A Multi-Outcome Result Model

Yoann A. MORELLO [a,1], Agata CIABATTONI [a], and Ilaria CANAVOTTO [b]
[a] *TU Wien (Vienna University of Technology)*
[b] *Maryland University*

Abstract. Formal models of precedential reasoning have focused on binary outcomes. We extend the dimensional result model to settings with an arbitrary set of outcomes and relations among them. Inspired by pairwise approaches to multiclass classification, we associate each pair of outcomes with a binary result model. For each focus case, the constraints induced by these pairwise models are represented by a directed graph over the possible outcomes. We introduce three semantics for this constraint graph and show that, when the outcomes are linearly ordered, all three coincide with the existing model for this setting.

Keywords. precedential reasoning; formal models; multiple outcomes

1. Introduction

Precedential reasoning plays a central role in legal decision making, where decisions in precedent cases constrain decisions in subsequent cases. The *result model* [1,2] captures a minimal view of this process, where only the outcome of a precedent case and its facts matter: a later court is constrained to follow the precedent when the facts before it are at least as strong for the prevailing party as were the facts of the precedent case. More sophisticated models enrich this basic picture by taking into account, for example, the reasons provided in a case [3], or a hierarchical organization of its facts [4,5]. In turn, the facts of a case can be represented either as *binary factors*—legally relevant considerations that favor one party or the other, as in CATO [6]—or as multi-valued dimensions—more nuanced considerations that can take multiple values, with different values favoring one party or the other to different extents, as in HYPO [7].

While existing models allow the facts of a case to be represented in a rich, multi-valued way, they still retain a binary notion of outcome: either the plaintiff or the defendant prevails. This paper investigates how the dimensional result model can be extended to cases with *multiple outcomes*. Such cases arise naturally in law and in other applications of precedential reasoning. One example, studied in [8], is the assessment of recidivism risk, as in the COMPAS system. Here the possible outcomes form a *linear order*, from lower to higher levels of risk, and precedents impose lower or upper bounds on the outcomes that future courts can select. Many decision problems, however, involve multiple outcomes that are not naturally represented by a single linear order. For example, regulatory violations may result in a warning, a fine, or the temporary or permanent suspension of a licence; immigration proceedings may result in removal, withholding of

[1] Corresponding Author: Yoann Morello yoann.morello@gmail.com

removal, asylum, or voluntary departure; and a nanny robot deciding a child's evening activity may send the child to bed or allow them to watch TV, play a video game, or read (example adapted from [9]).

Moving from binary outcomes to general sets of alternatives raises a fundamental question: *what kind of constraints can be derived from precedents decided for possibly unrelated outcomes?* In this paper, we investigate this question without assuming any prior relation among outcomes.

A useful perspective comes from machine learning: the original result model can be reformulated as a binary classifier [10], suggesting multiclass classification for multiple outcomes. While not all standard constructions [11] transfer to precedential reasoning, a natural extension is *pairwise*: for each pair of outcomes, we use the corresponding binary result model, preserving how dimensions support one outcome against another. Each such model generates constraints specifying which of the two outcomes takes priority in the focus case.

Combining these pairwise models may nevertheless yield a relation that cannot be captured by a single ordering of the outcomes. For each focus case, we represent this relation by a directed graph whose nodes are the outcomes and whose edges encode the pairwise constraints generated by precedents. Determining the outcome then requires aggregating these constraints, a problem extensively studied in different contexts, e.g., in multiclass and label ranking [12,13], tournament solutions [14], and choice from incomplete preferences [15,16]. These approaches, however, typically assume additional structure, such as quantified or complete comparisons, or resolve inconsistencies by retaining selected subsets of the constraints. In our setting, comparisons may be incomplete and the constraints are generated by precedents, giving reason to preserve them whenever possible.

We introduce three semantics for the constraint graph arising from the pairwise model, capturing different ways of resolving possibly incomplete and cyclic constraints: (1) a semantics interpreting edges as preferences among outcomes or as obligations to select one outcome over another; (2) the CF2 semantics from abstract argumentation [17]; and (3) a semantics closely related to a choice rule studied in [16]. We study their relationships and show that, when the *outcomes are linearly ordered*, all three coincide with the existing model for this setting [18].

2. Preliminaries

In order to represent cases with multiple outcomes, we start by fixing:

- A finite set $D = \{d_1, d_2, \ldots, d_n\}$ of dimensions, each associated with a set $\mathit{values}(d)$ of possible values.[2]
- A finite set $O = \{o_1, o_2, \ldots, o_m\}$ of outcomes.

We define a *situation* as a function $X : D \to \bigcup_{d \in D} \mathit{values}(d)$ such that $X(d) \in \mathit{values}(d)$ for each $d \in D$; a *case* as a pair $\langle X, o \rangle$, where $o \in O$; and a *case base* as a set of cases. Intuitively, a case base represents the set of precedents that constrain future decisions.

[2]The values of a dimension need not be numerical: for example, the number of disclosures made by the defendant is a numerical dimension in trade secret law, whereas the type of security measure taken by the plaintiff is a non-numerical one.

Before defining the notion of constraint in the general setting, we briefly recall its definition in the binary dimensional result model [19]. In this model, $O = \{\pi, \delta\}$, where π stands for the plaintiff and δ for the defendant. Each dimension $d \in D$ is equipped with a partial order $\leq_d^{\delta\,\pi}$ over $values(d)$, where $p \leq_d^{\delta\,\pi} q$ means that q provides at least as much support for π as p, or equivalently that p provides at least as much support for δ as q. This order is lifted componentwise to fact situations: $Y \leq^{\delta\,\pi} X$ iff $Y(d) \leq_d^{\delta\,\pi} X(d)$ for all $d \in D$. Thus, $Y \leq^{\delta\,\pi} X$ means that the fact situation X is at least as strong for π as Y, or equivalently that Y is at least as strong for δ as X. A fortiori reasoning requires a new case to follow a precedent whenever it is at least as favourable to the precedent's prevailing party on every dimension.

Definition 1. A case base Γ is inconsistent iff it contains two cases $\langle X, \pi\rangle$ and $\langle Y, \delta\rangle$ such that $X \leq^{\delta\,\pi} Y$. It is consistent otherwise.

Obligation is usually defined for consistent case bases. Following [20], we define obligations also for inconsistent case bases. For $s \in \pi, \delta$, let $\bar{s}$ denote the other side:

$$\Gamma \models \mathbf{O}_X(s) \iff \exists\, \langle Y, s\rangle \in \Gamma : (Y \leq^{\bar{s},s} X) \text{ and } \nexists\, \langle Z, \bar{s}\rangle \in \Gamma : (Z \leq^{s,\bar{s}} X).$$

where $\Gamma \models \mathbf{O}_X(\pi)$ is read "in the context of Γ, X ought to be decided for π," and similarly for $\Gamma \models \mathbf{O}_X(\delta)$.

For consistent case bases, this coincides with the usual definition; when precedents constrain X in both directions, the two constraints cancel each other out.

3. Generalization of the Model

We now turn to the general setting with multiple outcomes. Given its natural use in multiclass machine learning [11], one might consider a one-vs-rest construction, building a binary model for each outcome o_i against $O \setminus o_i$. For the result model, however, this approach already fails with binary dimensions, which may lack a well-defined direction of support in a one-vs-rest comparison.

Example 1. Let $O = \{o_1, o_2, o_3\}$, and consider a dimension d with values $\{0, 1\}$. Suppose that $d = 1$ favours o_1 over o_2, but favours o_3 over o_1. Then, in the comparison o_1 versus $\{o_2, o_3\}$, $d = 1$ has no well-defined direction of support.

This motivates a different construction. Rather than comparing each outcome with all the others, we compare outcomes pairwise: for each pair of outcomes and each dimension, values are ordered according to how strongly they support one outcome over the other.

Formally, for every $d \in D$ and $o, o' \in O$, let $\leq_d^{oo'}$ be a preorder over $values(d)$. We read $p \leq_d^{oo'} q$ as saying that q supports o' over o at least as much as p, and define the reverse comparison by

$$p \leq_d^{o'o} q \quad \text{iff} \quad q \leq_d^{oo'} p.$$

We lift these preorders componentwise to fact situations:

$$X \leq^{oo'} Y \quad \text{iff} \quad X(d) \leq_d^{oo'} Y(d) \text{ for every } d \in D.$$

Definition 2. A case base Γ is inconsistent iff it contains two cases $\langle X, o\rangle$ and $\langle Y, o'\rangle$, with $o \neq o'$, such that $Y \leq^{oo'} X$. It is consistent otherwise.

We define pairwise obligations for possibly inconsistent case bases as follows.[3]

Definition 3 (Pairwise obligations)**.**

$$\Gamma \models \mathbf{O}_X(o' > o) \quad \text{iff} \quad \begin{array}{l}\text{(i) there is } \langle Y, o'\rangle \in \Gamma \text{ such that } Y \leq^{oo'} X, \text{ and} \\ \text{(ii) there is no } \langle Z, o\rangle \in \Gamma \text{ such that } Z \leq^{o'o} X.\end{array}$$

As before, for consistent case bases, Condition (ii) is redundant.

Definition 4 (Constraint Graph)**.** The *constraint graph* induced by Γ and X is the directed graph with outcomes as nodes and edges defined by

$$o_i \longrightarrow o_j \quad \text{iff} \quad \Gamma \models \mathbf{O}_X(o_i > o_j).$$

We illustrate the difficulties arising from aggregating pairwise constraints using a simple scenario, rather than a more involved legal example.

Example 2 (Free Evening)**.** A person should choose among *staying home alone* (o_1), *going to a party* (o_2), *going to the gym* (o_3) and *visiting her stepmother* (o_4). The dimensions d_1 (fatigue), d_2 (desire for social interaction) and d_3 (exercise deficit) have values 1–5; d_4 is binary, with $d_4 = 1$ indicating a promise to visit her stepmother that evening.

The directions of support and cases are summarized in Tables 1 and 2.

Table 1. Directions of support.

Pair	d_1(*fatigue*)	d_2(*interaction*)	d_3(*sport*)	d_4(*promise*)
o_1/o_2	$\uparrow o_1$	$\uparrow o_2$	–	–
o_2/o_3	–	$\uparrow o_2$	$\uparrow o_3$	–
o_3/o_1	$\uparrow o_1$	–	$\uparrow o_3$	–
o_1/o_4	–	–	–	$\uparrow o_4$
o_2/o_4	–	$\uparrow o_2$	–	–
o_3/o_4	–	–	$\uparrow o_3$	–

Table 2. Cases and focus situations.

	d_1	d_2	d_3	d_4	o
c_1	2	1	5	0	o_1
c_2	3	2	3	0	o_2
c_3	1	3	4	0	o_3
c_4	2	1	1	1	o_4
X_a	2	2	4	0	–
X_c	2	2	4	1	–

In Table 1, $\uparrow o_i$ means that a larger value of the dimension supports o_i in the corresponding pairwise comparison, while "–" means neutral, i.e., all values are equivalent in the corresponding preorder; e.g., $d_4 = 1$ favours o_4 over o_1, but has no influence on the choice between o_4 and going to a party or to the gym.

To illustrate the constraints, consider $\Gamma = \{c_2\}$ and the focus situation X_a from Table 2, and compare the outcomes o_1 (home) and o_2 (party). Increasing d_1 favours o_1 (vs o_2) hence decreasing it favours o_2, while increasing d_2 favours o_2. From c_2 to X_a, d_1 decreases from 3 to 2 and d_2 remains equal to 2; d_3 is neutral for this comparison. Hence X_a is at least as favourable to o_2 over o_1 as c_2, so c_2 witnesses $\Gamma \models \mathbf{O}_{X_a}(o_2 > o_1)$.

We consider three examples, with different outcomes, case bases, and focus cases. All arrows in the constraint graphs in Figure 1 are obtained in the same way: each outgoing edge $o_i \to o_j$ is witnessed by the precedent c_i with the same index.

[3] The reason for allowing inconsistent case bases will become clear in Section 4 (Remark 1).

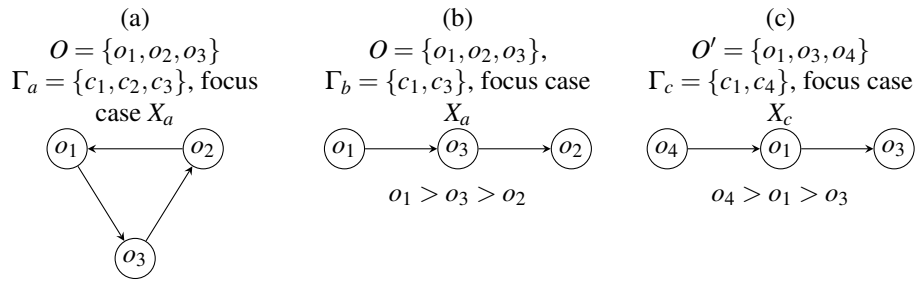


Figure 1. Constraint graphs for the Free Evening example: (a) a cycle; (b) a chain illustrating the preference reading; (c) a chain illustrating the exclusion reading.

Figure 1(a) shows that a consistent case base may generate a cycle. A cycle represents an unresolved comparative conflict, and intuitively such a conflict should not eliminate all decisions. This will rule out some aggregation methods below.

Figures 1(b) and 1(c) suggest two possible interpretations of pairwise constraints. On a *preference reading*, $o_i \to o_j$ means that o_i ought to be preferred to o_j; on an *exclusion reading*, it means that o_j ought to be excluded in favor of o_i. Which reading is appropriate depends, we think, on the specific domain and the outcomes under consideration. For instance, in Figure 1(b), the constraints can be read as expressing preferences: staying home ought to be preferred to going to the gym, which in turn ought to be preferred to going to the party. Figure 1(c), however, illustrates a possible motivation for the exclusion reading: going to the gym might be excluded in favor of staying home, which in turn might be excluded in favor of visiting one's stepmother. The distinction matters because, unlike the preference reading, the exclusion reading can support reinstatement: if o_1 is excluded, it no longer attacks o_3, which is then reinstated. Since different readings may call for different aggregation rules, we now examine how to aggregate pairwise constraints.

4. Choice of an Aggregation Method: Our Semantics

We start by considering a particularly simple aggregation rule, which we call the *preference semantics*, and show its limitations in the presence of cycles. We propose two alternatives: the Conflict-Free version 2 (CF2) argumentation semantics [17], which captures the exclusion reading, and a generalized preference semantics, corresponding to a choice rule studied by Joseph [16], which captures the preference reading while avoiding the problem created by cyclic constraint graphs like the one in Figure 1(a). For a semantics S, we write $\Gamma \models_{\mathsf{S}} \mathbf{P}_X(o)$ to mean that, given Γ, o is permissible for X under S.

Preference Semantics. A natural choice is to interpret the constraints as preferences and permit the maximal outcomes of the constraint graph.

Definition 5 (Preference semantics)**.** $\Gamma \models_{\mathsf{P}} \mathbf{P}_X(o)$ iff there is no o' such that $\Gamma \models \mathbf{O}_X(o' > o)$.

This semantics also preserves consistency: if Γ is consistent and X is decided for a permissible outcome, adding the resulting case to Γ does not create an inconsistency. It fails, however, on Figure 1(a). Under a transitive reading of preferences, the three-cycle

entails opposing preferences between every pair of outcomes. This mirrors an inconsistency in the binary model, where opposing constraints cancel and leave both outcomes permissible. By the same reasoning, all three outcomes should remain permissible in the cycle, whereas the preference semantics permits none.

CF2 Argumentation Semantics. The exclusion reading suggests viewing the constraint graph as an abstract argumentation framework. Such a framework is a directed graph whose nodes are arguments and whose edges represent attacks. Here, each outcome is treated as an argument for choosing it, and an edge from o_i to o_j means that o_i provides a reason to exclude o_j. A set of outcomes is conflict-free if no edge connects two outcomes in that set. Among the available semantics, CF2 [17] is particularly suitable for our examples, as it combines features of naive and grounded semantics.[4] Naive semantics permits the inclusion-maximal conflict-free sets. It therefore matches our intuition on the three-cycle of Figure 1(a): its extensions are $\{o_1\}$, $\{o_2\}$, and $\{o_3\}$, so all three outcomes are credulously permissible.[5] It behaves poorly, however, on a chain of attacks. For instance, from

$$o_4 \longrightarrow o_1 \longrightarrow o_3$$

naive semantics also admits $\{o_1\}$ as a maximal conflict-free set, and therefore permits o_1, although o_1 is itself defeated by o_4. Grounded semantics has the opposite strengths and weaknesses: it gives the expected reinstatement $\{o_4, o_3\}$ on this chain, but gives the empty extension on the three-cycle.

CF2 provides, in this sense, the best of both worlds. It uses naive semantics inside strongly connected components, while resolving the relations between components recursively. We give here a simplified description sufficient for our purposes; the fully formal recursive definition can be found in [17].

Definition 6 (CF2 extensions)**.** Let $F = (A, R)$ be the constraint graph viewed as an argumentation framework. If F is strongly connected, its CF2 extensions are its maximal conflict-free sets. Otherwise, CF2 recursively processes its strongly connected components, removing arguments defeated by those accepted in preceding components.

Definition 7 (Credulous CF2 semantics)**.** $\Gamma \models_{\mathsf{CF2}} \mathbf{P}_X(o)$ iff o belongs to some CF2 extension.

Thus, CF2 permits all outcomes in Figure 1(a), while in Figure 1(c) it accepts o_4 and reinstates o_3 after o_1 is defeated; Figure 1(b) is similar. This makes CF2 well suited to the exclusion reading: an outcome excluded by a stronger alternative need no longer exclude the outcomes that it attacks.

Generalized Preference Semantics. A third option is to retain the preference interpretation and take the transitive closure of the constraint relation. Thus, if there is a path from o_i to o_j, then o_i is treated as preferred to o_j. Outcomes that reach each other are therefore grouped into the same strongly connected component.

[4] The standard Dung semantics [21] discussed in [17] do not fit our intended treatment of odd cycles: they yield no accepted argument (or, for stable semantics, no extension). This behaviour is one of the motivations for introducing CF2.

[5] An outcome is credulously permissible if it belongs to at least one extension.

Definition 8 (Generalized preference semantics)**.** On the outcomes of the constraint graph, define

$$o \sim o' \quad \text{iff} \quad o \text{ and } o' \text{ are mutually reachable.}$$

The equivalence classes are the strongly connected components; quotienting the graph by $\sim$ yields a directed acyclic graph (DAG). We define $\Gamma \models_{\mathsf{GP}} \mathbf{P}_X(o)$ iff the class of o is maximal in this DAG, i.e., it has no incoming edge.

This preserves the transitive preference reading while solving the problem of the preference semantics on cycles. In Figure 1(b), $o_1 \to o_3 \to o_2$ yields $o_1 > o_2$ by transitivity, and only o_1 is permissible. In Figure 1(a), the three outcomes belong to the same strongly connected component, which is maximal in the quotient DAG; hence all three are permissible. For the strict preference relation represented by our arrows, this construction coincides with Joseph's relative-maximality rule [16].[6]

4.1. Properties and Comparison

The choice of semantics depends on how the pairwise constraints should be used. Preference semantics is appropriate when every incoming constraint must rule out an outcome; as we show below, it also preserves consistency when new decisions are added to a consistent case base. Generalized preference semantics treats the constraints as preferences and allows a choice within a cycle when no outcome outside it takes priority. CF2 is appropriate when excluding one outcome can remove its objection to another, allowing that other outcome to be chosen. The choice therefore depends on whether this reinstatement is intended in the application.

The three semantics also differ with respect to consistency preservation.

Remark 1. Unlike preference semantics, CF2 and generalized preference do not in general preserve consistency: adding a new case with a permissible outcome to a consistent case base may make the resulting case base inconsistent. This motivated the definition of the pairwise model for inconsistent case bases. Condition (ii) in Definition 3 ensures that opposing constraints cancel.

The binary model behaves well under this extension to inconsistent case bases:

Proposition 1. *Suppose that* Γ *contains precedents* $c = \langle Z, \pi \rangle$ *and* $c' = \langle Z', \delta \rangle$ *s.t.*

$$Z \leq^{\delta\,\pi} X \leq^{\delta\,\pi} Z'.$$

The two opposing constraints cancel, making both π *and* δ *permissible for* X*. Adding* X *to* Γ *with either decision does not change the effective constraints induced by* Γ *on any future case.*

[6] Joseph selects the inclusion-minimal sets that are not strictly dominated from outside. For a strict relation (as ours), these are precisely the maximal strongly connected components.

Proof. Assume that X is decided for π. If it later constrains a situation Y, then $X \leq^{\delta\,\pi} Y$. Since some existing precedent $c = \langle Z, \pi \rangle$ satisfies $Z \leq^{\delta\,\pi} X$, transitivity gives $Z \leq^{\delta\,\pi} X \leq^{\delta\,\pi} Y$. Hence c already imposed the same constraint on Y; adding X changes nothing. The argument for δ is symmetric. □

In Figure 1(a), deciding X_a for o_1 makes the case base inconsistent: the new precedent and c_2 support opposite outcomes in their pairwise comparison. Their constraints cancel, turning the cycle into a chain and changing the permissible outcomes. This cancellation does not restore consistency or remove either precedent; it allows the model to derive constraints from the resulting inconsistent case base. Both precedents remain available for reasoning about later cases. If new decisions must instead preserve consistency, preference semantics should be used.

We identify below a special case in which the three semantics coincide, which will also help relate the pairwise model to linearly ordered outcomes.

Proposition 2. *If the relation of pairwise constraints induced by* Γ *on a situation X is acyclic and transitive, then the following three semantics yield the same set of permissible outcomes:*

(i) the preference semantics,
(ii) the CF2 argumentation semantics, under credulous acceptance,
(iii) the generalized preference semantics

Proof. Since the graph is acyclic, all its strongly connected components are singletons. Hence the maximal-SCC semantics permits exactly the maximal, equivalently undefeated, outcomes. Thus (i) and (iii) coincide. It remains to compare (i) and (ii). As the graph is acyclic, CF2 coincides with grounded semantics [17].

(i) $\Rightarrow$ (ii) Suppose that o is permissible by (i). Then o is undefeated, hence unattacked in the associated argumentation framework. Every unattacked argument belongs to the grounded extension. Thus o is permissible according to (ii).

(ii) $\Rightarrow$ (i) Suppose that o is permissible according to (ii). Assume, for contradiction, that o is not maximal. Then there is some outcome o' such that o' attacks o. Since o belongs to the grounded extension, it must be defended, so there is some accepted outcome o'' that attacks o'. By transitivity, o'' also attacks o. But the grounded extension is conflict-free, contradicting the fact that both o'' and o are accepted. Hence o is maximal, and therefore permissible according to (i). □

5. Linearly Ordered Outcomes

We now consider linearly ordered outcomes and relate our pairwise model to the construction suggested in [18]. The latter is simple, intuitive, and easy to compute, making it a natural sanity check for our pairwise approach. As shown below all three aggregation semantics recover the same decisions in this special setting. In [18], as part of a hierarchical model, precedents below a new situation provide lower bounds on its outcome, while precedents above it provide upper bounds. We start by formalizing this mechanism.[7] Let

[7] In [18], directions of support are derived by combining an ordering of dimension values with whether the dimension is correlated or anticorrelated with the ordering of the outcomes. Here we take these directions as given by the orders $\leq_d^{-+}$.

$$O = \{o_1, \ldots, o_m\}, \qquad o_1 \prec \cdots \prec o_m.$$

Here $\prec$ is a given linear order on outcomes. For every $d \in D$, let $\leq_d^{-+}$ be an order over $values(d)$, oriented towards increasingly good outcomes, and define

$$X \leq^{-+} Y \quad \text{iff} \quad X(d) \leq_d^{-+} Y(d) \text{ for all } d \in D.$$

Definition 9. A case base Γ is inconsistent iff it contains $\langle X, o\rangle$ and $\langle Y, o'\rangle$ with $o \prec o'$ and $Y \leq^{-+} X$.

For a consistent Γ and situation X, define, when non-empty,

$$\ell_X = \max\{o : \langle Y, o\rangle \in \Gamma,\ Y \leq^{-+} X\}, \qquad u_X = \min\{o : \langle Y, o\rangle \in \Gamma,\ X \leq^{-+} Y\}.$$

We write

$$\Gamma \models \mathbf{O}_X(\textit{at-least}(o)) \iff o = \ell_X, \qquad \Gamma \models \mathbf{O}_X(\textit{at-most}(o)) \iff o = u_X,$$

$$\text{and} \quad \Gamma \models \mathbf{P}_X(o) \quad \iff \quad \ell_X \preccurlyeq o \preccurlyeq u_X,$$

with the obvious interpretation when one of the bounds is absent.

Consistency guarantees $\ell_X \preccurlyeq u_X$ whenever both bounds exist: otherwise witnesses for the two bounds would form an inconsistent pair.

We now compare this construction with the pairwise model.

Assume that the two models have the same outcomes and dimensions, that $o_1 \prec \cdots \prec o_m$, and that, for every $d \in D$, $p, q \in values(d)$ and $i < j$,

$$p \leq_d^{o_i o_j} q \quad \text{iff} \quad p \leq_d^{-+} q.$$

Proposition 3. *Under the previous assumptions, for every consistent case base Γ and situation X, the pairwise constraint graph is transitive and acyclic.*

Proof. Suppose that $\Gamma \models \mathbf{O}_X(o_i > o_j)$ and $\Gamma \models \mathbf{O}_X(o_j > o_k)$. Consistency rules out $i > j < k$ and $i < j > k$: in either case, the precedents witnessing the two constraints lie on opposite sides of X and would form an inconsistent pair. Hence either $i > j > k$ or $i < j < k$. By compatibility of the dimensional orders, w.l.o.g. the precedent witnessing the first constraint then also witnesses $\Gamma \models \mathbf{O}_X(o_i > o_k)$. Thus the graph is transitive. If it contained a cycle, transitivity would yield constraints in both directions between two outcomes, contradicting consistency. Hence it is acyclic. □

By Propositions 3 and 2, for linearly ordered outcomes, the three semantics coincide, so we use the simplest one: the preference semantics. Since this semantics preserves consistency, starting from a consistent case base and adding only permissible decisions keeps the case base consistent. This allows us to compare with the linearly ordered model, which is defined only for consistent case bases; consistency is also used in the equivalence proof.

Proposition 4. *For every consistent case base* Γ*, situation* X*, and outcome* $o \in O$*,*

$$\Gamma \models_S \boldsymbol{P}_X(o) \text{ in the pairwise model with } S \in \{P, CF2, GP\} \text{ iff } \ell_X \preccurlyeq o \preccurlyeq u_X,$$

with the lower (resp. upper) inequality omitted when ℓ_X *(resp.* u_X*) is undefined.*

Proof. ($\Longrightarrow$) Let $\Gamma \models_P \mathbf{P}_X(o_i)$ for the pairwise model. If $\ell_X = o_j$ with $j > i$, a precedent $\langle Y, o_j \rangle$ witnessing the lower bound satisfies $Y \leq^{-+} X$ and, by compatibility, yields $\Gamma \models \mathbf{O}_X(o_j > o_i)$, contradicting the permissibility of o_i. Hence $\ell_X \preccurlyeq o_i$ whenever ℓ_X exists. Symmetrically, $o_i \preccurlyeq u_X$ whenever u_X exists.

($\Longleftarrow$) Conversely, let o_i lie between the two bounds, with the corresponding inequality omitted when a bound is absent. Suppose that some o_j constrains o_i.

If $j > i$, the witnessing precedent for $\Gamma \models \mathbf{O}_X(o_j > o_i)$ lies below X. Hence it contributes a lower bound at least o_j, contradicting $\ell_X \preccurlyeq o_i$. If no lower bound exists, the existence of such a precedent is itself a contradiction.

If $j < i$, the witnessing precedent lies above X and contributes an upper bound at most o_j, contradicting $o_i \preccurlyeq u_X$; again, if no upper bound exists, such a precedent cannot exist. Thus no outcome constrains o_i, so o_i is permissible under the preference semantics, and therefore under all three semantics. □

6. Conclusions

Extending the dimensional result model to multiple outcomes reveals that there is no unique way of aggregating the constraints obtained from pairwise comparisons. We considered three semantics. In particular, *generalized preference* reflects a transitive preference reading of the pairwise constraints, while *CF2* reflects an elimination reading, in which one outcome can rule out another. The third, direct preference semantics, is simpler but may fail in the presence of cycles. Despite these differences, the three semantics coincide in the case of linearly ordered outcomes and recover the model introduced in [18] in this setting.

This work provides a first foundation for extending the dimensional result model to multiple outcomes, while leaving several questions open. First, the proposed semantics should be compared empirically on realistic case bases. Second, the pairwise directions of support, assumed as given here, could be learned or extracted automatically from data. Third, while our pairwise model makes no prior assumptions on the relations among outcomes, particular outcome structures may admit simpler constructions; for example, two linear orders, as when choosing between different amounts of a fine and terms of imprisonment.

Acknowledgments: Work supported by WWTF [Grant ID: 10.47379/ICT23030].

References

[1] Alexander L. Constrained by precedent. Southern California Law Review. 1989;63:1.
[2] Horty JF. The result model of precedent. Legal Theory. 2004;10(1):19-31.
[3] Horty JF, Bench-Capon TJM. A factor-based definition of precedential constraint. Artificial Intelligence and Law. 2012;20(2):181-214.

[4] Canavotto I, Horty JF. Reasoning with hierarchies of open-textured predicates. In: Proceedings of the 19th International Conference on Artificial Intelligence and Law. Association for Computing Machinery; 2023. p. 52-61.
[5] van Woerkom W, Grossi D, Prakken H, Verheij B. Hierarchical models of precedential constraint. Artificial Intelligence and Law. 2026.
[6] Aleven V. Teaching case-based argumentation through a model and examples [PhD thesis]. University of Pittsburgh; 1997.
[7] Ashley K. Modeling legal argument: Reasoning with cases and hypotheticals. MIT Press; 1990.
[8] van Woerkom WK. Formal results on case-base consistency: A COMPAS case study. In: Proceedings of the 20th International Conference on Artificial Intelligence and Law. Association for Computing Machinery; 2025. p. 298-307.
[9] Canavotto I, Horty JF. Piecemeal knowledge acquisition for computational normative reasoning. In: Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society. Association for Computing Machinery; 2022. p. 171-80.
[10] Liu X, Lorini E. A unified logical framework for explanations in classifier systems. Journal of Logic and Computation. 2023;33(2):485-515.
[11] Allwein EL, Schapire RE, Singer Y. Reducing multiclass to binary: A unifying approach for margin classifiers. Journal of Machine Learning Research. 2000;1:113-41.
[12] Hüllermeier E, Fürnkranz J, Cheng W, Brinker K. Label ranking by learning pairwise preferences. Artificial Intelligence. 2008;172(16–17):1897-916.
[13] Rajkumar A, Ghoshal S, Lim LH, Agarwal S. Ranking from stochastic pairwise preferences: Recovering Condorcet winners and tournament solution sets at the top. In: Proceedings of the 32nd International Conference on Machine Learning. vol. 37 of Proceedings of Machine Learning Research. PMLR; 2015. p. 665-73.
[14] Brandt F, Brill M, Harrenstein P. Tournament solutions. In: Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD, editors. Handbook of Computational Social Choice. Cambridge University Press; 2016. p. 57-84.
[15] Duggan J. A systematic approach to the construction of non-empty choice sets. Social Choice and Welfare. 2007;28(3):491-506.
[16] Joseph RR. Making choices with a binary relation: Relative choice axioms and transitive closures. European Journal of Operational Research. 2010;207(2):865-77.
[17] Baroni P, Giacomin M, Guida G. SCC-recursiveness: A general schema for argumentation semantics. Artificial Intelligence. 2005;168(1–2):162-210.
[18] van Woerkom WK, Grossi D, Prakken H, Verheij B. A fortiori case-based reasoning: From theory to data. Journal of Artificial Intelligence Research. 2024;81:401-41.
[19] Horty JF. Reasoning with dimensions and magnitudes. Artificial Intelligence and Law. 2019;27.
[20] Canavotto I. Reasoning with inconsistent precedents. Artificial Intelligence and Law. 2025;33(1):137-66.
[21] Dung PM. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and *n*-person games. Artificial Intelligence. 1995;77(2):321-57.