



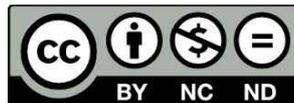
Newcastle University ePrints

O'Flynn I, Sood G. [What Would Dahl Say?: An Appraisal of the Democratic Credentials of Deliberative Polls and Other Mini-Publics](#). In: Grönlund, K; Bächtiger, A; Setälä, M, ed. *Deliberative Mini-Publics: Practices, Promises, Pitfalls*. Colchester, London: ECPR Press, 2014, pp.41-58.

Copyright:

© Ian O'Flynn and Gaurav Sood

Date deposited: 1 August 2014



This work is licensed under a
[Creative Commons Attribution-NonCommercial-NoDerivs 3.0 Unported License](http://creativecommons.org/licenses/by-nc-nd/3.0/)

ePrints – Newcastle University ePrints

<http://eprint.ncl.ac.uk>

What would Dahl Say?: An Appraisal of the Democratic Credentials of Deliberative Polls and other Mini-publics

Ian O'Flynn

Newcastle University, UK

Gaurav Sood

Stanford University, USA

Introduction

The concept of deliberation is both rich and complex. Without attempting too precise a definition, we take it to be a form of reasoned, open-minded discussion. Those who engage in it may have very different views but they must still be willing to listen to and reflect upon opposing arguments and respond to them seriously. Of course, deliberation is practiced in many different spheres, and to a variety of ends. In this chapter we focus on deliberation that contributes to democracy. In a democracy, important decisions of law and policy depend on the views and opinions of those who will be bound by them. Insofar as deliberation helps those bound by the decisions to refine their views and opinions, more thoughtful and informed policy decisions will result.

That, at least, is the claim. In practice, however, much will depend on the conditions under which deliberation occurs. As democrats, we want people to deliberate under conditions that can themselves be described as 'democratic'. But there is plenty of disagreement about precisely which conditions (embodied in an institution, procedure, practice etc.) ought to apply (Weale 1999: 40-42). In this chapter, we take our lead from Robert Dahl and in particular from his *Democracy and Its Critics* (1989). As he argues in that book, the most basic democratic standard of all is the principle of political equality. That principle supposes not just that each person is to have his interests treated with equal consideration, but also that each person is to have an equal right to say what his interests are. That is, of course, a relatively abstract formulation. Yet according to Dahl, the extent to which a political institution satisfies this principle can be assessed by judging its adherence to five more specific procedural criteria, namely inclusion, effective participation, enlightened understanding, voting equality at the decisive stage, and control of the agenda (Dahl 1989: chaps. 8-9).

While other criteria can be invoked, we think that these five offer a fairly objective set of standards by which to assess the democratic credentials of mini-publics in general and deliberative polls in particular.¹ Before proceeding with our analysis, however, a number of provisos need to be entered. Dahl claims that we could not call a political process democratic unless it satisfied the five criteria. Yet that claim needs to be appropriately understood. Since these are criteria of an ideal political order, we might in practice allow some deviation from them and still call that order democratic. Nor would we necessarily expect every democratic institution to satisfy all five – for example, participants in deliberative polls do not vote, let alone vote at the decisive stage (more on which below). However, the fact that a particular institution does not satisfy one or other of them can still be telling. It can remind us that a certain ‘sequencing’ of the democratic process is probably inevitable (Goodin 2005).

Inclusion

The demos should include all adults subject to the binding collective decisions of the demos (Dahl 1989: 120).

One of the most fundamental questions in democratic theory is who to include in the ‘demos’ (see, e.g., Goodin 2007; Miller 2009). The answer, according to Dahl, lies in a form of inclusive egalitarianism: if one is to be bound by a decision, one should be included on equal terms in the process by which that decision is made. So, to what extent do deliberative polls and other mini-publics conceptualise the demos as one that includes all who are potentially bound by its decisions, and to what extent do they succeed in turning concept into reality?

The answer to the first part of this question is complicated by difficulties in identifying who is bound by the decision. Take, for example, the Northern Ireland deliberative poll on education reform (Luskin et al. 2014) which limited participation to parents of school-aged children. On the one hand, this restriction seems appropriate since parents of school-aged children constitute the current pool of adults who would be the bound by the reforms (if there were to be any). Yet on the other hand, one might think it odd to exclude adults who plan to have children in the future and who may therefore find themselves bound by decisions in which they had no part. Time therefore complicates assessments of who is bound by a decision.

Such assessments are also complicated by the fact that a decision may affect some people directly while affecting others only indirectly. For example, since decisions about education policy inevitably have implications for the public purse, one might reasonably argue that all adults should be included, whether or not they have children in school. Indeed, even if it were the case that some people were in no way bound by the decision, one might want to include them anyway. For instance, even if the effects of the policy were limited to parents of school-aged children, one might still want to include those adults whose children have already gone through the education system so as to learn from their experiences. Those experiences may suggest different ways of ‘framing’ education policy, and hence suggest very different categories of assessment. In this vein, some deliberative theorists have suggested focusing on the representation of ‘discursive frames’ as well as persons or groups (Dryzek and Niemeyer 2008). Implementing this suggestion, however, poses novel challenges. How, for example, could we be sure that all of the relevant discourses were expressed or included in the course of a deliberative poll? Since discourses are not person-centred, conventional sampling, focused on people, may not do the trick.

Besides questions about who to include are questions about how to include. By definition, literally including all is not an option for mini-publics. However, not including all comes at a cost. The cost depends on the diversity of interests, opinions, private knowledge, social or demographic groups, or some combination thereof in the population. For instance, a single person can ‘represent’ all if everyone is the same. But if everyone is different, then nothing less than including all will do. More generally, more diversity means that we lose more when a few hundred speak for a few million. Yet while mini-publics do not include all (or even attempt to recruit larger samples commensurate with greater diversity of opinion on certain issues), their designers continue to claim that they are inclusive (Goodin and Dryzek 2006: 221). Typically, they justify such claims by arguing either (1) that the mini-public does not restrict access to anyone or (2) that the mini-public gives everyone an equal chance to be included.

For example, in the United States, executive agencies such as the Food and Drug Administration and the Environmental Protection Agency are required to hold public hearings before adopting any major regulation (Fung 2003: 342). Since public hearings are technically open to all, it also allows these agencies to claim that they are inclusive. Yet while taking an open-to-all approach is certainly a convenient way of

'recruiting' a sample, such samples of convenience often pose inconvenient questions. Who exactly attends public hearings? Are they mainly drawn from the ranks of those who feel strongly about an issue, the more educated, the representatives of well organised and richly resourced interest groups?² If there were any doubt, data amply confirm such distortions (Golden 1998).

The second justification has been used by the various Citizens' Assemblies, deliberative polls and more recently, by the Australian Citizens' Parliament (Dryzek 2009: 2). However, claims about giving everyone an equal chance can falter because of weaknesses in the theory underlying recruitment procedures (sampling) and due to failures in implementation. For instance, the British Columbia (BC) Citizens' Assembly on Electoral Reform used quota sampling. Quota sampling can be understood as stratified random sampling without (or weakened adherence to) randomisation: instead of randomisation, quota samples rely on interviewer discretion. However, reliance on interviewer discretion tends to yield samples that have significant biases in who is included in the sample (see Moser and Stuart 1953; Jowell et al. 1993).³

By contrast, deliberative polls typically rely on 'random' sampling.⁴ While in theory random sampling stands on firmer ground than quota sampling (see, e.g., Lynn and Jowell 1996), it can still yield samples from which entire groups (or segments of the population) are missing. In practice, this caveat rarely needs to be applied and can be solved by stratified random sampling (as practiced by the Australian Citizens' Parliament). Doing so, however, requires prior knowledge of the relevant groups and their proportions in the population at large.

The far more important and vexing issue with respect to random sampling is the slippage between theory and practice. In practice, random sampling often reduces to Random Digit Dialling (RDD).⁵ RDD frequently fails to ensure that each member of the relevant population has an equal chance to be part of the participant pool. Such failures are a consequence of either non-coverage or non-response.

Non-coverage is the failure to include all in the sampling frame. For example, the homeless, prisoners, those in hospitals and active-duty soldiers often have an exactly zero chance of being a participant. More generally, sub-populations with lower rates of ownership of fixed landline phones are liable to be under-represented. To be sure, the sub-populations that lack access to fixed landlines may vary over time. Whereas in the past household income was an important variable (Groves 1989: 117-

119), the problem today is mostly due to the rise of cell phone-only households: in some countries, cell phone directories do not exist while in other countries the auto-dialling of cell phones is not allowed. Omission of cell phone-only households means that certain subgroups, especially the young, are underrepresented in RDD surveys (see Keeter et al. 2007).

By contrast, non-response is the failure to contact certain (kinds of) people. It is usually due to a lack of resources. Lack of time, money and interviewer training can all exacerbate non-response bias. For instance, notable gender asymmetries in the Northern Ireland deliberative poll were attributed to a ‘compressed field period’ which did not leave enough time to contact enough men (Luskin et al. 2014); the fact that many Northern Irish women stay at home to rear their children while their husbands go out to work meant that men were harder to reach. In general, the net effect of such failures may be such that, in practice, not a lot may separate poorly done surveys based on probability sampling from survey samples recruited via vastly inferior techniques.

Recruitment of the initial sample is but one half of the process that determines who participates in a deliberative poll. The surveyed respondents (or a random subset) are invited to participate, and often enough a great many of the invited – roughly 75% on average – do not accept the invitation.⁶ However, the proportion of the invitees who eventually participate varies widely across polls, and across sub-populations. For instance, while nearly 92% of those invited participated in a local deliberative poll in China (Fishkin et al. 2010) only about 12% did so in a deliberative poll in Greece (Fishkin et al. 2008), and only 4.1% did so in a deliberative poll in Argentina.⁷ More troublingly, in the Greek deliberative poll, the percentage of women in the participant sample was 23% lower than their percentage in the recruitment sample.

Little is known of why participation rates differ so widely across polls. However we have some knowledge of the biases in who turns up. For instance, more knowledgeable invitees are significantly more likely to participate (Westwood and Sood 2010). That suggests that political interest plays a role in the decision to (not) attend deliberative polls. Regularly failing to include those not interested in politics along with the chasmic differences that sometimes appear on socio-demographic variables undercut claims to a reliable process that guarantees inclusion. In fact, a variety of inducements (such as honorariums) that can potentially reduce asymmetries in participation are applied patchily, and from what we gather, offered when resources

are available, and not when researchers think it necessary. If practitioners are serious about turning concepts into reality, they will need to be frank about acknowledging failures⁸, serious about studying them, and punctilious about implementing lessons learned from analyses.

Effective participation

Throughout the process of making binding decisions, citizens ought to have an adequate opportunity, and an equal opportunity, for expressing their preferences as to the final outcome. They must have adequate and equal opportunities for placing questions on the agenda and for expressing reasons for endorsing one outcome rather than another. (Dahl 1989: 109)

According to Dahl, a democratic process must give everyone an equal opportunity to have his say. But even where the opportunity is formally the same for everyone, in practice major disparities can still occur. For reasons already discussed, imperfections in recruiting mean that some people have less of an opportunity to participate than others. But even those who participate may not be able to participate on an equal footing. The garrulous, those strongly attached to their views, those who think the issue is important, the self-righteously knowledgeable, among many other species, all prefer talking to listening, often at the expense of giving others the chance to air their views. Then there are those who, even when they have the opportunity to talk, talk very little. In short, to ensure equality of opportunity, one typically cannot rely on formal rules alone.

Unfortunately, pro-active measures can fail. For instance, deliberative polls use trained moderators to facilitate the small group discussions.⁹ Their purpose is to ‘maintain an atmosphere of civility and mutual respect, encourage the diffident, restrain the loquacious and ensure that all the major arguments for and against in the briefing document get aired’ (Fishkin and Luskin 2005: 288). Even so, the fact remains that many participants hardly speak at all (Kim, Siu and Sood 2009).¹⁰ Insofar as they do not speak because they are crowded out by other participants, deliberative polls fail, in practice, to facilitate effective participation.

This latter worry extends the plenary-session component of a deliberative poll. Unlike, say, the participants in a citizens’ jury, participants in a deliberative poll are not asked to reach a decision, let alone a consensus, on the policy issue under discussion. Rather, what they are asked to do is to come up with an agreed set of questions that are then put the expert panel in the plenary sessions. Yet insofar as

some people hardly speak at all, the questions that are put to the panel are generated under conditions in which actual equality of opportunity falls short of the ideal.

The near silence of so many participants would be yet more troubling if members of socially disadvantaged groups were especially less likely to talk. At first glance, this fear appears well founded. Women and the less-educated do indeed speak less than men and the more highly educated. But after adjusting for prior levels of knowledge, and some other background characteristics, speaking-time is not predicted by socio-demographic characteristics (Siu 2009).¹¹ Even so, one cannot simply say that deliberative polls meet the criterion of equal participation. While on average participants from socially disadvantaged groups may be no less likely to speak, they may speak less in certain kinds of small group. Random assignment to small groups generally produces a fair amount of ‘variation in variation’ – some groups are more internally diverse than others. So, while certain minorities are well represented in some small groups, they are not well represented in others. Consequently, in some cases they may lack the ‘critical mass’ and hence the confidence to voice their own concerns (James 2008: 120-123). Again, we do not know to what extent this is so, but we recommend that scholars investigate this issue further.

What we do know, however, is that the more knowledgeable participants do speak vastly more than the less knowledgeable participants. On the face of it one might not see this as especially troubling. For although there is no necessary relation between a person’s being informed and his offering an informed argument, it seems likely that the more informed participants will, on average, offer a greater number of relevant facts when presenting arguments than the less informed participants.

On the other hand, one cannot simply assume that the more knowledgeable will speak for the ignorant or for those who are too diffident to speak. But even if we could assume this, democrats could still not rest content. Democrats do not want one group to speak for another but instead want each group to speak for itself.¹² In short, while there may be nothing about the design of deliberative polls or other mini-publics that biases the discussion process in favour of some participants, in practice much more needs to be done to ensure greater – that is, genuine, as opposed to merely formal – equality of opportunity.

A similar concern arises with respect to adequacy of opportunity. Deliberative polls usually take place over a weekend. But since some time must be spent organising the participants into groups, showing them to their rooms, explaining how

the sessions will work, introducing them to one another, hearing from the moderator, and so forth, the actual time available to participants for formal discussions is far less than one might imagine. A small group of 20 participants allowed five hours to discuss multiple issues leaves only ten minutes for each participant on average.¹³ We do not know if this is adequate but we would like to signal it as yet another issue that deserves closer attention. One way to study the problem would be to simply survey the participants about whether they thought the time available for discussion was adequate. There are, however, potential complications, including the fact that ‘adequate’ time is likely to vary by issue and by diversity of opinions (and information) on the issue.

Enlightened understanding

Each citizen ought to have adequate and equal opportunities for discovering and validating (within the time permitted by the need for a decision) the choice on the matter to be decided that would best serve the citizen’s interests. (Dahl 1989: 112)

As we explained above, the principle of political equality entails that each person has the right to say what her interests are. The right is predicated on the assumption that each person is the best judge of his own interest, which is, of course, what many critics of democracy doubt. Yet deliberative democrats are not only committed to treating each person as the best judge of his own interest, but are also committed to ensuring that each person has the opportunity to learn what he needs to know to arrive at an informed judgement (Dahl 1998: 39).

So what is it that each person needs to know to make an informed judgement? Unsurprisingly, the list is dauntingly long: policy relevant facts on all sides of the issue, logic and evidence behind conjectures about the probability of success of various alternatives, potential implications of the policy for oneself and others, the policy preferences of others and the reasons for their preferences, opportunity costs of adopting one policy as opposed to another, and much, much more. Not only is the list long, but each item on it is potentially overwhelmingly large – for instance, the domain of facts relevant to a single piece of legislation can easily run into thousands of pieces of information, if not more.

It follows, then, that any deliberative forum can no more than *aspire* to provide all the information that participants need to know. Yet it is also likely that

some mini-publics go further than others in satisfying this *desideratum*. Unfortunately, the task of comparing different mini-publics is hampered by lack of common measures, so again we restrict our analysis here to evidence from deliberative polls. In particular, we investigate which particular features of the deliberative poll design appear to bolster learning and which appear to be problematic.

Deliberative polls provide especially rich opportunities for acquiring information about the different sides to a policy issue, including balanced briefing materials, moderated discussions and access to experts during plenary sessions. As one might therefore expect, people on average learn a fair amount of policy-relevant factual information (Luskin et al. 2009). However, closer inspection reveals a fair bit of systematic variation around the mean – for example, better-educated participants who start out ill-informed learn more than similarly ill-informed but less-educated participants (Luskin et al. 2009).¹⁴

Some of this variation is explained by the different design elements of the deliberative poll. For instance, data suggest that roughly two-thirds of the acquisition of ‘public’ information, including that contained in the briefing materials, happens *before* the small group deliberations (see Hansen 2004; Luskin et al. 2009). Hence, it is likely that mini-publics that do not provide briefing materials to participants come up short in this respect. Of course, briefing materials are not a panacea. A non-trivial proportion of the deliberative poll participants report that they do not read the briefing materials completely.¹⁵ There is, moreover, a pattern to who reads the briefing materials, with women and the better-educated reading more of them. This matters because the extent to which people read briefing materials affects how much they learn (Luskin et al. 2009).

Selective attention, processing and retention of information are common human traits. This means that even where participants are provided with balanced briefing materials, most of the learning is limited to attitudinally congenial information (Hansen 2004). While this bias in information acquisition is mostly corrected by small-group deliberations (Hansen 2004), the pattern of adverse side-effects to good remedies continues. Varying composition of small groups creates variation in opportunities to learn. For instance, those assigned to groups whose members are more knowledgeable and those assigned to more attitudinally diverse groups learn more policy-relevant facts than those assigned to groups whose members are less knowledgeable or those assigned to groups whose members are less

attitudinally heterogeneous (Luskin et al. 2009).¹⁶ This problem affects not just the learning of ‘public’ information but also the learning of ‘private’ information that only emerges during the course of the small-group discussions; the learning of private information also varies across groups because different people have different private information.¹⁷ In all, the opportunities for learning are unequal and, courtesy of randomisation (where that is practiced), arbitrarily so.

There are a variety of other, hitherto unexplored, concerns. Here we only mention two. First, participants who feel strongly about a particular side of an issue may act strategically to raise doubts about correct information or, if it better suits their purposes, withhold correct information that they are privy to. Hence, participants may fail to learn correct information or end up doubting its veracity. Secondly, and arguably more worryingly, since at least some participants are liable to be misinformed, and since information offered in small groups is not vetted, it is likely that incorrect information is offered in group discussions. Correcting misinformation offered by others in small groups depends on those holding correct information speaking up, which may or may not happen. Consequently, it is possible (indeed, likely) that small group discussions impart both correct *and* incorrect information. To date, we do not know to what degree this occurs.

Some of the points that we have just raised readily translate into guidance for the design of a mini-public. In our view, all mini-publics should consider providing briefing materials in advance of deliberation. Ideally, such materials should be as accessible as possible—for example, in response to worries about citizen competence, briefing materials could contain illustrations to help elucidate complex points. Their designers may also want to explore instructor-led (online or in-person) sessions as a way of addressing the fact that some participants do not (for one reason or another) read the briefing materials on their own. In order to reduce variation across small groups one could employ stratified random assignment, with the relevant strata being ‘knowledge’ and ‘policy preferences’. Where learning is hampered by unequal ability, ‘affirmative action’ may be necessary to produce equal learning. For example, the less cognitively-able could be provided with additional help in the form of supplementary video materials. Finally, in order to judge whether people really are learning (or gaining the information that they need to make an informed judgement), designers could survey the participants multiple times over the course of their discussions to see how their attitudes change across survey times (see Goodin and Niemeyer 2003). Here

the thought is that the rate of change will decline as people proximate their ‘fully informed preferences’. Of course, the proposal assumes that the information that is presented, both in the briefing materials and in the small-group discussions, is balanced. Failure on this count means that lack of change in attitudes may not indicate a normatively defensible stasis, but in fact reflect an end result of one-sided persuasion.

Voting Equality at the Decisive Stage

At the decisive stage of collective decisions, each citizen must be ensured an equal opportunity to express a choice that will be counted as equal in weight to the choice expressed by any other citizen. In determining outcomes at the decisive stage, these choices, and only these choices, must be taken into account. (Dahl 1989: 109)

While deliberative polls and other mini-publics are not decisive in the way Dahl intends – that is, they do not usually result in universally binding policy decisions – we can nevertheless ask about the extent to which they guarantee participants ‘an equal opportunity to express a choice’.

While those responsible for designing deliberative polls and other mini-publics have mainly been concerned about the conditions that ought to precede the making of a choice, there is a need to think further about the implications of how exactly choices are registered. Some deliberative fora – for example, town hall meetings – use public voting. Yet those who hold unpopular preferences may lack the courage to express them openly and hence either stay silent, or adapt their preferences in ways that bring them more into tune with those of the majority (see Elster 1983).

Worries of this sort may be attenuated by means of the secret ballot, the approach adopted by the BC Citizens’ Assembly (see Warren and Pearse 2008), or the confidential questionnaire, the approach used by deliberative polls. However, confidential questionnaires bring problems of their own. At least partly as a result of lower social pressure to register a choice, some people may leave a question unanswered on the questionnaire, despite having a view about that question. The numbers here may be small (see, e.g., Luskin et al. 2006: 187; Hansen 2004) but they are likely to come disproportionately from the ranks of the less educated, racial minorities and other socially disadvantaged groups (see Berinsky 2004).

There are, therefore, worries about how choices are registered. There are also worries about what exactly gets counted. While participants in a deliberative poll are surveyed both before and after the event, there are problems in interpreting changes. Due to lack of motivation, social desirability pressures and other such problems linked to the fact that initial surveys are typically done over the phone (see Holbrook, Green, and Krosnick 2003), measurement error can be much greater in pre-deliberation surveys than post-deliberation surveys. (The latter are typically ‘self-complete’ questionnaires and respondents are often highly motivated to respond accurately, having gone through an extensive process.) For example, when survey questions are asked in ‘agree or disagree’ format, ‘yea-saying’ tendencies can disproportionately contaminate pre-deliberation responses, thereby rendering any pre-post movement suspect (see Weiksner 2008). Insofar as this occurs, the worry is that bogus ‘votes’ are being tallied.

In addition to questions about the processes of registering and counting preferences, imperfections in the participant sample must again be taken into account. As we noted above, even if self-selection is precluded, the sample may still differ from the population. It is common to correct such imperfections using post-stratification, reverse entropy weighting, regression, matching, and so forth. However, as far as voting equality is concerned, such techniques may not be enough. Suppose that a mini-public assigns people to small groups randomly and that composition of small group composition has some influence on the outcome we plan to tally. It naturally follows then that for the same sample, different random assignments will produce different outcomes. The potential outcomes (decisions, preferences, or what not) of each person are now a great many more than two, which means that the traditional way to estimate the effect of the ‘treatment’ – simple difference in means between the treatment and control group - may be inadequate. In response to such concerns, we suggest the following: either work to make small groups similar to each other¹⁸, making any talk of group composition moot, or estimate a model for how small-group composition matters for all variables of interest, and then simulate all random assignments to small groups. This will give us a distribution of effects, with each potential treatment effect having its own standard error. Of course, there may be other approaches that one could take. But the point remains that we need to work harder to ensure that mini-publics do more to ensure equality at the decisive stage.

Control of the Agenda

The demos must have the exclusive opportunity to decide how matters are to be placed on the agenda of matters that are to be decided by means of the democratic process. (Dahl 1989: 113)

The power to control the political agenda is the power to determine which political issues will be discussed, the order in which they will be discussed, and the options that will be decided on (Bachrach and Baratz 1962). While in principle that power should be distributed equally amongst the members of a democratic society, popular control over the agenda is notoriously difficult to secure (see, e.g., Riker 1982).

It will therefore come as little surprise that deliberative polls and other mini-publics have, to date at least, fared poorly on this criterion (Bua 2012). The Danish model of the consensus conference seems to allow the freest rein: participants get to participate in setting the agenda and also have a say in the selection of the experts to be heard (Hendricks 2005). Yet even here there is little control over the basic terms of reference or, more specifically, which area of technological development is to be discussed (Co-Intelligence Institute 2003). To take another example, participants in the BC Citizens' Assembly were charged with reviewing the provincial electoral system and, if necessary, recommending an alternative system which would then be put to a referendum. On the face of it, this was a serious democratic responsibility involving a great deal of political agency and control. Yet participants were severely hamstrung from the start. While they could recommend a new electoral system, they could not recommend any attendant changes to the composition of the provincial legislature (which was to remain at 79 seats). That ruled out electoral systems that needed a relatively large number of legislative seats to work well – including MPP, which, as a number of commentators have argued, might have delivered the stronger sense of local representation that most participants seemingly desired (Lang 2008: 92; Smith 2009: 89).¹⁹

In fairness, part of the failure to devolve control of agenda to participants is simply a consequence of extant funding realities. For the most part, deliberative polls and other mini-publics are funded by organizations (governments, political parties, NGOs etc.) interested in answering a particular question. It is our sense that most if not all mini-public designers would happily hand some control of the agenda over to participants if the right sort of organisational support were to become available (see, e.g., Dryzek 2009: 5). But even if mini-publics were to devolve control over the

choice of issue to be discussed, other concerns would remain – including concerns about the framing of the issue in the briefing materials, the choice of alternatives on which participants may offer their opinions, and the order in which issues are considered.

Lack of control over the crafting of the briefing materials naturally gives rise to concerns about their neutrality, comprehensiveness and correctness, not least of all because the process by which briefing materials are constructed tends to be woefully opaque. We need to know a lot more about where the information they contain comes from, what resources were expended in drawing them up, how the order in which information is presented was determined, which organisations (or ‘experts’) were consulted and how they were consulted (e.g., did any organisation have a veto), etc. The lack of data (a consequence of opaque procedures) means that it is hard to make progress on these issues, other than by relying on reports from participants, which may not be very diagnostic. Unsurprisingly, then, our recommendation is for greater transparency.

The construction of the questionnaire is another area of concern. Deliberative polls typically ask participants to mark their preferences on closed-ended scales, such as ‘on a scale of 0 to 10, where 0 is opposing something as strongly as possible, 10 is supporting it as strongly as possible, and 5 is exactly in-between, how strongly would you support or oppose policy X’. Yet, partly as a result of the structure of the questionnaire and partly as a matter of constraints on available time, closed-ended scales have the potential to poll people only on a subset of concerns about any one topic – concerns chosen by the creators of the questionnaire. As such, closed-ended scales provide inadequate opportunities for participants to note their own concerns about the issue, which one might view as an analogue of failure of control over the agenda. This is not to denigrate instruments of this sort; closed-ended measures are extremely valuable – they are efficient and cost-effective, and can capture a fair bit of important information.²⁰ But it does suggest that deliberative poll designers ought to consider making greater use of open-format measures, soliciting people’s opinions about the issues under discussion in more general terms, especially as we continue to make gains in analysis of textual data.

Finally, when multiple issues are discussed, one must worry whether a participant’s judgment about one issue will be inappropriately affected by his judgement about another. The graver worry that emerges here concerns the possibility

of manipulating outcomes by carefully choosing which issues get discussed. For instance, an organisation with a hidden anti-immigration agenda might seek to pair immigrant rights with terrorism, while a pro-immigration organisation might seek to pair immigrant rights with globalisation of capital. And then, of course, there are worries about the order in which issues are discussed: which issue is discussed first may be consequential. Indeed, there is evidence from a deliberative poll to suggest that the order of discussion may have an effect on how much people learn about each issue and on how their attitudes change (see the tables in Farrar et al. 2010).

Clearly, the issues here are vast and multifaceted. Yet for whatever reason, mini-public designers have been slow to confront them. This is surprising, given the seriousness with which political scientists in general regard the power that attaches to being able to control the policy agenda.

Conclusion

Deliberative polls and other mini-publics have been widely trumpeted as a panacea for various democratic ills, and in particular for the ‘democratic deficit’ in liberal democracies (Dryzek and Dunleavy 2009: ch. 9). Yet while deliberative polls and other mini-publics hold out much promise, their democratic credentials have been largely left untried. In this chapter we have sought to redress this lacuna. We find that the deliberative poll scores poorly on some of Dahl’s criteria and hence, to that extent at least, fails to create suitable conditions for *democratic* deliberation. In particular, there are serious concerns about effective participation (at least in online deliberative polls) and agenda control. Yet, as we pointed out in our introductory remarks, one would not necessarily expect every democratic institution to satisfy every democratic standard to the same degree. What ultimately matters is the democratic framework as a whole: Is it sufficiently balanced? Is institutional reform a realistic possibility? And so forth.

Unfortunately, such questions far exceed the scope of this chapter. But it is worth pointing out that deliberative democrats interested in the prospects for, and design of, mini-publics are acutely aware of such problems. In particular, there is a growing sense that what is required is a more panoptic view of their place within the democratic system as a whole (e.g., Smith 2009; Newton and Geissel 2011; Parkinson and Mansbridge 2012). We need to know much more about their possible contribution to the policy process and where that contribution might be best placed.

The sheer breadth of these questions (and more besides) necessarily means that our effort here has been a modest one; we have tried to identify places where there is some slippage between claims and practice, between assumptions and empirical reality, and between different assumptions. We have given our reasons and, where possible, proposed some solutions.

Surprisingly, we have found the going hard even on this modest venture. Both lack of data and incommensurability of measures (either due to lack of coordination or different underlying theoretical assumptions) have occluded our analytical lens. Naturally, then, our first recommendation is to urge designers and practitioners to collect more data and to make it more broadly available. To this end, they might seek to learn from the extensive efforts of political scientists to jointly collect data on different democratic systems (e.g., Comparative Constitutions Project, Comparative Study of Electoral Systems, etc.). For more foundational differences across mini-publics, it is important to be as clear as possible about criteria, principles, values and guiding assumptions – which, among other things, means that normative and empirical deliberative democrats need to work closely together. But if it is important to be explicit about how we understand each criterion, principle or value, we also need to be clearer about the relationships between them – and in particular about the difficult trade-offs that will sometimes need to be made (Thompson 2008). This matters not just for scholarly reasons, but for ensuring accountability. Mini-publics can seem artificial; they can also seem designed to further the interests of those who promote them. To some they will seem too liberal, to others they will seem too ‘other worldly’, which may explain why ‘uptake’ often proves so patchy (Goodin and Dryzek 2006). If such problems are to be addressed, then greater transparency will be needed.

Mini-publics have long been criticised by advocates of the status quo. They argue that deliberative alternatives are at best unnecessary and at worst distinctly poorer than current institutional arrangements. We are not part of that cabal which requires its members to elide over inconveniences such as the large chasm between foundational assumptions and reality (e.g., the U.S. congress was conceived of as a deliberative body, but most observers would agree that today it is anything but that). In fact, we are card-carrying members of that other cabal that sees distinct hope in deliberation as a way to solve some of the ills of contemporary democratic institutions. And none of the deficiencies we find in mini-publics have dimmed our

optimism, although they have made us more keenly aware of the work that needs to be done to fulfil their promise.

References

- Ansolabehere, S. and Schaffner, B. (2011) 'Re-examining the validity of different survey modes for measuring public opinion in the US: findings from a 2010 multi-mode comparison', working paper.
- Bachrach, P. and Baratz, M. (1962) 'Two faces of power', *American Political Science Review*, 56(4): 947–52.
- Berinsky, A. J. (2004) *Silent Voices: Public opinion and political participation in America*, Princeton, NJ: Princeton University Press.
- Bua, A. (2012) 'Agenda setting and democratic innovation: the case of the Sustainable Communities Act', *Politics* 32(1): 10–20.
- Co-Intelligence Institute (2003) Consensus Conference: a Danish description. Online. Available <http://www.co-intelligence.org/PConsensusConference1.html> (accessed 1 October 2013).
- Dahl, R. (1989) *Democracy and Its Critics*, New Haven, CT: Yale University Press.
- (1998) *On Democracy*, New Haven, CT: Yale University Press.
- Dryzek, J. (2009) 'The Australian citizens' parliament: a world first', *Journal of Public Deliberation* 5(1). Available <http://services.bepress.com/jpd/vol5/iss1/art9> (accessed 1 October 2013).
- Dryzek, J. and Niemeyer, S. (2008) 'Discursive representation', *American Political Science Review* 102(4): 481–93.
- Dryzek, J. and Dunleavy, P. (2009) *Theories of the Democratic State*, Basingstoke: Palgrave Macmillan.
- Elster, J. (1983) *Sour Grapes: Studies in the subversion of rationality*, Cambridge: Cambridge University Press.
- Farrar, C., Fishkin, J. S., Green, D. P., List, C., Luskin, R. C. and Levy Paluck, E. (2010) 'Disaggregating deliberation's effects: an experiment within a deliberative poll', *British Journal of Political Science* 40(2): 333–47.
- Fishkin, J. S. (1995) *The Voice of the People: Public opinion and democracy*, New Haven, CT: Yale University Press.
- Fishkin, J. S. and Luskin, R. C. (2005) 'Experimenting with a democratic ideal: deliberative polling and public opinion', *Acta Politica* 40(3): 284–98.

- Fishkin, J. S., Luskin, R. C., Panaretos, J., Siu, A., Xekalaki, E. 'Returning deliberative democracy to Athens: Deliberative polling for candidate selection', paper presented at the American Political Science Association Annual Conference, Boston, August 2008.
- Fishkin, J. S., He, B., Luskin, R. C., and Siu, A. (2010) 'Deliberative democracy in an unlikely place: deliberative polling in China', *British Journal of Political Science* 40(2): 435–48.
- Fung, A. (2003) 'Recipes for public spheres: eight institutional design choices and their consequences', *Journal of Political Philosophy* 11(3): 338–67.
- Golden, M. M. (1998) 'Interest groups in the rule-making process: Who participates? Who gets heard?', *Journal of Public Administration Research and Theory*, 8(2): 245–70.
- Goodin, R. E. (2005) 'Sequencing deliberative moments', *Acta Politica* 40(2): 182–96.
- Goodin, R. E. (2007) 'Enfranchising all affected interests, and its alternatives', *Philosophy & Public Affairs*, 35(1): 40–68.
- Goodin, R. E. and Dryzek, J. S. (2006) 'Deliberative impacts: the macro-political uptake of mini-publics', *Politics & Society* 34(2): 219–44.
- Goodin, R.E. and Niemeyer, S. (2003) 'When does deliberation begin? Internal reflection versus public discussion in deliberative democracy', *Political Studies* 51(4): 627–49.
- Groves, R. (2004) *Survey Errors and Survey Costs*, London: Wiley.
- Hansen, M.H. (2004) *Deliberative Democracy and Opinion Formation*. PhD dissertation, University of Southern Denmark.
- Hendriks, C. M. (2005) 'Consensus conferences and planning cells: lay citizen deliberations', in J. Gastil and P. Levine (eds) *The Deliberative Democracy Handbook: Strategies for effective civic engagement in the 21st century*, San Francisco: Jossey-Bass.
- Holbrook, A. L., Green, M. C. and Krosnick, J. A. (2003) 'Telephone vs. face-to-face interviewing of national probability samples with long questionnaires: comparisons of respondent satisficing and social desirability response bias', *Public Opinion Quarterly* 67(1): 79–125.
- Iyengar, S., Luskin, R. C. and Fishkin, J. S. (2003) 'Facilitating informed public opinion: evidence from face-to-face and online deliberative polls', working paper,

The Center for Deliberative Democracy, Stanford University. Online. Available <http://pcl.stanford.edu/common/docs/research/iyengar/2003/facilitating.pdf> (accessed 1 October 2013).

James, M. R. (2008) 'Descriptive representation in the British Columbia citizens' assembly', in M. Warren and H. Pearse (eds) *Designing Deliberative Democracy: The British Columbia Citizens' Assembly*, Cambridge: Cambridge University Press.

Jowell, R., Hedges, B., Lynn, P., Farrant, G. and Heath, A. (1993) 'The 1992 British election: the failure of the polls', *Public Opinion Quarterly*, 57 (2): 238–63.

Keeter, S., Courtney Kennedy, A. C., Tompson, T. and Mokrzycki, M. (2007) 'What's missing from national landline RDD surveys? The impact of the growing cell-only population', *Public Opinion Quarterly* 71(5): 772–92.

Kim, N., Siu, A. and Sood, G. 'Minority report: impact of opinion minorities on deliberation', paper presented at the Annual Meeting of the International Communication Association, Singapore, 2010.

Lang, A. (2008) 'Agenda-setting in deliberative forums: expert influence and citizen autonomy in the BC Citizens' Assembly', in M. Warren and H. Pearse (eds) *Designing Deliberative Democracy: The British Columbia Citizens' Assembly*, Cambridge: Cambridge University Press.

Luskin, R. C., Hahn, K. S., Fishkin, J. S. and Iyengar, S. (2006) 'The deliberative voter', working paper, The Center for Deliberative Democracy, Stanford University. Online. Available <http://cdd.stanford.edu/research/papers/2006/deliberative-voter.pdf> (accessed 1 October 2013).

Luskin, R. C., Sood, G., Fishkin, J. S. and Kim, N. 'Deliberation and learning: Evidence from the deliberative polls', paper presented at the American Political Science Association Annual Conference, Chicago, 2009.

Luskin, R. C., O'Flynn, I., Fishkin, J. S. and Russell, D. (2014) 'Deliberating across deep divides', *Political Studies*, 62(1).

Lynn, P. and Jowell, R. (1996) 'How might opinion polls be improved? The case for probability sampling', *Journal of the Royal Statistical Society* 159 (1): 21–8.

Mill, J. S. (1991 [1861]) 'Considerations on representative government' in J. Gray (ed.) *John Stuart Mill: On Liberty and Other Essays*, Oxford: Oxford University Press.

Miller, D. (2009) 'Democracy's domain', *Philosophy and Public Affairs* 37(3): 201–28.

- Moser, C. A. and Stuart, A. (1953) 'An experimental study of quota sampling', *Journal of the Royal Statistical Society Series A*, 116(4): 349–406.
- Newton, K. and Geissel, B. (eds) (2011) *Evaluating Democratic Innovations: Curing the democratic malaise?*, London: Routledge.
- Parkinson, J. and Mansbridge, J. (2012) *Deliberative Systems: Deliberative democracy at the large scale*, Cambridge: Cambridge University Press.
- Riker, W. H. (1982) *Liberalism Against Populism: A confrontation between the theory of democracy and the theory of social choice*, San Francisco: W. H. Freeman & Co Ltd.
- Setälä, M., Grönlund, K. and Herne, K. (2010) 'Citizen deliberation on nuclear power: a comparison of two decision-making methods', *Political Studies* 58(4): 688–714.
- Siu, A. (2009) *Look Who's Talking: Examining social influence, opinion change, and argument quality in deliberation*, Ph.D. dissertation, Stanford University.
- Smith, G. (2009) *Democratic Innovations: Designing institutions for citizen participation*, Cambridge: Cambridge University Press.
- Thompson, D. (2008) 'Deliberative democratic theory and empirical political science', *Annual Review of Political Science*, 11: 497–520.
- Warren, M. and Pearse, H. (eds) (2008) *Designing Deliberative Democracy: The British Columbia Citizens' Assembly*, Cambridge: Cambridge University Press.
- Weale, A. (1999) *Democracy*, Basingstoke: Palgrave Macmillan.
- Weiksner, M. 'Measurement error as a threat to causal inference: acquiescence bias and deliberative polling', paper presented at the Society for Political Methodology Annual Summer Conference, Ann Arbor, MI, 2008.
- Westwood, S. and Sood, G. 'Assessing quality of deliberation and its consequences: evidence from deliberative polling', paper presented at the Annual Meeting of the International Communication Association Meeting, Singapore, 2010.
- Yeager, D., Krosnick, J. A., Chang, L., Javitz, H. S., Levendusky, M. S., Simpsen, A. and Wang, R. (2011) 'Comparing the accuracy of RDD telephone surveys and internet surveys conducted with probability and nonprobability samples', 75(5): 861–71.

¹ It might be thought that the standards of assessment that need to be applied in the case of an analysis such as ours are those that derive from deliberative theory itself. Yet that would plainly bias our analysis from the start. By contrast, Dahl's five

criteria strike us as appropriately ‘neutral’ with respect to the content of different democratic models, deliberative or otherwise.

² Can institutions that use rely on a participant pool that is mostly composed of ‘stakeholders’ be classed as ‘democratic’? From our perspective, violation of the principle of political equality is too great for them to be classed as such.

³ Quota sampling has over the years attracted such opprobrium in academic circles that even studies comparing it to random sampling have all but gone out of fashion.

⁴ In some deliberative polls, the initial sample is recruited via non-probability marketing panels – for example, the sample for the ‘Power 2010’ DP was recruited by YouGov from its online panel. Samples from these online panels often fare badly compared to probability samples (see Yeager et al. 2011; cf. Ansolabehere and Schaffner 2011).

⁵ In some countries (e.g., Finland and Sweden) random samples are drawn from population registries and hence are much less susceptible to non-coverage issues.

⁶ We base our analyses on publicly available data from the following 14 polls: New Haven 2002, Greece 2004, Northern Ireland 2007, EU 2007, US GE Online 2004, NIC 1996, UK Crime 1994, Australian Referendum 1999, China Zeguo 2005, Argentina La Plata 2009, Brazil Porto Alegre 2010, Bulgaria 2008, Japan 2012, and Poznan 2009

⁷ <http://cdd.stanford.edu/polls/argentina/>

⁸ For instance, Argentinian La Plata 2009 seems like a good candidate for a poll where things did not go according to plan. Keeping in mind that statistical significance is a function of the size of the sample, and that there were 62 participants, the researchers find that 10 of the 55 attitude items have significant differences, a batting average of nearly 1 in 5. <http://cdd.stanford.edu/polls/argentina/>.

⁹ The training they receive usually lasts no more than a few hours.

¹⁰ The evidence is from two online deliberative polls, and may not generalise to face-to-face deliberative polls.

¹¹ Evidence for participation rates in Siu (2009) comes from two online deliberative polls. Cf. Setälä et al. (2010) who suggest that women may in fact speak less than men. Whereas Siu bases her argument on coded transcripts, Setälä et al. rely on self-reports which are vulnerable to social desirability biases. Indeed, the authors themselves modestly interpret their findings as ‘Women seem to play down their own activity compared to men’ (Setälä et al 2010: 710).

¹² As John Stuart Mill argued, we ‘need not suppose that when power resides in an exclusive class, that class will knowingly and deliberately sacrifice the other classes to themselves: it suffices that, in the absence of its natural defenders, the interest of the excluded is always in danger of being overlooked; and, when looked at, is seen with very different eyes from those of the persons whom it directly concerns’ (Mill 1991 [1861]: 246).

¹³ Although deliberative polls normally take place over a weekend, administrative and organisational matters take up a fair bit of time. For evaluation purposes, it would be useful to know how much time exactly.

¹⁴ This has potentially crucial implications. Given that knowledge gain is positively correlated with attitude change (Luskin et al. 2003), it is likely that the less educated change their opinions less so than the more educated.

¹⁵ The true proportion of those who read the briefing materials fully is likely to be even lower, given the prevalence of social desirability-induced inflation on these kinds of questions.

¹⁶ While the evidence comes from deliberative polls, it is likely that it generalizes to other mini-publics.

¹⁷ Given at least some information is only available to a particular person, it also means that pertinent private information available to be shared varies by the sample. We would like note that this point has deeper implications for statistical inference when members of the group ‘treat’ each other (i.e., person *A* is persuaded by information offered by person *B* in the group).

¹⁸ Of particular interest are differences in group composition that have an impact on the outcome. For instance, if the mean knowledge level of a group has a uniform impact on knowledge gain, one may want to make groups as similar as possible with regards to mean knowledge levels.

¹⁹ A sceptic might say that, in restricting the agenda in this way, political elites in BC were seeking to ensure that whatever the Assembly finally recommended, that recommendation would not be contrary to their interests (see Dahl 1998: 40).

²⁰ Efficiency comes from the fact that the information that they capture is structured and already reduced to numbers.