Aggregating preferences cannot be fair

Francesca Rossi[1], Kristen Brent Venable[1], Toby Walsh[2]
1: Department of Pure and Applied Mathematics, University of Padova, Italy.
Email: {frossi,kvenable}@math.unipd.it
2: NICTA,
School of Computer Science and Engineering,
University of New South Wales, Sydney, Australia.
Email: tw@cse.unsw.edu.au

SOMMARIO/*ABSTRACT*

Preferences typically define a partial ordering over outcomes. A number of formalisms like soft constraints and CP-nets exists to specify such partial orderings. In situations involving multiple agents, we need to combine the preferences of several individuals. In this paper, we consider each agent as voting on whether they prefer one outcome to another. We prove that, under certain conditions on the kind of partial orders that are allowed to express the preferences of the agents and of the result, if there are at least two agents and three outcomes to order, no preference aggregation system can be fair. That is, no preference aggregation system can be free (give any possible result), monotonic (improving a vote for an outcome only ever helps), independent to irrelevant assumptions (the result between two outcomes only depends on how the agents vote on these two outcomes), and non-dictatorial (there is not one agent who is never contradicted). This result generalizes Arrow's impossibility theorem for combining total orders [1].

1 Introduction and Motivation

Constraints and preferences occur in many real-world problems. For example, when rostering nurses, we will have hard constraints (such as "Nobody is allowed to work three consecutive night shifts" and "Each shift must have at least two intensive care nurses") as well as preferences on outcomes (such as "Alice prefers to work with Bob" and "Bob prefers not to work with Carol"). Preferences can be quantitative or qualitative (e.g. "Carol's preference for day shifts is 0.8 and night shifts is 0.2" versus "Carol prefers days shifts to night shifts") as well as conditional or unconditional (e.g. "If it is a night shift, Alice prefers the maternity ward to the ER ward")

A number of formalisms exist for representing preferences. For example, soft constraints [2, 10] can model quantitative preferences whilst CP nets model conditional qualitative preferences [3, 4]. Each provides an ordering on outcomes. In general, this ordering is partial as outcomes may be incomparable. For example, when comparing wines, we might prefer a white wine grape like chardonnay to the sauvignon blanc grape, but we might not want to order chardonnay compared to a red wine grape like merlot. In addition, we will often want to reason about the preferences of multiple agents. For example, each nurse can have different preferences over the shifts. As a second example, when choosing a wine, each person at the table may have different preferences. We therefore need to consider mechanisms for aggregating preferences. The result of aggregating the preferences of multiple agents is itself naturally a partial order. If two outcomes are incomparable for each agent, it is reasonable for them to remain incomparable in the final order. Incomparability can also help us deal with disagreement between the agents. If some agents prefer A to B and others prefer B to A, then it may be best to say that A and B are incomparable. In this paper, we consider this kind of scenario. We assume each agent has a preference ordering on outcomes represented via soft constraints, CP nets or any other mechanism. A preference aggregation procedure then combines these partial orders to produce an overall preference ordering, and this again can be a partial order. The question we address here is: can we combine such preferences fairly? Suppose we have 100 agents, and 50 of them prefer A to B, whilst the other 50 prefer B to A. It might seem reasonable to decide that A is incomparable to B. But what if 99 agents prefer A to B, and only one prefers B to A? Rather than always declare A to be incomparable to B whenever someone votes against A being better to B, is there not a more sophisticated voting system that will decide A is better than B and compensate the agent who preferred B to A on some other outcomes? We will show that, if each agent can order the outcomes via a partial order with unique top and bottom, and if the result has to be a partial order with a unique top or a unique bottom, then any preference aggregation procedure is ultimately unfair. This question has already been addressed in the context of total orders. Arrow's impossibility theorem demonstrates that there is no fair mechanism for combining total orders [1]. We show that this result can be generalized to certain partial orders. In brief, we will say that a preference aggregation system is fair if it has the following four properties:

Freeness: there is no restriction on the result;

Independence to irrelevant alternatives: the relation between A and B in the result depends only on the preference relation between A and B given by the agents (and not on their preferences over other elements);

Monotonicity: if an agent moves up the position of one outcome in his preference ordering, then (all else being equal) such an outcome cannot move down in the resulting preference ordering;

Non-dictatorial: there is no agent such that, no matter what the others vote, she is never contradicted in the result.

This is a straightforward generalization of the definition of fairness used by Arrow on total orders. The main result of this paper is a proof that, if there are at least two agents and three outcomes to order, no preference aggregation system (where preferences are described as said above) can be fair. This result extends a fundamental theorem obtained by Arrow [1], which demonstrates that voting systems in which both the voters' orders and the resulting order are total cannot be fair. This result is both disappointing and a little surprising. By moving from total order to a class of partial orders, we enrich greatly our ability to combine outcomes fairly. If agents disagree on two outcomes, we can always declare them to be incomparable. In addition, a partial ordering can have multiple outcomes which are optimal. Unlike an election, we need not declare a single winner. Nevertheless, we still do not escape the reach of Arrow's theorem. Any voting system will have one or more agents who dictate the overall preferences. However, preference may need more relaxed orders than the one considered in our results. For example CP nets, even acyclic, may produce partial orders with more than one bottom. Soft constraints produce arbitrary partial orders. So we still hope there is a voting semantics for preference aggregation which is fair.

2 Formal background

A preference ordering can be described by a binary relation on outcomes where x is preferred to y iff (x, y) is in the relation. Such relations may satisfy a number of properties. A binary relation R on a set S (that is, $R \subseteq S \times S$) is:

- **reflexive** iff for all $x \in S$, $(x, x) \in R$;
- **transitive** iff for all $x, y, z \in S$, $(x, y) \in R$ and $(y, z) \in R$ implies $(x, z) \in R$;
- **antisymmetric** iff for all $x, y \in S$, $(x, y) \in R$ and $(y, x) \in R$ implies $x = y$;
- **complete** iff for all $x, y \in S$, either $(x, y) \in R$ or $(y, x) \in R$.

A **total order** over satisfies all four of these properties. A total order has an unique **optimal** element, that is an element $o \in S$ such that $\forall x \in S, (o, x) \notin S$. We say that this element is **undominated.**

By comparison, a **partial order** over a set of elements S is a binary relation R on S which is reflexive, transitive and antisymmetric but may not be complete. There may be pairs of elements (x, y) of S which are not in the partial order relation, that is, such that neither $(x, y) \in R$ nor $(y, x) \in R$. Such elements are **incomparable**. Thus any partial order R induces a binary relation $I(R)$ which represents the incomparability among some of the elements of S. In general, given a partial order R, relation $I(R)$ is symmetric, not reflexive, and not transitive. In fact, all pairs (x, x) are in R and thus not in $I(R)$. Also, it could be that $(x, y) \in I(R)$, $(y, z) \in I(R)$ and $(x, z) \in R$. A partial order can have several optimal and mutually incomparable elements. We say that these elements are **undominated**. Given any relation R which is either a total or a partial order, if $(x, y) \in R$, it can be that $x = y$ or that $x \neq y$. If R is such that $(x, y) \in R$ implies $x \neq y$, then R is said to be strict. This means that reflexivity does not hold. Both total and partial orders can be extended to deal with ties, that is, sets of elements which are equally positioned in the ordering. This situation can be described via a binary relation R on the powerset of S rather than on S, such that the subsets are disjoint and they cover the whole S. Thus such subsets form a partition of S. Given R, we will write $Part(R)$ to denote the partition over S induced by R. As with total and partial orders, such a relation can be complete or not. In the following we will call such relations as a total order (resp., partial order) with ties. In a total or partial order with ties, say R, another relation can be derived, which we will call **indifference**: two elements $x, y \in S$ are in the indifference relation $Ind(R)$ iff $x, y \in S1$ and $S1 \in Part(R)$. Summarizing, in a total order with ties, two elements can be either ordered or indifferent. On the other hand, in a partial order with ties, two elements can be either ordered, or indifferent or incomparable. Notice that, while incomparability is not transitive in general, indifference is transitive, reflexive, and symmetric.

The usual strict total order ($<_{\mathbb{Z}}$ and total order ($\leq_{\mathbb{Z}}$) defined on the set of integers $\mathbb{Z}$ are classical examples. The following orderings are, instead respectively a strict partial order and a partial order over pairs of integers. Consider the set of pairs of integers $\mathbb{Z} \times \mathbb{Z}$ and the orders defined as follows: $\langle x, y\rangle <_{\mathbb{Z}\times\mathbb{Z}} \langle z, w\rangle$ iff $x <_{\mathbb{Z}} z$ and $y <_{\mathbb{Z}} w$ and $\langle x, y\rangle \leq_{\mathbb{Z}\times\mathbb{Z}} \langle z, w\rangle$ iff $x \leq_{\mathbb{Z}} z$ and $y \leq_{\mathbb{Z}} w$. It is easy to see that $<_{\mathbb{Z}\times\mathbb{Z}}$ is not reflexive nor complete and is transitive and antisymmetric while $\leq_{\mathbb{Z}\times\mathbb{Z}}$ is reflexive, transitive and antisymmetric but not complete. According to such orderings some elements of $\mathbb{Z} \times \mathbb{Z}$ are ordered (e.g. $\langle 1, 2\rangle <_{\mathbb{Z}\times\mathbb{Z}} \langle 3, 4\rangle$) while others are incomparable (e.g. $\langle 5, 3\rangle$ is incomparable with $\langle 2, 8\rangle$).

As an example of partial order with ties consider the following defined on triples in $\mathbb{Z} \times \mathbb{Z} \times \mathbb{Z}$. Given $\langle x_1, y_1, z_1\rangle$ and $\langle x_2, y_2, z_2\rangle$ in $\mathbb{Z} \times \mathbb{Z} \times \mathbb{Z}$, we say $\langle x_1, y_1, z_1\rangle \leq \langle x_2, y_2, z_2\rangle$ iff $\langle y_1, z_1\rangle \leq_{\mathbb{Z}\times\mathbb{Z}} \langle y_2, z_2\rangle$. In other words we are indifferent to the value of the first component. In this case we will have elements that are respectively ordered , as $\langle 2, 3, 4\rangle < \langle 1, 5, 6\rangle$, or incomparable, as $\langle 2, 3, 4\rangle$ and $\langle 2, 1, 5\rangle$ or indifferent, as $\langle 1, 2, 3\rangle$ and $\langle 2, 2, 3\rangle$.

3 Preferences

A number of formalisms have been proposed for representing and reasoning about preferences of a single agent. Common to all is that they induce some sort of partial or total ordering, possibly with ties, on the outcomes. For example, soft constraints can model quantitative preferences [2, 10]. Each constraint associates a preference value to each assignment of its variables. To model preference ordering and aggregation, the set of possible preference values is the carrier of a semiring, whose two operations state how to order values in the set and how to combine values to obtain new preferences. A complete assignment of values to variables is associated to a preference value by combining the preferences of each partial assignment in each constraint via the combination operation of the semiring. In general, the order induced on the preferences via this approach is a partial order with ties. Assignments with the same preference are naturally interpreted as ties.

Soft constraints can also represent hard statements, as in "I need to be back before 8pm": it is enough to take a set with just two preferences values (that can be interpreted as true and false), order them via logical or (thus true is better than false and we have a total order), and combine them via logical and (so an assignment has preference true if all constraints have preference true, and it is said to be consistent; an assignment has preference false, and it is said to be inconsistent, if some of the constraints have preference false). In this case, the ordering induced over the complete assignments is a total order with ties: all consistent assignments have preference true (thus they are all indifferent) and are better than all inconsistent assignments (which again are indifferent among them).

Another formalism for representing preferences is CP nets [3, 4]. They are a compact mechanism to model conditional qualitative preferences (as in "If I take the fish course, I prefer white wine over red") which satisfy the *ceteris paribus* or "all other things being equal" property. A dependency graph in a CP net states the relation among the features of the problem. Each feature X has a domain of possible values and some parent features $Pa(X)$ on which it depends on: given any complete assignment to $Pa(X)$, CP nets state a total order for the values in the domain of X (in a structure called a CP conditional preference table). Such a total order represents the preference order on the values of X given the values of its parents, all else being equal. A CP net induces an ordering over the complete assignments of all its features: an assignment O is better than another one O' if there is a chain of improving flips from O to O', where an improving flip is a change of the value of one feature that improves the preference according to some preference table in the CP net. Such an ordering is in general partial and does not have ties.

Partial CP nets [9] do not require that all features are ranked. This allows one to represent situations as in "I am indifferent to the color of the car". This means that the ordering induced by a partial CP net over its outcomes is in general a partial ordering with ties. In fact, there could be flips which are neither improving nor worsening, since they change the value of a non-ranked feature.

4 Aggregating preferences

There are many situations in which we need to combine the preferences of multiple agents. A number of mechanisms exist for preference aggregation [5, 11]. One possibility is to run an election in which each agent votes on how they rank every pair of outcomes. In [9], we assume that each agent represents their preferences with a partial CP net and then votes on how outcomes should be ordered. However, the agents can just as easily represent their individual preferences with soft constraints or any other formalism for representing preferences. We need, however, to specify how their votes are collected together into a result.

As in voting theory, the orderings of the agents is called a **profile**. A voting system is then a function mapping profiles onto a result (a partial ordering). In [9], we discuss a number of different voting rules for when agents vote with partial orders.

Pareto: One outcome α is better than another β (written $\alpha \succ_p \beta$) iff every agent says that $\alpha \succ \beta$ or $\alpha \approx \beta$. Two outcomes are incomparable iff they are not ordered either way. An outcome is Pareto optimal iff no other outcome is better.

Majority: One outcome α is majority better than another β (written $\alpha \succ_{maj} \beta$) iff the number of agents which say that α is better than β is greater than the number of agents which say the opposite plus the number of those that say that α and β are incomparable. Two outcomes are majority incomparable iff they are not ordered either way. An outcome is majority optimal iff no other outcome is majority better.

Max: One outcome α is max better than another β (written $\alpha \succ_{max} \beta$) iff more agents vote in favor than against or for incomparability. Two outcomes are max incomparable iff they are not ordered either way. An outcome is max optimal iff no other outcome is max better.

Lex: This rule assumes the agents are ordered in importance. One outcome α is lexicographically better than another β (written $\alpha \succ_{lex} \beta$) iff there exists some agent A such that all agents higher in the order say $\alpha \approx \beta$ and A says $\alpha \succ \beta$. Two outcomes are lexicographically incomparable iff there exists some distinguished agent such that all agents higher in the ordered are indifferent between the two outcomes and the outcomes are incomparable to the distinguished agent. Finally, an outcome is lexicographically optimal iff no other outcome is lexicographically better.

Rank: Each agent gives a numerical rank to each outcome. For example, in a partial CP net, the rank of an outcome is zero if the outcome is optimal, otherwise it is the length of the shortest chain of worsening flips between one of the optimal outcomes and it. We say that one outcome α is rank better than another β (written $\alpha \succ_r \beta$) iff the sum of the ranks assigned to α is smaller than that assigned to β. Two outcomes are rank indifferent iff the sum of the ranks assigned to them are equal. Either two outcomes are rank indifferent or one must be rank better than the other. Finally, an outcome is rank optimal iff no other outcome is rank better.

The Pareto, and Lex rules define strict partial orderings if the agents have a strict partial order, while if the agents have a partial order with ties then these rules define a partial order without ties. The Rank rule, instead, no matter what the agents use to represent their preferences (strict partial order or partial order with ties) defines a total order with ties. Maj and Max are irreflexive and antisymmetric but may be not transitive. However, they all have at least one optimal element. Notice that in all the five rules, except Rank, it is not possible for two outcomes to be indifferent, since we assume that each feature is ranked by at least one of the partial CP nets, while indifference in the qualitative relations (Pareto, Max, Majority, and Lex) means indifference for everybody. In all these voting rules, except Rank, the result of aggregating preferences is itself a partial order.

5 Arrow's impossibility theorem

The voting rules described in the previous section show that one can combine preferences in many different ways, obtaining a result which itself may be a partial or total order. How can we be sure that we are accurately and fairly combining together the agents' preferences? In voting theory, one property of an election which is highly desirable is **fairness**. Given a set of voters and a set of outcomes, and assuming each voter gives a total order with ties of the outcomes and the result is a total order with ties, the fairness of a voting system is defined [1] as the coexistence of the following properties:

Freeness: it is possible to obtain any possible result;

Independence to irrelevant alternatives: the relation between A and B in the resulting ordering depends only on the relation between A and B given by the agents;

Monotonicity: if an agent moves up the position of one outcome in her ordering, then (all else being equal) such an outcome cannot move down in the result;

Non-dictatorial: there is no voter such that, no matter what the others vote, her ordering is the result.

These properties are all very reasonable and desirable also for preference aggregation. Unfortunately, a fundamental result in voting theory is Arrow's impossibility theorem [1] which shows that no voting system on total orders with ties can be fair. In particular, given at least two voters and three outcomes, freeness, monotonicity, and independence of irrelevant assumptions, implies there must be at least one dictator.

A related notion is the concept of **unanimity** for a voting system: if all agents agree about the relation between A and B (A can be either better, worse, or indifferent to B), then the result must agree as well. It is possible to show that monotonicity and independence to irrelevant alternatives imply unanimity, whilst voting systems can be free, unanimous and independent to irrelevant alternatives but not monotonic [11]. Therefore a stronger version of Arrow's result can be obtained by proving that freeness, unanimity and independence of irrelevant assumptions, implies that there must be at least one dictator [6].

6 Fairness of preference aggregation

Arrow's theorem does not directly apply to aggregating preferences as voters are assumed to have a total ordering with ties. As we described earlier, preference orderings may be partial. In addition, they may have other restrictions. For example, CP nets can only represent orderings which decompose into independent conditional CP statements, whereas voters in an election can order their votes in any way. We first generalize the definition of dictatorship to deal with the introduction of incomparability when we go from total to partial orders. We will consider three different ways to define a dictator for a voting system involving partial orders.

Strong dictator: there is a voter such that, no matter what the others vote, her ordering is the result;

Dictator: there is a voter such that, no matter what the others vote, if she says A is better than B then the resulting ordering has A better than B;

Weak dictator: there is a voter such that, no matter what the others vote, if she says A is better than B then the resulting ordering does not have B better than A.

If a strong dictator says A is incomparable to B then A is also incomparable to B in the result. This is not necessarily the case with a dictator or weak dictator. On the other hand, if a weak dictator says A is better than B, then in the result A may be better than B or A incomparable to B. Clearly a strong dictator is a dictator, and a dictator is a weak dictator. Note also that whilst there can only be one strong dictator or dictator, there can be any number of weak dictators.

We first show that the absence of a strong dictator is a very weak property to demand of a voting system in which the resulting ordering is a partial order. Even an "unfair" voting system like the Lex rule, which explicitly favours a particular agent, is not strong dictatorial.

Lemma 1 *Given a set of voters and a set of outcomes, and assuming each voter gives a partial order with ties of the outcomes, and the result is a partial order with ties, a voting system can be free, monotonic, independent of irrelevant assumptions, and not have a strong dictator.*

The Lex rule on partial orders with ties is free, monotonic, and independent to irrelevant assumptions. The absence of a strong dictator might seem to be in contradiction to the nature of a Lex rule as there is a most important agent. However, such an agent does not dictate incomparability, since whatever is left incomparable by this agent can then be ordered by some less important agent. Notice that the first agent in the Lex ordering is always a dictator since if she states that A is better than B then A is better than B in the result. However, there are fairer voting systems on partial orders which are not dictatorial.

Lemma 2 *Given a set of voters and a set of outcomes, and assuming each voter gives a partial order with ties of the outcomes, and the result is a partial order with ties, a voting system can be free, monotonic, independent of irrelevant assumptions, and is not dictatorial.*

For example, the Pareto rule on partial orders with ties is free, monotonic, transitive, independent to irrelevant assumptions, and is not dictatorial. A particular agent can only force the result by stating the incomparability of all possible outcomes. However, this is not considered to be dictatorial according to our definition of dictatorship. Unfortunately, this appears to be as fair as a voting system on partial orders can be. In the next section, we will show that any voting system on partial orders must have one or more weak dictators.

7 Existence of weak dictators

We will prove a more general result, from which follows the existence of weak dictators in voting systems on a certain class of partial orders, and the existence of dictators in voting systems on total orders (Arrow's impossibility theorem). We will show that it is impossible for a voting system where voters give preferences as partial orders with a unique top and bottom, and the result is a partial order with a unique top or unique bottom, to be fair. This result holds also if we restrict the result to be a total order, since a total order is just a special case of a partial order with unique top or bottom. Moreover, the same result holds also when agents use just total orders, since the proof is very similar. Similar arguments can be used to show that these last two results imply Arrow's impossibility theorem. We thus have a lattice of four results.

Let us call uPO a partial order with unique top or bottom, uuPO a partial order with unique top and unique bottom, and TO a total order. Then, let us

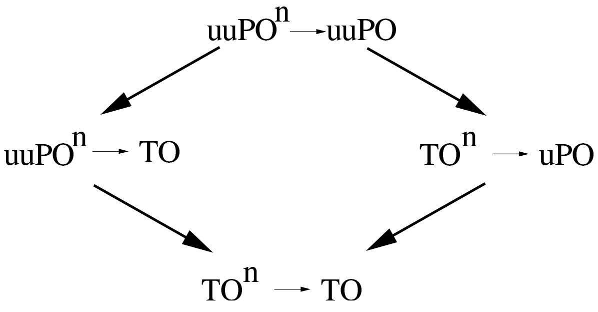

Figure 1: **Lattice of impossibility results.** uPO stands for partial order with unique top or bottom, $uuPO$ stands for partial order with unique top and unique bottom, TO stands for total order. Arrow's theorem applies to $TO^n \mapsto TO$. $\swarrow$ and $\searrow$ stand for the lattice ordering.

denote $A^n \to B$, where $A, B \in \{uPO, uuPO, TO\}$, a voting system where agents can give an ordering of type A and the result can be of type B. Then we have the lattice of impossibility results for the voting systems as described by Figure 1. The lattice ordering can be defined as follows: $A \to B \leq A' \to B'$ iff $A \subseteq A'$ or $B \subseteq B'$.

We show now that the top element of this lattice must be unfair. It follows then by the arguments given above that any other element in the lattice must be unfair.

We will prove that freeness, unanimity, and independence of irrelevant assumptions implies that there must be at least one weak dictator. Since monotonicity and independence of irrelevant alternatives implies unanimity, it follows that freeness, monotonicity, and independence of irrelevant alternative implies that there must be at least one weak dictator.

Theorem 1 *Consider a set of voters and a set of outcomes, and assume that:*

- *each voter gives a uuPO of the outcomes,*
- *the result is a uPO,*
- *there are at least 2 voters and 3 outcomes,*
- *the voting system is free, unanimous, and independent to irrelevant alternatives*

then there is at least one weak dictator.

Proof: The proof is similar in outline to that of Arrow's theorem. However we must adapt each step to this more general context. We first assume the resulting ordering is a PO with a unique bottom. The proof can then be repeated very similarly for the other case in which the resulting ordering is a PO with unique top.

1. First we prove that, if a element B is top or bottom in all uuPOs of the voters, then it must be a top or bottom element in the resulting uPO. If B is not a top or bottom element in the result, there must be other elements A and C such that $A > B > C$. We will now adjust the votes so that C is above A for all voters. Since we have uuPOs, we can always do that while keeping B at the extreme position and not changing the ordering relation between A and B and between C and B.

   By unanimity we must have C above A in the result. By independence, we still have $A > B$ and $B > C$. By transitivity, we have $A > C$ which is a contradiction.

2. There is a pivotal voter n^* such that, when he moves B from bottom to top, the same change happens in the result.

   Assume all voters have B as the bottom. Then, B must be at the bottom in the result by unanimity. Let n^* be the first voter such that, when B moves from bottom to top, this happens in the result. n^* must exist, because when all voters move B from bottom to top, by unanimity in the result we must have B at the top.

3. n^* is a weak dictator for pairs of elements not involving B.

   Let us consider the following scenarios in the context of moving B from the bottom to the top of each voters ranking.

   Π_1: B is still the bottom of n^*. In the result, B is the bottom element so we must have, $A > B$ for all A.

   Π_2: B has been moved to the top of n^*. In the result, B is a top element so we must have, $B > C$ or B incomparable to C for all C.

   Π_3: As in Π_2 but A has now been moved above B in n^* (and thus also above C), and all other voters move freely A and C leaving B in the top or bottom position.

   By independence to irrelevant alternatives, $A > B$ must be the result of Π_3, since all the AB votes are the same as in Π_1. Also, $B > C$ or B incomparable to C must be in the result of Π_3, since all BC votes are the same as in Π_2. By transitivity, the result of Π_3 cannot have $C > A$ since it would imply $C > B$ which is contradictory with the assumption that B and C are either incomparable or $B > C$. Thus n^* is a weak dictator.

4. There exists n' which is a weak dictator for pairs with no C. We can use the same construction as above but with C in place of B.

5. We show now that $n^* = n'$. On total orders, there can be only one dictator so it follows immediately that $n^* = n'$. With partial orders, there can be more than one weak dictator.

   Without loss of generality, assume $n^* \leq n'$. Suppose that $n^* < n'$. Let us consider the following scenarios: we start with all voters having B at the bottom and C at the top. Then we swap B and C in each orderings, going through the voters in the same order as in the previous constructions. When we move B up for n^*, B goes to the top in the result (by the previous part of the proof). C goes down for n^*, and in the result it can go down down as well, in which case we would have $n^* = n'$ and thus a contradiction. Otherwise, C could stay at the top together with B. Recall that the result is a PO with a unique bottom, so we can have several top elements. If this is the case, take any A between B and C in n^*. Now $A > C$ in n^*, so it must be $A > C$ or A incomparable to C in the result, since n^* is a weak dictator for pairs with no B. But if C is at the top, this can happen only if A is either incomparable to C and B, or $A < B$. So, it cannot be $B < A$. Therefore n^* is a weak dictator for pairs not involving C. We can therefore conclude that $n^* = n'$, which is a contradiction with the fact that $n^* < n'$.

□

As with total orders, we can prove that freeness, monotonicity and independence to irrelevant alternatives implies unanimity.

Theorem 2 *Consider a set of voters and a set of outcomes, and assume that:*

- *each voter gives a uuPO of the outcomes,*
- *the result is a uPO with ties,*
- *there are at least 2 voters and 3 outcomes.*
- *the voting system is free, monotonic, and independent to irrelevant alternatives*

then there is at least one weak dictator.

Proof: Suppose the voting system is free and monotonic, and that $A \geq B$ for all voters. If A is moved to the top of the ordering for all voters then, by independence to irrelevant alternatives, this leaves the result between A and B unchanged. Suppose in the result $A < B$ or A is incomparable to B. By monotonicity, any change to the votes of A over B will not help ensure $A \geq B$. Hence, the election cannot be free. Thus it must be $A \geq B$ in the result. The voting system is therefore unanimous. By the last theorem, there must be at least one weak dictator. □

We consider again the five voting rules described in Section 4 and identify which are affected by this result.

- Maj and Max are not transitive, so they do not generate a uPO as required by the theorem;
- Lex has a weak dictator, which is the most important agent;
- Rank is not independent to irrelevant alternatives;
- in the Pareto rule every agent is a weak dictator.

Note that we could consider also another class of voting rules which modifies Pareto by applying the Pareto rule only to a strict subset of the agents, and ignore the rest. The agents in the subset will then all be weak dictators, so also this voting system is not weakly fair.

One way around the limitations of Arrow's theorem is to restrict the type of election to one which can be fair. For example, one of the hypotheses of Arrow's theorem is that there are three or more outcomes. With just two outcomes, a voting system on total orders can be fair. For instance, the majority rule on total orders with a tie-breaking rule for an even number of voters is free, monotonic, independent to irrelevant alternatives and non-dictatorial.

8 Constrained preferences and Social choice

In many situations, we will have constraints as well as preferences. For example, we may wish to find the most preferred roster for the nurses which satisfies all the hard constraints like union rules. We then need to find those feasible outcomes which are undominated. One method to do this is to collect all the feasible outcomes, and order them. From a computational perspective, such a strategy may be very expensive. We will have to collect all the feasible outcomes (which is computationally expensive), ask each agent to compare them (which can be computationally expensive in formalisms like CP-nets), and then combine these votes. In collaboration with Steve Prestwich, we are currently exploring an alternative and computationally more attractive solution in which we compile the preferences into additional hard constraints. The solutions of the compiled problem are guaranteed to be feasible and undominated.

If we do not have any additional constraints, we may only be interested in the best outcome for all the agents. Such a situation can be described by means of a social choice function. A **social choice** function is a mapping from a profile to one outcome, the optimal outcome. The Muller-Satterthwaite theorem is a generalization of Arrow's theorem on total orders which shows that a dictator is inevitable if we have two agents, three or more outcomes and the social choice function collecting votes is unanimous and monotonic [8]. With a partial order, there can be several outcomes which are incomparable and optimal. We can therefore consider a generalization of social choice functions. A **social choices** function f is a mapping from a profile to a non-empty set of outcomes, the optimal outcomes.

We say that a social choices function f is **unanimous** iff when the outcome A is optimal for each agent in Π then $A \in f(\Pi)$, is **monotonic** iff given two profiles Π and Π' in which $A \in f(\Pi)$, and for any B O'_i in Π' ranks A better than B whenever O_i in Π does then $A \in f(\Pi')$, and is **weak dictatorial** iff there is an order O_i in Π and A is optimal in O_i implies $A \in f(\Pi)$. For example, consider the social choices function $f(\Pi)$ which returns the set of optimals for each order O_i in Π. This is unanimous and monotonic. In addition, every agent in the profile is a weak dictator. It is an interesting open question if the Muller-Satterthwaite theorem can be generalized to social choices functions. That is, are weak dictators inevitable if we have two or more agents, three or more outcomes and the social choices function is unanimous and monotonic? A further extension would be to the generalization of the Gibbard-Statterthwaite theorem [7]. That is, are weak dictators inevitable if we have at least two agents and three outcomes, and the social choices function is strategy proof and onto. A social choice function is strategy proof if it is best for each agent to order outcomes as they prefer and not to try to vote tactically.

9 Conclusions and future work

Many real-world problems involve constraints and preferences. Preferences typically define a partial ordering over outcomes. A number of formalisms like soft constraints and CP-nets exists to specify such partial orderings. In situations involving multiple agents, we need a mechanism to combine the preferences of several individuals. We have considered each agent as voting on whether they prefer one outcome to another. We have proved that if there are at least two agents and three outcomes to order, a preference aggregation system cannot be fair if agents use partial orders with a unique top and unique bottom, and the result is a partial order with a unique top or bottom. This result generalizes Arrow's impossibility theorem for combining total orders [1].

Fairness is just one of the desirable properties for a preference aggregation system. As we discussed earlier, an interesting open question is whether voting systems on partial orders can have other desirable properties. For example, can they encourage non-tactical voting? Are they non-manipulable? Another direction is to identify restricted types of preferences which can fairly combined. For example, is there a generalization of "single-peakedness" from total to partial orders which would permit preference aggregation to be fair? A third direction is developing methods to reason about preferences and constraints simultaneously. We are currently exploring mechanisms which compile preferences into additional hard constraints. The solutions of such compiled problems are guaranteed to be feasible and undominated.

10 Acknowledgements

The third author is supported by the Science Foundation Ireland. He wishes to thank Krzysztof Apt for stimulating discussion. This work is partially spported by ASI (Italian Space Agency) under project ARISCOM (Contract I/R/215/02).

BIBLIOGRAFIA

[1] K. Arrow. *Social Choice and Individual Values*. John Wiley and Sons, 1951.

[2] Stefano Bistarelli, Ugo Montanari, and Francesca Rossi. Semiring-based constraint solving and optimization. *Journal of the ACM*, 44(2):201–236, March 1997.

[3] C. Boutilier, R. Brafman, H. Hoos, and D. Poole. Reasoning with conditional ceteris paribus preference statements. In *UAI-99*, pages 71–80, 1999.

[4] C. Domshlak and R. Brafman. Cp-nets - reasoning and consistency testing. In *KR-02*, pages 121–132, 2002.

[5] Jon Doyle. Rational belief revision. In James F. Allen, Richard Fikes, and Erik Sandewall, editors, *KR'91: Principles of Knowledge Representation and Reasoning*, pages 163–174. Morgan Kaufmann, San Mateo, California, 1991.

[6] J. Geanakoplos. Three brief proofs of arrow's impossibility theorem. *Economic Theory*, 2001.

[7] A. Gibbard. manipulation of voting schemes: A general result. *Econimetrica*, 41, 1973.

[8] E. Muller and M.A. Satterthwaite. The equivalence of strong positive association and strategy-proofness. *Econimic Theory*, 14, 1977.

[9] F. Rossi, K. B. Venable, and T. Walsh. mcp nets: representing and reasoning with preferences of multiple agents. In *AAAI-2004*.

[10] T. Schiex. Possibilistic constraint satisfaction problems, or "how to handle soft constraints?". In *8th Conf. of Uncertainty in AI*, pages 269–275, 1992.

[11] A. Sen. *Collective Choice and Social Wellfare*. Holden-Day, 1970.