Article

Do Polls Influence Opinions? Investigating Poll Feedback Loops Using the Novel Dynamic Response Feedback Experimental Procedure

Sveinung Arnesen[1], Mikael P. Johannesson[1], Jonas Linde[1], and Stefan Dahlberg[2]

Abstract
Opinion polls may inadvertently affect public opinion, as people may change their attitudes after learning what others think. A disconcerting possibility is that opinion polls have the ability to create information cascades, wherein the majority opinion becomes increasingly larger over time. Testing poll influence on attitudes toward Syrian refugees and mandatory measles vaccination, we field survey experiments on a probability-based online survey panel. Through a novel automated procedure labeled the dynamic response feedback, we measure whether the answers from early poll respondents can influence the opinions of subsequent respondents who learn the answers of the previous respondents. Using this procedure, no feedback loops are identified.

Opinion polls inform people about what others think about different issues and where the majority opinion lies. This information may itself change the public opinion as people adapt change their views in a reaction to learning the opinions of others, potentially creating a self-fulfilling prophecy of public opinion (Rothschild & Malhotra, 2014). There is a disconcerting possibility that polls through feedback loops create "spirals of silence," where fear of isolation leads the minority to become increasingly silent and diminish in numbers over time (Noelle-Neumann, 1974). If polls have the potential to initiate such dynamics, they will be counterproductive to the ideals of democratic deliberation, which presuppose the existence of public discourse among free and equal individuals, and wherein political positions must be justified by arguments, not peer pressure (Elster, 1998).

[1] University of Bergen, Bergen, Norway
[2] University of Gothenburg, Gothenburg, Sweden

Corresponding Author:
Sveinung Arnesen, University of Bergen, Rosenbergsgaten 35, Bergen N-5007, Norway.
Email: sveinung.arnesen@uib.no

This study investigates the existence of poll feedback loops. We construct a survey experimental design that makes it possible to track the influence of an initial poll distribution across several subsequent poll iterations through the innovation of an experimental procedure labeled the *dynamic response feedback* (DRF). This automated procedure divides respondents in one survey wave into several groups; each group is exposed to the distribution of opinion from the previous group at the same time, as they are asked about their own opinion on the issue. If polls actually trigger self-fulfilling prophecies, this will be captured by our design. We here present experiments on attitudes toward two controversial issues: accepting more Syrian refugees into Norway and introducing mandatory measles vaccination for all children. The experiments were fielded in two separate waves of the Norwegian Citizen Panel (NCP); a probability-based online survey panel established for academic purposes.

The results from the DRF experiments show that when the respondents are exposed to polls, the aggregate level impact on public opinion is negligible, thus providing little reason to worry that polls are disruptive to public debate and deliberative processes.

Poll Influence in Nonelectoral Settings

Poll effects—and in particular the bandwagon effect—have been debated for decades (Hardmeier, 2008). If exposure to polls is actually affecting public opinion in the direction of the majority, it may be occurring through this mechanism (Mutz, 1998). The bandwagon effect originates from electoral studies and refers to voters who decide to vote for the expected winner (Lazarsfeld, Berelson, & Gaudet, 1948; Morton, Muller, Page, & Torgler, 2015). The concept has since been applied beyond the electoral context. It now generally refers to a situation wherein a person acquires and/or expresses a preference that is in accordance with the preference of the majority because it is intrinsically gratifying to be on the winning side.

There are relatively few studies on the influence of polls in a nonelection setting. Those that do exist demonstrate contradictory results (e.g., Marsh, 1985; Nadeau, Cloutier, & Guay, 1993; Ragozzino & Hartman, 2014; Sonck & Loosveldt, 2010). Recent experiments that presented respondents with opinion polls showing varying levels of support for different political issues found that respondents did react to the treatment and moved in the direction of the perceived majority on some—though not all—of the issues (Rothschild & Malhotra, 2014). Although the type of issue and the strength of the treatment mattered for the effect of the polls, the results provided evidence of the fact that a poll can influence political attitudes under certain circumstances.

In a worst-case scenario, opinion polls can—by their own power—have a permanent impact on the public opinion. The spiral of silence hypothesis claims that citizens who perceive their attitudes to be in the minority refrain from expressing their views publicly out of fear of social isolation; consequently, the dominant majority will therefore become more dominant and louder over time, while the losing minority will become increasingly silent (Matthes, 2015; Moy & Scheufele, 2000; Noelle-Neumann, 1974).

We investigate to which extent polls have the capacity to in and of themselves to trigger such dynamics. Does an initial poll influence subsequent distributions of public opinion? Is this first poll able to set public opinion on a dynamic path that would not have occurred otherwise?

Setting the Scene: Experimental Design and the DRF Procedure

The experimental design is as follows: Within a survey wave of several thousand respondents, 425 respondents are randomly distributed into one control group and one treatment group. In the control group, the respondents are asked their opinion about their attitudes toward a specific political issue. The treatment group is asked the same question, but they are also presented with the results from an

earlier poll and shown a pie chart with the share of respondents who disagree and agree on the issue in question. This first pie chart serves as the initial treatment condition. The initial treatment condition is generated from the first 10–15 responses in the survey. Hence, it is a real poll, but one that is statistically more likely to deviate from the population compared with a poll of several hundred respondents. In this way, we amplify the likelihood of observing a poll outlier, creating a situation that is less frequent yet not uncommon in the real-world. The respondents are informed that the pie chart is generated from previous responses in the survey wave, while the number of responses that the chart is based on is not provided.

When all respondents in the treatment group have answered, their responses are processed, and the responses then replace the initial poll, serving as the treatment condition for the treatment group at time t_2. Then, 425 new respondents are randomly assigned to a treatment or control group at time t_2.[1] The treatment group is asked the same question, but with the addition of being presented with the results of the poll at time t_1 and being shown a pie chart with the share of respondents who disagree and agree on the issue in question. When all respondents in the second treatment group have answered, their responses are processed and then replace the previous poll, serving as the treatment condition for the treatment group at time t_3. The procedure continues until the field period of the survey wave is ended.[2] This procedure is automated in the web survey through a script, which we label the dynamic response feedback (DRF) experimental procedure.

The DRF fetches previous answers given by respondents who have already completed the survey, and it presents—in real time—the current distribution to the respondent who is about to answer the question. For every ith response, the distribution of support for the issue is recalculated, thus creating several "mini polls" within the same survey wave. For each new treatment group, there is a complementary control group, ensuring that the only thing that varies between the treatment and the control group is the poll information and not, for example, external events that distort the level of support for the issue in question. In this way, we are able to investigate whether the effect of the initial treatment conditions vanishes or continues to have an impact on the aggregate distribution over several iterations.[3]

Issues and Questions

The first question concerned mandatory measles vaccination for children. During the field period (March 9 to March 30, 2015), this issue had received attention in the media and in public debate. From the media coverage, it could be inferred that a substantial share of citizens was opting out of vaccinating their children; one way of maintaining a high vaccination rate in society would be to make vaccination mandatory (as was later proposed by the main opposition party, the Labour party). The respondents in our experiment were asked to answer the following question, using a 7-point scale ranging from strongly agree to strongly disagree: "The vaccination of children has been heavily debated by the media recently. Some people think that it should be mandatory for all children to have measles vaccinations. To what extent do you agree or disagree with this?"

The second question asked about the accommodation of Syrian refugees in Norway. During the summer of 2015, the parliamentary majority had decided to accommodate 8,000 extra refugees over the next 3 years. This sparked a heated debate. The question was fielded during the period of October 28 to November 16, 2015, at a time when the refugee crisis dominated the news. Using the same response scale, respondents were asked "To what extent do you agree or disagree that, over the next 3 years, Norway should accept more Syrian refugees than was previously decided?" These two concrete issues were relatively fresh in the public's mind at the time. Hence, they arguably serve as cases in which people would "most likely" be influenced by what others think, as there were fewer available social cues about the opinions of others on these issues compared with issues that were more settled.

For each experiment, both the treatment groups and the control groups received four questions about the issue. In addition to the main question about how much they agreed or disagreed with the

statement, the respondents were asked about the strength of their opinion on the issue, how much knowledge they felt they had, and what they perceived public opinion to be on the issue.[4]

Data

We implemented the experiments in waves four and five of the NCP conducted in 2015 (Ivarsflaten, Arnesen, Böhm, et al., 2015; Ivarsflaten, Arnesen, Bjånesøy, et al., 2015). The NCP is a probability-based research-purpose online survey panel administered by the Digital Social Science Core Facility at the University of Bergen. See Blom et al. (2016) and Bosnjak, Das, and Lynn (2016) for discussions on probability-based online survey panels. A total of 4,582 respondents participated in at least one of the experiments, with 3,759 respondents participating in one experiment and 823 participating in two. For more details about response rates and other methodological issues, we refer to the reader to the NCP methodology reports (Skjervheim & Høgestøl 2015a, 2015b). The data are freely available for scholars via the Norwegian Centre for Research Data.

Identification

For each iteration, the respondents are either treated (by being shown a poll) or not. The average treatment effects are estimated to show the expected difference in answers between the respondents who received a poll (in that iteration) and the respondents who did not. If polls influence public opinion to agree with the majority, we should expect (a) an initial difference in attitudes between our treatment and control groups on the aggregate level and (b) that this initial difference increases over time. When the respondents, for instance, see a poll that shows higher support for an issue than there really is, even more respondents support the issue, creating an even more skewed distribution that is then presented to the next group of respondents, and so on.

We also want to see whether the effect is being mediated via respondents' stated strength of opinion, knowledge of the issue, or perception of others' attitudes. Therefore, we conduct an average causal mediation analysis (see Imai, Keele, Tingley, & Yamamoto, 2011), which can estimate the above but assumes that there are no confounding pre- or posttreatment variables that might affect either the respondents' main answer or the mediator. We do not estimate the average causal mediation effect (ACME) for each poll (i.e., iteration) but rather for all polls combined.

The analyses are conducted using the R programming environment with the mediation package (Tingley, Yamamoto, Hirose, Keele, & Imai, 2014).

Results

In the following, we present the most relevant results from the survey experiments. All results are available from the authors by request. The experimental procedures of the Syrian refugee issue (a) and the mandatory vaccine issue (b), respectively, are presented in Figure 1. The t_0 iteration generates the initial treatment condition, which is as a pie chart poll by the treated respondents at iteration t_1. For the Syrian refugee issue, this initial treatment condition shows a poll in which 40% of the respondents agree with the question statement. The corresponding poll for the vaccine issue is 93%. The circled numbers represent the proportion of the respondents who agree with the question statement at each iteration. The dotted arrows indicate how the responses at iteration t_k are, in turn, displayed as opinion polls to the following respondents at iteration t_{k+1}. For example, for the Syrian refugee issue, the proportion that agrees to allow more Syrian refugees into the country is 52% at iteration t_1. The treated respondents at iteration t_2 thus see a poll where 52% agree with the statement. The treated respondents at iteration t_3 see a poll where 61% agree and so on. The control
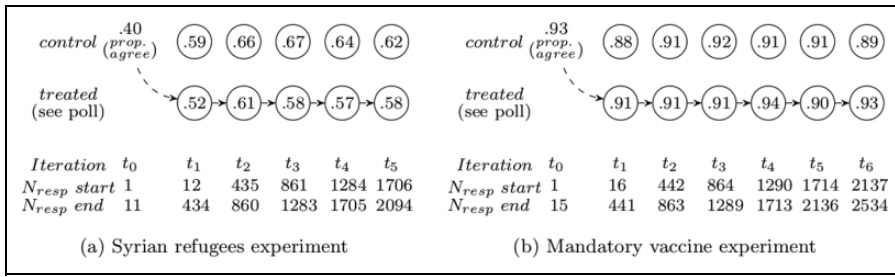
Figure 1. The dynamic response feedback experimental procedure implemented on the issues of Syrian refugees and mandatory vaccination.

group respondents receive no information about previous answers. The two bottom rows in Figure 1 show the cumulative number of respondents taking part in the experiments.

Figure 2 summarizes the average treatment effects on the treated groups for each of the post-treatment measures for the Syrian refugee experiment. The figure shows the (expected) difference in answers between those who are exposed to a poll and those who are not (the vertical axis) for each treatment iteration (the horizontal axis). The dotted line is the expected answer of the control group at the corresponding iteration.

In Figure 2, the main question refers to the acceptance of Syrian refugees. The t_1 group saw the t_0 distribution—a poll with 40% agreeing on accepting Syrian refugees. This is clearly lower than the true distribution. Yet, the average treatment effect at time t_1 is statistically insignificant. This pattern continues through all iterations.
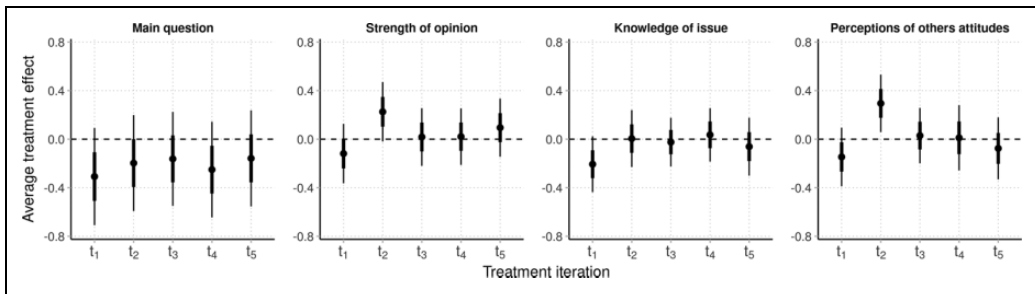

Figure 2. Average treatment effects over time between treatment and control groups for the Syrian refugee issue.

For the vaccine issue, the initial poll (t_0) to which the respondents in the first iteration were exposed showed 93% agreeing to make measles vaccination mandatory for children. Figure 3 shows that after being exposed to information stating that the vast majority agrees on mandatory vaccination, the respondents agree slightly more strongly than those who did not see that information. The treatment effect remains positive throughout the iterations, but again, the effect is not statistically significant from the control groups.

The respondents clearly shift their perceptions regarding what others think about the vaccine issue. This serves as a manipulation check, confirming that the treated respondents view the distribution as a signal of true public opinion. For the Syrian refugee issue, the respondents who are exposed to the poll do not significantly differ from the control group in this regard, possibly because the control group's perception is close to the distribution with which the treated respondents are presented. In sum, the experiments using the DRF procedure show negligible poll effects at the
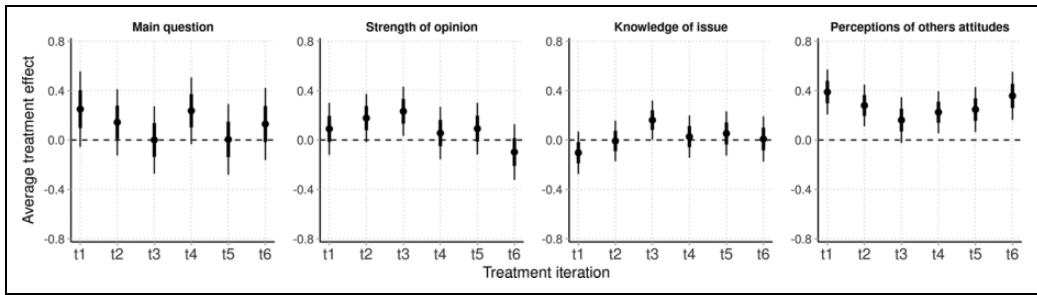
Figure 3. Average treatment effects over time between treatment and control groups for the mandatory vaccination issue.
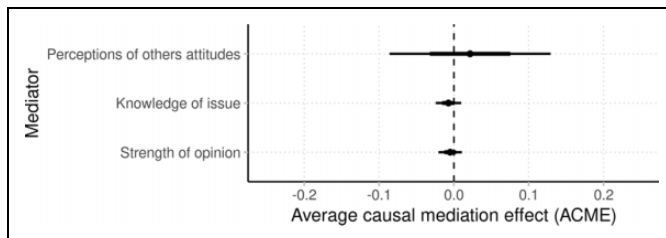

Figure 4. Average causal mediation effect of being shown a poll about what other people think about refugees coming to Norway.
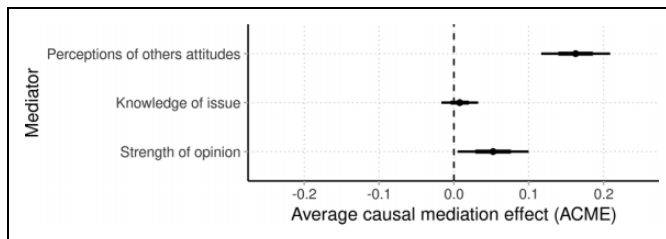

Figure 5. Average causal mediation effect of being shown a poll about what other people think about mandatory vaccination.

aggregate level.[5] It is evident that the polls in our experiments are not able to create information cascades wherein public opinion increasingly diverges from "true" public opinion. Based on these results, there is therefore little reason to fear that polls in and of themselves disrupt the public's deliberative process and become self-fulfilling prophecies. That said, the results from these experiments may not be representative of all kinds of issues. Similarly designed experiments for other issues in other contexts may produce different results. For the vaccine issue, the treatment groups increase from an already high level; in the control group, the mean response is 5.9 on a scale from 1 to 7. We cannot rule out that a ceiling effect reduces the impact of the poll treatment. The Syrian refugee issue is a high-salient issue, and perhaps, a low-salient issue that also had a high degree of polarization would have been more susceptible to changes in the aggregate opinion distribution.

While public opinion is, overall, little affected by individuals' exposure to information about the majority opinion, the causal mediation analysis is somewhat more supportive of poll effects. Figures 4 and 5 show the ACME of opinion strength as mediator, knowledge of issue as mediator,

and the respondents' perceptions of the attitudes of others as mediator on the Syrian refugee and the vaccine experiments, respectively.

Here, we observe that there is a significant positive mediating effect of the latter variable on the vaccine issue. In other words, this shows that those who see the polls on vaccination adjust their perceptions of others' opinions, which, in turn, changes their own attitudes on the issue. And, they change their opinions in the same direction as their perceptions of public opinion, which is an indication of a bandwagon effect.

Interestingly, there is also a large drop in the share of respondents who select the middle option on the answer scale (i.e., 4 = *neither agree nor disagree*). On both issues, between 25% and 29% fewer of the treated respondents reply that they neither agree nor disagree. Taken together, these results indicate that seeing a poll about the opinion of others makes the exposed respondents more opinionated and more likely to also take a stand themselves.

Conclusion: The Role of Polls in the Deliberative Process

In contemporary democracies, citizens are continuously exposed to the political opinions of other citizens, often through opinion polls presented in different types of media and carried out by a wide array of political actors, such as political parties, think tanks, and interest organizations. Polls inform political decision makers about citizens' views on political issues, information that helps them make policies that are responsive to the will of the people. The political consequences of opinion polls have been subject to debate, and both potentially positive and negative effects of polls have been proposed. On the one hand, it could be argued that the frequent use and presentation of opinion polls may lead to increased opportunities for the public to learn about important political issues and public opinion. It has also been argued that polls inform political decision makers about citizens' views on political issues, which helps in making policies that are responsive to the will of the people. On the other hand, there are concerns that opinion polls may become self-fulfilling prophecies because the polls themselves may drive public opinion.

In the real-world, using observational data, it is impossible to isolate opinion polls and their effects on public opinion. In this article, therefore, we present a novel experimental design to investigate the dynamics of the effect of opinion polls on public opinion. Rather than limiting the study to the one-time effect of these information treatments, we track their potential effect over time by applying the DRF procedure during the survey data collection. This procedure fetches previous answers given by respondents who have already completed the survey, and it presents—in real time—the current distribution to the respondent who is about to answer the question.

Given the prominent position of polls in the public political arena, there have been surprisingly few experimental studies on polls and how they may affect public opinion on their own. The results from our experiments at least indicate that there is little reason to worry about negative side effects from publishing public opinion polls. More research is still warranted before making strong inferences about the influence of polls on opinions. Rothschild and Malhotra (2014) found that people's susceptibility to change attitudes after poll exposure were highly issue dependent. Their study revealed that people seemed be more influenced by polls when they had weak pretreatment attitudes, knew little about other people's attitudes toward the issue, and when their attitudes were not hardened by partisan predispositions. Our results support their second finding, in the sense that we find a mediating effect of the perception of others' attitudes. For the vaccination issue, the respondents were not aware of the strong support for measles vaccination. Those who saw the polls adjusted their perceptions in the right direction, and we also observed a mediating effect on attitudes through the changed perceptions about what the opinions of others were. Summing up, people are to some extent affected by learning what others think about the issue in question, but the effect is not strong enough to have a significant effect on the aggregate public opinion.

Notes

1. While 425 respondents are assigned to a treatment or control group, not all may answer. Typically, between zero and five respondents are assigned to a group without participating. Hence, the attrition rate is negligible at less than 1%. The number of iterations depends on the number of survey respondents who take the survey: When 425 respondents have been assigned to a treatment or control group at time t_k, a new group is assigned to a treatment or control group at time t_{k+1}. The last iteration group therefore consists of fewer than 425 respondents.
2. The responses do not accumulate; so, the treatment group at time t is only shown the response distribution of the previous group at time $t-1$.
3. The experimental design is inspired by Jacobs and Campbell's (1961) successive iteration experiment, in which they measured whether initial misleading information by a confederate would influence not only the responses of the subject that was directly exposed to it but also indirectly affect the responses of subsequent subjects. When a confederate first expressed his or her untruthful opinion, some subjects would feel the social pressure to conform, modify their estimates in the direction of the confederate, and alter their responses. Their responses would again affect the subsequent responses of the next subject, and so on. Yet, after a few such iterations, the trace of the confederate's misleading information had evaporated.
4. All posttreatment questions and the plan for analysis were designed before the results were in. In all three experiments, the three follow-up questions were posed to the respondents using a 5-point unipolar answer scale: (1) "How strong are your views about this question?" (2) "How good do you feel that your knowledge about this subject is?" and (3) "If you were to guess, how many of Norway's citizens do you think would agree that [main issue statement repeated]."
5. A third experiment conducted in a similar fashion reveals similar results. The results from this third experiment will be made available upon request.

References

Blom, A. G., Bosnjak, M., Cornilleau, A., Cousteaux, A., Das, M., Douhou, S., & Krieger, U. (2016). A comparison of four probability-based online and mixed-mode panels in Europe. *Social Science Computer Review*, *34*, 8–25.

Bosnjak, M., Das, M., & Lynn, P. (2016). Methods for probability-based online and mixed-mode panels: Selected recent trends and future perspectives. *Social Science Computer Review*, *34*, 3–7.

Elster, J. (1998). *Deliberative democracy*. Cambridge, England: Cambridge University Press.

Hardmeier, S. (2008). The effects of published polls on citizens. In W. Donsbach & M. W. Traugott (Eds.), *The SAGE handbook of public opinion research* (pp. 504–513). Los Angeles, CA: Sage.

Imai, K., Keele, L., Tingley, D., & Yamamoto, T. (2011). Unpacking the black box of causality: Learning about causal mechanisms from experimental and observational studies. *American Political Science Review*, *105*, 765–789.

Ivarsflaten, E., Arnesen, S., Bjånesøy, L., Johannesson, M. P., Nordø, Å., & Tvinnereim, E. (2015). *Norwegian citizen panel wave 5*. Bergen, Norway: University of Bergen.

Ivarsflaten, E., Arnesen, S., Böhm, G., Elgesem, D., Gåsdal, O., Strømsnes, K., Tvinnereim, E., & Aars, J. (2015). *Norwegian citizen panel wave 4. Bergen*, Norway: University of Bergen.

Jacobs, R. C., & Campbell, D. T. (1961.) Generations of a laboratory microculture. *The Journal of Abnormal and Social Psychology*, *62*, 649–658.

Lazarsfeld, P. F., Berelson, B., & Gaudet, H. (1948). *The people's choice: How the voter makes up his mind in a presidential campaign*. New York, NY: Columbia University Press.

Marsh, C. (1985). Back on the bandwagon: The effect of opinion polls on public opinion. *British Journal of Political Science*, *15*, 51–74.

Matthes, J. (2015). Observing the "spiral" in the spiral of silence. *International Journal of Public Opinion Research*, *27*, 155–176.

Morton, R. B, Muller, D., Page, L., & Torgler, B. (2015). Exit polls, turnout, and bandwagon voting: Evidence from a natural experiment. *European Economic Review*, *77*, 65–81.

Moy, P., & Scheufele, D. A. (2000). Media effects on political and social trust. *Journalism & Mass Communication Quarterly*, *77*, 744–759.

Mutz, D. C. (1998). *Impersonal influence: How perceptions of mass collectives affect political attitudes*. Cambridge, UK: Cambridge University Press.

Nadeau, R., Cloutier, E., & Guay, J.-H. (1993). New evidence about the existence of a bandwagon effect in the opinion formation process. *International Political Science Review*, *14*, 203–213.

Noelle-Neumann, E. (1974). The spiral of silence: A theory of public opinion. *Journal of Communication*, *24*, 43–51.

Ragozzino, M., & Hartman, T. K. (2014). The influence of public opinion polls on issue preferencesRetrieved November 30, 2014 from http://dx.doi.org/10.2139/ssrn.2532324

Rothschild, D., & Malhotra, N. (2014). Are public opinion polls self-fulfilling prophecies? *Research & Politics*, *1*, 1–10.

Skjervheim, Ø., & Høgestøl, A. (2015a). Norwegian citizen panel 2015 fifth wave methodology report. Retrieved from http://www.nsd.uib.no/nsddata/serier/norsk_medborgerpanel_eng.html

Skjervheim, Ø., & Høgestøl, A. (2015b). Norwegian citizen panel 2015 fourth wave methodology report. Retrieved from Bergen, Norway, http://www.nsd.uib.no/nsddata/serier/norsk_medborgerpanel_eng.html

Sonck, N., & Loosveldt, G. (2010). Impact of poll results on personal opinions and perceptions of collective opinion. *International Journal of Public Opinion Research*, *22*, 230–255.

Tingley, D., Yamamoto, T., Hirose, K., Keele, L., & Imai, K. (2014). Mediation: R package for causal mediation analysis. *Journal of Statistical Software 59*, 1–38.

Author Biographies

Sveinung Arnesen is a Postdoctoral Fellow at the Department of Comparative Politics, University of Bergen. His research focuses on public opinion, democratic theory, and voting behavior. His work appears among other in *Comparative Political Studies, Electoral Studies, Political Studies, International Journal of Forecasting, the Journal of Prediction Markets*, and *the Journal of Social Research*.

Mikael P. Johannesson is a PhD candidate at the Department of Comparative Politics, University of Bergen, specializing in quantitative methodology and political behavior.

Jonas Linde is a Professor in political science at the Department of Comparative Politics, University of Bergen. His research on public opinion, political support, and e-government has been published in journals such as *Governance, Government Information Quarterly, European Journal of Political Research*, and *Political Studies*.

Stefan Dahlberg is an Associate Professor at the Department of Political Science, and Database Manager at the Quality of Government Institute (QoG), University of Gothenburg. Dahlberg's research focuses on representative democracy, democratic legitimacy, political parties and voting behavior. His work appears in journals such as *British Journal of Political Science, Public Opinion Quarterly, Electoral Studies*, and *European Journal of Political Research*.