

How to reach legitimate decisions when the procedure is controversial

Franz Dietrich¹

forthcoming in *Social Choice and Welfare*

Abstract

Imagine a group that faces a decision problem but does not agree on which decision procedure is appropriate. In that case, can a decision be reached that respects the procedural concerns of the group? There is a sense in which legitimate decisions are possible even if people disagree on which procedure to use. I propose to decide in favour of an option which maximizes the number of persons whose judged-right procedure happens to entail this decision given the profile. This decision rule is based not only on a profile in the standard sense, but in addition on a profile of judged-right procedures. To justify this decision rule, I present a set of simple axioms leading to it as the only solution.

1 Introduction and overview

For a collective choice problem, a ‘procedural judgment’ is a judgment about how group profiles should be aggregated into collective decisions, i.e. what is the right mapping from profiles to decisions. Procedural judgments are often controversial, both within the group and among social choice theoreticians. The concept of legitimacy developed here is based on the following premise:

Procedural Autonomy (premise).² *The manner in which the profile is aggregated into a collective decision should be determined by the procedural judgments within the group.*

In short, if the group wants a monarchy it gets a monarchy, if it wants a particular form of democracy it gets this form of democracy – whatever the recommendations of social choice theory. But what if the group disagrees on the decision procedure – surely the typical case? Do we then have to use a *procedure to elect a procedure*,

¹I wish to express my thanks to various people, including in particular Luc Bovens, Christian List, Boris Rotenberg, Christoph Schmidt-Petri and Paul Thorn. I also thank the Alexander von Humboldt Foundation, the Federal Ministry of Education and Research, and the Program for the Investment in the Future (ZIP) of the German Government, for supporting this research. I have presented this paper at the Public Choice Annual Meeting 2003 (Nashville, Tennessee, USA). Address for correspondence: Center For Junior Research Fellows, University of Konstanz, 78457 Konstanz (Germany). Email: franz.dietrich@uni-konstanz.de. Web: www.uni-konstanz.de/ppm/Dietrich.

²The entire paper can be translated so as to apply to the modified premise in which "within the group" is replaced by "within the committee". So, here not the procedural judgments within the group are relevant, but those within some given committee of persons. This committee could for instance have been elected by the group, and might or might not consist of social choice theoreticians.

which would raise a new and perhaps even harder problem of choice of procedure on a higher level? Can Procedural Autonomy then be respected at all? Yes, or so will be argued.

Note that the goal of respecting Procedural Autonomy contrasts with the standard social-choice-theoretic goal of determining the ‘right’ procedure independently of procedural judgments within the group. One can try to justify Procedural Autonomy either by the pragmatic argument that a group might not agree to use a procedure that it does not favour, or by different normative arguments. But, although we will discuss some of these arguments, the premise is an assumption, not a claim.

The obvious means to respect Procedural Autonomy – the use a ("legitimate") procedure favoured by the group – is limited by the possibility of diverging procedural judgments within the group. This observation motivates us to develop a new concept of legitimate *decisions*, which does not require the existence of a legitimate *procedure*. (The term "legitimate" is used as a technical term to indicate that Procedural Autonomy has been respected.) Let the group's profile be x , and imagine that each person i judges some procedure f_i as being the (normatively) right procedure. We consider a decision y as "legitimate" if the number of persons i for whom $f_i(x) = y$ is maximal. So, a legitimate decision is one that is the outcome of as many as possible of the judged-right procedures f_i in the group. Note that the procedures f_i that yield y on x may still be of quite different natures; hence, legitimate decisions may exist even given a strong disagreement over the procedure, i.e. without a legitimate procedure.

But does this concept of legitimate decisions really follow from Procedural Autonomy? To answer this question, I introduce a set of axioms which will be defended by appeal to Procedural Autonomy (and to an additional assumption about the structure of the decision problem). These axioms can be jointly fulfilled only by our legitimate decisions. This possibility theorem is partly analogous to May's Theorem (1952) on the majority rule in binary decision problems.

One may consider as an appealing feature of legitimate decisions that no normative views regarding the procedure choice are imposed upon the group, since the group's own procedural views are taken over. Under the standard approach, it seems that only in special cases such as in some binary decision problems it is possible to reach decisions without having to impose any controversial normative views regarding the procedure choice. But do our legitimate decisions really avoid this problem? The answer is split. On the one hand, if the available information consists in a standard profile together with the judged-right procedure of each person, our legitimate decisions are indeed inescapable (given Procedural Autonomy), or so I shall argue. On the other hand, our legitimate decisions neglect any information about personal procedural judgments other than judged-right procedures, thus for instance neglecting judged-worst procedures. As soon as such additional information is taken into account, legitimate decisions would have to be redefined, and there would be more than one plausible way to do so, leading to the same kind of normative dilemmas known from the standard social-choice-theoretic approach.

In Section 2 I discuss Procedural Autonomy and possible justifications of it; and I discuss the difficulty of finding a legitimate, i.e. reasonably non-controversial procedure. In Section 3 I introduce and discuss the present concept of legitimate decisions. For instance, it is observed that legitimate decisions may be hard to reach in practice because the submission of judged-right procedures is highly manipulable. In Section

4 we turn to the axiomatic justification of the present concept of legitimate decisions on the basis of Procedural Autonomy, of the form of a set of simple necessary and sufficient conditions. The challenge will be not just to find appropriate axioms, but also to carefully justify them by appeal to Procedural Autonomy. Finally, in Section 5, I summarize and make some concluding remarks. Appendix A contains proofs.

2 The premise and a first attempt

First, in Section 2.1, I discuss Procedural Autonomy and mention some possible justifications of it. Next, in Section 2.2, I discuss the difficulty of finding a standard social-choice-theoretic procedure favoured by the group. In later sections, this fact will not lead us to discard Procedural Autonomy, but will be a motivation to take the different approach developed by this article.

2.1 Procedural Autonomy and possible justifications

Procedural Autonomy states that the way in which profiles are aggregated into decisions should be determined by the procedural judgments within that group. We shall use the technical terms "legitimate" or "legitimacy" whenever this premise has been respected. Precise technical definitions of legitimate procedures and legitimate decisions will be given later.

In Procedural Autonomy, the word "judgments" is to be given a purely normative interpretation. Judging that a certain procedure is right has to be clearly distinguished from liking the procedure because it probably generates a collective decision that serves the person's private interests. For instance, someone may judge Borda count as being the right procedure while realising that his or her preferred candidate is more likely to win under plurality voting.

By contrast to the standard social-choice-theoretic goal of finding the independently "right" aggregation procedure, we assume that people within the group have their own judgments about how to aggregate profiles, and that these procedural judgments should be respected whatever their nature, i.e. whether or not people favour procedures that are democratic, egalitarian, anonymous, authoritarian, etc. So, decisions should be a function both of a standard profile and of a "profile of procedural judgments".

One could try to justify Procedural Autonomy in different ways, to be divided into pragmatic and normative arguments. This paper is not committed to any particular justification of Procedural Autonomy (nor to its justification at all, since a premise is an assumption). However it is worthwhile mentioning some of the possible arguments.

Pragmatic argument. One may argue that, for certain types of groups and/or decision problems, it is a factual necessity to respect Procedural Autonomy, because the group would not agree to use any procedure that is incompatible with their own procedural judgments. Indeed, the body or organization implementing a decision procedure needs to have at least some support within the group, to ensure that people submit information and respect outcomes. Even if people did agree, e.g. by force, to use a procedure against their will, in view of long-term stability it might still be a factual necessity to respect people's procedural judgments.

Note, however, that it is not always a factual necessity to respect Procedural Autonomy: it may well be that, although many people do not strongly favour the given procedure, they still agree to use it, without this generating major instability. This seems to be the case in many (democratic and non-democratic) societies.

Normative arguments. The other type of arguments for Procedural Autonomy is normative arguments. A variety of normative arguments could be thought of, some of which are indicated now.

First, Procedural Autonomy might be justified by some radical form of proceduralism or "metaproceduralism" whereby decisions should be a democratic and fair response not just to the profile of interests and views in the standard sense, but also to the profile of judgments about procedures. A related justification would be to take a *pluralistic* perspective and argue that procedures should not be imposed upon a group because the procedure choice usually involves normative commitments, and normative views should not be imposed upon people.

Or, one might argue from a *populist* point of view and postulate that the "will of the people" is best respected if people's procedural judgments are taken into account. If "will" is given a comprehensive meaning, it seems that procedural judgments should be part of it, and hence they should be respected in order for decisions to reflect the true "will of the people".

Further, one might try an *epistemic* justification by postulating that the best way to reach a "correct" decision (by an independent standard of truth) is by adhering to the group's own procedural judgments. Defending such a position will not be easy inasmuch as the group members may not be experts on truth-tracking decision procedures, or may not even be able to distinguish between their private interests and the independently "correct" decision.

Many attempts to provide a normative justification – including in particular the epistemic one – may draw plausibility from the diversity of collective decision problems. Indeed, collective decision problems are often so specific in nature that, arguably, no general social-choice-theoretic advice from outside may be appropriate. For instance, one might argue that any general statement of the form "in every decision problem of choosing one candidate out of four candidates, procedure f is best (from a "metaprocedural", pluralistic, populist, epistemic, or other perspective)" is false since the procedure choice must always be based on much more detailed information than just the information of a one-out-of-four-candidates choice problem. The relevant information might include the type of group (size, age, sex, etc.), the exact type of position for which the candidates are elected, the amount to which the four candidates differ, possible negative impacts of a wrong choice, etc. And, so the argument continues, since the specifics of the groups and the decision problem are best known to the group members themselves, they are in a better position to form a well-founded procedural judgment than external persons, possibly even if the external persons are social choice theoreticians!

Arguments against accepting the premise. On the other hand, it should be pointed out that one may plausibly argue against accepting Procedural Autonomy. For instance, one may hold that people's "procedural competence" is too low, perhaps even with a tendency to non-democratic procedures. Or, one may argue that, rather than

being a question of low competence, people are in fact not even willing to develop any non-self-interested procedural judgments. And further, even if they had developed them, in practice Procedural Autonomy could perhaps not be respected simply because a person would not be willing to reveal his or her genuine procedural judgments unless the judged-right procedure happens to help the person's favourite option win.

Certainly, Procedural Autonomy can only be accepted when having, at least to some extent, a positive view of people's ability and readiness of developing non-self-interested (procedural) judgments, and of their readiness to communicate them even if against their personal interests.

2.2 First attempt: looking for a legitimate procedure

From now on we assume that the premise of Procedural Autonomy is accepted. The first, rather obvious approach to respecting Procedural Autonomy is to search for a procedure that is being favoured by the group members. As will be discussed now, the limits of this approach are that in many cases, given sufficient procedural disagreement, such a legitimate procedure may not actually exist or be unambiguously identifiable.

Following Procedural Autonomy, we call a procedure *legitimate* (for a given decision problem) if the group agrees in some reasonably strong sense on this procedure. Here and in the following, "agreement" on the procedure is taken to mean not just agreement *to use* this procedure, but agreement that this procedure *is right (to use)*. Note that one may doubt that the decision procedures used in many societies, whether democratic or not, are legitimate in the present sense: although people tend to agree *to use* the actual decision rules (indeed, they use them), the agreement on the *rightness* of these procedures often seems much smaller. Note also that a legitimate procedure need not fulfil any of the requirements or criteria developed by social choice theory; indeed, it might even be dictatorial if that is what the group wants.

The first problem is that for many groups and/or decision problems a legitimate procedure may be non-existent due to insufficient procedural agreement. Note that if the procedural judgments are not sufficiently known it may be necessary to formally run a preliminary decision process to establish the (non-)existence and nature of a legitimate procedure. (However, for a procedure to be legitimate it is not here required that such a formal "legitimation" has taken place, as long as there is agreement on that procedure.)

Beside the problem of the possible non-existence of legitimate procedures, a second problem lies in the normative ambiguity of this concept, and the vagueness of its definition. When is an agreement on a procedure strong enough to provide legitimacy? Is a majority of more than 75% required, or of more than 50%, or does it even suffice that the procedure be pairwise preferred to any other procedure by a majority? These are difficult normative questions. In this sense the present concept of a legitimate procedure is inherently ambiguous.

If no legitimate procedure can be identified, one might try the escape-route of first electing a procedure with the hope that at least the procedure of electing a procedure is legitimate, i.e. reasonably non-controversial. This attempt is problematic

for different reasons.³

3 Solution: legitimate decisions rather than procedures

I have to be very clear about what I assume, and what I claim or wish to defend. Procedural Autonomy is an assumption, but it is a claim that *if* this premise is accepted *then* there are good reasons for accepting the relevance of the concepts of legitimate decisions to be presented in this section.

Can Procedural Autonomy be respected if no legitimate procedure exists or can be unambiguously identified? Possibly yes, or so I will argue. The present section has the following structure. I begin by making some general model assumptions and introducing an axiom ("Procedural Judgments") needed for the existence of legitimate decisions (Sections 3.1 and 3.2). Then, I define the present notion of legitimate decisions (the axiomatic justification of which will be provided later); and I discuss some aspects, such as the fact that legitimate decisions are a generalisation of legitimate procedures (Sections 3.3 and 3.4). Finally, we turn to the question of how to reach legitimate decisions in practice; the latter requires an additional sincerity assumption (Sections 3.5 and 3.6).

3.1 The general setting

Throughout, we consider a group of persons labelled $i = 1, \dots, n$, with $n \geq 2$, facing an entirely arbitrary collective decision problem.

Decision space and profiles. We assume a given decision space \mathcal{Y} , consisting of options y , one of which has to be chosen. \mathcal{Y} may have two or more elements, possibly even infinitely many. \mathcal{Y} might consist of candidates, measures (building one big building, or many small buildings), or of judgments (the defendant is guilty or the defendant is innocent), etc. A profile is a vector $x = (x_1, \dots, x_n)$, where x_i is person i 's input, which belongs to a given set \mathcal{X} of possible individual inputs. \mathcal{X} may consist of single votes, or of complete preference rankings, or indeed of any other kind of informational input that social choice theory may think of. The set of logically possible profiles is:

$$\mathcal{X}^n := \{(x_1, \dots, x_n) \mid x_i \in \mathcal{X} \text{ for all persons } i = 1, \dots, n\}.$$

³When no legitimate procedure can be identified, the attempt of identifying a legitimate "meta"-procedure for the "meta"-problem of selecting a procedure is problematic for the following reasons. Firstly, even if the attempt was successful and one had found a legitimate meta-procedure, our premise would only have been met on a meta-level, since only the meta-procedure is legitimate, while the procedure might have been elected "narrowly" out of a class of many procedures that are all nearly equally controversial, and hence nearly equally illegitimate. Secondly, even if one accepts to search for legitimacy (premise fulfilment) only on a meta-level, it is unlikely that this search will be successful. The meta-problem of choosing a procedure is often more complex than the original problem, since its alternatives (which are procedures) may exceed the original alternatives by complexity and number. So it seems unlikely to reach an agreement on the meta-procedure if there was no agreement on the procedure in first place. Thirdly, of course, one may continue indefinitely, by trying to meet the premise on a meta-meta-level, or on a meta-meta-meta-level, etc. Infeasibility aside, the same objections apply analogously.

Set of allowed procedures. A procedure is a function f mapping $\mathcal{D} \mapsto \mathcal{Y}$, where $\mathcal{D} \subseteq \mathcal{X}^n$ is any (non-empty) domain of profiles. We do not assume an agreement as to what decision procedure should be used. However, we assume an agreement on some set \mathcal{F} of procedures considered as allowed or possible, all of which are at least defined on the group's actual profile x . \mathcal{F} might be the set of all procedures defined at least on x , or the set of all anonymous procedures defined at least on x , or the set of all anonymous procedures defined on the universal domain $\mathcal{D} = \mathcal{X}^n$, etc.

The unknown distance assumption. Our only restriction is that it be desirable that the decision rule⁴ should not presuppose any prior knowledge as to how close different options in \mathcal{Y} are from each other – the "unknown distance" assumption. Our formal definition of legitimate decisions works perfectly without the unknown distance assumption, but it is only under this assumption that I wish to defend that this definition follows from Procedural Autonomy. So, let us postpone the precise definition and discussion of the unknown distance assumption until we justify our legitimacy concept (see Section 4.7). For now, an example should suffice: when electing a candidate, the unknown distance assumption is justified since a fair and neutral decision rule⁴ should not a priori impose a (political) judgment as to how close different candidates are from each other.⁵

3.2 Procedural Judgments

If a person $i \in \{1, \dots, n\}$ considers one particular procedure in \mathcal{F} as normatively strictly superior to all other procedures in \mathcal{F} , then we call it person i 's "*judged-right*" procedure (in \mathcal{F}). We assume the axiom of

Procedural Judgments (J).⁶ *Each person $i \in \{1, \dots, n\}$ has a judged-right procedure in \mathcal{F} , written f_i .*

It is essential to note the normative interpretation of "right", by contrast to liking a procedure because it serves one's personal interests.⁷ Apart from this normative interpretation, we are entirely non-restrictive regarding the reasons that lead a person i to believe that f_i is the right procedure. Some persons might choose their f_i on procedural grounds (f_i is the most democratic procedure). Others might choose their f_i on epistemic grounds (f_i is the best "truth-tracker"). Even others might choose their f_i on "populist" grounds (f_i is best at deciding in accordance with the "will of the people"). Also, the persons do not need to be able to compare *all* pairs of two procedures in \mathcal{F} : it suffices that one particular procedure f_i is pairwise comparable with all other procedures $f \in \mathcal{F}$, and that f_i beats f .

⁴By a "decision rule" I do not mean a procedure f ; rather, I mean a rule by which a decision is reached based not only on the profile x , but also on the vector of procedural judgments within the group; see Section 4.

⁵The unknown distance assumption excludes an important type of decision problem, namely those problems in which the options are collective rankings over some underlying set. Indeed, two rankings can be very close or very far away from each other; see the remarks in Section 4.7.

⁶For the case of the modified premise (see the footnote "²"), each person *of the committee* is required to have a judged-right procedure – a probably more realistic assumption. The definition of legitimate decisions can be formulated analogously.

⁷See our remarks in Section 2.1 on the normative interpretation of the procedural judgments in Procedural Autonomy.

3.3 Legitimate decisions

Throughout, let x denote the actually submitted profile. Procedural Judgments (J) allows us to introduce our central concept of the legitimate decisions (which is proposed only if an "unknown distance" assumption is satisfied, as discussed later in Section 4.7). Our legitimate decision can be seen as the best compromise given the (possibly diverging) procedural judgments. The axiomatic justification of our definition by appeal to the premise of Procedural Autonomy will appear in Section 4.

Definition 1 *If Procedural Judgments (J) is satisfied, then*

- we call "degree of legitimacy" or just "legitimacy" of an option (decision) $y \in \mathcal{Y}$ the number $L(y) := L_{x,f_1,\dots,f_n}(y) := \#\{i | f_i(x) = y\}$ of persons whose judged-right procedure leads to y ;⁸

- we call an option (decision) $y \in \mathcal{Y}$ "most legitimate" or just "legitimate" if it has maximal legitimacy $L_{x,f_1,\dots,f_n}(y)$.

First, note that we do *not* define as legitimate what wins under plurality voting, because the option $f_i(x)$ cannot be interpreted as person i 's vote under any standard interpretation of voting. Indeed, the option $f_i(x)$ may differ both from person i 's private interests and from person i 's altruistic judgment about what serves best the group's interests. Person i may judge Borda count as being the right procedure f_i , and at the same time Borda count may produce a winner $f_i(x)$ that neither serves person i 's interests, nor even corresponds to person i 's judgment about what decision would best serve common interests. Indeed, if person i judges the option y as best serving the common good, but the group does not like y , then the profile x is likely to produce an outcome $f_i(x) \neq y$ (provided that f_i is a democratic procedure). Hence, in a plurality voting person i need not submit the vote $f_i(x)$, whether the motivation is private interests or the common good.

Second, Procedural Autonomy has here been interpreted in the narrow sense that respecting person i 's procedural judgments means no more than deciding in favour of $y_i = f_i(x)$, whether or not it is *because* the procedure f_i has been applied. In other words, we define the term "procedural judgment" in the narrow sense of a judgment about what is the right *mapping* or *function* between profiles and decisions, regardless of the *means* by which this correspondence is established. As the referee rightly points out, this interpretation may be problematic, since a group of democrats may for instance not approve of a benevolent dictatorship that always decides according to the will of the majority. This example shows that our narrow definition of 'procedural judgments' may fail to capture important other types of procedural judgments or concerns, and accordingly that our legitimate decisions may violate other types of procedural concerns.

Third, the function $L(y)$ ($y \in \mathcal{Y}$) is intended as a purely *ordinal* (i.e. not *cardinal*) measure of the amount of legitimacy of options, and so any positive transformation of it would do exactly the same job. For instance, an equivalent definition of $L(y)$ would be to divide our definition by the group size n , so that $L(y)$ would always belong to the interval $[0, 1]$, whatever the group size n .

⁸If one allowed that the procedures in \mathcal{F} result in ties, i.e. choose non-empty sets of options rather than single options, then one could redefine $L_{x,f_1,\dots,f_n}(y)$ as $\#\{i | y \in f_i(x)\}$, the number of individuals i such that y is among the options chosen by f_i .

Example. Consider a group of $n = 30$ persons who must elect one out of four candidates:

- 10 persons judge plurality voting as right,
- 10 persons judge the Borda count as right, and
- 10 persons judge the Hare system as right.

Given a profile (x_1, \dots, x_{30}) of complete person preference rankings over the four candidates, it turns out that

- candidate 1 wins under plurality voting,
- candidate 2 wins both under Borda count and the Hare system, and
- candidates 3 and 4 do not win under any of the three aggregation rules.

Then

- candidate 1 has legitimacy $L(1) = 10$,
- candidate 2 has legitimacy $L(2) = 10 + 10 = 20$, and
- candidates 3 and 4 both have legitimacy $L(3) = L(4) = 0$.

Hence the (only) legitimate decision is candidate 2. Candidate 2 would also be the legitimate winner if 14 of the 30 persons supported plurality voting, 8 persons supported the Borda count and 8 persons supported the Hare system. This shows that a legitimate decision may differ from a decision chosen by the following two-step rule: first, a procedure is elected by plurality voting over possible procedures (here, plurality voting would easily win), then the winning procedure is applied to reach a decision (here, the winner would be candidate 1, although not the legitimate decision). Our present approach is different since it does not choose a legitimate *procedure*, but it achieves legitimate *decisions* without using any (legitimate?) procedure.

Note that the sum-total of the legitimacies of candidates is the group size:

$$L(1) + L(2) + L(3) + L(4) = 30.$$

As can easily be seen, the latter relation holds in general, i.e.⁹

$$\text{the sum-total of the legitimacies is the group size: } \sum_{y \in \mathcal{Y}} L(y) = n.$$

Also, note that there always exists a legitimate decision, because among the integers $L(y)$, $y \in \mathcal{Y}$, there is of course a largest one. If there exists a *unique* legitimate decision y , then y is *the* legitimate decision. A case of non-uniqueness is the situation where all of the outcomes $f_1(x), \dots, f_n(x)$ are different. Then the legitimacy $L(y)$ of any $y \in \mathcal{Y}$ is at most 1, and each y with legitimacy 1 is a legitimate decision. However, when the group size n exceeds well the number of possible decisions ($n \gg |\mathcal{Y}|$), there usually is a unique legitimate decision. The two extreme cases are:

1. *Fully concentrated legitimacy:* One particular decision $y \in \mathcal{Y}$ has maximal legitimacy $L(y) = n$ and all other decisions have legitimacy 0. This happens when $y = f_1(x) = \dots = f_n(x)$. It may happen that a group is far from a consensus on the procedure, but still one decision is y is legitimate in the strongest possible sense that $L(y) = n$.

⁹In the case that the decision space \mathcal{Y} is infinite, all except finitely many terms of the sum $\sum_{y \in \mathcal{Y}} L(y)$ are zero, and hence this sum is in fact one of finitely many terms.

2. *Uniformly distributed legitimacy:* All possible decisions have the same legitimacy, and hence per definition each option is a legitimate decision. For instance, each option $y \in \mathcal{Y}$ may have legitimacy $L(y) = 1$ since each y wins under exactly one of f_1, \dots, f_n . This is only possible if the group size n equals the number of options $\#\mathcal{Y}$. If there are twice as many persons as options $y \in \mathcal{Y}$, then it may happen that each option y has legitimacy $L(y) = 2$; if there are three times more persons than options, then possibly each option y has legitimacy $L(y) = 3$, etc.

An alternative. A different definition of legitimate decisions y would have been to require not only that the legitimacy $L(y)$ be maximal, but in addition that $L(y)$ exceeds a certain proportion of the group size. For instance, one might require that $L(y) > n/2$, i.e. that an absolute majority of the persons i favour procedures f_i leading to decision y . The latter is a stronger sense of legitimate decisions, but contains some arbitrariness due to the choice of the threshold. Under this stronger definition legitimate decisions need not exist, just as legitimate procedures need not exist.¹⁰

3.4 Why legitimate decisions are a generalisation of legitimate procedures

At first sight one might be sceptical about distinguishing between legitimacy of procedures and legitimacy of decisions. Do we really need two concepts? Or would we not better have defined legitimate decisions in terms of legitimate procedures by considering a decision as "legitimate" *if and only if* it is reached through a legitimate procedure? This article's concept of legitimacy is that Procedural Autonomy has been respected. Although the axiomatic justification for our legitimate decisions is still to come, one might see from an example that there seem to be cases in which Procedural Autonomy can be maximally fulfilled by a (legitimate) decision without there being a legitimate procedure whatsoever: imagine that every person of the group judges a different procedure as right, but that all of these different procedures happen to yield the same decision.

Given that we have two concepts, it would be desirable that at least there is no conflict between them. Indeed, a natural question is whether a legitimate procedure – if there happens to be one – will actually result in a legitimate decision. We have defined a procedure as legitimate if it is supported by the group in a reasonably strong sense, without specifying what "reasonably strong" exactly means. Assuming that it

¹⁰The stronger definition of legitimate decisions comes closer to our definition of legitimate *procedures*. Indeed, for a procedure to be legitimate it is not sufficient that this procedure obtains maximal support among all procedures, but we required that this support be *reasonably strong*, i.e. in a sense be stronger than some threshold. Note that one could also adapt the definition of legitimate procedures so as to match that of legitimate decisions, by considering a procedure as legitimate already if it obtains maximal support, i.e. more support than any other procedure. But this "maximal support" criterion would have been perhaps even more ambiguous than the "reasonably strong support" criterion: if all procedures obtain little support, then it should be obvious that no procedure has "reasonably strong support", but it may be very hard (and normatively ambiguous) to determine which one has "maximal support". To bring the terminologies for decisions and for procedures into line, one might prefer to talk on the one hand of "most legitimate" decisions respectively procedures if the legitimacy of the decision is maximal respectively if the consensus on the procedure is maximal in some reasonable sense (where "maximal" does not imply "high"), and on the other hand talk of "legitimate" decisions respectively procedures only if in addition the legitimacy of the decision respectively the consensus on the procedure is "reasonably large", i.e. exceeds some threshold.

means that an absolute majority in favour of the procedure is required, or that some higher (qualified) majority is required, legitimate decisions are indeed a generalisation of legitimate procedures:

Theorem 2 (*Legitimate procedures entail legitimate decisions*) *Assume Procedural Judgments (J). If there exists a procedure f that is the judged-right procedure of more than half of the persons, then the (only) legitimate decision is the outcome $y = f(x)$ of this procedure.*

PROOF. By definition, the legitimacy $L(y)$ of the option $y = f(x)$ is the number of persons who favour procedures that leads to this decision y , for instance who favour the procedure f . So the legitimacy $L(y)$ is at least as high as the number of persons who favour the procedure f . Since by assumption an absolute majority favours f , the legitimacy of y satisfies $L(y) > n/2$. It follows that any other option $y' \neq y$ has legitimacy $L(y') < n/2$, because by the sum-total of the legitimacies of options is n , implying that $L(y') < L(y)$. **QED.**

Theorem 2 shows that legitimate decisions are a generalisation of legitimate procedures. If an outcome is supported by a legitimate procedure, then it is a legitimate decision, but not vice versa. But what kind of generalisation is it? One may think of this as follows: while legitimate procedures aim at maximizing agreement on a *single* (legitimate) procedure, legitimate decisions maximize agreement on a *set* of procedures, namely on the set of those procedures which happen to entail the same (legitimate) decision.¹¹

3.5 Can legitimate decisions be reached in practice?

We have defined legitimate decisions in terms of personally judged-right procedures f_1, \dots, f_n . But can legitimate decisions actually be reached in practice? From now on, let f_i denote a procedure *submitted* by persons i , whether or not this is his or her judged-right procedure. The obvious decision rule aimed at reaching legitimate decisions consists in

- collecting from each person $i \in \{1, \dots, n\}$ both an x_i and an f_i , and
- choosing the option $y \in \mathcal{Y}$ that maximizes the number $\#\{i | f_i(x) = y\}$, or choosing the tie¹² between all options having this property if there are many.

Call this the "counting rule", which may be defined on any domain of inputs (x, f_1, \dots, f_n) for which each of f_1, \dots, f_n is defined at least on x . The counting rule is highly manipulable. Indeed, for this rule to always generate legitimate decisions it is crucial that people submit their judged-right procedures rather than some other procedures more likely to help their preferred options win. Let us assume

¹¹More formally, calling \mathcal{F}_y the set of those procedures which entail y (under the given profile x), each of the sets of procedures \mathcal{F}_y ($y \in \mathcal{Y}$) obtains a certain amount of support in the groups in the sense of a certain number $L(y)$ of persons believing that the right procedure belongs to \mathcal{F}_y . By choosing y so as to maximize $L(y)$, one maximizes the support for \mathcal{F}_y . So, like legitimate procedures, legitimate decisions aim at maximizing support for procedures – but now on the level of sets of procedures rather than of single procedures.

¹²In the case of a tie between many options, a final decision might be reached by running a probability experiment, e.g. where each of the options has the same probability of being chosen.

Sincerity (S). *For each person $i \in \{1, \dots, n\}$, if person i has a judged-right procedure then person i submits this procedure.*

If people both have judged-right procedures *and* are honest about them – our two strong axioms (J) and (S) –, then the counting rule obviously has the desired effect:

Proposition 3 *Suppose Procedural Judgments (J) and Sincerity (S). Then the counting rule entails the legitimate decision, or entails the tie between all legitimate decisions if there are many.*

Without Sincerity (S), the procedures f_1, \dots, f_n may be just another (strategic) expression of preferences or interests *on options* (not on procedures); and hence the number $\#\{i | f_i(x) = y\}$ may have to do little with the degree of legitimacy, in which case the counting rule maximizes the wrong quantity.

Note that, while restrictive about the motivation behind the submitted f_i s, we are completely unrestrictive about the motivation behind the submitted x_i s.

3.6 A way to achieve Sincerity (S)

It need not be the case that the judged-right procedures f_1, \dots, f_n are newly collected for each decision problem faced by the group. Given a class of decision problems of a similar in type (e.g. all decisions of choosing one out of four candidates), it is plausible that a person has the same judged-right procedure for all of them. Hence, the procedures f_1, \dots, f_n may be collected once and for all.

Such a separation of the collection of f_1, \dots, f_n from the concrete decision problem might have the positive effect that persons are led to fulfil Sincerity (S): indeed, in the absence of a concrete decision problem and perhaps in the ignorance of the future decision problems, a person has little grounds for forming strategic motives.

4 A set of necessary and sufficient conditions

In the last section I have postulated a concept of legitimate decisions based on the premise of Procedural Autonomy whereby the aggregation of profiles should be determined by people's own procedural judgments. We now come to a formal justification of our definition of legitimate decisions and the counting rule, based on a quite specific and strong interpretation of Procedural Autonomy (and based on the additional "unknown distance" assumption to be discussed in Section 4.7). This justification takes the form of a possibility result: the counting rule is the *only* decision rule that satisfies a set of axioms that will be developed. Some of these axioms reflect a precisification of Procedural Autonomy, other axioms correspond to more standard social-choice-theoretic requirements such as anonymity.

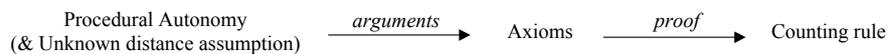


Figure 1: The way from Procedural Autonomy to the counting rule

Figure 1 illustrates the structure of our justification for the counting rule resp. our concept of legitimate decisions: arguments will lead us from Procedural Autonomy to precise axioms, and a proof will lead from the axioms to the counting rule as the only solution.

Throughout we assume Procedural Judgments (J) and Sincerity (S). So the submitted procedures f_1, \dots, f_n are the judged-right procedures of the persons. Accordingly, the number $L(y) = L_{x, f_1, \dots, f_n}(y) = \#\{i | f_i(x) = y\}$ coincides with what we have defined as the legitimacy of y .

After giving some definitions and conventions (Section 4.1), and stating our two theorems characterising the counting rule (Section 4.2), I present five axioms in the following chronological order: Domain Condition (Section 4.3), Anonymous Procedure Submission (Section 4.4), Strict Monotonicity (Section 4.5), Procedural Neutrality (Section 4.6), and Neutrality (Section 4.7). The two theorems differ in their domain assumptions: the first theorem assumes Universal Domain and is partly analogous to May's Theorem on the absolute majority rule; the second theorem is based on Domain Condition. The latter condition is more general and also applies to cases where one wants to restrict the range of decision procedures people may submit, for instance by excluding the submission of dictatorial procedures. The main challenge will be to justify these axioms by appeal to Procedural Autonomy. Here, by far the hardest work will have to be done for Procedural Neutrality (PN) and Neutrality (N). Both of these axioms express that F is neutral regarding how x should be aggregated, forcing F to rely only on people's proposals f_1, \dots, f_n . Much less controversial seem to be our anonymity and monotonicity requirements.

4.1 Definitions and conventions

We carefully distinguish between "procedures" and "decision rules": f_1, \dots, f_n are procedures since they make decisions on the basis of just the profile x , but the counting rule is a decision rule since its decisions also depend on f_1, \dots, f_n . A decision rule will be denoted by the capital letter F . To give a name to the input of a decision rule, let us call a vector $(x, f_1, \dots, f_n) = (x_1, \dots, x_n, f_1, \dots, f_n)$ a "metaprofile", which is simply the concatenation of a profile x and a vector of procedures (f_1, \dots, f_n) . The "domain" of a decision rule F is the set of metaprofiles (x, f_1, \dots, f_n) on which it is defined. We will require that decision rules be defined only on "compatible" metaprofiles in this sense.¹³

Definition 4 *A metaprofile (x, f_1, \dots, f_n) is called **compatible** if each of the procedures f_1, \dots, f_n is defined at least on the profile x .*

¹³Our restriction that decision rules be defined only on compatible metaprofiles is less a consequence of Procedural Autonomy itself than of us wanting to be able to apply it without having to ignore anybody's procedural concerns. Indeed, if the a decision rule were defined for a metaprofile (x, f_1, \dots, f_n) for which, say, f_1 is undefined on x , then it would seem unclear how to meet Procedural Autonomy. Should f_1 be ignored because the procedural judgment of person 1 "does not apply"? Perhaps. However, mainly for technical convenience, we consider only the case of compatible metaprofiles. A generalisation of our discussion to general domains of metaprofiles is possible (with some additional assumptions); the counting rule would then have to be defined on non-compatible metaprofiles by counting only among those outcomes $f_i(x)$ that are defined.

Recall that the counting rule is a decision rule that may lead to ties, namely if $L(y)$ can be maximized by more than one option y . In general, we wish to allow for such ties in decision rules (but not in the procedures f_1, \dots, f_n), and hence decision rules should be technically result in a *set of options*:^{14,15}

Definition 5¹⁶ A **decision rule** is a function F , defined on some domain of compatible metaprofiles, written $Dom(F)$, and whose values $F(x, f_1, \dots, f_n)$ are non-empty subsets of \mathcal{Y} . For any set \mathcal{M} of compatible metaprofiles, the **counting rule on \mathcal{M}** is the decision rule F defined on the domain \mathcal{M} by

$$\begin{aligned} F(x, f_1, \dots, f_n) &:= \{y \in \mathcal{Y} | y \text{ occurs maximally often among } f_1(x), \dots, f_n(x)\} \\ &= \left\{ y \in \mathcal{Y} | L(y) = \max_{y' \in \mathcal{Y}} L(y') \right\}. \end{aligned}$$

For instance, one of our theorems will impose

Universal Domain (UD). $Dom(F)$ consists of all compatible metaprofiles (x, f_1, \dots, f_n) .

In a slight abuse of language, we shall use the word "decision" for $F(x, f_1, \dots, f_n)$, although really this is a set of options, and the actual decision will be an element of this set.¹⁴

Note the important difference between the counting rule and the decision rule F consisting in first electing a procedure by plurality voting over procedures and then applying the winning procedure(s), formally defined by:

$$F(x, f_1, \dots, f_n) := \{f(x) | f \text{ occurs maximally often among } f_1, \dots, f_n\}.$$

4.2 The two possibility theorems

In anticipation of the following subsections, where our axioms are introduced, let us begin by stating our two theorems that uniquely characterise the counting rule. This provides a justification for the counting rule and of the present concept of legitimate decisions on the basis of Procedural Autonomy.

The first theorem requires that F be defined on the universal domain (of metaprofiles), which yields the analogue to May's Theorem:

¹⁴For a decision rule F , if $F(x, f_1, \dots, f_n)$ is a one-element set $\{y\}$ (the typical case), y is the decision to take. If $F(x, f_1, \dots, f_n)$ happens to be a many-element set (the case of ties), then a decision might be obtained by randomly selecting one element of that set, e.g. according to a uniform probability distribution.

¹⁵If ties were to be allowed also for the submitted procedures, i.e. if each of f_1, \dots, f_n map the profile into a non-empty subset of \mathcal{Y} , the counting rule would have to be redefined by

$$F(x, f_1, \dots, f_n) := \{y \in \mathcal{Y} | y \text{ is an element of maximally many of } f_1(x), \dots, f_n(x)\}.$$

This decision rule can be axiomatised like the counting rule without ties in f_1, \dots, f_n ; in fact, most of our axioms can be taken over literally.

¹⁶For the case of the modified premise (see the footnote "2"), the input of a decision rule should be a vector (x, f_1, \dots, f_k) , where now f_1, \dots, f_k are the procedures submitted by the k committee members. The counting rule can be defined analogously, and this section's axiomatisation of the counting rule is also possible analogously.

Theorem 6 *Let F satisfy Universal Domain (UD). Then F is the counting rule on its domain if and only if F satisfies Anonymous Procedure Submission (APS), Strict Monotonicity (SM) and Neutrality (N).*

As will be discussed in Section 4.3, Universal Domain (UD) might be undesirable. For instance one may want to exclude dictatorial procedures among f_1, \dots, f_n , which is a special case of the more general Domain Condition (DC). For the latter, an analogous theorem holds:

Theorem 7 *Let the domain of F satisfy Domain Condition (DC). Then F is the counting rule on its domain if and only if F satisfies Anonymous Procedure Submission (APS), Strict Monotonicity (SM) and Neutrality (N).*

In fact, Theorem 7 implies Theorem 6 since Universal Domain (UD) is a special case of Domain Condition (DC). In both theorems, what is trivial is that the counting rule satisfies all axioms. The converses are proven in Appendix A.

4.3 The domain of F

What metaprofiles (x, f_1, \dots, f_n) should be allowed? Our two theorems assume different answers to this question. Procedural Autonomy, if taken at its strongest, clearly requires Universal Domain (UD) (Section 4.1), since people's freedom of submitting procedures should not be a priori restricted. The metaprofile (x, f_1, \dots, f_n) might be collected in two steps: first, the profile x is collected, and then each person is asked to submit a procedure defined at least on x .

However, there might be good reasons for rejecting Universal Domain (UD), even if in contradiction with Procedural Autonomy taken strongly. For one, we may want to exclude unwanted procedures, such as dictatorial ones, or non-anonymous ones. This may be seen as an attempt to achieve Sincerity (S): if dictatorial procedures were to be allowed, the temptation would be just too great to submit the dictatorial procedure with oneself as dictator rather than one's judged-right procedure. And here is another reason for restricting the domain of F . As discussed in Section 3.6, it may be good to collect f_1, \dots, f_n before collecting x , because Sincerity (S) is more likely to be achieved if the procedure submission is disconnected with the specific decision problem. Now, if f_1, \dots, f_n are collected first and chosen freely, it may happen that the later collected x does not belong to all of the domains of the procedures f_1, \dots, f_n . To prevent this problem, one could impose that each of f_1, \dots, f_n be defined on the unrestricted domain \mathcal{X}^n . Then, obviously, the collected profile x belongs to the domain of each of f_1, \dots, f_n .

All of the restricted domains of F mentioned (and Universal Domain (UD)) are special cases of the following type of domain (where \mathcal{P} may for instance be the set of all non-dictatorial procedures, or of all procedures with unrestricted domain \mathcal{X}^n):

Domain Condition (DC). *There exists a set of procedures \mathcal{P} such that*

- (i) *Dom(F) is the set of all compatible metaprofiles (x, f_1, \dots, f_n) for which $f_1, \dots, f_n \in \mathcal{P}$;*
- (ii) *for all $y \in \mathcal{Y}$ and $x \in \mathcal{X}^n$, there is an $f \in \mathcal{P}$ defined on x and satisfying $f(x) = y$.*

Condition (ii) means that \mathcal{P} is sufficiently large to ensure that (whatever the profile $x \in \mathcal{X}^n$) procedures may be submitted that generate any given option $y \in \mathcal{Y}$. Perhaps this requirement is too strong. For instance it makes it impossible to define \mathcal{P} as the set of procedures that respect the (weak or strong) Pareto principle. So, a domain subject to (DC) might have to allow people to submit quite "bad" procedures. Such problems can be avoided by replacing (DC) by an even weaker domain constraint; but this will not be discussed here.

The reader will have recognised that the set of procedures \mathcal{F} used in Section 3 (the set of allowed procedures for the profile x) corresponds to the set of those procedures in \mathcal{P} defined at least on x .

4.4 Anonymous procedure submission

We need the largely non-controversial requirement that outcomes should not depend on who has submitted what procedure. This amounts to respecting the procedural judgments of every person equally, rather than considering some persons as having more "procedural competence" than others, or as having a higher right to impose their procedural views than others.

Anonymous Procedure Submission (APS). *For any two metaprofiles in $\text{Dom}(F)$ differing only in the order of the procedures, (x, f_1, \dots, f_n) and $(x, f_{\pi(1)}, \dots, f_{\pi(n)})$, where $\pi : \{1, \dots, n\} \mapsto \{1, \dots, n\}$ is any permutation,*

$$F(x, f_1, \dots, f_n) = F(x, f_{\pi(1)}, \dots, f_{\pi(n)}).$$

Note that, although we require anonymous procedure submission, we do not require anonymity with regards to the profile $x = (x_1, \dots, x_n)$. For instance, if all of the submitted f_1, \dots, f_n are identical to the dictatorial procedure with person 1 as the dictator, the counting rule will result in a decision which depends only on x_1 , a clear case of non-anonymity. However, F will be anonymous in both x and f_1, \dots, f_n if one imposes that all of the submitted f_1, \dots, f_n be anonymous procedures (together with some of our other axioms). This may be achieved by in Domain Condition (DC) choosing a set \mathcal{P} containing only anonymous procedures.

4.5 Strict Monotonicity

Next, we have to ensure that an option $y \in \mathcal{Y}$ is more likely to be chosen when more people aggregate into it, i.e. when it is more often the outcome $f_i(x)$ of submitted procedures f_i . Without imposing such a requirement, all of our other axioms would be met by the absurd rule which selects the option(s) $y \in \mathcal{Y}$ that occur *minimally* often (and hence perhaps never) among $f_1(x), \dots, f_n(x)$. But of course, Procedural Autonomy not only requires that decisions should depend on people's aggregations, but also that they depend on them in a positive way. So, in analogy to May's corresponding axiom, we impose

Strict Monotonicity (SM). *For any person $i \in \{1, \dots, n\}$ and any two metaprofiles in $\text{Dom}(F)$ differing only in person i 's submitted procedure, $m := (x, f_1, \dots, f_i, \dots, f_n)$*

and $m' := (x, f_1, \dots, f'_i, \dots, f_n)$, writing $y := f_i(x)$ and $y' := f'_i(x)$,

if $y' \neq y$ and y' is a chosen option under m (i.e. $y' \in F(m)$),
then y' is the unique chosen option under m' (i.e. $\{y'\} = F(m')$).

Loosely speaking, (SM) requires that if an option y' obtains more support, then it should be the *unique* choice in case it used to be a choice. What is "strict" about Strict Monotonicity is that y' should be the unique choice, not just a choice. This amounts to a strong responsiveness of collective choices to the procedural judgments of every person: a single person's change in favour of y' should already cause a "best among others" option y' to become a "better than all others" option.

4.6 Procedural Neutrality

We now come to the more controversial axiom of "Procedural Neutrality", which essentially states that F should be entirely neutral with regard to how information in the profile x is aggregated into decisions. F should fully relying on the group's proposals f_1, \dots, f_n . Arguably, Procedural Neutrality is a consequence of Procedural Autonomy provided that this premise is given a quite particular and strong interpretation.

Procedural Neutrality (PN). For any two metaprofiles in $Dom(F)$, (x, f_1, \dots, f_n) and (x', f'_1, \dots, f'_n) ,

if $f_i(x) = f'_i(x')$ for all $i \in \{1, \dots, n\}$,
then $F(x, f_1, \dots, f_n) = F(x', f'_1, \dots, f'_n)$.

This means that decisions taken by F should depend only on the outcomes of the submitted procedures f_1, \dots, f_n on the profile x , regardless of the particular nature of these procedures and of the profile (see Appendix A for the simple proof):

Proposition 8 F satisfies Procedural Neutrality (PN) if and only if there exists a function g defined on \mathcal{Y}^n such that

$$F(x, f_1, \dots, f_n) = g(f_1(x), \dots, f_n(x)), \quad \text{for all } (x, f_1, \dots, f_n) \in Dom(F). \quad (1)$$

The counting rule indeed satisfies (NP), the function g being here given by

$$g(y_1, \dots, y_n) := \{y \in \mathcal{Y} \mid y \text{ occurs maximally often among } y_1, \dots, y_n\}.$$

Because we want to justify Procedural Neutrality (PN) by appeal to our premise, this premise is now given the following precisification:

Procedural Autonomy (premise – second version). The manner in which F aggregates the profile into a decision should be entirely determined by the procedural judgments within the group, as far as information about these is strictly contained in the submitted metaprofile (x, f_1, \dots, f_n) .

(This is not yet the final version of the premise.) I call this a strong reading of the premise because of the words "entirely" and "strictly". We are more specific now due to the addition "as far as information...". By information "strictly" contained (in a statement, a profile, a metaprofile, etc.) we mean information explicitly contained¹⁷ or information undoubtedly implied by information explicitly contained. "Strictly" contained information has to be contrasted with information that follows by appeal to a particular interpretation, or follows by appeal to speculation, or follows "with high probability". Some examples: the statement "I ate ten slices of bread" does not strictly contain the information that I was hungry because I might have eaten for another reason; a profile x of individual preference rankings strictly contains no more than ordinal (non-cardinal) rankings.

The argument to justify (PN) is as follows. The strict informational content of a metaprofile (x, f_1, \dots, f_n) with regard to people's judgment about how the actual profile x should be aggregated is *fully contained* in the vector of outcomes $f_1(x), \dots, f_n(x)$. So by Procedural Autonomy (second version) it is correct that the decision rule F should use from (x, f_1, \dots, f_n) only the information of the outcomes $f_1(x), \dots, f_n(x)$. For this argument to work, the words "entirely" and "strictly" in the premise are crucial:

"entirely". Under a weaker reading of Procedural Autonomy whereby the aggregation method should be determined just *partly* by people's procedural judgments, the decision rule F would not only have to look at the outcomes $f_1(x), \dots, f_n(x)$, but would for instance also have to check whether "good" procedures f_1, \dots, f_n have been applied, where "good" would have to be an imported concept reflecting some "procedure bias", i.e. bias in favour of certain procedures believed to be more democratic, fairer, better truth-trackers, etc.

"strictly". By our premise, it is justified to henceforth call information "relevant" if it is information about people's judgments about how the actual profile x should be aggregated. Write $y_i := f_i(x)$ and consider the following statement:¹⁸

$$\bigwedge_{i=1}^n \text{"according to person } i, x \text{ should be aggregated into } y_i\text{"}. \quad (3)$$

This statement is fully encoded in the vector of outcomes $(y_1, \dots, y_n) = (f_1(x), \dots, f_n(x))$. So, what needs to be defended is the following

¹⁷Since we assume Procedural Judgments (J) and Sincerity (S), a person i 's submission of procedure f_i explicitly contains the information that this is person i 's judged-right procedure.

¹⁸One may argue that the strict informational content of the submitted procedures f_1, \dots, f_n is the statement

$$\bigwedge_{i=1}^n \left\{ \left[\bigwedge_{x' \in \mathcal{D}_i} \text{"according to person } i, x' \text{ should be aggregated into } f_i(x')\text{"} \right] \wedge \left[\bigwedge_{x'' \in \mathcal{X}^n \setminus \mathcal{D}_i} \text{"according to person } i, x'' \text{ is not an allowed profile"} \right] \right\}, \quad (2)$$

where " \bigwedge " denotes conjunction and \mathcal{D}_i denotes the domain of procedure f_i . As one easily verifies, statement (3) arises from statement (2) by removing any information about profiles other than the actual profile x .

Claim 1:¹⁹ *The relevant information strictly contained in the metaprofile (x, f_1, \dots, f_n) is the statement (3).*

At first sight, one might object that (x, f_1, \dots, f_n) in fact contains more relevant information. Indeed, person 1's procedure f_1 might tell us not just that option $y_1 = f_1(x)$ should win according to person 1 (and hence that all other options should lose), but perhaps also by "how much" y_1 is the right winner, and what options should nearly win or should absolutely not win. In summary, f_1 might tell us a lot about how strongly person 1 judges that different options should win or lose under x . For instance, if f_1 is the Borda count then one might try to infer from the Borda scores of options $y \in \mathcal{Y}$ how good or bad it would be for each option to win according to person 1 – information which gets lost in the statement (3) which only tells the winner. However, such information is not *strictly* contained in the metaprofile. The strict informational content of person 1's submitted f_1 is really just a mapping from profiles to decisions and hence no more than a statement as to which option should be declared winner on different profiles. Additional information could be guessed, perhaps even on good grounds, but this would never be more than guessing or speculating. It is pure speculation that person 1's second choice after the Borda winner would be the option with second highest Borda score – indeed, person 1 might instead believe that that the Condorcet winner should be chosen if not the Borda winner. Apart from the information that $y_1 = f_1(x)$ should win, the only other *strictly* contained information about person 1's procedural judgments is the information of the hypothetical winners $f_1(x')$ under counterfactual profiles $x' \neq x$ in the domain of f_1 ; but this is information about counterfactuals and is irrelevant for aggregating the actual profile x .

It might be worthwhile to cast more light on Procedural Neutrality (PN) by considering two implications of (PN) obtained by taking $f_i = f'_i$, respectively $x = x'$ in (PN). These implications express two slightly different kinds of neutrality relative to procedures, both of which can be separately motivate from Procedural Autonomy.

(PN1). *For any two metaprofiles in $\text{Dom}(F)$ containing identical procedures, (x, f_1, \dots, f_n) and (x', f_1, \dots, f_n) ,*

$$\begin{array}{l} \text{if} \quad f_i(x) = f_i(x') \text{ for all } i \in \{1, \dots, n\}, \\ \text{then} \quad F(x, f_1, \dots, f_n) = F(x', f_1, \dots, f_n). \end{array}$$

The justification of (PN1) is, loosely speaking, that the decision rule F should look at the profiles x only through the eyes of the submitted procedures f_1, \dots, f_n (because of the word "entirely"), and see from the profile no more than the outcomes $f_1(x), \dots, f_n(x)$ (because of the word "strictly"). Thus F cannot see the difference between the profiles x and x' in (PN1).

¹⁹If in addition using information about the meaning of options (as will be allowed by the final version of Procedural Autonomy), the relevant information still depends on the metaprofile only through the vector of outcomes y_1, \dots, y_n , so that (PN) is still an appropriate requirement. See the later footnote ²⁰.

(PN2). For any two metaprofiles in $\text{Dom}(F)$ containing identical profiles, (x, f_1, \dots, f_n) and $(x, f'_1, \dots, f'_n) \in \text{Dom}(F)$,

$$\begin{aligned} \text{if} \quad & f_i(x) = f'_i(x) \text{ for all } i \in \{1, \dots, n\}, \\ \text{then} \quad & F(x, f_1, \dots, f_n) = F(x, f'_1, \dots, f'_n). \end{aligned}$$

(PN2) expresses a different kind of procedural neutrality than (PN1). While (PN1) says that F should use no aggregation procedures other than f_1, \dots, f_n , (PN2) expresses that the manner in which F uses the procedures f_1, \dots, f_n should not depend on what these are (whether anonymous or not, Pareto-efficient or not, etc.). In other words, F should stay neutral relative to the kind of submitted procedures, which clearly corresponds to Procedural Autonomy taken strongly.

The reader has surely noticed the logical link between (PN), (PN1) and (PN2): (PN) implies (PN1)&(PN2), but the converse need not be true, and hence (PN) is really more than just the conjunction (PN1)&(PN2).

4.7 Neutrality relative to both procedures and options

In fact, we need a generalisation of Procedural Neutrality (PN), namely an axiom covering neutrality both relative to procedures and relative to options. We will see that this is an adequate assumption under the second version of Procedural Autonomy, but that under our final version of Procedural Autonomy the adequacy will crucially rely on an "unknown distance" assumption.

In effect, our stronger neutrality axiom says that for F the options $y \in \mathcal{Y}$ are no more than symbols, with unknown meaning or interpretation, treated entirely symmetrically (as in May's neutrality). More precisely, our stronger neutrality axiom requires the following. Already by Procedural Neutrality the decision $F(x, f_1, \dots, f_n)$ should only depend on the outcomes y_1, \dots, y_n of f_1, \dots, f_n on x . Now let $\pi : \mathcal{Y} \mapsto \mathcal{Y}$ be a permutation of the option space and let us ask how F would have to decide if the outcomes were $\pi(y_1), \dots, \pi(y_n)$ instead of y_1, \dots, y_n . Our stronger neutrality axiom prescribes that if on the outcomes y_1, \dots, y_n F decides y , then on the outcomes $\pi(y_1), \dots, \pi(y_n)$ F decides $\pi(y)$:

Neutrality (N). For any two metaprofiles in $\text{Dom}(F)$, (x, f_1, \dots, f_n) and (x', f'_1, \dots, f'_n) , and any permutation of the decision space $\pi : \mathcal{Y} \mapsto \mathcal{Y}$,

$$\begin{aligned} \text{if} \quad & f_i(x) = \pi(f'_i(x')) \text{ for all } i \in \{1, \dots, n\}, \\ \text{then} \quad & F(x, f_1, \dots, f_n) = \pi(F(x', f'_1, \dots, f'_n)). \end{aligned}$$

Note that, since $F(x', f'_1, \dots, f'_n)$ is a subset of \mathcal{Y} , not an element, in writing $\pi(F(x', f'_1, \dots, f'_n))$ we use π as a function defined on *subsets* of \mathcal{Y} , namely $\pi(Z) := \{\pi(y) | y \in Z\}$ for all $Z \subseteq \mathcal{Y}$.

By taking π to be the identical permutation $\pi(y) = y$, we see that Neutrality (N) is a stronger requirement than Procedural Neutrality (PN):

Proposition 9 *If F satisfies Neutrality (N), then it satisfies Procedural Neutrality (PN).*

Justification of Neutrality (N) on the basis of Procedural Autonomy (second version). By Claim 1, Procedural Autonomy (second version) implies that the statement (3) is the only information that F may use in aggregating x . But statement (3) does not tell anything about the meaning of the outcomes y_1, \dots, y_n ; these are just symbols with unknown interpretation. So, since names or symbols should not influence decisions, it follows that Neutrality (N) is justified. Indeed, violation of (N) would mean that decisions change under renaming each option $y \in \mathcal{Y}$ into $\pi(y)$ where π is any permutation of \mathcal{Y} .

Final version of Procedural Autonomy. As two examples will illustrate, the second version of Procedural Autonomy seems ultimately untenable and should be replaced by the following more adequate principle:

Procedural Autonomy (premise – final version). *The manner in which F aggregates the profile into a decision should be entirely determined by the procedural judgments within the group, as far as information about these is strictly contained in the submitted metaprofile (x, f_1, \dots, f_n) or can be deduced in an appropriate way from the meaning of options.*

What has changed compared to the second version is the addition of the clause "or can be deduced...". As will be seen, the only relevant information that can possibly be deduced from the meaning of options is information about the "distance" between options. Often, no such information can be deduced *in an appropriate way*, because "distance" is a matter of subjective assessment and/or for reasons of fairness or neutrality relative to options. For instance, in an election of a candidate, it would seem *not appropriate* if F deduced, say, from the political orientation of the different candidates that certain candidates are closer to each other than others.

The following two examples contain cases where information about the distance of options *can* be deduced in an appropriate way from the meaning of options. This information will lead us to reject Neutrality (N) for these examples. So, the examples suggest that an additional "unknown distance" assumption has to be satisfied in order for Neutrality (N) to be justified, as will be discussed after the examples.

First example of not unknown distance. Assume that \mathcal{Y} consists of the options y_1 = "building two houses", y_2 = "building one house", and y_3 = "building no house". Obviously, y_1 is closer to y_2 than to y_3 . The group consists of just $n = 3$ persons, where person 1 aggregates into $y_1 = f_1(x)$, person 2 into $y_2 = f_2(x)$, and person 3 into $y_3 = f_3(x)$. So, each option occurs once among $f_1(x), f_2(x), f_3(x)$, and hence the counting rule results in a tie between all three options: $F(x, f_1, f_2, f_3) = \{y_1, y_2, y_3\}$. On the other hand, instead of a tie it would seem reasonable to decide in favour of y_2 (building one house), as the "best compromise" between the outcomes $f_1(x), f_2(x), f_3(x)$. The latter seems the way to maximally respect people's procedural concerns. Let us see precisely why

- (i) the choice of building one house is incompatible with (N), and
- (ii) this choice is in accordance with the final but not with the second version of Procedural Autonomy,

As a consequence, by (i) Neutrality (N) is an inappropriate axiom for this decision problem with certain knowledge about distances, and by (ii) the second version of Procedural Autonomy cannot be a generally appropriate principle.

(i): Let F be the decision rule that decides in favour of y_2 (building one house) if each of y_1, y_2, y_3 occurs once among $f_1(x), f_2(x), f_3(x)$, and otherwise chooses the option occurring two or three times. Consider again a metaprofile (x, f_1, f_2, f_3) such that $f_1(x) = y_1, f_2(x) = y_2, f_3(x) = y_3$. Then $F(x, f_1, f_2, f_3) = \{y_2\}$ (building one house). Of course, the same decision would be made if persons 1 and 2 "switch procedures": $F(x', f'_1, f'_2, f'_3) = \{y_2\}$ where $(x', f'_1, f'_2, f'_3) := (x, f_2, f_1, f_3)$. This is incompatible with Neutrality (N): by taking $\pi : \mathcal{Y} \mapsto \mathcal{Y}$ to be a permutation that swaps y_1 and y_2 and leaves y_3 unchanged ($\pi(y_1) = y_2, \pi(y_2) = \pi(y_1), \pi(y_3) = y_3$), we have

$$\begin{aligned} f_i(x) &= \pi(f'_i(x')) \text{ for all } i \in \{1, 2, 3\}, \\ \text{but } F(x, f_1, f_2, f_3) &\neq \pi(F(x', f'_1, f'_2, f'_3)), \end{aligned}$$

since $F(x, f_1, f_2, f_3) = \{y_2\}$ and $\pi(F(x', f'_1, f'_2, f'_3)) = \pi(\{y_2\}) = \{y_1\}$.

(ii): Why is it in accordance with (or even required by) Procedural Autonomy (final version) that F chooses y_2 = "build one house" on the the metaprofile (x, f_1, f_2, f_3) ? As discussed in the previous section, the relevant information strictly contained in the metaprofile is the conjunction

$$\bigwedge_{i=1}^3 \text{"according to person } i, x \text{ should be aggregated into } y_i\text{"}. \quad (4)$$

But then we have "opened" the options y_i , thus discovering that y_2 means a compromise between y_1 and y_3 . This is clearly incompatible with the second version of Procedural Autonomy since we use information about the meaning of options, and more precisely about "distances" between options. From these "distances" we have inferred that by person 1 it is second-best to aggregate x into y_2 and worst to aggregate x into y_3 , and that by person 3 it is second-best to aggregate x into y_2 and worst to aggregate x into y_1 . This lead us to conclude that people's procedural judgments are better respected by choosing y_2 than by a tie (the outcome of the counting rule).

Second example of not unknown distance. An important type of decision problem that violates the unknown distance assumption is problems in which the options $y \in \mathcal{Y}$ are collective preference rankings over some underlying set of issues, candidates, measures, etc. Indeed, if \mathcal{Y} were the set of all transitive and complete rankings over the set $\{A, B, C\}$, it seems that it is appropriate to consider the option $(A \succ B \succ C)$ as being closer to the option $(A \sim B \sim C)$ than to the option $(A \prec B \prec C)$. By contrast, the counting rule ignores such relations. Again, we have used information about the meaning of options, which is compatible with the final but not with the second version of Procedural Autonomy.

Justification of Neutrality (N) on the basis of Procedural Autonomy (final version) and the unknown distance assumption. In the two above examples, what has lead us to reject Neutrality (N) was knowledge about the "distance" between options. This suggests that we need precisely the following

Unknown distance assumption: *For any three distinct options $y, y', y'' \in \mathcal{Y}$, if y is the outcome of a person's judged-right procedure, it is not appropriate that F deduces from the meaning of the options y, y', y'' any information about the person's judgment of whether an aggregation into y' is better than, worse than or equally good as an aggregation into y'' .*

Informally, F should not "know" how the distance from y to y' compares to the distance from y to y'' , for any three distinct options $y, y', y'' \in \mathcal{Y}$. This seems justified in many cases, because the distance between options is a matter of subjective assessment and/or for reasons of fairness or neutrality relative to options (e.g. the case of electing a political candidate mentioned before the two above examples).

Why is Neutrality (N) an appropriate axiom under the unknown distance assumption? By Claim 1 in the previous section, the relevant information strictly contained in the metaprofile (x, f_1, \dots, f_n) is the statement (3), i.e. the conjunction

$$\bigwedge_{i=1}^n \text{"according to person } i, x \text{ should be aggregated into } y_i \text{"}$$

where $y_i := f_i(x)$ and (because of Procedural Autonomy) information is called "relevant" if it is information about people's judgments about how the actual profile x should be aggregated. Procedural Autonomy (final version) also allows us to deduce relevant information from the meaning of options. Now, do such properties enable us to go beyond the information (3), i.e. do they add any relevant information to (3)? I shall argue that this is not so:

Claim 2: *Under the unknown distance assumption, the relevant information that is strictly contained in the metaprofile (x, f_1, \dots, f_n) or can be deduced in an appropriate way from the meaning of options is the statement (3).*

If this claim is true, then Procedural Autonomy (final version) implies that the statement (3) is the only information that F may use in aggregating x . So Neutrality (N) is an adequate requirement by exactly the same argument as that given above when justifying Neutrality (N) on the basis of the second version of Procedural Autonomy.

Now let us justify Claim 2. In principle, the relevant information that could possibly be added to (3) by using the meaning of options is, for any person i (who aggregates x into $y := y_i$), information about²⁰

²⁰In general, one might imagine that, for each $y \in \mathcal{Y}$, the meaning of options imply an (ordinal or cardinal) measure U_y of the degrees to which a person who aggregates into y judges that it is right to aggregate into the various options in \mathcal{Y} . In the case of an ordinal measure, U_y is a binary order relation on \mathcal{Y} , and in the case of a cardinal measure U_y is a utility function on \mathcal{Y} (with a special interpretation of "utility"). Besides ranking y highest, U_y may give a lot more information of the type of (a) or (b). In summary, the relevant information that is strictly contained in the metaprofile (x, f_1, \dots, f_n) or can be deduced in an appropriate way from the meaning of options would not be the statement (3), but the statement

$$\bigwedge_{i=1}^n \text{"according to person } i, \text{ the degree to which it is right to aggregate } x \text{ into the various options in } \mathcal{Y} \text{ is given by } U_{y_i} \text{" .} \quad (5)$$

(a) *how strongly* the person judges that x should not be aggregated into the various other options $y' \neq y$,

(b) *how strongly* the person judges that x should be aggregated into y .

Regarding (a), the unknown distance assumption is there precisely to exclude the existence of such additional relevant information. What F knows from (3) is that, by the person, x should not be aggregated into any option $y' \neq y$, but the meaning of options does not allow F to distinguish between "more inappropriate" and "less inappropriate".

Regarding (b), the answer is that the meaning of options does not (or does not in an appropriate way) entail information about how strongly the person judges that x *should* be aggregated into y (with or without the unknown distance assumption). The argument is best illustrated with an example. Consider the binary problem of either convicting or acquitting the defendant. One might argue that false convictions are considered by all persons as worse than false acquittals, and hence that a person believes more strongly in the rightness of his or her aggregation of x if the aggregation is into "acquit" than if the aggregation is into "convict". However, this argument is problematic, because a person's aggregation into "acquit" may have occurred by a very narrow majority (if f_i nearly resulted in a tie on x), thus overall causing the person to believe *little strongly* in the rightness of his or her aggregation of x *although* this aggregation is into "acquit". In general, without knowing by how narrow a margin a person's judged-right procedure reached decision y , the meaning of options can not tell certain information about the strength of this judgment. But knowledge of how "narrowly" the person aggregates into y is unavailable, as has been argued in the previous section.²¹

On the difference to May's neutrality axiom. It should be noted that the unknown distance assumption necessary for Neutrality (N) to be appropriate differs from assumptions needed for May's neutrality axiom to be appropriate. In May's binary case, $\mathcal{Y} = \{a, b\}$, the unknown distance assumption is always fulfilled (and hence Neutrality (N) is always justified), because in the absence of three distinct options $y, y', y' \in \mathcal{Y}$ the unknown distance assumption is vacuously true (no two distances have to be compared to each other). But if a were "conviction of the defendant" and b were "acquittal of the defendant" then May's neutrality axiom would perhaps not be justified, given the different nature of options and the different status of a false conviction and a false acquittal. May's neutrality axiom is indeed based on an assumption that options should be treated as similar objects, which is an assumption on the nature of options and not on their distance to each other. In the convict/acquit example, May's majority rule may not be appropriate, but our counting rule may well apply. Indeed,

What we argue is that under the unknown distance assumption U_y is the (highly incomplete) ordinal measure that strictly prefers y to any $y' \neq y$ and is silent on the order between any pair not involving y – so that (5) is equivalent to (3). Note that without the unknown distance assumption the weaker axiom of Procedural Neutrality (PN) is still justified, since the available relevant information (5) depends on (x, f_1, \dots, f_n) only through the vector of outcomes y_1, \dots, y_n .

²¹Only probabilistic statements may perhaps be deduced from the meaning of options. For instance, in the convict/acquit it might be true that, given that false convictions are worse than false acquittals, *on average* persons who aggregate x into "acquit" believe more strongly in the rightness of their aggregation than persons who aggregate x into "convict". We here assume that the derivation of such probabilistic information is not appropriate, and hence that F should ignore such information by Procedural Autonomy (final version).

the dissymmetry between the two options is likely to have already been accounted for in the submitted procedures f_1, \dots, f_n , and, by Procedural Autonomy, F should rely on these procedural judgments rather than imposing additional dissymmetry.

5 Summary

Let us summarise the essential points. The entire discussion was based on the premise of Procedural Autonomy whereby the information in the profile should be aggregated according to the group's own procedural judgments. This contrasts with the standard social-choice-theoretic aim of searching for an independently "good" procedure. One could try to give different, pragmatic or normative, justifications in support of Procedural Autonomy. But, rather than defending Procedural Autonomy, our aim has been to explore what follows *if* this premise is accepted.

Following Procedural Autonomy, we have called a procedure "legitimate" if it is reasonably non-controversial within the group, whatever the nature of this procedure. Divergent procedural judgments may imply that no procedure is legitimate. This fact motivated us to develop a new concept of legitimate decisions. This concept was proposed in great generality, making no assumptions on the type of profile or decision space, apart from the unknown distance assumption. Legitimate decisions do not conflict with, i.e. are a generalisation of legitimate procedures, because *if* a legitimate procedure exists *then* it entails a legitimate decision (Theorem 2). The problematic assumption on which the existence of legitimate decisions are based is Procedural Judgments (J) whereby each person has a judged-right procedure. The latter seems unrealistic even in societies with a high level of education and awareness. The assumption seems more realistic in committees where committee members openly discuss and collectively deliberate over the appropriateness of possible decision procedures. The obvious decision rule to reach legitimate decisions, called the counting rule, is highly manipulable. For the counting rule to always generate legitimate decisions one needs both Procedural Judgments (J) and Sincerity (S) (Proposition 3). The latter is a second tough assumption which can perhaps be achieved by disconnecting the procedure submission from concrete decision problems.

The axiomatic justification for the counting rule and the present concept of legitimate decisions was given in the form of two theorems that characterise the counting rule as the only decision rule subject to certain requirements. These requirements follow, arguably, from a quite particular precisification of Procedural Autonomy, together with the unknown distance assumption. The two theorems differ in their domain requirements. The first theorem is based on the universal domain assumption and is partly analogous to May's Theorem: if a decision rule F satisfies Universal Domain (UD), then it is the counting rule if and only if it satisfies Anonymous Procedure Submission (APS), Strict Monotonicity (SM) and Neutrality (N). But a universal domain may not be desirable, since one might want to exclude some procedures, such as dictatorial procedures, even if not in line with Procedural Autonomy taken strongly. The second theorem relaxes the domain assumption to Domain Condition (DC). If one wants to achieve anonymity of the counting rule relative to the profile $x = (x_1, \dots, x_n)$, it is sufficient to impose that all submitted procedures f_i be anonymous, a special case of Domain Condition (DC).

To conclude, one might wonder how it was possible to obtain possibility and uniqueness theorems which – quite untypically for social choice theory – apply to very general decision problems. The important remark here is that this was possible *only* because we have taken into account no more than people’s judged-right procedures, ignoring any other information about procedural judgments such as judged-worst procedures or approved versus not-approved procedures. If one wanted to take additional procedural information into account (which would require a stronger and hence even less realistic version of the axiom of Procedural Judgments (J)), the definition of legitimate decisions would have to be revised, raising difficult normative questions on how to "aggregate the aggregations". Uniqueness theorems would be much harder to obtain; and impossibility results may follow too because the decision rule would no longer have to decide on the basis of a vector of outcomes $(f_1(x), \dots, f_n(x))$, but possibly on the basis of a vector of rankings of outcomes.

A Proofs

PROOF OF THEOREM 6. Assume that F satisfies (UD), (APS), (SM) and (N). To prove that F is the counting rule, consider any given metaprofile $m := (x, f_1, \dots, f_n) \in \text{Dom}(F)$. Recall that $L_m(y)$ is the number of times y occurs among $f_1(x), \dots, f_n(x)$. We have to show that

$$F(m) = \left\{ y \in \mathcal{Y} \mid L_m(y) = \max_{y' \in \mathcal{Y}} L_m(y') \right\}. \quad (6)$$

" \subseteq ": First, suppose $y \in F(m)$ and assume for contradiction that $L_m(y) < \max_{y' \in \mathcal{Y}} L_m(y')$. Let $y^* \in \mathcal{Y}$ be such that $L_m(y^*) = \max_{y' \in \mathcal{Y}} L_m(y')$. Let f be any procedure defined at least on x such that $f(x) = y$. Consider the new metaprofile m_1 obtained from m by replacing by f as many procedures f_i for which $f_i(x) = y^*$ as necessary to achieve that y and y^* have switched legitimacy: $L_{m_1}(y) = L_m(y^*)$. By Strict Monotonicity, $F(m_1) = \{y\}$. Now modify m_1 into a new metaprofile m_2 by the following permutation of the procedure vector, which by Anonymous Procedure Submission (APS) leaves the result unchanged: $F(m_2) = F(m_1) = \{y\}$. Let N_y be the set of persons i for whom $f_i(x) = y$, and let N_{y^*} be the set of persons i for whom $f_i(x) = y^*$ and f_i is still the person’s procedure in m_1 . Then $\#N_y = \#N_{y^*}$, and so there exists a one-to-one correspondence between N_y and N_{y^*} . To get from m_1 to m_2 , each person in N_y switches procedure with a person in N_{y^*} according to a one-to-one correspondence. Comparing m_2 to m , what has changed is that every person i whose procedure in m lead to y^* now has a procedure leading to y , and conversely every person i whose procedure in m lead to y now has a procedure leading to y^* . Thus, using Neutrality (N) with the permutation $\pi : \mathcal{Y} \mapsto \mathcal{Y}$ defined by $\pi(y^*) := y$, $\pi(y) := y^*$, and $\pi(y') := y'$ for all $y' \notin \{y, y^*\}$, since $F(m_2) = \{y\}$ we deduce $F(m) = \pi(F(m_2)) = \{y^*\}$, in contradiction with $y \in F(m)$.

" \supseteq ": Now suppose $y \in \{y \in \mathcal{Y} \mid L_m(y) = \max_{y' \in \mathcal{Y}} L_m(y')\}$, and assume for contradiction that $y \notin F(m)$. By definition of a decision rule, $F(m)$ is non-empty, so there exists a $y^* \in F(m)$. By the inclusion " \subseteq " proven above, $L_m(y^*) = \max_{y' \in \mathcal{Y}} L_m(y')$, and hence $L_m(y^*) = L_m(y)$, i.e. there are as many individuals i with $f_i(x) = y$ as there are individuals j with $f_j(x) = y^*$. So we can modify m into m_1 by letting each individual i with $f_i(x) = y$ switch procedure with an individual j with $f_j(x) = y^*$,

according to a one-to-one correspondence. By Anonymous Procedure Submission (APS), $F(m_1) = F(m)$. On the other hand, we may again apply Neutrality (N): letting $\pi : \mathcal{Y} \mapsto \mathcal{Y}$ be the permutation defined by $\pi(y^*) := y$, $\pi(y) := y^*$, and $\pi(y') := y'$ for all $y' \notin \{y, y^*\}$, we deduce that $F(m_1) = \pi(F(m))$, so that by $y^* \in F(m)$ we have $y = \pi(y^*) \in F(m_1) = F(m)$, in contradiction with $y \notin F(m)$. **QED.**

PROOF OF THEOREM 7. The proof works exactly like for Theorem 6, where (DC) come in as follows:

- by (i) in (DC), any "switching" of procedures between persons is allowed, i.e. one stays within $Dom(F)$;
- by (ii) in (DC), in the proof of " \subseteq " one may assume that $f \in \mathcal{P}$, so that again by (i) one stays within $Dom(F)$. **QED.**

PROOF OF PROPOSITION 8. First, assume (PN). The function g defined as follows obviously satisfies the desired relation. Let $(y_1, \dots, y_n) \in \mathcal{Y}^n$. If there exist metaprofiles (x, f_1, \dots, f_n) in $Dom(F)$ such that $y_i = f_i(x)$ for all $i = 1, \dots, n$, then by (PN) the outcome $F(x, f_1, \dots, f_n)$ is the same for each of these metaprofiles, and hence we may define $g(y_1, \dots, y_n)$ as being precisely this outcome. If no such metaprofiles (x, f_1, \dots, f_n) exists in $Dom(F)$, then $g(y_1, \dots, y_n)$ may be defined arbitrarily.

Conversely, if there exists a function g satisfying (1) for all metaprofiles in $Dom(F)$, then, clearly, (PN) is satisfied. **QED.**

REFERENCES

- May, K. O (1952). A Set of Independent, Necessary and Sufficient Conditions for Simple Majority Decision. *Econometrica* 20 (4), p. 680-684.