Soc Choice Welf (2009) 32:221–241
DOI 10.1007/s00355-008-0320-1

ORIGINAL PAPER

Aggregation of binary evaluations for truth-functional agendas

Elad Dokow · Ron Holzman

Received: 6 December 2007 / Accepted: 19 May 2008 / Published online: 18 June 2008

Abstract In the problem of judgment aggregation, a panel of judges has to evaluate each proposition in a given agenda as true or false, based on their individual evaluations and subject to the constraint of logical consistency. We elaborate on the relation between this and the problem of aggregating abstract binary evaluations. For the special case of truth-functional agendas we have the following main contributions: (1) a syntactical characterization of agendas for which the analogs of Arrow's aggregation conditions force dictatorship; (2) a complete classification of all aggregators that satisfy those conditions; (3) an analysis of the effect of weakening the Pareto condition to surjectivity.

1 Introduction

The problem of judgment aggregation has received a significant amount of attention recently; see, e.g., the survey by List and Puppe (2007). It can be described as follows. There is a panel of n judges that faces an agenda $\mathcal{P}$ of m logical propositions, whose

This is a sequel to the paper "Aggregation of binary evaluations." The contents of both papers were presented, under the title "An Arrovian impossibility theorem for social truth functions," at the Second World Congress of the Game Theory Society, Marseille, July 2004. The first version of "Aggregation of binary evaluations" was completed in June 2005. That working paper was subsequently split into two parts, of which this is the second. The comments of an anonymous referee are gratefully acknowledged. Part of R. Holzman's work was done while he was a Fellow of the Institute for Advanced Studies at the Hebrew University of Jerusalem.

E. Dokow · R. Holzman (✉)
Department of Mathematics, Technion-Israel Institute of Technology, 32000 Haifa, Israel
e-mail: holzman@techunix.technion.ac.il

truth or falsehood has to be determined.[1] The propositions are interrelated, and so only a certain subset of all 2^m such determinations are logically consistent. The problem is to aggregate the n individual evaluations, each of which is assumed to be logically consistent, into a joint evaluation that needs to be logically consistent. Starting from the first axiomatic treatment by List and Pettit (2002), the literature has identified various combinations of conditions on the agenda and requirements from the aggregator that force the aggregator to be dictatorial. Certain ways out of this predicament have also been explored.

There exists in the literature also a related, but different model of abstract aggregation. In that model, there are n individuals facing m issues, each admitting two possible positions denoted by 0 and 1. The issues are abstract, not represented by logical propositions. Rather, there is an exogenously given subset X of $\{0, 1\}^m$, with the interpretation that only the evaluations in X are feasible. The problem is to aggregate each n-tuple of individual feasible evaluations into a joint feasible evaluation. This framework for abstract aggregation theory was introduced by Wilson (1975) and further developed by Rubinstein and Fishburn (1986a) and Fishburn and Rubinstein (1986b). These authors showed that the problem of (strict) preference aggregation may also be cast in this framework, and Arrow's (1951) impossibility theorem for that problem may be obtained as a special case of certain results for the abstract framework. In an earlier paper (Dokow and Holzman 2005), we identified two conditions on the set X of feasible evaluations, called total blockedness and nonaffineness, which together characterize those sets X for which the natural analog of Arrow's theorem holds true. Namely, they characterize the sets X such that every IIA and Paretian[2] aggregator $f : X^n \to X$ must be dictatorial.[3]

We proceed to describe the contributions of the current paper. In Sect. 2, after giving the formal definitions of the logic-based and the abstract frameworks, we show that they are equivalent in the following sense. Clearly, from an instance of the logic-based framework with agenda $\mathcal{P}$, we can construct a corresponding instance of the abstract framework by taking $X = X(\mathcal{P})$ to be the set of logically consistent evaluations of the propositions in $\mathcal{P}$. Conversely, given a nonempty subset X of $\{0, 1\}^m$, we show how to construct an agenda $\mathcal{P}$ of m propositions so that $X(\mathcal{P}) = X$. The construction is canonical but is not unique: there may be different agendas $\mathcal{P}$ and $\mathcal{P}'$ so that $X(\mathcal{P}) = X(\mathcal{P}')$. In this sense, passing from the logic-based framework to the abstract one entails a loss of information. This loss may be significant for certain aspects of the aggregation problem, as pointed out by List and Puppe (2007), but for the questions of possibility versus impossibility that we consider here, the set X contains all the relevant information.

[1] The literature has considered also some more general models, allowing for infinite agendas, or for many-valued logic, or for incomplete determinations of truth. We will not be concerned with these relaxations here.

[2] We say that an aggregator is IIA (independent of irrelevant alternatives) if the joint position on any given issue depends only on the individual positions on that same issue. We say that an aggregator is Paretian if any unanimously held position is adopted as the joint position.

[3] The condition of total blockedness was introduced by Nehring and Puppe (2002) in a different, but equivalent framework of property spaces. In our terminology, they proved that this condition characterizes the sets X such that every *monotone* IIA and Paretian aggregator $f : X^n \to X$ must be dictatorial.

Nevertheless, it is desirable to have a procedure to determine whether or not a given agenda $\mathcal{P}$ is an impossibility agenda, in the sense that every IIA and Paretian aggregator for it must be dictatorial. In principle, this could be done by passing to $X = X(\mathcal{P})$ and then checking the conditions on X, but being able to check this directly on $\mathcal{P}$, in terms of the syntactical structure of the propositions, would be more satisfactory. After recalling in Sect. 3, for the reader's convenience, the characterization of impossibility in terms of X from Dokow and Holzman (2005), we give a corresponding characterization in terms of $\mathcal{P}$ in Sect. 4. This version of the characterization, however, is restricted to a class of agendas called truth-functional. An agenda is truth-functional if it can be split into two parts, so that the evaluation of the propositions in the first part (the premises) is constraint-free, and it entirely determines the evaluation of the propositions in the second part (the conclusions).[4] A pleasing aspect of our syntactical characterization of impossibility for truth-functional agendas is that the conditions are easy to verify: it yields a computationally efficient procedure to classify $\mathcal{P}$ as a possibility or an impossibility agenda.[5]

In Sect. 5, we characterize all IIA and Paretian aggregators for any given truth-functional agenda $\mathcal{P}$. Essentially (and excluding some trivial exceptions), all such aggregators are combinations of oligarchic rules and parity rules.[6] Thus, even if $\mathcal{P}$ is a possibility agenda, only some very specific (and arguably unattractive) IIA and Paretian aggregators are available for it.[7] The conclusion is that, at least for nontrivial truth-functional agendas, the IIA and Pareto requirements from the aggregator are indeed extremely demanding and need to be relaxed.

In Sect. 6, we explore the relaxation of the Pareto condition.[8] We replace it with the requirement that the aggregator be surjective, so that every logically consistent evaluation is attained as the joint evaluation for some profile of individual evaluations. It turns out that if a truth-functional agenda $\mathcal{P}$ satisfies the conditions of Sect. 4 (under which IIA and Pareto allow only dictatorship), this weakening ushers in only aggregators, which are dictatorial up to the reversal of the position on some prescribed subagenda.[9] Such aggregators may or may not be available, depending on the agenda, and in any case, they are not attractive.

[4] This is essentially equivalent to assuming that every atomic proposition that appears in some proposition in $\mathcal{P}$ is itself a member of $\mathcal{P}$. When this is the case, we can take the atomic propositions to be the premises, and the composite propositions to be the conclusions.

[5] Nehring and Puppe (2008) independently obtained a characterization of truth-functional agendas that are impossibility agendas in their sense (which requires that every monotone IIA and Paretian aggregator be dictatorial). They also extended their characterization to handle situations where the evaluation of the premises is not free, but subject to certain constraints.

[6] The part of this result that singles out oligarchic rules has its monotone counterpart in Nehring and Puppe (2008).

[7] We also characterize when one can find, among the available aggregators, some that satisfy additional desirable properties such as no veto power, neutrality, and anonymity. These extra requirements further restrict the realm of possibility.

[8] Some relaxations of IIA were recently studied by Mongin (2005) and Dietrich (2006).

[9] In a related result, Dietrich and List (2007) gave conditions that single out two kinds of aggregators: dictatorship and inverse dictatorship. The more intricate possibility of reversing the dictator's positions only on some propositions did not come up in their analysis, because they assumed systematicity, which requires equal treatment of all propositions. Their result was not limited to truth-functional agendas.

2 Two frameworks and the equivalence between them

In the logic-based framework for judgment aggregation, we consider the language of propositional logic that contains the atomic propositions $p_1, \ldots, p_k$ and the connectives $\neg, \wedge, \vee, \rightarrow, \leftrightarrow$.[10] Let $\mathcal{L}(p_1, \ldots, p_k)$ be the set of all propositions in this language, that is, the smallest set $\mathcal{L}$ of expressions that includes $p_1, \ldots, p_k$ and satisfies: $\varphi \in \mathcal{L}$ implies $(\neg\varphi) \in \mathcal{L}$, and $\varphi, \psi \in \mathcal{L}$ implies $(\varphi \wedge \psi), (\varphi \vee \psi), (\varphi \rightarrow \psi), (\varphi \leftrightarrow \psi) \in \mathcal{L}$. We use the binary digits 0 and 1 as truth values (meaning "false" and "true" respectively). For $u \in \{0, 1\}$, we let $\overline{u}$ be the other element of $\{0, 1\}$, and for $u, v \in \{0, 1\}$, we denote by $u + v$ their sum modulo 2. Given an assignment $t = (t_1, \ldots, t_k) \in \{0, 1\}^k$ of truth values to the atomic propositions, we define inductively the truth value $T_t(\varphi) \in \{0, 1\}$ of $\varphi \in \mathcal{L}(p_1, \ldots, p_k)$ by the following rules: (1) $T_t(p_i) = t_i$, (2) $T_t(\neg\varphi) = \overline{T_t(\varphi)}$, (3) $T_t(\varphi \wedge \psi) = \min\{T_t(\varphi), T_t(\psi)\}$, (4) $T_t(\varphi \vee \psi) = \max\{T_t(\varphi), T_t(\psi)\}$, (5) $T_t(\varphi \rightarrow \psi) = \max\{\overline{T_t(\varphi)}, T_t(\psi)\}$, and (6) $T_t(\varphi \leftrightarrow \psi) = \overline{T_t(\varphi) + T_t(\psi)}$.

The *agenda* in a judgment aggregation problem is the set of propositions whose truth values need to be determined; in our set-up, it is a nonempty finite subset $\mathcal{P}$ of $\mathcal{L}(p_1, \ldots, p_k)$. For convenience, we assume that the propositions in $\mathcal{P}$ are listed as $\mathcal{P} = \{\varphi_1, \ldots, \varphi_m\}$, but the order is arbitrary and does not carry any substantial meaning (such as priority).[11] The set $\{0, 1\}^m$ is the set of *evaluations* of $\mathcal{P}$, with the jth coordinate representing the truth value assigned to φ_j. Its subset

$$X(\mathcal{P}) = \{(T_t(\varphi_1), \ldots, T_t(\varphi_m)) \,|\, t \in \{0, 1\}^k\}$$

is the set of all *logically consistent* evaluations of $\mathcal{P}$.

In the abstract aggregation framework, we consider a nonempty finite set of issues J. For convenience, we identify J with the set $\{1, \ldots, m\}$ of coordinates of vectors of length $m = |J|$. We think of issues as abstract entities on which two positions are possible, and denote those positions by 0 and 1. Thus, $\{0, 1\}^m$ is the set of *evaluations* of J, with the jth coordinate representing the position taken on issue j. We assume that some nonempty subset X of $\{0, 1\}^m$ is given. The evaluations in X are called *feasible*; the others are infeasible. Feasibility is viewed as the primitive notion in this framework, and it may have different interpretations in different applications.

One concrete application of the abstract framework is to the logic-based framework described above. Given an agenda $\mathcal{P} = \{\varphi_1, \ldots, \varphi_m\}$, we can view $\varphi_1, \ldots, \varphi_m$ as the issues, and take $X = X(\mathcal{P})$, so that feasibility means logical consistency. Thus,

---

[10] We follow the standard construction of the propositional calculus, except that we assume that there are only finitely many atomic propositions. Since we are going to consider only finite agendas, this will suffice for our purposes.

[11] We could allow the list $\varphi_1, \ldots, \varphi_m$ to include repetitions, but there is no gain in that; so, we assume that $|\mathcal{P}| = m$. (We do allow, however, two propositions in $\mathcal{P}$ to be logically equivalent.) We note that most authors represent such an agenda as consisting of $\varphi_1, \neg\varphi_1, \ldots, \varphi_m, \neg\varphi_m$, and consider complete judgment sets that contain one member of every pair $\varphi_j, \neg\varphi_j$; this is equivalent to assigning the value 1 or 0 to each φ_j in our set-up. Incomplete judgment sets, containing at most one member of every pair $\varphi_j, \neg\varphi_j$, are not captured by the framework described here, but this can be done by extending the set $\{0, 1\}$ to $\{0, 1, *\}$; see Dokow and Holzman (2006).

any instance of the logic-based framework may be viewed, by abstracting from the propositions, as an instance of the abstract framework. We show next that all instances of the abstract framework are obtained in this way, which implies that the logic-based framework is as general as the abstract one. Indeed, the following proposition asserts that any abstract instance with m issues may be represented by an agenda in a language with m atomic propositions, i.e., $k = m$. Of course, in some cases, fewer atomic propositions could suffice.

Proposition 2.1 *For every nonempty subset X of $\{0,1\}^m$, there exists an agenda $\mathcal{P} = \{\varphi_1, \ldots, \varphi_m\} \subseteq \mathcal{L}(p_1, \ldots, p_m)$ so that $X(\mathcal{P}) = X$.*

Proof Let $\emptyset \neq X \subseteq \{0,1\}^m$ be given. We denote by X_ℓ the projection of X on the first ℓ coordinates. We are going to construct the propositions $\varphi_1, \ldots, \varphi_m$ inductively, so that by denoting $\mathcal{P}_\ell = \{\varphi_1, \ldots, \varphi_\ell\}$, we will have after the ℓth stage of the construction $\mathcal{P}_\ell \subseteq \mathcal{L}(p_1, \ldots, p_\ell)$ and $X(\mathcal{P}_\ell) = X_\ell$.

At the first stage, we let

$$\varphi_1 = \begin{cases} p_1 & \text{if } X_1 = \{0,1\} \\ p_1 \wedge \neg p_1 & \text{if } X_1 = \{0\} \\ p_1 \vee \neg p_1 & \text{if } X_1 = \{1\} \end{cases}$$

Suppose that $1 \leq \ell \leq m-1$ and we have already constructed $\mathcal{P}_\ell$ satisfying the requirements. For each $u \in \{0,1\}$, let

$$X_{\ell,u} = \{(x_1, \ldots, x_\ell) \in X_\ell |\, (x_1, \ldots, x_\ell, \overline{u}) \notin X_{\ell+1}\}$$

Note that $X_{\ell,0}$ and $X_{\ell,1}$ are two (possibly empty) disjoint subsets of X_ℓ. For every $x = (x_1, \ldots, x_\ell) \in X_{\ell,0} \cup X_{\ell,1}$, let

$$\phi_x = \varphi_1^{x_1} \wedge \cdots \wedge \varphi_\ell^{x_\ell}$$

where $\varphi_j^{x_j}$ denotes φ_j if $x_j = 1$ and $\neg\varphi_j$ if $x_j = 0$. Note that, for any assignment $t = (t_1, \ldots, t_\ell) \in \{0,1\}^\ell$ of truth values to $p_1, \ldots, p_\ell$, we have

$$T_t(\phi_x) = 1 \Leftrightarrow T_t(\varphi_j) = x_j \text{ for } j = 1, \ldots, \ell$$

Now we define

$$\varphi_{\ell+1} = \left(p_{\ell+1} \wedge \neg \left(\bigvee_{x \in X_{\ell,0}} \phi_x\right)\right) \vee \left(\bigvee_{x \in X_{\ell,1}} \phi_x\right)$$

with the provision that if any of the sets $X_{\ell,u}$ is empty then the corresponding part of $\varphi_{\ell+1}$ is dropped. It is easily checked that $X(\mathcal{P}_{\ell+1}) = X_{\ell+1}$, as required for the inductive construction. □

In view of Proposition 2.1, it is clear that any aggregation problem that can be modeled by the logic-based framework may also be modeled by the abstract framework, and vice versa, but the model in the logic-based framework has an extra ingredient, the agenda, that is not present in the abstract framework, and cannot be uniquely reconstructed from the abstract model. Since the problem focuses on aggregating evaluations in X [or $X(\mathcal{P})$], it seems that for most purposes the abstract framework is adequate and the extra ingredient of $\mathcal{P}$ is redundant. Nevertheless, when the aggregation problem is naturally formulated in logical terms by an agenda $\mathcal{P}$, one may want to answer natural questions in logical terms, thus keeping $\mathcal{P}$ in the picture; this is what we do for most of this paper.

Another comment concerns the choice of logical language. In choosing the propositional calculus, we follow the practice of the first papers on judgment aggregation. More recently, other logics were considered, and results were obtained for any logic satisfying certain criteria (Dietrich 2007). The trivial direction of the equivalence explained above, i.e., the fact that any logical agenda gives rise to an instance of the abstract framework, does extend to general logics, provided they are two-valued and the concept of logical consistency is well defined. The other direction, supplied by Proposition 2.1, will depend on the properties of the logic under consideration. Without going into a detailed analysis, we note intuitively that the equivalence will hold for any sufficiently expressive logic.

In the rest of this section, we introduce some types of agendas and some conventions on how we write propositions. First, in the interest of readability, we denote atomic propositions by $p, q, \ldots$ (rather than $p_1, p_2, \ldots$). Secondly, since different syntactical forms such as $p \to q$ and $\neg p \vee q$ may represent logically equivalent propositions, we choose a particular form of writing propositions. Every proposition may be expressed in *disjunctive normal form* (abbreviated DNF), that is, as a disjunction of one or more clauses, each of which is a conjunction of one or more literals (a literal is an atomic proposition or its negation). For example, the expression

$$(\neg p \wedge \neg q \wedge \neg r) \vee (\neg q \wedge r) \vee (p \wedge r) \vee (p \wedge q)$$

is in DNF, with the expressions in parentheses as its clauses. A DNF may not be minimal in one or both of the following senses. It may happen that a literal may be dropped from a clause without altering the set of satisfying assignments. For instance, the literal $\neg r$ may be dropped from the first clause in the example above, because when p and q are false and r is true, the proposition is also true due to its second clause. Also, it may be the case that an entire clause may be dropped without altering the set of satisfying assignments. For instance, the third clause in the example above may be dropped, because any satisfying assignment for it also satisfies either the second or the fourth clause. Thus, an expression in minimal DNF for our example would be

$$(\neg p \wedge \neg q) \vee (\neg q \wedge r) \vee (p \wedge q).$$

Note also that after dropping $\neg r$ from the first clause, we could have dropped the second clause instead of the third, thus obtaining the logically equivalent minimal

DNF

$$(\neg p \wedge \neg q) \vee (p \wedge r) \vee (p \wedge q).$$

The nonuniqueness of minimal DNF expression will not cause any trouble.

We will consider agendas $\mathcal{P}$, which do not contain any proposition that is a tautology or a contradiction, and in which each proposition is in minimal DNF. Such an agenda $\mathcal{P}$ will be called *standard*. Note that restricting our attention to standard agendas entails no loss of generality.[12]

We will adopt the following notation and terminology. The literals p and $\neg p$ will also be referred to as p^1 and p^0, respectively. When we say, for some specific $u \in \{0, 1\}$, that p^u appears in a proposition φ, we mean that it has an appearance in φ in this particular form. When we just say that p appears in φ, we mean that it has an appearance in φ, which may be either as p or as $\neg p$. (e.g., p^1 appears in $p \wedge q$ but not in $\neg p \wedge q$, while p appears in both.)

A proposition φ in DNF is of exactly one of the following four types. It may be a literal, or it may be a conjunction of two or more literals (we call this an AND proposition), or it may be a disjunction of two or more literals (we call this an OR proposition), or it may be a disjunction of two or more clauses, at least one of which is a conjunction of two or more literals (we call this an OR- AND proposition).

An agenda $\mathcal{P} = \{\varphi_1, \ldots, \varphi_m\}$ is called *truth-functional*, if it can be split into two parts, say $\mathcal{P}_1 = \{\varphi_1, \ldots, \varphi_\ell\}$ (the premises) and $\mathcal{P}_2 = \{\varphi_{\ell+1}, \ldots, \varphi_m\}$ (the conclusions), so that $X(\mathcal{P}_1) = \{0, 1\}^\ell$ and for each $\ell + 1 \leq j \leq m$ there exists a function $g_j \colon \{0, 1\}^\ell \to \{0, 1\}$ so that $T_t(\varphi_j) = g_j(T_t(\varphi_1), \ldots, T_t(\varphi_\ell))$ for every assignment t of truth values to the atomic propositions. Clearly, not all agendas are truth-functional,[13] but some of the judgment aggregation literature has focused on truth-functional agendas, and so do we in this paper. The pragmatic reason for that is the ability to use the aggregation of evaluations of the premises as a stepping stone for dealing with the full aggregation problem.

For an agenda $\mathcal{P}$, we denote by $\mathcal{P}_{\mathrm{a}}$ the set of atomic propositions that appear in some proposition in $\mathcal{P}$. A sufficient condition for $\mathcal{P}$ to be truth-functional is that $\mathcal{P}_{\mathrm{a}} \subseteq \mathcal{P}$, because then we can take $\mathcal{P}_1 = \mathcal{P}_{\mathrm{a}}$. This condition is not necessary for truth-functionality: for example,

$$\mathcal{P} = \{p, q \wedge r, p \vee (q \wedge r)\}$$

is truth-functional (with premises p and $q \wedge r$) even though $q, r \in \mathcal{P}_{\mathrm{a}} \backslash \mathcal{P}$, but it is easy to see that any truth-functional agenda $\mathcal{P}$ may be transformed into an agenda $\mathcal{P}'$ satisfying $\mathcal{P}'_{\mathrm{a}} \subseteq \mathcal{P}'$ without affecting $X(\mathcal{P})$. For instance, $\mathcal{P}$ of the above example

[12] Indeed, suppose we are given an arbitrary agenda $\mathcal{P}$. Then, we can pass from $\mathcal{P}$ to a standard agenda $\mathcal{P}'$ by deleting any tautologies or contradictions, and replacing each remaining proposition that is not in minimal DNF by a logically equivalent one in minimal DNF. This change does not affect the aggregation problem as it is treated here. (Although, on a conceptual level, one may assign importance to the syntax of the propositions in the agenda, in which case the assumption of a standard agenda would not be innocuous.)

[13] One way to see this is to note that the definition of truth-functionality implies that $|X(\mathcal{P})|$ is of the form 2^ℓ, whereas for general agendas there is no such restriction.

may be transformed into

$$\mathcal{P}' = \{p, s, p \vee s\}.$$

Thus, when dealing with truth-functional agendas, there is no loss of generality in assuming that $\mathcal{P}_a \subseteq \mathcal{P}$. An agenda $\mathcal{P}$ will be called *standard truth-functional* if it is standard (in the sense defined above) and satisfies $\mathcal{P}_a \subseteq \mathcal{P}$.[14]

3 A general characterization of impossibility domains

In this section, we recall the definitions and the main result of Dokow and Holzman (2005). They are stated in the abstract framework of aggregation, for a given subset X of $\{0, 1\}^m$, but they apply also to the logic-based framework upon setting $X = X(\mathcal{P})$.

A *society* is a finite, nonempty set N. For convenience, if there are n individuals in N, we identify N with the set $\{1, \ldots, n\}$. If we specify a feasible evaluation $x^i = (x^i_1, \ldots, x^i_m) \in X$ for each individual $i \in N$, we obtain a *profile* of feasible evaluations $\mathbf{x} = (x^i_j) \in X^n$. We may view a profile as an $n \times m$ matrix all of whose rows lie in X. We use superscripts to indicate individuals (rows) and subscripts to indicate issues (columns).

An *aggregator* for N over X is a mapping $f: X^n \to X$. It assigns to every possible profile of individual feasible evaluations, a social evaluation, which is also feasible. Any aggregator f may be written in the form $f = (f_1, \ldots, f_m)$, where f_j is the jth component of f. That is, $f_j: X^n \to \{0, 1\}$ assigns to every profile the social position on the jth issue.

An aggregator $f: X^n \to X$ is *independent of irrelevant alternatives* (abbreviated IIA) if for every $j \in J = \{1, \ldots, m\}$ and any two profiles $\mathbf{x}$ and $\mathbf{y}$ satisfying $x^i_j = y^i_j$ for all $i \in N$, we have $f_j(\mathbf{x}) = f_j(\mathbf{y})$. This means that the social position on a given issue is determined entirely by the individual positions on that same issue. Viewing profiles as matrices, this means that the aggregation is done column-by-column.

An aggregator $f : X^n \to X$ is *Paretian* if we have $f(\mathbf{x}) = x$ whenever the profile $\mathbf{x}$ is such that $x^i = x$ for all $i \in N$. Note that in the presence of IIA, this is equivalent to demanding that whenever all individuals agree on any one issue, the society adopts this position on that issue.

An aggregator $f : X^n \to X$ is *dictatorial* if there exists an individual $d \in N$ such that $f(\mathbf{x}) = x^d$ for every $\mathbf{x} \in X^n$. That is to say, the society always adopts the dictator's evaluation. A dictatorial aggregator is trivially IIA and Paretian.

We say that X is an *impossibility domain* if for every society N, every IIA and Paretian aggregator for N over X is dictatorial. Otherwise, we say that X is a possibility domain. When dealing with the logic-based framework, we shall refer to an agenda $\mathcal{P}$ as an impossibility or a possibility agenda, meaning that $X(\mathcal{P})$ is an impossibility or a possibility domain, respectively.

[14] By a process similar to that explained in footnote 12, we can pass from an arbitrary truth-functional agenda $\mathcal{P}$ to a standard truth-functional agenda $\mathcal{P}'$. Then, we can check the relevant syntactical conditions (as they appear in our results) on $\mathcal{P}'$, and deduce the corresponding results for $\mathcal{P}$. In fact, this two-step method extends also to truth-functional agendas that are originally expressed in other (two-valued) logics.

For the purpose of characterizing impossibility domains, there is no loss of generality in assuming nondegeneracy in the following sense. We say that X is *nondegenerate*, if for every issue $j \in J$ and every $u \in \{0, 1\}$, there exists $x \in X$ with $x_j = u$. In terms of $\mathcal{P}$, this means that we do not allow it to contain a tautology or a contradiction (this condition is included in the definition of a standard agenda).

We turn now to the presentation of the first condition that appears in our characterization. The condition, named total blockedness, was introduced by Nehring and Puppe (2002). Let X be a nondegenerate subset of $\{0, 1\}^m$. If K is a subset of J, a vector $x = (x_j)_{j \in K} \in \{0, 1\}^K$ with entries for issues in K only is called a K-evaluation. Such a partial evaluation is said to be feasible if it lies in the projection of X on the coordinates in K, and infeasible otherwise. A *minimally infeasible partial evaluation* (abbreviated MIPE) is a K-evaluation $x = (x_j)_{j \in K}$ for some $K \subseteq J$, which is infeasible, but such that every restriction of x to a proper subset of K is feasible. By nondegeneracy, the length of any MIPE (i.e., the size of K) is at least two. We use the MIPEs to construct a directed graph associated with X, denoted by G_X. It has $2m$ vertices, labelled $0_1, 1_1, 0_2, 1_2, \ldots, 0_m, 1_m$. The vertex u_j is to be interpreted as holding the position u on issue j. There is an arc in G_X from vertex u_k to vertex v_ℓ (written $u_k \to v_\ell$) if and only if $k \neq \ell$ and there exists a MIPE $x = (x_j)_{j \in K}$ such that $\{k, \ell\} \subseteq K$ and $x_k = u, x_\ell = \overline{v}$. The interpretation of $u_k \to v_\ell$ is that u_k conditionally entails v_ℓ in the following sense: conditional on holding the positions prescribed in the MIPE x on all issues in $K \backslash \{k, \ell\}$, holding position u on issue k entails holding position v on issue ℓ (since x is infeasible). If $u_k \to v_\ell$ by virtue of a MIPE of length two, then u_k entails v_ℓ in the usual sense, without conditions. Note that the arcs obey the logical law of contrapositives: $u_k \to v_\ell$ if and only if $\overline{v}_\ell \to \overline{u}_k$. We write $u_k \to\to v_\ell$ if there exists a directed path in G_X from u_k to v_ℓ. Finally, we say that X is *totally blocked* if G_X is strongly connected, that is, for any two vertices u_k and v_ℓ we have $u_k \to\to v_\ell$.

The second condition that appears in our characterization comes from linear algebra. The set $\{0, 1\}^m$ may be viewed as a vector space over the field $\{0, 1\}$. In this space, addition is preformed modulo 2, and subtraction is the same as addition. A linear subspace is a nonempty subset closed under addition. An *affine subspace* is a subset obtained from a linear subspace by adding a fixed vector to each of its elements.

Theorem 3.1 (Dokow and Holzman 2005) *Let X be a nondegenerate subset of $\{0, 1\}^m$. Then X is an impossibility domain if and only if X is totally blocked and is not an affine subspace.*

4 A syntactical characterization of impossibility for truth-functional agendas

In this section, we give a necessary and sufficient list of conditions for a standard truth-functional agenda $\mathcal{P}$ to be an impossibility agenda. The conditions are expressed directly in terms of $\mathcal{P}$ (not X, as in Theorem 3.1), and refer to the actual form in which the propositions are written (their syntax).

Before stating the characterization, we introduce each condition separately. First, we associate with $\mathcal{P}$ an undirected graph $G_{\mathcal{P}}$. Its vertex set is $\mathcal{P}_{\mathrm{a}}$, and two atomic propositions p and q are joined by an edge in $G_{\mathcal{P}}$ if and only if there exists a proposition φ in $\mathcal{P}$ so that both p and q appear in φ. We say that $\mathcal{P}$ is *connected* if the graph $G_{\mathcal{P}}$

is connected. For instance, $\mathcal{P} = \{p, q, r, s, (p \wedge \neg r) \vee s, \neg q \wedge r\}$ is connected, but $\mathcal{P} = \{p, q, r, s, p \vee \neg q, r \wedge s, r \vee s\}$ is not.

We say that $\mathcal{P}$ is *heterogeneous* if it either contains at least one OR- AND proposition or, failing that, there is at least one reversal in $\mathcal{P}$, in the following sense: for some $p \in \mathcal{P}_{\mathrm{a}}$, some $u \in \{0, 1\}$, and some $\varphi, \psi \in \mathcal{P}$, either φ is an AND proposition, ψ is an OR proposition, and p^u appears in both; or φ and ψ are both AND propositions or both OR propositions, and p^u appears in φ while $p^{\overline{u}}$ appears in ψ. For instance, $\mathcal{P} = \{p, q, r, p \wedge q, q \vee (p \wedge r)\}$ is heterogeneous and so is $\mathcal{P} = \{p, q, r, p \wedge \neg q, \neg q \vee r\}$, but $\mathcal{P} = \{p, q, r, p \vee \neg q, \neg p \wedge r\}$ is not.

We say that a proposition φ in $\mathcal{P}$ is *parity-dependent* if for some subset $\{p_1, \ldots, p_h\}$ of $\mathcal{P}_{\mathrm{a}}$ and some $u \in \{0, 1\}$, φ is the disjunction of all clauses of the form $p_1^{u_1} \wedge \cdots \wedge p_h^{u_h}$ with $\sum_{j=1}^{h} u_j \equiv u (\mathrm{mod}\ 2)$. We say that $\mathcal{P}$ is parity-dependent if every proposition φ in $\mathcal{P}$ is parity-dependent. For instance, $\mathcal{P} = \{p, q, r, (\neg p \wedge \neg q) \vee (p \wedge q), (q \wedge \neg r) \vee (\neg q \wedge r)\}$ is parity-dependent, but $\mathcal{P} = \{p, q, r, (\neg p \wedge \neg r) \vee (p \wedge r), q \wedge r\}$ is not.

We are now ready to state our characterization.

Theorem 4.1 *Let $\mathcal{P}$ be a standard truth-functional agenda. Then $\mathcal{P}$ is an impossibility agenda if and only if $\mathcal{P}$ is connected, heterogeneous, and not parity-dependent.*

Before proving the theorem, we discuss its relation to some other impossibility results for truth-functional agendas. Pauly and van Hees (2006) and Mongin (2005) placed weaker requirements than ours on the aggregator, from which they derived that it must be dictatorial. Pauly and van Hees weakened the Pareto property to weak responsiveness (meaning that the aggregator is nonconstant), while Mongin weakened the IIA property by requiring it only for atomic propositions. In both cases, the authors required agenda conditions that are stronger than ours, and proved only their sufficiency for impossibility in the respective sense to hold, not their necessity.[15] Nehring and Puppe (2008) placed a stronger requirement than ours on the aggregator, replacing the IIA property by monotone independence (requiring that if a proposition is deemed true by the society for a given profile, and in another profile that proposition is supported by the same or more individuals, then it is still deemed true by the society). For the corresponding notion of impossibility, they gave necessary and sufficient agenda conditions: irreducibility and nonconjunctiveness. These two correspond, respectively, to our conditions of connectedness and heterogeneity.[16]

We will show in the proof that these two conditions on $\mathcal{P}$ amount to total blockedness of the corresponding set $X = X(\mathcal{P})$ of logically consistent evaluations, and parity dependence of $\mathcal{P}$ is equivalent to X being an affine subspace. Theorem 4.1 will then follow from Theorem 3.1.

We break the proof into several claims. In all claims, we assume, without explicitly repeating this, that $\mathcal{P}$ is a standard truth-functional agenda and $X = X(\mathcal{P})$. The

[15] It is interesting to note that for $|\mathcal{P}_{\mathrm{a}}| = \ell \geq 3$ the agenda condition of Pauly and van Hees required that all $2\ell(\ell - 1)$ propositions of the form $p^u \wedge q^v$ belong to $\mathcal{P}$, whereas our theorem shows that as few as $\ell - 1$ of them can suffice for impossibility in our sense. Moreover, our agenda conditions may be satisfied also by propositions that are not of this form, and in fact one OR- AND proposition can suffice.

[16] Nehring and Puppe used different conventions than ours for writing propositions, and phrased their agenda conditions in different terms, which are not purely syntactical like ours. Nevertheless, their conditions are essentially equivalent to our first two conditions.

subscripts used for the coordinates of evaluations and for the vertices of the directed graph G_X (defined in Sect. 3), rather than being the numbers $1, \ldots, m$, will be the corresponding propositions in $\mathcal{P}$.

Claim 4.2 Suppose that φ is a proposition in $\mathcal{P}$, which is not atomic, and p^u is a literal that appears in φ. Then,

1. $u_p \to 1_\varphi$ in G_X.
2. If q^v is another literal that appears in the same clause of φ as p^u, then $u_p \to \overline{v}_q$ in G_X.
3. $1_\varphi \to u_p$ in G_X.
4. If p^u does not appear in every clause of φ, then there exists a literal r^w that appears in φ so that $\overline{u}_p \to w_r$ in G_X.

Proof Let φ_1 be a clause of φ in which p^u appears. Consider the partial evaluation x in which $x_p = u$, for every other q^v that appears in φ_1 we have $x_q = v$, and $x_\varphi = 0$. Clearly x is infeasible, and the minimality of the expression φ implies that x is a MIPE. Parts 1 and 2 of the claim follow from this.

By the minimality of the expression φ, there exists an assignment of truth values to the atomic propositions, so that the value of p is $\overline{u}$, the value of every $q \neq p$ such that q^v appears in φ_1 is v, and φ is false for this assignment. This implies that the $\mathcal{P}_\mathrm{a} \cup \{\varphi\}$-evaluation y, which coincides with this assignment on $\mathcal{P}_\mathrm{a}$ and has $y_\varphi = 1$ is infeasible. Let y' be a restriction of y to some $K \subseteq \mathcal{P}_\mathrm{a} \cup \{\varphi\}$, which is a MIPE. Clearly $\{p, \varphi\} \subseteq K$, and so part 3 of the claim follows. Now, let φ_2 be a clause of φ in which p^u does not appear. As y' is infeasible, it follows that every assignment of truth values to the atomic propositions that agrees with y' on $K \cap \mathcal{P}_\mathrm{a}$ makes φ_2 false. This requires the existence of some $r \in K \cap \mathcal{P}_\mathrm{a}$, $r \neq p$, and some $w \in \{0, 1\}$ so that r^w appears in φ_2 and $y'_r = \overline{w}$. Part 4 of the claim follows. □

Claim 4.3 If $\mathcal{P}$ is heterogeneous, then there exists an atomic proposition $p \in \mathcal{P}_\mathrm{a}$ such that $0_p \to\to 1_p$ and $1_p \to\to 0_p$ in G_X.

Proof Suppose first that $\mathcal{P}$ contains an OR- AND proposition φ. Let φ_1 be a clause of φ that consists of at least two literals. Let φ_2 be another clause of φ. There must exist a literal p^u that appears in φ_1 but not in φ_2, for otherwise φ_2 would be redundant in φ. We show that such a p satisfies the requirements of the claim. Indeed, let q^v be another literal that appears in φ_1. Then, we have $u_p \to \overline{v}_q \to 0_\varphi \to \overline{u}_p$ by parts 2, 3, and 1 of Claim 4.2, respectively (using also the law of contrapositives). In the other direction, for a suitable r^w, we have $\overline{u}_p \to w_r \to 1_\varphi \to u_p$ by parts 4, 1, and 3 of Claim 4.2, respectively.

Next, suppose that $\mathcal{P}$ qualifies as heterogeneous by virtue of a reversal that occurs for the atomic proposition p. The same arguments as above show that if p^u appears in an AND proposition then $u_p \to\to \overline{u}_p$, and if p^u appears in an OR proposition then $\overline{u}_p \to\to u_p$. Applying this to each of the possible kinds of reversal shows that p satisfies the requirements of the claim. □

Claim 4.4 If $\mathcal{P}$ is connected and heterogeneous, then X is totally blocked.

Proof We have to show the existence of directed paths in G_X between any two vertices. By parts 1 and 3 of Claim 4.2 (and the law of contrapositives), vertices that correspond to propositions that are not atomic have arcs joining them in both directions to vertices that correspond to atomic propositions. Therefore, it suffices to show the existence of directed paths between any two vertices that correspond to atomic propositions.

Now, if both q^v and r^w appear in some proposition φ in $\mathcal{P}$, then again by parts 1 and 3 of Claim 4.2 we have $v_q \to\to w_r$ and $w_r \to\to v_q$, and also by taking contrapositives $\overline{v}_q \to\to \overline{w}_r$ and $\overline{w}_r \to\to \overline{v}_q$. Let p be an atomic proposition satisfying the requirements of Claim 4.3. Then it follows by the connectedness of $G_{\mathcal{P}}$ that in G_X any vertex of the form 0_q or 1_q, for $q \in \mathcal{P}_{\mathrm{a}}$, has directed paths joining it in both directions to either 0_p or 1_p. As 0_p and 1_p are themselves joined to each other by directed paths, we are done. □

Claim 4.5 If $\mathcal{P}$ is not connected, then X is not totally blocked.

Proof Suppose $\mathcal{P}$ is not connected. Then there exists a partition of $\mathcal{P}_{\mathrm{a}}$ into two nonempty sets V_1 and V_2 so that no $p \in V_1$ and $q \in V_2$ are joined by an edge in $G_{\mathcal{P}}$. Hence, there exists a partition of $\mathcal{P}$ into two nonempty sets $\mathcal{P}_1$ and $\mathcal{P}_2$ so that propositions in $\mathcal{P}_i$ use only atomic propositions in V_i, $i = 1, 2$. Since propositions in $\mathcal{P}_1$ are logically independent from propositions in $\mathcal{P}_2$, there are no arcs in G_X from any u_φ, $\varphi \in \mathcal{P}_1$, to any v_ψ, $\psi \in \mathcal{P}_2$, or vice versa. Therefore, X is not totally blocked. □

Claim 4.6 If $\mathcal{P}$ is not heterogeneous, then X is not totally blocked.

Proof We produce a partition of the vertex set of G_X into two parts of equal size, V_0 and V_1, by deciding for each pair of vertices 0_φ, 1_φ to put one of them in V_0 and the other in V_1, according to the following rules. If φ is an AND proposition in $\mathcal{P}$, then we put 0_φ in V_0 and 1_φ in V_1. If φ is an OR proposition in $\mathcal{P}$, then we put 0_φ in V_1 and 1_φ in V_0. As $\mathcal{P}$ is not heterogeneous, any remaining φ in $\mathcal{P}$ is a literal, say p^u, and satisfies at most one of the following two conditions: (1) p^u appears in some AND proposition in $\mathcal{P}$ or $p^{\overline{u}}$ appears in some OR proposition in $\mathcal{P}$; (2) p^u appears in some OR proposition in $\mathcal{P}$ or $p^{\overline{u}}$ appears in some AND proposition in $\mathcal{P}$. In case (1), we put 0_φ in V_0 and 1_φ in V_1, while in case (2) we put 0_φ in V_1 and 1_φ in V_0. If neither p^u nor $p^{\overline{u}}$ appears in any AND or OR proposition in $\mathcal{P}$, then we decide arbitrarily, subject to the restriction that if both p^u and $p^{\overline{u}}$ are in $\mathcal{P}$, then we make opposite decisions for them.

Now, suppose that $x = (x_\varphi)_{\varphi \in K}$ is a partial evaluation. For $i = 0, 1$, let K_i be the set of those $\varphi \in K$ for which the corresponding vertex $(x_\varphi)_\varphi$ is in V_i. It can be verified that x is infeasible if and only if there exists $\varphi \in K_0$ so that for every p that appears in φ there exists $\psi \in K_1$ in which p appears. Therefore, if x is a MIPE, we must have $|K_0| = 1$. This implies that there are no arcs in G_X from a vertex in V_0 to one in V_1. Hence, X is not totally blocked. □

Claim 4.7 $\mathcal{P}$ is parity-dependent if and only if X is an affine subspace.

Proof Let $\mathcal{P}_{\mathrm{a}} = \{p_1, \ldots, p_\ell\}$. By truth-functionality, there exist functions $g_\varphi\colon \{0,1\}^\ell \to \{0,1\}$ for $\varphi \in \mathcal{P}$, so that $x \in X$ if and only if $x_\varphi = g_\varphi(x_{p_1}, \ldots, x_{p_\ell})$

for every $\varphi \in \mathcal{P}$. Now, X is an affine subspace if and only if each of the functions g_φ is affine, that is, of the form $g_\varphi(x_{p_1}, \ldots, x_{p_\ell}) = \sum_{j=1}^{\ell} a_j x_{p_j} + b$ (with coefficients $a_1, \ldots, a_\ell, b \in \{0, 1\}$ which may depend on φ). This is equivalent to saying that for every $\varphi \in \mathcal{P}$ there exist a subset of $\mathcal{P}_{\mathrm{a}}$ (those p_j for which $a_j = 1$) and a parity (determined by b), so that φ is true if and only if the number of true atomic propositions in that subset has that parity. In other words, this amounts to $\mathcal{P}$ being parity-dependent. $\square$

5 All IIA and Paretian aggregators for truth-functional agendas

In this section, we look at the possibility side. If $\mathcal{P}$ is a possibility agenda, then there do exist nondictatorial IIA and Paretian aggregators. Here, we give a full description of those aggregators for any given standard truth-functional agenda $\mathcal{P}$.

First, we introduce some terminology and notation that will be used to describe aggregators. Let $\mathcal{P}$ be an agenda and let $X = X(\mathcal{P})$ be nondegenerate. Suppose that an IIA aggregator $f\colon X^n \to X$ is given. For each proposition $\varphi \in \mathcal{P}$ and each position $u \in \{0, 1\}$, we say that a subset S of N is a u_φ*-winning coalition* if

$$x_\varphi^i = \begin{cases} u & \text{if } i \in S \\ \bar{u} & \text{if } i \in N \backslash S \end{cases} \quad \Rightarrow \quad f_\varphi(\mathbf{x}) = u.$$

Thus, S is u_φ-winning if it prevails on proposition φ when its members, and only they, hold the position u. We denote by $\mathcal{W}_\varphi^u$ the collection of all u_φ-winning coalitions. It follows from the definition and the IIA property that for each φ the two collections $\mathcal{W}_\varphi^0$ and $\mathcal{W}_\varphi^1$ are *dual* to each other, in the sense that $S \in \mathcal{W}_\varphi^0 \Leftrightarrow N \backslash S \notin \mathcal{W}_\varphi^1$. Conversely, if we arbitrarily specify collections of coalitions $\mathcal{W}_\varphi^u$ for every φ and u that satisfy the duality condition, then we have implicitly defined the components $(f_\varphi)_{\varphi \in \mathcal{P}}$. The resulting function f may not map X^n into X, but if it does then it is an IIA aggregator. Note that f is Paretian if and only if $N \in \mathcal{W}_\varphi^u$ for every φ and u.

If a standard truth-functional agenda $\mathcal{P}$ is a possibility agenda, then it violates one or more of the three conditions of Theorem 4.1. For the moment, we will assume that $\mathcal{P}$ is connected (otherwise, we can break the aggregation problem into two or more independent problems and treat each of them separately). If $|\mathcal{P}_{\mathrm{a}}| \geq 2$ and $\mathcal{P}$ is connected, then it cannot violate both the other conditions of Theorem 4.1, because parity dependence allows only literals and OR- AND propositions, and nonheterogeneity rules out the latter. We therefore essentially have only two cases to look at: when only heterogeneity is violated, and when only nonparity dependence is violated.

Let $\mathcal{P}$ be a standard truth-functional agenda that is connected and nonheterogeneous, with $|\mathcal{P}_{\mathrm{a}}| \geq 2$. In the proof of Claim 4.6, we described a partition of the vertex set of G_X into two parts V_0 and V_1. Under the current assumptions of connectedness and $|\mathcal{P}_{\mathrm{a}}| \geq 2$, this partition is uniquely determined (the case when the decision was arbitrary cannot occur). We will use this partition to define the *oligarchic* rules for such $\mathcal{P}$. Let N be a society and let R be a nonempty subset of N. We define the R-oligarchic rule by specifying the collections of coalitions $\mathcal{W}^u_\varphi$ as follows:

$$\mathcal{W}^u_\varphi = \begin{cases} \{S \subseteq N \mid S \cap R \neq \emptyset\} & \text{if } u_\varphi \in V_0 \\ \{S \subseteq N \mid R \subseteq S\} & \text{if } u_\varphi \in V_1 \end{cases}$$

Thus, for every proposition φ, there is a position (that u for which $u_\varphi \in V_1$) whose adoption by the society requires the support of all members of R, and in the absence of such support, the opposite position is adopted.[17] The individuals in $N \backslash R$ are ignored. The two extreme cases, $R = N$ and $R = \{d\}$, correspond to unanimity rule and dictatorial rule, respectively.

Theorem 5.1 *Let $\mathcal{P}$ be a standard truth-functional agenda that is connected and nonheterogeneous, with $|\mathcal{P}_a| \geq 2$. Let $X = X(\mathcal{P})$ and let N be a society of n individuals. Then for every nonempty $R \subseteq N$, the R-oligarchic rule defines an IIA and Paretian aggregator $f\colon X^n \to X$. Conversely, every IIA and Paretian aggregator $f\colon X^n \to X$ is of this form, for some nonempty $R \subseteq N$.*

Proof To show that the R-oligarchic rule defines an IIA and Paretian aggregator, it suffices to verify that it maps X^n into X. Let f be the function whose components f_φ are implicitly defined by the collections $\mathcal{W}^u_\varphi$ above. Suppose, for the sake of contradiction, that $f(\mathbf{x}) \notin X$ for some $\mathbf{x} \in X^n$. Then some restriction of $f(\mathbf{x})$, say $y = (y_\varphi)_{\varphi \in K}$, is a MIPE. As shown in the proof of Claim 4.6, there is a unique $\varphi_0 \in K$ so that $(y_{\varphi_0})_{\varphi_0} \in V_0$. There is an individual $i \in R$ for whom $x^i_{\varphi_0} = y_{\varphi_0}$, and for this individual we have $x^i_\varphi = y_\varphi$ for all $\varphi \in K$, contradicting the feasibility of x^i.

Conversely, suppose that $f\colon X^n \to X$ is any IIA and Paretian aggregator, with associated collections $\mathcal{W}^u_\varphi$ of winning coalitions. We have to show that these collections coincide with those defined above for a suitable nonempty $R \subseteq N$. It follows from the definition of V_0, V_1, from parts 1 and 3 of Claim 4.2, and from the connectedness of $\mathcal{P}$ that any two vertices of G_X that lie in the same V_i are joined in G_X by a directed path. Therefore, by Claim 3.1 in Dokow and Holzman (2005), any two such vertices have the same collection $\mathcal{W}^u_\varphi$. Let $\mathcal{W}_0$ be the common collection for V_0, that is, $\mathcal{W}^u_\varphi = \mathcal{W}_0$ for every $u_\varphi \in V_0$. We have to show that $\mathcal{W}_0 = \{S \subseteq N \mid S \cap R \neq \emptyset\}$ for some nonempty $R \subseteq N$. This will suffice, because the collection for V_1 must be dual to this collection.

As $\mathcal{P}$ is nonheterogeneous and connected, and $|\mathcal{P}_a| \geq 2$, there is at least one proposition φ in $\mathcal{P}$ that is an AND or an OR proposition. We present the argument for the case when we have φ in $\mathcal{P}$ of the form $p_1 \wedge \cdots \wedge p_h$. By suitably interchanging 0s and 1s the argument can be adapted to the case when some literals are negated, and also to the case of an OR proposition.

For any two disjoint coalitions S and T with union $U = S \cup T$, consider the construction in Table 1. Note that the individual rows are logically consistent, and the social positions on $p_3, \ldots, p_h$ follow from the Pareto property. Observe also that since 0_{p_1}, 0_{p_2}, and $0_{p_1 \wedge \cdots \wedge p_h}$ are all in V_0, the social positions on p_1, p_2, and

[17] This is the default position according to the terminology of Nehring and Puppe (2008). They proved essentially a monotone version of our Theorem 5.1.

Table 1 Construction for Theorem 5.1

	p_1	p_2	p_3	$\cdots$	p_h	$p_1 \wedge \cdots \wedge p_h$
S	0	1	1	$\cdots$	1	0
T	1	0	1	$\cdots$	1	0
$N \backslash U$	1	1	1	$\cdots$	1	1
			1	$\cdots$	1	

$p_1 \wedge \cdots \wedge p_h$ are 0 if and only if S, T, and U, respectively, belong to $\mathcal{W}_0$. Now the logical consistency of the social evaluation implies that $U \in \mathcal{W}_0$ if and only if either $S \in \mathcal{W}_0$ or $T \in \mathcal{W}_0$ (or both).

Starting from $N \in \mathcal{W}_0$ and using this property repeatedly, we conclude that there exists at least one $i \in N$ so that $\{i\} \in \mathcal{W}_0$. Let R be the set of all such $i \in N$. Then it follows from the above property that $\mathcal{W}_0 = \{S \subseteq N \mid S \cap R \neq \emptyset\}$, as required. □

We turn our attention now to the case when the violated condition of Theorem 4.1 is that of not being parity-dependent. For such $\mathcal{P}$, we define the *parity* rules. Let N be a society and let R be a subset of N of odd cardinality. The R-parity rule maps each profile $\mathbf{x}$ to $\sum_{i \in R} x^i$ (with addition mod 2). Thus, the social position on any proposition φ is the position supported by an odd number of members of R. The individuals in $N \backslash R$ are ignored. When R is a singleton, we get the dictatorial rule.

Theorem 5.2 *Let $\mathcal{P}$ be a standard truth-functional agenda that is connected and parity-dependent, with $|\mathcal{P}_a| \geq 2$. Let $X = X(\mathcal{P})$ and let N be a society of n individuals. Then for every $R \subseteq N$ of odd cardinality, the R-parity rule defines an IIA and Paretian aggregator $f\colon X^n \to X$. Conversely, every IIA and Paretian aggregator $f\colon X^n \to X$ is of this form, for some $R \subseteq N$ of odd cardinality.*

Proof As noted above, since $\mathcal{P}$ is connected and parity-dependent and $|\mathcal{P}_a| \geq 2$, it is also heterogeneous. Hence, by Claims 4.4 and 4.7, X is totally blocked and is an affine subspace. For such X, the assertions of the theorem were proved in Proposition 4.3 of Dokow and Holzman (2005). □

By applying the last two theorems to each connected component, we obtain the following complete description of the available IIA and Paretian aggregators for any given standard truth-functional agenda $\mathcal{P}$.

Corollary 5.3 *Let $\mathcal{P}$ be a standard truth-functional agenda. Let $\mathcal{P}_a^1, \ldots, \mathcal{P}_a^\ell$ be the partition of $\mathcal{P}_a$ into connected components of the graph $G_{\mathcal{P}}$. Let $\mathcal{P}^1, \ldots, \mathcal{P}^\ell$ be the corresponding partition of $\mathcal{P}$. Let N be a society. The class of all IIA and Paretian aggregators for N over $\mathcal{P}$ is the Cartesian product of the corresponding classes for $\mathcal{P}^1, \ldots, \mathcal{P}^\ell$, which are described in the following mutually exclusive and exhaustive list of cases (where f^k denotes the aggregator over $\mathcal{P}^k$, $k = 1, \ldots, \ell$):*

1. *If $\mathcal{P}^k$ is heterogeneous and not parity-dependent, then f^k is the dictatorial rule for some dictator $d^k \in N$.*

2. *If* $|\mathcal{P}_{\mathrm{a}}^k| \geq 2$ *and* $\mathcal{P}^k$ *is not heterogeneous, then* f^k *is the* R^k*-oligarchic rule for some nonempty* $R^k \subseteq N$.
3. *If* $|\mathcal{P}_{\mathrm{a}}^k| \geq 2$ *and* $\mathcal{P}^k$ *is parity-dependent, then* f^k *is the* R^k*-parity rule for some* $R^k \subseteq N$ *of odd cardinality.*
4. *If* $|\mathcal{P}_{\mathrm{a}}^k| = 1$*, then* f^k *is an arbitrary Paretian rule.*

From this description, we can easily derive the conditions under which there exist nondictatorial, IIA and Paretian aggregators that satisfy some additional desirable properties. We do so below for several such properties. When stating the conditions we refer to the components of $\mathcal{P}$ as being of types 1, 2, 3, or 4, according to the four cases listed in Corollary 5.3.

We say that an aggregator f gives *no veto power*, if we have $f(\mathbf{x}) = x$ whenever the profile $\mathbf{x}$ is such that $x^i = x$ for at least $|N| - 1$ individuals $i \in N$. This is a common strengthening of the Pareto property and of nondictatorship.

Corollary 5.4 *Under the conditions of Corollary 5.3, there exists an IIA aggregator for N over* $\mathcal{P}$ *that gives no veto power if and only if* $|N| \geq 3$ *and every component of* $\mathcal{P}$ *is of type 4, or equivalently every proposition in* $\mathcal{P}$ *is a literal.*

In particular, aggregation by majority rule on each proposition works only in the trivial case when every proposition is a literal. We recall that the starting point of research on judgment aggregation was the observation, called the doctrinal paradox, that aggregation by propositionwise majority may fail to preserve logical consistency. Our conclusion is that the doctrinal paradox occurs for *every* nontrivial truth-functional agenda.

We say that an IIA aggregator f is *monotone* if all of the associated collections of winning coalitions $\mathcal{W}_\varphi^u$ are closed under taking supersets.[18]

Corollary 5.5 *Under the conditions of Corollary 5.3, there exists an IIA and Paretian aggregator for N over* $\mathcal{P}$ *that is monotone and nondictatorial if and only if* $|N| \geq 2$ *and either* (1) $\mathcal{P}$ *is connected and the unique component is of types 2 or 4, or* (2) $\mathcal{P}$ *is not connected.*

Finally, we turn to symmetry properties of the aggregator. We say that an IIA aggregator f is *neutral* if it treats the various propositions and their negations equally, that is, all of the associated collections of winning coalitions $\mathcal{W}_\varphi^u$ coincide.[19] We say that an aggregator f is *anonymous* if it treats the individuals equally, that is, $f(\mathbf{x})$ is invariant under permutations of the individuals in N. This is a far-reaching strengthening of nondictatorship.

Corollary 5.6 *Under the conditions of Corollary 5.3, and assuming* $|N| \geq 2$*, there exists an IIA and Paretian aggregator for N over* $\mathcal{P}$ *that is…*

1. *…neutral and nondictatorial if and only if* $|N| \geq 3$ *and every component of* $\mathcal{P}$ *is of types 3 or 4.*

---

18 This is the extra property that Nehring and Puppe (2008) assumed. Our corollary agrees with their result.

19 This property was assumed, under the name systematicity, in a number of papers on judgment aggregation.

2. *…anonymous if and only if either (1) every component of $\mathcal{P}$ is of types 2 or 4, or (2) $|N|$ is odd and every component of $\mathcal{P}$ is of types 2, 3, or 4.*
3. *…neutral and anonymous if and only if $|N|$ is odd and every component of $\mathcal{P}$ is of types 3 or 4.*

6 Relaxing the Pareto condition for truth-functional agendas

In this section, we show that if the Pareto condition is replaced by a weaker requirement of sovereignty, our main results remain valid with only minor modifications. In doing so, we are motivated by some of the literature on aggregation of preferences and on judgment aggregation, which either dropped the Pareto condition or relaxed it. In particular, Wilson (1972) showed that dropping the Pareto condition in Arrow's impossibility theorem results in adding to the dictatorial rules only the antidictatorial rules (where there is an individual whose preferences are reversed by the society) and the imposed rules (yielding a fixed social preference).

Sovereignty means that the aggregator should not rule out any outcome. Given an agenda $\mathcal{P}$ and $X = X(\mathcal{P})$, there are two natural requirements of this kind for an aggregator $f\colon X^n \to X$. We say that f is *globally surjective* if its image is the entire set X. We say that f is *locally surjective* if the image of each of its components f_φ is the entire set $\{0, 1\}$. Clearly, under our nondegeneracy assumption on X, global surjectivity implies local surjectivity. Conversely, for truth-functional agendas and IIA aggregators, local surjectivity implies global surjectivity (to see this, start by suitably choosing the individual positions on the premises). Thus, under the assumptions that we maintain here, the two versions of surjectivity are equivalent. We will henceforth use the term *surjective* without danger of ambiguity.

We start with an example showing that our impossibility results are not preserved in their current form if the Pareto condition is replaced by the weaker condition of surjectivity.

Example Let $\mathcal{P} = \{p, q, r, (\neg p \wedge \neg q) \vee (p \wedge r), (q \wedge r) \vee (\neg q \wedge \neg r)\}$. By Theorem 4.1, this is an impossibility agenda. Now consider the following rule of aggregation for a society N. Fix an individual in N, say 1, and let the social position on each of the propositions p and $(q \wedge r) \vee (\neg q \wedge \neg r)$ be that of individual 1, and the social position on each of the propositions q, r, and $(\neg p \wedge \neg q) \vee (p \wedge r)$ be the opposite of individual 1's position. Note that under the operation of switching the truth values of q and r (while keeping that of p), the propositions in the former list always preserve their truth values, while those in the latter list always reverse their truth values. This implies that the above rule always yields a logically consistent evaluation. Thus, we have an IIA and surjective aggregator for this $\mathcal{P}$, which is not dictatorial.

We proceed to define in general aggregators of the form described in the above example. Let $\mathcal{P}$ be an agenda, and let $\mathcal{R}$ be a subset of $\mathcal{P}$. Let $X = X(\mathcal{P})$ and let N be a society. An aggregator $f\colon X^n \to X$ is *$\mathcal{R}$-reverse dictatorial* if there exists an individual $d \in N$ such that for every $\varphi \in \mathcal{P}$ and every $\mathbf{x} \in X^n$

$$f_\varphi(\mathbf{x}) = \begin{cases} x_\varphi^d & \text{if } \varphi \in \mathcal{P}\backslash\mathcal{R} \\ \overline{x_\varphi^d} & \text{if } \varphi \in \mathcal{R} \end{cases}$$

In the case $\mathcal{R} = \emptyset$ we get the dictatorial rule, and the case $\mathcal{R} = \mathcal{P}$ is the analog of the antidictatorial rule in our framework. Clearly, an $\mathcal{R}$-reverse dictatorial aggregator is well defined (that is, maps X^n into X) if and only if X is closed under the operation of complementing the entries in the $\mathcal{R}$ coordinates. For some agendas $\mathcal{P}$, the only well-defined $\mathcal{R}$-reverse dictatorial aggregators are those with $\mathcal{R} = \emptyset$, but for other agendas, such as the one in the above example, nontrivial $\mathcal{R}$-reverse dictatorial aggregators do exist. For an agenda $\mathcal{P}$ satisfying $\mathcal{P}_a \subseteq \mathcal{P}$, an $\mathcal{R}$-reverse dictatorial aggregator with $\mathcal{R} \neq \emptyset$ will exist if and only if there exists a nonempty subset $\mathcal{R}_a$ of $\mathcal{P}_a$ with the following property: with respect to the operation of switching the truth values of the atomic propositions in $\mathcal{R}_a$ (while keeping those of the remaining atomic propositions), every proposition in $\mathcal{P}$ has a constant reaction—either it always preserves its truth value, or it always reverses its truth value. Finally, we observe that whenever an $\mathcal{R}$-reverse dictatorial aggregator is well defined, it is IIA and surjective.

Theorem 6.1 *Let $\mathcal{P}$ be a standard truth-functional agenda that is connected and nonparity-dependent. Let $X = X(\mathcal{P})$ and let N be a society of n individuals. Let $f\colon X^n \to X$ be an IIA and surjective aggregator. Then,*

1. *If $\mathcal{P}$ is heterogeneous, then f is an $\mathcal{R}$-reverse dictatorial aggregator for some $\mathcal{R} \subseteq \mathcal{P}$.*
2. *If $\mathcal{P}$ is not heterogeneous, then f is an R-oligarchic rule for some nonempty $R \subseteq N$.*

Part 1 of this theorem is the counterpart of Theorem 4.1 when the Pareto condition is replaced by surjectivity. It shows that, under the same conditions on $\mathcal{P}$, this relaxation enlarges the class of available IIA aggregators from the dictatorial rules to the $\mathcal{R}$-reverse dictatorial rules.[20] As explained above, whether or not this is a proper enlargement depends on the agenda $\mathcal{P}$. For many agendas $\mathcal{P}$, there will be no $\mathcal{R}$-reverse dictatorial aggregators with $\mathcal{R} \neq \emptyset$, and in these cases, the original impossibility carries over with the Pareto condition replaced by surjectivity. Part 2 of the theorem is the counterpart of Theorem 5.1. It shows that, under the same conditions on $\mathcal{P}$, weakening the Pareto condition to surjectivity makes no difference for the class of available IIA aggregators.[21]

We introduce now some notation and terminology that will be used in the proof of Theorem 6.1. Let $f : X^n \to X$ be as in the theorem. We associate with f the

[20] It is interesting to note that the result does not remain true if instead of surjectivity we assume the even weaker property of nonconstancy [this is Pauly and van Hees' (2006) weak responsiveness, which sufficed for their result under a much stronger agenda condition]. Consider for example the agenda $\mathcal{P} = \{p, q, p \wedge q, p \wedge \neg q\}$, which satisfies all the conditions of Theorem 6.1, part 1. We may define an aggregator $f\colon X^n \to X$ by letting f_q be any nonconstant function that depends only on the individual evaluations of q, and letting f_φ be the constant function 0 for every other φ in $\mathcal{P}$. Then f is well defined, IIA, and nonconstant, but it is not $\mathcal{R}$-reverse dictatorial for any $\mathcal{R} \subseteq \mathcal{P}$.

[21] We observe that the initial impossibility theorem for judgment aggregation, due to List and Pettit (2002), can be deduced from this result. Indeed, suppose that $\mathcal{P} \supseteq \{p, q, p \wedge q\}$ and $f\colon X^n \to X$, $n \geq 2$, is an IIA aggregator that is anonymous and neutral. We may assume that $\mathcal{P} = \{p, q, p \wedge q\}$, for otherwise we can consider the restriction of f to evaluations of only those three propositions. Clearly, neutrality implies surjectivity of f (this is obvious if one considers local surjectivity). Hence, by Theorem 6.1, part 2, f is an R-oligarchic rule. However, if $|R| < n$ then such a rule is not anonymous, and if $|R| > 1$ then it is not neutral.

collections $\mathcal{W}^u_\varphi$ of winning coalitions, just as we did in Sect. 5. The only difference is that now we do not know that $N \in \mathcal{W}^u_\varphi$ for every φ and u. We only know that $\mathcal{W}^u_\varphi \neq \emptyset$ for every φ and u (this is the surjectivity condition).

We partition the agenda $\mathcal{P}$ into four (possibly empty) parts, according to the behavior of f_φ on unanimous n-tuples of positions on φ. The four parts are denoted $\mathcal{P}_{00}$, $\mathcal{P}_{01}$, $\mathcal{P}_{10}$, and $\mathcal{P}_{11}$, and are defined by the following convention: we put a proposition φ in $\mathcal{P}_{uv}$ if u is the output of f_φ on the all 0 n-tuple and v is its output on the all 1 n-tuple. Clearly, f is Paretian if and only if $\mathcal{P}_{01} = \mathcal{P}$.

Our arguments will make use of a new aggregator, denoted by f^2, obtained from f by applying it twice in the following sense: $f^2(\mathbf{x}) = f(f(\mathbf{x}), \ldots, f(\mathbf{x}))$. Clearly, $f^2 \colon X^n \to X$ is a well-defined IIA aggregator. We may consider a partition of $\mathcal{P}$ into four parts, denoted $\mathcal{P}^2_{00}$, $\mathcal{P}^2_{01}$, $\mathcal{P}^2_{10}$, and $\mathcal{P}^2_{11}$, which are defined in the same way as above but with respect to f^2. Note that if φ is in one of the parts $\mathcal{P}_{00}$, $\mathcal{P}_{01}$, or $\mathcal{P}_{11}$, then it is in the same corresponding part with respect to f^2. However, if φ is in $\mathcal{P}_{10}$ then it is in $\mathcal{P}^2_{01}$. It follows that if $\mathcal{P}_{00} \cup \mathcal{P}_{11} = \emptyset$ then f^2 is Paretian. The plan of the proof of Theorem 6.1 is to show that indeed, under its assumptions, $\mathcal{P}_{00} \cup \mathcal{P}_{11} = \emptyset$, and then apply our earlier results to f^2.

In preparation for the proof of the theorem, we prove a number of claims. In all these claims, the agenda $\mathcal{P}$ is assumed to be standard truth-functional, the aggregator f is assumed to be IIA and surjective, and φ is a proposition in $\mathcal{P}$. We also use the following terminology. If the literal p^u appears in a clause φ_1 of φ, we say that p agrees to φ_1 if $p \in \mathcal{P}_{uu}$, and we say that p is opposed to φ_1 if $p \in \mathcal{P}_{\overline{uu}}$.

Claim 6.2 If $\varphi \in \mathcal{P}_{00}$, then for every atomic proposition p that appears in φ, there exists a clause φ_1 of φ to which p is opposed, and so that every other atomic proposition q that appears in φ_1 agrees to it.

Proof Let $\varphi \in \mathcal{P}_{00}$. We first show that in every clause φ_1 of φ there exists an atomic proposition r that is opposed to φ_1. Suppose, for the sake of contradiction, that φ_1 is a clause of φ to which no r is opposed. So, if r^u appears in φ_1, then $r \notin \mathcal{P}_{\overline{uu}}$, and therefore $r \in \mathcal{P}^2_{uu} \cup \mathcal{P}^2_{01}$. On the other hand, $\varphi \in \mathcal{P}^2_{00}$. Let x be a logically consistent evaluation under which φ_1 is true. Let $\mathbf{x}$ be the profile in which every individual holds the evaluation x. Then in $f^2(\mathbf{x})$, for every r^u that appears in φ_1, the truth value of r is u, whereas φ is false. This contradicts the consistency of $f^2(\mathbf{x})$.

Now we prove the assertion of the claim. Let p be an atomic proposition that appears in φ. By the minimality of the expression φ, there exist two logically consistent evaluations x and y that coincide on every atomic proposition other than p, so that $x_\varphi = 1$ but $y_\varphi = 0$. By surjectivity of f, we can choose some coalition $S \in \mathcal{W}^1_\varphi$. Let $\mathbf{z} = (x^S, y^{N \setminus S})$, that is, the profile in which each member of S holds the evaluation x, and each member of $N \setminus S$ holds the evaluation y. Then, $f_\varphi(\mathbf{z}) = 1$ by our choice of S. Hence, there exists a clause φ_1 of φ, which is true under $f(\mathbf{z})$. We show that φ_1 satisfies the requirements of the claim.

Note first that for every $q \neq p$ that appears in φ_1, all individuals hold the same position on q in $\mathbf{z}$, and therefore, as φ_1 is true under $f(\mathbf{z})$, such q cannot be opposed to φ_1. By the first part of the proof, some r is opposed to φ_1, and hence, p must be opposed to φ_1. Now, let q be any other atomic proposition that appears in φ_1. By what we have shown so far (applied to q instead of p), there exists a clause φ_2 of φ to which q is opposed. Hence, q either agrees to φ_1 or is opposed to it, but, as we have pointed out, q cannot be opposed to φ_1; so, it must agree to it. □

Claim 6.3 Suppose $\varphi \in \mathcal{P}_{00} \cup \mathcal{P}_{11}$. Then

1. Every atomic proposition that appears in φ is in $\mathcal{P}_{00} \cup \mathcal{P}_{11}$.
2. φ is parity-dependent.

Proof We will show that the assertions in the claim hold true in the case $\varphi \in \mathcal{P}_{00}$. The case $\varphi \in \mathcal{P}_{11}$ can be reduced to the other case as follows. Let φ' be a proposition in minimal DNF that is logically equivalent to $\neg\varphi$. Consider the agenda $\mathcal{P}' = \mathcal{P} \cup \{\varphi'\}$. This agenda satisfies all our assumptions, and the aggregator f can be naturally extended to an IIA and surjective aggregator f' for $\mathcal{P}'$. If $\varphi \in \mathcal{P}_{11}$ then $\varphi' \in \mathcal{P}'_{00}$, and once we know that the assertions of the claim hold true for φ', we can deduce that they hold true for φ as well.

So, let $\varphi \in \mathcal{P}_{00}$. Part 1 of the claim is implied directly by Claim 6.2. To verify the second part, it suffices to show that switching the truth value of any atomic proposition p that appears in φ (while keeping the values of all others) always changes the truth value of φ. Suppose, for the sake of contradiction, that this fails for some p that appears in φ. Let φ_1 be a clause of φ that satisfies the requirements of Claim 6.2 for p. Let u be such that p^u appears in φ_1, and let S be a coalition in $\mathcal{W}_p^u$. By our assumption, there exist two logically consistent evaluations x and y that coincide on every atomic proposition other than p and on φ, and satisfy $x_p = u$ and $y_p = \overline{u}$. Let $\mathbf{z} = (x^S, y^{N\setminus S})$. Then, $f_p(\mathbf{z}) = u$ by our choice of S. This, together with the fact from Claim 6.2 that every atomic proposition $q \neq p$ that appears in φ_1 agrees to it, implies that φ_1 is true under $f(\mathbf{z})$. However, as $\varphi \in \mathcal{P}_{00}$, we get $f_\varphi(\mathbf{z}) = 0$, which contradicts the logical consistency of $f(\mathbf{z})$. □

Claim 6.4 If $\varphi \notin \mathcal{P}_{00} \cup \mathcal{P}_{11}$, then no atomic proposition that appears in φ is in $\mathcal{P}_{00} \cup \mathcal{P}_{11}$.

Proof Suppose that $\varphi \notin \mathcal{P}_{00} \cup \mathcal{P}_{11}$. We first show that for no clause φ_1 of φ it is the case that one of its atomic propositions, say r, agrees to φ_1, and none of its atomic propositions is opposed to it. Suppose, for the sake of contradiction, that this is the case for φ_1 of the form $r^u \wedge q_1^{u_1} \wedge \cdots \wedge q_h^{u_h}$. By the minimality of the expression φ, there exists a logically consistent evaluation x so that $x_r = \overline{u}$, $x_{q_j} = u_j$ for $j = 1, \ldots, h$, and $x_\varphi = 0$. Let $\mathbf{x}$ be the profile in which every individual holds the evaluation x. It follows from our assumptions that under $f^2(\mathbf{x})$ the clause φ_1 is true, but φ is false, which contradicts the logical consistency of $f^2(\mathbf{x})$.

Now we proceed to prove the assertion of the claim. Let p be an atomic proposition that appears in some clause φ_1 of φ. We have to show that p neither agrees to φ_1 nor is opposed to it. Suppose first that p is opposed to φ_1. By the minimality of the expression φ, there exists a logically consistent evaluation x under which φ_1 is the unique clause of φ that is true. Let $\mathbf{x}$ be the profile in which every individual holds the evaluation x. It follows from our assumptions that under $f^2(\mathbf{x})$ the clause φ_1 is false, yet φ is true. Hence, there must exist another clause φ_2 of φ that is true under $f^2(\mathbf{x})$. Given that φ_2 is false under x, for it to become true under $f^2(\mathbf{x})$, it must be the case that one of its atomic propositions agrees to φ_2 and none of its atomic propositions is opposed to it. This, however, contradicts what we showed earlier.

Next, suppose that p agrees to φ_1. We already know, by the previous paragraph, that none of the atomic propositions in φ_1 is opposed to it. Again, this contradicts what we showed earlier. $\square$

We are now ready to prove Theorem 6.1. We show first that its assumptions imply that $\mathcal{P}_{00} \cup \mathcal{P}_{11} = \emptyset$. Indeed, suppose that $\mathcal{P}_{00} \cup \mathcal{P}_{11} \neq \emptyset$. Note that if $\mathcal{P}_{00} \cup \mathcal{P}_{11} = \mathcal{P}$, then, by part 2 of Claim 6.3, $\mathcal{P}$ is parity-dependent, contradicting our assumption. Hence, there exist propositions in $\mathcal{P}$, which are in $\mathcal{P}_{00} \cup \mathcal{P}_{11}$, and there exist propositions in $\mathcal{P}$, which are not in $\mathcal{P}_{00} \cup \mathcal{P}_{11}$. By Claim 6.3, part 1, and Claim 6.4, there exist atomic propositions, which are in $\mathcal{P}_{00} \cup \mathcal{P}_{11}$, and there exist atomic propositions, which are not in $\mathcal{P}_{00} \cup \mathcal{P}_{11}$. Moreover, no proposition in $\mathcal{P}$ can contain atomic propositions of both kinds. This contradicts our assumption that $\mathcal{P}$ is connected.

Now, suppose that $\mathcal{P}$ is heterogeneous. As remarked earlier, it follows from $\mathcal{P}_{00} \cup \mathcal{P}_{11} = \emptyset$ that f^2 is Paretian. Hence, by Theorem 4.1, f^2 is dictatorial, but then f is $\mathcal{P}_{10}$-reverse dictatorial.

Next, suppose that $\mathcal{P}$ is not heterogeneous. By considering the action of f on profiles in which all individuals hold the same evaluation, we know that X is closed under the operation of complementing the entries in the $\mathcal{P}_{10}$ coordinates. However, since every atomic proposition appears in some AND or OR proposition in $\mathcal{P}$, there is no way for X to be closed under that operation unless $\mathcal{P}_{10} = \emptyset$. We conclude that $\mathcal{P}_{01} = \mathcal{P}$ and f is Paretian. Hence, by Theorem 5.1, f is R-oligarchic for some nonempty $R \subseteq N$.

References

Arrow KJ (1951) Social choice and individual values. Wiley, New York

Dietrich F (2006) Aggregation theory and the relevance of some issues to others. Working paper, London School of Economics

Dietrich F (2007) A generalised model of judgment aggregation. Soc Choice Welf 28:529–565

Dietrich F, List C (2007) Arrow's theorem in judgment aggregation. Soc Choice Welf 29:19–33

Dokow E, Holzman R (2005) Aggregation of binary evaluations. Working paper, Technion-Israel Institute of Technology (forthcoming in J Econ Theory)

Dokow E, Holzman R (2006) Aggregation of binary evaluations with abstentions. Working paper, Technion-Israel Institute of Technology

Fishburn PC, Rubinstein A (1986b) Aggregation of equivalence relations. J Classif 3:61–65

List C, Pettit P (2002) Aggregating sets of judgments: an impossibility result. Econ Philos 18:89–110

List C, Puppe C (2007) Judgment aggregation: a survey. Working paper, London School of Economics (forthcoming in: Anand P, Pattanaik P, Puppe C (eds) The handbook of rational and social choice. Oxford University Press)

Mongin P (2005) Factoring out the impossibility of logical aggregation. Working paper, École Polytechnique (forthcoming in J Econ Theory)

Nehring K, Puppe C (2002) Strategy-proof social choice on single-peaked domains: possibility, impossibility and the space between. Working paper, University of California at Davis

Nehring K, Puppe C (2008) Consistent judgement aggregation: the truth-functional case. Soc Choice Welf 31:41–57

Pauly M, van Hees M (2006) Logical constraints on judgment aggregation. J Philos Log 35:569–585

Rubinstein A, Fishburn PC (1986a) Algebraic aggregation theory. J Econ Theory 38:63–77

Wilson R (1972) Social choice theory without the Pareto principle. J Econ Theory 5:478–486

Wilson R (1975) On the theory of aggregation. J Econ Theory 10:89–99