Dependence and Independence in Social Choice: Arrow's Theorem

Eric Pacuit and Fan Yang

Abstract One of the goals of social choice theory is to study the group decision methods that satisfy two types of desiderata. The first type ensures that the group decision depends in the right way on the voters' opinions. The second type ensures that the voters are free to express any opinion, as long as it is an admissible input to the group decision method. Impossibility theorems, such as Arrow's Theorem, point to an interesting tension between these two desiderata. In this paper, we argue that dependence and independence logic offer an interesting new perspective on this aspect of social choice theory. To that end, we develop a version of independence logic that can express Arrow's properties of preference aggregation functions. We then prove that Arrow's Theorem is derivable in a natural deduction system for the first-order consequences of our logic.

1 Introduction

The modern era in social choice theory started with Kenneth Arrow's ground-breaking *impossibility theorem* [3]. Arrow showed that there is no method that a group can use to rank a set of alternatives satisfying a minimal set of desirable properties. Much has been written about this theorem (see, for instance, [20, 30, 43]) and its implications for theories of democracy [9, 27, 36] and beyond [29, 31, 33, 44]. Social choice theory has since grown into a large and multi-faceted research area (see [26] for an overview). In this chapter, we focus on one type of theorem studied by social choice theorists: *axiomatic characterizations* of group decision methods. We will present a version of independence logic [14] that we use to formalize these theorems. This is not merely an exercise in applying a logical framework to a new area. We will argue that dependence and independence logic offers an interesting new perspective on the axiomatic characterization of group decision methods.

E. Pacuit (✉)
Department of Philosophy, University of Maryland, College Park, MD, USA
e-mail: epacuit@umd.edu

F. Yang
Department of Philosophy and Religious Studies, Utrecht University, Janskerkhof 13, 3512 BL Utrecht, The Netherlands
e-mail: fan.yang.c@gmail.com

S. Abramsky et al. (eds.), *Dependence Logic*, DOI 10.1007/978-3-319-31803-5_11

One of the main goals of social choice theory is to identify principles of group decision making that ensure that group decisions depend *in the right way* on the voters' preferences.[1] That is, group decision methods should be designed in such a way that no individual voter should have any undue influence over the group decision. At the same time, it is important to devise a group decision method without placing any restrictions on the inputs. That is, group decision methods should be designed under the assumption that the voters' opinions are *independent*. From this perspective, the so-called impossibility theorems in social choice theory highlight an interesting conflict between dependence and independence.

This chapter is organized as follows. Section 2 briefly recounts the basic mathematical framework used in social choice theory. Sections 3 and 4 are extended discussions of the notions of dependence and independence found in social choice theory. In Section 5, we present a version of independence logic that we use to formalize Arrow's Theorem. We show in Section 5.3 that Arrow's Theorem can be derived in a natural deduction system for this logic [16, 25]. Section 6 contains some concluding remarks.

2 The Social Choice Framework

Let us start by recalling some notions concerning relations.

Definition 1. A *relation* R on X is a subset of $X \times X$. We write $a\ R\ b$ when $(a, b) \in R$ and write $a \not R\ b$ when $(a, b) \notin R$. We write $a\ R\ b\ R\ c$ when $a\ R\ b$, $b\ R\ c$, and $a\ R\ c$.

Definition 2. Let $R \subseteq X \times X$ be a relation. R is said to be

- *transitive* provided for all $a, b, c \in X$, if $a\ R\ b$ and $b\ R\ c$, then $a\ R\ c$;
- *complete* provided for all $a, b \in X$, either $a\ R\ b$ or $b\ R\ a$;
- *antisymmetric* provided for all $a, b \in X$, if $a\ R\ b$ and $b\ R\ a$, then $a = b$;
- *linear* provided that it is transitive, complete, and antisymmetric.

Throughout this chapter, we fix a set $V = \{x_1, x_2, x_3, \ldots, x_n\}$ of n *voters* (or *individuals*) and a (finite[2]) set X of *alternatives* (e.g., candidates, restaurants, social states, etc.). Each voter $x_i \in V$ is asked to rank the elements of X, where a **ranking** is a transitive and complete relation on X. Let $O(X)$ denote the set of all rankings of X. Each ranking $R \in O(X)$ is associated with two special subrelations: The **strict subrelation** defined as

$$P_R = \{(a, b) \in X \times X \mid a\ R\ b \text{ and } b \not R\ a\},$$

[1]In formal work on social choice theory, it is common to identify a voter's *preference* over a set of alternatives X with her *ranking* over the set of alternatives. In general, a ranking of the alternatives is only one way in which a voter may express her preference over the set of alternatives. Consult [17] for a discussion of the main philosophical issues here.

[2]For simplicity, we restrict attention to a finite set of alternatives. This restriction is not necessary for what follows, though it does have some implications on the design of the formal language used to describe a social choice model.

and the **indifference subrelation** defined as

$$I_R = \{(a, b) \in X \times X \mid a\ R\ b \text{ and } b\ R\ a\}.$$

Let $R \in O(X)$ represents voter x_i's ranking of X. If $a\ R\ b$, then we say that "x_i weakly prefers a to b"; if $a\ P_R\ b$, then we say that "x_i strictly prefers a to b"; if $a\ I_R\ b$, then we say that "x_i is indifferent between a and b." In case a voter is indifferent between two alternatives, we say that the two alternatives are *tied* in the voter's ranking. A **linear ranking** is a ranking that is a linear relation and let $L(X)$ denote the set of all linear rankings of X. Clearly, ties are not allowed in linear rankings.

A **profile** for the set of voters V is a sequence of (linear) rankings of X that assigns to each voter x_i a ranking R_i, denoted $\mathbf{R} = (R_1, \ldots, R_n)$. The set of all profiles of rankings for n voters is denoted $O(X)^n$ (similarly for $L(X)^n$). If $\mathbf{R} = (R_1, \ldots, R_n) \in O(X)^n$, then P_i denotes the strict subrelation P_{R_i} of R_i. Similarly, I_i denotes the indifference subrelation I_{R_i} of R_i. For a profile $\mathbf{R} = (R_1, \ldots, R_n) \in O(X)^n$, let $\mathbf{V}_{\mathbf{R}}(a\ P\ b) = \{x_i \in V \mid a\ P_i\ b\}$ be the set of voters that rank a strictly above b (similarly for $\mathbf{V}_{\mathbf{R}}(a\ I\ b)$ and $\mathbf{V}_{\mathbf{R}}(a\ R\ b)$).

A group decision method associates an *outcome* ("the group decision") with each profile of *ballots*. Typically, the ballots are (linear) rankings of the alternatives. The social choice literature has largely focused on two types of outcomes. The first are (linear) rankings of X representing the overall group ranking of the alternatives. The second are non-empty subsets of X representing the "social choice" (or the "winning" alternatives). As we mentioned in Section 1, the starting assumption in the social choice literature is that group decisions should be *completely determined* by the voters' reported[3] rankings. This means that group decision methods should be represented by *functions* from sets of profiles to the possible outcomes. There are two types of functions corresponding to the two types of group outcomes:

- A **preference aggregation function** is a function $F : \mathscr{B} \to \mathscr{O}$, where $\mathscr{O}$ is a set of relations on X (typically, $\mathscr{O}$ is either $L(X)$ or $O(X)$) and $\mathscr{B} \subseteq L(X)^n$ or $\mathscr{B} \subseteq O(X)^n$.
- A **social choice function** is a function $F : \mathscr{B} \to \mathscr{O}$ where $\mathscr{O} = \wp(X) \setminus \{\emptyset\}$ and $\mathscr{B} \subseteq L(X)^n$ or $\mathscr{B} \subseteq O(X)^n$.

We illustrate the above definitions with the following examples. For each profile of linear rankings $\mathbf{P} = (P_1, \ldots, P_n)$, define a relation $R^{\mathbf{P}}_{maj}$ on the set X of candidates as follows: for all $a, b \in X$,

$$a\ R^{\mathbf{P}}_{maj}\ b \text{ iff } |\mathbf{V}_{\mathbf{P}}(a\ P\ b)| > |\mathbf{V}_{\mathbf{P}}(b\ P\ a)|.$$

The relation $R^{\mathbf{P}}_{maj}$ is called the *majority ordering*, ranking candidate a above candidate b provided more voters rank a above b than b above a. Note that

[3] In this article, we set aside any game-theoretic issues around whether voters have an incentive to report their *true* preferences.

$R^{\mathbf{P}}_{maj} \notin O(X)$, since $R^{\mathbf{P}}_{maj}$ is not necessarily transitive.[4] This is illustrated by the famous *Condorcet Paradox*: Consider a profile $\mathbf{P} = (P_1, P_2, P_3)$ for three voters $V = \{x_1, x_2, x_3\}$ and three candidates $X = \{a, b, c\}$. Suppose that ranking for voter x_1 is $a\ P_1\ b\ P_1\ c$; the ranking for voter x_2 is $b\ P_2\ c\ P_2\ a$; and the ranking for voter x_3 is $c\ P_3\ a\ P_3\ b$. Then, $\mathbf{V_P}(a\ P_1\ b) = \{x_1, x_3\}$, $\mathbf{V_P}(b\ P_1\ c) = \{x_1, x_2\}$, and $\mathbf{V_P}(a\ P_1\ c) = \{x_1\}$. Thus, $a\ R^{\mathbf{P}}_{maj}\ b$ and $b\ R^{\mathbf{P}}_{maj}\ c$, but it is not the case that $a\ R^{\mathbf{P}}_{maj}\ c$ (in fact, we have $c\ R^{\mathbf{P}}_{maj}\ a$ producing a cycle: $a\ R^{\mathbf{P}}_{maj}\ b\ R^{\mathbf{P}}_{maj}\ c\ R^{\mathbf{P}}_{maj}\ a$).

The **Borda ranking** is an example of a social ranking that is transitive. Let $\mathbf{P} = (P_1, \ldots, P_n)$ be an arbitrary profile of linear rankings of a k-element set X of candidates. For each voter x_i and each $m = 1, \ldots, k$, let $P_i(m)$ be the candidate ranked in the mth-position by the voter x_i. For example, if $a\ P_i\ b\ P_i\ c$, then $P_i(1) = a$, $P_i(2) = b$, and $P_i(3) = c$. For each $d \in X$, define the *Borda score* $BS(\mathbf{P}, d)$ as

$$BS(\mathbf{P}, d) = \sum_{m=1}^{k} (k - m) \cdot |\{x_i \in V \mid P_i(m) = d\}|.$$

Now, the **Borda ranking** $R^{\mathbf{P}}_b$ is defined as follows:

$$a\ R^{\mathbf{P}}_B\ b \ \text{ iff } \ BS(\mathbf{P}, a) \geq BS(\mathbf{P}, b).$$

The function $F_B : L(X)^n \to O(X)$, defined as $F_B(\mathbf{P}) = R^{\mathbf{P}}_B$, is a preference aggregation function. An example of a social choice function is plurality rule: $F_{pl} : L(X)^n \to \wp(X) \setminus \{\emptyset\}$, defined as: for each $\mathbf{P} = (P_1, \ldots, P_n) \in L(X)^n$,

$$F_{pl}(\mathbf{P}) = \{c \in X : |\{x_i \mid P_i(1) = c\}| \geq |\{x_i \mid P_i(1) = b\}| \text{ for all } b \in X\}.$$

Thus, F_{pl} selects the candidate(s) ranked first by the most voters. Consult [34] for an overview of preference aggregation and social choice rules and their properties.

3 Dependence in Social Choice Theory

Generally speaking, axiomatic characterization results proceed in two steps. The first step is to identify an interesting class of functions, each of which is intended to represent a possible group decision method. Different classes of functions build in different assumptions about the *structure* of the group decision problem. For instance, fix a set X of at least three candidates and a set $V = \{x_1, \ldots, x_n\}$ of n

[4] Also, $R^{\mathbf{P}}_{maj}$ may not be complete if there is an even number of voters. There are a variety of ways to modify the definition of the majority ordering to ensure completeness when there are an even number of voters.

voters. Then, the set

$$\mathfrak{F}(n, X) = \{F \mid F : O(X)^n \to O(X)\}$$

represents group decision problems in which the n voters are asked to rank the alternatives, and, based on the voters' rankings, identify the group ranking of the alternatives. Furthermore, since the domain of each function $F \in \mathfrak{F}(n, X)$, denoted $dom(F)$, is the set $O(X)^n$ of *all* profiles of rankings (i.e., all functions are assumed to be *total*), the voters' input is not restricted in any way. In particular, a voter's choice of ranking is *independent* of the other voters' choice of rankings (this will be discussed in more detail in Section 4). The second step is to characterize the desired set of group decision methods in terms of principles expressible as properties of the given class of functions. The goal is to find a set of properties of group decision rules that makes the group decision depend *in the right way* on the voters' inputs. Many different principles of group decision making have been discussed in the social choice literature. We discuss some of these properties in this section (see [20, 26, 30, 34] for discussions of additional properties). The statement of these properties will be tailored to the class $\mathfrak{F}(n, X)$ of preference aggregation functions. We leave it to the reader to adapt the principles to different classes of preference aggregation functions.

Since group decision methods are assumed to be *functions*, the output (a group ranking) does *functionally depend* on the rankings of the voters. However, this dependence is much too weak. There are many functions in $\mathfrak{F}(n, X)$ that are defective in some way. For instance, the constant function $F_R : O(X)^n \to O(X)$ for a fixed $R \in O(X)$, defined as $F_R(\mathbf{R}) = R$, is in $\mathfrak{F}(n, X)$. An obvious problem with a constant function is that the group decision is insensitive to any unanimous agreement among the voters. Suppose that $X = \{a, b, c\}$ and $a\ R\ b\ R\ c$, and consider the constant function F_R. If $\mathbf{R}$ is a profile in which every voter ranks b strictly above a (i.e., $\mathbf{V_R}(b\ P\ a) = V$), then $F_R(\mathbf{R}) = R$ is an outcome that does not truly reflect the voters' opinions (at least with respect to alternatives a and b). This suggests the following property:

Unanimity For all alternatives $a, b \in X$, for all profiles $\mathbf{R} = (R_1, \ldots, R_n) \in O(X)^n$, if $a\ P_i\ b$ for all $x_i \in V$ (i.e., $\mathbf{V_R}(a\ P\ b) = V$), then $a\ P_{F(\mathbf{R})}\ b$.

This principle ensures that any unanimous agreement among the voters' strict rankings is reflected in the group ranking.[5] Unanimity is a fundamental principle of group decision methods.

An important distinction that was prevalent early on in the burgeoning social choice literature is between *single-profile* and *multi-profile* properties [21, 35, 37, 38]. Unanimity is an example of a single-profile property. It rules out specious pairings of group rankings with profiles in terms of properties of

[5]One can also explore alternative definitions of *Unanimity* of varying strengths. For example, if all voters weakly rank candidate a above candidate b, then society does so as well.

the given profile (e.g., the property that all voters rank candidate a strictly above candidate b). Thus, as will become clear in Section 5, Unanimity can be formalized in dependence logic using just a first-order formula without *dependency atoms*.

The notion of dependency found in dependence and independence logic is best exemplified by *multi-profile* properties. The general form of a multi-profile property runs as follows: If (two or more) profiles are related in a certain way, then the outcomes associated with these profiles must be related in some way. The most prominent multi-profile property is *independence of irrelevant alternatives* (IIA).

Independence of Irrelevant Alternatives For all profiles, $\mathbf{R}, \mathbf{R}' \in O(X)^n$, for all $a, b \in X$, if $\mathbf{V}_{\mathbf{R}}(a\ R\ b) = \mathbf{V}_{\mathbf{R}'}(a\ R\ b)$ and $\mathbf{V}_{\mathbf{R}}(b\ R\ a) = \mathbf{V}_{\mathbf{R}'}(b\ R\ a)$, then $a\ F(\mathbf{R})\ b$ iff $a\ F(\mathbf{R}')\ b$.

IIA ensures that the group ranking of two candidates depends only on how the individual voters rank those candidates. That is, if all voters agree on the relative rankings of a and b in two profiles, then the group rankings under each profile must rank a and b in the same way. IIA plays a crucial role in Arrow's impossibility theorem and many other results in social choice theory. Intuitions vary about the reasonableness of the IIA requirement for group decision methods. Many well-known voting methods do *not* satisfy IIA (the most prominent such example is the *Borda's ranking* that we defined in Section 2, see [39, 40] for an extensive argument in favor of using the Borda score to make group decisions, and see [34] for a discussion and further examples). Nonetheless, there are persuasive arguments that IIA is a natural requirement for a group decision method (see, for instance, [51, pg. 58] and [3, Chapter III, section 3]). Furthermore, Muller and Satterthwaite [32, 41] showed that IIA is equivalent to *strategy proofness* (strategy proofness means that voters do not have an incentive to misrepresent their preferences[6]).

Various authors have explored the implications of weakening IIA. An important result along these lines is from Blau [4]. Given a profile $\mathbf{R} \in O(X)^n$ and $Y \subseteq X$, let $\mathbf{R}_Y = ((R_1)_Y, \dots, (R_n)_Y)$, where each $(R_i)_Y = R_i \cap (Y \times Y)$, i.e., $(R_i)_Y$ is the restriction of R_i to Y. Then, IIA can be reformulated as follows:

Binary Independence For all profiles, $\mathbf{R}, \mathbf{R}' \in O(X)^n$, for all $a, b \in X$, if $\mathbf{R}_{\{a,b\}} = \mathbf{R}'_{\{a,b\}}$, then $F(\mathbf{R})_{\{a,b\}} = F(\mathbf{R}')_{\{a,b\}}$.

Blau studied the following generalization of IIA:

***m*-ary Independence** For all profiles, $\mathbf{R}, \mathbf{R}' \in O(X)^n$, for all m-element sets $Y \subseteq X$, if $\mathbf{R}_Y = \mathbf{R}'_Y$, then $F(\mathbf{R})_Y = F(\mathbf{R}')_Y$.

Of course, if $m = |X|$, then m-ary independence simply amounts to the usual requirement for any function. Blau showed that if a preference aggregation function satisfies m-ary independence (where $2 \leq m < |X|$), then it must also satisfy binary independence (the converse is obvious). An alternative way to define m-ary

[6] A full discussion of this result is beyond the scope of this article. See [46] for a precise statement of the Müller-Satterthwaite Theorem (including the additional assumptions needed to prove the equivalence) and a discussion of the relevant literature.

independence is as follows: Let $\mathscr{S}_m = \{Y \subseteq X \mid |Y| = m\}$ be the collection of all sets of m candidates from X. Then, m-ary independence says that for any $S \in \mathscr{S}_m$, if the relative rankings of all candidates in S are the same in two profiles, then the outcomes associated with these profiles must agree on the rankings of all candidates in S. Recently, Susumo Cato [7] showed that Blau's results hold for any collection $\mathscr{S}$ of sets of candidates satisfying the following connectedness property: For all candidates $a, b \in X$ there are sets $S_1, \ldots, S_k \in \mathscr{S}$ such that $\{a, b\} = \bigcap_{i=1}^{k} S_i$.

Preference aggregation rules that do not satisfy IIA take a global perspective when determining the social ranking of the candidates. For example, the Borda ranking of candidates a and b depends on the voters' rankings of *all* the candidates. A weaker version of IIA requires that the group ranking of candidates a and b depends on the voters' rankings of *some* subset of candidates (which may contain more candidates than just a and b). Campbell and Kelly studied preference aggregation methods that satisfy this weaker version of IIA together with additional multi-profile properties [6, 20].

A second multi-profile property, Neutrality, requires that the aggregation method treats all the candidates equally. To state this formally, suppose that $\mu : X \to X$ is a permutation of the candidates (i.e., a one-to-one function from X onto X). Given a relation R on X, define the relation R^μ as follows: For all $a, b \in X$,

$$\mu(a) \; R^\mu \; \mu(b) \text{ iff } a \; R \; b$$

For any profile $\mathbf{R} = (R_1, \ldots, R_n) \in O(X)^n$, a permutation μ applied to $\mathbf{R}$ is defined as $\mathbf{R}^\mu = (R_1^\mu, \ldots, R_n^\mu)$. Now, we define Neutrality as follows:

Neutrality For all profiles $\mathbf{R} \in dom(F)$ and all permutations $\mu : X \to X$, $F(\mathbf{R}^\mu) = F(\mathbf{R})^\mu$.

Neutrality ensures that a social ranking of the candidates depends only on where the candidates fall in the rankings in a given profile. Suppose that there are two profiles $\mathbf{R}$ and $\mathbf{R}'$ and two candidates a and b such that the positions that a occupies in the rankings in $\mathbf{R}$ are the same as the positions that b occupies in $\mathbf{R}'$. That is $\mathbf{R}' = \mathbf{R}^\mu$ where μ is the permutation such that $\mu(a) = b$, $\mu(b) = a$ and for all $c \neq a, b$, $\mu(c) = c$. Then, the social ranking of a given the profile $\mathbf{R}$ must be the same as the social ranking of b given the profile $\mathbf{R}'$.

While Neutrality requires that the candidates are treated equally, another property, Anonymity, requires that the voters are treated equally. A permutation of the voters is a one-to-one function $\pi : V \to V$. Anonymity requires that the group decision does not depend on the *name* of the voters.

Anonymity For all profiles $\mathbf{R} \in dom(F)$ and all permutations of $\pi : V \to V$, $F(\mathbf{R}) = F(\pi(\mathbf{R}))$, where $\pi(\mathbf{R}) = (R_{\pi(1)}, \ldots, R_{\pi(n)})$.

Anonymity is a fundamental requirement of the democratic process and is strictly enforced in most elections. The overall tally of a ranking $R \in O(X)$ in a profile $\mathbf{R}$ is the number of voters that submitted the ranking R (i.e., the tally of R in

$\mathbf{R} = (R_1, \ldots, R_n)$ is $|\{x_i \in V \mid R_i = R\}|$). Anonymity requires that group decisions depend on the tallies of the rankings in a profile rather than the profiles themselves (which identify the voter associated with each ranking).

In many voting situations, anonymity is dropped when the group decision results in a tie. Often, one voter is chosen (perhaps at random) to be the designated "tie-breaker." In such a case, the tie-breaker imposes her strict ranking of the candidates on the rest of the group. An egregious failure of anonymity occurs when there is a voter that imposes her strict rankings of the candidates on the group *no matter what rankings the other voters submit*. Of course, in any given profile, there will often be a number of voters that completely agree with the social ranking. There is nothing wrong with this. Indeed, it may very well be that, for every profile, there is some voter that completely agrees with the group ranking associated with that profile. It is a problem only when the quantifiers are reversed: there is a voter x_d such that for all profiles, voter x_d's strict rankings of the candidates agree with the social ranking. Such a voter is called an **Arrovian dictator**.

Non-Dictator There is no $x_d \in V$ such that for all profiles $\mathbf{R} \in dom(F)$, for all $a, b \in X$, if $a\ P_d\ b$, then $a\ P_{F(\mathbf{R})}\ b$.

Non-Dictatorship ensures that the strict social ranking does not depend on only one voter (cf. the discussion in Section 6).

****

The properties introduced in this section ensure that group decisions depend *in the right way* on the voters' reported rankings. Arrow's ground-breaking theorem identified a surprising conflict between these principles:

Theorem 1 (Arrow [3]). *There are no preference aggregation functions* $F : O(X)^n \to O(X)$*, with* $|X| > 2$*, satisfying Independence of Irrelevant Alternatives, Unanimity, and Non-Dictatorship.*

Much of the subsequent work in social choice theory has focused on finding properties[7] that characterize interesting group decision rules. Amartya Sen adeptly explains the social choice problem in his Nobel Prize lecture:

> When a set of axioms regarding social choice can all be simultaneously satisfied, there may be several possible procedures that work, among which we have to choose. In order to choose between different possibilities through the use of discriminating axioms, we have to introduce further axioms, until one and only one possible procedure remains. This is something of an exercise in brinkmanship. We have to go on and on cutting alternative possibilities, moving—implicitly—towards an impossibility, but then stop just before all possibilities are eliminated, to wit, when one and only one option remains. [42, pg. 354]

[7]Properties of group decision methods are often called "axioms" in the social choice literature. However, the principles studied in the social choice literature do not have the same status as the axioms of, for example, Peano arithmetic or the axioms defining a group. As should be clear from the discussion in this section, many of the so-called axioms of social choice are certainly not "self-evident," and may require extensive justification.

There is much more to say about Arrow's Theorem (cf. [5, 20, 30]). We return to this theorem in Section 5.3, showing how it can be formalized in independence logic.

4 Independence in Social Choice Theory

Arrow's Theorem is directed at preference aggregation functions $F : O(X)^n \to O(X)$. A key assumption, which we only briefly mentioned in the previous section, is that the domain of F is $O(X)^n$. This is the **Universal Domain** (UD) assumption. Thus, F must assign a group ranking to any possible profile of rankings. Arrow argued that, without specialized knowledge about the group decision problem, preference aggregation functions must be designed to handle any possible input:

> If we do not wish to require any prior knowledge of the tastes of individuals before specifying our social welfare function, that function will have to be defined for every logically possible set of individual orderings. [3, pg. 24]

There are two aspects of UD that can be studied separately. The first is that there are no restrictions on the rankings available to a voter. This imposes a richness condition on the domain of a preference aggregation function $F : \mathscr{B} \to \mathscr{O}$, where $\mathscr{B}$ is a set of profiles of (linear) rankings of X:

All rankings For any voter $x_i \in V$ and any (linear) ranking R, there is a profile $\mathbf{R} = (R_1, \ldots, R_n) \in dom(F)$ such that $R_i = R$.

Consider the following set of profiles of linear rankings for three candidates $X = \{a, b, c\}$ and three voters $V = \{x_1, x_2, x_3\}$. To simplify the notation, we write $a\ b\ c$ to denote the ranking $a\ P\ b\ P\ c$. Let $\mathscr{E}$ be the profiles displayed in Table 1 (each row corresponds to a profile).

As the reader is invited to check, $\mathscr{E}$ satisfies the **all rankings** property. Of course, $\mathscr{E}$ does not contain all possible profiles of linear rankings (i.e., $\mathscr{E} \subsetneq L(X)^n$). In particular, voters x_1 and x_3 have the same ranking in each profile. This means that for each $\mathbf{P} \in \mathscr{E}$, the majority ordering $R^{\mathbf{P}}_{maj}$ is transitive. Thus, $F_{maj} : \mathscr{E} \to L(X)$ is a well-defined preference aggregation rule. The problem with this domain is that voters x_1's and x_3's rankings are not chosen *independently*. They form a winning coalition ensuring that the group decision always agrees with their rankings. Thus, both x_1 and x_3 are Arrovian dictators. This suggests an additional constraint on domains of preference aggregation functions.

Independence For any profiles $\mathbf{R} = (R_1, \ldots, R_n), \mathbf{R}' = (R'_1, \ldots, R'_n) \in dom(F)$ and any voter $x_i \in V$, there is a profile $\mathbf{R}'' = (R''_1, \ldots, R''_n) \in dom(F)$ such that $R''_i = R_i$ and $R''_j = R'_j$ for all $j \neq i$.

This constraint ensures that the voters' choice of rankings is not correlated in any way. The domain in Table 1 does not satisfy the independence property: There are profiles $\mathbf{P} = (a\ b\ c,\ a\ b\ c,\ a\ b\ c)$ and $\mathbf{P}' = (a\ c\ b,\ a\ b\ c,\ a\ c\ b)$, but no profile $\mathbf{P}'' = (P''_1, P''_2, P''_3)$ such that $P''_1 = a\ b\ c = P''_2$ and $P''_3 = a\ c\ b$.

Table 1 A set $\mathscr{E}$ of profiles satisfying the **All rankings** property.

x_1	x_2	x_3
$a\ b\ c$	$a\ b\ c$	$a\ b\ c$
$a\ c\ b$	$a\ b\ c$	$a\ c\ b$
$b\ a\ c$	$a\ b\ c$	$b\ a\ c$
$b\ c\ a$	$a\ b\ c$	$b\ c\ a$
$c\ a\ b$	$a\ b\ c$	$c\ a\ b$
$c\ b\ a$	$a\ b\ c$	$c\ b\ a$
$a\ b\ c$	$b\ a\ c$	$a\ b\ c$
$a\ c\ b$	$b\ a\ c$	$a\ c\ b$
$b\ a\ c$	$b\ a\ c$	$b\ a\ c$
$b\ c\ a$	$b\ a\ c$	$b\ c\ a$
$c\ a\ b$	$b\ a\ c$	$c\ a\ b$
$c\ b\ a$	$b\ a\ c$	$c\ b\ a$
$a\ b\ c$	$c\ a\ b$	$a\ b\ c$
$a\ c\ b$	$c\ a\ b$	$a\ c\ b$
$b\ a\ c$	$c\ a\ b$	$b\ a\ c$
$b\ c\ a$	$c\ a\ b$	$b\ c\ a$
$c\ a\ b$	$c\ a\ b$	$c\ a\ b$
$c\ b\ a$	$c\ a\ b$	$c\ b\ a$

x_1	x_2	x_3
$a\ b\ c$	$a\ c\ b$	$a\ b\ c$
$a\ c\ b$	$a\ c\ b$	$a\ c\ b$
$b\ a\ c$	$a\ c\ b$	$b\ a\ c$
$b\ c\ a$	$a\ c\ b$	$b\ c\ a$
$c\ a\ b$	$a\ c\ b$	$c\ a\ b$
$c\ b\ a$	$a\ c\ b$	$c\ b\ a$
$a\ b\ c$	$b\ c\ a$	$a\ b\ c$
$a\ c\ b$	$b\ c\ a$	$a\ c\ b$
$b\ a\ c$	$b\ c\ a$	$b\ a\ c$
$b\ c\ a$	$b\ c\ a$	$b\ c\ a$
$c\ a\ b$	$b\ c\ a$	$c\ a\ b$
$c\ b\ a$	$b\ c\ a$	$c\ b\ a$
$a\ b\ c$	$c\ b\ a$	$a\ b\ c$
$a\ c\ b$	$c\ b\ a$	$a\ c\ b$
$b\ a\ c$	$c\ b\ a$	$b\ a\ c$
$b\ c\ a$	$c\ b\ a$	$b\ c\ a$
$c\ a\ b$	$c\ b\ a$	$c\ a\ b$
$c\ b\ a$	$c\ b\ a$	$c\ b\ a$

It is not hard to see that imposing both **All rankings** and **Independence** ensures that the domain of the preference aggregation function is the set of *all* profiles of (linear) rankings. However, the **All rankings** and **Independence** constraints are stronger than what is needed to prove Arrow's Theorem. A weaker constraint on the domain that is sufficient to prove Arrow's Theorem was identified by Kalai, Muller, and Satterthwaite [19]. Their approach is to weaken the **All rankings** property while maintaining the **Independence** property.

We say that the Independence property is satisfied for a domain $\mathscr{B}$ whenever there is a set $\Omega \subseteq O(X)$ of "admissible" rankings for each voter and $\mathscr{B} = \Omega^n$. The following example from [19] illustrates this. Suppose that $Y = A \cup B$, where $A = \{a_1, a_2, a_3\}$ and $B = \{b_1, b_2, b_3\}$, and let

$$\Omega = \{R \in O(Y) \mid a\ P_R\ b \text{ for } a \in A \text{ and } b \in B\}.$$

So, Ω is the set of all rankings that rank all candidates in A strictly above all candidates in B. Then, $\mathscr{E} = \Omega^n$ satisfies **Independence**. Note that there is a preference aggregation function $F : \mathscr{E} \to O(Y)$ satisfying Unanimity, IIA, and Non-Dictatorship: For each $\mathbf{R} = (R_1, \ldots, R_n) \in \mathscr{E}$, let $F(\mathbf{R}) = (R_1)_A \cup (R_2)_B \cup \{(a, b) \mid a \in A \text{ and } b \in B\}$. That is, F ranks all candidates in A strictly above all candidates in B, ranks the candidates in A according to voter x_1, and ranks the candidates in B according to voter x_2. Thus, voter x_1 is an Arrovian dictator over the set A and x_2 is an Arrovian dictator over the set B, but there is no dictator for the entire set of candidates. This example shows that **Independence** alone is not sufficient to prove Arrow's Theorem.

To see that **All rankings** is not necessary to prove Arrow's Theorem, consider the following variant of the above example. Define $\Omega' \subseteq O(Y)$ as follows:

$$\Omega' = \{R \in O(Y) \mid a\ P_R\ b \text{ for } a \in A \text{ and } b \in B, \text{ and } b_1\ P_R\ b_2\ P_R\ b_3\}.$$

Thus, Ω' contains any ranking R that ranks all candidates in A strictly above all candidates in B, ranks b_1 strictly above b_2 and b_3, and ranks b_2 above b_3. Suppose that $\mathscr{E}' = (\Omega')^n$. Then, $\mathscr{E}'$ satisfies **Independence** but not **All rankings**. By inspecting the proof of Arrow's Theorem, it is not hard to show that there is no preference aggregation function $F : \mathscr{E}' \to O(Y)$ satisfying **Unanimity**, **IIA**, and **Non-Dictatorship**.

The difference between the two domains is that $\mathscr{E}'$ is based on a set of admissible rankings that satisfies a "saturation" property. We need some notation to formally state this property. Suppose that $\Omega \subseteq O(X)$ is a set of rankings for a set X of at least three candidates. A set of three candidates $\{a, b, c\} \subseteq X$ is called a **free triple** for Ω provided a, b, and c are distinct (i.e., $a \neq b \neq c \neq a$) and for each $R \in O(\{a, b, c\})$ there is a $R' \in \Omega$ such that $R'_{\{a,b,c\}} = R$. A pair of distinct candidates $\{a, b\}$ is said to be **trivial** in Ω provided for all $R, R' \in \Omega$, $R_{\{a,b\}} = R'_{\{a,b\}}$ (i.e., all rankings in Ω agree on the ranking of a and b). Two non-trivial pairs of candidates $A = \{a, b\}$ and $B = \{c, d\}$ are **strongly connected** in Ω provided $|A \cup B| = 3$ and $A \cup B$ is a free triple for Ω. Two pairs of candidates A and B are said to be **connected** provided there is a sequence of $B_1, \ldots, B_k$ of pairs of candidates such that $A = B_1$, $B = B_n$, and for all $i = 1, \ldots k - 1$, B_i and B_{i+1} are strongly connected. Finally, say that Ω is **saturated** provided there are at least two non-trivial pairs of candidates and every two non-trivial pairs of candidates are connected. Saturated domains are sufficient to prove Arrow's Theorem:

Theorem 2 (Kalai, Muller, and Satterthwaite [19]). *There is no $F : \mathscr{E} \to O(X)$, where $\mathscr{E} = \Omega^n$ and Ω is saturated, satisfying Unanimity, Independence of Irrelevant Alternatives, and Non-Dictatorship.*

5 Dependence and Independence Logic for Social Choice Theory

In this section, we use dependence and independence logic to formalize the notions of dependence and independence discussed in Sections 3 and 4. The initial idea to use dependence and independence logic to formalize results from social choice theory, such as Arrow's Theorem, is from Jouko Väänänen [49].

We think of the set of voters $V = \{x_1, \ldots, x_n\}$ as a set of distinguished first-order variables. In addition, we include a fresh first-order variable y that is intended to represent the group decision. Suppose that $\mathbf{R} = (R_1, \ldots, R_n) \in O(X)^n$ is a profile for

V and $F : \mathscr{B} \to \mathscr{O}$ is a preference aggregation function with $\mathbf{R} \in \mathscr{B}$. The pair $(\mathbf{R}, F)$ induces an assignment on $V^+ = \{x_1, \ldots, x_n, y\}$, denoted $s_{\mathbf{R},F} : V^+ \to \mathscr{B} \cup \mathscr{O}$, defined as follows:

$$s_{\mathbf{R},F}(x_1) = R_1, \ \ldots, \ s_{\mathbf{R},F}(x_n) = R_n \text{ and } s_{\mathbf{R},F}(y) = F(\mathbf{R}). \tag{1}$$

Then, any group decision function F is associated with a set of assignments:

$$S_F = \{s_{\mathbf{R},F} \mid \mathbf{R} \in dom(F)\} \tag{2}$$

Such a set of assignments is called a *team*, which is the central object of study in dependence and independence logic. Thus, there is a natural link with social choice theory: The properties of preference aggregation functions discussed in Sections 3 and 4 can be viewed as properties of teams, expressible in the language of dependence and independence logic.

Teams of assignments for the variables $\{x_1, \ldots, x_n, y\}$ are intended to represent election scenarios. Each assignment in the team represents a choice of ballot (typically, a ranking of the set of candidates) for each voter and the resulting group decision. Of course, not every team corresponds to some preference aggregation function. In particular, the rankings associated with y must be a function of the rankings associated with $\{x_1, \ldots, x_n\}$. In the language of dependence logic, this means that y *depends* on $\{x_1, \ldots, x_n\}$. Consider the team of assignments displayed in Table 2 assigning to $\{x_1, x_2, y\}$ linear rankings over the set $X = \{a, b, c\}$. That is, each assignment is a map $s : \{x_1, x_2, y\} \to L(X)$. (Recall that we write $a\ \ b\ \ c$ for the ranking $a\ P\ b\ P\ c$.)

Since the rankings associated with y do functionally depend on the rankings associated with the variables $\{x_1, x_2\}$, this team does represent possible election scenarios for 2 voters and 3 candidates. In the remainder of this section, we show how to use dependence and independence logic to reason about group decision methods.

Table 2 An example of a team for 2 voters.

	x_1	x_2	y
s_1	$a\ b\ c$	$c\ b\ a$	$b\ a\ c$
s_2	$a\ c\ b$	$b\ c\ a$	$c\ b\ a$
s_3	$c\ a\ b$	$b\ a\ c$	$a\ c\ b$
s_4	$b\ c\ a$	$a\ c\ b$	$c\ a\ b$
s_5	$a\ b\ c$	$b\ c\ a$	$b\ a\ c$

5.1 The Logic

In this section, we define the syntax and semantics of the version of independence logic (IndS) that we use to formalize Arrow's Theorem.

Suppose that Var is an infinite set of variables with distinguished elements $x_1, \ldots, x_n$ and y (where n is the number of voters). We use $u, v, w, \ldots$ (with or without subscripts) as metalanguage symbols that stand for first-order variables. Suppose that X is a finite set of candidates containing at least three elements. To simplify the presentation of the logical framework, we focus on logics for reasoning about preference aggregation functions. Note that both n (the number of voters) and X (the set of candidates) are parameters in the definition of our language. For simplicity, in this section, we restrict attention to linear orders. Recall that $L(X)$ is the (finite) set of all linear relations on X. Thus, our language is intended to describe properties of functions of the form $F : L(X)^n \to L(X)$. The definitions below can be adapted to reason about other types of group decision functions, such as social choice functions or functions where the domain and/or range is $O(X)$ (the set of all rankings on X).

The signature $\mathscr{L}_X$ contains an equality symbol $=$, unary predicate symbols E_R for each ranking $R \in L(X)$ and unary predicate symbols R_{ab} for each pair (a, b) of elements from X. Since our signature does not contain function symbols or constant symbols, variables are the only $\mathscr{L}$-terms. A **first-order atomic $\mathscr{L}$-formula** is a string of the form $u = v$, $E_R(w)$ or $R_{ab}(w)$.

Definition 3 (Syntax). A well-formed $\mathscr{L}$-formula of **independence logic** for social choice theory (IndS) is a string generated by the following grammar:

$$\varphi ::= \alpha \mid \neg\alpha \mid \bot \mid =\!\!(w_1, \ldots, w_k, u) \mid w_1 \ldots w_k \perp u_1 \ldots u_m \mid w_1 \ldots w_k \subseteq u_1 \ldots u_k$$
$$\mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \forall x\varphi \mid \exists x\varphi,$$

where α is an $\mathscr{L}_X$-atomic first-order formula.

The formulas $=\!\!(w_1, \ldots, w_k, u)$, $w_1 \ldots w_k \perp u_1 \ldots u_m$, and $w_1 \ldots w_k \subseteq u_1 \ldots u_k$ are called *dependence atom*, *independence atom* and *inclusion atom*, respectively. We refer to them in this chapter as *atoms of dependence and independence*. The original independence logic as introduced in [14] does not have dependence atoms or inclusion atoms in the language. Since these two atoms are definable in Ind ([11, 14]), we will include these atoms in the language of our logic IndS.

The set $\mathrm{Fv}(\varphi)$ of free variables of a formula φ of IndS is defined in the standard way except that we have the new cases for the atoms of dependence and independence:

- $\mathrm{Fv}(=\!\!(w_1, \ldots, w_k, u)) = \{w_1, \ldots, w_k, u\}$
- $\mathrm{Fv}(w_1 \ldots w_k \perp u_1 \ldots u_m) = \{w_1, \ldots, w_k, u_1, \ldots, u_m\}$
- $\mathrm{Fv}(w_1 \ldots w_k \subseteq u_1 \ldots u_k) = \{w_1, \ldots, w_k, u_1, \ldots, u_k\}$

A formula φ is called a *sentence* if $\mathrm{Fv}(\varphi) = \emptyset$.

Formulas of IndS are interpreted in standard first-order models M. We assume that the domain of a model M, denoted $dom(M)$, always has at least two elements.

Our formalization of Arrow's Theorem requires that the domain contains all linear rankings of (at least three) candidates. An **intended $\mathscr{L}_X$-model** for IndS is an $\mathscr{L}_X$-model M where $dom(M) = L(X)$. The set of intended models is first-order definable using the unary predicates E_R (this will be shown in Section 5.2). For any $e \in dom(M)$ and any linear ranking $R \in L(X)$, the intended interpretation of $E_R^M(e)$ is that e is the linear ranking R, i.e., $E_R^M = \{e \in dom(M) \mid e = R\}$. For each $e \in dome(M)$, the intended interpretation of $R_{ab}^M(e)$ is that the ranking associated with the element e ranks a above b. More formally, for $a, b \in X$, the unary $R_{ab}^M = \{R \in L(X) \mid a \ R \ b\}$. For example, if $X = \{a, b, c\}$, then $M = (L(X), \{E_R \mid R \in L(X)\}, \{R_{de}^M \mid d, e \in X\})$ is an intended $\mathscr{L}_X$-model for IndS. Suppose that R_1 is the relation $a \ R \ b \ R \ c$; R_2 is the relation $a \ R \ c \ R \ b$; R_3 is the relation $b \ R \ a \ R \ c$; R_4 is the relation $b \ R \ c \ R \ a$; R_5 is the relation $c \ R \ a \ R \ b$; and R_6 is the relation $c \ R \ b \ R \ a$. Then:

- $E_{R_i}^M = \{R_i\}$, for $i = 1, \ldots, 6$; and
- $R_{ab}^M = \{R_1, R_2, R_5\}$, $R_{ac}^M = \{R_1, R_2, R_3\}$, $R_{bc}^M = \{R_1, R_3, R_4\}$, ...

Definition 4 (Assignments, Teams). An **assignment** on M is a map $s : \mathsf{Var} \longrightarrow dom(M)$. A **team** S on M is a set of assignments on M.

For any assignment s and any element $a \in dom(M)$, we write $s(a/w)$ for the assignment defined as $s(a/w)(w) = a$ and $s(a/w)(u) = s(u)$ if $u \neq w$. We now define the team semantics for our logic IndS.

Definition 5 (Semantics). Suppose that M is an $\mathscr{L}$-model for IndS and S is a team on M. For each $\mathscr{L}_X$-formula φ of IndS, we define $M \models_S \varphi$ inductively as follows:

- $M \models_S \alpha$ with α a first-order atomic formula iff for all $s \in S$, $M \models_s \alpha$ in the usual sense;
- $M \models_S \neg\alpha$ with α a first-order atomic formula iff for all $s \in S$, $M \not\models_s \alpha$ in the usual sense;
- $M \models_S \bot$ iff $S = \emptyset$;
- $M \models_S =\!\!(w_1, \ldots, w_k, u)$ iff for all $s, s' \in S$,

$$\text{if } \langle s(w_1), \ldots, s(w_k)\rangle = \langle s'(w_1), \ldots, s'(w_k)\rangle, \text{ then } s(u) = s'(u);$$

- $M \models_S w_1 \ldots w_k \perp u_1 \ldots u_m$ iff for all $s, s' \in S$, there is $s'' \in S$ such that

$$\langle s''(w_1), \ldots, s''(w_k)\rangle = \langle s(w_1), \ldots, s(w_k)\rangle$$

and

$$\langle s''(u_1), \ldots, s''(u_m)\rangle = \langle s'(u_1), \ldots, s'(u_m)\rangle;$$

- $M \models_S w_1 \ldots w_k \subseteq u_1 \ldots u_k$ iff for all $s \in S$, there is $s' \in S$ such that

$$\langle s'(w_1), \ldots, s'(w_k)\rangle = \langle s(u_1), \ldots, s(u_k)\rangle;$$

- $M \models_S \varphi \wedge \psi$ iff $M \models_S \varphi$ and $M \models_S \psi$;
- $M \models_S \varphi \vee \psi$ iff there exist teams $S_0, S_1 \subseteq S$ with $S = S_0 \cup S_1$ such that $M \models_{S_0} \varphi$ and $M \models_{S_1} \psi$;
- $M \models_S \forall w\varphi$ iff $M \models_{S(M/w)} \varphi$, where $S(M/w) = \{s(a/w) \mid s \in S \text{ and } a \in M\}$;
- $M \models_S \exists w\varphi$ iff $M \models_{S[F/w]} \varphi$ for some function $F : S \longrightarrow \wp(M) \setminus \{\emptyset\}$, where $S[F/w] = \{s(a/w) \mid s \in S \text{ and } a \in F(s)\}$;

A sentence φ is said to be *true* in M if the team $\{\emptyset\}$ of the empty assignment satisfies φ, i.e., $M \models_{\{\emptyset\}} \varphi$. We say that a formula φ is a **logical consequence** of a set Γ of formulas provided, for all models M and all teams S on M, if $M \models_S \psi$ for all $\psi \in \Gamma$, then $M \models_S \varphi$. We write $\psi \models \varphi$ for $\{\psi\} \models \varphi$. If $\varphi \models \psi$ and $\psi \models \varphi$, then we say that φ and ψ are *logically equivalent*, in symbols $\varphi \equiv \psi$.

For any team S on a model M and any set $V \subseteq \mathsf{Var}$ of variables, the set $S \upharpoonright V = \{s \upharpoonright V : s \in S\}$ is called a *team on* V. It is straightforward to check that our logic IndS has the Locality Property and the Empty Team Property:

(Locality Property) If $S \upharpoonright \mathrm{Fv}(\varphi) = S' \upharpoonright \mathrm{Fv}(\varphi)$, then $M \models_S \varphi \iff M \models_{S'} \varphi$.
(Empty Team Property) $M \models_\emptyset \varphi$ for all models M.

We refer the reader to [11, 16, 24, 25, 48] for other properties of the logic. In our formalization of Arrow's Theorem, most formulas will have free variables only from the set $V^+ = \{x_1, \ldots, x_n, y\}$ of the distinguished variables that we fixed. By the Locality Property, in most cases, it is then sufficient to consider teams on the set V^+ only. These teams, as discussed, are in one-to-one correspondence to the sets of profiles together with a preference aggregation rule (which may or may not be a function).

We say that a formula of IndS is **first-order**, if it does not contain any atoms of dependence and independence. First-order formulas have the Flatness Property:

(Flatness Property) $M \models_S \varphi$ if, and only if, $M \models_{\{s\}} \varphi$ for all $s \in S$.

For any first-order formula φ, we write $\neg\varphi$ for the (first-order) formula inductively defined as follows:

$$\begin{array}{lll} \neg(\alpha) := \neg\alpha & \neg(\psi \wedge \chi) := \neg\psi \vee \neg\chi & \neg(\forall x\varphi) := \exists x\neg\varphi \\ \neg(\neg\alpha) := \alpha & \neg(\psi \vee \chi) := \neg\psi \wedge \neg\chi & \neg(\exists x\varphi) := \forall x\neg\varphi \end{array}$$

where α is an atomic first-order formula. The reader is invited to check that for any first-order formulas φ and ψ, the following clause holds:

$$M \models_S \neg\varphi \vee \psi \iff \text{for all } s \in S, \text{ if } M \models_{\{s\}} \varphi, \text{ then } M \models_{\{s\}} \psi. \tag{3}$$

In this sense, the formula $\neg\varphi \vee \psi$, abbreviated as $\varphi \supset \psi$, expresses a type of classical material implication that will play a role in the sequel.

It is known that independence logic has the same expressive power as *existential second-order logic* (Σ_1^1) [11]. Therefore, all Σ_1^1-properties of social choice theory can be expressed in our logic. In what follows, we will demonstrate how to express the properties needed to prove Arrow's Theorem.

Our goal is to find a set of formulas Γ_{Arrow} expressing the assumptions of Arrow's Theorem and a first-order formula θ_D expressing that there is an Arrovian dictator such that $\Gamma_{Arrow} \vdash \theta_D$. That is, θ_D is **derivable** in independence logic using the assumptions in Γ_{Arrow}. Due to its strong expressive power, the full independence logic is not axiomatizable (see [16] and also [25]). However, the first-order consequences of IndS are axiomatizable. A complete natural deduction system for the first-order consequence relation over sentences of IndS was given in [16]. More recently, Kontinen [23] generalized this result to open formulas by adding an extra predicate symbol to the signature. Our main goal in this section is to demonstrate that Arrow's Theorem not only can be formalized in IndS, but also can be derived syntactically using the system of [16] and [23].

5.2 *Expressing Arrow's Conditions*

There are three types of properties that we need to express in order to formalize Arrow's Theorem. The first type consists of properties that do not involve any atoms of dependence or independence. These are expressible using first-order formulas only. The second type is intended to capture the notion of dependence from Section 3. The third type captures the notion of independence from Section 4.

First-Order Properties

The first step is to find formulas guaranteeing that the domain contains all linear rankings of the set of candidates X. Let Γ_{DM} of be the following set of sentences:

$$\begin{aligned} & \big\{\exists w(E_R(w)) \mid R \in L(X)\big\} \cup \big\{\forall w \bigvee\{E_R(w) \mid R \in L(X)\}\big\} \\ (\textit{Domain Requirement}) \quad & \cup \big\{\forall w \forall u\big((E_R(w) \wedge E_R(u)) \supset (w = u)\big) \mid R \in L(X)\big\} \\ & \cup \big\{\forall w \bigwedge\{\neg E_R(w) \vee \neg E_{R'}(w) \mid R, R' \in L(X) \text{ and } R \neq R'\}\big\} \end{aligned}$$

Any model M of Γ_{DM} has the property that (1) each linear ranking $R \in L(X)$ corresponds to a unique element e in the model and (2) each element e of the model corresponds to a unique ranking $R \in L(X)$.

The next step is to characterize the intended meaning of the unary predicates R_{ab}. Recall that the atomic formula $R_{ab}(x_i)$ is intended to express the property that voter x_i ranks a above b. Let Γ_{RK} be the following set of sentences:

$$(\textit{Ranking}) \quad \bigcup_{P \in L(X)} \{\forall w\big((E_P(w) \supset R_{ab}(w)) \wedge (R_{ab}(w) \supset E_P(w))\big) \mid a\ P\ b\}$$

It is not hard to see that for any model M of $\Gamma_{DM} \cup \Gamma_{RK}$ and any $a, b \in X$, the interpretation of the predicate R_{ab} is $R_{ab}^M = \{P \in L(X) \mid a\ P\ b\}$. Thus, any model M of $\Gamma_{DM} \cup \Gamma_{RK}$ is an intended $\mathscr{L}_X$-model for IndS. Note that the order-theoretic properties of the relations are logical consequences of $\Gamma_{DM} \cup \Gamma_{RK}$. For instance, if M is a model of $\Gamma_{DM} \cup \Gamma_{RK}$, then the following formula that defines transitivity is true:

$$\forall w \left(\bigwedge \{(R_{ab}(w) \wedge R_{bc}(w)) \supset R_{ac}(w) \mid a, b, c \in X\} \right)$$

We introduce the following notation to express strict preference and indifference, respectively:

(*Strict preference*) For each $a, b \in X$, let $P_{ab}(w) := R_{ab}(w) \wedge \neg R_{ba}(w)$
(*Indifference*) For each $a, b \in X$, let $I_{ab}(w) := R_{ab}(w) \wedge R_{ba}(w)$

Thus, $P_{ab}(x_i)$ means that voter x_i strictly ranks a above b and $I_{ab}(x_i)$ means that voter x_i is indifferent between a and b. Similarly, $P_{ab}(y)$ means that society strictly ranks a above b and $I_{ab}(y)$ means that the society is indifferent between a and b.

Another property that is expressible using only first-order formulas is Unanimity.

(*Unanimity*) $\theta_U := \bigwedge \{(P_{ab}(x_1) \wedge \cdots \wedge P_{ab}(x_n)) \supset P_{ab}(y) \mid a, b \in X\}$.

To see why the above formula expresses Unanimity, suppose that S_F is a team induced by a preference aggregation function F. If $M \models_{S_F} \theta_U$, then for each $a, b \in X$, we have $M \models_{S_F} (P_{ab}(x_1) \wedge \cdots \wedge P_{ab}(x_n)) \supset P_{ab}(y)$. According to equation (3), this means that for all $a, b \in X$ and all $s_{\mathbf{R},F} \in S_F$,

$$\text{if } M \models_{\{s_{\mathbf{R},F}\}} P_{ab}(x_1) \wedge \cdots \wedge P_{ab}(x_n), \text{ then } M \models_{\{s_{\mathbf{R},F}\}} P_{ab}(y).$$

Unpacking the above definitions gives us the definition of Unanimity for a preference aggregation function F: For all candidates $a, b \in X$, and all profiles $\mathbf{R} = (R_1, \ldots, R_n) \in dom(F)$, if $a\ P_i\ b$ for all voters x_i, then $a\ P_{F(\mathbf{R})}\ b$.

Dependence Properties

The first dependence property we will express concerns the functional dependence of the group decision on the voters' rankings. In our setting, this non-trivial property is easily expressed using a simple dependence atom:

(*Functionality of Preference Aggregation Rule*) $\theta_F := {=}(x_1, \ldots, x_n, y)$

Recall that a team S on a model M satisfies θ_F iff for any two assignments $s, s' \in S$,

$$\text{if } s(x_i) = s'(x_i) \text{ for all } 1 \leq i \leq n, \text{ then } s(y) = s'(y).$$

To see that any team S_G induced by a preference aggregation function G satisfies θ_F, the key observation is that if $s, s' \in S_G$, then $s = s_{\mathbf{R},G}$ and $s' = s_{\mathbf{R}',G}$ for some profiles $\mathbf{R}, \mathbf{R}' \in dom(G)$. If $s_{\mathbf{R},G}(x_i) = s_{\mathbf{R}',G}(x_i)$ for all $1 \leq i \leq n$, then $\mathbf{R}$ and $\mathbf{R}'$ are the same profile, and, since G is a function, $s_{\mathbf{R},G}(y) = G(\mathbf{R}) = G(\mathbf{R}') = s_{\mathbf{R}',G}(y)$. We leave it for the reader to check that, conversely, any team satisfying θ_F is associated with a preference aggregation function.

As we argued in Section 3, the notion of dependence found in social choice theory goes beyond simple functional dependence of the group decision on the voters' inputs. The principles from Section 3 are defined to ensure that *properties* of the group ranking depend only on *properties* of the voters' rankings. To express this stronger form of dependence, for any first-order formulas $\varphi_1, \ldots, \varphi_k, \psi$, we introduce a new formula $=\!\!(\varphi_1, \ldots, \varphi_k, \psi)$. To define the semantics of this formula we need some notations. Suppose that Γ is a set of first-order formulas and s, s' are two assignments for a model M. We write $s \sim_\Gamma s'$ when

$$\text{for all } \gamma \in \Gamma, M \models_{\{s\}} \gamma \text{ if, and only if, } M \models_{\{s'\}} \gamma.$$

The semantics for $=\!\!(\varphi_1, \ldots, \varphi_k, \psi)$ is given by the clause:

- $M \models_S =\!\!(\varphi_1, \ldots, \varphi_k, \psi)$ iff for all $s, s' \in S$, if $s \sim_{\{\varphi_1,\ldots,\varphi_k\}} s'$, then $s \sim_{\{\psi\}} s'$.

Without going into any detail, we remark that this new formula $=\!\!(\varphi_1, \ldots, \varphi_k, \psi)$ is definable in our logic IndS, as

$$=\!\!(\varphi_1, \ldots, \varphi_k, \psi) \equiv \exists w_1 \ldots \exists w_k \exists u \exists v_0 \exists v_1 \Big(=\!\!(w_1, \ldots, w_k, u) \wedge =\!\!(v_0) \wedge =\!\!(v_1)$$

$$\wedge (v_0 \neq v_1) \wedge \bigwedge_{i=1}^{k} \big(\theta(w_i, v_0, v_1) \wedge \delta(w_i, \varphi_i, v_0, v_1)\big) \wedge \theta(u, v_0, v_1) \wedge \delta(u, \psi, v_0, v_1)\Big),$$

where $\theta(v, v_0, v_1) := (v = v_0) \vee (v = v_1)$ and

$$\delta(v, \chi, v_0, v_1) := \big(\chi \supset (v = v_1)\big) \wedge \big(\neg\chi \supset (v = v_0)\big).$$

Now, using this generalized dependence formula, we can state IIA in our logic.

(*Independence of Irrelevant Alternatives*)

$$\theta_{IIA} := \bigwedge \{ =\!\!(R_{ab}(x_1), R_{ba}(x_1) \ldots, R_{ab}(x_n), R_{ba}(x_n), R_{ab}(y)) \mid a, b \in X\}.$$

To see that this corresponds to binary independence, note that

$$\text{if } s_{\mathbf{R},F}(x_i) = R_i \in R_{ab}^M \Leftrightarrow s_{\mathbf{R}',F}(x_i) = R'_i \in R_{ab}^M \text{ and}$$

$$s_{\mathbf{R},F}(x_i) = R_i \in R_{ba}^M \Leftrightarrow s_{\mathbf{R}',F}(x_i) = R'_i \in R_{ba}^M, \text{ then } (R_i)_{\{a,b\}} = (R'_i)_{\{a,b\}},$$

where $\mathbf{R} = (R_1, \ldots, R_n)$ and $\mathbf{R}' = (R'_1, \ldots, R'_n)$.

Remark 1 (Alternative Definitions of IIA). One may be tempted to simplify the definition of IIA as follows:

$$\theta'_{IIA} := \bigwedge \{ =\!\!(R_{ab}(x_1), \ldots, R_{ab}(x_n), R_{ab}(y)) \mid a, b \in X\}$$

This formulas says that, for each pair of alternatives $a, b \in X$, the truth of $R_{ab}(y)$ depends only on the truth of $R_{ab}(x_1), \ldots R_{ab}(x_n)$.

Suppose that S_F is a team on a model M induced by a preference aggregation function F that satisfies the formula $=\!\!(R_{ab}(x_1), \ldots, R_{ab}(x_n), R_{ab}(y))$. Then, for any $s_{\mathbf{R},F}, s_{\mathbf{R}',F} \in S_F$,

$$\text{if } s_{\mathbf{R},F}(x_i) = R_i \in R^M_{ab} \Leftrightarrow s_{\mathbf{R}',F}(x_i) = R'_i \in R^M_{ab} \text{ for all } 1 \leq i \leq n,$$

$$\text{then } F(\mathbf{R}) \in R^M_{ab} \Leftrightarrow F(\mathbf{R}') \in R^M_{ab},$$

where $\mathbf{R} = (R_1, \ldots, R_n)$ and $\mathbf{R}' = (R'_1, \ldots, R'_n)$. While this does express a sense in which the social ranking of a and b depends on the individual rankings of a and b, it does *not* express Arrow's IIA property from Section 3. In particular, the above property is more demanding than binary independence. To see why, suppose that a team S contains two assignments $s_{\mathbf{R},F}$ and $s_{\mathbf{R}',F}$, where

- $\mathbf{R} = (R_1, \ldots, R_n)$ with for all i, $a\ I_i\ b$; and
- $\mathbf{R}' = (R'_1, \ldots, R'_n)$ with for all i, $a\ P'_i\ b$.

Since $a\ I_i\ b$ is defined as $a\ R_i\ b$ and $b\ R_i\ a$, $a\ P'_i\ b$ is defined as $a\ R'_i\ b$ and $b\ \not R'_i\ a$, it is true that for all i, $R_i \in R^M_{ab}$ iff $R'_i \in R^M_{ab}$. If θ_{IIA} is satisfied in the team, this would require that $F(\mathbf{R}) \in R^M_{ab}$ iff $F(\mathbf{R}') \in R^M_{ab}$. However, since for all $1 \leq i \leq n$,

$$(R_i)_{\{a,b\}} = \{(a, b), (b, a)\} \neq \{(a, b)\} = (R'_i)_{\{a,b\}},$$

binary independence does not impose any constraints on the social ranking of a and b.

We leave a full discussion of different versions of IIA, including a formalization of m**-ary independence** and a derivation in our logic of Blau's Theorem mentioned in Section 3, for an extended version of this paper.

We conclude this subsection by finding a formula that expresses the existence of an Arrovian dictator. The existence of an Arrovian dictator means that there is a strong form of dependence of the social outcome on a single voter. In particular, if x_d is an Arrovian dictator, then all of x_d's strict rankings are reflected in the social ranking. This is characterized by the following first-order formula:

- $\theta_{D_0}(x_d) := \bigwedge_{a,b \in X} (P_{ab}(x_d) \supset P_{ab}(y))$.

To express that there exists a dictator among the n voters, we need a new connective: The intuitionistic disjunction, denoted by $\mathbin{\vee\!\!\!\vee}$, whose semantics is given by the clause:

- $M \models_S \varphi \mathbin{\vee\!\!\!\vee} \psi$ iff $M \models_S \varphi$ or $M \models_S \psi$.

Without going into any detail, we remark that the intuitionistic disjunction is definable in our logic IndS:

$$\varphi \mathbin{\vee\!\!\!\vee} \psi \equiv \exists w \exists u \Big(=\!\!(w) \wedge =\!\!(u) \wedge \big((w = u) \vee \varphi\big) \wedge \big((w \neq u) \vee \psi\big)\Big),$$

where $w, u \notin \mathrm{Fv}(\varphi) \cup \mathrm{Fv}(\psi)$. The following formula expresses the existence of an Arrovian dictator among the n voters:

$$(\textit{Dictator}) \quad \theta_D := \bigvee\!\!\!\bigvee_{i=1}^{n} \theta_{D_0}(x_i).$$

Independence Properties

As explained in Section 4, a key assumption in Arrow's Theorem is the **Universal Domain** condition. This is characterized by the All Rankings condition and the Independence condition. Our logic IndS can express these two properties:

$$(\textit{All Rankings}) \quad \theta_{AR} := \bigwedge \{\forall u(u \subseteq x_i) : 1 \leq i \leq n\}$$

$$(\textit{Independence}) \quad \theta_I := \bigwedge \big\{\langle x_j\rangle_{j \neq i} \perp x_i : 1 \leq i \leq n\big\}$$

To see that θ_{AR} corresponds to the All Rankings condition, let M be a model and S_F be a team on M induced by a preference aggregation function F. Suppose that $M \models_{S_F} \theta_{AR}$. Then, for each voter x_i, we have $M \models_{S_F(M/u)} u \subseteq x_i$. The value of u ranges over all possible elements of the domain of M. Since the domain of (an intended model) M is the set of all (linear) rankings, the values of u range over all (linear) rankings. The inclusion atom ensures that each such (linear) ranking must occur in the team S_F as a value for the voter x_i. That is, for each ranking $R \in dom(M)$, there is an $s \in S_F$ such that $s(x_i) = R$. This is exactly the **All ranking** property. The correspondence between θ_I and the **Independence** condition is more straightforward, so we leave it for the reader to verify. Note that in our formalization of the Universal Domain assumption, we make essential use of the dependence and independence atoms.

5.3 Arrow's Theorem

There are two additional necessary assumptions for the proof of Arrow's Theorem. The first is that there are at least three candidates (i.e., $|X| \geq 3$). Indeed, if there are only two candidates, then majority rule satisfies Unanimity, Independence of Irrelevant Alternatives, Non-Dictatorship, and Universal Domain (cf. May's Theorem [28] for a characterization of majority rule). This is not a property that can be expressed in our logic. Rather, it is an implicit assumption built into the definition of our logic, as we have fixed a set X containing at least three alternatives and assumed that our signature $\mathscr{L}_X$ has predicate symbols R_{ab} for each pair $a, b \in X$.

The second assumption is that there are only finitely many voters. It can be shown that Arrow's Theorem does not hold (if the Axiom of Choice is assumed) when there are infinitely many voters (see [10, 22]). However, there are analogues of Arrow's Theorem for countably many voters ([46, Chapter 6] and [18]). Again, this is an assumption that is built into the definition of our logic. In the above presentation of our logic IndS, we started by distinguishing a finite set $V^+ = \{x_1, \ldots, x_n, y\}$ of variables. The fact that V^+ is a finite set of variables was implicitly used when we showed that the Arrow conditions are expressible in our logic. In particular, our logic IndS is finitary and we do not see a way to define the propositional dependence formula $=\!\!(\bar{\varphi}, \psi)$ when $\bar{\varphi}$ is an infinite sequence of formulas.

A complete discussion of the different ways in which these last two assumptions can be formally represented in independence logic will be left for an extended version of this paper. For the remainder of this paper, we assume that our logic IndS satisfies the above two assumptions which are needed to prove Arrow's Theorem.

Theorem 3 (Arrow's Theorem, semantic version). *$\Gamma_{Arrow} \models \theta_D$, where $\Gamma_{Arrow} = \Gamma_{DM} \cup \Gamma_{RK} \cup \{\theta_U, \theta_F, \theta_{IIA}, \theta_{AR}, \theta_I\}$.*

The proof of this theorem follows by adapting the standard proofs of Arrow's Theorem (see, for instance, [12] or [1] for a category-theoretic perspective). In the remainder of this section, we will demonstrate that Arrow's Theorem can also be derived syntactically in the natural deduction system of [16] and [23].

We write $\Gamma \vdash \varphi$ if the formula φ can be derived from the set Γ of formulas in the natural deduction system given in [16] and [23]. [8]

Theorem 4 ([16, 23]). *If φ is a first-order formula and Γ a set of formulas of independence logic, then we have $\Gamma \models \varphi$ iff $\Gamma \vdash \varphi$.*

Unfortunately, the above completeness theorem cannot be directly applied to Theorem 3 to show that Arrow's Theorem is derivable in our logic IndS. The problem is that the formula θ_D, which expresses the existence of an Arrovian dictator, is not a first-order formula. Nonetheless, it is possible to transform the formalization of Arrow's Theorem so that we can apply Theorem 4.

Consider a unary connective $\sim$, called **weak classical negation**, whose semantics is given by the clause:

- $M \models_S \sim \varphi$ iff $M \not\models_S \varphi$ whenever $S \neq \emptyset$.[9]

We invite the reader to check the following crucial fact. Note that since our logic has the empty team property, the additional condition "whenever $S \neq \emptyset$" in the semantics of $\sim$ is essential to establish this fact.

Fact 5. $\Gamma_{Arrow} \models \theta_D \iff \Gamma_{Arrow}, \sim \theta_D \models \bot$.

Since the atom $\bot$ (**falsum**) is a first-order formula, we are almost ready to apply Theorem 4. The remaining issue is that we need to make sure that the formula $\sim \theta_D$ is definable in our original independence logic Â IndS, or, equivalently, that it is Σ^1_1.

[8]The interested reader can consult [16] and [23] for the details of the natural deduction system. We do not include the system here since we are only proving the existence of a derivation of Arrow's Theorem rather than providing a derivation. We will take up this challenge in the extended version of this paper.

[9]Note that a slightly different connective $\sim'$ with the semantics $M \models_S \sim' \varphi$ iff $M \not\models_S \varphi$ is known as **classical negation** in the dependence logic literature.

To establish this, let us take a closer look at the formula $\sim \theta_D$, which expresses the fact that there is no dictator. Unpacking the definitions, we obtain

$$\sim \theta_D = \sim \bigvee\!\!\!\bigvee_{i=1}^{n} \theta_{D_0}(x_i) \equiv \bigwedge_{i=1}^{n} \sim \theta_{D_0}(x_i) \equiv \bigwedge_{i=1}^{n} \bigvee\!\!\!\bigvee_{a,b\in X} \sim (P_{ab}(x_i) \supset P_{ab}(y)).$$

This means that the problem reduces to defining the formula $\sim (P_{ab}(x_i) \supset P_{ab}(y))$ in our logic IndS. We present this translation in the following proposition, whose proof is left to the reader.

Proposition 6. $\sim (P_{ab}(x_i) \supset P_{ab}(y)) \equiv \exists w \exists u \big((wu \subseteq x_i y) \wedge P_{ab}(x_i) \wedge \neg P_{ab}(y)\big)$.

Now, since $\Gamma_A, \sim \theta_D \models \bot$ and all the formulas in the set $\Gamma_A \cup \{\sim \theta_D\}$ are expressible in IndS, by Theorem 4 we conclude that $\Gamma_A, \sim \theta_D \vdash \bot$. In order to derive that $\Gamma_A \vdash \theta_D$, we need a weak classical negation elimination rule ($\sim \mathsf{E}$) defined as follows:

$$\dfrac{\varphi \qquad \begin{matrix}[\sim \psi] \\ \vdots \\ \bot\end{matrix}}{\psi} \sim \mathsf{E}$$

We do not see how to derive this rule from the natural deduction system given in [16, 23]. Our solution is to add this rule (which is sound) to the natural deduction system of [16, 23]. We write $\Gamma \vdash^{\sim \mathsf{E}} \varphi$ if φ can be derived from Γ in this extended system. This gives us a syntactic version of Arrow's Theorem:

Theorem 7 (Arrow's Theorem, syntactic version). $\Gamma_{Arrow} \vdash^{\sim \mathsf{E}} \theta_D$.

6 Concluding Remarks

One of the goals of social choice theory is to develop group decision methods that satisfy two main desiderata. The first is that the group decision should depend in the right way on the voters' opinions. The second is that the voters should be free to express any opinion, as long as it is an admissible input to the group decision method. Impossibility theorems, such as Arrow's Theorem, point to an interesting tension between these two desiderata. Properties of group decision methods that ensure that group decisions depend on voters' opinions and that the voters' opinions are independent cannot be simultaneously satisfied. We argued that dependence and independence logic offers an interesting new perspective on this aspect of social choice theory.

Our main focus in the chapter was Arrow's ground-breaking theorem. We developed a version of independence logic that can express Arrow's properties of

Table 3 An example of the reasoning of Arrow's Theorem.

	x_1	x_2	y
s_1	$a\ P_1\ b\ P_1\ c$	$c\ P_2\ b\ P_2\ a$	$b\ P\ a\ I\ c$
s_2	$a\ P_1'\ c\ P_1'\ b$	$c\ P_2'\ b\ P_2'\ a$	??

preference aggregation functions. We then proved that Arrow's Theorem is derivable in a natural deduction system for the first-order consequences of our logic. Our work highlights a number of topics that deserve further study.

The most pressing topic is to find a derivation of Arrow's Theorem in the natural deduction system for our logic. This would not only lead to a potentially new proof of Arrow's Theorem, but it could also identify interesting patterns of reasoning used throughout the social choice literature. To illustrate, consider the following example. Suppose that $S = \{s_1, s_2\}$ is a team on a model M for the set $V = \{x_1, x_2\}$ of two voters and the set $X = \{a, b, c\}$ of three candidates. The assignments are given in Table 3.

Assuming that S satisfies Unanimity and IIA, the question is: What are the possible social ranking for s_2 (i.e., what are the possible values for $s_2(y)$?). Since Unanimity holds for S, i.e., $M \models_S \theta_U$, we have $M \models_{\{s_2\}} P_{cb}(x_1) \wedge P_{cb}(x_2) \supset P_{cb}(y)$. From Table 3 we know that $M \models_{\{s_2\}} P_{cb}(x_1) \wedge P_{cb}(x_2)$, thus we must conclude that $M \models_{\{s_2\}} P_{cb}(y)$, i.e., $s_2(y) \in P^M_{cb}$.

Now, since IIA holds for S, i.e., $M \models_S \theta_{IIA}$, we have

$$M \models_S =(R_{ac}(x_1), R_{ca}(x_1), R_{ac}(x_2), R_{ca}(x_2), R_{ac}(y)) \tag{4}$$

and

$$M \models_S =(R_{ba}(x_1), R_{ab}(x_1), R_{ba}(x_2), R_{ab}(x_2), R_{ba}(y))\,. \tag{5}$$

Let us examine (4). By examining Table 3, we have

$$M \models_{\{s_1\}} R_{ac}(x_1) \wedge \neg R_{ca}(x_1), \quad M \models_{\{s_2\}} R_{ac}(x_1) \wedge \neg R_{ca}(x_1),$$
$$M \models_{\{s_1\}} \neg R_{ac}(x_2) \wedge R_{ca}(x_2) \text{ and } M \models_{\{s_2\}} \neg R_{ac}(x_2) \wedge R_{ca}(x_2).$$

Thus, $s_1 \sim_{\{R_{ac}(x_1), R_{ca}(x_1), R_{ac}(x_2), R_{ca}(x_2)\}} s_2$. Hence, we conclude that $s_1 \sim_{\{R_{ac}(y)\}} s_2$. Now, since $M \models_{\{s_1\}} R_{ac}(y)$, we obtain $M \models_{\{s_2\}} R_{ac}(y)$, meaning $s_2(y) \in R^M_{ac}$. By a similar reasoning, since, $s_1 \sim_{\{R_{ba}(x_1), R_{ab}(x_1), R_{ba}(x_2), R_{ab}(x_2)\}} s_2$, we conclude from (5) that $s_1 \sim_{\{R_{ba}(y)\}} s_2$. Thus, $s_2(y) \in R^M_{ba}$.

Putting everything together, we have $s_2(y) \in P^M_{cb}$ and $s_2(y) \in R^M_{ac} \cap R^M_{ba}$. If, in addition, the team satisfies the transitivity axiom, then $s_2(y)$ cannot be assigned any element of the domain of M. The general approach is to use the dependence and independence properties to generate constraints on the group decision. These constraints may or may not be jointly satisfiable, depending on the form of the group decision (e.g., whether the group decision is a ranking).

A second topic for further investigation is to explore to what extent our logic can be a unifying framework to reason about principles of group decision making. Our analysis has identified three types of dependence found in the social choice literature. Suppose that $w_1, \ldots, w_k$ and v are variables, and φ is a first-order formula.

1. $=\!\!(w_1, \ldots, w_k, v)$: The value assigned to v is completely determined by the values assigned to the w_i.
2. $=\!\!(\varphi(w_1), \ldots, \varphi(w_k), \varphi(v))$: The truth value of $\varphi(v)$ is completely determined by the truth values of the $\varphi(w_i)$.
3. $(\bigwedge_{i=1}^{k} \varphi(w_i)) \supset \varphi(v)$: If each of the w_i satisfy φ, then v must also satisfy φ.

The logic from Section 5 is ideally suited to explore the relationship between these different levels of dependence, especially in conjunction with the independence properties discussed in Section 4. We further conjecture that our logic can capture the reasoning underlying many results related to Arrow's Theorem (e.g., the Muller-Satterthwaite Theorem [32], Wilson's Theorem [50], the Gibbard-Sattherthwaite Theorem [13, 41], and versions of Arrow's Theorem for an infinite population [10, 22]).

Finally, it is important to compare our formalization of Arrow's Theorem with other approaches using modal logic [2, 8, 47], first-order logic [15], and computer-aided proofs [45]. A complete comparison with these different logics for social choice will be left for future work.

References

1. Abramsky, S.: Arrow's theorem by arrow theory. In: Villaveces, A., Kossak, R., Kontinen, J. Hirvonen, Å. (eds.) Logic Without Borders: Essays on Set Theory, Model Theory, Philosophical Logic and Philosophy of Mathematics, pp. 15–30. De Gruyter, Boston, MA (2015)
2. Agotnes, T., van der Hoek, W., Wooldridge, M.: On the logic of preference and judgment aggregation. Auton. Agent. Multi-Agent Syst. **22**(1), 4–30 (2009)
3. Arrow, K.J.: Social Choice and Individual Values. Yale University Press, New Haven, CT (1951)
4. Blau, J.H.: Arrow's theorem with weak independence. Economica **38**(152), 413–420 (1971)
5. Campbell, D.E., Kelly, J.S.: Impossibility theorems in the Arrovian framework. In: Arrow, K.J., Suzumura, A.S.K. (eds.) Handbook of Social Choice and Welfare, Elsevier, vol. 1, pp. 35–94 (2002)
6. Campbell, D.E., Kelly, J.S.: Social welfare functions that satisfy Pareto, anonymity, and neutrality, but not independence of irrelevant alternatives. Soc. Choice Welf. **29**(1), 69–82 (2007)
7. Cato, S.: Independence of irrelevant alternatives revisited. Theor. Decis. **76**(4), 511–527 (2014)
8. Ciná, G., Endriss, U.: A syntactic proof of Arrow's theorem in a modal logic of social choice functions. In: Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2015) (2015)
9. Coleman, J.L., Ferejohn, J.: Democracy and social choice. Ethics **97**(1), 6–25 (1986)

10. Fishburn, P.: Arrow's impossibility theorem: concise proof and infinitely many voters. J. Econ. Theory **2**, 103–106 (1970)
11. Galliani, P.: Inclusion and exclusion in team semantics: on some logics of imperfect information. Ann. Pure Appl. Log. **163**(1), 68–84 (2012)
12. Geanakoplos, J.: Three brief proofs of Arrow's theorem. Econ. Theory **26**(1), 211–215 (2005)
13. Gibbard, A.: Manipulation of voting schemes: a general result. Econometrica **41**(3), 587–601 (1973)
14. Grädel, E., Väänänen, J.: Dependence and independence. Stud. Logica **101**(2), 399–410 (2013)
15. Grandi, U., Endriss, U.: First-order logic formalisation of impossibility theorems in preference aggregation. J. Philos. Log. **42**(4), 595–618 (2013)
16. Hannula, M.: Axiomatizing first-order consequences in independence logic. Ann. Pure Appl. Log. **166**(1), 61–91 (2015)
17. Hausman, D.M.: Preference, Value, Choice and Welfare. Cambridge University Press, Cambridge (2011)
18. Herzberg, F., Eckert, D.: The model-theoretic approach to aggregation: impossibility results for finite and infinite electorates. Math. Soc. Sci. **1**(2012), 41–47 (64)
19. Kalai, E., Muller, E., Satterthwaite, M.A.: Social welfare functions when preferences are convex and continuous: Impossibility results. Public Choice **34**, 87–97 (1979)
20. Kelly, J.S.: Arrow Impossibility Theorems. Academic, New York (1978)
21. Kemp, M.C., Ng, Y.K.: On the existence of social welfare functions, social orderings and social decision functions. Economica **43**, 59–66 (1976)
22. Kirman, A.P., Sondermann, D.: Arrow's theorem, many agents and invisible dictators. J. Econ. Theory **5**, 267–277 (1972)
23. Kontinen, J.: On natural deduction in dependence logic. In: Villaveces, A., Kossak, R., Kontinen, J., Hirvonen, A. (eds.) Logic Without Borders: Essays on Set Theory, Model Theory, Philosophical Logic and Philosophy of Mathematics, pp. 297–304. De Gruyter, Boston, MA (2015)
24. Kontinen, J., Väänänen, J.: On definability in dependence logic. J. Log. Lang. Inf. **18(3)**, 317–332 (2009)
25. Kontinen, J., Väänänen, J.: Axiomatizing first-order consequences in dependence logic. Ann. Pure Appl. Log. **164**(11), 1101–1117 (2013)
26. List, C.: Social choice theory. In: Zalta, E.N. (ed.) The Stanford Encyclopedia of Philosophy, Winter 2013 edn. (2013)
27. Mackie, G.: Democracy Defended. Cambridge University Press, Cambridge (2003)
28. May, K.: A set of independent necessary and sufficient conditions for simply majority decision. Econometrica **20**(4), 680–684 (1952)
29. Morreau, M.: It simply does not add up: trouble with overall similarity. J. Philos. **107**, 469–490 (2010)
30. Morreau, M.: Arrow's theorem. In: Zalta, E.N. (ed.) The Stanford Encyclopedia of Philosophy, Winter 2014 edn. (2014)
31. Morreau, M.: Mr. Fit, Mr. Simplicity and Mr. Scope: from social choice to theory choice. Erkenntnis **79**, 1253–1268 (2014)
32. Muller, E., Satterthwaite, M.A.: The equivalence of strong positive association and strategy-proofness. J. Econ. Theory **14**, 412–418 (1977)
33. Okasha, S.: Theory choice and social choice: Kuhn versus Arrow. Mind **120**, 83–115 (2011)
34. Pacuit, E.: Voting methods. In: Zalta, E.N. (ed.) The Stanford Encyclopedia of Philosophy, Winter 2012 edn. (2012)
35. Parks, R.P.: An impossibility theorem for fixed preferences: a dictatorial Bergson-Samuelson welfare function. Rev. Econ. Stud. **43**(3), 447–450 (1976)
36. Riker, W.H.: Liberalism Against Populism. W. H. Freeman, San Francisco, CA (1982)
37. Roberts, K.W.S.: Social choice theory: the single-profile and multi-profile approaches. Rev. Econ. Stud. **47**(2), 441–450 (1980)
38. Rubinstein, A.: The single profile analogues to multi profile theorems: mathematical logic's approach. Int. Econ. Rev. **25**(3), 719–730 (1984)

39. Saari, D.: Mathematical structure of voting paradoxes: II. Positional voting. J. Econ. Theory **15**(1), 55–102 (2000)
40. Saari, D.: Chaotic Elections! A Mathematician Looks at Voting. American Mathematical Society, Providence, RI (2001)
41. Satterthwaite, M.A.: Strategy-proofness and Arrow's conditions: existence and correspondence theorems for voting procedures and social welfare functions. J. Econ. Theory **10**(2), 187–217 (1975)
42. Sen, A.: The possibility of social choice. Am. Econ. Rev. **89**(3), 349–378 (1999)
43. Sen, A., Maskin, E.: The Arrow Impossibility Theorem. Columbia University Press, New York (2014)
44. Stegenga, J.: An impossibility theorem for amalgamating evidence. Synthese **190**(2), 2391–2411 (2013)
45. Tang, P., Lin, F.: Computer-aided proofs of Arrow's and other impossibility theorems. Artif. Intell. **173**(11), 1041–1053 (2009)
46. Taylor, A.D.: Social Choice and the Mathematics of Manipulation. Cambridge University Press, Cambridge (2005)
47. Troquard, N., van der Hoek, W., Wooldridge, M.: Reasoning about social choice functions. J. Philos. Log. **40**(4), 473–498 (2011)
48. Väänänen, J.: Dependence Logic: A New Approach to Independence Friendly Logic. Cambridge University Press, Cambridge (2007)
49. Väänänen, J.: Dependence logic. Talk — Dagstuhl Seminar (2010)
50. Wilson, R.: Social choice theory without the Pareto principle. J. Econ. Theory **5**, 478–486 (1972)
51. Young, H.P.: Optimal voting rules. J. Econ. Perspect. **9**(1), 51–64 (1995)