On the Robustness of Social-Circle Surveys: Respondent Selection Issues, Egocentrism, and Homophily

Kristoffer Ahlstrom-Vij
Birkbeck, University of London
k.ahlstrom-vij@bbk.ac.uk

Abstract: Asking people about the preferences of people in their social circles tends to yield more accurate estimates of population preference distributions than does asking each respondent about their own preference. This is likely because the former approach taps into people's knowledge about others and thereby generates an *implicit super sample* that includes non-sampled members of participants' social circles. The present paper makes two contributions. First, it uses a set of simulation studies to argue that the superiority of social-circle surveys can be expected to be robust in the face of respondent selection issues (e.g., non-response and coverage bias), people being highly fallible about other people's preferences (egocentric bias), and people largely surrounding themselves with those who share their preferences (homophily). Second, it reports on a survey experiment offering preliminary evidence that egocentric bias in particular can be reduced significantly through a simple survey prompt. In closing, the paper also discusses the relationship between social-circle questions and the type of closely related expectation questions (e.g., "Who do you expect will win the election?") typically found on prediction markets – markets for placing bets on future or otherwise unknown events – which also tend to outperform traditional polls.

1. Social-circle surveys

Who will win the election? Is there public support for some particular governmental policy? Will some particular product find a market among some specific segment of the population, or in some new geographic region? To answer questions such as these, both public and private bodies regularly need to make judgments about preference distributions in different populations — a demand that in turn is catered to by a large sector of market research and opinion poll providers, relying on by-now familiar and well-established sampling and survey practices.

However, a growing body of evidence suggests that surveys asking people about the preferences of people in their social circles yield more accurate estimates of the preference distribution in the population than do established survey practices, that involve asking people only about their own preferences (Galesic et al. 2018; Galesic and Bruine de Bruin 2020). In some cases, this might be explained in part by social-circle surveys likely being less prone to respondents censoring their true preferences, as in the case of the supposed 'hidden/shy Trump voter' (Enns, Lagodny, and Schuldt 2017), since disclosing controversial preferences of friends carry fewer social costs than disclosing those same preferences on the part of oneself. More generally, however, social-circle surveys can be expected to tap into people's non-trivial amount of knowledge about non-sampled members of the population (Nisbett and Kunda 1985), generating what we can refer to as an *implicit super sample*. This sample is larger than the set of those surveyed (hence, a *super* sample) in also including all members of participants' social circles, who are thereby *implicitly* sampled as well. This has the potential of mitigating respondent selection issues that typically affect surveys (Weisberg 2018), including non-response and coverage bias (Galesic et al. 2018).

The present paper uses a combination of simulation studies and experimental evidence to contribute to our understanding of such social-circle surveys. Specifically, Section 2 uses a simulation study to show that social-circle surveys can be expected to outperform traditional surveys with respect to levels of sampling error in small samples, in line with the idea of an implicit super sample, while Section 3 shows that the advantage of social-circle surveys becomes particularly pronounced in cases of respondent selection biases, such as non-response and coverage bias.

More significantly, Section 4 shows that the superiority of social-circle surveys is robust even if we assume that respondents are highly fallible about other people's preferences due to egocentric bias (Ross, Greene, and House 1977; Van Boven and Loewenstein 2003; Bernstein et al. 2004). It is also shown by way of an experimental survey that, in the general case involving people estimating the preferences of others — whether in their social circles or not — egocentric bias can be reduced through a simple survey prompt, informing people about the prevalence of egocentric bias, and asking them not to fall prey to it in their subsequent estimations.

Of course, there is little use to reducing egocentric bias if people barely engage with those who do not share their preferences to begin with. This gets to the issue of homophily (McPherson, Smith-Lovin, and Cook 2001). As is shown in Section 5, however, even if people both exhibit high levels of egocentric bias and almost exclusively surround themselves with people who share their preferences, social-circle surveys can still be expected to outperform traditional surveys. This is particularly noteworthy since such tendencies for "social bubbles" would seem to strike at the very heart of the value of a social-circle approach to surveying.

Finally, Section 6 discusses the relationship between social-circle questions and the closely related types of expectation questions (e.g., "Who do you expect will win the election?") studied by, e.g., Rothschild and Wolfers (2012) and Murr, Stegmaier, and Lewis-Beck (2021). It is suggested that the latter are likely successful by tapping into the same types of implicit super samples that social-circle questions do, and that this moreover offers a particularly parsimonious explanation of the relative success of prediction markets – markets for placing bets on future or otherwise unknown events – over traditional polls (J. E. Berg and Rietz 2014; Rothschild 2009; J. E. Berg, Nelson, and Rietz 2008). This, in turn, can help shed further light on the merits — and potential problems — with implicit super samples, and suggests that future work on the interrelations between social-circle surveys, expectation questions, and prediction markets in particular would be fruitful and worthwhile.

2. Harnessing implicit super samples

It was noted at the outset that social-circle questions likely outperform traditional survey questions asking participants exclusively about their own preference on account of the former tapping into implicit super samples. To illustrate the idea of such a sample, consider a simple case, involving three individuals, A, B, and C, with some preference on some binary policy matter X (grey for "for" and white for "against", say), as in Fig. 1.

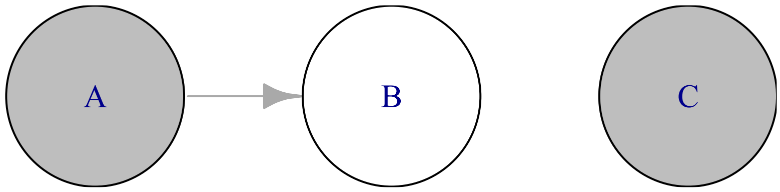

Figure 1: Three people, two of whom are connected.

Let the edge between A and B in Fig. 1 designate that A knows B's policy preference on X, but not C's. Moreover, imagine that we ask, not the traditional traditional survey question, but a social-circle question: "What proportion of people in your social circle are supportive of policy X?" Because of the connection between A and B, were we to sample A, we would also get B's preference "for free". Moreover, while on the traditional survey approach we can trust that sampling the entire population (i.e., A, B, and C) will guarantee an accurate estimate, this is interestingly not the case on the social-circle approach. It would have been the case, had all "connections" between people been bi-directional – that is, if my knowing your preferences invariably meant your knowing mine. (The reader may verify this for themselves in the example above.) Assuming such bidirectionality would be implausible, however, so it will not be assumed in what follows.

Moving beyond the simplistic example in Fig. 1, we may use a set of simulation studies to investigate under what specific conditions social-circle questions generate smaller errors in estimates of a population preference distribution compared to a traditional survey approach that asks people only about their own preferences.[1] To that end, the R script (R Core Team 2018) included in the Online Appendix does the following:

1. It generates 500 populations (set by *iterations*) of 1,000 persons each (set by *n*).
2. For each population, it randomly defines some distribution of support (set by *prop*) for an imagined policy (1 = for; 0 = against), in the range of 0-100%.
3. For each person in a population, it randomly finds some pre-specified (by *num_neighbours*) number of "neighbours" in their social circle, where each such person knows their own preference as well as that of their neighbours.
4. For each of the 500 populations, and for each sample size between 1% and 100%, it then draws 100 random samples (*number_of_samples*) and looks at both the mean level of support in that sample, and the mean perceived level of support in that sample. The former simulates asking each person sampled whether they are for or against the policy (and taking the mean response); the latter simulates asking each person about the proportion of their social circle (themselves included) that is for the policy. This corresponds to the traditional survey approach and the social-circle approach, respectively.
5. It then measures the mean absolute error in percentage points across those 100 samples, for each of the 100 sample sizes (i.e., 1% to 100%), and compares the average error for each sample size for the traditional survey approach and the social-circle approach, respectively.
6. Steps 1-5 are repeated for 1, 2, 5, 10, 50, and 100 neighbours, corresponding to each person knowing the preference of 0.1% to 1% of the population.

Fig. 2 plots the average error across the 100 samples drawn from each of the 500 populations for each sample size (1-100%), for a total of 100 x 500 x 100 = 5 million samples per *num_neighbours* value. This enables us to compare in particular the error of the social-circle approach to the traditional survey approach.

[1] In what follows, only binary preferences (i.e., 'for' vs. 'against') will be addressed. This is a non-consequential simplification for two reasons. First, most outcomes can in practice be transformed to binary ones in a straightforward manner. Second, while modeling non-binary outcomes would make the presentation more complex, it would not alter the substance of the conclusion. Many thanks to one of the reviewers for this journal for raising this question.

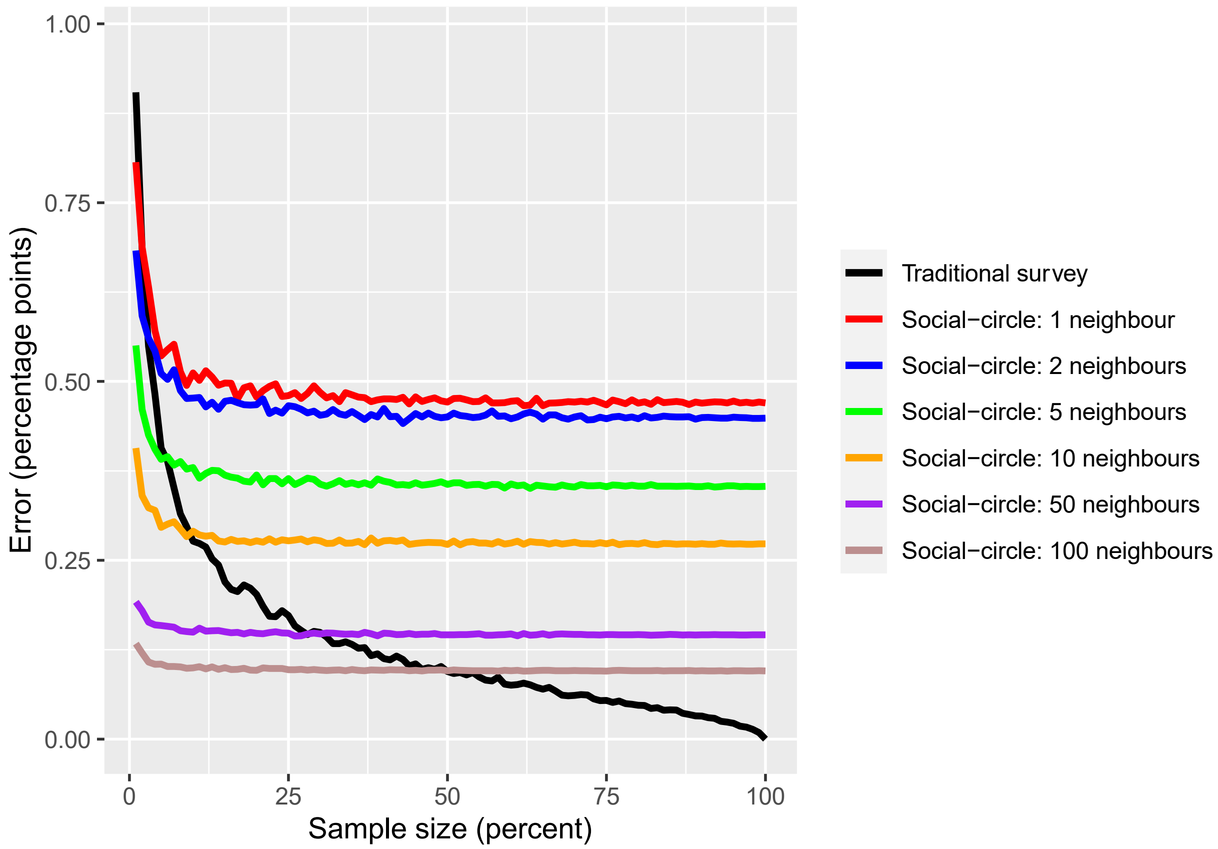

Figure 2: Average error across 100 samples drawn from each of the 500 populations for each sample size (1-100% of population). Shaded areas designate 95% confidence bands.

Given random sampling (as per step 4 above), the error measured in Fig. 2 is simply the sampling error, which can be seen to go down as the sample size increases for the traditional survey approach. We also see that the social-circle approach outperforms the traditional survey approach in small samples, in line with the idea of implicit super samples, and more decisively so the higher the level of connectedness. However, unlike on the traditional survey approach, sampling the entire population does not necessarily remove all error on the social-circle approach, since the relation of knowing another's preference is not assumed to be symmetric, as discussed at the outset of this section.

On the point about small samples in particular, it helps to look specifically at samples sizes in the range of 1-20%, as in Fig. 3.

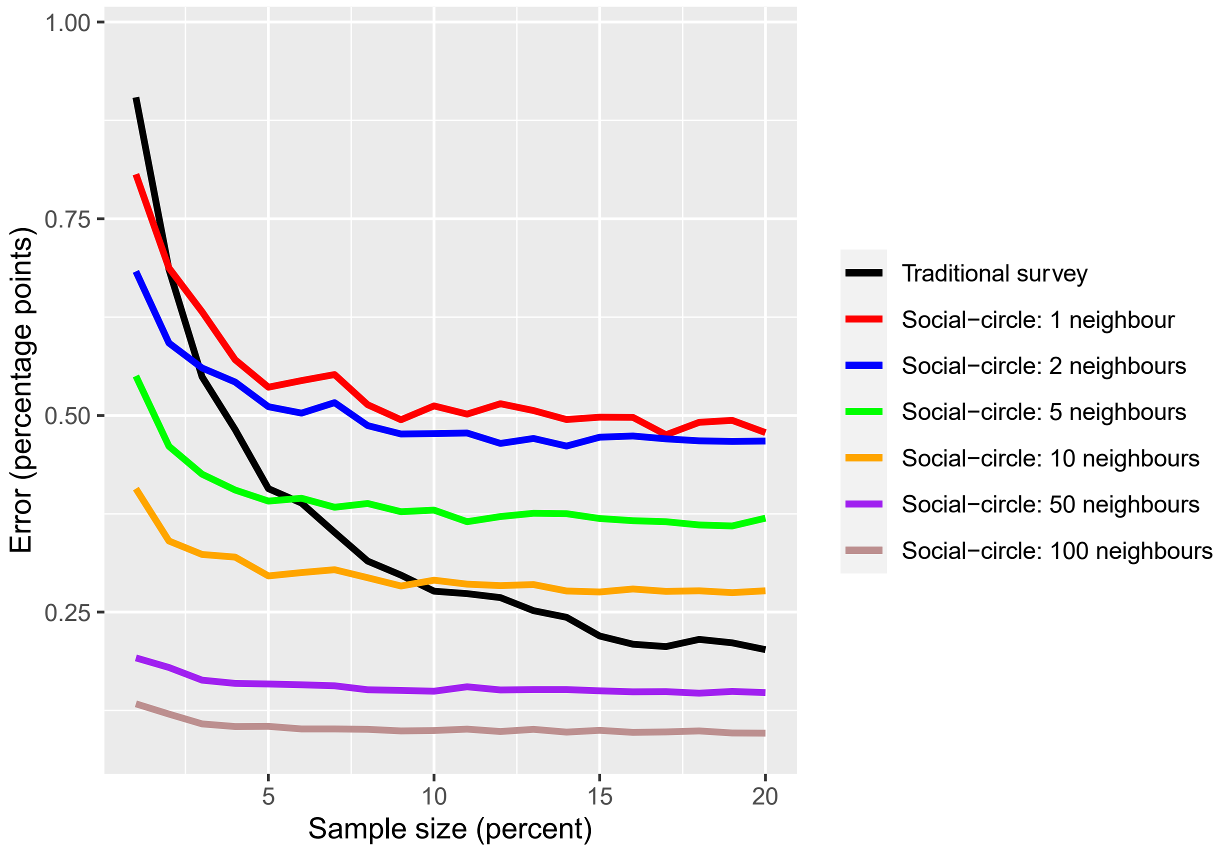

Figure 3: Average error across 100 samples drawn from each of the 500 populations for sample sizes of 1-20% of population. Shaded areas designate 95% confidence bands.

Even such small degrees of connectedness as each person having one "neighbour" have the social-circle approach outperform the traditional survey approach up to samples of about 2% of the population (although note the overlap in confidence bands); the same goes for samples of up to about 2.5% for 2 "neighbours", 5% for 5 "neighbours", 10% for 10 "neighbours", 30% for 50 "neighbours" (see Fig. 2), and 50% for 100 "neighbours" (also Fig. 2).

That bodes well for the social-circle approach, given that we are rarely able to achieve large samples in survey contexts. At the same time, this simulation assumes (perfectly) random sampling. This is not a realistic assumption. Any given survey will face some respondent selection issues (Weisberg 2018), on account of some groups of respondents systematically being more difficult to reach or not wanting to cooperate (unit-level non-response bias), or some groups being systematically unavailable through accessible recruitment channels (coverage bias). Such issues mean that sampling will not be (perfectly) random, as some groups will be more likely to be sampled than others. In the next section, we therefore relax the assumption about random sampling to see what then happens to the comparative merits of the two approaches.

3. Respondent selection bias

As noted in the previous section, we are rarely able to achieve perfectly random sampling. Let us factor that in, by adding the following to step 4 in our simulation (the full script of which is available in the Online Appendix):

- Prior to sampling from a population, we randomly define some degree of respondent selection *bias*, in the range of 0-99%, such that every person not supporting the policy (*pref*==0) has a probability of *bias* of being included in the sample, and everyone supporting it (*pref*==1) has a probability of 1-*bias* of being included in the sample.[2]

Fig. 4 compares the levels of accuracy between the traditional survey and the social-circle approach.

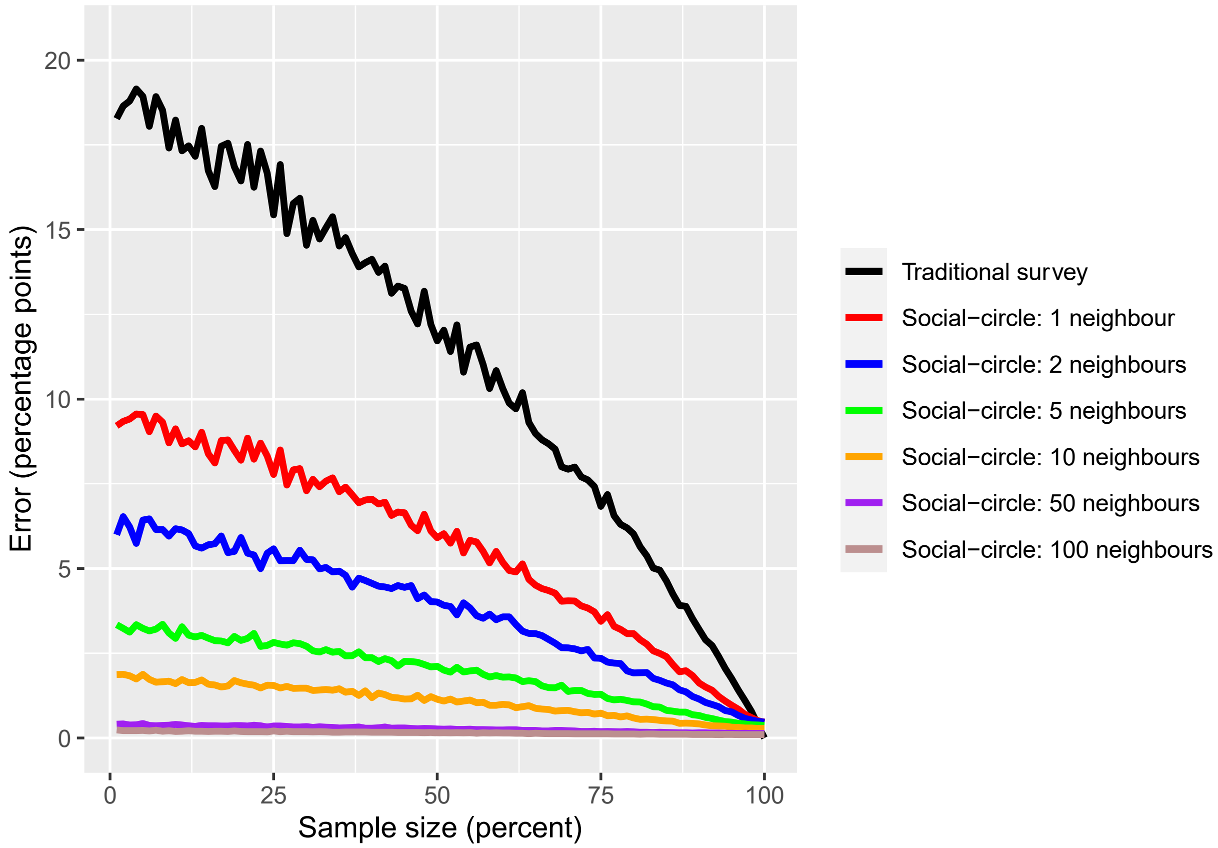

Figure 4: Average error across 100 samples drawn from each of the 500 populations for each sample size (1-100% of population), factoring in sampling bias. Shaded areas designate 95% confidence bands.

We see that the level of error goes up substantially across the board, as is to be expected. However, in small samples, the social-circle approach cuts the error of the traditional survey approach in half, even for cases where each person only has a single "neighbour". More generally, we see that, under the more realistic assumption of biased as opposed to random sampling due to respondent selection issues, the social-circle approach clearly outperforms the traditional survey approach, and it does so almost all the way up to samples encompassing the entire population.

4. Egocentric bias

It might be objected that what is driving the results in the previous section is the assumption in step 3 of our simulation that each person has perfect knowledge of the preferences of their "neighbours". In reality, we tend to exhibit an egocentric bias (Ross, Greene, and House 1977) in identifying others' preferences, meaning that we tend to assume that others prefer more or less what we do. Indeed, the tendency to project our own states onto others is not restricted to preferences, but extends to our knowledge (Bernstein et al. 2004), and even to our thirst (Van Boven and Loewenstein 2003), and is more generally predicted by simulation theory, i.e., the theory that mind-reading involves simulating the mental states of others based on one's own using perspective taking and pretense (Shanton and Goldman 2010).

---

[2] Note that, since we are measuring *absolute* percentage point error, randomizing the level of error from one sample to the next does not lead to that bias canceling out.

4.1. Factoring in egocentric bias
In light of this, let us add the following to step 3 of our simulation (the full script of which is, again, available in the Online Appendix):

- Each person will ascribe to each of their neighbours the correct preference with a probability of 51% (their level of *discernment*), and otherwise simply ascribe to them their own preference.

In other words, we are simulating a very high degree of egocentric bias; for each neighbour, a participant barely beats chance in correctly identifying their preferences. Fig. 5 shows the impact of this high degree of egocentric bias on the comparative merits of the traditional and the social-circle approach.

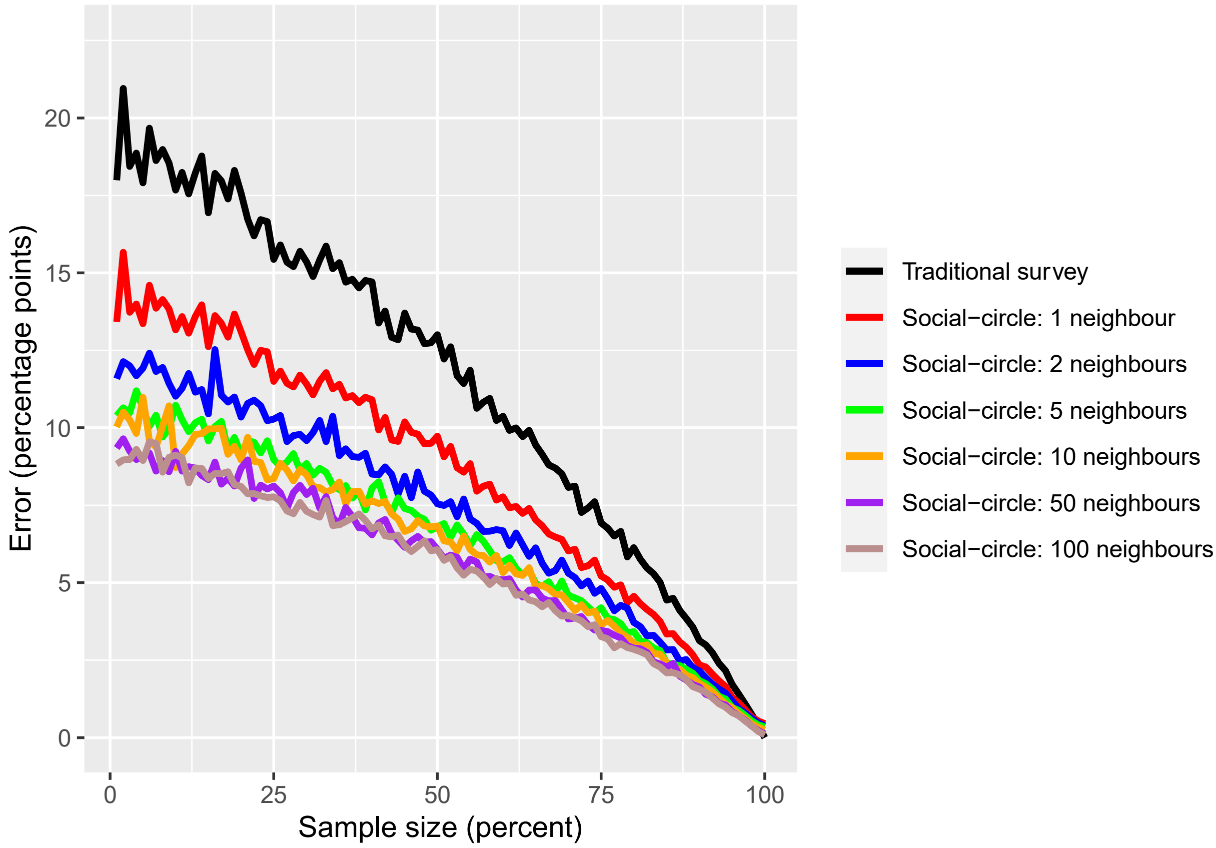

Figure 5: Average error across 100 samples drawn from each of the 500 populations for each sample size (1-100% of population), factoring in both sampling and egocentric bias. Shaded areas designate 95% confidence bands.

As is to be expected, the degree of error goes up for the social-circle approach under these circumstances. However, that error is still lower than for the traditional survey approach, suggesting that even widespread and substantial egocentric bias does not undo the advantage that the social-circle approach has over the traditional survey approach under circumstances of respondent selection bias.

4.2. Mitigating egocentric bias through survey prompts
Moreover, there is some reason to believe that egocentric bias can be mitigated. After all, such bias is a perfectly general bias, not unique to the context where someone is asked specifically about their social circle. Rather, as noted at the beginning of this section, it is a bias exhibited in the general case where we are asked about the preferences of others, whether ones located in our social circle or not. With this point in mind, an online survey experiment was fielded in the period of March 8-13, 2019

with a US sample (*N* = 352)[3], where participants were assigned to one of three conditions[4]:

In the *egocentric condition* (*N* = 118), participants were asked, first, about their political preference in the (at that point) upcoming US Presidential election – "In your view, should Donald J Trump be given a second term in next year's Presidential election?" – and, then, to estimate a preference distribution in a sample – "Using Prolific.ac, we are asking the previous question of 300 people in the US. In your estimation, what proportion of those people will answer that Donald J. Trump should be given a second term in next year's US Presidential election?"[5]

In the *treatment condition* (*N* = 115), participants were asked the same questions as in the egocentric condition, but also shown the following prompt prior to being presented with the estimation question:

> *Psychologists have shown that, when we try to figure out what other people feel, think, or want, we tend to project onto them our own feelings, thoughts, and desires. So, for example, if you like chocolate, you tend to overestimate the extent to which other people like chocolate, too.*
>
> *The same thing happens in politics: we tend to overestimate the extent to which other people hold the political views we ourselves hold. To prevent this bias from affecting your judgment, we now want you to keep this in mind when answering the question on the following page.*

Finally, in the *non-egocentric condition* (*N* = 119), the estimation question from the other two conditions was replaced with the following: "Using Prolific.ac, we are asking 300 people in the US the following question: 'Do you own a car?' In your estimation, what proportion of those people will answer 'Yes'?"

The experiment was set up to compare the levels of egocentric bias in the three conditions, with that bias being quantified as the point-biserial correlation between reported political preferences (with "Yes" coded as 1, and "No" coded as 0) and estimates. Specifically, two hypotheses were tested. The first hypothesis set out to replicate the phenomenon of an egocentric bias:

> H1. The correlation between reported preferences and estimates will be significantly stronger in the egocentric condition than in the non-egocentric condition.

The second hypothesis tested the effectiveness of the above intervention in reducing egocentric bias:

> H2. The correlation between reported preferences and estimates will be significantly weaker in the treatment condition than in the egocentric condition.

As can be seen from Table 1, the difference in the mean estimate between the two preference groups

---

[3] Ethical approval was obtained from the College Ethics Committee at Birkbeck, University of London, prior to recruitment, with full details available on request. Participants were recruited through Prolific.ac. Prolific offers several advantages over other established options such as Amazon Mechanical Turk, in virtue of a more diverse pool of respondents (Peer et al. 2017) and more reliable pre-screening (Palan and Schitter 2018).
[4] Six calls for participants (two per experimental condition) were set up and opened to participants in random order, to approximate random sampling. Table 1 in the Online Appendix provides details on the make-up of the sample, and shows the three groups to be balanced across demographic variables, suggesting that sampling was approximately random. 21 participants were dropped for answering "Don't know" to the main question about their political preference. Since disqualifying respondents based on criteria such as time spent on treatments might introduce post-treatment bias (Montgomery, Nyhan, and Torres 2018), no further respondents were dropped.
[5] 300 was the number of completes expected when setting up the study. A total of 352 participants ended up completing the study.

in the non-egocentric condition was small and statistically insignificant ($W = 1092.5$, $p = 0.174$, using a Wilcoxon rank sum test with continuity correction). The correlation between reported preferences and estimates was weak but not non-existent, possibly revealing some level of anchoring from the reported preferences on the subsequent estimates (Tversky and Kahneman 1974).

Condition	Mean estimate among 'Yes'	Mean estimate among 'No'	Difference	p-value (Wilcoxon)	r_{pb}	N
Non-egocentric	0.791	0.735	0.056	0.174	0.141	119
Egocentric	0.520	0.337	0.183	< 0.001	0.440	118
Treatment	0.352	0.365	-0.012	0.765	-0.031	115

Table 1.

By contrast, the difference in mean estimate between the preference groups in the egocentric condition was substantial and also statistically significant ($W = 1805$, $p < 0.001$, using a Wilcoxon rank sum test with continuity correction). Moreover, as for H1, the correlation between reported preferences and estimates was strong (0.44), and moreover significantly stronger than in the non-egocentric condition ($z = 2.511$ and $p = 0.012$, using the Fisher r-to-z transformation). This supports H1.

The difference between estimates was small and statistically insignificant in the treatment condition ($W = 754.5$, $p = 0.765$, using a Wilcoxon rank sum test with continuity correction). There was no correlation between answers to the preference and the estimation question in the treatment condition (-0.031), and that correlation was moreover significantly weaker than in the egocentric condition condition ($z = 3.792$ and $p < 0.001$, using the Fisher r-to-z transformation). This supports H2.

These results suggest that the challenge posed by egocentric bias in the estimation of the preferences of others can potentially be met by way of a simple survey prompt, along the lines of the one provided as part of this survey experiment. Moreover, since any egocentric bias involved in asking about someone's social circle in particular would, as already noted, just be a special case of that more general bias, this result also offers tentative evidence that that the level of error found in the simulations reported on by way of Fig. 5 is too pessimistic, given both the high level of egocentric bias assumed, and the prospects for mitigating the effects of such bias through survey prompts.

5. Homophily

Reducing egocentric bias will only be of use to the social-circle approach if people engage with people who do not share their preferences to begin with. (If they do not, any assumption on their part that everyone in their social circle shares their views will be largely accurate.) It is well-known, however, that people tend to surround themselves with people who are like them – as the saying goes, "birds of a feather flock together". In people, such flocking tends to happen around demographic factors like gender, race, ethnicity, religion, education, social class, and occupation (McPherson, Smith-Lovin, and Cook 2001), as well as social characteristics like political affiliation (Ackland and Shorish 2014). Indeed, people even date (Huber and Malhotra 2017) and mate (Alford et al. 2011) in political clusters, and the resulting homogeneity contributes to both political polarization (Butters and Hare 2020) and reduced sympathy for alternative perspectives (Klar and Shmargad 2017).

This has implications for the simulations above, since they sample neighbours randomly from the population, as per step 3. If people cluster together in ways that are not random, this is too much of a simplification. In light of this, let us add the following to step 3 of our simulation (the full script of which is, again, available in the Online Appendix):

- Define a *bubble*, designating the proportion of like-minded neighbours (i.e., neighbours of the same preference) that each person will tend to have. So, if *bubble* = 1, for example, then each person only interacts with like-minded people.

- When sampling neighbours for any given person, sample like-minded ones with a probability of *bubble* and not like-minded ones with a probability of 1-*bubble*.

To get a feel for the impact of different levels of homophily on the errors of the traditional and the social-circle approach, we run our simulation for *bubble* values of 99%, 90%, 75%, and 50%. (50% corresponds to no homophily; *bubble* values below 50% would designate heterophily, whereby 'opposites attract'.) Fig. 6 shows the impact of these four different levels of homophily on the comparative merits of the traditional and the social-circle approach.

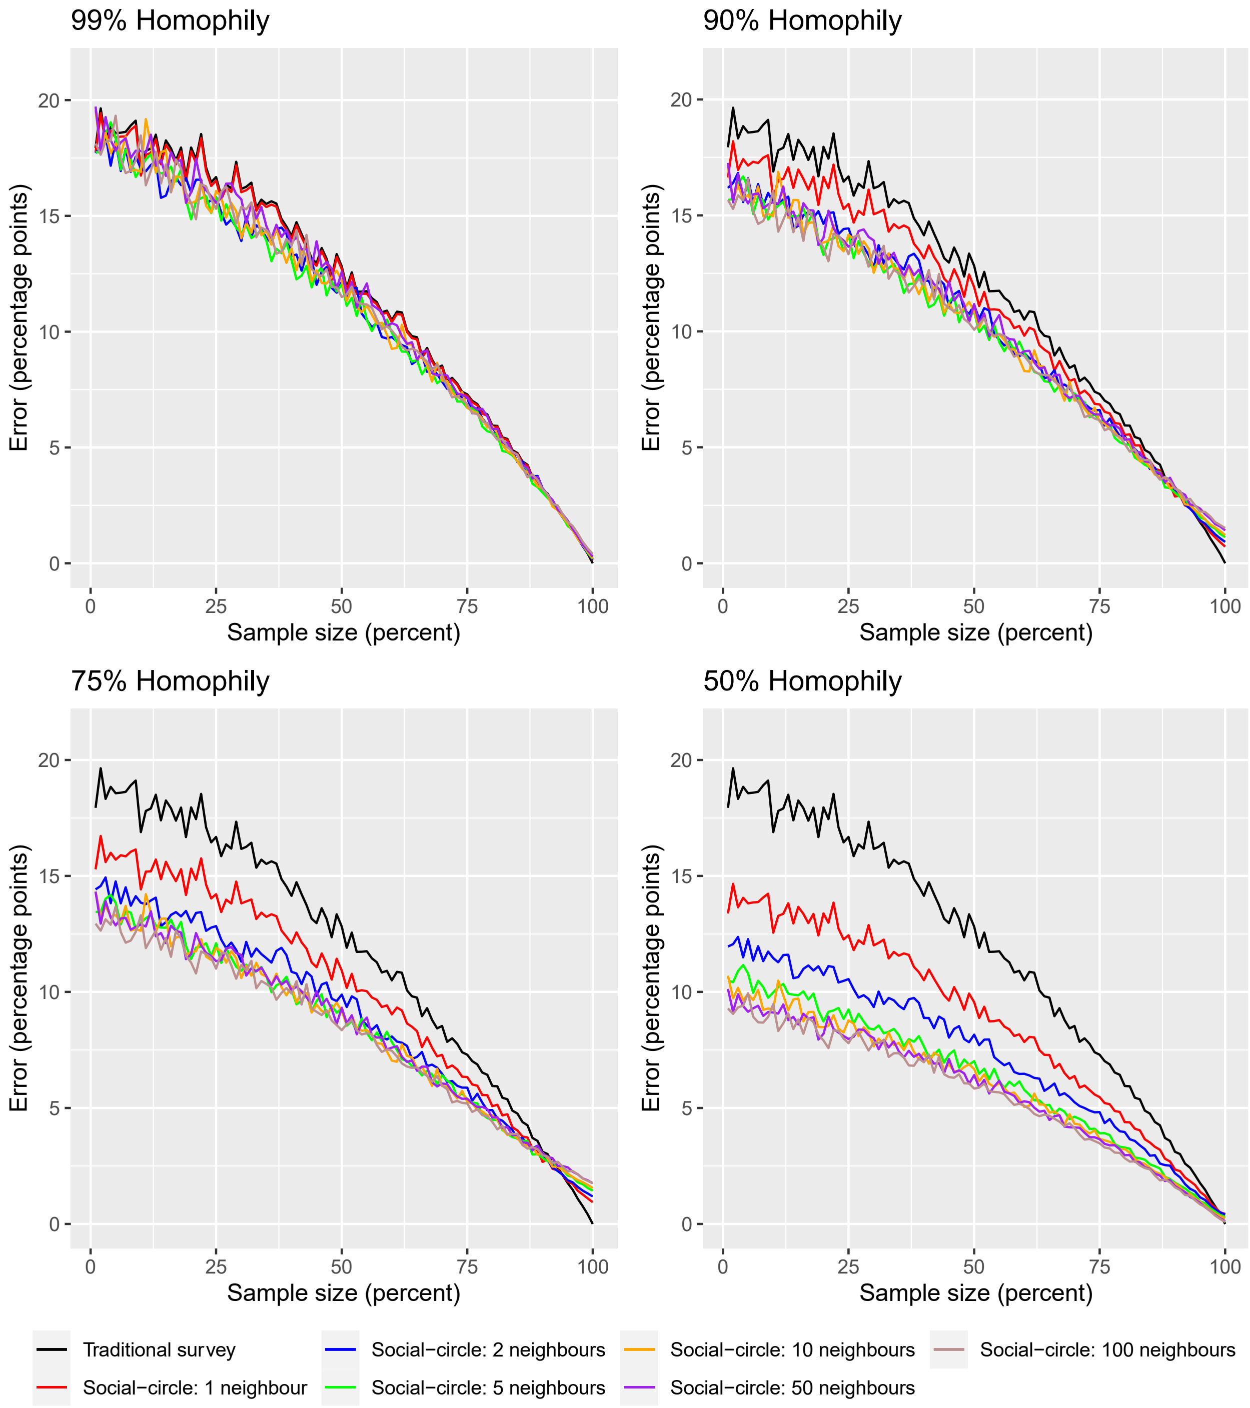

Figure 6: Average error across 100 samples drawn from each of the 500 populations for each sample size (1-100% of population), factoring in sampling bias, egocentric bias, and four different levels of homophily. Shaded areas designate 95% confidence bands.

Looking at the top-left panel of Fig. 6 first, we see that the difference between the traditional and the social-circle approach disappears once homophily reaches 99%, i.e., at the point where there is virtually no engagement whatsoever between persons that hold different preferences. This makes sense, because it would be at exactly that point that any edge that the social-circle approach has by virtue of tapping into the connections between people of different preferences vanishes.

In terms of the other three panels, we see that lower levels of homophily mean a greater reduction in error for the social-circle approach, compared to the traditional survey approach, with 50% homophily (i.e., *no* homophily, since people are equally likely to associate with like-minded as with not like-minded people) being identical to the results in the previous section. What level of homophily we can expect in real life likely depends on the context. By way of example, in analyzing the existence of 'echo chambers' in the 2016 US Presidential election, Galesic et al. (2018) found that the social circles of Trump and Clinton voters consisted of 71% and 68% like-minded individuals, respectively, prior to the election, and 77% and 68% after the election. If these number are anything to go by, the bottom left panel (75% homophily) might be the most realistic one, at least when it comes to political preferences.

However, for purposes of putting pressure on the social-circle approach, it is worth looking at the top-right panel, with homophily at 90%, and to focus on small sample sizes (0-20%) in particular, as in Fig. 7.

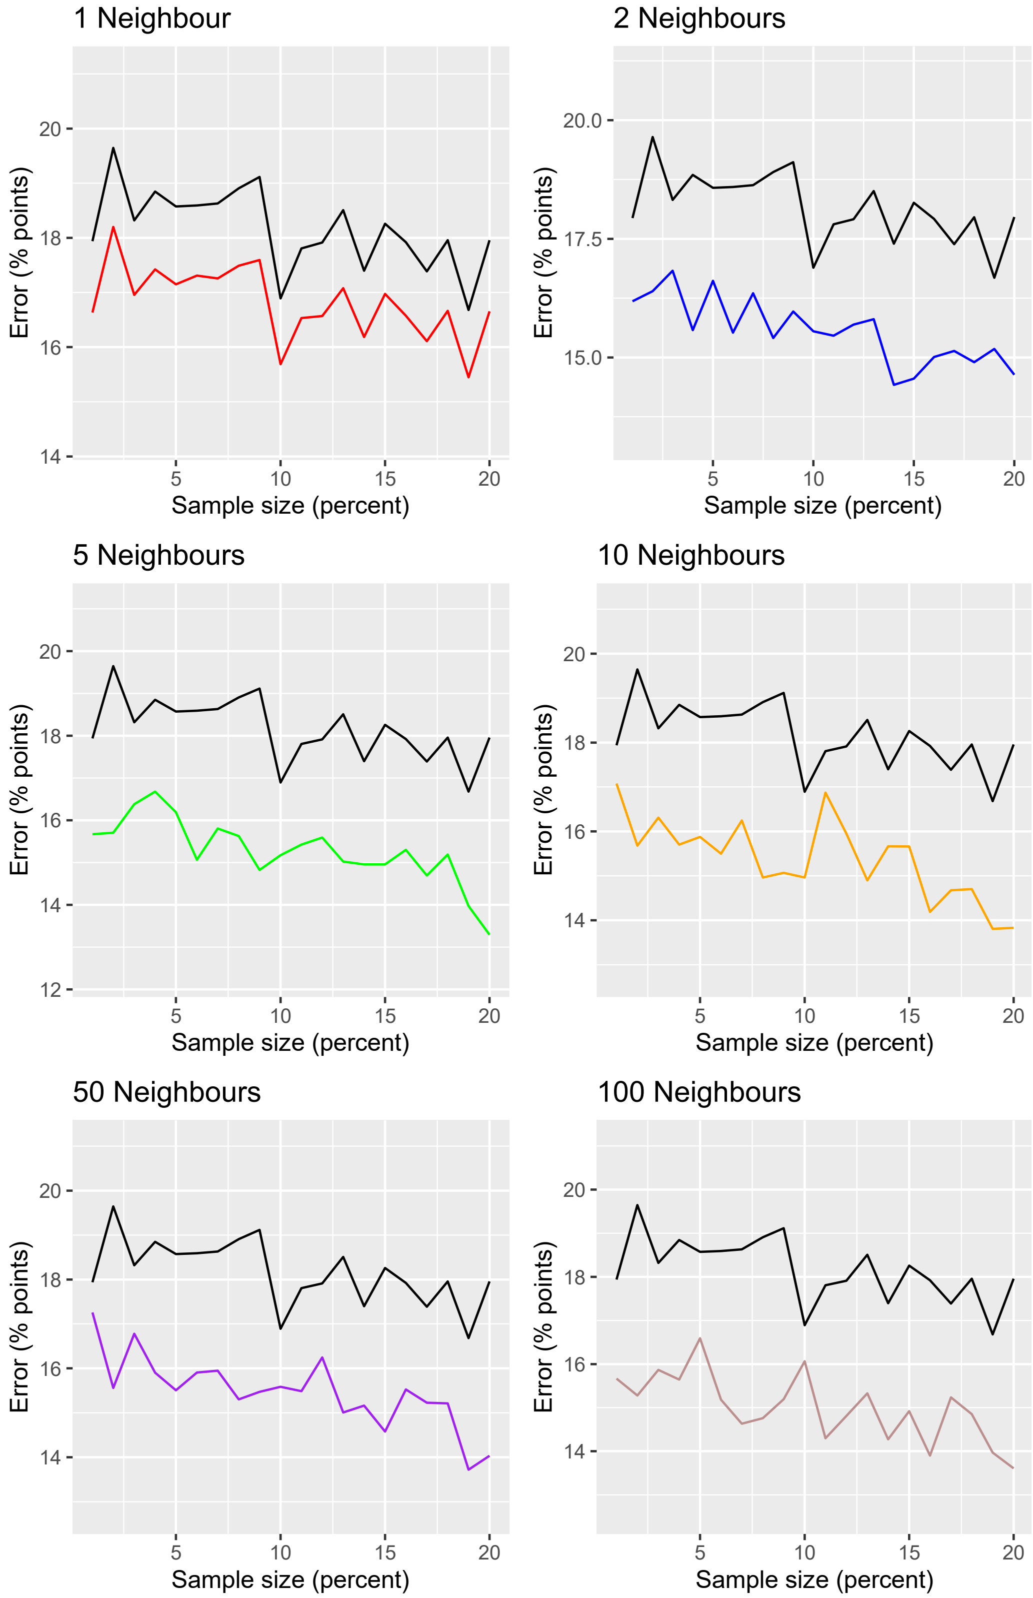

Figure 7: Average error across 100 samples drawn from each of the 500 populations for sample sizes of 1-20%, factoring in sampling bias, egocentric bias, and 90% homophily. Shaded areas designate 95% confidence bands.

What we see in Fig. 7 are two things. First, the number of neighbours does not make a substantial difference to the level of error. This is to be expected: given the high level of homophily, being connected to more people still means largely being connected with people who share your preferences. Second, and more to the point, despite (*a*) agents barely beating chance in estimating the preferences of others, on account of a strong egocentric bias, and (*b*) their social circles (i.e., their neighbours) being almost exclusively people who share their preferences owing to a very high level of homophily – two facts that, in combination, will more or less guarantee that people will estimate their neighbours' preferences to be identical to their own – the social-circle approach still outperforms the traditional survey approach. Indeed, at least for neighbour values greater than 1, there is a non-trivial amount of non-overlap in the confidence bands between the social circle and the traditional survey approach.

In parsing these results, it is worth noting exactly how hostile the modeling assumptions made here are to the social-circle approach. Generally, the more hostile the assumptions made in modeling, the more robust we can expect any results to be, since it reduces the chance that they are a mere artifact of conveniently set parameter values. In our particular case, while a realistic level of homophily in political contexts might be closer to 75%, as per Galesic et al. (2018), the fact that the social-circle approach does not 'break' under far less favourable assumptions is rather noteworthy and speaks to the aggregate power of even very low levels of social connection.

6. Discussion

The simulations reported on in the previous sections offer evidence that the superiority of social-circle surveys over traditional surveys is robust in the face of respondent selection bias (e.g., non-response and coverage bias), people being highly fallible about other people's preferences (egocentric bias), and people largely surrounding themselves with those who share their preferences (homophily). Since these conditions are quite hostile to any approach looking to harness the predictive power of people's judgments about the preferences of others – and potentially more hostile than what can reasonably be expected under realistic circumstances – this is an encouraging result for the social-circle approach.

Additional encouragement comes in the form of the experimental result offering tentative evidence that people's egocentric bias also can be reduced through a simple survey prompt, making them aware of the fact that such bias exists and asking them not to fall prey to it. The context in which this result was attained was the general one, in which participants were asked about the preferences of others, not specifically about their social circles. Given that questions about the preferences of people in your social circle is just a special case of this more general one, there is some reason to believe that such a prompt could also be incorporated in social-circle surveys to increase the advantage that the above simulations suggest that the social-circle approach already has over traditional surveys even in the absence of such an intervention.

On a related note, it is worth noting in closing that the social-circle approach is not the only alternative to the traditional survey approach. For example, Rothschild and Wolfers (2012) and Murr, Stegmaier, and Lewis-Beck (2021) investigate closely related types of expectation questions (e.g., "Who do you expect will win the election?") that have people estimate directly the preference distribution in a population. Using historical ANES data from fifteen US election in the period of 1952 to 2008, Rothschild and Wolfers (2012) in particular find that aggregated expectations are more accurate than the type of aggregated intentions that figure in traditional opinion polling. Their explanation appeals to what we have referred as implicit super samples, with people factoring in not only their own preferences, but also the preferences of anyone else about whom they might have relevant information. Rothschild and Wolfers estimate that a single voter expectation is equivalent to twenty voter intentions, in terms of its predictive value.

Unfortunately, however, we still know little about the *relative* merits of expectation versus social-circle questions. In a USC poll on the 2016 US Presidential election, Galesic et al. (2018) found that a social-circle question ("Of all of your social contacts who are likely to vote, what percentage do you think will vote for [candidate]?") gave Trump both a higher share of the popular vote than Clinton (incorrectly) and a higher number of electoral votes based on the state-level predictions (correctly). By contrast, the expectation question ("What is the percent chance that Clinton, Trump, or someone else will win?") gave Clinton a 53.4% chance of winning the election compared to 42.5% for Trump (and 4.1% for other candidates), which Galesic et al. speculate might be due to prominent media forecasts favorable to Clinton.[6]

Further work is therefore needed to better understand the relationship between social-circle questions and expectations questions. Such work would benefit from studying expectation questions as they figure specifically on prediction markets, i.e., markets for placing bets on future or otherwise unknown events. While by no means infallible (Ahlstrom-Vij 2016; Graefe 2017), the predictive success of such markets has been demonstrated across a variety of domains, including in relation to electoral outcomes (J. E. Berg and Rietz 2014; Rothschild 2009) where participants are drawn from an unrepresentative pool (J. Berg and Rietz 2006).

However, the reason for their relative accuracy is less well-understood. Standard "wisdom of crowds" accounts (Surowiecki 2005) fail to explain the full range of scenarios in which prediction markets generate accurate outputs (Ahlstrom-Vij 2016; Hanson 2013). Moreover, while it is tempting to simply point to financial incentives as the main driver of accuracy on such markets, the accuracy difference between play- and real-money markets turns out to be either non-existent (Servan-Schreiber et al. 2004) or small and context dependent (Mchugh and Jackson 2012). Relatedly, the particular way in which prediction markets are typically resolved – by rewarding participants with explicit reference to some external outcome (e.g., of an election) – does not explain their accuracy, as self-resolving markets, which reward participants with reference to some fact internal to the market (e.g., the market price at a closing time unknown to participants in advance) seem to perform equally well (Ahlstrom-Vij 2019).

By contrast, the very fact that prediction markets ask participants an expectation question, which in turn means that the sample of participants is likely a subset of a much larger, implicit super sample, offers a particularly promising and parsimonious explanation of the accuracy of prediction market estimates (Rothschild and Wolfers 2012), at least compared to the traditional survey paradigm of aggregating intentions. This, in turn, can help shed further light on the merits — and potential problems — with implicit super samples, as well as with the type of social-circle questions and expectation questions that presumably tap into them. All the more reason, then, to expect that future work probing further into the prospects for and interrelations between social-circle surveys, expectation question, and prediction markets in particular would be fruitful and worthwhile.[7]

References

Ackland, Robert, and Jamsheed Shorish. 2014. "Political Homophily on the Web." In *Analyzing Social Media Data and Web Networks*, edited by Marta Cantijoch, Rachel Gibson, and Stephen Ward, 25–46. London: Palgrave Macmillan UK. https://doi.org/10.1057/9781137276773_2.

Ahlstrom-Vij, Kristoffer. 2016. "Information Markets." In *A Companion to Applied Philosophy*, 89–102. John Wiley & Sons, Ltd. https://doi.org/https://doi.org/10.1002/9781118869109.ch7.

———. 2019. "Self-Resolving Information Markets: A Comparative Study." *Journal of Prediction Markets*, February. https://doi.org/10.5750/jpm.v13i1.1687.

Alford, John R., Peter K. Hatemi, John R. Hibbing, Nicholas G. Martin, and Lindon J. Eaves. 2011. "The Politics of Mate Choice." *The Journal of Politics* 73 (2): 362–79. https://doi.org/10.1017/S0022381611000016.

[6] Note that events that are 42.5% probable happen fairly regularly, and that this estimate therefore might very well have been completely accurate.

[7] Disclosure Statement: The author has consulted for software companies that develop forecasting technologies and holds shares in one company that develops such software. None of these companies have provided any funding for or otherwise played a role in the studies reported on in this paper. Funding for the experimental study was provided by the Department of Philosophy at Birkbeck, University of London, where the author is employed as a full-time academic researcher. (The simulation studies required no funding.)

Berg, Joyce E., Forrest D. Nelson, and Thomas A. Rietz. 2008. "Prediction Market Accuracy in the Long Run." *International Journal of Forecasting* 24 (2): 285–300. https://EconPapers.repec.org/RePEc:eee:intfor:v:24:y:2008:i:2:p:285-300.
Berg, Joyce E., and Thomas A. Rietz. 2014. "Market Design, Manipulation, and Accuracy in Political Prediction Markets: Lessons from the Iowa Electronic Markets." *PS: Political Science and Politics* 47 (2): 293–96. https://doi.org/10.1017/S1049096514000043.
Berg, Joyce, and Thomas Rietz. 2006. "The Iowa Electronic Markets: Stylized Facts and Open Issues." In *Information Markets: A New Way of Making Decisions*, 142–69. AEI-Brookings Joint Center for Regulatory Studies.
Bernstein, Daniel M, Cristina Atance, Geoffrey R Loftus, and Andrew Meltzoff. 2004. "We Saw It All Along: Visual Hindsight Bias in Children and Adults." *Psychol Sci* 15 (4): 264–67. https://doi.org/10.1111/j.0963-7214.2004.00663.x.
Butters, Ross, and Christopher Hare. 2020. "Polarized Networks? New Evidence on American Voters'political Discussion Networks." *Political Behavior*. https://doi.org/10.1007/s11109-020-09647-w.
Enns, Peter K., Julius Lagodny, and Jonathon P. Schuldt. 2017. "Understanding the 2016 US Presidential Polls: The Importance of Hidden Trump Supporters." *Statistics, Politics and Policy* 8 (1): 41–63. https://ideas.repec.org/a/bpj/statpp/v8y2017i1p41-63n5.html.
Galesic, M., and W. Bruine de Bruin. 2020. "Election Polls Are More Accurate If They Ask Participants How Others Will Vote." *The Conversation*, November. https://theconversation.com/election-polls-are-more-accurate-if-they-ask-participants-how-others-will-vote-150121.
Galesic, M., Bruine de Bruin W., Dumas M., Kapteyn A., Darling J. E., and E. Meijer. 2018. "Asking about Social Circles Improves Election Predictions." *Nature Human Behaviour* 2 (3): 187–93. https://doi.org/10.1038/s41562-018-0302-y.
Graefe, Andreas. 2017. "Prediction Market Performance in the 2016 U.S. Presidential Election." *Foresight: The International Journal of Applied Forecasting*, no. 45: 38–42. https://ideas.repec.org/a/for/ijafaa/y2017i45p38-42.html.
Hanson, Robin. 2013. "Shall We Vote on Values, but Bet on Beliefs?" *Journal of Political Philosophy* 21 (2): 151–78. https://doi.org/https://doi.org/10.1111/jopp.12008.
Huber, Gregory A., and Neil Malhotra. 2017. "Political Homophily in Social Relationships: Evidence from Online Dating Behavior." *The Journal of Politics* 79 (1): 269–83. https://doi.org/10.1086/687533.
Klar, Samara, and Yotam Shmargad. 2017. "The Effect of Network Structure on Preference Formation." *The Journal of Politics* 79 (2): 717–21. https://doi.org/10.1086/689972.
Mchugh, Patrick, and Aaron Jackson. 2012. "Prediction Market Accuracy: The Impact of Size, Incentives, Context and Interpretation." *Journal of Prediction Markets* 6 (2): 22–46. https://EconPapers.repec.org/RePEc:buc:jpredm:v:6:y:2012:i:2:p:22-46.
McPherson, M., L. Smith-Lovin, and J. M. Cook. 2001. "Birds of a Feather: Homophily in Social Networks." *Review of Sociology* 27: 415–44.
Montgomery, Jacob M., Brendan Nyhan, and Michelle Torres. 2018. "How Conditioning on Posttreatment Variables Can Ruin Your Experiment and What to Do about It." *American Journal of Political Science* 62 (3): 760–75. https://doi.org/https://doi.org/10.1111/ajps.12357.
Murr, A., M. Stegmaier, and A. Lewis-Beck. 2021. "Vote Expectations Versus Vote Intentions: Rival Forecasting Strategies." *British Journal of Political Science* 51 (1).
Nisbett, RE, and Z Kunda. 1985. "Perception of Social Distributions." *Journal of Personality and Social Psychology* 48 (2): 297—311. https://doi.org/10.1037//0022-3514.48.2.297.
Palan, Stefan, and Christian Schitter. 2018. "Prolific.ac—a Subject Pool for Online Experiments." *Journal of Behavioral and Experimental Finance* 17: 22–27. https://doi.org/https://doi.org/10.1016/j.jbef.2017.12.004.
Peer, Eyal, Laura Brandimarte, Sonam Samat, and Alessandro Acquisti. 2017. "Beyond the Turk: Alternative Platforms for Crowdsourcing Behavioral Research." *Journal of Experimental Social Psychology* 70: 153–63. https://doi.org/https://doi.org/10.1016/j.jesp.2017.01.006.
R Core Team. 2018. *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. https://www.R-project.org/.

Ross, Lee, David Greene, and Pamela House. 1977. "The 'False Consensus Effect': An Egocentric Bias in Social Perception and Attribution Processes." *Journal of Experimental Social Psychology* 13 (3): 279–301. https://doi.org/https://doi.org/10.1016/0022-1031(77)90049-X.
Rothschild, David. 2009. "Forecasting Elections: Comparing Prediction Markets, Polls, and Their Biases." *The Public Opinion Quarterly* 73 (5): 895–916. http://www.jstor.org/stable/40467652.
Rothschild, David, and Justin Wolfers. 2012. "Forecasting Elections: Voter Intentions Versus Expectations." *Brookings Report, November 1*.
Servan-Schreiber, Emile, Justin Wolfers, David M. Pennock, and Brian Galebach. 2004. "Prediction Markets: Does Money Matter?" *Electronic Markets* 14 (3): 243–51. https://doi.org/10.1080/1019678042000245254.
Shanton, Karen, and Alvin Goldman. 2010. "Simulation Theory." *WIREs Cognitive Science* 1 (4): 527–38. https://doi.org/https://doi.org/10.1002/wcs.33.
Surowiecki, James. 2005. *The Wisdom of Crowds*. New York: Anchor Books.
Tversky, Amos, and Daniel Kahneman. 1974. "Judgment Under Uncertainty: Heuristics and Biases." *Science* 185 (4157): 1124–31. https://doi.org/10.1126/science.185.4157.1124.
Van Boven, Leaf, and George Loewenstein. 2003. "Social Projection of Transient Drive States." *Pers Soc Psychol Bull* 29 (9): 1159–68. https://doi.org/10.1177/0146167203254597.
Weisberg, Herbert. P. 2018. "Total Survey Error." In *The Oxford Handbook of Polling and Survey Methods*, edited by Atkeson Lonna and R. Michael Alvarez. New York, NY: Oxford University Press.