

Constrained Consequence

Katarina Britz, Johannes Heidema and Ivan Varzinczak

Abstract. There are various contexts in which it is not pertinent to generate and attend to *all* the classical consequences of a given premiss — or to trace *all* the premisses which classically entail a given consequence. Such contexts may involve limited resources of an agent or inferential engine, contextual relevance or irrelevance of certain consequences or premisses, modelling everyday human reasoning, the search for plausible abduced hypotheses or potential causes, etc. In this paper we propose and explicate one formal framework for a whole spectrum of consequence relations, flexible enough to be tailored for choices from a variety of contexts. We do so by investigating semantic constraints on classical entailment which give rise to a family of infra-classical logics with appealing properties. More specifically, our infra-classical reasoning demands (beyond $\alpha \models \beta$) that $Mod(\beta)$ does not run wild, but lies within the scope (whatever that may mean in some specific context) of $Mod(\alpha)$, and which can be described by a sentence $\bullet\alpha$ with $\beta \models \bullet\alpha$. Besides being infra-classical, the resulting logic is also non-monotonic and allows for non-trivial reasoning in the presence of inconsistencies.

Mathematics Subject Classification (2010). Primary 03B20; Secondary 03B22, 03B45, 03B47, 03B53, 03B60, 03A05.

Keywords. Constrained consequences, weakening operator, infra-classical entailment, non-monotonic reasoning, modal logic, action, causality, abduction, diagnosis, paradoxes.

1. Introduction and Motivation

In a classical entailment $\alpha \models \beta$, no information beyond that encapsulated locally in α and β plays any role at all. Extra information may be employed to construct altered entailment relations, which sometimes allow *more* pairs (α, β) in the relation, going *supra-classical*, or *fewer*, going *infra-classical*, or just going *non-classical*.

We are interested in contexts in which it is expedient to go infra-classical. To do that in a disciplined way, extra information deriving from

the specific context is needed, be it syntactic (like an extra connective) or, at a deeper level, semantic (like additional structure on models).

In this paper we venture beyond the prevailing assumption that non-monotonic entailment is usually supra-classical. We do so by investigating semantic constraints on classical entailment which give rise to a family of infra-classical logics with appealing properties. These constraints may also be viewed as inducing a form of *bounded reasoning*, where the bounds are imposed by constraints that block undesirable consequences, eschew pairs (α, β) which, in a specific context, are not pertinent.

Syntactically, the extra semantic constraints can be represented by a unary *uniform weakening operator* \bullet on sentences. A special instance of this operator that we shall introduce is the modal operator \diamond , for which the corresponding semantic constraint is a reflexive accessibility relation.

This leads to our definition of constrained entailment as an infra-classical relation from which all pairs (α, β) , where β goes beyond the bounds set by α , have been culled. An immediate property of constrained entailment that we single out here is that it clamps constraints down on the valid entailments of some existing underlying logic — for present purposes the valid semantic entailments $\alpha \models \beta$ of classical logic. Classical \models is then in this paper an upper bound on the set of constrained entailment pairs (α, β) . (It could also have been, e.g., some supra-classical entailment \sim .)

The remainder of the paper is organized as follows: In the next section we introduce the notions of uniform weakening and constrained entailment, and discuss their basic properties and relationship. In Section 3 we present a modal instance of these constructions which also illustrates with an example how to reason abductively with constrained entailment in a causal or action oriented context. In Section 4 we discuss our work in the contexts of two well-known existing research platforms, namely preferential reasoning and relevance logics. We conclude with a recapitulation of the contributions of the present paper and some directions for further investigation.

2. Classical Entailment Constrained

We start by introducing some logical preliminaries and setting up the notation we are going to use throughout the paper. In this section we work in a propositional language over a set of propositional *atoms* \mathfrak{A} , together with the distinguished atom \top (*verum*). (In Section 3 we shall adopt a richer language.) Atoms are denoted by p, q, \dots . The formulas (or sentences) of our language are denoted by α, β, \dots . Those are recursively defined as follows:

$$\alpha ::= p \mid \top \mid \neg\alpha \mid \alpha \wedge \alpha \mid \bullet\alpha$$

(Here \bullet denotes a unary operator of the type we will define and fit with a semantics in Subsection 2.1.) All the other connectives ($\vee, \rightarrow, \leftrightarrow, \dots$) and the special atom \perp (*falsum*) are defined in terms of \neg and \wedge in the usual way. With \mathfrak{F} we denote the set of all formulas of the language.

We denote by W a set of *worlds* (alias *states* or propositional *valuations* or *interpretations*) $w : \mathfrak{P} \rightarrow \{0, 1\}$, with 0 denoting falsity and 1 truth. With $Mod(\alpha)$ we denote the set of all *models* of α (worlds satisfying α).

Classical logical consequence (semantic entailment) and logical equivalence are denoted by \models and \equiv respectively. Given sentences α and β , the meta-statement $\alpha \models \beta$ means $Mod(\alpha) \subseteq Mod(\beta)$. $\alpha \equiv \beta$ abbreviates $\alpha \models \beta$ and $\beta \models \alpha$.

2.1. Weakening Operators

Given a sentence α (together with its associated set of models $Mod(\alpha)$), we are interested in expanding $Mod(\alpha)$ by admitting meritorious worlds not satisfying α . Logically speaking, this can be achieved by *weakening* the sentence α .

We denote by \bullet any operator which performs this weakening on sentences. A very simple example thereof is the operator \vee_β , defined as $\vee_\beta \alpha \equiv_{\text{def}} \alpha \vee \beta$ (for a fixed β): given α , $\vee_\beta \alpha$ delivers a logically weaker sentence. We note that although $\bullet\alpha$ is a specific sentence of our language, this operator \bullet is not necessarily a truth functional connective of the language (cf. Section 3). Its semantics is flexible, but not arbitrary.

We identify two desirable properties that a ‘well-behaved’ weakening operator is here expected to satisfy:

Weakening. $\alpha \models \bullet\alpha$ (W)

Intuitively, when expanding $Mod(\alpha)$, one should not exclude any of the α -worlds. Therefore, $Mod(\alpha) \subseteq Mod(\bullet\alpha)$.

Uniformity. If $\alpha \models \beta$, then $\bullet\alpha \models \bullet\beta$ (U)

Intuitively, if $Mod(\alpha) \subseteq Mod(\beta)$, then any inflation of $Mod(\alpha)$ should fit within the corresponding inflation of the larger set $Mod(\beta)$, i.e., weakening behaves *isotonically* with respect to entailment.

Definition 2.1. A *uniform weakening operator* is a function $\bullet : \mathfrak{F} \rightarrow \mathfrak{F}$ satisfying Properties W and U.

Given a uniform weakening operator \bullet , we have that the following hold:

$$\bullet\alpha \vee \bullet\beta \models \bullet(\alpha \vee \beta) \tag{2.1}$$

$$\bullet(\alpha \wedge \beta) \models \bullet\alpha \wedge \bullet\beta \tag{2.2}$$

From $\alpha \models \alpha \vee \beta$ and U follows $\bullet\alpha \models \bullet(\alpha \vee \beta)$. Similarly, from $\beta \models \alpha \vee \beta$ and U follows $\bullet\beta \models \bullet(\alpha \vee \beta)$. Putting both results together, it follows that $\bullet\alpha \vee \bullet\beta \models \bullet(\alpha \vee \beta)$. By a similar argument, we can derive $\bullet(\alpha \wedge \beta) \models \bullet\alpha \wedge \bullet\beta$.

Another property that follows from Uniformity is the following:

$$\text{If } \alpha \equiv \beta, \text{ then } \bullet\alpha \equiv \bullet\beta \tag{2.3}$$

Therefore, as already expected, weakening is *syntax independent*. It is easy to see, however, that the converse of (2.3) above does not hold.

If \bullet is a uniform weakening operator, then the following holds:

$$\bullet\top \equiv \top \tag{2.4}$$

However, in general it is not the case that $\bullet\perp \equiv \perp$, since the expansion of $\emptyset = \text{Mod}(\perp)$ is not necessarily the empty set. (Cf. Section 3.2 for an instance of weakening which does force this.) Note that $\neg\bullet\alpha \models \neg\alpha \models \bullet\neg\alpha$.

Example 2.2. A simple example of a uniform weakening operator is the following forgetting operator [34]. Let $\text{forget}(\alpha, p) \equiv_{\text{def}} \alpha_p^+ \vee \alpha_p^-$, where α_p^+ is the result of replacing all occurrences of p in α by \top , and α_p^- is the result of replacing all occurrences of p in α by \perp . Then $\alpha \models \text{forget}(\alpha, p)$. Also, if $\alpha \models \beta$, then $\text{forget}(\alpha, p) \models \text{forget}(\beta, p)$.

2.2. Constrained Entailment Relations

We define a type of general and abstract entailment relation with which to capture the idea that a given consequence does not go too far beyond its premiss. Let \prec denote an entailment relation on the set of formulas \mathfrak{F} . We shall read $\alpha \prec \beta$ as “ α *constrainedly entails* β ”. Since \prec is intended not to rival but to generalize \models , properties that we select to characterize \prec should also hold for \models .

We specify that in order for \prec to instantiate the notion of a ‘well-behaved’ constrained entailment as studied here, it has to satisfy the following five properties:

Reflexivity. $\alpha \prec \alpha$ (R)

The intuition behind R is clear: an entailment should not be constrained up to the point where a sentence does not entail itself anymore.

Infra-Classicality. If $\alpha \prec \beta$, then $\alpha \models \beta$ (IC)

In the symbol \prec , the ‘ \prec ’ refers to the infra-classical aspect of the new entailment, as opposed to the ‘ \models ’ in \models , since what we want to do, in a sense, is to ‘cull down’ some of the pairs in the relation \models , obtaining a subset thereof.

Equivalence. If $\alpha \equiv \beta$ and $\alpha \prec \gamma$, then $\beta \prec \gamma$. If $\beta \equiv \gamma$ and $\alpha \prec \beta$, then $\alpha \prec \gamma$ (E)

Just like Property (2.3) for \bullet , E makes \prec *syntax independent*.

Semantic Interpolation. If $\alpha \models \beta$, $\beta \models \gamma$, and $\alpha \prec \gamma$, then $\alpha \prec \beta$ and $\beta \prec \gamma$ (Int)

If α constrainedly entails γ , then, for any β which is *classically* semantically interpolated between α and γ , we also have that β is *constrainedly* semantically interpolated between them. The set of constrained consequences of a fixed premiss (α) is “convex”.

Generalized Disjunction. If $\alpha \prec \beta$ and $\gamma \prec \delta$, then $\alpha \vee \gamma \prec \beta \vee \delta$ (GD)

This property says that \prec is isotonic with respect to bilateral disjunctive weakening, somewhat similar to Uniformity for \bullet -weakening in the case of \models .

Definition 2.3. A *constrained entailment relation* is a relation $\prec \subseteq \mathfrak{F} \times \mathfrak{F}$ satisfying Properties R, IC, E, Int and GD.

In what follows we mention a few other derived properties which our constrained entailment relation satisfies.

Modus Ponens. If $\alpha \prec \beta$ and $\alpha \prec \beta \rightarrow \gamma$, then $\alpha \prec \gamma$ (MP)

Suppose that $\alpha \prec \beta$ and $\alpha \prec \beta \rightarrow \gamma$. By IC, $\alpha \models \beta$ and $\alpha \models \beta \rightarrow \gamma$. Hence we have $\alpha \models \gamma$, $\gamma \models \beta \rightarrow \gamma$, and $\alpha \prec \beta \rightarrow \gamma$. By Int it follows that $\alpha \prec \gamma$.

Right Disjunction. If $\alpha \prec \beta$ and $\alpha \prec \gamma$, then $\alpha \prec \beta \vee \gamma$ (RD)

The proof of this property follows straightforwardly from GD.

Right Conjunction. If $\alpha \prec \beta$ and $\alpha \prec \gamma$, then $\alpha \prec \beta \wedge \gamma$ (RC)

Suppose that $\alpha \prec \beta$ and $\alpha \prec \gamma$. Then we have that $\alpha \models \beta$ and $\alpha \models \gamma$, by IC. Therefore $\alpha \models \beta \wedge \gamma$, $\beta \wedge \gamma \models \beta$, and $\alpha \prec \beta$. By Int we have $\alpha \prec \beta \wedge \gamma$.

Note that $\alpha \prec \beta \wedge \gamma$ on its own does *not* imply $\alpha \prec \beta \vee \gamma$.

Left Disjunction. If $\alpha \prec \gamma$ and $\beta \prec \gamma$, then $\alpha \vee \beta \prec \gamma$ (LD)

The proof of this property also follows straightforwardly from GD.

Finally, it can also be verified easily that \prec satisfies the following version of Cut:

Cut. If $\alpha \wedge \beta \prec \gamma$ and $\alpha \prec \beta$ (or even only $\alpha \models \beta$), then $\alpha \prec \gamma$ (Cut)

We shall address additional properties of \prec in Subsection 2.4. In Subsection 4.1 we shall look at the properties of \prec in the light of those of preferential reasoning [30].

2.3. Representation Results

In this subsection we investigate the relationship between our weakening operator \bullet and the constrained entailment relation \prec . In particular, here we establish the mutual representability or inter-definability of \bullet and \prec , each in terms of the other one with its concomitant properties.

Given some \bullet which satisfies Properties W and U, we define an entailment relation in the following way (anticipating Theorem 2.6, which justifies our notation):

Definition 2.4. $\alpha \prec \beta$ if and only if $\alpha \models \beta$ and $\beta \models \bullet\alpha$.

We note that \prec can be defined equivalently, but more concisely, as in the following result:

Theorem 2.5. $\alpha \prec \beta$ if and only if $\alpha \vee \beta \models \beta \wedge \bullet\alpha$.

Proof.

(\Rightarrow): If $\alpha \prec \beta$, then $\alpha \models \beta$ and $\beta \models \bullet\alpha$. Hence $\alpha \vee \beta \models \beta$ and $\beta \models \beta \wedge \bullet\alpha$. Therefore $\alpha \vee \beta \models \beta \wedge \bullet\alpha$.

(\Leftarrow): If $\alpha \vee \beta \models \beta \wedge \bullet\alpha$, then $\alpha \models \beta$ and $\beta \models \bullet\alpha$, hence $\alpha \prec \beta$. \square

We also note that, both in the context of the present Definition 2.4 as well as in the modal context of the later Definition 3.7 in Section 3.2, the semantically defined entailment relation \prec has a sound and complete syntactic counterpart in those cases when $\bullet\alpha$ is a Boolean combination of atomic symbols. This stays true when later, in Theorem 3.6, the weakening operator is an ordinary modal diamond. Let namely \vdash denote any one of

the syntactically defined entailment relations which is sound and complete with respect to \models ; then $\alpha \prec \beta$ if and only if $\alpha \vee \beta \models \beta \wedge \bullet\alpha$, if and only if $\alpha \vee \beta \vdash \beta \wedge \bullet\alpha$. Should, however, the weakening operator in some specific context have a radically different type of semantics, the problem of defining a suitable \vdash for that particular context remains open.

Given that, by Theorem 2.5, \prec can be expressed in terms of \models and \bullet , we may ask whether \prec , like \models , is preserved under uniform substitutions of arbitrary sentences for atomic letters. Well, $\bullet\alpha$ may in general not even be equivalent to a Boolean combination of atoms, so the answer is ‘no’, similar to the case of preferential entailment.

Theorem 2.6. *The entailment relation \prec from Definition 2.4 has Properties R, IC, E, Int and GD, and hence is a constrained entailment relation.*

Proof.

R: $\alpha \models \alpha$ and, by W, $\alpha \models \bullet\alpha$. Therefore $\alpha \prec \alpha$.

IC: This follows from the definition of \prec .

E: Let $\alpha \equiv \beta$ and $\alpha \prec \gamma$. Then $\alpha \models \gamma$ and $\gamma \models \bullet\alpha$. From this and Property (2.3) it follows that $\beta \models \gamma$ and $\gamma \models \bullet\beta$, which gives us $\beta \prec \gamma$. Now, assume $\beta \equiv \gamma$ and $\alpha \prec \beta$. Then $\alpha \models \beta$ and $\beta \models \bullet\alpha$, which, given $\beta \equiv \gamma$, implies $\alpha \models \gamma$ and $\gamma \models \bullet\alpha$. Thus $\alpha \prec \gamma$.

Int: Suppose that $\alpha \models \beta$, $\beta \models \gamma$, and $\alpha \prec \gamma$, which, given our present definition of \prec , is equivalent to $\alpha \models \beta$, $\beta \models \gamma$, and $\gamma \models \bullet\alpha$. Applying U to the chain $\alpha \models \beta \models \gamma$, we get $\bullet\alpha \models \bullet\beta \models \bullet\gamma$, hence $\alpha \models \beta \models \gamma \models \bullet\alpha \models \bullet\beta \models \bullet\gamma$, from which $\alpha \models \beta$ and $\beta \models \bullet\alpha$, i.e., $\alpha \prec \beta$, as well as $\beta \models \gamma$ and $\gamma \models \bullet\beta$, i.e., $\beta \prec \gamma$, follow.

GD: Suppose that $\alpha \prec \beta$ and $\gamma \prec \delta$, which means that $\alpha \models \beta$, $\gamma \models \delta$, $\beta \models \bullet\alpha$, and $\delta \models \bullet\gamma$. Then $\alpha \vee \gamma \models \beta \vee \delta$ and $\beta \vee \delta \models \bullet\alpha \vee \bullet\gamma$. From the latter statement, invoking Property 2.1, we get $\beta \vee \delta \models \bullet(\alpha \vee \gamma)$, so that, finally, $\alpha \vee \gamma \prec \beta \vee \delta$. \square

Given an entailment relation defined as in Definition 2.4, a corollary of Properties RD, RC and Theorem 2.6 above is that, for a fixed premiss α , the set of consequences that α entails in our new relation are all the β s that, with respect to \models , lie between that particular α -premiss and $\bullet\alpha$, and hence form a sub-lattice (closed under conjunction and disjunction) of the Lindenbaum-Tarski algebra of the language [16], as depicted in Figure 1 below (with entailment going ‘up’). Here we see clearly how the semantic constraint $\beta \models \bullet\alpha$ on β establishes an upper bound on the amount of semantic content or information that may be lost in passing constrainedly from premiss α to consequence β .

Given a consequence β , the set of all those premisses α such that $\alpha \prec \beta$ does not always constitute a sub-lattice of the Lindenbaum-Tarski algebra, since it is not, in general, closed under conjunction (cf. the non-monotonicity of \prec in the next Subsection 2.4). But it is closed under disjunction according to Property LD.




FIGURE 1. The sub-lattice induced by constrained entailment from premiss α .

Now, given some \llcorner which satisfies Properties R, IC, E, Int and GD, we define (while looking again at Figure 1) $\bullet : \mathfrak{F} \longrightarrow \mathfrak{F}$ by the following:

Definition 2.7. $\bullet\alpha \equiv_{\text{def}} \bigvee \{\beta \mid \alpha \llcorner \beta\}$.

We note that, should our language have infinitely many generating propositional symbols in \mathfrak{P} , then α may have infinitely many non-equivalent constrained consequences of incomparable logical strength, and hence the construction of $\bullet\alpha$ may require infinitary disjunctions (and conjunctions) in the language [29, 18]. For those who prefer not to allow infinitary connectives (even though they allow infinitary generation) we could avoid $\bullet\alpha$ in the object language — or even as a meta-name for a sentence of the object language — and give a purely semantic treatment of uniform weakening:

$$\bullet\text{Mod}(\alpha) = \bigcup \{\text{Mod}(\beta) \mid \alpha \llcorner \beta\}$$

This remark also explains why here we consider only single premisses and not (in Tarski style) *sets* of premisses in an entailment: the set is logically equivalent to its conjunction — or you can do everything semantically.

Lemma 2.8. *Let \bullet be as in Definition 2.7. Then $\bullet\alpha \vee \bullet\beta \models \bullet(\alpha \vee \beta)$.*

Proof. $\bullet\alpha \vee \bullet\beta \equiv \bigvee \{\gamma \mid \alpha \llcorner \gamma\} \vee \bigvee \{\delta \mid \beta \llcorner \delta\} \equiv \bigvee \{\gamma \vee \delta \mid \alpha \llcorner \gamma \text{ and } \beta \llcorner \delta\}$. Then, by GD, we have that $\bullet\alpha \vee \bullet\beta \models \bigvee \{\gamma \vee \delta \mid \alpha \vee \beta \llcorner \gamma \vee \delta\}$. Hence, $\bullet\alpha \vee \bullet\beta \models \bullet(\alpha \vee \beta)$. \square

Lemma 2.9. *Let \bullet be as in Definition 2.7. Then $\alpha \equiv \beta$ implies $\bullet\alpha \equiv \bullet\beta$.*

Proof. Let α, β be such that $\alpha \equiv \beta$. Since \llcorner satisfies E, we have that $\bullet\alpha \equiv \bigvee \{\gamma \mid \alpha \llcorner \gamma\} \equiv \bigvee \{\gamma \mid \beta \llcorner \gamma\} \equiv \bullet\beta$. \square

Theorem 2.10. *The operator \bullet from Definition 2.7 is a uniform weakening operator.*

Proof. We shall show that \bullet satisfies Properties W and U.

W: By Reflexivity, $\alpha \llcorner \alpha$, so α is one of the β 's considered in the definition of $\bullet\alpha$. Therefore we have $\alpha \models \bullet\alpha$.

U: Suppose that $\alpha \models \beta$, which is equivalent to $\alpha \vee \beta \equiv \beta$. Then we have the following chain of logical relationships: $\alpha \models \bullet\alpha \models \bullet\alpha \vee \bullet\beta \models \bullet(\alpha \vee \beta) \equiv \bullet\beta$. The first \models -link is guaranteed by W, already proven, and the third one by Lemma 2.8. Finally, the \equiv -link is guaranteed by Lemma 2.9. Therefore we have $\bullet\alpha \models \bullet\beta$. \square

Suppose now we start with \bullet satisfying its required properties, then define \prec in terms of \bullet as in Definition 2.4, and finally define \bullet' in terms of \prec as in Definition 2.7.

Theorem 2.11. $\bullet' = \bullet$.

Proof. $\bullet'\alpha \equiv \bigvee\{\beta \mid \alpha \prec \beta\} \equiv \bigvee\{\beta \mid \alpha \models \beta \text{ and } \beta \models \bullet\alpha\} \equiv \bullet\alpha$, since $\bullet\alpha$ is the logically weakest of the β 's. \square

Suppose we start with \prec satisfying its required properties, then define \bullet in terms of \prec as in Definition 2.7, and finally define \prec' in terms of \bullet as in Definition 2.4.

Theorem 2.12. $\prec' = \prec$.

Proof.

(\subseteq): Suppose $\alpha \prec' \beta$, which now means $\alpha \models \beta$ and $\beta \models \gamma$, where $\gamma = \bullet\alpha = \bigvee\{\delta \mid \alpha \prec \delta\}$. We have that $\alpha \models \beta \models \gamma$ and $\alpha \prec \gamma$, the latter by GD from all the $\alpha \prec \delta$. By Int we have that $\alpha \prec \beta$.

(\supseteq): Suppose that $\alpha \prec \beta$. Then $\alpha \models \beta$, by IC, and β is one of the δ 's considered in $\bullet\alpha = \bigvee\{\delta \mid \alpha \prec \delta\}$, so that $\beta \models \bullet\alpha$, completing the conditions for $\alpha \prec' \beta$. \square

2.4. Additional Properties of \bullet and \prec

Suppose that we have a weakening operator \bullet and a constrained entailment relation \prec having their respective required sets of properties and related in the required way: $\alpha \prec \beta$ if and only if $\alpha \models \beta$ and $\beta \models \bullet\alpha$; and $\bullet\alpha \equiv \bigvee\{\beta \mid \alpha \prec \beta\}$. Then we have the following:

Theorem 2.13. \bullet is idempotent ($\bullet\bullet\alpha \equiv \bullet\alpha$) if and only if \prec is transitive.

Proof.

(\Rightarrow): Suppose that \bullet is idempotent and let $\alpha \prec \beta$ and $\beta \prec \gamma$. Then $\alpha \models \beta$, $\beta \models \gamma$, $\beta \models \bullet\alpha$, and $\gamma \models \bullet\beta$. From $\beta \models \bullet\alpha$ and Property U, we get $\bullet\beta \models \bullet\bullet\alpha$. This and idempotence of \bullet gives us $\bullet\beta \models \bullet\alpha$. Then $\alpha \models \gamma$ and $\gamma \models \bullet\alpha$, and therefore $\alpha \prec \gamma$.

(\Leftarrow): $\alpha \models \bullet\alpha$ (from Property W). $\bullet\alpha \models \bullet\alpha$ (from reflexivity of classical \models). Therefore, $\alpha \prec \bullet\alpha$. Similarly we show that $\bullet\alpha \models \bullet\bullet\alpha$ and $\bullet\bullet\alpha \models \bullet\bullet\alpha$, which gives us $\bullet\alpha \prec \bullet\bullet\alpha$. Now, from $\alpha \prec \bullet\alpha$, $\bullet\alpha \prec \bullet\bullet\alpha$, and transitivity of \prec , it follows that $\alpha \prec \bullet\bullet\alpha$, i.e., $\alpha \models \bullet\bullet\alpha$ and $\bullet\bullet\alpha \models \bullet\alpha$. The first result is now irrelevant. The second together with W gives us $\bullet\bullet\alpha \equiv \bullet\alpha$. Hence \bullet is idempotent. \square

Corollary 2.14. *Let \bullet be idempotent and \llcorner be defined in terms of \bullet . If $\alpha \llcorner \beta$, then $\bullet\alpha \equiv \bullet\beta$.*

Consider (just for simplicity in the next corollary), the case of a finitely generated propositional language \mathcal{L} . Then every subset X of the finite set W is axiomatisable by its *theory*, a sentence $\tau(X)$ of \mathcal{L} . Define the function $C : \mathcal{P}(W) \rightarrow \mathcal{P}(W)$; $X \mapsto \text{Mod}(\bullet\tau(X)) = C(X)$. Then we have:

Corollary 2.15. *If \bullet is idempotent, then C is a closure operator on $\mathcal{P}(W)$ [16], i.e.,*

1. $X \subseteq C(X)$;
2. if $X \subseteq Y$, then $C(X) \subseteq C(Y)$;
3. $C(C(X)) = C(X)$.

Are there weakening operators which are *not* idempotent? The weakening $\bullet\alpha \equiv \alpha \vee \beta$ (for a fixed β) mentioned at the start of Section 2.1, as well as the forgetting weakening operator of Example 2.2, are both surely idempotent. The following example (granted to be rather artificial) demonstrates that there are non-idempotent weakening operators — and hence non-transitive constrained entailments.

Example 2.16. Language \mathcal{L} is generated by the atoms p_1, p_2, \dots, p_n together with \top , and W has 2^n elements. Each $w \in W$ is seen as a truth assignment on the atoms: $w(p_i) \in \{0, 1\}$, with $w(\top) = 1$, which links w uniquely to the natural number with binary expansion $n(w) = \sum_{i=1}^n w(p_i)2^{i-1}$, $0 \leq n(w) < 2^n$. We index w as w_j if $n(w) = j - 1$, inducing the linear order $w_1 < w_2 < \dots < w_{2^n}$ on W by stipulating that $w < w'$ if and only if $n(w) < n(w')$ in the natural numbers. For any sentence $\alpha \in \mathcal{L}$ we now define $\bullet\alpha$ to be the theory of a certain $<$ -induced expansion of $\text{Mod}(\alpha)$ within W :

$$\bullet\alpha \equiv_{\text{def}} \tau(\{w \in W \mid (\text{there is } w' \in \text{Mod}(\alpha) \text{ such that } w \leq w') \text{ or } w = w_{m+1}\}),$$

where w_m is the maximum element of $\text{Mod}(\alpha)$. It is not onerous to see that this \bullet is a uniform weakening operator which is not idempotent. (Later, at the end of Section 3.2, we shall see that in the special modal construal of \llcorner it is transitive only in the logic **S4**, or stronger.)

For two uniform weakening operators \bullet_1 and \bullet_2 we write $\bullet_1 \leq \bullet_2$ to indicate that for every α we have that $\bullet_1\alpha \models \bullet_2\alpha$. Let \llcorner_1 and \llcorner_2 be the respectively corresponding constrained entailment relations. From Definitions 2.4 and 2.7 now easily follows that:

$$\bullet_1 \leq \bullet_2 \text{ if and only if } \llcorner_1 \subseteq \llcorner_2 \tag{2.5}$$

Let $(\alpha, \beta) \in \llcorner_1$. Then $\alpha \models \beta$ and $\beta \models \bullet_1\alpha$. From this and the hypothesis $\bullet_1\alpha \models \bullet_2\alpha$ we get $\beta \models \bullet_2\alpha$. Putting the results together gives us $\alpha \llcorner_2 \beta$. For the converse, suppose now that $\llcorner_1 \subseteq \llcorner_2$. Then $\bullet_1\alpha \equiv \bigvee\{\beta \mid \alpha \llcorner_1 \beta\} \models \bigvee\{\beta \mid \alpha \llcorner_2 \beta\} \equiv \bullet_2\alpha$. Therefore, we have $\bullet_1 \leq \bullet_2$.

The minimum, with respect to \leq , most frugal, uniform weakening operator is given by $\bullet\alpha \equiv \alpha$. This corresponds to the *maximum restriction* of

the relation \llcorner , namely the case $\llcorner = \equiv$ (i.e., logical equivalence), since now $\beta \models \bullet\alpha$ is $\beta \models \alpha$, and β is a constrained consequence of α if and only if $\alpha \equiv \beta$. On the other hand, the maximum, most prodigal, operator, $\bullet\alpha \equiv \top$, delivers the *minimum restriction* of \llcorner , namely when $\llcorner = \models$ (no constraint at all). Hence \llcorner can, depending on the different \bullet , be any one of a whole *spectrum* of constrained entailment relations, but for all of them we have the following property, of which the proof is straightforward:

$$\equiv \subseteq \llcorner \subseteq \models \quad (2.6)$$

One of the Artificial Intelligence contexts in which constrained consequence (moving in the spectrum away from \models towards somewhere nearer \equiv) may be useful occurs when we attempt the formal modelling of aspects of everyday human reasoning as precipitated in natural language [26].

One expects the common-sense reasoning of people untrained in formal logic to be much more ‘pertinent’ than reasoning based on \models and \rightarrow . Note how, psychologically speaking, with increased pertinence between premiss and consequence (or antecedent and consequent) ‘if’ tends to drift in the direction of ‘if and only if’. (The “if” in “I’ll buy you a bicycle if you pass the exam” usually means “if and only if”.) And indeed, researchers in cognitive science find a strong tendency to interpret “if” as “if and only if” [28, 12, 27]. (See also Chapter 21, “The Puzzles of *If*” [26, pp. 295–310].)

Should we define \bullet by $\bullet\alpha \equiv \top$ when $\alpha \not\equiv \perp$, and $\bullet\perp \equiv \perp$, we get that $\alpha \llcorner \beta$ if and only if $\alpha \models \beta$ when $\alpha \not\equiv \perp$; but $\perp \llcorner \beta$ if and only if $\beta \equiv \perp$. This particular \llcorner then is *non-explosive*: any sentence \llcorner -entailed by a contradiction is itself a contradiction, which is a weak form of *paraconsistency* [39]. (In Subsection 3.3 we shall give an example of such \llcorner .)

Whereas \llcorner is infra-classical, “constrained logical equivalence” turns out to be just classical logical equivalence — similar to what holds for the “constrained biconditional” later in Property (2.10). It is easy to verify the following property:

$$\alpha \llcorner \beta \text{ and } \beta \llcorner \alpha \text{ if and only if } \alpha \equiv \beta \quad (2.7)$$

Paratriviality. Our constrained entailment relation is *paratrivial* in the sense that *verum* is not omnigenerated: the constrained entailment of a tautology from an arbitrary sentence is not trivial. Consider $\alpha \llcorner \top$ with a contingent α . From $\alpha \llcorner \top$, we get $\top \models \bullet\alpha$, and then it follows that $\bullet\alpha$ is a tautology — a rather strong stricture on $\bullet\alpha$. Intuitively, only the assumption that the α -worlds collectively must be expanded to the whole of W justifies the entailment of \top from α . The intuition, more generally, appreciates that, although a tautology yields no information, exploring whether and how it may be reached via a specific inferential relation from a specific premiss may be informative. Even while approaching the issue from a very different angle, Duží also formulates the insight that “. . . although analytically true sentences provide no *empirical information* about the state of the world, they convey *analytic information*, in the shape of constructions prescribing how to arrive at the truths in question” [20, p. 473].

Another interesting property of \llcorner is that the set of *constrained tautologies* of the language is identical to the set of all tautologies:

$$\top \llcorner \alpha \text{ if and only if } \top \models \alpha \quad (2.8)$$

$\top \llcorner \alpha$ if and only if $\top \models \alpha$ and $\alpha \models \bullet\top$ if and only if $\top \models \alpha$ and $\alpha \models \top$, by Property (2.4), if and only if $\top \models \alpha$.

No Contraposition. Classically we have *contraposition*: $\alpha \models \beta$ is equivalent to $\neg\beta \models \neg\alpha$. Not so for \llcorner , and proof by contradiction (if $\alpha \wedge \neg\beta$ entails \perp , then α entails β) does not hold in general either. $\neg\beta \llcorner \neg\alpha$ says that $\alpha \models \beta$ and $\neg\alpha \models \bullet\neg\beta$: Every α -world is a β -world and every $\neg\alpha$ -world is included in the expansion of the $\neg\beta$ -worlds. This may be an entailment relation worthy of study, but which we shall not pursue further in this paper.

No Deduction Theorem for \rightarrow . Now one question that naturally arises is whether the classical meta-theorem called *deduction*, or by some authors the *Ramsey test* for conditionals [13] ($\alpha \models \beta$ is equivalent to $\top \models \alpha \rightarrow \beta$), also holds for \llcorner . So, is it the case that $\alpha \llcorner \beta$ if and only if $\top \llcorner \alpha \rightarrow \beta$?

For the left-to-right direction, suppose that $\alpha \llcorner \beta$, i.e., $\alpha \models \beta$ and $\beta \models \bullet\alpha$. Then $\top \models \alpha \rightarrow \beta$ and surely $\alpha \rightarrow \beta \models \bullet\top$, since $\alpha \rightarrow \beta \models \top$ and $\bullet\top \equiv \top$, by Property (2.4). Now, for the right-to-left direction, let us assume that $\top \llcorner \alpha \rightarrow \beta$, i.e., $\top \models \alpha \rightarrow \beta$ and $\alpha \rightarrow \beta \models \bullet\top$. The second statement is just the triviality $\alpha \rightarrow \beta \models \top$. We do not (in general) get the needed $\beta \models \bullet\alpha$.

Hence, $\alpha \llcorner \beta$ implies $\top \llcorner \alpha \rightarrow \beta$, but not conversely — unless every β -world is included in the weakening of all α -worlds, which is precisely the constraint aspect of the definition of the relation \llcorner .

It turns out that in our constrained approach it is not difficult to define a conditional connective which *does* satisfy the Ramsey test. Here we define the new binary connective $\dot{\rightarrow}$, called the *constrained conditional*, as follows:

Definition 2.17. $\alpha \dot{\rightarrow} \beta \equiv_{\text{def}} (\alpha \rightarrow \beta) \wedge (\beta \rightarrow \bullet\alpha)$.

Then we have the following straightforward result:

$$\alpha \llcorner \beta \text{ if and only if } \top \llcorner \alpha \dot{\rightarrow} \beta \quad (2.9)$$

The proof of the following property is also easy:

$$(\alpha \dot{\rightarrow} \beta) \wedge (\beta \dot{\rightarrow} \alpha) \equiv \alpha \leftrightarrow \beta \quad (2.10)$$

So the “constrained bi-conditional” (say $\dot{\leftrightarrow}$) is just the classical bi-conditional! Remember Property (2.7): “constrained logical equivalence” is just classical logical equivalence.

Non-Monotonicity. For entailment \llcorner , the following monotonicity rule *fails*:

$$\frac{\alpha \llcorner \beta, \gamma \models \alpha}{\gamma \llcorner \beta}$$

So, assuming $\alpha \llcorner \beta$, we have *no* guarantee that $\alpha \wedge \alpha' \llcorner \beta$: some β -world may not be included in the weakening of the $\alpha \wedge \alpha'$ -worlds, even though it is included in the weakening of the α -worlds. From Theorem 2.13 and

Example 2.16 we see that even the weaker rule, where “ $\gamma \models \alpha$ ” is replaced by “ $\gamma \prec \alpha$ ”, does not hold in general. This result for infra-classical \prec stands in contrast to one of the tacit expectations in the non-monotonic reasoning literature [8, 30, 35], viz. that non-monotonic entailment relations are *a priori* supra-classical, or, at least, are ampliative in the sense that they are obtained by relaxing the underlying (possibly non-classical) monotonic entailment [3].

3. Constrained Modal Entailment

In a direct proof of an entailment there is a step-by-step ‘logical movement’ from premiss to consequence; in an indirect proof of β from α , such as *reductio ad absurdum* or by contraposition, there is, implicitly, a notion of transition from $\alpha \wedge \neg\beta$ to \perp or from $\neg\beta$ to $\neg\alpha$ — ‘directed movement’, to and fro.

This intuitive notion of entailment as a species of access relation between sentences or propositions — starting at the premiss *access to* the consequence, or starting at the consequence *access from* the premiss — this idea of entailment as ‘access’ seems appealing as a way of constraining classical entailment and has a natural analogue in the *accessibility relation* between *worlds* in modal logic. In the present section we investigate how these ideas can be used in the definition of a concrete uniform weakening operator and its associated constrained entailment.

3.1. Modal Logic

To get to a (propositional) multi-modal logic, we extend our propositional language with a family of normal modal operators \Box_i [7, 38], $1 \leq i \leq n$ for a given n . The formulas will then be recursively defined by:

$$\alpha ::= p \mid \top \mid \neg\alpha \mid \alpha \wedge \alpha \mid \Box_i \alpha$$

(As before, the other connectives and the special atom \perp are defined in terms of \neg and \wedge in the usual way.) As expected, the dual of each \Box_i , namely \Diamond_i , is defined by $\Diamond_i \alpha \equiv_{\text{def}} \neg \Box_i \neg \alpha$.

Definition 3.1. A *model* is a tuple $\mathcal{M} = \langle W, R, V \rangle$, where

- W is a set of *worlds* (or *states*);
- $R = \langle R_1, \dots, R_n \rangle$, where each $R_i \subseteq W \times W$ is an *accessibility relation* on W ; and
- $V : \mathfrak{P} \times W \longrightarrow \{0, 1\}$ is a *valuation*.

Definition 3.2. Given a model $\mathcal{M} = \langle W, R, V \rangle$ and a world $w \in W$

- $w \Vdash^{\mathcal{M}} \top$;
- $w \Vdash^{\mathcal{M}} p$ if and only if $V(p, w) = 1$;
- $w \Vdash^{\mathcal{M}} \neg\alpha$ if and only if $w \not\Vdash^{\mathcal{M}} \alpha$;
- $w \Vdash^{\mathcal{M}} \alpha \wedge \beta$ if and only if $w \Vdash^{\mathcal{M}} \alpha$ and $w \Vdash^{\mathcal{M}} \beta$;
- $w \Vdash^{\mathcal{M}} \Box_i \alpha$ if and only if $w' \Vdash^{\mathcal{M}} \alpha$ for every w' such that $(w, w') \in R_i$, $1 \leq i \leq n$;
- truth conditions for the other connectives are as usual.

Definition 3.3. Given a model $\mathcal{M} = \langle W, R, V \rangle$ and a formula α ,

- If $w \Vdash^{\mathcal{M}} \alpha$ for a given $w \in W$, then we say that w *satisfies* α , or *is a model of* α with respect to \mathcal{M} ;
- If $w \Vdash^{\mathcal{M}} \alpha$ for every $w \in W$, then we say that α is *valid* in \mathcal{M} , denoted $\Vdash^{\mathcal{M}} \alpha$.

Among all possible models, one might want to choose some with specific properties to work with. This defines a *class of models*. A class of models \mathcal{C} can be determined by imposing additional properties on the accessibility relations (e.g. transitivity, reflexivity, etc., which is usually done by stating *axiom schemas*), or by means of *global axioms* (instances of formulas one wants to be valid in the class) imposing restrictions on each W [6, 21]. In Section 3.2 we show how to define a class of models with the former approach, whereas in Section 3.3 we give an example illustrating ways in which \mathcal{C} can be defined with both.

Here we employ the following versions of *local consequence*:

Definition 3.4. Given a model $\mathcal{M} = \langle W, R, V \rangle$ and formulas α and β , we say that α *entails* β in \mathcal{M} (noted $\alpha \Vdash^{\mathcal{M}} \beta$) if and only if for every $w \in W$, if $w \Vdash^{\mathcal{M}} \alpha$, then $w \Vdash^{\mathcal{M}} \beta$.

Definition 3.5. Given a class \mathcal{C} of models and formulas α and β , if $\alpha \Vdash^{\mathcal{M}} \beta$ for every $\mathcal{M} \in \mathcal{C}$, we say that α *entails* β in \mathcal{C} (noted $\alpha \Vdash^{\mathcal{C}} \beta$), i.e., we have $\Vdash^{\mathcal{C}} = \bigcap \{ \Vdash^{\mathcal{M}} \mid \mathcal{M} \in \mathcal{C} \}$.

If $\Vdash^{\mathcal{M}} \alpha$ for every $\mathcal{M} \in \mathcal{C}$, we say that α is *valid* in \mathcal{C} (noted $\Vdash^{\mathcal{C}} \alpha$). If $\neg \alpha$ is *not* valid in \mathcal{C} , we say that α is *satisfiable* in \mathcal{C} . Since the specific class of models we are working with will be made clear from the context, for the sake of readability we shall dispense with superscripts and just write $\alpha \Vdash \beta$ instead of $\alpha \Vdash^{\mathcal{C}} \beta$.

3.2. A Modal Construal of Constrained Consequence

We now anchor our notion of constrained entailment in the *use* of the accessibility relation on worlds. To be specific: In our new entailment of β by α , we impose upon β -worlds the condition that each of them should be *accessible* from *some* α -world. This is how those β -worlds which are not α -worlds in a modal entailment $\alpha \Vdash \beta$ are ‘disciplined’.

First of all, in order for us to be able to talk about worlds from and to which can be accessed, we will need a *bi-modal* logic: one modal operator for talking about successor worlds, and one for talking about predecessor ones, similar to what one has in a temporal logic with a Future and a Past operator. Let \Box denote the former, and \Boxleftarrow the latter. (One could have chosen \Box_1 and \Box_2 , but, as will become clear in the sequel, the notation \Box and \Boxleftarrow suits better the underlying intuition we have of these operators.)

The idea here is to link \Box and \Boxleftarrow such that each one is the *converse* of the other: whenever one can go from a world w to w' via the accessibility

relation associated to \Box , one should be able to ‘come back’ from w' to w via the accessibility relation corresponding to \Box .

Models of this bi-modal language will be structures of the form $\mathcal{M} = \langle W, \langle R, \check{R} \rangle, V \rangle$, where \check{R} is the *converse* of R . So, from now on, R denotes a single accessibility relation (not a sequence $\langle R_1, \dots, R_n \rangle$).

To make sure that \Box has the intended behaviour of being \Box 's converse, we need to add two axiom schemas which then axiomatize the fact that \check{R} is the converse of R . This is done by requiring that all instances of the following axiom schemas be valid (remembering, \Diamond is the dual operator of \Box):

$$\alpha \rightarrow \Box \check{\Diamond} \alpha \quad (3.1)$$

$$\alpha \rightarrow \check{\Box} \Diamond \alpha \quad (3.2)$$

Therefore, here we will be interested in the class of models in which Schemas (3.1) and (3.2) above are valid. For reasons that will become clear in the sequel, we also want our models here to have *reflexive* accessibility relations, i.e., given a model $\mathcal{M} = \langle W, \langle R, \check{R} \rangle, V \rangle$ we want $id_W \subseteq R$, where id_W is the *identity relation* on W . (Since \check{R} is R 's converse, ensuring that only for R is enough.) We axiomatize that semantic condition by stating the following extra axiom schema:

$$T : \Box \alpha \rightarrow \alpha$$

This defines the modal logic KT [15]. Hence, in the rest of this section we work in the class of KT-models with converse. (It is easy to check that $\Box \alpha \rightarrow \alpha$ is also valid in this class. More generally, and leaving aside intuition, it is clear that from a technical point of view \Diamond and $\check{\Diamond}$ may be interchanged, and results obtained for the one also hold for the other.)

Having settled the underlying modal formalism, one might now wonder which definition of a weakening operator would be a suitable one in this specific modal setting. Given our discussion in the beginning of this subsection, it turns out that $\check{\Diamond}$ is a good candidate:

Theorem 3.6. $\check{\Diamond}$ is a uniform weakening operator.

Proof. All we need to show is that $\check{\Diamond}$ satisfies Properties W and U.

W: Let $\mathcal{M} = \langle W, \langle R, \check{R} \rangle, V \rangle$ be a KT-model. For every $\alpha \in \mathfrak{F}$, from the hypothesis that Schema T is valid, we have that $w \Vdash^{\mathcal{M}} \check{\Box} \neg \alpha \rightarrow \neg \alpha$ for every $w \in W$, i.e., $w \Vdash^{\mathcal{M}} \alpha \rightarrow \check{\Diamond} \alpha$ for every $w \in W$. This means that for every $w \in W$, if $w \Vdash^{\mathcal{M}} \alpha$ then $w \Vdash^{\mathcal{M}} \check{\Diamond} \alpha$, and then $\alpha \models^{\mathcal{M}} \check{\Diamond} \alpha$. Since \mathcal{M} is an arbitrary KT-model, the result follows.

U: Let $\alpha, \beta \in \mathfrak{F}$ be such that $\alpha \models \beta$ (in the class of KT-models). Let now $\mathcal{M} = \langle W, \langle R, \check{R} \rangle, V \rangle$ be a KT-model, and let $w \in W$ be a world such that $w \Vdash^{\mathcal{M}} \check{\Diamond} \alpha$. Then, there is $w' \in W$ such that $(w, w') \in \check{R}$ and $w' \Vdash^{\mathcal{M}} \alpha$. From the hypothesis $\alpha \models \beta$, it follows that $w' \Vdash^{\mathcal{M}} \alpha$ implies $w' \Vdash^{\mathcal{M}} \beta$. Therefore, there is $w' \in W$ such that $(w, w') \in \check{R}$ and $w' \Vdash^{\mathcal{M}} \beta$, and hence $w \Vdash^{\mathcal{M}} \check{\Diamond} \beta$.

Since w is arbitrary, it follows that $\diamond\check{\alpha} \models^{\mathcal{M}} \diamond\check{\beta}$. Because \mathcal{M} is also arbitrary, the result follows. \square

An immediate consequence of Theorem 3.6 is that defining a constrained entailment relation according to Definition 2.4 in terms of $\diamond\check{}$ will give us an entailment relation with all the properties we discussed in Section 2.

Definition 3.7. α *constrainedly entails* β in the KT-model \mathcal{M} (noted $\alpha \prec^{\mathcal{M}} \beta$) if and only if $\alpha \models^{\mathcal{M}} \beta$ and $\beta \models^{\mathcal{M}} \diamond\check{\alpha}$. α *constrainedly entails* β in the class \mathcal{C} of KT-models (noted $\alpha \prec^{\mathcal{C}} \beta$) if and only if for every $\mathcal{M} \in \mathcal{C}$, $\alpha \prec^{\mathcal{M}} \beta$: $\prec^{\mathcal{C}} = \bigcap \{\prec^{\mathcal{M}} \mid \mathcal{M} \in \mathcal{C}\}$.

When the KT-model or the class of KT-models we are working with is clear from the context, we shall dispense with superscripts and write $\alpha \prec \beta$ instead of $\alpha \prec^{\mathcal{M}} \beta$ and $\alpha \prec^{\mathcal{C}} \beta$.

Intuitively, Definition 3.7 states that premiss α constrainedly entails consequence β if and only if α entails β and every β -world is accessible from *some* α -world — importantly, the $\beta \wedge \neg\alpha$ -worlds (Figure 2). (The α -worlds are each accessible from themselves.)




FIGURE 2. Constrained consequence of β from α : α -worlds are β -worlds and *any* β -world is accessible from *some* α -world.

The second part in Definition 3.7 adds an important restriction to the traditional (permissive) modal entailment. It says: from every β -world we can look back to some world, possibly different from where we are, and from which we could have come, in which α is true. Obviously, \prec is an *infra-modal* entailment relation: if $\alpha \prec \beta$, then $\alpha \models \beta$ (cf. the discussion in Sections 1 and 2). (In this section \models denotes modal entailment, of course.)

Let $\mathcal{M} = \langle W, \langle R, \check{R} \rangle, V \rangle$ be such that $id_W \subseteq R \subseteq W \times W$. Then, recalling our discussion just before Property (2.6) in Section 2.4, the case of the *minimum* accessibility relation (with respect to \subseteq), i.e., in any subclass \mathcal{C} of KT-models $\mathcal{M} = \langle W, \langle id_W, id_W \rangle, V \rangle$, corresponds to the *maximum restriction* of the relation \prec , namely the case $\prec = \equiv$, since now $\beta \models \diamond\check{\alpha}$ says that $\beta \models \alpha$. On the other hand, let $\models_{<} \equiv \models \setminus \{(\perp, \beta) \mid \beta \not\models \perp\}$. Then the *maximum* case, i.e., in any subclass \mathcal{C} of KT-models $\mathcal{M} = \langle W, \langle R, \check{R} \rangle, V \rangle$ such that $R = W \times W$, corresponds to the *minimum constriction* of \prec , namely when

$\llcorner = \vDash_{\llcorner}$ (since now $\beta \vDash \diamond\alpha$ says that $\beta \not\equiv \perp$ implies $\alpha \not\equiv \perp$). Therefore, our modally construed \llcorner also gives us a whole spectrum of entailment relations, this time ranging between \equiv and \vDash_{\llcorner} :

$$\equiv \subseteq \llcorner \subseteq \vDash_{\llcorner} \quad (3.3)$$

Notice that reflexivity of R is required in the proof of $\equiv \subseteq \llcorner$. That is why we have chosen to work with KT-models.

We now briefly investigate two rather special features of modal constrained entailment. While constrained entailment in general (as in Section 2) can be *explosive* (think of the case $\llcorner = \vDash$), Property (3.3) above already suggests that here things are not as permissive.

Non-Explosiveness. Our modal instance of \llcorner is *non-explosive* in the strong sense that *falsum* is not omnigenerating, in fact, it is only self-generating: if $\perp \llcorner \beta$, then $\beta \equiv \perp$. No contingent or tautological sentence is \llcorner -entailed by a contradiction, i.e., by a sentence logically equivalent to \perp . This follows from the fact that for $\perp \llcorner \beta$ to hold, $\beta \vDash \diamond\perp$ has to be the case, which holds only when $\beta \equiv \perp$. More generally, given a class of models \mathcal{C} , we have the following result, of which the proof is easy to show:

Theorem 3.8. *Let $\alpha \llcorner^{\mathcal{C}} \beta$. If $\vDash^{\mathcal{C}} \alpha \rightarrow \perp$, then $\vDash^{\mathcal{C}} \beta \rightarrow \perp$.*

In other words, no sentence satisfiable in a class \mathcal{C} of models is $\llcorner^{\mathcal{C}}$ -entailed by a sentence *unsatisfiable* in that class.

We conclude this subsection with a useful observation: By also requiring the underlying class of models to be *transitive*, i.e., if instead of KT we work in the modal logic S4 [15], where \diamond is idempotent (and we remember Theorem 2.13), then we get a constrained entailment that satisfies some additional, and contextually desirable rules, notably:

Transitivity.

$$\frac{\alpha \llcorner \beta, \beta \llcorner \gamma}{\alpha \llcorner \gamma}$$

The strong special properties of modal \llcorner (non-explosiveness in KT, transitivity in S4) means that there can be no “converse” to Theorem 3.6: not every uniform weakening operator can be construed as a modal \diamond . Example 2.16 concurs.

3.3. Causal Action and Constrained Consequence

The following example illustrates the potential links between modal constrained consequence and a notion of *causation* by *action*. But first we have to consider whether it seems intuitively plausible that there may be such links. Causality resulting from action involves the flow of time, earlier states of a system evolving into later states. Effects are simultaneous with or later in time than causes, never earlier (except, perhaps, in the mind of some theoretical quantum physicist). The causal flow of states over time, prompted by some action, can be modelled by an accessibility relation on states [44, 14].

Let us sketch how $\alpha \prec^{\mathcal{M}} \beta$, with $\mathcal{M} = \langle W, \langle R, \check{R} \rangle, V \rangle$, may be deemed one reasonable construal of “ α causes β ”. Think of the worlds in W (as interpreted under V) as the systemically possible states of some system S . When w and w' are possible states of S , read wRw' as “subject to the laws of nature applicable to S , it is possible for S in state w , subjected to some action, to evolve over time into S in state w' .” (Such is the reading adopted in many causal approaches to reasoning about action and change [33, 42, 44].) Then $\text{Mod}(\check{\diamond}\alpha)$ is the set of all those states of S which are, or can over time evolve from, α -states; i.e., all α -states together with all their possible future states. Now it seems apt when claiming that “ α causes β ” to be asserting that (i) in all α -states β holds: $\alpha \models \beta$, a ‘now’ perspective in which “no action yet” is seen as a special case of ‘action’; and (ii) every β -state, even if it is not an α -state, could, with action, over time have evolved from some α -state: $\beta \models \check{\diamond}\alpha$, a ‘time flow’ causal constraint.

Example 3.9. Let us now consider the following variant of the Yale Shooting Problem in reasoning about actions, called the Walking Turkey Scenario [4]: Assume that we want to hunt a turkey, which may be alive or not, and which may either be walking around or not. In such a scenario we have one action, namely that of shooting the turkey with a gun. Our language has the atomic propositions $\mathfrak{P} = \{s, a, w, \top\}$. Let s be interpreted as “the turkey is shot” (with the intuitive meaning that the action shoot has taken place); a as “the turkey is alive”; and w as “the turkey is walking”.

Assume we have the (fairly rudimentary) formalization of this scenario given by the set of global axioms $\mathcal{K} = \{w \rightarrow a, s \rightarrow \neg a, \diamond s, s \rightarrow \Box s\}$. (Alternative specifications can be found in the literature [42, 25].) The intuition behind the formulas in \mathcal{K} is that “a walking turkey is alive”; “a shot turkey is dead”; “it is possible to shoot the turkey”; and “shooting cannot be undone”.

Now suppose that we want to define a class of transitive models (cf. end of Section 3.2) in which every axiom in \mathcal{K} is valid. First we make sure that the axiom schema 4 ($\Box\alpha \rightarrow \Box\Box\alpha$) [15] holds, and then we cull down the transitive models in which some formula in \mathcal{K} is *not* valid. One of the resulting models is depicted in Figure 3. (Since we consider local consequence, in this example we use just a single model as illustration. Constrained entailment in the subclass \mathcal{C} of S4-models consistent with the conjunction of \mathcal{K} then follows by generalizing over all models in \mathcal{C} — cf. Definition 3.7.) For the sake of presentation, we only depict the accessibility relation R in the figure (the converse \check{R} can be inferred easily from R).

Here we are interested in entailments pertinent to the question: given that β is observed, is α a *cause* of β relative to the shooting action?

While classically we have $\neg a \wedge \neg w \models \neg a$ and $\neg a \wedge \neg w \models \neg w$, we now get $\neg a \wedge \neg w \prec \neg a$, but $\neg a \wedge \neg w \not\prec \neg w$: the turkey could be alive and still prefer not to walk! (Remember Ockham’s razor.) We do not have $a \wedge \Box \neg s \prec a$ (explanation incompatible with background assumption — cf. Theorem 3.8). On the other hand, we do have $a \wedge \Box \diamond s \prec a$ (substitution of equivalents, since $\Box \diamond s$ is valid in this class of models). Observing just that the turkey is dead




FIGURE 3. A model induced by transitivity of the accessibility relations and by global axioms $\mathcal{K} = \{w \rightarrow a, s \rightarrow \neg a, \diamond s, s \rightarrow \Box s\}$.

is not enough to postulate shooting as the cause of death: in the model \mathcal{M} of Figure 3, we do *not* have $s \prec \neg a$ (even though $s \models \neg a$ in \mathcal{M}). The reason is that in \mathcal{M} $\neg a \not\models \diamond s$, since $w_1 \models \neg a$ but $w_1 \not\models \diamond s$. This means that death happens in a world in which no shot is fired and which cannot even evolve from any world where there is shooting. Similarly, the “constrained contrapositive” of $s \prec \neg a$, namely $a \prec \neg s$, is also *not* valid: being alive is surely not a plausible cause of being not shot.

On the other hand, when the turkey is observed to be shot and dead, or, even more completely, shot and dead and not-walking, then we can abduce or diagnose being shot as a plausible cause of the observed poor state of the turkey: in \mathcal{M} both $s \prec s \wedge \neg a$ and $s \prec s \wedge \neg a \wedge \neg w$ are valid. Here we have illustrated how, in a suitable context of causality, constrained consequence may be employed as a framework for *abduction* and *diagnosis* — reasoning processes which move from an observed situation (a constrained consequence, the “effect”) to a premiss (a plausible “cause” of the effect). For α to be abductively postulated as a plausible cause of β , we require that, additional to $\alpha \models \beta$, β *stays* within the scope of what could possibly evolve from α in the presence of some action.

4. Discussion and Related Work

In this section we compare our general construction from Section 2 and, whenever appropriate, its modal instance from Section 3.2, to existing work in the literature sharing some of the properties of \prec . We start with a discussion of \prec in the light of preferential reasoning, followed by a comparison with relevance logics.

As already pointed out above, our work also has links with *forgetting* [34, 32] that we have not pursued here. Like *epistemic relevance* [31, 45], our form of constrained consequence is also infra-classical and non-monotonic. Contrary to the mentioned approaches, however, our constructions are not formulated in terms of prime implicates, and therefore we are not limited to the propositional case (as shown in Section 3).

4.1. Preferential Reasoning

In Artificial Intelligence, there has been a great deal of work done on non-monotonic consequence relations [30, 8, 35, 22, 10, 11]. Here we give a brief outline of propositional preferential and rational consequence, as initially defined by Kraus et al. [30], and analyze how constrained consequence \llcorner stands in comparison to them.

A propositional defeasible consequence relation \sim is defined as a binary relation on formulas of an underlying (possibly infinitely generated) propositional logic. \sim is said to be *preferential* if it satisfies the following set of properties:

$$\begin{array}{ll}
 \text{(Ref)} & \alpha \sim \alpha \\
 \text{(And)} & \frac{\alpha \sim \beta, \alpha \sim \gamma}{\alpha \sim \beta \wedge \gamma} \\
 \text{(Or)} & \frac{\alpha \sim \gamma, \beta \sim \gamma}{\alpha \vee \beta \sim \gamma} \\
 \text{(LLE)} & \frac{\alpha \equiv \beta, \alpha \sim \gamma}{\beta \sim \gamma} \\
 \text{(RW)} & \frac{\alpha \sim \beta, \beta \models \gamma}{\alpha \sim \gamma} \\
 \text{(CM)} & \frac{\alpha \sim \beta, \alpha \sim \gamma}{\alpha \wedge \beta \sim \gamma}
 \end{array}$$

If, in addition to the properties of preferential consequence, \sim also satisfies the following Rational Monotonicity property, it is said to be a *rational* consequence relation:

$$\text{(RM)} \quad \frac{\alpha \sim \beta, \alpha \not\sim \neg\gamma}{\alpha \wedge \gamma \sim \beta}$$

Now we replace \sim with \llcorner . It follows from our definitions in Section 2.2 that \llcorner satisfies Ref above. Still in Section 2.2 we showed that \llcorner satisfies the derived properties RC and LD, which are, respectively, And and Or above. Property LLE follows from our E. Cautious Monotonicity (CM) can be shown to hold for \llcorner : From $\alpha \llcorner \beta$ and $\alpha \llcorner \gamma$ we get $\alpha \models \beta$, $\alpha \models \gamma$ and $\gamma \models \bullet\alpha$. $\alpha \models \gamma$ gives us $\alpha \wedge \beta \models \gamma$, while from $\alpha \models \beta$ (i.e., $\alpha \wedge \beta \equiv \alpha$) and $\gamma \models \bullet\alpha$ we get $\gamma \models \bullet(\alpha \wedge \beta)$. Putting both results together we get $\alpha \wedge \beta \llcorner \gamma$.

The constraint on \llcorner -consequences β of a given premiss α (to those β for which $\beta \models \bullet\alpha$) of course implies that \llcorner does *not* satisfy Right Weakening (RW). Therefore \llcorner is *not* a preferential relation in the sense of Kraus et al. Constrained entailment \llcorner does not satisfy the property of Rational Monotonicity either. On the other hand, we have seen (Section 2.2) that \llcorner does satisfy Cut. Hence, despite failing RW and RM, our non-monotonic inference relation shares a lot of the properties that are viewed as important in that

setting. A crucial difference between our work in this paper and the aforementioned is that our notion of constrained entailment is *infra*-classical, and therefore applicable in different contexts.

4.2. Relevance Logics

One intuitive connotation of entailment is that some additional relation of ‘relevance’ or ‘pertinence’, should hold between premiss and consequence. If rather specific, the extra information yielding the pertinence of premiss and consequence to each other is usually expressed either as syntactic rules or as semantic constraints, and typically involves an (often binary) relation on the set of sentences of the language. More generally and vaguely the ‘extra’ may be a desire to adapt classical entailment in order to obtain an entailment relation which more closely resembles everyday human reasoning as precipitated in natural language. (This was discussed briefly in Section 2.4, with reference to the cognitive science literature.) Starting with classical entailment, we invoke extra semantic information to trim down those entailment pairs in which premiss and consequence are not pertinent to each other. Syntactically, this extra semantic information may be induced by a weakening operator as introduced in Section 2.1. One reading of constrained consequence is therefore as a *pertinent entailment* relation [9].

Existing relevance logics [1, 2, 19] share some of the aims that we have with the present paper, for instance the elimination of some counter-intuitive entailment pairs.

Indeed, classical *disjunctive syllogism* — $(\neg\alpha \vee \beta) \wedge \alpha \models \beta$, which is classically equivalent to $\beta \wedge \alpha \models \beta$ — is a minor pet hate of some relevance logicians. Even though classically we have no problem with $(\neg\alpha \vee \beta) \wedge \alpha \models \beta$, one can appreciate that in $\beta \wedge \alpha \models \beta$ the α is rather irrelevant, at least in everyday human reasoning. In our setting, $\beta \wedge \alpha \prec \beta$ would mean that $\beta \wedge \alpha \models \beta$ and $\beta \models \bullet(\beta \wedge \alpha)$. This implies that when expanding the set of $\beta \wedge \alpha$ -worlds we have taken care of including at least *all* the β -worlds.

$(\neg\alpha \vee \beta) \wedge \alpha \models \beta$ is a version of *modus ponens*, viz. the *resolution rule* — while there are at least four different versions of *modus ponens* [13, p. 51]. For \prec we then saw that it holds only in a restricted and controlled way. A general form of *modus ponens* for \prec was proved in Section 2.2.

Some of the other *bêtes noires* of relevance logicians are the so-called paradoxes of material and strict implication. With the introduction of our (stricter) conditional $\dot{\rightarrow}$ (Definition 2.17), these are avoided, as shown below. Firstly, the positive paradoxes do not hold in general:

$$\top \not\prec \alpha \dot{\rightarrow} (\beta \dot{\rightarrow} \alpha), \quad \alpha \not\prec \beta \dot{\rightarrow} \alpha$$

Another paradox of material implication is $\models (\alpha \rightarrow \beta) \vee (\beta \rightarrow \alpha)$: given any two sentences, at least one implies the other. Its constrained version does not hold in general either:

$$\top \not\prec (\alpha \dot{\rightarrow} \beta) \vee (\beta \dot{\rightarrow} \alpha)$$

The following paradoxes of strict implication are also avoided:

$$\top \not\vdash (\alpha \wedge \neg\alpha) \dot{\rightarrow} \beta, \quad \top \not\vdash \alpha \dot{\rightarrow} (\beta \vee \neg\beta)$$

The discussion above invites the question whether our logic is a relevance logic. One of the consequences of adopting Property E (cf. Section 2.2) is that we do not have the *variable sharing property* required by most relevance logics [19]. For instance, in our approach we have that $p \wedge \neg p \not\vdash q \wedge \neg q$. In that sense, the logic we develop here is then (apart from other considerations) *not* a standard relevance logic.

Most of the existing work on substructural logics is proof-theoretic and has mainly focused on weakening the generating engine of *axioms* and *inference rules* to get rid of unwanted entailments [19, 40]. Other existing approaches are algebraic in nature [23].

Quite recently, after an early start [43], a few publications concentrated on further developing a proper semantics for relevant logics [37, 41, 5]. There, a possible worlds approach is also defined, but with the aid of a *ternary* accessibility relation between worlds. Having (w, w', w'') in the accessibility relation means different things for different authors. For Meyer “[w]orlds are best demythologized as theories”, and then, paraphrasing, theory w'' consists of all the outputs got by applying *modus ponens* in a certain way to major premisses from w and minor premisses from w' . Here we propose a simpler approach, in the modal case, viz. via a *binary* accessibility relation on plain W , for carrying out the required construction of constrained consequence.

Meyer claimed that in almost all standard relevance logics the *relevant* conditional (usually written as \rightarrow) *cannot* be defined as a modalized truth function [36]. More explicitly, Meyer proved that no standard *relevant* conditional can be represented as a “*strict*” $\Box\phi(\alpha, \beta)$, where $\phi(\alpha, \beta)$ is any truth functional combination of α and β . This prescription that “modalizing” in this context must mean “having \Box as main connective” is of course restrictive. A modal counterpart of our treatment in Definition 2.17 and Property 2.9 of the *constrained* conditional connective $\dot{\rightarrow}$ would show that we can achieve a similar objective in a quite elegant way, even if not with the main operator \Box . In a separate paper we are investigating, beyond the constrained conditional, other constrained connectives: negation, conjunction, and disjunction induced by different choices of the weakening operator, and their behaviour with respect to the corresponding constrained entailments.

Research on substructural logic usually adopts an *incremental* strategy (going from ‘nothing’ up to the entailments considered as relevant), and quite often via a proof-theoretic approach. Here we have followed a semantic-based *restrictive* strategy: we start from full classical logic (or full bi-modal logic KT) and then go down to infra-classical (or infra-modal) consequence by culling the irrelevant entailments. A similar strategy is that of ‘filtering out’ undesirable classical entailment pairs to prevent ‘explosion’ (*ex contradictione quodlibet*) in some paraconsistent logics [39, pp. 297–299].

5. Recapitulation

In this paper we defined general unary operators and binary relations, with a specific set of properties, on the set of formulae (sentences) of a (propositional) language, namely (i) uniform weakening functions $\bullet : \mathfrak{F} \rightarrow \mathfrak{F}$; and (ii) constrained entailment relations $\llcorner \subseteq \mathfrak{F} \times \mathfrak{F}$. We have shown that by suitable definitions, they describe inter-definable structures, representing constrained consequence in different inter-translatable ways.

We have seen in Properties (2.6) and (3.3) that our approach allows for a whole spectrum of constrained entailments, ranging between \equiv and \models (or $\models_{<}$), and offering potential reasoning tools for many different contexts (cf. Example 3.9 in Section 3.3).

Our constrained entailment relation \llcorner restricts some paradoxes shunned by relevance logics in an interesting way. Moreover, we have shown that \llcorner also possesses other non-classical properties, like strong non-explosiveness and non-monotonicity.

For \llcorner in general we remarked (following Theorem 2.5) that its semantic definition has sound and complete syntactic counterparts. With regards to a sound and complete proof theory for modal \llcorner , we can resort to existing decision procedures, notably tableaux [24] and resolution [17], for both conditions in Definition 3.7.

In the causal context of Example 3.9 we illustrated how entailment may be curtailed by background assumptions (i.e., a knowledge base) and by restricting the class of accessibility relations, thus preventing unwanted entailments of a formula from a premiss. An investigation on how to define a knowledge base and a class of accessibility relations given a particular type of application remains to be done.

Finally, while this paper deals with the case of infra-classical consequence, which shares some of the aims of traditional substructural logics and relevance logics, our constructions here can also be adapted to the context of an underlying *supra*-classical entailment, which relates to preferential reasoning and other forms of venturous reasoning.

References

- [1] A.R. Anderson and N.D. Belnap. *Entailment: the logic of relevance and necessity*, volume 1. Princeton University Press, 1975.
- [2] A.R. Anderson, N.D. Belnap, and J.M. Dunn. *Entailment: the logic of relevance and necessity*, volume 2. Princeton University Press, 1992.
- [3] O. Arieli and A. Avron. General patterns for nonmonotonic reasoning: From basic entailments to plausible relations. *Logic Journal of the IGPL*, 8:119–148, 2000.
- [4] A.B. Baker. Nonmonotonic reasoning in the framework of situation calculus. *Artificial Intelligence*, 49:5–23, 1991.

- [5] J. Beall, R. Brady, J.M. Dunn, A.P. Hazen, E. Mares, R.K. Meyer, G Priest, G. Restall, D. Ripley, J. Slaney, and R. Sylvan. On the ternary relation and conditionality. *Journal of Philosophical Logic*. Online: 7 May 2011.
- [6] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2001.
- [7] P. Blackburn, J. van Benthem, and F. Wolter. *Handbook of Modal Logic*. Elsevier North-Holland, 2006.
- [8] C. Boutilier. Conditional logics of normality: A modal approach. *Artificial Intelligence*, 68(1):87–154, 1994.
- [9] K. Britz, J. Heidema, and I. Varzinczak. Pertinent reasoning. In *13th International Workshop on Nonmonotonic Reasoning (NMR)*, 2010.
- [10] K. Britz, T. Meyer, and I. Varzinczak. Preferential reasoning for modal logic. In *Proc. of Methods for Modalities*, 2011.
- [11] K. Britz, T. Meyer, and I. Varzinczak. Semantic foundation for preferential description logics. In *Proc. 24th Australasian Conference on AI*, 2011.
- [12] R.M.J. Byrne and P.N. Johnson-Laird. ‘If’ and the problems of conditional reasoning. *Trends in Cognitive Sciences*, 13(7):282–287, 2009.
- [13] J. Cantwell. Conditionals in reasoning. *Synthese*, 171:47–75, 2009.
- [14] M.A. Castilho, A. Herzig, and I. Varzinczak. It depends on the context! A decidable logic of actions and plans based on a ternary dependence relation. In *9th International Workshop on Nonmonotonic Reasoning (NMR)*, 2002.
- [15] B. Chellas. *Modal logic: An introduction*. Cambridge University Press, 1980.
- [16] B.A. Davey and H.A. Priestley. *Introduction to Lattices and Order*. Cambridge Mathematical Textbooks. Cambridge University Press, 1990.
- [17] H. de Nivelle, R. Schmidt, and U. Hustadt. Resolution-based methods for modal logics. *Logic Journal of the IGPL*, 8(3):265–292, 2000.
- [18] M. Dickman. *Large Infinitary Languages*. North-Holland, 1975.
- [19] J.M. Dunn and G. Restall. Relevance logic. In D.M. Gabbay and F. Günthner, editors, *Handbook of Philosophical Logic*, volume 6, pages 1–128. Kluwer Academic Publishers, 2nd edition, 2002.
- [20] M. Duží. The paradox of inference and the non-triviality of analytic information. *Journal of Philosophical Logic*, 39:473–510, 2010.
- [21] D.M. Gabbay, A. Kurucz, M. Zakharyashev, and F. Wolter. *Many-Dimensional Modal Logics: Theory And Applications*. Number 148 in Studies in Logic and the Foundations of Mathematics. Elsevier North-Holland, 2003.
- [22] D.M. Gabbay and K. Schlechta. Roadmap for preferential logics. *Journal of Applied Non-Classical Logic*, 19(1):43–95, 2009.
- [23] N. Galatos, P. Jipsen, T. Kowalski, and H. Ono. *Residuated Lattices: An Algebraic Glimpse at Substructural Logics*. Studies in Logic and the Foundations of Mathematics. Elsevier, 2007.
- [24] R. Goré. Tableau methods for modal and temporal logics. In M. D’Agostino, D.M. Gabbay, R. Hähnle, and J. Posegga, editors, *Handbook of Tableau Methods*, pages 297–396. Kluwer Academic Publishers, 1999.
- [25] A. Herzig and I. Varzinczak. Metatheory of actions: beyond consistency. *Artificial Intelligence*, 171:951–984, 2007.
- [26] P.N. Johnson-Laird. *How we reason*. Oxford University Press, 2006.

- [27] P.N. Johnson-Laird, R.M.J. Byrne, and V. Girotto. The mental model theory of conditionals: A reply to Guy Politzer. *Topoi*, 28:75–80, 2009.
- [28] P.N. Johnson-Laird and F. Savary. Illusory inferences: a novel class of erroneous deductions. *Cognition*, 71:191–229, 1999.
- [29] C. Karp. *Languages with Expressions of Infinite Length*. North-Holland, 1964.
- [30] S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44:167–207, 1990.
- [31] G. Lakemeyer. Relevance from an epistemic perspective. *Artificial Intelligence*, 97:137–167, 1997.
- [32] J. Lang, P. Liberatore, and P. Marquis. Propositional independence: Formula-variable independence and forgetting. *Journal of Artificial Intelligence Research*, 18:391–443, 2003.
- [33] F. Lin. Embracing causality in specifying the indirect effects of actions. In C. Mellish, editor, *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1985–1991. Morgan Kaufmann Publishers, 1995.
- [34] F. Lin and R. Reiter. Forget it! In *Proceedings of the AAAI Fall Symposium on Relevance*, pages 154–159. AAAI, 1994.
- [35] D. Makinson. How to go nonmonotonic. In *Handbook of Philosophical Logic*, volume 12, pages 175–278. Springer, 2nd edition, 2005.
- [36] R. Meyer. Relevance is not reducible to modality. In A.R. Anderson and N.D. Belnap, editors, *Entailment: the logic of relevance and necessity*, volume 1, pages 462–471. Princeton University Press, 1975.
- [37] R. Meyer. Ternary relations and relevant semantics. *Annals of Pure and Applied Logic*, 127:195–217, 2004.
- [38] S. Popkorn. *First Steps in Modal Logic*. Cambridge University Press, 1994.
- [39] G. Priest. Paraconsistent logic. In D.M. Gabbay and F. Günthner, editors, *Handbook of Philosophical Logic*, volume 6, pages 287–393. Kluwer Academic Publishers, 2nd edition, 2002.
- [40] G. Restall. Relevant and substructural logics. In D.M. Gabbay and J. Woods, editors, *Handbook of the History of Logic*, volume 7: Logic and the Modalities in the Twentieth Century, pages 289–398. Elsevier North-Holland, 2006.
- [41] G. Robles and J. Méndez. A Routley-Meyer semantics for relevant logics including TWR plus the disjunctive syllogism. *Logic Journal of the IGPL*, 19(1):18–32, 2010.
- [42] M. Thielscher. Ramification and causality. *Artificial Intelligence*, 89(1–2):317–364, 1997.
- [43] A. Urquhart. Semantics for relevant logics. *Journal of Symbolic Logic*, 37:159–169, 1972.
- [44] D. Zhang and N. Foo. EPDL: A logic for causal reasoning. In B. Nebel, editor, *Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 131–138. Morgan Kaufmann Publishers, 2001.
- [45] Y. Zhou and Y. Zhang. A logical study of partial entailment. *Journal of Artificial Intelligence Research*, 40:25–56, 2011.

Katarina Britz
Centre for Artificial Intelligence Research
CSIR Meraka Institute and University of KwaZulu-Natal
South Africa
e-mail: arina.britz@meraka.org.za

Johannes Heidema
Department of Mathematical Sciences
University of South Africa
South Africa
e-mail: johannes.heidema@gmail.com

Ivan Varzinczak
Centre for Artificial Intelligence Research
CSIR Meraka Institute and University of KwaZulu-Natal
South Africa
e-mail: ivan.varzinczak@meraka.org.za