Synthese (2008) 163:227–243
DOI 10.1007/s11229-007-9200-y

On the role of language in social choice theory

Marc Pauly

Received: 15 February 2007 / Accepted: 23 May 2007 / Published online: 10 August 2007

Abstract Axiomatic characterization results in social choice theory are usually compared either regarding the normative plausibility or regarding the logical strength of the axioms involved. Here, instead, we propose to compare axiomatizations according to the language used for expressing the axioms. In order to carry out such a comparison, we suggest a formalist approach to axiomatization results which uses a restricted formal logical language to express axioms. Axiomatic characterization results in social choice theory then turn into definability results of formal logic. The advantages of this approach include the possibility of non-axiomatizability results, a distinction between absolute and relative axiomatizations, and the possibility to ask how rich a language needs to be to express certain axioms. We argue for formal minimalism, i.e., for favoring axiomatizations in the weakest language possible.

Keywords Social choice theory · Logic · Judgment aggregation

1 Introduction

The central results in social choice theory (SCT) are of an axiomatic nature. Arrow's theorem (1951), the Gibbard-Satterthwaite theorem (Gibbard 1973; Satterthwaite 1975) and May's theorem (May 1952) all characterize a particular social choice procedure—dictatorship in the case of Arrow and Gibbard-Satterthwaite, and majority voting in the case of May—by means of certain axioms such as neutrality, anonymity, etc. The question then arises how these particular axioms come to be chosen as a basis for any given result. The first most obvious answer is, naturally, that

M. Pauly (✉)
Department of Philosophy, Stanford University,
Stanford, CA 94305-2155, USA
e-mail: Marc.Pauly@stanford.edu

they allow one to prove the desired result. Second, the axioms are usually defended from a normative point of view, they describe properties of social choice procedures which seem normatively desirable. A voting procedure should treat all alternatives equal (neutrality), it should give no advantage to any one of the voters (anonymity), and so on. Third, there is the issue of logical strength. Ideally, one would like the theorem to be as general as possible, i.e., the axiomatic conditions should be as weak as possible.

There has been a lot of discussion in social choice theory concerning certain axioms in the classic results mentioned. Arrow's independence of irrelevant alternatives (IIA) condition has been criticized as being too strong, and so has May's condition of positive responsiveness (PR). In these discussions, it is sometimes not clear whether the strength of these conditions is felt to be objectionable on normative grounds, or simply on logical grounds. Normatively, one might dislike Arrow's IIA condition since it excludes arguably sensible social choice procedures such as the Borda count. Logically, on the other hand, it is always advisable to look for weaker conditions, thereby arriving at more general results.

A fourth criterion for evaluating axiomatizations that is occasionally invoked is methodological. As an example, one may wish to exclude counterfactual axioms which express not only how a social choice procedure should behave on a given preference profile, but also how it should have behaved if the preferences of the individuals had been different. We will have more to say on this discussion below.

Somewhat related to this methodological criterion, there is, a fifth criterion for comparing different axiomatizations which will be the central focus of this paper. This criterion can be applied to compare two axiomatizations whether or not these are logically comparable, and we shall call it the *language criterion*. SCT axioms are usually formulated in the general language of mathematics, involving relations, etc. with no restrictions placed on the set-theoretic complexity of the notions used, number of quantifier alternations, and so on. By looking at the language used in formulating the axiom, we can evaluate the logical complexity as well as the conceptual notions underlying the axiom.

In order to fully apply this language criterion in evaluating axiomatizations, we argue for an approach we call *formal minimalism*, i.e., we argue for formulating the axioms in a restricted formal logical language. Once this is done, SCT characterization results turn into definability theorems of formal logic. The advantages of this Gestalt-switch are fourfold: First, the formal logical language makes the conceptual ingredients of the axioms explicit by clearly specifying the concepts needed to formulate the axioms.

Second, since the formal language has a restricted syntax, one can investigate hierarchies of languages, allowing one to formulate more and more complicated axioms. In this way, we may not only compare axioms in terms of logical strength, but also in terms of the expressive power of the language used to formulate them. Other things being equal, we argue for preferring minimal axiomatizations, i.e., axiomatizations which use an axiomatic language with as few concepts as possible, with a minimal syntax. The axiomatization of majority voting given in the context of judgment aggregation in Sect. 3 will be an example of this *formal minimalism*.

Third, by using a precisely defined and limited formal language, one can obtain non-axiomatizability results, stating that a certain social choice procedure or class of procedures cannot be axiomatized within a given conceptual framework, i.e., within a given language. Fourth, even when a set of axioms fails to characterize fully a particular social choice procedure, the axioms may nonetheless fully characterize the procedure *with respect to the given language*. Hence, besides the standard language-independent notion of axiomatization, we also obtain a second weaker relative notion of axiomatization which is language dependent.

The plan of this paper is as follows: To introduce the reader to this language criterion, we start in Sect. 2 with May's classical characterization of majority voting, comparing its logical structure with a more recent characterization result of Maskin (1995). While these two results are formulated in the framework of preference aggregation, in Sect. 3, we will discuss a recent characterization of majority voting in the framework of judgment aggregation. Since this latter characterization result makes use of a formal logical language for expressing axioms, it is here that we can fully develop our argument for a language evaluation of axiomatization results in social choice theory. Section 4 presents an outline of the formal minimalist methodology we extrapolate from the example given. It also develops the necessary concepts such as absolute and relative axiomatization. Staying within the framework of judgment aggregation, Sect. 5 then presents a particular example of a non-axiomatizability result (in the sense of absolute axiomatization), together with a positive result on relative axiomatization. Finally, Sect. 6 will summarize the general methodological claims that are illustrated by this case study of majority voting.

2 Majority voting in preference aggregation

In order to make our case for formal minimalism, in Sect. 3 we shall look at an axiomatic characterization of majority voting in the framework of judgment aggregation. In this section, we shall review two axiomatizations of majority voting in the framework of preference aggregation. Later on, this will allow us to compare axiomatization results across frameworks. In this section, we will already make some informal remarks about axiomatic language differences within the same preference-based framework. We shall attempt to state these results in a way that the logical structure of the axioms becomes apparent.

The best-known axiomatic characterization of majority voting is due to May (1952). The set of voters is $N = \{1, \ldots, n\}$, where each of the n voters has a preference over two alternatives which we shall denote as 1 and -1. A voter can prefer alternative 1, alternative -1, or be indifferent between the two alternatives, denoted by 0. A *group decision function* $D\colon \{-1, 0, 1\}^n \to \{-1, 0, 1\}$ maps n individual preferences over these two alternatives into a group preference. The *majority rule* is the group decision function D_m defined as

$$D_m(x_1, \dots, x_n) = \begin{cases} 1 & \text{if } \sum_{i=1}^n x_i > 0 \\ 0 & \text{if } \sum_{i=1}^n x_i = 0 \\ -1 & \text{if } \sum_{i=1}^n x_i < 0 \end{cases}$$

Given some finite set X, let $Pm_X = \{f : X \rightarrow X \mid \text{f is injective}\}$ be the set of permutations over X. May then goes on to characterize the majority rule as the only group decision function satisfying the following properties:

Axiom 1 (*Anonymity*)

$$\forall p \in Pm_N \; \forall x_1, \dots, x_n \in \{-1, 0, 1\} : \; D(x_1, \dots, x_n) = D(x_{p(1)}, \dots, x_{p(n)})$$

Axiom 2 (*Neutrality*) $\forall x_1, \dots, x_n \in \{-1, 0, 1\} : -D(x_1, \dots, x_n) = D(-x_1, \dots, -x_n)$

Axiom 3 (*Positive Responsiveness*) $\forall x_1, \dots, x_n \in \{-1, 0, 1\} \; \forall i \in \{1, \dots, n\}$:

$$\big((D(x_1, \dots, x_i, \dots, x_n) = 0 \vee D(x_1, \dots, x_i, \dots, x_n) = 1) \wedge x_i < 1\big) \\ \rightarrow D(x_1, \dots, x_{i-1}, x_i + 1, x_{i+1}, \dots, x_n) = 1$$

Theorem 1 (May 1952) *A group decision function satisfies anonymity, neutrality and positive responsiveness if and only if it is the majority rule.*

Of the three axioms, positive responsiveness has received most criticism. While anonymity and neutrality are properties which appear quite regularly in characterization results in SCT, positive responsiveness has been criticized as too strong (Campbell and Kelly 2000). In particular, positive responsiveness demands that in cases where the group decision function is indifferent between the two alternatives -1 and 1, it suffices for a single voter to change from undecided to 1 for the group as well to accept alternative 1. Since the normative status of this axiom seems questionable, people have looked for alternative weaker axioms to replace positive responsiveness.

Regarding the axiomatic language employed by the three axioms, two things are worth noting. While neutrality is simple enough, anonymity does involve a rather complex quantification over the set of all permutation functions. Since the number of voters is finite, one can in principle eliminate this quantification and write out a formula containing all possible permutations, but the result would be a formula with length essentially exponential in the number of individual voters. Thus, while not much thought is usually given to anonymity, it is certainly more complex than neutrality.

A second observation will prove yet more relevant in comparison with the alternative axiomatization we shall present subsequently. All three axioms refer only to the group decision function under consideration. In other words, each axiom $A(D)$ states a property of a group decision function D, and this property does not refer to any other group decision function, only to D itself. This is in marked contrast to the following axiomatization by Maskin.

Maskin (1995) has obtained a characterization of majority voting in the case of preference aggregation over an arbitrary finite set of alternatives. Hence, the setting is more general than that of May's theorem, but for our purposes of comparing the axiomatic language informally, this difference will not turn out to be essential. Again, we have a finite set of voters $N = \{1, \ldots, n\}$. We will denote the finite set of alternatives as A, and let $L(A)$ be the set of linear orders over A, i.e., the set of all binary relations $\succeq \subseteq A \times A$ which are complete and transitive, and for which $x \succeq y \succeq x$ implies that $x = y$. We write $x \succ y$ to denote $x \succeq y \not\succeq x$. A *collective choice rule* (CCR) R maps a profile of n strict linear orders to a complete binary relation over A. Thus, the outcome of a CCR does not need to be transitive.

The *majority rule* in this context is the CCR R_m defined as follows: For all $\succeq_1, \ldots, \succeq_n$, we let $(a, b) \in R_m(\succeq_1, \ldots, \succeq_n)$ if and only if $|\{i \in N | a \succeq_i b\}| \geq |\{i \in N | b \succeq_i a\}|$. As Condorcet's paradox illustrates, R_m does not always produce a transitive ordering. We now consider the following axioms for a CCR R:

Axiom 4 (*Neutrality*) For $f \in Pm_A$, let $f(\succeq) = \{(f(a), f(b)) | (a, b) \in \succeq\}$. Then

$$\forall f \in Pm_A \forall \succeq_1, \ldots, \succeq_n \in L(A): \; R(f(\succeq_1), \ldots, f(\succeq_n)) = f(R(\succeq_1, \ldots, \succeq_n))$$

Axiom 5 (*Anonymity*)

$$\forall f \in Pm_N \forall \succeq_1, \ldots, \succeq_n \in L(A): \; R(\succeq_{f(1)}, \ldots, \succeq_{f(n)}) = R(\succeq_1, \ldots, \succeq_n)$$

Axiom 6 (*Pareto Property*) $\forall a, b \in A : \forall \succeq_1, \ldots, \succeq_n \in L(A):$

$$(\forall i \in N : a \succ_i b) \rightarrow \big((a, b) \in R(\succeq_1, \ldots, \succeq_n) \wedge (b, a) \notin R(\succeq_1, \ldots, \succeq_n)\big)$$

Axiom 7 (*IIA*) $\forall \succeq_1, \ldots, \succeq_n, \succeq'_1, \ldots, \succeq'_n \in L(A) \; \forall a, b \in A:$

$$(\forall i \in N : a \succeq_i b \leftrightarrow a \succeq'_i b) \rightarrow \big((a, b) \in R(\succeq_1, \ldots, \succeq_n) \leftrightarrow (a, b) \in R(\succeq'_1, \ldots, \succeq'_n)\big)$$

These axioms are all well-known from the SCT literature. Anonymity and neutrality are simple generalizations of the earlier axioms from 2 to more alternatives. The Pareto property expresses that whenever everyone agrees on strictly preferring a to b, then the group as a whole must also strictly prefer a to b. Independence of irrelevant alternatives (IIA) states that the group preference between a and b may only depend on the individual preferences between a and b.

The following axiom constitutes the conceptual novelty in Maskin's result. It states some kind of robustness of a CCR R. If R is maximally transitive over a set of CCRs $\mathcal{C}$, then R produces a transitive social choice relation more often than any other rule in $\mathcal{C}$. More precisely, whenever a rule R' in $\mathcal{C}$ produces a transitive social choice relation for some domain of preferences, then so does R, and there is at least one domain where R produces transitive social preferences and R' does not. Formally, a CCR R is *transitive on a domain* $L \subseteq L(A)$, denoted as $Trans(R, L)$, provided that for all $\succeq_1, \ldots, \succeq_n \in L$, and $a, b, c \in A$, we have that $(a, b) \in R(\succeq_1, \ldots, \succeq_n)$ and $(b, c) \in R(\succeq_1, \ldots, \succeq_n)$ imply $(a, c) \in R(\succeq_1, \ldots, \succeq_n)$.

Axiom 8 (*Maximal Transitivity among $\mathcal{C}$*) For all collective choice rules $R' \in \mathcal{C}$ such that $R \neq R'$,

$$(\forall L \subseteq L(A)\colon Trans(R', L) \rightarrow Trans(R, L)) \wedge (\exists L \subseteq L(A)\colon \neg Trans(R', L) \wedge Trans(R, L))$$

For an odd number of voters, Maskin's result then characterizes the majority rule as the maximally transitive rule among those satisfying the earlier axioms. The result has subsequently been generalized by weakening the axioms involved and by also allowing the set of voters to be even (Campbell and Kelly 2000). Since these generalizations do not alter our discussion regarding the point of axiomatic language, we shall stick to Maskin's original result.

Theorem 2 (Maskin 1995) *Let $\mathcal{C}$ be the set of collective choice rules which satisfy anonymity, neutrality, the Pareto property and independence of irrelevant alternatives. Then if n is odd, a CCR is maximally transitive among $\mathcal{C}$ if and only if it is the majority rule.*

Two of the axioms used in Maskin's result are of particular interest from the perspective of axiomatic language. In contrast to anonymity, neutrality and the Pareto condition, IIA needs to refer to two preference profiles rather than just one. In fact, this property of the IIA condition has received criticism within social choice theory. Proponents of the *intra-profile approach* held that axioms such as independence should not make reference to preference profiles other than the profile under consideration. Put differently, an axiom should only express properties of an aggregation rule on a single preference profile, rather than express also how the aggregation rule would behave if the profile were somewhat different. The IIA axiom clearly violates the intra-profile approach, and proponents of this approach had originally hoped that by sticking to intra-profile axioms one could avoid Arrow's negative results. This hope, however, turned out to be in vain. Roberts (1980) contains more details on the intra- versus. inter-profile approach in social choice theory, and Rubinstein (1984) explores a link with mathematical logic. In the present context, it suffices to note that our purely syntactic perspective has points of contact with a methodological discussion within social choice theory. Hence, focusing on axiomatic language can reveal properties of axioms which have an interest to the social choice theorist.

Besides the IIA axiom, the maximal transitivity axiom also involves more complex quantification. For IIA, the added complexity derived from an additional quantification over alternative preference profiles. In the case of the maximal transitivity axiom, however, it is the nature of the quantification itself which is complex, quantification over a set of collective choice rules, i.e., higher order quantification.

Comparing May's and Maskin's characterizations of the majority rule, Maskin's result is more abstract since it involves at least one complex axiom which makes use of higher-order quantification. The proof of Maskin's result is far more complex than the proof of May's result, and this is due in part to the more abstract central axiom, maximal transitivity. And naturally, one could imagine yet more abstract characterization results which go even beyond quantifying over choice rules.

What we propose is to first decide on a particular conceptual and syntactic apparatus, and then investigate possible axiomatizations within that limited conceptual apparatus. For example, our conceptual apparatus may or may not allow for referring to multiple preference profiles like the IIA axiom does, and it may or may not allow for higher order quantification over collective choice rules of the kind exhibited by the maximal transitivity axiom. However, while we could try to simply introduce restrictions of this kind, there would be an element of arbitrariness in these restrictions. One of the benefits of introducing a restricted formal languages for expressing axioms will be that it can overcome this problem. Furthermore, it will make the conceptual basis of an axiomatization fully explicit, and it will also allow us to investigate what characterization results are possible within a certain formal conceptual framework. The following section will give a concrete example of what we have in mind.

3 Majority voting in judgement aggregation

In judgment aggregation, a collective choice rule aggregates sets of individual judgments into a set of group judgments. Individual judgments are simply modeled as formulas of propositional logic, and we assume that individuals make logically consistent judgments. Condorcet seems to have been the first to observe that in the context of preference aggregation, majority voting can lead to a group preference which fails to be transitive, and it was this *Condorcet paradox* which was the basis of much of the work in preference aggregation. Similarly, in judgment aggregation, it was observed (List and Pettit 2002) that aggregating sets of logically consistent individual judgments can lead to a set of group judgments which is logically inconsistent. This *doctrinal paradox* or *discursive dilemma* has been the basis of work on judgment aggregation. It has subsequently been shown that judgment aggregation is general enough to encompass preference aggregation (Dietrich and list 2005, unpublished). Here, we are interested to axiomatize majority voting in the context of judgment aggregation. The axiomatization result to be presented, Theorem 3, is proved in Pauly (2007), together with axiomatization results for some other voting procedures like unanimity and dictatorship.

Given a finite nonempty set of propositional atoms Φ_0, we define the set of *individual formulas* Φ_I as the set of all formulas α generated by the following grammar, where $p \in \Phi_0$:

$$\alpha := p \mid \neg\alpha \mid \alpha_1 \wedge \alpha_2$$

Hence, Φ_I is simply the language of propositional logic. An *individual valuation* is a function $v\colon \Phi_0 \to \{0, 1\}$, and we let V_I be the set of all individual valuations. In the standard way, we extend an individual valuation v to a function $\hat{v}\colon \Phi_I \to \{0, 1\}$ which assigns truth values also to complex propositions:

$$\begin{aligned} \hat{v}(p) &= v(p) \quad \text{for } p \in \Phi_0 \\ \hat{v}(\neg\alpha) &= 1 - \hat{v}(\alpha) \\ \hat{v}(\alpha \wedge \beta) &= min(\hat{v}(\alpha), \hat{v}(\beta)) \end{aligned}$$

In general, we shall usually identify v and $\hat{v}$, simply writing $v(\alpha)$ instead of $\hat{v}(\alpha)$. Also, we shall let $V_I(\alpha) = \{v \in V_I | v(\alpha) = 1\}$ denote all the individual valuations satisfying $\alpha \in \Phi_I$. Analogously for a set of formulas $\Sigma \subseteq \Phi_I$, we let $V_I(\Sigma) = \{v \in V_I | \forall \sigma \in \Sigma\colon\ v(\sigma) = 1\}$. A formula $\alpha \in \Phi_I$ is a *tautology* iff $v(\alpha) = 1$ for all individual valuations v.

In order to talk about collective judgments, we shall use the $[>]$-operator. The formula $[>]\alpha$ will refer to the majority judgment on formula α, i.e., it will be true whenever a strict majority of individuals accepts α. Since in principle we want to allow for arbitrary collective judgments to be made, boxed formulas will be treated like atoms which can be assigned arbitrary truth values. Formally, we define the set of *collective formulas* Φ_C as the set of all formulas φ generated by the following grammar, where $\alpha \in \Phi_I$:

$$\varphi := [>]\alpha \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2$$

For both individual and collective formulas, we use the standard abbreviations for the remaining connectives: $\top := p \vee \neg p$, $\bot := \neg\top$, $\varphi \vee \psi := \neg(\neg\varphi \wedge \neg\psi)$, $\varphi \to \psi := \neg\varphi \vee \psi$ and $\varphi \leftrightarrow \psi := (\varphi \to \psi) \wedge (\psi \to \varphi)$.

To give some examples of aggregation procedures expressible in this language, note that the formula $[>](p \wedge q)$ will refer to the outcome of a majority vote on $p \wedge q$. In contrast, $[>]p \wedge [>]q$ refers to two majority votes, a majority vote on p and a majority vote on q. The formula will be true in a situation where a majority accepts p and a (possibly different) majority accepts q. The fact that a majority accepting $p \wedge q$ implies a majority accepting p and also a majority accepting q, given that individuals are assumed to be logically consistent, is expressed by the formula $[>](p \wedge q) \to ([>]p \wedge [>]q)$. Note that the converse implication does not hold. In general, we are interested in characterizing the complete logic of majority, i.e., all valid logical principles involving majority voting over arbitrary logical propositions.

We shall now proceed to formally define the semantics of these collective formulas. Let $\Phi_{[>]I} = \{[>]\alpha | \alpha \in \Phi_I\}$. Then we can also define a *collective valuation* or *model* as a function $v\colon \Phi_{[>]I} \to \{0, 1\}$, and we can analogously extend a collective valuation to a function $\hat{v}\colon \Phi_C \to \{0, 1\}$. Again, we shall usually identify v and $\hat{v}$. If $v(\varphi) = 1$, i.e., the collective accepts φ, we shall say that the collective valuation v *satisfies* φ, also to be denoted as $v \models \varphi$. We shall let V_C denote the set of all collective valuations, and analogous to individual formulas, we let $V_C(\varphi) = \{v \in V_C | v(\varphi) = 1\}$ for $\varphi \in \Phi_C$ and $V_C(\Sigma) = \{v \in V_C | \forall \sigma \in \Sigma\colon\ v(\sigma) = 1\}$ for $\Sigma \subseteq \Phi_C$. The notion of logical consequence is defined in the standard way: φ is a logical consequence of Γ, denoted as $\Gamma \models \varphi$, provided $V_C(\Gamma) \subseteq V_C(\varphi)$. We say that a set of collective formulas $\Delta \subseteq \Phi_C$ *axiomatizes* a class of models $\mathcal{C} \subseteq V_C$ iff $\mathcal{C} = V_C(\Delta)$.

We are interested in collective valuations which arise by means of majority voting. This is the class $\mathcal{MAJ}$ of collective formulas $v \in V_C$ for which there is some number n of voters with individual valuations $v_1, \ldots, v_n \in V_I$ such that for all $\alpha \in \Phi_I$, $v([>]\alpha) = 1$ if and only if $|\{i \leq n | v_i(\alpha) = 1\}| > n/2$. Thus, $\mathcal{MAJ}$ consists of all those collective valuations, or sets of collective judgments, which can be obtained based on majority voting for some number of individuals, i.e., based on majority voting applied to some number of logically consistent individual sets of judgments.

Our task of axiomatizing majority voting can now be expressed as the task of finding a set $\Delta \subseteq \Phi_C$ which axiomatizes $\mathcal{MAJ}$.

For the purposes of formulating our axiom system, it will be useful to define a number of abbreviations referring to non-strict majority, ties, etc.. We define $[=]\varphi$ as $\neg[>]\varphi \wedge \neg[>]\neg\varphi$, $[\geq]\varphi$ as $[>]\varphi \vee [=]\varphi$ and $[\leq]\varphi$ as $\neg[>]\varphi$.

Let STEM denote the set of collective formulas containing all instances of the following four axiom schemes:

E. $[>]\,\alpha \leftrightarrow [>]\beta$, where $\alpha \leftrightarrow \beta \in \Phi_I$ is a tautology
M. $[>]\,(\alpha \wedge \beta) \rightarrow ([>]\alpha \wedge [>]\beta)$
S. $[>]\,\alpha \rightarrow \neg[>]\neg\alpha$
T. $([\geq]\alpha_1 \wedge \ldots \wedge [\geq]\alpha_k \wedge [\leq]\beta_1 \wedge \ldots \wedge [\leq]\beta_k) \rightarrow \bigwedge_{1\leq i\leq k}([=]\alpha_i \wedge [=]\beta_i)$ where $\forall v \in V_I$: $|\{i\colon v(\alpha_i) = 1\}| = |\{i\colon v(\beta_i) = 1\}|$

Axiom E expresses that group judgments will not distinguish propositions which are logically equivalent. In other words, it is only semantic difference which matters, not syntactic difference. Axiom M is a monotonicity axiom. Taken together with the equivalence Axiom E, it implies that if a majority accepts a proposition, it will also accept any logically weaker proposition. Axiom S expresses that we are dealing with strict majority, i.e., it is impossible to have a strict majority for α and at the same time a strict majority for $\neg\alpha$. For this would mean that at least one voter would have to accept both α and its negation, which would violate our assumption that voters are logically consistent.

Finally, the rather complex axiom T expresses a property known in the theory of simple games as *trade-robustness*. Suppose that the β_is have been obtained from the α_is by trade, i.e., each possible individual valuation v satisfies exactly as many α_is as β_is. If we identify an individual proposition with the set of valuations satisfying it, then this means that each valuation occurs as often among the α_is as it does among the β_is. Hence, we can say that the β_is have been obtained from the α_is by trading valuations. The axiom states, that if the β_is have been obtained from the α_is by trading, then weak majorities for the α_is and weak minorities for the β_is can only arise in the case where there are ties on all propositions. Intuitively, for majority vote, if we had a strict majority for some α_0 and a tie for all the other α_is and β_is, this would mean that taken together, there is more support for the α_is then for the β_is. But since the β_is have resulted by trade from the α_is, the total support for the α_is must be the same as that for the β_is.

Theorem 3 STEM *axiomatizes* $\mathcal{MAJ}$*, i.e.,* $V_C(\text{STEM}) = \mathcal{MAJ}$.

How does Theorem 3 compare to Theorems 1 and 2 as a characterization of majority voting? In each theorem we have a particular class of semantic objects we wish to characterize axiomatically. In the case of Theorems 1 and 2, this is a class containing a single choice rule, whereas in the case of Theorem 3 this is a class containing all the collective valuations that can arise based on majority voting. But in the case of Theorems 1 and 2, the axioms do not belong to a particular axiomatic language, whereas in Theorem 3, they do. Furthermore, this axiomatic language of Theorem 3 is minimal in the sense that it contains only a single operator to refer to collective judgments, the outcome of a collective vote. Hence, Theorem 3 is an example of the formal minimalism approach we have in mind as a methodological principle. In the next section, we shall develop this approach in more generality.

4 General methodology: formal minimalism

Having seen a number of different axiomatic characterization results, this is a good point to explain in more detail the methodology we have in mind for a formalist theory of social choice. We advocate an approach that proceeds in three steps.

4.1 Semantics

We start by mathematically defining two classes of semantic objects that form the topic of our investigation. First, there is the *semantic domain* $\mathcal{D}$, the class of objects that we consider. Second, there is the *target class* $\mathcal{T}$, the class of objects we are aiming to characterize axiomatically in some way, where $\mathcal{T} \subseteq \mathcal{D}$. Both domain and target class are supposed to model naturally some intuitive concept we have in mind. In social choice theory, the semantic domain will usually consist of aggregation functions of some type, and the target class will consist of majority voting, dictatorship, and so on. Other alternatives are possible, like the domain of collective valuations we have seen in Theorem 3. Note that while this semantic stage is of course standard in social choice theory, it is nonetheless worth mentioning explicitly, because already at this stage, there are choices to be made. The semantic domains involved in the three characterization theorems we have seen so far are all different, and part of the novelty of a characterization result can lie precisely in an interesting new semantic domain.

4.2 Syntax

The crucial point of departure from standard SCT practice is the definition of a formal language. Like in the case of Theorem 3, this involves defining the syntax and semantics of the formal language, where the semantics will be defined in terms of the semantic domain identified at the previous stage. More formally, we have a set of formulas Φ, where each formula $\varphi \in \Phi$ is made up of symbols of some formal language. Furthermore, we specify the intended interpretation of formulas in Φ by means of a *satisfaction relation* $\models \subseteq \mathcal{D} \times \Phi$, i.e., we specify when the relation $M \models \varphi$ holds, for $M \in \mathcal{D}$ and $\varphi \in \Phi$. In this case, we will say that M satisfies formula φ. This definition provides the semantics of our formal language in terms of the semantic objects or models in our domain.

The formal language will allow one to describe properties of the semantic objects of interest. It is by defining this language of description and its expressive power that one formulates the *target region* for a possible characterization result. Not just any axiomatization will do. Rather, it has to be an axiomatization within the limits of the formal language specified. Choosing the *right* formal language will always be part of the art of axiomatization. Nonetheless, it is clear that the language should not be so rich as to trivialize the axiomatization problem. More generally, we advocate choosing a minimal formal language, using as little expressive power as possible.

As an alternative to minimalism, we might also take a *realist* approach to the choice of formal language. We may want to use our axiomatization to draw inferences from some given information about results from various voting processes. We may know that a majority has votes for α, but that there was no majority for $\beta \vee \gamma$. If all the information available to us is about outcomes of majority votes, then a minimal language like the one used in Theorem 3 will be appropriate. On the other hand, we may also know the exact voting outcome, i.e., how many votes exactly there were for α. In that case, a richer language would be more appropriate. Finally, if we have a vote by call where we know who voted for and who voted against, we can access individual valuations and thus might be better of using a yet richer language. In this way, the choice of formal language can become application-driven.

4.3 Axiomatics

Given the semantic objects of interest and a formal language specifying the allowable axiomatizations, we can then look for axiomatic characterizations. Consider a set of formulas $\Delta \subseteq \Phi$. We will say that a model $M \in \mathcal{D}$ satisfies Δ, denoted as $M \models \Delta$ iff for every $\delta \in \Delta$ we have $M \models \delta$. Then we say that Δ *axiomatizes* $\mathcal{T}$ *absolutely* just in case for all $M \in \mathcal{D}$, we have $M \in \mathcal{T}$ iff $M \models \Delta$. Note that Theorem 3 provides an instance of absolute axiomatization, where $\Delta = \mathtt{STEM}$, $\mathcal{T} = \mathcal{MAJ}$ and $\mathcal{D} = V_C$. In fact, also the earlier Theorems 1 and 2 are instances of absolute axiomatization: For May's theorem, $\mathcal{D}$ is the set of group decision functions, $\mathcal{T}$ is the majority rule. The set of formulas Φ is simply the set of all formulas of set-theoretic mathematics, and Δ consists of the axioms for anonymity, neutrality and positive responsiveness. Hence, we see that the principal difference between the theorems of Sect. 2 and Theorem 3 lies in the set Φ of allowable formulas.

For a set of formulas $\Delta \subseteq \Phi$, let $Mod\ \Delta = \{M \in \mathcal{D} | M \models \Delta\}$ be the class of models of Δ. Then Δ strongly axiomatizes $\mathcal{T}$ iff $\mathcal{T} = Mod\ \Delta$. In the terminology of formal logic, absolute axiomatization really refers to the definability of a class of structures; in the realm of first-order logic it establishes that a class of structures is an *elementary class* (in the wider sense), see Enderton (1972) Hintikka (1996) also uses the term *descriptive completeness*.

Besides absolute axiomatization, there is also an alternative weaker form of axiomatization relative to the formal language under consideration. For $\mathcal{C} \subseteq \mathcal{D}$, we write $\mathcal{C} \models \varphi$ to say that for all $M \in \mathcal{C}$ we have $M \models \varphi$. Given a set of formulas $\Gamma \subseteq \Phi$ and a formula $\varphi \in \Phi$, we say that φ is a *consequence* of Γ, denoted as $\Gamma \models \varphi$, provided that for all $M \in \mathcal{D}$, if $M \models \Gamma$ then $M \models \varphi$. In other words, all models of Γ will satisfy φ. Now we can define the weaker relative notion of axiomatization as follows: For $\Delta \subseteq \Phi$, we say that Δ *axiomatizes* $\mathcal{T}$ *relatively*, or *relative to* Φ, if for all $\varphi \in \Phi$, we have $\mathcal{T} \models \varphi$ iff $\Delta \models \varphi$.

Again, the notion of relative axiomatization is well-known in formal logic. For a class of models $\mathcal{C}$, let $Th\ \mathcal{C} = \{\varphi \in \Phi | \mathcal{C} \models \varphi\}$ be the *theory* of $\mathcal{C}$, i.e., the set of all formulas satisfied by all models of $\mathcal{C}$. For $\Delta \subseteq \Phi$, let $Cn\ \Delta = \{\varphi \in \Phi | \Delta \models \varphi\}$ be the set of logical consequences of Δ. Then Δ axiomatizes $\mathcal{T}$ relatively iff $Cn\,\Delta = Th\,\mathcal{T}$. In the terminology of formal logic, while an absolute axiomatization of $\mathcal{T}$ refers to the definability of $\mathcal{T}$ itself, a relative axiomatization refers to the axiomatizability of the theory $\mathcal{T}$.

Note that both absolute and relative axiomatizations make use of a restricted formal language in that the axioms need to be expressions of Φ. It is in this sense that both kinds of axiomatizations fit into our formal minimalist approach. On the other hand, in contrast to absolute axiomatizations, relative axiomatizations are language dependent in the sense that they only characterize the target class $\mathcal{T}$ with respect to the formulas Φ of our formal language. In an absolute axiomatization, the set of formulas Δ can be used to separate models in $\mathcal{T}$ from those outside of $\mathcal{T}$. In a relative axiomatization, this separation only goes as far as the formulas of Φ: the consequences of Δ are precisely those formulas true in all models of $\mathcal{T}$ and nothing else. Consider the situation where we have $M, M' \in \mathcal{D}$, $M \neq M'$ and $\mathcal{T} = \{M\}$. It may be the case that for all $\varphi \in \Phi$, $M \models \varphi$ iff $M' \models \varphi$. This means that the difference between the two models cannot be expressed using the formulas of Φ, but can only be expressed in a richer language. While such a situation is inconsistent with an absolute axiomatization based on formulas $\Delta \subseteq \Phi$, it is perfectly well consistent with a relative axiomatization. The following result expresses that absolute axiomatizations are stronger than weak axiomatizations, a simple consequence of the fact that $Cn\,\Delta = Th\,Mod\,\Delta$.

Theorem 4 *Let $\Delta \subseteq \Phi$ be a set of formulas and $\mathcal{T} \subseteq \mathcal{D}$ be a target class. If Δ axiomatizes $\mathcal{T}$ absolutely, then it also axiomatizes $\mathcal{T}$ relatively.*

Note that with respect to absolute axiomatizations, using a weaker language, i.e., a smaller set of formulas Φ, will make the axiomatization problem more difficult. Hence, aiming to obtain an axiomatization in the weakest language means making the problem as difficult as possible. On the other hand, for relative axiomatizations, things are not as clear cut. Enlarging Φ will allow for more possible axioms, on the other hand, since establishing a relative axiomatization result involves showing a universal claim covering all formulas in Φ, the task becomes harder. Still, we maintain that also here, looking for a relative axiomatization with minimal Φ would be a natural result to aim for, already because it presents a natural limit case for progressively richer languages.

Getting back to the question of methodology, one can aim either for absolute or for relative axiomatizations. SCT has only been concerned with absolute axiomatizations, so there seems little need to argue for these. But relative axiomatizations can also be of value. For conceptual reasons or due to the applications one has in mind, there may be a particular formal language one wants to investigate, in order to find out how much the formal language can express about the semantic domain. If it turns out that the language does not suffice to axiomatize the target class absolutely, one can aim for axioms characterizing the target class within the expressive limitations of the underlying formal language.

If a target class $\mathcal{T}$ cannot be axiomatized absolutely given a certain formal language Φ, there are two possible responses: Either one turns to relative axiomatizations, or one tries to obtain an absolute axiomatization using a richer formal language, by enlarging Φ. Hence, one may interpret an axiomatization failure as an indication that one's formal language is too poor. As a possible response, one may enrich the language by adding new non-logical expressions to the language, or one might move to a stronger logic, e.g. from propositional to first-order or higher-order logic. At the other extreme, a formal language may also be too rich. If my formal language includes all of set-theoretic mathematics, it is in principle always possible to write down the definition of a social choice procedure as an axiom, and thereby achieve an absolute axiomatization trivially. Hence, the formal language needs to be restricted in some way in order to avoid such triviality results.

Having laid out our general methodology, we now return to the particular case of majority voting as an illustration. Theorems 1–3 have illustrated a number of different semantic domains, target classes and absolute axiomatizations. Only in the case of Theorem 3, however, was the formal language restricted, and hence it is only the latter result which fits into the formal minimalist approach we have outlined in this section. In the next section, we will return to majority voting in the framework of judgment aggregation, but we will change our semantic domain from collective valuations to aggregation functions. This will allow us to illustrate one of the advantages that comes with a formal minimalist approach: it allows us to prove non-axiomatizability results as well as results about relative axiomatizations.

5 Non-axiomatizability and relative axiomatizations

In the framework of judgment aggregation, Theorem 3 presents an axiomatization of majority voting within a semantic domain V_C that consists of collective valuations. From a SCT perspective, a more standard semantic domain to consider is the domain of aggregation functions, and it is this semantic domain we shall consider in this section.

A *judgment aggregation function* $A : (V_I)^N \to V_C$ maps n individual valuations into a collective valuation. Judgement aggregation functions have been considered in List and Pettit (2002) and subsequent papers, but in contrast to our work, most of the results concerning judgment aggregation functions assume that the aggregated judgment is logically consistent (see List for a web site covering judgment aggregation and the discursive dilemma). We will say that an n-ary judgment aggregation function A satisfies a formula $\varphi \in \Phi_C$, denoted as $A \vdash \varphi$, provided that for all $v_1, \ldots, v_n \in V_I$ we have $A(v_1, \ldots, v_n)(\varphi) = 1$. A satisfies a set of formulas Σ (denoted as $A \vdash \Sigma$) if $A \vdash \sigma$ for all $\sigma \in \Sigma$. Let $\mathcal{A}$ be the class of all judgment aggregation functions, of arbitrary finite arities. The aggregation functions we are interested in are, of course, those representing majority voting. For $n \geq 1$, define the aggregation function A^n_{maj} by $A^n_{maj}(v_1, \ldots, v_n)([>]\alpha) = 1$ iff $|\{i \leq n \mid v_i(\alpha) = 1\}| > n/2$, for all $\alpha \in \Phi_I$.

In terms of the methodological terminology introduced in the last section, our semantic domain is $\mathcal{A}$ and our target class simply $\mathcal{A}_{maj} = \{A^n_{maj} | n \geq 1\}$. Our formal language is still the set of collective formulas Φ_C, but given the new semantic domain, the semantics and the satisfaction relation are defined somewhat differently.

We will now relate the semantic domain of aggregation functions to the collective valuations we used as a semantic domain in Sect. 3. To begin with, we can relate the two target classes $\mathcal{A}_{maj}$ and $\mathcal{MAJ}$ by observing that $\mathcal{MAJ} = \{v \in V_C | \exists v_1, \ldots, v_n \in V_I : A^n_{maj}(v_1, \ldots, v_n) = v\}$. Next, we can relate more generally how an absolute axiomatization in the domain of aggregation functions relates to an absolute axiomatization in the domain of collective valuations. Given a class of aggregation functions $\mathcal{C} \subseteq \mathcal{A}$, let

$$\mathcal{C}\downarrow = \{v \in V_C | \exists A \in \mathcal{C}\ \exists v_1, \ldots, v_n \in V_I :\ A(v_1, \ldots, v_n) = v\}.$$

Hence, $\mathcal{C}\downarrow$ includes all the collective valuations that can arise based upon aggregation functions in $\mathcal{C}$ under some set of individual valuations. The relation expressed earlier between majority voting as an aggregation function and the set of collective valuations that can arise based on majority voting can now be expressed by the equation $\mathcal{A}_{maj}\downarrow = \mathcal{MAJ}$.

Note that while the satisfaction relation $\models$ in the domain of collective valuations was simply the standard satisfaction relation of propositional logic, the $\vdash$ satisfaction relation for aggregation functions is different, it is of higher order. Formally, the two satisfaction relations are related by the fact that $A \vdash \varphi$ iff $\{A\}\downarrow \models \varphi$. In terms of absolute axiomatizations, it turns out that axiomatizing aggregation functions is harder than axiomatizing collective valuations, as the following result shows.

Theorem 5

(1) *If $\Delta \subseteq \Phi_C$ absolutely axiomatizes a class of aggregation functions $\mathcal{C} \subseteq \mathcal{A}$, then Δ also absolutely axiomatizes $\mathcal{C}\downarrow$.*
(2) *If Δ absolutely axiomatizes $\mathcal{C}\downarrow$ but not $\mathcal{C}$, then $\mathcal{C}$ cannot be axiomatized absolutely.*
(3) *$\mathcal{A}_{maj}$ cannot be absolutely axiomatized.*

Proof (1) If $v \in \mathcal{C}\downarrow$, there is some $A \in \mathcal{C}$ and some $v_1, \ldots, v_n \in V_I$ such that $A(v_1, \ldots, v_n) = v$. Since Δ axiomatizes $\mathcal{C}$, $A \vdash \Delta$, and so in particular $A(v_1, \ldots, v_n)(\Delta) = v(\Delta) = 1$. Conversely, if $v(\Delta) = 1$, consider the aggregation function A defined as $A(v_1, \ldots, v_n) = v$ for all $v_1, \ldots, v_n \in V_I$. Then by definition, $A \vdash \Delta$, and so $A \in \mathcal{C}$ which suffices to show that $v \in \mathcal{C}\downarrow$.

(2) Suppose that Δ absolutely axiomatizes $\mathcal{C}\downarrow$ but not $\mathcal{C}$, and suppose that Γ absolutely axiomatizes $\mathcal{C}$. Then

$$\begin{aligned} v(\Delta) = 1 &\Leftrightarrow \exists A \in \mathcal{C}\ \exists v_1, \ldots, v_n \in V_I : A(v_1, \ldots, v_n) = v \\ &\Leftrightarrow \exists A \vdash \Gamma\ \exists v_1, \ldots, v_n \in V_I : A(v_1, \ldots, v_n) = v \\ &\Leftrightarrow v(\Gamma) = 1 \end{aligned}$$

and hence Γ and Δ must be logically equivalent, i.e., for all $A \in \mathcal{A}$, we have $A \vdash \Gamma$ iff $A \vdash \Delta$. As a consequence, Δ also absolutely axiomatizes $\mathcal{C}$, a contradiction.

(3) It suffices to show that the absolute axiomatization already obtained in theorem 3 for $\mathcal{MAJ}$ does not axiomatize $\mathcal{A}_{maj}$. To see that STEM does not axiomatize $\mathcal{A}_{maj}$, consider the ternary aggregation function A_d where for all $v_1, v_2, v_3 \in V_I$ and $\alpha \in \Phi_I$, $A_d(v_1, v_2, v_3)([>]\alpha) = v_1(\alpha)$. Hence, A_d is a dictatorship. Now consider $\{A_d\}\downarrow$ which comprises all collective valuations that can arise based on this dictatorship. Note that $\{A_d\}\downarrow \subseteq \mathcal{MAJ}$, for any valuation that has arisen

based on a dictator can also have arisen based on a majority having that same valuation. Hence, by Theorem 3, $A_d \vdash \mathtt{STEM}$ but $A_d \notin \mathcal{A}_{maj}$. □

The result not only shows that an axiomatization of aggregation functions is stronger than an axiomatization of collective valuations, it also shows that we cannot hope to generalize the absolute axiomatization result obtained for majority voting from collective valuations to aggregation functions. In order to absolutely axiomatize majority voting as an aggregation function, we need a richer language. As mentioned, we cannot distinguish a situation of a dictatorship where the dictator has valuation v_i and the majority has a different valuation $v_j \neq v_i$ from a situation of majority voting where the majority shares the same valuation v_i. Since our logical language can only express properties of the collective decision, the difference between these two situations cannot be captured in our language. This example also allows us to further pinpoint the difference between collective valuations and aggregation functions in terms of the different semantics employed for the same set of expressions Φ_C: if a collective valuation v satisfies the axioms STEM, then it suffices that there are *some* individual valuations which produce the collective valuation under majority voting. On the other hand, any judgment aggregation function satisfying STEM has to only yield collective valuations satisfying STEM under *any* individual valuations, a much stronger requirement.

The connection between axiomatizations of aggregation functions and collective valuations presented in Theorem 5 refers to absolute axiomatizations, and hence there remains the question regarding an analogous connection for the case of relative axiomatization. The following result shows that with respect to relative axiomatization, there is no difference between aggregation functions and collective valuations. As a consequence, STEM does axiomatize $\mathcal{A}_{maj}$ relatively.

Theorem 6 $\Delta \subseteq \Phi_C$ *relatively axiomatizes a class of aggregation functions* $\mathcal{C} \subseteq \mathcal{A}$ *if and only if* Δ *relatively axiomatizes* $\mathcal{C}{\downarrow}$. *As a consequence,* STEM *relatively axiomatizes* $\mathcal{A}_{maj}$.

Proof It suffices to verify two claims. First, for every $\varphi \in \Phi_C$, $\mathcal{C} \vdash \varphi$ iff $\mathcal{C}{\downarrow} \models \varphi$, so $Th\,\mathcal{C} = Th\,\mathcal{C}{\downarrow}$. Second, for every $\Delta \subseteq \Phi_C$ and $\varphi \in \Phi_C$, we have $\Delta \models \varphi$ iff $\Delta \vdash \varphi$, so we can simply speak of $Cn\,\Delta$, without any need to make the underlying semantics explicit. From these two claims, the first part of the result follows. As a consequence, since $\mathcal{A}_{maj}{\downarrow} = \mathcal{MAJ}$, by Theorem 3, STEM relatively axiomatizes $\mathcal{A}_{maj}$. □

Theorems 5 and 6 serve both a substantive and a methodological purpose. On the one hand, they relate axiomatizations of two different semantic domains, be they absolute or relative. For absolute axiomatizations, the domain of aggregation functions is more challenging than the domain of collective valuations. In fact, with respect to majority voting, no axiomatization is possible in the domain of aggregation functions. Hence, either one needs to enrich the formal language available for formulating axioms, or one needs to turn to relative axiomatizations. For relative axiomatizations, the same axiomatization of majority voting works for both collective valuations and aggregation functions, in fact, in general, the axiomatization problem is the same in both domains.

On the other hand, Theorems 5 and 6 provide examples of the advantages of a methodology based on formal minimalism. They demonstrate that it becomes possible to prove results about non-axiomatizability, due to the fact that we have only a limited formal axiomatic language available. Furthermore, they also show that even in the case of non-axiomatizability with respect to absolute axiomatization, one may still obtain positive results for the weaker notion of relative axiomatization, showing that within the limits of the given formal language, we have succeeded in capturing the target class.

6 Conclusions

Social choice theory makes heavy use of the axiomatic method in obtaining its key results. Given the central importance of axiomatization results, it is somewhat surprising that these results have so far not been formalized in the way that mathematics has been. This paper has argued for taking a more formal approach to axiomatization in social choice theory.

To summarize, we have argued for evaluating axiomatizations according to the language used for expressing the axioms. Already applying this criterion on an informal level in Sect. 2, we saw that some observations concerning the quantificational structure of the axioms linked up with substantive discussions in SCT regarding what kinds of axioms to permit, intra- vs. inter-profile. But in order to fully apply this criterion, we need a precise definition of the language used to express axioms. This gave rise to our methodology of formal minimalism, the idea to use a very limited formal axiomatic language, with the aim to axiomatize a particular class of social choice procedures using only axioms from this limited language. Besides specifying the syntax of the formal language, we also need to define its semantics in terms of the social choice procedures of our domain. The richer the formal language, the easier the task of axiomatization becomes—hence, the desire for minimality. We also pointed out, however, that one may have reasons for picking a *particular* formal language, either because one has an interest in simply exploring how far certain concepts go in expressing interesting properties of social procedures, or because the application domain makes certain notions readily applicable and observable.

It may also be useful to summarize the advantages of the formal approach advocated. First, the use of a formal language clearly lays out the allowable tools in our axiomatization effort. Hence, the axiomatization problem becomes more clearly defined. Second, we obtain a partial order on axiomatizations. Since we can (partially) order axiomatic languages by expressivity, we obtain a map of different axiomatization results which is much more structured than without having a formal language. Third, we are able to prove non-axiomatizability results, showing that a particular axiomatic language is too weak to axiomatize a particular target class of social choice procedures. Fourth, we are able to distinguish different kinds of axiomatization. Absolute axiomatizations are the strongest kind, where we find axioms which suffice to isolate the target class over the chosen semantic domain. Relative axiomatizations, on the other hand, are weaker in that they only characterize the target class relative to the underlying axiomatic language. This distinction only becomes meaningful because our axiomatic language is formal. Both kinds of axiomatizations are language dependent in the sense that they make use of a particular formal language for expressing axioms. But relative axiomatization also is language dependent in the sense of aiming for an axiomatic characterization only within the bounds of the expressive power of the formal language.

Questions of axiomatization are the natural domain of the field of formal logic. Much of the work in logic can be described as *metamathematics*, the investigation of different mathematical theories and their properties. In order to do metamathematics, one needs a formal language in which mathematics is formulated. The metamathematical perspective has obtained a wealth of insights concerning questions such as axiomatizability, decidablity, complexity, etc. due to the seminal work of Gödel, Turing, and others. This investigation has shown, e.g., that the standard model of arithmetic is not axiomatizable (neither absolutely nor relatively) and that there is no computer algorithm which can decide for every proposition of arithmetic whether it is true or not. Once we have taken the step to formal languages, similar meta-questions can be asked in the domain of social choice. Is there an algorithm which decides for every formula whether it follows from the `STEM` axioms? How complex would such an algorithm be? These and other questions are typical questions of metamathematics. This paper has focused on the issue of axiomatization, since it is the topic in metamathematics which provides the most natural link with much of the work in social choice theory. For the future, the aim would be to pursue this agenda further, investigating the metamathematics of social choice.

References

Arrow, K. (1951). *Social choice and individual values*. Yale University Press.

Campbell, D., & Kelly, J. (2000). A simple characterization of majority rule. *Economic Theory, 15*, 689–700.

Enderton, H. (1972). *A mathematical introduction to logic*. Academic Press.

Gibbard, A. (1973). Manipulation of voting schemes: A general result. *Econometrica, 41*, 587–601.

Hintikka, J. (1996). *The principles of mathematics revisited.* Cambridge University Press.

List, C. http://personal.lse.ac.uk/LIST/doctrinalparadox.htm.

List, C., & Pettit, P. (2002). Aggregating sets of judgments: An impossibility result. *Economics and Philosophy, 18*, 89–110.

Maskin, E. (1995). Majority rule, social welfare functions, and game forms. In K. Basu, P. Pattanaik, & K. Suzumura (Eds.), *Choice, welfare, and development* (pp. 100–109). Oxford University Press.

May, K. O. (1952). A set of independent necessary and sufficient conditions for simple majority decision. *Econometrica, 20*, 680–684.

Pauly, M. (2007). Axiomatizing collective judgment sets in a minimal logical language. *Synthese*, *158*(2), 233–250.

Roberts, K. (1980). Social choice theory: The single-profile and multi-profile approaches. *Review of Economic Studies, 47*, 441–450.

Rubinstein, A. (1984). The single profile analogues to multi profile theorems: Mathematical logic's approach. *International Economic Review, 25*, 719–730.

Satterthwaite, M. (1975). Strategy-proofness and arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory, 10*, 187–217.