CrossMark

Economics & Philosophy (2022), **38**, 51–70
doi:10.1017/S0266267120000498

CAMBRIDGE UNIVERSITY PRESS

ARTICLE

Social choice problems with public reason proceduralism

Henrik D. Kugelberg

Department of Politics and International Relations, University of Oxford, Manor Road, Oxford OX1 3UQ, UK
Email: henrik.kugelberg@politics.ox.ac.uk

(Received 28 October 2019; revised 07 August 2020; accepted 22 November 2020; first published online 24 February 2021)

Abstract

Most political liberals argue that only rules, policies and institutions that are part of society's basic structure need to be justified with so-called public reasons. Laws enacted outside this set are legitimate if and when public reasons can justify the procedure that selects them. I argue that this view is susceptible to known problems from social choice theory. However, there are resources within political liberalism that could address them. If the scope of public reason is extended beyond the basic structure it could order people's preferences in a way that circumvents the identified issues.

Keywords: Arrow's theorem; collective self-legislation; John Rawls; political liberalism; single-peaked preferences

Introduction

Political liberals argue that the only justifications that can legitimize the state's policies, practices and institutions are those that do not rely on comprehensive conceptions of the good. Instead, legislators need to appeal to public reasons, a set of reasons that all sufficiently idealized citizens share. But not all laws need to be justified directly with public reasons. Most political liberals argue that the requirement only applies to laws, policies and institutions that are a part of the basic structure of society, things such as the constitution, and the legal and economic system. Even if their wider theories have significant differences, philosophers such as John Rawls (1993), Brian Barry (1995), T.M. Scanlon (2003), and Paul Weithman (2017) all share the view that laws enacted outside the basic structure in liberal societies are legitimate when public reason justifies the decision procedure with which they are made. This way of using public reasons to justify a decision procedure is contractualist: the procedure is legitimate if and only if it is one which sufficiently idealized citizens could agree to regardless of their comprehensive worldviews. From this, the argument goes, it follows that the law that the decision procedure 'generates' is also one which

© The Author(s), 2021. Published by Cambridge University Press

https://doi.org/10.1017/S0266267120000498 Published online by Cambridge University Press

citizens could have given themselves. Legitimacy, therefore, is procedural – since there are public reasons in favour of the procedure, public reason also justifies the substantive laws.

This *public reason proceduralism* has not received much critical attention.[1] In the following, I provide a structural critique of the paradigm, arguing that public reasons necessarily must be given beyond the basic structure. Legitimacy generated from a publicly justified decision procedure is implausible since public reason proceduralism fails to take seriously canonical results from social choice theory. Despite the centrality of these results for any discussion of procedural legitimacy, political liberals have thus far overlooked the implications for their views.[2]

I draw three critical conclusions. First, even if there are public reasons in favour of a chosen decision procedure, there will always be equally strong justifications for another. Hence, public reason is *inconclusive*. This inconclusiveness carries over to the 'higher' level, to the decision procedure employed for deciding what decision procedure to use. Second, this inconclusiveness is worrying since the aggregation of votes changes the laws that are generated, and different procedures will aggregate votes differently. Third, this creates different forms of arbitrariness in the system, something that entails that public reason proceduralism fails to satisfy basic desiderata for collective self-legislation, the standard of legitimacy that political liberals themselves aspire to. Taken together, this puts pressure on the link between the decision procedure and the legitimacy of the resulting decision.

Political liberals have thus failed to diagnose central issues. However, there are underexplored resources within their own framework that could mitigate them. If public reasoning is extended beyond the basic structure it could order people's preferences in a way that circumvents the problems that the social choice results demonstrate. I conclude that the laws generated under this regime could plausibly be described as given to the citizenry by the citizens themselves.

The problem identified is not unique to public reason views. Similar worries can be raised about any theory of legitimacy that relies heavily on proceduralism. And, conversely, public reasoning beyond the basic structure could resolve the problems also for views about legitimacy that are not based on public reason. The upshot of this is, somewhat paradoxically, that extending the scope of public reason can make political liberalism more appealing to the theory's critics. Resources from the theory could resolve some of democratic theory's central issues.

1. The scope of public reason

A basic premise of political liberalism is that people have mutually exclusive conceptions of the good – conceptions of what is 'valuable in human life' (Rawls 1985: 244). We should refrain from relying on these conceptions of the good when justifying coercive state power. Instead, we should use public reason. Public reason is a way of constructing arguments based on a shared understanding of what it is to be a citizen in a liberal democracy; it is an idea of 'the politically

[1]However, Quong (2004) is an important exception.

[2]Some arguments have been made about the relationship between public reason and social choice theory, see for instance List (2006), Kogelmann (2019) and Chung and Kogelmann (Forthcoming).

https://doi.org/10.1017/S0266267120000498 Published online by Cambridge University Press

reasonable addressed to citizens as citizens' (Rawls 1997: 766–769). Public officials and citizens should give public reasons, that is, reasons that all reasonable citizens are expected to endorse. There are competing views about how to delineate the set of public reasons, but all views hold that there is a consensus among citizens, sufficiently idealized, that these reasons are in some sense normatively relevant. Typically, public reasons are said to be based on the value of freedom, fairness and equality.

According to Rawls (1993: 227–230), not all political decisions need to meet the justificatory demands of public reason, only those that are a part of the *basic structure of society*, that is, constitutional essentials and matters of basic justice.[3] Constitutional essentials include fundamental principles for the set-up and structure of the political process and government, and the set of rights and liberties that have special constitutional protection. Importantly, this also includes rules for voting on or deciding questions outside the basic structure.

There are different ways of justifying this view. Crucially, one must find a morally salient feature that separates cases that are a part of the basic structure from those that are not. Rawls's justification for only demanding public reason in basic structure questions is somewhat underdeveloped (Quong 2011: 273). At times, he seems open to the idea that public reason may apply beyond it (e.g. Rawls 1993: 215). Quong (2011: 274) takes Rawls's position to mean that 'the idea of public reason must apply to constitutional essentials and matters of basic justice, but need not apply beyond this domain'.

At first glance, there is a contradiction here. If we accept, first, Rawls's view that public reasons need to justify the use of coercive political power, and, second, that 'all legislation involves the exercise of state power in some way – at least through the collection of taxes',[4] then it is puzzling that non-basic structure laws do not require public reasons in favour of them in order to be legitimate.[5] The standard solution to this puzzle is that if public reason can justify the decision procedure, the substantive decision, too, will be justified and legitimate. Scanlon (2003: 163), for instance, argues that if the political institutions 'have the right sort of justification, then this justification also supports legislation enacted through the procedures they define'. Weithman (2017: 105), in a recent restatement of Rawls's view, similarly argues that the settlement of a political dispute with a decision procedure is authoritative insofar as 'the decision procedure itself . . . is supportable by public reasons'.

Going beyond theories of public reason, something like this is the dominant view of where laws derive their legitimacy from. The idea is akin to what Stemplowska and Swift (2018: 6) call the *conventional view*, where a 'law or policy is permissibly enforceable if it was decided by a proper democratic procedure, even if the law or policy is somewhat unjust'. The core idea is the same: it is permissible to enforce a

[3]Even if there are apparent difficulties with successfully making a stable and theoretically sound separation of the two 'levels', between questions concerning the basic structure and other questions, I will within the confines of this paper assume that it is in principle possible to do so. For an influential critique, see Cohen (1997).

[4]According to Scanlon (2003: 162), Rawls is committed to both these claims.

[5]I will be using a simple definition: a law is legitimate if it is permissibly enforceable.

https://doi.org/10.1017/S0266267120000498 Published online by Cambridge University Press

law if and when it is enacted in accordance with the correct procedure. For political liberals, the correctness of a procedure is determined by whether the procedure can be justified with public reason. Hence, I will refer to the view as public reason proceduralism.

Previous attempts at rejecting public reason proceduralism have focused on the fact that the basic structure is not 'special', and that this view of legitimacy allows for 'limited forms of perfectionism' (Quong 2011: 274). There are two important problems with these arguments. If it could be shown that the basic structure *is* special, the first argument fails. And the second argument is not sufficient on its own, because determining whether perfectionism is wrong outside the basic structure is precisely what is at stake.

In the following, I provide a broader, structural, critique. I show that states adhering to public reason proceduralism do not meet political liberalism's own standard of legitimacy – regardless of whether the basic structure is special or not. This is because the framework cannot guarantee, firstly, that the decision procedure is selected through a fair procedure and that the choice of the procedure reflects the preferences of the citizenry, and, consequently and secondly, that the choice made with the procedure is determined by the preferences of the citizens.[6]

Thus, the alleged connection between proceduralism and legitimacy is weak. In making this argument, I draw on results from social choice theory. The first of these is Arrow's impossibility theorem.

2. Arrow's theorem

One way of understanding Arrow's theorem is by beginning with the Condorcet paradox. Condorcet showed that in pairwise majority voting, where the winning option should have at least as many votes as every other option, there is a (perhaps significant)[7] risk for cycles with intransitive results – even if the preferences of the individuals are transitive and complete. Imagine three individuals deciding on three options x, y and z:

Anna: $x \succ_A y \succ_A z$

Betty: $y \succ_B z \succ_B x$

Caroline: $z \succ_C x \succ_C y$

All individual preferences are transitive. However, on the aggregated level there is a Condorcet cycle. The result is intransitive: 'no matter which [option] is proposed to head the list, there is always a majority for another one. So majority rule does not let the group arrive at a decision' (Risse 2001: 709). The group prefers x to y (2/3),

[6]I outline what I take these conditions to require in section 3.

[7]The likelihood of cycles increases as the number of individuals increase. Hence, the preferences of an entire electorate, given certain assumptions, are likely to be cyclical. For discussions of this see Gehrlein (1983); List (2001); Riker (2003: 165–167). However, there are circumstances such that the likelihood of cycles decreases. I discuss this in section 5.

https://doi.org/10.1017/S0266267120000498 Published online by Cambridge University Press

y to z (2/3), but also z to x (2/3). Arrow generalized this result beyond the scope of pairwise majority voting, showing that when holding a set of reasonable conditions fixed, no aggregation rules exist that satisfy all conditions. The conditions are:

> WEAK PARETO PRINCIPLE (PP): if every individual (strictly) prefers x to y, then society should strictly prefer x to y.
>
> UNIVERSAL DOMAIN (U): all logically possible individual rankings of the options are permitted.
>
> NON-DICTATORSHIP (D): there is no individual such that if she prefers x to y, then society always prefers x to y.
>
> INDEPENDENCE OF IRRELEVANT ALTERNATIVES (I): the social ranking of x and y should depend only on the set of individual rankings of x and y.
>
> ORDERING (O): the social preference relations are transitive and complete.

The theorem states that if a finite number of individuals choose between at least three options, then there is no way of aggregating their individual preferences that satisfies all conditions. Any conceivable decision procedure will violate at least one of the conditions (Gaertner 2009: 19–36).

These results are significant. However, their implications are virtually unexplored for the consensus public reason literature.[8] This is regrettable because Arrow's theorem and other canonical results from social choice theory force us to question public reason proceduralism's understanding of legitimacy.

3. Procedural legitimacy and the constitutional level

3.1. Collective self-legislation, fairness and arbitrariness

As noted, the cornerstone of public reason proceduralism is that a decision procedure is legitimate if and only if it is justified with public reasons. This, then, confers legitimacy to the substantive decisions. To see how this works, suppose that the government decides between building a sports centre, an art gallery, and saving the money for next year. Clearly, this decision is not a part of the basic structure of society. Hence, the government uses a decision procedure that has public reasons in favour of it to decide what to do. By employing the legitimate procedure, it is determined that the art gallery should be built. For public reason proceduralists, this entails that the decision is legitimate, regardless of what reasons are given in favour of it. As Rawls (1993: 428) puts it, '[a] legitimate procedure gives rise to legitimate laws and policies made in accordance with it … [the] outcomes of a legitimate procedure are legitimate whatever they are.'

Weithman (2017) interprets Rawls as saying that laws and policies should be collectively self-legislated. That is, they should be given to the citizenry by itself.

[8]Gaus (2011) and Kogelmann (2017) discuss the implications for public reason. However, both endorse the convergence model, where a rule is justified if each member of the public has reason to endorse it – even if these are different, private, reasons.

https://doi.org/10.1017/S0266267120000498 Published online by Cambridge University Press

The idea is elusive, and it is not well-defined in Rawls's own writings. At the very least, however, we can be certain that the Rawlsian view of democracy is not minimal. Rawls (1995: 158–159) speaks of how the liberties of the moderns are not prior restrictions of 'the people's constituent will', he describes how his theory is compatible with 'popular sovereignty', the 'will of the people', and so on. I will not take a position on the deeper issue of how precisely these statements should be understood. It is possible to interpret them in several different ways – ranging from thick, Rousseauian, understandings, to thinner everyday democratic ones. For my purposes, it suffices to rely on only a thin understanding of democracy as collective self-legislation, because if public reason proceduralism fails to achieve it, then it also fails to achieve the thicker ones. Consequently, I will not give a full picture of what it would take for a law or policy to be collectively self-legislated in the thick sense. Rather, I offer two conditions such that if they are *not* met, it does not make sense to speak of *any* form of collective-self legislation.

> FAIRNESS (F): The procedure that selects law and policy treats citizens as free and equal.

For Rawlsians, a fair procedure is one that is soundly justified with public reasons. I take it that a fair procedure treats citizens as free and equal. An unfair procedure cannot lead to collectively self-legislated law and policy because it treats citizens as not equally worthy of concern. Dictatorial or aristocratic rules are paradigmatic examples of unfair procedures. Even if a dictator or an aristocracy only implemented publicly justified laws, those laws would not be collectively self-legislated. The laws were not given to the citizenry by itself. The citizens who are treated unfairly by being excluded from the decision-making are also part of the body of citizens who should be included if self-legislation is to be achieved. For Rawls, examples of fair procedures are voting procedures. Laws that are the result of a vote are 'binding on citizens by the majority principle' (Rawls 1997: 798).[9]

Collective self-legislation requires, moreover, that the chosen course of action tracks what the citizens want to do. If not, there is no clear link between the collective choice and the citizens' preferences. On a strong interpretation, this entails that the chosen law or policy needs to reflect the 'will of the people'. There is some evidence suggesting that something like this is what Rawls has in mind.[10] Statements such as these are dubious, however, and they are difficult to operationalize. For my purposes it suffices to rely on a much weaker interpretation of what it takes for a decision to track what the citizens want to do.

> POSITIVE PREFERENCE-TRACKING (P): The preferences of the citizenry over whether a law or policy should be enacted determine whether the law or policy is enacted or not.

[9]It is far from certain that lotteries can be ruled out as unfair procedures. There is an argument to be made that lottocratic decision procedures are fair in the sense that they treat citizens as free and equal. However, lotteries violate the second condition for collective self-legislation, so I will not discuss them further.

[10]See, for instance, Skorupski (2017) who ties Rawls's theory to radical democracy.

https://doi.org/10.1017/S0266267120000498 Published online by Cambridge University Press

Fair procedures are designed in order that the choices they generate track the preferences of those whose preferences are aggregated. One reason that dictatorial rules are not fair, is that the preferences of the citizenry are not the difference-makers. The preferences of the dictator are. However, even with a fair procedure in place it will sometimes be the case that the outcome does not track the preferences of the citizens in the right kind of way. These are the circumstances that I will focus on. That is, those that make it the case that (P) is not met, even when (F) is met.

The most important example of such a situation is if the choice is arbitrarily made. An arbitrary choice is incompatible with collective self-legislation, however it is defined, because the preferences of the citizens are not what makes the difference in determining what choice is made. I will focus on two states of affairs that lead to arbitrariness: (i) when there is no non-arbitrary way of choosing the decision procedure and different procedures generate different choices based on the same profile of individual preferences, and (ii) when there exist Condorcet cycles. As William Riker (2003: 165) puts this latter point:

> [A] social arrangement is indeterminate when a cycle exists. When the arrangement is indeterminate, the actual choice is arbitrarily made. The selection is not determined by the preference of the voters. Rather it is determined by the power of some chooser to dominate the choice or to manipulate the process to his or her advantage.

The two states of affairs can both lead to arbitrariness in two ways – it can arise either from chance or manipulation. If the difference-maker when a law is selected is chance or manipulation – as opposed to the preferences of the citizens – it is impossible to give a coherent explanation for in what way the chosen law is collectively self-legislated. A law that is the result of a 'fluke' is more accurately described as given to the citizenry by chance, than by the citizenry to itself. A law that is chosen because a group of citizens manipulated the outcome is better described as given by the manipulators to the citizenry. In these cases, the preferences of the citizenry will of course play a role in what course of action is taken. However, they do not *determine* it. When the choice is arbitrary, chance or manipulation is the real difference-maker.

Public reason proceduralism can meet (F) in many ways, but the framework cannot guarantee that (P) is met. As I will show next, this effectively severs the link between proceduralism and collective self-legislation.

3.2. Collectives choosing procedures for making collective choices

Let us begin by imagining that a constitutional assembly is deciding how to settle questions outside the basic structure of society. Three public reason proceduralists make up the assembly: Anna, Betty and Caroline. The first question they face is which decision procedures are legitimate. After a period of deliberation, they realize that it is possible to disregard several decision procedures, such as those that are dictatorial. Ultimately, they all agree that three procedures have sound public reasons in favour of them: Pairwise majority rule, the Borda count and

https://doi.org/10.1017/S0266267120000498 Published online by Cambridge University Press

Plurality rule.[11] Pairwise majority rule compares every option to every other option. The winner is the alternative that is at least weakly preferred to all others. In the Borda count, individuals assign a score for each option (k to the most preferred one, k−1 for the second-best option, k−2 for the third, and so on). The alternative that has the highest aggregated score wins. Plurality rule means that individuals vote for the option that they most prefer, and the option that gets the most votes wins.

All decision procedures are prima facie equally well-suited for choosing laws beyond the basic structure. It is impossible to conclusively determine that one of them uniquely meets (F) while the others violate it. This is so because the choice of decision procedure will inevitably involve trade-offs between different values. Arrow's theorem makes it easy to see this.[12] The Borda count and Plurality rule violate I (the ranking of x and y can depend on how high z is scored), and Pairwise majority rule violates O (it sometimes fails to produce a transitive outcome).

Public reason cannot conclusively point to one decision procedure for settling questions outside the basic structure of society. It is impossible to make an a priori choice between the procedures based on their compatibility with (F). As Gaus (2011: 328) puts a similar point, there exists no choice procedure that 'dominates the alternatives in the sense that it is best on all criteria for a good choice procedure: all possible procedures violate some reasonable condition'.

If public reason is inconclusive when it comes to choosing a decision procedure, how can we know which one to pick? A plausible suggestion is to, in line with the Rawlsian view of substantive decisions, use the decision procedure that we – as citizens – give ourselves. In other words, the decision-procedure should be collectively self-legislated. For this to obtain, the choice needs to, *mutatis mutandis*, satisfy (F) and (P).

To see if this approach works, let us return to the constitutional assembly. Suppose that they offer several public reasons for the different procedures. The members then weigh these differently, and end up ordering the procedures in the following way:

Anna: Pairwise majority $\succ_A$ Borda $\succ_A$ Plurality.

Betty: Borda $\succ_B$ Plurality $\succ_B$ Pairwise majority.

Caroline: Plurality $\succ_C$ Pairwise majority $\succ_C$ Borda.

Now, they need some way of deciding between the different procedures. In a pairwise majority voting, no alternative would be preferable to all others. 2/3 of

[11]In reality, this number would most likely be higher, see List (2011).

[12]Of course, there would still be disagreement over decision procedures even if Arrow's result did not hold. Indeed, it is not as if these types of discussions started with Arrow. The theorem makes clear, however, that in choosing a decision procedure we are always trading off different values. A reader who is certain that one of Arrow's conditions is dispensable might not be convinced by this argument. She might argue that the problem disappears if we just opt for her desired procedure. Nevertheless, it is then up to her to convince *everyone else* that, say, Plurality rule is the best decision procedure. Assuming that citizens will have different views about this, the argument still holds.

https://doi.org/10.1017/S0266267120000498 Published online by Cambridge University Press

the assembly prefers Pairwise majority to Borda, 2/3 prefers Borda to Plurality, but 2/3 also prefers Plurality to Pairwise majority. There is a Condorcet cycle. How should they decide between the alternatives? What decision procedure could they use for doing so? There are – at least – three methods for deciding what procedure to use, that are all equally well-suited (again, Pairwise majority, the Borda count and Plurality rule). To choose between these, they would need a third decision procedure, and so on. Hence, if the decision procedure should be collectively self-legislated, there would, under these circumstances, be an infinite regress. Even if all procedures are fair, no choice of decision procedure could satisfy (P), because that choice would be arbitrary.

Even without the existence of cyclical preferences over different procedures it will sometimes be impossible to non-arbitrarily settle on a decision procedure. To see this, note that what is required is a social choice function that selects one winner (one decision procedure) among the alternatives.[13] Different functions will often produce different winners. Depending on which decision procedure is employed for selecting the decision procedure to use beyond the basic structure, different winners can emerge (see also Gaus 2011: 329). If there is no non-arbitrary way of choosing between them, the resulting procedure is chosen by chance or manipulation. Contingent circumstances, not the preferences of the citizenry, determined what decision-procedure was used. This entails that (P) is not satisfied, and so collective self-legislation is not achieved.[14]

An obvious objection at this point is that the architects of the political systems of today managed to decide between the alternatives. One way to look at it is that non-

[13]The constitutional assembly do not need a full ranking of all options as the output of the decision procedure. They only need to settle on which procedure to use. In this sense, they need a social choice function, not a social welfare function. Similarly, legislative bodies typically do not need a full ranking of all options, only a conclusive top choice (should they decide to build the art gallery, the sports centre or save the money?). Hence, I will sometimes refer to 'winners', as shorthand for this. For my purposes, it does not matter whether this is understood as the top choice in a social welfare ranking or the social choice, because the problems of cyclical preferences and arbitrariness remain. I am grateful to an anonymous reviewer for asking me to clarify this point.

[14]An important objection is that political liberals need not limit themselves to Arrow's framework – they could take the intensity of preferences into account (e.g. Sen 1982). If so, the impossibility results can be avoided. I have two things to say in response. First, political liberals are typically not committed to such frameworks. Rawls, for instance, argues that a 'democratic constitution must include procedures of majority or other plurality-voting to reach decisions' (Rawls 1995: 148). I grant, however, that public reason is *compatible* with richer informational frameworks. However, and secondly, if the constitutional assembly is to make the choice to take the intensity of preferences over different decision procedures into account (in order to make a decision about which decision procedure to use), they need to make a collective choice. Suppose that Anna believes that a richer informational set-up is a good way of escaping Arrovian impossibilities. Betty and Caroline disagree – they hold that equality demands that all votes are counted equally – even if some voters have stronger preferences. Anna cannot convince the others that they should rely on the richer framework. If so, the only way to decide to have the richer set-up is through voting. How should they aggregate their votes in *this* contest? Should they take the intensity of preferences into account or not? What decision procedure should they use? The only way of settling these issues is, also, through voting, and so, the infinite regress re-emerges. Nevertheless, if these issues could be settled, that is, if it could be determined that a richer framework should be used to resolve questions outside of the basic structure, this could provide political liberals with a different escape-route than the one I will be proposing in section 5. This is an important avenue for further research. I am grateful to an anonymous reviewer for asking me to clarify this point. See also footnote 12.

https://doi.org/10.1017/S0266267120000498 Published online by Cambridge University Press

ideal contingencies determined how any given democratic society has chosen to handle political matters outside the basic structure of society. Gaus argues that evolutionary selection mechanisms, where '[a]ctual persons in society, through making choices, can end up coordinating on one moral rule through bandwagon effects, increasing returns, and path-dependency', offer plausible ways of settling the issue (Kogelmann 2017: 214). Because, according to Gaus (2011: 394), the evolutionary selection makes the decision procedure justified (and, I take it, legitimate), since 'consulting simply his or her own evaluative standards, each has decisive reason to freely endorse whichever moral requirement they have coordinated on'. If the rule ends up in this equilibrium, where everyone has reason to endorse it, everyone's independent choices delivers a rule that is uniquely justified (Gaus 2011: 402).

There are several problems with Gaus's solution, at least when viewed through the lens of political liberalism.[15] It is true that the evolutionary process could ultimately leave us with a single set of rules that all citizens have independent reasons to follow, but it does not say *anything* about the 'process of the mechanism itself' – the mechanism does not ensure that the equilibrium will have any moral status (D'Agostino 2013: 145–153; Kogelmann 2017: 214). This need not be a problem for Gaus's own theory, since he is careful not to moralize the idea, but when applied to the contractualist consensus public reason view discussed in this paper, an additional argument about why the procedure is justified is needed. This is so because the evolutionary process is not compatible with the perspective of seeing laws and decision procedures as collectively self-legislated. It does not satisfy (F), or, at least, we would need to show that it does, and the choice of the procedure fails to satisfy (P), because contingencies made it the case that it was this particular option that we coordinated on. Political liberals cannot be satisfied with seeing a rule as justified merely in virtue of it being the rule that we have. Because even if we look at the rule *ex post* and determine that it has public reasons in favour of it, neither (F) nor (P) was satisfied when the choice of decision-procedure was made.

The choice could also have been arrived at by other means. We could suppose that the constitutional assembly, after weighing up pros and cons, unanimously agreed that one decision procedure is superior for the conditions of their society. In other words, given that the members of the constitutional assembly were elected in the right way, (F) and (P) are satisfied. Then, it could be argued, the rule that we end up with is collectively self-legislated, and there is a clear explanation for why the laws and policies it selects, too, are justified and legitimate.

However, this only partially solves the problem. While the procedure is collectively self-legislated by the citizens at T_1, we need some further explanation for why it remains so for the citizens at T_2, perhaps centuries later. Even if there were a consensus among the constitutional assembly that the chosen procedure was the preferred option for settling political questions, this alignment of preferences is not 'carried over' to the preference profiles of the citizens using the decision procedure at T_2.

[15]From conservative, communitarian and at least some (non-public-reason) liberal points of view, this is not as obvious.

https://doi.org/10.1017/S0266267120000498 Published online by Cambridge University Press

To mitigate this concern, we could imagine a system in which we continuously decide whether or not to revise the constitution at regular intervals. Citizens could regularly vote on changing the procedure for another procedure.[16] This, however, re-creates the very same problems that the original constitutional assembly was faced with. If the preferences of the citizenry of T_2 are not unanimous, they would need to decide what decision procedure to use for aggregating their votes if there are several competing alternatives. Which procedure should they use to do so? There would be several, equally justified, alternatives – and we would need an additional decision procedure to choose between them. The infinite regress re-emerges.

The difficulties in establishing that our choice of decision-procedure is collectively self-legislated leads to further, substantial, problems. Whatever procedure is chosen will often produce different laws than the equally justified alternatives. This would not have been an issue if it were the case that we could be certain that the procedure was collectively self-legislated. If so, there is a strong case to be made that the chosen laws are not arbitrarily chosen, because they are the result of a non-arbitrarily chosen procedure. But when there are ambiguities at the constitutional level, the problem becomes pressing also at the substantive level. I discuss this next.

4. Collective self-legislation at the substantive level

Let us return to the government deciding on starting a building project or saving the money for next year. Suppose that the government consists of seven members, and that their preferences are distributed in the following way:[17]

Rank	Donna	Emma	Frank	George	Hanna	Irene	John
#1	Art Gallery	Sports Centre	Save Money	Art Gallery	Sports Centre	Save Money	Sports Centre
#2	Save Money	Save Money	Art Gallery	Save Money	Save Money	Art Gallery	Save Money
#3	Sports Centre	Art Gallery	Sports Centre	Sports Centre	Art Gallery	Sports Centre	Art Gallery

This illustrates a general point: different procedures – all of which are fair – can produce different outcomes. Pairwise majority rule leads to the building of the art gallery, the Borda count leads to saving money, and plurality rule picks Sports Centre. In Riker's (1982: 22) words, 'different methods of election and committee decision-making have produced different results from identical distributions of preferences'.

This is not a problem in and of itself. But since public reason proceduralism cannot conclusively establish that the choice of the procedure is collectively

[16]Helpful comments from an anonymous reviewer helped me formulate this suggestion.

[17]I am very grateful to an anonymous reviewer for help with constructing and thinking about this case.

https://doi.org/10.1017/S0266267120000498 Published online by Cambridge University Press

self-legislated, dissenters always have grounds for complaint. If the decision procedure used is not collectively self-legislated, the outcomes it generates are arbitrary – if we just as well could have used another decision-procedure, the laws and policies we chose were arbitrarily chosen. (P) is not satisfied, because the contingent fact of what decision procedure we used determined what the chosen law or policy was, not the preferences of the citizenry. The same distribution of preferences would have generated a completely different choice if the contingent circumstances would have been different.

It might be objected that if it could be shown that the system is not intentionally rigged to favour specific policies, the legitimacy of the outcome is intact. That is, that the decision procedure randomly favours some options over others is not a legitimacy problem. To see the force of the objection, consider an analogy.

A sprinting competition is to be held to determine who the best runner is. We are deciding between different distances and there are good reasons in favour of all of them: the longer test endurance, the shorter acceleration and top speed. Ultimately, it is decided that a 100-metre race should be held. We run the race, and a winner emerges. At this point, a long-distance runner complains that if we had held a longer race, a different winner would have emerged (here, he is thinking of himself of course). The organizers of the race could respond that this does not change the fact that the winner of the 100-metre race is the legitimate champion. If there are good reasons for the chosen procedure, the procedure determines who the best runner is – even if there are equally good reasons for running a longer race. At least as long as the race was not deliberately set up to favour one runner over another.

In response, the analogy does not carry all the way. True, we have satisfied (F) if a good balance of reasons supports the choice of the distance. If all runners agreed that the short-distance race was the best way to test our sprinting abilities, things would be uncontroversial. If there is unanimous support for the short distance and the runners take a vote to determine this, then (F) is also satisfied for the choice of the distance. Consequently, (P) is satisfied since the choice is a reflection of the preferences of the choosers. A winner emerges, and that winner is not arbitrarily chosen. She is the legitimate champion and should be celebrated as such.

Suppose, however, that a majority of the sprinters would have preferred a long-distance race, or, at least, that we do not know whether this is the case. Perhaps the decision to run a sprint was made by a committee centuries ago, and there are no clear routes for the runners to change the distance. Or, perhaps, the short-distance runners decided that they should settle which distance to run by holding a short-distance race. The winner of *that* race gets to decide the distance of the race that determines who the legitimate champion is. If so, things look very different. Now, the complaint is weightier because we cannot establish that (P) is satisfied.

The critic can press: why should we care about this? The decision procedures we employ for everyday legislation manage to find winning alternatives, even if it is difficult to say that they are collectively self-legislated. Similarly, the race does produce a winner. Perhaps this is just a basic fact of democratic decision making. Maybe all we need is stability, and as long as a political regime provides us with this (and somewhat just laws), we should be satisfied.

https://doi.org/10.1017/S0266267120000498 Published online by Cambridge University Press

I concede that some readers may be convinced by this objection. But the answer is not available to proponents of Rawlsian political liberalism. Paradigmatically, they are concerned with a thicker conception of stability – stability for the right reasons. And they care about collective self-legislation, in the non-trivial sense specified above. So, the chosen laws need to be the result of the preferences of the citizens.

There is a further disanalogy between the race and decision procedures. Democratic procedures will sometimes fail to reliably and non-arbitrarily select a winner. In legislative questions, we often do not need a full ranking of the options. What we do need, however, is for a winning alternative to emerge. Otherwise it is impossible to know what to do. In a race, typically, it is easy to see who comes first. But this is not the case for all democratic decisions.

Imagine that we are deciding where to go for dinner by using pairwise majority rule. If our aggregated preferences show that 2/3 of us prefer Chinese to Thai, 2/3 prefer Thai to Swedish, and 2/3 prefer Swedish to Chinese, then regardless of where we eventually end up going, a majority of us would rather have gone somewhere else. If we, for instance, go to a Chinese restaurant 2/3 of us would rather have Swedish food.

This kind of intransitivity is significant for reasons that should worry any democrat regardless of their commitment to public reason. Without it, it is impossible to ensure 'the independence of the final choice from the path to it' (Arrow 1963: 120). In a pairwise majority vote, the agent in control of setting the agenda will have disproportionate power to change the outcome of a vote. The outcome can depend on which two alternatives that are voted on first, or of the introduction of a new alternative. Thus, a difference between the vote and the race is that there potentially is space for agenda manipulation. This would effectively be like setting up the race to favour a specific winner.

Suppose that the chosen decision procedure for settling on which restaurant to go to is to first vote on two options and then eliminate the alternative with fewer votes. After the first vote, we vote on the two remaining alternatives, and the winner is the option that gets the most votes of these two. If we vote on Chinese-Swedish first, the winner is Thai, if we begin with Thai-Swedish the winner is Chinese and if the first vote is Thai-Chinese the winner is Swedish. If I – a lover of Swedish food – is in control of setting the sequence, I can make sure that this is what we will eat for dinner. If there is agenda manipulation, the manipulator chooses the winner for the citizenry. A clear violation of (P).

Manipulability cannot be avoided by instead choosing a procedure that violates (I). As McLean (2003: 16) puts it, if you take out (I) 'you have gross manipulability'. Similar arguments have been made by Vickrey (1960: 518), Patty and Penn (2014: 48), and Kogelmann (2019: Lemma 1). Thus, to push the sprinting analogy a bit further, violating I is allowing that the race could be rigged. And a rigged race does not produce a legitimate champion.[18]

And so, there are three major problems with the implemented decision procedure, whatever it is. First, we cannot ensure that the procedure is collectively

[18]Detailed and insightful comments from three anonymous reviewers helped me clarify the disanalogies between the sprint and public reason proceduralism. I also thank David Miller for putting the objection to me in the first place.

https://doi.org/10.1017/S0266267120000498 Published online by Cambridge University Press

self-legislated – we neither know that it is itself chosen by a fair procedure (F) nor that it is non-arbitrarily chosen (P). Second, when we use this procedure over another one, the outcome is arbitrary because other procedures would have given us different results, and there is no non-arbitrary justification for why the chosen procedure is the one we should use (¬P). Third, because of this, and because of the possibility of manipulation and cycles, we cannot ensure that the resulting laws are collectively self-legislated. Thus, public reason proceduralism fails on its own terms.

Next, I will argue that the legitimacy of laws beyond the basic structure can be rescued. To do so, public reasoning needs to be extended to matters of everyday legislation.

5. Admissibility thresholds and single-peaked preferences

The main claim of the proposed alternative to public reason proceduralism is simple: we should only select between laws with public reasons in favour of them. The first reason in support of this is negative – since there are no good reasons from the point of view of political liberalism to be satisfied with the standard view of legitimacy, public reason proceduralists will need a different account for showing that decisions outside the basic structure are publicly justified and legitimate (see also Quong 2011: 273–289). If they are unable to provide such an account, they should demand first-order public reasons beyond the basic structure.

The most apparent upshot of demanding that public reasons be given beyond the basic structure is that the laws that bind us would always have public reasons in favour of them. This is in line with some Rawlsian commitments, and it is surprising that Rawls did not hold this view himself. According to Neal (2009: 162), the Rawlsian fundamental concern is not that we arrive at a specific answer, but 'rather that *no illegitimate answers be allowed into the contest*'. This, of course, only concerns questions related to the basic structure. However, we have reason to want to extend this threshold also beyond it. A decision procedure requires public reasons in favour of it in order to be admitted into the set of permissibly enforceable procedures, and a law or a policy needs to have public reasons in favour of it to be admitted into the set of permissibly enforceable outcomes. Citizens have great claims, stakes or interests also in non-basic-structure questions. People care deeply about many political decisions, and these political decisions could have a significant impact on their life plans (Kugelberg Forthcoming). Ensuring that the resulting law at least has public reasons in favour of it would give those on the different sides of an issue a 'common currency' of disagreement (Quong 2011: 212).

However, there is also a different – stronger – case to be made. Public reasoning beyond the basic structure could order people's preferences in a way that entails that (F) and (P) are satisfied, making it possible to fulfil the demands of collective self-legislation.

https://doi.org/10.1017/S0266267120000498 Published online by Cambridge University Press

Recall that one of Arrow's conditions is Universal domain, the idea that all possible rankings of the alternatives are allowed. For deliberative democrats, this condition is not necessarily desirable. Deliberation could systematically shift people's preferences so that not all possible individual orderings are 'allowed'.[19] Specifically, people's preferences can become *single-peaked* after deliberation. When preferences are single-peaked, there cannot be any Condorcet cycles, and pairwise majority voting will always produce transitive and complete social preferences. The rest of Arrow's conditions are satisfied (Black 1948: 23–34; Dryzek and List 2003), and for the purposes of this paper, it also implies that it will be possible to determine what the winning alternative is, unless there is a tie. The winners generated will not be arbitrary, so if the procedure is fair both conditions for collective self-legislation are satisfied.

If the preferences of a group of individuals are single-peaked, they all share the same *scale* for ordering alternatives. Voters may disagree about what should be done, but they all agree on what the relevant evaluative dimension is for a given political question. An example of an evaluative dimension is the economic left-right scale. Preferences are single-peaked if there is agreement on three things: 'how the political spectrum is arranged', the ordering of the alternatives on this spectrum, and that every individual's ranking of the alternatives will be in accordance with the shared scale (Riker 2003: 171).

A simple example can illustrate this. Suppose that we are deciding between three border policies, Open, Semi-open and Closed. We agree that the best scale for evaluating a policy is based on how many immigrants that would be admitted when it is implemented. We also agree that Open would give the most immigration, Semi-open would be in the middle and Closed would be the most restrictive policy. Finally, we all use this scale when ranking the proposals ourselves. The idea is not that everyone will prefer one of the alternatives to the others. Instead, it merely excludes the rank orders Open $\succ_i$ Closed $\succ_i$ Semi-open and Closed $\succ_i$ Open $\succ_i$ Semi-open. Some people will prefer a restrictive policy, others one that allows for many immigrants, a third group the policy situated between the two extremes. No one would prefer the most restrictive policy to the least restrictive one while simultaneously preferring both these policies to the second-most restrictive policy. Single-peakedness only demands consistency. Voters can disagree, but their disagreement should be logical (Miller 1992: 64).[20]

It has been hypothesized that deliberation is one way of achieving single-peakedness. Miller provides two reasons. First, deliberation will likely reduce the number of people who are making faulty inferences and empirical mistakes. Second, and more importantly, when a political issue is discussed, it is possible to agree on a shared evaluative standard and divide decisions into different parts when several competing values are at stake: 'discussion amalgamates

[19]They would still be legally permissible, of course.

[20]There could be a certain logic to the preferences of someone who either wants no immigration or, for fairness reasons, completely open borders. For her, the worst option is semi-closed. However, this person does not use the shared scale for ordering the proposals, she applies another metric. Hence, it is not logical in relation to the shared scale.

https://doi.org/10.1017/S0266267120000498 Published online by Cambridge University Press

separate dimensions of choice to which different voters attach different weights'. The idea is that there should be a 'general willingness to break the decision down along its several dimensions, on each of which we should expect to find a winning position'. Then, the winning outcome will be a fair representation of the will of the majority since it 'follows the majority's judgement on each dimension of policy choice' (Miller 1992: 65). Miller's argument makes intuitive sense. And it has been empirically shown that small-scale deliberation moves the deliberating individuals towards single-peaked preferences (List *et al.* 2013). It is still an open question whether this result can be generalizable to a large-scale setting.

The public reason proceduralist might suggest that, to account for this, some version of deliberative democracy could be their proposed decision procedure beyond the basic structure. If so, several of the outlined problems could be avoided. While this would be preferable to resorting to a simple vote, the strategy is less convincing than deliberation with shared public reasons. A normative imperative to deliberate using public reasons beyond the basic structure should provide more fertile ground for inducing single-peakedness than unconstrained deliberation, because public reasons are, in some sense, shared or acceptable to all. When we order proposals after a public reason metric, a reasonable hypothesis is that people will accept the ranking as their preferred one. Often it will involve a trade-off between two political values, with freedom and equality as a paradigmatic example. If freedom-equality served as the shared scale for ordering proposals, these would be ranked from proposals that are most freedom-preserving and least egalitarian on one end, proposals with a mix between the values in the middle, and the most egalitarian proposals on the other end.

While empirical work remains to be done in order to determine whether deliberation using public reasons is more likely than non-public reasons to induce single-peaked preferences, the idea does have intuitive appeal. It is difficult to see how people could construct a shared scale around unshared or inaccessible reasons. There is a discrepancy between the shared scale and relying on reasons that many people do not see the normative relevance of.

Of course, we could picture that kind of case as entailing that those who do not agree on the relevance of the reason are on the same extreme end of the scale. When a proposal is presented and understood only in non-public terms, people with a wide range of moral commitments could disagree about it, to different extents and for different reasons. However, if their preferences are all grouped on the same extreme end of the spectrum, it does not capture the complexity of the decision at hand. In one sense, all of those who do not share the comprehensive worldview could agree – the proposal is wrong because it is based on faulty reasoning. At the same time, they disagree for different reasons and will prefer a whole range of different proposals in place of the non-public one. My hypothesis is that the absence of public reasons makes it more likely that we need several competing dimensions for understanding a political issue. It would then be more difficult – perhaps impossible – to achieve single-peakedness.

If it is true that public reasoning beyond the basic structure is likely to induce single-peaked preferences, that, in conjunction with a publicly justified decision procedure, will ensure that (F) and (P) holds. First, because there will not be any cycles, so the arbitrariness stemming from them is eliminated. Second, because it

https://doi.org/10.1017/S0266267120000498 Published online by Cambridge University Press

could also eliminate the arbitrariness arising from having different procedures generating different social choices from identical profiles of individual preferences. As Kogelmann points out, it has recently been shown that 'for all constant scoring rules, if preferences satisfy the *single-peaked* domain restriction, then, as population size increases, the Condorcet efficiency – the conditional probability that a rule will select the Condorcet winner, given that a Condorcet winner exists – of all such rules approaches one' (Kogelmann 2017: 225; see also Lepelly *et al.* 2000: table 1). Drawing on results such as this, Kogelmann (2017: 225) argues that given (assumptions like) single-peaked preferences and a large enough electorate, it will often be the case that 'all major social choice mechanisms end up selecting the same moral rule'. Thus, if it is true that public reasons beyond the basic structure could provide a fertile basis for inducing single-peaked preferences, (P) is likely to be satisfied. It then is not as worrying if the procedure is not collectively self-legislated, because its main competitors would generate the same winners.

Public reason proceduralists could not appeal to Kogelmann's argument unless they provide a more plausible story for why, without public reasons at the substantive level, the profiles of individual citizens would be likely to be single-peaked.

Moreover, even if we cannot fully achieve single-peaked preferences, deliberation with public reasons could produce more homogeneity and mutual coherence. There would be more agreement on the relevant rankings of options. As Gehrlein and Lepelly (2011: 190; see also Kogelmann 2017: 225–8) argue, this is promising because when 'voters have preferences on [options] that reflect increased levels of social homogeneity or group mutual coherence, voting rules should tend to show an increased level of Condorcet Efficiency'. Hence, if so, deliberation with public reasons is more likely than unrestricted public reason proceduralism to avoid arbitrariness.

Recent game-theoretical work also supports the conclusion. Chung and Duggan (2020) argue that deliberative democratic procedures need to have a certain form to confer legitimacy to decisions.[21] So-called 'competitive debate' outperforms 'myopic' and 'constructive' discussions in generating legitimate decisions. This entails that there are further, institutional questions to be resolved about precisely how the deliberations should take place. Crucially, however, the view presented here is compatible with competitive debate, the view that Chung and Duggan's (2020: 27) model shows to be successful. Competitive debate requires that 'positions should be supported by reasons that all endorse'. Hence, shared reasons play an important part in Chung and Duggan's model. If their view holds, if there are institutional ways of avoiding path-dependencies and arbitrariness by deliberating using shared reasons, this lends support to the idea of extending public reasoning beyond the basic structure.

6. Single-peakedness in the basic structure

A powerful objection to the above argument is to deploy precisely the same reasoning at the constitutional level. That is, if it could be established that deliberation based on public reasons could induce single-peaked preferences over

[21]I am grateful to an anonymous reviewer for raising this point.

https://doi.org/10.1017/S0266267120000498 Published online by Cambridge University Press

the choice of decision procedures, the resulting procedure would satisfy (F) and (P). That choice would then be collectively self-legislated, thereby transferring legitimacy to the outcomes it produces. In support of this argument, the critic could suggest that it is likely to hold because legislative questions typically are multi-dimensional and difficult to boil down to one dimension. Constitutional questions, on the other hand, may be less complex, allowing a structured ordering of the different positions.[22]

I readily grant that this is true for many constitutional questions. Different proposals in discussions about, for instance, the limits of free speech, the protection of religious freedom, and minority group rights could often be ordered after a shared public reason metric. However, what needs to be established is that different *decision procedures* could be more easily ordered on a shared scale than different legislative options. What needs to happen is that citizens must agree that, for instance, the Borda count is at the far-left end of the spectrum, that pairwise majority rule is in the middle, and that plurality rule is on the far-right end. Single-peaked preferences, then, demands that citizens not be 'allowed' to have the rankings Borda $\succ_i$ Plurality $\succ_i$ Pairwise majority and Pairwise majority $\succ_i$ Plurality $\succ_i$ Borda. If this happens, it is – by the same logic as the argument above – possible to non-arbitrarily choose between the procedures. (F) and (P) is satisfied.

Nevertheless, it is difficult to see on what basis such a ranking could be made. The public reason framework offers no resources for constructing the shared scale. All proposals respect equality, the inputs of every citizen are treated equally. They also respect freedom to the same degree, as citizens are free to use their votes as they see fit under all voting rules. And, as discussed, they are all examples of fair procedures. Thus, it does not make sense to say that the Borda count trades freedom for equality whereas the opposite is true for Plurality rule. We have no resources for saying that the Borda count, rather than plurality or pairwise majority rule, should be at the far-end of the spectrum.

A possible response is that the different ways the procedures violate Arrow's conditions could serve as a way of arranging the alternatives. However, this, too, would be difficult. Many decision procedures violate the same condition. How do we know in which order to put them? And, further, what should we think when several different conditions are violated? How do we order (I), (O) and (PP)? Relying on an ordering based on the Arrovian conditions is just a different way of recreating the same problem – there is no intuitive way of arranging different decision procedures along this shared scale either. Hence, it is very difficult to achieve single-peaked preferences over different decision procedures. All different rankings, Plurality $\succ_i$ Pairwise majority $\succ_i$ Borda; Borda $\succ_i$ Plurality $\succ_i$ Pairwise majority; ... and so on, are equally logical unless we can find some way of agreeing on a scale in which they are not.

At the level of substantive law, however, public reason can often function as a way of ordering alternatives. The previously mentioned left-right scale, a good example of a trade-off between different public reasons, is widely used by citizens when they make up their minds about what to vote for. Hence, even if

[22]Thoughtful, nuanced, and generous comments from an anonymous reviewer have aided me in developing this challenge.

https://doi.org/10.1017/S0266267120000498 Published online by Cambridge University Press

it is true that there are constitutional questions where single-peaked preferences are common, the ranking of decision procedures is not such a case.

7. Conclusion

The difficulties with legitimacy stemming from public reason proceduralism imply that it would be desirable if all options on the agenda have public reasons in favour of them – including in contests outside the basic structure. If so, even with a decision procedure that is not collectively self-legislated, the winning outcome will – consequently – have public reasons in favour of it. This does not imply that there is a consensus on the winning option, but it does mean that those who were against the winning proposal at least can appreciate the normative relevance of the reasons in favour of it. Of course, this is something that only looks plausible to someone who is already convinced that public reason generally is valuable.

The discussion above, however, also points to a broader conclusion. If the common currency of disagreement generates single-peaked preferences it could potentially play a previously unexplored part in dissolving several of the problems altogether. This implies that public reason's critics have good reason to revisit their scepticism. Promoting the use of public reasoning in public discourse could potentially provide an attractive escape-route from Arrow's theorem and other canonical social choice theory results. This should be an important concern for any democrat.

The solution I have outlined comes with a price. Universal domain is no longer satisfied – only public reason preference rankings would be permissible. Still, this restriction would not be imposed coercively. Rather, the goal is that the norms of public reason can appear appealing to all citizens so that they voluntarily conduct their political discussions in this ordered way. If this entails that we have traded universal domain for the ability to make legitimate decisions, the latter comes at a cheap price.

Acknowledgements. I am grateful to Shai Agmon, Paul Billingham, Ian Carroll, Niklas Dahlqvist, Gerald Gaus, Gina Gustavsson, Elsa Kugelberg, Cécile Laborde, C.M. Lim, David Miller, Lars Moen, Hallvard Sandven, Zofia Stemplowska, Collis Tahzib, Anthony Taylor and Jo Wolff for very helpful questions and comments. I also thank the audiences at the Nuffield Political Theory Workshop and the Themes from the Political Philosophy of Gerald Gaus Conference, University of Oxford.

References

Arrow K. 1963. *Social Choice and Individual Values*. New York, NY: John Wiley & Sons.

Barry B. 1995. *Justice as Impartiality*. Oxford: Oxford University Press.

Black D. 1948. On the rationale of group decision-making. *Journal of Political Economy* **56**, 23–34.

Chung H. and J. Duggan 2020. A formal theory of democratic deliberation. *American Political Science Review* **114**, 14–35.

Chung H. and B. Kogelmann Forthcoming. Diversity and rights: a social choice-theoretic analysis of the possibility of public reason. *Synthese*.

Cohen G.A. 1997. Where the action is: on the site of distributive justice. *Philosophy and Public Affairs* **26**, 3–30.

D'Agostino F. 2013. The orders of public reason. *Analytic Philosophy* **54**, 129–155.

https://doi.org/10.1017/S0266267120000498 Published online by Cambridge University Press

Dryzek J. and C. List 2003. Social choice theory and deliberative democracy: a reconciliation. *British Journal of Political Science* **33**, 1–28.
Gaus G. 2011. *The Order of Public Reason: A Theory of Freedom and Morality in a Diverse and Bounded World*. Cambridge: Cambridge University Press.
Gaertner W. 2009. *A Primer in Social Choice Theory (Revised Edition)*. Oxford: Oxford University Press.
Gehrlein W.V. 1983. Condorcet's paradox. *Theory and Decision* **15**, 161–197.
Gehrlein W.V. and D. Lepelly 2011. *Voting Paradoxes and Group Coherence. The Condorcet Efficiency of Voting Rules*. Heidelberg: Springer.
Kogelmann B. 2017. Aggregating out of indeterminacy: social choice theory to the rescue. *Politics Philosophy & Economics* **16**, 210–232.
Kogelmann B. 2019. Public reason's chaos theorem. *Episteme* **16**, 200–219.
Kugelberg H.D. Forthcoming. Opposing laws with religious reasons. *Journal of Social Philosophy*. https://doi.org/10.1111/josp.12363.
Lepelly D., P. Pierron and F. Valognes 2000. Scoring rules, Condorcet efficiency, and social homogeneity. *Theory and Decision* **49**, 175–196.
List C. 2001. Some remarks on the probability of cycles. Appendix to List, C. and Goodin, R. Epistemic Democracy: Generalizing the Condorcet Jury Theorem. *Journal of Political Philosophy* **9**, 277–306.
List C. 2006. The discursive dilemma and public reason. *Ethics* **116**, 362–402.
List C. 2011. The logical space of democracy. *Philosophy and Public Affairs* **39**, 262–297.
List C., J. Luskin, S. Fishkin and I. McLean 2013. Deliberation, single-peakedness, and the possibility of meaningful democracy: evidence from deliberative polls. *Journal of Politics* **75**, 80–95.
McLean I. 2003. The reasonableness of independence: a conversation from Condorcet and Borda to Arrow and Saari. *Nuffield College Working Paper* 2003-W6, University of Oxford.
Miller D. 1992. Deliberative democracy and social choice. *Political Studies* **40**, 54–67.
Neal P. 2009. Is Political Liberalism Hostile to Religion? In *Reflections on Rawls: An Assessment of his Legacy*, ed. S. Young, 153–176. London: Routledge.
Patty J.W. and E.M. Penn 2014. *Social Choice and Legitimacy: The Possibility of Impossibilities*. Cambridge: Cambridge University Press.
Quong J. 2004. The scope of public reason. *Political Studies* **52**, 233–250.
Quong J. 2011. *Liberalism Without Perfection*. Oxford: Oxford University Press.
Rawls J. 1985. Justice as fairness: political not metaphysical. *Philosophy & Public Affairs* **14**, 223–251.
Rawls J. 1993. *Political Liberalism*. New York, NY: Columbia University Press.
Rawls J. 1995. Political Liberalism: Reply to Habermas. *Journal of Philosophy* **92**, 132–180.
Rawls J. 1997. The idea of public reason revisited. *The University of Chicago Law Review* **64**, 765–807.
Riker W. 1982. *Liberalism Against Populism*. Long Grove: Waveland Press.
Riker W. 2003. Social choice theory and constitutional democracy. In *Philosophy and Democracy*, ed. T. Christiano, 161–194. Oxford: Oxford University Press.
Risse M. 2001. Arrow's theorem, indeterminacy, and multiplicity reconsidered. *Ethics* **111**, 706–734.
Scanlon T.M. 2003. Rawls on justification. In *The Cambridge Companion to Rawls*, ed. S. Freeman, 139–167. Cambridge: Cambridge University Press.
Sen A. 1982. *Choice, Welfare and Measurement*. Oxford: Blackwell.
Skorupski J. 2017. Rawls, liberalism, and democracy. *Ethics* **128**, 173–198.
Stemplowska Z. and A. Swift 2018. Dethroning democratic legitimacy. In *Oxford Studies in Political Philosophy (Vol. 4)*, ed. D. Sobel, P. Vallentyne and S. Wall, 4–27. Oxford: Oxford University Press.
Vickrey W. 1960. Utility, strategy, and social decision rules. *Quarterly Journal of Economics* **74**, 507–35.
Weithman P. 2017. Autonomy and disagreement about justice in political liberalism. *Ethics* **128**, 95–122.

Henrik D. Kugelberg is a DPhil Candidate in Political Theory at the University of Oxford. URL: http://www.henrikdkugelberg.net/

Cite this article: Kugelberg HD (2022). Social choice problems with public reason proceduralism. *Economics & Philosophy* **38**, 51–70. https://doi.org/10.1017/S0266267120000498

https://doi.org/10.1017/S0266267120000498 Published online by Cambridge University Press