). **Inherent Problem Difficulty** is a population quantity common to all three surveys which peaks when the benchmark vaccination rate approaches 50% in April 2021 (Fig. 1e). Therefore, the decomposition suggests that the increasing error in estimates of vaccine uptake in Delphi-Facebook and Census Household Pulse is primarily driven by increasing *ddc*, which captures the overall impact of the bias in coverage, selection, and response.

Equation 1 also yields a formula for the bias-adjusted effective sample size n_{eff}, which is the size of a simple random sample that we would expect to exhibit the same level of Mean Square Error (MSE) as what was actually observed in a given study with a given *ddc*. Unlike the classical effective sample size[23], this quantity captures the impact of bias as well as that of an increase in variance from weighting and sampling. Details for this calculation are in Methods ("Error decomposition with survey weights").

For estimating the US vaccination rate, Delphi-Facebook has a bias-adjusted effective sample size of less than 10 in April 2021, a 99.99% reduction from the raw average weekly

sample size of 250,000 (Fig. 2). The Census Household Pulse also suffers from over 99% reductions in effective sample size by May 2021. A simple random sample would have controlled estimation errors by controlling *ddc*. However, once this control is lost, small increases in *ddc* beyond what is expected in simple random samples can result in drastic reductions of effective sample sizes for large populations[1].

Comparing study designs

Understanding *why* bias occurs in some surveys but not others requires an understanding of the sampling strategy, modes, questionnaire, and weighting scheme of each survey. Table 1 compares the design of each survey (more details in the "Additional survey methodology" in Methods and Extended Data Table 1).

All three surveys are conducted online and target the US adult population, but vary in respondent recruitment methods[35]. The Delphi-Facebook survey recruits respondents from active Facebook users (the Facebook Active User Base, or FAUB) using daily unequal-probability stratified random sampling[2]. The Census Bureau uses a systematic random sample to select households from the subset of the Census' Master Address File (MAF) for which they have obtained either cell phone or email contact information (approximately 81% of all households on the MAF)[4].

In comparison, Axios-Ipsos relies on inverse response propensity sampling from Ipsos' online KnowledgePanel. Ipsos recruits panelists using an address-based probabilistic sample from USPS's Delivery Sequence File (DSF)[5]. The DSF is similar to the Census' MAF. Unlike the Census Household Pulse, potential respondents are not limited to the subset for whom email

and phone contact information is available. Furthermore, Ipsos provides internet access and tablets to recruited panelists who lack home internet access. In 2021, this "offline" group typically comprises 1% of the final survey (Extended Data Table 1).

All three surveys weight on age and gender, i.e., assign larger weights to respondents of underrepresented age by gender subgroups and smaller weights to those of overrepresented subgroups[2,4,5] (Table 1) Axios-Ipsos and Census Household Pulse also weight on education and race/ethnicity. Axios-Ipsos additionally weights to the composition of political partisanship measured with the ABC News/Washington Post poll in 6 of the 11 waves we study. Education, a known correlate of propensity to respond to surveys[36] and social media use[37], are notably absent from Delphi-Facebook's weighting scheme, as is race/ethnicity. As noted before, none of the surveys use the CDC benchmark to adjust or assess estimates of vaccine uptake.

Explanations for error

Table 2 illustrates some consequences of these design choices. Axios-Ipsos samples mimic the actual breakdown of education attainment among US adults even before weighting, while those of Census Household Pulse and Delphi-Facebook do not. After weighting, Axios-Ipsos and Census Household Pulse match the population benchmark, by design. Delphi-Facebook does not explicitly weight on education, and hence the education bias persists in their weighted estimates: those without a college degree are underrepresented by nearly 20 percentage points. The story is similar for race/ethnicity. Delphi-Facebook's weighting scheme does not adjust for race/ethnicity, and hence their weighted sample still over-represents White adults by 8 percentage points, and under-represents Black and Asian proportions by around 50 percent of their size in the population (Table 2).

The overrepresentation of White adults and people with college degrees explains part of the error of Delphi-Facebook. The racial groups that Delphi-Facebook under-represents tend to be more willing and less vaccinated in the samples (Table 2). In other words, re-weighting the Delphi-Facebook survey to upweight racial minorities will bring willingness estimates closer to Household Pulse and the vaccination rate closer to CDC. The three surveys also report that people *without* a 4-year college degree are less likely to have been vaccinated compared to those *with* a degree (Table 2 and Supplemental Information Table 1). If we assume that vaccination behaviors do not differ systematically between non-respondents and respondents *within* each demographic category, under-representation of less-vaccinated groups would contribute to the bias found here. However, this alone cannot explain the discrepancies in all the outcomes. Census Household Pulse weights on both race and education[4] and still over-estimates vaccine uptake by over ten points in late May of 2021 (Fig. 1b).

Delphi-Facebook and Census Household Pulse may be unrepresentative with respect to political partisanship, which has been found to be correlated with vaccine behavior[38] and with survey response[39], and thus may contribute to observed bias. However, neither Delphi-Facebook nor Census Household Pulse collects partisanship of respondents; Census agencies are prohibited from asking about political preference. Moreover, no unequivocal population benchmark for partisanship exists.

Rurality may also contribute to the errors because it correlates with vaccine status[8] and home internet access[40]. Neither the Census Household Pulse nor Delphi-Facebook weights on sub-state geography, which may mean that adults in more rural areas are less likely to be vaccinated and also underrepresented in the surveys, leading to overestimation of vaccine uptake.

Axios-Ipsos weights to metropolitan status and also recruits a fraction of its panelists from an "offline" population of individuals without Internet access. We find that *dropping* these offline respondents ($n = 21$, or 1 percent of the sample) in their March 22, 2021 wave *increases* Axios-Ipsos' overall estimate of the vaccination rate by 0.5 percentage points, thereby increasing the total error (Extended Data Table 2). However, this offline population is simply too small to explain the entirety of the difference in accuracy between Axios-Ipsos and either the Census Household Pulse (6 percentage points) or Delphi-Facebook (14 percentage points), in this time period.

Careful recruitment of panelists is at least as important as weighting. Weighting on observed covariates alone cannot explain or correct the discrepancies we observe. For example, reweighting Axios-Ipsos March 22, 2021 wave using only Delphi-Facebook's weighting variables (age group and gender) increased the error in their vaccination estimates by 1 percentage point, but this estimate with Axios-Ipsos data is still more accurate than that from Delphi-Facebook during the same period (Extended Data Table 2). The Axios-Ipsos estimate with Delphi-Facebook weighting overestimated vaccination by 2 percentage points, whereas Delphi-Facebook overestimated it by 11 percentage points.

The key implication is that there is no silver bullet: every small part of panel recruitment, sampling, and weighting matters for controlling the data quality measured as the correlation between an outcome and response, what we call the *ddc*. In multi-stage sampling, which includes for instance the selection of participants followed by non-response, bias in even a *single* step can substantially impact the final result (see Methods "Population size in multi-stage sampling", Extended Data Table 3). A *total quality control* approach – inspired by the Total Survey Error framework[41] – is a better strategy than trying to prioritize some components over others in order

to improve data quality. This emphasis is merely a reaffirmation of the best practice for survey research as advocated by the American Association for Public Opinion Research[6]: "The quality of a survey is best judged not by its size, scope, or prominence, but by how much attention is given to [preventing, measuring and] dealing with the many important problems that can arise."[42]

Addressing common misperceptions

The three surveys discussed in this article demonstrate a seemingly paradoxical phenomenon – the two larger surveys that we studied are far more statistically confident, yet also far more biased, than the smaller, more traditional Axios-Ipsos poll. These findings are paradoxical only when we fall into the trap of the long-held, but incorrect, intuition that estimation errors necessarily decrease in larger datasets[12].

A limitation of our vaccine uptake analysis is that we only examine *ddc* with respect to an outcome for which a benchmark is available: first-dose vaccine uptake. One might hope that surveys biased on vaccine uptake are not biased on other outcomes, for which there may not be benchmarks to expose their biases. However, the absence of evidence of bias for the remaining outcomes is not evidence of its absence. In fact, mathematically, when a survey is found to be biased with respect to one variable, it implies that the entire survey fails to be *statistically representative*. The theory of survey sampling relies on statistical representativeness for all variables achieved via probabilistic sampling[43]. Indeed, Neyman's original introduction of probabilistic sampling showed the limits of purposive sampling, which attempted to achieve overall representativeness via enforcing it only on a set of variables[18,44].

In other words, when a survey loses its overall statistical representativeness (e.g., through bias in coverage or nonresponse), which is difficult to repair (e.g., via weighting or modeling on observable characteristics) and almost impossible to verify[45], researchers who wish to use the survey for scientific studies must supply other reasons to justify the reliability of their survey estimates, such as evidence about the independence between the variable of interest and the factors that are responsible for the unrepresentativeness. Furthermore, scientific journals that wish to publish studies based on unrepresentative surveys[17], especially those with large sizes such as Delphi-Facebook (biased with respect to vaccination status (Fig. 1), race and education (Table 2)) need to ask for reasonable effort from the authors to address the unrepresentativeness. A simple acknowledgment of the potential bias is insufficient for alerting about potentially seriously flawed findings, as we reveal in this article.

Some may argue that bias is a necessary trade-off for having data that is sufficiently large for conducting highly granular analysis, such as county-level estimation of vaccine hesitancy[26]. While high-resolution inference is important, we warn that this is a double-edged argument. A highly biased estimate with a misleadingly small confidence interval can do more damage than having no estimate at all. We further note that bias is not limited to population point estimates, but also affects estimates of changes over time (contrary to published guidance[3]) – both Delphi-Facebook and Census Household Pulse significantly overestimate the *slope* of vaccine uptake relative to that of the CDC benchmark (Fig. 1b).

The accuracy of our analysis also relies on the accuracy of the CDC's estimates of COVID vaccine uptake. However, if the selection bias in the CDC's benchmark is significant enough to alter our results, then that itself would be yet another example of the Big Data Paradox.

Discussion

This is not the first time that the Big Data Paradox has reared its head: Google Trends predicted more than twice the number of influenza-like illnesses than the CDC in February 2013[46]. This analysis demonstrates that the Big Data Paradox applies not only to organically-collected Big Data, like Google Trends, but also to surveys. Delphi-Facebook is "the largest public health survey ever conducted in the United States"[47]. The Census Household Pulse is conducted in collaboration between the US Census Bureau and eleven statistical government partners, all with enormous resources and survey expertise. Both studies take steps to mitigate selection bias, yet overestimate vaccine uptake by double digits. As we demonstrated, the impact of bias is magnified as relative sample size increases.

In contrast, Axios-Ipsos records only about 1,000 responses per wave, but makes additional efforts to prevent selection bias. Small surveys can be just as wrong as large surveys in expectation – of the three other small-to-medium online surveys additionally analyzed, two also miss the CDC vaccination benchmark (Extended Data Fig. 5). The overall lesson is that investing in data quality (particularly during collection, but also in analysis) minimizes error more efficiently than does increasing data quantity. Of course, a sample size of 1,000 may be too small (i.e., leading to unhelpfully large confidence intervals) for the kind of 50-state estimates given by big surveys. However, small area methods that borrow information across subgroups[48] can perform better with better quality, albeit small, data, and it is an open question whether that approach would outperform the large, biased surveys.

There are approaches to correct for these biases in both probability and nonprobability samples alike. For COVID-19 surveys in particular, since June 2021, the AP-NORC multi-mode

panel has weighted their COVID-19 related surveys to the CDC benchmark, so that the weighted *ddc* for vaccine uptake is zero by design[49]. More generally, there is an extensive literature on approaches for making inferences from data collected from nonprobability samples[50–52]. Other promising approaches include integrating surveys of varying quality[53,54], and leveraging the estimated *ddc* in one outcome to correct bias in others under several scenarios (Supplemental Information D).

While more needs to be done to fully examine the nuances of large surveys, organically collected administrative datasets, and social media data, we hope this first comparative study of *ddc* highlights the alarming implications of the *Big Data Paradox* – how large sample sizes magnify the impact of seemingly small defects in data collection, leading to overconfidence in incorrect inferences.

Main Text Tables

	Axios-Ipsos	Census Household Pulse	Delphi-Facebook
Recruitment mode	Address-based mail sample to Ipsos KnowledgePanel	SMS and email	Facebook Newsfeed
Interview mode	Online	Online	Online
Average size	1,000/wave	75,000/wave	250,000/week
Sampling frame	Ipsos KnowledgePanel; internet/tablets provided to ~5% of panelists who lack home internet	Census Bureau's Master Address File (individuals for whom email / phone contact information is available)	Facebook active users
Vaccine uptake question	"Do you personally know anyone who has already received the COVID-19 vaccine?"	"Have you received a COVID-19 vaccine?"	"Have you had a COVID-19 vaccination?"
Vaccine uptake definition	"Yes, I have received the vaccine"	"Yes"	"Yes"
Other vaccine uptake response options	"Yes, a member of my immediate family," "Yes, someone else," "No"	"No'	"No," "I don't know"
Weighting variables	Gender by age, race, education, Census region, metropolitan status, household income, partisanship.	Education by age by sex by state, race/ethnicity by age by sex by state, household size	Stage 1: age, gender "other attributes which we have found in the past to correlate with survey outcomes" to FAUB; Stage 2: state by age by gender

Table 1

	Composition of US Adults							Survey Estimates		
	Axios-Ipsos		Household Pulse		Delphi-Facebook		ACS	Household Pulse		
	Raw	Weighted	Raw	Weighted	Raw	Weighted	Benchmark	Vax	Will	Hes
Education										
High School	35%	39%	14%	39%	19%	21%	39%	39%	40%	21%
Some College	29	30	32	30	36	36	30	44	38	18
4-Year College	19	17	29	17	25	25	19	54	36	10
Post-Graduate	17	14	26	13	20	18	11	67	26	7
Race/Ethnicity										
White	71%	63%	75%	62%	74%	68%	60%	50%	33%	17%
Black	10	12	7	11	6	6	12	42	39	19
Hispanic	11	16	10	17	11	16	16	38	48	14
Asian			5	5	2	3	6	51	43	5

Table 2

Figure Legends

Fig 1 | Errors in estimates of vaccine uptake. **a.** Estimates of vaccine uptake for US adults in 2021 compared to CDC benchmark data, plotted by end date of each survey wave. Points indicate each study's weighted estimate of first-dose vaccine uptake, and intervals are 95% CIs using reported standard errors and design effects. Delphi-Facebook has $n = 4{,}525{,}633$ across 19 waves, Census Household Pulse has $n = 606{,}615$ across 8 waves, and Axios-Ipsos has $n = 11{,}421$ across 11 waves. Delphi-Facebook's CIs are too small to be visible. **b.** Total error $\bar{Y}_n - \bar{Y}_N$, **c.** data defect correlation $\hat{\rho}_{Y,R}$, **d.** data scarcity $\sqrt{\frac{(N-n)}{n}}$, **e.** inherent problem difficulty σ_Y. Shaded bands represent scenarios of +/-5% (darker) and +/-10% (lighter) error in the CDC benchmark relative to reported values (points). **b** - **e** comprise the decomposition in Equation 1.

Fig 2 | Bias-adjusted effective sample size. The bias-adjusted effective sample size of an estimate (different from the classic Kish effective sample size) is the size of a simple random sample which would have the same Mean Square Error as the observed estimate. Effective sample sizes are shown on the $\log_{10}$ scale. The original sample size was $n = 4{,}525{,}633$ across 19 waves for Delphi-Facebook, $n = 606{,}615$ across 8 waves for Census Household Pulse, $n = 11{,}421$ across 11 waves for Axios-Ipsos. Shaded bands represent scenarios of +/-5% error in the CDC benchmark relative to point estimates based on actual reported values.

Table 1 | Comparison of survey designs. Comparison of key design choices across Axios-Ipsos, Census Household Pulse, and Delphi-Facebook studies. All surveys target the US adult population. See Extended Data Table 1 for additional comparisons.

Table 2 | Composition of survey respondents by educational attainment and race/ethnicity. Axios-Ipsos: wave ending March 22, 2021, n = 995. Census Household Pulse: wave ending March 29, 2021, n = 76,068. Delphi-Facebook: wave ending March 27, 2021, n = 181,949. Benchmark uses the 2019 US Census American Community Survey (ACS), composed of roughly 3 million responses. Right-most column shows estimates of vaccine uptake (Vax), willingness (Will) and hesitancy (Hes) from the Census Household Pulse of the same wave.

Methods

Calculation and interpretation of *ddc*

The mathematical expression for Equation 1 is given here for completeness:

$$\bar{Y}_n - \bar{Y}_N = \hat{\rho}_{Y,R} \times \sqrt{\frac{N-n}{n}} \times \sigma_Y \qquad (2)$$

The first factor $\hat{\rho}_{Y,R}$ is called the *data defect correlation* (*ddc*)[1]. It is a measure of data quality represented by the correlation between the recording indicator R ($R = 1$ if an answer is recorded and $R = 0$ otherwise) and its value, Y. Given a benchmark, the *ddc* $\hat{\rho}_{Y,R}$ can be calculated by substituting known quantities into Equation 2. In the case of a single survey wave of a COVID-19 survey, n is the sample size of the survey wave, N is the population size of US adults from US Census estimates[55], $\bar{Y}_n$ is the survey estimate of vaccine uptake, and $\bar{Y}_N$ is the estimate of vaccine uptake for the corresponding period taken from the CDC's report of the cumulative count of first doses administered to US adults[8,13]. We calculate $\sigma_Y = \sqrt{\bar{Y}_N(1 - \bar{Y}_N)}$ because Y is binary (but Equation 2 is not restricted to binary Y).

We calculate $\hat{\rho}_{Y,R}$ by using *total* error $\bar{Y}_n - \bar{Y}_N$, which captures not only selection bias but also any measurement bias (e.g., from question wording). However, with this calculation method, $\hat{\rho}_{Y,R}$ lacks the direct interpretation as a correlation between Y and R, and instead becomes a more general index of data quality directly related to classical design effects (see Methods section "Bias-adjusted effective sample size").

It is important to point out that the increase in *ddc* does not necessarily imply that the response mechanisms for Delphi-Facebook and Census Household Pulse have changed over time. The correlation between a changing *outcome* and a steady response mechanism could change over time, hence changing the value of *ddc*. For example, as more individuals become vaccinated, and vaccination status is driven by individual behavior rather than eligibility, the correlation between vaccination status and propensity to respond could increase even if propensity to respond for a given individual is constant. This would lead to large values of *ddc* over time, reflecting the *increased impact* of the same response mechanism.

Error decomposition with survey weights

The data quality framework given by Equations 1 and 2 are a special case of a more general framework for assessing the actual error of a weighted estimator $\bar{Y}_w = \frac{\sum_i w_i R_i Y_i}{\sum_i w_i R_i}$, where w_i is the survey weight assigned to individual i. It is shown in Meng[1] that

$$\bar{Y}_{\mathrm{w}} - \bar{Y}_N = \hat{\rho}_{Y,R_{\mathrm{w}}} \times \sqrt{\frac{N-n_{\mathrm{w}}}{n_{\mathrm{w}}}} \times \sigma_Y, \qquad (3)$$

where $\hat{\rho}_{Y,R_{\mathrm{w}}} = \mathrm{Corr}(Y, R_{\mathrm{w}})$ is the finite population correlation between Y_i and $R_{\mathrm{w},i} = w_i R_i$ (over $i = 1, \ldots, N$). The "hat" on ρ reminds us that this correlation depends on the specific realization of $\{R_i, i = 1, \ldots, N\}$. The term n_{w} is the classical "effective sample size" due to weighting[2,3], i.e., $n_{\mathrm{w}} = \frac{n}{(1+CV_{\mathrm{w}}^2)}$, where CV_{w} is the coefficient of variation of the weights for all individuals in the observed sample, that is, the standard deviation of weights normalized by their mean. It is common for surveys to rescale their weights to have mean 1, in which case CV_w^2 is simply the sample variance of W.

When all weights are the same, Equation 3 reduces to Equation 2. In other words, the *ddc* term $\hat{\rho}_{Y,R_w}$ now also takes into account the impact of the weights as a means to combat the selection bias represented by the recording indicator R. Intuitively, if $\hat{\rho}_{Y,R} = \text{Corr}(Y, R)$ is high (in magnitude), then some Y_i's have a higher chance of entering our data set than others, thus leading to a sample average that is a biased estimator for the population average. Incorporating appropriate weights can reduce $\hat{\rho}_{Y,R}$ to $\hat{\rho}_{Y,R_w}$, with the aim to reduce the impact of the selection bias. However, this reduction alone may not be sufficient to improve the accuracy of $\bar{Y}_w$ because the use of weight necessarily reduces the sampling fraction $f = \frac{n}{N}$ to $f_w = \frac{n_w}{N}$ as well since $n_w < n$. Equation 3 precisely describes this trade off, providing a formula to assess when the reduction of *ddc* is significant to outweigh the reduction of the effective sample size.

Measuring the correlation between Y and R is not a new idea in survey statistics (though note that *ddc* is the population correlation between Y and R, not the sample correlation), nor is the observation that as sample size increases, error is dominated by bias instead of variance[56,57]. The new insight is that *ddc* is a general metric to index the *lack of* representativeness of the data we observe, regardless of whether or not the sample is obtained via a probabilistic scheme, or weighted to mimic a probabilistic sample. As discussed in the the section on addressing common misperception, any single *ddc* deviating from what is expected under representative sampling (e.g., probabilistic sampling) is sufficient to establish the sample is not representative (but the converse is not true). Furthermore, the *ddc* framework refutes the common belief that increasing sample size necessarily improves statistical estimation[1,58].

Bias-adjusted effective sample size

By matching the mean-squared error of $\bar{Y}_w$ with the variance of the sample average from simple random sampling, Meng[1] derives the following formula for calculating a *bias-adjusted effective sample size*, or n_{eff}:

$$n_{\text{eff}} = \frac{n_{\text{w}}}{N - n_{\text{w}}} \times \frac{1}{E\left[\hat{\rho}^2_{Y,R_{\text{w}}}\right]}$$

Given an estimator $\bar{Y}_{\text{w}}$ with expected total Mean Squared Error (MSE) T due to data defect, sampling variability, and weighting, this quantity n_{eff} represents the size of a simple random sample such that its mean $\bar{Y}_N$, as an estimator for the same population mean $\bar{Y}_N$, would have the identical MSE T. The term $E\left[\hat{\rho}^2_{Y,R_{\text{w}}}\right]$ represents the amount of selection bias (square) expected on average from a particular recording mechanism R and a chosen weighting scheme.

For each survey wave, we use $\hat{\rho}^2_{Y,R_{\text{w}}}$ to approximate $E\left[\hat{\rho}^2_{Y,R_{\text{w}}}\right]$. This estimation is unbiased by design, since we use an estimator to estimate its expectation. Therefore, the only source of error is the sampling variation, which is typically negligible for large surveys, such as for Delphi-Facebook and the Census Household Pulse surveys. This estimation error may have more impact for smaller traditional surveys, such as Axios-Ipsos' survey, an issue we will investigate in subsequent work.

We compute $\hat{\rho}_{Y,R_{\text{w}}}$ by using the benchmark $\bar{Y}_N$, namely, via solving Equation 3 for $\hat{\rho}_{Y,R_{\text{w}}}$,

$$\hat{\rho}_{Y,R_{\text{w}}} = \frac{Z_w}{\sqrt{N}}, \quad \text{where} \quad Z_w = \frac{\bar{Y}_{\text{w}} - \bar{Y}_N}{\sqrt{\frac{1 - f_{\text{w}}}{n_{\text{w}}}}\sigma_Y}.$$

We introduce this notation Z_w because it is the quantity that determines the well-known survey efficiency measure, the so-called *design effect*, which is the variance of Z_w for a probabilistic sampling design[23] (when we assume the weights are fixed). For the more general setting where $\bar{Y}_w$ may be biased, we replace the variance by MSE, and hence the bias-adjusted design effect $D_e = E[Z_w^2]$, which is the MSE relative to the benchmark measured in the unit of the variance of an average from a simple random sample of size n_w. Hence $D_I \equiv E\left[\hat{\rho}_{Y,R_w}^2\right]$, which was termed as *data defect index*[1], is simply the bias-adjusted design effect *per unit*, because $D_I = \frac{D_e}{N}$.

Furthermore, because Z_w is the standardized actual error, it captures any kind of error inherited in $\bar{Y}_w$. This observation is important because when Y is subject to measurement errors, $\frac{Z_w}{\sqrt{N}}$ no longer has the simple interpretation as a correlation. But because we estimate D_I by $\frac{Z_w^2}{N}$ directly, our effective sample size calculation is still valid even when Equation 3 does not hold.

Asymptotic behavior of *ddc*

As shown in Meng[1], for any probabilistic sample without selection biases, the *ddc* is on the order of $\frac{1}{\sqrt{N}}$. Hence the magnitude of $\hat{\rho}_{Y,R}$ (or $\hat{\rho}_{Y,R_w}$) is small enough to cancel out the impact of $\sqrt{N-n}$ (or $\sqrt{N-n_w}$) in the data scarcity term on the actual error, as seen in Equation 2 (or Equation 3). However, when a sample is unrepresentative, e.g. when those with $Y = 1$ are more likely to enter the dataset than those with $Y = 0$, then $\hat{\rho}_{Y,R}$ can far exceed $\frac{1}{\sqrt{N}}$ in magnitude. In this case, error will increase with $\sqrt{N}$ for a fixed *ddc* and growing population size N (Equation 2). This result may be counter-intuitive in the traditional survey statistics framework, which often

considers how error changes as sample size n grows. The *ddc* framework considers a more general setup, taking into account individual response behavior, including its impact on sample size itself.

As an example of how response behavior can shape both total error and the number of respondents n, suppose individual response behavior is captured by a logistic regression model

$$\text{logit}[\Pr(R = 1|Y)] = \alpha + \beta Y. \qquad (4)$$

This is a model for a response propensity score. Its value is determined by α, which drives the overall sampling fraction $f = \frac{n}{N}$, and by β, which controls how strongly Y influences whether a participant will respond or not.

In this logit response model, when $\beta \neq 0$, $\hat{\rho}_{Y,R}$ is determined by individual behavior, not by population size N. In Supplemental Information B.1, we prove that *ddc* cannot vanish as N grows, nor can the observed sample size n ever approach 0 or N for a given set of (finite and plausible) values of $\{\alpha, \beta\}$, because there will always be a non-trivial percentage of non-respondents. For example, an f of 0.01 can be obtained under this model for either $\alpha = -0.46, \beta = 0$ (no influence of individual behavior on response propensity), or for $\alpha = -3.9, \beta = -4.84$. However, despite the same f, the implied *ddc* and consequently the MSE will differ. For example, the MSE for the former (no correlation with Y) is 0.0004, while the MSE for the latter (a -4.84 coefficient on Y) is 0.242, over 600 times larger.

See Supplemental Information B.2 for the connection between *ddc* and a well-studied non-response model from econometrics, the Heckman selection model[59].

Population size in multi-stage sampling

We have shown that the asymptotic behavior of error depends on whether the data collection process is driven by individual response behavior or by survey design. The reality is often a mix of both. Consequently, the relevant "population size" N depends on when and where the representativeness of the sample is destroyed, i.e., when the individual response behaviors come into play. Real-world surveys that are as complex as the three surveys we analyze here have multiple stages of sample selection.

Extended Data Table 3 takes as an example the sampling stages of the Census Household Pulse, which has the most extensive set of documentation among the three surveys we analyze. As we have summarized (Table 1 and Extended Data Table 1 the Census Household Pulse (1) first defines the sampling frame as the reachable subset of the Master Address File, (2) takes a random sample of that population to prompt (send a survey questionnaire), and (3) waits for individuals to respond to that survey. Each of these stages reduces the desired data size, and the corresponding *population size* is the intended sample size from the prior stage (in notation, $N_s = n_{s-1}$, for $s = 2,3$). For example, for stage 3, the population size N_3 is the size of the intended sample size n_2 from the second stage, i.e., the sampling stage, because only the sampled individuals have a chance to respond.

Although all stages contribute to the grand *ddc*, the stage that dominates is the *first stage at which the representativeness of our sample is destroyed*— whose size will be labeled as the *dominating population size (dps)*—when the relevant population size decreases dramatically at each step. However, we must bear in mind that *dps* refers to the worse case scenario, when biases accumulate, instead of (accidentally) canceling each other out.

For example, if the 20 percent of the MAF excluded from the Census Household Pulse sampling frame (because they had no cell phone or email contact information) is not representative of the US Adult population, then the *dps* is N_1, or 255 million adults contained in 144 million households. Then the increase in bias for given *ddc* is driven by the rate of $\sqrt{N_1}$ where $N_1 = 2.55 \times 10^8$ and is large indeed (with $\sqrt{2.5 \times 10^8} \approx 15{,}000$). In contrast, if the the sampling frame is representative of the target population and the outreach list is representative of the frame (and hence representative of the US adult population) but there is non-response bias, then *dps* is $N_3 = 10^6$ and the impact *ddc* is amplified by the square root of that number ($\sqrt{10^6} = 1{,}000$). In contrast, Axios-Ipsos reports a response rate of about 50%, and obtains a sample of $n = 1000$, so the *dps* could be as small as $N_3 = 2000$ (with $\sqrt{2000} \approx 45$).

This decomposition is why our comparison of the surveys is consistent with the *Law of Large Populations* (estimation error increases with $\sqrt{N}$), *even though all three surveys ultimately target the same US Adult Population*. Given our existing knowledge about online-offline populations[40] and our analysis of Axios-Ipsos' small "offline" population, Census Household Pulse may suffer from unrepresentativeness at Stage 1 of Extended Data Table 3 where $N = 255$ million, and Delphi-Facebook may suffer from unrepresentativeness at the initial stage of starting from the Facebook User Base. In contrast, the main source of unrepresentativeness for Axios-Ipsos maybe at a later stage where the relevant population size is orders of magnitude smaller.

CDC estimates of vaccination rates

The CDC benchmark data used in our analysis was downloaded from the CDC's COVID data tracker[13]. We employ the cumulative count of people who have received at least one dose of COVID-19 vaccine reported in the "Vaccination Trends" tab. This data set contains vaccine uptake counts for all US residents (not only adults). However, the surveys of interest estimate vaccine uptake among adults. The CDC receives age-group-specific data on vaccine uptake from all states except for Texas on a daily basis, which is also reported cumulatively over time.

Therefore, we must impute the number of adults who have received at least one dose on each day. We assume Texas is exchangeable with the rest of the states in terms of the age distribution for vaccine uptake. Under this assumption, for each day, we use the age group vaccine uptake data from all states except for Texas to calculate the proportion of cumulative vaccine recipients who are 18 or older, then we multiply that number by the total number of people who have had at least one dose to estimate the number of US *adults* who have received at least one dose.

The CDC performs a similar imputation for the 18+ numbers reported in their COVID data tracker. However the CDC's imputed 18+ number is available only as a snapshot and not a historical time series, hence the need for our imputation. See Supplemental Information for details of the imputation implementation.

Additional survey methodology

The Census Household Pulse and Delphi-Facebook surveys are the first of their kind for each organization, while Ipsos has maintained their online panel for 12 years.

Question wording

All three surveys ask whether respondents have received a COVID-19 vaccine. See Extended Data Table 1. Delphi-Facebook and Census Household Pulse ask similar questions ("Have you had / received a COVID-19 vaccination / vaccine?"). Axios-Ipsos asks "Do you personally know anyone who has already received the COVID-19 vaccine?," and respondents are given response options including "Yes, I have received the vaccine." The Axios-Ipsos question wording might pressure respondents to conform to their communities' modal behavior and thus misreport their true vaccination status, or may induce acquiescence bias from the multiple "yes" options presented. This pressure may exist both in high- and low-vaccination communities, so its net impact on Axios-Ipsos' results is unclear. Nonetheless, Axios-Ipsos' question wording does differ from that of the other two surveys, and may contribute the observed differences in estimates of vaccine uptake across surveys.

Population of Interest

All three surveys target US adult population, but with different sampling and weighting schemes. Household Pulse sets the denominator of their percentages as the household civilian, non-institutionalized population in the United States of 18 years of age or older, excluding Puerto Rico or the island areas. Axios-Ipsos designs samples to representative of the US general adult population 18 or older. For Facebook, the US target population reported in weekly contingency tables is the US adult population, excluding Puerto Rico and other US territories. For the CDC Benchmark, we define the denominator as the US 18+ population, excluding Puerto Rico and other US territories. To estimate the size of the total US population, we use the US Census Bureau Annual Estimates of the Resident Population for the United States and Puerto

Rico, 2019[55]. This is also what the CDC uses as the denominator in calculating rates and percentages of the US population[60].

Axios-Ipsos and Delphi-Facebook generate target distributions of the US adult population using the Current Population Survey (CPS), March Supplement, from 2019 and 2018, respectively. Census Household Pulse uses a combination of 2018 1-year American Community Survey (ACS) estimates and the Census Bureau's Population Estimates Program (PEP) from July 2020. Both the CPS and ACS are well-established large surveys by the Census and the choice between them is largely inconsequential.

Axios-Ipsos Data

The Axios-Ipsos Coronavirus tracker is an ongoing, bi-weekly tracker intended to measure attitudes towards COVID-19 of adults in the US. The tracker has been running since March 13, 2020 and has released results from 45 waves as of May 28, 2021. Each wave generally runs over a period of 4 days. The Axios-Ipsos data used in this analysis was scraped from the topline PDF reports released on the Ipsos website[5]. The PDF reports also contain Ipsos' design effects, which we have confirmed are calculated as 1 plus the variance of the (scaled) weights.

Census Household Pulse Data

The Census Household Pulse is an experimental product of the US Census Bureau in collaboration with eleven other federal statistical agencies. We use the point estimates presented in Data Tables, as well as the standard errors calculated by the Census Bureau using replicate weights. The design effects are not reported, however we can calculate it as $1 + CV_{\text{w}}^2$, where CV_{w} is the coefficient of variation of the individual-level weights included in the microdata[23].

Delphi-Facebook COVID symptom survey

The Delphi-Facebook COVID symptom survey is an ongoing survey collaboration between Facebook, the Delphi Group at Carnegie Mellon University (CMU), and the University of Maryland[2]. The survey is intended to track COVID-like symptoms over time in the US and in over 200 countries. We use only the US data in this analysis. The study recruits respondents using a daily stratified random samples recruiting a cross-section of Facebook Active Users. New respondents are obtained each day, and aggregates are reported publicly on weekly and monthly frequencies. The Delphi-Facebook data used here was downloaded directly from CMU's repository for weekly contingency tables with point estimates and standard errors.

Ethical compliance

According to HRA decision tools (http://www.hra-decisiontools.org.uk/research/), our study is considered Research, and according to the NHS REC review tool (http://www.hra-decisiontools.org.uk/ethics/), we do not need NHS Research Ethics Committee (REC) review, as we used only (1) publicly available, (2) anonymized, and (3) aggregated data outside of clinical settings.

Data availability

Raw data is deposited in the Harvard Dataverse, at https://doi.org/10.7910/DVN/GKBUUK. Data was collected from publicly available repositories of survey data by downloading it directly or using APIs.

Code availability

Code to replicate the findings is available in the repository https://github.com/vcbradley/ddc-vaccine-US. The main decomposition of the *ddc* is available on the package "ddi" from the Comprehensive R Archive Network (CRAN).

Acknowledgments

We thank Frauke Kreuter, Alex Reinhart, and the Delphi Group at Carnegie Mellon University, Facebook's Demography and Survey Science group; Frances Barlas, Chris Jackson, Catherine Morris, Mallory Newall, and the Public Affairs team at Ipsos; and Jason Fields and Jennifer Hunter Childs at the US Census Bureau for productive conversations about their surveys. We further thank the Delphi Group at CMU for their help in computing weekly design effects for their survey, the Ipsos team for providing data on their "offline" respondents, and the CDC for responding to our questions. Susan Paddock, other participants at the JPSM 2021 lecture (delivered by Meng), and Steve Finch provided helpful comments, which we greatly appreciate. We thank the anonymous reviewers for their constructive comments, which substantially improved our work. We thank Ariel Edwards-Levy for a tweet which originally inspired our interest in this topic, and Rick Born for suggesting more intuitive terms used in Equation 1. V.B. is funded by the University of Oxford's Clarendon Fund and the EPSRC and MRC through the OxWaSP CDT programme (EP/L016710/1). X-L. M acknowledges partial financial support by NSF. S.F. acknowledges the support of the EPSRC (EP/V002910/1).

Author contributions

V.B. and S.F. conceived and formulated the research questions. V.B. and S.K. contributed equally to data analysis, writing, and visualization. X-L.M. conceived and formulated the methodology. All authors contributed to methodology, writing, visualization, editing, and data analysis. S.F. supervised the work.

Competing Interests

Authors have no competing interests, financial or otherwise.

References

1. Meng, X.-L. Statistical paradises and paradoxes in big data (I): Law of large populations, big data paradox, and the 2016 US presidential election. *The Annals of Applied Statistics* **12**, 685–726 (2018).

2. Barkay, N. *et al.* Weights and methodology brief for the COVID-19 Symptom Survey by University of Maryland and Carnegie Mellon University, in partnership with Facebook. 1–7 (2020).

3. Kreuter, F. & et al. Partnering with Facebook on a university-based rapid turn-around global survey. *Survey Research Methods* **14**, 159–163 (2020).

4. Fields, J. *et al.* Design and operation of the 2020 Household Pulse survey. (2020).

5. Jackson, C., Newall, M. & Yi, J. Axios Ipsos Coronavirus Index. (2021).

6. Public Opinion Research (AAPOR), A. A. for. Best practices for survey research. (2021).

7. Hastak, M., Mazis, M. B. & Morris, L. A. The role of consumer surveys in public policy decision making. *Journal of Public Policy & Marketing* **20**, 170–185 (2001).

8. B.P. Murthy, et al. Disparities in COVID-19 vaccination coverage between urban and rural counties: United States, December 14, 2020 – April 10, 2021. *Morbidity and Mortality Weekly Report* (2021) doi:10.15585/mmwr.mm7020e3.

9. Arrieta, A., Gakidou, E., Larson, H., Mullany, E. & Troeger, C. Through understanding and empathy, we can convince women to get the COVID-19 vaccine. *Think Global Health* (2021).

10. Japec, L. *et al.* AAPOR report on Big Data. *American Association of Public Opinion Researchers* (2015).

11. Reinhart, A., Kim, E., Garcia, A. & LaRocca, S. Using the COVID-19 Symptom Survey to track vaccination uptake and sentiment in the United States. (2021).

12. Mayer-Schönberger, V. & Cukier, K. *Big data: A revolution that will transform how we live, work, and think*. (Houghton Mifflin Harcourt, 2013).

13. CDC. Trends in number of COVID-19 vaccinations. (2021).

14. Nguyen, K. H. *et al.* Comparison of COVID-19 vaccination coverage estimates from the Household Pulse Survey, Omnibus Panel Surveys, and COVID-19 vaccine administration data, United States, March 2021.

15. Santibanez, T. A. *et al.* Sociodemographic Factors Associated with Receipt of COVID-19 Vaccination and Intent to Definitely Get Vaccinated, Adults aged 18 Years or Above — Household Pulse Survey, United States, April 28–May 10, 2021. (2021).

16. Kruskal, W. & Mosteller, F. Representative sampling, I: Non-scientific literature. *International Statistical Review/Revue Internationale de Statistique* 13–24 (1979).

17. Kruskal, W. & Mosteller, F. Representative sampling, II: Scientific literature, excluding statistics. *International Statistical Review/Revue Internationale de Statistique* 111–127 (1979).

18. Kruskal, W. & Mosteller, F. Representative sampling, III: The current statistical literature.

19. Kruskal, W. & Mosteller, F. Representative sampling, IV: The history of the concept in statistics, 1895-1939. *International Statistical Review/Revue Internationale de Statistique* 169–195 (1980).

20. AAPOR. Margin of Sampling Error/Credibility Interval.

21. The Delphi Group at Carnegie Mellon University in partnership with Facebook. Topline Report on COVID-19 Vaccination in the United States. (2021).

22. Haas, E. J. *et al.* Impact and effectiveness of mRNA BNT162b2 vaccine against SARS-CoV-2 infections and COVID-19 cases, hospitalisations, and deaths following a nationwide vaccination campaign in Israel: An observational study using national surveillance data. *The Lancet* (2021).

23. Kish, L. *Survey Sampling*. (Wiley, 1965).

24. Institute for Health Metrics and Evaluation (IHME). COVID-19 vaccine hesitancy. (2021).

25. King, W. C., Rubinstein, M., Reinhart, A. & Mejia, R. J. Time trends and factors related to COVID-19 vaccine hesitancy from january-may 2021 among US adults: Findings from a large-scale national survey. *medRxiv* (2021) doi:10.1101/2021.07.20.21260795.

26. Centers for Disease Control. Estimates of vaccine hesitancy for COVID-19. (2021).

27. Groves, R. M. *et al. Survey methodology*. vol. 561 (John Wiley & Sons, 2011).

28. Dempsey, W. The hypothesis of testing: Paradoxes arising out of reported Coronavirus case-counts. *arXiv* 1–21 (2020).

29. Isakov, M. & Kuriwaki, S. Towards principled unskewing: Viewing 2020 election polls through a corrective lens from 2016. *Harvard Data Science Review* **2**, (2020).

30. Hartley, H. O. & Ross, A. Unbiased ratio estimators. *Nature* **174**, 270–271 (1954).

31. Tiu, A., Susswein, Z., Merritt, A. & Bansal, S. Characterizing the spatiotemporal heterogeneity of the COVID-19 vaccination landscape. (2021).

32. Groen, J. Sources of Error in Survey and Administrative Data: The Importance of Reporting Procedures. *Journal of Official Statistics* **28**, 173–198 (2012).

33. Tu, X. M., Meng, X.-L. & Pagano, M. The AIDS epidemic: estimating survival after AIDS diagnosis from surveillance data. *Journal of the American Statistical Association* **88**, 26–36 (1993).

34. Barnes, O. & Burn-Murdoch, J. COVID response hampered by population data glitches. *Financial Times* (2021).

35. Kennedy, C. *et al.* Evaluating online nonprobability surveys. *Pew Research Center* (2016).

36. Kennedy, C. *et al.* An evaluation of the 2016 election polls in the United States. *Public Opinion Quarterly* **82**, 1–33 (2018).

37. Auxier, B. & Anderson, M. Social media use in 2021. *Pew Research Center* (2021).

38. Gadarian, S. K., Goodman, S. W. & Pepinsky, T. B. Partisanship, health behavior, and policy attitudes in the early stages of the COVID-19 pandemic. *PLOS One* **16**, (2021).

39. Mercer, A., Lau, A. & Kennedy, C. For Weighting Online Opt-In Samples, What Matters Most? *Pew Research Center* (2018).

40. Ryan, C. Computer and internet use in the United States: 2016. *American Community Survey Reports* **ACS-39**, (U.S. Census Bureau, Washington, DC, 2017).

41. Biemer, P. P. & Lyberg, L. E. *Introduction to survey quality*. (John Wiley & Sons, 2003).

42. Scheuren, F. What is a survey? in (American Statistical Association Alexandria, 2004).

43. Sukhatme, P. V. Sampling theory of surveys with applications. (1954).

44. Neyman, J. On the two different aspects of the representative method: The method of stratified sampling and the method of purposive selection. *Journal of Royal Statistical Society* **97**, 558–625 (1934).

45. Groves, R. M. Nonresponse rates and nonresponse bias in household surveys. *Public opinion quarterly* **70**, 646–675 (2006).

46. Lazer, D., Kennedy, R., King, G. & Vespignani, A. The parable of Google Flu: Traps in big data analysis. *Science* **343**, 1203–1205 (2014).

47. Salomon, J. A. *et al.* The US COVID-19 Trends and Impact Survey, 2020-2021: Continuous real-time measurement of COVID-19 symptoms, risks, protective behaviors, testing and vaccination. (2021).

48. Park, D. K., Gelman, A. & Bafumi, J. Bayesian multilevel estimation with poststratification: State-level estimates from national polls. *Political Analysis* **12**, 375–385 (2004).

49. Associated Press-NORC Center for Public Affairs Research. The june 2021 AP-NORC center poll. (2021).

50. Wang, W., Rothschild, D., Goel, S. & Gelman, A. Forecasting elections with non-representative polls. *International Journal of Forecasting* **31**, 980–991 (2015).

51. Elliott, M. R. & Valliant, R. Inference for nonprobability samples. *Statistical Science* **32**, 249–264 (2017).

52. Little, R. J., West, B. T., Boonstra, P. S. & Hu, J. Measures of the degree of departure from ignorable sample selection. *Journal of survey statistics and methodology* **8**, 932–964 (2020).

53. Wiśniowski, A., Sakshaug, J. W., Perez Ruiz, D. A. & Blom, A. G. Integrating probability and nonprobability samples for survey inference. *Journal of Survey Statistics and Methodology* **8**, 120–147 (2020).

54. Yang, S., Kim, J. K. & Song, R. Doubly robust inference when combining probability and non-probability samples with high dimensional data. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **82**, 445–465 (2020).

55. Methodology for the United States population estimates: Vintage 2019.

56. Bethlehem, J. Weighting Nonresponse Adjustments Based on Auxiliary Information. 275--288 (New York: Wiley, 2002).

57. Meng, X.-L. A trio of inference problems that could win you a Nobel prize in statistics (if you help fund it). 537–562 (CRC Press, 2014).

58. Meng, X.-L. & Xie, X. I got more data, my model is more refined, but my estimator is getting worse! Am I just dumb? *Econometric Reviews* **33**, 218–250 (2014).

59. Heckman, J. J. Sample selection bias as a specification error. *Econometrica* 153–161 (1979).

60. CDC. Reporting COVID-19 vaccination demographic data. (2021).

Extended Data Legends

Extended Data Fig. 1 | **Comparisons of state-level vaccine uptake, hesitancy, and willingness across surveys and the CDC: March 2021.** Comparison of state-level point estimates (A-C) and rankings (D-F) for vaccine hesitancy, willingness, and uptake from Delphi-Facebook, and Census Household Pulse. Dotted black lines show agreement and red points show the average of 50 states. Panels G-J compare state-level point estimates and rankings for the same survey waves to CDC benchmark estimates from March 31, 2021. The Delphi-Facebook data is from the week ending March 27, 2021 and the Census Household Pulse is the wave ending March 29, 2021.

Extended Data Fig. 2 | **Comparisons of state-level vaccine uptake, hesitancy, and willingness across surveys and the CDC: May 2021.** Comparison of state-level point estimates (A-C) and rankings (D-F) for vaccine hesitancy, willingness, and uptake from Delphi-Facebook, and Census' Household Pulse. Dotted black lines show agreement and red points show the average of 50 states. Panels G-J compare state-level point estimates and rankings for the same survey waves to CDC benchmark estimates from May 15, 2021. The Delphi-Facebook data is from the wave week ending May 8, 2021 and the Census Household Pulse is the wave ending May 10, 2021.

Extended Data Table 1 | Methodologies of Axios-Ipsos, Census Household Pulse, and Delphi-Facebook studies

Extended Data Fig. 3 | **Retroactive adjustment of CDC vaccine uptake figures for April 3-12, 2021, over the 90 days from April 12.** Increase is shown as a percentage of the vaccine uptake reported on April 12. Most of the retroactive increases in reported estimates appear to occur in the first 10 days after an estimate is first reported. By about 40 days after the initial estimates for a particular day are reported, the upward adjustment plateaus at around 5-6\% of the initial estimate. We use this analysis to guide the choice of 5\% and 10\% error in the CDC benchmark for our robustness checks.

Extended Data Fig. 4 | Revised estimates of hesitancy and willingness after accounting for survey errors for vaccination uptake. The gray point shows the reported value at the last point of the time series. Each line shows a different scenario for what might be driving the error in uptake estimate, derived using hypothetical *ddc* values for willingness and hesitancy based on the observed *ddc* value for uptake. *Access* scenario: willingness suffers from at least as much, if not more, bias than uptake. *Hesitancy* scenario: hesitancy suffers from at least as much, if not more, bias than uptake. *Uptake* scenario: the error is split roughly equally between hesitancy and willingness. See Supplementary Information D for more details.

Extended Data Fig. 5 | Vaccination Rates compared with CDC benchmark for four online polls. Ribbons indicate traditional 95 percent confidence intervals which are twice the standard error reported by the poll. Data for Progress asks "As of today, have you been vaccinated for Covid-19?"; Morning Consult asks "Have you gotten the vaccine, or not?"; Harris Poll asks "Which of the following best describes your mindset when it comes to getting the COVID-19 vaccine when

it becomes available to you?". See Supplementary Information C.3 for more details on each survey and discussion of differences. Gray line is the CDC benchmark.

Extended Data Fig. 6 | Survey error by Age Group (18-64 year-olds, and those 65 and over. a. Estimates of vaccine uptake from Delphi-Facebook (blue) and Census Household Pulse (green) for each 18-64 year-olds (left) and those 65 or older (right). Bounds on the CDC's estimate of vaccine uptake for those groups are shown in gray. The CDC receives vaccination-by-age data only from some jurisdictions. We do know, however, the total number of vaccinations in the U.S. Therefore, we calculate the bounds by allocating all the vaccine doses for which age is unknown to either 18-64 or 65+. **b.** Unweighted *ddc* for each Delphi-Facebook and Census Household Pulse calculated for the 18-64 group using the bounds on the CDC's estimates of uptake. *ddc* for 65+ is not shown due to large uncertainty in the bounded CDC estimates of uptake.

Extended Data Table 2 | Contribution of offline recruitment and weighting schemes to discrepancies between surveys. A portion of each Axios-Ipsos wave is recruited from a population with no stable internet connection; Ipsos KnowledgePanel provides tablets to these respondents. In the Axios-Ipsos March 22, 2020 wave, the offline panelists (n = 21) were 24 percentage points less likely to be vaccinated than online panelists (n = 974). Weighting the same Axios-Ipsos data (n = 995) to the age and gender target distribution implied by Delphi-Facebook's weights make the vaccination estimates higher by 1 percentage point. However, this number is still lower than Delphi-Facebook's (responses from March 14-20 2020, n = 249,954) own estimate of 46%. During this time period, the CDC benchmark vaccination rate was 35.2%. This suggests that the recruitment of offline respondents and different weighting schemes each explains only a small portion of the discrepancy between the two data sources.

Extended Data Table 3 | Example of multi-stage population selection. The *Law of Large Populations* described in Methods section “Population size in multi-stage sampling” shows that the population size at the sampling stage where simple random sampling breaks down will dominate the error. This table explains these stages with a concrete example, using the Census Household Pulse. Population and sample sizes for three stages (stage number denoted s = 1, 2, or 3) of sampling of the Census Household Pulse survey data collection process. Approximate sample sizes based on the March 24, 2021 wave. "m" stands for millions and "hh" stands for household. The final row compares the total adult population in the US (255 million adults, made up of 144 million households) to the sample size in one wave of the household pulse. For illustration, we have ignored the impact of unequal sampling probabilities on the sample sizes at each stage.